

2014

A hybrid advection scheme for conserving angular momentum on a refined Cartesian mesh

Zachary Duncan Byerly

Louisiana State University and Agricultural and Mechanical College, zbyerly@gmail.com

Follow this and additional works at: https://digitalcommons.lsu.edu/gradschool_dissertations



Part of the [Physical Sciences and Mathematics Commons](#)

Recommended Citation

Byerly, Zachary Duncan, "A hybrid advection scheme for conserving angular momentum on a refined Cartesian mesh" (2014). *LSU Doctoral Dissertations*. 1548.

https://digitalcommons.lsu.edu/gradschool_dissertations/1548

This Dissertation is brought to you for free and open access by the Graduate School at LSU Digital Commons. It has been accepted for inclusion in LSU Doctoral Dissertations by an authorized graduate school editor of LSU Digital Commons. For more information, please contact gradetd@lsu.edu.

A HYBRID ADVECTION SCHEME FOR CONSERVING ANGULAR MOMENTUM ON A REFINED CARTESIAN MESH

Submitted to the Graduate Faculty of the
Louisiana State University and
Agricultural and Mechanical College
in partial fulfillment of the
requirements for the degree of
Doctor of Philosophy

in

The Department of Physics and Astronomy

by

Zachary Duncan Byerly

B.S., Louisiana State University and Agricultural and Mechanical College, 2007

M.S., Louisiana State University and Agricultural and Mechanical College, 2011

May 2014

This dissertation is dedicated to my friends and family, especially my parents Maud and Gary.

Acknowledgements

I would like to thank my major professor, Joel Tohline, for his support and guidance. It has been a great privilege to be his graduate student. This work would not have been possible without the advice – both professional and personal – I received from Juhan Frank. I remember fondly my first conversation with Juhan, back in the fall of 2001, before I had even graduated from high school. Robert Hynes has helped me in numerous ways over the last 8 years, including guiding me to complete my undergraduate research project studying the long term X-ray variability of the neutron star binary system EXO 0748-676.

I would like to thank Hwang Lee and William Adkins for joining Joel, Juhan, and Robert on my dissertation committee. Many other professors in the Department of Physics & Astronomy, especially Geoff Clayton and Brad Schaefer, have taught me well over the last 10 years. Hartmut Kaiser of the Center for Computation & Technology has been both a mentor and a friend since I began working with him in 2012.

During my time at LSU, my friends and colleagues Dominic Marcello, Wes Even, Ashley Pagnotta, Bryce Adelstein-Lelbach, Patrick Motl, Andrew Collazzi, Sarah Caudill, Jan Staff, Jay Call, Kundan Kadam, and Vayujeet Gokhale provided me with invaluable advice and support. Our graduate student secretary, Arnell Nelson, deserves many thanks for her hard work and patience. My high school science teachers, David Hopkins and Yasuo Namba, helped me realize my interest in science. I would especially like to thank my family – my mom, Maud Walsh, my dad, Gary Byerly, my brother, Ben, and my sister-in-law, Eli.

This work has been supported, in part, by grants ACI-1246443 and AST-1240655 from the U.S. National Science Foundation, in part, by NASA ATP grant NNX10AC72G, and, in part, by U.S. Department of Energy grant DE-SC0008714. This research also has been made possible by grants of high-performance computing time at HPC@LSU through the allocation `hpc_dwd_amr`, and across LONI (Louisiana Optical Network Initiative), especially award `loni_lsuaastro11`.

Table of Contents

Acknowledgements	iii
Abstract	v
1 Introduction	1
1.1 Context	1
1.2 Overview of this Work	4
2 Methods	6
2.1 Initial Models	6
2.2 Hydrodynamics Code	7
2.3 Description of Hybrid Angular Momentum Conservation Scheme	9
2.4 Postprocessing	12
2.5 Quantifying Results	13
3 Results	18
3.1 Uniform Specific Angular Momentum Tori	18
3.2 Near-Keplerian Disk Simulation	31
4 Conclusion	37
4.1 This Work	37
4.2 Future Goals	38
References	39
Appendix: Momentum Equation Formulation	42
A.1 Statements of Conservation	42
A.2 Rotating Frame of Reference	43
A.3 Eulerian Representation	43
A.4 Angular Momentum Conservation	44
A.5 Summary of Equations	45
A.5.1 Non-rotating frame	46
A.5.2 Rotating Frame, Rotating Frame Quantities	46
A.5.3 Rotating Frame, Inertial Frame Quantities	47
Vita	48

