From the holdings of the Department of Special Collections , Princeton University Library, Princeton, NJ

These documents are either in the public domain, or can only be used for educational and research purposes (Fair use) as per U.S. Copyright law (text below). By accessing this file, all users agree that their use falls in the public domain or within fair use as defined by the copyrightlaw. We do not charge any permission or usage fees for the publication of images of material inour collections, including those provided by Princeton University Library, via our website, catalog, or directly from staff.For more information regarding copyright, credit, and citation please review the information onthe Copyright, Credit and Citation Guidelines page of our website.
https://library.princeton.edu/special-collections/policies/copyright-credit-and-citation-guidelines. If you have any questions, including how to order higher-quality images from the collection, contact Special Collections Public Services via our Ask Us! Form. https://library.princeton.edu/special-collections/ask-us

Doctoral Dissertations, 2011-2026 Astrophysical Sciences

2016

Development and Application of Numerical Techniques for General-Relativistic Magnetohydrodynamics Simulations of Black Hole Accretion

White, Christopher

Princeton, NJ : Princeton University

http://arks.princeton.edu/ark:/88435/dsp01n296x160r

U.S. Copyright law: The copyright law of the United States (Title 17, §108, United States Code) governs the making of photocopies or other reproductions of copyrighted material. Under certain conditions specified in the law, libraries and archives are authorized to furnish a photocopy or other reproduction. One of these specified conditions is that the photocopy or reproduction is not to be used for any purpose other than private study, scholarship, or research. If a user makes a request for, or later uses, a photocopy or reproduction for purposes in excess of fair use that user may be liable for copyright infringement. This institution reserves the right to refuse to accept a copying order if, in its judgment, fulfillment of the order would involve a violation of copyright law. https://www.copyright.gov/title17/title17.pdf

Development and Application of Numerical Techniques for General-Relativistic Magnetohydrodynamics Simulations of Black Hole Accretion

Christopher Joseph White

A Dissertation

Presented to the Faculty

of Princeton University

in Candidacy for the Degree

of Doctor of Philosophy

Recommended for Acceptance

by the Department of

Astrophysical Sciences

Advisor: James M. Stone

September 2016

© Copyright by Christopher Joseph White, 2016. All rights reserved.

Abstract

We describe the implementation of sophisticated numerical techniques for general-relativistic magnetohydrodynamics simulations in the Athena++ code framework. Improvements over many existing codes include the use of advanced Riemann solvers and of staggered-mesh constrained transport. Combined with considerations for computational performance and parallel scalability, these allow us to investigate black hole accretion flows with unprecedented accuracy. The capability of the code is demonstrated by exploring magnetically arrested disks.

Acknowledgments

The success of a thesis project depends in no small way on the support and guidance received along the journey. From my parents I have received many years of encouragement, dating all the way back to the first telescope had as a child and the mindset of scientific inquiry it inspired. A number of friends have also helped me pursue my academic aspirations, at some times by listening to my ideas and at others by reminding me that taking breaks can be a good thing.

Many fruitful conversations have contributed to my knowledge and understanding in this particular field, especially at Princeton and amongst the members of the Horizon Collaboration theoretical and computational astrophysics network. In particular, Charles Gammie has provided invaluable insight into the development of general-relativistic numerical codes.

Finally, I would like to thank my advisor, Jim Stone. Three years ago he suggested that the time was right for numerical advances in relativity, fluid dynamics, and other fields of astrophysics to be applied to simulations, that black hole accretion in particular stood to gain from better computation and theory. My experience since then has shown me that this assessment was correct. This proved to be no simple undertaking – and indeed I would have it no other way – but Jim has ever been available to assist and counsel me.

Preface

There are many factors that motivate the writing of a doctoral dissertation, not the least of which is the requirement of doing so in order to obtain an advanced degree.

That said, one of the benefits of the dissertation format is the ability to expand upon ideas and to give details often omitted in other venues. This work was written with the intention of providing a clear and thorough reference for the research at hand, including not just the results but the motivation behind the project as well as a log of useful insights gained while executing it. The hope is that in the future a bright young graduate student can use this document to quickly learn enough to replicate the results. Furthermore, this serves to document the relativistic portions of the Athena++ code, whose development was a core component of this author's studies.

Especially in numerics, all too often there are critical tidbits left unspoken – ordering of computations, magic numbers found by trial and error, conditions under which the code is not expected to work properly. Almost assuredly there are some points that have been forgotten and useful facts that have been omitted from this work. Still, perhaps enough information is presented here to eliminate much of the frustration our hypothetical student might have in learning the subject. Patience is asked of established researchers reading this document; the verbosity is intentional.

Versions of some of this writing can be found elsewhere, often in a more condensed format. In particular, part of Chapter 1 and much of Chapters 2 and 3 and Appendix A can be found in White et al. (2016, in press). An extension of Chapter 4 is also forthcoming.

As the work here is very much done in a relativistic setting, factors of G and c have been omitted from most equations, retained only in cases where they lend clarity to discussion involving conventional units. The $(-,+,+,+)$ signature is used throughout. Greek indices

are understood to run over all spacetime components of tensors, while Latin indices run over only spatial components.

Contents

Abstract iii

Acknowledgments iv

Preface v

1 Introduction 1

1.1 Black Holes and Their Importance to Astrophysics 1

1.1.1 Stellar-Mass Black Holes 2

1.1.2 Supermassive Black Holes 3

1.1.3 Other Black Holes 4

1.2 Astrophysical Accretion 4

1.3 Numerical Considerations 7

1.3.1 Riemann Solvers 8

1.3.2 The Divergence-Free Constraint 10

1.4 Existing Codes 11

1.5 Developing a New Code 12

1.5.1 Riemann Solvers and Constrained Transport 13

1.5.2 Mesh Refinement 14

1.5.3 Polar Boundary Conditions 15

1.5.4 Speed and Scalability 17

2 Numerical Solution of GRMHD 21

2.1 Differential Forms of Equations 21

2.1.1 General Equations . . . 21
2.1.2 Rewriting Covariant Derivatives . . . 23
2.1.3 The Equations in Terms of Partial Derivatives . . . 24
2.2 Integral Forms of Equations . . . 26
2.2.1 General Equations . . . 26
2.2.2 Discretized Forms of Equations . . . 28
2.2.3 Magnetic Evolution . . . 32
2.3 Numerical Algorithm . . . 33
2.3.1 Variable Inversion . . . 34
2.3.2 Reconstruction . . . 36
2.3.3 Frame Transformation . . . 37
2.3.4 Riemann solver . . . 40
2.3.5 Inverse Frame Transformation . . . 42
2.3.6 Constrained Transport . . . 42
2.3.7 Source Terms . . . 46
2.3.8 Time Integrator . . . 46
2.3.9 Storage Requirements . . . 47

3 Testing Athena++ 49

3.1 Linear Waves . . . 50
3.2 Shock Tubes . . . 53
3.3 Spherical Blast Wave . . . 57
3.4 Bondi Accretion . . . 59
3.5 Magnetized Equatorial Inflow . . . 63
3.6 Torus Simulations . . . 63

4 Application to MAD 71

4.1 The Importance and Physics of MAD . . . 72
4.2 The Current State of Simulations . . . 73

4.3 MAD Simulations in Athena++ 74
4.3.1 Grid Parameters 74
4.3.2 Torus Parameters 75
4.4 Results of New MAD Simulations 79
4.4.1 Primary Results at Different Resolutions 80
4.4.2 Variations Near the Poles 83
4.4.3 Very High Resolution 87
4.5 Interpretation of Our Results 96

5 Conclusion **102**
5.1 Implications of MAD Studies 102
5.2 Future Scientific Investigations 104
5.3 Roadmap for Future Code Development 105
5.3.1 Algorithmic Improvements and Extensions 105
5.3.2 New Physics 106

A Frame Transformation **108**

B Wavespeeds in SRMHD **113**
B.1 Vanishing Velocity 114
B.2 Vanishing Longitudinal Magnetic Field 115
B.3 General Case 115

C Coordinate System Examples **119**
C.1 Minkowski Spacetime 119
C.1.1 Minkowski Coordinates 120
C.1.2 Spherical Coordinates 122
C.1.3 Sinusoidal Coordinates 124
C.1.4 Tilted Coordinates 127
C.2 Schwarzschild Spacetime 129

C.3 Kerr Spacetime 132
C.3.1 Boyer–Lindquist Coordinates 132
C.3.2 Kerr–Schild Coordinates 137

References **145**

Chapter 1

Introduction to the Study of Black Hole Accretion

When considering black holes, an overabundance of possibilities for research avenues presents itself. Without focus there would be too much material for a single work to do the subject justice. We choose then to consider black holes in the broader context of astrophysical phenomena. In particular, we will study accretion of matter onto them, as this is a line of investigation rich with both theoretical considerations and observational constraints.

1.1 Black Holes and Their Importance to Astrophysics

Black hole solutions to the equations of general relativity (GR) were first recognized shortly after the formulation of the theory a century ago (Schwarzschild, 1916). However, it was not until 1958 that these solutions were appreciated as describing actual physical phenomena (Finkelstein). Moreover, the crucial ingredient of spin was only introduced in 1963 (Kerr).

Since that time, however, black holes have been observed throughout the cosmos. Ranging from several to billions of times the mass of the Sun, they can be found scattered throughout our own galaxy and indeed in galaxies out to the far side of the observable universe.

1.1.1 Stellar-Mass Black Holes

The lighter, stellar-mass black holes are now recognized as one of the common endpoints of stellar evolution (Woosley and Weaver, 1995). The collapse of the iron cores of massive stars in some cases cannot be halted at nuclear densities, the pressure forces of the matter proving insufficient to overcome gravity.

Even were these black holes to do nothing else, the processes around their formation alone would justify considerable study. They are born not in the vacuum of space but rather inside a star. The large amount of surrounding matter is susceptible to falling into the black hole, as indeed some of it will. The small size of the hole, however, means the surrounding material will generally have too much angular momentum to fall inward, necessitating transport of angular momentum outward – a common theme that will recur in this work. Indeed some material will accrete while the rest will be ejected, all taking place within the violent death throes of a star that is both collapsing and exploding.

The accretion onto the black hole itself is thought to power the jets we observe as long-duration gamma-ray bursts (GRBs) (Woosley, 1993). In this case, modeling GRBs requires understanding both supernovae and general-relativistic accretion flows.

Once the progenitor star no longer exists, its matter either being consumed or thrown off to infinity, it is possible that a stellar-mass black hole will enter a quiescent state. With no accreting matter to emit radiation, such a black hole would be the very definition of invisible. In this case, the only way to detect such objects would be through microlensing of passing starlight (e.g. the MACHO Project, Bennett et al., 2002).

In fact, though, many stellar-mass black holes do not find themselves in such complete isolation. Stars often come in pairs and so there is often a stellar companion nearby. Though a black hole's gravity is no stronger than a star of equivalent mass, it can still find itself drawing material off the companion star. This infalling matter can form an accretion disk, and indeed such disks have been observed in what are known as X-ray binaries. This can happen either as it captures particles blown off the star in winds (high-mass X-ray binaries), or as the star evolves, swells, and overflows its Roche lobe (low-mass X-ray binaries). These

systems provide ideal laboratories to study black hole accretion unobscured by surrounding matter as is found in GRBs and often with supermassive black holes.

1.1.2 Supermassive Black Holes

Dynamical observations of stars at the center of our galaxy (Schödel et al., 2002; Gillessen et al., 2009) have shown that a very large mass, $4 \times 10^6\ M_\odot$, must be confined to a very small region, 2×10^{15} cm. While such constraints could technically be satisfied by diffuse matter that has not collapsed to a black hole, or by a large collection of stars and other condensed objects, these alternate arrangements could only exist for short times in contrived scenarios before self-gravity collapses them. Moreover, there is a distinct lack of the visible signatures one would expect from large amounts of ordinary matter. Thus the community has concluded there must be a supermassive black hole of lying at the center of the galaxy.

Our galaxy is not unique in this regard. It is now believed that nearly all major galaxies host supermassive black holes ranging in mass from approximately $10^5\ M_\odot$ to $10^9\ M_\odot$ (Ho, 2008). The most energetic systems are known as active galactic nuclei (AGNs).

The very features that allow us to observe such distant black holes – emission from hot accretion disks and relativistic jets – hint at an important role these objects play in the evolution of galaxies. Accreting black holes emit large amounts of energy and momentum, affecting the galaxy as a whole. Ultimately, black holes provide a mechanism for extracting work from the enormous reservoir of gravitational potential energy in a relatively loosely bound system such as a galaxy (compared to the extreme of tight binding as in a black hole). That is, having a small amount of mass fall all the way to the horizon can release enough energy and momentum to unbind a significantly larger mass from the galactic potential.

This black hole feedback is now recognized as an important part of the dynamics and evolution of many galaxies (Fabian, 2012).

1.1.3 Other Black Holes

Black holes of other mass ranges are discussed in the literature. For example, there is speculation that intermediate-mass black holes of hundreds to thousands of solar masses could exist (Miller and Colbert, 2004). If they do occur, there is no reason to expect their accretion processes could not be modeled with the same techniques developed for lesser and greater masses. Formed from either direct collapse, black hole mergers, or accretion in dense star clusters or near supermassive black holes, such black holes even have claimed observations. However, the community has yet to reach a consensus on the interpretation of these observations or even on whether such black holes are prevalent in nature.

Another mass range of interest are very small black holes with masses less than that of the Moon. Any such black hole would presently be losing mass to Hawking radiation faster than it gains mass from cosmic microwave background photons. In fact, black holes formed in the very early universe with masses of around 10^{15} g would be undergoing the final stages of evaporation in the present epoch (MacGibbon et al., 2008), leading to potentially observable outcomes at the ends of their lives. To date, no such phenomena have been unambiguously observed.

In the interest of connecting to observations, and recognizing that there is no shortage of science to be done with the commonly observed black holes, we restrict our attention to stellar-mass and supermassive black holes in the present work.

1.2 Astrophysical Accretion

Consider matter falling onto or orbiting around a black hole. Momentum transfer and/or initial conditions often work to make the angular momentum directions aligned between different parts of the flow. This results in a disk, whether it be thick or thin, at least whenever the infalling matter's density is large enough to have a large number of collisions between particles over the course of an orbit. We will be working in this regime, where the continuum fluid approximation holds.

As already mentioned, a generic feature of systems in which matter falls into black holes is that the infalling matter initially has too much angular momentum. Despite the effectiveness of angular momentum transfer in creating a disk, once such a disk is formed transfer via intermolecular viscous forces is far too slow to account for observed accretion rates. This is a general attribute of astrophysical accretion.

Despite the lack of a clear mechanism to transport angular momentum outward, models were developed based on the assumption that such a mechanism would be found. The most famous of these is the α-disk of Shakura and Sunyaev (1973). In that work they argue that in both turbulent and magnetic phenomena the $r\phi$-component of the stress tensor (i.e. the radial flux of angular momentum) is proportional to $\rho c_{\rm s}^2$, where $c_{\rm s}$ is the sound speed. The proportionality is effectively the turbulent Mach number in the former case and the square of the Alfvén Mach number in the latter. Calling this proportionality constant α and assuming it to be constant in a disk allows one to investigate models that capture the important features of the flow and allow for a connection to observations. Still, an ansatz is not a mechanism, and the explanation for the emergence of an effective viscosity took a number of years.

An extensive effort comprising both analytical and numerical work was dedicated to resolving this issue. A thorough review can be found in Balbus and Hawley (1998). We present only the briefest summary here. The ability of turbulence to transport angular momentum (though not guaranteed in all situations) was recognized early, and so many investigations focused on demonstrating the inevitability of turbulence in astrophysical disks. Extremely high Reynolds numbers combined with the shear inherent in Keplerian motion would seem to promote turbulence. However as discussed in Balbus et al. (1996) the opposite signs of the angular velocity and angular momentum gradients actually lead to linear and nonlinear stability against turbulence. The idea of convection driving turbulence was also considered, but as shown in Ryu and Goodman (1992) and verified in simulations this has the tendency to transport angular inward rather than outward.

Attention was also given to global features of accretion disks, those that rely on the

curvature and periodicity of the domain rather than just local shear, centrifugal, and Coriolis effects. The general results, though, were that such mechanisms do not solve the transport issue. The nonaxisymmetric instability of thick disks (Papaloizou and Pringle, 1984) was found to saturate without turbulence, and a tidally driven instability (Goodman, 1993) also did not show clear signs of the desired angular momentum transport. Another tidal effect, that of spiral shocks, was however shown to lead to angular momentum transport, though this mechanism only applies when there is a massive companion as in a binary system.

The solution to the transport problem was discovered by Balbus and Hawley (1991). It is not purely hydrodynamical in nature but rather relies on magnetic fields. Weak fields generically lead to a local instability, the magnetorotational instability (MRI). This results in turbulent flow, allowing effective angular momentum transport in disks.

The MRI can be pictured by imagining a vertical magnetic field line threading a disk. If the fluid containing the line at the disk midplane is perturbed outward, it will find itself rotating with a decreased angular velocity (by assumption; this instability depends on such an angular velocity profile). In ideal magnetohydrodynamics (MHD), flux freezing means we can consider the field line to be dragged with the fluid. The result is that the outwardly displaced part of the field line falls behind in rotation, leading to a magnetic tension that pulls the outer element forward in azimuth and the inner elements backward. This radial flux of azimuthal momentum is exactly the transport desired.

As discussed by Balbus and Hawley, the MRI is linearly unstable: small perturbations grow exponentially. Moreover, it relies on a weak rather than strong field, and the above argument generalizes to other orientations of the field. As a result one generically expects it to manifest in nearly any ionized plasma in a roughly Keplerian disk.

Since its discovery, the MRI has been modeled in a large number of simulations, both local shearing boxes and global disks. It is found to lead to a turbulent state in which angular momentum is transported outward, driving accretion. The MRI eventually saturates in disks, leading to an effective viscosity.

Though initially investigated in the Newtonian regime, the general principles behind

the MRI hold in a general-relativistic setting (Gammie, 2004). As a result, it becomes critical that electromagnetic fields be incorporated into our models. With a proper, relativistic treatment of hydrodynamics and electromagnetism, a diverse set of complex accretion phenomena can be understood.

Accretion in the regime we are considering is fundamentally a fluid dynamics problem (albeit one on a curved manifold). Moreover, we expect the systems we are considering to become turbulent. As such, purely analytic modeling can only take us so far, and numeric simulations must be utilized.

We note that for the purposes of studying black hole accretion, it is convenient that ideal systems are scale invariant. Everything scales with the mass of the central object, so for instance a given set of predictions or simulations for a black hole of mass M holds with the understanding that the units of length are GM/c^2 and those of time are GM/c^3, regardless of whether $M = 1\ M_\odot$ or $M = 10^6\ M_\odot$. As a result, the theoretical and numerical study of supermassive black hole accretion is by no means disjoint from that of accretion onto smaller black holes. This freedom to scale results only breaks down if other assumptions are violated in the rescaling, such as if the accreting material becomes nonnegligibly self-gravitating, or if the density becomes low enough that the continuum fluid approximation no longer applies.

1.3 Numerical Considerations

The complex, nonlinear dynamics of most astrophysical fluids make numerical simulations a key tool in understanding them, even in flat spacetime. A number of interesting phenomena, however, occur in regions of strong enough gravity that general relativity cannot be neglected, requiring numerical algorithms capable of evolving fluids in curved spacetimes. Examples include accretion onto black holes (Abramowicz and Fragile, 2013), collapsar models of long-duration gamma-ray bursts (Woosley, 1993), and merging neutron star binaries (Faber and Rasio, 2012), among others.

Some of the most successful numerical algorithms for solving the equations of com-

pressible fluid dynamics are finite-volume methods, primarily due to their superior accuracy and stability for shock capturing (van Leer, 1979). In such methods the domain is partitioned into discrete cells and the quantities stored and evolved are cell volume averages of the conserved quantities. In Godunov schemes, fluxes of conserved quantities are determined by solving Riemann problems at each interface. In particular, one considers the interface to be a plane separating two spatially constant fluids filling all space, and the solutions of the Riemann problem determine the time-averaged fluxes across the plane in this simplified problem. The overall scheme is described in §2.3.

Two key algorithmic components of a finite-volume Godunov scheme for MHD are the Riemann solver and the method by which the divergence-free constraint on the magnetic field is enforced. We now examine each of these in turn, reviewing the options that have been developed for implementing them.

1.3.1 Riemann Solvers

The accuracy of a simulation depends critically on the accuracy of the Riemann solver adopted. Methods that have been developed to solve the Riemann problem include:

1. Central. As described in Kurganov and Tadmor (2000), one can neglect the internal wave structure in the Riemann fan (equivalently the eigenstructure of the conservation equations) entirely, using only the conserved states, their corresponding fluxes, and possibly maximal wavespeeds. Conceptually central schemes proceed from the state at one time level to the state at the next without necessarily finding fluxes through interfaces; in this sense they are not true Riemann solvers. Gradients of both conserved quantities and fluxes are used to compute time evolution, but these gradients are taken componentwise, allowing the scheme to be independent of the (usually highly nonlinear) relationship between components.

2. Roe. Developed in Roe (1981) and improved in Harten and Hyman (1983), this is an exact solver in that it finds exact solutions to a particular linearized form of the

equations. Accuracy comes at the cost of computational complexity, as the fluxes are determined by solving an eigenvalue problem for a 7-dimensional system (in MHD).

3. HLL (Harten, Lax, and van Leer). These approximate methods are built upon the foundation laid by Harten et al. (1983). They assume a certain number of waves propagate from discontinuities, neglecting some of the waves. They then base the fluxes on which waves are bounding the region of interest, i.e. which are the slowest leftgoing and rightgoing waves. A number of different HLL Riemann solvers have been developed for relativistic fluids, including the following:

   (a) LLF (local Lax–Friedrichs). The largest-in-magnitude eigenvalue of the Roe matrix, i.e. the fastest linear wavespeed, is taken to be the signal speed on both sides of the interface. This solver is computationally inexpensive even in relativity, as it has only a single intermediate state and neglects the physical nature of the left and right states. This comes at the cost of being diffusive for flows slower than the extremal wavespeed (see §3.1).

   (b) HLLE (Harten–Lax–van Leer+Einfeldt). This solver is nearly identical to LLF, having just a single intermediate state, with the only difference being that the two signal speeds are allowed to be different. One uses the speed of the fastest leftgoing wave, considering both sides of the interface, as well as the fastest rightgoing wave, as suggested by Einfeldt (1988). In practice HLLE and LLF give very similar results and have only small performance differences.

   (c) HLLC (HLL+contact). By resolving not only the extremal waves but also the contact discontinuity in the Riemann fan, one can get better results. Such an HLLC Riemann solver is developed for special relativity in Mignone and Bodo (2005), and we consider that algorithm for hydrodynamics. Mignone and Bodo (2006) also develop an HLLC solver for relativistic MHD, though we forgo implementing this in favor of HLLD.

   (d) HLLD (HLL+discontinuities). This solver is developed for relativistic MHD

by Mignone et al. (2009). Given the extremal (fast magnetosonic) wavespeeds, it resolves not only the contact but also the Alfvén waves. It improves the accuracy of the fluxes by lessening the numerical diffusion, though at the expense of requiring nonlinear root finds. In particular, one uses the secant method to find the total pressure across the contact. In the case of a strong longitudinal magnetic field, the conserved HLLE state is constructed and inverted using a Newton–Raphson solver in order to find a total pressure to initialize the secant method.

Exact Riemann solvers are computationally expensive, especially in general-relativistic MHD (GRMHD). Most of the codes written to date employ simple solvers such as LLF or HLLE. However, the simplest approximate solvers tend to be much more diffusive than their more exact counterparts for subsonic flows. This has led to our decision to work to implement HLLC and HLLD in general relativity, as described in §2.3, following the suggestion by Pons et al. (1998) and Antón et al. (2006).

1.3.2 The Divergence-Free Constraint

In the case of MHD, enforcing the divergence-free constraint is important to prevent spurious production of magnetic monopoles. A number of algorithms have been developed to this end, including the following:

1. Divergence cleaning. Auxiliary equations are introduced in order to damp and/or advect to the boundaries any monopoles that are developed. These equations can be elliptic (Brackbill and Barnes, 1980; Ramshaw, 1983), parabolic (Marder, 1987), or hyperbolic (Dedner et al., 2002), with only the last naturally able to respect causality in relativistic settings.

2. Vector potential evolution. The vector potential components A^i are evolved instead of the magnetic field components B^i. This naturally obeys the divergence-free constraint if the vector potential is appropriately staggered, though one must be careful to choose

an appropriate gauge, especially when including mesh refinement (Etienne et al., 2012).

3. Constrained transport (CT). The magnetic field is staggered in such a way as to maintain a specific discretized version of the constraint. Evans and Hawley (1988) introduce staggered-mesh CT in a general-relativistic context.

4. Flux-CT. Tóth (2000) discusses how to alter the fluxes so as to preserve the divergence-free constraint on a particular stencil of cell-centered magnetic fields.

We choose to employ the Evans and Hawley method. The details of the algorithm are developed for Cartesian grids by Gardiner and Stone (2005), where only magnetic fields (and not velocities as in the original formulation) are located at interfaces between cells. Here we extend the Gardiner and Stone algorithm, used with great success in special relativity (Stone and Gardiner, 2008; Beckwith and Stone, 2011), to arbitrary stationary coordinate systems.

1.4 Existing Codes

A number of Godunov scheme codes have been developed with GRMHD capabilities on a stationary background spacetime, including Harm (Gammie et al., 2003), Komissarov (2004), the Valencia group's code (Antón et al., 2006), and ECHO (Del Zanna et al., 2007). In addition, some codes combine the GRMHD and Einstein equations, evolving a self-gravitating magnetized fluid. These numerical relativity (NR) codes include the Tokyo/Kyoto group's code (Shibata and Sekiguchi, 2005), Anderson et al. (2006), WhiskyMHD (Giacomazzo and Rezzolla, 2007), GRHydro (Mösta et al., 2014), SpEC (Muhlberger et al., 2014), and IllinoisGRMHD (Etienne et al., 2015).

Of these ten codes, only Komissarov (2004) and Antón et al. (2006) describe the use of Roe-type Riemann solvers (with the latter making the point that they can in principle use any Riemann solver), and only Shibata and Sekiguchi (2005) and Antón et al. (2006) describe the use of central, non-Godunov schemes. All others restrict themselves to HLLE and LLF.

Table 1.1: Properties of existing codes.

Code	NR?	Flux Scheme	Field Scheme
Harm	No	LLF, HLLE	Flux-CT
Komissarov	No	Roe	Staggered CT
Tokyo/Kyoto	Yes	Central	Staggered CT
Valencia	No	Central, Roe, HLLE	Staggered CT
Anderson	Yes	LLF, HLLE	Hyperbolic divergence cleaning
WhiskyMHD	Yes	HLLE	Staggered CT
ECHO	No	HLLE	Staggered CT
GRHydro	Yes	HLLE	Flux-CT
SpEC	Yes	HLLE	Staggered CT
IllinoisGRMHD	Yes	HLLE	Staggered potential

Of the aforementioned ten codes, Anderson et al. (2006) uses hyperbolic divergence cleaning, IllinoisGRMHD uses a staggered vector potential, Harm and GRHydro use flux-CT, and the remaining six use a staggered CT scheme.

The properties of these codes are summarized in Table 1.1. While we are not concerned with NR here, we highlight its presence to indicate the focus of a code.

1.5 Developing a New Code

The choice to develop a new code – even one that builds upon an established code – should not be made lightly. Often the most efficient way forward is to use an existing tool. However, we feel there are sufficient ways in which existing GRMHD codes are not ideally suited for our purposes that it becomes worthwhile to invest the effort in writing a new code. Advanced features are missing from some codes, and it should be noted that a number are not even public.

Three key considerations stand out as important for guiding the development of our code. These are *accuracy*, *performance*, and *modularity*. While any numerical procedure will have its approximations, and even the underlying physical models will be idealized in various ways, we nonetheless strive for as accurate a representation of nature as we can given the computational resources we have available. One simple way most simulations can gain more accuracy, at least to a certain extent, is to increase the resolution at which they model a scenario. This however leads to a greater computational cost, and so we are led to

performance as another guiding principle. Moreover, performance can help codes simulate systems for longer times, allowing more complex behavior to emerge. Finally, recognizing that new algorithms and models are continuously being developed, we want our code to be modular enough to be able to adopt newly developed methods. These can be both numerical insights and more detailed physics.

The new Athena++ code framework is the result of these considerations. It employs finite-volume Godunov methods to evolve the GRMHD equations on any stationary spacetime. With Athena++ in general and GRMHD in particular, we have applied our guiding principles to achieve what we consider to be valuable features. We will highlight several of these in the following sections, after which we will summarize what rewards we anticipate will follow.

1.5.1 Riemann Solvers and Constrained Transport

We have already detailed the choices of Riemann solver (§1.3.1) and method for updating the magnetic field (§1.3.2). Both considerations directly relate to the *accuracy* of the solution, minimizing numerical diffusion and maintaining the antisymmetric nature of the electromagnetic field tensor respectively. The adoption of approximate Riemann solvers also helps *performance*.

For constrained transport, there is little we can vary once the general method is chosen. For example, we do not support flux-CT or divergence cleaning schemes, as these require a fundamentally different (non-staggered) grid for the magnetic field. For Riemann solvers, however, a degree of *modularity* can be achieved. In a Godunov scheme, the solver has the well-defined role of taking a left state and a right state and returning the associated fluxes. The rest of the code does not care about the internal working of the solver, and so we are free to implement a variety of solvers. LLF (both frame-transforming and not), HLLE (also both frame-transforming and not), HLLC (for hydrodynamics), and HLLD (for MHD) are all implemented in Athena++, allowing one to choose on a per-simulation basis the method best suited for the task at hand.

1.5.2 Mesh Refinement

In many simulations the resolution requirements are not constant in space, and sometimes they even vary in time. For grids at a single, fixed resolution this leads to a choice of underresolving features of interest (harming *accuracy*) or overcomputing unnecessary values (harming *performance*). For example, when studying an accretion disk – especially a thin disk as occurs with radiative cooling – one often does not care about the near-vacuum near the poles, but one would like to see details near the midplane of the disk.

One solution is to allow for mesh refinement. Athena++ has such a scheme available, to be described more fully in a forthcoming general code method paper. Here we summarize the salient features:

1. The grid is already divided into blocks of equal logical size and shape (i.e. the numbers of cells in each direction are fixed constants) as part of domain decomposition for parallel computing.

2. Any block can be subdivided into 8 blocks of the same logical size (4 in two dimensions) recursively.

3. Neighboring blocks are restricted to differ by no more than one refinement level. Any blocks that share a face, an edge, or a corner are considered neighbors.

4. Once a block is refined, it is no longer simulated at the coarser level.

5. Prolongation and restriction operators translate information between levels. They are applied to populate ghost zones whose corresponding active zones are at a different level.

This scheme works with *modularity* in that all choices of algorithm and physics relevant to this study are compatible with such refinement. At present all refinement levels proceed in lockstep; in the future we intend to allow for hierarchical timestepping. For all simulations in this work static as opposed to adaptive mesh refinement has proved sufficient.

One particularly useful feature of mesh refinement presents itself in polar grids where the emphasis is on the equatorial regions. With polar grids the timestep is often limited by the azimuthal thickness of the cells closest to both the pole and the inner boundary. This limit in fact scales as the square of the linear resolution: doubling the number of azimuthal cells makes each cell half as wide, while doubling the number of polar cells moves the innermost cell centers to approximately half their distance from the pole (under the small angle approximation). (The radial coordinate of the cell centers also decreases, but because we only go as far as the horizon and not all the way to the coordinate origin, this effect is negligible.) Compare this limit to that of Cartesian grids, where the timestep decreases linearly with the linear resolution.

Without refinement, doubling the linear resolution of a polar simulation results in 32 times the computational cost, with at most a factor of 8 able to be weak-scaled away by using more processors in parallel (see §1.5.4). Thus one generally expects a factor of 4 increase in wall time for every factor of 2 in resolution when running polar simulations. With refinement, however, one can refine only areas away from the pole. In practice even multiple levels of refinement of such regions will leave the cells wider than the original cells near the pole, meaning the timestep does not decrease at all. That is, the computational cost only scales by a factor of about 8, all of which can be handled with more processors with no decrease in wall time.

1.5.3 Polar Boundary Conditions

As much of our interest is in systems well represented by polar coordinates, the coordinate singularity at the poles is an issue that must be confronted. The simple way of handling the singularity, found in a number of existing codes, is to excise a small cone around the pole from the domain, placing reflecting walls at the conical boundaries. Even in the limit that the cone opening angle vanishes, this leads to an unphysical boundary that impedes the flow of material and electromagnetic fields across the pole, harming *accuracy*.

Our solution is to connect cells across the pole much as we connect cells in periodic

domains. That is, we avoid having a domain boundary at the pole. There are four considerations that allow us to implement this feature:

1. For cells at the edge of a block along a pole, their ghost zones must match cells across the pole. In particular, a cell with azimuthal extent $\phi_- < \phi < \phi_+$ should have as its neighbor in the polar direction (θ decreasing through 0 or increasing through π) the cell with the same radial and polar extent but covering $\phi_- + \pi < \phi < \phi_+ + \pi$ (module 2π as necessary). These ghost zones are needed for determining left and right states in radial Riemann problems. Importantly, there are Riemann problems affected by these ghost zones that are not lying on the pole itself, and so their fluxes cannot be neglected.

2. Vector quantities must be appropriately reflected when copied from an active zone to a ghost zone belonging to a block on the other side of a pole. This is straightforward and little different from implementing a hard reflecting wall boundary condition.

3. Vanishing areas force fluxes to not directly contribute to updating cells. In particular, between two opposite cells straddling a pole, the constant-r interface has no area. Whatever hydrodynamical fluxes are computed from the Riemann problem, we know these will not contribute toward the update of either cell's conserved quantities. In practice, naive implementation of the update might lead to $0/0$ errors, and so one must guard against this in the code. The Riemann problem cannot be ignored, however, as the fluxes of magnetic field correspond to electric field along the pole, and this is used to update other interfaces, in particular constant-ϕ interfaces bordering the pole.

4. Physically coincident electric fields must be kept the same in their various representations. As mentioned in the previous point, the electric fields along the pole are important for self-consistent evolution; they cannot be neglected or held constant. However, there will be as many representations of such values as there are cells in a ring around the pole. Roundoff error can cause these representations to differ, leading to inconsistent evolution of the magnetic fields. This leads to the production of

magnetic monopoles and even numerical instability. The solution is to have all such representations of the same electric field averaged, with the average overwriting all the values, after every partial time step. Note this changes the topology of the connectivity graph for block communication.

The complexity of polar boundary bookkeeping is simplified by forcing all blocks bordering a particular pole to be at the same level of mesh refinement.

A physical effect of properly treating the poles is that the magnetic field does not get artificially wound up as it would around an excised cone. In simulations where jet launching is being studied such artifacts are a cause for concern, as field winding can contribute to creating a jet.

1.5.4 Speed and Scalability

Independent of the particular algorithms or problems considered, the *performance* of any scientific code is grounded by the serial speed and parallel scaling of computation.

Modern processors no longer improve performance by increasing the clock frequency. Instead, pipelining, vectorization, and more efficient hardware instructions are used to reduce the wall time of computations. This means performance is no longer tied to a single number beyond all control of the programmer. While ideally compilers will optimize code to the hardware, there are often performance benefits to be had by tuning a code.

In writing Athena++ we paid close attention to such issues. The following techniques proved useful in improving performance:

- Minimizing expensive evaluations and memory access. We precompute trigonometric and other such expensive functions related to the stationary metric as much as possible. At the same time, we minimize memory usage by storing only 1D and 2D arrays of precomputed values, as discussed in §2.3.9.

- Sweeping through arrays as they are laid out in memory. For example, the hydrodynamical primitive variables are stored in a $5 \times N_3 \times N_2 \times N_1$ array. All nested loops

range over N_1 in the innermost loop, avoiding strides greater than unity and thus minimizing cache misses. This includes calculating the fluxes, where even the x^2- and x^3-fluxes are evaluated in the x^1-direction order.

- Vectorizing all innermost loops. While modern compilers can often automatically vectorize loops, even when that entails inlining functions defined in other compilation units, there are several failure modes we sought to address. For one, complicated functions sometimes require manual inlining (e.g. in the HLLD solver, the function to evaluate the residual given a guess for the contact pressure). Additionally, vectorization cannot happen if the loop we seek to vectorize has indeterminate loops within it, as happens with algorithms that iterate until convergence. In such cases we examine how many iterations are typically used in cases were convergence happens at all, enabling us to fix a constant number of iterations.

We measure hydrodynamics and MHD performance as a function of both the metric and the Riemann solver. For special relativity, we run a shock tube in Minkowski coordinates. We then run the same shock tube using the full general relativity framework but still in Minkowski coordinates. Finally, we run a Fishbone–Moncrief torus (§3.6) in Kerr–Schild coordinates with nonzero spin. All tests are 3D. For the Riemann solvers we use a non-frame-transforming LLF solver (LLF-NT, as used in many GRMHD codes), a frame-transforming LLF solver (LLF-T), and either HLLC or HLLD (both frame-transforming) as appropriate. The distinction between using or not using a frame transformation (§2.3.3) only applies to general relativity.

Tests were performed on a single core of an Intel Xeon E5-2670 (2.6 GHz Sandy Bridge) processor. We report the number of cells updated per wall time second in Table 1.2.

The decrease in performance in going from special to general relativity with the Minkowski metric reflects the cost of using the general-relativistic variable inversion and wavespeed formulas. The full cost of general relativity in realistic but reasonable metrics is represented by the final column of numbers, where nontrivial Kerr–Schild geometric factors enter into most calculations.

Table 1.2: Athena++ single-core performance (10^5 cell updates per second).

Riemann Solver	SR	GR: Minkowski	GR: Kerr–Schild
Hydrodynamics			
LLF-NT	12.8	10.4	3.82
LLF-T	—	10.1	4.28
HLLC	8.57	7.36	3.28
MHD			
LLF-NT	3.24	4.70	2.94
LLF-T	—	2.60	1.97
HLLD	1.23	1.13	1.22

For a fixed geometry and Riemann solver, MHD problems run at 1/4 to 3/4 the speed of pure hydrodynamics problems in both special and general relativity. We also performed tests using HLLE solvers. These ran at $93 \pm 6\%$ the speed of the corresponding LLF tests, with similar accuracies.

Single-core optimization is however not sufficient to make use of cutting-edge computational systems, where parallel computing is the central paradigm. Athena++ is written to efficiently balance computation across over one million message-passing interface (MPI) ranks, and it can also utilize OpenMP for threading across shared memory.

The scaling of Athena++ was tested with a full 3D GRMHD computation using Kerr–Schild coordinates. The computation was performed on the NAOJ Cray XC30, which has 24 cores per node. We divide the domain into blocks of 64^3 cells, assigning one block to each core. The scaling results out to 6144 cores are shown in Figure 1.1. For both hydrodynamics and MHD, there is a 20% per-core performance penalty to go from one core to one full node, almost certainly associated with saturating the memory bus when all cores on a node are used. However, the cost of going to many nodes from one is negligible. The performance of 6144 cores is over 97% relative to 24 cores for hydrodynamics, and it is indistinguishable from 100% for MHD.

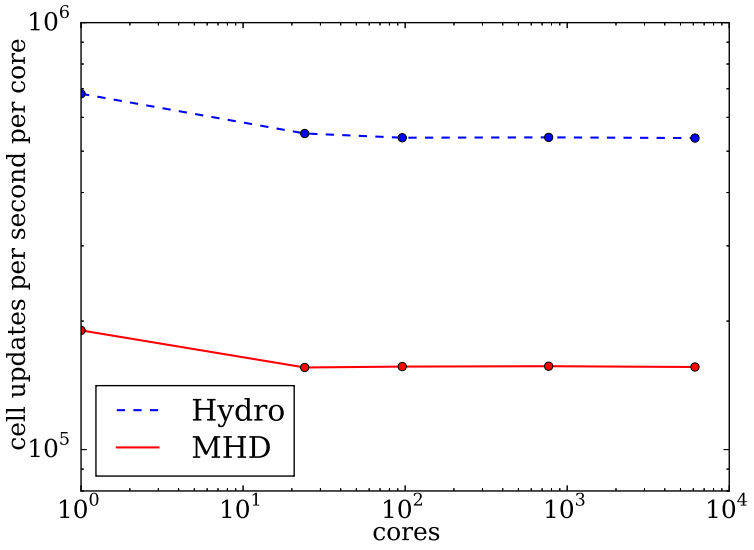

Figure 1.1: Performance per core running a 3D GRMHD simulation on a cluster.

Chapter 2

Numerical Solution of the Equations of General-Relativistic Magnetohydrodynamics

With our motivations and guiding considerations established, we now turn to the mathematical and numerical details necessary for our endeavor.

2.1 Differential Forms of Equations

The relevant equations for general-relativistic magnetohydrodynamics (GRMHD) are derived in numerous sources in the literature (e.g. Gammie et al., 2003). Here we omit the derivation of the governing differential equations, but we include all the necessary definitions in order to avoid notational ambiguity.

2.1.1 General Equations

Denote the covariant components of the spacetime metric by $g_{\mu\nu}$. Consider a perfect fluid with comoving mass density ρ, comoving gas pressure $p_{\rm gas}$, and comoving enthalpy per unit mass h. For an adiabatic gas of index Γ, we have $h = 1+\Gamma/(\Gamma-1)\times p_{\rm gas}/\rho$. The fluid will have coordinate-frame 4-velocity components u^μ. There will also be magnetic field components B^i in the coordinate frame.

The stress-energy tensor for the fluid has components

$$T^{\mu\nu} = (\rho h + b_\lambda b^\lambda) u^\mu u^\nu + \left(p_{\text{gas}} + \frac{1}{2} b_\lambda b^\lambda \right) g^{\mu\nu} - b^\mu b^\nu. \tag{2.1}$$

Here

$$b^0 = g_{i\mu} B^i u^\mu, \tag{2.2a}$$

$$b^i = \frac{1}{u^0}(B^i + b^0 u^i) \tag{2.2b}$$

are the components $b^\mu = u_\nu (*F)^{\nu\mu}$ of the contravariant comoving magnetic field in terms of the metric, 3-magnetic field, and 4-velocity. We note in passing that $b^\mu u_\mu = 0$. In terms of the electric and magnetic fields, the contravariant components of the dual electromagnetic field tensor are

$$(*F)^{\mu\nu} = \begin{pmatrix} 0 & -B^1 & -B^2 & -B^3 \\ B^1 & 0 & E^3 & -E^2 \\ B^2 & -E^3 & 0 & E^1 \\ B^3 & E^2 & -E^1 & 0 \end{pmatrix}, \tag{2.3}$$

where μ indexes rows and ν columns. Note that there are different sign conventions for the field tensor.

One can verify that

$$(*F)^{\mu\nu} = b^\mu u^\nu - b^\nu u^\mu. \tag{2.4}$$

As we consider only ideal MHD here, the electric fields can always be inferred from the magnetic fields and fluid velocities via (2.2) and (2.4).

The equations of ideal MHD are

$$\nabla_\mu (\rho u^\mu) = 0, \tag{2.5a}$$

$$\nabla_\mu T^{\mu\nu} = 0, \tag{2.5b}$$

$$\nabla_\mu (*F)^{\mu\nu} = 0 \tag{2.5c}$$

(cf. Gammie et al., 2003, (1), (3), (14)). These conservation laws for mass, momentum, and magnetic field govern the time evolution of a system consisting of a magnetized fluid in a stationary spacetime.

2.1.2 Rewriting Covariant Derivatives

In order to time-evolve our system, we need to have explicit temporal partial derivatives in our equations. Additionally, we prefer to have spatial partial derivatives. These considerations motivate the following lemma useful for rewriting covariant derivatives.

Suppose we have an arbitrary tensor with components ${A^{\mu\alpha_1\cdots\alpha_i}}_{\beta_1\cdots\beta_j}$. Then

$$\begin{aligned}\nabla_\nu {A^{\mu\alpha_1\cdots\alpha_i}}_{\beta_1\cdots\beta_j} &= \partial_\nu {A^{\mu\alpha_1\cdots\alpha_i}}_{\beta_1\cdots\beta_j} + {A^{\sigma\alpha_1\cdots\alpha_i}}_{\beta_1\cdots\beta_j}\Gamma^\mu_{\sigma\nu} \\ &\quad + \sum_{k=1}^{i} {A^{\mu\alpha_1\cdots\alpha_{k-1}\sigma\alpha_{k+1}\cdots\alpha_i}}_{\beta_1\cdots\beta_j}\Gamma^{\alpha_k}_{\sigma\nu} \\ &\quad - \sum_{k=1}^{j} {A^{\mu\alpha_1\cdots\alpha_i}}_{\beta_1\cdots\beta_{k-1}\sigma\beta_{k+1}\cdots\beta_j}\Gamma^{\sigma}_{\beta_k\nu}.\end{aligned} \tag{2.6}$$

The explicit summations directly result in the source terms in the equations we will be integrating.

In the case of calculating a divergence, it is possible to simplify the first two terms on the right-hand side of (2.6). For this we will suppress the spectator indices $\alpha_1, \ldots, \alpha_i$ and $\beta_1, \ldots, \beta_j$. Manipulating partial derivatives we can write

$$\begin{aligned}\nabla_\mu A^\mu &= \partial_\mu A^\mu + A^\nu \Gamma^\mu_{\nu\mu} \\ &= \frac{1}{\sqrt{-g}}\left(\partial_\mu(\sqrt{-g}A^\mu) - A^\mu \partial_\mu \sqrt{-g}\right) + A^\nu \Gamma^\mu_{\nu\mu} \\ &= \frac{1}{\sqrt{-g}}\partial_\mu(\sqrt{-g}A^\mu) + A^\mu\left(\Gamma^\nu_{\mu\nu} - \frac{1}{\sqrt{-g}}\partial_\mu\sqrt{-g}\right)\end{aligned} \tag{2.7}$$

for any suitably differentiable strictly negative function g. Note that

$$\frac{1}{\sqrt{-g}}\partial_\mu\sqrt{-g} = -\frac{1}{2g}\partial_\mu(-g) = \frac{1}{2g}\partial_\mu g. \tag{2.8}$$

Now take g to be the determinant of the metric. Jacobi's formula tells us that in the case of an invertible matrix M_t parameterized by t,

$$\partial_t \det(M_t) = \det(M_t)\,\mathrm{tr}(M_t^{-1}\partial_t M_t). \tag{2.9}$$

Therefore

$$\frac{1}{\sqrt{-g}}\partial_\mu\sqrt{-g} = \frac{1}{2}g^{\alpha\beta}\partial_\mu g_{\alpha\beta}. \tag{2.10}$$

We can then write

$$\nabla_\mu A^\mu = \frac{1}{\sqrt{-g}}\partial_\mu(\sqrt{-g}A^\mu) + A^\mu\left(\frac{1}{2}g^{\nu\sigma}(\partial_\mu g_{\nu\sigma} + \partial_\nu g_{\mu\sigma} - \partial_\sigma g_{\mu\nu}) - \frac{1}{2}g^{\alpha\beta}\partial_\mu g_{\alpha\beta}\right). \tag{2.11}$$

Changing summation indices and relying on the symmetry of the metric, one can see that the first and fourth terms in parentheses cancel, as do the second and third. Thus we are left with

$$\nabla_\mu A^\mu = \frac{1}{\sqrt{-g}}\partial_\mu(\sqrt{-g}A^\mu), \tag{2.12}$$

which is (8.51c) of Misner et al. (1973).

