Turbulent Transport in Global Models of Magnetized Accretion Disks

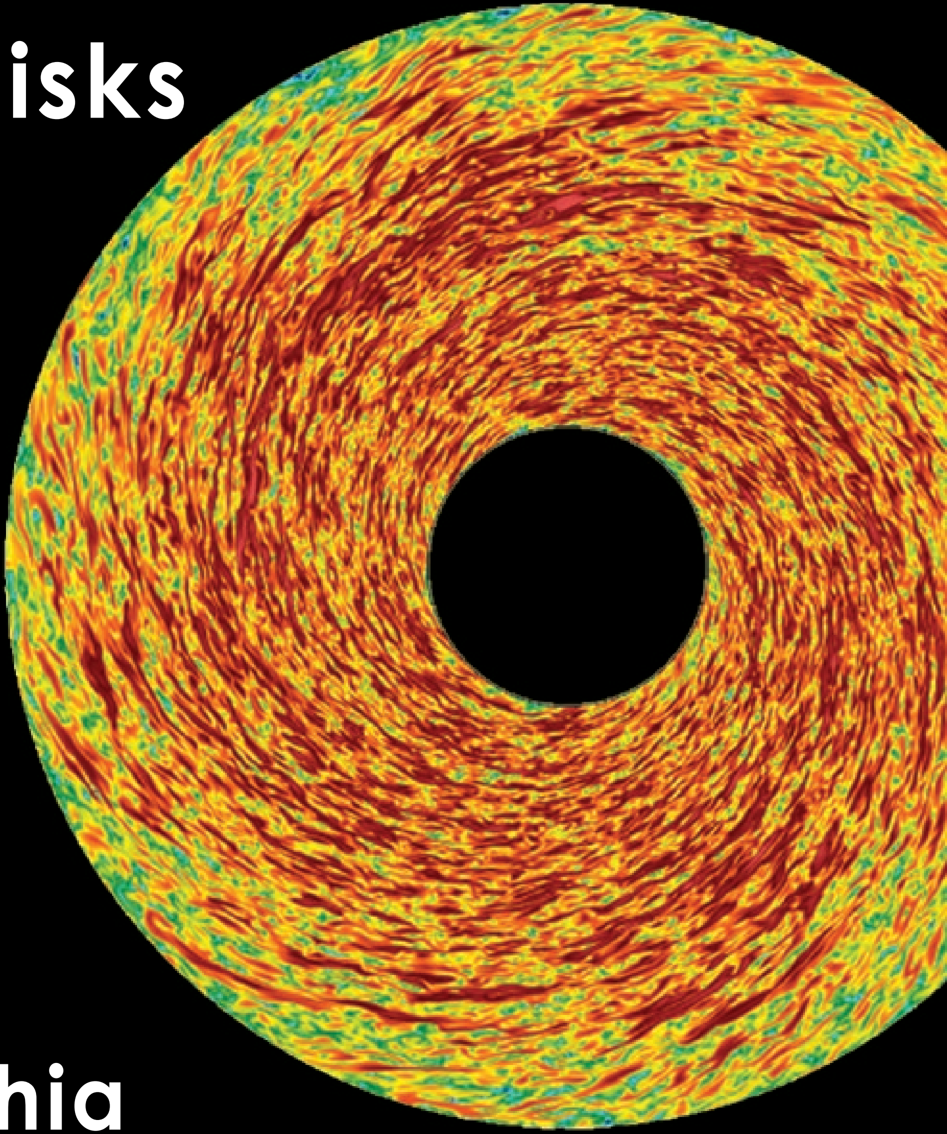

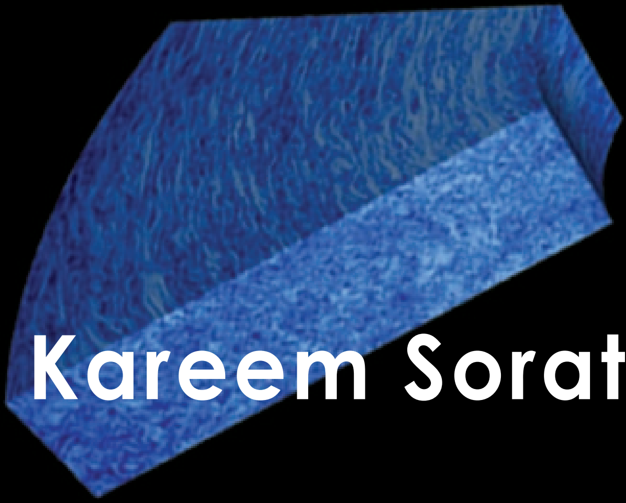

Kareem Sorathia

Advisory Committee:

Chris Reynolds (Chair)
William Dorland
David Levermore
Cole Miller
Eve Ostriker

ABSTRACT

Title of dissertation: Turbulent Transport In Global
Models of Magnetized Accretion Disks

Kareem Sorathia, Doctor of Philosophy, 2011

Dissertation directed by: Professor Christopher Reynolds
Department of Astronomy

The modern theory of accretion disks is dominated by the discovery of the magnetorotational instability (MRI). While hydrodynamic disks satisfy Rayleigh's criterion and there exists no known unambiguous route to turbulence in such disks, a weakly magnetized disk of plasma is subject to the MRI and will become turbulent. This MRI-driven magnetohydrodnamic turbulence generates a strong anisotropic correlation between the radial and azimuthal magnetic fields which drives angular momentum outwards. Accretion disks perform two vital functions in various astrophysical systems: an intermediate step in the gravitational collapse of a rotating gas, where the disk transfers angular momentum outwards and allows material to fall inwards; and as a power source, where the gravitational potential energy of infalling matter can be converted to luminosity. Accretion disks are important in astrophysical processes at all scales in the universe. Studying accretion from first principles is difficult, as analytic treatments of turbulent systems have proven quite limited. As such, computer simulations are at the forefront of studying systems this far into the non-linear regime.

While computational work is necessary to study accretion disks, it is no panacea. Fully three-dimensional simulations of turbulent astrophysical systems require an enormous amount of computational power that is inaccessible even to sophisticated modern supercomputers. These limitations have necessitated the use of local models, in which a small spatial region of the full disk is simulated, and constrain numerical resolution to what is feasible. These compromises, while necessary, have the potential to introduce numerical artifacts in the resulting simulations. Understanding how to disentangle these artifacts from genuine physical phenomena and to minimize their effect is vital to constructing simulations that can make reliable astrophysical predictions and is the primary concern of the work presented here.

The use of local models is predicated on the assumption that these models accurately capture the dynamics of a small patch of a global astrophysical disk. This assumption is tested in detail through the study of local regions of global simulations. To reach resolutions comparable to those used in local simulations an orbital advection algorithm, a semi-Lagrangian reformulation of the fluid equations, is used which allows an order of magnitude increase in computational efficiency. It is found that the turbulence in global simulations agrees at intermediate- and small-scales with local models and that the presence of magnetic flux stimulates angular momentum transport in global simulations in a similar manner to that observed for local ones. However, the importance of this flux-stress connection is shown to cast doubt on the validity of local models due to their inability to accurately capture the temporal evolution of the magnetic flux seen in global simulations.

The use of orbital advection allows the ability to probe previously-inaccessible resolutions in global simulations and is the basis for a rigorous resolution study presented here. Included are the results of a study utilizing a series of global simulations of varying resolutions and initial magnetic field topologies where a collection of proposed metrics of numerical convergence are explored. The resolution constraints necessary to establish numerical convergence of astrophysically-important measurements are presented along with evidence suggesting that the use of proper azimuthal resolution, while computationally-demanding, is vital to achieving convergence. The majority of the proposed metrics are found to be useful diagnostics of MRI-driven turbulence, however they suffer as metrics of convergence due to their dependence on the initial magnetic field topology. In contrast to this, the magnetic tilt angle, a measure of the planar anisotropy of the magnetic field, is found to be a powerful tool for diagnosing convergence independent of initial magnetic field topology.

Turbulent Transport In Global Models of Magnetized Accretion Disks

by

Kareem Sorathia

Dissertation submitted to the Faculty of the Graduate School of the
University of Maryland, College Park in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
2011

Advisory Committee:
Professor Christopher Reynolds, Chair/Advisor
Professor William Dorland
Professor David Levermore
Professor Cole Miller
Professor Eve Ostriker

Preface

Portions of the research described in this work have been published or submitted to The Astrophysical Journal. Chapter 3 appeared in a modified form in Sorathia et al. (2010). The entirety of Chapter 5 and material from Chapter 6 have been submitted as Sorathia et al. (2011) and includes a reduced version of Chapter 4 as an appendix. The remainder of the material of Chapter 6 will form the basis of a future work. The content of these papers has been modified and expanded where appropriate for consistency and clarity.

Dedication

To my parents, for supporting my seemingly unending education, and in memory of Vali and Ahmed Sorathia, whose love and support will be dearly missed.

Acknowledgments

Enumerating all of the people to whom I owe my gratitude for their support during my education would likely be impossible, however I would like to explicitly acknowledge some.

First and foremost, I would like to thank my advisor, Professor Christopher Reynolds, without whom this dissertation would not have been possible. His encouragement and advice has been instrumental in shaping me as a scientist and it has been a pleasure to have been mentored by such a kind and intelligent person. I would also like to thank my committee for reading my dissertation and offering valuable suggestions and comments that led to an improved manuscript. I have been quite lucky to have the opportunity to collaborate on research with many intelligent and thoughtful scientists and I have benefited greatly from their experience and knowledge. I would like to thank Professor James Stone, Professor Phil Armitage, Professor Cole Miller, Dr. Sean O'Neill, and Dr. Kris Beckwith. Finally, if it was not for the encouragement and mentorship I received as an undergraduate I likely would not have pursued higher education and I would like to thank Professor Thomas Armstrong, Professor Mark Henriksen, and Dr. Eric Tittley for igniting and encouraging my passion for mathematics and astrophysics.

Surviving six years of graduate school would not have been possible without the support structure provided by my friends and family. The friendships I've made with my fellow graduate students have broadened my horizons as a scientist and kept me sane through late nights of studying and setbacks in research. I would like to thank my colleagues and officemates to whom I'm indebted: Andy Sanders, Emily King, Nathaniel Strawn, Malina Kirn, Shelby Wilson, and Jeffrey Frazier. In particular, I would like to thank Aaron Skinner who has at times been an officemate, a roommate, and classmate but has always been an all around great mate. I would also like to thank Charlie Wyse for making sure I never worked too hard or forgot that school wasn't all that mattered in life. Throughout my life my sister, Sadia, has been a constant source of support and inspiration and I am looking forward to celebrating the completion of both our journeys in education together. During the most difficult parts of my education Sarah Buchanan has been a constant source of strength. The desire to live up to her opinion of me has made me a better person and the completion of this dissertation is in no small part her accomplishment as well as mine.

Finally, I would like to thank the National Science Foundation for financial support and providing computational resources through TeraGrid as well as NASA for its financial support.

Table of Contents

List of Tables

List of Figures

Chapter 1

Scientific Background

1.1 Introduction

Accretion, at its simplest, is merely the infall of matter onto a gravitational source. This simple definition, however, belies the complex mechanism that underlies this important astrophysical phenomenon. Understanding the details of this mechanism leads one to the intersection of three disciplines: plasma physics, astrophysics, and computational science. Accretion is of vital importance at a range of scales in the Universe: from its role in mediating gravitational collapse in the formation of stars and planetary systems to its role as a power source in active galaxies. The mechanism of accretion is reliant upon plasma turbulence, and as is common in turbulent systems, analytic treatments are limited. Simulations, therefore, are at the forefront of studying these systems.

While these details will be expanded upon below, for now an accretion disk will be defined through a simple illustration: a cloud of gas about a central gravitational source. In the singular case where all of the dynamics of the system are spherically symmetric (centered on the gravitating mass) and the gas far from the central mass is at rest, we have Bondi accretion (Bondi, 1952) in which the gas dynamical equations can be exactly solved using analytical methods. Difficulty arises, however, upon consideration of the more general situation in which the gas is rotating about the central source. In this case, the gas possesses centrifugal support against the inward gravitational force of the central source. However, vertical collapse proceeds until impeded by internal pressure support. This results in a thin, rotating disk of gas about the central object. Due to the inverse square spatial dependence of the gravitational force the disk is differentially-rotating, specifically that the orbital rotation profile is decreasing with radius. *The central problem of accretion disk theory is the precise manner in which gravitational contraction can proceed despite the centrifugal barrier imposed by the conservation of angular momentum.*

The primary function of an accretion disk is to mediate gravitational contraction by redistributing angular momentum away from the inner regions of the disk and thus allowing matter to fall inwards. As will be described, weakly magnetized differentially rotating disks are subject to an instability, the magnetorotational instability, that results in rapid and significant magnetic field amplification during the linear growth phase of the instability followed by saturation and transition into a fully turbulent flow. This turbulence, highly anisotropic due to the shear of differential rotation, will induce a radially outward angular momentum flux that allows matter infall.

Accretion is a vital mechanism of structure formation in the Universe; stars, planets, black holes, and even galaxies are formed through accretion processes. In many cases there is a sufficient presence of angular momentum for accretion disks to form as a intermediate step in the formation of astrophysical structure. Additionally, accretion onto compact objects, particularly black holes, is a tremendously efficient source of power and is indeed the operant mechanism behind the most luminous phenomena in the Universe. Accretion onto black holes is significantly more efficient than fusion at converting mass into energy, and represents the second most prolific source of radiative energy in the Universe.

Computational modeling, while a necessity of studying systems this far into the nonlinear regime, is no panacea. Choices one makes regarding algorithms, models, and initial conditions can leave imprints, or numerical artifacts, in the induced turbulence. Understanding how to disentangle numerical artifacts from genuine physics is a vital precondition towards the overall goal of being able to reconcile the predictions of computer simulations with observational data. It is towards this goal that the bulk of the research described here will be directed.

1.2 The Role of Accretion in Astrophysics

Astrophysical accretion disks can broadly be classified into three types: protoplanetary disks (PPDs), a precursor to the formation of stars and planetary systems; accretion disks in binary star systems, in which an accretion disk mediates mass transfer from a donor star onto a compact object (white dwarf, neutron star, or black hole); and disks in active galactic nuclei (AGN), where accretion disks facilitate matter inflow onto supermassive black holes (SMBHs). Accretion in binary systems generally involves a donor star that has expanded to fill its Roche lobe, the critical equipotential of the binary gravitational potential associated with the L_1 Lagrange point. Matter from the donor star streams through the Lagrange point, becoming gravitationally bound to the compact object, and forms an accretion disk.

At a much larger scale, AGN are associated with SMBHs at the centers of galaxies. Not all SMBHs are associated with AGN however; in the majority of galaxies, very little matter accretes onto the SMBH and thus very little "activity" results. Our own SMBH in the Milky Way is in this class.

While the mechanism underlying angular momentum transport is the same in all these systems, PPDs can be distinguished from binary and AGN disks due to their disparate thermodynamics and composition. The presence of dust and comparatively cold temperatures present in PPDs result in a complex chemistry and the need to treat resistive effects carefully. The significantly higher thermal energies characteristic of binary and AGN disks make these systems amenable to a treatment utilizing simpler and more well-posed mathematical models predicated on the assumptions of a fully ionized hydrogen plasma. For this reason, the results and simulations described here will be primarily relevant to these two types of systems. The outer regions of binary and AGN disks are likely to exhibit significant differences due to their differing geometric structure, however the interior regions are expected to be quite similar. The simulations described here will, in general, remain agnostic as to the specific context of the high-energy disk by restricting the simulation domain to these inner regions.

Interest in accretion in the context of high energy astrophysics began in the 1960's with the advent of X-ray telescopes. The ability to study X-ray emission opened a new window through which we could study the most powerful and luminous objects in the Universe. It was realized early on in the study of these objects that the nuclear processes that powered stars were wholly inadequate to explain the immense energetics of these systems. The search for a power source that was more efficient than even fusion eventually led to gravity, specifically accretion.

As a first step towards understanding the manner in which accretion can act as a power source we perform a simple calculation in which we consider the change in gravitational potential energy of a particle of mass m infalling from infinity onto a Newtonian gravitational source of mass and radius M_* and R_*, respectively. The change in gravitational potential energy is given by

$$\Delta E = Gm\frac{M_*}{R_*} = \left(\frac{GM_*}{R_* c^2}\right) mc^2. \tag{1.1}$$

In particular, the dependence on the dimensionless ratio of mass to radius suggests the importance of accretion onto compact objects: white dwarfs, neutron stars, and black holes. The obvious flaw in this simple calculation is that while we have identified a source of energy, the mechanism by which gravitational potential energy is converted to radiative emission remains unclear. Indeed, this example suggests that the entirety of the potential energy will be converted into kinetic energy which would be deposited onto the surface of a white dwarf and neutron star, or carried into the event horizon of a black hole. In addition to its role in mediating gravitational contraction through angular momentum transport, the accretion disk also functions to couple the loss of gravitational potential energy to thermal radiation emission. In reality, a particle infalling from infinity will not fall directly onto the object, it will slowly spiral inwards through the disk as its angular momentum is drained. During this spiral infall, turbulent dissipation will act to convert gravitational potential energy into thermal energy.

A common metric used to measure the efficiency of a power source for a luminous phenomenon is the parameter η which relates the luminosity, L, to the mass-flux, $\dot{M}$, converted to an energy-flux by scaling by the square of the speed of light through

$$L = \eta \dot{M} c^2. \tag{1.2}$$

For stellar processes, in which hydrogen is fused into helium the efficiency parameter is given by, $\eta \approx 0.007$, where in this case the mass flux corresponds to the mass of stellar material processed by the nuclear reactions per unit time. In contrast, estimates of the efficiency of a non-rotating, or Schwarzschild, black hole yields an efficiency of $\eta \approx 6\%$. For rotating, or Kerr, black holes values of η approach 40% for maximally-rotating black holes. The superior efficiency of accretion in converting rest-mass energy to radiative energy is vital to understanding observed ultra-luminous phenomena like AGN, where observed luminosities can reach $10^{40}W$.

1.3 Turbulent Nature of Accretion Disks

The turbulent nature of accretion disks is hinted at in attempts to explain angular momentum transport in the context of laminar flow. Towards this end, we begin with the development of the theory of thin accretion disks. The geometry of these structures is amenable to a treatment utilizing a cylindrical coordinate system, defined as (R, ϕ, z). For brevity and clarity, the cylindrical and spherical radius are distinguished with the former given as R, and the latter as $r = \sqrt{R^2 + z^2}$. Accretion disks are quite large, and in particular the length scales of interest are generally far in excess of the microscopic kinetic scales, thus the appropriate equations are those in the continuum limit. For hydrodynamic disks the relevant equations are the Navier-Stokes equations, a set of non-linear coupled partial differential equations relating the fluid variables: the scalar density, ρ, and the vectorial quantity velocity, $\mathbf{v}$. These quantities are functions of space and time, and are related through equations representing the conservation of mass and momentum with dissipative viscous terms given by

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho \mathbf{v}) = 0, \tag{1.3}$$

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho \mathbf{v} \cdot \nabla \mathbf{v} = -\nabla P - \rho \nabla \Phi_N + \rho \nu \left(\nabla^2 \mathbf{v} + \frac{1}{3} \nabla (\nabla \cdot \mathbf{v}) \right). \tag{1.4}$$

This set of equations is incomplete and requires a formulation of the scalar pressure, P. This can be accomplished through the addition of an equation encapsulating the conservation of energy and an adiabatic equation of state. Alternatively, a simpler treatment can be used in which an isothermal equation of state is used that omits the internal energy and simply relates the gas pressure to the density by

$$P = {c_s}^2 \rho, \tag{1.5}$$

where the speed of sound, c_s, is assumed constant throughout the disk. Physically, the assumption of an isothermal equation of state is the statement that the timescale associated with radiative equilibration is much smaller than the kinematic timescales of interest. The gravitational potential of the central source is included, represented by the static, axisymmetric gravitational potential Φ, however, the self-gravity of the disk itself will be neglected. The Navier-Stokes equations differ from the Euler fluid dynamic equations due the inclusion of a viscous dissipation term, ν.

An accretion disk can be thought of as a collection of rings of gas rotating at close to the Keplerian velocity. The Keplerian velocity, v_K, is defined as the rotational velocity associated with a circular orbit in the gravitational potential of the central source of mass M. For a Newtonian disk, in which the gravitational potential is given as

$$\Phi_N = -\frac{GM}{r}, \tag{1.6}$$

then the Keplerian velocity is given as

$$v_K = \sqrt{\frac{GM}{r}}. \tag{1.7}$$

That the disk is geometrically thin, $R \gg z$, allows a simple treatment of the disk structure in which the vertical and radial structure decouple. The azimuthal velocity is simply the Keplerian value with the vertical dependence neglected, *i.e.* $v_K \propto R^{-1/2}$. Similarly, we define the Keplerian angular velocity for a thin disk as

$$\Omega_K = \sqrt{\frac{GM}{R^3}}, \tag{1.8}$$

where the linear and angular velocities are related in the standard manner by the relationship $v_K = R\Omega_K$. The density structure is given by the solution of the equation representing vertical hydrostatic equilibrium, and is given by

$$\rho = \rho_0 \exp(-z^2/H^2), \tag{1.9}$$

where the pressure scale height is defined as

$$H = \sqrt{2}\frac{c_s}{\Omega_K}. \tag{1.10}$$

It is worth noting that the pressure scale height is a function of radius through its dependence on the Keplerian angular velocity, and that definitions in the literature vary regarding the inclusion of the constant coefficient, $\sqrt{2}$. In non-isothermal disks the temperature, and thus the sound speed, will vary with height and although hydrostatic equilibrium would still apply this will result in a non-Gaussian vertical density profile. For clarity, results from the literature presented here will be modified as necessary so as to be consistent with this definition. As a requirement for self-consistency with the assumption of a geometrically thin disk, one expects that $H/R \ll 1$. This is equivalent to the statement that the Keplerian velocity is supersonic, specifically that $M_K \gg 1$, where the Keplerian Mach number is defined as

$$M_K = \frac{v_K}{c_s} \approx \left(\frac{H}{R}\right)^{-1}. \tag{1.11}$$

This ratio, H/R, has both important physical and numerical consequences: it is a dimensionless measure of the temperature in the disk and defines the partition of disk energy between the thermal energy of the gas and the coherent rotational energy; it is also a measure of the disparity of scales in the disk, measuring the separation between the length scales H and R as well as the dominant signal speeds, v_K and c_s. Large separation of spatial and temporal scales present difficult numerical challenges, as simultaneously resolving all the important scales is computationally very demanding. Conversely, it is interesting to note that it is precisely this separation of scales that makes analytic progress on thin disks possible.

The disk structure described above can be used to construct a simple, one dimensional analytic model. The fluid variables are assumed to be axisymmetric and, with the exception of the density, to have no dependence on height. Further, it assumed that the fluid is constrained to orbit at the local Keplerian velocity, thus $v_\phi = v_K$. To treat the density in a one-dimensional manner, we define the surface density $\Sigma(R) = \int \rho dz$, as the vertically-integrated density. Rewriting equations 1.3 and 1.4 in their vertically-integrated form and in terms of the angular momentum yields

$$R\frac{\partial \Sigma}{\partial t} + \frac{\partial}{\partial R}(R\Sigma v_R) = 0, \tag{1.12}$$

$$R\frac{\partial}{\partial t}(R^2\Sigma\Omega) + \frac{\partial}{\partial R}(R^3\Sigma v_R \Omega) = \frac{\partial}{\partial R}\left(\nu R^3 \Sigma \frac{\partial \Omega}{\partial R}\right). \tag{1.13}$$

Using equation 1.12 to simplify equation 1.13 allows the radial velocity to be isolated and solved as

$$v_R = \frac{-3}{\sqrt{R}\Sigma}\frac{\partial}{\partial R}(\nu\sqrt{R}\Sigma), \tag{1.14}$$

where a Newtonian gravitational potential has been assumed ($\Omega_K \propto R^{-3/2}$).

Substituting the radial velocity into equation 1.12 yields the evolution equation for the surface density

$$\frac{\partial \Sigma}{\partial t} = \frac{3}{R}\frac{\partial}{\partial R}\left[\sqrt{R}\frac{\partial(\nu\Sigma\sqrt{R})}{\partial R}\right]. \tag{1.15}$$

Equation 1.15 encapsulates a simple and intuitive model of angular momentum transport in the form of a diffusion equation. The Keplerian velocity decreases with distance from the central source, and thus the inner rings will rotate faster than the outer rings. The action of a shear viscosity between a pair of radially adjacent rings will transfer angular momentum outwards. The viscous interaction between the faster inner ring and the slower outer ring induces a torque whose orientation is antiparallel to the rotation vector of the disk. This torque results in the transfer of angular momentum from the inner ring to the outer ring, and the iterative interactions of adjacent rings would allow sustained angular momentum transport radially outwards. The simplicity of this model makes it quite attractive, and indeed as a zeroth-order approximation it will turn out that this model has many useful applications. The difficulty, however, comes about when attempting to identify the nature of the viscosity.

The natural candidate for viscosity in an accretion disk is molecular/ionic viscosity caused by kinetic-scale effects. Consideration of the Reynolds number, $Re = Rv_K/\nu$, allows one to assess the importance of molecular/ionic viscosity. Physically, the quantity represents the ratio of the inertia of the flow to the viscous dissipation. Thus, small values of Re are associated with dissipative flows whereas larger values are associated with relatively inviscid flows. Estimates of this quantity using values appropriate to those in an accretion disk yields a Reynolds number associated with molecular/ionic viscosity of $Re_{mol} > 10^{14}$ (Frank et al., 2002). The high value associated with the Reynolds number suggests that molecular/ionic viscosity is negligible, and thus not a significant source of angular momentum transport. Consequently, the inclusion of microphysical viscous dissipation will be neglected moving forwards. However, the large value of the Reynolds number does suggest an underlying turbulent nature to astrophysical accretion disks.

Experience with laboratory fluids teaches one to associate large Reynolds numbers with turbulent flows, however one must be careful not to mistake this correlation with a causal relationship. The Reynolds number is the ratio of the inertia, represented by the convective nonlinearity term ($\rho\mathbf{v} \cdot \nabla\mathbf{v}$) in the fluid equations, and the viscous term ($\nu\rho\nabla^2\mathbf{v}$). Thus, a large Reynolds number is merely a measure of the relative importance of these terms. A fluid in equilibrium at high Reynolds numbers, already subject to a linear instability, would only require a small perturbation to undergo a transition to turbulence. At low Reynolds numbers, dissipation would dominate and smooth perturbations to stabilize the equilibrium. Neglecting, temporarily, the question of the existence of a linear instability, the presence of turbulence is often associated with enhanced transport. The presence of turbulence would result in a turbulent viscosity with a characteristic scale related to the size of the turbulent eddies, and this turbulent viscosity could be many orders of magnitude larger than the viscosity associated with molecular, kinetic-scale effects. Again, however, this is not to say that the presence of turbulence alone is sufficient to guarantee sustained, coherent momentum transport. As we will see, for the sustained radial transport of angular momentum to proceed the turbulence has to be highly anisotropic, with strong correlations between the radial and azimuthal components of the vector fields associated with the flow.

With a preliminary understanding of the structure of an accretion disk in place, we can proceed to the more pressing question as to the nature of the instability. The hydrodynamic stability of a differentially-rotating system to linear, axisymmetric perturbations is governed by the Rayleigh criterion. The criterion is given in terms of the epicyclic frequency,

$$\kappa^2 = \frac{1}{R^3}\frac{d(R^4\ {\Omega_K}^2)}{dR}, \tag{1.16}$$

the frequency at which a radially displaced fluid element will oscillate. The Rayleigh criterion for local stability is simply the condition that $\kappa^2 > 0$, or that the specific angular momentum, $\ell = R^2\Omega_K$, should increase with radius. Physically, this is the statement that if a fluid element is radially displaced, while conserving angular momentum, the centrifugal imbalance will act as a force opposed to the perturbation. While this condition may be violated in laboratory fluids, it will clearly be satisfied for astrophysical disks where $\ell = \sqrt{GMR}$.

The difficulty of the situation presented by the realization that turbulence was the likely cause of significant angular momentum transport while being unable to identify a cause can not be overstated and was the source of significant theoretical difficulty for several decades. While the magnetic character of the instability will be discussed in the following section, for now a brief description of the anomalous viscosity model of Shakura & Sunyaev (1973) will be presented. The anomalous viscosity, or α model, represents an important bridge and its simplicity has lent it a longevity that continues to influence accretion disk theory today.

Anomalous viscosity models owe their simplicity to their ability to treat complex turbulent phenomenon in the context of a one-dimensional hydrodynamic model like equation 1.15 with the addition of a prescribed form of enhanced viscosity. There are two, roughly equivalent, means in which to treat turbulent transport in these models. The first, assumes a turbulent viscosity of the form

$$\nu_T = \alpha_{SS} c_s H. \tag{1.17}$$

While this is written in the form of an isotropic Navier-Stokes viscosity, it is applied in the vertically-integrated, axisymmetric equations and thus this quantity is only meant to cause radial viscous diffusion. This form of viscosity is merely dimensional analysis used to quantify the viscosity given a characteristic velocity and turbulent eddy size and based on the observation that the relevant scales are given by the sound speed and scale height of the disk and therefore the viscosity must take this form. Intuitively, this is based on the expectation that turbulent velocity fluctuations will be, in a statistical sense, subsonic as supersonic velocity perturbations would become highly dissipative and that the size of a turbulent eddy will be limited by the vertical span of the disk. Thus, we expect the parameter α_{SS} to be bound between zero and unity.

The alternative formulation, assumes a prescribed form of the stress tensor in the Euler equation. One can write the equation representing conservation of momentum, absent viscous dissipation, in the manifestly hyperbolic form as

$$\frac{\partial(\rho\mathbf{v})}{\partial t} + \nabla\cdot\mathbf{W} = 0, \tag{1.18}$$

where $\mathbf{W}$ is the stress tensor, which for hydrodynamic systems is given by $\mathbf{W} = \rho\mathbf{v}\mathbf{v} + P\mathbf{I}$, where $\mathbf{I}$ is the identity tensor. Assuming a similar dimensional analysis to the formulation of the turbulent viscosity, one sets $\mathbf{W}_{R\phi} = \alpha_{SS}P$. This constant of proportionality, α_{SS}, that measures the relationship between the stress tensor and the gas pressure has become the most common scalar diagnostic of accretion efficacy. It is worth noting that while the same constant is used in both formulations of the anomalous viscosity model, strictly speaking these constants are not equal but are related by a coefficient of order unity. The treatment of a complex turbulent phenomenon in the context of these models is appealing as all of the uncertainty regarding the nature of the turbulence is confined to the choice of α_{SS}.

While a full treatment of these α-viscosity models will be neglected, particularly the inclusion of radiative transfer and its predictions for observational signatures of accretion disks, it is useful to consider the important timescales for these models as they will be considered again in the context of the full simulations presented. The most important timescales are the dynamical timescale, $T_d = {\Omega_K}^{-1}$, and the viscous timescale, $T_\nu = R^2/\nu$. The former is the fastest timescale as it is associated with the fastest signal speed, the Keplerian rotation. The latter represents the timescale on which the radial structure of the disk can vary, and can be thought of as a measure of the time it takes for matter to accrete through the disk. Under the assumptions of a thin disk we expect the viscous timescale to be significantly larger than the dynamical timescale, specifically $T_\nu = {M_K}^2 T_d/\alpha_{SS}$. Related to this, we define the timescale associated with changes in the vertical structure of the disk, $T_z = H/c_s$, and expect the relationship $T_d \approx T_z$.

While theory provides little guidance towards estimating values of α_{SS}, this parameter does represent an important intersection between simulation and observation. Estimates of this quantity from simulation will be presented in later chapters, for now we briefly discuss observational estimates. Observational constraints on α_{SS} are, in general, based on the estimation of the viscous (or in some cases thermal) timescale of an astrophysical accretion disk, from which the parameter α_{SS} can be calculated. A specific example of this occurs in a subclass of cataclysmic variables (CVs), binary systems in which a white dwarf accretes from a donor star. This subclass is subject to outbursts (referred to as “dwarf novae”), in which the accretion disk can switch from a cold to hot state. In the cold state, the ionization fraction of the plasma decreases and resistive effects suppress the instability that drives turbulence (Gammie & Menou, 1998). The transition to the hot state increases the ionization fraction, and accretion proceeds on a timescale given by T_ν. This timescale can be estimated observationally, and thus an approximate value of α_{SS} can be computed. In a survey paper, King et al. (2007) collect results using this method as well as several others to compute an estimated value of $\alpha_{SS} \approx 0.1-0.4$ for fully ionized accretion disks. As we will see, estimates of α_{SS} from simulations are generally significantly lower except when the MRI is driven by a particularly strong vertical magnetic field, however the existence of such strong vertical magnetic fields in astrophysical systems is contentious (van Ballegooijen, 1989).

Anomalous viscosity models of disk accretion are written in the guise of hydrodynamics, but that should not be taken as evidence that early treatments of accretion disk theory did not appreciate the importance of magnetic fields. Indeed, Shakura & Sunyaev (1973) discuss the likelihood that magnetic fields play an important role in accretion disk turbulence nearly two decades before the linear instability that induces turbulence, the magnetorotational instability (MRI) was fully appreciated. The final precursor to a full discussion of this instability will be a presentation of the relevant equations.

1.4 Magnetohydrodynamics

The equations of ideal magnetohydrodynamics (MHD), are at their simplest the coupling of the Euler equations of fluid dynamics to Maxwell's equations of electrodynamics and are used to study the behavior of plasmas at large scales. A full derivation will be neglected, however some of the implicit assumptions of the model will be detailed. The equations of MHD assume a continuum model, and thus are of no use towards exploring the behavior of a plasma at kinetic scales: the collisional mean free path and the cyclotron radius of the particles in the plasma. The equations utilize a one fluid model, in which hydrodynamic quantities simultaneously describe the properties of the electrons and ions. This is predicated on the assumption of a strongly collisional plasma, in which collisions ensure that the velocity distributions of the electrons and ions are Maxwellian and further are both described by the same characteristic temperature (local thermodynamic equilibrium, LTE). This is a classical model, in which the assumption that fluid velocities are non-relativistic is explicitly used to neglect the displacement current in Maxwell's equations. Finally, the fluid is assumed to have an infinite conductivity so that magnetic resistivity can be neglected.

Like the Euler equations, the equations of MHD include hyperbolic conservation laws expressing the conservation of mass, momentum, and energy. Additionally, an evolution equation for the magnetic field is necessary and represents the conservation of magnetic flux. The equations of ideal MHD are[1]:

$$\frac{\partial \rho}{\partial t} + \nabla\cdot(\rho \mathbf{v}) = 0, \tag{1.19}$$

$$\frac{\partial (\rho \mathbf{v})}{\partial t} + \nabla\cdot(\rho \mathbf{v}\mathbf{v} - \mathbf{B}\mathbf{B} + P^{*}\mathbf{I}) = -\rho\nabla\Phi, \tag{1.20}$$

$$\frac{\partial E}{\partial t} + \nabla\cdot[(E + P^{*})\mathbf{v} - \mathbf{B}(\mathbf{B}\cdot\mathbf{v})] = 0, \tag{1.21}$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla\times(\mathbf{v}\times\mathbf{B}) = 0. \tag{1.22}$$

The hydrodynamic variables are augmented by the inclusion of a magnetic vector field ($\mathbf{B}$), and total energy (E). The total energy is the sum of the internal, or thermal, energy of the gas, the kinetic energy, and the magnetic pressure. Formally,

$$E = e + \frac{\rho \mathbf{v}^2}{2} + \frac{\mathbf{B}^2}{2}. \tag{1.23}$$

The sum of the gas and magnetic pressure, the total pressure, is given as $P^{*} = P + P_b$, where $P_b = \mathbf{B}^2/2$. The identity tensor is represented by $\mathbf{I}$, and Φ represents the gravitational potential of the central object. An adiabatic equation of state is used to define a relationship between the internal energy and the pressure as required to close the system of equations, specifically $P = (\gamma - 1)e$. This formulation can be simplified through the use of an isothermal equation of state, in which case equation 1.21 is neglected in favor of a closure relation given by equation 1.5. A useful measure of the partition of energy in a plasma is given by the ratio of gas to magnetic pressure,

$$\beta = \frac{P}{P_b}. \tag{1.24}$$

---

[1] For consistency with Stone et al. (2008) we adopt a system of units where the magnetic permeability is unity.

The formulation of MHD given in equations 1.19 through 1.22 clearly demonstrates the mathematical structure of the coupled system of partial differential equations (PDEs); that the evolutions equations for the mass, momentum, and energy are given by hyperbolic conservation laws. However, this is at the cost of obfuscating many of the important physics. In contrast, one can rewrite equation 1.20 as

$$\rho \frac{\partial \mathbf{v}}{\partial t} + \rho \mathbf{v} \cdot \nabla \mathbf{v} = -\nabla \left(P + \frac{\mathbf{B}^2}{2} \right) + \mathbf{B} \cdot \nabla \mathbf{B} - \rho \nabla \Phi, \tag{1.25}$$

which clarifies the physical forces acting in a magnetized fluid. The individual forces are clearly identified, in the order they appear on the RHS as: the pressure (gas and magnetic) force; magnetic tension, associated with the geometric curvature of the magnetic field; and the gravity of the central source.

The inclusion of the magnetic field in the fluid equations adds new waves to the hyperbolic structure of the system. These waves are the magnetosonic waves, in which pressure (both magnetic and gas) act as a restoring force allowing a wave to propagate orthogonal to the magnetic field lines, and the Alfven wave, in which magnetic tension is the restoring force allowing a wave to propagate parallel to the magnetic field lines. The magnetosonic waves are separated into the slow and fast modes depending on whether the magnetic and gas pressures are acting in opposition or concert, respectively. The speed of the Alfven wave is given by

$$v_A = \frac{|\mathbf{B}|}{\sqrt{\rho}}, \tag{1.26}$$

and for later use we also define the component Alfven wave speed

$$v_{Ax} = \frac{|B_x|}{\sqrt{\rho}}. \tag{1.27}$$

The wave speeds obey the relationship, $v_s \leq v_A \leq v_f$, where v_s and v_f denote the speeds associated with the slow and fast magnetosonic wave.

The interconnected nature of the magnetic field and the fluid in MHD are vital to the dynamics of accretion disks. The magnetic flux through a closed surface that is allowed to evolve freely with the local fluid velocity is conserved. This constraint, often referred to as the "frozen field" condition, implies that in the weak-field regime the magnetic field will be advected with the fluid flow. Phrased another way, this means that fluid elements initially associated with a magnetic field line will remain so indefinitely.

As the radial transport of angular momentum is the primary function of an accretion disk, it is worthwhile to consider the relevant conservation equation directly. The angular momentum, $L = \rho R v_\phi$, evolves according to

$$\frac{\partial L}{\partial t} + \nabla \cdot \mathbf{F_L} = 0, \tag{1.28}$$

where our concern will primarily be on the radial component of the angular momentum flux

$$\mathbf{F}_{\mathbf{L},R} = R[\rho R \Omega_K v_R + \rho v_R v'_\phi - B_r B_\phi], \tag{1.29}$$

which makes use of the perturbation velocity, defined as

$$\mathbf{v}' = \mathbf{v} - v_K \hat{\phi}. \tag{1.30}$$

The first term on the RHS of Equation 1.29 represents the direct advection of angular momentum through accretion. The second ($T_{R\phi} = \rho v_R v'_\phi$) and third ($M_{R\phi} = -B_R B_\phi$) terms, the Reynolds and Maxwell stress terms, illustrate that correlations between the radial and azimuthal turbulent quantities allow the radial transport of angular momentum. The equation also illustrates that turbulence alone is insufficient to guarantee enhanced transport; if the turbulence were isotropic the correlation in these terms would, in a statistical sense, be zero and there would be no net transport. The anisotropic character of turbulence in accretion disks is vital for the consistent transport of angular momentum.

The efficiency of angular momentum transport is defined in a manner reminiscent of the anomalous viscosity models, specifically the stress scaled by the gas pressure. Two measurements that will be used extensively are

$$\alpha = \frac{T_{R\phi} + M_{R\phi}}{P}, \tag{1.31}$$

and

$$\alpha_M = \frac{M_{R\phi}}{P}. \tag{1.32}$$

1.5 The Magnetorotational Instability

The magnetorotational instability has a somewhat curious history; the first hints at an understanding of the instability are attributed to Velikhov (1959) and Chandresekhar (1960) regarding their work on the stability of Couette flow in the presence of a magnetic field. However, the first paper to explicitly connect the stability of a rotating magnetized fluid to angular momentum transport in accretion disks did not appear for several decades (Balbus & Hawley, 1991). The intuitive simplicity of the MRI is in stark contrast to the importance of the problem its discovery solved.

We begin with a simple calculation of the dispersion relation for the MRI, based on Balbus & Hawley (1991). The calculation is done in the cylindrical geometry described above, and the fluid variables are assumed to have perturbations of the form $\delta X e^{i(k_R R + k_z z - \omega t)}$, thus explicitly assuming axisymmetry. The linear perturbation analysis is done in the local sense in which $k_R R$ is assumed sufficiently large so that large-scale radial structure can be neglected. The rotation profile, Ω, is assumed to have only radial dependance and is given by the Keplerian value, and the zeroth order field is assumed to be $B_z(R)$, $B_\phi(R, z)$ and $B_R = 0$. Further, the perturbation analysis is done in the Boussinesq approximation, in which magnetoacoustic waves are excised from the system. This results in the dispersion relation

$$\frac{k^2}{k_z^2}\tilde{\omega}^4 - \left[\kappa^2 + \left(\frac{k_R}{k_z}N_z - N_R\right)^2\right]\tilde{\omega}^2 - 4\Omega^2 k_z^2 v_{Az}^2 = 0, \tag{1.33}$$

where $\tilde{\omega}^2 = \omega^2 - k_z^2 v_{Az}^2$, and $k^2 = k_z^2 + k_R^2$ (not to be confused with the epicyclic frequency, κ^2; equation 1.16). The Brunt-Väisälä (BV) frequency is given by

$$N_x^2 = -\frac{3}{5\rho}\frac{\partial P}{\partial x}\frac{\partial \ln(P\rho^{-5/3})}{\partial x}, \tag{1.34}$$

which represents the frequency associated with buoyant oscillations, and we define $N^2 = N_R^2 + N_z^2$.

The stability is investigated by setting $\omega = 0$, which results in a system that is a quadratic in k_R. Requiring the discriminant of this system to be negative ensures stability by requiring a complex value of k_R for ω^2 to pass through zero. This results in the stability criterion for an individual mode to be

$$k_z^4 v_{Az}^4 + k_z^2 v_{Az}^2\left(N^2 \frac{d\Omega^2}{d\ln R}\right) + N_z^2 \frac{d\Omega^2}{d\ln R} > 0 \quad \text{STABILITY.} \tag{1.35}$$

Working under the assumption that the disk is buoyantly stable ($N_R, N_z > 0$), the condition for all modes to be stable is

$$\frac{d\Omega^2}{dR} \geq 0 \quad \text{STABILITY}. \tag{1.36}$$

The criterion of equation 1.36 is quite different than that of the Rayleigh stability criterion. The former requires an angular velocity increasing with radius, a condition violated in astrophysical disks, while the latter requires an increasing specific angular momentum, which will be satisfied. Thus, this represents the desired linear instability that can drive accretion disk turbulence. Of note, is that the stability criterion associated with the MRI does not depend on the strength of the magnetic field. In the absence of resistivity, the stability of differentially rotating disks is fundamentally altered even in the limit of vanishing magnetic field. While the formal derivation presented above is vital to understand in a quantitative manner the action of the instability, a simpler, physical picture of the instability is also useful.

At its simplest, the MRI is an interchange instability in which the addition of a weak magnetic field to a differentially rotating system drastically changes the stability properties. In words, one considers two fluid elements at an equal distance from the central object tethered by a magnetic field. One of these fluid elements is perturbed to a smaller radius and will be constrained to orbit at the local, faster orbital velocity. The differential orbital velocity between the two fluid elements leads to a shearing of the magnetic field, thus allowing energy to be transferred from the differential rotation to the magnetic field. Note, the assumption of a weak magnetic field is vital, were the field strong it would act against the initial radial perturbation through the magnetic tension force. As the two fluid elements move further apart due to the differential orbital velocity the magnetic field acts to mediate a torque between the two elements. This torque transfers angular momentum from the inner element to the outer element, which allows the inner element to move further inwards thus enhancing the radial displacement and continuing the process.

Before moving on, it is worthwhile to return to the axisymmetric dispersion relation and calculate several important values. Assuming that N^2 and N_z^2 are negligible, the critical vertical wavelength becomes

$$\lambda_{crit} = 2\pi v_{Az} \left| \frac{d\Omega^2}{d\ln R} \right|^{-1/2}, \tag{1.37}$$

in the general case and

$$\lambda_{crit} = \frac{2\pi v_{Az}}{\sqrt{3}\Omega}, \tag{1.38}$$

for a Newtonian disk. Perturbations of spatial lengths larger than this critical wavelength will be unstable, and as the strength of the vertical magnetic field increases this critical wavelength becomes larger. A useful constraint for thin accretion disks is $\lambda_{crit} < 2H$, the requirement that the critical wavelength fit within the "body" of the disk, somewhat arbitrarily defined to reside within a distance H of the midplane. This constraint can be rewritten in terms of β and becomes

$$\beta = \frac{\pi^2}{3} \left(\frac{\lambda_{crit}}{2H} \right)^{-2}, \tag{1.39}$$

which implies that when β drops below approximately 3, the critical wavelength becomes larger than the body of the disk and the instability is suppressed.

In addition to the assumption that the BV frequency is negligible, we now also assume $k_R = 0$ and that the disk is Newtonian. Equation 1.33 can be used to solve for the vertical wavelength with the highest growth rate, denoted λ_{MRI} and given by

$$\lambda_{MRI} = 2\pi\sqrt{\frac{16}{15}}\frac{v_{Az}}{\Omega}, \tag{1.40}$$

with a growth rate of $\omega_{MRI} = 0.75\Omega$. For the specific case of a Newtonian disk, $\lambda_{MRI}/\lambda_{crit} = 4/\sqrt{5}$. We note that the growth rate of the fastest growing mode does not depend on the strength of the magnetic field. The significance of this growth rate is in stark contrast to the nature of the MRI as a weak field instability, ω_{MRI} suggests the potential for a millionfold amplification in only a few orbits. It has been conjectured (Balbus & Hawley, 1992b), that this growth rate represents the most efficient possible growth of an instability powered by differential rotation.