Abstract

We demonstrate the capabilities of a new hybrid scheme for simulating dynamical fluid flows in which cylindrical components of the momentum are advected across a rotating Cartesian coordinate mesh. This hybrid scheme allows us to conserve angular momentum to machine precision while capitalizing on the advantages offered by a Cartesian mesh, such as mesh refinement. The work presented here focuses on measuring the real and imaginary parts of the eigenfrequency of unstable axisymmetric modes that naturally arise in massless polytropic tori having a range of different aspect ratios, and quantifying the uncertainty in these measurements. Our measured eigenfrequencies show good agreement with the results obtained from the linear stability analysis of Kojima (1986) and from nonlinear hydrodynamic simulations performed on a cylindrical coordinate mesh by Woodward et al. (1994). When compared against results conducted with a traditional Cartesian advection scheme, the hybrid scheme achieves qualitative convergence at the same or, in some cases, much lower grid resolutions and conserves angular momentum to a much higher degree of precision. As a result, this hybrid scheme is much better suited for simulating astrophysical fluid flows, such as accretion disks and mass-transferring binary systems.

1 Introduction

1.1 Context

Binary star systems, especially those containing compact components, are of great current interest in astrophysics. Most binary systems have a separation much larger than the radii of either star, and thus only interact gravitationally as point masses throughout most of their lifetimes. While these systems are more interesting than single stars, they pale in comparison to the richness of information that can be gleaned from studying stars which, by some process, have come into contact (Frank, King, & Raine 2002).

In order to understand the interactions of close binaries, an understanding of an idealized potential field known as the Roche potential is crucial. When the orbital separation of a binary system becomes comparable to the radius of one of the components, it becomes useful to approximate the gravitational potential field as one given by two point masses in a circular orbit. Any matter co-rotating with the binary will feel this potential field, plus a Coriolis force. This is a good approximation for most stars, as they are usually centrally condensed.

The Roche lobes are equipotential surfaces inside of which matter co-rotating with the binary system is bound to the star it surrounds. The Roche lobes of the two binary companions meet at the first Lagrange point, L1. It is a saddle point in the potential surface. When a star fills its Roche lobe, matter will flow over the L1 point and toward the companion star. This process is known as Roche Lobe Overflow (RLOF) and it is the primary way mass can transfer between two binary companions.

Binaries can be classified in many different ways, but when studying the dynamics of close binary systems a classification system based on how the companion stars fill their Roche lobes is typically most useful. We can break these down into three categories. In a detached binary system, neither component fills its Roche lobe, and no RLOF mass transfer occurs. In a semi-detached binary system, one of the two components (known as the donor) fills its Roche lobe, and matter from the donor flows over the L1 point onto the other component (known as the accretor). If both stars fill their Roche lobes, a contact binary is formed.

Consider a binary system that starts as a pair of main-sequence stars. The evolutionary path of this system depends on several factors, including the total mass of the system, the ratio of the masses of the two component stars, and the distance between the two stars. Yungelson et al. (2004) review many possible evolutionary paths of binary systems. For example, when one of the components leaves the main sequence and expands beyond its Roche lobe, a common envelope is formed. The two cores of the stars orbit around each other inside this common envelope, and angular momentum is transferred from the orbit to the common envelope, reducing the separation of the binary. A second common envelope phase occurs when the second star evolves off of the main-sequence. The resulting binary has a much smaller separation than the original binary. At this smaller radius, angular momentum loss due to the general-relativistic effect of gravitational radiation becomes important. Eventually this process will drive the two degenerate binary companions together, causing

them to eventually begin mass transfer. Double white dwarf (DWD) binaries, in particular, are estimated to number $\sim 2.5 \times 10^8$ in our Galaxy (Nelemans et al. 2001).

While the Phillips relation (Phillips 1993) allows astronomers to use type Ia supernovae (SNeIa) as standard candles (which led to the Nobel prize winning discovery of the acceleration of the universe by Riess et al. (1998) and Perlmutter et al. (1999)), the exact mechanism of the explosion, and the nature of the progenitors, continue to be investigated. The consensus is that SNeIa occur when a carbon-oxygen (CO) white dwarf undergoes a runaway fusion reaction (Hillebrandt & Niemeyer 2000), as a result of either accretion from the companion star (single-degenerate model, Whelan & Iben 1973), or the merger of a DWD binary (double-degenerate model, Webbink 1984). This debate over which channel dominates continues (Geier et al. 2007; Rosswog et al. 2009; Fryer et al. 2010). Schaefer & Pagnotta (2012) in particular point to the lack of an ex-companion object in the vicinity of a 400 year old supernova remnant, which they argue rules out all published single-degenerate models. Efforts to reduce the intrinsic Hubble scatter (deviations in the magnitude-redshift relationship) have not been very successful. Better theoretical models of SNeIa are needed to understand the nature of this intrinsic scatter, which will ultimately help to reduce measured uncertainties in cosmological parameters (Diemer et al. 2013).