Restoring the suppressed indices yields the covariant-to-partial divergence rule

$$\begin{aligned}\nabla_\mu A^{\mu\alpha_1\cdots\alpha_i}{}_{\beta_1\cdots\beta_j} = {} & \frac{1}{\sqrt{-g}}\partial_\mu\left(\sqrt{-g}A^{\mu\alpha_1\cdots\alpha_i}{}_{\beta_1\cdots\beta_j}\right) \\ & + \sum_{k=1}^{i} A^{\mu\alpha_1\cdots\alpha_{k-1}\sigma\alpha_{k+1}\cdots\alpha_i}{}_{\beta_1\cdots\beta_j}\Gamma^{\alpha_k}_{\sigma\mu} \\ & - \sum_{k=1}^{j} A^{\mu\alpha_1\cdots\alpha_i}{}_{\beta_1\cdots\beta_{k-1}\sigma\beta_{k+1}\cdots\beta_j}\Gamma^{\sigma}_{\beta_k\mu}.\end{aligned} \tag{2.13}$$

2.1.3 The Equations in Terms of Partial Derivatives

We next express our equations in terms of partial derivatives, forgoing manifest covariance in favor of expediency in computational implementation.

Applying (2.12) to (2.5a) yields

$$\partial_0(\sqrt{-g}\rho u^0) + \partial_j(\sqrt{-g}\rho u^j) = 0. \tag{2.14}$$

For (2.5b), we choose to apply (2.13) to the lowered form of the equation:

$$\partial_0(\sqrt{-g}\,T^0{}_\mu) + \partial_i(\sqrt{-g}\,T^i{}_\mu) = \sqrt{-g}\,T^\nu{}_\sigma\,\Gamma^\sigma_{\mu\nu}. \tag{2.15}$$

The last conservation law, (2.5c), can also be expanded via (2.13):

$$\partial_0\Big(\sqrt{-g}\,(*F)^{0\mu}\Big) + \partial_j\Big(\sqrt{-g}\,(*F)^{j\mu}\Big) = -\sqrt{-g}\,(*F)^{\nu\sigma}\Gamma^\mu_{\sigma\nu}. \tag{2.16}$$

Now the right-hand side vanishes, as it is the contraction of symmetric indices with anti-symmetric ones. This is in agreement with Misner et al. (1973, 8.51c). Also (2.4) and (2.2b) tell us

$$(*F)^{0i} = b^0u^i - (B^i + b^0u^i) = -B^i \tag{2.17}$$

and

$$(*F)^{ji} = b^ju^i - b^iu^j. \tag{2.18}$$

Thus the spatial components of (2.16) are equivalent to

$$\partial_0(\sqrt{-g}B^i) + \partial_j\Big(\sqrt{-g}\,(b^iu^j - b^ju^i)\Big) = 0, \tag{2.19}$$

while the temporal component is the constraint

$$\partial_j(\sqrt{-g}B^j) = 0. \tag{2.20}$$

Summarizing, the evolution equations for GRMHD can be written

$$\partial_0(\sqrt{-g}\rho u^0) + \partial_j(\sqrt{-g}\rho u^j) = 0, \tag{2.21a}$$

$$\partial_0(\sqrt{-g}\,T^0{}_\mu) + \partial_j(\sqrt{-g}\,T^j{}_\mu) = \sqrt{-g}\,T^\nu{}_\sigma\,\Gamma^\sigma_{\mu\nu}, \tag{2.21b}$$

$$\partial_0(\sqrt{-g}B^i) + \partial_j\Big(\sqrt{-g}\,(b^iu^j - b^ju^i)\Big) = 0. \tag{2.21c}$$

These equations agree with (2), (4), and (18) of Gammie et al. (2003).

2.2 Integral Forms of Equations

The differential forms of the equations, written in terms of partial rather than covariant derivatives, would be sufficient for a finite difference scheme. However, we endeavor to develop a finite volume method, and so we must integrate our equations over regions of spacetime. This will naturally lead to equations for the evolution of volume-averaged quantities in terms of surface-averaged fluxes.

2.2.1 General Equations

Consider the generic form

$$\partial_0(\sqrt{-g}\,\mathbf{C}) + \partial_1(\sqrt{-g}\,\mathbf{F}) + \partial_2(\sqrt{-g}\,\mathbf{G}) + \partial_3(\sqrt{-g}\,\mathbf{H}) = \sqrt{-g}\,\mathbf{S} \tag{2.22}$$

of (2.21), where we group the conserved quantities, fluxes, and sources as

$$\mathbf{C} = (\rho u^0, T^0{}_\mu, B^i), \tag{2.23a}$$

$$\mathbf{F} = (\rho u^1, T^1{}_\mu, b^i u^1 - b^1 u^i), \tag{2.23b}$$

$$\mathbf{G} = (\rho u^2, T^2{}_\mu, b^i u^2 - b^2 u^i), \tag{2.23c}$$

$$\mathbf{H} = (\rho u^3, T^3{}_\mu, b^i u^3 - b^3 u^i), \tag{2.23d}$$

$$\mathbf{S} = (0, T^\rho{}_\sigma\,\Gamma^\sigma_{\rho\mu}, 0^i) \tag{2.23e}$$

(cf. Antón et al., 2006, (41)–(43)).

Fix particular coordinates x^μ_* and positive increments Δx^μ_*. Let Σ_μ be the surface $x^\mu = x^\mu_*$ for each μ, and similarly let Σ'_μ be the surface $x^\mu = x^\mu_* + \Delta x^\mu_*$. We refer to the parallelepiped in spacetime bounded by the 3-surfaces Σ_μ and Σ'_μ as Ω, and we understand the surfaces to not extend past Ω. This is illustrated in Figure 2.1

Following Banyuls et al. (1997), integrate (2.22) over Ω with the volume measure $\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3$ (which is not coordinate independent). Noting that the integrands on the

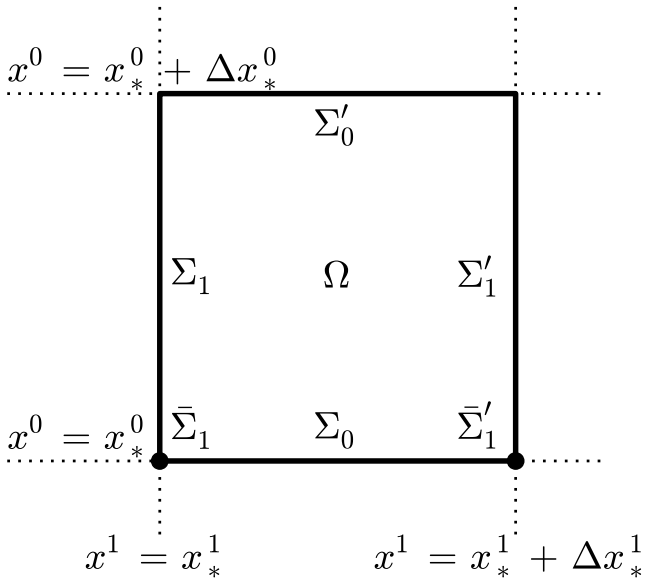

Figure 2.1: 1+1 spacetime diagram of a volume element Ω and its bounding surfaces $\Sigma^{(\prime)}_\mu$ of constant coordinate values.

left-hand side will be partial derivatives with respect to the coordinates, we can write

$$\begin{aligned}\int_{\Sigma'_0} \sqrt{-g}\,\mathbf{C}\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3 = &\int_{\Sigma_0} \sqrt{-g}\,\mathbf{C}\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3 + \int_{\Omega} \sqrt{-g}\,\mathbf{S}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3 \\ &+ \int_{\Sigma_1} \sqrt{-g}\,\mathbf{F}\,\mathrm{d}x^0\,\mathrm{d}x^2\,\mathrm{d}x^3 - \int_{\Sigma'_1} \sqrt{-g}\,\mathbf{F}\,\mathrm{d}x^0\,\mathrm{d}x^2\,\mathrm{d}x^3 \\ &+ \int_{\Sigma_2} \sqrt{-g}\,\mathbf{G}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^3 - \int_{\Sigma'_2} \sqrt{-g}\,\mathbf{G}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^3 \\ &+ \int_{\Sigma_3} \sqrt{-g}\,\mathbf{H}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^2 - \int_{\Sigma'_3} \sqrt{-g}\,\mathbf{H}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^2.\end{aligned} \tag{2.24}$$

To see that this is equivalent to (28) and (29) of Banyuls et al., we note two changes of notation. First, their $\mathrm{d}\Omega$ is the proper volume element $\sqrt{-g}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3$. Second, their $\sqrt{\gamma}\,\mathbf{F}^{\bar{0}} = \sqrt{-g}\,\mathbf{F}^{\bar{0}}/\alpha$ (here we place bars over indices in their basis) is equal to our $\mathbf{C}$, as their

basis has as its $\bar{0}$-component the unit timelike normal

$$\hat{n} = \frac{1}{\alpha}\Big(\vec{\partial}_0 - \beta^i \vec{\partial}_i\Big). \tag{2.25}$$

That is, we can write a vector

$$\vec{A} = A^0 \vec{\partial}_0 + A^i \vec{\partial}_i \tag{2.26}$$

as

$$\vec{A} = A^{\bar{0}} \hat{n} + A^{\bar{\imath}} \vec{\partial}_{\bar{\imath}} = \frac{1}{\alpha} A^{\bar{0}} \vec{\partial}_0 + \left(A^{\bar{\imath}} - \frac{\beta^{\bar{\imath}}}{\alpha} A^{\bar{0}} \right) \vec{\partial}_{\bar{\imath}}, \tag{2.27}$$

from which we conclude that $A^{\bar{0}}/\alpha = A^0$.

2.2.2 Discretized Forms of Equations

Let i index regions bounded by surfaces of constant x^1, and similarly let j and k relate to x^2 and x^3. Let n index surfaces of constant x^0. Fundamentally we choose to evolve the 3-volume-averaged quantities

$$\mathbf{C}^n_{i,j,k} = \frac{1}{\Delta V_{i,j,k}} \int_{\Sigma_0} \sqrt{-g}\, \mathbf{C}\, \mathrm{d}x^1\, \mathrm{d}x^2\, \mathrm{d}x^3, \tag{2.28}$$

where

$$\Delta V_{i,j,k} = \int_{\Sigma_0} \sqrt{-g}\, \mathrm{d}x^1\, \mathrm{d}x^2\, \mathrm{d}x^3. \tag{2.29}$$

As a result of having a stationary metric, $\Delta V_{i,j,k}$ is constant in "time" x^0.

We note that $\Delta V_{i,j,k}$ is not coordinate invariant; it would be so only if we replaced the $\sqrt{-g}$ with $\sqrt{\gamma}$, where γ is the determinant of the 3-metric induced on surfaces of constant x^0 ($\gamma_{ij} = g_{ij}$). It is however constant in "time" x^0, as a result of having a stationary metric. Also note that this formula leads to a definition of cell-centered positions. The x^1-coordinate of the center of cell i, j, k is the value $\bar{x}^1$ such that the integral (2.29) limited by x^1 going from its minimum value to $\bar{x}^1$ is $\Delta V_{i,j,k}/2$. Similarly we can define $\bar{x}^2$ and $\bar{x}^3$.

The fluxes found along the interfaces are taken to be constant along those interfaces,

denoted $\mathbf{F}^{n+1/2}_{i\pm1/2,j,k}$, $\mathbf{G}^{n+1/2}_{i,j\pm1/2,k}$, and $\mathbf{H}^{n+1/2}_{i,j,k\pm1/2}$. The constant (coordinate-dependent) 2-areas of the interfaces will be denoted

$$\Delta A_{i\pm1/2,j,k} = \int\limits_{\bar{\Sigma}_1^{(\prime)}} \sqrt{-g}\,\mathrm{d}x^2\,\mathrm{d}x^3, \tag{2.30a}$$

$$\Delta A_{i,j\pm1/2,k} = \int\limits_{\bar{\Sigma}_2^{(\prime)}} \sqrt{-g}\,\mathrm{d}x^1\,\mathrm{d}x^3, \tag{2.30b}$$

$$\Delta A_{i,j,k\pm1/2} = \int\limits_{\bar{\Sigma}_3^{(\prime)}} \sqrt{-g}\,\mathrm{d}x^1\,\mathrm{d}x^2, \tag{2.30c}$$

where primes are associated with plus signs and $\bar{\Sigma}$ is the projection of the 3-surface Σ along coordinate x^0 to a 2-surface of constant x^0.

We pause to also define the formulas for "lengths" of edges in a similar fashion:

$$\Delta L_{i,j\pm_a1/2,k\pm_b1/2} = \int\limits_{\bar{\Sigma}_2^{(\prime a)}\cap\bar{\Sigma}_3^{(\prime b)}} \sqrt{-g}\,\mathrm{d}x^1, \tag{2.31a}$$

$$\Delta L_{i\pm_a1/2,j,k\pm_b1/2} = \int\limits_{\bar{\Sigma}_1^{(\prime a)}\cap\bar{\Sigma}_3^{(\prime b)}} \sqrt{-g}\,\mathrm{d}x^2, \tag{2.31b}$$

$$\Delta L_{i\pm_a1/2,j\pm_b1/2,k} = \int\limits_{\bar{\Sigma}_1^{(\prime a)}\cap\bar{\Sigma}_2^{(\prime b)}} \sqrt{-g}\,\mathrm{d}x^3. \tag{2.31c}$$

The arrangement of volumes, areas, and lengths is illustrated in Figure 2.2. The left panel labels regions of space according to the notation used in Figure 2.1, while the right panel labels the corresponding sizes of those regions as given by (2.29), (2.30), and (2.31).

The lengths are distinct from the cell widths needed to set the maximum allowable time step for stability. These latter quantities are obtained by integrating the induced 3-metric line element on the constant-x^0 surface along curves of constant coordinate x^i. That is, let C^1 be the curve in Σ_0 that varies only in x^1, passes through the center of cell i,j,k, and is bounded by $\bar{\Sigma}_1$ and $\bar{\Sigma}'_1$, and similarly define C^2 and C^3. Since we have the line element ${}^{(3)}ds^2 = \gamma_{ij}\,\mathrm{d}x^i\,\mathrm{d}x^j$ on Σ_0, where $\gamma_{ij} = g_{ij}$ are the 3-metric components, we can take these

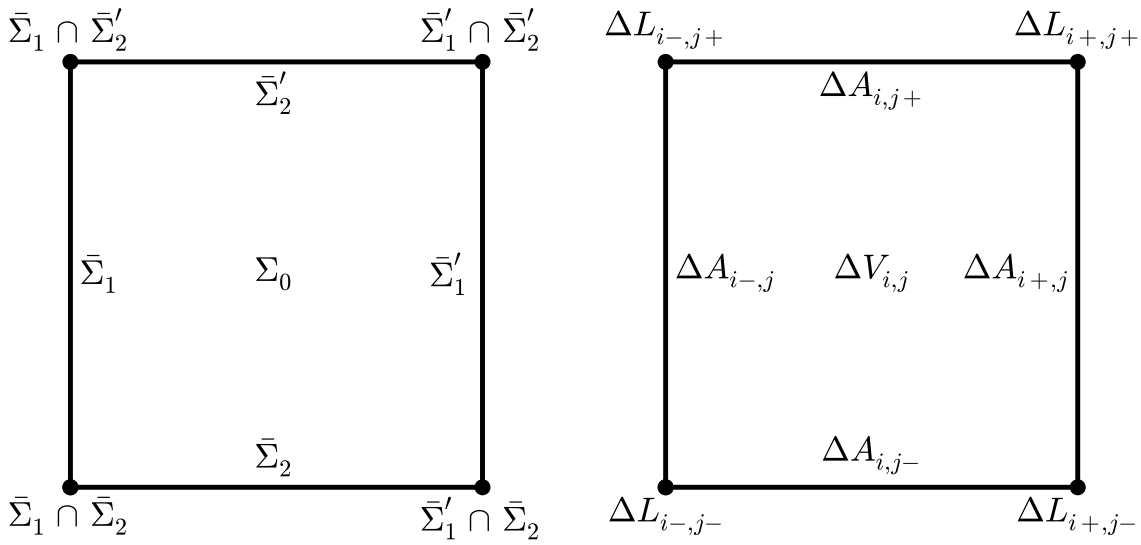

Figure 2.2: 2+0 spacetime diagram of a single cell and its bounding surfaces and edges. The spacetime regions are labeled on the left, while on the right their corresponding sizes (integrals over them of $\sqrt{-g}$) are shown.

cell widths to be

$$\Delta W_{x^1;i,j,k} = \int_{C^1} \sqrt{g_{11}}\,\mathrm{d}x^1, \tag{2.32a}$$

$$\Delta W_{x^2;i,j,k} = \int_{C^2} \sqrt{g_{22}}\,\mathrm{d}x^2, \tag{2.32b}$$

$$\Delta W_{x^3;i,j,k} = \int_{C^3} \sqrt{g_{33}}\,\mathrm{d}x^3. \tag{2.32c}$$

Dividing through by the volume and making the approximation that states are constant on timescales shorter than those resolved, the source term becomes

$$\frac{1}{\Delta V_{i,j,k}} \int_{\Omega} \sqrt{-g}\,\mathbf{S}\,\mathrm{d}x^0\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3 = \int_{x_*^0}^{x_*^0+\Delta x_*^0} \mathbf{S}_{i,j,k}^{n+1/2}\,\mathrm{d}x^0, \tag{2.33}$$

where

$$\mathbf{S}_{i,j,k}^{n+1/2} = \frac{1}{\Delta V_{i,k,k}} \int_{\Sigma_0^*} \sqrt{-g}\,\mathbf{S}\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3. \tag{2.34}$$

Here Σ_0^* is the surface between Σ_0 (index n) and Σ_0' (index $n+1$) corresponding to index $n+1/2$.

Now the nature of finite volume methods means we must at some point accept volume averages as the only quantities to which we have access. According to (2.23e), we can write the volume averages of the nonzero components of $\mathbf{S}$ – call them S_μ – as

$$(S_\mu)_{i,j,k}^{n+1/2} \approx T^\rho{}_\sigma(\mathbf{P}_{i,j,k}^{n+1/2})\,\Gamma^\sigma_{\rho\mu}, \tag{2.35}$$

where the parentheses on the right-hand side indicate the use of the pointwise function $\mathbf{P} \mapsto T^\mu{}_\nu$ defined by lowering (2.1), $\mathbf{P}$ is the vector of primitive variables

$$\mathbf{P} = (\rho, p_{\text{gas}}, v^i, B^i), \tag{2.36}$$

and we invoke the standard volume averaging notation to write

$$\mathbf{P}_{i,j,k}^{n+1/2} = \frac{1}{\Delta V_{i,j,k}} \int\limits_{\Sigma_0^*} \sqrt{-g}\,\mathbf{P}\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3 \tag{2.37}$$

(and similarly for n and Σ_0 and for $n+1$ and Σ_0'). Also in the same vein we will assume the pointwise correspondence $\mathbf{P} \leftrightarrow \mathbf{C}$ holds for $\mathbf{P}_{i,j,k}^n \leftrightarrow \mathbf{C}_{i,j,k}^n$ using the same cell-centered geometric quantities.

The connection coefficients appearing in (2.35) can be evaluated at the cell centers. The approximation of moving the coefficients outside the integral by assuming they have constant values equal to their cell-centered values is second-order accurate in space. One could imagine taking a volume average of the connection coefficients

$$\bar{\Gamma}^\sigma_{\rho\mu} = \frac{1}{\Delta V_{i,k,k}} \int\limits_{V_{i,k,k}} \sqrt{-g}\,\Gamma^\sigma_{\rho\mu}\,\mathrm{d}x^1\,\mathrm{d}x^2\,\mathrm{d}x^3, \tag{2.38}$$

though this would not improve spatial accuracy. (Because a stationary metric implies a stationary connection, these values can be precomputed in either method.) Fundamentally, calculating the source term integrals to better than second order requires more information

per cell. While it would be simple to evaluate the connection coefficients at more points, a higher-order quadrature would require higher-order reconstruction of either the primitives or the corresponding stress-energy tensor, utilizing information from neighboring cells.

All integrands consist of parts that are either exactly constant in x^0 (because the metric is assumed to be stationary) or are taken to be constant over the interval in question, and so we can perform the x^0-integrations separately. Piecing everything together, (2.24) can then be written

$$\begin{aligned}\mathbf{C}^{n+1}_{i,j,k} = \mathbf{C}^{n}_{i,j,k} &+ \Delta t \mathbf{S}^{n+1/2}_{i,j,k} \\ &+ \frac{\Delta t}{\Delta V_{i,j,k}}\left(\mathbf{F}^{n+1/2}_{i-1/2,j,k}\Delta A_{i-1/2,j,k} - \mathbf{F}^{n+1/2}_{i+1/2,j,k}\Delta A_{i+1/2,j,k}\right) \\ &+ \frac{\Delta t}{\Delta V_{i,j,k}}\left(\mathbf{G}^{n+1/2}_{i,j-1/2,k}\Delta A_{i,j-1/2,k} - \mathbf{G}^{n+1/2}_{i,j+1/2,k}\Delta A_{i,j+1/2,k}\right) \\ &+ \frac{\Delta t}{\Delta V_{i,j,k}}\left(\mathbf{H}^{n+1/2}_{i,j,k-1/2}\Delta A_{i,j,k-1/2} - \mathbf{H}^{n+1/2}_{i,j,k+1/2}\Delta A_{i,j,k+1/2}\right),\end{aligned} \tag{2.39}$$

where Δt is a more familiar symbol for what we earlier denoted Δx^0_*, the change in coordinate time from surface n to $n+1$. Note that suitable reinterpretations of n, $n+1/2$, and $n+1$ yield formulas for different substeps in various integration schemes.

2.2.3 Magnetic Evolution

Thus far we have treated the three components of the magnetic field in the same manner as the density and four components of momentum. Equation (2.39) could be used to update volume-averaged magnetic fields in the same way as the other quantities. However, this method is known to violate the divergence-free constraint of the magnetic field $\nabla_\mu(*F)^{0\mu} = 0$. That is, even initial conditions free of magnetic monopoles develop them numerically over time.

A number of schemes have been developed to correct for spurious magnetic monopole generation as discussed in §1.3.2. Rather than apply post hoc fixes, however, we choose a different discretization, one that naturally preserves the constraint.

In the case of magnetic fields, we take the fundamental variables to be area averages.

This is a crucial part of the constrained transport (CT) scheme of Evans and Hawley (1988) (see §2.3.6). For concreteness, consider the evolution of

$$(B^1)^n_{i-1/2,j,k} = \frac{1}{\Delta A_{i-1/2,j,k}} \int\limits_{A_1} B^1 \sqrt{-g}\, \mathrm{d}x^2\, \mathrm{d}x^3, \tag{2.40}$$

where the positioning of quantities and orientations of surfaces and edges is the same as Figure 1 of Stone and Gardiner (2008). Given the electric fields E^i defined along the edges (as discussed further in §2.3.6), we can update the magnetic field according to

$$\begin{aligned}\left(B^1\right)^{n+1}_{i-1/2,j,k} = \left(B^1\right)^{n}_{i-1/2,j,k} &\\ + \frac{\Delta t}{\Delta A_{i-1/2,j,k}} \Big(\Delta L_{i-1/2,j-1/2,k} \left(E^3\right)^{n+1/2}_{i-1/2,j-1/2,k} &\\ - \Delta L_{i-1/2,j+1/2,k} \left(E^3\right)^{n+1/2}_{i-1/2,j+1/2,k} \Big) &\\ + \frac{\Delta t}{\Delta A_{i-1/2,j,k}} \Big(\Delta L_{i-1/2,j,k-1/2} \left(-E^2\right)^{n+1/2}_{i-1/2,j,k-1/2} &\\ - \Delta L_{i-1/2,j,k+1/2} \left(-E^2\right)^{n+1/2}_{i-1/2,j,k+1/2} \Big), &\end{aligned} \tag{2.41}$$

which is the discretized form of the third evolution equation (2.21c). This replaces (2.39) in updating the magnetic field.

2.3 Numerical Algorithm

With a discretization scheme in hand, we proceed to detail how quantities are evolved on the grid. The general procedure we will follow is:

1. Invert the conserved variables to primitive variables.
2. Reconstruct the primitive variables to interfaces.
3. Transform the primitive variables into locally Minkowski frames.
4. Solve the Riemann problem at each interface, finding fluxes.
5. Transform the fluxes back into the global coordinate frame.

6. Apply the CT method to calculate electric fields.

7. Calculate the curvature source terms from the primitive variables.

8. Integrate forward in time.

2.3.1 Variable Inversion

One of the most challenging aspects of numerical methods for relativistic fluid dynamics is variable inversion, the process of extracting the primitive variables from the conserved variables. The difficulty lies in the conserved variables being given by nonlinear functions of the primitive variables with no simple inverses. It is compounded by the fact that only certain primitive states are physically admissible (for example there can be no negative densities or superluminal velocities), and so a given set of conserved variables might not exactly correspond to any allowable set of primitives. Our methods described here are largely based on the $1D_W$ inversion scheme of Noble et al. (2006).

If magnetic fields are present, they are interpolated to the volume-averaged centers as follows, using the B^1 case for concreteness. Given $B^1_{i\pm1/2}$ at interface coordinates $x^1_{i\pm1/2}$ and volume-averaged coordinate x^1_i, define $\lambda = (x^1_i - x^1_{i-1/2})/(x^1_{i+1/2} - x^1_{i-1/2})$. Then the volume-averaged value of B^1 is taken to be

$$B^1_i = (1-\lambda)B^1_{i-1/2} + \lambda B^1_{i+1/2}. \tag{2.42}$$

In the case of general relativity, we take the primitive variables to be $\{\rho, p_{\rm gas}, \tilde{u}^i\}$ and the conserved variables to be $\{\rho u^0, T^0_{\ \mu}\}$. Here we use the same velocities as in Noble et al. (2006): $\tilde{u}^\mu = (\delta^\mu_\nu + n^\mu n_\nu)u^\nu$, where $n_\mu = -\alpha\delta^0_\mu$ are the components of the future-pointing unit vector normal to constant-x^0 hypersurfaces, and $\alpha = (-g^{00})^{-1/2}$ is the lapse. These projections have the desirable property that they describe subluminal motion no matter what values they have. Additionally, the fourth component is $\tilde{u}^0 = 0$. The Lorentz factor of the fluid as seen by the normal observer is $\gamma = (1 + g_{ij}\tilde{u}^i\tilde{u}^j)^{1/2}$. Following the $1D_W$ scheme, the Newton–Raphson method is used to solve a nonlinear equation for the relativistic gas

enthalpy $W = \gamma^2 \rho h$. In the hydrodynamics case we use the same procedure as with MHD, simply setting all magnetic fields to 0.

In special relativity, we take the primitive variables to be $\{\rho, p_{\rm gas}, v^i\}$ and the conserved variables to be $\{D, E, M^i\} = \{\gamma\rho, T^{0\mu}\}$, where $\gamma = u^0$ is the fluid Lorentz factor in the coordinate frame. For MHD, we also use the Newton–Raphson method to find W, where we simply have $W = \gamma^2 \rho h$. The method is detailed in Mignone and McKinney (2007), though we iterate on the residual of E as a function of W rather than $E - D$ as a function of $W - D$ as they do.

In the particular case of special-relativistic hydrodynamics, we can avoid using an iterative method and instead directly solve a quartic equation (Schneider et al., 1993). In particular, the 3-velocity magnitude $|v|$ must satisfy

$$a_4 |v|^4 + a_3 |v|^3 + a_2 |v|^2 + a_1 |v| + a_0 = 0, \tag{2.43}$$

where in terms of $M^2 = \eta_{ij} M^i M^j$ and $|M| = \sqrt{M^2}$ we define

$$a_4 = (\Gamma - 1)^2 (D^2 - M^2), \tag{2.44a}$$

$$a_3 = -2\Gamma(\Gamma - 1)|M|E, \tag{2.44b}$$

$$a_2 = \Gamma^2 E^2 + 2(\Gamma - 1)M^2 - (\Gamma - 1)^2 D^2, \tag{2.44c}$$

$$a_1 = -2\Gamma |M| E, \tag{2.44d}$$

$$a_0 = M^2. \tag{2.44e}$$

Given the root $|v|$, we recover the primitives

$$\rho = D\sqrt{1 - |v|^2}, \tag{2.45a}$$

$$v^i = \frac{|v|}{|M|} M^i, \tag{2.45b}$$

$$p_{\rm gas} = (\Gamma - 1)(E - \eta_{ij} M^i v^j - \rho). \tag{2.45c}$$

As part of variable inversion we impose limits on the recovered variables. Since negative

densities and pressures are unphysical, and since vanishing values cause problems in finite-volume fluid codes, we ensure $\rho > \rho_{\rm min}$ and $p_{\rm gas} > p_{\rm gas}^{\rm min}$, where $\rho_{\rm min}$ and $p_{\rm gas}^{\rm min}$ are adjusted based on the problem. We also must obey the physical constraint $\gamma \geq 1$, and for numerical purposes we limit $\gamma < \gamma_{\rm max}$, where $\gamma_{\rm max}$ is often chosen to be 100.

If the recovered primitives fall outside these ranges, they are modified to the nearest admissible values. The primitives are also set to the floor values in cases where the iteration does not converge. Any such modifications are then followed by reevaluating the conserved quantities in the given cell.

2.3.2 Reconstruction

As input to the Riemann problem, one needs fluid states on either side of the interface. While one can simply take the states from the two adjacent cells, this method is only first-order accurate in the size of the cells. For smooth flows, information from a larger stencil can be used to infer left and right states in a higher-order way. Reconstruction is the process of interpolating these states while preserving any actual discontinuities and not introducing spurious extrema.

For a primitive quantity q, we define the values on the left and right sides of the $x^1_{i-1/2}$ interface using piecewise linear reconstruction. Specifically, we use the modified van Leer slope limiter algorithm of Mignone (2014), which reconstructs $q_{i-1/2,{\rm L/R}}$ from the values q_{i-2}, q_{i-1}, q_i, and q_{i+1}. This algorithm correctly takes into account nonuniform spacing, as well as the difference between the volume-averaged cell center coordinates ("centroids of volume" in that work) and arithmetic means of surface coordinates ("geometrical cell centers"). The slope limiter has been proved to be total variation diminishing (TVD) for all orthogonal coordinate systems. While this guarantee has not been extended to non-orthogonal coordinate systems, to our knowledge this has not been done with any other method, and indeed a number of methods in use are not even TVD with orthogonal but non-Cartesian or non-uniform grids.

For completeness, we summarize the Mignone algorithm here. Define

$$\Delta_{\mathrm{L}} = \frac{q_{i-1} - q_{i-2}}{x^1_{i-1} - x^1_{i-2}}, \tag{2.46a}$$

$$\Delta_{\mathrm{C}} = \frac{q_{i} - q_{i-1}}{x^1_{i} - x^1_{i-1}}, \tag{2.46b}$$

$$\Delta_{\mathrm{R}} = \frac{q_{i+1} - q_{i}}{x^1_{i+1} - x^1_{i}}, \tag{2.46c}$$

Define also

$$\Delta x_{\mathrm{L}} = x^1_{i-1/2} - x^1_{i-1}, \qquad \Delta x_{\mathrm{R}} = x^1_{i} - x^1_{i-1/2}, \tag{2.47a}$$

$$C^{\mathrm{F}}_{\mathrm{L}} = \frac{x^1_{i} - x^1_{i-1}}{\Delta x_{\mathrm{L}}}, \qquad C^{\mathrm{F}}_{\mathrm{R}} = \frac{x^1_{i+1} - x^1_{i}}{x^1_{i+1/2} - x^1_{i}}, \tag{2.47b}$$

$$C^{\mathrm{B}}_{\mathrm{L}} = \frac{x^1_{i-1} - x^1_{i-2}}{x^1_{i-1} - x^1_{i-3/2}}, \qquad C^{\mathrm{B}}_{\mathrm{R}} = \frac{x^1_{i} - x^1_{i-1}}{\Delta x_{\mathrm{R}}}. \tag{2.47c}$$

The left and right quantities are then given by

$$q_{i-1/2,\mathrm{L}} = \begin{cases} q_{i-1}, & \Delta_{\mathrm{L}}\Delta_{\mathrm{C}} \le 0, \\ q_{i-1} + \dfrac{\Delta x_{\mathrm{L}}\Delta_{\mathrm{L}}\Delta_{\mathrm{C}}(C^{\mathrm{F}}_{\mathrm{L}}\Delta_{\mathrm{L}} + C^{\mathrm{B}}_{\mathrm{L}}\Delta_{\mathrm{C}})}{\Delta^2_{\mathrm{L}} + (C^{\mathrm{F}}_{\mathrm{L}} + C^{\mathrm{B}}_{\mathrm{L}} - 2)\Delta_{\mathrm{L}}\Delta_{\mathrm{C}} + \Delta^2_{\mathrm{C}}}, & \Delta_{\mathrm{L}}\Delta_{\mathrm{C}} > 0; \end{cases} \tag{2.48a}$$

$$q_{i-1/2,\mathrm{R}} = \begin{cases} q_{i}, & \Delta_{\mathrm{C}}\Delta_{\mathrm{R}} \le 0, \\ q_{i} - \dfrac{\Delta x_{\mathrm{R}}\Delta_{\mathrm{C}}\Delta_{\mathrm{R}}(C^{\mathrm{F}}_{\mathrm{R}}\Delta_{\mathrm{C}} + C^{\mathrm{B}}_{\mathrm{R}}\Delta_{\mathrm{R}})}{\Delta^2_{\mathrm{C}} + (C^{\mathrm{F}}_{\mathrm{L}} + C^{\mathrm{B}}_{\mathrm{L}} - 2)\Delta_{\mathrm{C}}\Delta_{\mathrm{R}} + \Delta^2_{\mathrm{R}}}, & \Delta_{\mathrm{C}}\Delta_{\mathrm{R}} > 0. \end{cases} \tag{2.48b}$$

In the case of transverse magnetic field $q \in \{B^2, B^3\}$, the cell-centered values are used in the reconstruction algorithm. The longitudinal field B^1 is already defined at $x^1_{i-1/2}$, and so no reconstruction is needed for this quantity. Reconstruction for x^2- and x^3-surfaces proceeds analogously.

2.3.3 Frame Transformation

Rather than try to solve a Riemann problem directly in general coordinates, we choose to explicitly transform the left and right states at an interface to a locally Minkowski frame.

In this way, we can leverage the advances made in special-relativistic Riemann problems, as suggested by Pons et al. (1998) and Antón et al. (2006). While the scalar quantities ρ and p_{gas} transform between frames trivially, vector quantities rely on constructing a new basis $\{\vec{e}_{(\hat{\mu})}\}$ from the old coordinate basis $\{\vec{\partial}_{(\mu)}\}$, where the inner product with respect to the metric yields $\vec{\partial}_{(\mu)} \cdot \vec{\partial}_{(\nu)} = g_{\mu\nu}$ but $\vec{e}_{(\hat{\mu})} \cdot \vec{e}_{(\hat{\nu})} = \eta_{\mu\nu}$.

For a surface of constant x^i, we define the new basis to have the following properties:

1. $\vec{e}_{(\hat{\mu})}$ must be orthogonal to $\vec{e}_{(\hat{\nu})}$ for all $\mu \neq \nu$.
2. Each $\vec{e}_{(\hat{\mu})}$ must be normalized to have an inner product of ± 1 with itself, with $\vec{e}_{(\hat{t})}$ being timelike and $\vec{e}_{(\hat{j})}$ being spacelike.
3. $\vec{e}_{(\hat{t})}$ must be orthogonal to surfaces of constant x^0.
4. The projection of $\vec{e}_{(\hat{x})}$ onto the surface of constant x^0 must be orthogonal to the surface of constant x^i within that submanifold.

The first two properties are a restatement of the condition $\vec{e}_{(\hat{\mu})} \cdot \vec{e}_{(\hat{\nu})} = \eta_{\mu\nu}$. Without the third property, the fluxes returned by the Riemann solver would not necessarily correspond to time evolution in the understood x^0-direction. The fourth property allows us to ignore y- and z-fluxes in the Riemann problem, using only the x-fluxes returned by a 1-dimensional Riemann solver to infer the x^1-fluxes.

The details of the frame transformation are worked out in Appendix A. Here we quote the results. For an x^1-interface, define the components

$$M^{\mu}_{\ \hat{\nu}} = \begin{pmatrix} Ag^{00} & 0 & 0 & 0 \\ Ag^{01} & B(g^{01}g^{01} - g^{00}g^{11}) & 0 & 0 \\ Ag^{02} & B(g^{01}g^{02} - g^{00}g^{12}) & Dg_{33} & 0 \\ Ag^{03} & B(g^{01}g^{03} - g^{00}g^{13}) & -Dg_{23} & C \end{pmatrix}, \tag{2.49}$$

where we have the shorthands

$$A = -\left(-g^{00}\right)^{-1/2}, \tag{2.50a}$$

$$B = \left(g^{00}\left(g^{00}g^{11} - g^{01}g^{01}\right)\right)^{-1/2}, \tag{2.50b}$$

$$C = (g_{33})^{-1/2}, \tag{2.50c}$$

$$D = \left(g_{33}(g_{22}g_{33} - g_{23}g_{23})\right)^{-1/2}. \tag{2.50d}$$

This matrix transforms contravariant tensors in the locally Minkowski basis to the coordinate basis, so for example $T^1{}_\mu = g_{\mu\nu}M^1{}_{\hat{\mu}}M^\nu{}_{\hat{\nu}}T^{\hat{\mu}\hat{\nu}}$ are the x^1-fluxes of 4-momentum in the coordinate basis if $T^{\hat{\mu}\hat{\nu}}$ are the stress-energy tensor components in the orthonormal basis. The inverse transformation takes the form

$$M^{\hat{\mu}}{}_\nu = \begin{pmatrix} -A & 0 & 0 & 0 \\ Bg^{01} & -Bg^{00} & 0 & 0 \\ B^2Eg^{00}/Dg_{33} & B^2Fg^{00}/Dg_{33} & 1/Dg_{33} & 0 \\ I & J & (1/C)g_{23}/g_{33} & 1/C \end{pmatrix}, \tag{2.51}$$

where we define

$$E = g^{01}g^{12} - g^{11}g^{02}, \tag{2.52a}$$

$$F = g^{01}g^{02} - g^{00}g^{12}, \tag{2.52b}$$

$$G = g^{01}g^{13} - g^{11}g^{03}, \tag{2.52c}$$

$$H = g^{01}g^{03} - g^{00}g^{13}, \tag{2.52d}$$

$$I = \frac{B^2g^{00}}{C}\left(G + \frac{Eg_{23}}{g_{33}}\right), \tag{2.52e}$$

$$J = \frac{B^2g^{00}}{C}\left(H + \frac{Fg_{23}}{g_{33}}\right), \tag{2.52f}$$

With this matrix we can for example write the orthonormal frame fluid 4-velocity components in terms of the coordinate basis components: $u^{\hat{\mu}} = M^{\hat{\mu}}{}_\nu u^\nu$.

These formulas hold for the interfaces of constant x^2 (respectively x^3) under one (re-

spectively two) applications of the cyclic permutation $1 \mapsto 2 \mapsto 3$ in all indices, including the rows of $M^{\mu}{}_{\hat{\nu}}$ and the columns of $M^{\hat{\mu}}{}_{\nu}$.

One important detail to be noted is that the constant-coordinate interface across which we desire fluxes will generally appear to be moving in the orthonormal coordinates. To see this, consider an interface of constant x^1 and how it appears in orthonormal coordinates $(\hat{t}, \hat{x}, \hat{y}, \hat{z})$. According to (2.49) we know

$$x^1 = Ag^{01}\hat{t} + B(g^{01}g^{01} - g^{00}g^{11})\hat{x}. \tag{2.53}$$

Given that x^1 is constant, differentiation of (2.53) tells us the interface velocity can be calculated as

$$v^{\hat{x}} \equiv \frac{\mathrm{d}\hat{x}}{\mathrm{d}\hat{t}} = -\frac{Ag^{01}}{B(g^{01}g^{01} - g^{00}g^{11})} = \frac{g^{01}}{\sqrt{g^{01}g^{01} - g^{00}g^{11}}}. \tag{2.54}$$

In the language of the 3+1 formalism with lapse α, shift components β^i, and contravariant 3-metric components γ^{ij}, we can write $v^{\hat{x}} = \beta^1/\alpha\sqrt{\gamma^{11}}$, in agreement with Pons et al. (1998). That is, a nonzero shift in the direction of interest results in a nonzero interface velocity in the chosen orthonormal frame.

As this transformation yields physical quantities as measured by a normal observer, it can be applied even inside the ergosphere or, in the case of horizon-penetrating coordinates, inside the event horizon. This would not be the case if for example we attempted to transform into the frame of a stationary observer, since the spacetime might not everywhere admit timelike stationary worldlines. The transformation can be singular at coordinate singularities such as a polar axis, but such cases must be handled carefully no matter what method is being employed.

2.3.4 Riemann solver

The Riemann problem is that of determining time-averaged fluxes across an interface given constant left and right states. A wide variety of exact and approximate solvers have been developed. For example, Roe solvers (Roe, 1981) solve a full linearized set of equations,

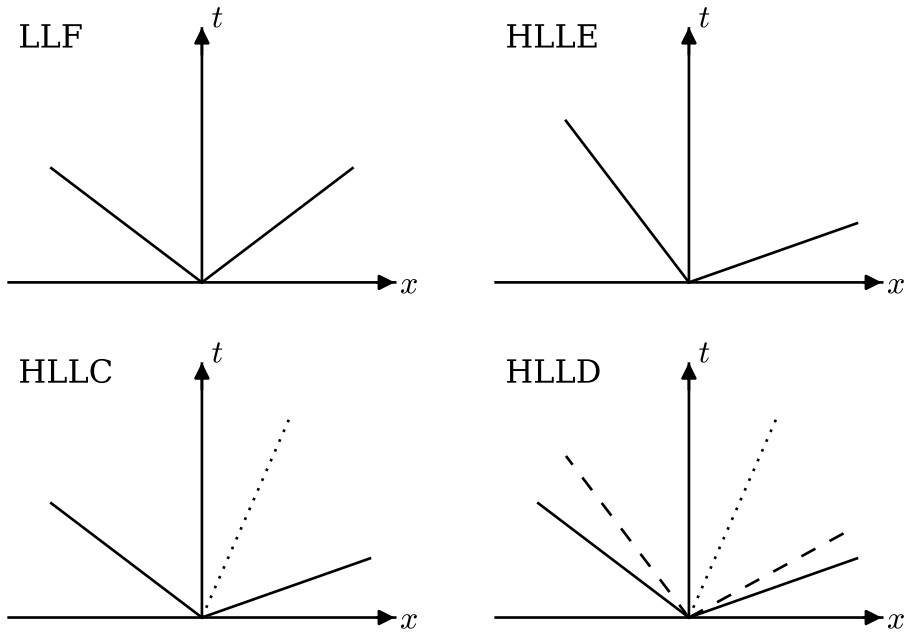

Figure 2.3: Schematic 1+1 spacetime diagrams of different HLL Riemann solvers' wavefans. The solid diagonal lines indicate the fastest moving waves in either direction, with LLF assuming the speeds are the same. The dotted lines in HLLC and HLLD represent the contact discontinuity, and the dashed lines in HLLD represent the Alfvén waves. The ordering of the waves shown is accurate, though waves can be degenerate. Fluxes are determined along the $x = 0$ surface for interfaces at rest, which can be in any of the 3, 4, or 6 regions of the wavefan. Slopes are qualitative; in particular all wavespeeds are subluminal.

obtaining high accuracy at the cost of performance.

An alternative is the HLL family of Riemann solvers (§1.3.1), based on the work of Harten et al. (1983), which have the advantage of being computationally simple while guaranteeing positivity (i.e. returning physically admissible intermediate states given physically admissible inputs). All HLL Riemann solvers take as input left and right states as well as leftgoing and rightgoing signal speeds, which we calculate as either the sound speeds in hydrodynamics or the fast magnetosonic speeds in MHD. The calculation of special-relativistic MHD wavespeeds, which are needed even in general relativity in the transformed frame, is detailed in Appendix B.

The HLL solvers we consider are introduced in §1.3.1. Figure 2.3 illustrates how the different solvers treat the internal shock structure.

Owing to the frame transformation, the same Riemann solvers can be used in both special and general relativity. As noted in §2.3.3, however, one must be careful to consider the Riemann problem with a moving interface. This is done by constructing the waves and intermediate states as usual and then determining the relevant portion of the wavefan according to where the generally nonzero interface velocity (2.54) falls relative to the wavespeeds. In the schematic sense of Figure 2.3, the fluxes sought are not necessarily along the vertical lines, but could be along angled lines also passing through the origin.

2.3.5 Inverse Frame Transformation

The transformation (2.49) back to global coordinates will have a nonzero component $M^1{}_{\hat{t}}$, in addition to the ever-present $M^1{}_{\hat{x}}$, whenever there is a moving interface. This necessitates having not just the components $T^{\hat{x}\hat{\nu}}$ but also $T^{\hat{t}\hat{\nu}}$ in the orthonormal frame. That is, we need the conserved state in the appropriate region of the wavefan, not just the associated fluxes. In practice even the more complicated Riemann solvers solve for the intermediate states by finding primitive variables, from which one can deduce consistent conserved states as well as fluxes.

The general procedure for obtaining fluxes proceeds as follows. First, the Riemann problem consisting of left and right states, extremal wavespeeds, and an interface velocity is transformed into the appropriate locally Minkowski frame. Next, the Riemann solver is used to solve for the wavefan, including the speeds of internal waves as well as the conserved quantities and their associated fluxes normal to the interface in each region separated by the waves. The appropriate region is then selected based on where the interface velocity falls inside the wavefan. Finally, the fluxes and conserved quantities in that region are transformed back to fluxes in the original coordinate system.

2.3.6 Constrained Transport

We employ a CT update of the face-centered magnetic fields in order to maintain the divergence-free constraint without resorting to divergence cleaning. The fundamentals of

CT were in fact developed within the context of general relativity (Evans and Hawley, 1988). Our implementation differs slightly from that original description, for example by having cell-centered rather than face-centered velocities. Ours is a simple extension of the algorithm detailed in Gardiner and Stone (2005) to relativity, and we summarize it here.

The goal of the algorithm is to determine edge-centered electric fields to use in the update (2.41) consistent with the fluxes returned by the Riemann solver. For concreteness, we will show how we calculate $E^3_{i-1/2,j-1/2}$, dropping all time superscripts. The algorithm proceeds in five steps:

1. Obtain the four face-centered electric fields $E^3_{i-1/2,j(-1)}$ and $E^3_{i(-1),j-1/2}$ from the Riemann solver.

2. Calculate the four cell-centered electric fields $E_{i(-1),j(-1)}$.

3. Find the eight electric field gradients

$$(\partial_1 E^3)_{i-1/4(-1/2),j(-1)} \text{ and } (\partial_2 E^3)_{i(-1),j-1/4(-1/2)}.$$

4. Upwind the electric field gradients to obtain the four gradients

$$(\partial_1 E^3)_{i-1/4(-1/2),j-1/2} \text{ and } (\partial_2 E^3)_{i-1/2,j-1/4(-1/2)}.$$

5. Combine the face-centered fields and upwinded gradients in a specific way to form the edge-centered electric field $E^3_{i-1/2,j-1/2}$.