A full calculation of the details of the non-axisymmetric dispersion relation (Balbus & Hawley, 1992a) can be quite cumbersome, and as such will be omitted here. Instead, the salient details will merely be summarized. Due to the differential rotation the calculation must be done in time-dependent variables, thus modifying the wavevector, $\mathbf{k} = (k_R(t), m/R, k_z)$ The time-dependence of the radial component of the wavevector is given as

$$k_R(t) = k_R(0) - mt\frac{d\Omega}{dR}. \tag{1.41}$$

Further, the differential rotation acts upon the zeroth order unperturbed magnetic field to produce azimuthal field from radial in the following manner,

$$B_\phi(t) = B_\phi(0)\left[1 + \frac{B_R}{B_\phi(0)}\frac{d\Omega}{d\ln R}t\right]. \tag{1.42}$$

Though individually both the wavevector and unperturbed magnetic field have time-dependence, the quantity $\mathbf{k}\cdot\mathbf{B}$ is constant, and consequently so is $\mathbf{k}\cdot\mathbf{v}_A$.

The nature of the instability is understood through its movement through the two dimensional parameter space, $(k/k_z, \mathbf{k}\cdot\mathbf{v}_A)$, where $k^2 = k_R^2 + m^2/R^2 + k_z^2$ is simply the magnitude of the wavevector. Because the second component is constant, this movement will be along horizontal lines in the plane. The time dependence of k_R causes the k/k_z term to initially decrease with time until reaching a value of approximately unity when k_R passes through zero and then increase again. During this horizontal oscillation, there will be exponential growth during the period

$$\frac{k}{k_z} < -\frac{d\Omega^2}{d\ln R}. \tag{1.43}$$

The toroidal MRI is thus a transient amplification as opposed to a true instability. The condition for instability is that there is a finite amount of time spent in the unstable region, which leads to the condition for the stability of a given mode

$$(\mathbf{k}\cdot\mathbf{v}_A)^2 > -\frac{d\Omega^2}{d\ln R} \quad \text{STABILITY.} \tag{1.44}$$

For a purely toroidal field and a Keplerian disk, the critical wavenumber becomes $m_{crit} = \sqrt{3}v_K/v_{A\phi}$. Wavenumbers smaller than m_{crit}, *i.e.* those associated with a larger spatial scale, will be unstable. The wavevector associated with the maximum growth rate of the non-axisymmetric instability is given by

$$\frac{(\mathbf{k}\cdot\mathbf{v}_A)}{\Omega} = \sqrt{1 - \frac{\kappa^4}{16\Omega^4}}. \tag{1.45}$$

In a Newtonian disk, $\kappa = \Omega$, and the RHS of equation 1.45 is approximately unity. For a purely toroidal field, this results in the most rapidly growing wavenumber, called the characteristic wavenumber, $m_c = v_K/v_{A\phi}$. The resultant wavelength is

$$\lambda_c = 2\pi\sqrt{\frac{16}{15}}\frac{v_{A\phi}}{\Omega}, \tag{1.46}$$

in analogy with the most unstable vertical wavelength associated with the axisymmetric MRI (equation 1.40). The presence of a purely toroidal field results in significant amplification on a timescale of dozens of orbits in contrast to the sub-dynamical timescale associated with amplification of the vertical-field MRI. However, the presence of even a weak poloidal field can allow amplification of a predominantly toroidal field configuration comparable to the vertical-field MRI.

While analytic theory is useful towards exploring the linear regime in which the MRI operates and some aspects of the early nonlinear behavior (Goodman & Xu, 1994), fully turbulent plasmas are too far into the nonlinear regime for analytic methods to be feasible. Because of this, computational work is at the forefront of efforts to understand the MRI-driven MHD turbulence in accretion disks. While analytic treatments are quite limited, simulations have proven invaluable to understanding the nonlinear resolution of the MRI. Simulations have shown that the MRI results in vigorous, sustained, anisotropic turbulence that drives the redistribution of angular momentum and makes consistent accretion possible. These simulations have also facilitated estimates of α, and by extension the rate of accretion, which is necessary to understand the timescales associated with accretion-driven structure formation and the energetics of ultra-luminous phenomena powered by accretion. Nonlinear processes in accretion disk turbulence, dynamo-driven field amplification for instance, can be studied in the context of simulations and has facilitated the ability to connect these processes to potentially observable astrophysical phenomena, *e.g.* O'Neill et al. (2010).

This reliance on simulation, however, has brought to the fore new concerns.

The measure of accretion, α, has been found to depend on numerical resolution, choice of model, and initial conditions. Further, there remains a significant order of magnitude discrepancy between values of α observed in simulations and observational estimates from astrophysical systems. Understanding the interplay between numerics and physics and constructing simulations from which reliable astrophysical implications can be inferred is a vital concern.

1.6 Outline

Prior to a thorough discussion of the research that will be described in this work, Chapter 2 will complete the background material and will describe the numerical methods and models that are traditionally used to study MRI-driven turbulence in accretion disks. In particular, this chapter will focus on finite-difference and finite-volume algorithms to solve the equations of MHD and the use of the local, or shearing-box approximation, to study the small-scale structure of accretion disk turbulence. Chapter 3 will describe a preliminary exploration of the methods, analysis and goals that will form the cornerstone of the dissertation as a whole. The content of Chapter 3 is largely material that was originally published as Sorathia et al. (2010).

With this preliminary exploration described and the goal of higher resolution simulations motivated, Chapter 4 will describe the algorithm development that allowed a significant improvement in the performance of the simulations whose analysis will form most of the remainder of the dissertation. Chapter 5 focuses on the numerical concerns that arise in simulations of this type and provides metrics to measure the reliability of the results against numerical artifacts. Following this, Chapter 6 presents an extension and refinement of the material presented in Chapter 3. In particular, a focus of this chapter is on the validity of the shearing-box approximation and the assumptions upon which it relies. The bulk of the material in Chapters 5 and 6 have been submitted for publication (Sorathia et al., 2011). Finally, Chapter 7 presents a summary of the results described throughout this dissertation and a discussion of the future work that the results and methodology presented here suggest.

Chapter 2

Simulating Astrophysical Accretion Disks

Computational simulations of accretion disk turbulence rely on two factors: an algorithm to solve the equations of MHD and a set of initial and boundary conditions with which to initialize the algorithm, referred to as a model. The algorithms to be discussed here are those used by the software packages `Zeus`, predominantly based on a finite-differencing approach, and `Athena`, utilizing a finite-volume algorithm. This is not to suggest that these approaches are representative of all the algorithms used in simulations of accretion disk turbulence, though they do represent two common astrophysical MHD algorithms, merely that these are the software packages that are used in the simulations that will be presented here.

The models that will be considered here are broadly separable into local simulations, in which a small Cartesian patch of the Keplerian flow is modeled, and global simulations, in which the full geometric extent of the disk is modeled in cylindrical or spherical coordinates. The assumptions of a geometrically thin disk will be used, in contrast to alternative models of thicker accretion tori. The formulation and salient results of local models will be summarized here, with a description of global models deferred to the following chapter and discussed in the context of the particular global models used in the simulations to be presented.

2.1 Numerical Magnetohydrodynamics

While the purely mathematical treatment of algorithms often focuses solely on order of accuracy, the use of algorithms in computational physics requires methods that identify and respect quantities of physical importance. As an example, a high-order method that poorly conserves energy would be of little practical use in many astrophysical applications. Further, the equations of fluid dynamics admit non-classical, or weak, solutions. These solutions may exhibit discontinuities or other points at which differentiability fails but solve the equations of MHD in an integral sense.

Prior to a discussion of the general numerics involved in solving the equations of MHD, it is worth introducing the often unwieldy notation associated with these methods. The algorithms that will be discussed here are grid-based, in the sense that they rely on a decomposition of the spatial domain of the equations into a collection of grid cells, denoted $\mathcal{V}_{i,j,k}$, centered about spatial nodes (x_i, y_j, z_k) (in Cartesian geometry) for a three-dimensional spatial domain. The grid spacing, referred to as the resolution, is denoted Δx for each given coordinate direction. In general this spacing may itself have spatial dependence. Initial conditions are translated to the grid by defining for each cell a representative value of each variable in one of two ways. For `Zeus`, the cell representative merely represents the pointwise value of the quantity at the spatial center of the cell, or in the case of vectorial quantities at the interface between cells. The finite volume algorithm of `Athena`, by contrast, defines the representative value to be the volume average of the quantity over the entire cell, or in the case of the magnetic field the area average over a specific surface. For smooth flows, this distinction is negligible, however the cell averaging method is able to naturally handle discontinuous flows. The function of the algorithm is to update the grid variables at a given time, t^n, to a new time, $t^{n+1} = t^n + \Delta t$, where the timestep Δt is constrained by the particular method and is discussed in more detail below. For a flow attribute, q, the cell-centered representative value of grid cell (i, j, k) at time t^n will be denoted $q^n_{i,j,k}$. Cell interfaces are denoted by integer indices displaced by 1/2, *e.g.* the magnetic field at an x-interface would be given as $B_{x,(i+1/2,j,k)}$.

In the MHD equations, the evolution equations for the conserved hydrodynamic variables (mass, momentum, and energy) take the form of a system of hyperbolic conservation laws, each of which can be written as

$$\frac{\partial q}{\partial t} + \nabla \cdot \mathbf{F}_q = 0, \tag{2.1}$$

where q and $\mathbf{F}_q$ represent the conserved quantity and its flux, respectively. Integrating over a fiducial cell denoted $\mathcal{V}$ results in

$$\frac{d}{dt}\int_{\mathcal{V}} q \; dV = -\oint_{\partial\mathcal{V}} \mathbf{F}_q \cdot d\mathbf{A} = -\sum_{i=1}^{6}\int_{\mathcal{S}_i} \mathbf{F}_q \cdot d\mathbf{A}, \tag{2.2}$$

where $\mathcal{S}_i$ represents the various cell faces and $d\mathbf{A}$ representing the differential surface area element with outwards facing normal . The formulation given in equation 2.2 represents the fundamental physical importance of the flux through the cell interface of a conserved quantity. A simple differencing scheme to solve equation 2.1 would allow variation in the total volume integral of the quantity q related to the order of the method. In contrast, a flux update, in which the conserved quantity is updated utilizing an appropriately temporally- and spatially-averaged interface flux, will preserve the total volume integral of the conserved quantity to machine precision. This is because the flux at each interface is used to update two cells and thus the flux from one cell into an adjacent cell is precisely tracked. In one dimension, a flux update of the quantity q_i^n takes the form

$$q_i^{n+1} = q_i^n - \frac{\Delta t}{\Delta x}\left[f_{i+1/2}^{n+1/2} - f_{i-1/2}^{n+1/2}\right], \tag{2.3}$$

where the interface fluxes are given by

$$f_{i-1/2}^{n+1/2} = \frac{1}{\Delta t}\int_{t^n}^{t^{n+1}} F_q(x_{i-1/2}, t)dt. \tag{2.4}$$

The induction equation (equation 1.22) is not a hyperbolic conservation law, but does represent an equally important physical principle, namely the solenoidal constraint. The induction equation ensures that the constraint, $\nabla\cdot\mathbf{B} = 0$, will be maintained so long as the initial magnetic field is itself divergence free. Again, a simple differencing treatment of the induction equation will not ensure the solenoidal constraint. Integrating the solenoidal constraint over a fiducial cell, $\mathcal{V}$, yields

$$\oint_{\partial\mathcal{V}} \mathbf{B}\cdot d\mathbf{A} = \sum_{i=1}^{6}\Phi_i = 0, \tag{2.5}$$

where $\Phi_i = \int_{\mathcal{S}_i}\mathbf{B}\cdot d\mathbf{A}$, represents the magnetic flux piercing a given cell interface. Equation 2.5 encapsulates the lack of magnetic source points, or monopoles, and thus the total magnetic flux through a given cell will be zero. Preserving this magnetic flux formulation of the solenoidal constraint is at the heart of the method of constrained transport (CT; Evans & Hawley 1988).

The approach of Evans & Hawley (1988) utilizes the observation that there is an analogous quantity to the interface-flux by which to update the magnetic field. This quantity is the electromotive force (EMF), defined by

$$\epsilon = \mathbf{v} \times \mathbf{B}. \tag{2.6}$$

Integrating the induction equation over a face of a given grid cell yields

$$\frac{d\Phi_i}{dt} = \oint_{\partial \mathcal{S}_i} \epsilon \cdot d\mathbf{l} = \sum_{j=1}^{4} \int_{\mathcal{E}_{i,j}} \epsilon \cdot d\mathbf{l}, \tag{2.7}$$

where $\mathcal{E}_{i,j}$ represents the edges of the cell interface $\mathcal{S}_i$.

As each edge is represented in the boundary of two cell surfaces, each with different orientation to the outward normal, pairwise cancellation gives

$$\frac{d}{dt} \sum_{i=1}^{6} \Phi_i = \sum_{i=1}^{6} \sum_{j=1}^{4} \int_{\mathcal{E}_{i,j}} \epsilon \cdot d\mathbf{l} = 0. \tag{2.8}$$

Algorithms utilizing CT often rely on a staggered mesh, in which the magnetic field quantities are spatially centered at the corresponding cell interface, *i.e.* B_x would be localized at the x-interface. Further, in the staggered mesh formalism the magnetic flux through the x-interface is approximately $B_x A_x$, where A_x is simply the area of the x-interface. Thus, equation 2.7 can be reformulated utilizing appropriately spatially- and temporally-averaged values of the EMF to construct a means of updating the magnetic field that precisely conserves the solenoidal constraint. Generally, this update uses the time-centered EMF, $\epsilon^{n+1/2}$, and utilizes a four-point update in which each directional component of the line integral is approximated by the value of $\epsilon^{n+1/2}$ at the edge center, as in `Zeus`, or an approximation to the time- and edge-integrated EMF, as in `Athena`.

Formulating algorithms around flux and EMF updates allows the creation of methods that naturally satisfy the fundamental physical principles encapsulated in the equations of MHD. Utilizing these types of updates will be of paramount importance when discussing the code development undertaken as part of the research presented here. Calculating the quantities necessary for these updates, however, presents challenges in and of itself.

The algorithms discussed here, explicit time-advancing schemes, require as a condition for stability that the timestep is limited by the time it takes for any signal to cross a grid cell. This is a consequence of the CFL condition which requires that the numerical domain of dependence, one cell in each direction, must contain the physical domain of dependence. For the MHD equations, the fastest wave speed, c_f, is that of the fast magnetosonic wave and the fastest signal is the velocity enhanced fast magnetosonic wave, $|v| + c_f$. The maximum timestep allowable given signals in a specific coordinate direction is (minimizing over all grid cells),

$$\Delta t_x = C_0 \frac{\Delta x}{|v_x| + c_{fx}}, \tag{2.9}$$

where C_0 is the Courant number, required to be less than $1/2$ for three dimensional simulations, and c_{fx} is the fast magnetosonic wave in the x-direction. The overall timestep is then

$$\Delta t = \min\left[\Delta t_x, \Delta t_y, \Delta t_z\right]. \tag{2.10}$$

For accretion disk simulations, in which the Keplerian rotation is highly supersonic and thus the dominant signal speed, the timestep will approximately be given by the Keplerian timestep defined as

$$\Delta t_K = C_0 \frac{\Delta\phi}{\Omega_K}, \tag{2.11}$$

where $\Delta\phi$ is simply the azimuthal grid spacing.

2.1.1 The Finite-Difference Algorithm of `Zeus`

This discussion of `Zeus` (Stone & Norman, 1992; Hayes et al., 2006) will refer to its underlying algorithm as opposed to one of the many individual implementations. The algorithm is a grid-based method that relies heavily on finite differencing, but with several extensions to improve the numerics of the method. The `Zeus` method uses a staggered mesh formalism in which scalar quantities are spatially localized to the center of the cell, and vectorial quantities are spatially centered at cell interfaces. This centering simplifies the use of CT to ensure maintenance of the solenoidal constraint on the magnetic field.

Finite difference approaches to the equations of fluid dynamics exhibit stability problems when faced with discontinuous flows; this necessitates the use of an artificial viscosity that acts to "smear" discontinuities over several grid cells so that they are resolvable. Additionally, utilizing differencing to treat the source terms precludes a full flux-update to the conserved quantities, which results in a method that is not fully conservative. The temporal accuracy of `Zeus` is limited to first order, however the relative simplicity of the method is conducive to the implementation of new physics.

The algorithm `Zeus` employs is operator split, and the time evolution of the equations is done in three steps: a source step, in which source terms of the momentum and energy equations are included; a separate treatment in which the magnetic tension force is included and the magnetic field is evolved using a CT update, denoted the method of characteristics-constrained transport (MOCCT;Hawley & Stone 1995); and finally, a transport step in which the advection of the fluid variables due to local flow is treated. For simplicity, the components of the algorithm will be discussed in the context of isothermal MHD although the `Zeus`-based simulations presented in Chapter 3 will employ an adiabatic equation of state.

The source step solves the equation

$$\rho\frac{\partial \mathbf{v}}{\partial t} = -\nabla\left(P + \frac{\mathbf{B}^2}{2}\right) - \nabla\cdot\mathbf{Q} - \rho\nabla\Phi, \tag{2.12}$$

where $\mathbf{Q}$ is the von Neumann-Richtmyer artificial viscosity tensor. The velocity is evolved using explicit second-order centered spatial differencing with the unadvanced values of the density and magnetic field. The inclusion of the pressure terms incorporates the magnetosonic waves, however the absence of the magnetic tension force implies that the Alfven wave has not yet been taken into account.

The MOCCT step includes the magnetic tension force, thus solving the equation

$$\rho\frac{\partial \mathbf{v}}{\partial t} = \mathbf{B}\cdot\nabla\mathbf{B}, \tag{2.13}$$

and evolves the magnetic field utilizing an EMF update. This is an operator split method in that the magnetic tension force is applied to the value of the velocity after the source step is complete. Incorporating the Alfven waves involves solving a series of one-dimensional characteristic equations, derived from the propagation of Alfven waves in an incompressible fluid; as the effects of compressibility having already been taken into account in the source step. Updating the magnetic field requires the time- and edge-centered EMFs, which in turn require the relevant velocity and magnetic field components with the same centering. For instance, $\epsilon_z^{n+1/2} = (v_x B_y - v_y B_x)^{n+1/2}$. Calculating the time-centered velocity and magnetic field components is accomplished by tracing the characteristics backwards in time to the "footpoints", identifying a spatial point from which information will propagate. The calculation of ϵ_z is illustrated in Figure 2.1. The necessary constituent information required for calculating the EMF is then interpolated at the footpoints of the characteristics, and the magnetic field is evolved consistent with CT. Finally, the magnetic field values calculated from the footpoints are used to include the magnetic tension force in the updated velocity by differencing equation 2.13.

The final step of `Zeus` algorithm is the transport step which incorporates fluid advection and solves the equations

$$\frac{d}{dt}\int_{\mathcal{V}} \rho \, dV = -\oint_{\partial\mathcal{V}} \rho \mathbf{v} \cdot d\mathbf{A} \tag{2.14}$$

$$\frac{d}{dt}\int_{\mathcal{V}} \rho \mathbf{v} \, dV = -\oint_{\partial\mathcal{V}} \rho \mathbf{v}\mathbf{v} \cdot d\mathbf{A}, \tag{2.15}$$

written in integral form. The transport step is directionally-split and involves a series of one-dimensional flux updates for each coordinate direction, referred to as directional sweeps. As the updates are of the form of equation 2.3, the only issue is to calculate the fluxes. For clarity, this will be discussed in the context of a sweep in the x-direction, however the general procedure does not depend on the direction chosen, save for appropriate scaling in curvilinear geometries. The first step in the procedure is a reconstruction of the conserved quantity (mass or momentum), in practice this is a linear reconstruction of the one-dimensional dependence of the conserved quantity in each cell. An arbitrary linear reconstruction would be problematic as the reconstructed profile may introduce new extrema to the profile which may induce growing oscillations that would lead to an instability of the method. To help ensure stability, the reconstruction is required to be total variation diminishing (TVD), or formally the constraint that $TV(q^{n+1}) \leq TV(q^n)$, where

$$TV(q) = \sum_i |q_i - q_{i-1}|. \tag{2.16}$$

The reconstruction used in `Zeus` is piecewise linear, thus second order, and satisfies the TVD constraint through slope-limiting.

With this reconstruction in place, as demonstrated in Figure 2.2, the next step is to calculate the flux through the interface $x_{i-1/2}$. The flux is upwinded, in that the reconstruction used, either from the left or right of the interface, depends on the direction of v_x at the $x_{i-1/2}$ interface. A positive velocity requires use of the reconstruction in cell $\mathcal{V}_{i-1}$, whereas a negative value will use the reconstruction of cell $\mathcal{V}_i$. The numerical flux at the interface is then calculated by approximating the area integral at the interface of the appropriate reconstruction moving at a velocity v_x.

While the formulation of the source step precludes a conservative flux update of the momentum, the absence of source terms in the continuity equation implies that this method will conserve mass. The completion of the directional sweeps in the transport step finalizes the evolution of the algorithm from one timestep

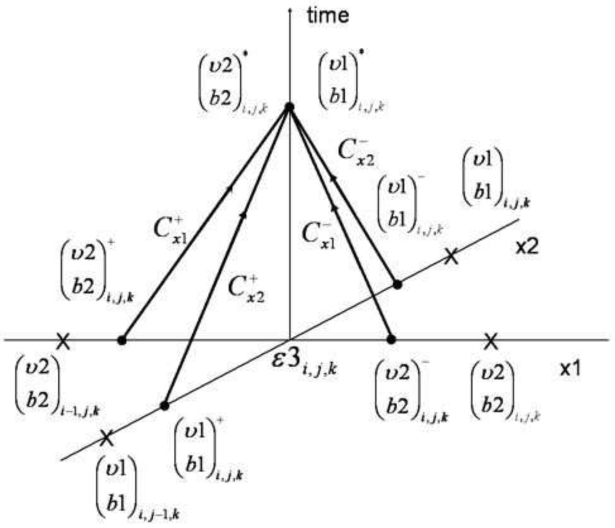

Figure 2.1: Schematic of characteristic tracing used in MOCCT step. Figure from Hayes et al. (2006). Characteristics are traced backwards from $\epsilon_z^{n+1/2}$ to footpoints at t^n.

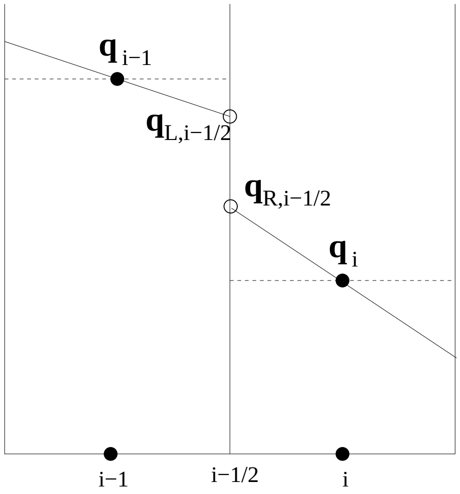

Figure 2.2: Linear Reconstruction in one dimension.

to the next. The simplicity, robustness, and flexibility of `Zeus` has made it an indispensable tool for modeling astrophysical plasmas for many years. However, advances in computational fluid dynamics have rendered it somewhat outdated.

2.1.2 The Finite-Volume Algorithm of `Athena`

The `Athena` algorithm (Stone et al., 2008) is a second-order finite-volume code based on the method described in Gardiner & Stone (2005) and Gardiner & Stone (2008). The algorithm represents an important step forward in astrophysical MHD as it was able to fuse the finite volume methods at the cutting edge of hydrodynamic work with CT to evolve the magnetic field consistent with the solenoidal constraint. Specifically, `Athena` augments the piecewise-parabolic method (PPM; Colella & Woodward 1984) to solve the equations of MHD.

The `Athena`-based simulations to be presented are all performed in the context of isothermal MHD, and so the discussion of the `Athena` algorithm will be limited to this context. The fluid variables are, depending on the context within the algorithm, formulated either in conservative form, $\mathbf{U} = \{\rho, \rho\mathbf{v}, \mathbf{B}\}$, or primitive form, $\mathbf{W} = \{\rho, \mathbf{v}, \mathbf{B}\}$. The variables are stored in conservative form, where the mass and momentum are volume-averaged, and the magnetic field is stored as both a volume-averaged quantity and as an area-averaged quantity at the interface as demonstrated in Figure 2.3. It is the area-averaged form of the magnetic field that is evolved while the volume-averaged magnetic field calculated using second-order averaging of the interface states. Using this notation, the isothermal MHD equations in Cartesian geometry can be written

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} + \frac{\partial \mathbf{G}}{\partial y} + \frac{\partial \mathbf{H}}{\partial z} = 0, \tag{2.17}$$

where

$$\mathbf{F} = \begin{bmatrix} \rho v_x \\ \rho v_x^2 + P^* - B_x^2 \\ \rho v_x v_y - B_x B_y \\ \rho v_x v_z - B_x B_z \\ 0 \\ B_y v_x - B_x v_y \\ B_z v_x - B_x v_z \end{bmatrix}, \tag{2.18}$$

and the fluxes along the other axes, **G** and **H**, are defined similarly with their spatial localizations shown in Figure 2.3.

As in PPM, `Athena` relies on a Riemann solver to compute interface fluxes. The Riemann problem in the most general form refers to a solution of the one-dimensional equations given piecewise constant initial data about an interface, *e.g.* in the x-direction the initial conditions take the form

$$\mathbf{q} = \left\{ \begin{array}{ll} \mathbf{q}_L & : x < x_{i+1/2}, \\ \mathbf{q}_R & : x > x_{i+1/2}, \end{array} \right\}. \tag{2.19}$$

The hyperbolic nature of the MHD equations yields a solution of the type demonstrated in Figure 2.4. The solution over all of spacetime is characterized by the propagation of a series of waves away from the interface that decompose spacetime into disjoint sets upon which the solution, given by the state vector $\mathbf{q}$, is generally constant. While hyperbolic the equations of MHD are not strictly so, and thus certain waves may be degenerate resulting in fewer intermediate states. A full solution of the Riemann problem yields more information than is actually required for the purposes of the `Athena` algorithm. All that is required for a conservative flux update of the variables is the behavior of the solution at the cell interface from which the flux can be calculated. Because of this, the use of an approximate Riemann solver is preferred. An approximate Riemann solver, given the left and right interface states and the interface magnetic field component orthogonal to the interface, returns the interface flux which can be used to perform a conservative update. The approximate Riemann solvers are less costly due to the reduced number of waves, and consequently intermediate states, considered. While the full details will be omitted, the simulations presented here use the HLLD solver (Miyoshi & Kusano, 2005), which utilizes four intermediate states bounded by the contact, fast magnetosonic, and Alfven waves. The heart of the `Athena` algorithm is the calculation of time-averaged left and right interface states ($\mathbf{q}^{n+1/2}_{Lx,i+1/2,j}$ and $\mathbf{q}^{n+1/2}_{Rx,i+1/2,j}$) as well as the parallel component of the magnetic field ($B^{n+1/2}_{x,i+1/2,j}$). With these values in place, the Riemann solver can be used to calculate fluxes with which to conservatively update the fluid variables. Further, the fluxes returned by the Riemann solver can be used to construct a CT update for the magnetic field.

Prior to a discussion of the calculation of the time-centered values to be passed to the Riemann solver, it is worth noting the manner in which the fluxes returned by the solver can be used to update the magnetic field. Note that the flux in the x-

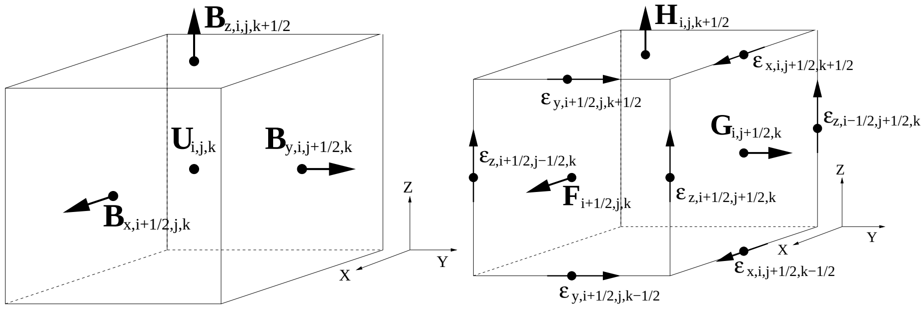

Figure 2.3: (*Left*) Centering of conservative variables and interface magnetic fields. (*Right*) Centering of interface fluxes computed from the Riemann solver. Figures taken from Stone et al. (2008)

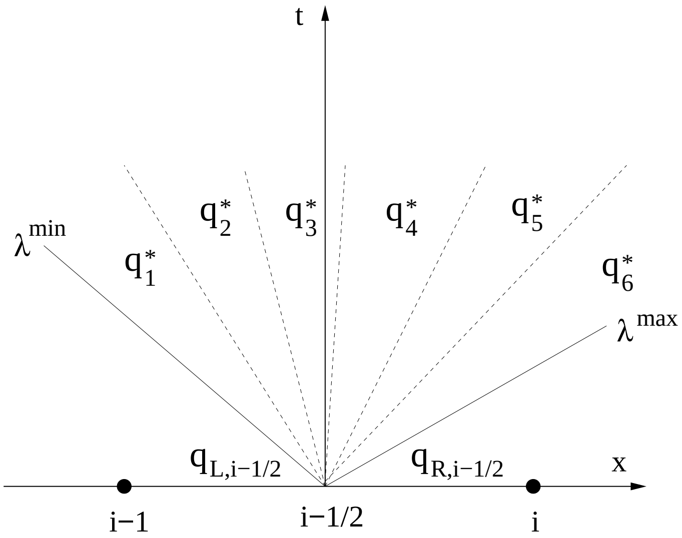

Figure 2.4: Illustration of the solution of the Riemann problem. Waves decompose spacetime into disjoint regions upon which the state vector, $\mathbf{q}$, is constant. Calculation of the interface flux requires only the the region which contains $x = x_{i-1/2}$ and corresponds to state $\mathbf{q}_3^*$ in the above diagram. Figure taken from Stone et al. (2008).

direction given by equation 2.18 includes ϵ_z and ϵ_y which are necessary to update B_x at the interface. However, these components of the EMFs are centered incorrectly, at the interface as opposed to the edges. Converting the EMFs to the proper edge centering is referred to in the algorithm as "integrating to the corner." This is accomplished through a gradient approximation using, in the case of ϵ_z, values from the fluxes in the x and y direction. This results in the formula

$$\begin{aligned}
\epsilon_{z,i-1/2,j-1/2,k} &= \frac{1}{4}\left(\epsilon_{z,i-1/2,j,k} + \epsilon_{z,i-1/2,j-1,k} + \epsilon_{z,i,j-1/2,k} + \epsilon_{z,i-1,j-1/2,k}\right) \\
&+ \frac{\delta y}{8}\left(\left(\frac{\partial \epsilon_z}{\partial y}\right)_{i-1/2,j-1/4,k} - \left(\frac{\partial \epsilon_z}{\partial y}\right)_{i-1/2,j-3/4,k}\right) \\
&+ \frac{\delta x}{8}\left(\left(\frac{\partial \epsilon_z}{\partial x}\right)_{i-1/4,j-1/2,k} - \left(\frac{\partial \epsilon_z}{\partial x}\right)_{i-3/4,j-1/2,k}\right) . \qquad (2.20)
\end{aligned}$$

The derivatives in equation 2.20 are "upwinded" in that they are approximated using backwards or forwards numerics based on the sign of the mass flux at the interface. Analogous equations for the other components of the EMFs allow the fluxes returned by the Riemann solver to be used to accurately update the interface magnetic fields consistent with CT.

Calculating the time-centered interface states is done in a manner similar to the corner-transport upwind (CTU;Colella 1990) method. For a hydrodynamic system, the method begins with a series of one-dimensional reconstructions to compute left and right states ($\mathbf{q}_L$ and $\mathbf{q}_R$) at each interface. This is similar to the reconstruction done in the transport step in the `Zeus` algorithm, demonstrated in Figure 2.2, however it is modified to utilize a slope-limited piecewise-parabolic reconstruction resulting in a third-order spatially accurate method. Characteristic tracing in one dimension, described below, is used to advance the left and right states by a timestep of $\Delta t/2$ along the interface direction. For instance, the left and right states about an x-interface ($\mathbf{q}_{Lx}$ and $\mathbf{q}_{Rx}$) are advanced by $\Delta t/2$ in the x-direction to compute the evolved states, $\mathbf{q}^*_{Lx}$ and $\mathbf{q}^*_{Rx}$. These "star" states are then passed to the Riemann solver to calculate the fluxes (f^*, g^*, h^*) in the x-, y-, and z-directions respectively. The transverse fluxes are used to evolve the star-states formally to $t_{n+1/2}$, *i.e.* $\mathbf{q}^*_{Lx}$ is evolved to $\mathbf{q}^{n+1/2}_{Lx}$ using g^* and h^*. The final time-centered left and right states are again passed to the Riemann solver to calculate temporally second-order accurate fluxes, $f^{n+1/2}$ *etc.* , with which to update the cell-centered conserved quantities in the manner of equation 2.3 and complete the hydrodynamic algorithm.

Extending the CTU method to MHD is complicated by the multidimensional nature of the solenoidal constraint, and by extension, the induction equation. The directional sweeps necessary for the characteristic tracing allow magnetic field components to evolve independently, whereas physically the solenoidal constraint requires multidimensional coupling. One of the primary innovations of the `Athena` algorithm is the formulation of magnetic source terms to include in the characteristic tracing step to incorporate this information. Defining $\mathbf{W}^x$ to be the set of primitive variables with B_x omitted allows the simple definition of the characteristic tracing in the x-direction to correspond to a solution of

$$\frac{\partial \mathbf{W}^x}{\partial t} + \mathbf{A}\frac{\partial \mathbf{W}^x}{\partial x} = \sigma, \tag{2.21}$$

where $\mathbf{A}$ is the hyperbolic matrix linearized about the state $\mathbf{W}^x$, and σ represents the source terms associated with the x-direction. These source terms will include the MHD source terms and potentially gravitational and, in the case of curvilinear geometry, geometric source terms. The hyperbolic nature of the system implies that the wave matrix $\mathbf{A}$ is diagonalizable, *i.e.* there exists $\mathbf{L}$ such that $\mathbf{L}\mathbf{A}\mathbf{L}^{-\mathbf{1}} = \mathbf{\Lambda}$, where $\mathbf{\Lambda}$ is a diagonal matrix whose entries correspond to, potentially non-distinct, wavespeeds. Projecting the primitive variables onto the characteristic variables, $\mathbf{a}^{\mathbf{x}} = \mathbf{L}\mathbf{W}^{\mathbf{x}}$, decouples the equations and results in a collection of linear advection equations which are advanced by $\Delta t/2$, including only the information from waves that reach the interface in the timeframe considered, which are then projected back onto the primitive variables.

This allows the computation of the star-states at the left and right of each interface and their associated fluxes from the Riemann solver as described above. Next, the left and right states are advanced to the half-timestep through the inclusion of transverse flux gradients, *e.g.* $\mathbf{q}_{Lx}$ is evolved according to

$$\begin{aligned}
\mathbf{q}^{n+1/2}_{Lx,i-1/2,j,k} &= \mathbf{q}^{*}_{Lx,i-1/2,j,k} - \frac{\Delta t}{2\Delta y}\left[\mathbf{g}^{*}_{i+1,j+1/2,k} - \mathbf{g}^{*}_{i+1,j-1/2,k}\right] \\
&- \frac{\Delta t}{2\Delta z}\left[\mathbf{h}^{*}_{i+1,j,k+1/2} - \mathbf{h}^{*}_{i+1,j,k-1/2}\right] + \frac{\Delta t}{2}\sigma_{yz}, \qquad (2.22)
\end{aligned}$$

where σ_{yz} refers to source terms associated with the y- and z-direction, and the remaining interface states are updated analogously. These time-centered states along with the time-centered magnetic field, computed by integrating to the corner with the intermediate fluxes, are again passed to the Riemann solver to compute second-order accurate fluxes, $[f, g, h]^{n+1/2}$. The final update of the conserved quantities then takes the form of the multidimensional analog of equation 2.3

$$\mathbf{U}^{n+1}_{i,j,k} = \mathbf{U}^{n}_{i,j,k} - \frac{\Delta t}{\Delta x}\left[\mathbf{f}^{n+1/2}_{i+1/2,j,k} - \mathbf{f}^{n+1/2}_{i-1/2,j,k}\right] - \frac{\Delta t}{\Delta y}\left[\mathbf{g}^{n+1/2}_{i,j+1/2,k} - \mathbf{g}^{n+1/2}_{i,j-1/2,k}\right]$$

$$- \quad \frac{\Delta t}{\Delta z}\left[\mathbf{h}^{n+1/2}_{i,j,k+1/2} - \mathbf{h}^{n+1/2}_{i,j,k-1/2}\right] + \mathbf{S}^{n+1/2}_{i,j,k}, \tag{2.23}$$

where $\mathbf{S}^{n+1/2}$ are the time-centered and volume-averaged gravitational, and if necessary, geometric source terms. Time-centering is necessary to ensure that the source terms are consistent with the finite-volume formulation.

The global simulations that will be presented that using `Athena` algorithm also make use of the cylindrical geometry extension described by Skinner & Ostriker (2010). The extension to cylindrical geometry omits modification of the Riemann solver itself at the cost of introducing source terms to the MHD equations. This requires modifications to the conservation of radial and azimuthal momentum given by

$$\frac{\partial(\rho v_R)}{\partial t} + \frac{1}{R}\frac{\partial(R\mathbf{T}_{RR})}{\partial R} + \frac{1}{R}\frac{\partial \mathbf{T}_{\phi R}}{\partial \phi} + \frac{\partial \mathbf{T}_{zR}}{\partial z} = \frac{\rho v_\phi^2 - B_\phi^2 + P^*}{R} \tag{2.24}$$

$$\frac{\partial(\rho v_\phi)}{\partial t} + \frac{1}{R^2}\frac{\partial(R^2\mathbf{T}_{R\phi})}{\partial R} + \frac{1}{R}\frac{\partial \mathbf{T}_{\phi\phi}}{\partial \phi} + \frac{\partial \mathbf{T}_{z\phi}}{\partial z} = 0 \tag{2.25}$$

where $\mathbf{T} = \rho\mathbf{v}\mathbf{v} - \mathbf{B}\mathbf{B} + P^*\mathbf{I}$ is the full stress tensor. The additional source term in equation 2.24 and modified form of the scaling factor in the radial tensor derivative in equation 2.25 will be of particular importance when describing the extensions to the cylindrical `Athena` algorithm. Additionally, the spatial dependence of cell volume necessitate the inclusion of scale factors at various points in the algorithm. The details of these will be omitted save for when they are relevant to the algorithmic extensions that will be detailed in Chapter 4.

2.2 Computational Models of Accretion Disks

With algorithms in place to solve the equations of MHD, the next concerns are models and initial conditions appropriate to simulating accretion disks. Modeling accretion disks can be quite challenging due to the multiscale nature of the problem. The three dominant spatial scales of an accretion disk are, in decreasing order, the global radial lengthscale, the vertical scale height, and the dissipation scale, λ_D. These spatial scales are each separated by orders of magnitude. The disparate temporal scales serve to only exacerbate the problem with the dynamical and viscous timescale also separated by several orders of magnitude. This, combined with the computational limitations of even the most sophisticated supercomputers, has led to a bifurcated approach towards modeling accretion disks. Local models focus on resolving the scale height, and more recent applications also resolve the dissipation scale. Global models, by contrast, focus on modeling the radial and vertical scales albeit at generally lower resolution.

There is a natural inclination to simulate accretion disks in two dimensions due the vastly reduced computational expense, but this must be avoided. Two dimensional turbulence exhibits a fundamentally different character than three dimensional turbulence. For instance, in inviscid, incompressible two dimensional hydrodynamic turbulence the enstrophy and energy are conserved quantities. A cascade in one quantity, in which power is driven to smaller spatial scales, necessitates an inverse cascade in the other, in which power is driven to larger spatial scales. Three dimensional turbulence is characterized by three conserved quantities, and as such there is no analogous simple relationship between the spectral structure of the conserved quantities. Further, the likely existence of an active dynamo in accretion disks makes a two dimensional treatment unphysical. Astrophysical disks are observed to accrete on timescales much larger than the resistive timescales, the time it takes for a magnetic field to be resistively dissipated. This suggests the existence of a dynamo, the process by which the turbulent fluctuations of the fluid act in a statistically self-consistent manner to amplify the magnetic field and compensate for resistive dissipation. Numerical evidence supports the existence of an accretion disk dynamo. The anti-dynamo theorem of Cowling states that a dynamo can not exist without symmetry breaking all three spatial dimensions.

A simplification that can be made is the use of unstratified models, in which the vertical component of gravity is omitted from the equations. This approximation suppresses the vertical structure of the disk which often presents numerical challenges. Regions away from the midplane take on the characteristics of a corona, a highly magnetized and low density structure. Strongly magnetized regions in the midplane of the disk are buoyantly unstable and will tend to float away thus forming the corona. The corona is home to supersonic Alfven waves which can act to constrain the allowable timestep. Removing the vertical dependance of the gravitational force, in essence replacing the spherical radius with the cylindrical, has the physical interpretation of modeling just the midplane of a realistic disk. Unstratified models have the added benefit of allowing a simple periodic treatment of the vertical boundary conditions. Neglecting the vertical component of gravity removes the vertical density and pressure structure of the disk, however it does not reduce the importance of the vertical scale height. Even in the absence of vertical structure, the effective scale height remains an important ratio representing the relative importance of thermal energy to rotational energy.

2.2.1 Local Models of Disk Turbulence

The necessity of three dimensional simulations presents a significant computational difficulty; the simultaneous resolution of the vertical and radial spatial scales. To alleviate this constraint local, or shearing box, models are often used. Local models simulate a Cartesian domain centered about a fiducial point, R_0, corotating with the disk at an orbital velocity $\Omega_0 = \Omega_K(R_0)$. Under the assumption that the radial extent of the simulation domain, L_x, is small compared to R_0, the radial dependence of the Keplerian velocity is linearized about the fiducial point. The limited radial extent significantly reduces the volume of the computational domain allowing significantly higher resolutions than would be otherwise feasible. The Cartesian domain is related the standard cylindrical geometry through the transformation $(x, y, z) = (R - R_0, R_0\phi, z)$. The Keplerian velocity in its linearized form is given by

$$v_K = -q\Omega_0 x, \tag{2.26}$$

where the shear parameter, q, is defined as

$$q = -\frac{d\ln \Omega_K}{d\ln R}. \tag{2.27}$$

The shear parameter is a measure of the strength of the differential rotation, and for disks in a Newtonian gravitational potential $q = 3/2$. While in general, the shear parameter is a function of radius and height from the midplane, in the context of local models it is taken to be a constant defined by its value at the fiducial point.

Simulating a rotating domain requires the inclusion of source terms in the momentum and energy equations accounting for the centrifugal and Coriolis force, while the mass and induction equation remain unchanged. The modified momentum and energy equations are given as

$$\frac{\partial(\rho\mathbf{v})}{\partial t} + \nabla\cdot(\rho\mathbf{v}\mathbf{v} - \mathbf{B}\mathbf{B} + P^{*}\mathbf{I}) = \rho\Omega_0^2(2qx\hat{\mathbf{x}} - z\hat{\mathbf{z}}) - 2\Omega_0\hat{\mathbf{z}} \times \rho\mathbf{v}, \tag{2.28}$$

$$\frac{\partial E}{\partial t} + \nabla\cdot[(E + P^{*})\mathbf{v} - \mathbf{B}(\mathbf{B}\cdot\mathbf{v})] = \Omega_0^2\rho\mathbf{v} \times (2qx\hat{\mathbf{x}} - z\hat{\mathbf{z}}). \tag{2.29}$$

Finally, boundary conditions appropriate for the rotating Cartesian patch must be defined. The azimuthal coordinate naturally lends itself to periodic boundary conditions. While less physical, the vertical coordinate is often also treated using periodic boundary conditions. Treatment of the radial boundaries is more difficult, and rely on the concept of shearing periodicity. For a fluid variable, X, not including the azimuthal velocity, v_y, the boundary conditions are defined as

$$X(x, y, z) = X(x + L_x, y - q\Omega_0 L_x t, z) \tag{2.30}$$

$$v_y = v_y(x + L_x, y - q\Omega_0 L_x t, z) + q\Omega_0 L_x. \tag{2.31}$$

At the beginning of the simulation, each ghost cell is mapped to its periodic counterpart, as the simulation evolves the time-dependent boundary condition maintains this initial mapping by taking into account the differential rotation over the domain. The azimuthal velocity must be further corrected to obtain the appropriate local orbital velocity. The imposition of these boundary conditions results in a “closed” simulation, in which matter accreting through the inner radial domain will come back into the simulation through the outer radial domain. Additionally, this constrains the evolution of the magnetic flux through the domain. In the case of an initial magnetic field topology lacking a net radial flux, as is common, the net magnetic flux through the domain will remain constant.

2.2.2 Results from Local Models

In addition to pioneering the use of the local model in accretion disk simulations, Hawley et al. (1995) also used the methodology to perform the first thorough exploration of the parameter space associated with the MRI. Utilizing a variety of unstratified local simulations initialized with differing initial magnetic fields they were able to demonstrate unambiguously that the nonlinear resolution of the MRI is a state of vigorous, anisotropic turbulence that results in sustained angular momentum transport. The saturated, or quasi-steady, state (QSS) refers to the fluid state after the initial growth phase of the linear MRI and the transition into a fully nonlinear state. The delicate interplay between numerics and physics was noted in Hawley et al. (1995), in particular the dependance of the saturated state on both the size of the numerical domain being simulated and on the presence of a net magnetic field threading the simulation domain.