DWD binary mergers are also thought to produce R Coronae Borealis (RCB) stars (Webbink 1984), although again the exact mechanism is not well understood. Extremely hydrogen deficient with an overabundance of carbon and nitrogen, these enigmatic objects have attracted the attention of astronomers since their discovery over 200 years ago (Clayton 1996). Efforts by our group (Staff et al. 2012) are underway to reproduce with hydrodynamic simulations the unusual abundances and isotopic ratios of these stars.

Our desire is to employ computational fluid techniques to model, in a self-consistent manner and to a high degree of accuracy, mass-transfer in a wide variety of interacting binary star systems over hundreds, if not thousands, of orbits. Such a capability would allow us not only to better understand the behavior of binaries that are dynamically unstable toward merger or tidal disruption of the donor, but also to examine how spin-orbit coupling – for example, the exchange of angular momentum between the donor star and a disk surrounding the accretor – facilitates dynamical stability and leads to long phases of quasi-steady mass transfer. With such a tool we could simulate how slow accretion can bring an initially sub-Chandrasekhar-mass accretor to the brink of critical collapse; how a transition from sub- to super-Eddington accretion rates affects common-envelope development and evolution; and the steady-state structure of mass-transferring AM CVn type binaries.

Examples of our group’s efforts, to date, include the code introduced in Motl et al. (2002), which conserves angular momentum, $\Delta J_z/J_z$, to $\approx 10^{-4}$ per orbit, and successfully evolved DWD binaries through over 5 orbits. Modifications to that code presented in D’Souza et al. (2006) include a corrected handling of the pressure gradients that accounts for gradients in the fluid’s specific entropy, which slightly improved the angular momentum conservation to $\Delta J_z/J_z \approx 5 \times 10^{-5}$ per orbit, allowing the code to evolve a binary for over 30 orbits. This research has been continued by Motl et al. (2007), Even & Tohline (2009), and Marcello & Tohline (2012). The work presented in this dissertation builds on this group experience and

offers a new numerical technique for simulating astrophysical fluid flows at even higher levels of precision.

In simulating interacting binary systems numerically, a faithful representation of the flow can be achieved only if the grid resolution is sufficiently high across a range of dynamically interesting flow regions. These regions can vary in structure as well as in identity over time, so the grid needs to adapt accordingly. For example, even when only considering double degenerate binaries, the smallest length scales (surface layers of both stars, scale height of the disk, and fluid flow through the L1 Lagrange point) can be tiny compared to the binary separation. At the other extreme, an envelope consisting of an optically thick atmosphere engulfing both stars can develop shortly after accretion begins, quickly filling the original computational domain. This “common envelope” structure may expand to a size many times larger than the binary separation. Adaptive mesh refinement (AMR) techniques can be called upon to provide an appropriately high degree of spatial resolution in various, as well as in time-varying, regions of the flow. Astrophysical simulation codes that employ AMR include FLASH (Fryxell et al. 2000), ZEUS (Hayes et al. 2006), PLUTO (Mignone et al. 2007), and Scorpio, recently developed by Marcello et al. (2014, in preparation). AMR techniques can be straightforwardly implemented on a Cartesian mesh but are quite difficult to employ across curvilinear grids.

On the other hand, advection schemes implemented on Cartesian meshes are most naturally designed to conserve linear momentum, rather than angular momentum. Because binary evolutions can be faithfully followed through hundreds of orbits only if the simulation conserves angular momentum to a high degree of accuracy, in the past we have chosen to use a cylindrical computational grid, which more naturally facilitates conservation of orbital angular momentum. However, even a cylindrical grid does not match the symmetry of each individual binary component. And in the absence of a mesh refinement capability we have not had the full freedom to distribute resolution where it is needed.