The relevant quantities are laid out schematically in Figure 2.4.

For the first step, we simply note that the electric fields in the 3-direction are nothing more than the $1,2$-components of the dual of the electromagnetic field tensor (2.3) and are thus the fluxes returned by the Riemann solver in the x^1- and x^2-directions (after transforming back to global coordinates). In particular, the electric fields $E^3_{i-1/2,j(-1)} = -(*F)^{21}_{i-1/2,j(-1)}$ are the negatives of the x^1-fluxes of B^2. Similarly, we use the positive x^2-fluxes of B^1 to obtain $E^3_{i(-1),j-1/2} = (*F)^{12}_{i(-1),j-1/2}$.

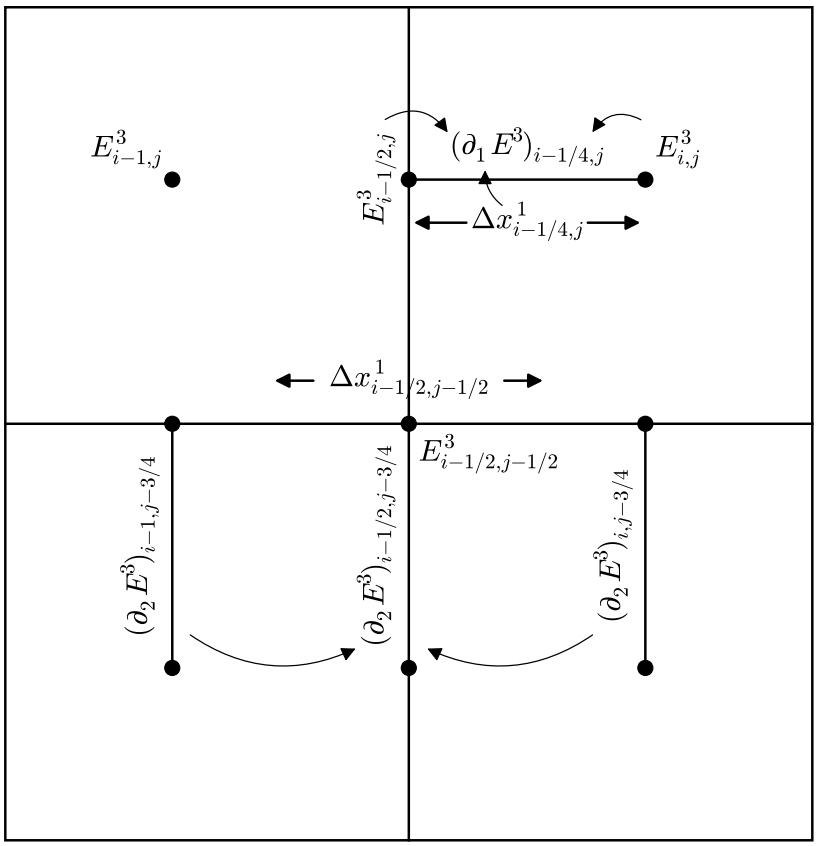

Figure 2.4: 2+0 spacetime diagram showing a timeslice of four cells. The x^3 direction is out of the plane of the figure, and indices k have been suppressed from all quantities. The top left cell shows what is known at the end of the Riemann problem: electric fields on the faces, as well as easily deducible cell-centered electric fields. The top right cell illustrates how a gradient is determined. The bottom two cells show gradients being upwinded back to the faces. Four face-centered fields, four upwinded gradients, and two edge-centered widths are combined to yield the edge-centered electric field $E^3_{i-1/2,j-1/2}$.

For the second step (which does not depend on the first), we calculate the cell-centered electric fields from the cell-centered velocities and the interpolated magnetic fields (e.g. (2.42)). This amounts to simply constructing the appropriate 4-vector components u^μ and b^μ and using (2.4) and (2.3).

In the third step we take gradients of electric field components from cell centers to faces, similar to (45) of Gardiner and Stone. In order to calculate such gradients, we require cell widths, though as we shall see these can be canceled from the final expression to a good approximation. Consider the gradient $\partial_1 E^3$ located between cell center i, j and face center $i - 1/2, j$, as in the top right cell of Figure 2.4. If we have

$$\Delta x^1_{i-1/4,j} = \int_C \sqrt{g_{11}}\, \mathrm{d}x^1 \tag{2.55}$$

with C the curve of constant x^0, x^2, and x^3 running from the cell center to the face, then we can take

$$(\partial_1 E^3)_{i-1/4,j} = \frac{E^3_{i,j} - E^3_{i-1/2,j}}{\Delta x^1_{i-1/4,j}}. \tag{2.56}$$

Similarly we can calculate the other seven such gradients.

Fourth, we shift the gradients onto the faces via upwinding. For example, as illustrated in Figure 2.4, we calculate

$$(\partial_2 E^3)_{i-1/2,j-3/4} = \begin{cases} (\partial_2 E^3)_{i-1,j-3/4}, & u^1_{i-1/2,j-1} > 0, \\ (\partial_2 E^3)_{i,j-3/4}, & u^1_{i-1/2,j-1} < 0, \end{cases} \tag{2.57}$$

according to the sign of the mass flux returned by the Riemann solver. This is (50) of Gardiner and Stone.

Finally, we compute the edge-centered value as is done in (41) of Gardiner and Stone. If we have an appropriate width $\Delta x^1_{i-1/2,j-1/2}$ defined similarly to $\Delta x^1_{i-1/4,j}$ – that is, by integrating $\sqrt{g_{11}}$ as x^1 varies over half a cell, as shown in Figure 2.4 – then the value we seek

is

$$E^3_{i-1/2,j-1/2} = \frac{1}{4}\Big(E^3_{i-1/2,j} + E^3_{i-1/2,j-1} + E^3_{i,j-1/2} + E^3_{i-1,j-1/2}\Big) + \frac{\Delta x^1_{i-1/2,j-1/2}}{4}\Big(E^3_{i-3/4,j-1/2} - E^3_{i-1/4,j-1/2}\Big) + \frac{\Delta x^2_{i-1/2,j-1/2}}{4}\Big(E^3_{i-1/2,j-3/4} - E^3_{i-1/2,j-1/4}\Big). \tag{2.58}$$

However, if the metric and grid spacing are varying smoothly, we can approximately cancel the Δx factors in (2.58) with those in (2.56).

2.3.7 Source Terms

In non-Cartesian coordinates one must consider not only fluxes but also geometric source terms when updating conserved quantities according to (2.39). As can be seen in (2.21), the continuity and magnetic field equations never have source terms but the energy-momentum equation generally does.

As noted in Gammie et al. (2003), with the adopted index placement on the stress-energy tensor the source terms for $T^0_{\ \mu}$ will vanish when the metric does not depend on x^μ. In particular, since the metric is stationary we never have geometric source terms in energy.

For the remaining momentum equations we evaluate the stress-energy components using the primitive variables at the appropriate timestep, taking the metric to have its cell-centered values. We then contract these components with the connection coefficients according to (2.35), resulting in the appropriate quantities for adding to the conserved variables.

Source terms for additional physics (e.g. heating and cooling functions) can be added to the conserved quantities along with the geometric terms.

2.3.8 Time Integrator

At the heart of the numerical implementation is the time integrator. We employ a temporally second-order van Leer integrator for all quantities. Given the conserved variables at time step

n, we first calculate the associated primitive variables (§2.3.1). Next, Riemann problems are set up (§2.3.2 and §2.3.3) and solved (§2.3.4), yielding fluxes. We also compute any necessary source terms from the primitives (§2.3.7).

The fluxes and source terms are used to update all conserved variables by half a timestep, including the magnetic fields (§2.3.6). The same procedure is repeated at time step $n+1/2$, except the resulting fluxes and source terms are used to update the step n state to step $n+1$. Schematically, we advance the grid by a single timestep via

$$C^{n+1/2} = C^n + \frac{\Delta t}{2} S^n + \frac{\Delta t}{2\Delta V} \sum_{\text{faces}} F^n \Delta A, \tag{2.59a}$$

$$C^{n+1} = C^n + \Delta t S^{n+1/2} + \frac{\Delta t}{\Delta V} \sum_{\text{faces}} F^{n+1/2} \Delta A. \tag{2.59b}$$

The van Leer integrator is TVD, and so it will not introduce spurious extrema.

Athena++ implements a variety of other temporal integration algorithms, including the TVD second-order and third-order Runge–Kutta methods (RK2 and RK3) of Shu and Osher (1988). The RK2 integrator is the same as Heun's method and we have found is no better than the van Leer method we implement. It suffers from the drawback that second-order spatial reconstruction in both substeps, whereas we can use first-order reconstruction for the first half step with the van Leer scheme and still obtain second-order convergence. RK3 and other higher-order schemes generally require extra storage of intermediate results, and for our present purposes we find second order in time to be sufficient. CT corner transport upwind (CTU) integrators have also been developed for MHD (Gardiner and Stone, 2005, 2008), but CTU methods require time-advanced variable estimates that are difficult to generalize to arbitrary coordinate systems.

2.3.9 Storage Requirements

At each cell, we must store 5 primitive variables and 5 conserved variables for each half step as required by the van Leer integrator. For MHD in 3 spatial dimensions we also store 3 each of face-centered and cell-centered magnetic fields and edge-centered, face-centered, and

cell-centered electric fields, all at each half step. This amounts to 20 numbers per cell in hydrodynamics, and 50 numbers per cell in MHD.

Any finite volume code must have access to cell volumes, interface areas, edge lengths (in the case of MHD), and cell widths (in order to determine maximum stable timesteps). In general relativity, the values of $g_{\mu\nu}$, $g^{\mu\nu}$, $\Gamma^{\sigma}_{\mu\nu}$, and the transformation matrices (2.49) and (2.51) are also required at various stages of the integration. Fortunately the assumption of a stationary metric means these terms can be precomputed. Moreover, symmetries of the metric combined with the restriction that cells be divided by interfaces of constant coordinate values often results in separable formulas for these values, obviating the need for 3D storage arrays.

For example, in the Kerr–Schild coordinates used in Chapters 3 and 4 and defined in §C.3.2, the volume (2.29) of a cell bounded by coordinates $r_\pm$, $\theta_\pm$, and $\phi_\pm$ is

$$\Delta V = \frac{1}{3}(r_+ - r_-)(\cos\theta_- - \cos\theta_+)(\phi_+ - \phi_-)\Big(f(r_-, r_+) + a^2 f(\cos\theta_-, \cos\theta_+)\Big), \qquad (2.60)$$

where $f(x, y) = x^2 + xy + y^2$. The three parenthesized prefactors and the two f terms can each be stored in 1D arrays, where ΔV is computed using just a small handful of additions and multiplications when needed. In fact, no more than 2D storage is required for any metric discussed in this work. For Kerr–Schild, the most complicated metric discussed here, we store $22N_r + 32N_\theta + 9N_\phi + 20N_rN_\theta$ values on an $N_r \times N_\theta \times N_\phi$ grid. This may seem substantial in comparison to the $4N_x+4N_y+4N_z$ coordinate-related values stored when using the Minkowski metric, but still pales in comparison to the (20–50)$N_1N_2N_3$ values holding evolving variables. Thus the memory footprint of our code is substantially smaller than that of codes that store all geometric factors without regard for symmetry or separability.

All the necessary expressions are cataloged in Appendix C for a number of different coordinate systems.

Chapter 3

Testing Athena++

As with any large code, there are many places where errors can be introduced, either in the underlying algorithm or the actual implementation. While some errors may lead to crashes or other such obvious symptoms, there exists a worrying class of issues wherein the code returns reasonable results that are nonetheless wrong. A battery of quantitative tests can be used to combat such insidious bugs.

The strategy, then, is to have the code solve problems that altogether exercise all of the components of the procedure. For example, small perturbations (small enough to be able to neglect nonlinear terms in the evolution) can be initialized matching the eigenfunctions of the linearized equations. These linear waves should hold their shape over time, with the errors decreasing with increasing resolution. Success indicates reconstruction and temporal integration are almost certainly operating correctly. Similarly, the nonlinear terms are important in the Riemann solver and variable inversion, and running shock tubes should detect problems with these parts of the code.

A suite of selected tests is described in the following sections. Because much of the machinery for our code is the same in both general and special relativity, several of our tests are special relativistic in nature. Some are run in both special- and general-relativistic settings, and others are purely general relativistic.

3.1 Linear Waves

Linear wave convergence is a strong test of a code, and so we present convergence results for special-relativistic hydrodynamics and magnetohydrodynamics (MHD).

Choose a background constant state and consider only perturbations in the x-direction. If we write the time evolution equations in terms of a matrix A depending on the background and the vector of primitives $\mathbf{P}$,

$$\partial_t \mathbf{P} + A \partial_x \mathbf{P} = 0, \tag{3.1}$$

then the linear waves are perturbations $\delta\mathbf{P}$ that are eigenvectors of A.

In order to quantify the error, we evolve a wave for exactly one period and calculate the error for each primitive quantity q as the L^1 norm of the difference between the initial and final states:

$$\epsilon_q = \frac{1}{N} \sum_{i=1}^{N} |q_i^{\text{init}} - q_i^{\text{fin}}|. \tag{3.2}$$

We then take the overall error to be the root-mean-square value of the set of errors ϵ_q.

The eigenvectors for hydrodynamics are given in Falle and Komissarov (1996, their Appendix A). We choose a background state consisting of $\rho = 4$, $p_{\text{gas}} = 1$, and $v^i = (0.1, 0.3, -0.05)$, with ratio of specific heats $\Gamma = 4/3$. The resulting wavespeeds are approximately -0.30, 0.1, and 0.47. We evolve the entropy and rightgoing sound wave, using both HLLE and HLLC Riemann solvers. The results are shown in Figure 3.1. Both methods converge at second order in the number of cells, as is expected since the reconstruction and time integration are both second order. Note that the error in the entropy wave is 1.7 times larger using HLLE compared to HLLC for these parameters, while the solvers have the same error for the sound wave. This reflects the fact that when inside the wavefan, using just the sound speeds leads to averaging over too large a domain of dependence, diffusing the solution.

For MHD we take the same hydrodynamic background state and add a magnetic field with components $B^i = (2.5, 1.8, -1.2)$, resulting in all wavespeeds (approximately -0.71, -0.55, -0.21, 0.1, 0.35, 0.54, and 0.81) being well separated. The eigenvectors are given

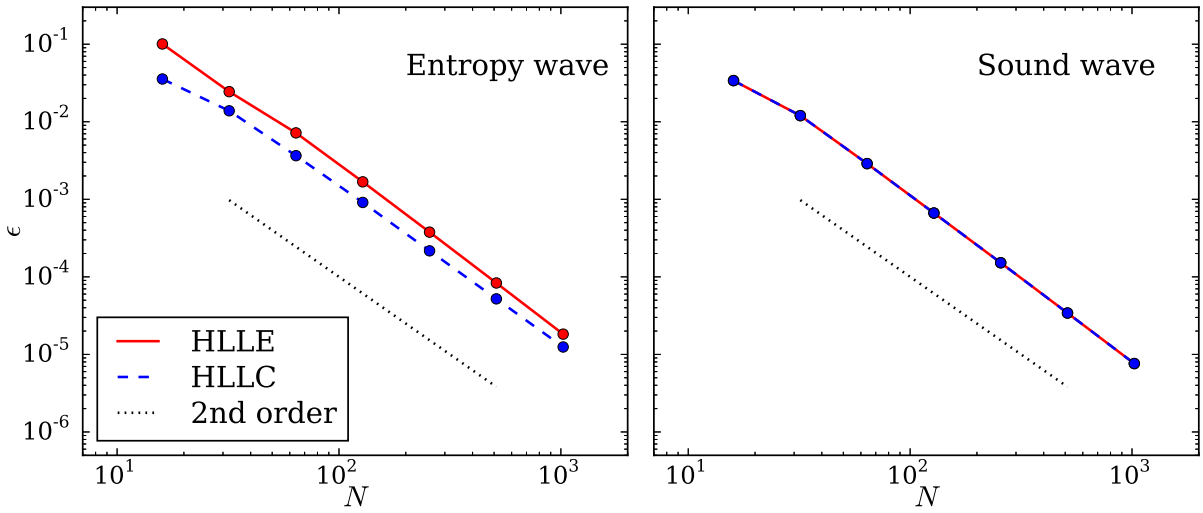

Figure 3.1: Convergence tests for special-relativistic hydrodynamics. Both Riemann solvers converge at second order, with HLLC having a smaller overall error compared to HLLE for the entropy wave. LLF results are indistinguishable from HLLE.

in Antón et al. (2010) as their (46) (entropy), (65) (Alfvén), and (71) (magnetosonic). The results for the four rightmost waves are given in Figure 3.2. Again we have second order convergence, and again the more advanced Riemann solver has smaller errors for the internal waves of the Riemann fan. Here HLLD's error is a factor of 3.1, 1.7, and 1.4 times lower than that of HLLE for the entropy, slow magnetosonic, and Alfvén waves, respectively.

We can also perform these same tests in a way that exercises the general-relativistic portions of the code. Consider the "tilted" coordinates (T, X, Y, Z) related to Minkowski coordinates (t, x, y, z) by

$$T = \frac{t + ax}{\sqrt{1 + a^2}}, \tag{3.3a}$$

$$X = \frac{x - at}{\sqrt{1 + a^2}}, \tag{3.3b}$$

$$Y = y, \tag{3.3c}$$

$$Z = z \tag{3.3d}$$

for some parameter a, $|a| < 1$. This coordinate system has its time axis corresponding to an observer moving with velocity a with respect to the Minkowski frame, while the surface of

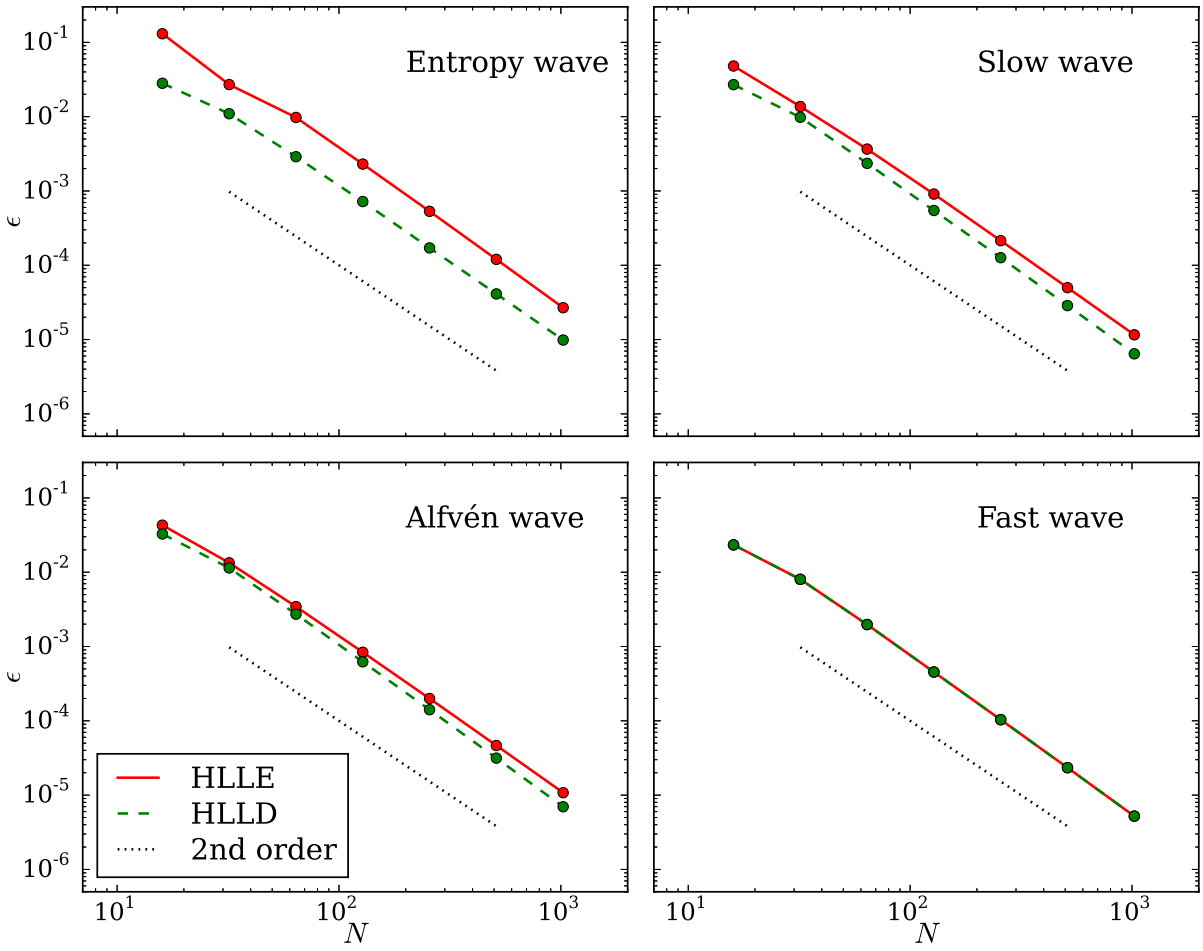

Figure 3.2: Convergence tests for special-relativistic MHD. Both Riemann solvers converge at second order, with HLLD having a smaller overall error compared to HLLE for all but the fast wave. LLF results are indistinguishable from HLLE.

constant time corresponds to that of an observer with velocity $-a$. Visually this corresponds to rigidly rotating the tx-axes into the TX-axes. While the connection coefficients and thus source terms vanish in this coordinate system, the TX metric coefficients are nonvanishing and so there are off-diagonal components in the transformation matrices (2.49) and (2.51). Additionally, the interface velocity (2.54) is nonzero. (See §C.1.4 for the details of this coordinate system.)

We can rerun the previous convergence tests in this new coordinate system. We use the same background states, with the given numbers understood to be in Minkowski coordinates. The linear solution is known throughout spacetime, and we initialize the perturbation over a single wavelength on the surface $T = 0$. This is evolved until $T = (1 + a\lambda)/|\lambda - a|$, which corresponds to one wavelength crossing the domain at Minkowski 3-velocity λ. For most waves we choose $a = 0.1$. For the entropy waves we choose $a = 0.05$, since if a matches λ the HLLC and HLLD solvers will exactly capture the wave, placing our comparison (3.2) at the machine precision noise floor even for low resolutions.

The results for hydrodynamics and MHD are given in Figures 3.3 and 3.4. For the hydrodynamic entropy wave, the HLLC solver's error is lower than that of HLLE by a factor of 2.8. In MHD, the HLLD solver's error beats that of HLLE by 4.8, 1.8, and 1.4 in the entropy, slow, and Alfvén waves. We note that in these particular tests the advanced Riemann solvers outperform HLLE by larger margins in general relativity compared to special relativity.

3.2 Shock Tubes

Another class of one-dimensional special-relativistic tests is that of shock tubes. While the above wave tests focus on linear perturbations, shock tubes emphasize the nonlinear behavior of the equations, as well as the ability to resolve discontinuities in the solutions.

We show the results for seven shock tubes from Komissarov (1999) (cf. Gammie et al., 2003, their Figures 4 and 5) in Figure 3.5. These results in black all use the HLLD solver

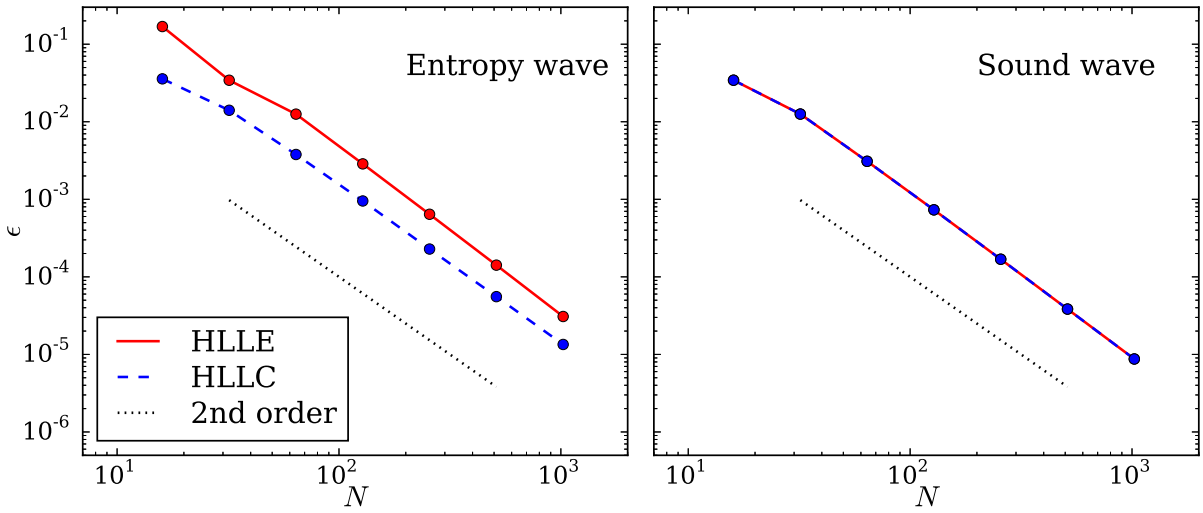

Figure 3.3: Convergence tests for general-relativistic hydrodynamics. Both Riemann solvers converge at second order, with HLLC improving upon HLLE for the entropy wave, as in Figure 3.1.

with 220 grid cells in the interval $[-1.1, 1.1]$, corresponding to the same fixed resolution used by Gammie et al. in all cases. The underplotted red line is a reference solution computed with ten times the resolution. The Courant–Friedrichs–Lewy (CFL) number is 0.8 except in the cases of the fast shock (0.2) and the collision (0.5).

Some of the same numerical artifacts observed by Komissarov can be seen in our results. For example, we see ringing upstream of the slow shock, a spurious bump at the low-density end of the switch-off rarefaction wave, and overshooting at the low-density end of the shock tube 2 rarefaction wave. In the case of shock tube 1 we do not observe the overshooting in density at the base of the density rarefaction wave, and we have greatly diminished overshooting at the top of the velocity discontinuity, despite using the same resolution as Komissarov. That HLLD in particular does well at capturing shocks is not surprising; Mignone et al. (2009) note this in comparing different Riemann solvers. We conclude that we are correctly capturing the nonlinear behavior expected in this special-relativistic test suite.

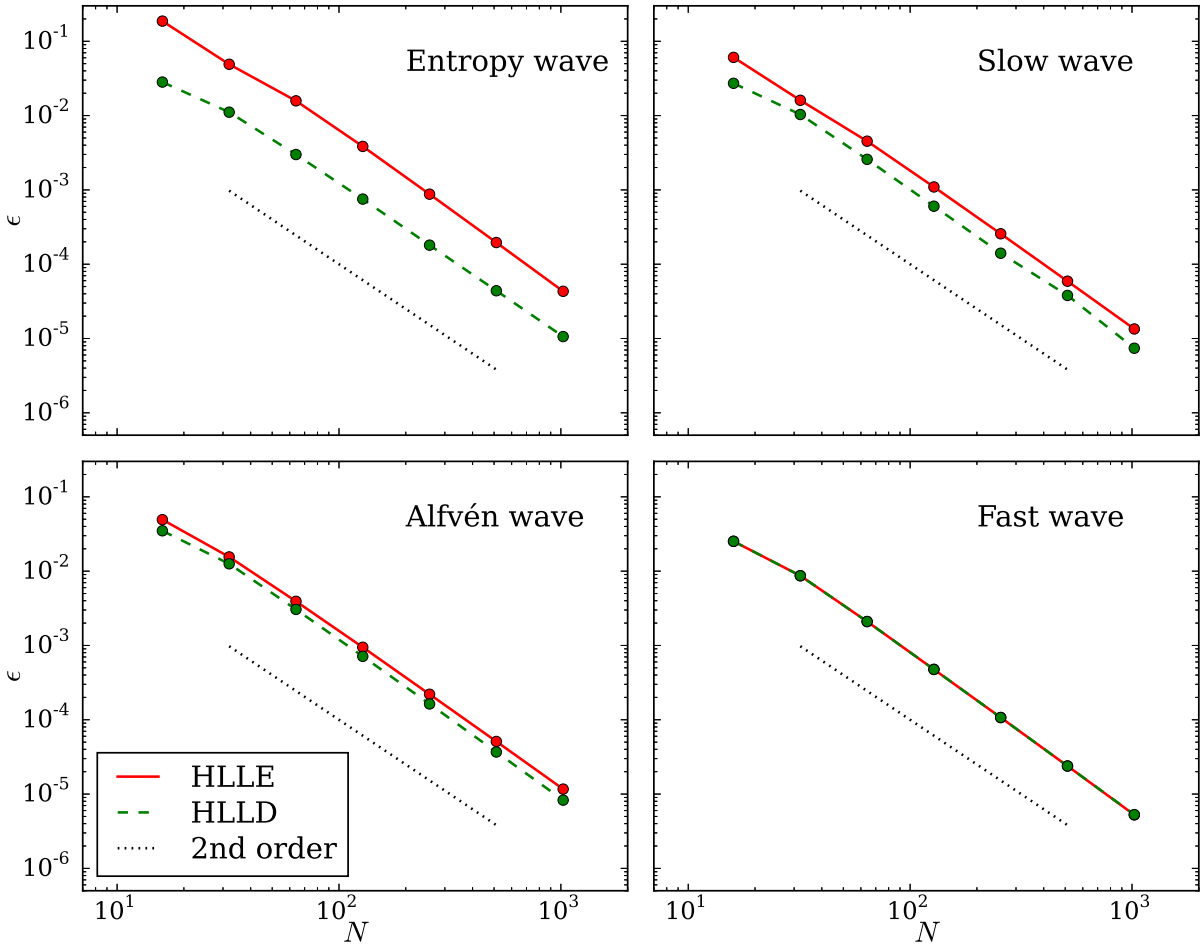

Figure 3.4: Convergence tests for general-relativistic MHD. Both Riemann solvers converge at second order, with HLLD improving upon HLLE for all but the fast wave, as in Figure 3.2.

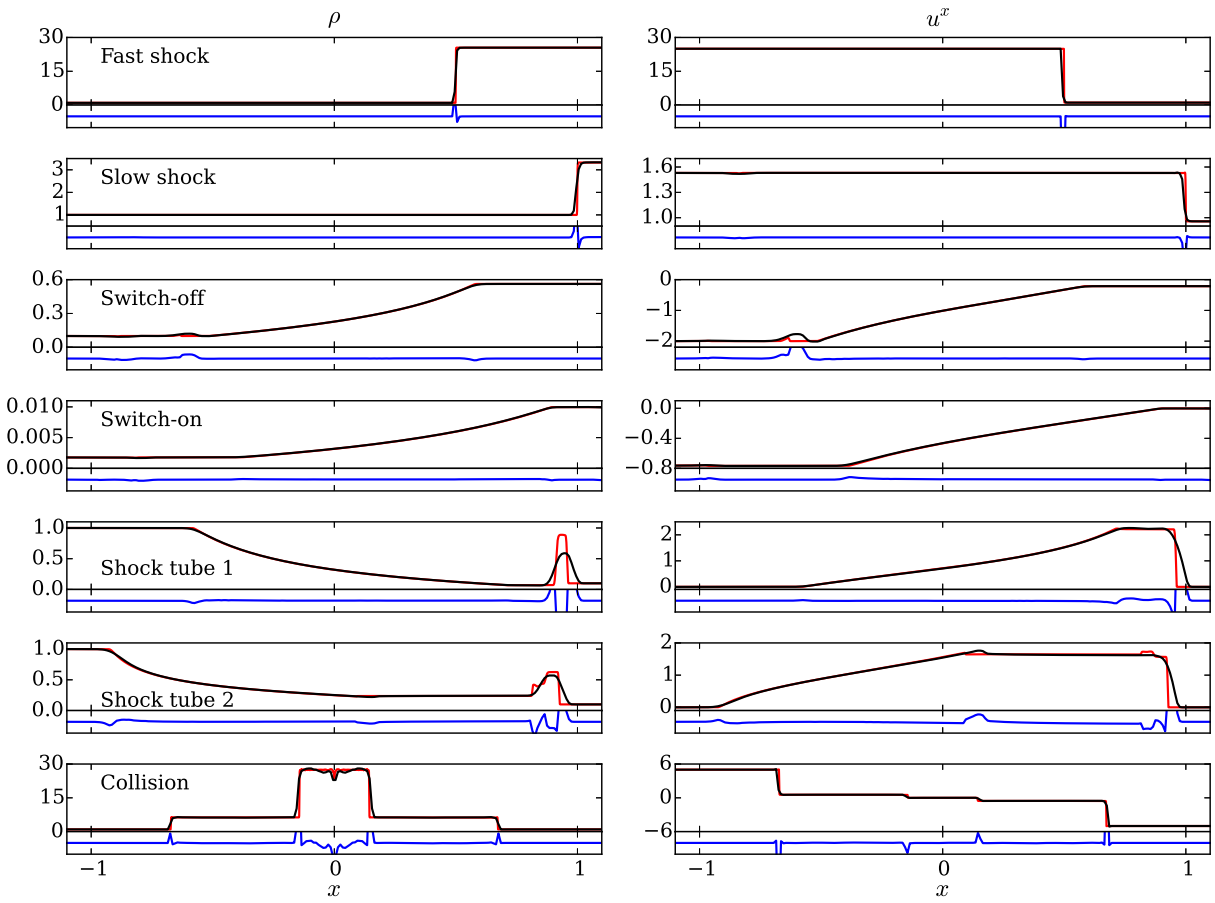

Figure 3.5: Komissarov shock tubes for special-relativistic MHD. Results are computed with HLLD on a grid of 220 (black) and 2200 (red) cells. The blue lines show the relative error between the two resolutions normalized by the maximum value attained by the variable over the domain, with the plot limits being ±10% in all cases.

3.3 Spherical Blast Wave

As a two-dimensional test of our general-relativistic formulation, we evolve a spherical blast wave. The initial conditions are uniform density ($\rho = 1$) with a central overpressure ($p_{\text{gas}} = 2.5$ inside a radius of 0.5; $p_{\text{gas}} = 0.1$ outside). We enforce periodic boundary conditions on a rectangular grid of width 4 and height 6. By time $t = 15$, the blast wave has intersected itself a number of times, inducing Richtmeyer–Meshkov instability features. The shock velocities here are not so high as to be in the regime where Richtmeyer–Meshkov instability is strongly suppressed, as discussed in Mohseni et al. (2014) and Zanotti and Dumbser (2015).

In order to test nontrivial spatial geometry, we evolve the system not only in Minkowski coordinates but also in "snake" coordinates that vary sinusoidally with Minkowski position. Explicitly, given Minkowski (t, x, y, z), define

$$x^0 = t, \tag{3.4a}$$

$$x^1 = x, \tag{3.4b}$$

$$x^2 = y + a\sin(kx), \tag{3.4c}$$

$$x^3 = z, \tag{3.4d}$$

where a and k are free parameters. If we define the parameter

$$\delta = ak\cos(kx^1), \tag{3.5}$$

the metric becomes

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & \sqrt{1+\delta^2} & -\delta & 0 \\ 0 & -\delta & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}. \tag{3.6}$$

The transformations (2.49) and (2.51) differ from the identity via $M^2_{\ \hat{x}} = -M^{\hat{y}}_{\ 1} = \delta$, and there is a source term due to Γ^2_{11}. (More details can be found in §C.1.3.)

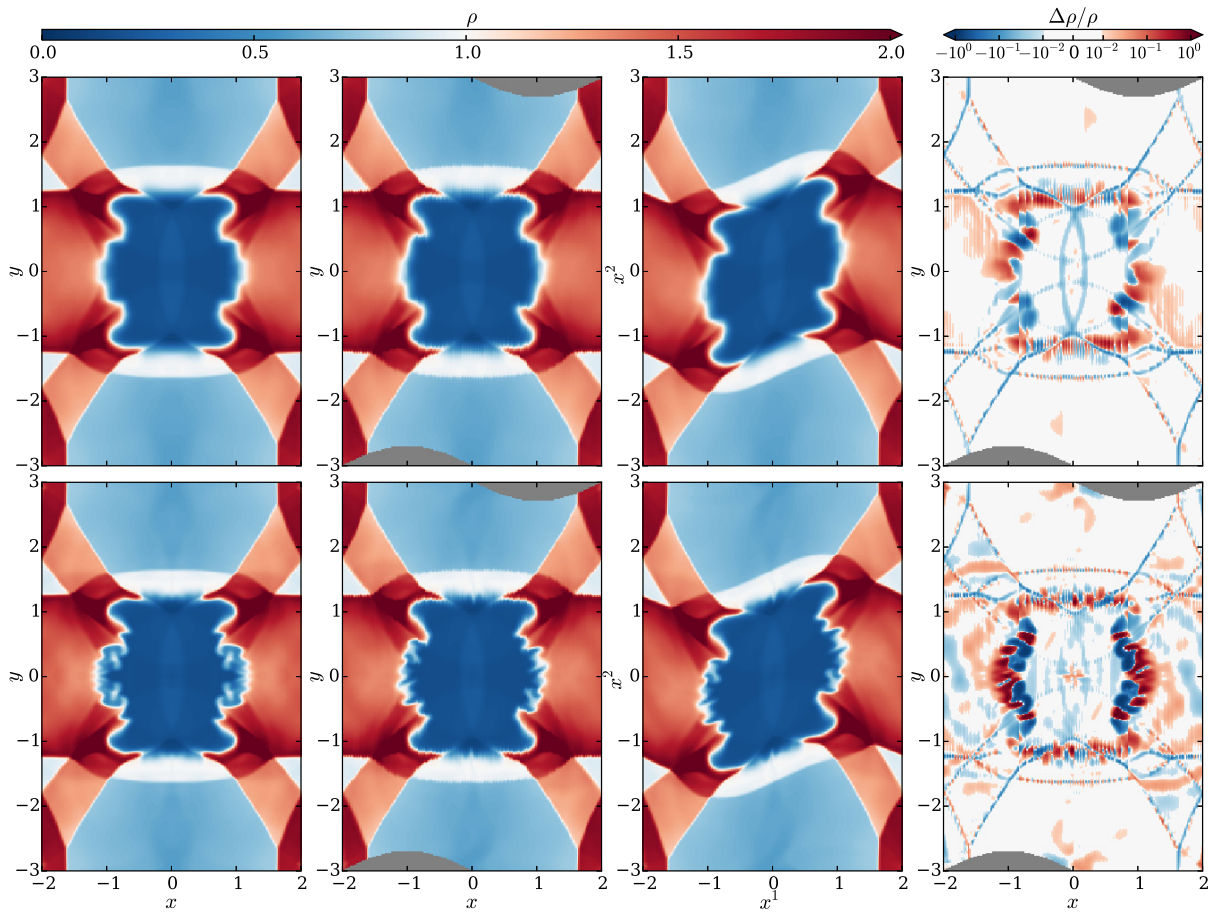

Figure 3.6: Density ρ for the spherical blast wave at $t = 15$ in Minkowski coordinates (left), snake coordinates transformed back to Minkowski (center left), and snake coordinates (center right). The right panels show the fractional difference $\Delta\rho/\rho$ between the first and second columns, where $\Delta\rho = \rho_{\text{snake}} - \rho_{\text{Minkowski}}$ and white regions have a fractional difference of less than 1%. The top panels use HLLE, while the bottom use HLLC.

The hydrodynamical blast wave density is shown in Figure 3.6. The top row displays results using the HLLE Riemann solver, while the bottom row uses HLLC. The first column shows the results at $t = 15$ as computed in Minkowski coordinates on a 200×300 grid with a CFL number of 0.4. The second column shows the same results as calculated in snake coordinates and transformed back to Minkowski at the end of the calculation. Here we also use a 200×300 grid, with $a = 0.3$ and $k = \pi/2$. Lines of constant x^2 in snake coordinates parallel the sinusoidal boundary of the shaded region.

One can see that the solutions remain qualitatively unchanged when the coordinates are varied, despite the fact that the fluxes, the source terms, and even the conserved quantities

themselves vary between coordinate systems. To give a sense of how the internal representation of the fluid changes between coordinate systems, the third column of Figure 3.6 plots density in snake coordinates (x^1, x^2).

For a more quantitative comparison, the last column shows the fractional difference in density between the Minkowski solution and the snake solution transformed to Minkowski coordinates. Colored areas indicate regions where the fractional difference exceeds 1%. Differences are generally confined to cell-level variations in the locations of shock fronts, or else to the details of the Richtmeyer–Meshkov fingering. This latter effect is to be expected with an instability seeded with grid noise. We also note that the HLLC Riemann solver leads to more sharply defined fingers, whereas using just HLLE the structures are notably more diffuse.

As a further test, we add an initially uniform magnetic field ($B^i = (1, 1, 0)$) and rerun the blast wave, this time with the HLLD solver. In this strongly magnetized regime, the blast wave should be guided by the magnetic field lines, breaking its initial symmetry.

The MHD results are shown in Figure 3.7, with color denoting ρ and streamlines showing the magnetic field. The same layout is used here as in both rows of Figure 3.6. Both MHD calculations are done with HLLD on a 200×300 periodic grid with a CFL number of 0.4.

As expected, the field guides the blast wave and suppress the Richtmeyer–Meshkov instability. Moreover, the same results are found in both coordinate systems. The self-interaction of the wave after wrapping around the boundaries induces a slight curvature in the boundary of the rarefaction region and in the magnetic field lines, and even this is the same in the two coordinate systems.

3.4 Bondi Accretion

As a test of the code in a nontrivial spacetime, we model Bondi accretion onto a Schwarzschild black hole. That is, we seek steady-state solutions to radial flow onto a black hole, where

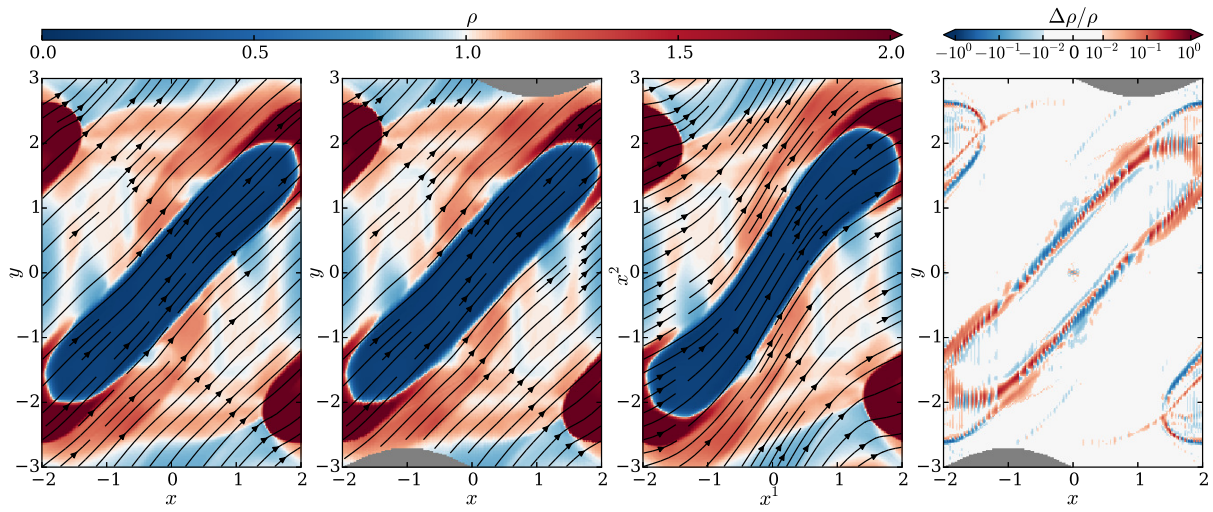

Figure 3.7: Similar to Figure 3.6 with MHD added. Density ρ (color) and magnetic field B^i (streamlines) are shown for the strongly magnetized blast wave at $t = 5$ in Minkowski coordinates (left), snake coordinates transformed back to Minkowski (center left), and snake coordinates (center right). The right panel shows the fractional difference between the first and second columns as in Figure 3.6. All panels use HLLD.

pressure can be nonnegligible.

The Bondi solution is derived in Hawley et al. (1984), and we summarize the pertinent formulas here. Begin with a black hole mass M, an adiabatic index Γ, an adiabat K, and a critical radius r_c. Define the polytropic index $n = 1/(\Gamma - 1)$. The critical velocity and temperature are then

$$u^1_c = -\sqrt{\frac{M}{2r_c}}, \tag{3.7}$$

$$T_c = \frac{n}{n+1} \cdot \frac{u^1_c u^1_c}{1 - (n+3)u^1_c u^1_c}. \tag{3.8}$$

Next we calculate the constants

$$C_1 = T_c^n u^1_c r_c^2, \tag{3.9}$$

$$C_2 = \Big(1 + (n+1)T_c\Big)^2 \left(1 - \frac{2M}{r_c} + u^2_c u^1_c\right). \tag{3.10}$$

The run of temperature with radius is given by solving the equation

$$\Big(1+(n+1)T\Big)^2\left(1-\frac{2M}{r}+\frac{C_1^2}{r^4T^{2n}}\right)=C_2. \tag{3.11}$$

There will generally be two roots to this equation. The lesser root is taken for $r < r_{\rm c}$, while the greater root is taken outside the critical point. We can then find the radial velocity according to

$$u^1=\frac{C_1}{r^2T^n}. \tag{3.12}$$

Finally, density and pressure are given by

$$\rho=\left(\frac{T}{K}\right)^n, \tag{3.13}$$

$$p_{\rm gas}=T\rho. \tag{3.14}$$

Note these formulas assume Schwarzschild coordinates.

We choose parameters $M=1$, $\Gamma=4/3$, $K=1$, and $r_{\rm c}=8$, and we simulate the region $3<r<10$, $\pi/4<\theta<3\pi/4$. In order to test MHD as well as pure hydrodynamics, we optionally include a purely radial monopole field with $B^r\propto 1/r^2$, normalized such that $b_\lambda b^\lambda/\rho=10$ at the inner radius. This corresponds to plasma $\beta=0.246$ at the critical point. The 2D domain is divided into N equally spaced cells in each direction. We initialize the Bondi solution and run the simulation until time $t=10$, computing the L^1 difference in $p_{\rm gas}$ between the beginning and end. We use the inner 3/4 of the grid in each direction, and normalize the errors ϵ by the L^1 norm of the initial $p_{\rm gas}$.

The steady-state solution is shown in Figure 3.8. The convergence is second order for both hydrodynamics using HLLC and MHD using HLLD, as shown in Figure 3.9. Since this is a stationary problem, convergence is not affected by the order of the integrator. However, it is limited by the order of reconstruction, and so these tests show we are correctly getting the full second-order convergence in space expected, even in a curved spacetime.

