

NUMERICAL SIMULATIONS OF GENERALLY RELATIVISTIC HYDRODYNAMIC SYSTEMS

A Dissertation

Presented to the Faculty of the Graduate School

of Cornell University

in Partial Fulfillment of the Requirements for the Degree of

Doctor of Philosophy

by

Curran Daniel Muhlberger

August 2014

© 2014 Curran Daniel Muhlberger

NUMERICAL SIMULATIONS OF GENERALLY RELATIVISTIC HYDRODYNAMIC SYSTEMS

Curran Daniel Muhlberger, Ph.D.

Cornell University 2014

The study of binary neutron star coalescence, one the most energetic classes of events in the universe, requires calculating the complicated interactions of strong gravity, relativistic fluids, and magnetic fields. The Spectral Einstein Code provides a framework for simulating the inspiral and merger of black holes and neutron stars, but its ability to model the behavior of binary neutron stars and magnetic fields is a recent development. This work describes the implementation of an initial data solver for neutron star binaries, a magnetohydrodynamics module for neutron star and accretion disk evolutions, a pair of basis functions well-suited to spectral representations of neutron star spacetimes, and a selection of other improvements to this research code. It also presents the results of early investigations using these new capabilities, including the effects of magnetic fields on shear instabilities in differentially rotating neutron stars. Such stars may be formed from core-collapse supernovae or low-mass binary neutron star mergers, and fluid instabilities in galactic sources can produce gravitational waves observable by detectors in the near future. We find that strong magnetic fields are capable of suppressing a shear instability, but they also trigger magnetic instabilities whose effects may be just as observable as the original signal.

BIOGRAPHICAL SKETCH

Curran Muhlberger was born in Columbia, MD, in December 1985 and moved to nearby Ellicott City at age five. There he received his education from kindergarten through high school, showing particular aptitude in mathematics and science. Harboring an interest in astronomy all his life (facilitated by books and posters brought home by his parents), in seventh grade he read Stephen Hawking's *A Brief History of Time*, and his focus shifted slightly to theoretical physics and cosmology. In parallel with his academic pursuits, he also took an early interest in computers, taking advantage of the Internet to teach himself programming and keep up-to-date with evolving technologies.

In 2004 he accepted the Banneker/Key scholarship to attend college at the University of Maryland, College Park, difficultly declining a competing offer from Cornell University. There he received bachelor's degrees in physics, mathematics, and astronomy, earning High Honors for research conducted with Dr. Douglas P. Hamilton on the origin of Pluto's inclined orbit. Along the way he continued to develop his computer skills through summer jobs and internships, culminating with the NASA Academy in Huntsville, AL, where he worked with Dr. Charles A. Meegan to evaluate the on-orbit performance of the newly launched Fermi satellite.

The following fall he enrolled in the physics Ph.D. program at Cornell University with an intent to combine his interests in high-performance computing and general relativity. Working with Dr. Saul A. Teukolsky from the summer of 2009, he has done just that, while also finding joy in his new town of Ithaca, NY. Here he took up the hobby of photography while continuing to take pleasure in music and the outdoors.

Upon earning his Ph.D., Curran will begin work as a flight software developer at Space Exploration Technologies, continuing to follow his interests in computing and the cosmos and hoping to help mankind reach a little closer to the stars.

For my parents, whose love and support have been constant every step of the way.

ACKNOWLEDGEMENTS

This dissertation, and my graduate education in general, is the product of my interactions with numerous talented and influential people both at Cornell and elsewhere. I would first like to thank Saul Teukolsky for accepting me into his research group, for giving me the freedom to pursue my diverse interests at my own pace, and for advising me over the course of my graduate career. His knowledge and patience were invaluable as I toiled to complete my projects to my own satisfaction. I would also like to Larry Kidder and Francois Foucart for their help on untold occasions, especially as I first learned to use and extend the Spectral Einstein Code, our primary research tool. Matt Duez provided direction and encouragement during the largest project represented here, made all the more valuable by the fact that he provided it from three time zones away.

My fellow graduate students under Saul also deserve thanks for providing me with ideas, entertainment, and feedback. Andy Bohn, François Hébert, Dan Hemberger, Kate Henriksson, and Will Throwe made the sixth floor of the Space Sciences Building a comfortable place to both work and play. Mike Boyle, Geoffrey Lovelace, and Leo Stein similarly enhanced my research environment from their postdoctoral positions.

The broader SXS collaboration has enabled all of us to accomplish far more than we could on our own, and I would like to extend particular thanks to Harald Pfeiffer, Mark Scheel, Béla Szilágyi, Christian D. Ott, and Fatemeh Hossein-Nouri for their contributions of both code and expertise. I would also like to thank Cornell physics professors Ira Wasserman and Matthias Liepe for serving on my special committee, as well as all of the professors, Matthias included, whose courses I had the pleasure of assistant-teaching.

While my graduate career spanned just the past six years, the skills and character traits that allowed me to complete it are the result of a much longer education and upbringing. There have been too many people who have had a positive influence on my life to mention them all here, but I would like to give special thanks to all of my teachers, supervisors, and

mentors, to my wonderful friends from my home state of Maryland, to my girlfriend Linda Huang, and most of all to my parents and siblings. Without all of you, I would not be who I am today.

Finally, I would like to acknowledge the funding I received from the National Science Foundation and the Sherman Fairchild Foundation, which helped finance my graduate education at Cornell.

TABLE OF CONTENTS

Biographical Sketch	iii
Dedication	iv
Acknowledgements	v
Table of Contents	vii
List of Figures	ix
List of Tables	x
Preface	xi
Previously published work	xi
Notation	xi
1 Introduction	1
2 Binary neutron star coalescence	6
2.1 Generally relativistic stars	7
2.2 Binary neutron star initial data	10
2.3 Evolutions and future work	14
3 An upwind constrained transport scheme for magnetohydrodynamics	16
3.1 Background on finite volume methods	17
3.2 Generally relativistic magnetohydrodynamics in SpEC	19
3.3 Electromagnetic gauges	27
3.4 Test problems	29
4 Shear instabilities in differentially rotating neutron stars	39
4.1 Motivation	39
4.2 Setup	41
4.3 Analysis	46
4.4 Results	50
4.5 Summary	70
5 Conclusion	72
A Basis functions for cylinders and spheres	74
A.1 Background on spectral methods	75
A.2 Zernike polynomials	77
A.3 Spectral method	79
A.4 Scalar wave test	84

A.5 Zernike analysis	86
B Quadratic Bézier smoothing for piecewise polytropes	90
Bibliography	92

LIST OF FIGURES

2.1	Convergence of NS–NS ID solver	12
2.2	Domain decomposition for NS–NS ID solver	13
2.3	Comparison of NS–NS ID results	15
3.1	Interpolation artifacts in algebraic gauge	29
3.2	Final rest-mass density of the shock tests	32
3.3	Final velocity of the shock tests	33
3.4	Error in the final value of u^y for the <i>wave</i> test	34
3.5	Error norm for the Bondi test	38
4.1	Illustrations of magnetic field lines at early and intermediate times	44
4.2	Illustration of x – z slice of domain decomposition	46
4.3	Consistency of the growth rate of the low- $T/ W $ instability	51
4.4	Growth and saturation of the unmagnetized low- $T/ W $ instability	52
4.5	Comparison of spectrograms of η_\times	54
4.6	Effects of reconstructor choice on the growth of $ \eta $	55
4.7	Relative power of ρ in azimuthal modes	56
4.8	Range of behavior of distortion parameter η at different magnetic field strengths for $n_s = 1$	57
4.9	Range of behavior of distortion parameter η at different magnetic field strengths for $n_s = 2$	58
4.10	Energy exchange for three magnetic field strengths	59
4.11	Lagrangian displacement of tracer particles	60
4.12	Spectrograms of the quadrupole moment I_{xy}	61
4.13	Magnitude of radial component of B -field	62
4.14	Magnitude of radial component of B -field vs. radius vs. time	63
4.15	Power of b^2 in azimuthal modes	65
4.16	Growth of the maximum of the cylindrical components of the B -field	67
A.1	Convergence and stability of spherical scalar wave	86

LIST OF TABLES

3.1	Initial data for shock tests	31
4.1	Basic properties of the neutron star	43
4.2	Summary of the magnetic configurations studied	45

PREFACE

Previously published work

Significant portions of this thesis have been submitted for publication in *Physical Review D* under the title *Magnetic effects on the low- $T/|W|$ instability in differentially rotating neutron stars*. That work was co-authored with Fatemeh Hossein Nouri, Matthew D. Duez, Francois Foucart, Lawrence E. Kidder, Christian D. Ott, Mark A. Scheel, Béla Szilágyi, and Saul A. Teukolsky. A preprint was submitted to the arXiv.org e-print archive on May 9, 2014, and is available under the identifier arXiv:1405.2144 [1].

Notation

Physical equations in this work are written in geometrized units where the speed of light c and the gravitational constant G are set equal to 1. Residual dimensions can be expressed as powers of mass, for which we often choose the mass of the Sun, M_\odot , as the unit. When discussing electromagnetic fields in the context of simulation formalism and stability analysis, we adopt the Lorentz-Heaviside convention, absorbing a factor of $1/\sqrt{4\pi}$ into the definition of the magnetic field \mathbf{B} . However, when presenting physical results, we express all quantities in CGS–Gaussian units. In particular, $\mathbf{B}_{\text{LH}} = \mathbf{B}_{\text{G}}/\sqrt{4\pi}$.

We denote the Cartesian coordinates of space by x, y, z . The coordinate distance from the origin is denoted by $r \equiv \sqrt{x^2 + y^2 + z^2}$. When cylindrical coordinates are used, $\varpi \equiv \sqrt{x^2 + y^2}$ represents the coordinate distance to the z -axis, and $\phi \equiv \tan^{-1}(y/x)$ defines a point's azimuthal angle. In spherical coordinates, the polar angle is represented by $\theta \equiv \cos^{-1}(z/r)$.

Tensor indices from the beginning of the Latin alphabet (a, b, \dots) represent spacetime components without reference to any particular coordinate system, while indices from

the Greek alphabet (μ, ν, \dots) range from 0 to 3 and correspond to components in our Cartesian coordinate system of (t, x, y, z) . Indices from the middle of the Latin alphabet (i, j, \dots) range from 1 to 3 and represent spatial Cartesian components. Spatial vectors may alternately be typeset in **boldface** when referred to as a whole.

1 INTRODUCTION

The study of high-energy astrophysical systems is entering an exciting era. We already observe events like supernovae, gamma-ray bursts, and binary neutron star inspirals across the electromagnetic spectrum, from radio waves to X-rays, and even in neutrinos. But the gravitational wave interferometers coming online in the near future, such as Advanced LIGO and Advanced Virgo, should provide a brand new way of looking at these events and others through their gravitational wave emission. Pulsar timing arrays are also poised to usher in the era of gravitational wave astronomy, though for a different selection of sources. With both varieties of experiment expecting a first detection around 2016, there is a present need for better theoretical modeling of expected sources to aid in both detection and subsequent data analysis.

Gravitational waves are a prediction of Einstein's 1916 theory of general relativity. In Newton's classical theory of gravity, two objects may orbit each other indefinitely, resulting in completely stable binary systems. General relativity, however, requires that such systems radiate away their gravitational potential energy in the form of propagating ripples in spacetime: gravitational waves. We have strong indirect evidence of this emission, most famously from the decaying orbit of the Hulse–Taylor binary pulsar [2]. Other predictions of general relativity have been confirmed as well in the *weak field* regime, where gravity is weak and velocities are small compared to the speed of light. However, gravitational waves have yet to be directly detected, and such detections could give us valuable insight into the behavior of gravity in *strong field* regimes near black holes or during the early universe.

As ripples in spacetime, one of the observational consequences of gravitational waves is a *strain*—a relative shortening or lengthening of the distance between two points when a wave passes through. Laser interferometer observatories like LIGO and Virgo look for such a strain by monitoring the distance between a pair of mirrors relative to that between another pair in an orthogonal direction. If one arm of this arrangement contracts while

the other expands, in an oscillatory fashion, then the resulting signal is consistent with a gravitational wave. Unfortunately, the expected strain from astrophysical sources is incredibly small, on the order of 10^{-22} . For LIGO, this means that the change in arm length is less than 1/1000 the size of a proton. Signals at that level are overwhelmed by noise, requiring matched filtering to pull out a detection. In other words, to find such a signal, one has to know exactly what one is looking for to begin with.

Theoretical modeling seeks to remedy this situation by providing templates of expected signals to search for. While some portions of these templates can be constructed analytically, the portions sensitive to strong gravity and high velocities can only be provided by numerical simulations, as Einstein's equations are highly nonlinear in that regime. When matter is present, these simulations must also incorporate the complexities of nuclear physics, magnetic fields, radiation transport, etc. This additional microphysics significantly complicates simulations and vastly widens the parameter space of potential systems that must be studied. On the other hand, when observables are sensitive to these effects, there is that much more we can infer about the astrophysical processes taking place. Determining the equation of state at nuclear densities is one exciting possibility, as it is currently an unknown facet of fundamental physics that cannot be probed directly in terrestrial laboratories.

The Spectral Einstein Code

The *Simulating eXtreme Spacetimes*¹ collaboration focuses on accurate simulations of compact binary coalescence, supernova collapse, rotating neutron stars, and accretion disks. For many of these systems, our primary tool is the Spectral Einstein Code (SpEC). Simulations generally consist of four phases:

1. Choose the parameters of the system.

¹<http://www.black-holes.org/>

2. Construct the initial data.
3. Evolve the system through the period of interest.
4. Analyze the results.

Constructing initial data consistent with Einstein’s equations requires solving a system of elliptic partial differential equations (PDEs), while its evolution is the solution to a system of hyperbolic PDEs. Away from the singularity of black holes, functions representing the curvature of spacetime are expected in general to be smooth. This makes spectral methods an appealing choice for solving the equations. While finite difference approaches have been used successfully by a number of groups, spectral methods converge to the true solution exponentially with resolution, making them more accurate and efficient for vacuum spacetimes. Unfortunately, adding matter to the system introduces the possibility of dynamic surfaces and shocks. These discontinuities are a poor match to high-order spectral methods, but they can be handled robustly in a finite volume framework. `SpEC` takes a hybrid approach, using spectral methods to evolve the spacetime while employing a finite volume grid, limited to non-vacuum regions, to evolve the matter in the system.

When solving a PDE, we divide the spatial domain of the problem into subdomains. The solution in each subdomain is then represented by a sum of basis functions adapted to the local geometry. When black holes are present, we choose to excise a causally-disconnected region of the domain enclosed by the event horizon, thus avoiding the need to solve the equations in the vicinity of the singularity (this is in contrast to “puncture” methods, which leave the singularity on the simulation grid). The regions surrounding these excision zones are covered with subdomains in the shape of spherical shells. A neutron star, on the other hand, requires a subdomain covering its center. A similar need exists for regions of space covered by cylindrical shells. Zernike polynomials (and their 3D generalizations) form a basis well-suited to these subdomain topologies, but their limited exposure in the literature

makes correct and efficient implementations non-trivial to construct. As an aid to future users, our implementation of these basis functions is discussed in Appendix A.

Scope of this work

One class of system of particular interest to gravitational wave observatories is binary neutron stars. Unlike binaries containing black holes, these systems have already been observed, in the form of binary pulsars, giving us tighter estimates for their abundance. The Advanced LIGO–Virgo network expects to observe between 0.4 and 400 such mergers per year [3], making them potentially the most likely source of gravitational wave detections. The inspiral phase of these systems can be modeled quite well by post-Newtonian methods, making expensive simulations unnecessary for detection. However, the late inspiral, disruption, and merger phases cannot be accurately modeled without simulations, and the details of the signal from these phases contain interesting information on the equation of state. It is also unclear how late into the inspiral the post-Newtonian predictions can be trusted. `SpEC` is well-suited to study these systems, but first requires accurate and physically-relevant initial data. The construction of this data is discussed in Chapter 2, along with preliminary results from its subsequent evolution.

While the gravitational waves from the late inspiral and plunge of neutron stars in binaries are well-predicted by just a few (albeit complicated) physical ingredients, their disruption, merger, and accretion disk behavior depend strongly on a number of additional microphysical effects. These include magnetic fields, neutrino emission, and nuclear reactions. Modeling these effects is essential for *multi-messenger astronomy*—correlating gravitational and electromagnetic signals from the same events. The effects of magnetic fields are particularly relevant to the hypothesis that short gamma-ray bursts may be driven by black hole–neutron star mergers. We therefore created a magnetohydrodynamics (MHD) module for `SpEC` capable of simulating the evolution of a magnetic field in a conducting fluid and its effects on that fluid’s motion. Combining techniques from several

existing codes, this module prevents the formation of spurious magnetic monopoles during evolution and limits the unphysical impact of interpolation artifacts. Code improvements made during the development of this module also benefited our simulations of systems with neutrinos [4]. We discuss the details of our MHD module in Chapter 3.

Magnetic fields are responsible for a number of important phenomena in astrophysical systems, from powering accretion to forming jets. They can also modify or suppress interesting behavior present in unmagnetized systems. One example we studied is their effect on a shear instability in differentially-rotating neutron stars, such as those formed after supernova collapse or low-mass binary neutron star coalescence. The shear instability is a potential source of gravitational radiation, but strong magnetic fields can interfere with the mechanism of the instability, suppressing the signal. As discussed in Chapter 4, we find that fields strong enough to have this effect are also susceptible to a range of magnetic instabilities, resulting in potentially accelerated signal growth and magnetized outflows of matter. We also find that the behavior of these systems depends strongly on the specific numerical methods used to simulate them, suggesting that some previous studies may have been under-resolved. This strong dependence is an interesting result in its own right, especially if it proves to be more general, and will likely be the topic of future investigation.

Other future projects involving the MHD module include a collaboration with the University of Illinois on simulating magnetized high-spin black hole–neutron star mergers. Code comparison tests such as this will help us understand the uncertainties in our simulations’ results and identify errors that may exist in the implementation of our solution methods. As we gain confidence in our current methods and models, we are better able to extend them to the larger computers and more complete microphysics needed to faithfully represent and ultimately understand the rich astrophysical phenomena we observe in the high-energy universe and the potential new sources we may discover through gravitational waves.

2 BINARY NEUTRON STAR COALESCENCE

One of the most promising sources of detectable gravitational waves is the merger of binary neutron stars. With an expected detection rate of about 40 per year for the Advanced LIGO–Virgo network [3], signals from these events should be observed in the near future. Additionally, most of the in-band signal will come from the inspiral phase, which is well-approximated by post-Newtonian theory; this facilitates template generation and thus detection. But while there is much to be learned from simply determining the rates of these events, signals from the late inspiral and merger hold the key to answering questions about the equation of state of dense nuclear matter, and modeling these signals requires full numerical relativity simulations. Additionally, these simulations can be used to calibrate and improve effective one-body models, which promise a cheaper way of generating signals more accurate than post-Newtonian theory provides. Finally, with the inclusion of additional microphysics, simulations of binary neutron stars can help determine whether they are the source of short gamma-ray bursts [5].

Performing simulations of binary neutron stars first requires appropriate and accurate initial data. There are two sides to this problem: first, the data must represent an astrophysically-relevant system, containing stars with the right masses, orbits with the right shape, etc.; and second, the data must satisfy the Hamiltonian and momentum constraint equations of general relativity. The latter are elliptic PDEs that can be solved numerically with established methods, but first their “free data” must be specified. This freedom is what allows the properties of the system to be specified, but the correspondence between the data and these properties is far from straightforward. We thus take an iterative approach: we start with a guess possessing the properties we want, then repeatedly solve the constraint equations and tweak the result to drive it back towards our desired specifications. The first step, then, is to specify the properties of the constituent neutron stars as they would appear in isolation.

2.1 Generally relativistic stars

The structure of relativistic stars in equilibrium is determined by the Tolman–Oppenheimer–Volkoff (TOV) equations. At a given coordinate distance r from the center of the star, let P denote the pressure, ρ_E the mass–energy density, and m the enclosed ADM mass¹. Then,

$$\partial_r P = \frac{(\rho_E + P)(m + 4\pi P r^3)}{r(2m - r)} \quad (2.1)$$

$$\partial_r m = 4\pi \rho_E r^2. \quad (2.2)$$

To close this system of equations, we need an equation of state relating P to ρ_E .

We might additionally be interested in M_0 , the total rest mass, a.k.a. baryon mass, of the star. It can be found by integrating

$$M_0 = \int_0^R \frac{4\pi \rho r^2}{\sqrt{1 - 2m/r}} dr, \quad (2.3)$$

where ρ is the rest mass density, related to P and ρ_E through the equation of state, and R is the final radius of the star. For numerical convenience, this can be cast as a third ordinary differential equation (ODE) to be integrated along with the TOV equations, yielding

$$\partial_r m_0 = \frac{4\pi \rho r^2}{\sqrt{1 - 2m/r}}. \quad (2.4)$$

While this system of ODEs is straightforward to integrate numerically, we do not know a priori how far to integrate; i.e. what the final radius of the star will be. But we do know that at the surface of the star (if it exists), the pressure will be zero. Therefore, if we could integrate the equations in terms of P instead of r , we could integrate to the final condition $P = 0$. In this form, the above equations become

$$\partial_P r = \frac{r(2m - r)}{(\rho_E + p)(m + 4\pi P r^3)} \quad (2.5)$$

$$\partial_P m = 4\pi \rho_E r^2 \partial_P r \quad (2.6)$$

$$\partial_P m_0 = \frac{4\pi \rho r^2}{\sqrt{1 - 2m/r}} \partial_P r. \quad (2.7)$$

¹The subscript “ADM” denotes quantities in the Arnowitt–Deser–Misner formalism; see, e.g., Baumgarte and Shapiro [6].

With these two forms of the equations, we can integrate in terms of r until some cutoff criterion (such as when the pressure reaches half its initial value), then integrate in terms of P starting from that state². Thus, we can specify initial conditions at $r = 0$ and terminate the solution when $P = 0$. If required, the metric inside the star can also be found with an additional integration; outside the star, it is equivalent to the Schwarzschild metric with energy $M \equiv m(R)$.

Given these equations and an ODE solver, we can solve for the structure and properties of neutron stars given two inputs: the central density and an equation of state. Often we are looking for specific final properties and need to determine which initial conditions will give rise to them. For example, we might want to know which central density and which member of a parameterized family of equations of state will give rise to a neutron star with a given mass and radius. By wrapping the solver in a nonlinear rootfinder, we can answer these questions, thus determining the fields of a neutron star spacetime for any desired neutron star properties, assuming a stable solution exists. This becomes the main ingredient in the initial guesses for binary solutions—once the properties of the constituent neutron stars are specified in isolation, their isolated solutions can be found and blended together in a corotating frame.

Polytropes

Polytropic equations of state are featured prominently in the study of classical stars, and they have useful properties when studying relativistic ones as well (though ultimately more realistic equations of state will be used when making comparisons to observations). A polytropic equation of state is one in which the pressure is related to the density via

$$P = \kappa \rho^\Gamma \tag{2.8}$$

²Equations 2.5–2.7 evaluate to an indeterminate $0/0$ at the center of the star, requiring Eqs. 2.1–2.2 & 2.4 to initialize the solution process.

for some constants κ and Γ . The mass–energy density is $\rho_E = \rho(1 + \epsilon)$, with the specific internal energy ϵ given by

$$\epsilon = \frac{P}{\rho(\Gamma - 1)} . \quad (2.9)$$

A polytropic equation of state yields a natural lengthscale R_{poly} :

$$R_{\text{poly}} \equiv \kappa^{\frac{1}{2(\Gamma-1)}} . \quad (2.10)$$

When combined with geometrized quantities (where $c = G = 1$), this means that simulations of polytropes can be performed in a dimensionless fashion and the results scaled to match a choice of physical units. In practice, this allows one to construct a neutron star of a given *compactness*: the ratio of its mass to its circumferential (a.k.a. areal) radius in geometrized units. Then, after simulating, the results can be scaled to correspond to a neutron star of any given physical mass.

Towards realistic equations of state

While the rescalability of polytropes is convenient, and while the simplicity of the equation of state makes simulations efficient (and while its smoothness improves the performance of spectral methods), ultimately they are an inaccurate representation of physical reality. More realistic initial data will use equations of state inspired by models of physics at nuclear densities. For simplicity, we assume the stars are uniformly cold (temperature $T = 0$) initially. Equations of state that depend on composition also require an assumption regarding the initial composition (such as β equilibrium) in order to be expressed in barotropic form (P as a function of ρ alone). Results are no longer rescalable with these equations of state, as they are tied to physical units. One consequence of this loss of scaling is that a star’s mass and compactness may no longer be independently specified.