Alternative curvilinear meshes designed to address some of these shortcomings have been used in previous disk and torus simulations. For example, Zink et al. (2008) discuss a multi-patch technique meant for spherically symmetric or axisymmetric simulations, and Fragile et al. (2009) introduce a “patched-sphere” mesh similar to the multi-patch technique. But, as with a cylindrical coordinate grid, such schemes do not easily accommodate mesh refinement techniques. Ultimately, Zink et al. (2008) say that Cartesian mesh-refinement is likely better suited for problems like binary mergers. Historically, therefore, it has been difficult to achieve both high – and adaptive – spatial resolution while at the same time achieving a high degree of angular momentum conservation. The hybrid advection scheme designed by Call et al. (2010) and implemented here, allows us to have our cake and eat it, too. It facilitates conservation of angular momentum to machine accuracy on a refined Cartesian mesh. Other Cartesian AMR codes have not demonstrated an ability to conserve angular momentum to the degree we believe necessary for simulating binary orbits. Hawley et al. (2012) perform zero impact parameter white dwarf collisions in FLASH. The lack of published simulations of binary orbits despite the recent evidence of DWD SNeIa progenitors

may indicate that existing Cartesian AMR astrophysics codes lack the ability to perform such simulations.

1.2 Overview of this Work

In this dissertation we demonstrate the utility of the hybrid scheme by focusing on a quantitative analysis of nonaxisymmetric, dynamical instabilities that arise spontaneously in Papaloizou-Pringle tori (Papaloizou & Pringle 1984), hereafter referred to as PP tori. Each PP torus is a non-self-gravitating, differentially rotating, geometrically thick, axisymmetric disk in orbit about a central point mass. Its internal structure is defined by a balance between gas pressure gradients and gradients in the effective potential. The vertical thickness of the disk/torus relative to its radial extent is determined by the choice of the polytropic index for the gas and an initial angular momentum distribution. (See §2 for details.) These configurations are suitable for demonstrating the capabilities of our hybrid scheme because: They each have a simple analytically definable initial state; while each initial state is axisymmetric, the system is unstable to the development of nonaxisymmetric structure, hence, its evolution has a fully three-dimensional character; the eigenvector of the most unstable mode for each chosen initial configuration, while not known analytically, should be well defined and its measured properties – for example, its complex eigenfrequency – should be reproducible and independent of the specific numerical scheme that is used to perform the dynamical simulation. At the same time, the PP torus provides a good test for hydro codes such as ours because of the challenges it provides. A Cartesian mesh is not ideally suited for the initially axisymmetric torus problem, and thus gives the hybrid scheme an opportunity to prove its worth, for example, by partially overcoming the spurious $m = 4$ modes that are excited by the structure of an underlying Cartesian grid.

Each of our simulations is carried out on a rotating and refined Cartesian grid. The hydrodynamic code we are using (see §§2.2-2.3) has AMR capabilities, but for simplicity we have chosen not to activate the AMR feature. Instead, for each simulation the volume of the grid that is occupied by the initial torus is resolved using a time-invariant, fixed level of refinement (LOR). The effect of grid resolution is assessed by repeating individual simulations several times, using a different (fixed) number of refinement levels. (Typically, we employ 4, 5, or 6 LOR.) In addition, for each initial state and for each specified LOR, the dynamical evolution is carried out using two separate advection schemes: (1) A traditional “Cartesian” scheme in which the x , y , and z components of the linear momentum are advected across the refined Cartesian grid; and (2) our new “hybrid” scheme in which radial momentum and angular momentum – instead of the x and y components of the momentum – are advected across the refined Cartesian grid. In total, results from 23 simulations are reported here.

For each simulation, the real and imaginary components of the complex eigenfrequency of the fastest growing unstable mode are measured and, as appropriate, compared with previously published results from the nonlinear hydrodynamic simulations performed on a

uniformly zoned cylindrical mesh by Woodward, Tohline, & Hachisu (1994) and from the linear stability analysis presented by Kojima (1986).

In an effort to eliminate any subjective bias that might be introduced into the measurement of these eigenfrequencies and, at the same time, to facilitate future efforts to reproduce our results, we introduce a mathematically prescriptive method for quantifying both the value of and uncertainty in each eigenfrequency measurement. In doing this we are able to meaningfully assess the performance of our new hybrid advection scheme, relative to the performance of the traditional Cartesian advection scheme. We show that qualitative convergence is achieved with the hybrid scheme at the same, or sometimes at significantly lower, grid resolutions.

At the same time, we show that the hybrid scheme allows conservation of the system's total angular momentum to machine accuracy. As explained above, this is a highly desirable feature that is not possible to achieve using a familiar Cartesian advection scheme. This is perhaps the most significant attribute of our hybrid scheme. Historically, the expectation has been that angular momentum conservation can be achieved when modeling an astrophysical disk or binary system only if one adopts a coordinate grid – for example, cylindrical or spherical coordinates – whose underlying basis vectors accommodate the curvilinear features of the flow. Our hybrid scheme facilitates conservation of angular momentum on a Cartesian grid.