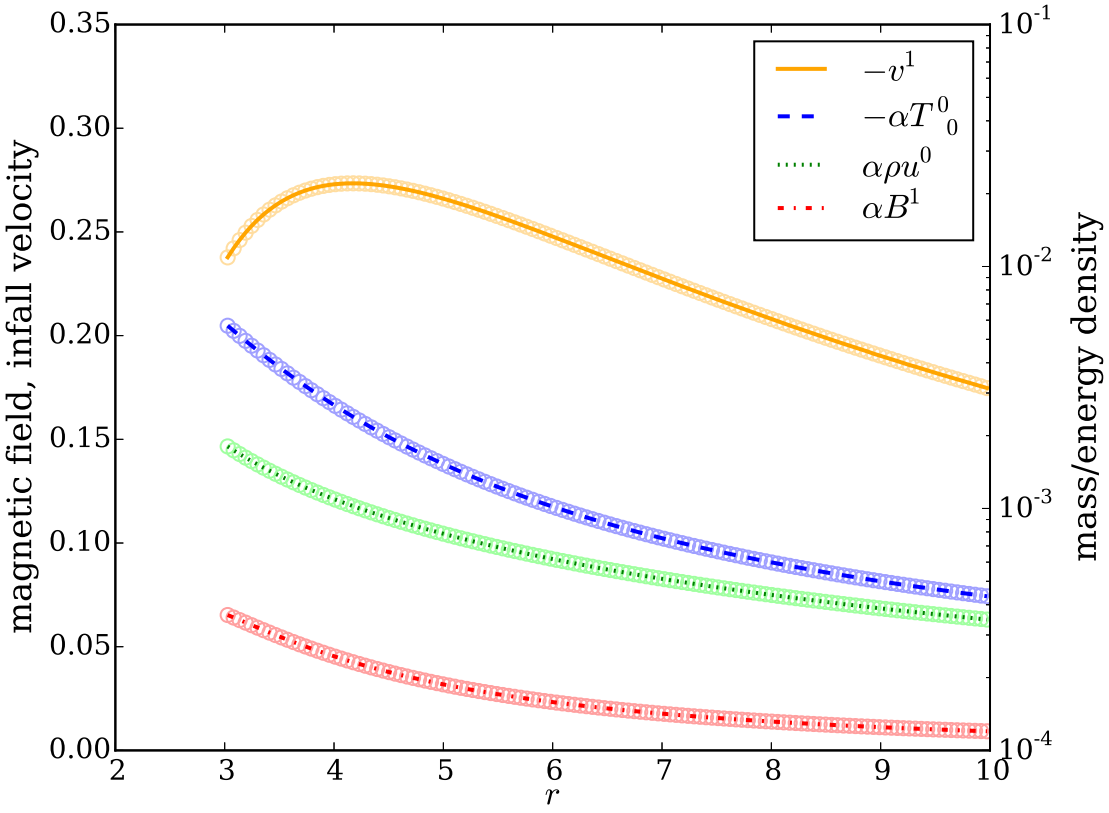

Figure 3.8: Magnetized Bondi flow in Schwarzschild coordinates with parameters as described in the text. Lines indicate the initial solution, while the circles indicate the values at each cell at $t = 10$. The conserved variables are scaled by the lapse α.

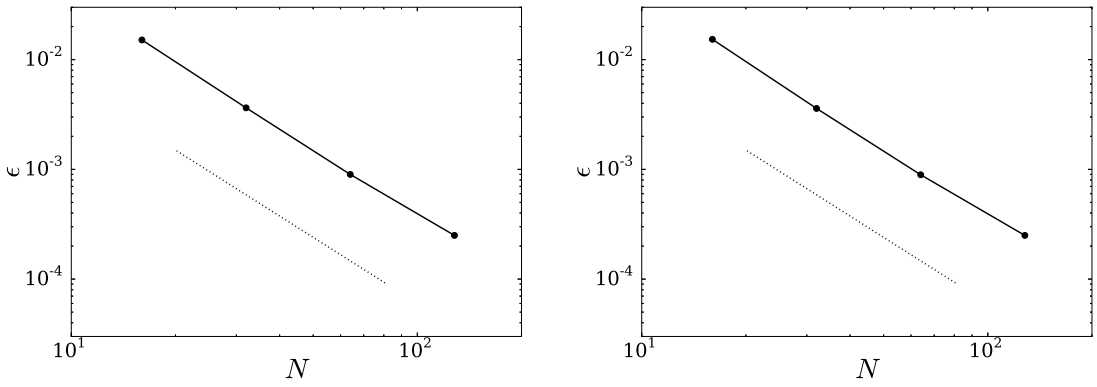

Figure 3.9: Errors in gas pressure for hydrodynamical (left) and MHD (right) Bondi flow on various $N \times N$ grids. In both cases we achieve second-order convergence, as indicated by the dotted lines.

3.5 Magnetized Equatorial Inflow

For a final test problem, we simulate the plunging region of a magnetized thin disk around a spinning black hole, as described in Gammie (1999). We choose the same values for the free parameters as in Gammie et al. (2003): $M = 1$, $a/M = 0.5$, $F_M = -1$, and $F_{\theta\phi} = 0.5$, with the flow matching onto the innermost stable circular orbit. The simulation covers the range $2 < r < 4$ (where the horizon is at $r = 1.86603$ and the innermost stable circular orbit is at $r = 4.23300$). We use the same Kerr–Schild coordinates $\{t, r, \theta, \phi\}$ as Gammie et al., defined in §C.3.2. Note that Kerr–Schild r and θ are the same as Boyer–Lindquist r and θ.

The flow is initialized to the steady-state solution, with gas pressure set to 10^{-10} times the density, and is run to time $t = 15$. A plot of the L^1 error in four quantities is shown in Figure 3.10. We find second-order convergence, demonstrating that the code is solving the correct equations even in a highly nontrivial spacetime. In particular, this test is performed with a frame-transforming HLLE Riemann solver, showing that the frame transformations work even when there is a large shift between the frames.

3.6 Torus Simulations

As a full demonstration of the code in a science application setting, we simulate a fluid torus around a spinning black hole. As proven in Fishbone and Moncrief (1976), there are steady-state solutions for hydrodynamical tori of an isentropic fluid with constant angular momentum orbiting in Kerr spacetime. Given a black hole mass and spin and a fluid equation of state, the solutions are defined by choosing the radius of the inner edge of the torus, $r_{\rm edge}$, and either the radius of the pressure maximum, $r_{\rm peak}$, or the angular momentum per unit mass, $\ell = u^t u_\phi$ (in Boyer–Lindquist coordinates).

A general-relativistic code should be able to maintain such a torus, balancing gravity, centrifugal force, and pressure gradients. We take a black hole with mass $M = 1$ and spin $a/M = 0.95$ and initialize a $\Gamma = 13/9$ fluid torus with $r_{\rm edge} = 3.7$ and $\ell = 3.85$ (so $r_{\rm peak} = 7.82$). The torus is evolved on a 2D grid with N_r cells in the radial direction

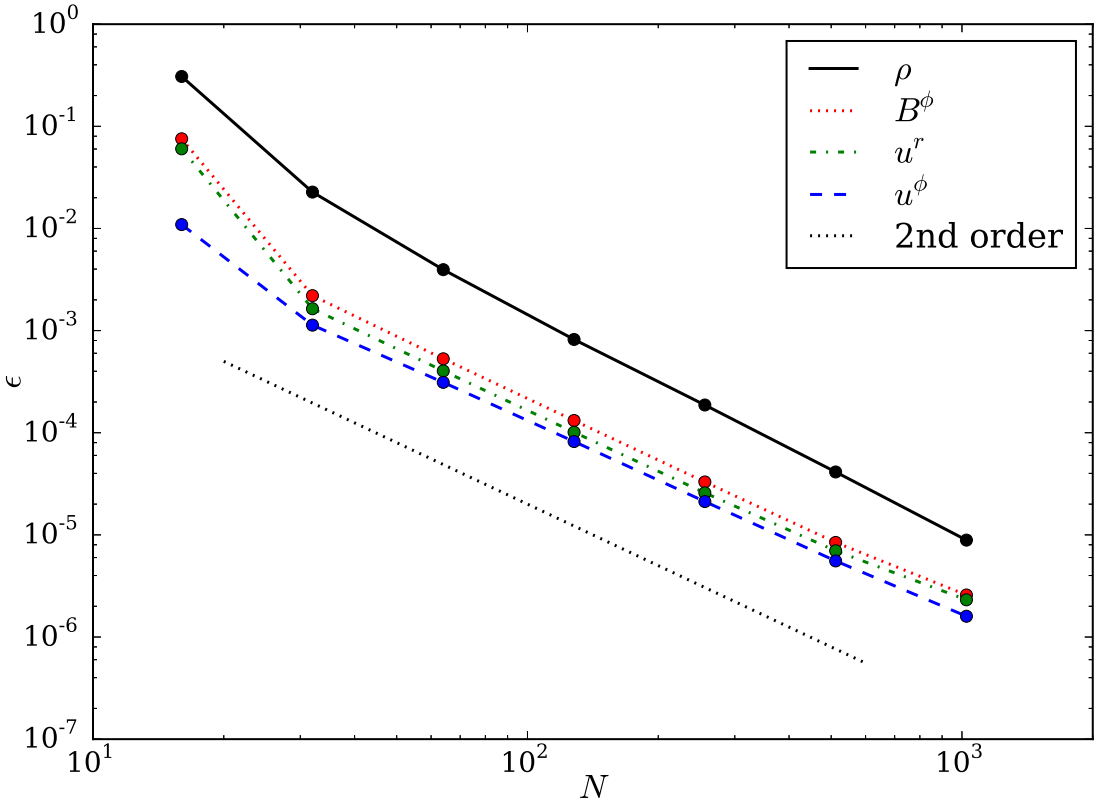

Figure 3.10: Errors in density, azimuthal magnetic field, and radial and azimuthal velocities as functions of the number of radial grid cells in the magnetized inflow test. As expected, all quantities converge at second order.

and $N_r/2$ cells in the polar direction using the aforementioned Kerr–Schild coordinates (see §C.3.2). Our grid extends radially from 0.98 times the outer horizon radius $M + \sqrt{M^2 - a^2}$ to $r = 20$, with geometric grid spacing such that each cell is 1.025 times as wide as its inner neighbor. We limit the polar angle to $\pi/4 < \theta < 3\pi/4$, excising parts of the necessarily non-stationary atmosphere that should not impact the torus proper. This simulation employs uniform spacing in θ. We impose outflow conditions on the inner and outer boundaries and reflecting boundary conditions on the constant-θ boundaries.

The simulation is evolved using the HLLC solver with a CFL number of 0.2 until a time $t = 1$, at which point the error accrued in ρ is calculated. Define the region T to be those cells for which ρ is initially at least 0.02 its peak value. We calculate the error ϵ over the region T (thus avoiding edge effects) in the L^1 sense:

$$\epsilon = \frac{\int_T |\rho_{\rm fin} - \rho_{\rm init}| \sqrt{-g}\, {\rm d}r\, {\rm d}\theta}{\int_T |\rho_{\rm init}| \sqrt{-g}\, {\rm d}r\, {\rm d}\theta}. \tag{3.15}$$

The errors are shown in Figure 3.11, where one can see that they are second order. This is the same test problem as done in Gammie et al. (2003, cf. their Figure 15) with the exception that we use a slightly different grid with no equatorial compression of the constant-θ surfaces.

Our torus is indeed able to maintain equilibrium, even over longer times. To demonstrate this we continue the 64×32 simulation to $t = 430$, corresponding to three orbital periods at the pressure maximum, using both HLLE and HLLC solvers. With the HLLE solver the error is $\epsilon = 0.030$, while it is $\epsilon = 0.021$ with HLLC.

We next consider a 3D Fishbone–Moncrief torus with an added magnetic field, upsetting equilibrium via the magnetorotational instability (MRI). Here we choose $M = 1$, $a/M = 0.9375$, $\Gamma = 13/9$, $r_{\rm edge} = 6$, and $r_{\rm peak} = 12$ (so $\ell = 4.28$). The grid is expanded to $0.98(M + \sqrt{M^2 - a^2}) < r < 40$ and $0 < \theta < \pi$ (with no artificial boundary at the poles), and it covers the entire azimuthal range. In one case we use the equivalent of $144 \times 128 \times 96$ cells in the radial, polar, and azimuthal directions, with geometric radial spacing (ratio 1.02397) and uniform spacing in the angles. Static mesh refinement is used to de-refine the polar

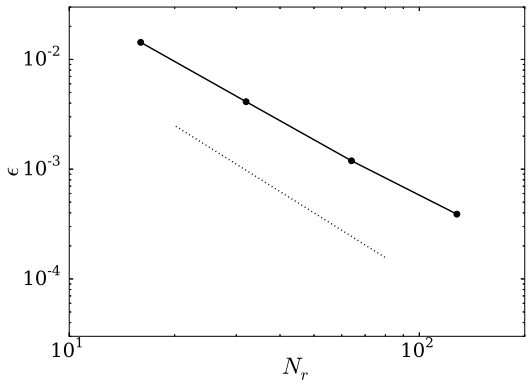

Figure 3.11: Errors in density for 2D Fishbone–Moncrief hydrodynamical tori on various $N_r \times N_r/2$ grids. Errors are defined by (3.15). The dotted line indicates second-order convergence.

regions outside the torus by a factor of 4 in each dimension, thus preventing the timestep from being limited by those regions. Explicitly, there are 80 cells equally spaced in the range $3\pi/16 < \theta < 13\pi/16$. For a higher resolution grid, we further double the number of cells in each dimension in the region $7\pi/32 < \theta < 25\pi/32$. This forms an effective $288 \times 256 \times 192$ grid.

To this torus we add a magnetic field with purely azimuthal vector potential $A^\phi \propto \max(\rho - 0.2, 0)$, calculated from a density normalized to unity at its maximum. The magnetic field is normalized such that the ratio of maximum gas pressure to maximum magnetic pressure (necessarily located at different points) is 100. These parameters are chosen to roughly match those of Shiokawa et al. (2012) except in not having an equatorially compressed grid like theirs. In particular, our low and high resolution runs have $\Delta\theta = 0.0245$ and $\Delta\theta = 0.0123$ in their respective most refined regions, while the $96 \times 96 \times 64$ and $144 \times 144 \times 96$ grids from Shiokawa et al. have $\Delta\theta = 0.00983$ and $\Delta\theta = 0.0065$ at the midplane.

We evolve the system using an LLF Riemann solver just as is used in Harm (Gammie et al., 2003). The simulation is run until a time $t = 12{,}000$. For comparison, a geodesic circular orbit at the pressure maximum has a period of $t = 267$. After several orbits the MRI sets in, causing turbulence that eventually disrupts the entire torus. The initial and final densities for the run are illustrated in Figure 3.12.

For a more quantitative analysis, we construct radial profiles of physical quantities.

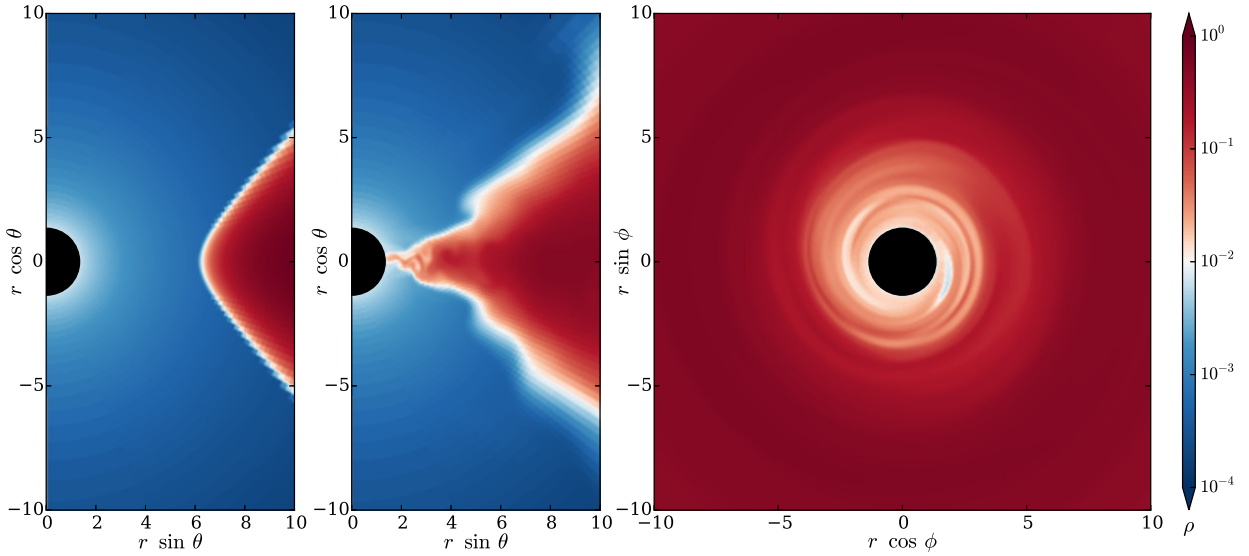

Figure 3.12: Density ρ of a 3D Fishbone–Moncrief MHD torus with effective resolution $288 \times 256 \times 192$. Shown are the initial conditions in the $r\theta$-plane (left), the final state at $t = 12{,}000$ in the same plane (center), and the final state in the $r\phi$-plane (right).

Following Shiokawa et al. (2012, cf. their (9) and (10)), we define a spherical shell average of a quantity q at a Kerr–Schild time t according to

$$\bar{q} = \frac{\int_0^{2\pi} \int_{\theta_{\min}}^{\theta_{\max}} q\rho\sqrt{-g}\,\mathrm{d}\theta\,\mathrm{d}\phi}{\int_0^{2\pi} \int_{\theta_{\min}}^{\theta_{\max}} \rho\sqrt{-g}\,\mathrm{d}\theta\,\mathrm{d}\phi}, \tag{3.16}$$

where in our case we take $(\theta_{\min}, \theta_{\max}) = (\pi/4, 3\pi/4)$ in order to exclude any effects of the atmosphere. We then average $\bar{q}$ over the time range 4000 to 12,000, noting the mean and standard deviation of values at each radius. This is done for plasma β, electron temperature $T = m_{\mathrm{p}} p_{\mathrm{gas}} / 2 m_{\mathrm{e}} \rho$ (called θ_{e} by Shiokawa et al.), and magnetic pressure $p_{\mathrm{mag}} = b^\lambda b_\lambda / 2$ in Figure 3.13. For a more direct comparison to Shiokawa et al. we also show the result of time-averaging $\bar{p}_{\mathrm{gas}} / \bar{p}_{\mathrm{mag}}$, the result of which is simply called β in that work.

The electron temperature roughly matches what Shiokawa et al. find, both in magnitude and in trend with radius, though they do not have the same kink at around $r = 5$–6 and our disk gets notably hotter inside the innermost stable circular orbit at $r = 2.04420$. We get the same shape as Shiokawa et al. for $p_{\mathrm{gas}} / p_{\mathrm{mag}}$, especially their low resolution "soft" polar run (cf. the upper left panel of their Figure 2; data reproduced in Figure 3.13). They

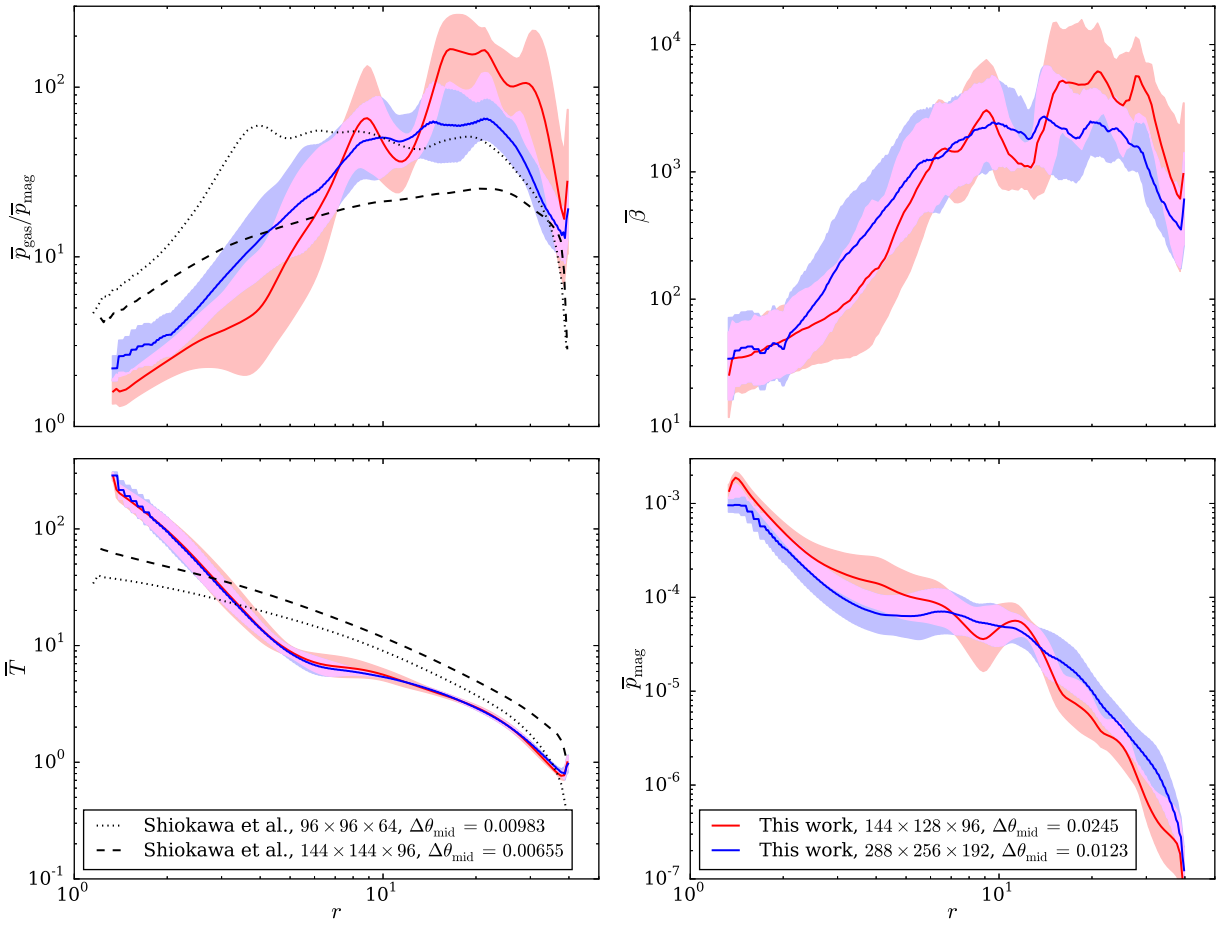

Figure 3.13: Spherically averaged, time-averaged profiles of various quantities in the 3D torus. These are the gas-to-magnetic pressure ratio (spherical averages taken before ratio, time averages take after; top left), plasma β (top right), electron temperature (bottom left), and magnetic pressure (bottom right). The left two panels show the same quantities as plotted in Figure 2 of Shiokawa et al. (2012), corresponding to their "soft" polar boundary condition. In all cases shaded regions indicate 1σ standard deviations from 800 samples at different times.

find $p_{\rm gas}/p_{\rm mag}$ to be less than about 60 in all cases, whereas our simulations get up to about 200 in the outer regions. Our ratio also approaches unity in the innermost regions, while theirs never drops below 4. The same shape is found when we perform spherical averages of β itself, rather than $p_{\rm gas}$ and $p_{\rm mag}$ separately, as shown in the upper right of Figure 3.13. While there are some differences, we also note that as we increase resolution our $p_{\rm gas}/p_{\rm mag}$ profile approaches theirs, and we expect even better agreement if we were to have midplane $\Delta\theta$ values as small as theirs.

Shiokawa et al. also examine azimuthal correlation lengths in the equatorial plane. We use the same definitions as they do (their (13) through (17)), which we summarize here. For a vector quantity with components q^μ define the time-dependent correlation function

$$R(r,\phi,t) = \frac{1}{r\Delta r\Delta\theta}\int_{\theta_{\rm min}}^{\theta_{\rm max}}\int_{r_{\rm min}}^{r_{\rm max}}\int_0^{2\pi}\delta q^\mu(r',\theta',\phi_0,t)\delta q_\mu(r',\theta',\phi_0+\phi,t)\,{\rm d}\phi\,r'\,{\rm d}r'\,{\rm d}\theta'. \quad (3.17)$$

The integrals run over a region of radial width $\Delta r = r_{\rm max} - r_{\rm min}$ of a single cell and polar width $\Delta\theta = \theta_{\rm max} - \theta_{\rm min}$ of two cells bordering the equatorial plane. The quantities δq^μ and δq_μ are deviations from averages of q^μ and q_μ over the domains of integration. For scalar quantities we simply use the integrand $\delta q(r',\theta',\phi_0,t)\delta q(r',\theta',\phi_0+\phi,t)$ instead. Define the time-averaged correlation function as

$$\bar{R}(r,\phi) = \int_{t_{\rm min}}^{t_{\rm max}}\frac{R(r,\phi,t)}{R(r,0,t)}\,{\rm d}t, \quad (3.18)$$

where we average from 4000 to 12,000. Then the correlation length λ at a given radius r is the value for which $\bar{R}(r,\lambda) = \bar{R}(r,0)/{\rm e}$.

We plot the runs of four correlation lengths in Figure 3.14, to be compared to Figure 4 of Shiokawa et al. (2012), from which select data has been replotted here. Compared to Shiokawa et al. our disks have larger correlation lengths, especially in the outer regions. This is a sign that we have not yet reached the fully saturated, turbulent state in that region. It is worth repeating, however, that our midplane resolution is not as high as in Shiokawa et al., and moreover our correlation lengths decrease as we increase resolution.

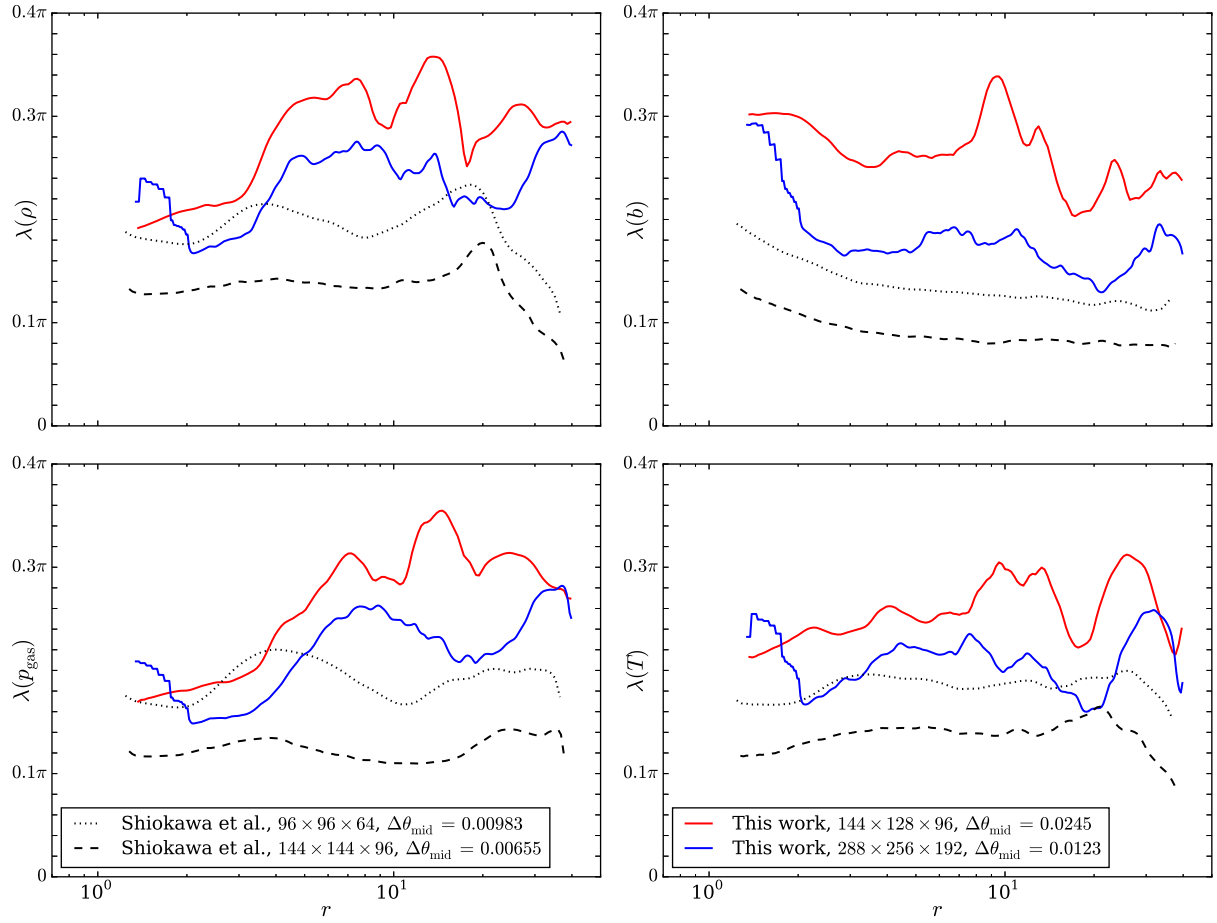

Figure 3.14: Azimuthal correlation lengths as functions of radius in the 3D torus runs. These are for density (top left), magnetic field (top right), gas pressure (equivalently internal energy; bottom left), and electron temperature (bottom right). The black lines show two select datasets from Shiokawa et al. (2012), both using the "soft" polar boundary condition.

Chapter 4

Application to Magnetically Arrested Disks

In selecting a first scientific application of Athena++, two considerations come to the forefront: an existing body of literature and room for improvement.

If we were to apply a new code to a situation previously lacking in simulations, results would be subject to a degree of doubt as to whether they are physical or peculiarities of the implementation at hand. Such forays into the unknown of course should be done, but we prefer to save them for later applications (see §5.2). As a result, we refrain from situations where the initial conditions and other parameters are largely unconstrained. Instead, we prefer a setting in which some numerical work has been done.

At the same time, merely reproducing old results, though arguably necessary for parts of the scientific method, is of little interest. It is important that we be able to say more about the situation afterward, whatever the result, than could have been said before our investigation.

We choose therefore to reexamine the situation of magnetically arrested disks (MAD). In particular, the question arises as to whether MAD states and their corresponding spin-powered jets survive in high-resolution simulations, where more instabilities are naturally present.

4.1 The Importance and Physics of MAD

Powerful relativistic jets have long been observed by astronomers. While such jets need not be associated with black holes – for example, pulsars produce jets – often there is no other source that can explain a particular jet. The connection between black holes – whose existence, recall, is only indirectly inferred, though with very strong evidence – and well-observed jets was solidified by Blandford and Znajek (1977), where a mechanism was outlined by which accretion could drive energy extraction from a spinning black hole along its polar axis.

In the Blandford–Znajek (BZ) mechanism, the magnetosphere of a spinning black hole can lend itself to extracting energy from the spin of the black hole itself. Under the assumption that pair production will lead to the presence of a diffuse plasma even if there was no matter originally, they approximate this magnetosphere in the force-free limit. This is what is modeled in our ideal MHD simulations in the low-β limit. The spin torques the field in such a way as to produce a net Poynting flux in a method analogous to that of Goldreich and Julian (1969) around pulsars. It was noted in Penrose (1969) that energy can be extracted to infinity at the cost of reducing the spin of the black hole using mechanisms that rely on the fact that a particle's energy as seen at infinity changes sign as it crosses into the ergosphere. Indeed the BZ mechanism can be cast in terms of energy-momentum transfer via the electromagnetic field between the ergosphere and distant regions. The key then for the BZ effect to occur is to have a spinning black hole with a low-β plasma hosting magnetic field ordered on large scales as was assumed in their model.

Under the right circumstances, one expects to be able to extract more energy from a black hole than is lost to it by infall, corresponding to an efficiency of greater than 100%. Tchekhovskoy et al. claim to see such efficiencies, attributing them to the simulations being in a magnetically arrested disk (MAD). As described in Narayan et al. (2003) and based on the ideas in Bisnovatyi-Kogan and Ruzmaikin (1974), MAD occurs when enough vertical magnetic flux has accumulated so as to prevent or at least restrict further mass accretion.

While a weak field threading a disk will destabilize it to the MRI, a strong enough field will disrupt this instability and thus reduce the effective viscosity. The overall effect is to limit the inflow of matter to a low rate and low velocity.

We know some black holes are launching jets. If MAD is to be part of the explanation for such observations, it must meet two conditions:

1. The conditions for the effect – the initial conditions of the flow, the feeding of the flow, the spin of the black hole, etc. – must be realized often enough in nature. It is conceivable that a mechanism could work but never actually be at play in the systems we observe.

2. The effect must be generic enough to occur given the right conditions. That is, we should be able to consistently see MAD states in the appropriate simulations, and we should see some of them producing jets as a result of their MAD states. Moreover, we should be confident that what we see in simulations is an accurate portrayal of reality.

We leave the first consideration for a more observationally oriented study. The second, though, is where there is room for improvement upon what is in the computational literature.

4.2 The Current State of Simulations

As already discussed, Tchekhovskoy et al. (2011) pioneered the study of MAD around black holes as a driving force behind jet launching. Their fiducial model A0.99f is of a black hole with dimensionless spin 0.99. The size of the simulation is $288 \times 128 \times 64$ in (r, θ, ϕ). The radial coordinate is roughly logarithmically spaced, extending from approximately 1 (the outer horizon is at $r \approx 1.14$) to 10^5, with most of the cells inside 10^3. (All distances here are in terms of the natural scale that is the mass M of the black hole.) The polar coordinate is not uniform but rather concentrated both at the poles and at the equator; it is also deformed to produce a more cylindrical grid near the poles. This results in a cell aspect ratio of $\Delta r : r\Delta\theta : r\Delta\phi \approx 2 : 1 : 7$, with approximately 96 cells per decade in radius. It should be noted that they have a model (A0.99fh) that doubles the number of cells in the

azimuthal direction, though this is not run from the initial conditions but rather starting with a turbulent state produced by the fiducial model.

The initial conditions they choose are those of an equilibrium torus as defined in Chakrabarti (1985), where the quantity $-u_\phi/u_t$ (Boyer–Lindquist coordinates) is taken to be constant. The inner edge is at $r = 15$, and the pressure maximum is at $r = 34$. For comparison, the innermost stable circular orbit is at $r \approx 1.45$. This equilibrium only holds in pure hydrodynamics, however the torus is seeded with a weak poloidal magnetic field. As described in Tchekhovskoy et al. (2011), this field is configured so as to make the parameter $\beta = p_{\rm gas}/p_{\rm mag}$ reach a minimum of 100, with a large portion of the inner part of the torus close to this minimum.

The key to the field configuration for the purposes of MAD is that the loops circulate about $r \approx 300$ rather than the pressure maximum. Often in magnetized torus simulations the field is initialized from a purely azimuthal vector potential whose magnitude is proportional to (some power of) density, causing the field lines to match isocontours of density and pressure. However this quickly leads to accretion of oppositely directed field and thus avoids the accumulation of ordered flux that defines MAD.

4.3 MAD Simulations in Athena++

We now describe the exact details of how we set up our MAD simulations.

4.3.1 Grid Parameters

We use spherical Kerr–Schild coordinates (see §C.3.2) for a black hole of mass $M = 1$ and spin $a = 0.99$. For these values the outer event horizon is at radius $r_{\rm hor} = M + \sqrt{M^2 - a^2} \approx 1.14107$ and the innermost stable circular orbit is at radius $r_{\rm isco} \approx 1.45450$.

The root grid covers $0.98 r_{\rm hor} < r < 1000$ with 96 cells spaced as close as possible to logarithmically. That is, we have a constant geometric ratio $\Delta r_i/\Delta r_{i-1} \approx 1.07336$ in adjacent cell widths, where this value is chosen such that the ratio of cell width to cell center

(taken to be the arithmetic mean of the inner and outer radii) is the same for the innermost and outermost cells. The root grid also covers both angular coordinates in their entirety, $0 < \theta < \pi$ and $0 < \phi < 2\pi$, using 32 cells for each.

We refer to this root grid as level 0. Level 1 mesh refinement involves dividing all cells in the range $\pi/8 < \theta < 7\pi/8$ into 8 new cells, doubling the resolution in each dimension. Note that the radial geometric ratio becomes the square root of its level 0 value here. Similarly, level 2 refinement is used in the region $3\pi/16 < \theta < 13\pi/16$, the resolution again doubling in each dimension. Finally, we consider level 3 refinement to take place in the region $7\pi/32 < \theta < 25\pi/32$, though here we leave the innermost 16 cells in radius at level 2 unrefined. That is, our level 3 grid only extends inward to $r \approx 1.48428 \approx 1.02047\ r_{\rm isco}$. Figure 4.1 illustrates the statically refined grid. With this scheme the timestep at any refinement level is limited by W_ϕ (see (2.32)) at the poles (at the innermost radii). A result of our refinement scheme and of the code's scaling (§1.5.4) is that we can increase resolution at the cost of only more computational cores but not wall time.

The outer radial boundary is held fixed, while a strict outflow condition is imposed at the inner boundary (outflow from the grid, which is inflow into the black hole). The grid is of course periodic in ϕ, and we stitch the grid together along the poles as well (§1.5.3).

4.3.2 Torus Parameters

On this grid we initialize an equilibrium hydrodynamical torus in the style of Fishbone and Moncrief (1976), which takes $\ell = u_\phi u^t$ (Boyer–Lindquist coordinates) to be constant. Such tori are parameterized by an inner edge radius $r_{\rm edge}$ and either a density peak radius $r_{\rm peak}$ or a specific angular momentum ℓ. We choose $r_{\rm edge} = 16.45$ and $r_{\rm peak} = 34$, resulting in $l \approx 6.29943$ (Fishbone and Moncrief, 1976, (3.8)). The orbital period at the pressure maximum is $P \approx 1251.88$ (Fishbone and Moncrief, 1976, (3.11)).

The equation of state is taken to be $p_{\rm gas} = \rho^\Gamma$ for $\Gamma = 13/9$, and the initial density is normalized to unity at its maximum. Rather than have vanishing density and pressure outside the torus, we impose radially varying floors of $\rho \geq \max\{10^{-2} r^{-3/2}, 10^{-6}\}$ and $p_{\rm gas} \geq$

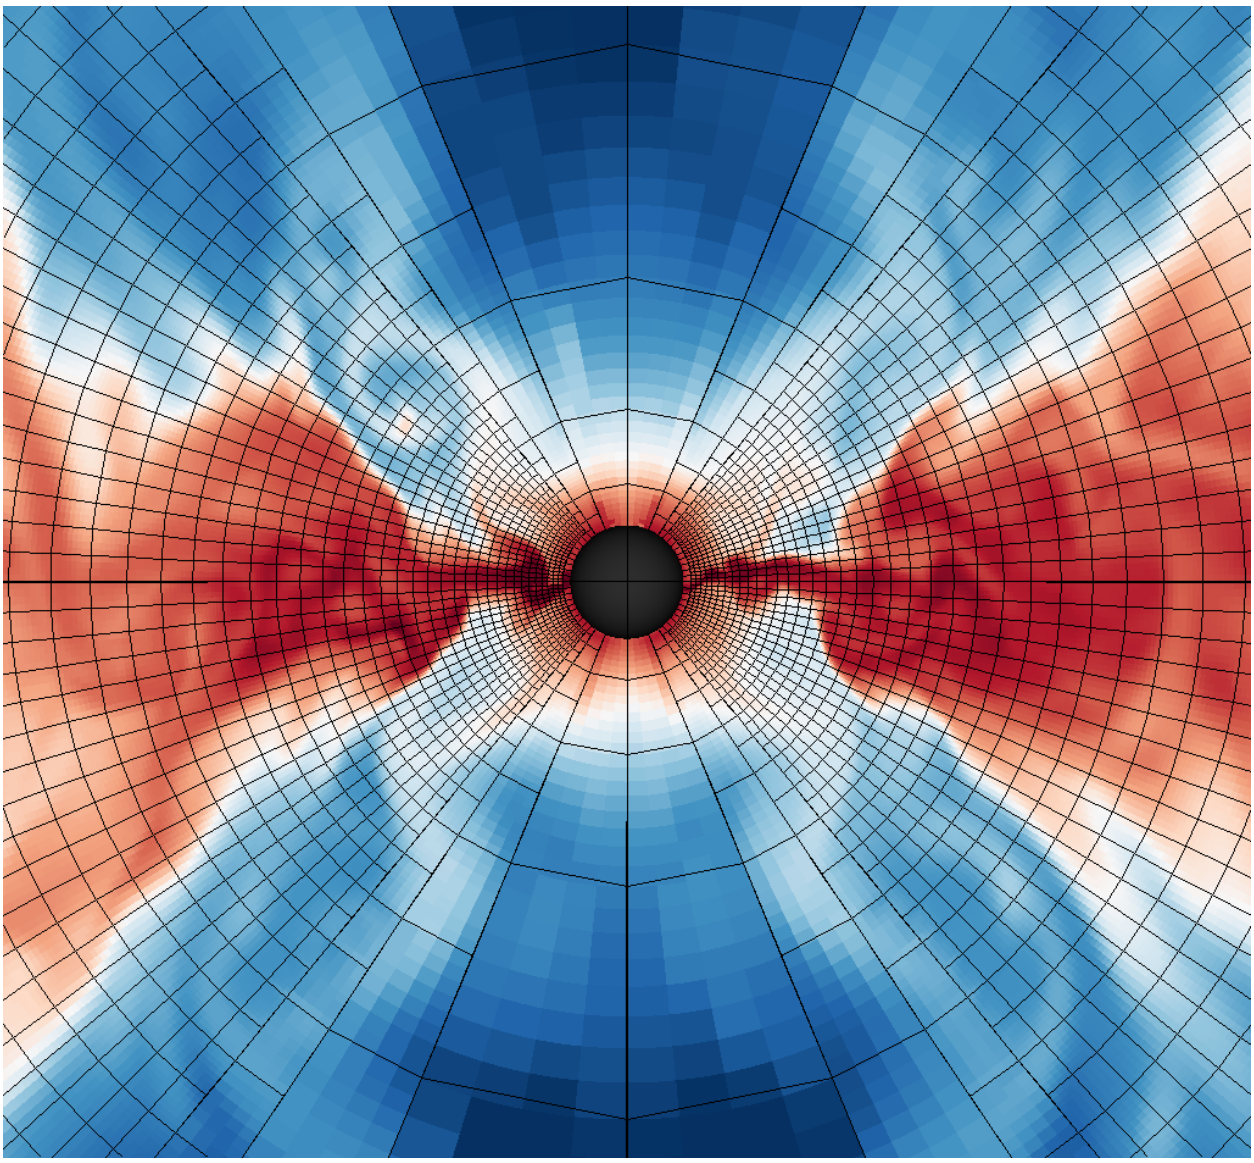

Figure 4.1: Grid used for MAD simulations, up to level 3 refinement. In this poloidal slice of the inner 11 M, each boxed region corresponds to 8 cells in r and 4 in θ. The background is the density ρ at time 5000 M, with the color scale running from 10^{-3} to 10^{1}.

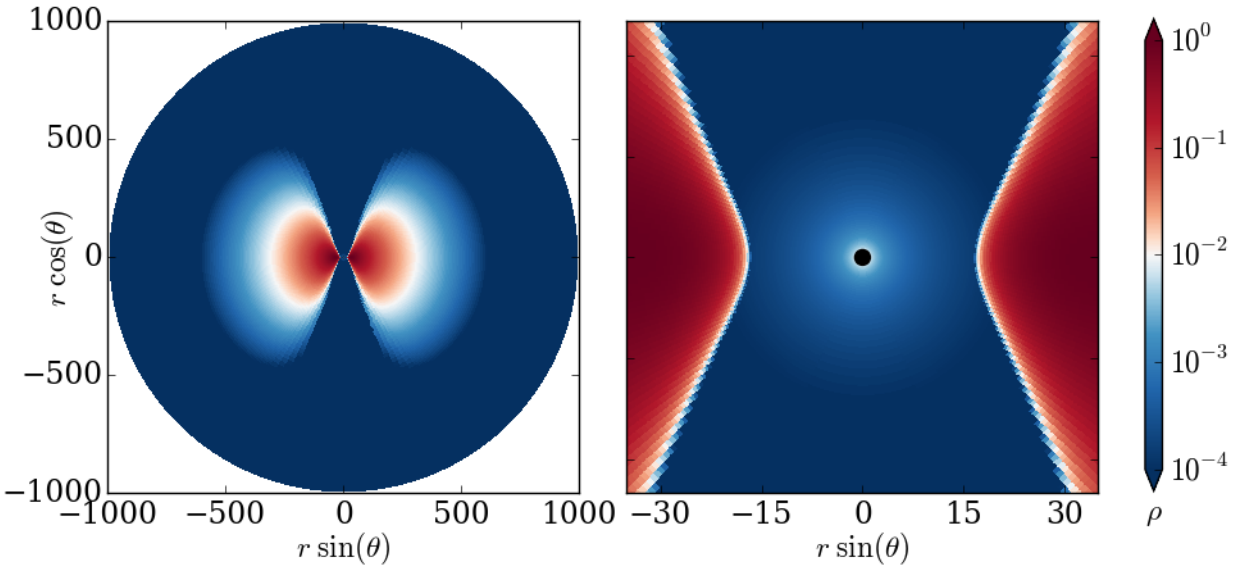

Figure 4.2: Initial density of MAD torus, shown at level 2 refinement and two zoom levels.

$\max\{10^{-2}(\Gamma-1)r^{-5/2}, 10^{-8}\}$. We also restrict the Lorentz factor of the fluid as seen by the normal observer to be $\gamma \leq \gamma_{\max} = 10$ throughout the simulation, where cells found to have higher velocities are adjusted according to $\tilde{u}^i \to \sqrt{(\gamma_{\max}^2 - 1)/(\gamma^2 - 1)}\tilde{u}^i$. The initial density, gas pressure, and Lorentz factor are shown in Figures 4.2–4.4.

In order to investigate MAD, we seed this torus with a weak magnetic field. Following the procedure given in Tchekhovskoy et al. (2011) we define an initial toroidal vector potential $\bar{A}_\phi = r^5\rho^2$ and from this we derive a poloidal magnetic field with components $\bar{B}^i$. Next we rescale this magnetic field locally to fix $\beta = p_{\rm gas}/p_{\rm mag}$ to be constant (thus violating the divergence-free constraint). We proceed to define the true vector potential as $A_\phi = \int_{\theta'=0}^{\theta'=\theta} \bar{B}^r\sqrt{-g}\,\mathrm{d}\theta'$, and from this we derive the true magnetic field components B^i (which obey the divergence-free constraint by construction). Finally, the field is rescaled globally such that β reaches a minimum value of 100 over the region $\rho > 0.05$. This cutoff is used to make the magnetic field strength minimally dictated by jitter in cell boundaries near the edge of the torus. We also perform the above calculations on a grid with 1536×512 cells in (r, θ), interpolating A_ϕ onto (the corners of) the actual grid just before differencing to find (the face-centered values of) B^r and B^θ. Thus even at different resolutions we have the same magnetic field.