This concept is encapsulated in the notion of a saturation predictor, in which the stress or magnetic pressure in the QSS is related to the initial conditions of the simulation. For simulations initialized with a constant vertical magnetic field profile it was found that

$$\langle P_b \rangle_{QSS} \approx 1.21\rho_0(L_z\Omega_0)(\lambda_{MRI}\Omega_0), \tag{2.32}$$

where $< X >_{QSS}$ refers to a spatial and temporal average of the quantity in the QSS, and ρ_0 is the value of the constant density used to initialize the simulation. The dependence on the size of the simulation domain is a consequence of the small domains considered and would not be expected to extend to simulations large enough to ensure that the domain size doesn't inhibit the formation and resolution of structure. Alternatively, for simulations initialized with a constant toroidal magnetic field, the following predictor is derived

$$\langle P_b \rangle_{QSS} \approx 0.01\rho_0(L_y\Omega_0)(\lambda_c\Omega_0), \tag{2.33}$$

where the characteristic wavelength, λ_c is given by equation 1.46. Further, independent of initial field topology, $< M_{R\phi} + T_{R\phi} >_{QSS} \propto < P_b >_{QSS}$. In particular, this relationship is encapsulated by the dimensionless product

$$\alpha\beta = \frac{M_{R\phi} + T_{R\phi}}{P_b} \approx \frac{1}{2}. \tag{2.34}$$

The relationship of equation 2.34 was further bolstered by (Blackman et al., 2008), who conducted a meta-analysis of a wide range of published local simulations, and showed that this result holds quite broadly.

The saturation predictors given in equations 2.32 and 2.33 demonstrate the dependence of the saturated state on both physical (initial field strength) and numerical (box-size) parameters. Increasing either will result in more angular momentum transport and a stronger characteristic magnetic energy. Further, as suggested by the linear theory of the MRI, poloidal magnetic fields are much more efficient at generating angular momentum transport than toroidal field configurations of the same energy. The dependence on domain size is even deeper than this; the simulations considered by Hawley et al. (1995) are all such that $L_x/L_z = 1$ and exhibit significant levels of intermittency characterized by impulsive spikes of accretion. It was found (Bodo et al., 2008) that more disparate values of L_x and L_z led to suppressed intermittency as compared to simulations with a radial aspect ratio of unity.

At first glance the dependence of the saturated state on the initial magnetic field strength and topology are genuinely physical, however even here there are numerical details lurking beneath the surface. It is known from the linear theory of the MRI that a sufficiently strong vertical magnetic field becomes a stabilizing, rather than destabilizing, force. In practice, sufficiently strong fields are associated with large wavelengths and will be eventually become larger than the vertical extent of the simulation. Conversely, as weak fields are associated with small wavelength disturbances, sufficiently weak fields will be driven to spatial values smaller than the grid scale and thus be unresolvable. In this case, even though there is a physical magnetic field present, numerically this field will be equivalent to a null field. These details were used by Pessah et al. (2007) to augment the vertical field saturation predictor as

$$< \alpha_M >_{QSS} \left(\frac{H}{L_z}\right)^{(5/3)} = 0.61 \times \left\{ \begin{array}{ll} \Delta_z/L_z & : \lambda_{MRI} \leq \Delta_z, \\ \lambda_{MRI}/L_z & : \Delta_z < \lambda_{MRI} \leq L_z, \\ 0 & : L_z < \lambda_{MRI}, \end{array} \right\} . \tag{2.35}$$

For simulations not initialized with a constant vertical field, λ_{MRI} represents the volume average of the pointwise value. The augmented saturation predictor includes three regions: an unresolved region in which λ_{MRI} is unresolvable and consequently $\alpha \propto \Delta$; the resolved region probed by the saturation predictor given in equation 2.32; and a stable region in which λ_{MRI} exceeds the vertical domain of the simulation and turbulence is absent. The structure of the saturation predictor given in equation 2.35 is shown in Figure 2.5 along with the results of a series of simulations of varying initial field strength.

The augmented saturation predictor of Pessah et al. (2007) incorporates the earlier convergence results of Fromang & Papaloizou (2007). Fromang & Papaloizou (2007) consider consider the resolution dependence of a series of unstratified local models, initialized with a vertical field with vanishing net flux. For these simulations initialized with a vertical field, it is common to discuss resolution in the context of the number of zones per vertical scale-height (ZPH). The simulations considered are of size $(L_x, L_y, L_z) = (H, \pi H, H)$ with vertical resolutions (in ZPH) of $H/\Delta_z = 64, 128, 256$. It was found that as the resolution increased, $< \alpha >_{QSS}$ tended towards zero.

The lack of convergence in these simulations was initially quite troubling, however a series of results shortly afterwards demonstrated that the lack of convergence required a combination of model parameters. The absence of any of these parameters fundamentally changed the result and convergence was achieved. Convergence towards a non-zero value of α can be achieved through the addition of stratification (Davis et al., 2010), the inclusion of an initial magnetic field with a non-zero net flux

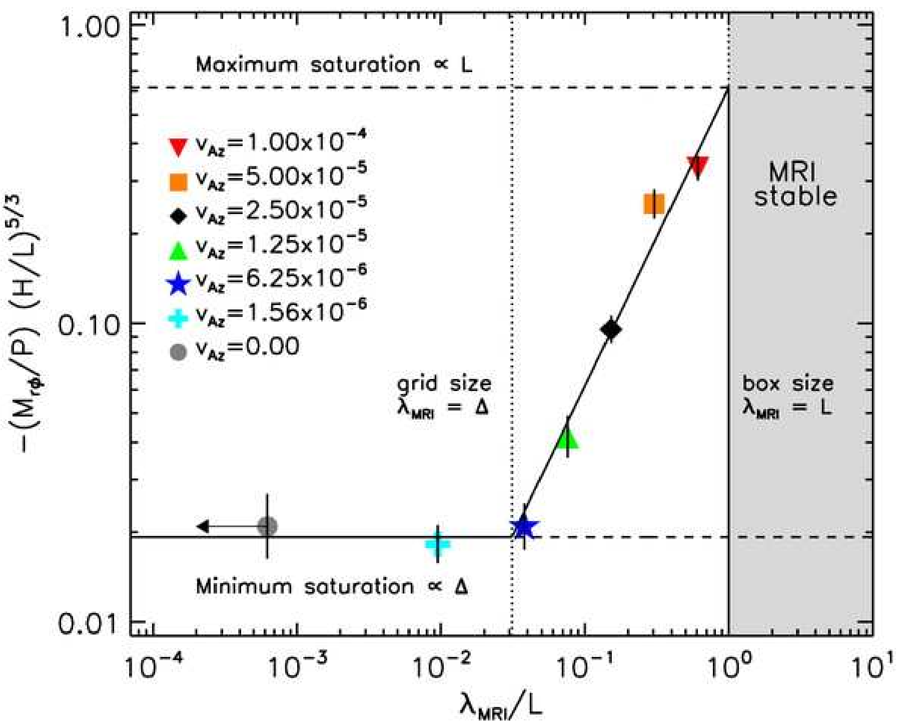

Figure 2.5: Series of simulations utilizing purely vertical initial magnetic fields of varying strengths, characterized by the Alfven velocity of the initial field. Figure taken from Pessah et al. (2007).

(Simon et al., 2009), the inclusion of explicit dissipation terms, or the presence of a sufficiently large L_z (Jim Stone, private communication). The presence of explicit dissipation in addition to contributing to convergence, also has a significant effect on the saturated state. It was found (Lesur & Longaretti, 2007) that in the saturated state $\alpha \propto P_m$, where $P_m = \nu/\eta$ represents the dimensionless Prandtl number, the ratio of the viscosity and the resistivity.

The lack of convergence in the models considered by Fromang & Papaloizou (2007) is often attributed to the lack of a physical scale in the problem. The absence of a physical scale gives the gridscale an unphysical importance in the problem, as evidenced by the behavior $\alpha_M \propto \Delta_z$ when the wavelength associated with the net vertical flux is unresolvable. The addition of stratification transforms the scale height from a thermodynamic measurement to an important vertical lengthscale. The presence of net flux imposes an alternative lengthscale, namely the wavelength associated with the net magnetic flux. Analogously, explicit dissipation incorporates a dissipation lengthscale. These imposed physical lengthscales take the place of the numerical gridscale and thus these models will be less sensitive to resolution. While the lack of convergence of unstratified, zero net flux models is a somewhat singular case the result does serve to highlight the importance of numerical effects in accretion disk simulations.

These results are indicative of the difficulty associated with simulations of MRI-driven MHD turbulence in accretion disks. Choices of initial parameters, particularly the box-size and initial magnetic field strength and topology, strongly influence the saturated state. The effect of this is a particularly large parameter space associated with local disk models.

2.3 Local Versus Global Models

The reduced computational expense associated with local models is a clear advantage, however this comes at a steep price. There are significant, well-known limitations (Regev & Umurhan 2008) to the local approximation. The closed nature of these simulations enforces the conservation of quantities on scales much smaller than would be astrophysically realistic. This constraint inhibits the evolution and formation of radial structure. In particular, the evolution of the magnetic field is constrained by the condition that the net magnetic flux is conserved. As we will see, the strong dependence of accretion on magnetic flux makes this a difficult approximation to justify. On the other hand, global models allow the study of the full evolution of the disk at the cost of significant computational resources. Clearly, these two models must be used in a complimentary manner with local models probing the structure of disk turbulence at small scales and global models studying large scale structure formation. Consequently, verifying that these two models agree at intermediate scales is a vital self-consistency check and an important step towards understanding the limitations of both types of models. The results from a preliminary investigation exploring the connection between local and global models is presented in the next chapter.

Chapter 3

Connecting Local and Global Models: Preliminary Investigation

3.1 Introduction

Local simulations have known limitations (Regev & Umurhan, 2008; Bodo et al., 2008). By construction, they enforce the local conservation of quantities (in particular the net magnetic flux) that in reality are only globally conserved. In many older implementations they also enforce periodicity (in radius, azimuth, and in some instances also height) on a scale that may be small enough to impact the results. Recently, a number of authors (Davis et al., 2010; Guan et al., 2009; Johansen et al., 2009) have used larger-than-usual shearing-box simulations to quantify whether these limitations matter for practical purposes. In this chapter we address the same problem from the other direction. We analyze small patches of *global* disk simulations in an attempt to determine whether the disk behaves as if it were a collection of shearing-boxes. Our specific goal is to ascertain whether the relationship between local $R\phi$-component of the magnetic stress and vertical magnetic flux that is found in local simulations (Pessah et al., 2007), and encapsulated in equation 2.35, is recovered in local patches of global simulations. As a result, we uncover the potential importance of the magnetic connectivity of the disk and the need to fully capture the vertical extent of the corona.

The work presented here is a preliminary investigation of the connection between local and global models of accretion and certain numerical concerns related to the geometry of the simulation domain. This investigation is preliminary in the sense that it uses a simulation model that was originally intended to serve a somewhat different astrophysical purpose (Reynolds & Miller, 2009). Because of this there are additional complications introduced by the physics associated with the pseudo-Newtonian potential and vertical stratification. The benefit of using this simulation is that it has been thoroughly explored and can be considered free of defects and has moderate resolution. The work presented in this chapter uses this model as an incubator for the ideas that will be explored in greater detail throughout the remainder of this dissertation utilizing models that have been tailored to address the questions raised here.

The plan of this chapter is as follows. Section §3.2 describes the simulations used in this work. Section §3.3 discusses the importance of vertical domain size to the accretion efficiency in the saturated state. The main body of work is described in Section §3.4 which discusses the correlation between magnetic flux and stress and how this relationship in global models is an extension of the saturation predictors used for local models. Finally, a discussion and concluding remarks is included in Section §3.6.

3.2 Simulations

The simulations used here are based upon those described by Reynolds & Miller (2009). We use the ZEUS-MP code (Stone & Norman, 1992; Hayes et al., 2006) to solve the equations of ideal MHD in three dimensions. The basic version of the code has been modified to incorporate a Paczynski-Wiita pseudo-Newtonian gravitational potential (as a first approximation of the gravitational field about a Schwarzschild black hole; Paczynsky & Wiita 1980) and the simulation is run using cylindrical polar coordinates. These simulations are ideal MHD in the sense that no explicit resistive or viscous dissipation is included; all dissipative processes are due to the discretization of the spatial domain and hence occur close to the grid scale. Furthermore, an internal energy equation assuming an adiabatic equation of state with $\gamma = 5/3$ is used. Energy is lost from the domain when magnetic fields undergo numerical reconnection due to the non-conservative nature of the treatment of the energy equation in `Zeus`.

The initial disk is in a state of Keplerian rotation (with respect to the pseudo-Newtonian potential), and is in vertical hydrostatic equilibrium with a constant scale height h. Thus, the initial density, pressure and velocity field is

$$\rho(R,z) = \rho_0(R)\exp(-\frac{z^2}{2h^2}), \tag{3.1}$$

$$p(R,z) = \frac{GMh^2}{(r-2r_g)^2 r}\rho(R,z), \tag{3.2}$$

$$v_\phi = R\Omega = \frac{\sqrt{GMR}}{R-2r_g}, \tag{3.3}$$

$$v_z = v_r = 0, \tag{3.4}$$

where R represents the cylindrical radius, $r = \sqrt{R^2+z^2}$, $r_g = GM/c^2$ and $h = 0.05r_{isco} = 0.3r_g$. We set the initial midplane density to be $\rho_0(R) = 1$ beyond the innermost stable circular orbit (ISCO) at $r_{isco} = 6r_g$, and $\rho_0(R) = 0$ within the ISCO.

The initial magnetic field is specified in terms of a vector potential of the form,

$$A_\phi = A_0 f(R,z) p^{1/2} \sin(\frac{2\pi R}{5h}), \tag{3.5}$$

$$A_R = A_z = 0. \tag{3.6}$$

Here f is an envelope function that is unity in the disk body and smoothly goes to zero away from the main body so as to avoid unphysical interactions with the boundaries. This results in a magnetic field topology consisting of distinct poloidal field loops of alternating orientation throughout the main body of the simulation. Of importance to the current discussion is that there is no net vertical magnetic flux threading the disk as a whole. The constant A_0 is chosen to ensure that the magnetic field strength is normalized so that the ratio of volume-integrated gas and magnetic pressure $\beta \approx 10^3$.

The boundary conditions used in global simulations often present greater difficulties than their local counterparts. The desire to allow material and magnetic flux to leave the domain through accretion processes require the use of appropriate outflow boundary conditions. Here, a zero-gradient outflow boundary is used in the radial direction in which the fluid variables in the ghost zones are set based on an enforced zero-gradient constraint. Outflow is ensured through the use of a "diode" condition, in which the radial velocity in the ghost zones are set to zero in the event that they would otherwise correspond to motion into the simulation domain. Along the ϕ-boundary periodic boundary conditions are the natural choice, however more unnaturally this choice is also made along the z-boundary. This latter choice is made to avoid common numerical issues associated with outflow along the vertical

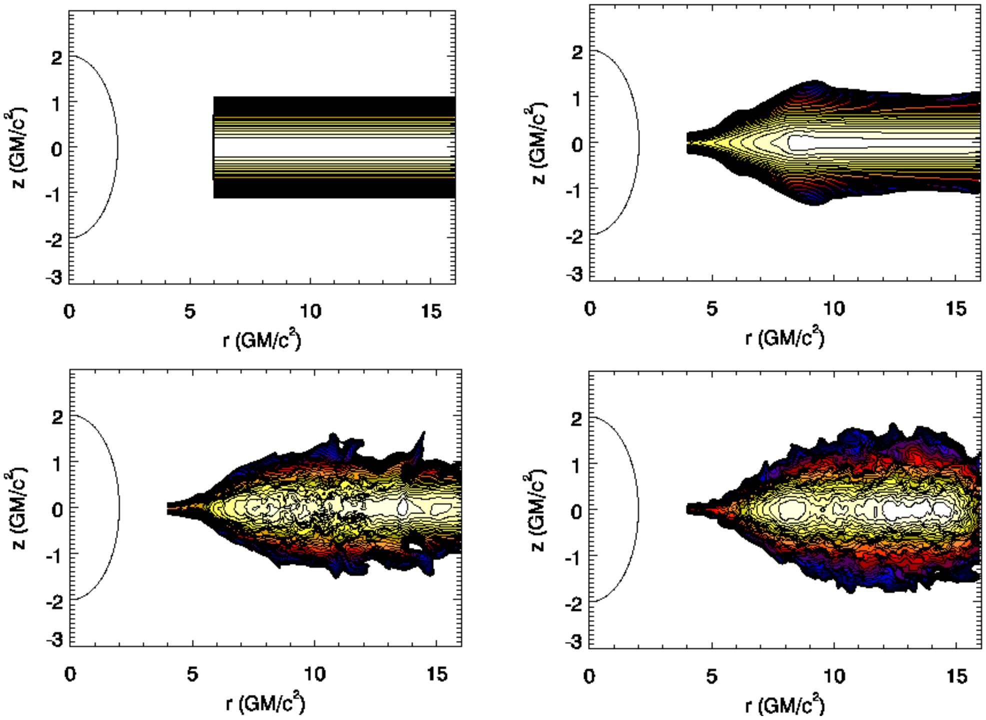

Figure 3.1: Images of the evolution of simulation `Thin.M-Res.6z`. Shown are poloidal contours of the logarithmic density utilizing 10 contours per decade over the range of three decades in density. The snapshots correspond to the initial conditions (**Top-Left**), and the evolution after one (**Top-Right**), ten (**Bottom-Left**), and one hundred (**Bottom-Right**) orbits at r_{isco}. Figure taken from Reynolds & Miller (2009).

boundary. Specifically, as noted by Stone et al. (1996) outflow boundary conditions can result in field-line “snapping”, in which a field loop traversing the boundary becomes disconnected and subject to potentially strong magnetic tension forces that can affect the entire disk. This is avoided here, as in Stone et al. (1996), through the use of periodic vertical boundary conditions where the vertical boundaries are placed sufficiently far from the turbulent disk to avoid potential artifacts.

A set of simulations were run to span a range of numerical parameters, specifically varying the vertical and radial resolution, as well as the vertical and azimuthal extent of the domain. A comparison of simulations with varying azimuthal extents suggest a negligible dependence on this parameter, and as such all the simulations considered here use the same 30° wedge-shaped azimuthal domain. The vertical domain size is found to be important, and will be discussed below. A detailed study of the dependence on resolution, requiring much greater computational expense, is deferred to a later work. All simulations presented here have a radial domain $r \in (4r_g, 16r_g)$. Details of the simulations considered are given below in Table 3.1. A series of still images is presented in Figure 3.1 to give a sense of the evolution of the fiducial disk model `Thin.M-Res.6z`. Shown are the contours of the logarithmic density in the poloidal plane and how they evolve over time, the horizon of the black hole is marked for reference as the curved line at $r = 2$. The turbulent nature of the disk is clear after 10 ISCO orbits from the small-scale structure of the density contours and from the snapshots the radial and indeed vertical evolution of the disk can be inferred.

Run ID	Resolution (R,ϕ,z)	Vertical Extent (in r_g)	Total Orbits (at r_{isco})	Mean α_M	Standard Deviation of α_M
Thin.M-Res.12z	(240,32,512)	12	122	0.008	0.0011
Thin.M-Res.6z	(240,32,256)	6	664	0.0086	0.0016
Thin.M-Res.3z	(240,32,128)	3	112	0.0061	0.001

Table 3.1: The spatial resolution, vertical domain size, duration of the simulations analyzed in this work. Also included are the mean and standard deviation of α_M, defined by equation 3.7, between 50 and 100 orbits at r_{isco}.

3.3 Dependence of Stress on Vertical Domain Size

Quantifying angular momentum transport in a global, stratified disk requires a suitable metric that highlights the dominant region of accretion in the midplane, while minimizing the importance of potentially problematic regions near the boundary. Towards this end, a density-weighted average of α_M (equation 1.32) is defined to be

$$\alpha_{\mathrm{M}} = \left\langle \frac{\int -\frac{\rho M_{r\phi}}{p} dz}{\int \rho dz} \right\rangle_{\phi, r \in (7r_g, 12r_g)}. \tag{3.7}$$

The restriction on the radial range of the averaging is designed to ignore the plunging region of the accretion flow ($r \leq 6r_g$) and any effects related to the outer radial boundary. Density weighting is used in the vertical direction to take into account the low density, highly magnetized regions while still allowing the dominant contribution to the integral to come from the denser mid-plane of the disk.

A comparison of α_M and its dependence on vertical domain is given in Figure 3.2. The initial growth phase of the MRI is unaffected by the vertical domain as expected, since all the simulations considered have the same vertical resolution and can thus resolve the same unstable MRI modes. While the two simulations with the largest vertical domains exhibit comparable values of α_M, simulation `Thin.M-Res.3z` produces a consistently smaller value. We attribute this effect to stifling of the growth of the magnetized regions in the smaller vertical extent simulation. However, the long-term effects of vertical extent are ambiguous. Whether the simulations converge to the same α_M or the apparent convergence is a result of short-term variability is unclear from the current simulations. Longer simulations will need to be carried out to determine the vertical domain size that is needed in order to reliably capture the dynamics of a global disk. These results do suggest the potential dynamical importance of the corona and this point is returned to in the context of the generation of stress in response to local magnetic flux.

3.4 Connection Between Flux and Stress

3.4.1 Vertical Structure of Stratified Thin Disks

Our primary goal here is to study the instantaneous correlation between stress and magnetic flux within the simulated disk. In a stratified disk this correlation may vary with height above the disk midplane, and hence we start by considering how the mean magnetic field structure varies vertically in our simulations.

To analyze the simulations the principal quantity of interest is the $R-\phi$ component of the Maxwell stress tensor which dominates MRI-driven angular momentum transport (Balbus & Hawley, 1998). In a turbulent stratified disk, both $M_{R\phi}$ and other physical quantities of interest are complicated functions of space and time. To make sense of them we use temporal and spatial averages. Run `Thin.M-Res.6z` has the longest duration of any of our simulations, and we use this run to construct a representative vertical profile of the magnetic structure of the disk. To reduce the effects of spatial intermittency in the turbulence, we azimuthally average over the entire domain and average over a small radial range centered about a fiducial radius in the body of the disk ($8r_g - h < r < 8r_g + h$) . To smooth out the temporal variability and isolate the behavior of the disk in a saturated turbulent state we time average over 400 ISCO orbits starting at orbit 50.

The results are given in Figure 3.3, which shows the vertical profiles of the relevant quantities scaled to their maximum values. Our interest in vertical structure is predominantly in the magnetic fields, and in particular the vertical magnetic flux, B_z, and the magnetic stress, $M_{R\phi}$. However, we also plot the density, ρ, to highlight the contrast between the relatively unmagnetized midplane of the disk and the sparse magnetized "corona" away from the midplane. The obvious reason for the formation of these two disparate regions is magnetic buoyancy resulting from the effect of vertical gravity, but this may be overly simplistic. Also of interest is the double-peak vertical profile of vertical flux and stress. Broadly similar results are seen in a subset of the stratified local simulations of Miller & Stone (2000) and a subset of the global simulations of Fromang & Nelson (2006). It is interesting that the vertical location of the peak field and stress seem to approach constant values rather than growing monotonically with time. Whether the region of strong flux is trapped, possibly due to magnetic tension from field lines connecting it to the midplane, or is continually dissipating while outflowing, is currently unclear.

The local flux-stress relationship to which we will compare our results is based on unstratified shearing box simulations and the rich vertical magnetic structure due to stratification in our global disk means that there is necessarily some ambiguity in the comparison. In what follows, we therefore analyze the flux-stress relationship within the global simulation not just at the midplane but also as a function of height. One should note that the offset of the peak vertical flux and stress from $z = 0$ means

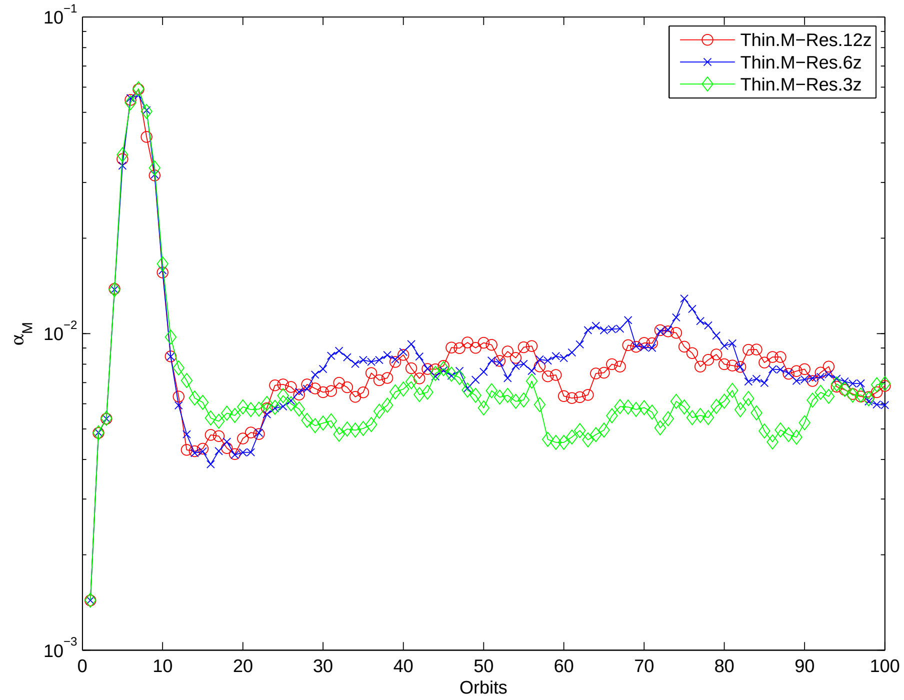

Figure 3.2: Behavior of α_M and its dependence on vertical domain.

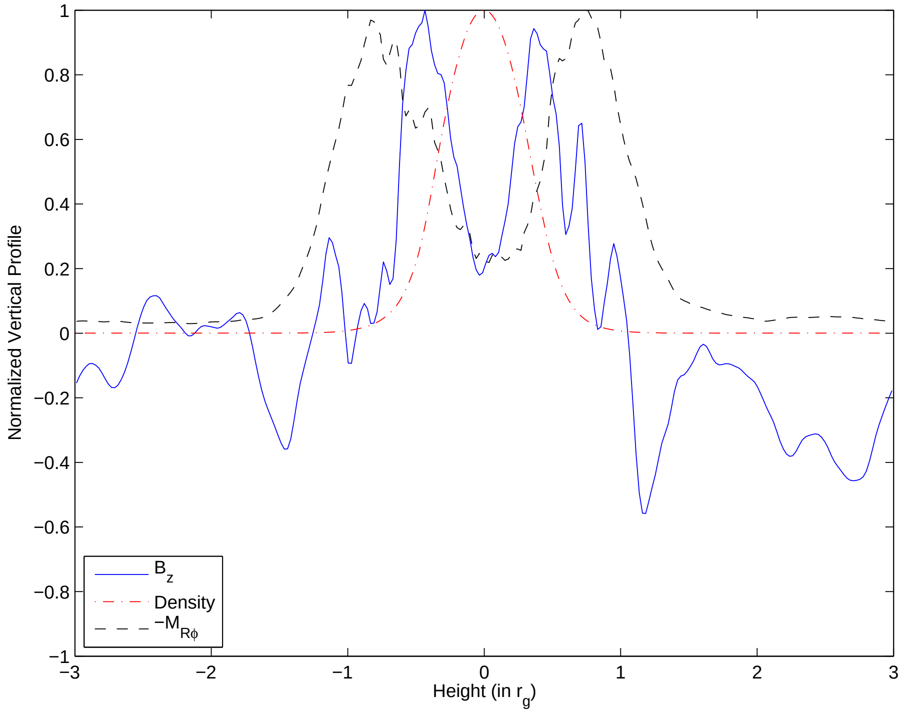

Figure 3.3: Vertical structure of physical quantities averaged over four hundred orbits for run `Thin.M-Res.6z`.

that high values of flux and stress are only accessible away from the midplane.

3.4.2 Local Flux-Stress Correlations in a Global Disk

Although our simulations have zero net vertical flux, small patches of the disk *are* instantaneously threaded by a vertical field. We seek to determine whether the Maxwell stress tracks this transient vertical field in the same way as it would in a local simulation where the vertical field is persistent (Hawley et al., 1995; Pessah et al., 2007). To proceed, we break up the global simulation domain at each timestep into several hundred small cylindrical wedges of size $\Delta z = \Delta r = h$ and $\Delta\phi \approx 0.1$. Within each wedge we average to obtain a single estimate of the magnetic stress ($M_{r\phi}/p$, normalized to the local gas pressure) and the local vertical flux, which we express in terms of the wavelength of the most unstable MRI mode, λ_{MRI} (equation 1.40). Note that because we scale the stress by the local pressure (a decreasing function of height), we immediately introduce a height dependence. Choosing instead to scale by the midplane pressure (only a function of radius) still yields a height dependence and as such we are confident that the height dependence seen is not solely a consequence of the pressure scaling.

To avoid early transients, our analysis excludes the period prior to the first 50 ISCO orbital periods. To improve our statistics (and to give a measure of the convergence of our results given the finite duration of the run) we consider wedges that are centered at $z = \pm[0, 1, 2, 3]h$ and plot results for the samples separately. The resulting pairs of flux-stress values from all of the wedges and all of the snapshots in time were binned according to (logarithmic) vertical flux in order to diagnose trends.

The resulting flux-stress relations for all the simulations considered are broadly similar. Our best statistics come from the long duration run `Thin.M-Res.6z`, and the flux-stress relation for this case is plotted in Figure 3.4. The stress is observed to be flat for weak vertical fields (small λ_{MRI}), while for larger field strengths we have approximately $M_{R\phi} \propto \lambda_{MRI}$. This may be compared with the local scaling relation derived from unstratified simulations, equation 2.35. Note the distinction between the locally defined H, used to calculate the dimensionless stress in each wedge, and the globally constant h, used in the initial conditions. In comparing our results to those obtained for local simulations, we consider L to be the size of the wedge and H to be the locally defined pressure scale height given by

$$H = \sqrt{\frac{2}{\gamma}} \frac{c_s}{\Omega_K}, \tag{3.8}$$

which is consistent with the previously-defined isothermal scale height (equation 1.10) for $\gamma = 1$.

Our simulations do not spontaneously develop vertical fields strong enough to quench the MRI (and hence we do not sample the $\lambda_{MRI} > L$ regime), but the behavior of patches at low and intermediate vertical fluxes is qualitatively the same as that found in local simulations. In contrast to the local (unstratified) results of Pessah et al. (2007) is the strong dependence on height in our stratified simulations. In addition to the fact that the largest values of flux are only accessible at large heights is the fact that the stress response to flux is also height dependent. There is a height-dependence of the transition point between low and intermediate flux, as well as in the slope of the intermediate flux regime. Of particular note is the location of the transition point itself. The vertical line marked in Figure 3.4 is given by $\lambda_{MRI} = \Delta/20$, and approximately marks the location of the transition point at $z = \pm H$. This stands in contrast to the transition point for local unstratified simulations, $\lambda_{\rm MRI} = \Delta_z$. We return to a discussion of the physics of this transition in Section 3.4.3.

The fact that transient self-generated vertical flux is able to stimulate the local stress in the same manner as occurs in local models is primarily a formal result, although it does lend some credence to models in which patches of vertical field are assumed to have a physical identity (Spruit & Uzdensky, 2005). Of greater import is the observation that, across much of the disk, the self-generated field is strong enough to fall into the linear regime of the flux-stress relation. Figure 3.5 shows the distribution of flux through patches in all three. The vertical flux distributions are approximately symmetric in log-space. We find that, at any instant in time, about half of the area of the disk is threaded by a field strong enough to control the stress. We interpret this to mean that in a zero-net field global simulation, much of the disk sees a field strong enough to control its dynamics. In other words, the dynamics of the disk is strongly influenced by the connectivity of the self-generated magnetic field between different patches of the disk.

In these results we again see the importance of the magnetized region away from the midplane. Not only is this the location of the largest values of flux, and

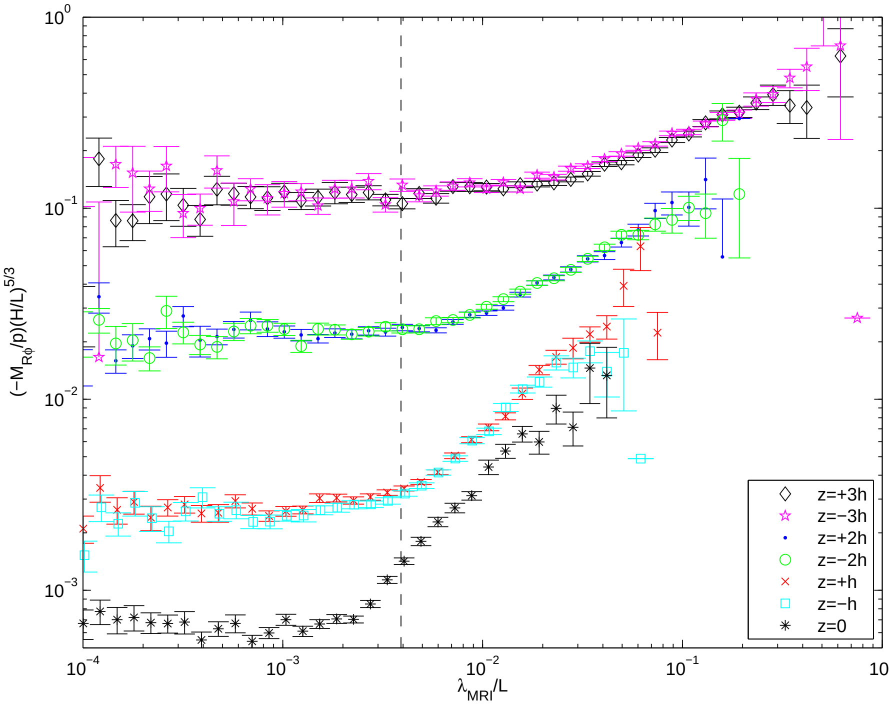

Figure 3.4: Flux versus stress averaged over four hundred orbits for run `Thin.M-Res.6z`.

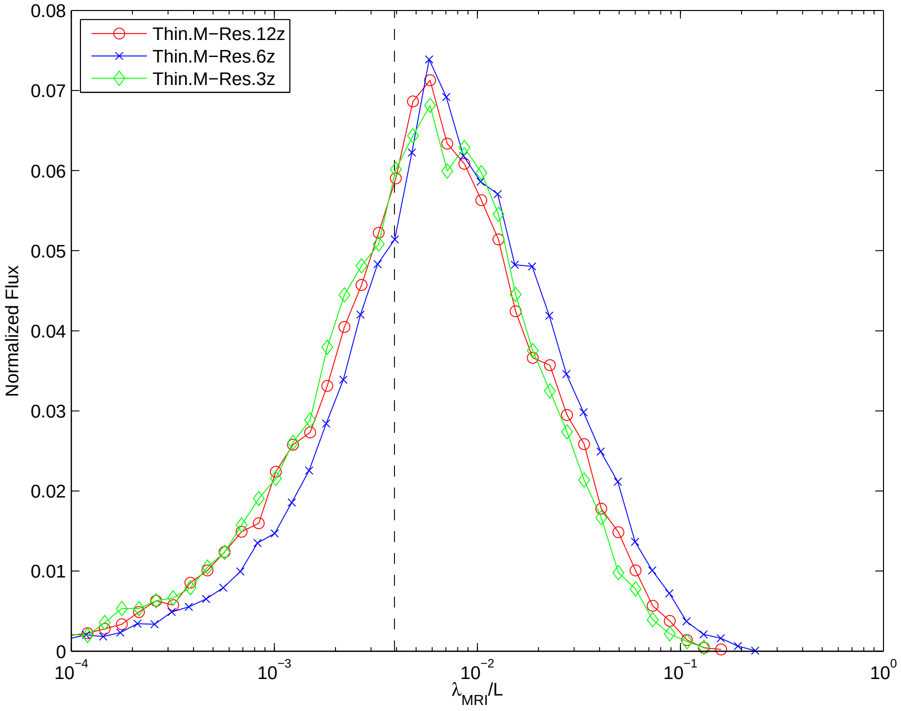

Figure 3.5: Normalized flux distributions for all simulations, computed between 50-100 ISCO orbits. The dashed vertical line indicates, approximately, the strength of vertical field above which magnetic stress in the disk is stimulated.

thus stress, but for the same flux the stress response is higher. This increased stress response to flux suggests that the corona is important not just as a warehouse of magnetic energy, but also has the ability to use this magnetic energy more efficiently to induce angular momentum transport. One possible explanation for this efficiency is the ability of the corona to mediate magnetic links through radially disparate regions of the disk due to the presence of coronal field loops. This is potentially connected to the results discussed in Section 3.3, in which the importance of vertical domain size is considered, and is possibly the result of the truncation of the corona affecting the dynamics of the disk. Considered next, is the nature of the transition point itself.

3.4.3 The Nature of the Transition Point in the Flux-Stress Relation

To reiterate, the flux-stress relation obtained from our global simulation shows a transition at approximately $\lambda_{MRI} \sim \Delta/20$, in contrast with $\lambda_{MRI} \sim \Delta$ found from local unstratified simulations (Pessah et al., 2007). Neither of these transitions can correspond straightforwardly to the condition that the fastest growing MRI mode is resolved. To resolve a mode in a ZEUS-like scheme requires that the wavelength is spanned by at least ~ 8 computational zones. Since λ_{MRI} is, by construction, the wavelength of the fastest growing mode corresponding to the net vertical magnetic field, the condition that the fastest growing MRI mode is resolved corresponds to $\lambda_{MRI} \sim 8\Delta$. Our result then implies that magnetic fields that are very weak – in the sense that their fastest growing mode cannot be resolved – nonetheless have an important influence on the dynamics of our global disk.

We do not have a quantitative explanation of why $\lambda_{MRI} \sim \Delta/20$. On general grounds, however, we note that we would expect that the transition point would lie at $\lambda_{MRI} \ll 8\Delta$. A given vertical field is unstable not just to the fastest growing mode, but also to a whole spectrum of slower-growing modes that have longer wavelengths that are more easily resolvable numerically. Plausibly, the transition point will then correspond to the condition that we resolve the slowest growing mode that grows appreciably before it is truncated by non-linear coupling to other MRI modes or some other aspect of the physics (e.g. a dynamo cycle). If this is the case, then it is unsurprising that the transition point varies between local and global simulations, since the time scale available for a mode to grow may well depend on the presence or absence of a low density disk corona within which the MRI is not active.

To consider this more quantitatively, consider purely vertical MRI modes ($k_r = k_\phi = 0$) in a thin accretion disk (so that radial gradients of pressure and entropy can be neglected). Let the spacetime dependence of the modes be $e^{i(\omega t - kz)}$. The dispersion relation for these modes (Balbus & Hawley, 1991) reads

$$\tilde{\omega}^4 - \kappa^2\tilde{\omega}^2 - 4\Omega_0^2 k^2 v_A^2 = 0, \tag{3.9}$$

where $\tilde{\omega}^2 = \omega^2 - k^2 v_A^2$ and κ is the radial epicyclic (angular) frequency. We wish to examine modes with wavelengths much longer than the fastest growing mode, i.e., with $|kv_A/\Omega_0| \ll 1$. Rewriting in terms of the growth rate, $\sigma = -i\omega$ and expanding the dispersion relation to lowest order in $k^2 v_A^2/\Omega_0^2$ gives

$$\sigma^2 = \left[4\left(\frac{\Omega_0}{\kappa}\right)^2 - 1\right] k^2 v_A^2. \tag{3.10}$$

Suppose that a given mode can grow exponentially for a time τ before it is truncated by mode coupling or some other unspecified physical process. Then, the slowest growing mode that actually experiences significant growth (hereafter, the slowest appreciably growing mode [SAGM]) has $\sigma_{sagm} = 2\pi/\tau$ and a wavenumber given by

$$k_{sagm}^2 v_A^2 = \frac{4\pi^2\kappa^2}{\tau^2(4\Omega_0^2 - \kappa^2)}. \tag{3.11}$$

Our hypothesis is that the transition point in the flux-stress relation corresponds to the point where the slowest appreciably growing mode is just resolvable, i.e., where $\lambda_{sagm} \equiv 2\pi/k_{sagm} \sim 8\Delta$. This predicts a transition point at

$$\lambda_{MRI} \sim \frac{16\kappa}{15^{1/2}\pi(4\Omega_0^2 - \kappa^2)^{1/2}} \frac{\Delta}{\tau'}, \tag{3.12}$$

where $\tau' \equiv \tau/t_{orb}$, t_{orb} being the orbital period at that radius.

Equation 3.12 offers some insight into the transition point found in the flux-stress relations that we have been considering. In the local unstratified simulations of Pessah et al. (2007), the implicit potential is Newtonian ($\kappa^2 = \Omega_0^2$) and their results indicate $\lambda_{MRI} \sim \Delta$, implying $\tau \sim t_{orb}$. In the global simulations presented here, we find the transition point at $\lambda_{MRI} \sim \Delta/20$ which (accounting for the fact that $\kappa^2 < \Omega_0^2$ in the pseudo-Newtonian potential) gives $\tau \sim 5 - 10 t_{orb}$. Thus, within the framework of this argument, the difference in the location of the transition point between the local unstratified and the global simulations is due to a difference in the robustness of the long wavelength and slowly growing modes; slowing growing modes appear to be able to grow for longer within the global simulation before being truncated. The nature of this difference, which must be closely related to the saturation of the turbulent state, is beyond the scope of this work. It will be demonstrated in Chapter 6 that the transition point itself is a function of resolution.

3.4.4 Temporal Flux-Stress Correlations

The results of §3.4.2 suggests that the (fluctuating) magnetic flux threading a local patch of the disk determines the $R-\phi$ component of the magnetic stress generated by the turbulence in that patch. If the vertical magnetic flux is indeed the causal agent in determining the stress, we expect a temporal lag between fluctuations in the magnetic flux and the resulting variations in the stress. On the basis of experiments with local simulations (Hawley et al., 1996), we expect this lag to be approximately two (local) orbital periods. Thus, we expect the temporal lag to increase with radius in the disk due to the increasing orbital period.

To search for this lag, we use `Thin.M-Res.6z` and output the 3-d structure of the disk once every 0.1 ISCO orbits during the interval between 50–90 ISCO orbits (this is 10 times the nominal data output rate). Using these 400 snapshots of the disk structure, we then computed the instantaneous vertical magnetic fluxes and magnetic $R-\phi$ stresses in families of co-moving wedges at three radii $R \in \{8r_g, 10r_g, 12r_g\}$. The azimuthally averaged value of v_ϕ at each radius was used to track a given co-moving wedge between timesteps. This procedure is not fully Lagrangian, because it does not account for the radial movement or fluctuating azimuthal velocity of a co-moving patch, but we expect these effects to be negligible for the short timeframe under consideration. The time-series of magnetic flux and stress for each wedge were then cross-correlated and, finally, the cross-correlations for all wedges at a given radius were averaged.

The resulting averaged temporal cross-correlations are shown in Fig. 3.6. At each radius we see a strong instantaneous correlation, likely due to the immediate shearing of perturbed vertical fields. However, in general, the cross-correlation is biased toward positive lag. This is consistent with what we would expect, namely that the presence of vertical flux will feed the MRI and result in enhanced transport. Of note is the fact that the inner-most radius considered, $R = 8r_g$ exhibits a double peak structure whereas this is unresolved at higher radii. Also peculiar is the fact that the outer-most radius, $R = 12r_g$ is significantly less biased towards positive lag than the other radii under consideration. A further exploration of these issues is beyond the scope of this paper, and will be explored in future work employing orbital advection algorithms and test-particle tracers in order to correctly follow the evolution of a local patch.

3.5 Long-Term Behavior

In the one case of run `Thin.M-Res.6z`, the disk was simulated for 664 ISCO orbits. This simulation allows us to search for long-term trends in the dynamics of the disk. As shown in Figure 3.7 there is a slight downward drift in α_M over time. The same temporal trend is also evident in the flux-stress relationship. Figure 3.8 shows the flux-stress relationship averaged in 100 orbit blocks starting at 50 ISCO orbits. During the first 300 orbits, there appears to be a secular drift in the flux-

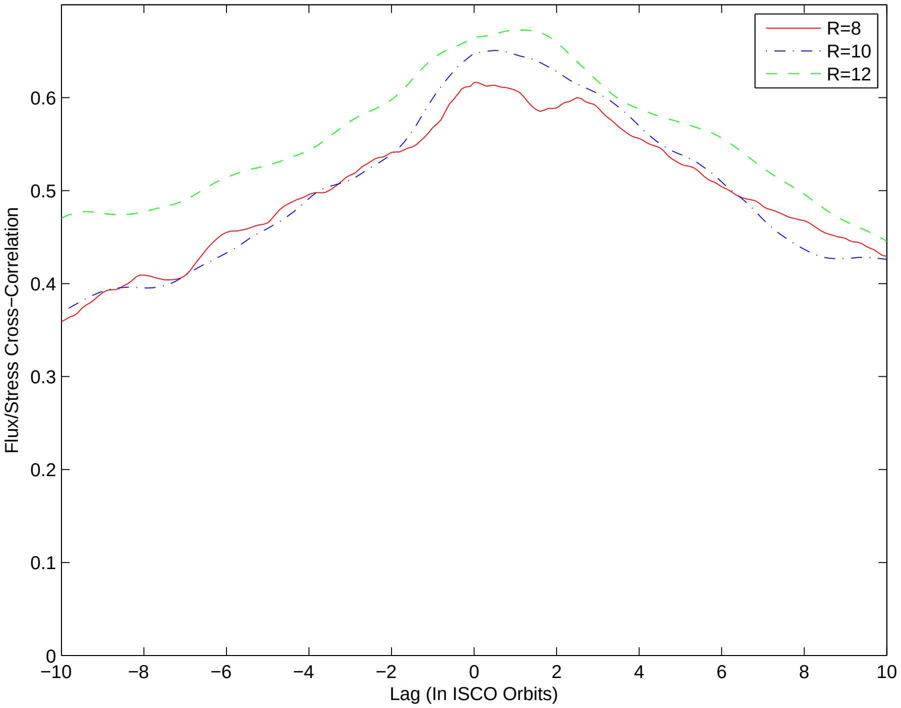

Figure 3.6: Cross-correlation of flux and stress evaluated in co-moving patches of the disk. Positive lags imply that changes in the vertical flux occur, on average, prior to changes in the magnetic stress.

stress relation. The linear (high-flux) part of the relation achieves a steady state relatively quickly (only the first time block between 50 and 150 ISCO orbits shows significant differences), but the flat (low-flux) part of the relation continues to fall until it too achieves a steady state at approximately 350 ISCO orbits into the run. Associated with this, the “knee” in the flux-stress relation appears to move to smaller fluxes.