The hydrocode that has been used to carry out the primary set of simulations reported in this dissertation employs OpenMP to enable multiple execution threads within a single, multi-core compute node. All of the models in our primary set of simulations – totaling 22 in number and using up to 6 LOR – fit within a single node of our high-performance computing system. In §3 of this dissertation we present results from one simulation that was conducted on a rotating Cartesian grid with 7 LOR. This single simulation was run using Octopus, a closely aligned hydrocode built on top of High Performance ParalleX (HPX), a newly emerging parallel runtime system.

2 Methods

2.1 Initial Models

As has already been stated, the initial axisymmetric equilibrium models used in this study were all geometrically thick (uniform specific angular momentum) massless tori with structures as derived by Papaloizou & Pringle (1984). There have been a number of published studies of the nonaxisymmetric stability of these massless tori. Some of these are linear stability analyses (Papaloizou & Pringle 1984; Kojima 1986; Frank & Robertson 1988), while others are nonlinear hydrodynamics simulations, either in two dimensions (Hawley 1987, 1990), or in full three dimensions (Zurek & Benz 1986; Hawley 1990; Woodward et al. 1994). Our initial models have been chosen to provide a direct comparison with the linear stability analysis of Kojima (1986) and the nonlinear hydrodynamic results of Woodward, Tohline, & Hachisu (1994, hereafter WTH).

We generated the same sequence of seven initial models used by WTH, as detailed here in Table 2.1; Figure 2.1 displays the cross-sectional surface of each. The models in this sequence vary in the ratio of the inner radius R_- to the outer radius R_+ .

Following PP, we assume a polytropic equation of state for the fluid,

$$p(\rho) = K\rho^{1+\frac{1}{n}}, \quad (2-1)$$

where, n is the polytropic index and K is the polytropic constant, and we impose a power-law rotation profile given by,

$$\Omega(R) = \Omega_0 \left(\frac{R_0}{R} \right)^q, \quad (2-2)$$

where $\Omega_0 = \Omega(R_0)$ is the angular velocity in the equatorial plane at the radius of pressure maximum, R_0 . Throughout this section we assume $n = 3$ and $q = 2$ (uniform specific angular momentum). As PP have shown, a solution of the hydrostatic balance equation,

$$-\frac{1}{\rho}\nabla p - \nabla\Phi_{\text{eff}} = 0, \quad (2-3)$$

where,

$$\Phi_{\text{eff}} \equiv -\frac{GM_{\text{pt}}}{(R^2 + Z^2)^{1/2}} + \frac{1}{2}\Omega_0^2(R)R^2, \quad (2-4)$$

yields the following axisymmetric equilibrium density distribution:

$$\rho_A(R, Z) = \left(\frac{GM_{\text{pt}}}{4K} \right)^3 \left[\frac{1}{(R^2 + Z^2)^{1/2}} - \frac{R_-}{R^2(1 + R_-/R_+)} - \frac{1}{(R_+ - R_-)} \right]^3, \quad (2-5)$$

where, M_{pt} is the specified mass of the central point mass. In our simulations, we normalize all lengths to the outer edge of the torus, such that $R_+ = 1.0$, and vary R_- . (For illustration purposes, Figure 2.1 is handled differently.) The mass M_{pt} is chosen such that $\rho_{\text{max}} = \rho_A(R_0, 0) = 1.0$.




Fig. 2.1.— Meridional-plane cross sections of the seven different axisymmetric equilibrium disks evolved in this dissertation, listed in Table 2.1. Each disk has a $n = 3$ polytropic index and a uniform specific angular momentum ($q = 2$). Lengths have been normalized to the radius of pressure maximum, R_0 .

Before each model was introduced into the hydrocode, the initial axisymmetric density distribution was perturbed by a single-mode “kick”,

$$\rho(R, Z, \theta) = \rho_A(R, Z)[1 + a \cos(m_{\text{kick}}\theta)],$$

where, m_{kick} is the specific azimuthal mode that we choose to excite, and $a = 10^{-2}$. No velocity perturbation was introduced.