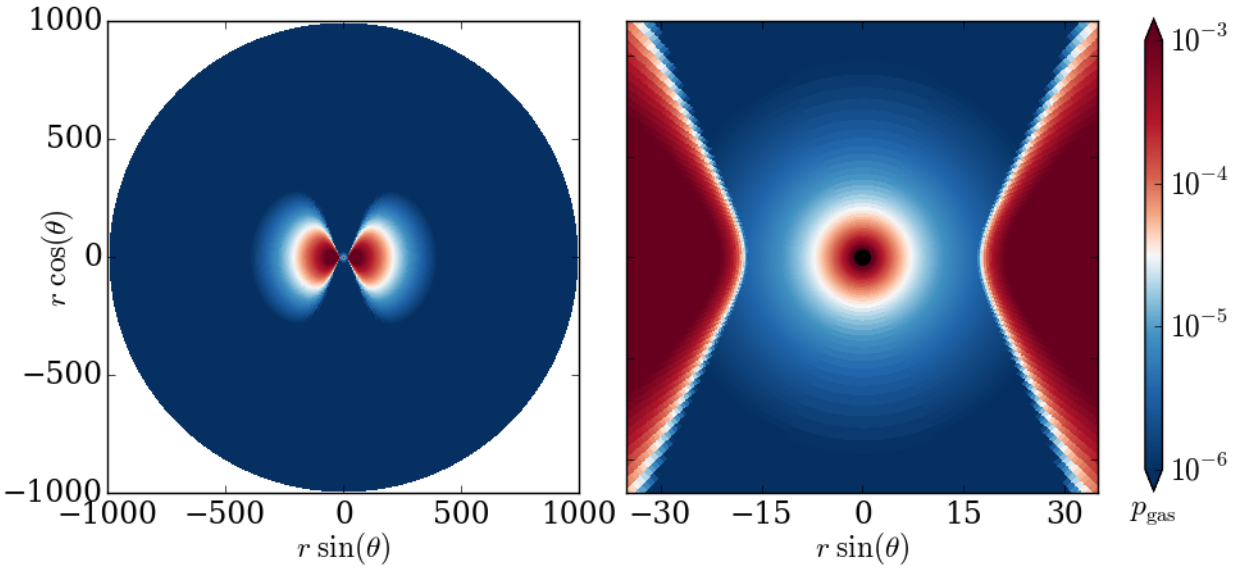

Figure 4.3: Initial gas pressure of MAD torus, shown at level 2 refinement and two zoom levels.

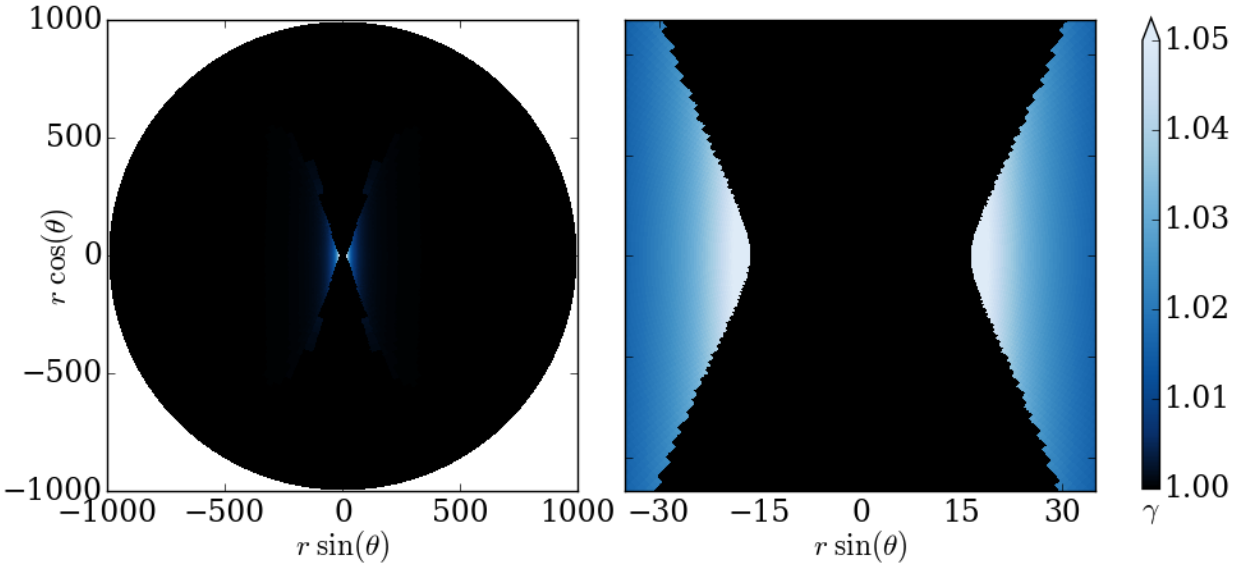

Figure 4.4: Initial normal observer Lorentz factor of MAD torus, shown at level 2 refinement and two zoom levels.

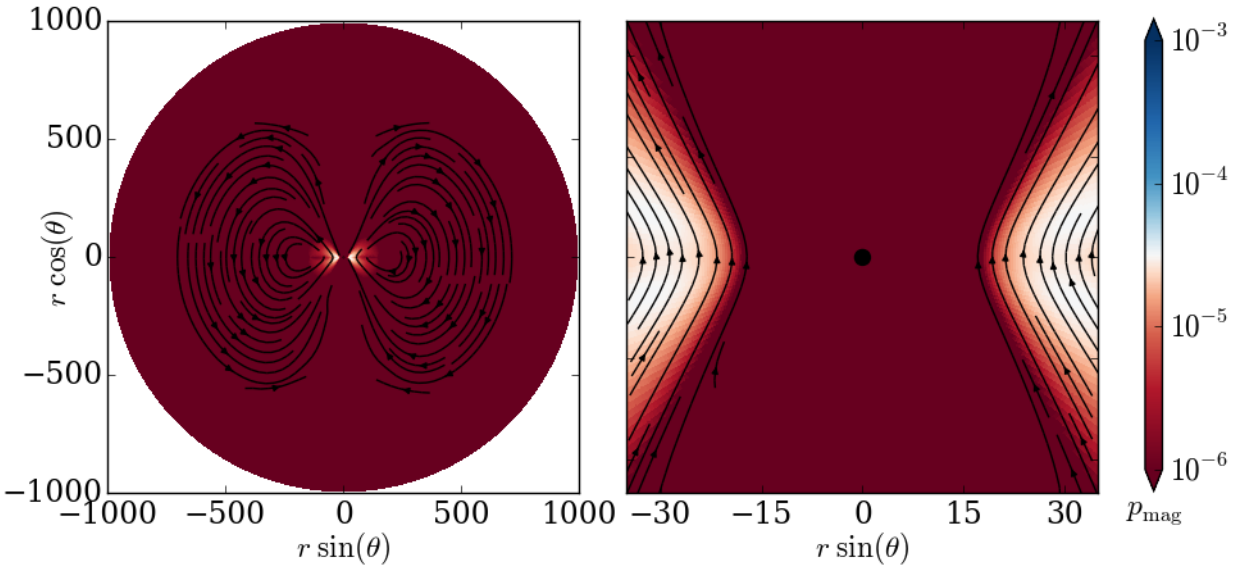

Figure 4.5: Initial magnetic pressure of MAD torus, shown at level 2 refinement and two zoom levels. Overlaid are streamlines for the magnetic field.

The complicated scheme has two effects. First, the scaling of $\bar{A}_\phi$ with r means the poloidal field circulates a ring well outside $r_{\rm peak}$. This ensures the flux being accreted is all of one alignment. Second, the local rescaling results in a magnetic field that has roughly constant β across the grid. The initial magnetic pressure $p_{\rm mag} = b_\lambda b^\lambda/2$ and β are shown in Figures 4.5 and 4.6, with the field configuration illustrated in the latter.

We note that two other limits are imposed on the variables throughout the simulation. The density must obey $\rho \geq p_{\rm mag}/\gamma_{\rm max}$, and the gas pressure must obey $p_{\rm gas} \geq (\Gamma - 1)p_{\rm mag}/10\gamma_{\rm max}$, just as in Tchekhovskoy et al. (2011). Because we cannot alter the magnetic field locally without violating the divergence-free constraint, these conditions are enforced via floors on ρ and $p_{\rm gas}$.

4.4 Results of New MAD Simulations

We ran three resolutions (refinement levels 0 through 2) with a Courant–Friedrichs–Lewy number of 0.25 out to a time of $t = 30{,}000\ M$. In all cases the flow reached a quasisteady state. We also ran variations on the level 2 simulation, changing polar resolution, polar boundary conditions, and numerical floors, in order to investigate what effects these have.

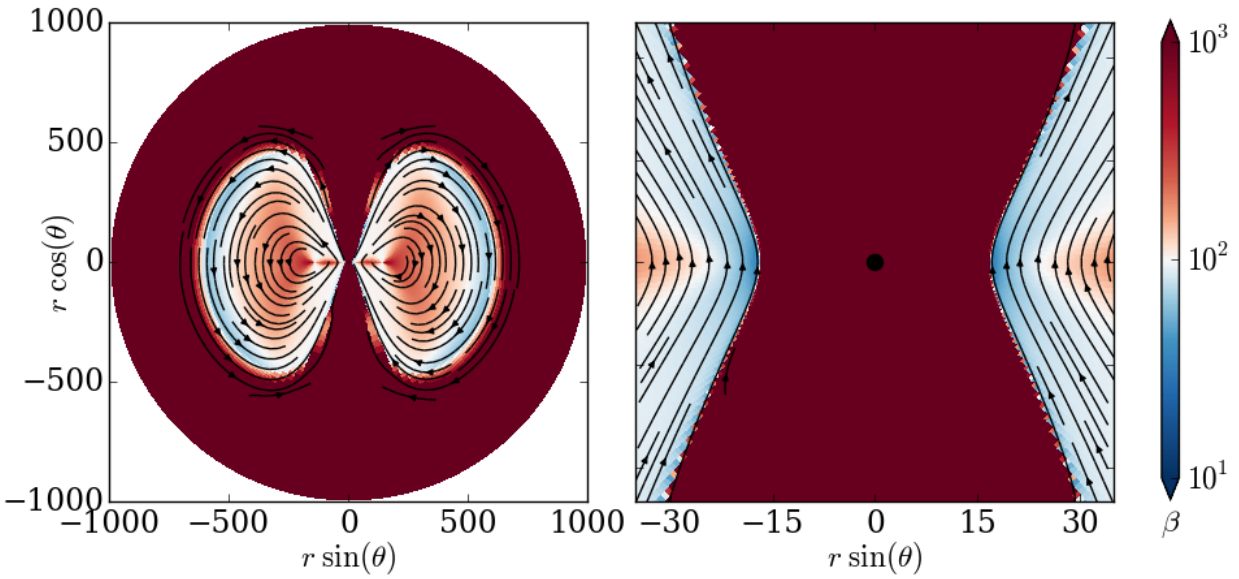

Figure 4.6: Initial plasma β of MAD torus, shown at level 2 refinement and two zoom levels. Overlaid are streamlines for the magnetic field.

Furthermore we ran a very high resolution simulation (level 3) to a time of $t = 4000\ M$ in order to further study convergence.

4.4.1 Primary Results at Different Resolutions

Figure 4.7 shows the density in the poloidal ($\phi = 0, \pi$) and equatorial ($\theta = \pi/2$) planes at $t = 15{,}000\ M$. Several features are immediately apparent. First, accretion has set in, with material funneling into the black hole in all cases. However, there is little evidence of turbulence at low resolution. Also, the level 0 simulation entered a warped, asymmetric state, and it stayed in this shape throughout the simulation. At high resolution, turbulence is visible in both the poloidal and equatorial slices.

Similar features can be seen in the gas pressure shown in Figure 4.8. In addition, there are regions of high pressure near the poles. These are not visible in the density plot, indicating they are near vacuum but very hot. These jet-like structures are indeed flowing outward from the black hole at high velocity. Figure 4.9 shows the Lorentz factor. From this figure the asymmetry at level 0 and even level 1 is quite apparent, with most of the outflow directed upward. This asymmetry is not seen at level 2, indicating that it probably is an

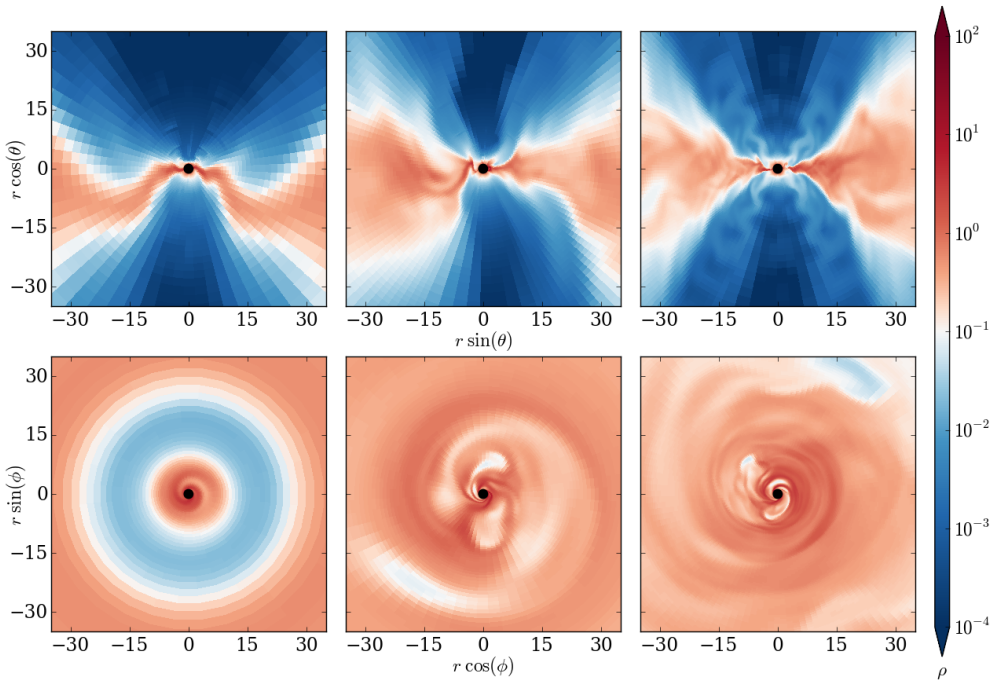

Figure 4.7: Poloidal (top) and equatorial (bottom) slices of density in MAD torus at time $t = 15{,}000\ M$. Refinement levels 0, 1, and 2 are shown left to right.

artifact of underresolving the flow.

Plots of magnetic pressure are shown in Figure 4.10. Here we can clearly see the effect of higher resolution on the amount of turbulence present. At level 0 there is only a single, smooth current sheet in the disk. At level 1 this sheet is contorted. At level 2 it is broken into multiple parts. Moreover, the region of high magnetic pressure seen in the equatorial slice becomes fragmented at higher resolution.

This turbulence is reflected in the magnetic field lines, which are more ordered at lower resolution. It is also worth noting that the field lines in the polar funnel region deviate significantly from vertical at lower resolutions, even crossing the polar axis horizontally. The plots show streamlines generated by taking only the slice shown and neglecting the toroidal component. In particular, they are not azimuthally averaged. When averaging is done, these deviations are significantly reduced, showing that even at a single timeslice the azimuthally average field is ordered while there are nonetheless nonaxisymmetric features.

The plots of β and $\sigma = b_\lambda b^\lambda/\rho$, Figures 4.11 and 4.12, tell a similar story. From them we can also see that the asymmetric jets seen at low resolution have relatively little magnetic

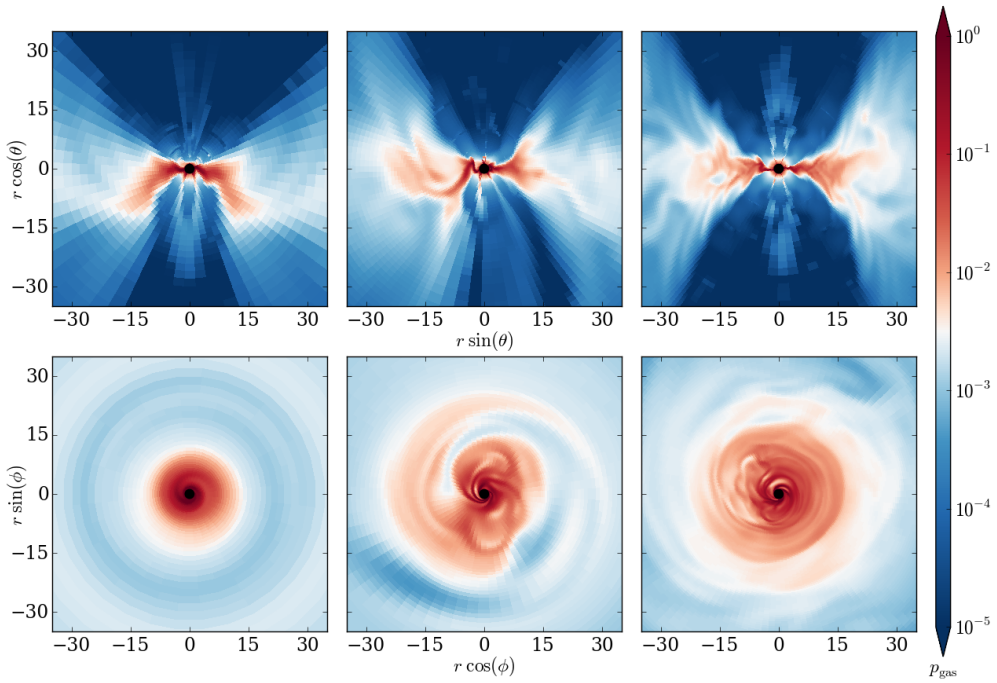

Figure 4.8: Poloidal (top) and equatorial (bottom) slices of gas pressure in MAD torus at time $t = 15{,}000\ M$. Refinement levels 0, 1, and 2 are shown left to right.

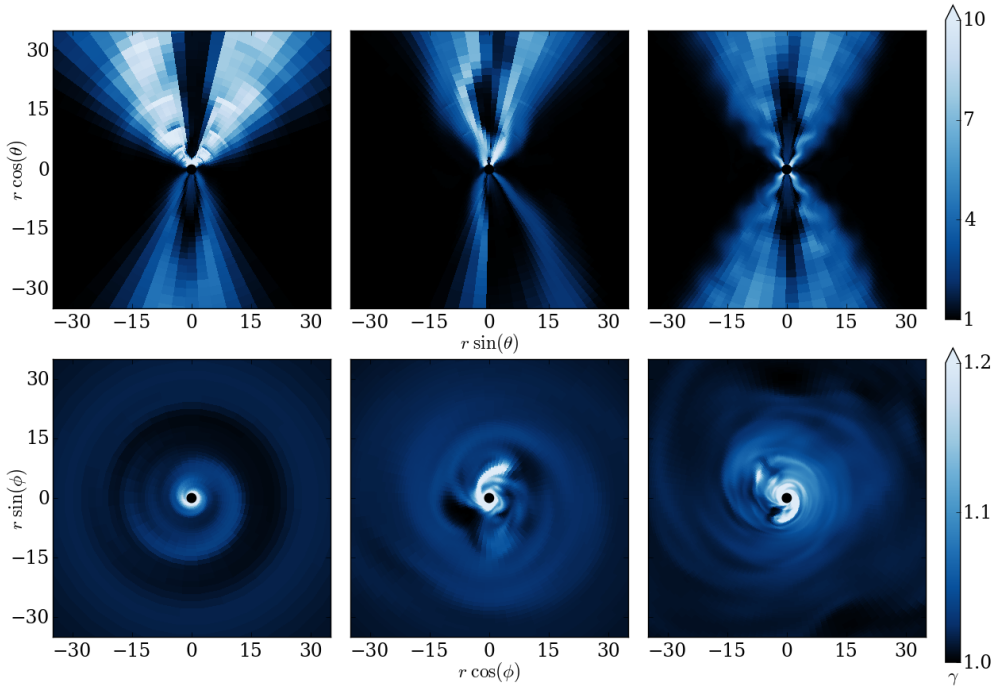

Figure 4.9: Poloidal (top) and equatorial (bottom) slices of normal observer Lorentz factor in MAD torus at time $t = 15{,}000\ M$. Refinement levels 0, 1, and 2 are shown left to right.

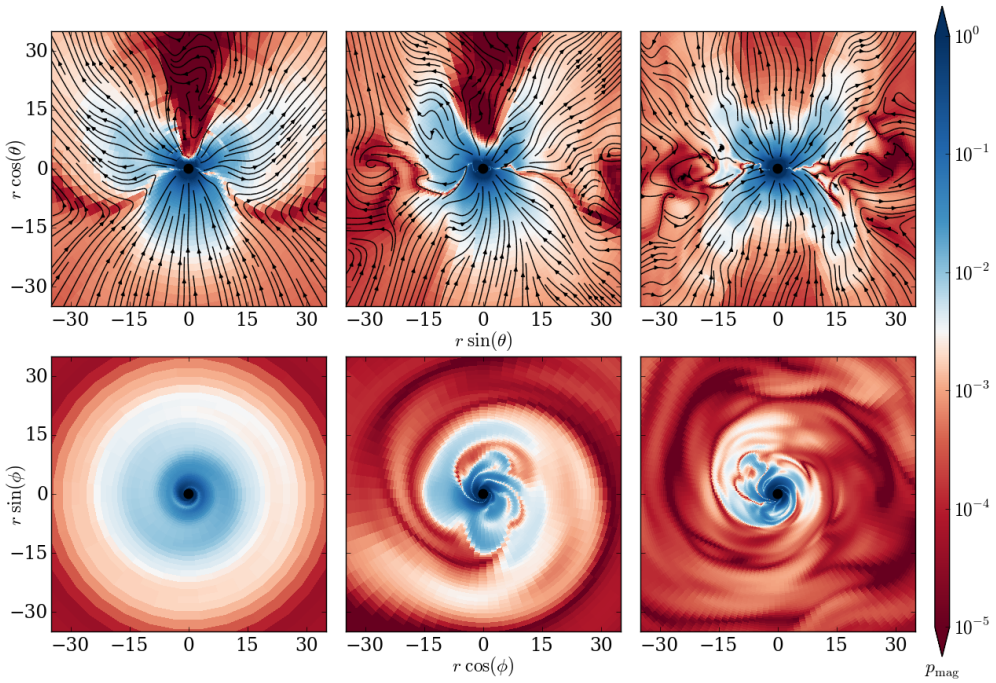

Figure 4.10: Poloidal (top) and equatorial (bottom) slices of magnetic pressure in MAD torus at time $t = 15{,}000\ M$. Refinement levels 0, 1, and 2 are shown left to right. Overlaid are streamlines for the magnetic field lying on their respective slices through the data.

pressure (equivalently energy density) given their gas pressure and rest mass density.

4.4.2 Variations Near the Poles

In order to run any numerical simulation, a number of choices must be made regarding implementation parameters. Here we take the refinement level 2 setup and change some of the numerical details in order to see what effect this has on the outcome. We focus on conditions that are likely to most strongly impact the funnel region near the poles.

While we have given much attention to resolving the accretion disk proper, we generally leave the polar regions at level 0 refinement in all simulations. Low azimuthal resolution here might be forgiven: nonaxisymmetric structures are likely to be advected away before significantly affecting the dynamics. However, low resolution in the polar angle could lead to diffusion of the magnetic field or momentum flux away from the pole.

In order to examine the effects of resolution in this region, we run a modified version of the level 2 grid in which level 1 refinement goes all the way to the poles (rather than

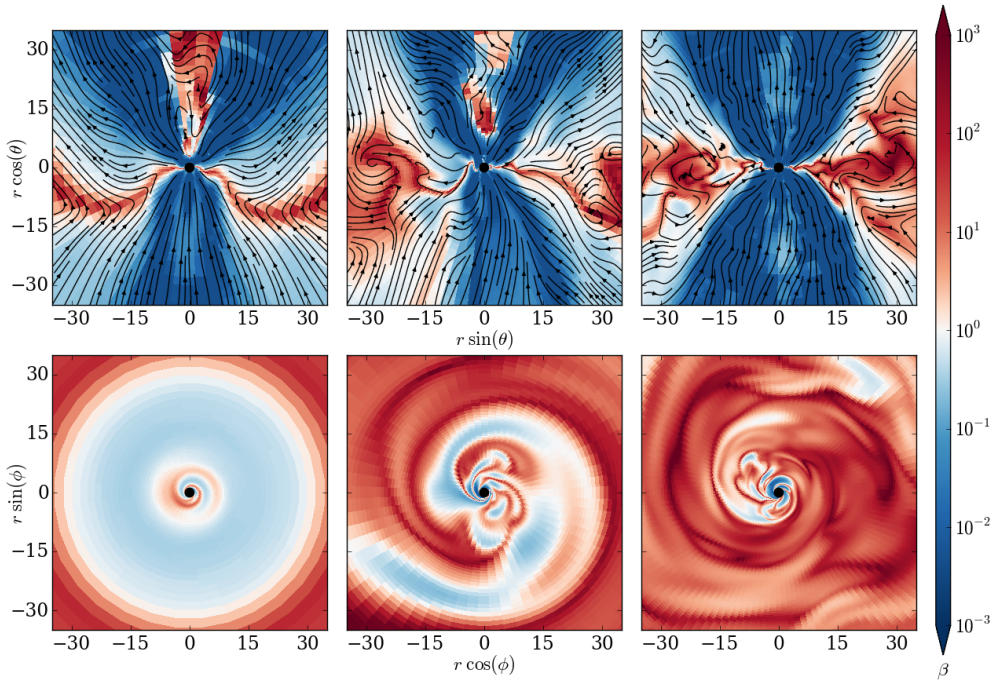

Figure 4.11: Poloidal (top) and equatorial (bottom) slices of plasma β in MAD torus at time $t = 15{,}000\ M$. Refinement levels 0, 1, and 2 are shown left to right. Overlaid are streamlines for the magnetic field lying on their respective slices through the data.

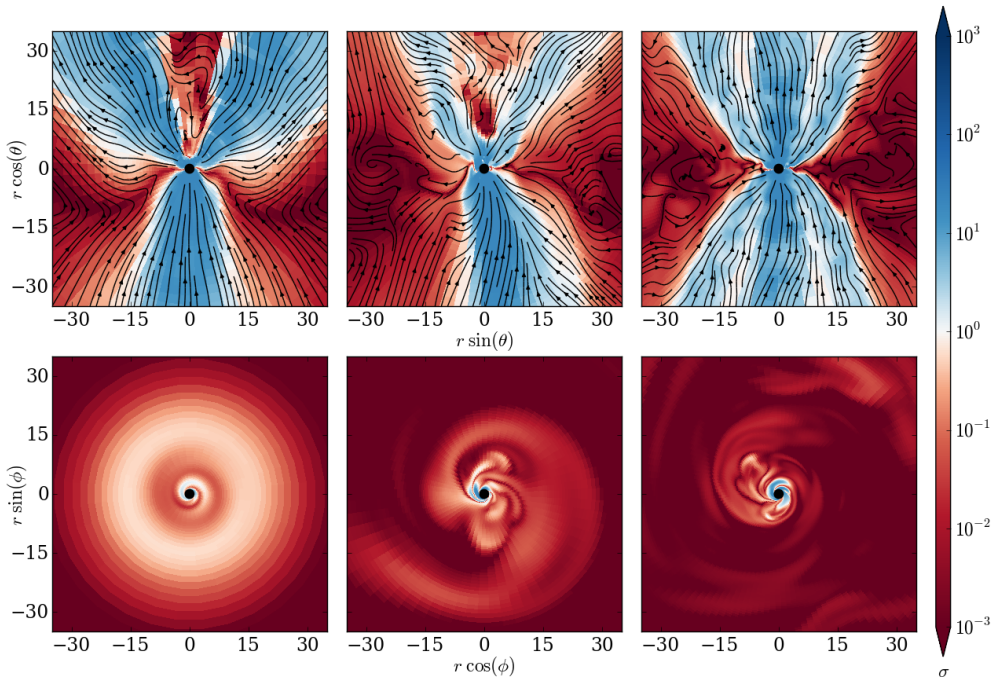

Figure 4.12: Poloidal (top) and equatorial (bottom) slices of plasma σ in MAD torus at time $t = 15{,}000\ M$. Refinement levels 0, 1, and 2 are shown left to right. Overlaid are streamlines for the magnetic field lying on their respective slices through the data.

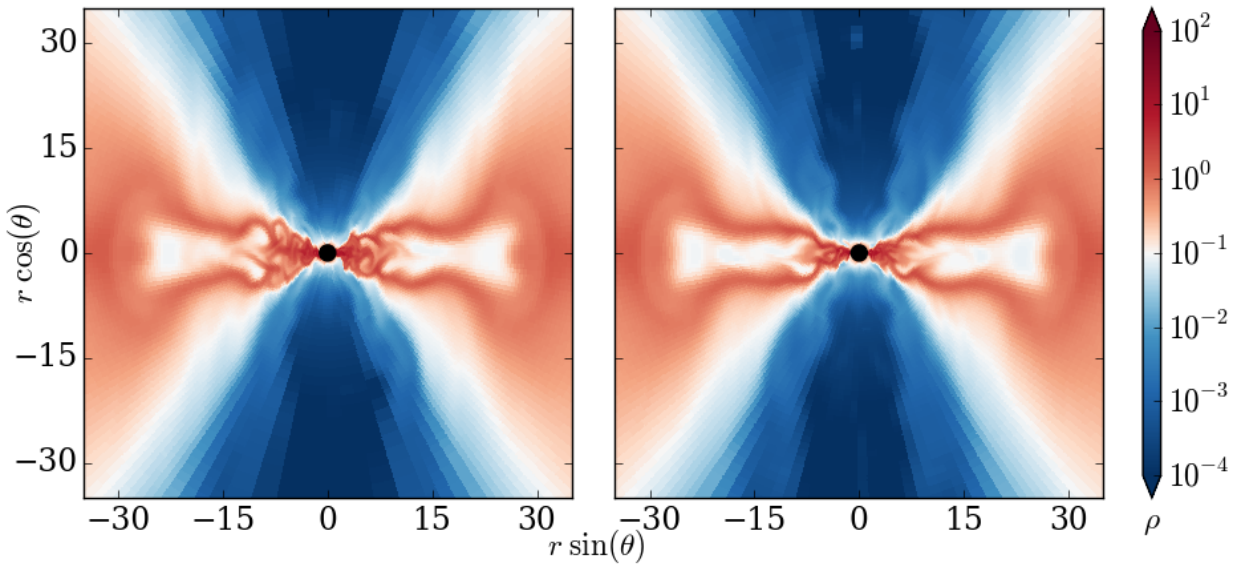

Figure 4.13: Poloidal slices of density in MAD torus at time $t = 2000\ M$. Shown are the fiducial run (left) and the variation with double resolution near the pole (right).

stopping at $\theta = \pi/8$ and $\theta = 7\pi/8$, see §4.3.1). In order to keep the timestep reasonable, we then halve the ϕ-resolution everywhere, meaning we have 32 cells in azimuth at the poles in both cases.

This simulation was run to a time of $t = 2000\ M$, after which numerical issues forced it to stop. Plots of ρ, β, and σ for the fiducial run and this variation at this time are shown in Figures 4.13 through 4.15. In all cases the turbulence in the disk is more developed in the fiducial case, however this is a result of the different number of azimuthal cells in the disk (128 versus 64). The jet region in the variation displays higher β and lower σ values, just as is seen in the fiducial runs at lower resolution (Figures 4.11 and 4.12), as well as hints of the field being less vertically oriented. Thus it appears that azimuthal resolution, even away from the poles, is more important than polar angle resolution for our purposes.

Two further variations were considered: imposing a reflecting wall and adjusting the floor values. For the wall we excise the region for which θ is within 10^{-8} of the poles, imposing a reflecting wall boundary condition along the newly created surface. For the new floor values, we change the position-dependent floors to $\rho \geq \max\{10^{-4}r^{-3/2}, 10^{-6}\}$ and $p_{\rm gas} \geq \max\{10^{-6}(\Gamma-1)r^{-5/2}, 10^{-8}\}$ (see §4.3.2 for the original floors). We also increase $\gamma_{\rm max}$

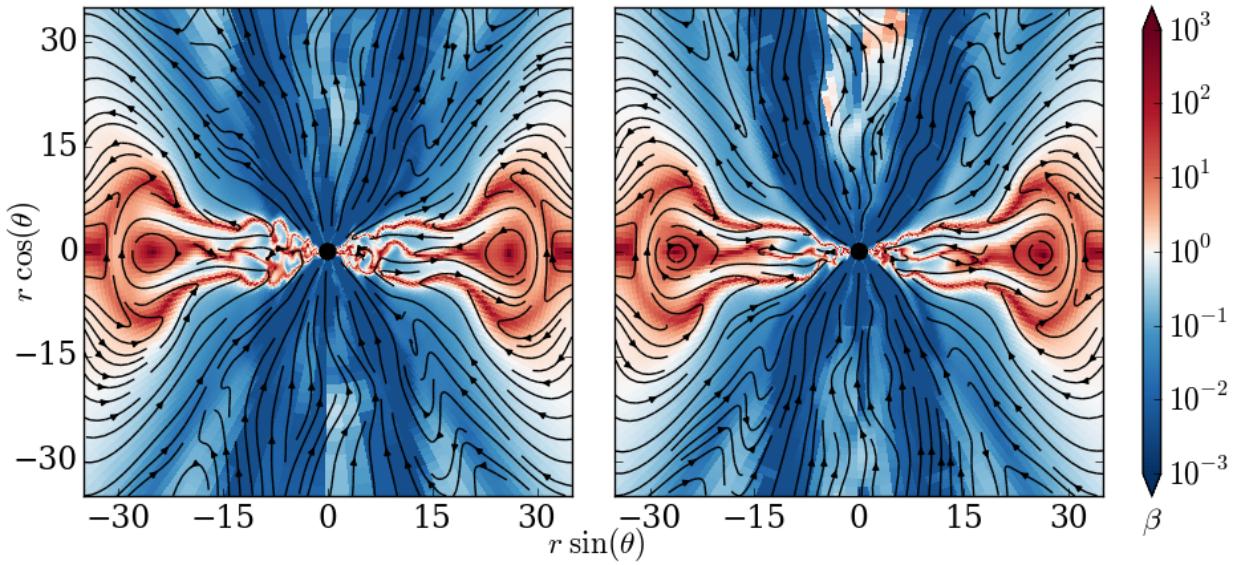

Figure 4.14: Poloidal slices of plasma β in MAD torus at time $t = 2000\ M$. Shown are the fiducial run (left) and the variation with double resolution near the pole (right). Overlaid are streamlines for the magnetic field lying on this slice through the data.

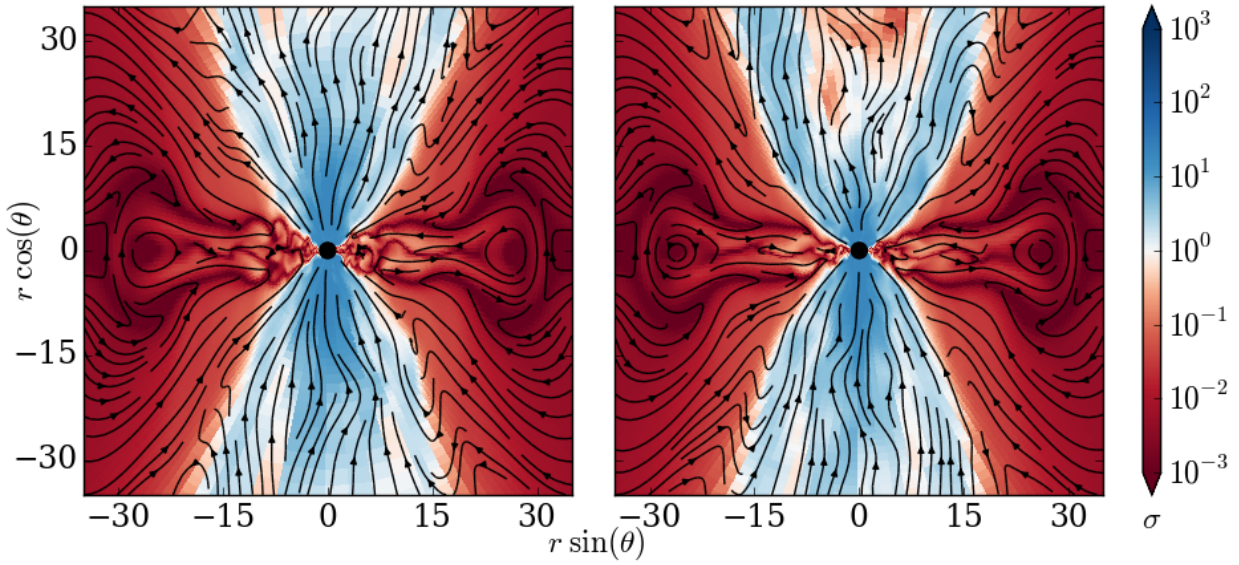

Figure 4.15: Poloidal slices of plasma σ in MAD torus at time $t = 2000\ M$. Shown are the fiducial run (left) and the variation with double resolution near the pole (right). Overlaid are streamlines for the magnetic field lying on this slice through the data.

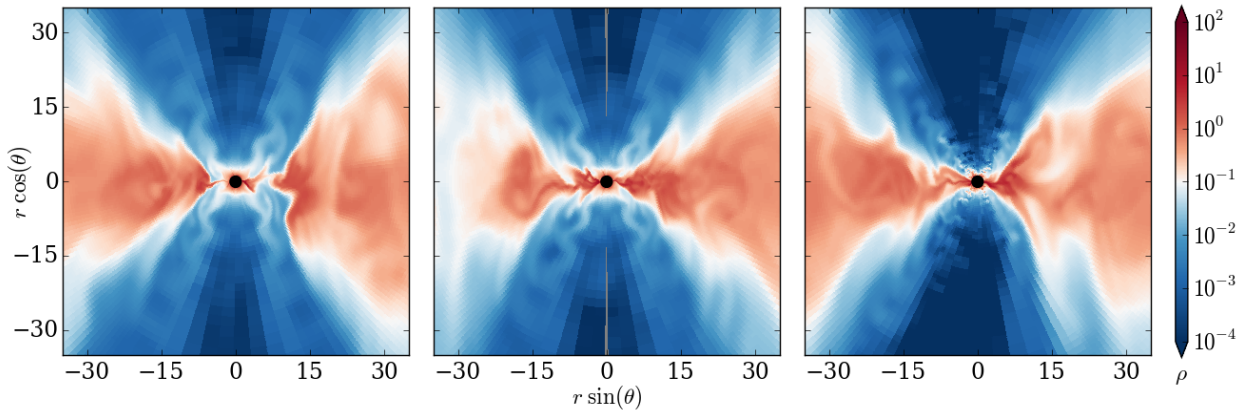

Figure 4.16: Poloidal slices of density in MAD torus at time $t = 10{,}000\ M$. Shown are the fiducial run (left), the variation with reflecting walls (center), and the variation with looser floors (right).

from 10 to 25, thus decreasing the limits $\rho \geq p_{\rm mag}/\gamma_{\rm max}$ and $p_{\rm gas} \geq (\Gamma - 1)p_{\rm mag}/10\gamma_{\rm max}$.

These variations were run to a time of $t = 10{,}000\ M$. Plots of ρ, β, and σ at this time are shown in Figures 4.16 through 4.18. The reflecting wall appears little different from the fiducial case. The looser floors, however, led to a significant amount of noise being generated, with large cell-to-cell variations apparent in all three figures. Interestingly, lowering the density a pressure floors, and indeed lowering these floors relative to the magnetic energy density, results in lower magnetization (higher β, lower σ) in the core of the funnel region, as is seen in the fiducial runs at lower resolution (Figures 4.11 and 4.12). The direction of the change indicates this likely unphysical demagnetization is not a simple effect of floors being activated in that region, but rather is a complex phenomenon emerging from the accumulation of numerical inaccuracies.

While some of the variations studied show pronounced differences in the outflow region, it is reassuring that they do not appear to affect the bulk of the accretion flow.

4.4.3 Very High Resolution

In addition to the runs with refinement levels 0 through 2, we ran a MAD simulation with refinement level 3. In the most refined region of this simulation the effective resolution (the resolution if the whole domain were at level 3 refinement) is $768 \times 256 \times 256$. In the radial direction this is 260 cells per decade.

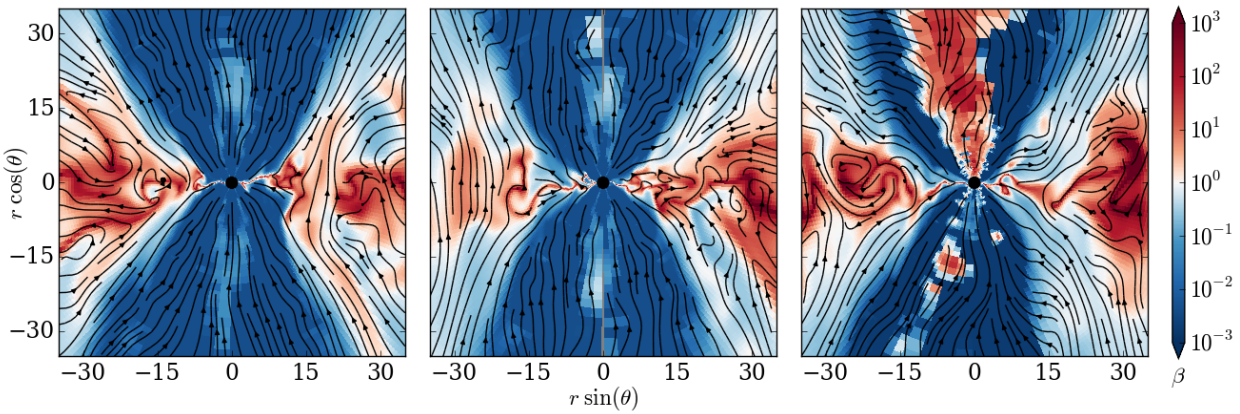

Figure 4.17: Poloidal slices of plasma β in MAD torus at time $t = 10{,}000\ M$. Shown are the fiducial run (left), the variation with reflecting walls (center), and the variation with looser floors (right). Overlaid are streamlines for the magnetic field lying on this slice through the data.

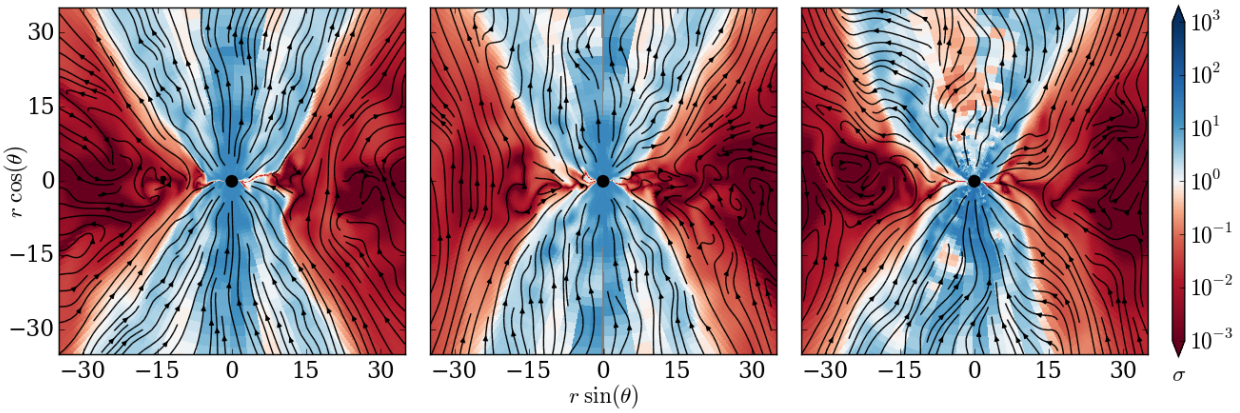

Figure 4.18: Poloidal slices of plasma σ in MAD torus at time $t = 10{,}000\ M$. Shown are the fiducial run (left), the variation with reflecting walls (center), and the variation with looser floors (right). Overlaid are streamlines for the magnetic field lying on this slice through the data.

This costly simulation was only run to $t = 4000\ M$ as of the time of this writing. We compare the fiducial level 2 simulation at this same time in the evolution. Figure 4.19 shows slices of the density in the two simulations. In the poloidal plane the high resolution disk is already disrupted by turbulence, whereas the lower resolution disk only shows mild distortions. It should be noted, though, that the turbulence has not saturated at level 2 by this point, as can be seen by considering the upper right panel of Figure 4.7, which resembles the upper right panel of Figure 4.19. Indeed turbulence develops more rapidly at higher resolution, as expected.

In the equatorial plane the effects of magnetized, low-density fluid bubbling up are more apparent at higher resolution. Indeed an $m = 4$ Rayleigh–Taylor mode is visually apparent, whereas at lower resolution only $m = 1$ and $m = 2$ spiral patterns are visible. Even at late times the lower resolution resolution simulation does not develop any readily apparent $m = 4$ mode, as can be seen in the lower right panel of Figure 4.7.

Similar patterns can be seen with β (Figure 4.20) and σ (Figure 4.21): the inner part of the disk is more turbulent at higher resolution. The low-magnetization streams feeding the black hole are thinner and more numerous at higher resolution. This holds even when considering later times at lower resolution, as shown in the lower right panels of Figures 4.11 and 4.12.

In order to more clearly show features very close to the horizon, Figures 4.22–4.24 are zoomed in versions of Figures 4.19–4.21. Of particular note are the disruption of the disk at high resolution in the former and the thinness of the accreting streams in the latter two.

Thus we see that at least in some aspects even the midplane flow is not entirely converged in our series of simulations. The polar funnel region, however, appears the same in all three displayed quantities between the two resolutions.