In essence, this result says that low-flux regions still support (small) stresses at early times but that those stresses decay over a period of several hundred ISCO orbits. We ascribe this to stresses associated with a sheared residual of the initial magnetic field configuration which are “mixed away” on a relatively long timescale. Our initial field configuration threads the midplane with regions of net magnetic flux which alternate with a radial periodicity of $5h$. Radial Fourier transforms of the midplane azimuthally-averaged B_z do indeed find a (weak) periodicity corresponding to the initial field even once the turbulence is fully developed. This periodic component grows weaker and is no longer detectable at approximately the same time that the flux-stress curve achieves steady-state and suggests that the correlation scale of the vertical field may affect the transition point in the flux-stress relationship. These observations further suggest that residual flux from the initial conditions is responsible for the long term variability.

Assuming that a long-lived residual of the initial magnetic field is the driving mechanism for this phenomenon, we can recover the time required to achieve the steady state from elementary arguments. The time needed to turbulently diffuse together two patches of oppositely directed flux separated by a radial distance $\Delta r =$

$2.5h$ is given by

$$t_{\rm mix} \sim \frac{\Delta r^2}{\eta_{\rm eff}}, \tag{3.13}$$

where $\eta_{\rm eff}$ is the effective turbulent resistivity. If we define $Pr_{\rm m,eff}$ as the effective turbulent magnetic Prandtl number (i.e. the ratio of the effective turbulent viscosity to the effective turbulent resistivity), we can write

$$\eta_{\rm eff} = Pr_{\rm m,eff}^{-1} \alpha_{\rm M} c_s h, \tag{3.14}$$

where c_s is the sound speed. We can then write the mixing time as

$$t_{\rm mix} \sim \frac{Pr_{\rm m,eff}}{2\pi\alpha_M} \left(\frac{\Delta r}{h}\right)^2 t_{\rm orb}, \tag{3.15}$$

where $t_{\rm orb}$ is the local orbital period and we have used the fact that $h/c_s \sim r/v_\phi \sim t_{\rm orb}/2\pi$. Using $Pr_{\rm m,eff} = 1$ (Guan et al., 2009; Lesur & Longaretti, 2009; Fromang & Stone, 2009) and $\alpha_M = 0.005$ suggests that the memory of the initial conditions will be lost on a timescale of $t_{\rm mix} \sim 200 t_{\rm orb}$. This crude estimate is in reasonable agreement with the timescale on which we see the flux-stress relationship achieve a stationary state.

3.6 Conclusions

It has been a long-held ansatz that one can extract and model the dynamics of a local patch of an accretion disk and obtain results (for the angular momentum transport, for example) that have meaning for the disk as a whole. By examining local patches of a moderately high resolution global disk simulation, we have provided a direct test of this notion. We have shown that MRI-driven turbulence in

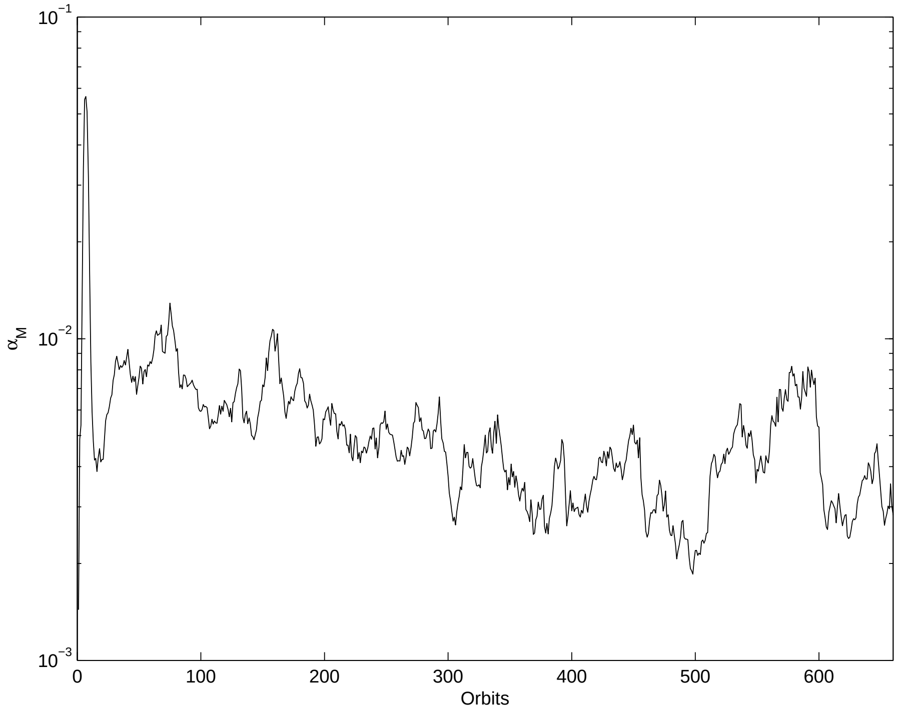

Figure 3.7: Long term behavior of α_M for run `Thin.M-Res.6z`.

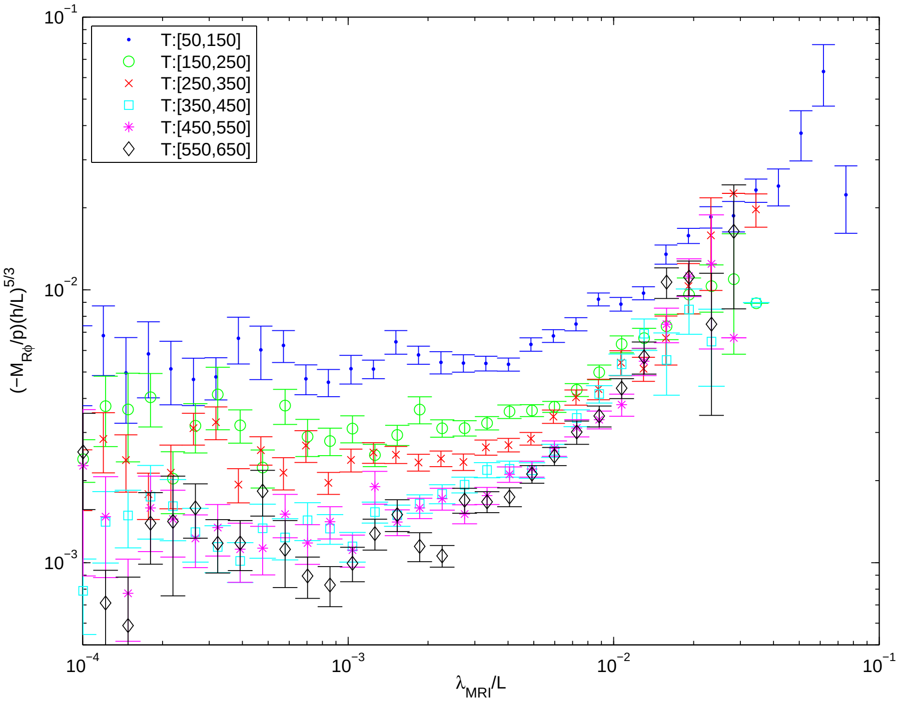

Figure 3.8: Flux versus stress averaged over 100 orbit blocks for run `Thin.M-Res.6z`. Vertical centering of $z = h$.

global geometrically thin accretion disks behaves in a way consistent with scaling laws derived for local simulations. In particular, we find that global disks display a local flux-stress relation qualitatively similar to that found in local simulations (Hawley et al., 1995; Pessah et al., 2007). However, other aspects of the global models are distinctly different.

Even though we model a global accretion disk that has zero net magnetic field, any given patch of the disk is threaded by a net magnetic flux resulting from the self-generated field in the MRI-dynamo. Across much of the disk, the local flux is strong enough to have a controlling effect on the local stress. Thus, our *zero-net field* global disk is behaving as a collection of *net field* local patches.

The normalization, slope, and location of the "knee" of the flux-stress relationship changes with vertical height in the accretion disk. This amplifies the role of the off-midplane region ($h < |z| < 2h$) of the disk; not only does this region have stronger vertical magnetic fields than the midplane, but a given vertical field induces stronger magnetic stresses. The result is a strong enhancement of magnetic stress well off the midplane of the disk. The transition point (or "knee") in the flux-stress relation occurs as significantly smaller fluxes in the global simulation as compared to the local unstratified simulations. We relate this transition point to the ability of the simulation to marginally resolved the slowest appreciably growing mode. Angular momentum transport (i.e. α_M) in the global disk appears to be impeded if the vertical domain size of the simulation is too small; we found significant differences between our $z = \pm 5h$ and $z = \pm 10h$ cases. On the other hand, the $z = \pm 10h$ and $z = \pm 20h$ cases appear very similar suggesting convergence has been achieved. Given that magnetic linkages between different patches of the disk appear to be crucial for determining the local flux (and hence the local stress), and that such linkages are made through the low-density corona of the disk, such a sensitivity to the vertical domain size is not surprising.

Analysis of our long simulation (which ran for 664 ISCO orbits) reveals long-term secular trends. In particular, there is a secular variation of the flux-stress relationship such that the knee of the relationship moves to smaller fluxes and the low-flux normalization decreases. This secular drift stabilizes after approximately 300 ISCO orbits. We attribute this to stresses associated with a long-lived residual of the initial magnetic field configuration.

The desire to make comparisons between local and global models is complicated due to the difficulty of appropriately resolving accretion disk turbulence in the significantly larger spatial domain associated with global models and by the need to disentangle the effects of stratification from the MRI. Toward the end of developing a suite of simulations that offer a better comparison to unstratified local models, further work is done using the superior `Athena` algorithm. Comparisons of local models run using `Zeus` and `Athena` (Stone, 2009) find that at a fixed resolution `Athena` offers superior accuracy and less dissipation while also precisely, *i.e.* to machine precision, respecting the conserved quantities. The next chapter will detail a semi-Lagrangian extension developed to complement the sophisticated algorithm of `Athena` that allows order of magnitude performance benefits when running accretion disk simulations. This combination of high-order Godunov algorithm and semi-Lagrangian extension has allowed the study of unstratified global disks at resolutions comparable to local models and unprecedented in the context of global simulations.

Chapter 4

Orbital Advection

4.1 Introduction

Orbital advection, first proposed for hydrodynamic systems by Masset (2000) and extended to MHD by Johnson et al. (2008), utilizes a semi-Lagrangian approach to alleviate the dominant timestep restriction associated with Eulerian simulations of supersonically rotating disks. The essence of orbital advection is the observation that in accretion disk simulations the dominant signal speed is associated with the Keplerian rotation. Thus, the timestep of an Eulerian code will be limited by $\Delta t \leq \Delta t_K$, with the Keplerian timestep given by equation 2.11. The background Keplerian rotation, however, is time-independent and has a simple spatial dependence which suggests the existence of a simple Lagrangian frame, referred to as the Keplerian frame. In this frame, assuming a weak magnetic field, the timestep will be limited by $\Delta t \leq C_0 R \Delta\phi / c_s$. Thus, solving the equations in the Keplerian frame results in a potential speedup corresponding to the maximum Keplerian mach number (M_K; equation 1.11) over the simulation domain. Orbital advection is an operator-split method in which the equations of MHD are solved in the Keplerian frame followed by a remap, or advection, step in which the Keplerian frame is mapped back onto the initial Eulerian grid. This section will describe an implementation of orbital advection for use in simulating global accretion disks utilizing the `Athena` algorithm in cylindrical geometry. The use of orbital advection and the typically order of magnitude speedup that it facilitates has allowed the exploration of accretion disk turbulence in global models at resolutions that were previously computationally inaccessible.

4.2 Derivation and Implementation

A necessary first step in developing orbital advection is to define the Lagrangian frame which is taken to be the frame induced by the rotation profile $\Omega_L(R)$. In what follows, it will be assumed that the Lagrangian frame is the Keplerian frame, *i.e.* $\Omega_L = \Omega_K$ (equation 1.8), and that the disk is unstratified. It is worth noting that the actual implementation is far more flexible than suggested by this assumption and that any user-specified rotation profile is admissible so long as it depends solely on the cylindrical radius. In the case of stratified disks $\Omega_L(R) = \Omega_K(R, z = 0)$, and the residual gravitational force and vertical gravity are taken into account.

We begin by noting that the equations of MHD are nearly linear in velocity and further, the Keplerian velocity is solenoidal. These factors suggest decomposing the velocity into its Keplerian component and remaining perturbation as in equation 1.30. Inserting this decomposition into the isothermal MHD equations and simplifying results in

$$\frac{\partial \rho}{\partial t} + \nabla\cdot(\rho \mathbf{v}') + \Omega_K \frac{\partial \rho}{\partial \phi} = 0 \tag{4.1}$$

$$\frac{\partial (\rho \mathbf{v}')}{\partial t} + \nabla\cdot(\rho \mathbf{v}'\mathbf{v}' - \mathbf{B}\mathbf{B} + P^{*}\mathbf{I}) + \Omega_K \frac{\partial (\rho \mathbf{v}')}{\partial \phi} = \mathbf{F_c} \tag{4.2}$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times \epsilon' - \nabla \times \epsilon_K = 0, \tag{4.3}$$

where $\epsilon' = \mathbf{v}' \times \mathbf{B}$ is the perturbation EMF and $\epsilon_K = v_K \times \mathbf{B}$ is the Keplerian EMF. The non-inertial forces associated with the Keplerian frame are given by $\mathbf{F_c}$ and defined as

$$\mathbf{F_c} = 2\Omega_K \rho v'_\phi \hat{R} + \Omega_K (q-2) \rho v'_R \hat{\phi}, \tag{4.4}$$

where q is the shear parameter (equation 2.27) and may, in general, have radial dependence. The assumption of Keplerian rotation results in the cancellation of the centrifugal force associated with the rotating frame and the radial gravitational force. Equation 4.4 includes both the Coriolis force associated with the Keplerian frame and centrifugal force associated with the perturbation velocity.

The structure of this system is suggestive, in that each of the initial conservation laws is now in the form of a conservation law with a flux set by the perturbation velocity, $\mathbf{v}'$, and an additional linear advection term. The induction equation is modified in an analogous manner, with the evolution of the magnetic field coming from the EMF in the rotating frame and a Keplerian EMF. To solve the resulting system numerically, we consider an operator-split method that decomposes the equations into a system of linear advection equations and an MHD system (with Coriolis and centrifugal source terms) with a characteristic velocity given by $\mathbf{v}'$. The advection system is given by

$$\frac{\partial \rho}{\partial t} + \Omega_K \frac{\partial \rho}{\partial \phi} = 0 \tag{4.5}$$

$$\frac{\partial (\rho \mathbf{v}')}{\partial t} + \Omega_K \frac{\partial (\rho \mathbf{v}')}{\partial \phi} = 0 \tag{4.6}$$

$$\frac{\partial \mathbf{B}}{\partial t} - \nabla \times \epsilon_K = 0, \tag{4.7}$$

while the MHD system is given by

$$\begin{aligned} \frac{\partial \rho}{\partial t} + \nabla\cdot(\rho \mathbf{v}') &= 0 && (4.8)\\ \frac{\partial (\rho \mathbf{v}')}{\partial t} + \nabla\cdot(\rho \mathbf{v}'\mathbf{v}' - \mathbf{B}\mathbf{B} + P^{*}\mathbf{I}) &= \mathbf{F_c} && (4.9)\\ \frac{\partial \mathbf{B}}{\partial t} - \nabla \times \epsilon' &= 0. && (4.10) \end{aligned}$$

Prior use of orbital advection focused on either local or global hydrodynamic disks, or local MHD disks in the context of the `Zeus` algorithm. Extending the use of orbital advection to conservative second-order algorithms like `Athena` requires care. When solving the advection system it is necessary to ensure that the algorithm is conservative and preserves the solenoidal constraint of the magnetic field. Further, introducing new source terms into the equations of MHD requires a treatment that is second-order accurate and includes the geometric source terms already present. This is done through an extension of the method described by Stone & Gardiner (2010), an orbital advection algorithm for use with shearing box simulations in `Athena`.

4.2.1 Splitting Methodology and Timestep Constraint

The implementation that will be described here uses Godunov splitting, in which the equations are first evolved for Δt via the advection step and the results are used as the initial conditions to evolve the MHD system in the Keplerian frame for Δt. At first glance, this is problematic as the temporal order of accuracy of the `Athena` algorithm is second-order whereas Godunov splitting is formally only first-order. This can be rectified by using Strang splitting, in which the advection system is evolved for $\Delta t/2$, followed by a full Δt update of the MHD system, and finalized with another $\Delta t/2$ update from the advection system. Given the small cost of solving the advection system compared to the expensive Riemann solvers used twice every timestep for each cell in the MHD system, Strang splitting performs almost as well as Godunov splitting. In practice, however, there appears to be little benefit to using Strang splitting as the coefficient associated with the first-order error of the Godunov splitting is so small.

To define the timestep constraint associated with the full algorithm we begin by considering the timestep limitations of the individual systems. The MHD system is governed by the standard CFL condition, however in this case it is the CFL condition associated with the perturbed velocity $\mathbf{v}'$ which will be considerably more lenient. The advection system, made up of a series of linear advection equations, has no associated timestep constraint as there is a simple analytic solution. While it is indeed true that the timestep constraint of the split method is given simply by the CFL condition associated with the perturbed velocity $\mathbf{v}'$, this is not a direct consequence. The constraint of splitting the equations must be taken into account, although as will be demonstrated the CFL condition remains the dominant timestep constraint.

As pointed out by Masset (2000) there is a timestep constraint associated with the shear itself. This can be derived by considering two radially adjacent cells, centered at R_i and $R_{i+1} = R_i + \Delta R$, and the action of the advection system upon them. The timestep must be constrained so that the shear in the Lagrangian frame will not disconnect two radially adjacent cells. This ensures that the physical domain of dependence, *i.e.* the backwards "light"-cone, does not become disjoint from the numerical domain of dependence, *i.e.* the discretized grid values used in the numerical update. Assuming uniform grid spacing this can be written formally as

$$\left[\Omega_K(R_i) - \Omega_K(R_{i+1})\right] \Delta t \leq \frac{\Delta\phi}{2} \tag{4.11}$$

where this must be true for all indices i. This can be seen as a simultaneous constraint on both the radial resolution and the shear parameter as coarse radial resolution and large values of the shear parameter would allow this condition to be violated. Using the estimate $\Delta t \leq 0.5R\Delta\phi/c_s$ a sufficient condition for equation 4.11 to be satisfied is for the following to hold over the entire simulation domain

$$\frac{H}{\Delta R} \geq \sqrt{2}q, \tag{4.12}$$

where, in general, the scale height (H; equation 1.10) and the shear parameter have radial dependence. For the case of a Newtonian disk, $q = 3/2$ and the condition becomes

$$\frac{H}{\Delta R} \geq 2.12. \tag{4.13}$$

This condition is quite lenient as having radial resolution coarser than this is unlikely to be sufficient to adequately resolve turbulent structure.

4.2.2 Advection System

Solving the advection system, equations 4.5 through 4.7, is done in the manner of Stone & Gardiner (2010). The conserved quantities, mass and momentum, are updated utilizing an interface flux and the magnetic field is updated consistent with CT using the Keplerian EMF, ϵ_K. The evolution equations are essentially one dimensional, as the solution of the linear advection equation is merely an azimuthal advance of the initial data by $\Omega_K \Delta t / \Delta\phi$ grid cells. In the case that this quantity is an integer, a simple cyclic permutation of the grid cells would be sufficient to solve the evolution equations. In general this will not be the case and a reconstruction of the azimuthal profiles of the quantities is necessary. This is done utilizing a third-order piecewise-parabolic spatial reconstruction used in the `Athena` algorithm, specifically the reconstruction described in (Colella & Sekora, 2008).

Given the one dimensional nature of the evolution equations, for simplicity, the update can be defined at one ring of fixed radius and height. As a matter of notation, we define the azimuthal reconstruction of a conserved quantity q at fixed indices i and k, where these indices will be suppressed, as $\hat{q}(\phi)$. The flux at each azimuthal interface is defined as the upstream integral in the manner

$$f_{q,j-1/2} = \int_{\phi_{j-1/2}-\Omega_K \Delta t}^{\phi_{j-1/2}} \hat{q} d\phi. \tag{4.14}$$

In general, the domain of dependence, $\phi_{j-1/2} - \Omega_K \Delta t$, will contain a number of full grid cells as well as a fractional region of a grid cell. For cells fully contained within the domain of dependence the volume-average is simply added to the flux; integration of the reconstructed profile is only necessary for cells that straddle the domain of dependence. With these interface fluxes calculated the conserved quantities can be updated in a conservative manner using equation 2.3.

Evolving the magnetic fields at each interface is done consistent with CT based on the Keplerian EMF, $\epsilon_K = v_K B_z \hat{R} - v_K B_R \hat{z}$. At each edge an appropriate time-advanced EMF is calculated using an upstream integral similar to the calculation of the flux. After reconstructing the interface magnetic fields, $\hat{B_R}(R_{i-1/2}, \phi, z_k)$ and $\hat{B}_z(R_i, \phi, z_{k-1/2})$, the edge-centered Keplerian EMF is (suppressing the subscript K)

$$\begin{aligned} \epsilon^{n+1/2}_{R,i,j-1/2,k-1/2} &= v_K(R_i) \int_{\phi_{j-1/2}-\Omega_K \Delta t}^{\phi_{j-1/2}} \hat{B}_z(R_i, \phi, z_{k-1/2}) d\phi \qquad (4.15) \\ \epsilon^{n+1/2}_{z,i-1/2,j-1/2,k} &= -v_K(R_{i-1/2}) \int_{\phi_{j-1/2}-\Omega_K \Delta t}^{\phi_{j-1/2}} \hat{B_R}(R_{i-1/2}, \phi, z_k) d\phi. \qquad (4.16) \end{aligned}$$

These edge-centered EMFs are then used to update the interface fields in the standard manner of CT described in Chapter 2, save for the addition of appropriate geometric scaling terms to account for cylindrical geometry.

Ensuring this implementation of orbital advection is capable of parallelization is paramount, as the improved performance would provide little benefit restricted to only one processor. Parallelization is accomplished in the standard manner, where the spatial domain is decomposed into tiles and each tile is assigned to a particular processor. Ensuring the proper evolution of the system is accomplished through the use of “ghost zones”, where variables from spatially-adjacent tiles are attached to the processor’s computational domain. This is a standard practice in parallelized codes, however the use of orbital advection requires a larger number of ghost zones to account for the potentially large azimuthal upstream domain. This requires, where $N_{g\phi}$ is the number of ghost cells in the azimuthal domain for a given tile, that

$$N_{g\phi} \geq \frac{\Omega_K \Delta t}{\Delta \phi}, \qquad (4.17)$$

hold over the entire tile. It is inefficient to reallocate memory for ghost zones at every timestep, so instead a bound on each MPI tile is calculated at the beginning of the simulation and used for the duration of the run. Using the fact that the timestep is bounded by the azimuthal sound-crossing time for a grid cell allows

$$N_{g\phi} = \lceil \max_{R}(C_0 M_K) \rceil + 1, \tag{4.18}$$

the maximum Keplerian Mach number over the given tile. It is worth noting that for very thin disks, where $H/R \ll 1$, the required number of ghost zones would be very large and require a great deal of memory. In this case using dynamic memory allocation to only gather needed data from neighboring MPI tiles may become necessary. This completes the description of the solution of the advection system, and we now turn our attention to the second operator in the algorithm.

4.2.3 Non-Inertial MHD System

Solving the equations of MHD in the non-inertial frame proceeds in a manner quite similar to the standard MHD integrator save for three points: the radial gravitational force is removed[1]; the perturbation velocity, $\mathbf{v}'$, is used instead of the standard Eulerian velocity; and source terms associated with the rotating frame must be included. Transitioning to the perturbation velocity is a simple matter of passing only the perturbed velocity to the MHD integrator. Including the source terms properly, however, requires greater care.

The rotating frame only introduces new source terms to the azimuthal and radial momentum equations of MHD, equations 2.24 and 2.25 in cylindrical coordinates. Defining $m = \rho v_R$ and $n = \rho v_\phi$ while letting f_m and f_n correspond to the

---

[1] Note, that this is under the assumptions of this section, namely that we are in the Keplerian frame. In practice, the radial gravitational force is modified to subtract the centrifugal force associated with the rotating frame.

appropriately scaled sum of tensor derivatives results in the evolution equations for the planar momenta becoming

$$\begin{aligned} \frac{\partial m}{\partial t} + f_m &= 2\Omega_K n + \frac{n^2}{\rho R} + \frac{P^* - B_\phi^2}{R}, & (4.19) \\ \frac{\partial n}{\partial t} + f_n &= \Omega_K (q-2) m. & (4.20) \end{aligned}$$

Properly including source terms into the predictor-corrector formalism of `Athena` can be difficult due to the numerous places within the algorithm where source terms must be applied. Source term corrections must be applied to the interface states, when evolving from t_n to $t_{n+1/2}$, and when using the time-centered fluxes to compute the final evolution from t_n to t_{n+1}. Save for the last case, applying the source terms is done at first-order accuracy using simple forward Euler mimicking the treatment of source terms by Stone & Gardiner (2010) and Skinner & Ostriker (2010).

The full timestep update, corresponding to equation 2.23, requires second-order temporal accuracy. Skinner & Ostriker (2010) treat the geometric source term in equation 2.24 using the time-centered source terms requiring the time-centered value of the azimuthal velocity, itself not subject to source terms. The coupled nature of equations 4.19 and 4.20 requires a unified approach. Stone & Gardiner (2010) apply the Cartesian non-inertial source terms, both linear in m and n, using the Crank-Nicholson method. The quadratic dependence on n in equation 4.19 makes this approach unwieldy. Instead, equations 4.19 and 4.20 are jointly evolved using a two-step Runge-Kutta method. This is a predictor-corrector method, where a Δt forward Euler advance is the predictor, and the trapezoidal method utilizing the arithmetic average of the initial data and predicted value is used to correct. While this does not explicitly conserve the energy of epicyclic oscillations in a stress-free medium in the manner of Crank-Nicholson, as pointed out by Stone & Gardiner (2010) this is relatively unimportant in the context of MHD as epicyclic oscillations are unstable.

4.3 Tests of Orbital Advection

The bulk of the tests discussed here are reproductions of tests used by Skinner & Ostriker (2010) to validate the cylindrical `Athena` algorithm. We begin by presenting a force-balance problem to demonstrate the formal order of accuracy of the method. Next, the advection of a magnetic field loop in cylindrical geometry is presented to demonstrate the accuracy of the constrained transport implemented in the orbital advection module. Finally, we present three somewhat more qualitative tests and compare both the results and speed with the standard cylindrical integrator. These qualitative tests are of particular interest in that they all deal with the stability of rotating disks. The first is a study of the Rayleigh stability criterion where we perturb uniform hydrodynamic disks (with varying values of q) and test that the Rayleigh criterion (stability when $q < 2$) is respected. Finally, we present two studies of the MRI in accretion disks. We consider the evolution and saturation of the MRI in both a uniform, unstratified Newtonian disk and an unstratified disk in equilibrium in a psuedo-Newtonian potential.

4.3.1 Force Balance

The first test we consider is a static force balance problem in MHD. We begin by noting that the purely azimuthal field,

$$B_\phi = \frac{B_0}{R}, \tag{4.21}$$

is such that the magnetic pressure and tension forces balance. We construct a uniform ($\rho_0 = 1$) disk rotating as a solid body, with potential

$$\Phi = \frac{(\Omega_0 R)^2}{2}, \tag{4.22}$$

where $\Omega_0 = 2\pi$ and introduce a uniform vertical velocity ($v_z = 0.1$). The resulting system is stationary, and as such the analytic solution is merely the initial condition. The magnetic field is initialized with $B_0 = 1$, and the isothermal sound speed is $c_s = 0.1$. The physical domain of the simulation is $(R, \phi, z) \in [1, 2] \times [0, \pi] \times [0, 1]$ and is evolved until $t = 10$.

We define the L_1 error of a physical quantity q by,

$$L_1(q) = \Sigma_{i,j,k} |q_{i,j,k}(t = 10) - q_{i,j,k}(t = 0)|. \tag{4.23}$$

The RMS error is the Euclidean norm of the vector of L_1 errors for each physical quantity. The goal of this test is to verify that the algorithm converges at second order, in other words that by doubling resolution the RMS error will decrease by a factor of four. The simulations are run with increasing resolutions of the form $N_R = N_\phi = N$ and $N_z = N/4$, where $N \in \{16, 32, 64, 128, 256\}$. The results are given in Figure 4.1 along with results using the cylindrical integrator kindly provided by Aaron Skinner. This verifies that convergence is achieved at the expected order and that the rate of convergence and accuracy is comparable to that of the cylindrical integrator alone.

4.3.2 Field Loop Advection

To test that the implementation of the constrained transport update to the magnetic field used in the orbital advection update has been correctly modified for cylindrical geometry, we consider the advection of a weak field loop. The physical domain is given by $(R, \phi, z) \in [1, 2] \times [-2/3, 2/3] \times [0, 1]$ using a resolution of $N_R = N_\phi = N_z = 128$. The simulation is initialized with a uniform density ($\rho_0 = 1$), a solid-body rotation profile ($v_\phi = \Omega_0 R$), and a uniform $v_z = 1$. The simulation is run for two orbits, where $\Omega_0 = 4/3$ so that $P_{orb} = 1$. The equation of state is given by setting the isothermal sound speed, $c_s = 1$. Note that this problem is not rotationally dominated, the timestep constraint is given by the uniform vertical velocity. The magnetic field is defined via a magnetic vector potential

$$A_z = \times \left\{ \begin{array}{ll} A_0(a_0 - r) & : r \leq a_0, \\ 0 & : r > a_0, \end{array} \right\} \tag{4.24}$$

with $A_R = A_\phi = 0$, $a_0 = 0.3$, and r representing the distance from the center of the simulation domain. To ensure that the field loop is passively advected, the magnetic pressure of the field loop is constrained to be much smaller than the ambient gas pressure by setting $A_0 = 10^{-3}$. This formulation is identical to that used by Skinner & Ostriker (2010).

A comparison of the results is given in Table 4.1. The L_1 errors in each of the physical quantities are presented, *i.e.* $L_1(q) = \Sigma|q(t = 2) - q(t = 0)|$, and the RMS error is given by the Euclidean norm of the individual L_1 errors. As this is fundamentally an advection problem, it is unsurprising that in each category orbital advection performs either comparably or superior to the standard cylindrical integrator. It is also worth noting that since the problem presented here is not a smooth flow (due to the discontinuous magnetic field), the RMS error does not converge at second order; the convergence is first order for both the cylindrical integrator and orbital advection. The next comparison is based on the dissipation of magnetic energy, the evolution of $< B^2/2 >$ is plotted in Figure 4.2. After two orbits the simulation run using the cylindrical integrator has dissipated 6.18% of its initial magnetic energy in comparison to the orbital advection run which has dissipated 6.92%. While the orbital advection algorithm is slightly more dissipative, the results are comparable.

4.3.3 Rayleigh Stability Criterion

Here we consider differentially rotating systems with $\Omega(R) = \Omega_0 R^{-q}$. Rayleigh's criterion states that systems of this form will be stable to small axisymmetric perturbations for $q < 2$ and unstable otherwise. Simulations of differentially rotating systems at the cusp of stability are used to demonstrate that the momentum source terms are being correctly implemented. To this end, the values of q considered are, $q \in \{1.95, 1.99, 2.01, 2.05\}$, using both orbital advection and the standard cylindrical

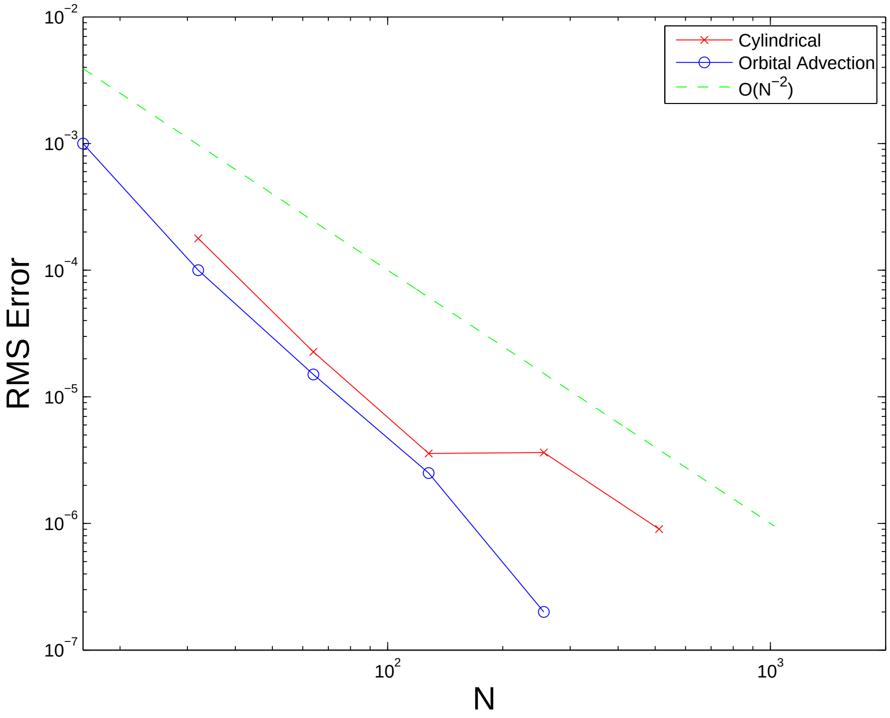

Figure 4.1: Convergence of RMS error of an MHD force balance problem with increasing resolution.

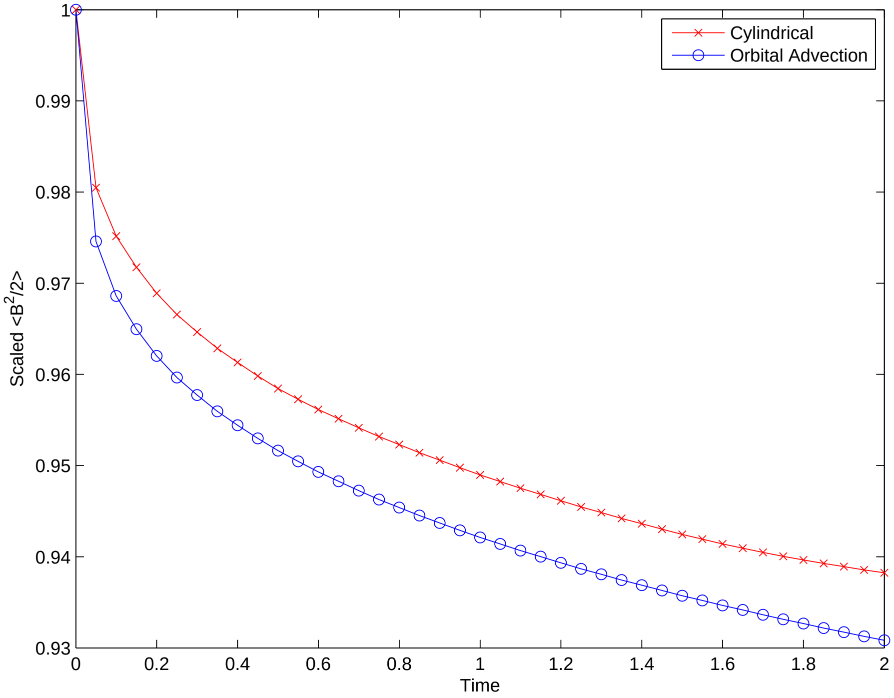

Figure 4.2: Comparison of the dissipation of magnetic energy (scaled to unity at $t = 0$) using the standard cylindrical integrator (red) and orbital advection (blue).

Method	$L_1(\rho)$	$L_1(M_R)$	$L_1(M_\phi)$	$L_1(M_z)$	$L_1(B_R)$	$L_1(B_\phi)$	$L_1(B_z)$	RMS Error
Orbital Advection	1.57×10^{-6}	1.22×10^{-7}	1.46×10^{-7}	1.57×10^{-6}	9.35×10^{-6}	5.97×10^{-6}	6.18×10^{-18}	4.53×10^{-5}
Cylindrical	1.67×10^{-6}	1.21×10^{-7}	3.37×10^{-6}	1.67×10^{-6}	8.38×10^{-6}	7.07×10^{-6}	1.78×10^{-16}	4.69×10^{-5}

Table 4.1: Comparison of the errors after two orbits in the field loop advection problem.

integrator.

The physical domain simulated is given by $(R, \phi, z) \in [3, 7] \times [0, \pi/2] \times [-0.5, 0.5]$, using a resolution of $(N_R, N_\phi, N_z) = (128, 200, 32)$. The simulation is initialized using a constant density $\rho_0 = 200$, with $\Omega_0 = 2\pi$, and an isothermal sound speed of $c_s = 0.1$ which results in a flow that is highly rotationally dominated. Note that this is just a three-dimensional extension of the Rayleigh stability test presented in Skinner & Ostriker (2010). As in Skinner & Ostriker (2010), the instability is seeded using perturbations to the azimuthal velocity of the form, $v_\phi = v_K(1 + \Delta)$, where $\Delta \in [-10^{-4}, 10^{-4}]$ and is uniformly distributed. The simulations are evolved to $t = 300$; the number of cycles taken and the wall-time to complete each simulation is given in Table 4.2.

Taking as an example the case $q = 1.95$, corresponding to a Newtonian disk, the simulation without orbital advection takes approximately 24 times as long to run and processes approximately the same number of cycles per minute. Even in the case of an unstable disk, *e.g.* $q = 2.05$, where mass is driven off the grid there is almost an order of magnitude performance increase.

To measure the instability of these systems, we consider a radial scaling of the Reynolds stress normalized to the gas pressure given by

$$\frac{< R\rho v_R v'_\phi >}{< RP >} = \frac{\int R\rho v_R (v_\phi - v_K) dV}{c_s^2 \int R\rho dV}, \tag{4.25}$$

where the integrals are performed over the entire domain. The time evolution of this quantity for all simulations is shown in Figure 4.3. As discussed by (Skinner & Ostriker, 2010), this stress should remain very low, for $q < 2$, if the method is to be sufficiently

q	Cylindrical Integrator (Cycles : Wall-time)	Orbital Advection (Cycles : Wall-time)
$q = 1.95$	72.9 : 4.7h	3.2 : 11.7m
$q = 1.99$	69.9 : 4.6h	3.2 : 12m
$q = 2.01$	68.6 : 4.5h	4.5 : 17.4m
$q = 2.05$	65.4 : 4.9h	9.6 : 35.8m

Table 4.2: Comparison of the performance of the cylindrical integrator versus orbital advection. Specifically shown are cycles (in thousands) and wall-time (in h[ours] and m[inutes]). The simulations were run on Ranger using 64 processors and an MPI topology of $(T_R, T_\phi, T_z) = (8, 4, 2)$.

accurate that only physical, rather than numerical, effects are transporting angular momentum. There is a clear difference in the evolution of the stress depending on the shear parameter as expected. For $q > 2$ (unstable) we see that the scaled stress grows by many orders of magnitude, whereas in the case of $q < 2$ (stable) we see that the scaled stress remains fairly uniform. Even near the cusp of the stability criterion, $q = 1.99$ and $q = 2.01$, orbital advection accurately discerns between stability and instability while providing a significant performance boost.

4.3.4 Evolution and Saturation of the MRI in a Newtonian Disk

The first accretion disk test is done in the context of an unstratified disk in a Newtonian potential. The model used for this test forms the basis of a suite of simulations that will be analyzed in detail throughout the remainder of this work. The simulation model is discussed in detail in the following chapter and corresponds to run `BzZ8` in that notation, and will merely be summarized here. The disk is initially threaded with a sinusoidal vertical magnetic field such that the total magnetic flux through the simulation domain is zero, and initialized with a constant density, a Keplerian rotation profile and small vertical and azimuthal velocity perturbations. The Keplerian Mach number, $M_K = 20$, at the inner edge of the simulation domain and thus is the expected limitation on the performance boost obtainable with orbital advection.

Figure 4.4 shows a comparison of the evolution of the volume-integral of α_M (equation 1.32) for simulations run with and without orbital advection. The linear growth of the instability is precisely matched and the agreement remains strong even into the strongly nonlinear regime. Finally, and perhaps most importantly, the quantitative benefit of orbital advection is addressed. This is quantified through a comparison of the timesteps used in the evolution of each simulation. The timesteps are sampled ten times per orbit, and denoted Δt_C and Δt_F for the cylindrical and orbital advection runs, respectively. The speedup is defined as

$$S(t) = \frac{\Delta t_F}{\Delta t_C}. \tag{4.26}$$

The time dependance of the speedup is shown in Figure 4.4. The timestep associated with orbital advection is significantly more volatile than without, as the timestep is associated with turbulent quantities as opposed to the time-invariant Keplerian rotation. In particular, the speedup drops significantly during the initial linear growth of the MRI due to the impulsive accretion and larger magnetosonic speeds associated with the exponential growth phase. Upon the saturation and breakup into turbulence, the speedup increases again steadily and overall maintains an order of magnitude increase in the allowable timestep.

4.3.5 Evolution and Saturation of the MRI in a Pseudo-Newtonian Disk

The next model considered is that of an unstratified accretion disk about a non-rotating black hole, and as such we use the psuedo-Newtonian potential of Paczynski and Witta given by,

$$\Phi = \frac{-1}{R - 2R_g}. \tag{4.27}$$

In the simulation units are normalized so that the gravitational radius, R_g, is unity. As in the Newtonian disk model, a system that is in hydrodynamic equilibrium is used. To this end, a disk with a gaussian surface density profile given by,

$$\rho = \rho_0 \exp\left(-\frac{(R - R_0)^2}{2L_R^2}\right), \tag{4.28}$$

is used. The radial dependence of this density profile leads to radial pressure forces which must be balanced to construct a disk in equilibrium. To this end a non-Keplerian velocity profile, u_ϕ, is constructed to balance out the pressure forces. Solving the radial force-balance equation yields,

$$\frac{u_\phi^2}{R} = \frac{1}{(R - 2R_g)^2} - \frac{c_s^2(R - R_0)}{L_R^2}. \tag{4.29}$$

The simulation is initialized using $\rho_0 = 1$, $L_R = 4$, $R_0 = 12$, and $c_s = 0.05$. The radial profiles of the density, u_ϕ, and v_K are given in Figure 4.5. In terms of (R, ϕ, z) the simulation domain is given by $[4, 19] \times [0, \pi/4] \times [-2, 2]$ with a resolution of $(N_R, N_\phi, N_z) = (128{,}64{,}64)$. The initial magnetic field is given by,

$$B_z = \frac{A_0 M_f(R; R_0, L_R)}{R} \sin\left(\frac{R - R_0}{L_R}\right), \tag{4.30}$$

where A_0 is a constant and M_f is a smooth mollifier function symmetric about R_0 that is unity at $R = R_0$ and zero for $|R - R_0| > L_R$. Specifically,

$$M_f(x; 0, 1) = \left\{ \begin{array}{ll} \exp\left(\frac{-x^2}{1-x^2}\right) & : |x| < 1 \\ 0 & : |x| \geq 1 \end{array} \right\}. \tag{4.31}$$

The vertical field is scaled by A_0/R so that $< B_z >= 0$ and A_0 is defined such that $\beta = 500$. Boundary conditions analogous to the simulations presented in Chapter 3 are used.

The instability is seeded using azimuthal velocity perturbations of the form

$$v_\phi = v_K[1 + \Delta M_f(R; R_0, L_R)\cos(\phi)\cos(z/\pi)], \tag{4.32}$$

where $\Delta = 0.01$. A comparison of the volume-integral of the Maxwell stress, $M_{R\phi}$, with and without orbital advection is given in Figure 4.6. Again, the agreement between the two is excellent, reproducing the double-peaked behavior of the initial transient almost precisely. The timestep evolution is compared in Figure 4.6, and shows a consistent boost in the allowable timestep. The enhancement of orbital advection is somewhat reduced in the psuedo-Newtonian case due to the existence of an ISCO, inside of which material will plunge inwards generally in excess of the sound speed.

4.4 Conclusion

The addition of the orbital advection algorithm to the cylindrical integrator leads to solutions with similar or better accuracy and an order of magnitude increase in performance. The combination of orbital advection with the cutting edge algorithm at the heart of `Athena` provides us with a new tool with which to explore accretion disk turbulence at unprecedented resolutions.

The significant cost associated with global models makes a rigorous treatment of resolution difficult. Often this leads to global simulations simply being run at the highest resolution affordable without appropriate consideration being paid to ensuring that the astrophysical implications of these simulations are robust and not strongly dependent on resolution. This has resulted in the impression, *e.g.* Hawley et al. (2011), that the vast majority of global simulations presented in the literature are unresolved or at best, under-resolved. In particular, Hawley et al. (2011) point to only one simulation (Noble et al., 2010) with adequate poloidal resolution and remark that sufficient azimuthal resolution is universally absent. Increasing azimuthal resolution is particularly costly due to the combination of more grid cells and the reduction of the Keplerian timestep. Because of this and the importance of resolving the vertical magnetic field, due to its association with the fastest growing MRI modes, the azimuthal resolution of global simulations has often been somewhat poor.

Orbital advection not only provides significant performance benefits, but because of its treatment of the MHD equations in the Keplerian frame it is amenable to the use of isotropic resolutions, meaning a grid with an aspect ratio of unity. The following chapter makes use of the model introduced in Section 4.3.4 as the basis of a convergence study utilizing a suite of simulations of varying initial magnetic field topologies and resolutions to diagnose the importance of resolution and provide metrics to diagnose converged MRI-driven accretion disk turbulence.