Uryū et al. have constructed binary neutron star (NS–NS) initial data using piecewise polytropic equations of state based on fits made by Read et al. [8] (but with a simpler treatment of the low-density crust)—these are efficient to compute and only need to be

implemented once to model a wide variety of nuclear physics models [7]. One might worry that the C^0 smoothness of such fits might interfere with the convergence of spectral methods. Quadratic Bézier curves can be used in the exponent to smooth these transitions to C^1 with only a minor impact on performance. This has additional benefits in evolutions, where quantities requiring thermodynamic derivatives, such as the sound speed, are now C^0 instead of discontinuous. Our implementation of quadratic Bézier smoothing is outlined in Appendix B.

Our TOV and NS–NS initial data solvers are compatible with any barotropic equation of state. Convergence is slower and final accuracy reduced for non-polytropic ones (in particular, the equations behave best at the surface when $P \propto \rho^2$), but for stars with compactness $\lesssim 0.17$, we can reliably produce initial data with sufficient accuracy for our current evolutions.

2.2 Binary neutron star initial data

To produce initial data for binary neutron star systems, I wrote an initial data solver based on the work of Foucart et al. for black hole–neutron star (BH–NS) systems [9]. As in that work, we start by considering systems in quasiequilibrium, where time derivatives vanish in a corotating frame (this neglect of the small radial velocity will be addressed later). We take the metric to be conformally flat,

$$ds^2 = -\alpha^2 dt^2 + \phi^4 \delta_{ij} (dx^i + \beta^i dt)(dx^j + \beta^j dt), \quad (2.11)$$

and solve for the lapse α , shift β^i , and conformal factor ϕ using the extended conformal thin sandwich equations [10]. The matter in the stars is modeled as a cold ($T = 0$) perfect fluid with an irrotational velocity profile. The irrotational limit allows a straightforward solution for the velocity and is a more realistic approximation than the corotating limit, as the effective viscosity of neutron star matter is insufficient to synchronize the stars’ spins with their orbital frequency [11, 12].

A particular binary neutron star system is specified in terms of the equation of state of neutron star matter, the baryon masses of both stars, and their coordinate separation. The solver then uses the above assumptions of quasiequilibrium and cold irrotational flow to determine the metric and matter content of the corresponding spacetime. Since the initial data problem consists of several coupled equations, the solver takes an iterative approach, with each iteration composed of a number of substeps (this procedure closely follows Sec. III.C of Foucart et al., which should be consulted for additional details).

First, given a trial matter distribution, we find an approximate solution to the elliptic thin sandwich equations by taking a single step of a nonlinear solver. By imposing force balance at the centers of the stars, we then adjust the orbital frequency of the binary. We also modify the enthalpy of the matter to drive the locations of its maxima to the specified stellar centers, thus controlling the stars' separation. Finally, we approximately solve the elliptic equations imposing irrotational flow (constrained to preserve the baryon masses of the stars) and feed the output to the next step of the iterative procedure. All of these updates are made using a relaxation scheme to aid convergence. The convergence of our solver, compared with that of Gourgoulhon et al.'s LORENE code [13], is illustrated in Fig. 2.1.

Throughout the solution process, data is represented on a spectral grid composed of hexahedra, cylindrical shells, and spherical shells (see Fig. 2.2), and approximate solutions to the elliptic equations are provided by the *spells* framework [14]. We periodically evaluate the grid and adjust it to better conform to the stars' surfaces. By placing subdomain boundaries at these surfaces, the discontinuities there will not affect the spectral convergence of the method. Additionally, we occasionally perturb the centers of the stars to control the ADM linear momentum of the system. During this procedure, the centers are *not* constrained to be collinear with the center of revolution, and the separation of the stars may deviate slightly from the initially specified value. Separations reported are therefore measured from the final solution.

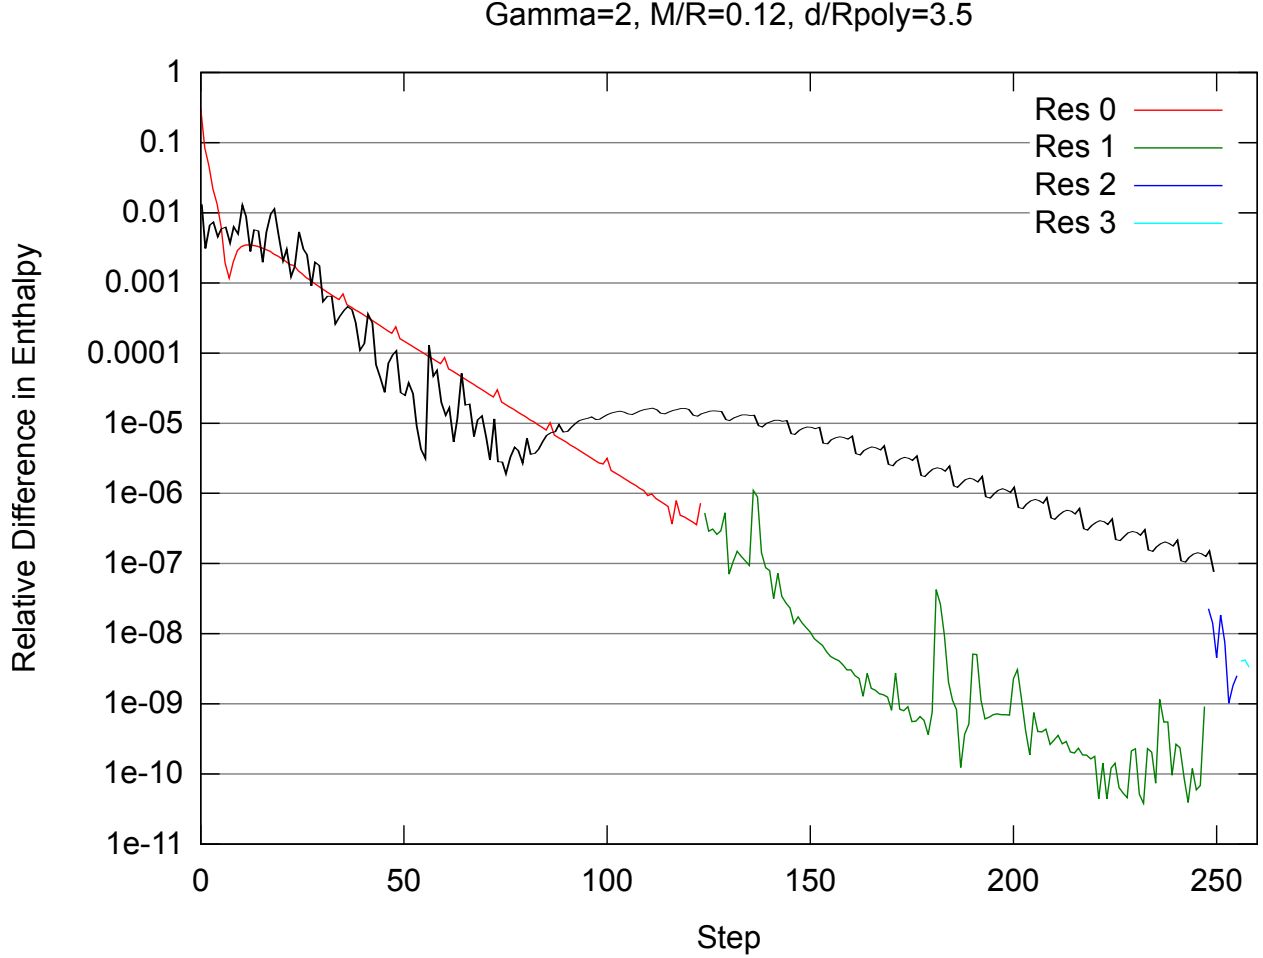


Figure 2.1: Convergence of NS-NS initial data compared to Gourgoulhon et al. [13] (black curve). Following that work, we monitor the sum of the absolute value of the difference in the specific enthalpy at each point between two successive iterations of the solver, normalized by the sum of the specific enthalpy over all points. The stars obey a polytropic equation of state with $\Gamma = 2$ and have a compactness of 0.12. They are separated by a coordinate distance of 3.5 in polytropic units.

The details of the two solvers are quite different—our solver employs a sequence of grids with increasing resolution (“Res 0”–“Res 3”), while Gourgoulhon et al. use a single resolution. Additionally, they do not attempt to fix the baryon mass until around iteration 70, while we fix the baryon mass at all times. Initial convergence rates are similar, but ultimately our solver achieves a more precise solution in fewer iterations. The periodic spikes in our “Res 0” (red) curve correspond to steps when we change our grid’s map to adapt to the neutron stars’ new surfaces.




Figure 2.2: Domain decomposition for NS-NS initial data problem. Subdomain types include rectangular prisms, spherical shells, and cylindrical shells and are allowed to overlap. The boundaries between the blue and green spheres are mapped to match the surfaces of the neutron stars. This portion of the domain is itself surrounded by a large spherical shell whose outer radius is mapped to a very large (effectively infinite) value.

When constructing strictly quasiequilibrium data, the solver chooses the orbital angular velocity Ω by requiring force balance at the centers of the stars. Later, when subsequently refining the initial data, Ω is fixed. By adding an initial radial velocity, we relax the quasiequilibrium assumption in order to more accurately model inspiral conditions and reduce the initial eccentricity of the orbits. Primordial compact object binaries are expected to have very low eccentricity by the time they enter LIGO's band [15], as gravitational wave emission will circularize their orbits over the course of their very long inspiral. The magnitude of the radial velocity is chosen by evolving each trial set of initial data for a short time in order to measure the eccentricity, then adjusting the (fixed) orbital frequency

and radial velocity according to a heuristic procedure based on the work of Pfeiffer et al. [16] and repeating until that eccentricity is below 10^{-3} .

Quasiequilibrium sequences

Results from our code closely match those of Gourgoulhon et al.. In particular, we can accurately reproduce the quasiequilibrium sequences of Taniguchi and Gourgoulhon [17, 18], achieving agreement to within 0.05% for the total ADM mass and orbital frequency over a range of mass ratios and EOSs (for additional comparisons, see Fig. 2.3). Using our implementation of piecewise polytropic equations of state, we can also reproduce the results of Taniguchi and Shibata [19]. Quasiequilibrium sequences were once a useful way to study the late inspiral of binary compact objects without the complications of evolving the full Einstein equations. However, now that such evolutions are reliable, we can study this and later phases of mergers with much higher fidelity.

2.3 Evolutions and future work

Initial data generated by this solver has been evolved by Haas et al. (in preparation) for more than 22 orbits. For this study, we considered a gamma-law equation of state of the form

$$P = \kappa \rho^\Gamma + \rho T \quad (2.12)$$

$$\epsilon = \frac{1}{\Gamma - 1} \frac{P}{\rho} \quad (2.13)$$

with $\Gamma = 2$ and $\kappa = 123.6 M_\odot^2$. Both neutron stars have a baryon mass of $M_0 = 1.779 M_\odot$, corresponding to an isolated TOV star with an ADM mass of $M = 1.64 M_\odot$, circumferential radius of $R_{\text{areal}} = 15.1 \text{ km}$, and a compactness of $M/R_{\text{areal}} = 0.16$. Because of the large separation, the binding energy is small, $E_b = 6.7 \times 10^{-3} M_\odot$, and the total ADM mass of the system is approximately twice the ADM mass of an isolated constituent star. In their

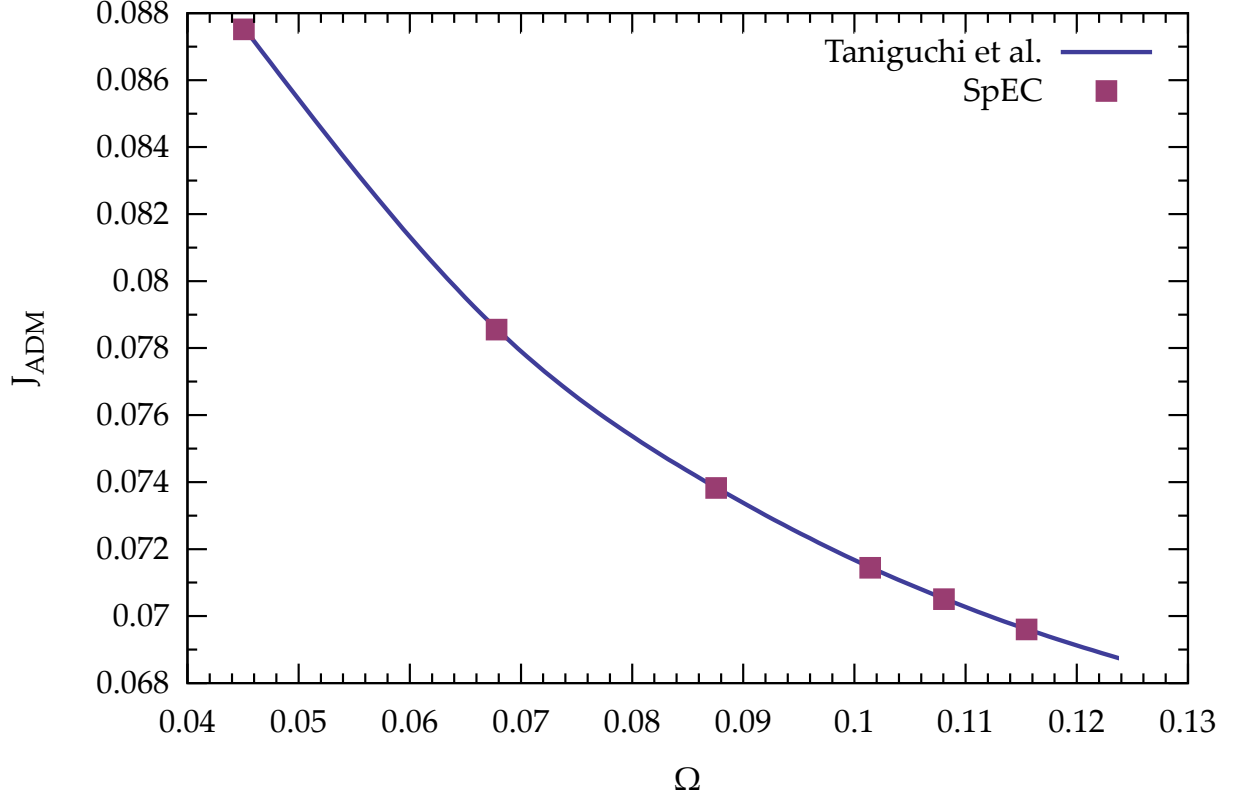


Figure 2.3: Comparison of quasiequilibrium sequence with Taniguchi et al. for a polytropic equation of state with $\Gamma = 2$ and two equal-mass neutron stars with a compactness of 0.14. Plotted is the ADM angular momentum J_{ADM} at various separation distances, parameterized by the angular frequency of the orbit Ω at each separation; both are expressed in units where $R_{\text{poly}} = 1$. We have made similar comparisons for non-equal-mass neutron stars (compactnesses of 0.14 vs. 0.16) and for $\Gamma = 2.5$ equations of state.

binary configuration (and in our initial gauge), they each extend to an average coordinate radius of 12 km, and their centers are separated by a coordinate distance of 81 km. This system has an orbital frequency of $\Omega/2\pi = 133$ Hz and an eccentricity of less than 9×10^{-4} .

Results from evolving this data will be used to estimate the minimum number of numerical wave cycles required before hybridization with post-Newtonian waveforms is sufficiently accurate for LIGO. The waveforms will also be compared with those of binary black holes, highlighting the effects of neutron star deformability on the resulting gravitational wave signal.

3 AN UPWIND CONSTRAINED TRANSPORT SCHEME FOR MAGNETOHYDRODYNAMICS

While spectral methods work well for evolving smooth spacetime fields or constructing initial data with sharp features at known locations, they are less well-suited for evolving dynamic matter fields, which can form discontinuous shocks. For this reason, SpEC has traditionally evolved matter on a separate finite volume grid, using interpolation to communicate between the spacetime and matter solutions. However, it was limited to evolving unmagnetized fluids, and magnetic fields are extremely important in high energy astrophysics. Therefore, I embarked on the development of a magnetohydrodynamics (MHD) module for SpEC.

Like the Einstein equations, the equations of electromagnetism both provide a rule for evolution and impose a set of constraints on the solution. Magnetic fields must form a divergence-free vector field; in other words, there may be no magnetic monopoles in the solution. While the analytic evolution equations will preserve this constraint if it is satisfied by the initial data, a numerical scheme may generate constraint violations, resulting in unphysical solutions. There are several approaches to preserving this constraint numerically; we concentrated our efforts on an upwind constrained transport scheme, though we have also experimented with other techniques. Here we describe the context in which we evolve magnetized fluids and outline the specific approach we take to model the magnetic fields.

3.1 Background on finite volume methods

Many laws of nature are conservation laws—some quantity is never created or destroyed, but merely moves from place to place. Written as a PDE, this would be represented as

$$\partial_t \psi + \partial_i F^i = 0, \quad (3.1)$$

where ψ is a conserved field and F is its flux. More generally, the right-hand side may consist of a non-zero source term and still be regarded as being expressed in “conservative form.”

In the finite volume framework, consider the original integral form of such an equation over a small volume V (i.e. a grid cell):

$$\partial_t \int_V \psi d\mathbf{x} + \int_{\partial V} \mathbf{F} \cdot d\mathbf{\Sigma} = 0. \quad (3.2)$$

That is, the amount by which ψ changes in the volume of the cell is equal to the net flux of ψ flowing through the cell’s boundaries. Dividing this by the volume of the cell yields an evolution equation for the cell average of ψ that depends on the divided differences of surface averages of the flux.

Evaluating the flux on cell faces requires interpolation of data from cell centers. The details here distinguish *finite volume* and *finite difference* schemes, but to second order they are mostly equivalent. Solutions to nonlinear conservation laws, such as the relativistic Euler equations, often contain discontinuities, or *shocks*, which interact poorly with simple polynomial interpolation, especially at high order. Instead, we use a procedure called *reconstruction*, which performs two interpolations for each face: one biased by data from the left, the other biased by data from the right (this enables *upwind* schemes where the final value of the flux takes into consideration the direction of information flow). These interpolations also incorporate some variety of slope limiting, reducing their order of accuracy in certain conditions to avoid creating unphysical results (a negative density, for instance). With two values for the flux on every face, we then compute a single *numerical*

flux F^* based on the values of the fluxes, solution, and characteristic speeds. The procedure for computing this numerical flux basically considers every face to be a potential shock front and produces some approximate solution to the corresponding Riemann problem. For most of our simulations, we choose the HLL flux [20]:

$$F^{*i} = \frac{c_+^i F^{iL} + c_-^i F^{iR} - c_+^i c_-^i (\psi^R - \psi^L)}{c_+^i + c_-^i}, \quad (3.3)$$

where the superscripts L and R indicate values reconstructed on the left and right sides of the face, respectively, and c_\pm^i represent the two fastest characteristic speeds at that face in the i direction.

Our favored reconstruction algorithm is the fifth-order weighted essentially nonoscillatory (WENO) scheme [21, 22]. It considers a five-point 1D stencil containing the face in question and constructs three second-degree polynomials through each of its three-point substencils. It then estimates the smoothness of each polynomial and weights their contributions to the desired face according to their smoothness. The weighting procedure is designed so that fifth-order accuracy is recovered when the solution is smooth. However, when smoothness is lacking, the effective order drops and severe oscillations are largely avoided. Shifting the stencil by one point provides a value biased from the opposite direction.

The choice of smoothness indicator and the construction of weights from its output allows for a family of fifth-order WENO reconstructors. To avoid division by zero when the solution is very smooth, a small offset is often added to the smoothness indicator when forming the weights¹. Shu recommends an offset of 10^{-6} [23]; however, this offset is not dimensionless and can negatively impact the quality of the reconstruction when the values being reconstructed are $\ll 1$. We instead add $10^{-17}(1 + \sum_i y_i)$, where y_i are the values being interpolated in the corresponding substencil. The resulting nonlinear weights perform much better in low-density regions.

¹This offset must be added after the rest of the smoothness indicator has been computed to avoid being lost in floating-point roundoff. Optimizing compilers that assume commutativity of floating-point addition can otherwise nullify the effect of this step.

We also occasionally use a second-order monotized centered (MC2) reconstructor [24]. This algorithm is generally more robust, but it is less accurate in smooth regions and especially at local extrema, where it always drops to first order. In addition to its use as a diagnostic tool, we also switch to MC2 in the vicinity of black hole horizons, which helps avoid horizon-crossing stencils. For a more complete discussion of this problem, see Hawke et al. [25].

Reconstruction and interpolation algorithms in SpEC are now all implemented as 1D kernels acting on continuous stripes of data. Higher-level logic can then apply any kernel to a variety of situations, including 3D domains, axisymmetric domains, domains with excision masks, and even staggered grids. Ghost zones—extra gridpoints at the boundary of a subdomain whose evolution is controlled by a neighboring subdomain or a boundary condition—ensure that full-accuracy results are available where they are needed and that parallel and serial executions produce identical results.

3.2 Generally relativistic magnetohydrodynamics in SpEC

Here we outline the full evolution scheme for a single magnetized neutron star. Our generally relativistic magnetohydrodynamics (GRMHD) code builds on top of the work of Duez et al. [26]; additional non-magnetic details are provided by Foucart et al. [27].

Metric evolution

As in previous studies using SpEC, the spacetime is evolved according to Einstein’s equations in generalized harmonic form [28], and the coordinates x^a are assumed to obey

$$g_{ab} \nabla^c \nabla_c x^b = H_a \quad (3.4)$$

for some gauge source function H_a (where ∇_a is the covariant derivative operator associated with g_{ab}). To reduce the equations to first-order form, we evolve the derivatives of the

spacetime metric g_{ab} , defined as

$$\Phi_{iab} \equiv \partial_i g_{ab} \quad (3.5)$$

$$\Pi_{ab} \equiv -n^c \partial_c g_{ab} , \quad (3.6)$$

where n^a is the normal to a spacelike slice. This slicing defines a 3 + 1 decomposition of the metric into a 3-metric γ_{ij} , lapse α , and shift vector β^i (see, e.g., Baumgarte and Shapiro [6]), with line element given by:

$$ds^2 = -\alpha^2 dt^2 + \gamma_{ij}(dx^i + \beta^i dt)(dx^j + \beta^j dt) . \quad (3.7)$$

The spacetime variables g_{ab} , Φ_{iab} , and Π_{ab} are evolved according to the principal parts and constraint damping terms in Appendix A of Foucart et al. [27] (augmented with the matter and magnetic source terms described below), and the gauge source H_a is evolved according to the “frozen” condition in that work. The damping parameters for a system containing a single neutron star are distributed according to:

$$\gamma_0(r) = \frac{0.1}{M} f(r) + \frac{0.1}{M} , \quad (3.8)$$

$$\gamma_1(r) = -1 , \quad (3.9)$$

$$\gamma_2(r) = \frac{1.5}{M} f(r) + \frac{0.1}{M} , \quad (3.10)$$

where $f(r)$ is given by:

$$f(r) = e^{-r/(6M)} \quad (3.11)$$

and M is the ADM mass of the star.

The presence of matter and magnetic fields results in a non-zero stress-energy tensor T_{ab} , and this shows up in additional source terms when evolving the spacetime fields. In particular, the vacuum evolution equation for Π_{ab} is modified as follows:

$$\partial_t \Pi_{ab} = \cdots - 2\alpha \left(T_{ab} - \frac{1}{2} g_{ab} T^{cd} g_{cd} \right) . \quad (3.12)$$

The stress-energy tensor for our treatment of MHD is given in Eq. (3.14). Note that we expect the magnetic contributions to T_{ab} to be small (even for the strongest field strengths considered in Chapter 4, magnetic pressure is at most 1% of fluid pressure at $t = 0$).

Magnetohydrodynamics

The matter in the system is modeled as a perfect fluid with rest-mass density ρ , specific internal energy ϵ , and 4-velocity u^a . An equation of state relates ρ and ϵ to the fluid's pressure P , and from these, the relativistic specific enthalpy is $h = 1 + \epsilon + P/\rho$. We denote the Lorentz factor corresponding to the fluid's velocity by $W_L \equiv \alpha u^t$.

To this we add an electromagnetic field with Faraday tensor F^{ab} , from which we define the magnetic field in a spatial slice to be $B^i = \alpha(\star F^{0i})$ (where $\star F^{\mu\nu}$ is the Hodge dual of the Faraday tensor). Several quantities of interest are naturally expressed in terms of b^a , the magnetic field in a frame co-moving with the fluid:

$$b^a = (\star F^{ab})u_b. \quad (3.13)$$

We adopt the assumptions of ideal MHD; namely, that the fluid is perfectly conducting.