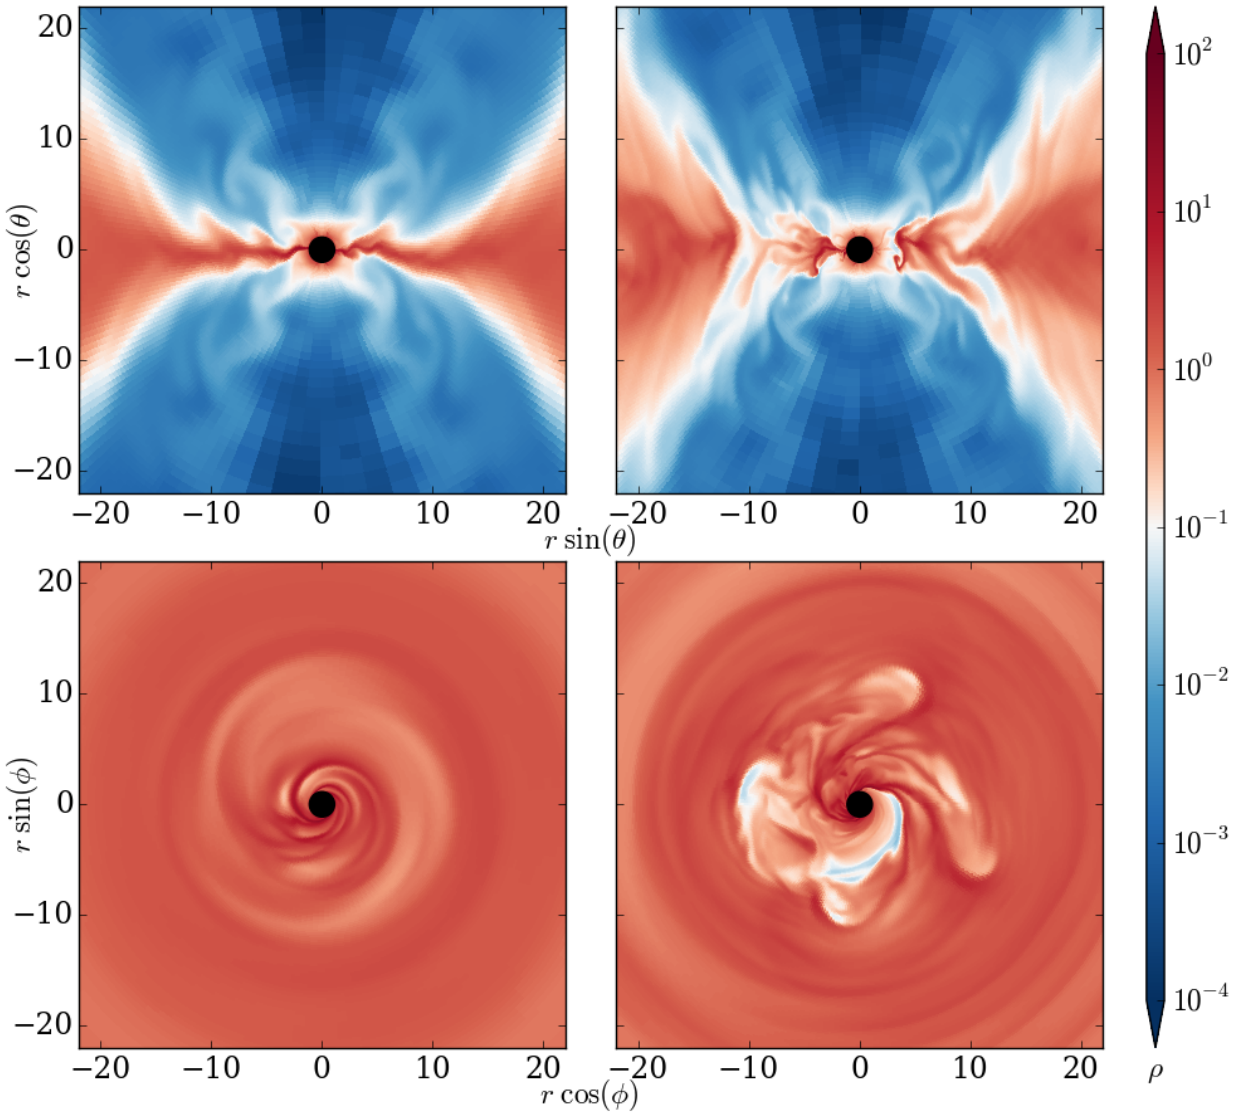

Figure 4.19: Poloidal (top) and equatorial (bottom) slices of density in MAD torus at time $t = 4000\ M$ with refinement levels 2 (left) and 3 (right). See Figure 4.22 for a zoomed in version.

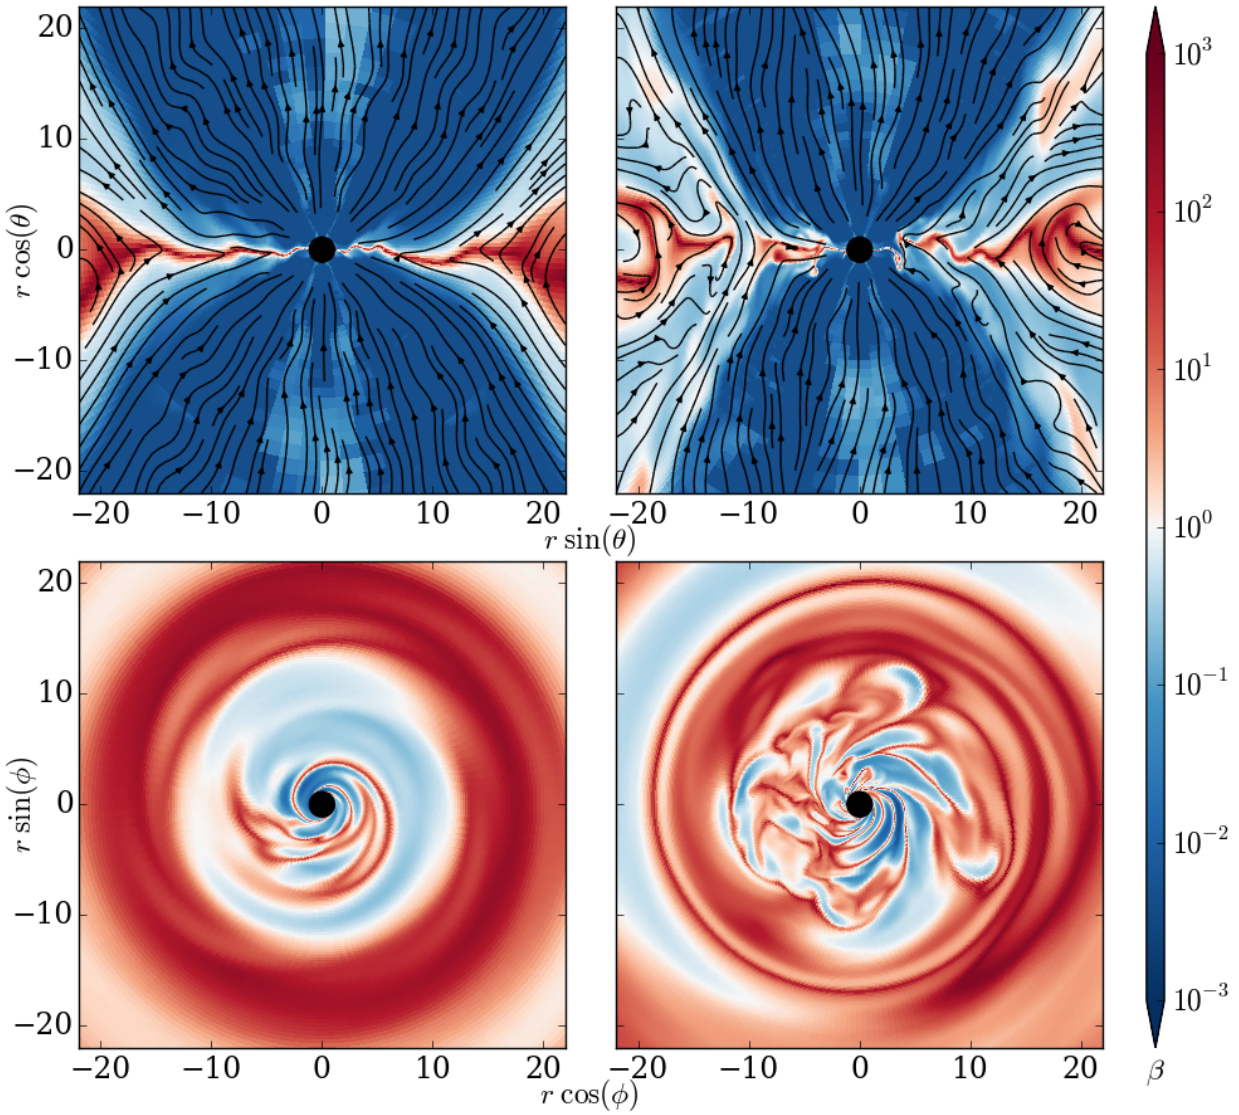

Figure 4.20: Poloidal (top) and equatorial (bottom) slices of plasma β in MAD torus at time $t = 4000\ M$ with refinement levels 2 (left) and 3 (right). Overlaid are streamlines for the magnetic field lying on their respective slices through the data. See Figure 4.23 for a zoomed in version.

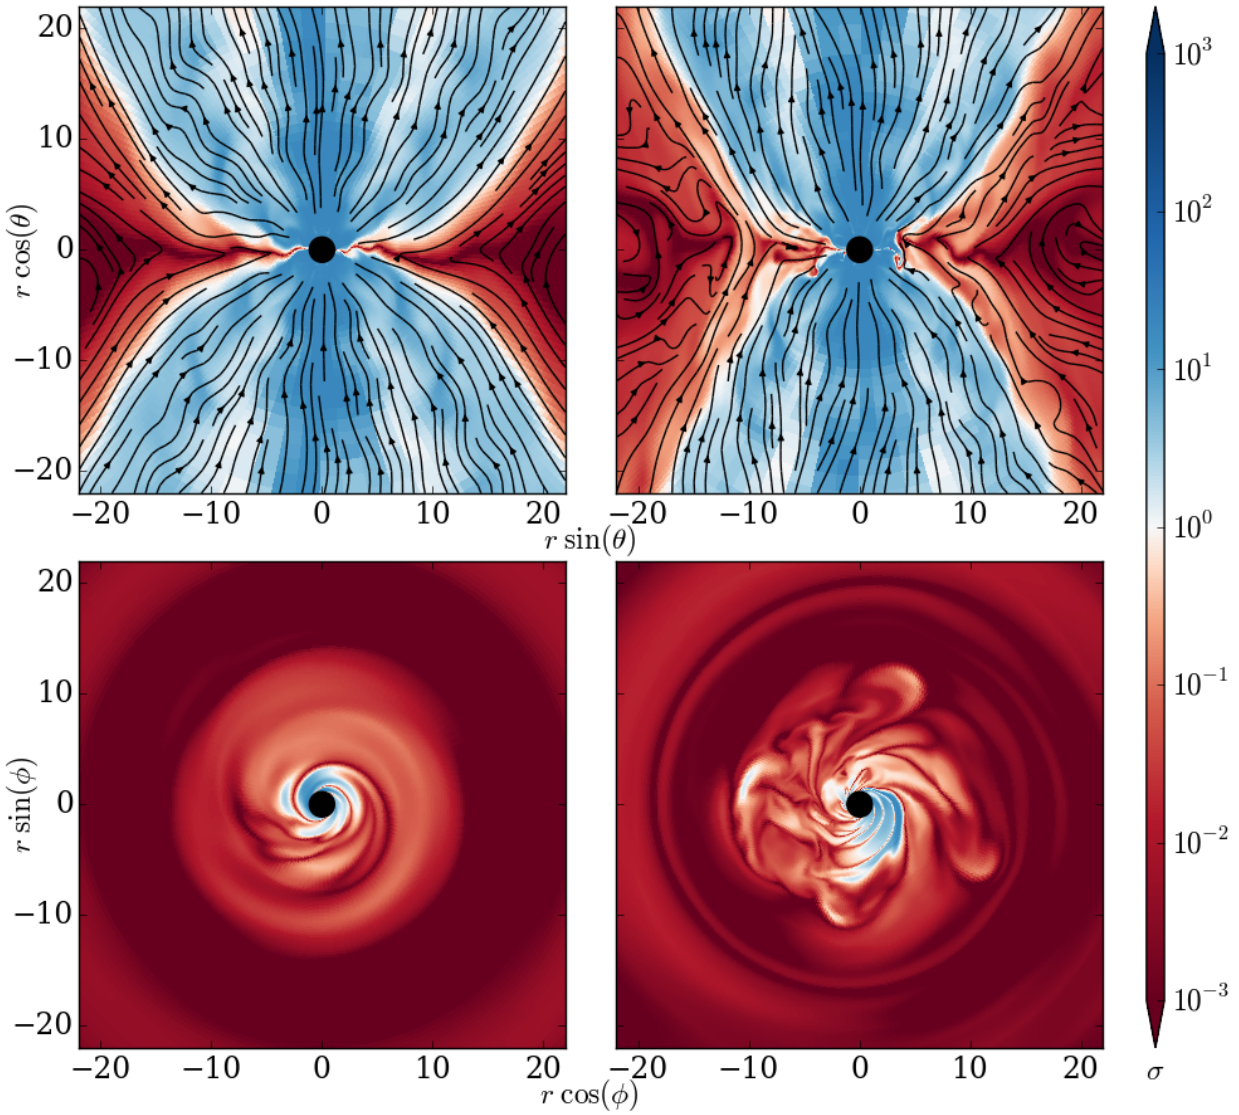

Figure 4.21: Poloidal (top) and equatorial (bottom) slices of plasma σ in MAD torus at time $t = 4000\ M$ with refinement levels 2 (left) and 3 (right). Overlaid are streamlines for the magnetic field lying on their respective slices through the data. See Figure 4.24 for a zoomed in version.

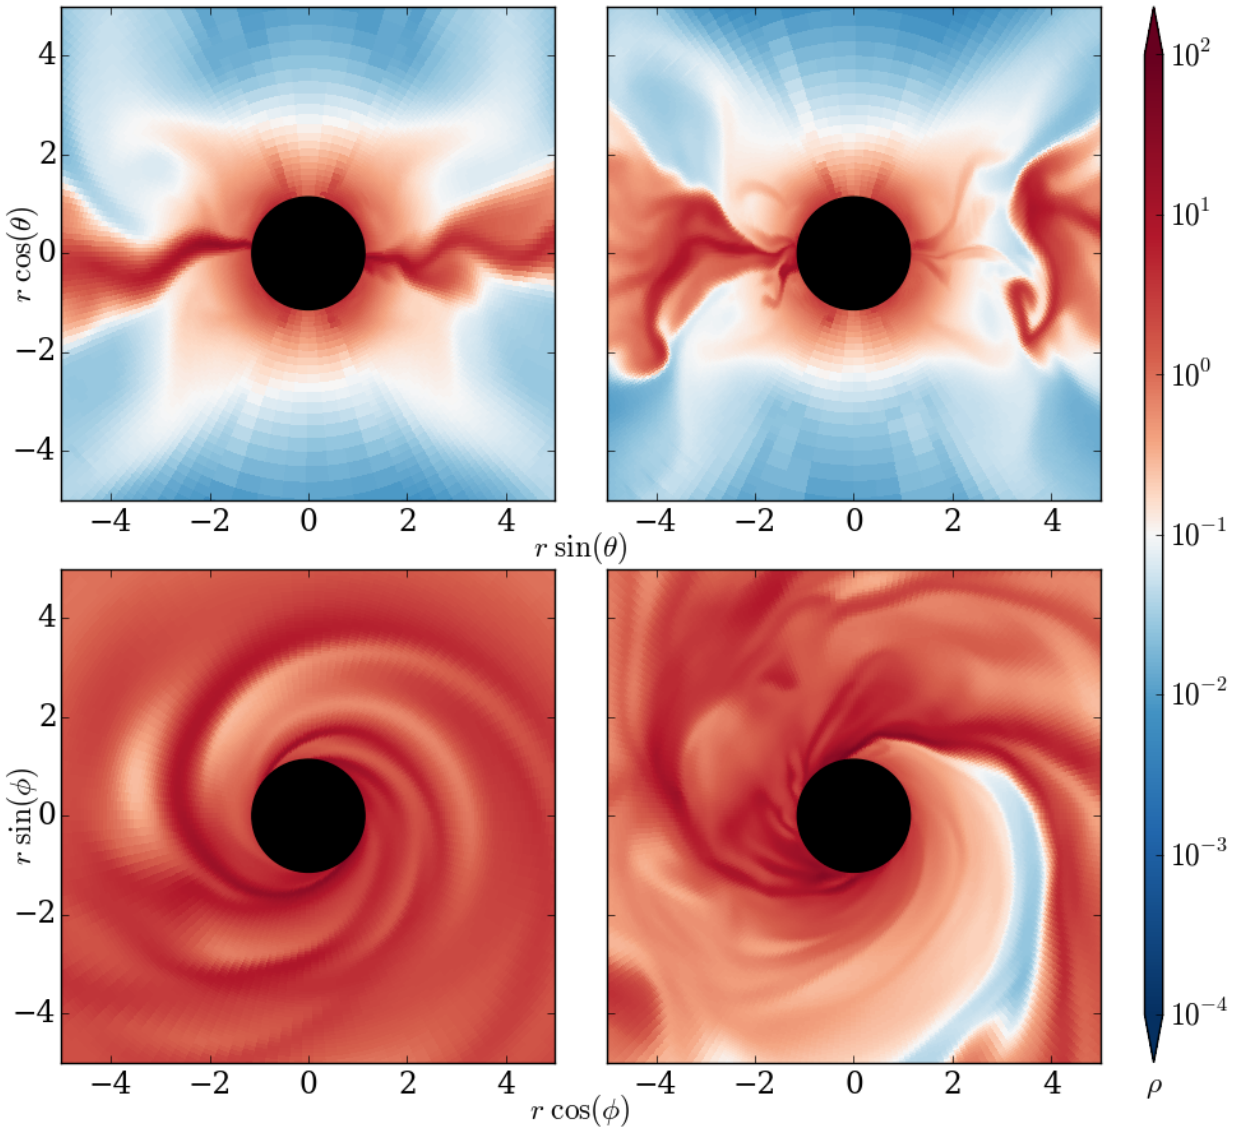

Figure 4.22: Zoomed in version of Figure 4.19, showing poloidal (top) and equatorial (bottom) slices of density in MAD torus at time $t = 4000\ M$ with refinement levels 2 (left) and 3 (right).

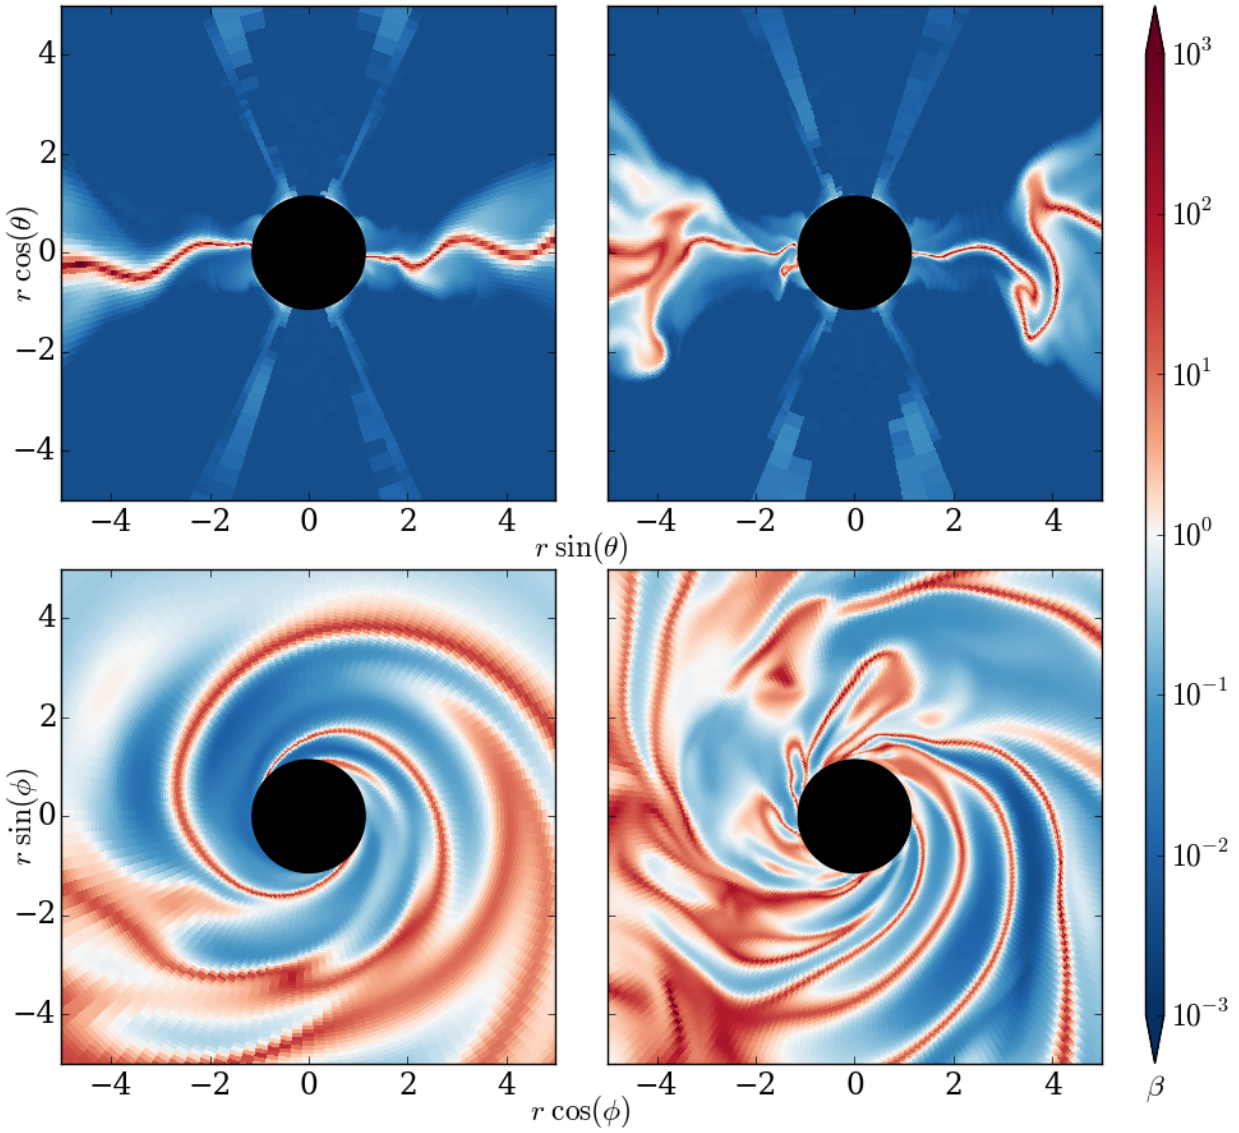

Figure 4.23: Zoomed in version of Figure 4.20, showing poloidal (top) and equatorial (bottom) slices of plasma β in MAD torus at time $t = 4000\ M$ with refinement levels 2 (left) and 3 (right).

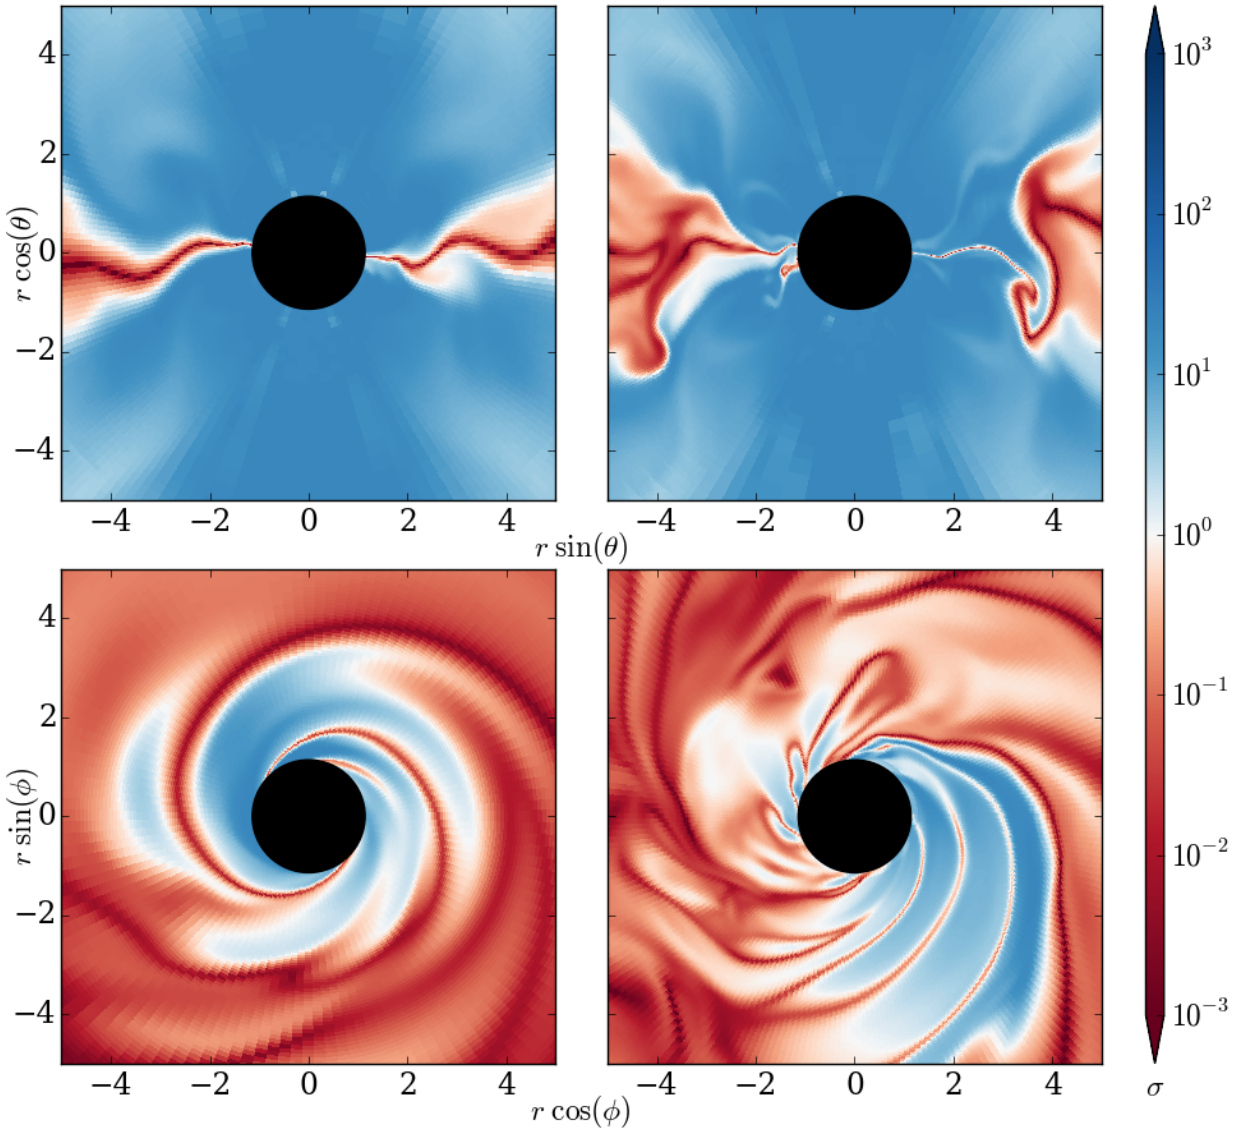

Figure 4.24: Zoomed in version of Figure 4.21, showing poloidal (top) and equatorial (bottom) slices of plasma σ in MAD torus at time $t = 4000\ M$ with refinement levels 2 (left) and 3 (right).

4.5 Interpretation of Our Results

While the lowest resolution simulation shows the large flux (Figure 4.10) and small amount of mass (Figure 4.7) in the central region one expects in a MAD scenario, the picture becomes more complicated as we go to higher resolution. The flow is more turbulent, and nonaxisymmetric structures allow for mass to flow inward through the region of large flux.

This behavior is not unexpected. A magnetic field can provide pressure support, but as it has no rest mass any attempt to hold a massive fluid against gravity is likely to be Rayleigh–Taylor unstable. This is not to say that there can be no arresting of the flow in nature; MAD can still act to slow down the mass accretion.

In order to analyze the situation more quantitatively, we measure accretion in several ways as a function of time. Following Tchekhovskoy et al. (2011) we define the mass accretion rate

$$\dot{M} = -\int_{r=r_0} \rho u^r \sqrt{-g}\, \mathrm{d}\theta\, \mathrm{d}\phi, \tag{4.1}$$

where the quantities are evaluated in Kerr–Schild coordinates and r_0 is a fixed radius. Similarly we define an energy accretion rate

$$\dot{E} = \int_{r=r_0} T^r{}_t \sqrt{-g}\, \mathrm{d}\theta\, \mathrm{d}\phi. \tag{4.2}$$

These quantities measure the inflow of rest mass and energy at infinity, respectively, and in both cases a positive sign represents a flow into the black hole. We analogously define a magnetic energy flux by neglecting the density and gas pressure terms in the stress-energy tensor:

$$\dot{E}_{\mathrm{mag}} = \int_{r=r_0} (b_\lambda b^\lambda u^r u_t - b^r b_t)\sqrt{-g}\, \mathrm{d}\theta\, \mathrm{d}\phi. \tag{4.3}$$

Tchekhovskoy et al. also define the magnetic flux threading half the horizon as

$$\Phi = \frac{1}{2}\int_{r=r_0} \left|\sqrt{4\pi} B^r\right| \sqrt{-g}\, \mathrm{d}\theta\, \mathrm{d}\phi \tag{4.4}$$

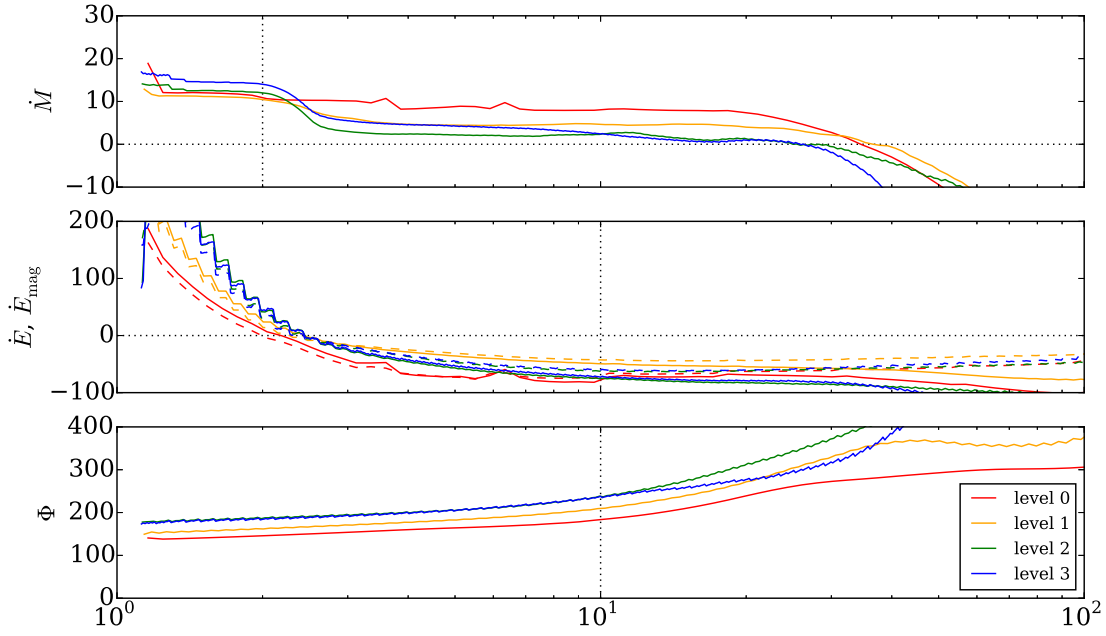

Figure 4.25: Mass accretion, energy accretion, and horizon-threading flux in the MAD simulations as functions of radius. Dashed lines show pure magnetic energy flux defined in (4.3). For refinement levels 0 through 2, values at each radius are averaged over times from 10,000 M to 15,000 M at intervals of 100 M; for level 3 the averaging is done from 3000 M to 4000 M, sampled every 10 M. Vertical dotted lines indicate selected values of r_0.

(where we insert the factor of $\sqrt{4\pi}$ to account for the difference between our units and theirs), which they normalize to the average accretion rate over a suitablly large interval:

$$\phi = \frac{\Phi}{\langle \dot{M} \rangle^{1/2} M}. \tag{4.5}$$

Before examining the time dependence of these quantities, we must choose r_0. Ideally we would choose $r_0 = r_{\rm hor}$, but in practice the numerics near the horizon (especially the imposition of floors) make fluxes evaluated here suspect. In order to aid this choice, we inspect the time-averaged radial dependence of the fluxes, plotted in Figure 4.25. This includes data from refinement levels 0 through 3, with the highest level not run all the way to late times.

The effect of density floors is to artificially add mass the near-vacuum regions, which in our case coincide with the outflow. In order to approximately negate this effect in the

analysis we zero the density of any cell within 15% of the floor at that location. This has the effect of raising the values of $\dot{M}$ inside approximately $r = 2$. Note that wherever the simulation has reached a steady turbulent state we expect $\dot{M}$ to be constant. This is generally true, although the values outside approximately $r = 3$ are likely lower than they ideally would be, given that artificial mass loading closer in will appear as a massive outflow, but that at further radii this injected mass will be well above the floors and so not removed by our procedure.

By using $T^r{}_t$ rather than T^{rt} we are essentially examining the radial flux of binding energy at infinity. This quantity should therefore also be independent of radius, at least out as far as steady state has been achieved. As can be seen in the second panel of Figure 4.25, we do see a rise in $\dot{E}$ as we move inward. This is presumably also an effect of the floors, though in a way that is less easy to correct. In particular, it is not an effect of the mass or gas pressure terms in the stress-energy tensor: the solid lines show $\dot{E}$ and the dashed lines show $\dot{E}_{\rm mag}$, and they are similar at all radii.

The magnetic flux Φ is rather constant inside $r = 10$, though it varies based on resolution. The flux is reduced at lower resolutions. Notably however the flux inside $r = 10$ is almost identical for the level 2 and level 3 simulations, indicating some degree of convergence.

Based on the above discussion, we choose to take $r_0 = 2$ for evaluating $\dot{M}$ in (4.1). Meanwhile we choose $r_0 = 10$ for $\dot{E}$, $\dot{E}_{\rm mag}$, and Φ in (4.2), (4.3), and (4.4). We believe this leads to an accurate representation of the fluxes. One can use Figure 4.25 to see how later results depend on this choice.

With suitable values of r_0 chosen, we can investigate the full time dependence of accretion. Figure 4.26 shows the evolution of $\dot{M}$, $\dot{E}$, and ϕ for the four resolutions. In order to normalize the flux ϕ we must choose an interval over which to average $\dot{M}$ (see (4.5)). For the lower three resolutions we choose the interval from 5000 M to 30,000 M, and for the highest resolution we choose the interval from 3000 M to 4000 M. Measures of the means and standard deviations of these quantities, as well as η defined in (4.6), over these intervals are presented in Table 4.1. Though a simplistic method, analyzing standard deviations has

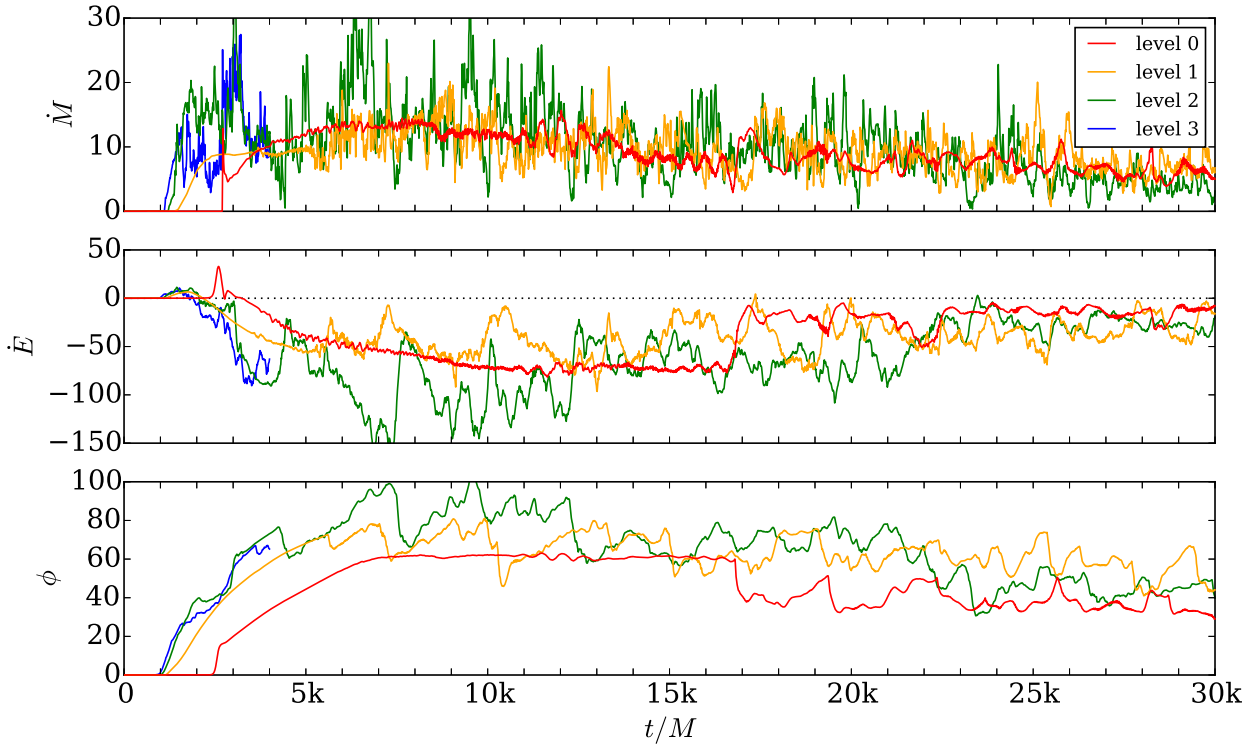

Figure 4.26: Mass accretion, energy accretion, and normalized horizon-threading flux in the MAD simulations as a function of time for four resolutions. All simulations are sampled at a rate of $\Delta t = 10\ M$.

Table 4.1: Means and standard deviations of accretion values.

Level	$\dot{M}$	$\dot{E}$	ϕ	η
0	9.2 ± 2.7	-40 ± 26	49 ± 12	5.4 ± 3.0
1	9.4 ± 3.2	-40 ± 17	63.9 ± 8.5	5.3 ± 1.9
2	10.0 ± 5.6	-62 ± 34	66 ± 17	7.2 ± 3.8
3	14.0 ± 5.3	-71 ± 11	62.8 ± 3.2	6.10 ± 0.68

the benefit of not being affected by the sampling rate in time at which we dump data.

The relative steadiness of the level 0 lines reflect the lack of turbulence at low resolution. We note that this resolution undergoes a qualitative jump in behavior at $t \approx 17{,}000\ M$. At this time a single large-amplitude spiral wave appears in the disk, introducing some degree of variability. With regard to $\dot{M}$ this variability is on longer timescales than seen at the other resolutions.

The level 1 and level 2 runs display similar accretion behavior, with $\dot{M} \approx 10$ and $\phi \approx 65$. There is around 50% more energy outflow at level 2 compared to level 1. In all three quantities level 2 displays approximately twice the variability (measured as the standard

deviation of the values). The level 3 data behaves very similarly to the level 2 data over the limited time in which the former was generated.

One way in which the level 1 and level 2 runs differ is in the time to the onset of turbulence. The higher resolution simulation shows variability in $\dot{M}$ earlier, by $t = 1500\ M$, whereas at lower resolution there is a buildup period until $t = 5000\ M$. In some sense the first 16,000 M time of the level 0 resolution simulation is this same buildup phase.

Another difference is in the strength of short-term variability. The level 2 run shows large variations between adjacent samples in $\dot{M}$. In all cases the sampling rate is fixed at $\Delta t = 10\ M$. This short-time variance is not unexpected at higher resolutions, as the finer grid allows the turbulence to cascade to smaller length scales and thus shorter characteristic times.

The fact that the level 3 data has the same $\dot{M}$ variability as the level 2 data (see Table 4.1) may indicate a saturation in the level of turbulence. We caution however that the high-resolution simulation should be run for longer to verify this claim.

Finally, we return to the consideration of energy extraction, which is one of the central motivations for the study of MAD accretion onto black holes. As in Tchekhovskoy et al. (2011) we define the efficiency

$$\eta = \frac{\dot{M} - \dot{E}}{\langle \dot{M} \rangle}. \tag{4.6}$$

Values greater than unity indicate more energy is being released than can be accounted for by the rest mass of the infalling material. The run of this efficiency in the various simulations is shown in Figure 4.27, with averages presented in the last column of Table 4.1.

We find notably larger efficiencies than Tchekhovskoy et al., who record values of approximately 1.4 for their run with similar parameters. To the extent that our mass accretion rates are similar to theirs, this indicates we are seeing significantly more energy outflow. Note that had we chosen a more interior r_0 to evaluate $\dot{E}$ (see Figure 4.25) we would have calculated lower values of η. In fact, with $r_0 = 2$ we would have found a net inflow of energy and $\eta < 0$.

There is good evidence that our accretion disks proper have converged with higher

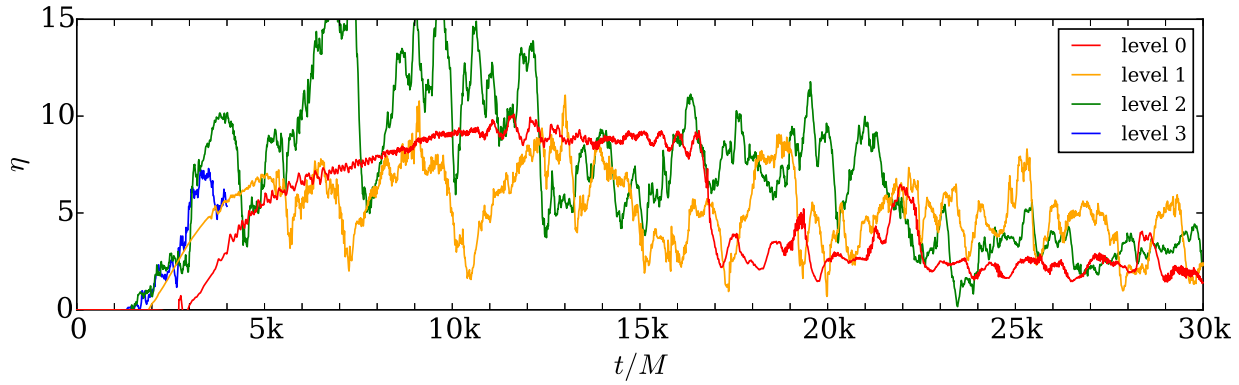

Figure 4.27: Efficiency parameter η in the MAD simulations as a function of time for four resolutions. All simulations are sampled at a rate of $\Delta t = 10\ M$.

resolution – if not in all the details then at least in important characteristics like magnetic flux and field geometry. However it is not clear that quantitative behavior of the near-vacuum outflow region is entirely free of numerical effects. This is not unexpected, given the difficulty finite-volume codes have when there is negligible matter. This is especially the case then that matter is highly magnetized, moving relativistically, and in a strongly curved region of spacetime.

Chapter 5

Conclusion

Here we offer some closing remarks on the preceding material. We will consider the broader implications of our investigation into magnetically arrested disk (MAD) accretion, followed by anticipations of further applications of Athena++. We will also mention code developments slated for the future.

5.1 Implications of MAD Studies

An important question about MAD accretion is how ubiquitous it is. MAD has been seen in several simulations, but it has heretofore been studied by a small number of groups using variants on the same code. We seek to understand whether the MAD scenario arises in our own simulations. As part of this we vary both resolution and numerical implementation details in order to ensure we are making statements about nature and not just about our code.

The key features of MAD as applied to black holes are the buildup of ordered poloidal magnetic field, the suppression of accretion in the inner regions, and the extraction of energy from the black hole spin. We do see this poloidal field accumulate, especially at higher resolutions (Figures 4.12 and 4.21). That is, although the turbulence in the disk would seem to disorder the field, there is a net flux (by design of the initial conditions) that cannot help but accumulate. Once it is in the near-vacuum region, the field dominates the dynamics and it is no longer susceptible to being distorted by the fluid.

More quantitatively, consider the steady-state flux Φ threading spheres of various radii, as shown in Figure 4.25. The sign of the slope of the curves indicates the field lines near the black hole are not closing back at small radii. If they were, small spheres would be threaded by field loops that remained interior to large spheres, and we would see Φ increasing inward. The fact that the slope is shallow indicates that whatever field threads a sphere at large radius also threads spheres at smaller radius. That is, the magnetic field is brought all the way to the horizon. This is consistent with the MAD scenario.

MAD then proposes that the accumulated field will prevent matter from freely accreting. Indeed we do see highly magnetized regions very close to the horizon in the top panels of Figures 4.23 and 4.24. In the lower panels however one can clearly see some matter streaming in. On the other hand there is a clear sign of disruption at the highest resolution in the upper right panels of Figures 4.19 and 4.22. A future investigation might focus on measuring accretion rates as a function of field geometry – would the inflow be faster and more continuous if not for the accumulated net flux?

The energy extraction issue is more uncertain. We do see values of η greater than 100% (Figure 4.27), but these are values large enough to be suspicious. Indeed they depend highly upon the measured radius r_0 chosen (cf. Figure 4.25, especially the middle panel).

The large dynamical range of densities, magnetizations, and velocities found in simulations that include both an accretion disk and a jet presents a great challenge for a magnetohydrodynamics (MHD) code. Even to the extent that MHD is applicable across the domain, numerical issues have the potential to alter the solution. As already discussed, the need for floors on density and gas pressure, as well as the ceiling on Lorentz factors, reduces the reliability of simulations in the very jet region we may wish to study. Improvements to robustness (see §5.3.1) can help to secure results in this area.

5.2 Future Scientific Investigations

Development of a new general-relativistic MHD (GRMHD) code would be a high price indeed if it were only to be applied to one particular accretion problem. Fortunately there are many problems that can be attacked with the code we have developed.

We pause to consider the comparative advantages offered by Athena++. It is fast and scalable; thus it can run models for long times covering large domains. Algorithmically it has the capability of minimizing diffusion with advanced Riemann solvers, as well as evolving the magnetic field in a natural divergence-free manner. Moreover it can accurately handle the polar singularity in spherical coordinate systems. Finally, Athena++ can also employ adaptive mesh refinement (AMR).

One potential application is the study of disks tilted relative to their black hole host spins, including the associated alignment, precession, and outflows. The vast majority of GRMHD simulations to date are done with aligned disks, to the extent that the black holes are even modeled with spin. Tilted disks challenge a number of codes by demanding resolution off the equatorial plane. This can easily be achieved with AMR without requiring large swaths of vacuum to be refined. Moreover, as momentum transport and thus alignment can be affected by standing shocks in the disk, it pays to have a code with minimal diffusion in order to get accurate results.

One of the most prominent observational features of black hole accretion is the powerful jet that can be launched. However, there is still something of a disconnect between the GRMHD simulations, focusing on accretion disk physics, and the large-scale jet propagation simulations, focusing on special relativity and emission from the jet. As a result, despite our understanding of jets themselves, we have no clear answer to the question of what accretion configurations lead to relativistic jets. A systematic study focusing on the parameter space covering plausible near-horizon conditions can yield insights into this connection. Key to such a study is the use of a coordinate system that does not contain a singularity along the very object of study, and thus our code is a prime candidate.

One last application we will mention is the relatively young field of tidal disruption events. The idea is that a star passing too close to a supermassive black hole will be tidally shredded, leading to a short-term burst of intense accretion. This scenario is believed to explain transient phenomena observed in active galactic nuclei. Again, we find an application that demands the range of scales accessible only with mesh refinement and a code that scales well computationally.

5.3 Roadmap for Future Code Development

The modular design of Athena++ allows it to support new components on top of the existing infrastructure. Here we briefly summarize some of the exciting prospects for the future of the code, both numerical and physical. Note that these need not be implemented by any one person, and indeed a number of them are applicable beyond the scope of GRMHD and so are likely to be included as a result of interest in other fields.

5.3.1 Algorithmic Improvements and Extensions

A unique feature of Athena++ is its ability to run with a number of different Riemann solvers. This allows for a fair comparison between different solvers where all else is kept the same. While some comparisons were presented in §3.1, §3.2, and §3.3, these are contrived examples rather than realistic turbulent flows. An easy future project then is to quantify the net benefit (or lack thereof) obtained by using advanced Riemann solvers.