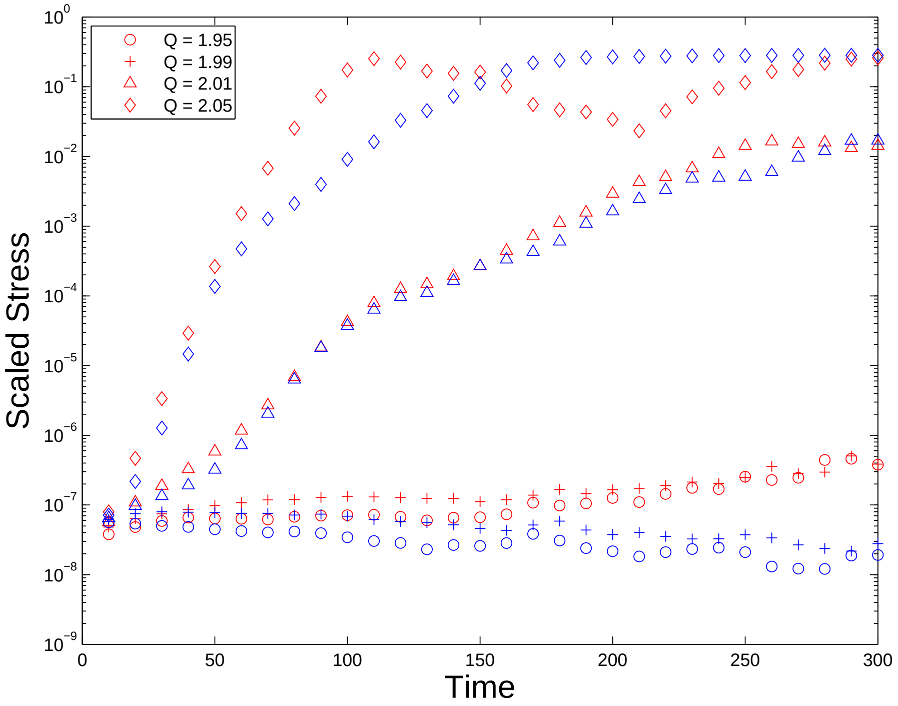

Figure 4.3: Evolution of the scaled stress (equation 4.25) using the cylindrical integrator (red) and orbital advection (blue).

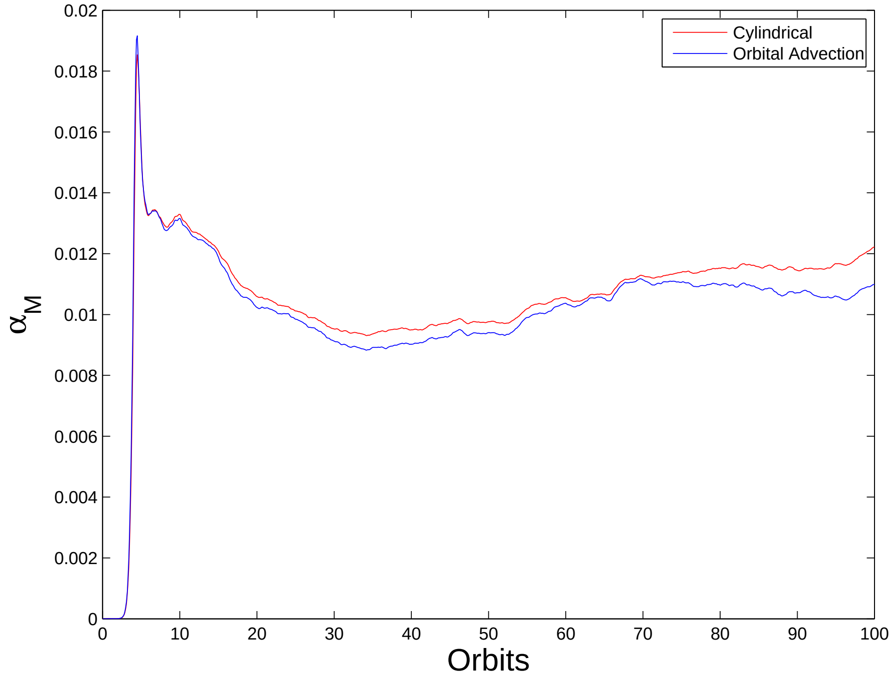

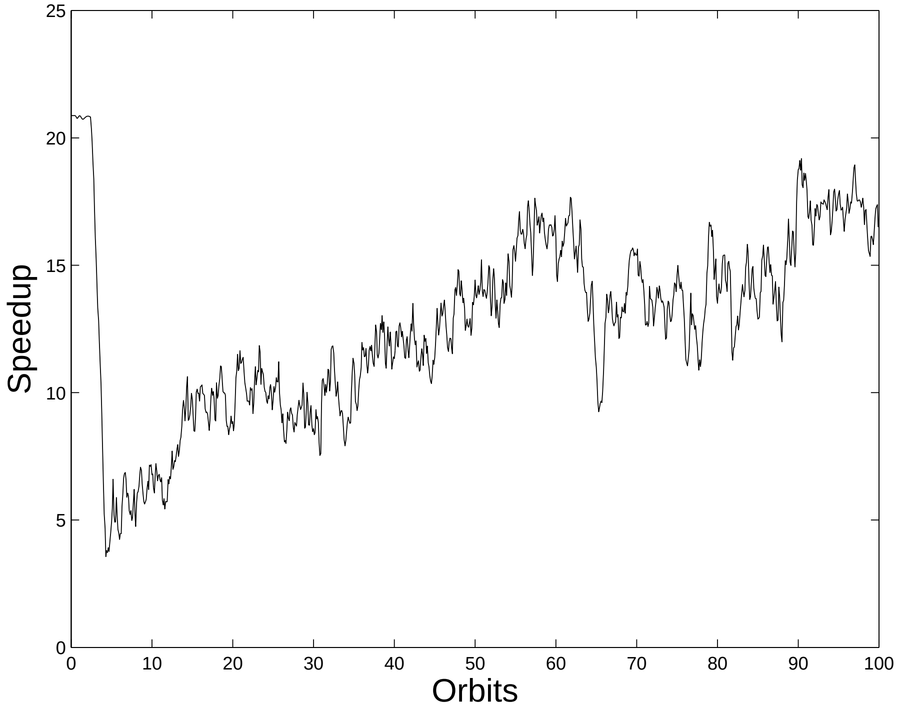

Figure 4.4: (**Top**) Comparison of disk evolution with and without orbital advection. (**Bottom**) Speedup using orbital advection.

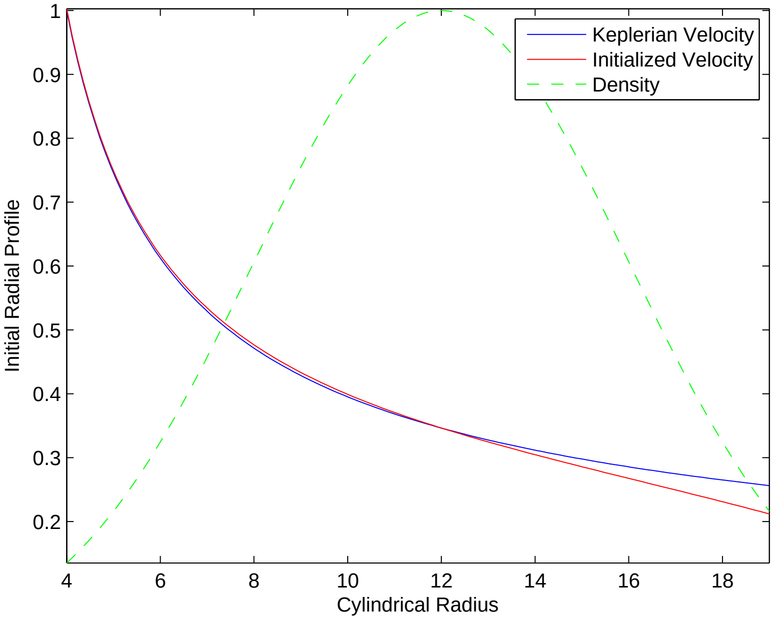

Figure 4.5: Initial radial density and velocity profile for the psuedo-Newtonian accretion disk model.

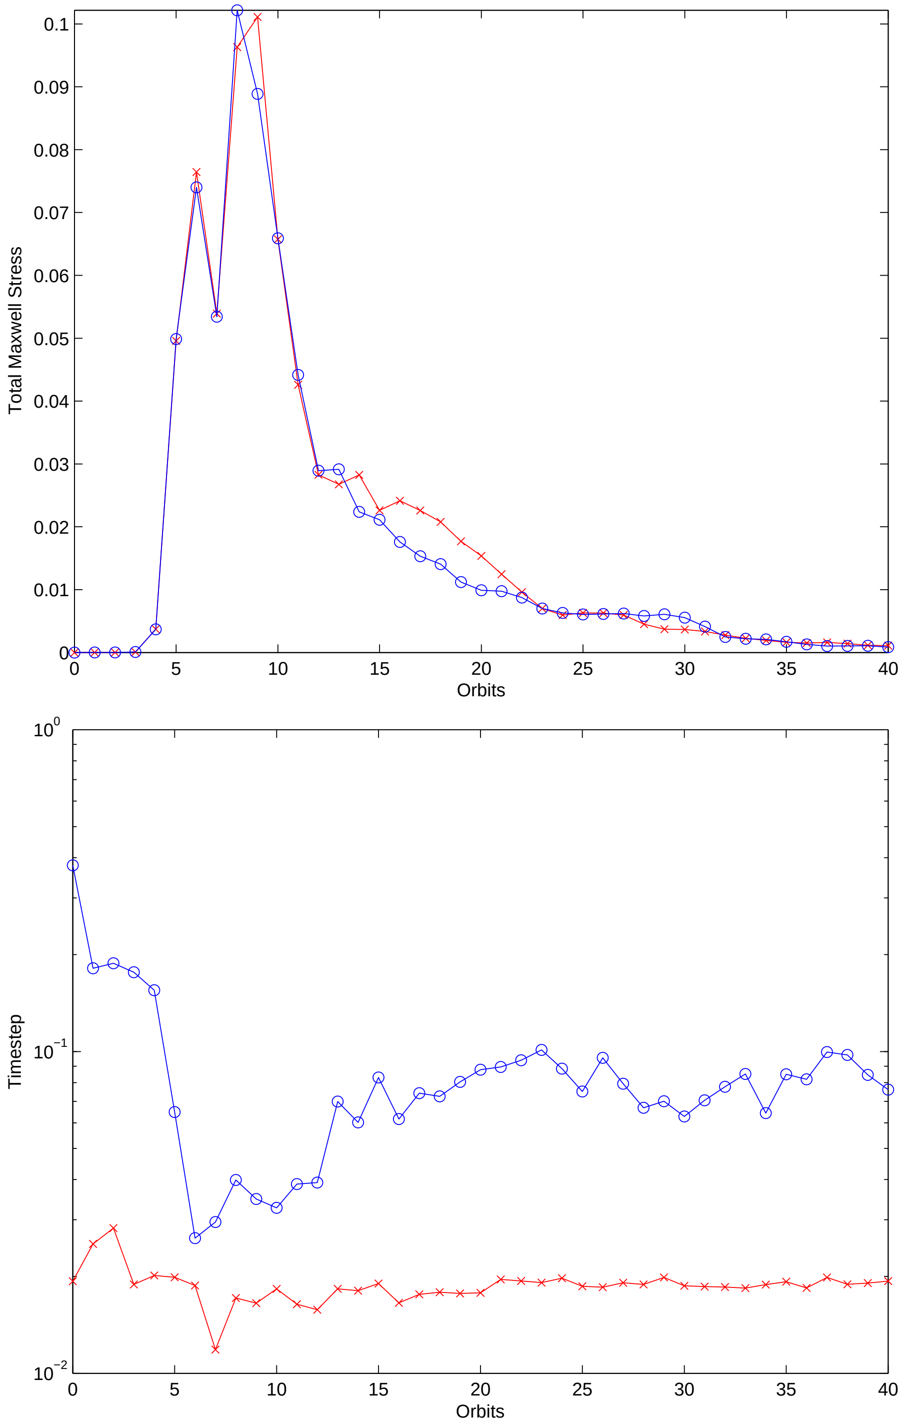

Figure 4.6: (**Top**) Comparison of disk evolution with (blue) and without orbital advection (red). (**Bottom**) Comparison of timesteps with (blue) and without orbital advection (red).

Chapter 5

Convergence of Global Models

5.1 Introduction

Numerical simulations are necessary to study accretion disk turbulence, but they are no panacea. Choices ones makes regarding the solution of the relevant equations and initial conditions can leave unintended and unphysical artifacts in the results. Understanding how to disentangle numerical artifacts from genuine physical phenomenon is vital to constructing simulations that can make reliable predictions. In particular, the discretization of the physical simulation domain, or resolution, will be of interest to us here. Concern was raised when it was discovered by Fromang & Papaloizou (2007) that local unstratified simulations without explicit dissipation, and without a net magnetic flux threading the domain, result in angular momentum transport that vanishes with increasing resolution. This lack of convergence is not robust as the inclusion of stratification, explicit dissipation, or a net magnetic field lead to a converged value of momentum transport (Davis et al. 2010, Fromang 2010). Further, Sorathia et al. (2010) (Chapter 3) observe that, to the extent that local models are predicated on being representative of a small patch of a global disk, the constraint of enforcing a zero magnetic flux over a small physical domain is likely unphysical, albeit an important theoretical pathology. These results highlight the importance of understanding the sometimes delicate nature of numerical convergence when dealing with simulations of turbulence.

The primary goal of this chapter is to study accretion disk turbulence induced by a variety of different initial field topologies in the context of global simulations of higher resolution than have previously been considered and to ascertain the importance of resolution. To reduce the significant expense associated with global, three-dimensional MHD simulations we use an implementation of orbital advection (Masset 2000, Johnson et al. 2008), which results in an *order of magnitude speed-up* for the simulations considered here. The implementation of this algorithm has been described in Chapter 4. These simulations are meant to be comparable with the surveys performed for local models by Hawley et al. (1995) and for global by Hawley (2001). As we will see, global simulations analogous to the local models considered by Fromang & Papaloizou (2007) do converge to non-zero angular momentum transport. Additionally, we consider the resolution requirements to attain convergence as well as explore the question of what convergence means in relation to non-dissipative simulations of turbulent systems.

The plan of this chapter is as follows. Sections §5.2.1 and §5.2.2 describes the simulations and diagnostics used in this and the following chapter. Section §5.2.3 compares the disk evolution under the influence of varying initial magnetic field topologies and discusses this evolution in the context of reduced anomalous viscosity models. Section §5.3 explores the use of a selection of convergence metrics: physical metrics (Section §5.3.1); numerical metrics (Section §5.3.2); and spectral metrics (Section §5.3.3). Section §5.3.4 presents results suggesting that the magnetic tilt angle is a robust indicator of convergence that is invariant of initial magnetic topology. A discussion of the convergence study with an emphasis on a comparison with the recent work of Hawley et al. (2011) is presented along with concluding remarks in §5.3.5.

5.2 Methodology and Diagnostics

5.2.1 Simulations

The work presented here is based on a series of simulations exploring the behavior of global accretion disks under a variety of resolutions and initial seed magnetic fields. Our simulations model an unstratified, isothermal, relatively cold Keplerian disk in a Newtonian potential. These simulations are run in the context of 3-d ideal isothermal magnetohydrodynamics; the equations are integrated using the cylindrical coordinate extension (Skinner & Ostriker, 2010) to the Athena code package (Stone et al., 2008). The equations are solved in cylindrical coordinates, denoted by (R, ϕ, z), with $r = \sqrt{R^2 + z^2}$ referring to the standard spherical radius. The accretion disks we evolve are Newtonian and unstratified, as in Armitage (1998) and Hawley (2001), meaning that the gravitational potential is independent of the z coordinate and given by $\Phi = -1/R$. Neglecting the vertical dependence of the gravitational potential physically means that our simulations are meant to model the midplane of a realistic disk. This eliminates the effects of magnetic buoyancy and allows a more pure probe of MRI-driven phenomena. Numerically, it allows us to remove the vertical variation of zones per scale-height without resorting to the use of a more complex graded mesh with poorly understood grid-scale dissipation.

As one of the goals of this study is to explore issues of convergence, a word about resolution is in order. We classify simulations based on the number of zones per vertical scale height, specifically $H_0/\Delta z$. Given a choice of zones per scale height, the resolution is constrained by the condition that the aspect ratio is $1:1:1$ at the fiducial radius, $R_0 = 2$. This results in the condition $\Delta z = \Delta R \approx R_0 \Delta\phi$. The ability to explore simulations with an aspect ratio of unity is a particular benefit of orbital advection, detailed in the preceding chapter, and stands in contrast to the vast majority of the current literature on simulations of global accretion disks. The expense of running these simulations using the standard cylindrical integrator would be extreme, for instance the highest resolution simulation considered would take well over 4M CPU-hours. The use of orbital advection allows the exploration of resolutions whose expense would otherwise be computationally prohibitive.

To give a sense of the simulations and their resolution, Figure 5.1 shows two still images of the MRI-driven turbulence in the saturated state. The first, Figure 5.1, is a volume rendering of the logarithmic magnetic pressure in a simulation with a vertical resolution of 64 cells per vertical scale height, denoted simulation `BzZ64W` in the terminology introduced below. The second, Figure 5.1, shows the same variable in the $z = 0$ plane at the same timestep from a simulation with a vertical resolution of 32 cells per vertical scale height, denoted simulation `BzZ32`.

In cylindrical coordinates the physical domain spans $(R, \phi, z) \in [1, 4] \times [0, 2\pi] \times [-2H_0, 2H_0]$, where $H_0 = 0.2$, with the exception of two reduced runs that use a smaller azimuthal domain. The sound speed is set so that, $H(R_0) = H_0$, where

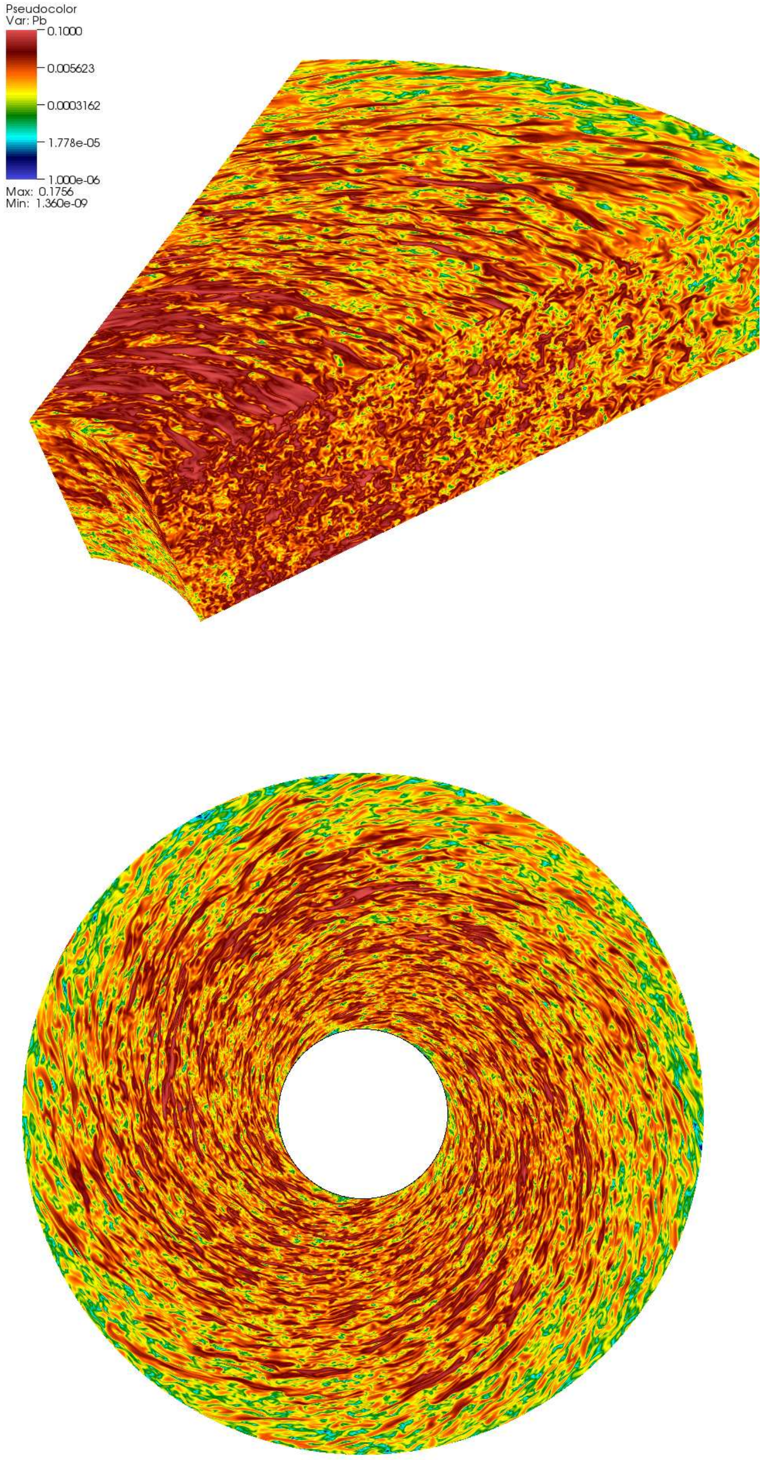

Figure 5.1: Still images of logarithmic magnetic pressure at orbit 75. Both images use the same color table. (**Top**) Volume rendering from simulation `BzZ64W`. (**Bottom**) Plane $z = 0$, from simulation `BzZ32`.

$R_0 = 2$ and $H(R)$ is given by equation 1.10 and the Keplerian angular velocity takes its Newtonian value $\Omega_K = R^{-1.5}$.

The hydrodynamic variables are initialized with a constant density, $\rho = 100$, and with velocities initialized as

$$v_R(R, \phi, z) = 0, \tag{5.1}$$

$$v_\phi(R, \phi, z) = R\Omega_K(1 + \delta), \tag{5.2}$$

$$v_z(R, \phi, z) = \delta, \tag{5.3}$$

where δ is a uniformly-distributed random perturbation such that $\delta \in [-10^{-2}, 10^{-2}]$.

To explore the effects of varying initial seed topologies, we consider here three distinct initial field configurations; a zero net flux vertical field (BzZ), a net-flux vertical field (BzN), and a net-flux azimuthal field (BpN). For simplicity we will often distinguish between the zero net flux (ZF) simulations and the net-flux simulations (NF). For the net-flux vertical field runs (BzN), the initial magnetic field is set such that the fastest-growing unstable mode of the MRI is equal to the fiducial scale height, *i.e.* $\lambda_{MRI} \equiv H(R_0)$, with λ_{MRI} given by equation 1.40. In our lowest-resolution runs, this gives a fastest growing mode that is marginally resolved, $\lambda_{MRI}/\Delta z = 8$. For the zero-net flux runs (BzZ), this vertical field configuration is modulated by a sinusoidal function with a wavelength equal to the effective scale height to ensure that net magnetic flux threading the simulation domain is zero. We do not explore the use of alternative ZF field topologies, but note that as the wavelength is increased we expect the behavior of the simulations to approach that of the BzN simulations. The toroidal field runs (BpN) are constructed to ensure that the critical azimuthal wavenumber of the toroidal MRI, $m_c = R\Omega/v_A$ (where $v_A = |\mathbf{B}|/\sqrt{\rho}$), is constant in radius with a value of 20. Finally, the field strengths are tapered to zero close to the radial boundaries of the calculation. These choices of seed fields are not meant to be representative of what may be present in an astrophysical accretion disk, but are chosen to allow controlled experiments to study MRI-driven disk turbulence.

Formally, the fields for the three cases are given by

$$B_z(R) = A_0 \frac{I(R) R_* \Omega_K(R_*)}{R} \sin\left[\frac{2\pi(R - R_0)}{H_0}\right] \quad \text{(BzZ)}, \tag{5.4}$$

$$B_z(R) = A_N S(R) \Omega_K(R) \quad \text{(BzN)}, \tag{5.5}$$

$$B_\phi(R) = \frac{S(R)\sqrt{\rho/R}}{M_c} \quad \text{(BpN)}. \tag{5.6}$$

Each initial field topology is defined by a single magnetic field component with the remaining components identically zero throughout the domain. In the above definitions, A_0 and A_N represent scaling terms, R_* is a piecewise constant function of radius that assumes the central radial value of each sinusoid, and $M_c = 20$. The function $I(R)$ is an indicator function on the domain $R \in [1.2, 3.8]$. The function $S(R)$ is a mollifier function, defined as in equation 4.31, with value unity on the domain $R \in [1.5, 3.5]$, zero on the complement of the domain $R \in [1.5 - H_0, 3.5 + H_0]$, and smoothly interpolates between the two domains. Figure 5.2 shows the resulting radial profiles of the initial magnetic field strengths.

As the simulations described here are unstratified, periodic boundary conditions are appropriate in the azimuthal and vertical direction. Radial boundary conditions are, as always, more difficult to implement appropriately. The boundary conditions used in the simulations presented here are particularly simple: the values of physical quantities in the ghost zones are simply copied over from the last physical zone. There are two exceptions to this: the radial velocity and the azimuthal velocity. The radial velocity uses what we refer to as an enforced diode condition to ensure that material does not enter the simulation. Formally, this means that if the radial velocity in the last physical zone is directed outwards then it is copied into the ghost values, otherwise the radial velocities in the ghost zone are set to zero. For the azimuthal velocity a Dirichlet boundary condition is used in which the ghost zone values are set to the appropriate Keplerian value. While some of the simulations described here are initialized to have a zero-net magnetic field, the boundary conditions are not constructed to enforce this flux constraint during the natural evolution of the disk. To ensure that these choices of radial boundary condition do not affect the physical behavior of the simulations several potential boundary conditions were considered. These include two choices for hydrodynamic variables and two choices for the magnetic field variables. In addition to the simple copy with diode condition for the hydrodynamic variables used, a variant in which the perturbed azimuthal velocity (i.e., the azimuthal velocity in the local Keplerian frame) was held constant in the ghost zones with an analogous diode condition. As an alternative the simple copy of the magnetic field variables, an alternative in which the magnetic field in the ghost zone was constrained to only have a divergence-free radial component was tested. These choices result in four possible radial boundary conditions, and all were tested to ensure that none resulted in an appreciable change in the evolution and behavior of the disk. It is worth noting that replicating a magnetic field configuration from one cell to another in a curvilinear coordinate system will not, in general, guarantee that it is divergence free. Indeed, the magnetic field in the ghost zones will not necessarily solve the solenoidal constraint. However, the ghost zones are necessary merely for the reconstruction of the magnetic field inside the physical domain and the use of constrained transport will guarantee that the magnetic field in the simulation domain will satisfy the solenoidal constraint.

The details of the suite of simulations are given below in Table 5.1. The simulations are classified according to their initial field topology and resolution (defined as $H_0/\Delta z$). These simulations are run for a period of time measured by the orbital period at the inner radial edge of the simulation. While the majority of the simulations presented use a full azimuthal domain and an aspect ratio of unity, extra control runs were performed to study the importance of these choices. Wedge runs, utilizing a reduced azimuthal domain of $\pi/4$, were run to assess the importance of low-m modes and to access resolutions higher than were possible with the restrictive constraints of the other simulations. These simulations are denoted `BzZ32W` and `BzZ64W`. In additional to using a truncated azimuthal domain, it is common in the literature to use a reduced azimuthal resolution. To assess the importance of azimuthal resolution, two runs (`BzZ32R` and `BzZ32RR`) were performed which halved and quartered, respectively, the number of azimuthal grid cells used in run `BzZ32` while maintaining the vertical and radial resolution. Also included in Table 5.1 are a series of scalar metrics whose definitions are given in §5.2.2.

Simulation	Azimuthal Range	Resolution (N_R, N_ϕ, N_z)	Orbits	$< \alpha >_{QSS}$	$< \beta >_{QSS}$	$T_{1/2}$
BzZ8	2π	(120,480,32)	200	0.013	26.55	184.4
BzZ16	2π	(240,960,64)	200	0.02	23.01	127.2
BzZ32	2π	(480,1920,128)	200	0.018	27.67	116.1
BzZ32W	$\pi/4$	(480,240,128)	100	0.023	23.35	N/A
BzZ64W	$\pi/4$	(960,480,256)	100	0.024	20.40	90.0
BzZ32R	2π	(480,960,128)	200	0.021	24.03	109.6
BzZ32RR	2π	(480,480,128)	200	0.015	30.01	127.5
BzN8	2π	(120,480,32)	200	0.058	8.18	20.0
BzN16	2π	(240,960,64)	200	0.076	8.26	19.4
BpN8	2π	(120,480,32)	200	0.055	6.37	56.2
BpN16	2π	(240,960,64)	200	0.064	7.37	40.2
BpN32	2π	(480,1920,128)	200	0.067	7.42	35.5

Table 5.1: Simulation Parameters

5.2.2 Diagnostics

As turbulence involves fluctuating quantities in both space and time, diagnostics will invariably involve some combination of spatial and temporal averaging. For simplicity, we define here the quantities we will use below. The most common quantity is simple volume-averaging and for a variable $X(R, \phi, z, t)$ this is defined as $\bar{X}(t) =< X > /|V|$, where V represents the full simulation domain.

The fiducial timescale of these simulations is the orbital period at the inner edge of the disk, $P_o = 2\pi$, and for brevity simply designated as an orbit without further qualification. Often we would like to compute scalar values of quantities that are representative of a simulation as a whole. The initial growth of certain physical quantities, particularly the stress and magnetic energy, is exponential and driven by the linear phase of the MRI. These quantities reach a peak value and quickly decrease as the linear phase of the MRI transitions into a fully turbulent state in which there is significantly reduced secular variation. We will refer to this stage as the quasi-steady state (QSS), which we define for simulations initialized with a vertical field as between 50 orbits and the end of the simulation. Due to the slower growth rate of the toroidal MRI we quantify the QSS as between 100 orbits and the end of the simulation. Scalar representative values, like those in Table 5.1, are denoted $\bar{X}_{QSS}$ and are defined as the temporal average over the timeframe defined as the QSS of the volume-average of the quantity.

The dynamics of accretion disk turbulence are of particular astrophysical interest, due to the anisotropic structure of the resultant turbulent magnetic field and its ability to drive angular momentum radially outward. To study the efficacy of angular momentum transport, many of our diagnostics will focus on the stress that allows this transport to happen. This stress is of two types: the Reynolds stress, $T_{R\phi} = \rho v_R \delta v_\phi$, and the Maxwell stress, $M_{R\phi} = -B_R B_\phi$, with the Maxwell stress generally dominating the Reynolds component. While these two quantities are of the most direct physical relevance, it is common to scale the stress by the gas pressure resulting in the diagnostics given by equations 1.31 and 1.32. Also of interest is the strength of the magnetic field in relation to the thermal energy of the gas, given by $\beta = P/P_b$, the ratio of gas and magnetic pressure. We define $< \alpha >$ and $< \beta >$ to refer to the volume-average of the numerator divided by the volume-average of the denominator of these quantities. The quantity $< X >_{QSS}$ is defined in the same manner as $\bar{X}_{QSS}$.

In an effort to study how well resolved the MRI is, we consider the following diagnostics

$$F_z = |V|^{-1} \int_V (\lambda_{MRI} \geq 8\Delta z) dV, \tag{5.7}$$

$$F_\phi = |V|^{-1} \int_V (\lambda_C \geq 8R\Delta\phi) dV, \tag{5.8}$$

where in the above equations logical statements refer to indicator functions that assume the value of unity or zero based on the truth or falsity of the statement. The characteristic wavelength of the toroidal field, λ_C, is defined by equation 1.46 and is analogous to λ_{MRI} save for the use of the toroidal Alfven speed in place of the vertical. These scaled integrals represent the fraction of the disk where the fastest-growing modes of the vertical and critical toroidal modes of the MRI are resolvable, using the often employed 8-zone criterion. These are related to the quality factors, first used by Noble et al. (2010), given as

$$Q_z = \lambda_{MRI}/\Delta z, \tag{5.9}$$

$$Q_\phi = \lambda_c/(R\Delta\phi), \tag{5.10}$$

where these values are in general functions of space and time. Use of the quality factor as a diagnostic represents an important step in appreciating the numerical resolvability of the MRI, but we feel that the resolvability fraction diagnostic we employ is an even more stringent resolvability requirement.

It is well-known that the anisotropy of the magnetic field is key to angular momentum transport, and indeed the correlation between the radial and azimuthal component of the magnetic field is an important diagnostic of angular momentum transport in disks. An alternative measure of the anisotropy, called the magnetic tilt angle, was first discussed by Guan et al. (2009). This measure is defined as,

$$\theta_B = \arcsin(\alpha_M \beta)/2. \tag{5.11}$$

Physically, this can be thought of as an approximation to the angle between the planar magnetic field and the azimuthal axis assuming a weak vertical magnetic field. An estimate of the tilt-angle is derived as $\theta_B \approx 15^\circ$, in Guan et al. (2009), by noting that all the local models considered satisfy the relationship $\alpha\beta \approx 1/2$. Strictly speaking, $\alpha_M < \alpha$, and thus we expect this estimate to be an upper bound.

Studying the spectral structure of turbulent flows is often more natural than alternative diagnostics in physical space. Our primary interest here will be studying the azimuthal structure of power spectra of physical quantities, in particular the density and magnetic pressure. To this end, we begin by defining the subdomain

$$\mathcal{S} = [R_0 - 2H_0, R_0 + 2H_0] \times [0, 2\pi] \times [-2H_0, 2H_0]. \tag{5.12}$$

For simulations that do not model the full 2π radian azimuthal domain the subdomain definition is modified to include the full azimuthal range modeled. We then consider, for a physical quantity $X(R, \phi, z)$, the azimuthal Fourier decomposition $X(R, m, z)$. Further, we define $X(m)$ as the volume-weighted radial and vertical mean over the subdomain $\mathcal{S}$. This removes spatial transients, but to remove temporal transients we also average between orbits 50 and 100, for simulations seeded with a vertical field, and between orbits 100 and 150 for simulations seeded with a toroidal field. We refer to this reduced temporal domain as RQSS. To ensure that these results are not adversely affected by secular trends within the subdomain related to the evolution of the disk we scale, at each timestep, by $< X >^{\mathcal{S}}$, the volume integral of the quantity over the subdomain. Formally,

$$\hat{X}(m) = \left\langle \frac{X(m)}{< X >^{\mathcal{S}}} \right\rangle_{RQSS}. \tag{5.13}$$

To more directly study important azimuthal scales we will employ the transformation of the wavenumber, m, to an effective azimuthal wavevector given by

$$k_\phi = \frac{m}{2\pi R_0}. \tag{5.14}$$

One final tool we employ when studying the spectral structure of disk turbulence is the use of a fiducial power spectra. To study the relative importance of large-scale versus small-scale features, we consider a fiducial power spectrum given by, $P_m = m^{-1}$. This fiducial spectra has the property that the inner and outer scales are equally important to the total integrated power over the accessible wavenumbers. We diagnose the dominant azimuthal scale by considering the location of the peak value of the spectra defined in equation 5.13 scaled by this fiducial power spectrum, namely $m\hat{X}(m)$.

The above formalism will be employed to study the structure of density and magnetic energy, and while it is common to study the structure of stress in a similar manner we choose a slight variant. Astrophysical interest in stress as a diagnostic is based on its ability to drive radial angular momentum transport. We note, however, that when projecting this quantity onto a basis of periodic functions all but the $m = 0$ mode will correspond to an azimuthal average of zero. Instead, we choose to analyze the structure of contributions to stress that result in net radial transport. Formally, we define

$$\tilde{\alpha_M}(m) = \left\langle \frac{-B_R(m)B_\phi(m)}{< P >^{\mathcal{S}}} \right\rangle_{RQSS}. \tag{5.15}$$

One of the goals of this work will be to explicitly compare local and global models, towards this end we proceed in a similar fashion to Sorathia et al. (2010) (Chapter 3) and decompose our global simulation into a set of subvolumes through which we can calculate "local" statistics. This is accomplished by considering subvolume of the physical domain $(R, \phi, z) \in \mathcal{S}$ and decomposing it into wedges of size $[H_0, 2\pi H_0, H_0]$. This yields, at each timestep, a set of 160 subvolumes in which relevant physical quantities can be volume-averaged. To calculate statistics, we simply take the mean, denoted $[X](t)$, and when relevant, the standard deviation of these quantities at each timestep.

5.2.3 Evolution of Global Disks

Prior to a full discussion of the resolution dependance of the simulations that will be presented in Section §5.3, we focus on the field topology dependance of several fiducial runs. For this we choose the highest resolution simulation of each field topology that was run for the full 200 orbits, these are: `BzZ32`, `BzN16`, and `BpN16`. The most significant difference between global and local simulations is the secular evolution of global simulations. Open boundary conditions allows mass to be accreted off the grid and the total magnetic flux to evolve in a dynamical manner. The development of radial structure of the mass profile adds radial pressure gradients to the dynamics of the turbulence.

To provide a sense of the evolution, Figure 5.3, shows the temporal evolution of several global quantities: the mass fraction, the volume integral of the mass in the simulation scaled by the initial total amount of mass; the efficiency of angular momentum transport, $< \alpha_M >$; and the dimensionless magnetic energy, $< \beta^{-1} >$. Most striking in all of these figures is the significant accretion and field amplification caused by the presence of a net magnetic field in the initial growth phase. Contrasting this, is the comparatively similar behavior of these quantities in the QSS. In this initial phase the net field runs exhibit magnetic fields with energies comparable to and even in excess of the thermal energy of the disk and accretion efficiencies an order of magnitude above those normally associated with zero-net field disks.

Indeed, a majority of the mass is accreted during the initial growth phase.

The significant accretion of mass is in contrast to what one would expect from a simple estimate of the viscous timescale, $T_\nu = R^2/\nu$, where the viscosity takes the common form $\nu = \alpha_{SS} c_s H$. Calculation of the viscous timescale at the fiducial radius, $R_0 = 2$, and converting to orbits yields $T_\nu \approx 63.6/\alpha_{SS}$. Comparison with the mass evolution (Figure 5.3, Top) and $T_{1/2}$ (Table 5.1) suggests a value of α_{SS} considerably in excess of the value of α measured in the QSS (Table 5.1) or even during the initial evolution (Figure 5.3, Middle). For instance, equating T_ν and $T_{1/2}$ suggests a value of $\alpha_{SS} \approx 3$ for run `BzN16` ($< \alpha >_{QSS} = 0.076$) and $\alpha_{SS} \approx 0.5$ for `BzZ32` ($< \alpha >_{QSS} = 0.018$).

The viscous timescale is, of course, a crude estimate. Formally, it should be valid only in the case of a radially localized mass distribution with a viscosity exhibiting minimal radial dependance. Attempting to apply this estimate to global disks of the type described here stretches the approximation far beyond its area of applicability. Understanding the evolution of global disks in the context of an anomalous viscosity model requires a more sophisticated treatment including the radial structure. Towards this end, we compare our simulated disks with a 1-dimensional reduced model for the time-evolution of the surface density $\Sigma = \int \rho dz$ based upon the anomalous viscosity model (e.g., see Pringle 1981) and given by equation 1.15.

We solve Eqn 1.15, with $\nu = \alpha_{SS} c_s H$, using the values of α calculated from the full simulation. We use the vertically and azimuthally averaged values of α, temporally spaced every tenth of an orbit, in place of α_{SS} in the formulation of the turbulent viscosity. We note that α_{SS} and α are related by a factor of $3/\sqrt{2}$. The radial dependance of the viscosity is important, as the evolution of the surface density depends on its spatial derivatives. The early evolution of the disk is characterized by small-scale structure in the stress that is likely to form an important contribution to the evolution of the surface density, which itself exhibits weak spatial variability during the initial phase of the evolution. The evolution equation is solved utilizing a simple implicit, finite-differencing scheme. For simplicity, the boundary conditions are set by the constraint that the mass profile agrees with the vertically and azimuthally averaged mass in the full simulation at the inner and outer boundary.

Figure 5.4 compares the evolution predicted by the reduced model to that obtained in the full MHD simulation. A comparison of the broad features of the evolution, specifically the mass fraction, is given in Figure 5.4 (Top). In all cases the reduced model accurately predicts the mass evolution of the disk. More detail is given in Figure 5.4 (Bottom), in which the radial mass profile computed by the reduced model is compared to that from the full simulation at orbit 75. The overall features of the profile are captured, although we notice a larger discrepancy than observed in the mass fraction. The observed discrepancy appears to be largely caused by the limitations of the temporal discretization used in which values of α from the simulation are calculated ten times every orbit. The initial growth of the MRI induces fluctuations with significant variability in both space and time that are not adequately captured by the temporal discretization. The model showed in Figure 5.4 must agree with the full simulation if the latter conserves mass and angular momentum in the limit as the temporal spacing of α from the simulation vanishes. Indeed, taking as the initial condition the density profile after the saturation of the turbulence and evolving this profile using the one-dimensional model and data from the simulation removes much of the discrepancy.

While this section has focused on a formal comparison of our MHD simulations with reduced models, these results may have more direct astrophysical implications. The most direct way of determining α in real accretion disks is the analysis of dwarf nova outbursts and X-ray transient outbursts. As pointed out by King et al. (2007), these estimates suggest $\alpha \sim 0.1 - 0.4$ whereas numerical simulations (including those presented here) typically obtain steady-state values of α that are an order of magnitude smaller. In the light of our results, we note that an accretion disk which has just entered an outburst state may well go through a period of field-growth that resembles the early transients seen in our simulations. During these transients, the effective value of α is substantially enhanced, and spatio-temporal gradients in α further enhance the angular momentum transport. Thus, it is interesting to conjecture that the large values of α inferred from outburst systems correspond to these transient phenomena. These issues will be explored in a future publication.

5.3 Convergence of Global Disk Simulations

Standard tests of convergence rely on running simulations with increasing resolution while leaving the underlying physical problem unchanged. Convergence in the case of turbulent non-explicitly-dissipative systems is inherently ill-defined for two reasons. The first is that when increasing resolution there will invariably be minor differences in the seed perturbations that feed the instability[1], and as a result we cannot expect precise agreement between simulations. The second and more fundamental reason is that in ideal MHD the dissipation scale is set by the grid scale, and thus when increasing resolution we are not leaving the underlying physical problem unchanged. Further, changes in resolution alters the evolution of the disk; changes in mass accretion and mass distribution results in a fundamentally different disk. In light of these complexities, we will hereafter take the notion of convergence to mean that an increase in resolution will leave relatively unchanged spatially- and temporally-averaged measurements.

We consider three broad categories of convergence metrics: physical, numerical, and spectral. The physical metrics, α_M and β, are perhaps the most natural and directly physically relevant and therefore have the longest history of use as a diagnostic of accretion disk turbulence. However, as we will demonstrate these metrics are often ambiguous and display non-monotonic resolution dependance. Numerical metrics, the resolvability fractions and quality factors, directly measure how well the linear MRI is resolved and were a focus of the convergence study described by Hawley et al. (2011). The spectral metrics we employ study the azimuthal structure of the turbulent flow, and specifically seek to identify the dominant azimuthal scale and its dependance on resolution. We will demonstrate that all of these metrics are useful diagnostics towards studying the nature of the turbulent flow in accretion

---

[1] This effect can be removed by initializing a simulation with perturbations defined by an explicit power spectrum, however in the QSS power will be spread amongst all available modes regardless.

disks, however as a convergence criterion the magnetic tilt angle appears unambiguous and robust.

5.3.1 Physical Metrics

The simplest and most astrophysically relevant convergence criterion is accretion efficacy as measured by the dimensionless stress α_M. Figure 5.5 (Top) shows the evolution of the Maxwell stress over the course of the simulation for all three standard ZF models. As is expected, initial peaks in the stress associated with the linear growth of the vertical MRI are resolution dependent. While the initial magnetic field is constructed so that the most unstable mode, λ_{MRI}, is resolvable in each simulation there are other slower-growing unstable modes whose resolvability will vary depending on the resolution of the simulation. The true test of convergence is the behavior of the stress in the saturated quasi-steady state. The lowest resolution simulation, `BzZ8`, exhibits $< \alpha_M >_{QSS} \approx 0.01$, while both higher resolution simulations exhibit a comparable stress to each other that is roughly 50% greater than `BzZ8`.

While this fundamental criterion of convergence in stress is satisfied, other diagnostics paint a more subtle picture. Figure 5.5 (Bottom) illustrates the evolution of the scaled magnetic energy, β^{-1}. Again, as expected, we see that the initial field amplification is monotonic with resolution and dominated by the growth of the toroidal magnetic field, however the saturation and transition into the fully nonlinear state is more complex. While `BzZ32` peaks at a higher value due to the larger number of resolvable, unstable MRI modes the lower resolution simulations, `BzZ8` and `BzZ16`, maintain a stronger magnetic field proportionally. These lower resolution simulations also lack the steep drop-off in magnetic field energy often associated with the saturation of the MRI. This may be a consequence of the inability to resolve the parasitic instabilities associated with saturation of the linear MRI. Also of note, is the late-time field growth associated with `BzZ8`. The nature of this growth is unclear, but may be suggestive of a very low-frequency temporal behavior.

Our analysis of the physical metrics of models initialized with a seed field possessing net flux proceeds in much the same way as our analysis of the ZF models. As above, we begin by considering the time evolution of the global quantities, α_M and β^{-1}. The results are shown in Figure 5.6 in the Top and Bottom panel respectively. The evolution of the linear MRI displays a significant resolution dependence, as would be expected, but in all cases the values of α_M in the saturated state are comparable for the same initial field topologies. Due to the significantly higher values of α_M in the net field simulations we see a corresponding significant increase in mass loss of the disk (Table 5.1) compared to the ZF runs. Field amplification for the runs initialized with a vertical field is monotonic with resolution. However, regarding field amplification the net toroidal runs behave more analogously to the net-zero field runs discussed above, in which lower-resolution simulations seem to exhibit greater field amplification in the QSS. The presence of a net field, regardless of topology, results in order of magnitude increases in the peak values of α_M and β^{-1} compared to the ZF simulations considered above. Insofar as angular momentum transport is concerned, there is only a weak dependance on resolution in the saturated state even for the most poorly resolved simulations. The evidence from these global metrics suggests that net-flux simulations converge more quickly, however we'll see through the consideration of the other metrics that the picture is more subtle.