The stress-energy tensor of a magnetized perfect fluid is given by

$$T_{ab} = \rho h u_a u_b + P g_{ab} + F_{ac} F_b{}^c - \frac{1}{4} F^{cd} F_{cd} g_{ab}. \quad (3.14)$$

Additionally, we adopt the assumption of ideal MHD that the fluid is perfectly conducting:

$$F^{ab}u_b = 0 \quad (3.15)$$

(that is, the electric field vanishes in a frame co-moving with the fluid). This eliminates the electric field as an independent quantity and leaves eight degrees of freedom: five for the fluid and three for the magnetic field.

The state of the fluid at each gridpoint is represented in the code by the “primitive variables” ρ , T , u_i , and B^i , where T is a variable, related to the temperature, parameterizing the thermal pressure². Given ρ and T , the equation of state specifies the pressure $P(\rho, T)$ and specific internal energy $\epsilon(\rho, T)$ ³.

²The precise relationship of T to the temperature and thermal pressure is allowed to vary with the equation of state.

³More generally, we also allow these to depend on the composition of the fluid through the electron fraction Y_e , but the evolution of composition is not considered in this work.

In order to express the equations of their evolution in conservative form, we recompose them into the following set of “conservative” variables:

$$\rho_* = \sqrt{\gamma} W_L \rho \quad (3.16)$$

$$\tilde{\tau} = \sqrt{\gamma} \left(W_L \rho (W_L h - 1) - P + B^2 - \frac{1}{2} \frac{B^2 + (B^i u_i)^2}{W_L^2} \right) \quad (3.17)$$

$$\tilde{S}_i = \sqrt{\gamma} \left(W_L \rho h u_i + \frac{1}{W_L} (B^2 u_i - B^j u_j B^k \gamma_{ik}) \right) \quad (3.18)$$

$$\mathcal{B}^i = \sqrt{\gamma} B^i \quad (3.19)$$

(see also, e.g., Refs. [6, 29]). Here, γ is the determinant of the 3-metric, $W_L \equiv \alpha u^t$ is the Lorentz factor corresponding to the fluid’s velocity, and $B^2 \equiv B^i B^j \gamma_{ij}$. These “conservative” evolved variables map to the set of “primitive” variables through an inversion procedure described later.

The conservative variables are evolved according to:

$$\partial_t \rho_* + \partial_i (\rho_* v^i) = 0, \quad (3.20)$$

$$\partial_t \tilde{\tau} + \partial_i (\alpha^2 \sqrt{\gamma} T_{0i} - \rho_* v^i) = -\alpha \sqrt{\gamma} T^{\mu\nu} \nabla_\nu n_\mu, \quad (3.21)$$

$$\partial_t \tilde{S}_i + \partial_i (\alpha \sqrt{\gamma} T^j_i) = \frac{1}{2} \alpha \sqrt{\gamma} T^{\mu\nu} \partial_i g_{\mu\nu}, \quad (3.22)$$

where $v^i = u^i / u^t$ is the “transport velocity” of the fluid.

To compute the behavior of the magnetic field in the ideal MHD limit, we define an analog to the electric field,

$$\mathcal{E}_i \equiv -[ijk] v^j \mathcal{B}^k, \quad (3.23)$$

and then evolve the magnetic field according to

$$\partial_t \mathcal{B}^i = -[ijk] \partial_j \mathcal{E}_k, \quad (3.24)$$

where $[ijk]$ is +1 for an even permutation of the indices and −1 for an odd permutation.

This evolution is constrained by the zero-monopole criterion,

$$\nabla^{(3)} \cdot \mathbf{B} = \partial_i \mathcal{B}^i = 0 \quad (3.25)$$

(where $\nabla^{(3)}$ is the covariant derivative operator corresponding to the 3-metric γ_{ij}). In general, a numerical evolution scheme for the magnetic field will not preserve this constraint, so we adopt a constrained transport framework (first used by Yee [30] and later for GRMHD by Evans & Hawley [31]) to do so.

Our constrained transport implementation follows the prescription for *upwind constrained transport* proposed by Londrillo & Del Zanna [32] and described in detail by Del Zanna et al. as implemented in the ECHO code [33]. In particular, the longitudinal components of \mathcal{B}^i are evolved at cell faces. This presents a convenient definition of magnetic divergence at cell centers as the second-order divided difference of \mathcal{B}^i . The constrained transport algorithm guarantees that the time derivative of this quantity will be zero to machine precision. When the B -field itself is needed at cell centers, fourth-order polynomial interpolation is used, since discontinuities in the longitudinal direction are forbidden. Such interpolation is also used when metric quantities are needed at cell faces, as these fields are expected to be smooth.

In order to compute the fluxes of the evolution variables, non-smooth matter quantities must be reconstructed at cell faces and edges. Our code allows a choice of reconstructors, including a second-order monotonized centered (MC2) limiter and a fifth-order weighted essentially non-oscillatory (WENO5) scheme. The HLL approximate Riemann solver determines a single value for the flux on each interface. Flux derivatives are computed as second-order divided differences, making our scheme formally second-order accurate (that is, we do not perform the DER operation employed by the ECHO code). However, higher-order reconstructors, while not affecting the convergence rate, can greatly improve the accuracy of the code (see Sec. 4.4) at the expense of parallelization efficiency (their larger stencils require additional ghost zones).

In common with other high-resolution shock-capturing codes, `SpEC` requires procedures for inverting the relationship between primitive and conservative variables, along with a prescription for maintaining a tenuous atmosphere around the star. The addition of a

magnetic field necessitates changes to these algorithms, the details of which we describe below.

Primitive variable recovery in MHD

We generally follow the prescription of Noble et al. [34] for recovering primitive variables from the evolved conservative variables; i.e., the task of numerically inverting equations 3.16–3.19. We define

$$\tilde{S}^2 = \gamma^{ij} \tilde{S}_i \tilde{S}_j, \quad (3.26)$$

$$H = h(\rho, T) \rho W_L^2, \quad (3.27)$$

so that the relations between primitive and conservative variables can be written as

$$\tilde{S}^2 W_L^2 = \gamma (W_L^2 - 1) (B^2 + H)^2 - W_L^2 \frac{(\tilde{S}_i B^i)^2 (B^2 + 2H)}{H^2}, \quad (3.28)$$

$$-\frac{\rho_* W_L^2 + \tilde{\tau} W_L^2}{\sqrt{\gamma}} = \frac{B^2}{2} + W_L^2 \left(\frac{(\tilde{S}_i B^i)^2}{2\gamma H^2} - B^2 - H + P(\rho, T) \right) \quad (3.29)$$

(our formulation of this inversion procedure was originally described by Foucart [35]).

We solve these equations for (T, W_L^2) using the `gnewton` method as implemented by the GNU Scientific Library [36], subject to the constraint $W_L^2 \geq 1$. These equations are more challenging for the root-finding algorithm than the $B = 0$ case, especially in cases where the magnetic and/or kinetic energy of the fluid is large compared to its rest mass energy. When the 2D root-finder for (T, W_L^2) fails, we switch to a simple 1D bracketing algorithm solving for H (W_L is then considered as a known function of H). Once T and W_L are found, ρ and u_i can be computed analytically. The trivial computation of B^i from \mathcal{B}^i completes the primitive variable recovery.

Low density force-free primitive variable recovery⁴

Recovery of the full set of primitive variables can be difficult or impossible at low-density, magnetically-dominated gridpoints. Fortunately, it is also unnecessary when those points

⁴This section is primarily the work of Fatemeh Hossein Nouri.

are in the force-free regime. Our treatment of such points is similar to that of Etienne et al. [37]. For each gridpoint, the code first attempts to solve the full 2D system for (T, W_L^2) as described above. If a root cannot be found via either rootfinding procedure, we determine whether the failing gridpoint is in the force-free regime by checking the following conditions:

1. $\rho W_L/B^2 < 0.001$;
2. $B^2 > \sqrt{\tilde{S}^2/\gamma}$, which is necessary to have $B^2 > \mathcal{E}^2$;
3. $(\tilde{S}_j B^j)^2/(B^2 \rho_*^2) < 10$ to prevent very large velocities along field lines.

If the point satisfies these conditions, then we attempt a simpler 1D primitive variable recovery that ignores the internal energy of the gas (and if the point is not in the force-free regime, the code halts). First, we solve for the 4-velocity:

$$u_i = \frac{W_L}{B^2} \left(-\frac{\epsilon_{ijk}(\epsilon^{jlm} \tilde{S}_l B_m) B^k}{\sqrt{\gamma} B^2 + \rho_* h W_L} + \frac{(\tilde{S}_j B^j) B_i}{W_L \rho_* h} \right). \quad (3.30)$$

Assuming $T = 0$, $h = 1$, and using the normalization condition $W_L^2 = 1 + \gamma^{ij} u_i u_j$, we find

$$W_L^2 = 1 + \frac{W_L^2}{B^4} \left(\frac{\epsilon_{ijk}(\epsilon^{jlm} \tilde{S}_l B_m) B^k}{\sqrt{\gamma} B^2 + \rho_* W_L} \right)^2 + \frac{(\tilde{S}_j B^j)^2}{B^2 \rho_*^2}. \quad (3.31)$$

The velocity \mathbf{u} is composed of a parallel (to the magnetic field) part and a perpendicular part $W_L^2 = 1 + u_{\parallel}^2 + u_{\perp}^2$, so we have

$$u_{\perp}^2 = \frac{W_L^2}{B^4} \left(\frac{\epsilon_{ijk}(\epsilon^{jlm} \tilde{S}_l B_m) B^k}{\sqrt{\gamma} B^2 + \rho_* W_L} \right)^2, \quad (3.32)$$

$$u_{\parallel}^2 = \frac{(\tilde{S}_j B^j)^2}{B^2 \rho_*^2}. \quad (3.33)$$

Equation (3.31) is solved for W_L^2 with a 1D Newton–Raphson root solver; the other variables can be inferred from the solved W_L and the assumed $T = 0$. For force-free points with very low densities, or force-free points where we fail to solve Eq. (3.31), we remove the density-dependent terms in Eq. (3.31) and set u_i to the drift velocity $(u_{\perp})_i$. We note that the

$h = 1$ approximation used above would have to be adjusted when using a nuclear equation of state in which $h(\rho \rightarrow 0, T \rightarrow 0)$ is slightly less than one (i.e. when the binding energy of nucleons is taken into account and the specific internal energy of the fluid becomes negative as ρ approaches zero).

Special treatment of low-density matter

The methods we use to evolve relativistic fluids assume that the fluid density is strictly positive everywhere. To meet this requirement, we surround our stars with a low-density “atmosphere” (rather than vacuum) and require that the density always remain above this value. Additionally, numerical errors in low-density regions can take the solution into a regime where there exists no set of primitive variables corresponding to the evolved conservative variables. A robust code must employ techniques to handle these situations while minimizing unphysical effects. Our treatment of low-density regions when magnetic fields are present is described by Foucart [35] and Muhlberger et al. [1].

Because we evolve the spacetime metric on a separate grid, the fluid grid need not extend to the outer boundaries of our evolution domain. At the boundary of our fluid grid, a half-stencil’s worth of points are frozen at atmosphere levels. This “boundary condition” avoids the complexities of one-sided differencing and has no effect on the bulk evolution of the matter provided that the grid is large enough. When magnetic fields are present, they are initially confined to high-density regions far from grid boundaries, so this freezing procedure does not conflict with our constrained transport scheme. However, as the simulation progresses, outflows may develop. If left unchecked, these can ruin the accuracy of the simulation, as matter will permanently leave the fluid grid and magnetic divergences can be created at the boundaries. In these situations, we either extend the grid in the direction of the outflow or halt the simulation before these boundary effects have a chance to affect the behavior we are studying.

When the magnetic field is evolved via a vector potential rather than the field itself (see

Sec. 3.3), this boundary condition is no longer valid, even without magnetized outflows, unless one drives the potential to zero at the boundary of the grid. Robust solutions to this problem, including spectral evolution of the potential, are currently being considered. In practice, however, we find that potentials corresponding to initial poloidal fields in stars tend to remain zero away from the star for long enough to evolve several orbits of inspiral or to witness growth and suppression of dynamical instabilities such as those considered in Chapter 4.

3.3 Electromagnetic gauges

Instead of evolving the B -field directly, the scheme can be re-written in terms of the vector potential A (where $B = \nabla \times A$). The result is numerically equivalent, but requires the initial data to be expressed in terms of A instead of B . The advantage is that the B -field will remain divergence-free even if errors are introduced to A from outside of the evolution kernel. The most common source of such errors is interpolation—if the location of gridpoints changes during the evolution, e.g. during grid refinement, then the evolved fields will need to be interpolated to the new locations, introducing errors. Any divergence introduced to B will be preserved by the upwind constrained transport scheme. It guarantees that $\partial_t(\nabla \cdot B) = 0$, but does not drive $\nabla \cdot B$ itself to 0. Divergence cleaning techniques [38] can be used to provide such a driving force and to propagate the constraint violations off the grid, but we found that they act too slowly relative to the frequency of our grid refinement to adequately remove the divergence created by interpolation. If instead we interpolate A , the interpolation error does not manifest itself in the form of unphysical B -field divergence.

Evolving the vector potential requires a choice of gauge, encapsulated in the relationship between the vector potential A and a scalar potential Φ , which can be merged into a 4-vector potential $\mathcal{A}_\mu \equiv \Phi n_\mu + A_\mu$. In general, the evolution of A is governed by

$$\partial_t A_i = \epsilon_{ijk} v^j B^k - \partial_i (\alpha \Phi - \beta^j A_j). \quad (3.34)$$

The straightforward generalization of our upwind constrained transport scheme implies an “algebraic gauge” of

$$\alpha\Phi = \beta^i A_i . \quad (3.35)$$

Unfortunately, as shown by Etienne et al. [39], this gauge has a zero-speed characteristic mode. One effect of this that we observed is that the initial structure of the vector potential leaves an imprint at late times. While this has no effect within the evolution kernel (which is numerically identical to a direct evolution of B), once interpolated it results in large errors, as shown in Fig. 3.1. By adopting Lorenz gauge ($\nabla_a \mathcal{A}^a = 0$) and evolving the scalar potential Φ explicitly, the characteristic speeds are non-zero and the initial configuration does not leave an imprint. Farris et al. [40] also generalize the Lorenz gauge to damp the gauge modes, further reducing their negative effects on interpolation (though this is less essential for our code, which interpolates relatively rarely compared to the Illinois AMR code). Their generalized Lorenz gauge evolves Φ according to

$$\partial_t(\sqrt{\gamma}\Phi) + \partial_i(\alpha\sqrt{\gamma}\gamma^{ij}A_j - \sqrt{\gamma}\beta^i\Phi) = -\xi\alpha\sqrt{\gamma}\Phi , \quad (3.36)$$

or, more compactly,

$$\nabla_\mu \mathcal{A}^\mu = \xi n_\mu \mathcal{A}^\mu . \quad (3.37)$$

Here, ξ is the damping parameter; choices for ξ in BH–NS systems are discussed by Etienne et al. [41] and Paschalidis et al. [42].

When evolving the vector potential instead of the B -field directly, the natural location of A is for transverse components to be stored on cell edges. Given this, it is convenient to represent Φ at cell vertices; divided differences then represent the gradient operator on cell edges, which is required to evolve A according to Eq. 3.34. When evolving Φ in Lorenz gauge according to Eq. 3.36, Φ is reconstructed onto cell edges while A is first interpolated to cell centers (so that all three components are cosituated), then contracted with the metric and interpolated back to cell edges. The use of data at cell vertices requires an additional

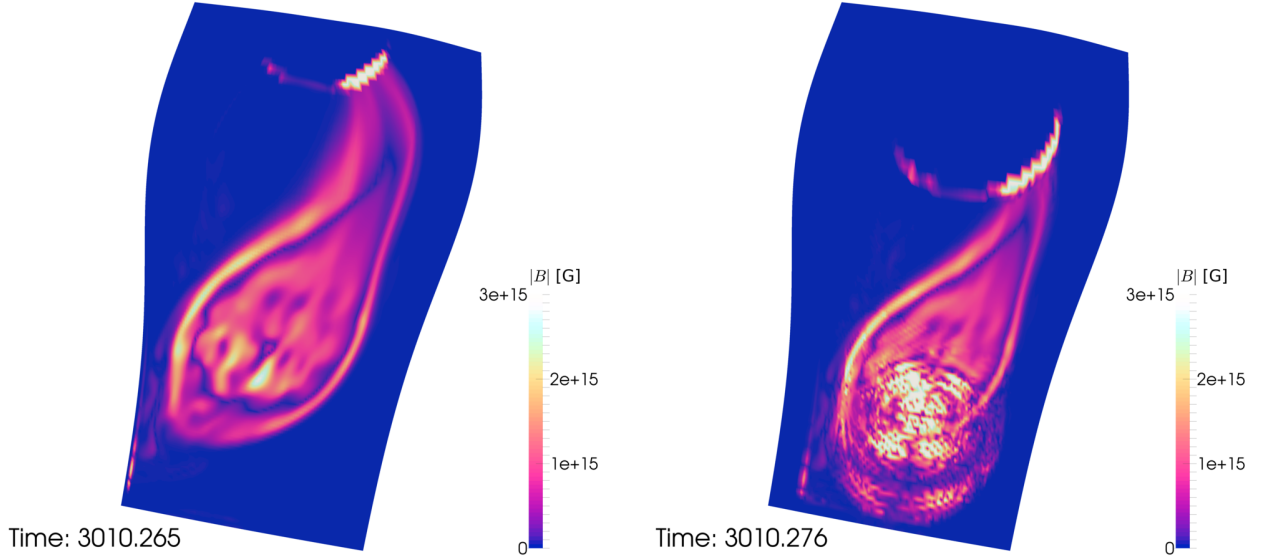


Figure 3.1: Interpolation artifacts caused by gauge imprint of initial data in the vector potential, as exhibited in the (physically-relevant) B -field. These images depict results from a BN–NS merger when the disrupted neutron star first starts accreting onto the black hole. The interpolation was performed as the result of a “regrid” operation, which expanded the grid to the left of the star. The circular ringing in the B -field corresponds to the original location of the neutron star in the corotating frame, prior to disruption. The boundary of the excised region, roughly corresponding to an apparent horizon of the black hole, is visible in the blocky buildup of magnetic field above the star.

level of staggered grid routines built around the same 1D interpolation and reconstruction kernels as the rest of our finite volume code.

3.4 Test problems

As with any new simulation code, we first confirm its ability to reproduce results from well-understood test problems before applying it to new astrophysical systems. The spacetime and hydrodynamics components of `SpEC` have been tested previously [26, 43]. Here, we check the performance of our new MHD module, using a similar test suite as Duez et al. [44]. In particular, we study its accuracy and convergence by comparing results to known analytical solutions exhibiting a range of non-trivial behaviors, including shocks and strong gravity.

One-dimensional relativistic tests⁵

To test the shock-capturing methods used in SpEC, we evolve a set of one-dimensional problems first proposed by Komissarov [45]. The initial data consist of two homogeneous states separated by a discontinuity at $x = 0$. The initial conditions for each test are listed in Table 3.1. We integrate the relativistic MHD equations from $t = 0$ to $t = t_{\text{final}}$ (also given in Table 3.1). The fluid follows a Γ -law equation of state with $\Gamma = 4/3$:

$$P = \rho^{4/3} + \rho T, \quad (3.38)$$

$$\epsilon = 3 \frac{P}{\rho}, \quad (3.39)$$

where we have now defined the code’s internal temperature variable T for the Γ -law case such that ρT is the thermal pressure of the fluid. To facilitate comparisons with previously published results, we use the same resolution as in Duez et al. [44], where the same tests were performed (see Figs. 7–8 and Table II of that work): our numerical domain covers the region $x = [-2, 2]$, and uses 400 grid points (higher resolution results are also provided to test the convergence of our code). The tests are performed with both the MC2 reconstructor used by Duez et al. and the WENO5 reconstructor that we prefer in most of our simulations. We use fourth-order Runge-Kutta time stepping, with a Courant factor of 0.5 ($dt = 0.005$ for $dx = 0.01$), except for the *fast shock* problem using WENO5 reconstruction, for which we use a Courant factor of 0.25 (the evolution is unstable for a Courant factor of 0.5, an issue which was also noted by Duez et al. when using the third-order piecewise parabolic method for reconstruction).

Fast and slow shocks

For these two tests, the shock front satisfies the relativistic Rankine-Hugoniot jump conditions [46]. The exact solution to the evolution of the fluid equation is known, with the shock propagating at constant speed while the fluid variables on each side of the shock

⁵The execution and analysis of these 1D tests was largely the work of Francois Foucart.

Table 3.1: Initial data for the shock tests.

Test	Initial state for $x < 0$	Initial state for $x > 0$
Fast shock ($t_{\text{final}} = 2.5$)	$\rho = 1, P = 1$ $u_i = (25, 0, 0), B^i = (20, 25.02, 0)$	$\rho = 25.48, P = 367.5$ $u_i = (1.091, 0.3923, 0), B^i = (20, 49, 0)$
Slow shock ($t_{\text{final}} = 2.0$)	$\rho = 1, P = 10$ $u_i = (1.53, 0, 0), B^i = (10, 18.28, 0)$	$\rho = 3.323, P = 55.36$ $u_i = (0.9571, -0.6822, 0), B^i = (10, 14.49, 0)$
Switch-off ($t_{\text{final}} = 1.0$)	$\rho = 0.1, P = 1$ $u_i = (-2, 0, 0), B^i = (2, 0, 0)$	$\rho = 0.562, P = 10$ $u_i = (-0.212, -0.590, 0), B^i = (2, 4.71, 0)$
Switch-on ($t_{\text{final}} = 2.0$)	$\rho = 0.00178, P = 0.1$ $u_i = (-0.765, -1.386, 0), B^i = (1, 1.022, 0)$	$\rho = 0.01, P = 1$ $u_i = (0, 0, 0), B^i = (1, 0, 0)$
Shock tube 1 ($t_{\text{final}} = 1.0$)	$\rho = 1, P = 1000$ $u_i = (0, 0, 0), B^i = (1, 0, 0)$	$\rho = 0.1, P = 1$ $u_i = (0, 0, 0), B^i = (1, 0, 0)$
Shock tube 2 ($t_{\text{final}} = 1.0$)	$\rho = 1, P = 30$ $u_i = (0, 0, 0), B^i = (0, 20, 0)$	$\rho = 0.1, P = 1$ $u_i = (0, 0, 0), B^i = (0, 0, 0)$
Collision ($t_{\text{final}} = 1.22$)	$\rho = 1, P = 1$ $u_i = (5, 0, 0), B^i = (10, 10, 0)$	$\rho = 1, P = 1$ $u_i = (-5, 0, 0), B^i = (10, -10, 0)$
Wave ($t_{\text{final}} = 2.5$)	$\rho = 1, P = 1$ $u_i = -0.4133 \cdot (0, \cos x, \sin x), B^i = (1, \cos x, \sin x)$	$\rho = 1, P = 1$ $u_i = -0.4133 \cdot (0, \cos x, \sin x), B^i = (1, \cos x, \sin x)$

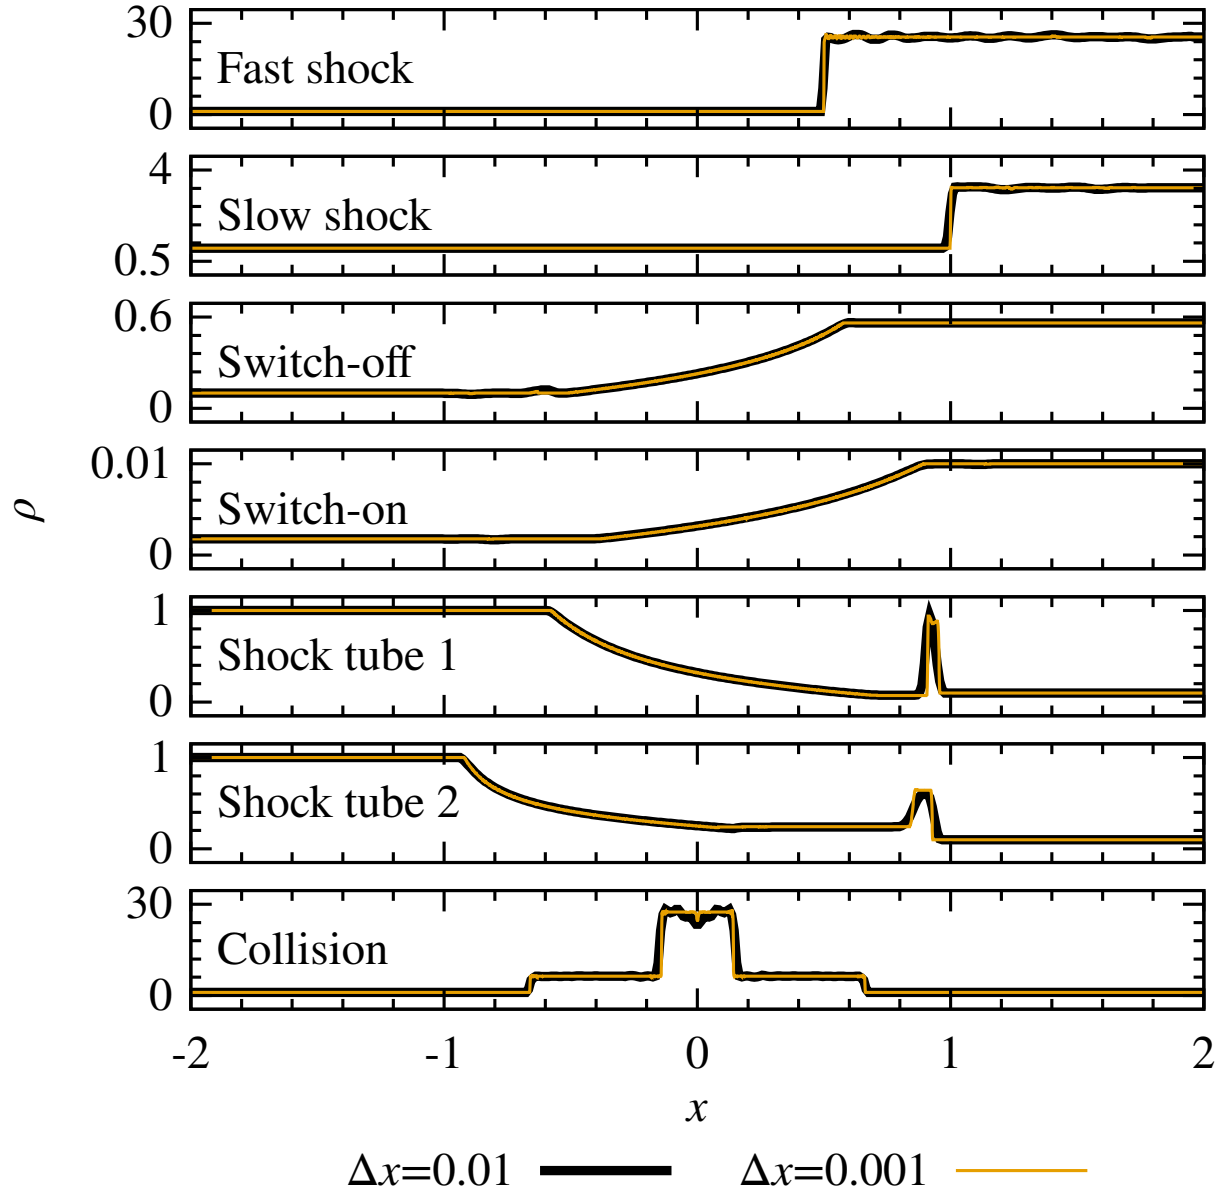


Figure 3.2: Rest-mass density at $t = t_{\text{final}}$ for the shock tests described in Table 3.1, shown for two resolutions ($N = 400$ and $N = 4000$ points).