2.2 Hydrodynamics Code

The perturbed models are evolved on a rotating Cartesian grid using Scorpio (Marcello et al. 2014, in preparation) with modifications as detailed here. Following the description provided in that work, the code’s mesh structure is a refined oct-tree of sub-grids as also introduced by MacNeice et al. (2000) in PARAMESH. Each sub-grid is composed of an

$8 \times 8 \times 8$ three-dimensional mesh. Each node contains its own sub-grid and up to eight child sub-grids. Child sub-grids have one-half the grid spacing (twice the spatial resolution) of their parents. The code adheres to “proper nesting,” which requires that there be no more than one jump in refinement level across a sub-grid boundary.

We use the piecewise parabolic (PPM) reconstruction scheme of Colella & Woodward (1984), the semi-discrete central scheme described in Kurganov & Tadmor (2000), and the explicit 3rd order Runge Kutta integrator of Shu & Osher (1988). We enforce a polytropic equation of state that is evolved with the same polytropic index, $n = 3$, that was used to construct the initial models, and a mass density floor of $10^{-10} \rho_{\max}$. Since the PP torus is massless, we evolve the system in a time-independent gravitational potential given by a central object of mass M_{pt} .

This work presents simulations run at different resolutions. Because we use an octree-based grid, it is convenient to speak in terms of the number of levels of refinement of the coarsest $8 \times 8 \times 8$ grid. Each additional level of refinement doubles the resolution in the most refined region. A simulation run at “ N ” LOR has a maximally refined grid spacing $\Delta x = (x_{\max} - x_{\min}) / (8 \cdot 2^N)$, where x_{\max} and x_{\min} are the maximum and minimum extents of the simulation domain, respectively. For the purposes of this study, we use a fixed mesh refinement based on the geometry of the torus, rather than an adaptive mesh. As illustrated in Figure 2.2, the grid is fully refined throughout the volume occupied by the initially axisymmetric torus.

Table 2.2 gives three different numbers for each model at each LOR in an effort to fully reveal the grid properties. For each model, the top row repeats the ratio of the inner radius to the outer radius, R_- / R_+ , shown in Table 2.1; a measure of the slenderness of the torus. The next three rows give three different numbers for 4 LOR: the number (in millions) of fully-refined (smallest grid spacing) zones, the number of fully-refined (leaf) subgrids (each containing $8 \times 8 \times 8$ zones), and the ratio of the radius at pressure maximum to the finest grid spacing, $R_0 / \Delta x$. The next six rows give the same set of numbers for 5 and 6 LOR. The first two numbers are related by a factor of $8^3 = 512$, so Table 2.2 gives two potential

Table 2.1. Torus Geometries

Model	$\frac{R_-}{R_+}$	$\frac{R_0}{R_+}$	$\frac{R_-}{R_0}$
1	0.1	0.1818	0.55
2	0.2	0.3333	0.60
3	0.3	0.4615	0.65
4	0.4	0.5714	0.70
5	0.5	0.6666	0.75
6	0.6	0.7500	0.80
7	0.7	0.8235	0.85




Fig. 2.2.— A comparison of the mesh structure of a $R_-/R_+ = 0.7$ torus simulation taken from three different resolutions. From left to right are 4, 5, and 6 LOR. All simulations run using Cartesian momentum advection. 4 LOR simulation shown at $t = 1.5t_{orb}$, 5 LOR simulation shown at $t = 2.0t_{orb}$, and 6 LOR simulation shown at $t = 2.8t_{orb}$.

metrics for quantifying how well a torus is resolved. The first is simply the total number of fully-refined zones inside of the torus, which decreases with increasing R_-/R_+ ; while the ratio $R_0/\Delta x$ increases with R_-/R_+ . These two metrics seem to contradict each other. The first (number of fully-refined zones) suggests that at any given LOR, a fatter (lower R_-/R_+) torus is better resolved, while the second ($R_0/\Delta x$) would indicate that a skinnier (higher R_-/R_+) torus is better resolved. This illustrates that comparing resolution between different models is not straightforward.

It seems clear that determining whether or not the resolution used to simulate a particular model is sufficient must be done on a model-by-model basis. Hawley et al. (2013) distinguish between true numerical convergence and “qualitative” convergence. In our hydrodynamic simulations, we do not expect to be able to achieve true numerical convergence, characterized by errors in the solution approaching zero. Rather, we strive to achieve qualitative convergence, a condition that is obtained when physically important macroscopic quantities do not change by a significant amount with an increase in resolution. However, one must be cautious when using this criterion, as the lack of change from one resolution to another may only be an indication that neither resolution is resolving some important physical process. In the case of our rather idealized toroidal models, the physically relevant dynamic processes that determine the eigenvectors of unstable modes should be resolved sufficiently well that this will not be a problem.