Overall, we see that while the behavior of α_M paints a simple picture and is suggestive of convergence, the magnetic energy is far more volatile and does not follow a clear pattern in its dependance on resolution. Further, as a convergence criterion α_M suffers from its dependance on initial field topology.

The use of orbital advection allows us a unique opportunity of exploring isotropic resolutions in a cost-effective manner. Without orbital advection, the timestep constraint is set by v_K/Δ_ϕ and thus doubling azimuthal resolution results in a steep price, specifically a quadrupling of the necessary operations from the doubled number of cells to time-advance and the half-timestep being used. The result is that in the majority of the literature, azimuthal resolution is compromised for computational expediency. The goal of runs `BzZ32R` and `BzZ32RR` are to study the importance of azimuthal resolution in a controlled way in comparison to run `BzZ32` and deduce a maximum aspect ratio from which converged turbulence is guaranteed. Figure 5.7 shows the evolution of α_M (Top) and the scaled magnetic energy (Bottom) for these runs. We note that in both cases we see non-monotonic behavior, specifically reducing the resolution by a factor of two results in increased accretion efficacy and toroidal field amplification. Further reduction results in a significant drop-off in both of these quantities. While increased field amplification with reduced resolution is expected from Figure 5.5 (Bottom), that a reduction in resolution could increase angular momentum transport is unexpected. Assessing convergence from these global quantities is murky, at best. For now, we merely note the importance of azimuthal resolution by pointing out that reducing azimuthal resolution can severely impact accretion efficiency. The convergence of these reduced azimuthal runs is returned to in §5.3.4 where it is demonstrated that run `BzZ32R` is converged, whereas `BzZ32RR` is not.

5.3.2 Numerical Metrics

Our consideration of numerical metrics begins with the resolvability fractions defined in equations 5.7 and 5.8 and shown, for the ZF runs, in Figure 5.8 (Top panels). The initial radial profile of the vertical seed field is constructed so that $\lambda_{MRI}/\Delta z \geq 8$ for simulation `BzZ8` in a small neighborhood around each maximal value of the sinusoid. The variations in the initial values of F_z are due to the size of the neighborhood about each maximal value of the initial sinusoid in which the resolvability criterion is satisfied. Of all the diagnostics considered, the most startling behavior is seen when considering F_z. The two lower resolution simulations, `BzZ8` and `BzZ16`, both show a significant initial drop in the resolvability of the vertical MRI correlated with the linear growth and saturation phase of the evolution. Following the transition, both of these simulations show a slight increase in the resolvability fraction but clearly in the most of the disk the vertical MRI is not adequately resolved. In contrast to this, `BzZ32` seems to exhibit substantially differing behavior and shows a value of F_z that is roughly constant during the full evolution of the simulation. The evolution of F_ϕ, shown in Figure 5.8 (Top-Right) is by comparison much simpler and monotonic in nature. As the toroidal field grows due to shear amplification driven by the vertical MRI, the toroidal MRI becomes resolvable throughout a significant portion of the disk.

The behavior of the resolvability fractions paints an interesting picture. Runs `BzZ16` and `BzZ32` exhibit similar saturated values of stress, but their ability to resolve the vertical MRI seem to be significantly different. This suggests the intriguing possibility that run `BzZ16`, while initially seeded with a vertical field is actually reliant on the resolvability of the toroidal field to reach a comparable stress to run `BzZ32`. However, if these two simulations are actually taking differing routes to turbulence there then must be some additional mechanism that accounts for the similar values of stress achieved in the quasi-steady state. The resolvability fractions suggest that run `BzZ32` has reached a categorically different state in its resolvability of the vertical MRI, and the toroidal MRI is resolvable throughout almost all of the disk.

Next we consider the resolvability fractions of the NF runs, given in Figure 5.8 (Bottom panels). The resolvability of the vertical MRI is, at first glance, generally somewhat poor. Of all the simulations considered, only `BpN32` resolves the vertical MRI in a majority of the disk over the full evolution. As was the case with the ZF runs considered, all of the simulations resolve the toroidal MRI in the majority of the disk (Figure 5.8, Bottom-Right). From these resolvability fractions, only `BpN32` is clearly well-resolved. Of interest is that the resolvability fraction in the saturated state appears to have stronger dependence on resolution than initial field topology.

The resolvability fractions employ the often used 8-zone resolvability criterion; however this is somewhat arbitrary. They allow us to get a sense of how much of the disk is "well"-resolved. Now we employ the quality factors to understand how "well"-resolved the disk is. Figure 5.9 (Top panels) illustrate the evolution of the vertical and toroidal quality factors for the ZF runs. The vertical quality factors roughly double with corresponding resolution doublings, starting at $Q_z \approx 2$ for the lowest resolution simulation. This results in a situation in which run `BzZ16` is crudely resolved with an average of four zones per λ_{MRI} and `BzZ32` is demonstrably resolved. Consideration of the toroidal quality factor results in a reiteration of the point made regarding the toroidal resolvability fraction, F_ϕ, specifically that the two highest resolution simulations clearly are resolving the toroidal MRI. In particular, run `BzZ32` has an average of 40 zones per critical wavelength.

Consideration of the quality factors for the NF runs (Figure 5.9, Bottom panels) reveal a similar resolution dependance as that seen in the resolvability fractions. With the exception of `BpN32`, all of the net-flux runs evolve to a point with an average of a vertical quality factor less than eight, the fiducial criterion. The toroidal quality factors, in contrast, are quite large and the simulations with all but the lowest resolution simulations exhibit toroidal quality factors above 20. Overall, consideration of the resolvability fractions and quality factors are taken as evidence that runs `BzZ32` and `BpN32` are resolved whereas `BzZ16` is barely resolved.

5.3.3 Spectral Metrics

While integrated global quantities are undoubtedly important when diagnosing turbulence, it is often the case that the spectral structure of turbulence is more amenable to study than is the physical structure. Convergence in this domain is again inherently ill-defined due to the non-dissipative nature of these simulations. Increasing the resolution of a simulation increases the number of modes in which power can reside, and it would be unphysical to expect that these newly opened modes would remain free of power. We adopt as a definition of convergence that the mode at which power peaks remains constant with increasing resolution. To probe the spectral structure at the smallest scales, we include in consideration the wedge runs, `BzZ32W` and `BzZ64W`, where the function of the former is primarily as a control to ensure that the small-scale behavior is not altered by the reduction of azimuthal domain.

Figure 5.10 shows the time-averaged azimuthal power spectra of the density, magnetic pressure, and stress scaled to remove the secular evolution and temporally averaged for 50 orbits beginning at orbit 50. It is clear from these figures that reducing the azimuthal domain of the simulation, as done in `BzZ32W`, does not significantly change the small-scale distribution of power. The importance of large-scale structure to the overall density profile of the disk is clear from Figure 5.10 (Top). Comparing the structure of the magnetic pressure and stress, Figure 5.10 Middle and Bottom panel, suggests that the magnetic pressure is clearly dominated by intermediate scales while the stress is more equally distributed at large and intermediate scales. From a visual inspection, it is clear that all of the power spectra are self-similar and, for resolutions above 32 zones per scale height, all peak at approximately the same scale. At small scales the power spectra look quite like those of the non-converged, zero net-flux models presented by Fromang & Papaloizou (2007). The large-scale behavior, however, is quite different as the power at large scales is at most weakly dependent on resolution in stark contrast to the non-converged models of Fromang & Papaloizou (2007).

Next, we consider the azimuthal spectral structure of the turbulence for the NF runs (Figure 5.11). Broadly, we note again that the structure is dependent largely on resolution and only minimally on field topology. Simulations of lower resolution are more dependent on large-scale features, and as resolution increases we see a more even distribution of power between the intermediate and small scales. While the density and pressure have a clear peak at an intermediate scale, the scale-distribution of stress is more evenly distributed at large to intermediate scales.

To quantify the dominant scale, we consider in Table 5.2 the wavenumber at which the wavenumber-scaled power peaks for each simulation and for each of the variables considered. For the eventual comparison of these results with local models, it is useful to consider this measurement both as wavenumber, appropriate for global simulations, and as an effective azimuthal wavevector, $k_\phi H_0$, which is more directly analogous to local models. The dominant density scale is much larger than the corresponding dominant magnetic scales. All of the simulations considered exhibit a dominant density scale of approximately $4H_0$, and in this regard resolution does not appear to play a significant role save for the absolute lowest resolution simulation. Reducing the azimuthal domain results in a increase of the effective scale associated with the magnetic quantities. While the magnetic quantities do not clearly converge with increasing resolution, the peak of the highest resolution simulation, `BzZ64W`, is almost precisely displaced in wavenumber by eight. The reduced azimuthal domain simulations use only an eighth of the full 2π domain, which means that the accessible wavenumbers are limited to modes that are integer multiples of eight. In this regard, the peak of runs `BzZ32W` and `BzZ64W` are within one accessible wavenumber of the peak of `BzZ32`. With this in mind, we again take this as evidence of the convergence of `BzZ32`.

The dominant density scale is roughly halved when going from run `BpN8` to `BpN16`, however the transition from `BpN16` to `BpN32` results in only a minor ($\sim 10\%$) decrease and suggests that this latter run is near convergence. This is contrasted with the dominant scales of the magnetic energy and stress which decrease much more significantly in the transition from `BpN16` to `BpN32`. While the dominant density scale of run `BzZ32` is twice that of the dominant stress scale, the toroidal field runs display a much more comparable value of these two scales. Overall, we take this evidence as suggesting that run `BpN32` is approaching convergence but not fully so.

The most striking feature of the net vertical field simulations is the significantly larger dominant stress scale ($\sim 10H_0$). Also larger than expected is the dominant magnetic energy scale, which is again in excess of the other topologies considered. These large scales may suggest the existence of a memory of the initial field configuration. In particular, this large structure may be a remnant of the linear shearing phase in the initial growth of the MRI. This linear shearing phase is particularly strong in the simulations initialized with a net vertical field, and exhibit magnetic pressures in excess of the gas pressure.

Simulation	Density Peak		P_b Peak		Stress Peak	
	m	$k_\phi H_0$	m	$k_\phi H_0$	m	$k_\phi H_0$
`BzZ8`	7	0.111	13	0.207	11	0.176
`BzZ16`	14	0.223	25	0.398	19	0.302
`BzZ32`	14	0.223	48	0.764	33	0.525
`BzZ32W`	16	0.255	40	0.637	24	0.382
`BzZ64W`	16	0.255	56	0.891	24	0.382
`BpN8`	8	0.127	12	0.191	8	0.127
`BpN16`	17	0.271	17	0.271	12	0.191
`BpN32`	19	0.302	32	0.509	16	0.255
`BzN8`	10	0.160	13	0.207	5	0.080
`BzN16`	15	0.234	15	0.234	6	0.095

Table 5.2: Comparison of dominant azimuthal mode in the power spectra of density, magnetic pressure, and stress.

5.3.4 Tilt Angle

The metrics discussed thus far are useful tools to measure convergence, but they suffer from the limitation that their behavior must be compared against other simulations using identical field topologies. Meaningful convergence studies can be computationally quite expensive and therefore a more robust indicator of convergence that is independent of field topology would be useful. The evidence we will present here suggest that the magnetic tilt-angle, defined in Eqn 5.11, may indeed be this diagnostic. As a precursor to this, we consider the evolution of the tilt angle in all of the ZF simulations (Figure 5.12, Top panel). The magnetic tilt angle reaches an approximately steady-state value that is almost identical for all of the higher resolution simulations (above 32 zones/H_0). This value, $\theta_B \approx 13°$, is remarkably close to the estimated value of 15° (Guan et al. 2009). Also of note is that while the initial peaks in both stress and magnetic energy are strongly resolution dependent, the tilt angle, a function of the ratio of the two, has an initial peak that is seemingly independent of resolution. This may suggest that the transition from the linear growth of the MRI into saturated turbulence depends on a precise relationship between the quantities α_M and β.

Next, we consider the behavior of the magnetic tilt angle in the NF runs, given in Figure 5.12 (Middle). As resolution increases, we see a corresponding increase in the tilt angle. Regarding the toroidal field simulations, we note the strong similarity between runs `BpN16` and `BpN32`. This suggests that a further doubling in resolution may indeed prove convergence. The vertical field simulations display a similar behavior and ambiguity. Unfortunately, with the simulations available we are not able to conclusively demonstrate convergence for these net field runs, however we do argue that `BpN32` is likely converged.

At this point, we recall the difficulty of defining convergence using the physical metrics in the context of the reduced azimuthal resolution runs, `BzZ32R` and `BzZ32RR`. The physical metrics are ambiguous but, when we consider the tilt angle (Figure 5.12, Bottom), we see a clear monotonicity with resolution. A reduction in the azimuthal resolution by a factor of two results in a very minor change in the tilt angle, however further reduction of the azimuthal resolution results in a much more significant alteration. We take this as evidence that run `BzZ32R` is converged, whereas `BzZ32RR` is not. This demonstrates the importance of azimuthal resolution, and indeed suggests that treating azimuthal resolution on nearly equal footing with vertical resolution is of vital importance towards ensuring simulations that are numerically converged.

The results we have presented indicate that the magnetic tilt angle is a powerful diagnostic tool towards demonstrating convergence of MRI-driven accretion disk simulations. The fact that it is monotonic with resolution, exhibits minimal variation in the saturated state, and appears to converge suggest that fiducial values of converged tilt angle can be computed and compared against simulations. Additionally, as demonstrated in Figure 5.13, the value in the saturated state appears to be almost independent of initial field topology and that there may exist a single fiducial value of tilt angle, although a much wider range of simulations including differing physical domains and physics will be necessary to verify this rigorously. We conjecture that higher resolution simulations utilizing the same model initialized with a net field will indeed definitively demonstrate that the tilt angle is a robust convergence metric. Assuming this can be done, this would prove a powerful tool towards verifying convergence without resorting to the significant expense of running higher resolution control simulations. Further, the existence of a single fiducial tilt angle would imply the ability to unify the often disparate phenomenology associated with varying initial field topologies.

The natural question raised by the existence of a single scalar that appears to characterize the saturated, fully non-linear state of MRI-driven MHD turbulence is whether this can be derived from theory. The relationship, $\alpha\beta \approx 1/2$, was first observed by Hawley et al. (1995) and discussed at length by Blackman et al. (2008). The latter work includes a heuristic argument that α and β should be related by a constant, however predicting that constant is beyond the means of the dimensional analysis used. Recent work by Pessah (2010) provides the potential for a more quantitative understanding of this relationship. Through an analysis of the growth and subsequent saturation of a single vertical MRI mode by parasitic instabilities, as initially studied by (Goodman & Xu, 1994), they find that when Kelvin-Helmholtz parasitic modes dominate that the saturation of the primary MRI mode occurs when $\alpha\beta \approx 0.4$ ($\theta_B \approx 12^\circ$). The regime in which Kelvin-Helmholtz parasites dominates is when the Elsasser number, $\Lambda_\eta = v_A^2/\eta\Omega_K$, is larger than unity and it is this regime that we expect our ideal simulations to correspond to. This work is encouraging as it represents an important point of connection between our numerical results and theory. The analysis of Pessah (2010) finds that the magnetic tilt angle will depend on dissipative coefficients and finding agreement between this dependence and future numerical simulations would bolster the importance of this quantity. Additionally, the fact that simulations initialized with a toroidal field are characterized by the same tilt angle as simulations initialized with a vertical field may suggest that the formation and subsequent parasitic destabilization of channel solutions may play an important role in the saturation of the non-axisymmetric MRI.

5.3.5 Discussion and Conclusions

Hawley et al. (2011) survey a series of local simulations to test various proposed convergence metrics, and then compare these metrics with a series of pseudo-Newtonian simulations of thick accretion disks. Our comparison to their work begins with their choice of metrics: $\alpha_{Mag} = M_{R\phi}/P_B$; the quality factors, Q_z and Q_ϕ (though the definitions given there differ from ours by a numerically negligible factor of $\sqrt{16/15}$); and the correlations, $< B_R^2/B_\phi^2 >$ and $< B_z^2/B_R^2 >$. We note that the measure α_{Mag} used in Hawley et al. (2011) is not the same as our metric α_M; while both involve a pressure scaling of the magnetic stress the former uses the magnetic pressure whereas the latter uses the gas pressure. We also note that two of the metrics proposed are in an information content sense equivalent to the magnetic tilt angle, θ_B. Due to the nature of the magnetic tilt angle being a measure of anisotropy in the planar field, we can rewrite $\alpha_{Mag} = \alpha_M \beta$ which implies $\sin(2\theta_B) = \alpha_{Mag}$. Similarly, the correlation $B_R^2/B_\phi^2 = \tan^2(\theta_B)$. Regarding the latter correlation, B_z^2/B_R^2, the authors note that this term does not appear to exhibit a strong trend with increasing resolution and as such we will not consider it further. To aid in our comparison we include a summary of convergence metrics in Table 5.3.

For the local models reviewed in Hawley et al. (2011), simulations of 64 zones per scale-height result in quality factors of $Q_z \approx 10$ and $Q_\phi \approx 40$. These are fairly comparable to each other, and run `BzZ32` (Table 5.3), despite the disparate initial field topologies. A zero-net vertical field is used in Davis et al. (2010) whereas a net toroidal field is used in Simon et al. (2011). Of note, is that the net toroidal field model considered here, `BpN32`, shows significantly higher quality factors at lower resolutions. While Davis et al. (2010) also consider a model utilizing 128 zones per scale-height which results in an increase in the quality factors by approximately 150%, this does not result in a change in $\alpha_{Mag} = 0.36$ from the 64 zone per scale-height run. This value of α_{Mag} corresponds to $\theta_B \approx 10.5^\circ$. The simulations discussed in Simon et al. (2011) result in a largest value of tilt angle of $\theta_B \approx 11.8^\circ$. These are both comparable to the values seen in our global runs with a slightly larger value of $\theta_B \approx 12^\circ - 13^\circ$.

In their discussion of stratified local models, Hawley et al. (2011) also note the importance of azimuthal resolution which is bolstered by the work presented here. In §5.3.4 we present evidence that when the azimuthal resolution is reduced by a factor of four from an aspect ratio of unity, the resultant simulations exhibit significantly lower accretion and magnetic tilt angles. Also discussed, is the ability for large toroidal quality factors to compensate for a poorly resolved vertical MRI. Indeed, this likely explains the results for simulation `BzZ16` in which a comparable value of α in the steady-state is achieved despite the significant discrepancy between the resolvability fractions of simulations `BzZ16` and `BzZ32`.

The global simulations presented by Hawley et al. (2011) are diagnosed to have lower quality factors and smaller tilt angles than the ones presented here and to the local models they consider. However, we observe that the stratified local models also seem to involve a smaller tilt angle than what may be expected from our unstratified global models. Beckwith et al. (2011) measure a tilt angle of approximately 9° in a stratified global simulation. This suggests that stratification itself may suppress tilt angle somewhat and a future resolution study of stratified global disks using orbital advection would be useful. Overall, the work suggests that convergence is attained when $Q_z \gtrsim 10 - 15$ for poorly resolved toroidal quality factors ($Q_\phi \approx 10$), and that larger values of toroidal quality factor ($Q_\phi \gtrsim 25$) can alleviate the constraint on the vertical quality factor. Based on these constraints, we find that all of the simulations utilizing a resolution exceeding or equal to 32 zones per scale-height are converged as well as the simulations seeded with a net field and using resolutions above 16 zones per scale-height. The criterion presented here of a converged tilt angle would not consider the runs `BzN16` and `BpN16` converged, but would consider run `BzZ32R` converged in contrast to the quality factor criterion.

Simulation	$< \theta_B >_{QSS}$	$\bar{Q}_{z,QSS}$	$\bar{Q}_{\phi,QSS}$
`BzZ8`	8.61°	1.68	8.38
`BzZ16`	11.44°	4.57	16.82
`BzZ32`	12.86°	9.91	41.22
`BzZ32W`	12.75°	9.98	33.95
`BzZ64W`	12.76°	24.51	73.77
`BzZ32RR`	11.12°	6.64	7.38
`BzZ32R`	12.59°	9.08	16.13
`BpN8`	8.08°	4.46	18.86
`BpN16`	10.98°	9.23	31.92
`BpN32`	12.09°	22.2	64.22
`BzN8`	9.37°	3.54	16.24
`BzN16`	11.83°	8.97	32.41

Table 5.3: Summary of Convergence Metrics

In this chapter, we have presented results from a series of simulations utilizing, for the first time, orbital advection in global cylindrical coordinates. The use of orbital advection, and the order of magnitude performance enhancement it provides, has allowed the exploration of global disks at resolutions not only comparable to local models, but in many cases exceeding the resolution of shearing box simulations in the literature. Removing the constraint of the Keplerian timestep has also allowed the exploration of simulations with isotropic resolution ($\Delta z = \Delta R \approx R\Delta\phi$) without the significant computational cost normally associated with azimuthal resolution.

The primary distinction between local and global simulations is the significant degree of temporal evolution of the latter. Simple estimates of accretion relying on the viscous timescale from anomalous viscosity disk theory vastly underestimate the degree of accretion observed during the initial transient of global simulations. However, a more sophisticated treatment utilizing a one-dimensional model (equation 1.15) based on the measured stress in the simulation accurately reproduce the temporal evolution and radial distribution of the mass. Estimates of α based on the viscous timescale inferred from the initial transient result in values well over an order of magnitude above the values of α that characterize the simulation in the QSS and even in excess with the measured values of α associated with the transient. This discrepancy may have potentially important astrophysical implications regarding observational estimates of α based on the viscous timescale.

Understanding and minimizing the grid-scale dependance exhibited by simulations of accretion disk turbulence is a vital precursor to constructing computational models that can be confidently compared against observational data. The ability to quantify complex, spatially and temporally varying turbulence into a simple scalar, the magnetic tilt angle, and use that to verify the presence of converged turbulence is an important step towards that goal. While several convergence metrics have been explored here focusing on either the physical, numerical, or spectral nature of the simulations, the magnetic tilt angle stands out as being the most important indicator of convergence. In the context of unstratified global simulations presented here, it appears that numerically-converged MRI-driven MHD turbulence is characterized by $\theta_B \approx 13°$. While all the metrics considered here are useful towards understanding the physics and numerics of MRI-driven turbulence, the magnetic tilt angle alone possesses all the qualities we desire in a convergence metric. The tilt angle is monotonic with resolution, and is independent of initial field topology, while exhibiting weak dependence on stratification and local versus global formalism. The data presented here suggests the following resolution requirements for convergence (assuming that $\lambda_{MRI} < H$): a vertical resolution of $H/\Delta z \geq 32$, $\Delta R = \Delta z$, and an azimuthal resolution that satisfies $R\Delta\phi \leq 2\Delta z$. This final requirement is particularly constraining given the computational difficulty associated with nearly isotropic resolution in simulations without orbital advection. The importance of orbital advection in future simulations is clear.

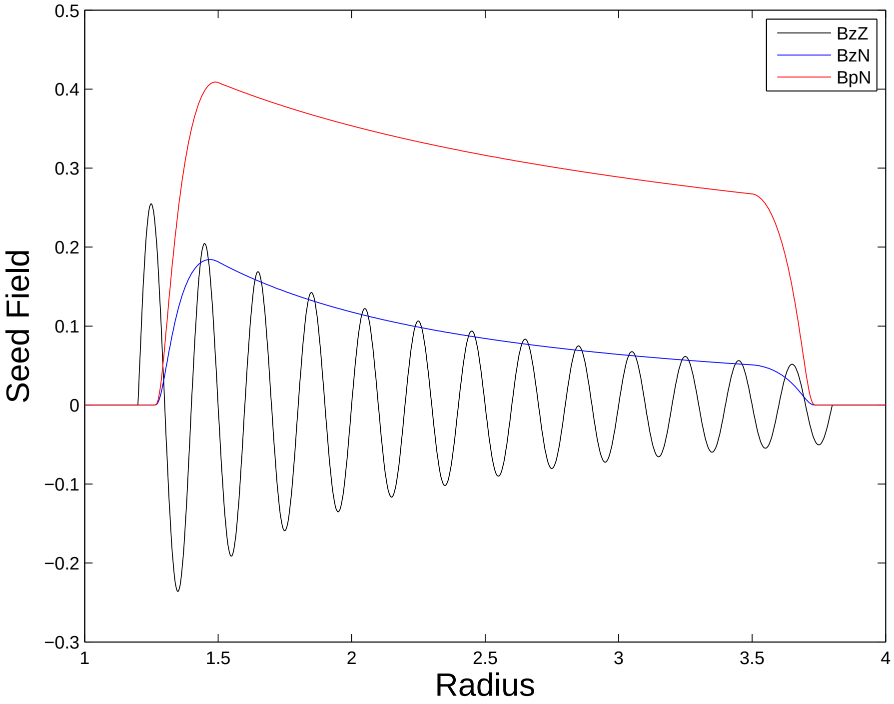

Figure 5.2: Radial profiles of initial magnetic field topologies.

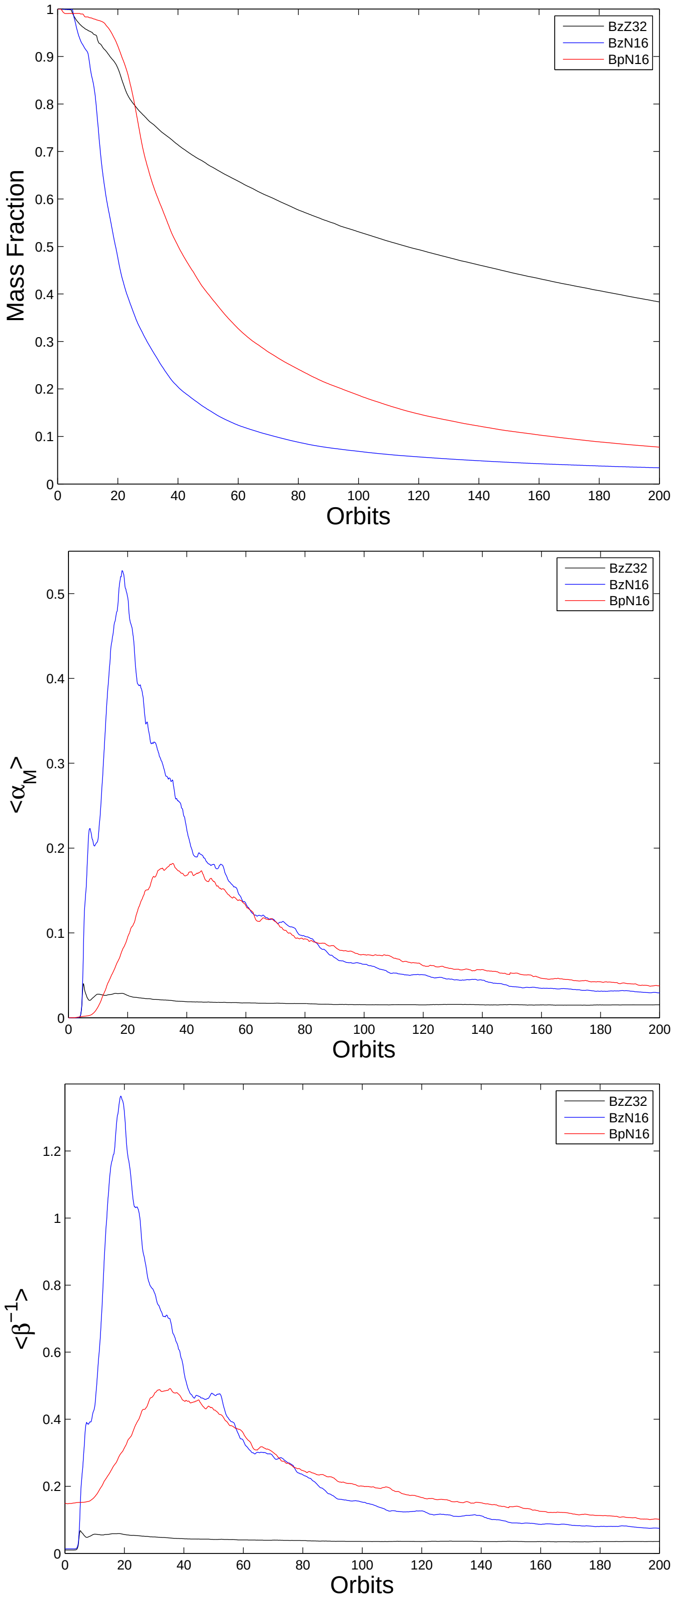

Figure 5.3: Comparison of the evolution of the fiducial simulations. (**Top**) Mass fraction (**Middle**) Accretion efficiency (**Bottom**) Dimensionless magnetic energy

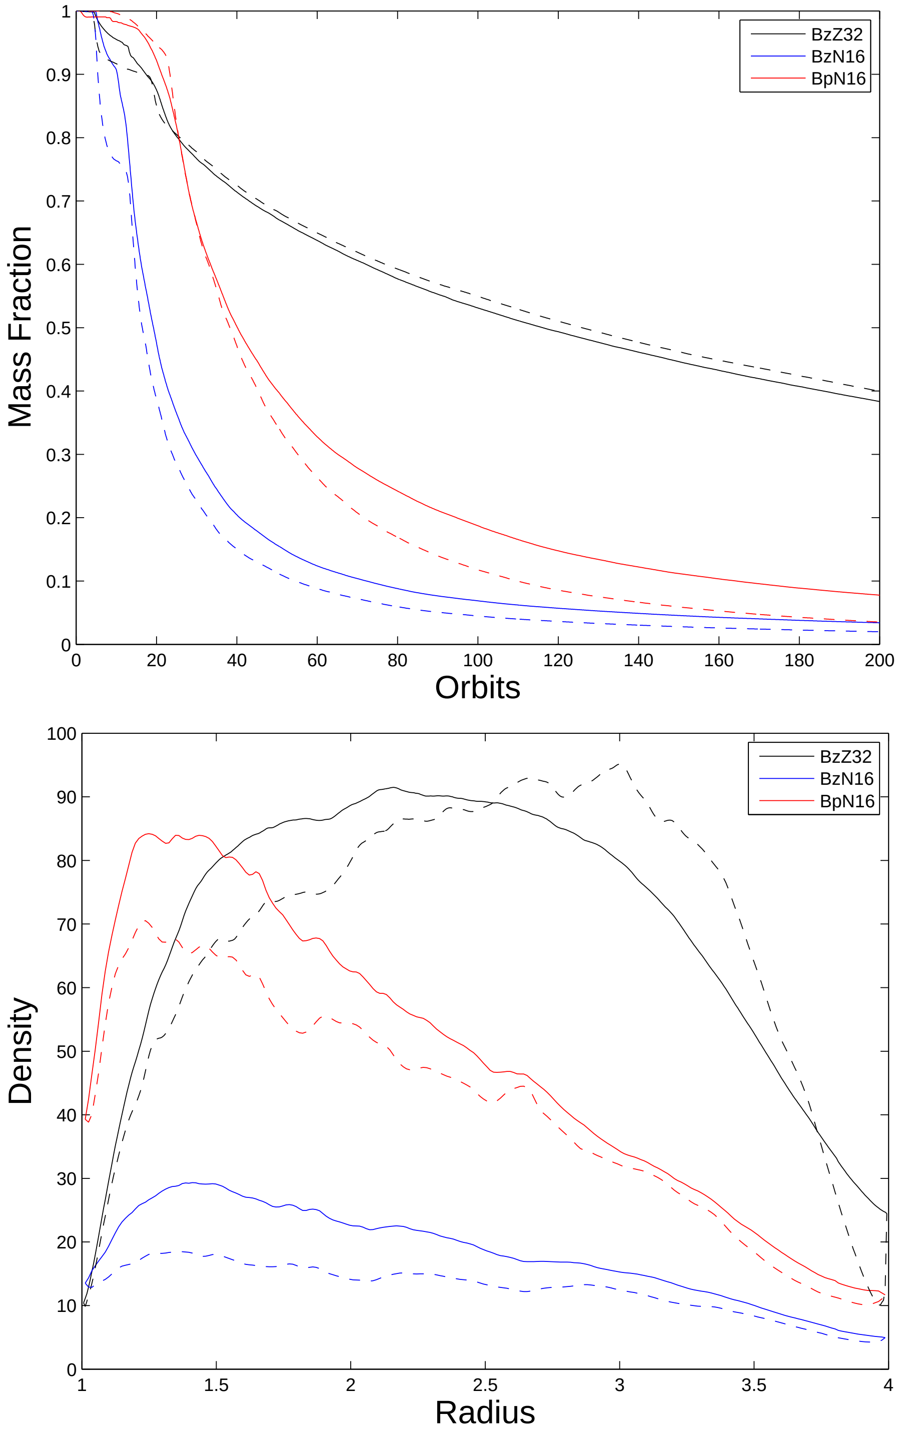

Figure 5.4: Comparison of simulations (solid) and reduced model (dashed). (**Top**) Evolution of mass fraction. (**Bottom**) Radial mass profile at orbit 75.

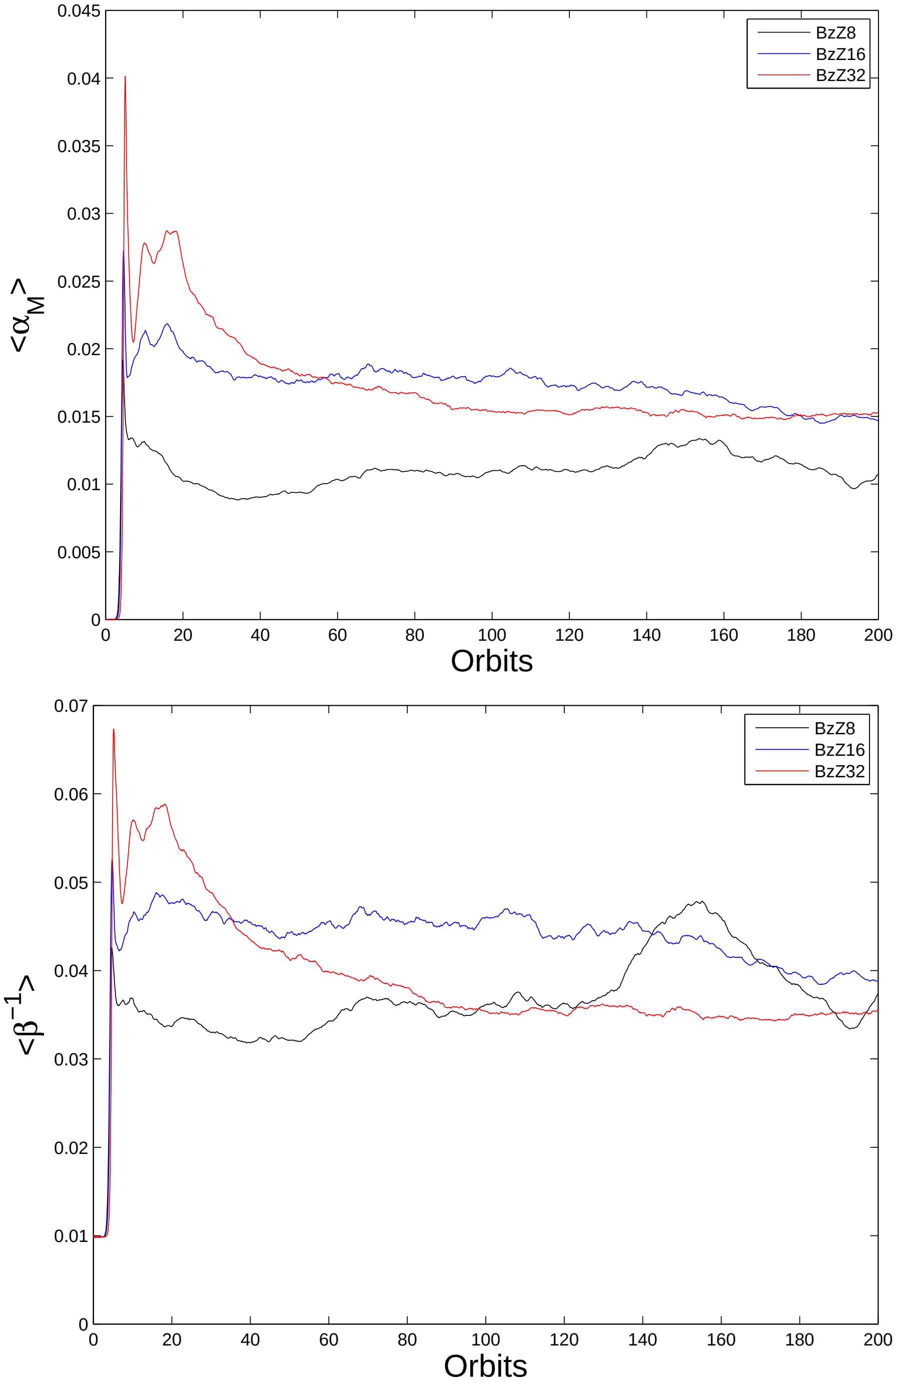

Figure 5.5: Physical convergence metrics, ZF simulations. (**Top**) Accretion efficiency (**Bottom**) Dimensionless magnetic energy

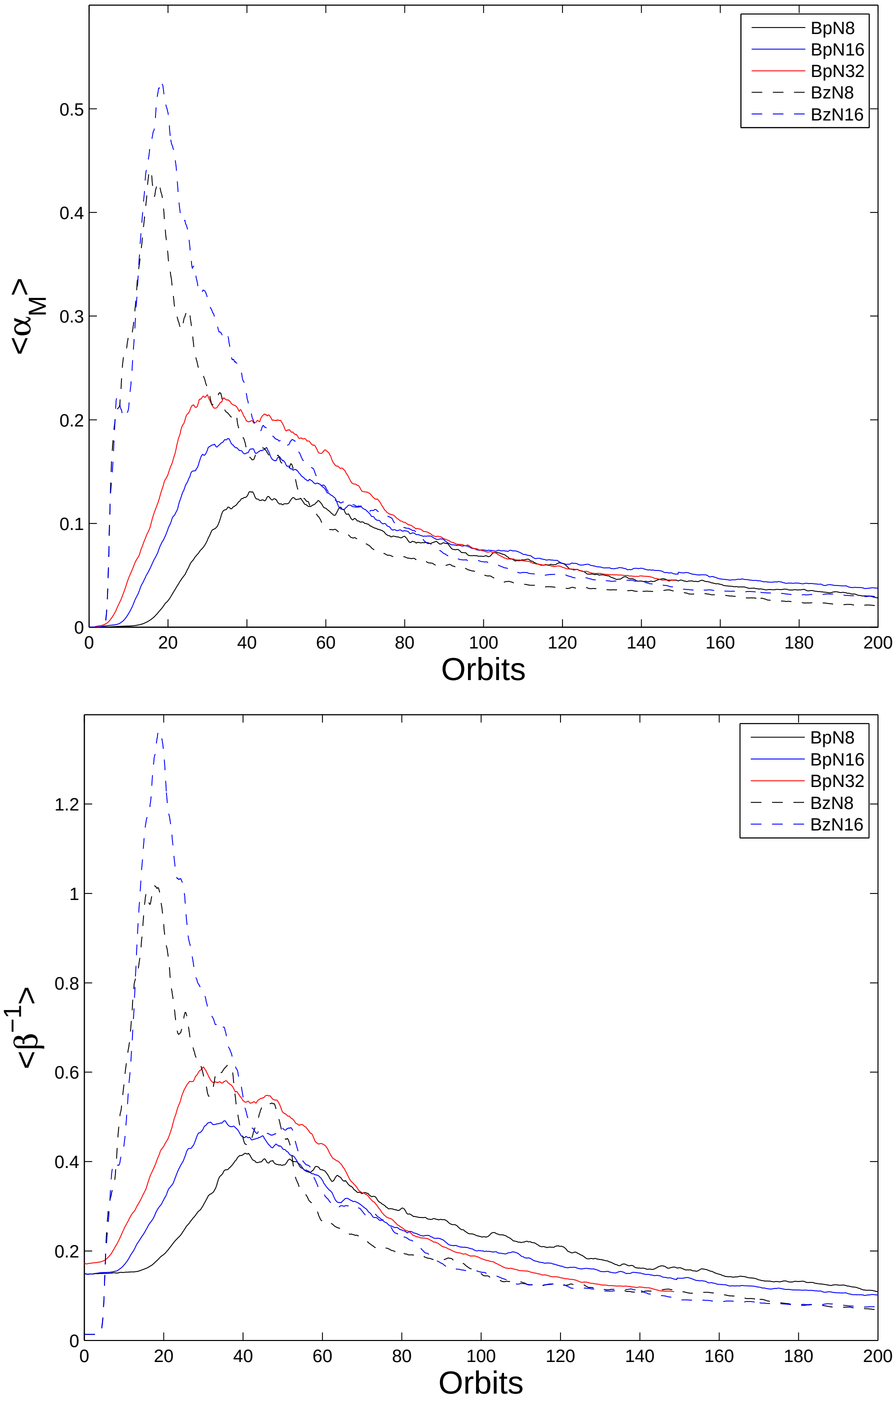

Figure 5.6: Physical convergence metrics, NF simulations. (**Top**) Accretion efficiency (**Bottom**) Dimensionless magnetic energy

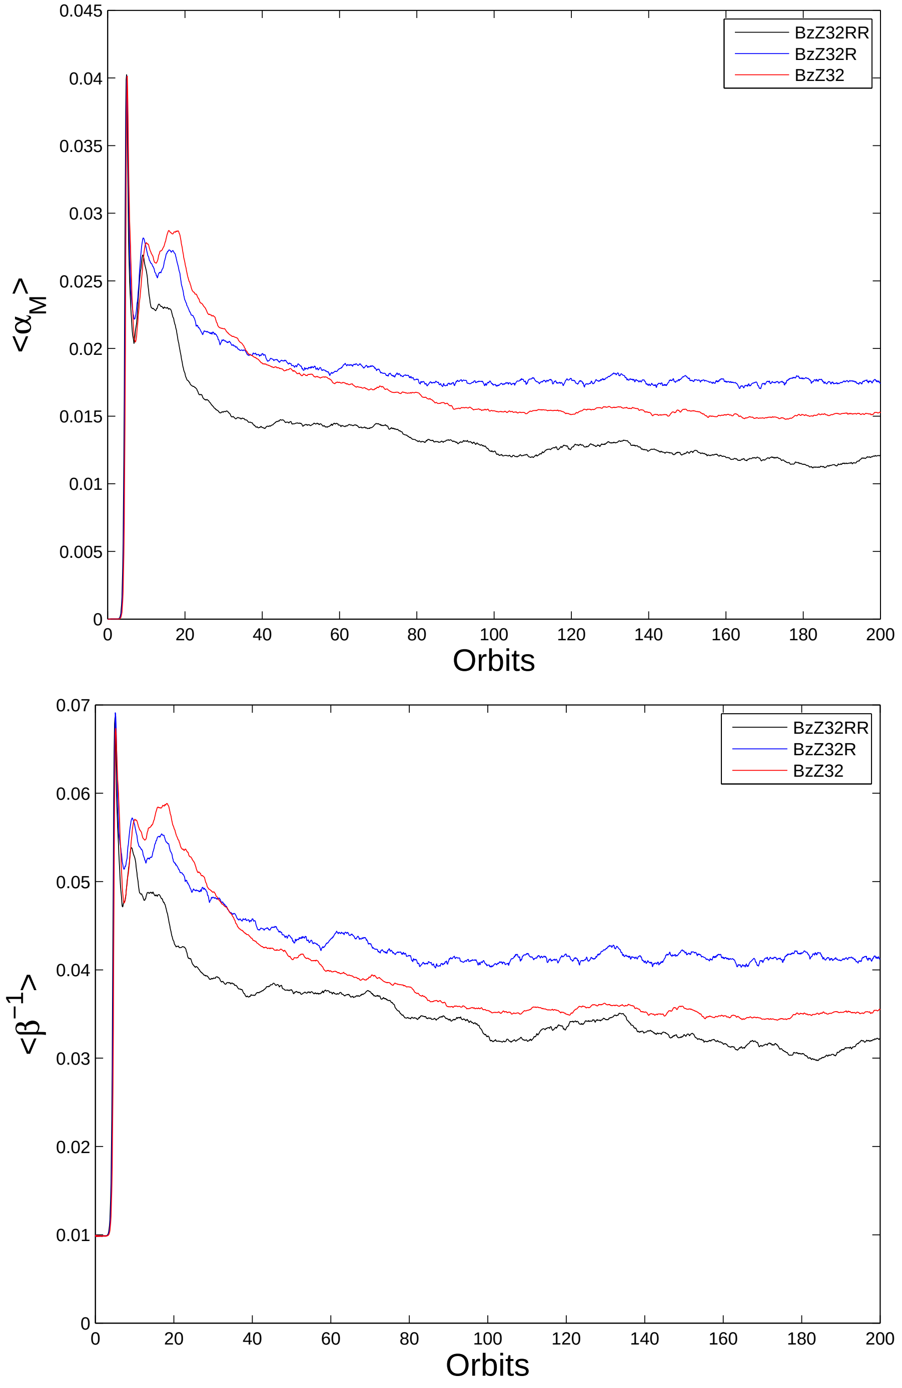

Figure 5.7: Physical convergence metrics for reduced azimuthal resolution runs. (**Top**) Accretion efficiency (**Bottom**) Dimensionless magnetic energy