Figure 3.3: Velocity at $t = t_{\text{final}}$ for the shock tests described in Table 3.1, shown for two resolutions ($N = 400$ and $N = 4000$ points).




Figure 3.4: Error in the final value of u^y for the “wave” test at 3 resolutions ($N = 50$, $N = 100$, $N = 200$), rescaled for the expected second-order convergence.

remain constant [45, 47]. The *fast shock* test is the hardest test for our code: it evolves a strong shock, with the shock front moving relatively slowly on the grid ($0.2c$) but the fluid being highly relativistic (Lorentz factor $W_L = 25.02$). As already noted, it is the only test that is unstable when using a Courant factor of 0.5 (for WENO5 reconstruction). It is also fairly sensitive to the choice of variables that are interpolated from cell centers to cell faces when computing the fluxes entering the conservative hydrodynamics equations: if we interpolate the transport velocity v^i , the shock evolves as expected, while if we interpolate the spatial components of the 4-velocity u_i the shock immediately stalls. Considering that in practice, in 3-dimensional evolutions of neutron stars or binary mergers, we do not reliably evolve fluid elements with $W_L \sim 25$ (the occurrence of such high Lorentz factors is

prevented by the corrections applied to the velocity and temperature of low-density points in the atmosphere), this difference is unimportant in practice. The *fast shock* test is mostly evolved in order to verify that our implementation of the MHD equations is correct in the limit of ultra-relativistic fluids. In fact, because of the practical advantages of using u_i instead of v^i , we usually reconstruct the former ($W_L = \sqrt{1 + g^{ij}u_iu_j}$ is always well-defined while $W_L = 1/\sqrt{1 - g_{ij}v^iv^j}$ is not if numerical errors in the low-density regions cause v^i to satisfy $g_{ij}v^iv^j > 1$). In Figs. 3.2 and 3.3, we show the result of that test when using the MC2 reconstruction method (and reconstructing v^i), for 400 and 4000 grid points. The results converge towards the solution at the expected first-order rate. The *slow shock* test is generally less extreme. As in previous studies [44, 45, 48], we observe that the evolution is very accurate on the left side of the shock, while oscillations are visible on the right side of the shock (see Fig. 3.2). Although these oscillations converge away as we increase the resolution, they do so more slowly than expected past 200-400 points in the evolution domain (convergence order of ~ 0.6). This is the only test for which we do not observe at least first-order convergence.

Other shock tests

The five other one-dimensional shock tests, for which results are presented in Figs. 3.2 and 3.3, are comparable to previously published results in accuracy (for the simulations using 400 points), and convergent when the resolution is increased to 4000 points. As expected, the convergence is fairly slow (first order), which explains why sharp features remain visible even at high resolution. These tests cover a wide range of potential behaviors (shock waves, rarefaction waves, contact discontinuities), and indicate that the shock capturing methods implemented in SpEC are capable of handling the discontinuities which are likely to arise in our simulations.

Wave

The last one-dimensional test to which we submit our code is the propagation of a wave on a periodic grid. In this case, all variables are continuous, and the error in the simulations should be second-order convergent. In the exact solution, the initial profile (given in Table 3.1) simply propagates with velocity $v = 0.3820$. The error in the density ρ at the end of the simulation for 3 different resolutions (50, 100 and 200 points per wavelength) is shown in Fig. 3.4, rescaled for the assumed second-order convergence. Our results also appear in good agreement with the theoretical predictions for this smooth configuration.

Bondi accretion

We also test the ability of our code to evolve a magnetized fluid in full 3D in the presence of strong gravitational fields. We check its ability to maintain stationary and spherically symmetric accretion onto a Schwarzschild black hole according to the relativistic Bondi accretion solution, but with the addition of a radial B -field. This provides a nontrivial test of the GR terms in our MHD evolution scheme while still possessing an exact solution to which we can compare our numerical results.

We write the metric in the Kerr–Schild coordinates; as a result, all of the variables are well-behaved at the horizon (horizon penetrating). For this test, we fix the metric and evolve the fluid equations only (this is often referred to as the Cowling approximation).

We evolve the same configuration used by Duez et al. [44]. The accretion rate is $\dot{M} = 1$, the sonic radius is at $r = 8M$ (where M is the mass of the black hole), and the equation of state obeys a $\Gamma = 4/3$ power law [see Eqs. (3.38)–(3.39)]. We freeze the hydro evolution variables at the inner and outer boundaries. We set the inner boundary radius outside of the horizon at $r = 2.8M$ (the horizon is at $r = 2M$), and the outer boundary is placed at $r = 9M$; the Cartesian grid extends $\pm 10M$ along each axis.

We simulate this accretion flow at three different resolutions, with grid extents of 64^3 ,

96^3 and 128^3 . The initial magnetic field is radial, with strength $b^2/\rho = 1$, resulting in a stationary solution (while the black hole is effectively a large magnetic monopole, the divergence of the B -field is zero everywhere on the grid). Reconstruction is performed using a fifth-order WENO kernel. We also add Kreiss-Oliger dissipation [49] to the evolution of all conservative variables. This removes short-wavelength noise that would otherwise interfere with clean convergence for this system.

To measure the convergence of our code, we compute the volume L_2 norm of the deviation of the conservative variables from their exact Bondi solutions:

$$\delta u = \left(\frac{\int |u - u_{\text{exact}}|^2 \sqrt{\gamma} d^3x}{\int \sqrt{\gamma} d^3x} \right)^{1/2}. \quad (3.40)$$

In Fig. 3.5 we plot the error norm measured by Eq. (3.40) for all conservative variables after $100M$ of evolution for three different resolutions. These show that our results are converging at second order, as expected (and also as observed in previous studies of this problem, e.g. Refs. [50, 51]).




Figure 3.5: Error norm for the Bondi test at three resolutions, rescaled for second-order convergence.

4 SHEAR INSTABILITIES IN DIFFERENTIALLY ROTATING NEUTRON STARS

4.1 Motivation

Stellar core collapse, accretion-induced white dwarf collapse, and binary neutron star merger all naturally produce rapidly spinning neutron stars with strong differential rotation. The resulting neutron stars could be subject to well-known dynamical instabilities, and the resulting stellar deformations could produce a strong gravitational wave signal which, if detected, would provide invaluable information on these violent phenomena.

Global $m = 2$ instabilities (perturbations with an azimuthal dependence of $e^{im\phi}$) are particularly relevant for gravitational wave production. One source of such modes is the dynamical bar mode instability. However, this instability only sets in for extremely high values of the ratio of the rotational kinetic energy T to the gravitational potential energy W : $T/|W| \geq 0.27$ (with small variations depending on the equation of state and ratio of mass to radius [52–55]). Simulations have revealed another dynamical nonaxisymmetric instability that can appear at much lower $T/|W|$ if sufficient differential rotation is present [56–67]. Watts, Andersson, and Jones [68] have given compelling arguments for identifying this “low- $T/|W|$ instability,” as it was called, as a form of corotation shear instability, similar in basic principle to the better-known Papaloizou-Pringle instability in thick accretion disks [69]. Namely, nonaxisymmetric modes trapped in a resonant cavity make multiple passes across a corotation radius (the radius where the mode pattern speed matches the local fluid angular speed) and are amplified on each pass. A local minimum of the radial vortensity profile has been suggested as the mechanism for mode trapping [61]. Simulations of protoneutron stars indicate that realistic core collapse scenarios can produce stars subject to this instability [70]. Indeed, the gravitational waves from this instability

have been proposed as a distinctive signal from hypothesized magnetorotationally-driven galactic supernovae with rapidly rotating cores [71].

Magnetohydrodynamic simulations have shown that the dynamical bar mode instability can be suppressed by magnetic forces, although only for unrealistically high magnetic field strengths [72, 73]. Fu & Lai have investigated the effect of a toroidal magnetic field on the low- $T/|W|$ instability using an analytic model, treating the star as an infinite cylinder with no vertical structure [74]. Because of the strong differential rotation, a more modest poloidal seed field ($\sim 10^{14}$ G) could wind up to a sufficiently strong toroidal field ($\sim 10^{16}$ G) within the growth time of the instability (around 30 ms). The protoneutron stars most likely subject to the low- $T/|W|$ instability have strong differential rotation and potential for magnetorotational dynamo action, and in such stars magnetic fields of this magnitude are plausible [75]. Magnetic suppression could therefore eliminate the potential gravitational wave signal of core-collapse supernovae. However, Fu & Lai’s model makes a number of strong simplifying assumptions: cylindrical stars, a polytropic equation of state, and purely toroidal fields. These could lead to the neglect of other important magnetohydrodynamical effects and instabilities. Thus, simulations of more realistic configurations in full 3D are needed to evaluate the robustness of the suppression mechanism.

Here we simulate the effects of magnetic fields on differentially-rotating neutron stars susceptible to the low- $T/|W|$ instability using our new magnetohydrodynamics module for `SpEC`. The instability is indeed suppressed for a narrow range of strong seed magnetic fields, but the more commonly observed behavior is for either magnetic fields to be too weak to affect the global quadrupole mode or for them to be sufficiently strong for magnetic instabilities to set in and actually amplify the mode. In general, we find gravitational waves comparable in magnitude to the unmagnetized case.

4.2 Setup

Physical system

Since our purpose is to study the effect of magnetic field strength and configuration on the low- $T/|W|$ instability, we focus here on one system that, in the unmagnetized case, is subject to this instability. We choose one of the differentially rotating neutron star models studied by Corvino et al. [66], namely their configuration M. 1. 200, which they indeed find to be unstable. The star has a baryon mass of $M_b = 2.44 M_\odot$, a central density of $\rho_c = 1.16 \times 10^{-3} M_\odot^{-2}$, and a ratio of kinetic to gravitational potential energy of $T/|W| = 0.2$ (low enough to avoid the high- $T/|W|$ dynamical bar mode instability, which becomes accessible for $T/|W| \gtrsim 0.24$ [54, 76]). The degenerate component of the equation of state is given by the SLy model [77], which we implement via the fitting formula introduced by Shibata et al. [78]. Thermal contributions to the pressure and internal energy are included by a simple Γ -law addition to the equation of state (see Shibata et al., Duez et al. [26]), where we have chosen $\Gamma_{\text{th}} = 2$. At the start of simulations, the temperature of the star is set to zero. Thus, we ignore for the purposes of this study the significant thermal energy that would be found in a realistic protoneutron star or binary post-merger remnant scenario, but we do model the dominant cold nuclear physics component of the equation of state.

For the initial state of the star, we create an axisymmetric nonmagnetized equilibrium solution of the Einstein equations. The spacetime metric, set in quasi-isotropic coordinates, takes the form

$$ds^2 = -e^{\mu+\nu} dt^2 + e^{\mu-\nu} r^2 \sin^2(\theta) (d\phi - \omega dt)^2 + e^{2\xi} (dr^2 + r^2 d\theta^2), \quad (4.1)$$

where μ , ν , ω , and ξ are arbitrary functions of axisymmetric space. Differential rotation is a key requirement for the instability and is incorporated by setting the initial angular velocity, $\Omega \equiv v^\phi$, according to

$$\Omega_c - \Omega = \hat{A}^{-2} u^t u_\phi = \frac{1}{\hat{A}^2 R_e^2} \left[\frac{(\Omega - \omega) r^2 \sin^2(\theta) e^{-2\nu}}{1 - (\Omega - \omega)^2 r^2 \sin^2(\theta) e^{-2\nu}} \right], \quad (4.2)$$

where R_e is the coordinate equatorial radius, Ω_c is the central angular velocity, and \hat{A} is a dimensionless parameter characterizing the strength of differential rotation. For the initial state of the system under study, $R_e = 7.8 M_\odot$, $\Omega_c = 2\pi \times 3.0 \text{ kHz}$, and $\hat{A} = 1$. The ratio of polar to equatorial coordinate radii is $R_p/R_e = 0.414$. We compute the equilibrium configuration using the code of Cook, Shapiro, and Teukolsky [79].

Since the equilibrium data are axisymmetric to numerical precision, we seed the star with a small $m = 2$ perturbation in order to make the initial perturbation resolution-independent and its subsequent growth numerically convergent. This perturbation is applied to the rest-mass density and takes the form

$$\rho \rightarrow \rho \left(1 + \delta_2 \frac{x^2 - y^2}{R_e^2} \right). \quad (4.3)$$

The size of the initial perturbation is $\delta_2 = 2 \times 10^{-5}$. This yields an initial distortion [see Eq. (4.11)] of $\eta_+ = 4.08 \times 10^{-6}$.

The properties of the star in its initial state are summarized in Table 4.1. While the mass is considerably higher than would be expected for a protoneutron star (though not implausible for a binary neutron star merger remnant), we expect our conclusions regarding the interaction of magnetic fields and the low- $T/|W|$ instability to apply qualitatively to lower-mass systems. Several properties differ slightly from those of Corvino et al.'s **M. 1. 200**, so while we expect the overall evolution to be quite similar, we should not expect perfect correspondence in quantitative measurements.

Finally, we introduce a seed poloidal magnetic field. Following a standard practice in the numerical literature (e.g., [80–82]), we introduce a toroidal vector potential with strength

$$A_\phi = A_b \omega^2 \max(P - P_{\text{cut}}, 0)^{n_s}, \quad (4.4)$$

where A_b sets the overall strength of the resulting B -field, n_s controls the smoothness of the field, and the cutoff pressure P_{cut} (set to 4% of the central pressure) confines the initial field to regions of high-density matter. The vector potential is evaluated at cell

Table 4.1: Basic properties of the neutron star. R_e is the equatorial coordinate radius, and R_p is the polar coordinate radius. $\Delta\Omega$ is the angular frequency range—the difference between the central and equatorial rotation frequencies.

	$G, c, M_\odot = 1 \quad \text{cgs}$	
M_0	2.44	$4.85 \times 10^{33} \text{ g}$
M_{ADM}	2.19	$4.35 \times 10^{33} \text{ g}$
R_p/R_e	0.414	0.414
ρ_c	0.00116	$0.717 \times 10^{15} \text{ g cm}^{-3}$
Ω_c	0.0922	$2.98 \times 2\pi \text{ kHz}$
$\Delta\Omega$	0.0650	$2.10 \times 2\pi \text{ kHz}$

edges, with a fourth-order curl operator producing the initial B -field at cell faces. This field is then superimposed on top of the unmagnetized equilibrium solution. While not formally self-consistent, at the field strengths we consider we expect both the deviation from equilibrium and the constraint violations in the equations of general relativity to have negligible effects on our conclusions. Specifically, the norm of the generalized harmonic constraint energy increased by $< 1\%$ with the addition of the magnetic field. Selected field lines for the initial and evolved states of the star are illustrated in Fig. 4.1.

We explored a region of the two-parameter space $A_b \times n_s$. However, it is more intuitive to talk about magnetic field strengths measured in Gauss than the poloidal coefficient A_b . The magnetic configurations studied are summarized in Table 4.2, which reports both the maximum strength of the B -field at $t = 0$ as well as a representative initial field strength B_0 that more closely reflects the average field in the star. We assign this representative strength to each magnetic field configuration by measuring the early growth of the magnetic energy within the star, hereafter labeled H_B [see Eq. (4.9)], and fitting to it the formula

$$H_B \approx B_0^2 \left(\frac{\Delta\Omega^2 R^3}{6} \right) t^2 \quad (4.5)$$

to solve for B_0 . Here we take $\Delta\Omega = 2.1 \times 2\pi \text{ kHz}$ and $R = 15.3 \text{ km}$ (the proper equatorial radius, as opposed to the isotropic coordinate radius reported earlier). This formula was also used by Fu & Lai in their analysis [74], easing comparisons with that work.

The dynamical importance of the magnetic field can be inferred from the ratio of the




Figure 4.1: Illustrations of magnetic field lines at early ($t = 0$, above) and intermediate ($t = 2160$, below) times. Contours represent regions of similar rest-mass density. Magnetic field lines are seeded at coordinate radii of $2 M_{\odot}$ (yellow) and $4 M_{\odot}$ (pink).

Table 4.2: Summary of the magnetic configurations studied. B_{\max} is the maximum strength of the initial poloidal magnetic field, B_0 is its “representative” strength as defined in the text, and β_{\min} is the minimum ratio of fluid pressure to magnetic pressure found initially in the interior of the star.

$A_b [G, c, M_{\odot} = 1]$	n_s	B_{\max}/G	B_0/G	β_{\min}
0	n/a	0	0	∞
0.00768	1	2.5×10^{14}	4×10^{13}	1.1×10^6
0.0379	1	1.3×10^{15}	2×10^{14}	5.2×10^4
0.0892	1	2.9×10^{15}	5×10^{14}	9.5×10^3
0.444	1	1.5×10^{16}	2×10^{15}	3.8×10^2
424	2	1.8×10^{15}	2×10^{14}	5.9×10^5
1000	2	4.1×10^{15}	5×10^{14}	1.1×10^5

gas to magnetic pressure $\beta = 2P/b^2$. For our strongest initial field, β starts no lower than 3.8×10^2 .

Simulation parameters

We used several evolution grids over the course of this investigation, but our final results were achieved on a “reference” finite volume grid with $\Delta x = \Delta y = 0.17 M_{\odot} = 250$ m and $\Delta z = 0.10 M_{\odot} = 150$ m. Grids employed during the exploratory phase (discussed in Sec. 4.4) used uniform resolution and are detailed where mentioned.

Our spectral grid (for evolving the spacetime; see Fig. 4.2) consists of a filled sphere (using a basis of three-dimensional generalizations of Zernike polynomials; see Sec. A.2) surrounded by layers of “cubed spheres”—products of Chebyshev polynomials distorted to conform to 1/6 of a spherical shell. These encompass the entire finite volume grid and are in turn surrounded by true spherical shells (a product of Chebyshev polynomials and spherical harmonics) extending to 300 stellar equatorial radii. The spectral resolution of our reference grid corresponds to spherical harmonics out to $l = 21$ for the central sphere and $l = 17$ for the outer spheres. The radial dimensions of these spheres are resolved by 12 and 11 collocation points, respectively. The cubed spheres contain 12 radial points and 20 transverse points.




Figure 4.2: Illustration of x - z slice of domain decomposition. The shaded region with a bold outline represents the initial star. The dashed rectangle represents the finite-difference domain, which has a coordinate width of $25 M_{\odot}$ and a coordinate height of $14.5 M_{\odot}$. For spectral subdomains, the actual reference grid has twice as many collocation points in each direction as are shown in the figure.

4.3 Analysis

To study the low- $T/|W|$ instability in our simulations and the effects that magnetic fields have on it, we consider several global measures of the simulation results as functions of time. These include various energy integrals, defined as follows:

Rest mass:

$$M_b = \int \rho W_L \sqrt{\gamma} d^3x . \quad (4.6)$$

Kinetic energy:

$$T = \frac{1}{2} \int \rho h W_L u_i v^i \sqrt{\gamma} d^3x, \quad (4.7)$$

where $v^i \equiv u^i/u^0$.

Internal energy:

$$U = \int \rho W_L \epsilon \sqrt{\gamma} d^3x. \quad (4.8)$$

Magnetic energy:

$$H_B = \frac{1}{2} \int b^2 W_L \sqrt{\gamma} d^3x. \quad (4.9)$$

Since total energy is conserved (and our hydrodynamic evolution is conservative), we can infer the change in gravitational energy from the sum of the changes in these non-vacuum energies. Some of this is lost in the form of gravitational waves, which emit $2.1 \times 10^{-4} M_\odot$ of energy over the duration of the simulation in the unmagnetized case. Any remaining difference must therefore be a change in the gravitational binding energy of the star.

Following previous studies, we consider the quadrupole moment of the rest mass density about the origin (which is the initial center-of-mass):

$$I^{ij} = \int \rho W_L x^i x^j \sqrt{\gamma} d^3x. \quad (4.10)$$

To reduce this to a scalar measure, we consider two polarizations of the x and y components of the quadrupole tensor,

$$\eta_+(t) \equiv \frac{I^{xx}(t) - I^{yy}(t)}{I^{xx}(0) + I^{yy}(0)} \quad (4.11)$$

$$\eta_\times(t) \equiv \frac{2I^{xy}(t)}{I^{xx}(0) + I^{yy}(0)}, \quad (4.12)$$

and, following Corvino et al. [66], take their magnitude to define the “distortion parameter” η :

$$|\eta(t)| = \sqrt{\eta_+^2(t) + \eta_\times^2(t)}. \quad (4.13)$$

Note that the numerical atmosphere surrounding the star (see Sec. 3.2) has the potential to bias integral measurements like those above. A common solution is to impose density

or radius thresholds when summing the integrand. However, because our fluid grid only covers the region immediately around the star and does not extend into the wave zone, the effect of the atmosphere on these measurements is negligible.

The invariant strength of the magnetic field is simply the magnitude of b^a , whose square is equal to

$$b^2 = \frac{B^2}{W_L^2} + \left[B^i \left(\frac{u^j}{W_L} + \frac{\beta^j}{\alpha} \right) \gamma_{ij} \right]^2. \quad (4.14)$$

To report physical results, we convert this strength to CGS-Gaussian units via

$$\begin{aligned} |B_{\text{CGS}}| &= \frac{\sqrt{4\pi} b^2}{1 \text{ M}_\odot} \left(\frac{c^2}{GM_\odot} \right) \left(\frac{c}{\sqrt{4\pi\epsilon_0 G}} \right) \times 10^4 \text{ G} \\ &= \sqrt{b^2} \times 8.352 \times 10^{19} \text{ G}. \end{aligned} \quad (4.15)$$

We also consider the evolution of some quantities in a Lagrangian frame of reference. To do this, we seed “tracer” particles in the fluid and evolve their positions according to the fluid velocity in our Eulerian evolution frame. The resulting trajectories provide useful information in their own right, and observing quantities along those trajectories allows for their Lagrangian analysis.