2.3 Description of Hybrid Angular Momentum Conservation Scheme

We evolve the following two coupled equations,

$$\frac{\partial}{\partial t}\rho + \nabla \cdot \rho \mathbf{u} = 0, \tag{2-6}$$

and

$$\frac{\partial}{\partial t}(\rho \mathbf{v}) + \nabla \cdot (\rho \mathbf{v} \mathbf{u}) = -\nabla p - \rho \nabla \Phi, \quad (2-7)$$

where ρ is the mass density, p is the gas pressure, and both \mathbf{v} and \mathbf{u} are the same fluid velocity (that is, $\mathbf{u} = \mathbf{v}$), and Φ is the gravitational potential generated by a central point mass, M_{pt} , specifically,

$$\Phi = -\frac{GM_{\text{pt}}}{(x^2 + y^2 + z^2)^{1/2}}. \quad (2-8)$$

Following Call et al. (2010), we have the freedom to choose different coordinate bases for each of the velocity terms, \mathbf{u} and \mathbf{v} , that appear in the dyadic tensor product $\mathbf{v} \mathbf{u}$ in the nonlinear advection term of equation (2-7). Our hybrid scheme, which combines the advantages of a Cartesian mesh (e.g., easily refinable) with the advantage of angular momentum conservation offered by a cylindrical coordinate representation, employs a Cartesian base for the transport velocity, \mathbf{u} , and a cylindrical-coordinate base for the velocity vector, \mathbf{v} , associated with the momentum. In our hybrid scheme we furthermore choose to advect components of the momentum as viewed from the inertial frame through a Cartesian grid that is rotating with angular velocity Ω_0 . See the Appendix for a more detailed description. More specifically, in our simulations that are carried out using the hybrid scheme, the chosen components of the momentum density are the radial momentum density,

$$s_R \equiv \frac{\rho}{R} (xu_x + yu_y) = \frac{\rho}{R} (xu'_x + yu'_y), \quad (2-9)$$

Table 2.2. Resolution

LOR	R_-/R_+	0.1	0.2	0.3	0.4	0.5	0.6	0.7
4	zones ($\times 10^6$)	0.303	0.106
	subgrids	592	208
	$R_0/\Delta x$	19.23	34.31
5	zones ($\times 10^6$)	2.43	2.07	1.64	1.38	0.987	0.844	0.668
	subgrids	4,744	4,040	3,208	2,696	1,928	1,648	1,304
	$R_0/\Delta x$	15.15	27.78	38.46	47.62	55.56	62.50	68.63
6	zones ($\times 10^6$)	10.73	3.232
	subgrids	20,952	6,312
	$R_0/\Delta x$	78.22	139.58
7	zones ($\times 10^6$)	19.31
	subgrids	37,712
	$R_0/\Delta x$	279.16

Note. — Information is shown only for simulations performed here.

the z component of the *inertial*-frame angular momentum density,

$$\ell_z \equiv \rho(xu_y - yu_x) = \rho(xu'_y - yu'_x) + \rho\Omega_0 R^2, \quad (2-10)$$

and the vertical momentum density,

$$s_z \equiv \rho u_z = \rho u'_z, \quad (2-11)$$

where u_x , u_y , and u_z are the x , y , and z components of the fluid velocity as viewed from the inertial frame, respectively; u'_x , u'_y , and u'_z are the components of the fluid velocity in the frame rotating with angular velocity Ω_0 . In our simulations that are carried out using a traditional Cartesian advection scheme, we advect inertial frame Cartesian momentum, that is, we advect the x , y , and z components of the momentum density, defined as $s_x \equiv \rho u_x$, $s_y \equiv \rho u_y$ and $s_z \equiv \rho u_z$.

When rewriting equation (2-7) in individual components, we choose between two sets of momentum equations:

$$\frac{\partial}{\partial t} s_x + \nabla \cdot s_x \mathbf{u}' = -\frac{\partial}{\partial x} p - \rho \frac{\partial}{\partial x} \Phi + s_y \Omega_0, \quad (2-12)$$

$$\frac{\partial}{\partial t} s_y + \nabla \cdot s_y \mathbf{u}' = -\frac{\partial}{\partial y} p - \rho \frac{\partial}{\partial y} \Phi - s_x \Omega_0, \quad (2-13)$$

$$\frac{\partial}{\partial t} s_z + \nabla \cdot s_z \mathbf{u}' = -\frac{\partial}{\partial z} p - \rho \frac{\partial}{\partial z} \Phi, \quad (2-14)$$