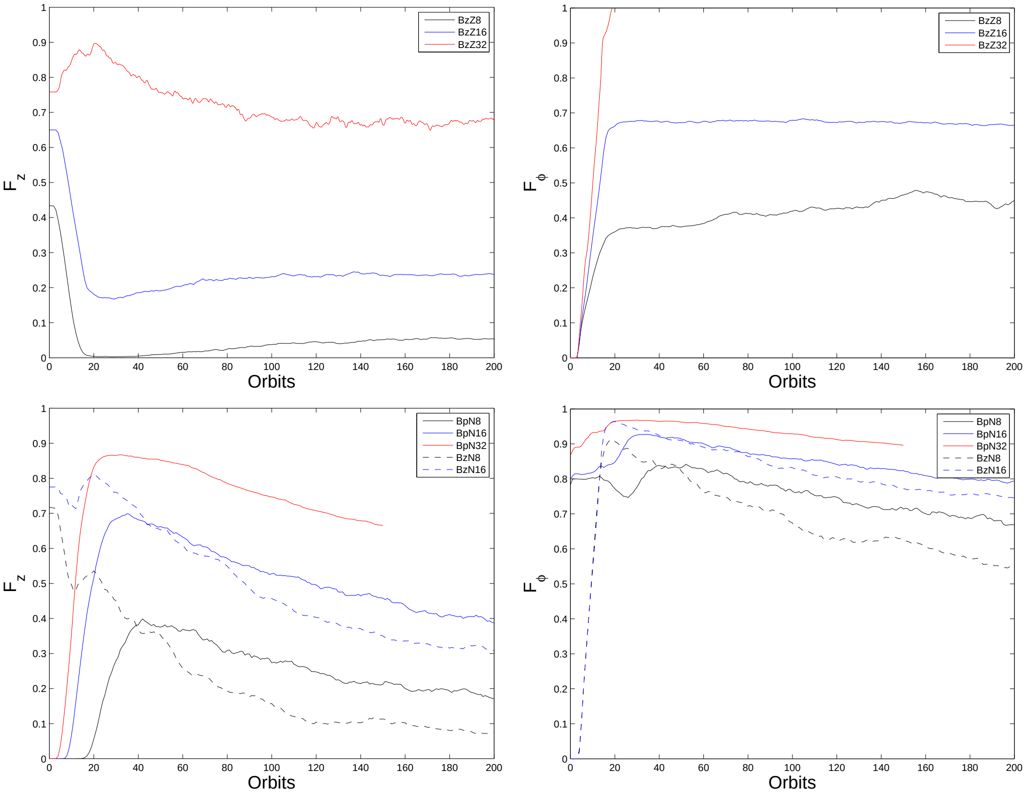

Figure 5.8: Vertical and toroidal resolvability fractions for ZF and NF simulations. (**Top-Left**) Vertical Resolvability, ZF (**Top-Right**) Toroidal Resolvability, ZF (**Bottom-Left**) Vertical Resolvability, NF (**Bottom-Right**) Toroidal Resolvability, NF

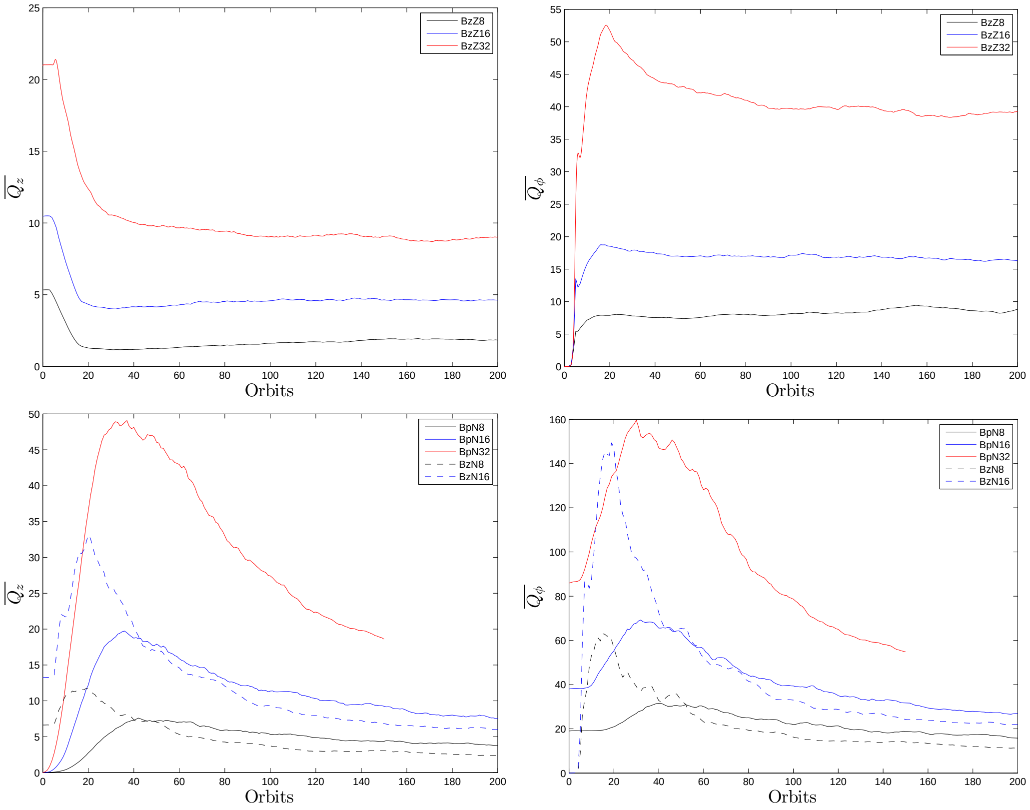

Figure 5.9: Vertical and toroidal quality factors for ZF and NF simulations. (**Top-Left**) Vertical Quality, ZF (**Top-Right**) Toroidal Quality, ZF (**Bottom-Left**) Vertical Quality, NF (**Bottom-Right**) Toroidal Quality, NF

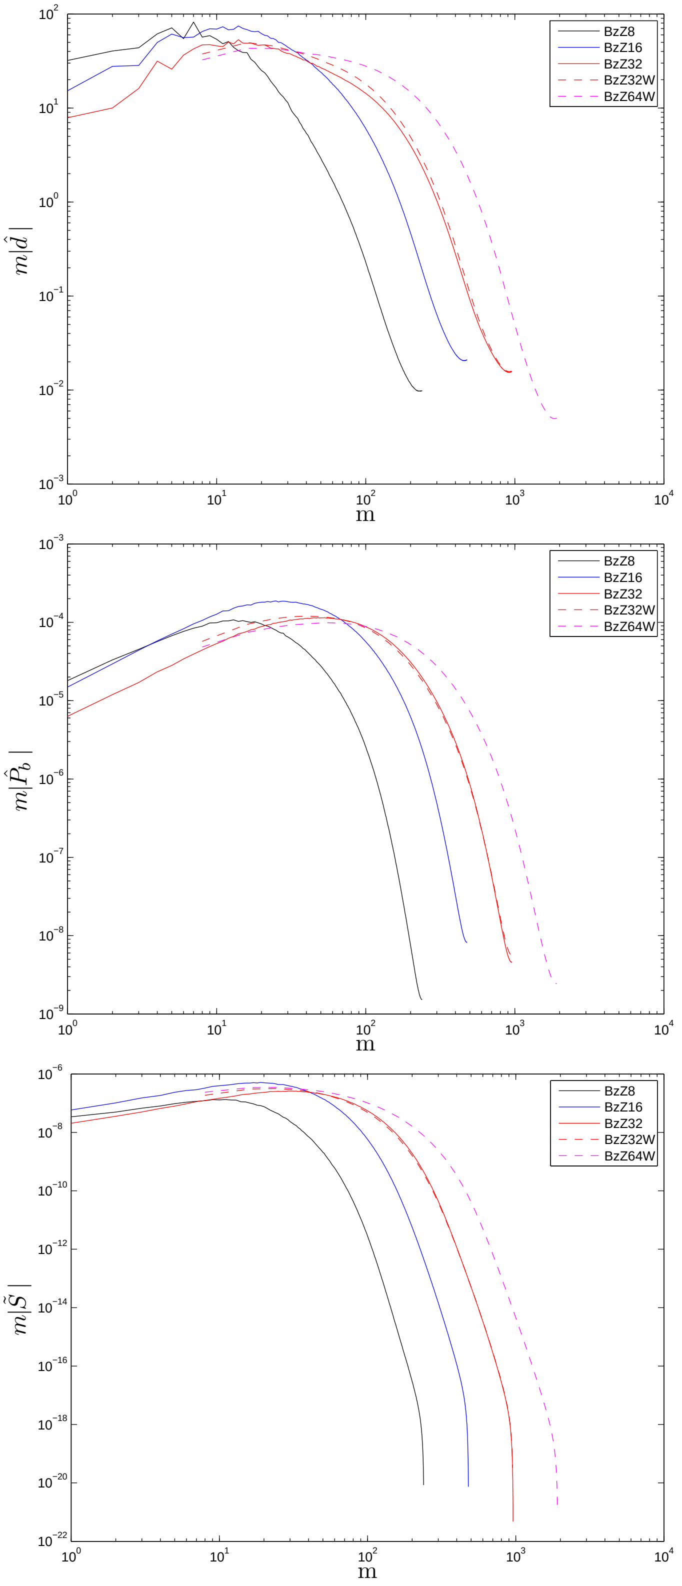

Figure 5.10: Azimuthal distribution of power for the ZF simulations. (**Top**) Density (**Middle**) Magnetic pressure (**Bottom**) Stress ($M_{R\phi}$)

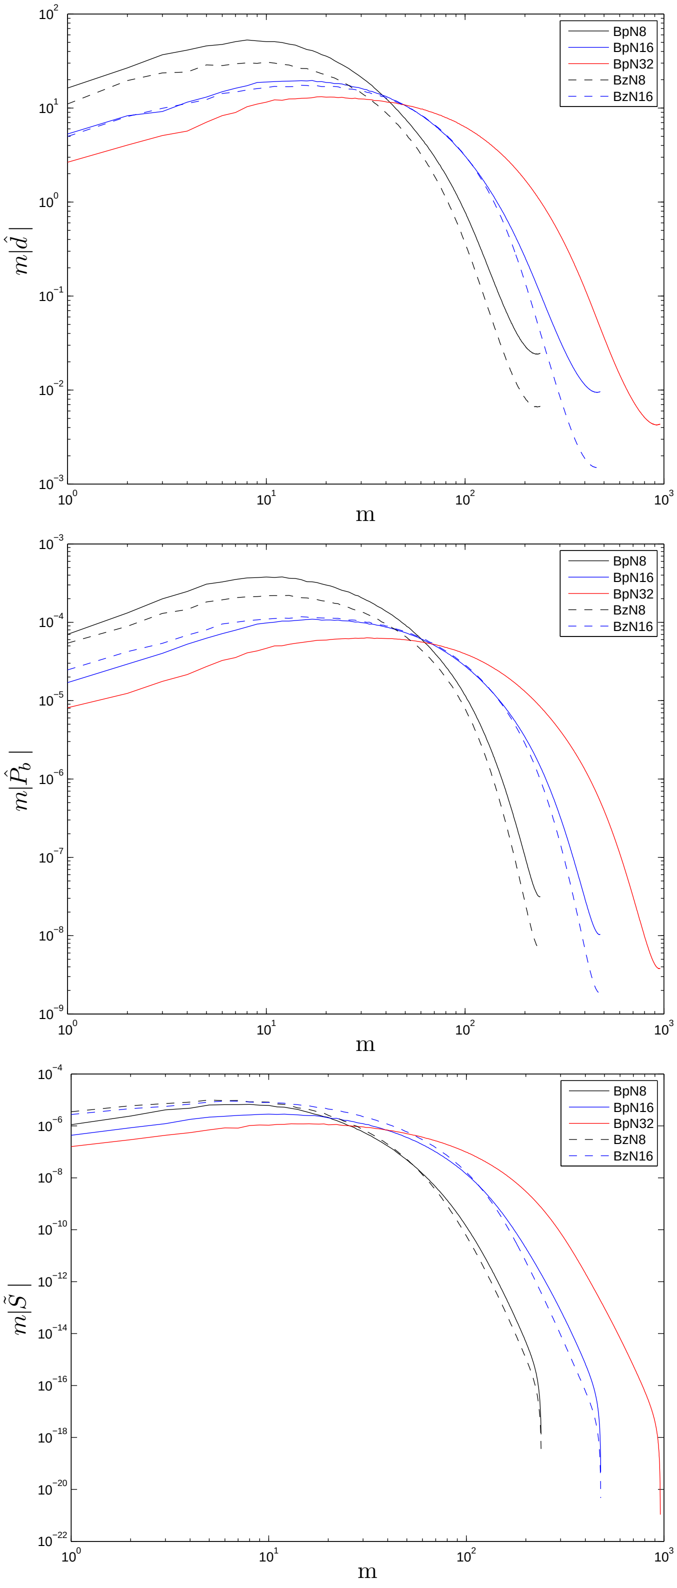

Figure 5.11: Azimuthal distribution of power for the NF simulations. (**Top**) Density (**Middle**) Magnetic pressure (**Bottom**) Stress ($M_{R\phi}$)

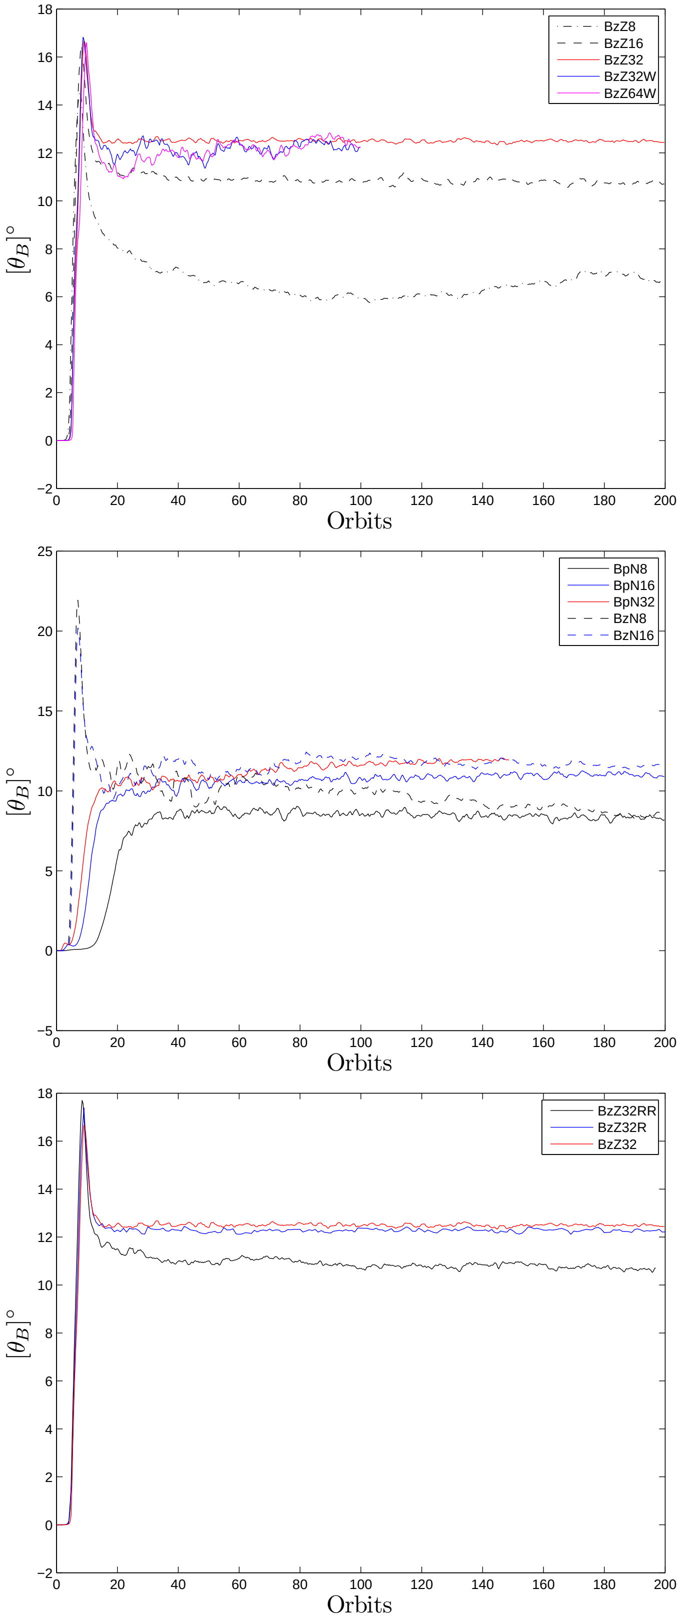

Figure 5.12: Temporal evolution of the magnetic tilt angle. (**Top**) ZF Models (**Middle**) NF Models (**Bottom**) Reduced azimuthal resolution models

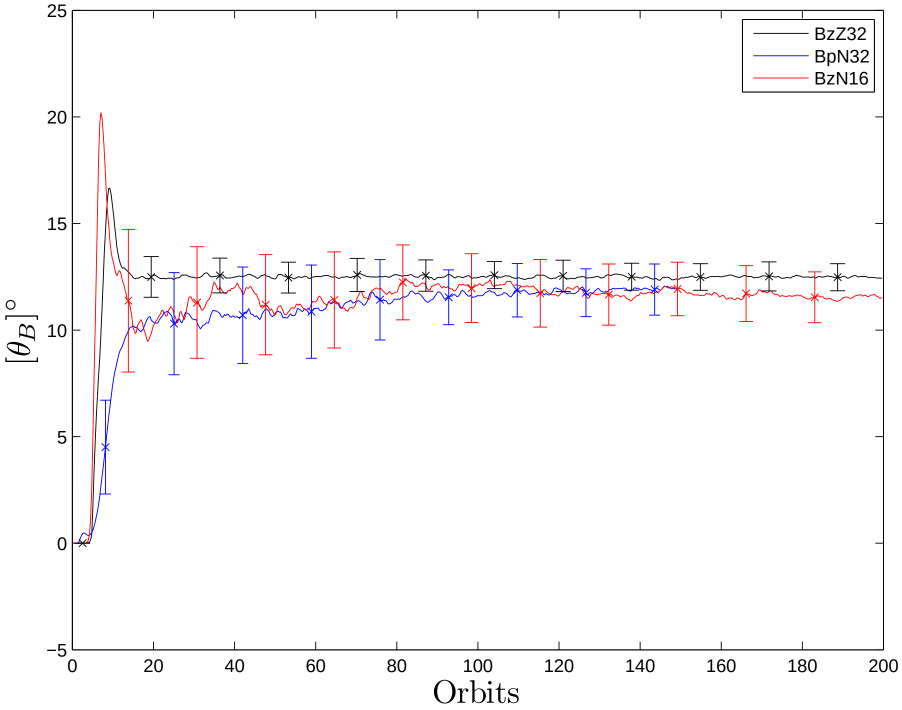

Figure 5.13: Comparison of the tilt angle evolution of the fiducial models. Error bars given by standard deviation of $[\theta_B]$.

Chapter 6

Locality of Global Accretion Disk Turbulence

6.1 Introduction

While the preceding chapter focused primary on numerical issues, specifically the reliability of astrophysically-relevant metrics derived from simulations subject to finite-resolution effects this chapter will return to the material of Chapter 3. In this earlier chapter the physics of stratification and coarse resolution made the comparison of local and global models problematic. Running simulations with orbital advection has facilitated the exploration of resolutions comparable to those used in local models and as such these models provide a vastly superior testing ground to study the issues brought to light in Chapter 3.

Returning to the concept of generating "local" statistics through the decomposition of a global simulation into local analogs we begin by performing an analysis of the instantaneous correlation of magnetic flux and stress in patches of the disk. This analysis is extended from the one presented in Chapter 3 with the inclusion of a study of the toroidal magnetic flux and its connection to local stress. Additionally, the use of a fiducial lengthscale to quantify the flux-stress relationship is introduced and discussed. The ability to predict local stress from the presence of magnetic flux is complemented with a thorough analysis of the distribution and evolution of magnetic flux in the fiducial subdomain. The significant evolution of magnetic flux observed casts doubt on the validity of the flux-constraining boundary conditions commonly employed in local models.

Finally, the comparison of local and global models is extended to smaller lengthscales through the use of two-point correlation functions in the manner of Guan et al. (2009) and Beckwith et al. (2011). The partition of turbulent energy into its components is explored and the resulting hierarchy of scales and rotated principal axes is found to be in strong agreement with previously published works. In particular, agreement between these simulations and local models is found at intermediate- and small-scales. The resolvability of the smallest important scale of turbulence, identified by Guan et al. (2009), is discussed and while not achieved the resolution requirements for the resolvability of this scale is estimated.

The structure of this chapter is as follows. Section §6.2 expands upon the diagnostics introduced in §5.2.2 to define the appropriate metrics of interest for the remainder of the study. Section §6.3 begins the comparison between local and global models with a qualitative set of figures of the fiducial subdomain transformed into an appropriate local geometry. This is replaced in Section §6.3.1 with a more quantitative analysis of the derived local statistics with a focus on their variability. Next, Section §6.3.2 returns to the question of the instantaneous relationship between magnetic flux and stress. This discussion is augmented with an analysis of the distribution and evolution of magnetic flux over the course of the simulation. Section §6.4 focuses on the intermediate- and small-scale structure of MRI-driven turbulence through the use of two-point correlation functions. Finally, Section §6.5 presents a brief summary of the results of this chapter and concluding remarks.

6.2 Diagnostics

Unless otherwise noted the notations and definitions introduced in the previous chapter will be maintained here. The goal of the metrics introduced here will be to facilitate a comparison between local and global models in a quantitative manner. The focus will be on the fiducial subvolume $\mathcal{S}$ and its decomposition defined in the previous chapter.

First, we consider the instantaneous correlation of flux and stress as a means of comparing local and global models. To study the global analogue to these local saturation predictors, we proceed in a manner similar to Sorathia et al. (2010) (Chapter 3). We calculate the instantaneous flux and stress in each subvolume at each timestep after orbit 50, to remove transients associated with the linear growth phase of the MRI. The resulting flux-stress pairs are logarithmically binned according to flux in order to diagnose trends. In contrast to the local model and Sorathia et al. (2010), where flux-stress pairs are scaled by L and $(H/L)^{5/3}$ respectively, we proceed in a manner more appropriate to global simulations in which the size of the subdomain has no physical meaning as the boundaries of the subdomain can not act to suppress the formation of structure on scales larger than the size of the subdomain. The flux-stress pairs calculated here are of the form $(\lambda_{MRI}/\Delta z,\alpha_M)$, as in Beckwith et al. (2011). Scaling the local flux to the grid scale is more directly meaningful and facilitates a closer inspection of the transition point discussed in Sorathia et al. (2010) (Chapter 3).

In addition to studying the relationship between vertical flux and stress, we also consider the analogous connection in the presence of toroidal flux, encapsulated by equation 2.33. Arguably, the presence of strong net toroidal field is more astrophysically relevant than strong vertical field (van Ballegooijen 1989). The evolution of the toroidal MRI is more complex than its vertical counterpart, and there is no simple form for its most unstable mode. We consider flux-stress pairs of the form $(\lambda_c/R\Delta\phi,\ \alpha_M)$, with λ_c given by equation 1.46, and perform the same logarithmic binning as described for the vertical flux-stress pairs.

To further quantify the relationship between the subvolume $\mathcal{S}$ and local models we introduce the mapping

$$(R, \phi, z) \rightarrow (x, y, z) = (R - R_0, R_0\phi, z), \tag{6.1}$$

through which the geometry of the subvolume can be explored in a local manner. To study the structure of turbulence, we follow the model of Guan et al. (2009) and consider the two-point correlation, or autocorrelation, function (ACF) for a fluid variable $f(x, y)$ defined as

$$C_f(\Delta x, \Delta y) = \int_{\mathcal{S}} \delta f(x + \Delta x, y + \Delta y)\delta f(x, y) dx dy, \tag{6.2}$$

where the scaled perturbation δf is given by

$$\delta f(x, y) = \frac{f(x, y) - < f >_y (x)}{< f >_y (x)}, \tag{6.3}$$

specifically the fluctuation about the annulus average scaled by the annulus average. As pointed out by Beckwith et al. (2011), removing radial structure from the fluctuation quantities is important as global simulations can exhibit radial structure due to their evolution that would be missing from local models.

Though the definition of the ACF is given by equation 6.2, the actual manner in which it is calculated is somewhat more circuitous. The calculation utilizes the Weiner-Khinchin theorem which states that the Fourier transform of the two-point correlation is the power spectral density, $|\delta f(k_x, k_y)|^2$. Formulating the calculation in Fourier space is far more efficient due to the significantly reduced computation involved in the calculation of the fast Fourier transform (FFT) compared to the full integral necessary in the initial formulation of the ACF. Utilizing Fourier transforms implicitly assumes periodicity which is satisfied in the azimuthal and vertical direction but will be violated in the radial. This has the potential to adversely affect the ability for the ACFs to accurately capture large radial structure, however the strong shear in these systems suppresses this structure regardless. As will be shown, radial structure decorrelates on a lengthscale much smaller than the radial extent of the subdomain used for these calculations.

The computation is done through the calculation of the full three dimensional FFT over the subvolume followed by a temporal averaging of the power from orbits 65 through 85, resulting in $< |\delta f(k_x, k_y, k_z)|^2 >_T$. The ACF is then simply the inverse FFT of $< |\delta f(k_x, k_y, k_z = 0)|^2 >_T$, as in Beckwith et al. (2011), which are then scaled to have a maximum of unity. An example is shown in Figure 6.1 (Top), which shows C_ρ for simulation `BzZ32`. The structure of the two-point correlation of density, equivalently the gas pressure, is representative of all the correlation functions. Specifically, it is characterized by contours associated with ellipses having major and minor axes rotated from the Cartesian. To characterize the major and minor axes we define

$$\hat{x'} = \hat{x}\cos(\theta) + \hat{y}\sin(\theta), \tag{6.4}$$

$$\hat{y'} = -\hat{x}\sin(\theta) + \hat{y}\cos(\theta), \tag{6.5}$$

where θ is calculated based on the major and minor axes of a fiducial contour. Calculating θ is done by taking the numerically computed contour associated with $1/e$ (assuming a maximum value of unity at the origin) and fitting the general equation of an ellipse. To further quantify the correlation function, we employ the approximation

$$C_f \approx \exp\left[-\frac{|x'|}{\lambda_m} - \frac{|y'|}{\lambda_M}\right], \tag{6.6}$$

where the values of λ_M and λ_m are simply the semi-major and semi-minor axes of the ellipse fit to the $1/e$ contour of the correlation function. As a demonstration of the validity of using an exponential model for the correlation function, Figure 6.1 (Bottom) compares the value of the correlation function and its exponential approximation along both the major and minor axes for the example correlation function shown in Figure 6.1 (Top). Thus the correlation function, C_f, is determined solely by the parameters θ, λ_m, and λ_M.

6.3 Comparing Local and Global Models

The restricted geometric domain of local models makes them ideal to explore accretion disk turbulence at resolutions far in excess than what is feasible in global models. However ensuring that these local models accurately represent the small-scale dynamics of accretion disk turbulence is a vital validation of the wealth of

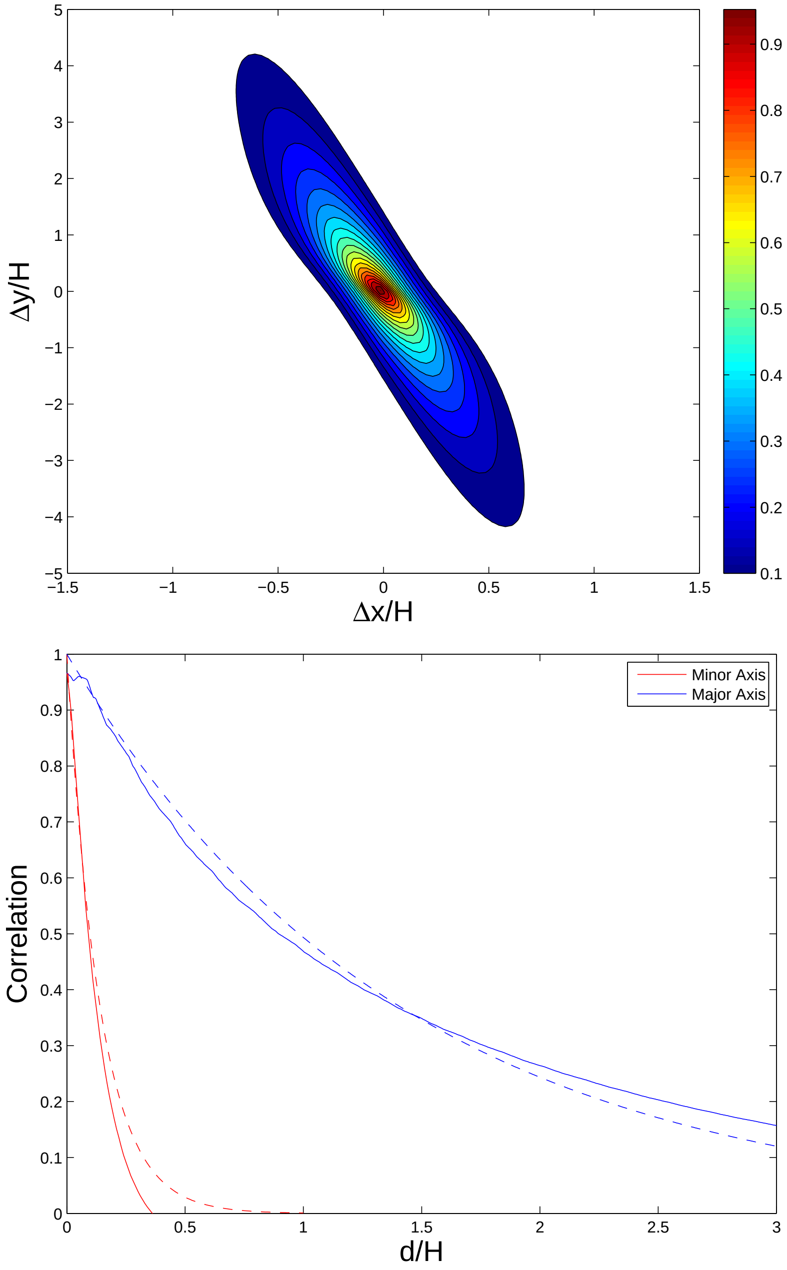

Figure 6.1: (**Top**) Example correlation function, C_ρ from simulation `BzZ32`. (**Bottom**) Comparison of the correlation function (equation 6.2, solid) and its exponential approximation (equation 6.6, dashed) along the major and minor axes of C_ρ from simulation `BzZ32`.

results learned from local models. The simulations presented here provide a unique opportunity to accomplish as the resolutions of these global models are not only comparable to, but greater, than the majority of local simulations in the literature.

Towards this end, we proceed in a manner similar to Sorathia et al. (2010) (Chapter 3) and decompose the global simulation domain into an ensemble of small regions that are comparable to shearing box models as described in detail in §5.2.2. To provide a sense of the structure of the subdomain, $\mathcal{S}$, we employ the transformation to a Cartesian domain through the mapping given in equation 6.1. To further facilitate the comparison to a shearing box, we consider the quantity

$$v_y = v_\phi - v_K(R_0). \tag{6.7}$$

Figure 6.2 (Top) shows the quantity v_y/c_s in the transformed subdomain, $\mathcal{S}$, at 100 orbits in simulation `BzZ64W`. Similarly, Figure 6.2 (Bottom), illustrates the toroidal magnetic field for the same time and simulation. We note the visual similarity between these images and those presented from shearing box simulations.

6.3.1 Evolution of Local Ensemble

The benefit of treating a global simulation as a local ensemble is that it allows us to study the local dynamics of accretion disk turbulence over a large range of parameter space simultaneously. Additionally, it offers the opportunity to directly test the importance of curvature terms and the degree to which local flux is truly conserved. Local statistics can be used not only to study the average value of a quantity but also the variation in the quantity and how it evolves. Figure 6.3 illustrates the

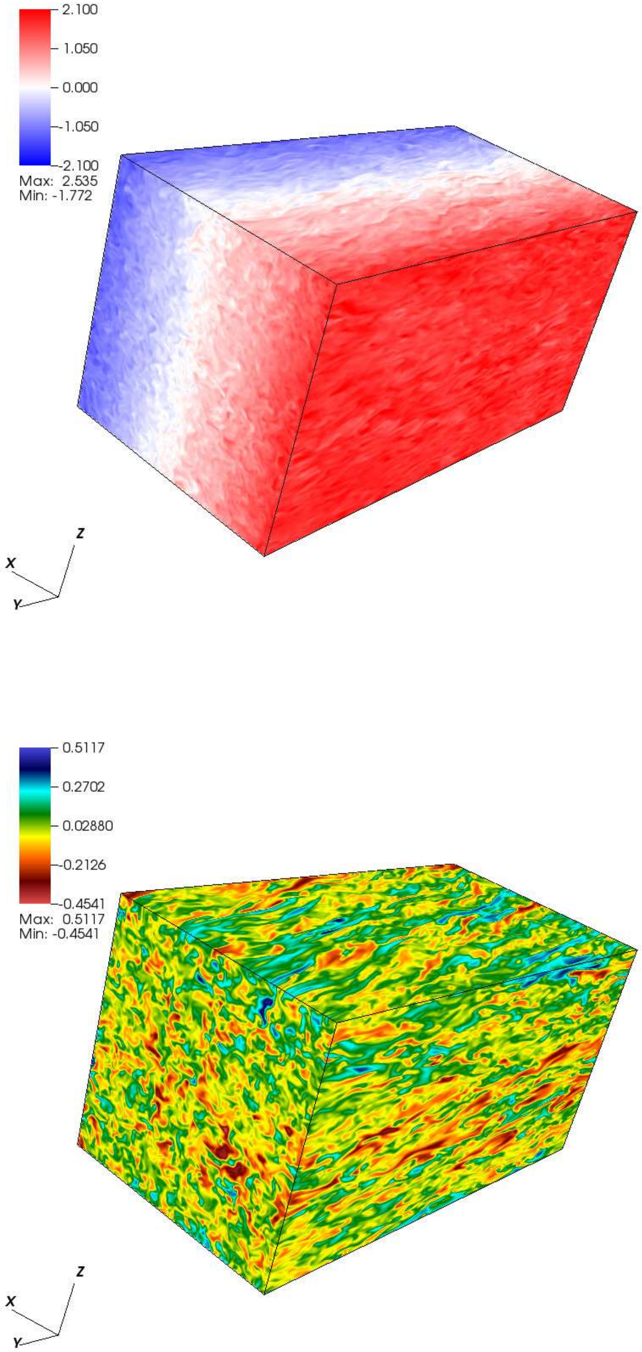

Figure 6.2: Still images of subvolume of simulation `BzZ64W` at orbit 100 mapped to the local Cartesian frame. (**Top**) Azimuthal velocity, v_y/c_s (**Bottom**) Azimuthal magnetic field

time evolution of $[\alpha_M]$ and $[\beta^{-1}]$ for a selection of ZF and NF simulations. Comparing $[\alpha_M]$ and $[\beta^{-1}]$ in the ZF simulations, Figure 6.3 (Top panels), illustrates the resolution dependence analogous to the physical metrics considered above, however the variability of these quantities is inversely related to the resolution. The behavior of run `BzZ8` is statistically separated from the behavior of the higher-resolution simulations, with runs `BzZ16` and `BzZ32` demonstrating more comparable values of these quantities.

The behavior and evolution of the accretion efficiency and scaled magnetic energy for a selection of NF runs is given in Figure 6.3 (Bottom panels). Most striking is the extreme initial transient associated with run `BzN16`, characterized by values of α_M well in excess of unity and strongly magnetized regions with Alfven speeds several times the sound speed. This transient, however, is short lived with the long-term evolution comparable to that of runs `BpN16` and `BpN32`. The toroidal runs, like the ZF simulations, display an inverse relationship between variability and resolution. The resolution dependence of the variability is itself interesting, while the global treatments of these quantities discussed above can demonstrate their convergence that the variability weakens with resolution suggests a separate resolution requirement.

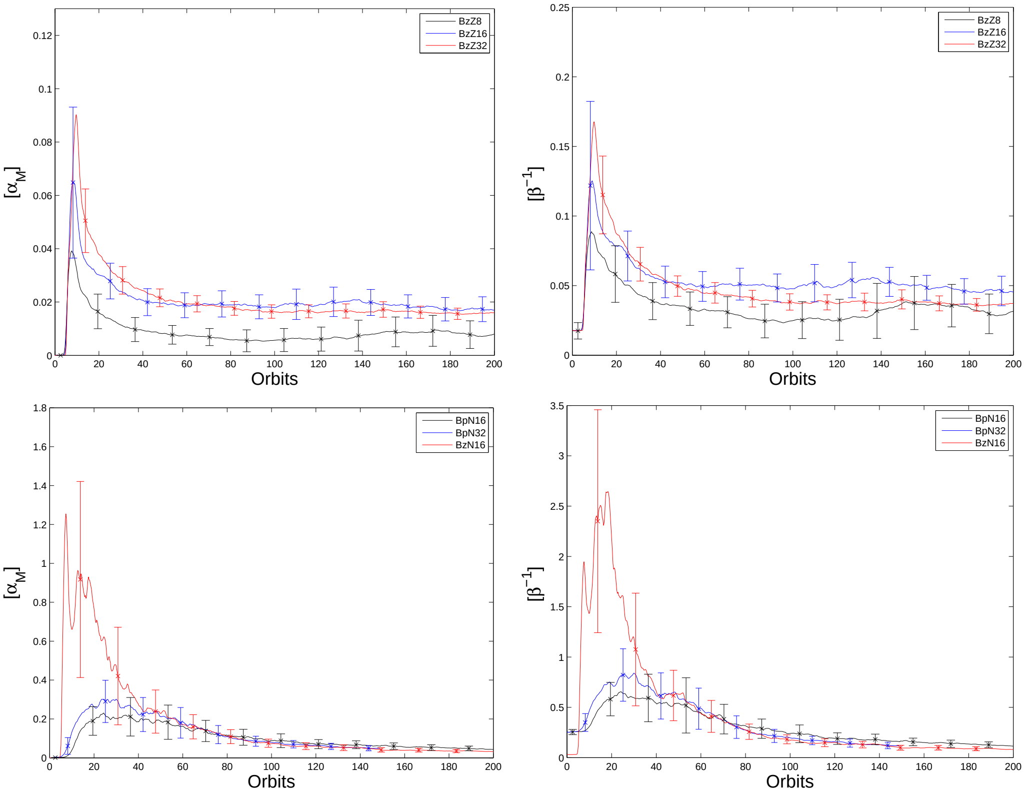

Figure 6.3: Evolution of the local ensemble. (**Top-Left**) Accretion efficiency, ZF (**Top-Right**) Magnetic energy, ZF (**Bottom-Left**) Accretion efficiency, NF (**Bottom-Right**) Magnetic energy, NF

6.3.2 Distribution and Evolution of Magnetic Flux and its Relation to Stress

In Sorathia et al. (2010), we make a preliminary effort towards connecting local and global models of accretion disks using the method described in §5.2.2 and §6.2 to generate "local" statistics. The simulations considered were useful but suffered due to the added difficulty of separating effects of stratification and resolvability from the desired comparison. However that the resolvability of the linear MRI in saturated turbulence is important or sufficient is not clear from analytic theory. The simulations considered in this work are designed from first principles to be more directly applicable to the comparisons we would like to make. We have previously considered the instantaneous correlation between vertical flux and stress; here we also study the relationship between toroidal flux and stress. That there is a structure to the relationship between toroidal flux and stress is interesting (Figure 6.4, Right panels). At first glance, one might not expect a transition in the correlation, since the comparison is between toroidal field and stress, a term dominated by the toroidal field. The nature of this correlation is unclear, as in contrast to the vertical MRI which grows quite rapidly and would be expected to correlate to the presence of field, it is not clear that the toroidal MRI would be associated with timescales small enough to correlate to the presence of field. Assuming that the connection between toroidal flux and stress is causal, as is believed to be the case in the vertical flux-stress correlation, it may suggest that toroidal flux generates stress in the non-linear regime in a manner much faster than it would in the linear. An alternative possibility is that the correlation between toroidal flux and stress is a consequence of the vertical flux-stress relationship. In this case, the presence of vertical flux would drive stress with toroidal flux amplification as an intermediate step. This possibility would suggest that the relationship between the turning points in the two flux-stress relationships is meaningful. We note, that while the fluxes are scaled to the grid resolution these are not the same as the quality factors presented earlier as the flux is averaged over each wedge as opposed to taken on a cell by cell basis. This averaging reduces the strength of the net field which results in lower values than the quality factors.

The calculated flux-stress relationship for the ZF simulations is given for both vertical flux and toroidal flux in Figure 6.4 (Top panels). It was noted in Chapter 3 that the transition between the flat unresolved region and the linear resolved region suggested a seeming super-resolvability, specifically that we saw this transition at $\lambda_{MRI} \approx \Delta z/20$. We note similar behavior in run `BzZ8`, in Figure 6.4 (Top-Left), in which the transition occurs at $\lambda_{MRI} \approx \Delta z/10$. The remaining simulations exhibit a transition at $\lambda_{MRI} \approx \Delta z$, in agreement with the scaling law given by equation 2.35 (Pessah et al. 2007) and the simulations of Beckwith et al. (2011). Related to this, we see a qualitatively similar resolvability threshold in the toroidal flux-stress relationship. While the two highest resolution simulations exhibit a transition at approximately $10R\Delta\phi$, the transition for `BzZ8` occurs at approximately half of that.

The exact cause of the differences between run `BzZ8` and its higher resolution counterparts, namely the weaker stress in the unresolved flux-stress regime and the seeming ability to resolve a sub-gridscale λ_{MRI}, is unclear. The former is interpreted as a consequence of the "density" of resolved regions, in a manner similar to the resolvability fractions presented in Chapter 5. In run `BzZ8` the fraction of the computational domain in which the MRI is properly resolved is quite small and thus one would expect the stress in the unresolved flux-stress regime to be a remnant of when the subvolume last encountered resolvable flux, which would likely be much longer in the lowest-resolution simulation than in the higher-resolution counterparts. Equivalently, the stress in this regime is a measure of the global properties of the disk and should depend on the timescale required for the presence of magnetic flux to drive stress. The nature of the transition point is interpreted in the context of the slowest appreciably growing mode (SAGM) discussed in Section §3.4.3. The SAGM argument supposes that the transition point is governed by more slowly growing MRI modes that are able to grow significantly before being truncated through some aspect of the nonlinear physics of the saturation. If this truncation is caused by the action of parasitic modes, as suggested by Goodman & Xu (1994), the timescale τ used in Section §3.4.3 may itself have a resolution dependence due to the resolvability of the parasitic modes. This suggests the counterintuitive result that the lower resolution simulation `BzZ8` super-resolves the MRI due to its inability to properly resolve parasitic modes. It is worth noting, that both of these interpretations are speculative and more detailed analysis will likely be needed to settle the issue. It is clear, however, that the seeming super-resolvability of run `BzZ8` is unphysical and an attribute of poor resolution. More physically, we interpret the transition point of simulations `BzZ16` and `BzZ32` as also being a consequence of the SAGM as we do not expect $\lambda_{MRI} = \Delta z$ to be resolved.

Next we consider the flux-stress relationship for the net field simulations for both vertical flux and toroidal flux in Figure 6.4 (Bottom panels). We note that in general the net flux runs exhibit a very similar general structure to the net-zero flux runs. Specifically, the transition points between the flat and linear stress-response regimes are independent of initial field topology. None of the net-field simulations considered exhibit the super-resolvability seen in `BzZ8` which suggests that even low resolution net-field runs are more resolved than their zero-flux counterparts. The general trend is that, with increasing resolution, the stress response to unresolved flux (*i.e.* flux corresponding to wavelengths above the transition point) is generally increased and the slope of the linear regime, in which flux is resolved, becomes shallower. An interesting feature is seen in the stress response to vertical flux for the net toroidal field runs. For vertical flux $\lambda_{MRI}/\Delta z \approx 6$ we see a flat stress response. This region occurs well below what we would expect for the MRI-stable region in which the most unstable vertical mode exceeds the vertical domain of the simulation, and we also note that the net vertical field runs maintain a linear slope in this region.

A further comparison we can make between local and global models is based not just on the structure of the flux-stress relationships, but on the precise manner in which increased flux generates an increased stress response. While the dependence of local saturation predictors on box-size prevent direct comparisons, we can work backwards and use the slope of the resolved region to define an effective local box

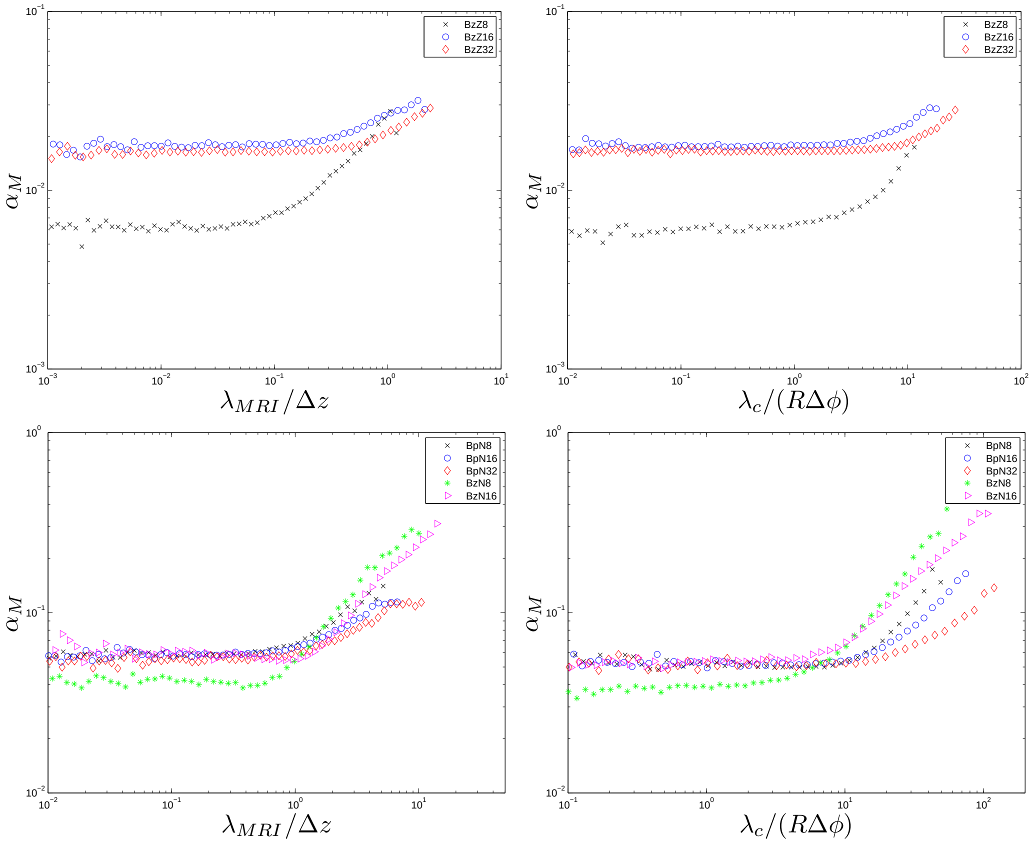

Figure 6.4: Correlations between vertical and toroidal magnetic flux and stress, ZF and NF runs. (**Top-Left**) Vertical flux, ZF (**Top-Right**) Toroidal flux, ZF (**Bottom-Left**) Vertical flux, NF (**Bottom-Right**) Toroidal flux, NF

size. The physical meaning of this length scale is unclear, and is presented here as merely a convenient manner in which to quantify the slope of the linear-response regime. Formally, we use the saturation predictors given by Eqn 2.35 and Eqn 2.33 to constrain the value of the box-size to fit the instantaneous correlations we find. The values of box-size from the vertical flux-stress and toroidal flux-stress are denoted ℓ_z and ℓ_ϕ respectively.