Finally, in order to accurately monitor the growth of instabilities of arbitrary m in a robust manner, we consider an additional measure of non-axisymmetry that differs from diagnostics used in previous investigations. Our approach is discussed below.

Azimuthal modes

Previous studies have analyzed the “Fourier power” of m -modes of a field ψ by integrating the quantity $\psi e^{im\phi}$. Some have performed this integral over a ring, capturing the power at a single radius and height within the system [61, 70]. Others, including Corvino et al., have performed a volume integral. While the latter approach incorporates contributions from the entire system, it has several disadvantages. The integrand is in general discontinuous at the origin for $m > 0$, and thus naive numerical computations of $|P_m|$ can produce spurious results (for example, computing a finite volume integral with a gridpoint at the origin

will result in non-zero $m > 0$ power for axisymmetric data). Additionally, m -modes of ψ whose phase changes with radius or height will be biased (for instance, a tightly wound spiral structure will produce canceling contributions to the integral for each infinitesimal annulus). Diagnostics defined in terms of multipole moments, like η , do not suffer the discontinuity problem, but radial cancellations still cause, for instance, the quadrupole moment to be a potentially poor representation for what one would intuitively call “ $m = 2$ power.”

A hybrid approach is to sum the power of ψ in several rings, thus sampling the field at multiple heights and radii. More generally, ψ can be multiplied by a set of orthogonal window functions isolating particular subsets of the domain, with volume integrals used to compute the power of each product. These functions would approach the origin as ω^m , ensuring smoothness there, and would be localized at various radii, avoiding cancellation from spiral structure. A natural choice for such a set of functions are the radial and vertical cardinal functions associated with a basis for functions over a cylinder (for example, the product of Zernike polynomials over a disk with Legendre polynomials in z). These functions are smooth, orthogonal, and generally localized around their corresponding node.

In fact, this approach is equivalent to a spectral measure of m -power, defined in Eq. A.29, where the Fourier components of ψ are decomposed into a set of basis functions, and the squared magnitude of the spectral coefficients are summed (see Sec. A.5 for proof). It is this definition of m -power, which we denote with $P_m[\psi]$, that we employ in our analysis. To account for possible center-of-mass motion, the origin is chosen to follow the measured center-of-mass ($\int \mathbf{x} \rho W_L \sqrt{\gamma} d^3 \mathbf{x} / \int \rho W_L \sqrt{\gamma} d^3 \mathbf{x}$) of the system.

4.4 Results

Having established the accuracy and convergence of our code on standard test problems (see Sec. 3.4), we can now compare our findings regarding the unmagnetized low- $T/|W|$ instability with previous simulations of the same system, confirming the baseline against which magnetized results will be compared.

Unmagnetized instability

When simulating the unmagnetized system, we find the behavior of the low- $T/|W|$ instability to depend sensitively on the reconstruction algorithm employed by the code (see Sec. 3.2 for the role and implementation of reconstruction in our evolution scheme). In particular, the growth of the distortion parameter $|\eta|$ was not convergent with resolution for the majority of reconstructors considered (a more thorough investigation is the subject of ongoing work). We are, however, able to obtain consistent results using WENO5 reconstruction, as shown in Fig. 4.3.

Even when using WENO5 reconstruction, insufficient resolution, particularly in the vertical direction, can introduce spurious features in the distortion parameter's evolution at intermediate times and otherwise increase the simulation's sensitivity to other choices in numerical methods. We see long-term consistency in the growth of η when $\Delta z \lesssim 0.1 M_\odot$.

We follow the unmagnetized system through the saturation and initial decay of the instability, as shown in Fig. 4.4. The growth exhibits a single exponential mode with a time constant of $\tau \approx 3.6$ ms, and the amplitude of the instability saturates when the distortion parameter reaches $|\eta|_{\max} \approx 0.035$. This is the reference against which our magnetized results will be measured. As noted by Cerdá-Durán et al., this is less than the $O(1)$ (fully nonlinear) magnitude usually achieved by the classical bar-mode instability [63]. They attribute this lower saturation value to the accumulation of angular momentum in the outer layers of the star, beyond the corotation radius, which excites Kelvin–Helmholtz-like

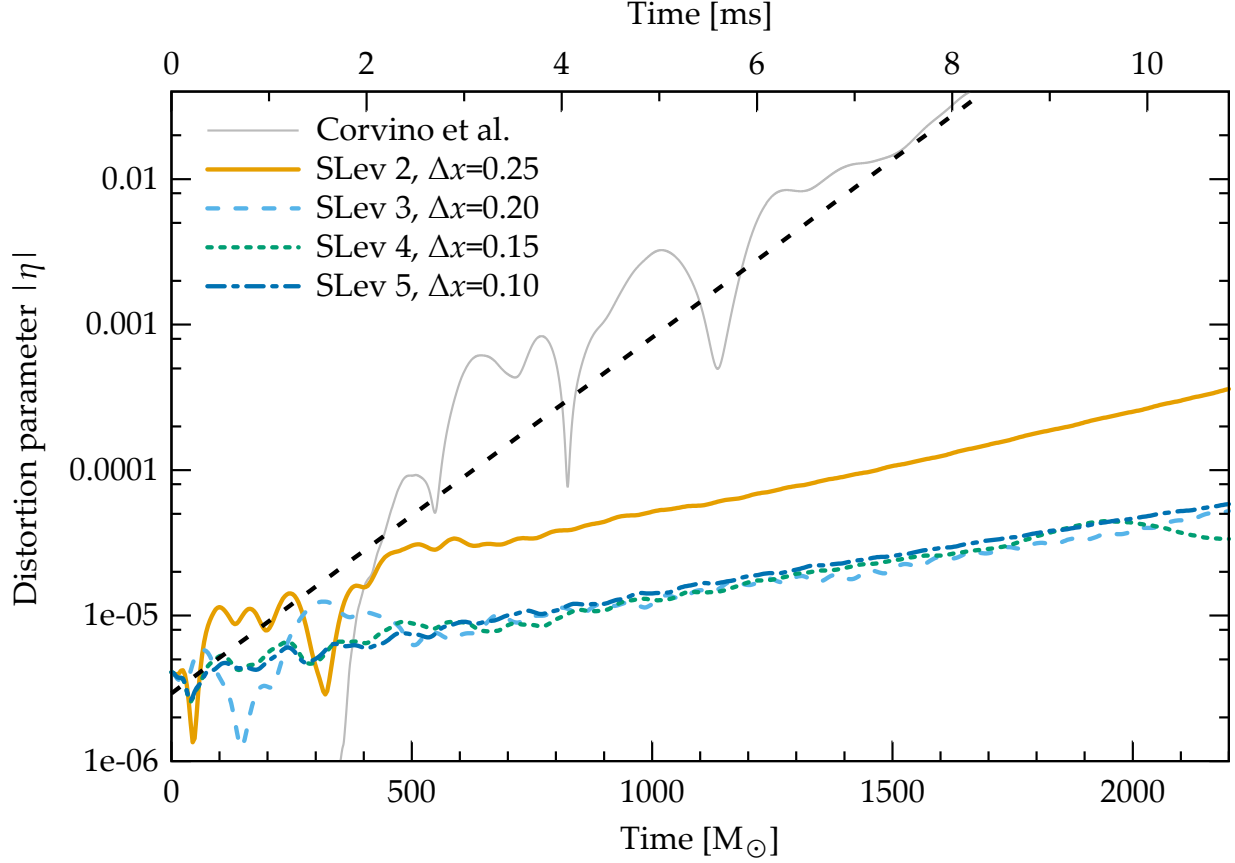


Figure 4.3: Consistency of the net growth rate of the low- $T/|W|$ instability when using WENO5 reconstruction at various resolutions (no magnetic field is present). The black dashed line represents our approximation to the growth rate found by Corvino et al. for $M.1.200$. Results from resolutions of $\Delta x \lesssim 0.2 M_{\odot}$, while not formally convergent, are in good agreement and are clearly distinct from those of Corvino et al. “SLev” indicates the spectral resolution level, with higher levels corresponding to finer resolution (the “reference” grid uses SLev 4), and grid spacings are measured in solar masses.

instabilities that prevent further angular momentum extraction. Our data is consistent with their observations, with late-time $m = 2$ density perturbations most pronounced at large radii.

Comparing to the results of Corvino et al. [66], who used the piecewise parabolic method (PPM) for reconstruction, we find a large disagreement in the growth rate of η . Our simulations exhibited clean exponential growth for over 30 ms with a characteristic time of $\tau \approx 3.6$ ms. For comparison, from Fig. 3 in Corvino et al.’s work we estimate a

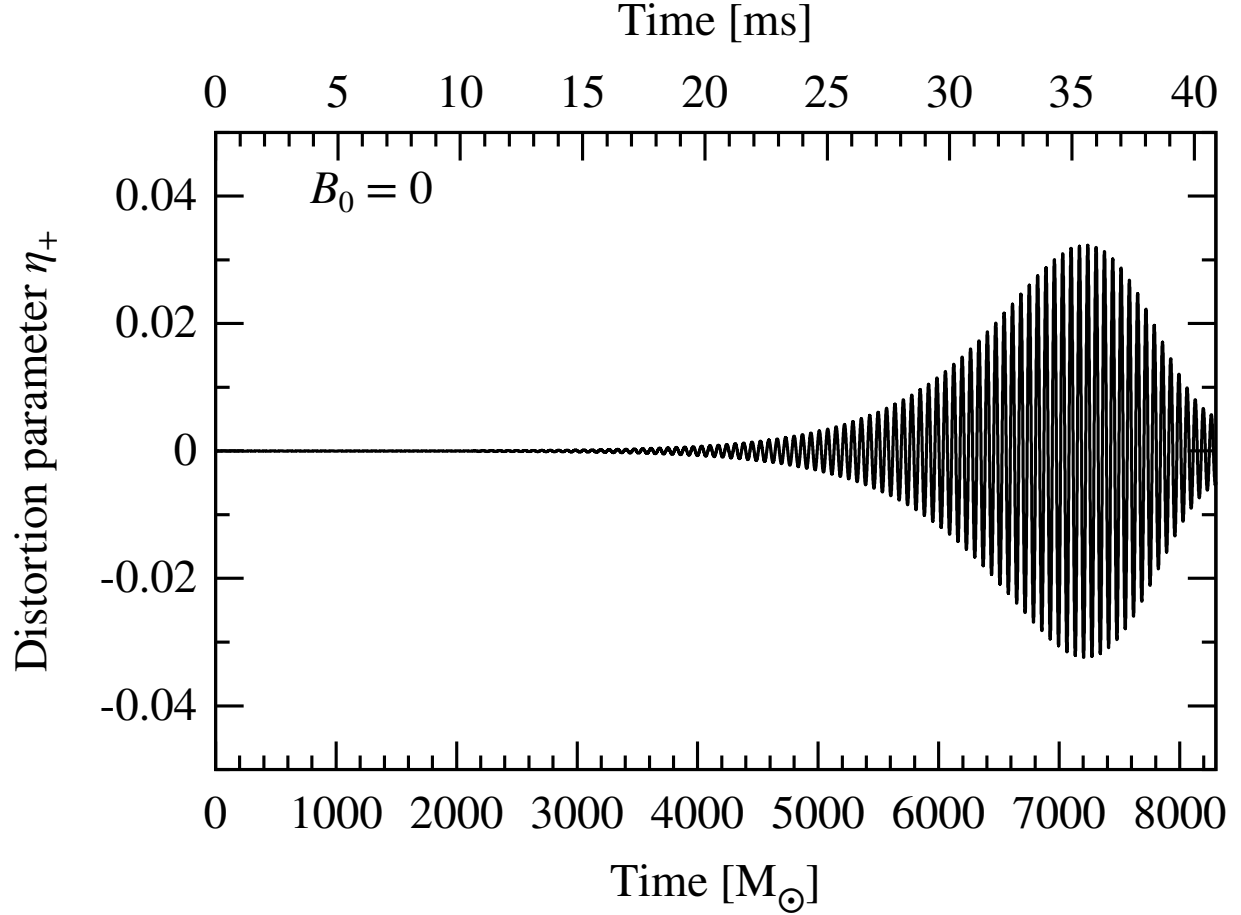


Figure 4.4: Growth and saturation of the unmagnetized low- $T/|W|$ instability as expressed in the “plus” polarization of the distortion parameter η . The “cross” polarization exhibits the same behavior with a phase shift. Compare to Corvino et al. Fig. 3.

growth time of $\tau \approx 0.88 \text{ ms}^1$. This rate is illustrated by the dashed line in Fig. 4.3 and results in saturation of the instability considerably sooner than in our simulations. Saturation amplitudes, however, agree to within a factor of two (0.035 vs. 0.055). Overall, the growth profile we observe for η is much more similar to those Corvino et al. report for stars with even lower values of $T/|W|$ (0.15 and 0.16), showing smooth exponential growth followed by decay, than what they report for $T/|W| = 0.2$, where the growth is comprised of multiple unstable modes growing at different rates. Figure 4.5 shows the spectrum of the distortion

¹This value encapsulates the speed with which the distortion of the star grows to saturation, but it does not necessarily represent the growth rate of the same unstable mode we see in our simulations. In particular, Corvino et al. see a spectrum of unstable modes with different growth rates.

parameter for their case compared to ours, highlighting the fact that our codes disagree on the stability and/or excitation of these modes.

While the reason for the disagreement is currently unknown (and will be the subject of future investigation), it appears to be linked to the reconstruction algorithm employed by the simulation. We have preliminary results from evolving this same system with the Zelmani code [83]², some of which are shown in Fig. 4.6. While the growth profile found by Zelmani when using PPM roughly matches that of Corvino et al., additional tests suggest it may not be convergent with resolution, and the Zelmani results using different reconstructors yield different growth profiles still.

Returning to the reference results of this work, the relative power of the density perturbation in the lowest few Fourier modes is shown in Fig. 4.7. Unlike Ott et al. [70], but consistent with Scheidegger et al. [65] and Corvino et al., we find $m = 2$ to be the dominant mode. This is also the mode whose interaction with magnetic fields was analyzed in detail by Fu & Lai [74].

Magnetic effects

We find that the presence of a magnetic field could have two competing effects on the growth of the $m = 2$ fluid instability. Simulations with fields of 4×10^{13} G and greater demonstrate suppression of the instability, with the distortion parameter saturating at a significantly smaller value (3–50× lower) than in an unmagnetized star. Even stronger fields (starting at 5×10^{14} G), however, made the star susceptible to a small-scale (few gridpoints per wavelength) magnetic instability that rapidly amplified the $m = 2$ distortion of the star (in addition to other modes). This instability may operate at lower field strengths as well, but there its effects would not be resolvable at our current resolution. The net behavior for all simulated cases is plotted in Figs. 4.8 and 4.9 and is qualitatively independent of the seed field geometry (parameterized by n_s ; in particular, the threshold for instability

²These partial simulations were performed by Christian D. Ott and Philipp Mösta.

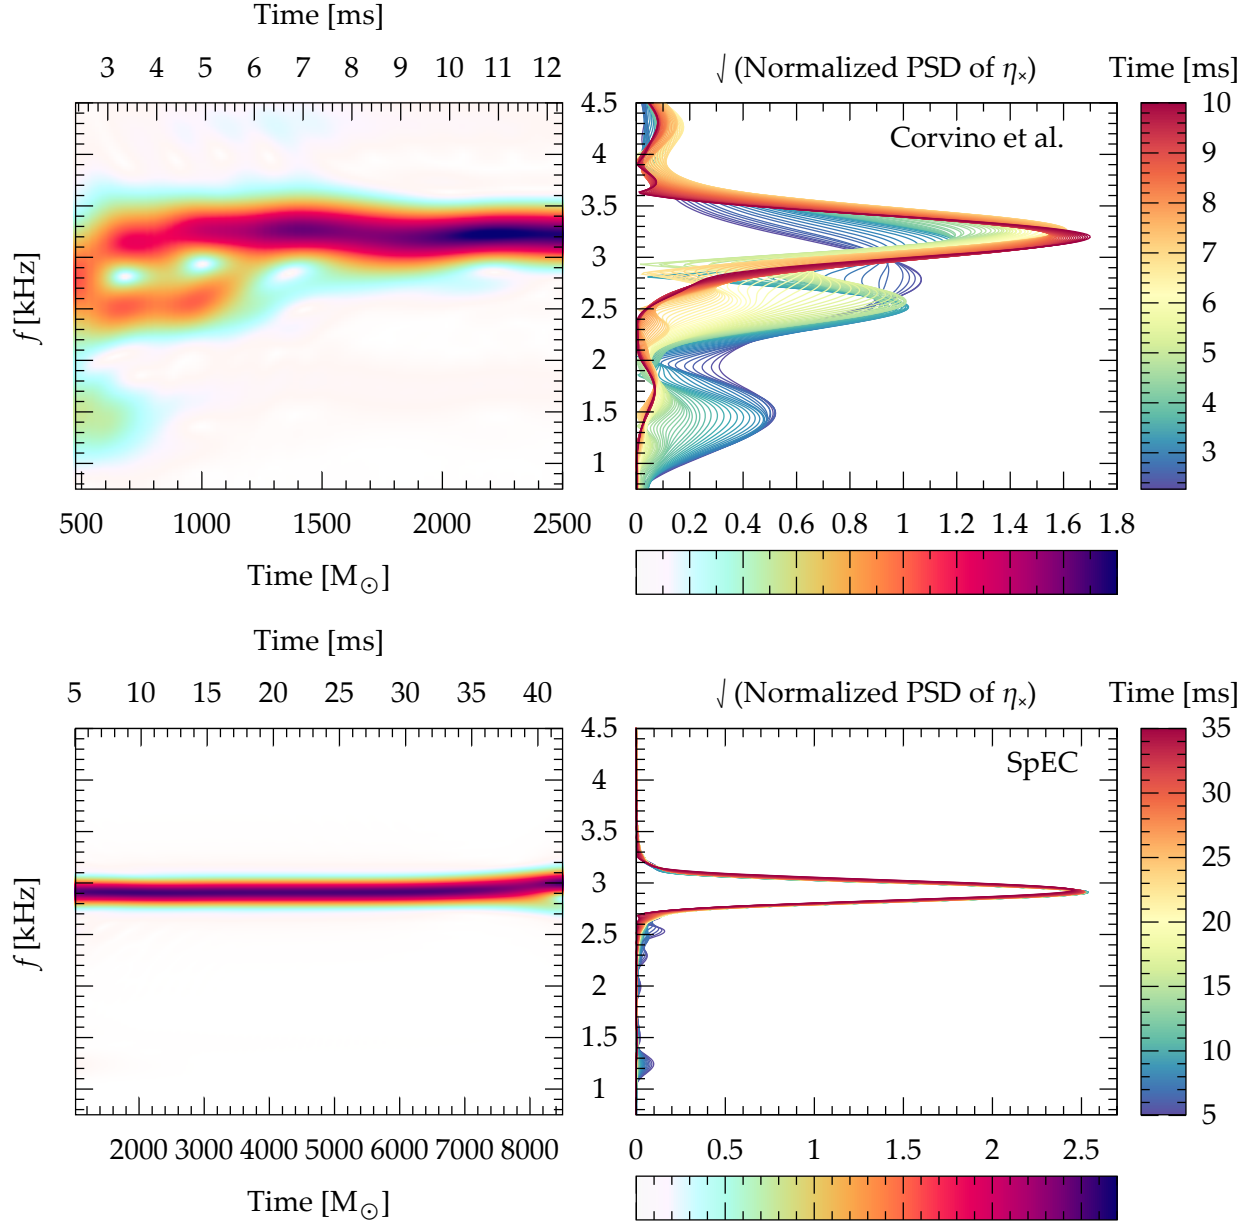


Figure 4.5: Comparison of the normalized power spectral densities (PSDs) of η_{\times} as simulated by Corvino et al. (top) and by SpEC (bottom). In Corvino et al.’s PPM simulation, the growth of the distortion parameter is the results of several unstable modes growing with different timescales. In contrast, SpEC’s WENO simulation exhibits a single unstable mode at a fixed frequency. The PSDs are estimated via FFT periodograms using Welch’s method with a Hann window and are normalized by the total power at each time.

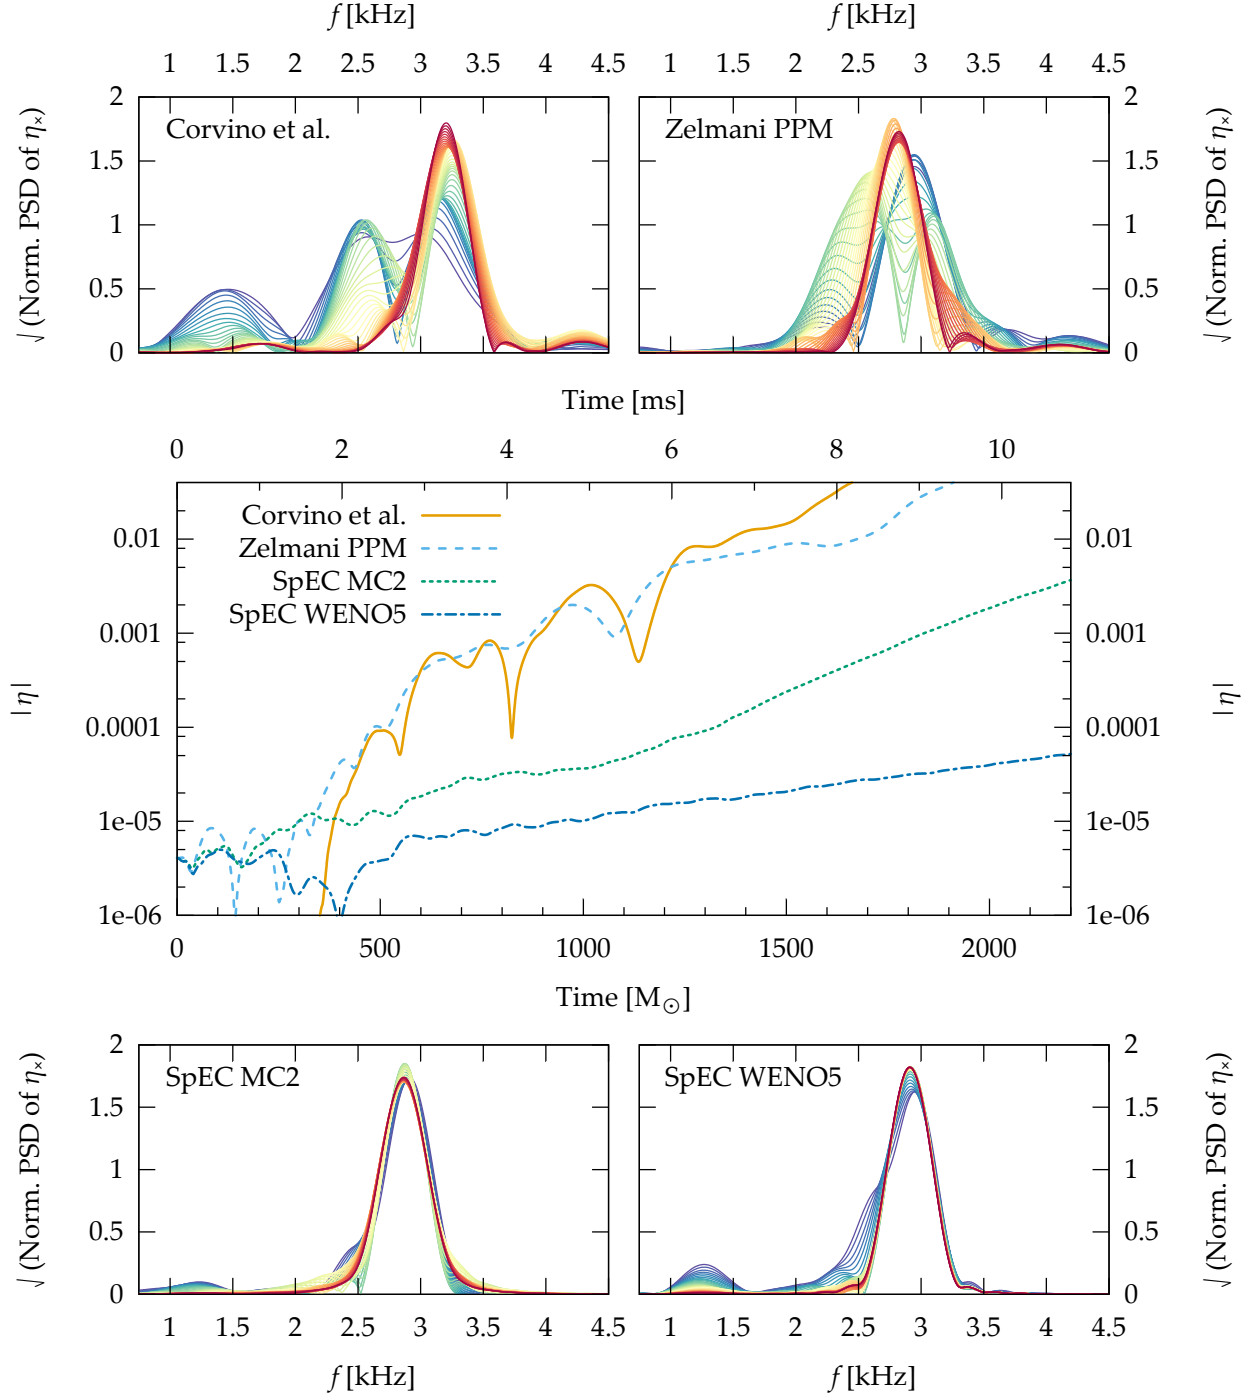


Figure 4.6: Effects of different reconstructors and different simulation codes on the spectrum and growth rate of the distortion parameter $|\eta|$ for $t < 10$ ms. Both PPM results exhibit the most complicated spectra and the fastest growth rates; however, the spectra show notable differences, and convergence for these cases has not been demonstrated. Both SpEC results possess clean spectra with lower growth rates, and the WENO5 results are consistent with different resolutions. Preliminary Zelmani results using WENO look most similar to SpEC’s WENO results. The Zelmani and SpEC results plotted here were evolved with a resolution of $\Delta x = 0.2 M_\odot$.