Part of the difficulty in using less diffusive methods is that sharp gradients can wreak havoc with variable inversion. Sudden changes of conserved values can make the initial guess used in inversion procedures (§2.3.1) too far from the new solution to converge. More robust methods, such as possibly that of Newman and Hamlin (2014), should be explored.

All simulations discussed here are second order in space (for smooth enough flows away from coordinate singularities) and time. However the code is designed to be able to incorporate higher-order methods. These can potentially achieve a desired accuracy with

less overall cost by reducing the number of cells needed. Higher-order spatial reconstruction is already implemented in some codes. There is the caveat briefly mentioned in §2.2.2, in that the source terms are only computed at second-order spatial accuracy. However even a method that is still formally second order often in practice stands to gain by having higher-order constituent parts where possible. For temporal accuracy there is no hindrance to going to higher orders in any scheme, like the van Leer integrator, that uses partial forward-Euler substeps.

5.3.2 New Physics

The frontier of black hole accretion simulations is rapidly evolving, with ever more accurate physical models being incorporated. One area that has seen much recent development is that of fluid dynamics beyond ideal MHD. For example, Chandra et al. (2015) describes how to modify MHD in a way amenable to general relativity so as to incorporate weakly collisional effects such as anisotropic heat conduction guided by the magnetic field. Treating electrons separately from protons in the thermodynamic sense is explored in Ressler et al. (2015). Without modifications such as these traditional GRMHD methods would be unsuitable for studying the extremely diffuse accretion in low-luminosity systems such as the supermassive black hole in our own galaxy.

At the other end of the spectrum of accretion rates, high-luminosity systems result in radiation being dynamically important. When optical depths are large across a single cell, radiation merely acts as a relativistically hot fluid and can be treated easily enough; when they are small across the entire domain the radiation streams out of the system, which can be evolved without considering the effects of the photons. The challenge is in the intermediate regime, where the equations of radiative transport must be incorporated at some level into the simulation. This is now being done in GRMHD with both closure schemes (McKinney et al., 2014) and Monte Carlo methods (Ryan et al., 2015). By incorporating such a scheme, Athena++ can be used in new regions of the parameter space of accretion systems. In addition to the aforementioned methods, work is currently underway for implementing a

method of short characteristics compatible with general relativity into the code.

The aforementioned physical models are examples of what can be incorporated directly into the code. At the same time, we note that simulations do not exist in a vacuum: work must be done to connect the results of simulations to observations. Future work with Athena++ should involve the use of post-processing methods such as the Monte Carlo radiative transport described in Dolence et al. (2009). In this way we can make predictions more directly related to observations of electromagnetic emission from real accreting black holes.

Appendix A

Frame Transformation

Given a coordinate basis $\{\vec{\partial}_{(\mu)}\}$, we desire a new basis $\{\vec{e}_{(\hat{\mu})}\}$ satisfying the enumerated list of properties in §2.3.3, which we repeat here:

1. $\vec{e}_{(\hat{\mu})}$ must be orthogonal to $\vec{e}_{(\hat{\nu})}$ for all $\mu \neq \nu$.
2. Each $\vec{e}_{(\hat{\mu})}$ must be normalized to have an inner product of ± 1 with itself, with $\vec{e}_{(\hat{t})}$ being timelike and $\vec{e}_{(\hat{\jmath})}$ being spacelike.
3. $\vec{e}_{(\hat{t})}$ must be orthogonal to surfaces of constant x^0.
4. The projection of $\vec{e}_{(\hat{x})}$ onto the surface of constant x^0 is orthogonal to the surface of constant x^i within that submanifold.

As in the main text, we note that tensor components are indicated with unparenthesized subscripts and superscripts, with parentheses reserved for indices into generic sets. For concreteness we will work with an interface of constant x^1; results in the other directions follow from cyclic permutations $1 \mapsto 2 \mapsto 3$.

Property 3 tells us

$$e_{\mu}^{(\hat{t})} = A(1, 0, 0, 0) \tag{A.1}$$

for some A, which is determined by property 2:

$$-1 = g^{\mu\nu} e_{\mu}^{(\hat{t})} e_{\nu}^{(\hat{t})} = A^2 g^{00}. \tag{A.2}$$

Thus we have

$$A = -(-g^{00})^{-1/2}, \tag{A.3}$$

where the sign is chosen such that the $\hat{t}$-direction is forward in time, as can be easily checked in the case $g_{\mu\nu} = \eta_{\mu\nu}$. Our first basis vector has components

$$e^{\mu}_{(\hat{t})} = g^{\mu\nu} e^{(\hat{t})}_{\nu} = A g^{0\mu}, \tag{A.4}$$

and this is nothing more than the unit timelike normal $\hat{n}$ used in 3+1 foliations of spacetime.

Next we set the $\hat{x}$ direction. It might be tempting to simply have it be parallel to $\vec{\partial}_{(1)}$, since this vector is already orthogonal to $\vec{e}_{(\hat{t})}$ by construction of the latter. However, this complicates the transformation of fluxes back to the global frame, as it results in $\hat{y}$- and $\hat{z}$-fluxes entering into the formula for x^1-fluxes. It is a more serious error to make $\vec{e}_{(\hat{x})}$ orthogonal to the 3-surface of constant x^1. Though this would lead to a very simple formula for the flux conversion (requiring no knowledge of conserved quantities inside the wavefan), this vector can only be orthogonal to $\vec{e}_{(\hat{t})}$ if g^{01} vanishes.

Instead, we enforce property 4. This amounts to saying the projection of $\vec{e}_{(\hat{x})}$ onto the constant-x^0 subspace must have components

$${}^{(3)}e^{(\hat{x})}_{i} \propto (1, 0, 0). \tag{A.5}$$

Raising indices with the appropriate 3-metric components $\gamma^{ij} = g^{ij} - g^{0i} g^{0j} / g^{00}$ and noting $e^{0}_{(\hat{x})} = 0$ (by property 1 and (A.1)), we have

$$e^{\mu}_{(\hat{x})} = B(g^{01} g^{0\mu} - g^{00} g^{1\mu}). \tag{A.6}$$

The normalization is again set by property 2:

$$1 = g_{\mu\nu} e^{\mu}_{(\hat{x})} e^{\nu}_{(\hat{x})} = B^2 g^{00} (g^{00} g^{11} - g^{01} g^{01}). \tag{A.7}$$

Choosing the sign such that $e^{\mu}_{(\hat{x})} = (0, +1, 0, 0)$ if $g_{\mu\nu} = \eta_{\mu\nu}$, we take

$$B = \left(g^{00}(g^{00}g^{11} - g^{01}g^{01})\right)^{-1/2}. \tag{A.8}$$

Property 1, together with (A.1), forces

$$e^{0}_{(\hat{y})} = e^{0}_{(\hat{z})} = 0; \tag{A.9}$$

that is, $\vec{e}_{(\hat{y})}$ and $\vec{e}_{(\hat{z})}$ must lie in the constant-x^0 plane. From (A.5) and property 1 we have

$$e^{1}_{(\hat{y})} = e^{1}_{(\hat{z})} = 0, \tag{A.10}$$

which conveniently implies no $\hat{y}$- or $\hat{z}$-fluxes will be needed from the special-relativistic Riemann solver.

For slight convenience we choose $e^{2}_{(\hat{z})}$ to be 0. This corresponds to choosing an orientation of $\vec{e}_{(\hat{y})}$ and $\vec{e}_{(\hat{z})}$ in the plane they must span. Then we have

$$e^{\mu}_{(\hat{z})} = C(0, 0, 0, 1). \tag{A.11}$$

From property 2 we know

$$1 = g_{\mu\nu}e^{\mu}_{(\hat{z})}e^{\nu}_{(\hat{z})} = C^2 g_{33}, \tag{A.12}$$

so

$$C = (g_{33})^{-1/2}, \tag{A.13}$$

where again we check that we have a sensible definition when $g_{\mu\nu} = \eta_{\mu\nu}$: $e^{\mu}_{(\hat{z})} = (0, 0, 0, +1)$.

Finally, we solve for the two nonzero components of $\vec{e}_{(\hat{y})}$. Orthogonality (property 1) with $\vec{e}_{(\hat{z})}$, together with normalization (property 2), constrain the solution to be

$$e^{\mu}_{(\hat{y})} = D(0, 0, g_{33}, -g_{23}), \tag{A.14}$$

where

$$D = \Big(g_{33}(g_{22}g_{33} - g_{23}g_{23})\Big)^{-1/2}. \tag{A.15}$$

Once again we check the sign choice: $e^{\mu}_{(\hat{y})} = (0, 0, +1, 0)$ if $g_{\mu\nu} = \eta_{\mu\nu}$.

Given the formulas (A.4), (A.6), (A.11), (A.14), it is a simple matter to assemble the transformation matrix (2.49):

$$M^{\mu}{}_{\hat{\nu}} = e^{\mu}_{(\hat{\nu})} = \begin{pmatrix} Ag^{00} & 0 & 0 & 0 \\ Ag^{01} & B(g^{01}g^{01} - g^{00}g^{11}) & 0 & 0 \\ Ag^{02} & B(g^{01}g^{02} - g^{00}g^{12}) & Dg_{33} & 0 \\ Ag^{03} & B(g^{01}g^{03} - g^{00}g^{13}) & -Dg_{23} & C \end{pmatrix}, \tag{A.16}$$

with

$$A = -(-g^{00})^{-1/2}, \tag{A.17a}$$

$$B = \Big(g^{00}(g^{00}g^{11} - g^{01}g^{01})\Big)^{-1/2}, \tag{A.17b}$$

$$C = (g_{33})^{-1/2}, \tag{A.17c}$$

$$D = \Big(g_{33}(g_{22}g_{33} - g_{23}g_{23})\Big)^{-1/2}. \tag{A.17d}$$

The inverse transformation (2.51) is the inverse matrix

$$M^{\hat{\mu}}{}_{\nu} = \begin{pmatrix} -A & 0 & 0 & 0 \\ Bg^{01} & -Bg^{00} & 0 & 0 \\ B^2Eg^{00}/Dg_{33} & B^2Fg^{00}/Dg_{33} & 1/Dg_{33} & 0 \\ I & J & (1/C)g_{23}/g_{33} & 1/C \end{pmatrix}, \tag{A.18}$$

with

$$E = g^{01}g^{12} - g^{11}g^{02}, \tag{A.19a}$$

$$F = g^{01}g^{02} - g^{00}g^{12}, \tag{A.19b}$$

$$G = g^{01}g^{13} - g^{11}g^{03}, \tag{A.19c}$$

$$H = g^{01}g^{03} - g^{00}g^{13}, \tag{A.19d}$$

$$I = \frac{B^2 g^{00}}{C}\left(G + \frac{E g_{23}}{g_{33}}\right), \tag{A.19e}$$

$$J = \frac{B^2 g^{00}}{C}\left(H + \frac{F g_{23}}{g_{33}}\right), \tag{A.19f}$$

Both matrices are lower-diagonal, mostly as a result of our enumerated properties but also thanks to our choice of rotation $e^2_{(\hat{z})} = 0$.

Appendix B

Extremal Wavespeeds in Special-Relativistic Magnetohydrodynamics

Our Riemann solvers all require knowing extremal leftgoing and rightgoing signal speeds. If information propagates faster than the values we use, numerical instability ensues. On the other hand, if we overestimate these speeds (for example by always using the speed of light) we introduce unnecessary numerical dissipation.

In the case of magnetohydrodynamics, the speeds to use are the fast magnetosonic wavespeeds. As these are nontrivial to find, we detail here the method in the case of special relativity. Note that when using frame transformations in general relativity we will be computing wavespeeds in the transformed system, so the following procedure still applies.

Our procedure is based on the equations given in Mignone and Bodo (2006) for the magnetosonic speeds. We consider only a constant-Γ equation of state. Let x denote the direction in which we are interested in calculating propagation speeds (e.g. for interfaces of constant x). We assume the density ρ, gas pressure $p_{\rm gas}$, 3-velocity components v^i, magnetic field components B^i, contravariant magnetic field components b^t and b^x, and magnetic pressure $p_{\rm mag}$ are known. We can easily calculate the gas and total enthalpies,

$$w_{\rm gas} = \rho + \frac{\Gamma}{\Gamma - 1} p_{\rm gas}, \tag{B.1a}$$

$$w_{\rm tot} = w_{\rm gas} + 2p_{\rm mag}, \tag{B.1b}$$

as well as the square of the sound speed,

$$c_{\rm s}^2 = \frac{\Gamma p_{\rm gas}}{w_{\rm gas}} \tag{B.2}$$

(Mignone and Bodo, 2005), and the square of the Lorentz factor,

$$\gamma^2 = \frac{1}{1 - v_i v^i}. \tag{B.3}$$

Before solving the general case, we consider two degenerate limits. In all that follows, explicit square roots are understood to return 0 if their arguments are negative.

B.1 Vanishing Velocity

If there is no velocity, then the magnetosonic speeds will be symmetrically distributed around 0. To be precise, if $v_i v^i < 10^{-12}$, we apply (57) from Mignone and Bodo (2006). Rewritten, this becomes

$$\lambda^4 + a_2\lambda^2 + a_0 = 0, \tag{B.4}$$

where we define

$$a_2 = -\frac{1}{w_{\rm tot}}\Big(2p_{\rm mag} + c_{\rm s}^2(w_{\rm gas} + B_x^2)\Big), \tag{B.5a}$$

$$a_0 = \frac{c_{\rm s}^2 B_x^2}{w_{\rm tot}}. \tag{B.5b}$$

We can then calculate the wavespeeds of interest according to

$$s^2 = a_2^2 - 4a_0, \tag{B.6}$$

$$s = \sqrt{s^2}, \tag{B.7}$$

$$\lambda_\pm = \pm\sqrt{\frac{s - a_2}{2}}. \tag{B.8}$$

B.2 Vanishing Longitudinal Magnetic Field

If the fluid velocity does not vanish but the longitudinal magnetic field effectively does ($|B^x| < 10^{-7}$ in our code), then (58) from Mignone and Bodo applies. Like them, we define

$$Q = 2p_{\mathrm{mag}} - c_{\mathrm{s}}^2(v^y B^y + v^z B^z)^2. \tag{B.9}$$

Then the fast magnetosonic speeds satisfy

$$\lambda^2 + a_1\lambda + a_0 = 0, \tag{B.10}$$

where we have

$$a_1 = -\frac{2w_{\mathrm{gas}}\gamma^2 v^x(1-c_{\mathrm{s}}^2)}{w_{\mathrm{gas}}(c_{\mathrm{s}}^2 + \gamma^2(1-c_{\mathrm{s}}^2)) + Q}, \tag{B.11a}$$

$$a_0 = \frac{w_{\mathrm{gas}}(\gamma^2 v_x^2(1-c_{\mathrm{s}}^2) - c_{\mathrm{s}}^2) - Q}{w_{\mathrm{gas}}(c_{\mathrm{s}}^2 + \gamma^2(1-c_{\mathrm{s}}^2)) + Q}. \tag{B.11b}$$

Selectively applying the alternate form of the quadratic formula to avoid cancellation errors (Press et al., 1997), we find

$$s^2 = a_1^2 - 4a_0; \tag{B.12}$$

$$s = \sqrt{s^2}; \tag{B.13}$$

$$\lambda_{\pm} = \begin{cases} -\dfrac{2a_0}{a_1 \pm s}, \\ -\dfrac{1}{2}(a_1 \mp s); \end{cases} \tag{B.14}$$

where the first case is taken if $s^2 >= 0$ and either $a_1 >= 0$ (+ case) or $a_1 < 0$ (− case).

B.3 General Case

If neither of the former two cases apply, we are forced to solve the quartic equation (56) from Mignone and Bodo (2006). Here we detail the method used to solve this equation, which

has the advantage of using only real arithmetic. This particular procedure is justified on the physical grounds that the roots of our equation are magnetosonic speeds and so must be real.

Begin by rewriting the equation as

$$\lambda^4 + a_3\lambda^3 + a_2\lambda^2 + a_1\lambda + a_0 = 0, \tag{B.15}$$

with

$$a_3 = -\frac{1}{C}\Big((4A + 2B)v^x - 2c_{\rm s}^2 b^t b^x\Big), \tag{B.16a}$$

$$a_2 = \frac{1}{C}\Big(6Av_x^2 + B(v_x^2 - 1) + c_{\rm s}^2(b_t^2 - b_x^2)\Big), \tag{B.16b}$$

$$a_1 = -\frac{1}{C}(4Av_x^3 - 2Bv^x + 2c_{\rm s}^2 b^t b^x), \tag{B.16c}$$

$$a_0 = \frac{1}{C}(Av_x^4 - Bv_x^2 + c_{\rm s}^2 b_x^2). \tag{B.16d}$$

Here we define

$$A = \gamma^4 w_{\rm gas}(1 - c_{\rm s}^2), \tag{B.17a}$$

$$B = \gamma^2(2p_{\rm mag} + w_{\rm gas}c_{\rm s}^2), \tag{B.17b}$$

$$C = A + B - c_{\rm s}^2 b_t^2. \tag{B.17c}$$

Next transform to the reduced quartic equation in $\lambda' = \lambda + a_3/4$:

$$(\lambda')^4 + b_2(\lambda')^2 + b_1\lambda' + b_0 = 0, \tag{B.18}$$

with

$$b_2 = a_2 - \frac{3}{8}a_3^2, \tag{B.19a}$$

$$b_1 = a_1 - \frac{1}{2}a_2a_3 + \frac{1}{8}a_3^3, \tag{B.19b}$$

$$b_0 = a_0 - \frac{1}{4}a_1a_3 + \frac{1}{16}a_2a_3^2 - \frac{3}{256}a_3^4. \tag{B.19c}$$

In order to write (B.18) in terms of quadratic equations, we must find a real root x to the resolvent cubic equation

$$x^3 + c_2 x^2 + c_1 x + c_0 = 0, \tag{B.20}$$

with

$$c_2 = -b_2, \tag{B.21a}$$

$$c_1 = -4b_0, \tag{B.21b}$$

$$c_0 = 4b_0 b_2 - b_1^2. \tag{B.21c}$$

Press et al. (1997) provide just such a method. Define the quantities

$$D = \frac{1}{9}(c_2^2 - 3c_1), \tag{B.22a}$$

$$E = \frac{1}{54}(2c_2^3 - 9c_1 c_2 + 27c_0), \tag{B.22b}$$

$$s^2 = E^2 - D^3. \tag{B.22c}$$

If $s^2 < 0$, then our root x_0 is given by

$$\theta = \cos^{-1}\left(\frac{E}{\sqrt{D^3}}\right), \tag{B.23}$$

$$x_0 = -2\sqrt{D}\cos\left(\frac{\theta}{3}\right) - \frac{1}{3}c_2. \tag{B.24}$$

Otherwise, we have

$$s = \sqrt{s^2}, \tag{B.25}$$

$$F = -\operatorname{sgn}(E)\sqrt[3]{|E| + s}, \tag{B.26}$$

$$G = \frac{D}{F}, \tag{B.27}$$

$$x_0 = F + G - \frac{1}{3}c_2. \tag{B.28}$$

where G is taken to be 0 if F vanishes. These are essentially the same formula with different

branch cut considerations.

Using the real cubic root x_0, we define quadratic coefficients

$$d_1 = +\sqrt{x_0 - b_2}, \qquad e_1 = -\sqrt{x_0 - b_2}, \tag{B.29a}$$

$$d_0 = \frac{1}{2}x_0 - \mathrm{sgn}(b_1)t, \qquad e_0 = \frac{1}{2}x_0 + \mathrm{sgn}(b_1)t, \tag{B.29b}$$

where

$$t = \sqrt{\frac{1}{4}x_0^2 - b_0}. \tag{B.30}$$

The roots of (B.18) are then given by

$$u^2 = d_1^2 - 4d_0; \qquad v^2 = e_1^2 - 4e_0; \tag{B.31}$$

$$u = \sqrt{u^2}; \qquad v = \sqrt{v^2}; \tag{B.32}$$

$$\lambda_1' = \begin{cases} -\dfrac{2d_0}{d_1 - u}, \\ -\dfrac{1}{2}(d_1 + u); \end{cases} \qquad \lambda_3' = \begin{cases} -\dfrac{2e_0}{e_1 - v}, \\ -\dfrac{1}{2}(e_1 + v); \end{cases} \tag{B.33}$$

$$\lambda_2' = \begin{cases} -\dfrac{2d_0}{d_1 + u}, \\ -\dfrac{1}{2}(d_1 - u); \end{cases} \qquad \lambda_4' = \begin{cases} -\dfrac{2e_0}{e_1 + v}, \\ -\dfrac{1}{2}(e_1 - v). \end{cases} \tag{B.34}$$

In the above equations the first case is taken if the linear coefficient (d_1 or e_1 as appropriate) is nonnegative.

Finally, we take the fast magnetosonic speeds of interest to be

$$\lambda_- = \max\left\{\min\{\lambda_1', \lambda_3'\} - \frac{1}{4}a_3, -1\right\}, \tag{B.35a}$$

$$\lambda_+ = \min\left\{\max\{\lambda_2', \lambda_4'\} - \frac{1}{4}a_3, +1\right\}. \tag{B.35b}$$

We note that a very similar procedure can be used for variable inversion in the case of special-relativistic hydrodynamics, as this is also given by a quartic equation, (2.43).

Appendix C

Coordinate System Examples

Here we explicitly provide the core formulas for several metrics. For each metric we are interested in three sets of quantities:

- There are the fundamental, pointwise-defined geometric quantities $g_{\mu\nu}$, $g^{\mu\nu}$, $\sqrt{-g}$, and $\Gamma^{\sigma}_{\mu\nu}$, recorded here in their entirety.
- Next there are the finite-volume quantities ΔV, ΔA, ΔL, and Δx^i defined in (2.29), (2.30), (2.31), and (2.32), as well as the volume-averaged midpoints $\bar{x}^i$ at which various quantities are defined.
- Finally, there are the transformation matrices with components $M^{\hat{\mu}}{}_{\nu}$ and $M^{\mu}{}_{\hat{\nu}}$. There is a pair of such matrices for each of the three types of interfaces, with the formulas for the x^1-interface versions given in (2.51) and (2.49).

The i, j, k index of cells is understood if not specified. Interfaces and edges will be indexed with i_+ between i and $i+1$ and i_- between $i-1$ and i as appropriate, and likewise for j and k. Coordinates will be subscripted with $-$ or $+$ when they refer to interface values.

C.1 Minkowski Spacetime

The simplest spacetime we can have is the flat Minkowski spacetime. Several coordinate systems are given.

C.1.1 Minkowski Coordinates

First we consider the Minkowski (Cartesian) coordinates of special relativity: $x^\mu = (t, x, y, z)$. The metric and its inverse are

$$g_{\mu\nu} = g^{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.1}$$

and we have

$$\sqrt{-g} = 1. \tag{C.2}$$

The connection coefficients all vanish:

$$\Gamma^\sigma_{\mu\nu} = 0. \tag{C.3}$$

The volumes are given by

$$\Delta V = (x_+ - x_-)(y_+ - y_-)(z_+ - z_-), \tag{C.4}$$

the areas by

$$\Delta A_{i_-,j,k} = (y_+ - y_-)(z_+ - z_-), \tag{C.5a}$$

$$\Delta A_{i,j_-,k} = (x_+ - x_-)(z_+ - z_-), \tag{C.5b}$$

$$\Delta A_{i,j,k_-} = (x_+ - x_-)(y_+ - y_-), \tag{C.5c}$$

the lengths by

$$\Delta L_{i,j_-,k_-} = x_+ - x_i, \tag{C.6a}$$

$$\Delta L_{i_-,j,k_-} = y_+ - y_i, \tag{C.6b}$$

$$\Delta L_{i_-,j_-,k} = z_+ - z_i, \tag{C.6c}$$

and the widths by

$$\Delta W_x = x_+ - x_-, \tag{C.7a}$$

$$\Delta W_y = y_+ - y_-, \tag{C.7b}$$

$$\Delta W_z = z_+ - z_-. \tag{C.7c}$$

The volume-averaged midpoints are

$$\bar{x} = \frac{1}{2}(x_- + x_+), \tag{C.8a}$$

$$\bar{y} = \frac{1}{2}(y_- + y_+), \tag{C.8b}$$

$$\bar{z} = \frac{1}{2}(z_- + z_+). \tag{C.8c}$$

The transformation matrices are very simple:

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i_-,j,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.9a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j_-,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{C.9b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j,k_-} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \tag{C.9c}$$

C.1.2 Spherical Coordinates

Next we consider flat spacetime but recast in terms of spherical spatial variables r, θ, and ϕ. The metric and its inverse are diagonal:

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & r^2 & 0 \\ 0 & 0 & 0 & r^2 \sin^2\theta \end{pmatrix}, \qquad g^{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/r^2 & 0 \\ 0 & 0 & 0 & \csc^2\theta/r^2 \end{pmatrix}. \tag{C.10}$$

We have the factor

$$\sqrt{-g} = r^2 |\sin\theta|. \tag{C.11}$$

The nonzero connection coefficients are

$$\Gamma^1_{22} = -r, \tag{C.12a}$$

$$\Gamma^1_{33} = -r\sin^2\theta, \tag{C.12b}$$

$$\Gamma^2_{12} = \Gamma^2_{21} = \frac{1}{r}, \tag{C.12c}$$

$$\Gamma^2_{33} = -\sin\theta\cos\theta, \tag{C.12d}$$

$$\Gamma^3_{13} = \Gamma^3_{31} = \frac{1}{r}, \tag{C.12e}$$

$$\Gamma^3_{23} = \Gamma^3_{32} = \cot\theta. \tag{C.12f}$$

The volumes are given by

$$\Delta V = \frac{1}{3}(r_+^3 - r_-^3)|\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-), \tag{C.13}$$

the areas by

$$\Delta A_{i_-,j,k} = r_-^2 |\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-), \tag{C.14a}$$

$$\Delta A_{i,j_-,k} = \frac{1}{3}(r_+^3 - r_-^3)|\sin\theta_-|(\phi_+ - \phi_-), \tag{C.14b}$$

$$\Delta A_{i,j,k_-} = \frac{1}{3}(r_+^3 - r_-^3)|\cos\theta_- - \cos\theta_+|, \tag{C.14c}$$

the lengths by

$$\Delta L_{i,j_-,k_-} = \frac{1}{3}(r_+^3 - r_-^3)|\sin\theta_-|, \tag{C.15a}$$

$$\Delta L_{i_-,j,k_-} = r_-^2|\cos\theta_- - \cos\theta_+|, \tag{C.15b}$$

$$\Delta L_{i_-,j_-,k} = r_-^2|\sin\theta_-|(\phi_+ - \phi_-), \tag{C.15c}$$

and the widths by

$$\Delta W_r = r_+ - r_-, \tag{C.16a}$$

$$\Delta W_\theta = \bar{r}(\theta_+ - \theta_-), \tag{C.16b}$$

$$\Delta W_\phi = \bar{r}|\sin\bar{\theta}|(\phi_+ - \phi_-). \tag{C.16c}$$

The volume-averaged midpoints are

$$\bar{r} = \left(\frac{r_-^3 + r_+^3}{2}\right)^{1/3}, \tag{C.17a}$$

$$\bar{\theta} = \cos^{-1}\left(\frac{\cos\theta_- + \cos\theta_+}{2}\right), \tag{C.17b}$$

$$\bar{\phi} = \frac{1}{2}(\phi_- + \phi_+). \tag{C.17c}$$

The transformation matrices are as follows:

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & r & 0 \\ 0 & 0 & 0 & r|\sin\theta| \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i_-,j,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1/r & 0 \\ 0 & 0 & 0 & |\csc\theta|/r \end{pmatrix}, \tag{C.18a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & r & 0 \\ 0 & 0 & 0 & r|\sin\theta| \\ 0 & 1 & 0 & 0 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j_-,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1/r & 0 & 0 \\ 0 & 0 & |\csc\theta|/r & 0 \end{pmatrix}, \tag{C.18b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & r|\sin\theta| \\ 0 & 1 & 0 & 0 \\ 0 & 0 & r & 0 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j,k_-} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1/r \\ 0 & |\csc\theta|/r & 0 & 0 \end{pmatrix}. \tag{C.18c}$$

C.1.3 Sinusoidal Coordinates

A useful coordinate system for testing the correctness of the code is obtained from Minkowski coordinates (t, x, y, z) according to

$$T = t, \tag{C.19a}$$

$$X = x, \tag{C.19b}$$

$$Y = y + a\sin(kx), \tag{C.19c}$$

$$Z = z, \tag{C.19d}$$

where a and k are parameters for the transformation. These "snake" coordinates (T, X, Y, Z) have lines of constant Y that are sinusoidal in the xy-plane. For these coordinates define

$$\alpha = \sqrt{1 + a^2k^2\cos^2(kX)}, \tag{C.20a}$$

$$\beta = ak\cos(kX), \tag{C.20b}$$

where bars and subscripts applied to these variables are inherited by the X's in the definition.

The metric and its inverse are given by

$$g_{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & \alpha^2 & -\beta & 0 \\ 0 & -\beta & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad g^{\mu\nu} = \begin{pmatrix} -1 & 0 & 0 & 0 \\ 0 & 1 & \beta & 0 \\ 0 & \beta & \alpha^2 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.21}$$

and we have

$$\sqrt{-g} = 1. \tag{C.22}$$

The only nonzero connection coefficient is

$$\Gamma^2_{11} = ak^2 \sin(kX). \tag{C.23}$$

The volumes are given by

$$\Delta V = (X_+ - X_-)(Y_+ - Y_-)(Z_+ - Z_-), \tag{C.24}$$

the areas by

$$\Delta A_{i_-,j,k} = (Y_+ - Y_-)(Z_+ - Z_-), \tag{C.25a}$$

$$\Delta A_{i,j_-,k} = (X_+ - X_-)(Z_+ - Z_-), \tag{C.25b}$$

$$\Delta A_{i,j,k_-} = (X_+ - X_-)(Y_+ - Y_-), \tag{C.25c}$$

the lengths by

$$\Delta L_{i,j_-,k_-} = X_+ - X_i, \tag{C.26a}$$

$$\Delta L_{i_-,j,k_-} = Y_+ - Y_i, \tag{C.26b}$$

$$\Delta L_{i_-,j_-,k} = Z_+ - Z_i, \tag{C.26c}$$

and the widths by

$$\Delta W_X = \int_{X_-}^{X_+} \alpha \, \mathrm{d}X \geq \int_{X_-}^{X_+} \mathrm{d}X = X_+ - X_-, \tag{C.27a}$$

$$\Delta W_Y = Y_+ - Y_-, \tag{C.27b}$$

$$\Delta W_Z = Z_+ - Z_-. \tag{C.27c}$$

The bound on the elliptic integral can be used to conservatively limit the step sizes (ΔX being used for no other purpose). The volume-averaged midpoints are

$$\bar{X} = \frac{1}{2}(X_- + X_+), \tag{C.28a}$$

$$\bar{Y} = \frac{1}{2}(Y_- + Y_+), \tag{C.28b}$$

$$\bar{Z} = \frac{1}{2}(Z_- + Z_+). \tag{C.28c}$$

The transformation matrices are given by

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & -\beta & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i_-,j,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & \beta & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.29a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1/\alpha & 0 \\ 0 & 0 & 0 & 1 \\ 0 & \alpha & -\beta/\alpha & 0 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j_-,k} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & \beta/\alpha & 0 & 1/\alpha \\ 0 & \alpha & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{C.29b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \\ 0 & -\beta & 1 & 0 \end{pmatrix}, \qquad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j,k_-} = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & \beta & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \tag{C.29c}$$

C.1.4 Tilted Coordinates

Just as sinusoidal coordinates can be used to test space–space terms in the metric, we use tilted coordinates to test time–space terms. For a shift parameterized by a, $|a| < 1$, define

$$\alpha = \sqrt{1 + a^2}, \tag{C.30a}$$

$$\beta = \sqrt{1 - a^2}. \tag{C.30b}$$

Then we have the transformation from Minkowski coordinates (t, x, y, z) according to

$$T = \frac{1}{\alpha}(t + ax), \tag{C.31a}$$

$$X = \frac{1}{\alpha}(x - at), \tag{C.31b}$$

$$Y = y, \tag{C.31c}$$

$$Z = z. \tag{C.31d}$$

Observers moving with x-velocity a have worldlines in the T-direction, while surfaces of constant T correspond to surfaces of constant time for observers with x-velocity $-a$.

The metric is its own inverse,

$$g_{\mu\nu} = g^{\mu\nu} = \begin{pmatrix} -\beta^2/\alpha^2 & 2a/\alpha^2 & 0 & 0 \\ 2a/\alpha^2 & \beta^2/\alpha^2 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.32}$$

and we have

$$\sqrt{-g} = 1. \tag{C.33}$$

The connection coefficients all vanish:

$$\Gamma^{\sigma}_{\mu\nu} = 0. \tag{C.34}$$

The volumes are given by

$$\Delta V = (X_+ - X_-)(Y_+ - Y_-)(Z_+ - Z_-), \tag{C.35}$$

the areas by

$$\Delta A_{i_-,j,k} = (Y_+ - Y_-)(Z_+ - Z_-), \tag{C.36a}$$

$$\Delta A_{i,j_-,k} = (X_+ - X_-)(Z_+ - Z_-), \tag{C.36b}$$

$$\Delta A_{i,j,k_-} = (X_+ - X_-)(Y_+ - Y_-), \tag{C.36c}$$

the lengths by

$$\Delta L_{i,j_-,k_-} = X_+ - X_i, \tag{C.37a}$$

$$\Delta L_{i_-,j,k_-} = Y_+ - Y_i, \tag{C.37b}$$

$$\Delta L_{i_-,j_-,k} = Z_+ - Z_i, \tag{C.37c}$$

and the widths by

$$\Delta W_X = \frac{\beta}{\alpha}(X_+ - X_-), \tag{C.38a}$$

$$\Delta W_Y = Y_+ - Y_-, \tag{C.38b}$$

$$\Delta W_Z = Z_+ - Z_-. \tag{C.38c}$$

The volume-averaged midpoints are

$$\bar{X} = \frac{1}{2}(X_- + X_+), \tag{C.39a}$$

$$\bar{Y} = \frac{1}{2}(Y_- + Y_+), \tag{C.39b}$$

$$\bar{Z} = \frac{1}{2}(Z_- + Z_+). \tag{C.39c}$$

The transformation matrices are given by

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} \alpha/\beta & 0 & 0 & 0 \\ 2a/\alpha\beta & \beta/\alpha & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i_-,j,k} = \begin{pmatrix} \beta/\alpha & 0 & 0 & 0 \\ -2a/\alpha\beta & \alpha/\beta & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix}, \tag{C.40a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} \alpha/\beta & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \\ 2a/\alpha\beta & \beta/\alpha & 0 & 0 \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j_-,k} = \begin{pmatrix} \beta/\alpha & 0 & 0 & 0 \\ -2a/\alpha\beta & 0 & 0 & \alpha/\beta \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \tag{C.40b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} \alpha/\beta & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 \\ 2a/\alpha\beta & \beta/\alpha & 0 & 0 \\ 0 & 0 & 1 & 0 \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j,k_-} = \begin{pmatrix} \beta/\alpha & 0 & 0 & 0 \\ -2a/\alpha\beta & 0 & \alpha/\beta & 0 \\ 0 & 0 & 0 & 1 \\ 0 & 1 & 0 & 0 \end{pmatrix}. \tag{C.40c}$$

C.2 Schwarzschild Spacetime

Perhaps the simplest nontrivial metric used in practical applications is the Schwarzschild metric. The only parameter is the mass M, and we will use coordinates $x^\mu = (t, r, \theta, \phi)$. For convenience we introduce the lapse function

$$\alpha = \sqrt{1 - \frac{2M}{r}}. \tag{C.41}$$

The metric and its inverse are

$$g_{\mu\nu} = \begin{pmatrix} -\alpha^2 & 0 & 0 & 0 \\ 0 & 1/\alpha^2 & 0 & 0 \\ 0 & 0 & r^2 & 0 \\ 0 & 0 & 0 & r^2 \sin^2\theta \end{pmatrix}, \qquad g^{\mu\nu} = \begin{pmatrix} -1/\alpha^2 & 0 & 0 & 0 \\ 0 & \alpha^2 & 0 & 0 \\ 0 & 0 & 1/r^2 & 0 \\ 0 & 0 & 0 & \csc^2\theta/r^2 \end{pmatrix}, \tag{C.42}$$

and we have

$$\sqrt{-g} = r^2|\sin\theta|. \tag{C.43}$$

The nonzero connection coefficients are

$$\Gamma^0_{01} = \Gamma^0_{10} = \frac{M}{r^2\alpha^2}, \tag{C.44a}$$

$$\Gamma^1_{00} = \frac{M\alpha^2}{r^2}, \tag{C.44b}$$

$$\Gamma^1_{11} = -\frac{M}{r^2\alpha^2}, \tag{C.44c}$$

$$\Gamma^1_{22} = -r\alpha^2, \tag{C.44d}$$

$$\Gamma^1_{33} = -r\alpha^2\sin^2\theta, \tag{C.44e}$$

$$\Gamma^2_{12} = \Gamma^2_{21} = \frac{1}{r}, \tag{C.44f}$$

$$\Gamma^2_{33} = -\sin\theta\cos\theta, \tag{C.44g}$$

$$\Gamma^3_{13} = \Gamma^3_{31} = \frac{1}{r}, \tag{C.44h}$$

$$\Gamma^3_{23} = \Gamma^3_{32} = \cot\theta. \tag{C.44i}$$

So long as a cell does not include the polar axis in its interior, the volumes are given by

$$\Delta V = \frac{1}{3}(r_+^3 - r_-^3)|\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-), \tag{C.45}$$

the areas by

$$\Delta A_{i_-,j,k} = r_-^2|\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-), \tag{C.46a}$$

$$\Delta A_{i,j_-,k} = \frac{1}{3}(r_+^3 - r_-^3)|\sin\theta_-|(\phi_+ - \phi_-), \tag{C.46b}$$

$$\Delta A_{i,j,k_-} = \frac{1}{3}(r_+^3 - r_-^3)|\cos\theta_- - \cos\theta_+|, \tag{C.46c}$$

and the lengths by

$$\Delta L_{i,j_-,k_-} = \frac{1}{3}(r_+^3 - r_-^3)|\sin\theta_-|, \tag{C.47a}$$

$$\Delta L_{i_-,j,k_-} = r_-^2 |\cos\theta_- - \cos\theta_+|, \tag{C.47b}$$

$$\Delta L_{i_-,j_-,k} = r_-^2 |\sin\theta_-|(\phi_+ - \phi_-), \tag{C.47c}$$

just as in spherical coordinates in Minkowski spacetime. The widths are given by

$$\Delta W_r = r_+\alpha_+ - r_-\alpha_- + M\log\left(\frac{r_+(1+\alpha_+) - M}{r_-(1+\alpha_-) - M}\right), \tag{C.48a}$$

$$\Delta W_\theta = \bar{r}(\theta_+ - \theta_-), \tag{C.48b}$$

$$\Delta W_\phi = \bar{r}|\sin\bar{\theta}|(\phi_+ - \phi_-). \tag{C.48c}$$

The volume-averaged midpoints are

$$\bar{r} = \left(\frac{r_-^3 + r_+^3}{2}\right)^{1/3}, \tag{C.49a}$$

$$\bar{\theta} = \cos^{-1}\left(\frac{\cos\theta_- + \cos\theta_+}{2}\right), \tag{C.49b}$$

$$\bar{\phi} = \frac{1}{2}(\phi_- + \phi_+), \tag{C.49c}$$

again just as in spherical coordinates.

The transformation matrices are very simple:

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} \alpha & 0 & 0 & 0 \\ 0 & 1/\alpha & 0 & 0 \\ 0 & 0 & r & 0 \\ 0 & 0 & 0 & r|\sin\theta| \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i_-,j,k} = \begin{pmatrix} 1/\alpha & 0 & 0 & 0 \\ 0 & \alpha & 0 & 0 \\ 0 & 0 & 1/r & 0 \\ 0 & 0 & 0 & |\csc\theta|/r \end{pmatrix}, \tag{C.50a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} \alpha & 0 & 0 & 0 \\ 0 & 0 & r & 0 \\ 0 & 0 & 0 & r|\sin\theta| \\ 0 & 1/\alpha & 0 & 0 \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j_-,k} = \begin{pmatrix} 1/\alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & \alpha \\ 0 & 1/r & 0 & 0 \\ 0 & 0 & |\csc\theta|/r & 0 \end{pmatrix}, \tag{C.50b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} \alpha & 0 & 0 & 0 \\ 0 & 0 & 0 & r|\sin\theta| \\ 0 & 1/\alpha & 0 & 0 \\ 0 & 0 & r & 0 \end{pmatrix}, \quad \left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j,k_-} = \begin{pmatrix} 1/\alpha & 0 & 0 & 0 \\ 0 & 0 & \alpha & 0 \\ 0 & 0 & 0 & 1/r \\ 0 & |\csc\theta|/r & 0 & 0 \end{pmatrix}. \quad \text{(C.50c)}$$

C.3 Kerr Spacetime

Kerr spacetime applies to spinning black holes, and more generally Kerr–Newman spacetime allows black holes to carry charge as well. Because of the ubiquity and importance of this spacetime, we present several different forms for the metric, some of which may be better than others in numerical applications.

In all cases the black hole is parameterized by mass M, spin $a = J/M$, and charge Q. We employ the conventional shorthand symbols

$$\Delta = r^2 - 2Mr + a^2 + Q^2, \quad \text{(C.51)}$$

$$\Sigma = r^2 + a^2\cos^2\theta, \quad \text{(C.52)}$$

where r and θ are the same radial and polar coordinates in all of the following examples. The timelike and azimuthal coordinates, however, may vary between coordinate systems.

In what follows we will always take $Q = 0$.

C.3.1 Boyer–Lindquist Coordinates

In traditional Boyer–Lindquist coordinates $x^\mu = (t, r, \theta, \phi)$ the metric and its inverse are written

$$g_{\mu\nu} = \begin{pmatrix} -(1 - 2Mr/\Sigma) & 0 & 0 & -(2Mar/\Sigma)\sin^2\theta \\ 0 & \Sigma/\Delta & 0 & 0 \\ 0 & 0 & \Sigma & 0 \\ -(2Mar/\Sigma)\sin^2\theta & 0 & 0 & (r^2 + a^2 + (2Ma^2r/\Sigma)\sin^2\theta)\sin^2\theta \end{pmatrix}, \quad \text{(C.53a)}$$

$$g^{\mu\nu} = \begin{pmatrix} -1 - 2Mr(r^2+a^2)/\Delta\Sigma & 0 & 0 & -2Mar/\Delta\Sigma \\ 0 & \Delta/\Sigma & 0 & 0 \\ 0 & 0 & 1/\Sigma & 0 \\ -2Mar/\Delta\Sigma & 0 & 0 & (\Sigma - 2Mr)/\Delta\Sigma\sin^2\theta \end{pmatrix} \tag{C.53b}$$

(Misner et al., 1973, 33.2). We have the factor

$$\sqrt{-g} = \Sigma|\sin\theta|. \tag{C.54}$$

There are a large number of nonzero connection coefficients; we have the following formulas for them:

$$\Gamma^0_{01} = \Gamma^0_{10} = \frac{1}{\Delta\rho^4} M(r^2+a^2)(r^2 - a^2\cos^2\theta), \tag{C.55a}$$

$$\Gamma^0_{02} = \Gamma^0_{20} = -\frac{Ma^2 r}{\rho^4}\sin(2\theta), \tag{C.55b}$$

$$\Gamma^0_{13} = \Gamma^0_{31} = -\frac{Ma}{\Delta\rho^4}\Big(3r^4 + a^2r^2(1+\cos^2\theta) - a^4\cos^2\theta\Big)\sin^2\theta, \tag{C.55c}$$

$$\Gamma^0_{23} = \Gamma^0_{32} = \frac{Ma^3 r}{\rho^4}\sin(2\theta)\sin^2\theta, \tag{C.55d}$$

$$\Gamma^1_{00} = \frac{M\Delta}{\rho^6}(r^2 - a^2\cos^2\theta), \tag{C.55e}$$

$$\Gamma^1_{03} = \Gamma^1_{30} = -\frac{Ma\Delta}{\rho^6}(r^2 - a^2\cos^2\theta)\sin^2\theta, \tag{C.55f}$$

$$\Gamma^1_{11} = \frac{r}{\Sigma} - \frac{r-M}{\Delta}, \tag{C.55g}$$

$$\Gamma^1_{12} = \Gamma^1_{21} = -\frac{a^2}{2\Sigma}\sin(2\theta), \tag{C.55h}$$

$$\Gamma^1_{22} = -\frac{\Delta r}{\Sigma}, \tag{C.55i}$$

$$\Gamma^1_{33} = \frac{\Delta}{\Sigma}\left(\frac{Ma^2}{\rho^4}(r^2 - a^2\cos^2\theta)\sin^2\theta - r\right)\sin^2(\theta), \tag{C.55j}$$

$$\Gamma^2_{00} = -\frac{Ma^2 r}{\rho^6}\sin(2\theta), \tag{C.55k}$$

$$\Gamma^2_{03} = \Gamma^2_{30} = \frac{Mar}{\rho^6}(r^2+a^2)\sin(2\theta), \tag{C.55l}$$

$$\Gamma^2_{11} = \frac{a^2}{2\Delta\Sigma}\sin(2\theta), \tag{C.55m}$$

$$\Gamma^2_{12} = \Gamma^2_{21} = \frac{r}{\Sigma}, \tag{C.55n}$$

$$\Gamma^2_{22} = -\frac{a^2}{2\Sigma}\sin(2\theta), \tag{C.55o}$$

$$\Gamma^2_{33} = -\frac{1}{2\Sigma}\left(\Delta + \frac{2Mr}{\rho^4}(r^2+a^2)^2\right)\sin(2\theta), \tag{C.55p}$$

$$\Gamma^3_{01} = \Gamma^3_{10} = \frac{Ma}{\Delta\rho^4}(r^2 - a^2\cos^2\theta), \tag{C.55q}$$

$$\Gamma^3_{02} = \Gamma^3_{20} = -\frac{2Mar}{\rho^4}\cot\theta, \tag{C.55r}$$

$$\Gamma^3_{13} = \Gamma^3_{31} = \frac{1}{\Delta\Sigma}\left(r(\Sigma - 2Mr) - \frac{Ma^2}{\Sigma}(r^2 - a^2\cos^2\theta)\sin^2\theta\right), \tag{C.55s}$$

$$\Gamma^3_{23} = \Gamma^3_{32} = \frac{Ma^2r}{\rho^4}\sin(2\theta) + \cot\theta. \tag{C.55t}$$

The volumes are given by

$$\begin{aligned}\Delta V = \frac{1}{3}(r_+^3 - r_-^3)&|\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-)\\ &\times\Big(r_-^2 + r_-r_+ + r_+^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big),\end{aligned} \tag{C.56}$$

the areas by

$$\Delta A_{i_-,j,k} = \frac{1}{3}(\phi_+ - \phi_-)|\cos\theta_- - \cos\theta_+|\Big(3r_-^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big),$$

(C.57a)

$$\Delta A_{i,j_-,k} = \frac{1}{3}(r_+ - r_-)|\sin\theta_-|(\phi_+ - \phi_-)(r_-^2 + r_-r_+ + r_+^2 + 3a^2\cos^2\theta_-), \tag{C.57b}$$

$$\begin{aligned}\Delta A_{i,j,k_-} = \frac{1}{3}(r_+ - r_-)&|\cos\theta_- - \cos\theta_+|\\ &\times\Big(r_-^2 + r_-r_+ + r_+^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big),\end{aligned} \tag{C.57c}$$

and the lengths by

$$\Delta L_{i,j_-,k_-} = \frac{1}{3}(r_+ - r_-)|\sin\theta_-|(r_-^2 + r_-r_+ + r_+^2 + 3a^2\cos^2\theta_-), \tag{C.58a}$$

$$\Delta L_{i_-,j,k_-} = r_-^2|\cos\theta_- - \cos\theta_+| + \frac{1}{3}a^2(\cos^3\theta_- - \cos^3\theta_+), \tag{C.58b}$$

$$\Delta L_{i_-,j_-,k} = \Sigma_-|\sin\theta_-|(\phi_+ - \phi_-). \tag{C.58c}$$

Note that while these cannot be factored into the schematic form $f(r)g(\theta)h(\phi)$, they can still be separated into a sum of such terms, and therefore only require 1D storage. As was the case with sinusoidal coordinates, we find it convenient to bound some of the cell widths. By replacing the integrands $\sqrt{g_{ii}}$ with more tractable forms, we find

$$\Delta W_r \geq \int_{r_-}^{r_+} \frac{r}{\sqrt{\Delta}}\,\mathrm{d}r = \sqrt{\Delta_+} - \sqrt{\Delta_-} + M\log\left(\frac{r_+ + \sqrt{\Delta_+} - M}{r_- + \sqrt{\Delta_-} - M}\right), \tag{C.59a}$$

$$\Delta W_\theta \geq \int_{\theta_-}^{\theta_+} r\,\mathrm{d}\theta = \bar{r}(\theta_+ - \theta_-), \tag{C.59b}$$

$$\Delta W_\phi = |\sin\bar{\theta}|(\phi_+ - \phi_-)\sqrt{\bar{r}^2 + a^2 + \frac{2Ma^2\bar{r}}{\bar{\Sigma}}\sin^2\bar{\theta}}. \tag{C.59c}$$

The midpoints somewhat more complicated. While the azimuthal coordinate is trivial, the other two obey cubic equations:

$$\bar{r}^3 + \beta\bar{r} - \frac{1}{2}\Big(r_-^3 + r_+^3 + \beta(r_- + r_+)\Big) = 0, \tag{C.60a}$$

$$\cos^3(\theta_j) + \frac{4\alpha}{3a^2}\cos\bar{\theta} - \frac{2}{3}(\cos^3\theta_- + \cos^3\theta_+) - \frac{2\alpha}{3a^2}(\cos\theta_- + \cos\theta_+) = 0, \tag{C.60b}$$

$$\bar{\phi} = \frac{1}{2}(\phi_- + \phi_+), \tag{C.60c}$$

with

$$\alpha = r_-^2 + r_-r_+ + r_+^2, \tag{C.61}$$

$$\beta = a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+). \tag{C.62}$$

While solving these equations numerically is not difficult, the problem is in the coupling between r and θ; the exact midpoint formulas are not separable. For simplicity of code design, then, the averages

$$\bar{r} \approx \frac{1}{2}(r_- + r_+), \tag{C.63a}$$

$$\bar{\theta} \approx \frac{1}{2}(\theta_- + \theta_+), \tag{C.63b}$$

$$\bar{\phi} = \frac{1}{2}(\phi_- + \phi_+) \tag{C.63c}$$

are used instead. These are still second-order accurate in space.