The values of this effective box-size are given in Table 6.1. Both the vertical and azimuthal length scales are monotonic with resolution, with the net field simulations corresponding to larger effective box-sizes. Of interest is the approximate relationship, $\ell_z \approx 2\pi\ell_\phi$, in particular due to its connection with commonly used local domain sizes. Also of note is that by reducing the azimuthal domain we see a significant drop in ℓ_z, while ℓ_ϕ remains roughly constant.

Simulation	ℓ_z/H_0	ℓ_ϕ/H_0
`BzZ8`	0.348	1.838
`BzZ16`	0.899	4.956
`BzZ32`	0.964	5.040
`BzZ32W`	0.283	5.970
`BzZ64W`	0.372	7.696
`BpN8`	0.487	3.377
`BpN16`	0.674	4.612
`BpN32`	0.905	4.953
`BzN8`	0.472	5.552
`BzN16`	0.731	8.255

Table 6.1: Effective box-size for global simulations.

The manner in which a localized region of the disk responds to the instantaneous presence of flux makes up an important part of the dynamics of the disk, the natural analog to this is the distribution of magnetic flux through the localized regions of the disk. Here, we focus on the latter. We compute a distribution function of the vertical and toroidal magnetic flux through each wedge during the QSS. As in the flux-stress diagrams, these fluxes are scaled to the grid resolution. The distribution is calculated using logarithmic binning in the flux (80 bins per decade) and the result is smoothed using a ten bin moving window. The results for simulations `BzZ32`, `BpN32`, and `BzN16` are displayed in Figure 6.5.

Of interest is the fact that for simulation `BzZ32`, the majority of the disk is below the threshold in which there is a linear stress response to flux. In contrast to this, simulations `BpN32` and `BzN16` have the majority of the disk well above the resolvability threshold of $(Q_z, Q_\phi) \approx (1, 10)$. Related to the initial conditions, `BpN32` is biased towards toroidal flux and `BzN16` is biased towards vertical. Forming an intermediate case, `BzZ32` possesses a more even distribution of flux albeit weaker than the net field cases.

An alternative manner by which to consider the distribution of magnetic flux is through the consideration of the local statistics as representative of the trajectory of an ensemble. That is, to consider the path $([Q_z](t), [Q_\phi](t))$ and its temporal evolution. Figure 6.6 illustrates the trajectory for simulations `BzZ32`, `BpN32`, and `BzN16` where markers are placed on the trajectory approximately every 8 orbits. For the net field simulations the evolution of the flux during the initial transient is quite stark, resulting in migration from the axes to large values of both azimuthal and vertical flux, specifically large enough to exceed the transition point to linear

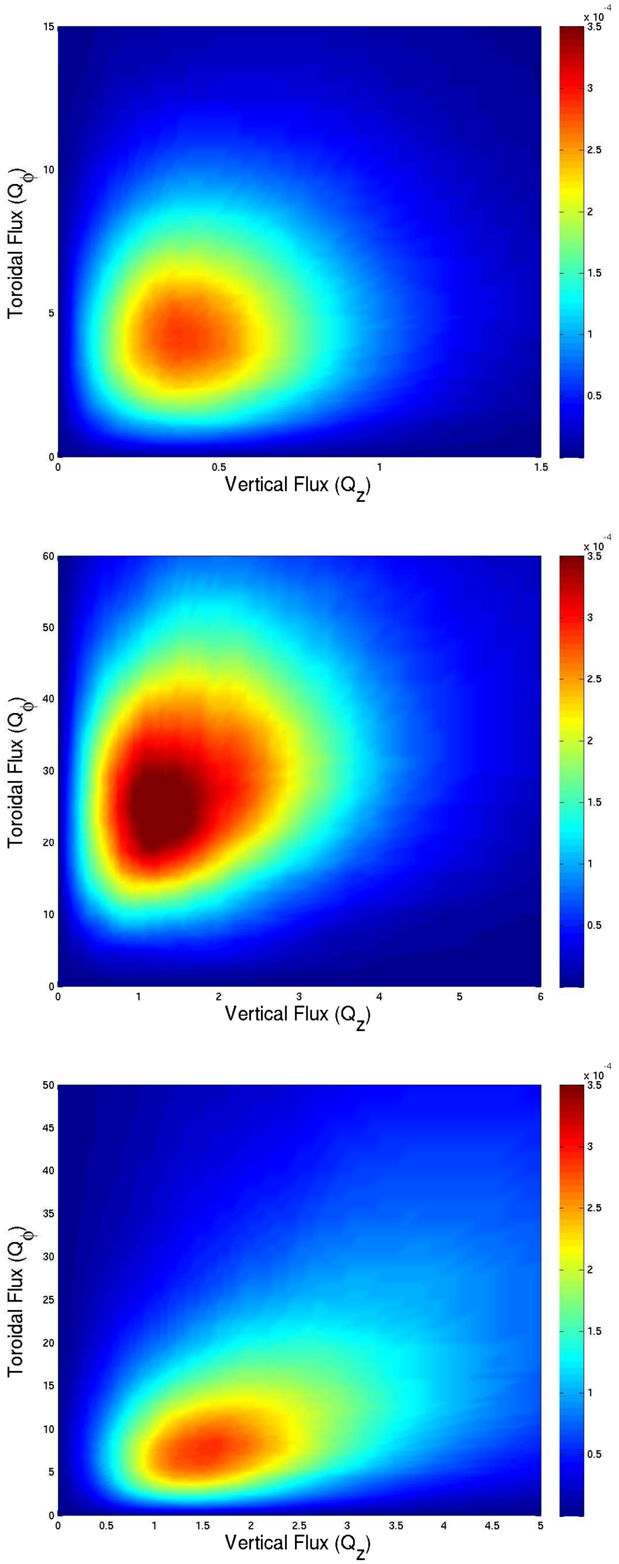

Figure 6.5: Distribution of magnetic flux in the QSS. (**Top**) `BzZ32` (**Middle**) `BpN32` (**Bottom**) `BzN16`

stress response. After the initial transient, the net flux simulations are characterized by a general decrease in the mean flux. In contrast to the net flux simulations, the trajectory of `BzZ32` is constrained to a much smaller region of the flux-space. The trajectory only briefly crosses into the region associated with a linear stress response, unlike the net flux simulations whose trajectories spend the majority of the simulation in this region. Compared to truly local simulations, in which the boundary conditions prevent migration in the $Q_z \times Q_\phi$ space, global simulations are characterized by significant evolution particularly those initialized with a net flux. This suggests the possibility that local simulations initialized with a net flux could result in an artificially-maintained angular momentum transport due to the boundary conditions preventing the advection of flux off the grid. The trajectory of simulations `BzZ32` and `BzN16` evolve to possess comparable fluxes despite exhibiting significantly different behavior during much of the simulation. The limited temporal domain for which simulation `BpN32` was evolved prevents making a similar comparison, however the behavior is consistent with the possibility of evolving to a state of magnetic flux comparable to simulations `BzZ32` and `BzN16`. Run `BpN16`, while run for 50 orbits more than `BpN32`, displays similar behavior that is consistent with this possibility but does not achieve it during the time domain simulated.

6.4 Spatial Scales and Orientations of MRI-Driven Turbulence

In Chapter 5 the dependence of large-scale quantities like the global angular momentum transport and dimensionless magnetic energy upon resolution and initial

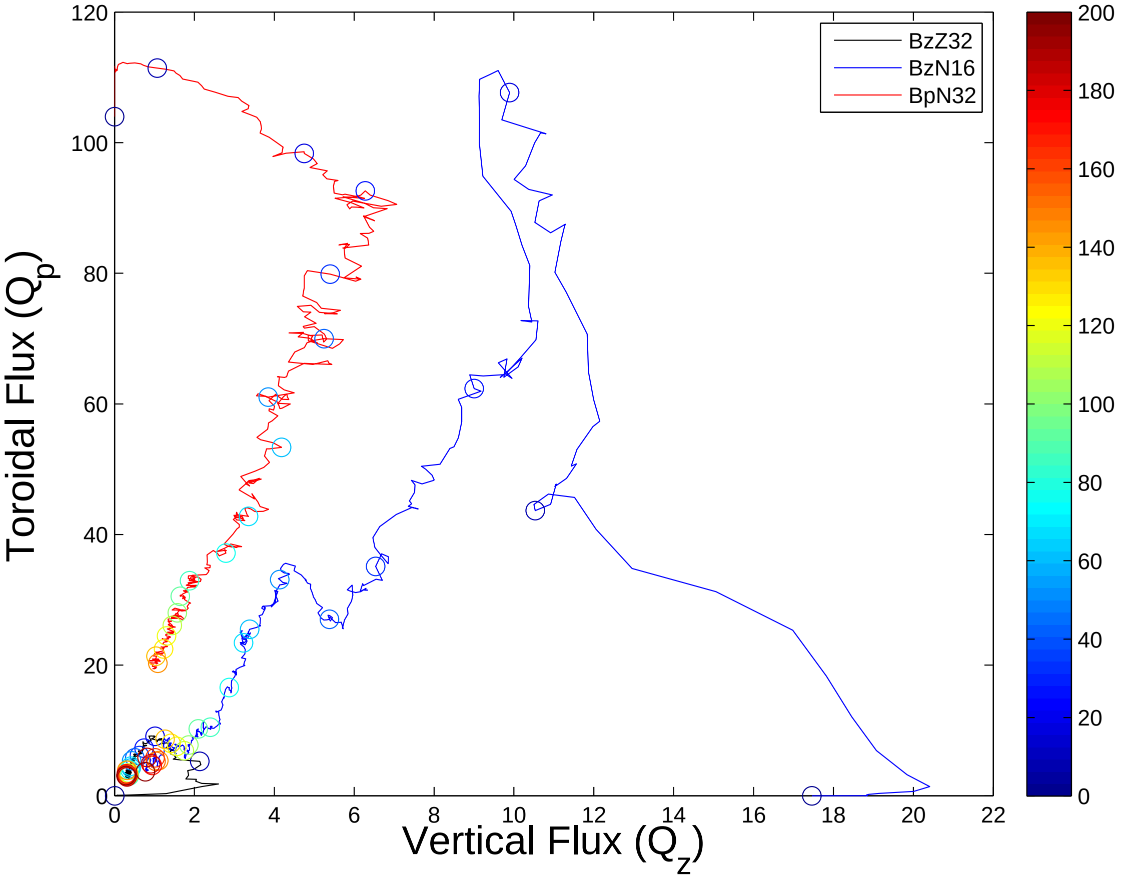

Figure 6.6: Evolution of mean quality factors of the local ensemble. Circular markers placed approximately every 8 orbits, with color determined by the orbit associated with the marker.

magnetic topology is explored. Here, the use of the two-point correlation functions allows a complementary analysis of the small-scale structure of the MRI-driven turbulence. The choice of quantities in the analysis that will be presented is based on the investigations of Guan et al. (2009), in the context of unstratified local models. Their analysis has subsequently been extended to large shearing boxes (Davis et al., 2010) and global simulations (Beckwith et al. 2011, Nelson & Gressel 2010). The analysis will focus on a comparison of the small-scale structure encapsulated by the ACF with prior works.

The energetics of an accretion disk are dominated by the energy associated with the Keplerian rotation which acts as a energy reservoir that powers turbulence. The energy of the turbulence itself is partitioned into three sources: the gas pressure, P, which for an isothermal disk is equivalent to the density structure; the turbulent kinetic energy, $K' = \rho|\mathbf{v}'|^2$; and the magnetic pressure, P_b. The associated correlation functions are denoted C_P, $C_{K'}$, and C_B respectively.

Prior to the presentation of the full data, we consider the two-point correlation functions for simulation `BzZ32`. Figure 6.7 shows the correlation functions C_P (Top), C_B (Middle), and $C_{K'}$ (Bottom). Immediately obvious is the hierarchy of scales exhibited in the turbulent energy partition in which the decorrelation lengthscale of the gas pressure is much larger than that of the magnetic pressure, which in turn is much larger than that of the turbulent kinetic energy. Additionally, while the contours of the gas and magnetic pressure are highly eccentric, those of the kinetic energy are nearer to circular in nature. As we will see, this hierarchy of scales and differing eccentricity between the pressure and kinetic energy terms is characteristic of the behavior over all the simulations considered and identical to the results presented by Beckwith et al. (2011) (*cf.* Figure 20). Also of note is the difference in structure of C_P compared to $C_{K'}$. Whereas the latter exhibits quite uniform contours, the former exhibits a "twisting" at large displacement. This point is returned to below.

The correlation functions associated with the components of the turbulent energy are shown in Table 6.2 for all the simulations using an isotropic resolution. The behavior exhibited in Figure 6.7 is broadly applicable; the major axes associated with the gas pressure are larger than those associated with the magnetic pressure and similarly, the magnetic pressure exhibits a larger lengthscale than the turbulent kinetic energy. While the ratio of major to minor axis (λ_M/λ_m) is roughly ten for the pressure terms, this ratio is approximately between four and five for the turbulent kinetic energy. Guan et al. (2009) interpret the large scale correlations associated with the gas pressure as a manifestation of acoustic waves that act to set the largest scale of the turbulence.

Independent of the eccentricity of the two-point correlation functions is the orientation of the principal axis, specifically the angle θ. In this case, the comparison to be made is between the gas pressure and perturbed kinetic energy and there is a contrast between these and the magnetic pressure. The principal axis of the gas pressure and perturbed kinetic energy are, in general, comparable and (at the highest resolutions) differ appreciably from the principal axis associated with the magnetic pressure. This is in contrast to the results of Guan et al. (2009) in which the gas pressure exhibits a larger angle than the perturbed kinetic energy, although their results also indicate the largest angle is associated with the magnetic pressure. The cause of this discrepancy is unclear, but may be a consequence of the reduced domain considered in their local models. Overall, with increased resolution the principal axis is rotated further from $\hat{y}$ and appears to be slowly converging although this can not be conclusively demonstrated with the data at hand. It is noted by Guan et al. (2009) that the angle associated with the magnetic pressure is consistent with the magnetic tilt angle θ_B and this is approximately the case for the results presented here. While the magnetic tilt angle appears to converge to $\theta_B \approx 13^\circ$, the principal axis of the magnetic pressure is $\theta \approx 15^\circ$ which is closer to the value expected from the estimate $\alpha\beta \approx 1/2$.

Next, we consider in more detail the magnitude of λ_M/H_0 given in Table 6.2. Beckwith et al. (2011) consider only one model for which $\lambda_M = (3, 2, 1)H$ for the gas, magnetic pressure, and kinetic energy respectively and note the similarity of their results to those of Nelson & Gressel (2010). The results here, however, suggest that resolution is likely a concern. The values of λ_M observed by Beckwith et al. (2011) and Nelson & Gressel (2010) are comparable to run `BzZ8` and would likely decrease further with greater resolution although there remains the possibility that stratification in the simulations they consider may act to elongate the planar structure of the turbulence. Guan et al. (2009) consider a series of local models with both zero net vertical field and net toroidal field with resolutions (in zones per scale height) ranging from 32 to 256. The correlation scales associated with the former field topology fail to converge as does the stress. They note, for their non-converging simulations that both λ_M and λ_m decrease in a manner proportional to N_x ($N_x = L_x/\Delta x$). Due to the lack of convergence for this field topology we instead compare our results to their converged net toroidal field simulations. For their lowest resolution simulation, they find $\lambda_M/H \approx 0.63$ and $\lambda_m/H \approx 0.11$ for the magnetic pressure which is almost identical to the values observed for simulation `BpN32`. That these values are comparable is quite encouraging as it suggests that these global simulations are not only able to access the large-scale structure inaccessible to local models but also match local models in the intermediate- and small-scale regimes.

Regarding the resolution dependence of the correlation lengthscales we compare our results to Guan et al. (2009) and their thorough resolution study. Overall, we find broad agreement between their results and the data presented here at intersecting resolutions (Guan et al. 2009, Table 2). Of particular concern is the resolvability of λ_m, which they find to converge (using the 8-zone criterion) at a resolution of 128 zones per scale height. Using a fiducial value of $\lambda_m/H_0 \approx 0.10$ from our global simulations we find that this corresponds to (for a 32 zone per scale height simulation) $\lambda_m/\Delta \approx 3$. Assuming a negligible reduction in λ_m for increasing resolution, this suggests that a resolution of 128 zones per scale height would also resolve the minor axes of the correlation functions for our global simulations.

Finally, our attention turns to the turbulent structure of the Maxwell stress, and as always specifically the component $M_{R\phi}$. The correlation function associated with this will be denoted C_M and a fiducial example is shown in Figure 6.8 with data taken from simulation `BzZ32`. Similar to the energy quantities, save for the gas pressure, the stress is ellipsoidal in nature. The two-point correlation functions

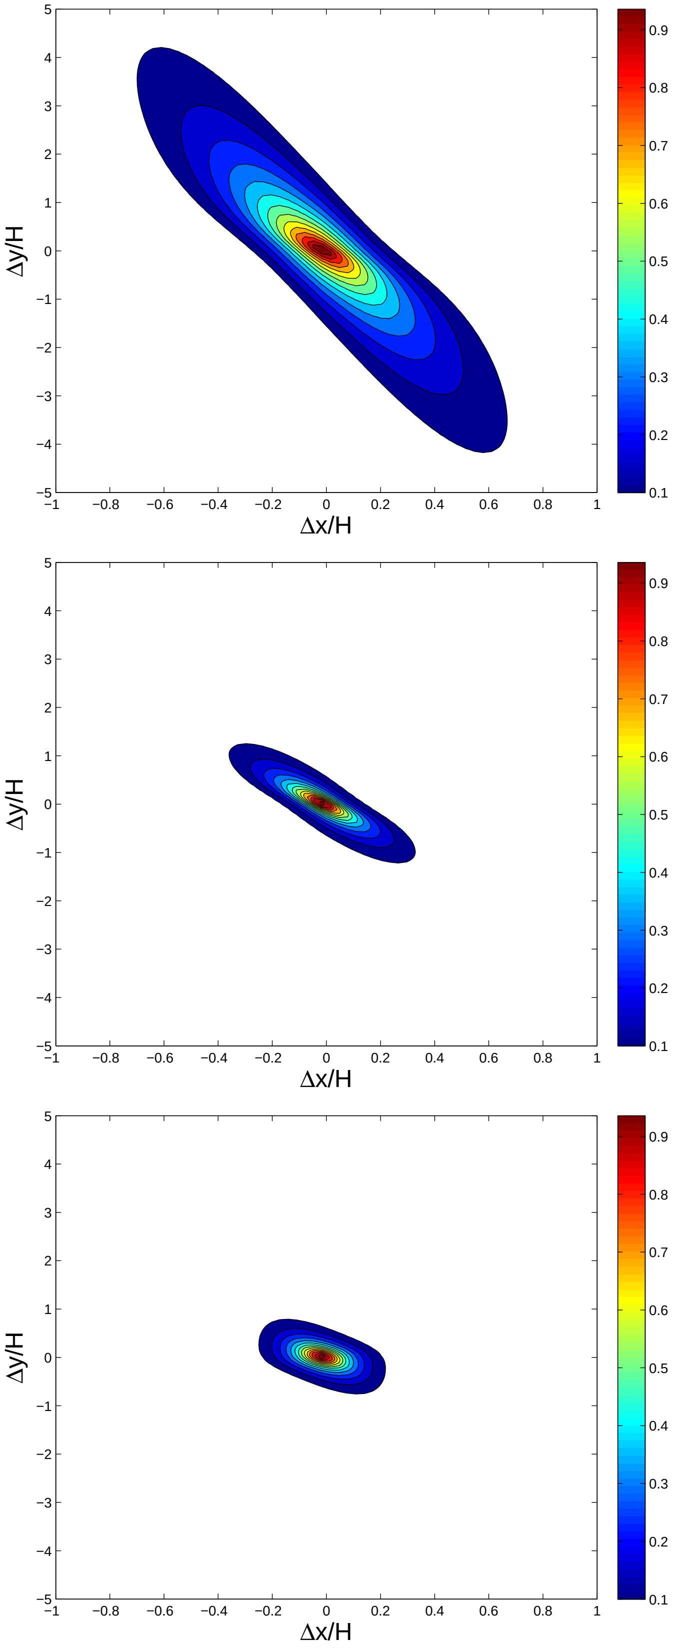

Figure 6.7: Two-point correlation functions of the turbulent energy components for simulation `BzZ32`. (**Top**) Gas pressure (**Middle**) Magnetic pressure (**Bottom**) Perturbation kinetic energy (K')

Simulation	Gas Pressure			Magnetic Pressure			Kinetic Energy (K')		
	λ_M	λ_m	θ	λ_M	λ_m	θ	λ_M	λ_m	θ
BzZ8	3.78	0.35	8.37°	1.74	0.17	8.37°	1.04	0.22	7.66°
BzZ16	1.86	0.20	10.30°	0.86	0.12	12.62°	0.52	0.16	9.20°
BzZ32	1.42	0.14	10.25°	0.51	0.07	14.98°	0.32	0.10	10.10°
BzZ32W	1.03	0.12	11.62°	0.55	0.08	14.90°	0.31	0.10	10.38°
BzZ64W	0.81	0.10	11.31°	0.44	0.07	15.60°	0.23	0.07	10.66°
BzN8	2.06	0.30	10.16°	1.63	0.23	11.39°	0.86	0.26	9.36°
BzN16	1.27	0.21	12.39°	1.09	0.18	13.66°	0.59	0.19	10.56°
BpN8	2.19	0.30	9.27°	1.84	0.24	10.18°	0.88	0.23	9.18°
BpN16	1.37	0.20	11.61°	1.05	0.16	12.82°	0.55	0.17	10.37°
BpN32	0.96	0.14	12.61°	0.69	0.11	14.21°	0.35	0.11	10.75°

Table 6.2: Structure of the correlation functions associated with the partition of turbulent energy for all simulations. Note, the major and minor axes of the correlation function are scaled by H_0.

of the stress for all the isotropic resolution simulations are quantified in Table 6.3. Broadly, the stress structure is similar to the other quantities considered in that as resolution increases both the major and minor axes decrease and the angle of the principal axis increases. Regarding the specific lengthscale associated with the stress, it can be placed in the previously described hierarchy between the magnetic pressure and perturbed kinetic energy with an angle larger than the other quantities considered. It is interesting to note that the calculated angle of the stress correlations is approximately 19°, while this is much larger than the fiducial magnetic tilt angle ($\theta_B \approx 13^\circ$) it is approximately consistent with the peak magnetic tilt angle at saturation (Figure 5.13) for the models initialized with vertical magnetic fields. This may suggest that the saturation of the MRI imposes a specific anisotropy on the spectral structure of the magnetic stress, however connecting the principal axis of the spectral structure to the planar anisotropy of the magnetic field is murky and this analysis is well beyond the scope of this work.

Finally, we return to the point regarding the disparate character of the gas pressure ACF compared to the other quantities considered. As a first step, the variation of the principal axis with displacement is quantified in Figure 6.9. Using a series of contours of the ACF, spaced between 0.2 and 0.8, the angle of the ellipse fit to each contour is plotted. From Figure 6.9 we note that with the exception of the gas pressure all the quantities considered exhibit minimal variation of principal axis with displacement. In contrast to this, the gas pressure at large displacements exhibits a similar principal axis to the kinetic energy whereas at small displacement

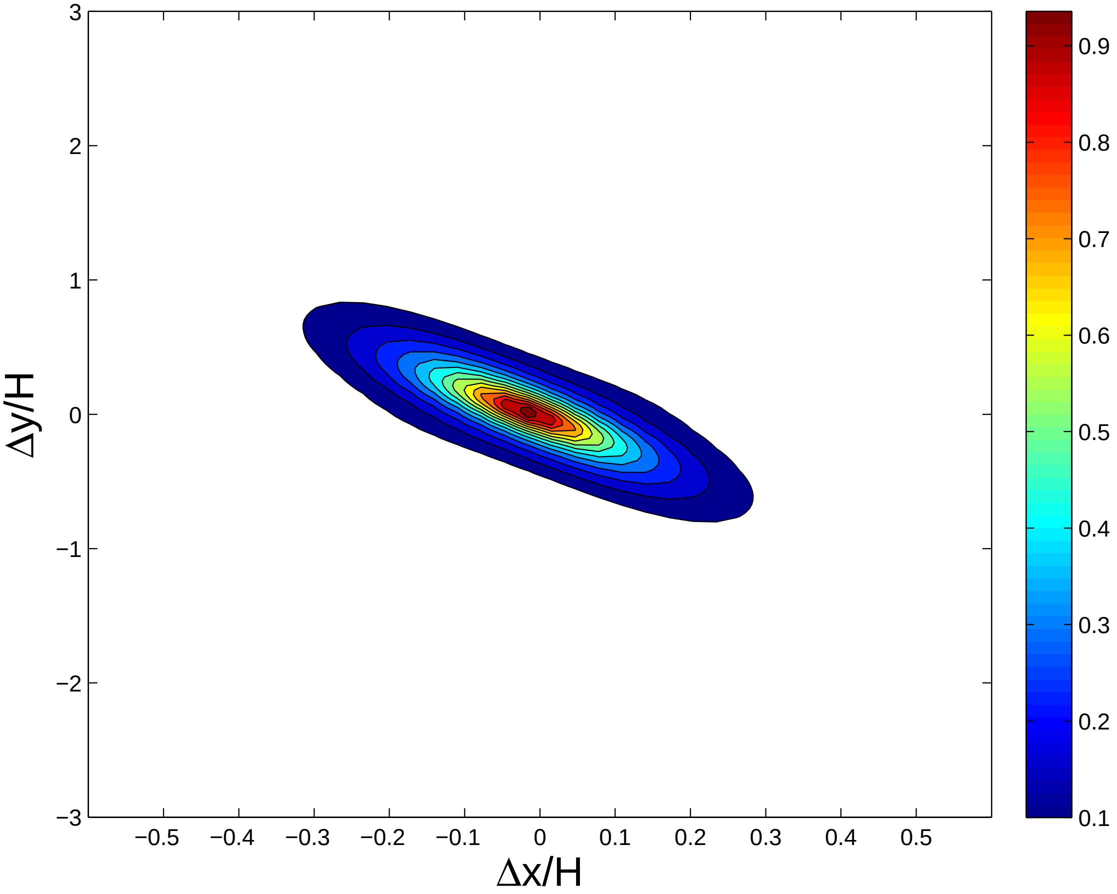

Figure 6.8: Two-point correlation function of the Maxwell stress ($M_{R\phi}$), denoted C_M for simulation `BzZ32`.

Simulation	Maxwell stress ($M_{R\phi}$)		
	λ_M	λ_m	θ
BzZ8	1.00	0.18	11.83°
BzZ16	0.65	0.12	16.02°
BzZ32	0.40	0.08	18.85°
BzZ32W	0.42	0.08	19.02°
BzZ64W	0.32	0.07	20.32°
BzN8	1.12	0.26	15.09°
BzN16	0.76	0.19	18.28°
BpN8	1.10	0.25	14.41°
BpN16	0.71	0.16	17.49°
BpN32	0.48	0.11	19.42°

Table 6.3: Two-point correlation functions of the Maxwell stress for all simulations.

the ACF has a more tilted principal axis. This behavior is well-documented, showing up in previous studies of correlation functions of MRI-driven turbulence except in the work of Guan et al. (2009). Beckwith et al. (2011) attribute these features to the excitation of spiral density waves as described by Heinemann & Papaloizou (2009) and further note that this excitation of spiral density waves is predicted to be suppressed in simulations with an azimuthal domain of less than $6H$, like those considered by Guan et al. (2009).

6.5 Discussion and Conclusions

The use of shearing box models, while necessary to explore turbulence at small scales, is predicated on a series of significant assumptions. Testing that the predictions of local simulations hold in the context of global models is an important validation of the local model. Towards this end, we have demonstrated that saturation predictors from the local regime correspond to instantaneous correlations between magnetic flux and stress in the global regime. This, combined with an understanding of flux distributions in global simulations would allow simpler statistical treatments of disk turbulence. However, the conservation of magnetic flux on small scales implicit in local models is not realized in global disks. Indeed, local statistics derived from global simulations seeded with a magnetic field topology possessing net flux are characterized by significant migrations in the space $Q_z \times Q_\phi$ in contrast to the assumption of conserved flux implicit in the shearing box formalism. Additionally, these migrations appear to be consistent with the magnetic flux evolving

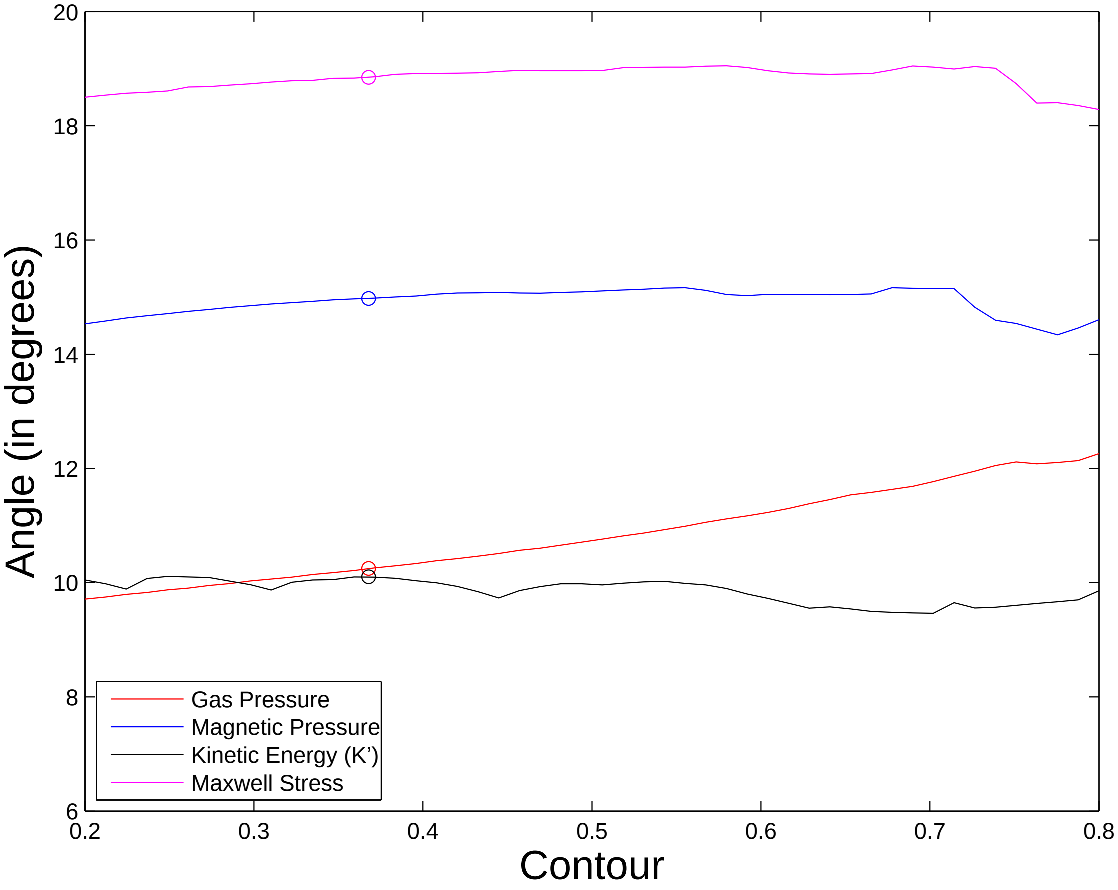

Figure 6.9: The dependence of the principal axis of the ACF on displacement. Plotted is the calculated tilt angle as a function of fiducial contour. Circles mark the angle associated with the e^{-1} contour.

to values independent of initial field topology. Simulations with larger temporal domains will need to be done to verify this point.

The intermediate- and small-scale structure of the MRI-driven turbulence is probed through the use of two-point correlation functions. It is found that the partition of turbulent energy between the gas and magnetic pressures and the kinetic energy in the rotating frame exhibits a hierarchy of associated scales and tilted principal axes. The observed hierarchy is in agreement with previous results exploring both local unstratified models and global stratified disks. Further, the intermediate- (λ_M) and small- (λ_m) scale structure agree quite well with results obtained from local models. Resolvability of the smallest scales, the minor axes of the correlation functions, can be estimated to occur for simulations using an isotropic resolution of 128 zones per scale height. Given the cost of simulation `BzZ64W`, using a reduced temporal and azimuthal simulation domain, at approximately 150×10^3 CPU-hours this would suggest a cost of approximately 2.5M CPU-hours to run `BzZ128W` (almost 40M CPU-hours without orbital advection). While this expense is fairly high, it is not outside the bounds of the current generation of supercomputers. This suggests that the resolvability of the minor axes, suggested by Guan et al. (2009) as a precondition for the formation of an inertial range, is within reach.

Chapter 7

Conclusions and Future Work

While the majority of the issues discussed here are primarily technical in nature, the connection to astrophysical systems should not be forgotten. The use of the magnetic tilt angle as a convergence metric has the potential to greatly simplify the often difficult task of running simulations from which astrophysically-relevant quantities can be confidently measured. That the magnetic tilt angle is independent of initial field topology, in the unstratified simulations presented, suggests an important point of unification between the often disparate phenomenologies observed for MRI-driven turbulence induced by various field topologies. However, it is unlikely that MRI-driven turbulence can be characterized by a simple scalar in the presence of more complex physics. The smaller volume-integrated value of θ_B obtained from stratified simulations (*e.g.* Beckwith et al. 2011) suggests that the magnetic anisotropy is suppressed in the corona. Additionally, recent numerical evidence (John Hawley, private communication) indicates that the use of a pseudo-Newtonian potential results in a magnetic tilt angle that exhibits radial dependence. The existence of a connection between the fiducial magnetic tilt angle and the saturation of the MRI by parasitic modes (Pessah, 2010) would further imply that the tilt angle depends on dissipative coefficients. The development of a new generation of simulations will have to be undertaken that can simultaneously include these more complicated physics at resolutions comparable to those used in local models to study the converged behavior of the spatial and dissipative dependence of θ_B. The difficulty inherent in these types of simulations will be returned to below.

The use of the shearing box model has been tremendously influential and important in advancing the goal of studying MRI-driven turbulence in differentially-rotating systems through the use of computational simulations. The use of this approximation relies on the assumption that a shearing box accurately models a small co-rotating region of a global accretion disk. The results present here suggest that the veracity of this assumption is subtle. The small-scale structure of MRI-driven turbulence more efficiently modeled using the shearing box formalism is found, through the study of two-point correlation functions, to be accurately reproduced in high-resolution global disks. We also find that the presence of magnetic flux stimulates stress in a manner analogous to that understood from shearing box simulations. Quantifying this relationship between magnetic flux and stress results in a lengthscale, ℓ_z and ℓ_ϕ, that can be interpreted as measuring the "locality" of global turbulence. These lengthscales are found to be remarkably similar to the computational domain sizes that are often used in local studies of accretion disk turbulence. While this is encouraging, it is precisely this relationship between magnetic flux and stress that has the potential to cause difficulties. Magnetic flux is conserved on small scales in local models, whereas consideration of the local statistics derived from global models suggests that the magnetic flux threading a small subdomain is quite volatile and displays considerable temporal evolution. This evolution, while volatile, is in general towards smaller values of magnetic flux and thus, due to the flux-stress relationship, towards smaller values of stress. Because the accurate modeling of this evolution is impossible in local models, one would expect them to exhibit an artificially-enhanced amount of stress. It should be noted, however, that the addition of stratification could potentially affect this last point as the use of open boundary conditions along with magnetic buoyancy would allow the expulsion of magnetic flux from the computational domain. Local models will undoubtedly remain important to facilitate the study of MRI-driven turbulence and its small-scale properties, but care must be taken when making astrophysical inferences from these models.

As computational power increases and the sophistication of algorithms grows global models will likely come to largely displace local models. The use of orbital advection for the first time in simulations of global disks represents an important step towards that. Its use has allowed the simulation of global disks at comparable resolutions to local models and demonstrated that global simulations can accurately model small-scale features of MRI-driven turbulence while capturing large-scale features and temporal evolution inaccessible to local models. In addition to the simple performance enhancement orbital advection provides there is the additional benefit that orbital advection allows the use of high azimuthal resolution without the significant cost that azimuthal resolution normally entails. It has often been argued that the strong shear of accretion disks acts to "smear" azimuthal features over a larger lengthscale and thus coarse azimuthal resolution is adequate. The evidence presented here contradicts this belief and demonstrates that an azimuthal resolution nearly identical to the vertical resolution is required to adequately capture the details of MRI-driven turbulence.

The use of orbital advection provides numerous benefits with almost no associated costs and as such will likely become an important feature of the landscape of global accretion disk simulations. However, there are complications involved in using orbital advection in the context of more astrophysically-relevant accretion disk simulations, specifically the inclusion of stratification or the use of a pseudo-Newtonian gravitational potential. The former causes the development of a strongly magnetized corona that exhibits supersonic Alfven speeds whereas the latter includes a "plunging" region within the ISCO in which matter falls into the central potential at supersonic speeds. Because of these regions, the performance benefit of orbital advection is limited, although preliminary evidence suggests that there remains a performance boost of a factor of approximately four. The `Athena` algorithm is currently limited to linear grids, in which the grid spacing is uniform throughout the computational domain. Current work is near completion to extend the Cartesian geometry to allow non-uniform grids, and the next phase of this work will be to adapt this to cylindrical geometry through the extension of the prolongation and restriction operators (Tóth & Roe, 2002) required to reconstruct magnetic fields while preserving the solenoidal constraint. Using coarser grids within the plunge region or the corona will allow a more significant speed-up when combined with orbital advection while still allowing appropriate resolution in the body of the disk. The sophisticated algorithm of `Athena` coupled to the flexibility of a non-uniform grid and the efficiency of orbital advection combined with next-generation supercomputers will provide an unprecedented platform through which to study the details of

MRI-driven turbulence at all spatial and temporal scales simultaneously. Leveraging this platform and the lessons learned here regarding proper resolution to simulate more astrophysical accretion disks will be the focus of the next phase of this work.

Bibliography

Armitage, P. J. 1998, Astrophysical Journal Letters v.501, 501, L189

Balbus, S. A., & Hawley, J. F. 1991, Astrophysical Journal, 376, 214

—. 1992a, Astrophysical Journal v.400, 400, 610

—. 1992b, Astrophysical Journal, 392, 662

—. 1998, Reviews of Modern Physics, 70, 1

Beckwith, K., Armitage, P. J., & Simon, J. B. 2011, arXiv.org, 1105, 1789, 24 pages and 25 figures. MNRAS in press. Version with high resolution figures available from http://jila.colorado.edu/ krb3u/Thin_Disk/thin_disk_turbulence.pdf

Blackman, E. G., Penna, R. F., & Varnière, P. 2008, New Astronomy, 13, 244

Bodo, G., Mignone, A., Cattaneo, F., Rossi, P., & Ferrari, A. 2008, Astronomy and Astrophysics, 487, 1

Bondi, H. 1952, Monthly Notices of the Royal Astronomical Society, 112, 195

Colella, P. 1990, Journal of Computational Physics (ISSN 0021-9991), 87, 171

Colella, P., & Sekora, M. D. 2008, Journal of Computational Physics, 227, 7069

Colella, P., & Woodward, P. R. 1984, Journal of Computational Physics (ISSN 0021-9991), 54, 174

Davis, S. W., Stone, J. M., & Pessah, M. E. 2010, The Astrophysical Journal, 713, 52

Evans, C. R., & Hawley, J. F. 1988, Astrophysical Journal, 332, 659

Frank, J., King, A., & Raine, D. J. 2002, Accretion Power in Astrophysics, iSBN: 0521620538

Fromang, S. 2010, Astronomy and Astrophysics, 514, L5

Fromang, S., & Nelson, R. P. 2006, Astronomy and Astrophysics, 457, 343

Fromang, S., & Papaloizou, J. 2007, Astronomy and Astrophysics, 476, 1113

Fromang, S., & Stone, J. M. 2009, Astronomy and Astrophysics, 507, 19

Gammie, C. F., & Menou, K. 1998, Astrophysical Journal Letters v.492, 492, L75

Gardiner, T. A., & Stone, J. M. 2005, Journal of Computational Physics, 205, 509

—. 2008, Journal of Computational Physics, 227, 4123, elsevier Inc.

Goodman, J., & Xu, G. 1994, Astrophysical Journal, 432, 213

Guan, X., Gammie, C. F., Simon, J. B., & Johnson, B. M. 2009, Astrophysical Journal, 694, 1010

Hawley, J. F. 2001, The Astrophysical Journal, 554, 534

Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1995, Astrophysical Journal v.440, 440, 742

—. 1996, Astrophysical Journal v.464, 464, 690

Hawley, J. F., Guan, X., & Krolik, J. H. 2011, arXiv.org, 1103, 5987, submitted to the Astrophysical Journal

Hawley, J. F., & Stone, J. M. 1995, Computer Physics Communications, 89, 127

Hayes, J. C., Norman, M. L., Fiedler, R. A., Bordner, J. O., Li, P. S., Clark, S. E., ud Doula, A., & Mac Low, M.-M. 2006, The Astrophysical Journal Supplement Series, 165, 188

Heinemann, T., & Papaloizou, J. C. B. 2009, Monthly Notices of the Royal Astronomical Society, 397, 52

Johansen, A., Youdin, A., & Klahr, H. 2009, The Astrophysical Journal, 697, 1269

Johnson, B. M., Guan, X., & Gammie, C. F. 2008, The Astrophysical Journal Supplement Series, 177, 373

King, A. R., Pringle, J. E., & Livio, M. 2007, Monthly Notices of the Royal Astronomical Society, 376, 1740

Lesur, G., & Longaretti, P.-Y. 2007, Monthly Notices of the Royal Astronomical Society, 378, 1471

—. 2009, Astronomy and Astrophysics, 504, 309

Masset, F. 2000, Astronomy and Astrophysics Supplement, 141, 165

Miller, K. A., & Stone, J. M. 2000, The Astrophysical Journal, 534, 398

Miyoshi, T., & Kusano, K. 2005, Journal of Computational Physics, 208, 315

Nelson, R. P., & Gressel, O. 2010, Monthly Notices of the Royal Astronomical Society, 1392, (c) Journal compilation © 2010 RAS

Noble, S. C., Krolik, J. H., & Hawley, J. F. 2010, The Astrophysical Journal, 711, 959

O'Neill, S. M., Reynolds, C. S., Miller, M. C., & Sorathia, K. A. 2010, arXiv.org, astro-ph.GA, submitted to The Astrophysical Journal, 6 pages, 5 figures

Paczynsky, B., & Wiita, P. J. 1980, Astronomy and Astrophysics, 88, 23

Pessah, M. E. 2010, The Astrophysical Journal, 716, 1012

Pessah, M. E., Chan, C.-k., & Psaltis, D. 2007, The Astrophysical Journal, 668, L51

Pringle, J. E. 1981, In: Annual review of astronomy and astrophysics. Volume 19. (A82-11551 02-90) Palo Alto, 19, 137, a&AA ID. AAA030.064.057

Regev, O., & Umurhan, O. M. 2008, Astronomy and Astrophysics, 481, 21

Reynolds, C. S., & Miller, M. C. 2009, The Astrophysical Journal, 692, 869

Shakura, N. I., & Sunyaev, R. A. 1973, Astron. Astrophys., 24, 337, a&AA ID. AAA009.066.049

Simon, J. B., Hawley, J. F., & Beckwith, K. 2009, The Astrophysical Journal, 690, 974

—. 2011, The Astrophysical Journal, 730, 94

Skinner, M. A., & Ostriker, E. C. 2010, The Astrophysical Journal Supplement, 188, 290

Sorathia, K. A., Reynolds, C. S., & Armitage, P. J. 2010, Astrophysical Journal, 712, 1241

Sorathia, K. A., Reynolds, C. S., Stone, J. M., & Beckwith, K. 2011, arXiv

Spruit, H. C., & Uzdensky, D. A. 2005, The Astrophysical Journal, 629, 960

Stone, J. M. 2009, in Proceedings, 277–286

Stone, J. M., & Gardiner, T. A. 2010, The Astrophysical Journal Supplement, 189, 142

Stone, J. M., Gardiner, T. A., Teuben, P., Hawley, J. F., & Simon, J. B. 2008, The Astrophysical Journal Supplement Series, 178, 137

Stone, J. M., Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1996, Astrophysical Journal v.463, 463, 656

Stone, J. M., & Norman, M. L. 1992, The Astrophysical Journal Supplement Series, 80, 753

Tóth, G., & Roe, P. L. 2002, Journal of Computational Physics, 180, 736

van Ballegooijen, A. A. 1989, in IN: Accretion disks and magnetic fields in astrophysics; Proceedings of the European Physical Society Study Conference, Harvard-Smithsonian Center for Astrophysics, Cambridge, MA, Harvard-Smithsonian Center for Astrophysics, Cambridge, MA, 99–106