Figure 4.7: Relative power of ρ in azimuthal modes for $m = 1-4$. Note that measurements of $m = 4$ power have a noise floor of 10^{-3} due to the Cartesian nature of the grid.

appears to be the same).

Simulations of these magnetically unstable cases were halted prior to the original saturation time, as magnetized outflows of matter began to leave the grid. Both magnetically-dominated and pressure-dominated matter leave the star relatively isotropically with mildly relativistic velocities ($W_L \lesssim 0.15$). The stronger the magnetic field, the sooner these outflows develop. Similar outflows have been noted in previous investigations [84, 85], though due to the small size of our grid, we cannot make quantitative comparisons.

Suppression of the low- $T/|W|$ instability

When we observe suppression, we would like to determine whether the mechanism is consistent with that proposed by Fu & Lai. Unfortunately, the correspondence is far from




Figure 4.8: Range of behavior of distortion parameter η at different magnetic field strengths for $n_s = 1$. Curves that terminate at early times developed significant outflows, making further evolution impractical on our grid.

clear. In particular, while magnetic winding produces peak toroidal field strengths comparable to those considered in their work (and surpassing their threshold for suppression of 2×10^{16} G), the total magnetic energy saturates at much lower values than they deem necessary for suppression to take place. Our runs with initial poloidal field strengths on the order of $B_0 \approx 2 \times 10^{14}$ G wind up toroidal fields as strong as 10^{17} G but with magnetic energies of only half a percent of the star's kinetic energy. For comparison, their model implies that such fields would possess magnetic energy equivalent to 20% of T , which they find is the minimum energy ratio for suppression to occur.

We see that magnetic winding increases the magnetic energy in the star at the expense of




Figure 4.9: Range of behavior of distortion parameter η at different magnetic field strengths for $n_s = 2$, showing same classes of behavior as when $n_s = 1$ (see Fig. 4.8).

gravitational potential energy, as shown in Fig. 4.10, but saturates within 30 ms in the cases we considered (prior to the saturation of the low- $T/|W|$ instability). Matter near the core of the star is compacted, increasing the central density. The internal energy of the matter also increases in magnetized scenarios, but the kinetic energy is barely affected in most cases. For the magnetically-unstable systems, however, kinetic energy from non-azimuthal fluid velocities grows exponentially at late times as the rotational kinetic energy begins to decrease at an amplified rate (the separation of rotational and non-rotational kinetic energy is not shown in the figure). This likely corresponds to small-scale fluid oscillations associated with the magnetic turbulence described below.

Other comparisons are difficult as well. In Fig. 5 of their paper, Fu & Lai show that the




Figure 4.10: Energy exchange for three magnetic field strengths ($n_s = 1$ for each case). The change in gravitational energy is inferred from the sum of the changes in the other energies.




Figure 4.11: Lagrangian displacement of tracer particles seeded at various cylindrical radii for an unmagnetized star. For each initial radius, 12 tracers were distributed uniformly in azimuth. The corotation radius for this system is at $\varpi \approx 4.25 M_\odot$.

Lagrangian displacement of fluid elements should diverge at the corotation radius during the low- $T/|W|$ instability, but that this resonance should split in the presence of a strong toroidal magnetic field. Using tracers, we do see an amplification in radial displacement in the vicinity of the corotation radius in the unmagnetized case (see Fig. 4.11), but the response is so broad that we cannot resolve any splitting when magnetic fields are added.

Nevertheless, there are clues pointing to a resonance splitting. In particular, spectrograms of the distortion parameter show a split peak when magnetic suppression is




Figure 4.12: Spectrograms of the quadrupole moment I_{xy} for six cases. The power spectral density (PSD) is estimated via an FFT periodogram using Welch’s method with a Hann window. The PSD is normalized by the total power (in arbitrary units) at each time to highlight relative changes in the frequency spectrum.

observed (see Fig. 4.12). The magnitude of splitting for $B_0 = 2 \times 10^{14} \text{ G}$, $n_s = 2$ is about $\Delta\omega \approx 2\pi \times 0.1 \text{ kHz}$. Defining the angular Alfvén speed,

$$\omega_A \equiv B^\phi / (\bar{\omega} \sqrt{\rho}) , \quad (4.16)$$

and the slow magnetosonic wave frequency,

$$\omega_s \equiv \sqrt{\frac{c_s^2}{c_s^2 + (B^\phi)^2/\rho}} m \omega_A , \quad (4.17)$$

(where c_s is the adiabatic sound speed), resonances are expected at $\Delta\omega = \omega_s$ and (in the full 3D case) $\Delta\omega = m\omega_A$. In the strongly magnetized regions of the star, the observed splitting agrees with the values of ω_s and $2\omega_A$ to within a factor of four. Given the differences in the particular systems under study, this is reasonably consistent with Fu & Lai’s proposed mechanism.




Figure 4.13: Magnitude of radial component of B -field in the y - z plane at $t = 3760 M_\odot$ for $B_0 = 5 \times 10^{14} \text{ G}$, $n_s = 2$.

Magnetic instability

When the initial magnetic field exceeds $B_0 = 5 \times 10^{14} \text{ G}$, our simulations start to exhibit strong magnetic instability. This instability results both in the amplification of low- m global modes in the star and in turbulence at the smallest scales we can resolve on our grid. The marginally-resolved nature of this instability complicates its identification and interpretation.

The growth of small-scale features is most visible in poloidal field components, as illustrated in Figs. 4.13 & 4.14, while large-scale nonaxisymmetric structure is easily seen in the much stronger toroidal field (see Fig. 4.15). The crest-to-crest separation of the poloidal perturbations is measured to be approximately $\lambda \sim 1 M_\odot$, which is resolved by roughly five gridpoints. This suggests that the unstable modes are only marginally resolved, so we




Figure 4.14: Magnitude of radial component of B -field vs. radius vs. time in the $z = 1$ plane for three configurations ($n_s = 1$ in all cases), illustrating the onset of turbulence. Colorbars are scaled relative to the initial B -field strength. Plot inspired by the analysis of Franci et al. [73].

cannot expect their subsequent evolution to be more than qualitatively correct (at best).

In fact, magnetically-driven instabilities in the fluid are not unexpected. Magnetic winding generates a strong toroidal field in the interior of the star, and toroidal field gradients are potentially unstable to kink (Taylor) and buoyancy (Parker) instabilities [86–90]. For a toroidal field centered on the rotation axis, the Taylor instability can occur at cylindrical radii ϖ less than the radial pressure scale height H_P (defined as in [88, 91] as $2c_s^2/g_\varpi$, with g_ϖ denoting the radial acceleration) for positive $dB^\phi/d\varpi$. Kink instabilities have in fact recently been identified in 3D magnetized core-collapse simulations [85]. The Parker instability can be triggered by radial or vertical field gradients (negative $dB^\phi/d\varpi$ for $\varpi > H_P$ or negative dB^ϕ/dz). The growth rate of the Taylor instability is of order the angular Alfvén speed ω_A for weak rotation and ω_A^2/Ω for strong rotation, where $\Omega \gg \omega_A$ is the condition for strong rotation [92]. Growth timescales for the Parker instability are similar. Although much analytic work on field-gradient instabilities assumes weak differential rotation, the Parker instability has been found to be operable even in some flows with strong shear [93]. In our magnetically-unstable cases, ω_A/Ω is $O(1/2)$ at the corotation radius, suggesting an intermediate regime between weak and strong rotation.

In addition to the above-mentioned field gradient-driven instabilities, differential rotation will also trigger shear-driven instabilities. The most famous is the classic magnetorotational instability (MRI), an axisymmetric instability triggered by a nonzero (but arbitrarily small) poloidal field and an outward-decreasing rotation rate [94]. More generally, the MRI can also be found in nonaxisymmetric configurations [95, 96], in which case the background toroidal field can also contribute to seeding the instability [88, 95]. The fastest-growing unstable mode grows on a timescale of $\sim \Omega^{-1}$ and has a wavenumber given by

$$\Omega/\sqrt{-g_{00}} \sim \mathbf{k} \cdot \mathbf{v}_A \approx \frac{k^\varpi B^\varpi + k^z B^z + m B^\phi/\varpi}{\sqrt{\rho h + b^2}} \quad (4.18)$$

(on the relativistic factor, see Siegel et al. [97]). The main challenge for numerical MHD simulations is to resolve the MRI wavelength $\lambda_{\text{MRI}} = 2\pi|\mathbf{k}|^{-1}$. Since the field is usually azimuthally-dominated, we see that $m \neq 0$ modes are potentially easier to resolve, a




Figure 4.15: Power of b^2 in azimuthal modes for $m = 1$ – 4 . Except in the most strongly magnetized systems, the $m = 4$ power does not rise above that of the ambient grid mode. The growing strength of the magnetic field is factored out by normalizing by the $m = 0$ power; thus, trends shown here represent growth of the proportional power of nonaxisymmetric modes.

fact also recently noted by Franci et al. [73], who resolve MRI-like field growth only in nonaxisymmetrically-unstable stars. On the other hand, the growth of a given nonaxisymmetric mode will be expected to terminate when the mode becomes too tightly wound [95]. In fact, it has long been known that even a purely toroidal field can seed a shear instability [88, 95, 96], although the growth timescales tend to be longer than those associated with poloidal seed fields, except for the case of very high m , and in that case even a small poloidal field would be expected to radically alter the flow [96].

Given the presence of differential rotation and a poloidal magnetic field, our system is certainly susceptible to the MRI; what is less clear is our ability to resolve it. Siegel et al. [97] state that a minimum of five gridpoints per wavelength was required to resolve the MRI in their simulations. Using Eq. (4.18), we can estimate what the wavelength of the fastest-growing unstable mode would be at any point in our simulation, optimizing over propagation directions. Comparing this to our effective grid resolution in those directions, we find that when turbulence starts to develop in our systems, there are $O(\text{few})$ gridpoints per wavelength in the unstable regions of the star even for $m = 0$ modes, and when considering higher m , these unstable regions begin to meet the criterion of five gridpoints per wavelength. Therefore, resolving the MRI, if only marginally, is conceivable given our resolution and magnetic field strengths.

One approach to diagnosing the source of turbulence is to measure the growth rates of observed instabilities and match them to linear predictions. As mentioned above, the Tayler and Parker instabilities should grow at a rate between ω_A and ω_A^2/Ω , while the MRI's growth rate is Ω , independent of the B -field magnitude. The rotational frequency of the star in the region of magnetic instability (which occurs in the vicinity of the corotation radius) is about $\Omega \approx 1.45 \times 2\pi$ kHz.

Looking at the growth of the most magnetized point on the grid (see Fig. 4.16) reveals exponential behavior at rates that increase with the magnetic field strength. This scaling, in addition to the magnitude of the rates, is incompatible with the MRI (while the expected

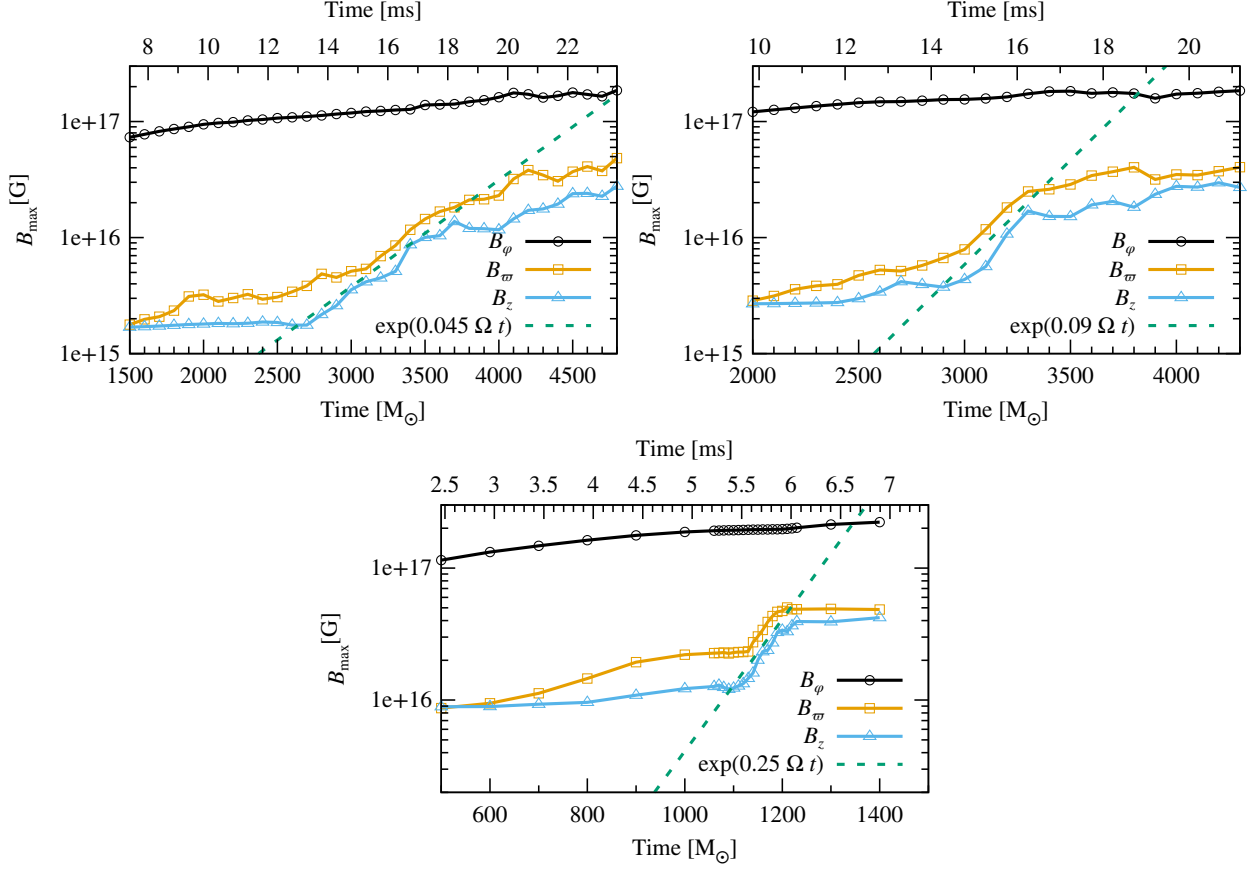


Figure 4.16: Growth of the maximum of the cylindrical components of the B -field for three cases: $B_0 = 5 \times 10^{14} \text{ G}$, $n_s = 1$, $B_0 = 5 \times 10^{14} \text{ G}$, $n_s = 2$, and $B_0 = 2 \times 10^{15} \text{ G}$, $n_s = 1$. The temporal resolution during the period of rapid growth for the last case is $10\times$ finer than our default.

rate of Ω is an approximation derived from accretion disks, the numerical prefactor for our system is expected to be $O(3/4)$, insufficient to explain the discrepancy).

Considering the field gradient-driven instabilities, the “weak rotation” rate of ω_A is too large as well and also does not match the observed scaling with B -field strength. The “strong rotation” prediction, however, while still larger than observed, is only off by a factor of a few and is the closest match to the data in terms of scaling. This suggests that, while the MRI is potentially resolvable with our techniques, the observed local maximum B -field growth is most attributable to field gradient instabilities. Shear instabilities are almost certainly still present and impacting the dynamics, however, and likely play a large role

in less-magnetized cases where we currently cannot resolve them. In fact, their expected growth rates suggest that they would dominate the dynamics on relevant timescales were they resolved.

Detectability

To help put these results in an astrophysical context, we consider the detectability of gravitational waves produced by the (unmagnetized) low- $T/|W|$ instability for this system. We follow the procedure outlined by Sutton [98]. Given both polarizations of the gravitational wave strain, h_+ and h_\times , at some distance from the source, define the root-sum-square amplitude h_{rss} to be

$$h_{\text{rss}} = \sqrt{\int \left(h_+^2(t) + h_\times^2(t) \right) dt}. \quad (4.19)$$

For a narrow-band signal from a rotating system like ours, we expect the emitted gravitational wave energy E_{GW} to be well-approximated by

$$E_{\text{GW}} \approx \frac{2}{5} \frac{\pi^2 c^3}{G} f_0^2 r^2 h_{\text{rss}}^2, \quad (4.20)$$

where f_0 is the central frequency of the signal. The effective detection range \mathcal{R}_{eff} for a narrow-band burst signal is given by

$$\mathcal{R}_{\text{eff}} = \beta \sqrt{\frac{G}{\pi^2 c^3} \frac{E_{\text{GW}}}{S(f_0) f_0^2 \rho_{\text{det}}^2}}, \quad (4.21)$$

where $S(f)$ is the one-sided noise power spectrum for the target detector, ρ_{det} is the threshold signal-to-noise ratio for detection, and β is a geometrical factor related to the polarization of the waves. Specializing to rotating sources, this becomes

$$\mathcal{R}_{\text{eff}} = 0.698 \frac{r h_{\text{rss}}}{\rho_{\text{det}}} \sqrt{\frac{2}{5} \frac{1}{S(f_0)}}. \quad (4.22)$$

We extract gravitational waves from our simulations at a radius of $400 M_\odot$ using Regge-Wheeler-Zerilli techniques [99] and consider the strains h_+ and h_\times for an observer above the axis of rotation. For the unmagnetized star considered in this work, the gravitational wave

frequency is sharply peaked at 2.9 kHz (this is slightly lower than the 3.2 kHz primary peak observed by Corvino et al. [66]). If we consider only the instability's initial growth through saturation, the total emitted gravitational wave energy is 3.68×10^{50} erg ($2.06 \times 10^{-4} M_{\odot}$). Using the ZERO_DET_high_P noise curve for Advanced LIGO [100] and a signal-to-noise threshold of $\rho_{\text{det}} = 20$, this instability would be detectable out to 92 kpc.

The emitted gravitational wave energy is significantly larger than what was found in core-collapse supernovae simulations [64, 65] ($E_{\text{GW}} \sim 10^{46}$ – 10^{47} erg for a similar simulation length). However, the difference can easily be understood by noting that the neutron star considered in this work rotates significantly more rapidly (with the wave signal peaking at 2.9 kHz vs. ~ 0.9 kHz in the core-collapse results) and is also more massive than protoneutron stars are expected to be. Since $E_{\text{GW}} \propto M^2 \Omega^6$, this accounts for most of the difference in the emitted gravitational wave energy. On the other hand, the more slowly rotating neutron stars emit waves at a more favorable frequency, improving their detectability.

The effect of magnetic fields on detectability is difficult to discern from our data, as outflows prevented us from evolving the most highly magnetized systems long enough to see the instability saturate. For $B_0 = 5 \times 10^{14}$ G, $n_s = 2$, the distortion parameter peaks nearly as high as the saturation value in the unmagnetized case while the frequency spectrum at that time peaks at a slightly lower (and more favorable) value, suggesting that a gravitational wave signal from magnetic instabilities could be just as detectable as that of the unmagnetized low- $T/|W|$ instability. On the other hand, mildly magnetized cases exhibit a suppressed distortion parameter with an unchanged frequency spectrum. Using the quadrupole approximation, and the fact that \mathcal{R}_{eff} is linear in h_{rss} , this means that the effective detection range is decreased by factor of ~ 2.4 for $B_0 = 4 \times 10^{13}$ G, $n_s = 1$, and by a factor of ~ 34 for $B_0 = 2 \times 10^{14}$ G, $n_s = 1$, for an observer above the axis of rotation.

4.5 Summary

In writing an MHD module for `SpEC`, we have expanded the range and fidelity of astrophysical systems that can be simulated while still taking advantage of its highly accurate spacetime evolution. The future scope of this code includes many systems of contemporary interest, including magnetized compact binary coalescence, but here we focus our attention on instabilities in differentially rotating neutron stars.

Of significant relevance to existing literature regarding these stars is the variability in simulated instability spectra and their corresponding growth rates when using different resolutions and reconstruction methods. It appears that the choice of reconstructor can have a significant effect on the stability and excitement of potentially unstable modes. We find qualitative convergence when using high resolution and high-order reconstruction, but these results differ significantly from those of lower-accuracy techniques and of some previous studies. Further investigation of such instabilities' delicate dependence on simulation methods is warranted.

Regarding the low- $T/|W|$ instability, it is clear that poloidal magnetic fields on the order of 10^{14} G can have a strong effect on the distribution of mass in differentially rotating neutron stars and therefore on their gravitational wave signatures. However, while suppression of the instability is feasible, it occurs in a small region of parameter space. B -fields strong enough to enable the suppression mechanism are likely also strong enough to trigger magnetic instabilities, accelerating the growth of a mass quadrupole moment rather than suppressing it.

In our simulations, with clean poloidal initial fields, the window between the onsets of magnetic suppression and magnetic instability – roughly 4×10^{13} G– 5×10^{14} G – is rather small, and future runs with increased resolution may lower the upper bound still further. Therefore, amplification of matter perturbations seems to be the more likely magnetic effect, with peak amplitudes comparable to those in the unmagnetized case. The spectrum of the gravitational waves, while perhaps possessing more structure, will also remain peaked

near the same frequency. As a result, even with such extreme field strengths, the net effect on burst detectability is likely minor.

Regarding Fu & Lai's conclusions, we find some disagreement between their predictions for cylindrical stars and our simulations of realistic ones. In particular, they concluded that suppression would occur once the magnetic energy H_B reached about 20% of the kinetic energy T . However, the magnetic energy in our simulations peaks at 0.56% of T , yet we still find suppression in some cases. Despite this, we agree on the minimum strength of the poloidal seed field, roughly 10^{14} G. Additionally, the frequency spectrum of the instability is consistent with their proposed mechanism for suppression.

Uncertainties in our investigation include the details of the formation of the star and its seed field, as nature will not be nearly as clean as the system we considered. Additionally, we expect that if the MRI were fully resolved, it would grow on such a short timescale that it would dominate the effects observed here.

Future work to understand the details of the suppression mechanism could investigate the effects of purely toroidal fields, removing the complications of magnetic winding and the MRI. On the other hand, the impact of the magnetic instabilities could be better understood by increasing resolution and by extending the simulations to observe their saturation behavior. Additionally, the systematic effects of reconstruction order and grid resolution on the growth rate of this particular instability warrant further investigation. Lastly, while this chapter has limited itself to studying the growth of instabilities, the later evolution of such stars, after the commencement of magnetically-driven winds, would be a very astrophysically interesting subject for future numerical modeling.

5 CONCLUSION

With its new capabilities to simulate binary neutron stars and magnetic fields, along with numerous efficiency improvements, SpEC is poised to contribute a wealth of new information about the behavior of gravitational wave sources and high-energy astrophysical events. Our NS–NS initial data solver allows NS–NS systems to benefit from the advantages of SpEC’s spectral methods, in particular, the ability to simulate long, high-accuracy inspirals, as has recently been done by Roland Haas and collaborators. It also opens the door to comparisons with independent NS–NS codes, improving our understanding of the accuracy of numerical relativity simulations. Additionally, it forms the basis of work by Nick Tacik to construct neutron star initial data with arbitrary spins, lifting the irrotational or corotational restriction of the current solver and similar solvers used by other groups.. This capability expands the parameter space we can simulate and improves the realism of our models.