The transformations (2.51) are

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} M^{\hat{t}}{}_0 & 0 & 0 & 0 \\ 0 & \sqrt{\Sigma/\Delta} & 0 & 0 \\ 0 & 0 & \sqrt{\Sigma} & 0 \\ M^{\hat{z}}{}_0 & 0 & 0 & \sqrt{g_{33}} \end{pmatrix}, \tag{C.64a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} M^{\hat{t}}{}_0 & 0 & 0 & 0 \\ 0 & 0 & \sqrt{\Sigma} & 0 \\ M^{\hat{y}}{}_0 & 0 & 0 & \sqrt{g_{33}} \\ 0 & \sqrt{\Sigma/\Delta} & 0 & 0 \end{pmatrix}, \tag{C.64b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} M^{\hat{t}}{}_0 & 0 & 0 & 0 \\ M^{\hat{x}}{}_0 & 0 & 0 & \sin\theta/M^{\hat{t}}{}_0 \\ 0 & \sqrt{\Sigma/\Delta} & 0 & 0 \\ 0 & 0 & \sqrt{\Sigma} & 0 \end{pmatrix}, \tag{C.64c}$$

where we have

$$M^{\hat{t}}{}_0 = \left(\frac{\Delta\Sigma}{\Delta\Sigma + 2Mr(r^2+a^2)}\right)^{1/2}, \tag{C.65a}$$

$$M^{\hat{z}}{}_0 = -\frac{2Mar\sqrt{g_{33}}}{r^4 + (r^2 + 2Mr + \Delta\cos^2\theta)a^2}, \tag{C.65b}$$

$$M^{\hat{y}}{}_0 = -\frac{2Mar\sin\theta}{r^4 + (r^2 + 2Mr + \Delta\cos^2\theta)a^2}\left(r^2 + a^2 + \frac{2Ma^2r}{\Sigma}\sin^2\theta\right)^{1/2}, \tag{C.65c}$$

$$M^{\hat{x}}{}_0 = -\frac{2Mar\sin\theta}{\sqrt{\Delta\Sigma^2 + 2Mr\Sigma(r^2+a^2)}}. \tag{C.65d}$$

The inverse transformations (2.49) are

$$\left(M^{\mu}{}_{\hat{\nu}}\right)_{i_-,j,k} = \begin{pmatrix} M^{0}{}_{\hat{t}} & 0 & 0 & 0 \\ 0 & \sqrt{\Delta/\Sigma} & 0 & 0 \\ 0 & 0 & 1/\sqrt{\Sigma} & 0 \\ M^{3}{}_{\hat{t}} & 0 & 0 & M^{3}{}_{\hat{z}} \end{pmatrix}, \tag{C.66a}$$

$$\left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j_-,k} = \begin{pmatrix} M^{0}{}_{\hat{t}} & 0 & 0 & 0 \\ 0 & 0 & 0 & \sqrt{\Delta/\Sigma} \\ 0 & 1/\sqrt{\Sigma} & 0 & 0 \\ M^{3}{}_{\hat{t}} & 0 & M^{3}{}_{\hat{y}} & 0 \end{pmatrix}, \tag{C.66b}$$

$$\left(M^{\mu}{}_{\hat{\nu}}\right)_{i,j,k_-} = \begin{pmatrix} M^{0}{}_{\hat{t}} & 0 & 0 & 0 \\ 0 & 0 & \sqrt{\Delta/\Sigma} & 0 \\ 0 & 0 & 0 & 1/\sqrt{\Sigma} \\ M^{3}{}_{\hat{t}} & M^{3}{}_{\hat{x}} & 0 & 0 \end{pmatrix}, \tag{C.66c}$$

where we have

$$M^{0}{}_{\hat{t}} = \left(\frac{\Delta\Sigma + 2Mr(r^2+a^2)}{\Delta\Sigma}\right)^{1/2}, \tag{C.67a}$$

$$M^{3}{}_{\hat{t}} = \frac{2Mar}{\Delta\Sigma}\left(\frac{\Delta\Sigma}{\Delta\Sigma + 2Mr(r^2+a^2)}\right)^{1/2}, \tag{C.67b}$$

$$M^{3}{}_{\hat{z}} = M^{3}{}_{\hat{y}} = M^{3}{}_{\hat{x}} = \csc\theta\left(\frac{\Sigma}{r^2+a^2+2Ma^2r\sin^2\theta}\right)^{1/2}. \tag{C.67c}$$

C.3.2 Kerr–Schild Coordinates

From Boyer–Lindquist coordinates, define ingoing Kerr coordinates $(\tilde{V}, r, \theta, \tilde{\phi})$ according to

$$\mathrm{d}\tilde{V} = \mathrm{d}t + \frac{r^2+a^2}{\Delta}\mathrm{d}r, \tag{C.68a}$$

$$\mathrm{d}\tilde{\phi} = \mathrm{d}\phi + \frac{a}{\Delta}\mathrm{d}r \tag{C.68b}$$

(Misner et al., 1973, Box 33.2.4). We then define the so-called Kerr–Schild coordinates $(\tilde{t}, r, \theta, \tilde{\phi})$ with

$$\mathrm{d}\tilde{t} = \mathrm{d}\tilde{V} - \mathrm{d}r \tag{C.69}$$

(cf. Gammie et al., 2003, (35)). Note that there are multiple coordinate systems in the literature named "Kerr–Schild."

Hereafter in this section we will omit tildes distinguishing Kerr–Schild from Boyer–Lindquist quantities. For convenience we also introduce the shorthand

$$\Xi = r^2 - a^2 \cos^2\theta. \tag{C.70}$$

The metric and its inverse are

$$g_{\mu\nu} = \begin{pmatrix} -(1 - 2Mr/\Sigma) & 2Mr/\Sigma & 0 & -(2Mar/\Sigma)\sin^2\theta \\ 2Mr/\Sigma & 1 + 2Mr/\Sigma & 0 & -(1 + 2Mr/\Sigma)a\sin^2\theta \\ 0 & 0 & \Sigma & 0 \\ -(2Mar/\Sigma)\sin^2\theta & -(1 + 2Mr/\Sigma)a\sin^2\theta & 0 & g_{33} \end{pmatrix}, \tag{C.71a}$$

$$g^{\mu\nu} = \begin{pmatrix} -(1 + 2Mr/\Sigma) & 2Mr/\Sigma & 0 & 0 \\ 2Mr/\Sigma & \Delta/\Sigma & 0 & a/\Sigma \\ 0 & 0 & 1/\Sigma & 0 \\ 0 & a/\Sigma & 0 & 1/\Sigma\sin^2\theta \end{pmatrix}, \tag{C.71b}$$

with

$$g_{33} = \left(r^2 + a^2 + \left(\frac{2Ma^2 r}{\Sigma}\right)\sin^2\theta\right)\sin^2\theta, \tag{C.72}$$

and we still have

$$\sqrt{-g} = \Sigma|\sin\theta|. \tag{C.73}$$

None of the connection coefficients vanish. The full set of them is

$$\Gamma^0_{00} = \frac{2M^2\Xi r}{\Sigma^3}, \tag{C.74a}$$

$$\Gamma^0_{01} = \Gamma^0_{10} = \frac{M\Xi}{\Sigma^2}\left(1 + \frac{2Mr}{\Sigma}\right), \tag{C.74b}$$

$$\Gamma^0_{02} = \Gamma^0_{20} = -\frac{2Ma^2 r}{\Sigma^2}\sin\theta\cos\theta, \tag{C.74c}$$

$$\Gamma^0_{03} = \Gamma^0_{30} = -\frac{2M^2 ar\Xi}{\Sigma^3}\sin^2\theta, \tag{C.74d}$$

$$\Gamma^0_{11} = \frac{2M\Xi}{\Sigma^2}\left(1 + \frac{Mr}{\Sigma}\right), \tag{C.74e}$$

$$\Gamma^0_{12} = \Gamma^0_{21} = -\frac{2Ma^2 r}{\Sigma^2}\sin\theta\cos\theta, \tag{C.74f}$$

$$\Gamma^0_{13} = \Gamma^0_{31} = -\frac{Ma\Xi}{\Sigma^2}\left(1 + \frac{2Mr}{\Sigma}\right)\sin^2\theta, \tag{C.74g}$$

$$\Gamma^0_{22} = -\frac{2Mr^2}{\Sigma}, \tag{C.74h}$$

$$\Gamma^0_{23} = \Gamma^0_{32} = \frac{2Ma^3 r}{\Sigma^2}\sin^3\theta\cos\theta, \tag{C.74i}$$

$$\Gamma^0_{33} = -\frac{2Mr}{\Sigma}\left(r - \frac{Ma^2\Xi}{\Sigma^2}\sin^2\theta\right)\sin^2\theta; \tag{C.74j}$$

$$\Gamma^1_{00} = \frac{M\Delta\Xi}{\Sigma^3}, \tag{C.75a}$$

$$\Gamma^1_{01} = \Gamma^1_{10} = \frac{M\Xi(\Delta - \Sigma)}{\Sigma^3}, \tag{C.75b}$$

$$\Gamma^1_{02} = \Gamma^1_{20} = 0, \tag{C.75c}$$

$$\Gamma^1_{03} = \Gamma^1_{30} = -\frac{Ma\Delta\Xi}{\Sigma^3}\sin^2\theta, \tag{C.75d}$$

$$\Gamma^1_{11} = \frac{M\Xi(\Delta - 2\Sigma)}{\Sigma^3}, \tag{C.75e}$$

$$\Gamma^1_{12} = \Gamma^1_{21} = -\frac{a^2}{\Sigma}\sin\theta\cos\theta, \tag{C.75f}$$

$$\Gamma^1_{13} = \Gamma^1_{31} = \frac{a}{\Sigma^3}\Bigg(r^3(r^2 + 2M^2) + a^2\Big(ra^2\cos^4\theta + \Big(2r^3 + M(\Delta - \Sigma)\Big)\cos^2\theta - Mr^2\sin^2\theta\Big)\Bigg), \tag{C.75g}$$

$$\Gamma^1_{22} = -\frac{r\Delta}{\Sigma}, \tag{C.75h}$$

$$\Gamma^1_{23} = \Gamma^1_{32} = 0, \tag{C.75i}$$

$$\Gamma^1_{33} = -\frac{\Delta}{\Sigma}\left(r - \frac{Ma^2\Xi}{\Sigma^2}\sin^2\theta\right)\sin^2\theta; \tag{C.75j}$$

$$\Gamma^2_{00} = -\frac{2Ma^2r}{\Sigma^3}\sin\theta\cos\theta, \tag{C.76a}$$

$$\Gamma^2_{01} = \Gamma^2_{10} = -\frac{2Ma^2r}{\Sigma^3}\sin\theta\cos\theta, \tag{C.76b}$$

$$\Gamma^2_{02} = \Gamma^2_{20} = 0, \tag{C.76c}$$

$$\Gamma^2_{03} = \Gamma^2_{30} = \frac{2Mar}{\Sigma^3}(r^2+a^2)\sin\theta\cos\theta, \tag{C.76d}$$

$$\Gamma^2_{11} = -\frac{2Ma^2r}{\Sigma^3}\sin\theta\cos\theta, \tag{C.76e}$$

$$\Gamma^2_{12} = \Gamma^2_{21} = \frac{r}{\Sigma}, \tag{C.76f}$$

$$\Gamma^2_{13} = \Gamma^2_{31} = \frac{a}{\Sigma}\left(1 + \frac{2Mr}{\Sigma^2}(r^2+a^2)\right)\sin\theta\cos\theta, \tag{C.76g}$$

$$\Gamma^2_{22} = -\frac{a^2}{\Sigma}\sin\theta\cos\theta, \tag{C.76h}$$

$$\Gamma^2_{23} = \Gamma^2_{32} = 0, \tag{C.76i}$$

$$\Gamma^2_{33} = -\frac{1}{\Sigma}\left(\Delta + \frac{2Mr}{\Sigma^2}(r^2+a^2)^2\right)\sin\theta\cos\theta; \tag{C.76j}$$

and

$$\Gamma^3_{00} = \frac{Ma\Xi}{\Sigma^3}, \tag{C.77a}$$

$$\Gamma^3_{01} = \Gamma^3_{10} = \frac{Ma\Xi}{\Sigma^3}, \tag{C.77b}$$

$$\Gamma^3_{02} = \Gamma^3_{20} = -\frac{2Mar}{\Sigma^2}\cot\theta, \tag{C.77c}$$

$$\Gamma^3_{03} = \Gamma^3_{30} = -\frac{Ma^2\Xi}{\Sigma^3}\sin^2\theta, \tag{C.77d}$$

$$\Gamma^3_{11} = \frac{Ma\Xi}{\Sigma^3}, \tag{C.77e}$$

$$\Gamma^3_{12} = \Gamma^3_{21} = -\frac{a}{\Sigma^2}(2Mr+\Sigma)\cot\theta, \tag{C.77f}$$

$$\Gamma^3_{13} = \Gamma^3_{31} = \frac{1}{\Sigma^3}(r\Sigma^2 - Ma^2\Xi\sin^2\theta), \tag{C.77g}$$

$$\Gamma^3_{22} = -\frac{ar}{\Sigma}, \tag{C.77h}$$

$$\Gamma^3_{23} = \Gamma^3_{32} = \frac{2Ma^2 r}{\Sigma^2}\sin\theta\cos\theta + \cot\theta, \tag{C.77i}$$

$$\Gamma^3_{33} = \frac{a}{\Sigma^3}(Ma^2\Xi\sin^2\theta - r\Sigma^2)\sin^2\theta. \tag{C.77j}$$

The volumes, areas, and lengths have the same formulas as in Boyer–Lindquist coordinates. For completeness we record them here. The volumes are given by

$$\begin{aligned}\Delta V = {} & \frac{1}{3}(r_+ - r_-)|\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-) \\ & \times \Big(r_-^2 + r_- r_+ + r_+^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big),\end{aligned} \tag{C.78}$$

the areas by

$$\Delta A_{i_-,j,k} = \frac{1}{3}|\cos\theta_- - \cos\theta_+|(\phi_+ - \phi_-)\Big(3r_-^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big), \tag{C.79a}$$

$$\Delta A_{i,j_-,k} = \frac{1}{3}(r_+ - r_-)|\sin\theta_-|(\phi_+ - \phi_-)(r_-^2 + r_- r_+ + r_+^2 + 3a^2\cos^2\theta_-), \tag{C.79b}$$

$$\begin{aligned}\Delta A_{i,j,k_-} = {} & \frac{1}{3}(r_+ - r_-)|\cos\theta_- - \cos\theta_+| \\ & \times \Big(r_-^2 + r_- r_+ + r_+^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big),\end{aligned} \tag{C.79c}$$

and the lengths by

$$\Delta L_{i,j_-,k_-} = \frac{1}{3}(r_+ - r_-)|\sin\theta_-|(r_-^2 + r_- r_+ + r_+^2 + 3a^2\cos^2\theta_-), \tag{C.80a}$$

$$\Delta L_{i_-,j,k_-} = \frac{1}{3}|\cos\theta_- - \cos\theta_+|\Big(3r_-^2 + a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+)\Big), \tag{C.80b}$$

$$\Delta L_{i_-,j_-,k} = (\phi_+ - \phi_-)|\sin\theta_-|(r_-^2 + a^2\cos^2\theta_-). \tag{C.80c}$$

Again we bound two of the cell widths:

$$\Delta W_r \geq \int_{r_-}^{r_+} \frac{r+M}{\sqrt{r^2+M^2}}\,\mathrm{d}r = \sqrt{r_+^2+M^2} - \sqrt{r_-^2+M^2} + M\log\left(\frac{\sqrt{r_+^2+M^2}+r_+}{\sqrt{r_-^2+M^2}+r_-}\right), \tag{C.81a}$$

$$\Delta W_\theta \geq \int_{\theta_-}^{\theta_+} r \, \mathrm{d}\theta = \bar{r}(\theta_+ - \theta_-), \tag{C.81b}$$

$$\Delta W_\phi = |\sin\bar{\theta}|(\phi_+ - \phi_-)\sqrt{\bar{r}^2 + a^2 + \frac{2Ma^2\bar{r}}{\bar{r}^2 + a^2\cos^2\bar{\theta}}\sin^2\bar{\theta}}. \tag{C.81c}$$

The exact midpoints obey the same relations as in Boyer–Lindquist coordinates:

$$\bar{r}^3 + \beta\bar{r} - \frac{1}{2}\Big(r_-^3 + r_+^3 + \beta(r_- + r_+)\Big) = 0, \tag{C.82a}$$

$$\cos^3(\theta_j) + \frac{4\alpha}{3a^2}\cos\bar{\theta} - \frac{2}{3}(\cos^3\theta_- + \cos^3\theta_+) - \frac{2\alpha}{3a^2}(\cos\theta_- + \cos\theta_+) = 0, \tag{C.82b}$$

$$\bar{\phi} = \frac{1}{2}(\phi_- + \phi_+), \tag{C.82c}$$

with

$$\alpha = r_-^2 + r_-r_+ + r_+^2, \tag{C.83}$$

$$\beta = a^2(\cos^2\theta_- + \cos\theta_-\cos\theta_+ + \cos^2\theta_+). \tag{C.84}$$

In practice we therefore use the formulas

$$\bar{r} \approx \frac{1}{2}(r_- + r_+), \tag{C.85a}$$

$$\bar{\theta} \approx \frac{1}{2}(\theta_- + \theta_+), \tag{C.85b}$$

$$\bar{\phi} = \frac{1}{2}(\phi_- + \phi_+). \tag{C.85c}$$

The transformations (2.51) are given by

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i_-,j,k} = \begin{pmatrix} 1/\sqrt{g_{11}} & 0 & 0 & 0 \\ M^{\hat{x}}{}_t & M^{\hat{x}}{}_r & 0 & 0 \\ 0 & 0 & \sqrt{\Sigma} & 0 \\ M^{\hat{z}}{}_t & M^{\hat{z}}{}_r & 0 & M^{\hat{z}}{}_\phi \end{pmatrix}, \tag{C.86a}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j_-,k} = \begin{pmatrix} 1/\sqrt{g_{11}} & 0 & 0 & 0 \\ 0 & 0 & \sqrt{\Sigma} & 0 \\ 0 & 0 & 0 & \sqrt{\Sigma}|\sin\theta| \\ 2Mr/\Sigma\sqrt{g_{11}} & \sqrt{g_{11}} & 0 & -a\sin^2\theta\sqrt{g_{11}} \end{pmatrix}, \tag{C.86b}$$

$$\left(M^{\hat{\mu}}{}_{\nu}\right)_{i,j,k_-} = \begin{pmatrix} 1/\sqrt{g_{11}} & 0 & 0 & 0 \\ 0 & 0 & 0 & \sqrt{\Sigma}|\sin\theta| \\ 2Mr/\Sigma\sqrt{g_{11}} & \sqrt{g_{11}} & 0 & -a\sin^2\theta\sqrt{g_{11}} \\ 0 & 0 & \sqrt{\Sigma} & 0 \end{pmatrix}, \tag{C.86c}$$

where we have

$$M^{\hat{x}}{}_{t} = 2Mr\sqrt{\Sigma}\Big((\Sigma+2Mr)(r^4+a^2r^2+2Ma^2r+\Delta a^2\cos^2\theta)\Big)^{-1/2}, \tag{C.87a}$$

$$M^{\hat{x}}{}_{r} = \sqrt{\frac{\Sigma+2Mr}{r^2+a^2+(2Ma^2r/\Sigma)\sin^2\theta}}, \tag{C.87b}$$

$$M^{\hat{z}}{}_{t} = -\frac{2Mar|\sin\theta|\sqrt{r^2+a^2+(2Ma^2r/\Sigma)\sin^2\theta}}{r^4+a^2r^2+2Ma^2r+\Delta a^2\cos^2\theta}, \tag{C.87c}$$

$$M^{\hat{z}}{}_{r} = -\frac{a(\Sigma+2Mr)|\sin\theta|\sqrt{r^2+a^2+(2Ma^2r/\Sigma)\sin^2\theta}}{r^4+a^2r^2+2Ma^2r+\Delta a^2\cos^2\theta}, \tag{C.87d}$$

$$M^{\hat{z}}{}_{\phi} = |\sin\theta|\sqrt{r^2+a^2+(2Ma^2r/\Sigma)\sin^2\theta}. \tag{C.87e}$$

The inverse transformations (2.49) are given by

$$(M^{\mu}{}_{\hat{\nu}})_{i_-,j,k} = \begin{pmatrix} 1/M^{\hat{t}}{}_{t} & 0 & 0 & 0 \\ -M^{\hat{x}}{}_{t}/M^{\hat{t}}{}_{t}M^{\hat{x}}{}_{r} & 1/M^{\hat{x}}{}_{r} & 0 & 0 \\ 0 & 0 & 1/M^{\hat{y}}{}_{\theta} & 0 \\ 0 & -M^{\hat{z}}{}_{r}/M^{\hat{x}}{}_{r}M^{\hat{z}}{}_{\phi} & 0 & 1/M^{\hat{z}}{}_{\phi} \end{pmatrix}, \tag{C.88a}$$

$$(M^{\mu}{}_{\hat{\nu}})_{i,j_-,k} = \begin{pmatrix} 1/M^{\hat{t}}{}_{t} & 0 & 0 & 0 \\ -M^{\hat{z}}{}_{t}/M^{\hat{t}}{}_{t}M^{\hat{z}}{}_{r} & 0 & -M^{\hat{z}}{}_{\phi}/M^{\hat{z}}{}_{r}M^{\hat{y}}{}_{\phi} & 1/M^{\hat{z}}{}_{r} \\ 0 & 1/M^{\hat{x}}{}_{\theta} & 0 & 0 \\ 0 & 0 & 1/M^{\hat{y}}{}_{\phi} & 0 \end{pmatrix}, \tag{C.88b}$$

$$(M^{\mu}{}_{\hat{\nu}})_{i,j,k_-} = \begin{pmatrix} 1/M^{\hat{t}}{}_{t} & 0 & 0 & 0 \\ -M^{\hat{y}}{}_{t}/M^{\hat{t}}{}_{t}M^{\hat{y}}{}_{r} & -M^{\hat{y}}{}_{\phi}/M^{\hat{y}}{}_{r}M^{\hat{x}}{}_{\phi} & 1/M^{\hat{y}}{}_{r} & 0 \\ 0 & 0 & 0 & 1/M^{\hat{z}}{}_{\theta} \\ 0 & 1/M^{\hat{x}}{}_{\phi} & 0 & 0 \end{pmatrix}, \tag{C.88c}$$

where the i, j, k indices in the matrix entries are understood to match those on the left-hand side of their respective equation.

References

Marek A. Abramowicz and P. Chris Fragile. Foundations of Black Hole Accretion Disk Theory. *Living Reviews in Relativity*, 16(1), January 2013. doi: 10.12942/lrr-2013-1.

Matthew Anderson, Eric W Hirschmann, Steven L Liebling, and David Neilsen. Relativistic MHD with adaptive mesh refinement. *Classical and Quantum Gravity*, 23:6503–6524, November 2006. doi: 10.1088/0264-9381/23/22/025.

Luis Antón, Olindo Zanotti, Juan A. Miralles, José M. Martí, José M. Ibáñez, José A. Font, and José A. Pons. Numerical 3+1 general relativistic magnetohydrodynamics: a local characteristic approach. *The Astrophysical Journal*, 637:296–312, January 2006. doi: 10.1086/498238.

Luis Antón, Juan A. Miralles, José M Martí, José M Ibáñez, Miguel A. Aloy, and Petar Mimica. Relativistic Magnetohydrodynamics: Renormalized Eigenvectors and Full Wave Decomposition Riemann Solver. *The Astrophysical Journal Supplement Series*, 188:1–31, May 2010. doi: 10.1088/0067-0049/188/1/1.

Steven A. Balbus and John F. Hawley. A Powerful Local Shear Instability in Weakly Magnetized Disks. I. Linear Analysis. *The Astrophysical Journal*, 376:214–222, July 1991. doi: 10.1086/170270.

Steven A. Balbus and John F. Hawley. Instability, turbulence, and enhanced transport in accretion disks. *Reviews of Modern Physics*, 70(1):1–53, January 1998. doi: 10.1103/RevModPhys.70.1.

Steven A. Balbus, John F. Hawley, and James M. Stone. Nonlinear Stability, Hydrodynamical Turbulence, and Transport in Disks. *The Astrophysical Journal*, 467:76–86, August 1996. doi: 10.1086/177585.

Francesc Banyuls, José A. Font, José Ma. Ibáñez, José Ma. Martí, and Juan A. Miralles. Numerical $\{3+1\}$ General Relativistic Hydrodynamics: A Local Characteristic Approach. *The Astrophysical Journal*, 476: 221–231, February 1997. doi: 10.1086/303604.

Kris Beckwith and James M. Stone. A Second-Order Godunov Method for Multi-Dimensional Relativistic Magnetohydrodynamics. *The Astrophysical Journal Supplement Series*, 193:6, March 2011. doi: 10.1088/0067-0049/193/1/6.

D. P. Bennett, A. C. Becker, J. L. Quinn, A. B. Tomaney, C. Alcock, R. A. Allsman, D. R. Alves, T. S. Axelrod, J. J. Calitz, K. H. Cook, A. J. Drake, P. C. Fragile, K. C. Freeman, M. Geha, K. Griest, B. R. Johnson, S. C. Keller, C. Laws, M. J. Lehner, S. L. Marshall, D.Minniti, C. A. Nelson, B. A. Peterson, P. Popowski, M. R. Pratt, P. J. Quinn, S. H. Rhie, C. W. Stubbs, W. Sutherland, T. Vandehei, and D. Welch. Gravitational Microlensing Events Due to Stellar-Mass Black Holes. *The Astrophysical Journal*, 579:639–659, November 2002. doi: 10.1086/342225.

G. S. Bisnovatyi-Kogan and A. A. Ruzmaikin. The Accretion of Matter by a Collapsing Star in the Presence of a Magnetic Field. *Astrophysics and Space Science*, 28:45–59, May 1974. doi: 10.1007/BF00642237.

R. D. Blandford and R. L. Znajek. Electromagnetic extraction of energy from Kerr black holes. *Monthly Notices of the Royal Astronomical Society*, 179:433–456, May 1977. doi: 10.1093/mnras/179.3.433.

J. U. Brackbill and D. C. Barnes. The Effect of Nonzer $\nabla \cdot \mathbf{B}$ on the Numerical Solution of the Magnetohydrodynamic Equations. *Journal of Computational Physics*, 35:426–430, May 1980. doi: 10.1016/0021-9991(80)90079-0.

Sandip K. Chakrabarti. The Natural Angular Momentum Distribution in the Study of Thick Disks Around Black Holes. *The Astrophysical Journal*, 288:1–6, January 1985. doi: 10.1086/162755.

Mani Chandra, Charles F. Gammie, Francois Foucart, and Eliot Quataert. An Extended Magnetohydrodynamics Model for Relativistic Weakly Collisional Plasmas. *The Astrophysical Journal*, 810:1–16, September 2015. doi: 10.1088/0004-637X/810/2/162.

A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, and M. Wesenberg. Hyperbolic Divergence Cleaning for the MHD Equations. *Journal of Computational Physics*, 175:645–673, January 2002. doi: 10.1006/jcph.2001.6961.

L. Del Zanna, O. Zanotti, N. Bucciantini, and P. Londrillo. ECHO: a Eulerian conservative high-order scheme for general relativistic magnetohydrodynamics and magnetodynamics. *Astronomy and Astrophysics*, 473: 11–30, October 2007. doi: 10.1051/0004-6361:20077093.

Joshua C. Dolence, Charles F. Gammie, Monika Mościbrodzka, and Po Kin Leung. grmonty: A Monte Carlo Code for Relativistic Radiative Transport. *The Astrophysical Journal Supplement Series*, 184:387–397, October 2009. doi: 10.1088/0067-0049/184/2/387.

Bernd Einfeldt. On Godunov-Type Methods for Gas Dynamics. *SIAM Journal on Numerical Analysis*, 25 (2):294–318, April 1988. doi: 10.1137/0725021.

Zachariah B. Etienne, Vasileios Paschalidis, Yuk Tung Liu, and Stuart L. Shapiro. Relativistic magnetohydrodynamics in dynamical spacetimes: Improved electromagnetic gauge condition for adaptive mesh refinement grids. *Physical Review D*, 85:024013, January 2012. doi: 10.1103/PhysRevD.85.024013.

Zachariah B Etienne, Vasileios Paschalidis, Roland Haas, Philipp Mösta, and Stuart L Shapiro. Illinois-GRMHD: an open-source, user-friendly GRMHD code for dynamical spacetimes. *Classical and Quantum Gravity*, 32:175009, September 2015. doi: 10.1088/0264-9381/32/17/175009.

Charles R. Evans and John F. Hawley. Simulation of Magnetohydrodynamic Flows: A Constrained Transport Method. *The Astrophysical Journal*, 322:659–677, September 1988. doi: 10.1086/166684.

Joshua A. Faber and Frederic A. Rasio. Binary Neutron Star Mergers. *Living Reviews in Relativity*, 15(8), July 2012. doi: 10.12942/lrr-2012-8.

A. C. Fabian. Observational Evidence of Active Galactic Nuclei Feedback. *Annual Review of Astronomy and Astrophysics*, 50:455–489, June 2012. doi: 10.1146/annurev-astro-081811-125521.

S. A. E. G. Falle and S. S. Komissarov. An upwind numerical scheme for relativistic hydrodynamics with a general equation of state. *Monthly Notices of the Royal Astronomical Society*, 278:586–602, January 1996.

David Finkelstein. Past–Future Asymmetry of the Gravitational Field of a Point Particle. *Physical Review*, 110(4):965–967, May 1958. doi: 10.1103/PhysRev.110.965.

Leslie G. Fishbone and Vincent Moncrief. Relativistic Fluid Disks in Orbit around Kerr Black Holes. *The Astrophysical Journal*, 207:962–976, August 1976. doi: 10.1086/154565.

Charles F. Gammie. Efficiency of Magnetized Thin Accretion Disks in the Kerr Metric. *The Astrophysical Journal*, 522:L57–L60, September 1999. doi: 10.1086/312207.

Charles F. Gammie. The Magnetorotational Instability in the Kerr Metric. *The Astrophysical Journal*, 614: 309–313, October 2004. doi: 10.1086/423443.

Charles F. Gammie, Jonathan C. McKinney, and Gábor Tóth. HARM: A Numerical Scheme for General Relativistic Magnetohydrodynamics. *The Astrophysical Journal*, 589:444–457, May 2003. doi: 10.1086/374594.

Thomas A. Gardiner and James M. Stone. An unsplit Godunov method for ideal MHD via constrained transport. *Journal of Computational Physics*, 205:509–539, May 2005. doi: 10.1016/j.jcp.2004.11.016.

Thomas A. Gardiner and James M. Stone. An unsplit Godunov method for ideal MHD via constrained transport in three dimensions. *Journal of Computational Physics*, 227:4123–4141, April 2008. doi: 10.1016/j.jcp.2007.12.017.

Bruno Giacomazzo and Luciano Rezzolla. WhiskyMHD: a new numerical code for general relativistic magnetohydrodynamics. *Classical and Quantum Gravity*, 24:S235–S258, June 2007. doi: 10.1088/0264-9381/24/12/S16.

S. Gillessen, F. Eisenhauer, S. Trippe, T. Alexander, R. Genzel, F. Martins, and T. Ott. Monitoring Stellar Orbits Around the Massive Black Hole in the Galactic Center. *The Astrophysical Journal*, 692:1075–1109, February 2009. doi: 10.1088/0004-637X/692/2/1075.

Peter Goldreich and William H. Julian. Pulsar Electrodynamics. *The Astrophysical Journal*, 157:869–880, August 1969. doi: 10.1086/150119.

Jeremy Goodman. A Local Instability of Tidally Distorted Accretion Disks. *The Astrophysical Journal*, 406:596–613, April 1993. doi: 10.1086/172472.

Ami Harten and James M. Hyman. Self-Adjusting Grid Methods for One-Dimensional Hyperbolic Conservation Laws. *Journal of Computational Physics*, 50:235–269, May 1983. doi: 10.1016/0021-9991(83)90066-9.

Amiram Harten, Peter D. Lax, and Bram van Leer. On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws. *SIAM Review*, 25(1):35–61, January 1983. doi: 10.1137/1025002.

John F. Hawley, Larry L. Smarr, and James R. Wilson. A Numerical Study of Nonspherical Black Hole Accretion. I. Equations and Test Problems. *The Astrophysical Journal*, 277:296–311, February 1984. doi: 10.1086/161696.

Luis C. Ho. Nuclear Activity in Nearby Galaxies. *Annual Review of Astronomy and Astrophysics*, 46:475–539, September 2008. doi: 10.1146/annurev.astro.45.051806.110546.

Roy P. Kerr. Gravitational Field of a Spinning Mass as an Example of Algebraically Special Metrics. *Physical Review Letters*, 11(5):237–238, September 1963. doi: 10.1103/PhysRevLett.11.237.

S. S. Komissarov. A Godunov-type scheme for relativistic magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 303:343–366, February 1999. doi: 10.1046/j.1365-8711.1999.02244.x.

S. S. Komissarov. General relativistic magnetohydrodynamic simulations of monopole magnetospheres of black holes. *Monthly Notices of the Royal Astronomical Society*, 350:1431–1436, June 2004. doi: 10.1111/j.1365-2966.2004.07738.x.

Alexander Kurganov and Eitan Tadmor. New High-Resolution Central Schemes for Nonlinear Conservation Laws and Convection-Diffusion Equations. *Journal of Computational Physics*, 160(1):241–282, May 2000. doi: 10.1006/jcph.2000.6459.

Jane H. MacGibbon, B. J. Carr, and Don N. Page. Do evaporating black holes form photospheres? *Physical Review D*, 78:064043, September 2008. doi: 10.1103/PhysRevD.78.064043.

Barry Marder. A Method for Incorporating Gauss' Law into Electromagnetic PIC Codes. *Journal of Computational Physics*, 68:48–55, January 1987. doi: 10.1016/0021-9991(87)90043-X.

Jonathan C. McKinney, Alexander Tchekhovskoy, Aleksander Sadowski, and Ramesh Narayan. Three-dimensional general relativistic radiation magnetohydrodynamical simulation of super-Eddington accretion, using a new code HARMRAD with M1 closure. *Monthly Notices of the Royal Astronomical Society*, 441:3177–3208, July 2014. doi: 10.1093/mnras/stu762.

A. Mignone. High-order conservative reconstruction schemes for finite volume methods in cylindrical and spherical coordinates. *Journal of Computational Physics*, 270:784–814, August 2014. doi: 10.1016/j.jcp.2014.04.001.

A. Mignone and G. Bodo. An HLLC Riemann solver for relativistic flows – I. Hydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 364:126–136, November 2005. doi: 10.1111/j.1365-2966.2005.09546.x.

A. Mignone and G. Bodo. An HLLC Riemann solver for relativistic flows – II. Magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 368:1040–1054, May 2006. doi: 10.1111/j.1365-2966.2006.10162.x.

A. Mignone and Jonathan C. McKinney. Equation of state in relativistic magnetohydrodynamics: variable versus constant adiabatic index. *Monthly Notices of the Royal Astronomical Society*, 378:1118–1130, July 2007. doi: 10.1111/j.1365-2966.2007.11849.x.

A. Mignone, M. Ugliano, and G. Bodo. A five-wave Harten–Lax–van Leer Riemann solver for relativistic magnetohydrodynamics. *Monthly Notices of the Royal Astronomical Society*, 393:1141–1156, March 2009. doi: 10.1111/j.1365-2966.2008.14221.x.

M. Coleman Miller and E. J. M. Colbert. Intermediate-Mass Black Holes. *International Journal of Modern Physics D*, 13(1):1–64, January 2004. doi: 10.1142/S0218271804004426.

Charles W. Misner, Kip S. Thorne, and John Archibald Wheeler. *Gravitation*. W. H. Freeman and Company, New York, 1973.

F. Mohseni, M. Mendoza, S. Succi, and H. J. Herrmann. Relativistic effects on the Richtmyer–Meshkov instability. *Physical Review D*, 90:125028, December 2014. doi: 10.1103/PhysRevD.90.125028.

Philipp Mösta, Bruno C Mundim, Joshua A Faber, Roland Haas, Scott C Noble, Tanja Bode, Frank Löffler, Christian D Ott, Christian Reisswig, and Erik Schnetter. GRHydro: a new open-source general-relativistic magnetohydrodynamics code for the Einstein toolkit. *Classical and Quantum Gravity*, 31:015005, January 2014. doi: 10.1088/0264-9381/31/1/015005.

Curran D. Muhlberger, Fatemeh Hossein Nouri, Matthew D. Duez, Francois Foucart, Lawrence E. Kidder, Christian D. Ott, Mark A. Scheel, Béla Szilágyi, and Saul A. Teukolsky. Magnetic effects on the low-$T/|W|$ instability in differentially rotating neutron stars. *Physical Review D*, 90:104014, November 2014. doi: 10.1103/PhysRevD.90.104014.

Ramesh Narayan, Igor V. Igumenshchev, and Marek A. Abramowicz. Magnetically Arrested Disk: an Energetically Efficient Accretion Flow. *Publications of the Astronomical Society of Japan*, 55:L69–L72, December 2003. doi: 10.1093/pasj/55.6.L69.

William I. Newman and Nathaniel D. Hamlin. Primitive Variable Determination in Conservative Relativistic Magnetohydrodynamic Simulations. *SIAM Journal on Scientific Computing*, 36(4):B661–B683, July 2014. doi: 10.1137/140956749.

Scott C. Noble, Charles F. Gammie, Jonathan C. McKinney, and Luca Del Zanna. Primitive Variable Solvers for Conservative General Relativistic Magnetohydrodynamics. *The Astrophysical Journal*, 641:626–637, April 2006. doi: 10.1086/500349.

J. C. B. Papaloizou and J. E. Pringle. The dynamical stability of differentially rotating discs with constant specific angular momentum. *Monthly Notices of the Royal Astronomical Society*, 208:721–750, June 1984. doi: 10.1093/mnras/208.4.721.

R. Penrose. Gravitational Collapse: The Role of General Relativity. *Rivista del Nuovo Cimento*, 1:252, 1969.

José A. Pons, José A. Font, José Mª. Ibáñez, José Mª. Martí, and Juan A. Miralles. General relativistic hydrodynamics with special relativistic Riemann solvers. *Astronomy and Astrophysics*, 339:638–642, November 1998.

William H. Press, Saul A. Teukolsky, William T. Vetterling, and Brian P. Flannery. *Numerical Recipes: The Art of Scientific Computing*. Cambridge University Press, third edition, 1997.

John D. Ramshaw. A Method for Enforcing the Solenoidal Condition on Magnetic Field in Numerical Calculations. *Journal of Computational Physics*, 52:592–596, December 1983. doi: 10.1016/0021-9991(83)90009-8.

S. M. Ressler, A. Tchekhovskoy, E. Quataert, M. Chandra, and C. F. Gammie. Electron thermodynamics in GRMHD simulations of low-luminosity black hole accretion. *Monthly Notices of the Royal Astronomical Society*, 454:1848–1870, December 2015. doi: 10.1093/mnras/stv2084.

P. L. Roe. Approximate Riemann Solvers, Parameter Vectors, and Difference Schemes. *Journal of Computational Physics*, 43:357–372, October 1981. doi: 10.1016/0021-9991(81)90128-5.

B. R. Ryan, J. C. Dolence, and C. F. Gammie. bhlight: General Relativistic Radiation Magnetohydrodynamics with Monte Carlo Transport. *The Astrophysical Journal*, 807:31, July 2015. doi: 10.1088/0004-637X/807/1/31.

Dongsu Ryu and Jeremy Goodman. Convective Instability in Differentially Rotating Disks. *The Astrophysical Journal*, 388:438–450, April 1992. doi: 10.1086/171165.

V. Schneider, U. Katscher, D. H. Rischke, B. Waldhauser, and J. A. Maruhn. New Algorithms for Ultrarelativistic Numerical Hydrodyamics. *Journal of Computational Physics*, 105:92–107, March 1993. doi: 10.1006/jcph.1993.1056.

R. Schödel, T. Ott, R. Genzel, R. Hofmann, M. Lehnert, A. Eckart, N. Mouawad, T. Alexander, M. J. Reid, R. Lenzen, M. Hartung, F. Lacombe, D. Rouan, E. Gendron, G. Rousset, A.-M. Lagrange, W. Brandner, N. Ageorges, C. Lidman, A. F. M. Moorwood, J. Spyromillo, N. Hubin, and K. M. Menten. A star in a 15.2-year orbit around the supermassive black hole at the centre of the Milky Way. *Nature*, 419:694–696, October 2002. doi: 10.1038/nature01121.

K. Schwarzschild. Über das Gravitationsfeld eines Massenpunktes nach der Einsteinchen Theorie. *Sitzungsberichte der Königlich Preußischen Akademie der Wissenschaften*, pages 189–196, 1916.

N. I. Shakura and R. A. Sunyaev. Black Holes in Binary Systems. Observational Appearance. *Astronomy and Astrophysics*, 24:337–355, 1973.

Masaru Shibata and Yu-ichirou Sekiguchi. Magnetohydrodynamics in full general relativity: Formulation and tests. *Physical Review D*, 72:044014, August 2005. doi: 10.1103/PhysRevD.72.044014.

Hotaka Shiokawa, Joshua C. Dolence, Charles F. Gammie, and Scott C. Noble. Global General Relativistic Magnetohydrodynamical Simulations of Black Hole Accretion Flows: A Convergence Study. *The Astrophysical Journal*, 744:187, January 2012. doi: 10.1088/0004-637X/744/2/187.

Chi-Wang Shu and Stanley Osher. Efficient Implementation of Essentially Non-oscillatory Shock-Capturing Schemes. *Journal of Computational Physics*, 77:439–471, August 1988. doi: 10.1016/0021-9991(88)90177-5.

James M. Stone and Thomas A. Gardiner. Athena: A New Code for Astrophysical MHD. *The Astrophysical Journal Supplement Series*, 178:137–177, September 2008. doi: 10.1086/588755.

Alexander Tchekhovskoy, Ramesh Narayan, and Jonathan C. McKinney. Efficient generation of jets from magnetically arrested accretion on a rapidly spinning black hole. *Monthly Notices of the Royal Astronomical Society*, 418:L79–L83, September 2011. doi: 10.1111/j.1745-3933.2011.01147.x.

Gábor Tóth. The $\nabla \cdot B = 0$ Constraint in Shock-Capturing Magnetohydrodynamics Codes. *Journal of Computational Physics*, 161:605–652, July 2000. doi: 10.1006/jcph.2000.6519.

Bram van Leer. Towards the Ultimate Conservative Difference Scheme. V. A Second-Order Sequel to Godunov's Method. *Journal of Computational Physics*, 32:101–136, July 1979. doi: 10.1016/0021-9991(79)90145-1.

Christopher J. White, James M. Stone, and Charles F. Gammie. An Extension of the Athena++ Code Framework for GRMHD Based on Advanced Riemann Solvers and Staggered-Mesh Constrained Transport. *The Astrophysical Journal Supplement Series*, 2016.

S. E. Woosley. Gamma-Ray Bursts from Stellar Mass Accretion Disks around Black Holes. *The Astrophysical Journal*, 405:273–277, March 1993. doi: 10.1086/172359.

S. E. Woosley and Thomas A. Weaver. The Evolution and Explosion of Massive Stars. II. Explosive Hydrodynamics and Nucleosynthesis. *The Astrophysical Journal Supplement Series*, 101:181–235, November 1995. doi: 10.1086/192237.

O. Zanotti and M. Dumbser. High order numerical simulations of the Richtmyer–Meshkov instability in a relativistic fluid. *Physics of Fluids*, 27:074105, July 2015. doi: 10.1063/1.4926585.