Our new ability to simulate magnetic fields similarly extends the field of systems we can simulate, both in terms of comparisons with other groups (such as the University of Illinois) and in terms of more accurately modeling real systems. It also synergizes with the other advantages of SpEC, including our ability, thanks largely to Geoffrey Lovelace and Francois Foucart, to construct and evolve BH–NS systems containing rapidly-rotating black holes.

There is also much still to explore in the context of magnetized differentially-rotating neutron stars. The magnetized outflows observed by us and other groups deserve further study, and our understanding of the magnetic instabilities at work would be greatly improved by simulations with additional resolution (though computational limits make this impractical in the near-term). The strong dependence we observed of the growth of fluid instabilities on the reconstruction algorithm used also serves as a warning sign that simulations can easily lead us astray, especially when working at the limits of what

can be resolved. Understanding the precise relationship between reconstructors and fluid instabilities is a project that I think would be valuable to our community.

The age of gravitational wave astronomy is at our doorstep, and simulations have a central role to play. While the need for additional computing power and more complete microphysics may never be completely satisfied, the incremental improvements described in this thesis, along with many others, are paving the way for a greater understanding of high-energy events in the universe.

A BASIS FUNCTIONS FOR CYLINDERS AND SPHERES

The Spectral Einstein Code solves elliptic and hyperbolic PDEs using multidomain pseudospectral methods. The spatial domain of the problem is decomposed into multiple subdomains of various shapes, including cubes, cylinders, and spheres. Through coordinate maps, these shapes can be distorted into more general shapes, such as wedges, allowing them to fill the domain with minimal overlap. Within each subdomain, the solution is represented using a local basis of functions tailored to the topology of the shape, and the solution is not required to be continuous across subdomain boundaries. This multidomain approach has several advantages over using a single global domain, including improved parallelism, the ability to refine resolution locally, and the ability to accommodate discontinuities in the solution (provided they occur at subdomain boundaries).

For reasons of efficiency, it is desirable to minimize the distortion introduced by coordinate maps, for the basis is in effect responsible for resolving the product of the solution and the map. Furthermore, it is best to choose shapes well-adapted to any symmetries expected in the solution (for example, the density in a spherical star is more efficiently represented by the basis functions within spheres and spherical shells than by those of a collection of cubes). Having a variety of subdomain shapes available is therefore beneficial. Unfortunately, the centers of spheres and cylinders require special treatment to handle the coordinate singularities there, making implementing a well-suited basis for these shapes non-trivial. While the requisite mathematics are well-described in the literature, names and notation are inconsistent, and there is considerable room for errors when putting all the pieces together.

Because of instabilities encountered with a previous implementation, SpEC for a time “plugged the holes” in cylindrical and spherical shells using cubes. However, given the

symmetries often encountered in astrophysical systems, having a reliable implementation of basis functions for cylinders and spheres is valuable for both evolutions and subsequent analysis. Here I describe the theory and practice of these functions in full detail.

A.1 Background on spectral methods

Consider a function space spanned by a set of N basis functions $\phi_n(x)$ that are orthonormal with respect to a weight function $w(x)$; that is,

$$\int \phi_m(x)\phi_n(x)w(x)dx = \delta_{mn} . \quad (\text{A.1})$$

Further, assume the existence of a quadrature rule on a set of N collocation points x_i that is exact for all products of two functions in this space weighted by $w(x)$. In other words,

$$\sum_{i=0}^{N-1} \phi_m(x_i)\phi_n(x_i)w_i = \delta_{mn} , \quad (\text{A.2})$$

where w_i are the quadrature weights. Note that Gaussian quadrature meets this criterion for polynomial bases: it integrates exactly for polynomials of order up to $2N - 1$, while the order of the integrand is at most $2N - 2$.

Let $f(x)$ be a member of this space, which we write as a linear combination of the basis functions:

$$f(x) = \sum_{n=0}^{N-1} f_n \phi_n(x) , \quad (\text{A.3})$$

where the spectral coefficients f_n can be computed via

$$f_n = \int f(x)\phi_n(x)w(x)dx = \sum_{i=0}^{N-1} f(x_i)\phi_n(x_i)w_i . \quad (\text{A.4})$$

Cardinal functions

There exists a unique set of cardinal function $C_i(x)$ in this space with the property that

$$f(x) = \sum_{i=0}^{N-1} f(x_i)C_i(x) , \quad (\text{A.5})$$

which we can solve for as follows: First, expand each $C_i(x)$ into its spectral coefficients $c_{n,i}$. Then,

$$f(x) = \sum_{i=0}^{N-1} f(x_i) C_i(x) = \sum_{i=0}^{N-1} f(x_i) \sum_{n=0}^{N-1} c_{n,i} \phi_n(x),$$

which implies that

$$\sum_{n=0}^{N-1} f_n \phi_n(x) = \sum_{n=0}^{N-1} \left(\sum_{i=0}^{N-1} f(x_i) c_{n,i} \right) \phi_n(x),$$

and thus that

$$f_n = \sum_{i=0}^{N-1} f(x_i) \phi_n(x_i) w_i = \sum_{i=0}^{N-1} f(x_i) c_{n,i}.$$

This means that

$$c_{n,i} = \phi_n(x_i) w_i,$$

and therefore

$$C_i(x) = w_i \sum_{n=0}^{N-1} \phi_n(x_i) \phi_n(x). \quad (\text{A.6})$$

Observe that the cardinal functions obey the property

$$C_i(x_j) = \delta_{ij} \quad (\text{A.7})$$

and are orthogonal to one another with norm $\sqrt{w_i}$:

$$\int C_i(x) C_j(x) w(x) dx = w_i \delta_{ij}. \quad (\text{A.8})$$

Thus, the functions $\tilde{C}_i(x) \equiv C_i(x)/\sqrt{w_i}$ form another orthonormal basis for the space. (Note that this also provides a convenient way of computing the quadrature weights via $1/w_i = \sum_n \phi_n^2(x_i)$.) For polynomial bases of degree less than the number of collocation points, the terms of the Lagrange interpolating polynomial,

$$P_i(x) = \prod_{\substack{j=1 \\ j \neq i}}^{N-1} \frac{x - x_j}{x_i - x_j}, \quad (\text{A.9})$$

provide a set of cardinal functions.

Lower quadrature orders

Key to the above analysis is the fact that the quadrature rule was of sufficiently high order to perfectly project any function in the space onto the basis functions. While Fourier and Gaussian quadratures satisfy this requirement, it is often convenient to employ a lower-order quadrature rule when implementing a spectral method. Gauss–Lobatto quadrature, for example, places a collocation point at each endpoint of the interval, while Gauss–Radau quadrature places a node at a single endpoint. Having collocation points at the edges of the interval is convenient for imposing boundary conditions. Unfortunately, Gauss–Lobatto quadrature on N collocation points is only exact for polynomials of order $2N - 3$, and while Gauss–Radau is exact up to order $2N - 2$, there are situations (such as with the Zernike polynomials discussed below) where this is insufficient. Thus, when using the procedure above, the highest-order modes will be analyzed incorrectly.

There are two ways to work around this problem. One is to reduce the rank of the basis without reducing the number of collocation points. The other is to normalize the basis functions with respect to the quadrature rule, rather than analytically. Failure to adopt one of these approaches will result in an incorrect and often unstable implementation, as transforming a sampled function from “physical” space to spectral space and back again will not correspond to the identity operation.

A.2 Zernike polynomials

When evolving the spacetime metric on a spectral grid, we try to adapt the domain decomposition to the geometry of the evolved fields. This often means using sections of a sphere, in the form of spherical shells or “cubed spheres.” In black hole spacetimes, this is sufficient to cover the area surrounding the excised region within the apparent horizon. However, for neutron star spacetimes, a different approach is taken to cover the center of the star.

Polar and spherical coordinates are singular at the origin, creating difficulties if one tries to use tensor products of one-dimensional function bases. This same problem exists at the poles of a spherical surface. Spherical harmonics, $Y_l^m(\theta, \phi)$, provide a clean solution in that case, able to represent smooth functions without artificial boundaries and without severely restricting the timestep allowed by the Courant–Friedrichs–Lewy stability limit [101]. For the radial “pole problem,” Zernike polynomials and their higher dimensional generalizations provide a similar solution.

The use of Zernike polynomials in spectral methods over the unit disk, B^2 , was explored independently by Matsushima and Marcus [102] and by Verkley [103]. Notation varies throughout the literature, so we summarize ours here:

Denote an orthonormal azimuthal (Fourier) basis as

$$F_m(\phi) \equiv \begin{cases} \frac{1}{\sqrt{2\pi}} & m = 0 \\ \frac{1}{\sqrt{\pi}} e^{im\phi} & m > 0 \end{cases}. \quad (\text{A.10})$$

Then an arbitrary smooth function $f(\omega, \phi)$ over B^2 can be decomposed into its Fourier coefficients $f_m(\omega)$:

$$f(\omega, \phi) = \Re \sum_{m=0}^{m_{\max}} f_m(\omega) F_m(\phi), \quad (\text{A.11})$$

where $m_{\max} = \lfloor N_\phi/2 \rfloor$, N_ϕ being the number of azimuthal collocation points. (Note that if N_ϕ is odd, the highest mode will lack a sine component.)

These Fourier coefficients can be further decomposed into a radial sub-basis $R_n^m(\omega)$, composed of one-sided Jacobi polynomials multiplied by ω^m :

$$R_n^m(\omega) \equiv \sqrt{2n+2} \omega^m P_{(n-m)/2}^{(0,m)}(2\omega^2 - 1), \quad (\text{A.12})$$

where $P_k^{(\alpha,\beta)}(x)$ represents the Jacobi polynomial of degree k . In this notation, the radial functions are only defined for $n \geq m$, $2 \mid (n - m)$. For smooth functions, the $f_m(\omega)$ must satisfy the pole condition: $f_m(\omega) \rightarrow \omega^m$ as $\omega \rightarrow 0$. This basis manifestly respects that condition.

The Zernike polynomials are then defined as

$$Z_m^n(\varpi, \phi) \equiv R_n^m(\varpi)F_m(\phi). \quad (\text{A.13})$$

They form an orthonormal basis for smooth functions over the unit disk:

$$f(\varpi, \phi) = \sum_{m=0}^{m_{\max}} \sum_{\substack{n=m \\ n+=2}}^{n_{\max}} f_{nm} Z_m^n(\varpi, \phi), \quad (\text{A.14})$$

where $n_{\max} = 2N_{\varpi} - 1$, N_{ϖ} being the number of radial collocation points. Note that if Gauss–Radau quadrature is used (placing collocation points on the outer boundary of the disk), then the highest-order radial basis functions should be normalized with respect to the quadrature rule (rather than analytically) or else omitted entirely. Specifications for the quadrature nodes and weights can be found in the references.

As mentioned by Livermore et al. [104], this basis can be generalized to filled spheres (B^3). In that case, a function $f(r, \theta, \phi)$ is decomposed into f_{nlm} such that

$$f(r, \theta, \phi) = \sum_{m=-m_{\max}}^{m_{\max}} \sum_{l=|m|}^{l_{\max}} \sum_{\substack{n=l \\ n+=2}}^{n_{\max}} f_{nlm} R_n^l(r) Y_l^m(\theta, \phi), \quad (\text{A.15})$$

where now $R_n^l(r)$ is given by

$$R_n^l(r) = \sqrt{2n+3} r^l P_{(n-l)/2}^{(0, l+1/2)}(2r^2 - 1), \quad (\text{A.16})$$

which corresponds to an integration weight of r^2 instead of ϖ . Here, $Y_l^m(\theta, \phi)$ are the spherical harmonics, and $l_{\max} = N_{\theta} - 1$ for N_{θ} latitudinal collocation points.

A.3 Spectral method

To use Zernike polynomials as the basis for a subdomain in SpEC, one needs to provide efficient routines for spectral transforms, interpolation, differentiation, integration, and filtering. As the coupling of the radial polynomials with the angular basis is not as simple as an tensor product, and as the order of the radial polynomials exceeds the number of radial collocation points, here I outline the implementation of these routines.

Quadrature

Because the radial polynomials have definite parity for a given m or l , Gaussian quadrature can integrate products of them exactly up to a combined order of $4N_{\omega,r} - 2$ (twice the order for arbitrary polynomials sampled at the same number of points). This is just sufficient to project any polynomial of order $\leq n_{\max}$ onto the basis. However, Gauss–Radau quadrature, which would place a collocation point on the boundary of the disk or sphere, can only integrate up to a combined order of $4N_{\omega,r} - 4$. This results in a normalization error when integrating terms of order n_{\max} . While this normalization could be corrected by modifying the projection procedure for the highest mode, for simplicity we instead opt to add an extra collocation point while only incrementing n_{\max} by 1 instead of 2 (or equivalently, reduce n_{\max} by 1 for the same number of points). This is also the approach outlined by Matsushima and Marcus [102].

For a Gauss–Radau grid, we therefore have the following requirements and relationships:

Filled disks (B^2):

- $m_{\max} = \lfloor N_{\phi}/2 \rfloor$, $n_{\max} = 2N_{\omega} - 2$
- $2N_r \geq \lfloor N_{\phi}/2 \rfloor + 2$

Filled spheres (B^3):

- $m_{\max} = \lfloor N_{\phi}/2 \rfloor$, $l_{\max} = N_{\theta} - 1$, $n_{\max} = 2N_{\omega} - 2$
- $N_{\theta} \geq \lfloor N_{\phi}/2 \rfloor + 1$, $2N_r \geq N_{\theta} + 1$

Failure to meet the resolution inequalities means that the order of accuracy implied by the definitions above them will not be met—the high angular modes will not be resolved, as there will be no radial (or polar) modes of sufficient order to resolve them while still satisfying the pole condition.

Spectral transforms

Transforming data from physical to spectral space is a two-step process. First, angular transforms are performed on rings or shells of data at each collocation radius ϖ_i or r_i . For B^2 , this requires a set of N_ϖ discrete Fourier transforms, each of which can be performed in $O(N_\phi \log(N_\phi))$ time using the fast Fourier transform (FFT). For B^3 , this involves N_r spherical harmonic transforms. Once the angular modes are known at each radius, each mode can be re-represented in terms of radial modes by projecting onto the radial basis corresponding to that m or l . This projection is easily performed using a matrix multiplication transform, an $O(N^2)$ operation. While not as asymptotically efficient as the FFT, this transform is often not a bottleneck in 3D multidomain spectral codes, as we rarely encounter extremely high radial resolution in a single subdomain.

Note that this transformation is not invertible. By choosing a constant n_{\max} , we have adopted a triangular truncation of the radial spectrum, as opposed to a rhomboidal truncation where the number of radial modes is independent of m . A triangular truncation is commonly used for spherical harmonics as well, since this guarantees isotropic resolution over the surface of a sphere [105]. A consequence of this is that there are fewer spectral coefficients than collocation points—roughly half as many for B^2 and a quarter as many for B^3 . Thus, more functions can be represented in physical space than in spectral space; these extra functions contain power in high radial (or polar) modes for $m, l > 0$ which are resolvable by the grid but not representable by our truncation.. Transforming to and from the spectral domain effectively filters these extra modes. This has implications for elliptic problems—solving for collocation values yields an overdetermined system, while solving for spectral coefficients complicates preconditioning. See Pfeiffer [106] for ways to proceed in this situation.

Even in a pseudospectral code, where the physical space values are evolved directly, transforms to and from the spectral domain can be useful. The spectral coefficients are a valuable diagnostic tool, and some operations, such as grid refinement, are most easily

performed in spectral space. Adaptive mesh refinement and filtering are both currently implemented using spectral coefficients in SpEC.

Interpolation and differentiation

Livermore et al. provide an appendix that may be useful when computing spectral derivatives in the radial direction, but it is faster to compute them in “Fourier collocation space” (where each Fourier mode is represented by its value at each of the radial quadrature points).

Fornberg’s differentiation matrices can be used to interpolate and compute derivatives in Fourier collocation space, but since we sample on only half as many collocation points as the maximum degree of our radial polynomials, the method must be modified to account for the parity of the angular modes.¹ The abscissa to use in Fornberg’s algorithm are not $\{\omega_i\}$, but rather $\{x_i\} = \{\omega_i^2\}$, and we will use the resulting matrices to act on a function $g_m(x)$ instead of directly on $f_m(\omega)$. The definition of $g_m(x)$ depends on the parity of $f_m(\omega)$ as follows:

$$g_m(x_i) = f_m(\omega_i) \quad m \text{ even} \quad (\text{A.17})$$

$$g_m(x_i) = f_m(\omega_i)/\omega_i \quad m \text{ odd} . \quad (\text{A.18})$$

Using Fornberg matrices, we can now compute $g'_m(x)$ and $g''_m(x)$. To recover the desired quantities $f'_m(\omega)$ and $f''_m(\omega)$, we compute the following combinations:

$$f'_m(\omega_i) = 2\omega g'_m(x_i) \quad f''_m(\omega_i) = 2[g'_m(x_i) + 2\omega^2 g''_m(x_i)] \quad m \text{ even} \quad (\text{A.19})$$

$$f'_m(\omega_i) = g_m(x_i) + 2\omega^2 g'_m(x_i) \quad f''_m(\omega_i) = 2\omega[3g'_m(x_i) + 2\omega^2 g''_m(x_i)] \quad m \text{ odd} . \quad (\text{A.20})$$

One important subtlety when working with derivatives in these bases is that, if Zernike polynomials or spherical harmonics are involved, coordinate derivatives of representable

¹ Alternately, to achieve full-order accuracy using Fornberg’s method, we could provide additional points on the reflected interval $[-1, 0)$. This is possible because each Fourier mode $f_m(\omega)$ is either an even or odd function of ω (depending on whether m is even or odd).

functions are *not* themselves representable in the basis. Other quantities related to coordinate derivatives, such as Cartesian derivatives, are representable, and if the coordinate derivatives are given in collocation space, these quantities can be computed from them without trouble. However, coordinate derivatives should not be spectrally decomposed or have further coordinate derivatives taken of them.

In the case of B^2 , the implications of this are that the second radial derivative must be computed directly and not by taking a single radial derivative of the first radial derivative. The problem is easy to spot: a function represented in this basis has Fourier modes whose radial dependence matches the parity of the mode. But computing ∂_{ω} changes the parity of this radial dependence, and the result cannot be represented in the basis. In fact, the radial derivative is not well-defined at the origin, since ω does not have a well-defined direction there. One solution is to compute the quantity $\omega \partial_{\omega} f$, which can be expressed in the basis. Alternatively, we can promise to only use the collocation values of $\partial_{\omega} f$ in the process of computing Cartesian derivatives, which can also be expressed in the basis. This latter approach is what we do in SpEC.

Integration

Because Zernike polynomials are orthogonal with respect to the weight function $w(\omega) = \omega$ (for B^2), computing the definite integral of a function over the unit disk is straightforward and efficient. One property of Fourier series is that only the zero-frequency DC coefficient is needed to compute the angular integral. Our procedure is then to compute $f_0(\omega_i)$ at each radius, then to take the dot product of this with the vector of quadrature weights w_i . The integral is simply this dot product times a normalization factor. In summary,

$$\int_0^1 \int_0^{2\pi} f(\omega, \phi) \omega d\phi d\omega = \frac{2\pi}{N_{\phi}} \sum_{i=0}^{N_r-1} w_i f_0(\omega_i). \quad (\text{A.21})$$

While the DC coefficients f_0 could be extracted from the first elements of the Fourier transforms, it is more efficient to compute them directly rather than allowing the FFT

to unnecessarily compute all of the other coefficients as well. The (unnormalized) DC coefficient is simply the sum of the function samples at that radius, so

$$f_0(\omega_i) = \sum_{j=0}^{N_\phi-1} f(\omega_i, \phi_j) . \quad (\text{A.22})$$

We handle the case of B^3 analogously.

Filtering

Spectral methods can be susceptible to aliasing instabilities when, for instance, non-linear interactions allow the creation of higher spectral modes through the mixing of lower ones. Appropriate filtering of the solution is therefore required for stable evolutions [107]. When using cylindrical and spherical domains in SpEC, we have found filtering to be unnecessary in the radial direction. Filtering in angular directions, meanwhile, is performed as for spherical shells [108]. We found it sufficient to perform filtering only once every full timestep, rather than at every substep.

A.4 Scalar wave test

Two important properties of any method for solving PDEs are its stability and rate of convergence. For a spectral method, we expect exponential convergence to smooth solutions. To test my implementation of basis functions for spherical subdomains, I evolved the scalar wave equations, using a Gaussian pulse as initial data. Since the properties of the spherical harmonics used in the angular directions are already well-established, this initial data was spherically symmetric, focusing the test on the radial direction.

The scalar wave equation is also a good analogy for how SpEC evolves Einstein's equations in generalized harmonic gauge. In both cases, a second-order wave equation must be expressed in first-order form in both time and space (for work on maintaining

second order in space, see Taylor [109]). The wave equation for a scalar field $\psi(x, t)$ is

$$\partial_t^2 \psi = \nabla^2 \psi . \quad (\text{A.23})$$

Define the auxiliary fields $\pi(x, t)$ and $\phi(x, t)$ as follows:

$$\pi \equiv -\partial_t \psi \quad (\text{A.24})$$

$$\phi \equiv \nabla \psi . \quad (\text{A.25})$$

The original PDE is then equivalent to the following system of first-order equations:

$$\partial_t \psi = -\pi \quad (\text{A.26})$$

$$\partial_t \pi = -\nabla \cdot \phi \quad (\text{A.27})$$

$$\partial_t \phi = -\nabla \pi . \quad (\text{A.28})$$

In a semi-discrete view (where spatial derivatives are approximate, but time derivatives are assumed to be continuous), this reduction to first-order form introduces a “constraint” to the system; namely, that $\nabla \psi - \phi = 0$. While this is trivially true analytically by the definition in Eq. (A.25), poor initial data or errors during evolution can cause this difference to diverge from zero. In this case the constraint is marginally stable [110] and will not be discussed further, but the constraints of GRMHD require special attention in order to preserve them.

The wave equation is simple enough that analytic initial data and solutions in time are easy to derive. For my stability and convergence test, I seeded a stationary Gaussian pulse at the center of the domain and allowed it to expand. The results at several different radial resolutions are shown in Fig. A.1. The evenly-spaced lines on a logarithmic axis show that the convergence is exponential, as expected. Additionally, the evolution is stable.




Figure A.1: Convergence and stability of spherical scalar wave evolved using B3 basis functions. Evolving a spherical scalar wave with a Gaussian profile is both stable and exponentially convergent. Computational cost per timestep is less than that of I1×S2 in our implementation.

A.5 Zernike analysis

Azimuthal power

Within the space of smooth functions defined in a cylindrical volume, consider the subspace spanned by a finite number of orthonormal basis functions of the form $P_l(z)Z_m^n(\omega, \phi)$, where $P_l(z)$ is a basis for functions on a finite interval (such as Legendre polynomials) and $Z_m^n(\omega, \phi) = R_n^m(\omega)F_m(\phi)$ are the Zernike polynomials (see Sec. A.2 for notation). Any

function f in this subspace can be decomposed into spectral coefficients f_{lmn} . The amount of power in a given azimuthal mode m is defined to be

$$P_m[f] = \sum_l \sum_n |f_{lmn}|^2. \quad (\text{A.29})$$

One approach to computing this power for an arbitrary f is to compute each f_{lmn} by integrating $f(z, \phi, r)$ against the corresponding product of basis functions. If f is band-limited and the integration is of sufficiently high order, this will produce the exact result. Alternatively, f can be integrated against the set of cardinal functions along z and r . Here we show the equivalence of this nodal approach to the aforementioned modal one.

Let us denote our nodal power measurement by $Q_m[f]$:

$$Q_m[f] \equiv \sum_{i,j} \left| \iiint dz d\phi d\omega f(z, \phi, \omega) \tilde{C}_i(z) \tilde{C}_j^m(\omega) F_m(\phi) \right|^2; \quad (\text{A.30})$$

here, $\tilde{C}_i(z)$ are the normalized cardinal functions associated with $P_l(z)$ and $\tilde{C}_j^m(\omega)$ are the normalized cardinal functions associated with $R_n^m(\omega)$. Expanding those cardinal functions in terms of their associated basis functions yields

$$Q_m[f] = \sum_{i,j} \left| \iiint dz d\phi d\omega f(z, \phi, \omega) \left(\sqrt{w_i^P} \sum_l P_l(z_i) P_l(z) \right) \times \right. \\ \left. \left(\sqrt{w_j^R} \sum_n R_n^m(\omega_j) R_n^m(\omega) \right) F_m(\phi) \right|^2. \quad (\text{A.31})$$

The presence of the weights suggests that the outer sums can be interpreted as integrals (note that the corresponding integrands are products of two basis functions and therefore exactly integrable by quadrature). And since the basis functions are orthonormal, the integral of a product of sums is equal to a sum of products. This simplifies the above expression to

$$Q_m[f] = \sum_{l,n} \left| \iiint dz d\phi d\omega f(z, \phi, \omega) P_l(z) R_n^m(\omega) F_m(\phi) \right|^2. \quad (\text{A.32})$$

But the integral above is merely the projection of f onto the basis function indexed by l, m, n ; thus

$$Q_m[f] = \sum_{l,n} |f_{lmn}|^2 = P_m[f]. \quad (\text{A.33})$$

This gives us two formally equivalent ways to measure the azimuthal power in f : one involving projections onto the modal basis, the other projecting onto the nodal (cardinal) basis. The latter matches an intuitive approach to avoiding the problem of power cancellation due to phase changes at different ω and z .

Error floor

Unfortunately, when performing these integrations on a finite volume domain, the Cartesian nature of the grid results in spurious power in $m = 4, 8, \dots$ modes proportional to the error of the integration scheme (these “ambient grid modes” are also noted in studies where mode measurement is restricted to rings [65, 70]). If the function does not approach zero at the boundary of the reference cylinder, then this spurious power will be significant because of the “Lego circle” approximation to the boundary.

This effect can be mitigated by windowing the data with a smooth function that transitions between one at the center and zero at the boundary. We have achieved good results using the window

$$W(\omega) = \frac{1}{2} \{1 - \tanh [\tan (\pi (\omega + 1/2))]\} . \quad (\text{A.34})$$

The effect of the windowing on the power spectrum can then be undone via a deconvolution (made robust by using a truncated singular value decomposition). Expressing the convolution of the spectrum as

$$C_{ij} \lambda_j = \lambda'_i, \quad (\text{A.35})$$

the elements of \mathbf{C} are given by

$$C_{ij} = \int W(\omega) R_i^m(\omega) R_j^m(\omega) \omega d\omega . \quad (\text{A.36})$$

However, if the function being analyzed is entirely contained within the reference cylinder (by making its radius larger than that of the star, for instance), then this windowing technique offers minimal improvement to the error floor. Additionally, for our setup, evolved data exhibits 100× more spurious power than initial data. The net result is that, at our resolution, $m = 4$ perturbations can only be measured if they are larger than 10^{-5} relative to the background. The act of windowing does make this procedure more robust, however, should the data expand beyond the chosen reference cylinder.

B QUADRATIC BÉZIER SMOOTHING FOR PIECEWISE POLYTROPES

Even for cold matter, polytropic equations of state are an inaccurate representation of the true relationship between pressure and density. As the actual equation of state at nuclear densities is unknown, a variety of theoretical candidates have been proposed based on models of nuclear physics and constrained by both laboratory and astrophysical observations [111]. To study the impact these models have on simulation results (with the hope of then constraining them with future observations), one could implement each candidate in the simulation code or, more practically, tabulate the pressure–density relation in each case and interpolate between the samples at runtime. However, a more systematic approach to constraining this relation is to use a parameterization of the equation-of-state space where a small number of parameters can distinguish between the proposed models.

Along these lines, Lindblom suggests a thermodynamically consistent spectral representation [112], while Read et al. show that a piecewise polytropic equation of state can accurately approximate many models with a small number of parameters. To specify a piecewise polytrope, one must provide a sequence of breakpoint densities ρ_i , a sequence of polytropic exponents Γ_i for each density interval $(\rho_{i-1}, \rho_i]$, and an initial coefficient κ_0 applying to the interval $(0, \rho_0]$. Then, in each density interval, the pressure–density relation is given by

$$P(\rho) = \kappa_i \rho^{\Gamma_i}, \quad (\text{B.1})$$

where, ensuring continuity, the remaining coefficients are determined by

$$\kappa_i = \kappa_{i-1} \rho_{i-1}^{\Gamma_{i-1} - \Gamma_i}. \quad (\text{B.2})$$

This relation is only C^0 at the ρ_i , and this non-smoothness could potentially affect the convergence of spectral methods. As a simple workaround, we propose that quadratic

Bézier smoothing be used in the exponents near the breakpoints. This adds one degree of smoothness to the quantities provided by the equation of state, and the only free parameter in this construction is the width of the smoothed region.

To take advantage of smoothing, one chooses an interval around each ρ_i over which the smoothing should take place. In our implementation, we smooth over intervals of the form $(e^{x_{i-}}, e^{x_{i+}}) \equiv (\rho_i/(1+s), \rho_i(1+s))$, where s parameterizes the size of the smoothing region. The pressure corresponding to a density ρ in a smoothed region is computed as follows: First, define the auxiliary variables

$$y_{i-} \equiv \Gamma_i x_{i-} + \log(\kappa_i) \quad (\text{B.3})$$

$$y_{i+} \equiv \Gamma_{i+1} x_{i+} + \log(\kappa_{i+1}) \quad (\text{B.4})$$

$$t \equiv \frac{\log(\rho) - x_{i-}}{2 \log(1+s)} \quad (\text{B.5})$$

$$P_i \equiv \kappa_i \rho_i^{\Gamma_i}. \quad (\text{B.6})$$

Then,

$$P(\rho) = \exp \left[(1-t)^2 y_{i-} + 2t(1-t) \log(P_i) + t^2 y_{i+} \right]. \quad (\text{B.7})$$

The specific internal energy ϵ is computed by integrating the relation

$$d\epsilon = \frac{P}{\rho^2} d\rho. \quad (\text{B.8})$$

The contribution from a polytropic region can be computed analytically, with antiderivative $P/[\rho(\Gamma-1)]$, but contributions from smoothed regions must be integrated numerically. These contributions can be pre-computed and scanned¹ to speed up computations in polytropic regions, meaning that runtime integration only needs to be performed when querying points in smoothed regions (which generally represent a small fraction of total EOS queries).

Finally, the specific enthalpy is computed straightforwardly as

$$h = 1 + \epsilon + \frac{P}{\rho}. \quad (\text{B.9})$$

¹A *scan* is a cumulative sum of a sequence of numbers.

BIBLIOGRAPHY

- [1] C. D. Muhlberger, F. H. Nouri, M. D. Duez, F. Foucart, L. E. Kidder, C. D. Ott, M. A. Scheel, B. Szilágyi, and S. A. Teukolsky, (2014), arXiv:1405.2144 [astro-ph.HE] .
- [2] J. M. Weisberg and J. H. Taylor, in *Binary Radio Pulsars*, ASP Conf. Ser., Vol. 328, edited by F. A. Rasio and I. H. Stairs (2005) p. 25, astro-ph/0407149 .
- [3] The LIGO Scientific Collaboration and the Virgo Collaboration and K. Belczynski, *Class. Quant. Grav.* **27**, 173001 (2010), arXiv:1003.2480 [astro-ph.HE] .
- [4] M. B. Deaton, M. D. Duez, F. Foucart, E. O'Connor, C. D. Ott, L. E. Kidder, C. D. Muhlberger, M. A. Scheel, and B. Szilagyi, *Astrophys. J.* **776**, 47 (2013), arXiv:1304.3384 [astro-ph.HE] .
- [5] L. Rezzolla, B. Giacomazzo, L. Baiotti, J. Granot, C. Kouveliotou, and M. A. Aloy, *Astrophys. J. Lett.* **732**, L6 (2011), arXiv:1101.4298 [astro-ph.HE] .
- [6] T. W. Baumgarte and S. L. Shapiro, *Numerical Relativity: Solving Einstein's Equations on the Computer* (Cambridge University Press, 2010).
- [7] K. Uryū, F. Limousin, J. L. Friedman, E. Gourgoulhon, and M. Shibata, *Phys. Rev. D* **80**, 124004 (2009), arXiv:0908.0579 [gr-qc] .
- [8] J. S. Read, B. D. Lackey, B. J. Owen, and J. L. Friedman, *Phys. Rev. D* **79**, 124032 (2009), arXiv:0812.2163 [astro-ph] .
- [9] F. Foucart, L. E. Kidder, H. P. Pfeiffer, and S. A. Teukolsky, *Phys. Rev. D* **77**, 124051 (2008), arXiv:0804.3787 [gr-qc] .
- [10] H. P. Pfeiffer and J. W. York, Jr., *Phys. Rev. D* **67**, 044022 (2003), gr-qc/0207095 .
- [11] C. S. Kochanek, *Astrophys. J.* **398**, 234 (1992).

- [12] L. Bildsten and C. Cutler, *Astrophys. J.* **400**, 175 (1992).
- [13] E. Gourgoulhon, P. Grandclément, K. Taniguchi, J.-A. Marck, and S. Bonazzola, *Phys. Rev. D* **63**, 064029 (2001), gr-qc/0007028 .
- [14] H. P. Pfeiffer, L. E. Kidder, M. A. Scheel, and S. A. Teukolsky, *Comp. Phys. Comm.* **152**, 253 (2003), gr-qc/0202096 .
- [15] I. Kowalska, T. Bulik, K. Belczynski, M. Dominik, and D. Gondek-Rosinska, *Astron. & Astrophys.* **527**, A70 (2011), arXiv:1010.0511 [astro-ph.CO] .
- [16] H. P. Pfeiffer, D. A. Brown, L. E. Kidder, L. Lindblom, G. Lovelace, and M. A. Scheel, *Class. Quant. Grav.* **24**, S59 (2007), gr-qc/0702106 .
- [17] K. Taniguchi and E. Gourgoulhon, *Phys. Rev. D* **66**, 104019 (2002), gr-qc/0207098 .
- [18] K. Taniguchi and E. Gourgoulhon, *Phys. Rev. D* **68**, 124025 (2003), gr-qc/0309045 .
- [19] K. Taniguchi and M. Shibata, *Astrophys. J. Suppl. Ser.* **188**, 187 (2010), arXiv:1005.0958 [astro-ph.SR] .
- [20] A. Harten, P. D. Lax, and B. van Leer, *SIAM Rev.* **25**, 35 (1983).
- [21] X.-D. Liu, S. Osher, and T. Chan, *J. Comp. Phys.* **115**, 200 (1994).
- [22] G.-S. Jiang and C.-W. Shu, *J. Comp. Phys.* **126**, 202 (1996).
- [23] C.-W. Shu, *SIAM Rev.* **51**, 82 (2009).
- [24] B. van Leer, *J. Comp. Phys.* **23**, 276 (1977).
- [25] I. Hawke, F. Löffler, and A. Nerozzi, *Phys. Rev. D* **71**, 104006 (2005), arXiv:gr-

qc/0501054 .

- [26] M. D. Duez, F. Foucart, L. E. Kidder, H. P. Pfeiffer, M. A. Scheel, and S. A. Teukolsky, Phys. Rev. D **78**, 104015 (2008), arXiv:0809.0002 [gr-qc] .
- [27] F. Foucart, M. B. Deaton, M. D. Duez, L. E. Kidder, I. MacDonald, C. D. Ott, H. P. Pfeiffer, M. A. Scheel, B. Szilagyi, and S. A. Teukolsky, Phys. Rev. D **87**, 084006 (2013), arXiv:1212.4810 [gr-qc] .
- [28] L. Lindblom, M. A. Scheel, L. E. Kidder, R. Owen, and O. Rinne, Class. Quant. Grav. **23**, S447 (2006), gr-qc/0512093 .
- [29] M. Shibata and Y.-i. Sekiguchi, Phys. Rev. D **72**, 044014 (2005), arXiv:astro-ph/0507383 .
- [30] K. S. Yee, IEEE Trans. Antenn. Propag. **14**, 302 (1966).
- [31] C. R. Evans and J. F. Hawley, Astrophys. J. **332**, 659 (1988).
- [32] P. Londrillo and L. Del Zanna, J. Comp. Phys. **195**, 17 (2004), astro-ph/0310183 .
- [33] L. Del Zanna, O. Zanotti, N. Bucciantini, and P. Londrillo, Astron. & Astrophys. **473**, 11 (2007), arXiv:0704.3206 [astro-ph] .
- [34] S. C. Noble, C. F. Gammie, J. C. McKinney, and L. Del Zanna, Astrophys. J. **641**, 626 (2006), astro-ph/0512420 .
- [35] F. Foucart, *Numerical studies of black hole-neutron star binaries*, Ph.D. thesis, Cornell University (2011).
- [36] M. Galassi, J. Davies, J. Theiler, B. Gough, G. Jungman, P. Alken, M. Booth, and F. Rossi, *GNU Scientific Library Reference Manual*, 3rd ed. (Network Theory Limited, 2009).

- [37] Z. B. Etienne, Y. T. Liu, V. Paschalidis, and S. L. Shapiro, Phys. Rev. D **85**, 064029 (2012), arXiv:1112.0568 [astro-ph.HE] .
- [38] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg, J. Comp. Phys. **175**, 645 (2002).
- [39] Z. B. Etienne, V. Paschalidis, Y. Liu, and S. L. Shapiro, Phys. Rev. D **85**, 024013 (2012), arXiv:1110.4633 [astro-ph.HE] .
- [40] B. D. Farris, R. Gold, V. Paschalidis, Z. B. Etienne, and S. L. Shapiro, Phys. Rev. Lett. **109**, 221102 (2012), arXiv:1207.3354 [astro-ph.HE] .
- [41] Z. B. Etienne, V. Paschalidis, and S. L. Shapiro, Phys. Rev. D **86**, 084026 (2012), arXiv:1209.1632 [astro-ph.HE] .
- [42] V. Paschalidis, Z. B. Etienne, and S. L. Shapiro, Phys. Rev. D **88**, 021504 (2013), arXiv:1304.1805 [astro-ph.HE] .
- [43] M. Boyle, L. Lindblom, H. P. Pfeiffer, M. A. Scheel, and L. E. Kidder, Phys. Rev. D **75**, 024006 (2007), gr-qc/0609047 .
- [44] M. D. Duez, Y. T. Liu, S. L. Shapiro, and B. C. Stephens, Phys. Rev. D **72**, 024028 (2005), astro-ph/0503420 .
- [45] S. S. Komissarov, Mon. Not. R. Astron. Soc. **303**, 343 (1999).
- [46] A. Lichnerowicz, J. Math. Phys. **17**, 2135 (1976).
- [47] A. Majorana and A. M. Anile, Phys. Fluids **30**, 3045 (1987).
- [48] C. F. Gammie, J. C. McKinney, and G. Tóth, Astrophys. J. **589**, 444 (2003), astro-ph/0301509 .

- [49] H. Kreiss and J. Oliger, *Methods for the Approximate Solution of Time Dependent Problems*, Global Atmospheric Research Programme (GARP): GARP Publication Series, Vol. 10 (GARP Publication, 1973).
- [50] B. Giacomazzo and L. Rezzolla, *Class. Quant. Grav.* **24**, S235 (2007), gr-qc/0701109 .
- [51] P. Mösta, B. C. Mundim, J. A. Faber, R. Haas, S. C. Noble, T. Bode, F. Löffler, C. D. Ott, C. Reisswig, and E. Schnetter, *Class. Quant. Grav.* **31**, 015005 (2014), arXiv:1304.5544 [gr-qc] .
- [52] S. Chandrasekhar, *The Silliman Foundation Lectures, New Haven: Yale University Press*, 1969 (1969).
- [53] J. Toman, J. N. Imamura, B. K. Pickett, and R. H. Durisen, *Astrophys. J.* **497**, 370 (1998).
- [54] M. Shibata, T. W. Baumgarte, and S. L. Shapiro, *Astrophys. J.* **542**, 453 (2000), astro-ph/0005378 .
- [55] M. Saijo, M. Shibata, T. W. Baumgarte, and S. L. Shapiro, *Astrophys. J.* **548**, 919 (2001), astro-ph/0010201 .
- [56] J. M. Centrella, K. C. B. New, L. L. Lowe, and J. D. Brown, *Astrophys. J.* **550**, L193 (2001), astro-ph/0010574 .
- [57] M. Shibata, S. Karino, and Y. Eriguchi, *Mon. Not. R. Astron. Soc.* **334**, L27 (2002), gr-qc/0206002 .
- [58] M. Shibata, S. Karino, and Y. Eriguchi, *Mon. Not. R. Astron. Soc.* **343**, 619 (2003), astro-ph/0304298 .
- [59] M. Saijo, T. W. Baumgarte, and S. L. Shapiro, *Astrophys. J.* **595**, 352 (2003), astro-ph/0302436 .

- [60] C. D. Ott, S. Ou, J. E. Tohline, and A. Burrows, *Astrophys. J. Lett.* **625**, L119 (2005), astro-ph/0503187 .
- [61] S. Ou and J. E. Tohline, *Astrophys. J.* **651**, 1068 (2006), astro-ph/0604099 .
- [62] M. Saijo and S. Yoshida, *Mon. Not. R. Astron. Soc.* **368**, 1429 (2006), astro-ph/0505543 .
- [63] P. Cerdá-Durán, V. Quilis, and J. A. Font, *Comput. Phys. Commun.* **177**, 288 (2007), arXiv:0704.0356 [astro-ph] .
- [64] C. D. Ott, H. Dimmelman, A. Marek, H.-T. Janka, I. Hawke, B. Zink, and E. Schnetter, *Phys. Rev. Lett.* **98**, 261101 (2007), astro-ph/0609819 .
- [65] S. Scheidegger, T. Fischer, S. C. Whitehouse, and M. Liebendörfer, *Astron. & Astrophys.* **490**, 231 (2008), arXiv:0709.0168 [astro-ph] .
- [66] G. Corvino, L. Rezzolla, S. Bernuzzi, R. De Pietri, and B. Giacomazzo, *Class. Quant. Grav.* **27**, 114104 (2010), arXiv:1001.5281 [gr-qc] .
- [67] T. Kuroda and H. Umeda, *Astrophys. J. Suppl. Ser.* **191**, 439 (2010), arXiv:1008.1370 [astro-ph.SR] .
- [68] A. L. Watts, N. Andersson, and D. I. Jones, *Astrophys. J. Lett.* **618**, L37 (2005), astro-ph/0309554 .
- [69] J. C. B. Papaloizou and J. E. Pringle, *Mon. Not. R. Astron. Soc.* **208**, 721 (1984).
- [70] C. D. Ott, H. Dimmelman, A. Marek, H.-T. Janka, B. Zink, I. Hawke, and E. Schnetter, *Class. Quant. Grav.* **24**, S139 (2007), astro-ph/0612638 .
- [71] C. D. Ott, *Class. Quant. Grav.* **26**, 063001 (2009), arXiv:0809.0695 [astro-ph] .

- [72] K. D. Camarda, P. Anninos, P. C. Fragile, and J. A. Font, *Astrophys. J.* **707**, 1610 (2009), arXiv:0911.0670 [astro-ph.SR] .
- [73] L. Franci, R. De Pietri, K. Dionysopoulou, and L. Rezzolla, *Phys. Rev. D* **88**, 104028 (2013), arXiv:1308.3989 [gr-qc] .
- [74] W. Fu and D. Lai, *Mon. Not. R. Astron. Soc.* **413**, 2207 (2011), arXiv:1011.4887 [astro-ph.SR] .
- [75] M. Obergaulinger, P. Cerdá-Durán, E. Müller, and M. A. Aloy, *Astron. & Astrophys.* **498**, 241 (2009), arXiv:0811.1652 [astro-ph] .
- [76] L. Baiotti, R. De Pietri, G. M. Manca, and L. Rezzolla, *Phys. Rev. D* **75**, 044023 (2007), astro-ph/0609473 .
- [77] F. Douchin and P. Haensel, *Astron. & Astrophys.* **380**, 151 (2001), astro-ph/0111092 .
- [78] M. Shibata, K. Taniguchi, and K. Uryū, *Phys. Rev. D* **71**, 084021 (2005), gr-qc/0503119 .
- [79] G. B. Cook, S. L. Shapiro, and S. A. Teukolsky, *Astrophys. J.* **398**, 203 (1992).
- [80] J. De Villiers and J. F. Hawley, *Astrophys. J.* **592**, 1060 (2003), astro-ph/0303241 .
- [81] J. C. McKinney and C. F. Gammie, *Astrophys. J.* **611**, 977 (2004), astro-ph/0404512 .
- [82] B. Giacomazzo, L. Rezzolla, and L. Baiotti, *Phys. Rev. D* **83**, 044014 (2011), arXiv:1009.2468 [gr-qc] .
- [83] C. D. Ott, E. Schnetter, A. Burrows, E. Livne, E. O’Connor, and F. Löffler, *J. Phys.: Conf. Ser.* **180**, 012022 (2009), arXiv:0907.4043 [astro-ph.HE] .
- [84] D. M. Siegel, R. Ciolfi, and L. Rezzolla, *Astrophys. J. Lett.* **785**, L6 (2014),

- arXiv:1401.4544 [astro-ph.HE] .
- [85] P. Mösta, S. Richers, C. D. Ott, R. Haas, A. L. Piro, K. Boydstun, E. Abdikamalov, C. Reisswig, and E. Schnetter, *Astrophys. J. Lett.* **785**, L29 (2014), arXiv:1403.1230 [astro-ph.HE] .
 - [86] E. N. Parker, *Astrophys. J.* **145**, 811 (1966).
 - [87] R. J. Tayler, *Mon. Not. R. Astron. Soc.* **161**, 365 (1973).
 - [88] D. J. Acheson and M. P. Gibbons, *Phil. Trans. R. Soc. A.* **289**, 459 (1978).
 - [89] M. C. Begelman, *Astrophys. J.* **493**, 291 (1998), astro-ph/9708142 .
 - [90] H. C. Spruit, *Astron. & Astrophys.* **349**, 189 (1999), astro-ph/9907138 .
 - [91] K. Kiuchi, S. Yoshida, and M. Shibata, *Astron. & Astrophys.* **532**, A30 (2011), arXiv:1104.5561 [astro-ph.HE] .
 - [92] E. Pitts and R. J. Tayler, *Mon. Not. R. Astron. Soc.* **216**, 139 (1985).
 - [93] T. Foglizzo and M. Tagger, *Astron. & Astrophys.* **287**, 297 (1994), astro-ph/9403019 .
 - [94] S. A. Balbus and J. F. Hawley, *Astrophys. J.* **376**, 214 (1991).
 - [95] S. A. Balbus and J. F. Hawley, *Astrophys. J.* **400**, 610 (1992).
 - [96] S. A. Balbus and J. F. Hawley, *Rev. Mod. Phys.* **70**, 1 (1998).
 - [97] D. M. Siegel, R. Ciolfi, A. I. Harte, and L. Rezzolla, *Phys. Rev. D* **87**, 121302 (2013), arXiv:1302.4368 [gr-qc] .

- [98] P. J. Sutton, arXiv:1304.0210 [gr-qc] .
- [99] O. Rinne, L. T. Buchman, M. A. Scheel, and H. P. Pfeiffer, *Class. Quant. Grav.* **26**, 075009 (2009), arXiv:0811.3593 [gr-qc] .
- [100] D. Shoemaker, *LIGO-T0900288-v3: Advanced LIGO anticipated sensitivity curves*, Tech. Rep. (LIGO, 2010).
- [101] R. Courant, K. Friedrichs, and H. Lewy, *IBM Journal of Research and Development* **11**, 215 (1967).
- [102] T. Matsushima and P. S. Marcus, *J. Comp. Phys.* **120**, 365 (1995).
- [103] W. T. M. Verkley, *J. Comp. Phys.* **136**, 100 (1997).
- [104] P. W. Livermore, C. A. Jones, and S. J. Worland, *J. Comp. Phys.* **227**, 1209 (2007).
- [105] J. P. Boyd, *Chebyshev and Fourier Spectral Methods*, Dover books on mathematics (Dover Publications, 2001).
- [106] H. P. Pfeiffer, *Initial data for black hole evolutions*, Ph.D. thesis, Cornell University (2003), gr-qc/0510016 .
- [107] B. Szilágyi, L. Lindblom, and M. A. Scheel, *Phys. Rev. D* **80**, 124010 (2009), arXiv:0909.3557 [gr-qc] .
- [108] L. E. Kidder, L. Lindblom, M. A. Scheel, L. T. Buchman, and H. P. Pfeiffer, *Phys. Rev. D* **71**, 064020 (2005), gr-qc/0412116 .
- [109] N. W. Taylor, *Numerical simulations of black hole binaries: Second order spectral methods*, Ph.D. thesis, Cornell University (2010).
- [110] R. Owen, *Phys. Rev. D* **76**, 044019 (2007), gr-qc/0703145 .

[111] J. M. Lattimer, *Annu. Rev. Nucl. Part. Sci.* **62**, 485 (2012), arXiv:1305.3510 [nucl-th] .

[112] L. Lindblom, *Phys. Rev. D* **82**, 103011 (2010), arXiv:1009.0738 [astro-ph.HE] .