when we choose to advect Cartesian momentum; and

$$\frac{\partial}{\partial t} s_R + \nabla \cdot s_R \mathbf{u}' = -\frac{1}{R} \left(x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \right) p - \frac{\rho}{R} \left(x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \right) \Phi + \frac{\ell_z^2}{\rho R^3}, \quad (2-15)$$

$$\frac{\partial}{\partial t} \ell_z + \nabla \cdot \ell_z \mathbf{u}' = \left(y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y} \right) p + \rho \left(y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y} \right) \Phi. \quad (2-16)$$

$$\frac{\partial}{\partial t} s_z + \nabla \cdot s_z \mathbf{u}' = -\frac{\partial}{\partial z} p - \rho \frac{\partial}{\partial z} \Phi, \quad (2-17)$$

when we choose to use the hybrid scheme, where again \mathbf{u}' is the velocity expressed in a Cartesian coordinate frame that is rotating with constant angular frequency, Ω_0 , about the z -axis.

It is important to note exactly what quantities are being reconstructed at the cell faces, as this can affect the results of the simulation, as discussed for example in Norman et al. (1980). The momentum density quantities (s_x , s_y , s_z , s_r) are divided by the mass density before reconstruction. Those reconstructed quantities are then multiplied by the reconstructed mass density. The rotating frame tangential velocity that is required to determine the components of the transport velocity, \mathbf{u}' , are extracted from the inertial frame angular momentum, ℓ_z , by first dividing ℓ_z by ρR , then subtracting the velocity of the frame, $\Omega_0 R$. This reconstructed quantity is then converted back to ℓ_z by reversing the process, now using the reconstructed mass density.

2.4 Postprocessing

We want to quantify small deviations from the initial axisymmetric density distribution, which can be described by,

$$\frac{\delta\rho}{\rho} = f(R, Z)e^{-i[\omega t - m\theta]}, \quad (2-18)$$

where, m is the azimuthal mode number and ω is a complex frequency. In order to do this, we describe the mass density in the system in terms of a discrete Fourier series,

$$\rho(J, L, K, t) = \frac{1}{2}c_0(J, K, t) + \sum_{m=1}^{L_{\max}} c_m(J, K, t) \cos\left[m\frac{2\pi L}{L_{\max}} + \phi_m(J, K, t)\right], \quad (2-19)$$

where J is the radial index, K is the vertical index, and L is the azimuthal index. For each Z we use a two-dimensional bivariate spline technique to interpolate mass density values from the Cartesian grid (x, y) to a polar coordinate grid (R, θ) . We then employ a Fourier transform to determine the coefficient of the discrete Fourier series as follows:

$$\begin{aligned} a_m(J, K) &= \frac{2}{\pi} \sum_{L=1}^{L_{\max}} \rho(J, L, K) \cos(mL\delta\theta)\delta\theta, \\ b_m(J, K) &= \frac{2}{\pi} \sum_{L=1}^{L_{\max}} \rho(J, L, K) \sin(mL\delta\theta)\delta\theta, \\ c_m(J, K) &= [a_m(J, K)^2 + b_m(J, K)^2]^{1/2}, \\ \phi_m(J, K) &= \arctan[-b_m(J, K)/a_m(J, K)], \\ D_m &\equiv \frac{c_m}{\sum_{i=0}^{L_{\max}} c_i}. \end{aligned}$$

After running each simulation and plotting the time variation of D_m and ϕ_m for various modes, the components of the complex frequency ω for a dynamically growing eigenmode of the system can be written as,

$$\begin{aligned} \text{Re}(\omega) &= \frac{d\phi_m}{dt}, \\ \text{Im}(\omega) &= \frac{d \ln c_m}{dt}. \end{aligned}$$

In order to compare with Kojima (1986), we use the quantities,

$$y_1 \equiv [\text{Re}(\omega)/\Omega_0 - m] = \left[\frac{1}{\Omega_0} \frac{d\phi_m}{dt} - m \right], \quad (2-20)$$

and

$$y_2 \equiv \text{Im}(\omega)/\Omega_0 = \left[\frac{1}{\Omega_0} \frac{d \ln c_m}{dt} \right], \quad (2-21)$$

to describe the properties of the unstable modes. When plotting $\ln D_m$ vs time (a “ $D_m - t$ ” diagram, as shown, for example, in Figure 2.3a), where t is normalized to $t_{\text{orb}} \equiv 2\pi/\Omega_0$, the value of y_2 is the slope of the line divided by 2π . Similarly, y_1 is straightforwardly obtained when plotting ϕ_m versus time (a “ $\phi_m - t$ ” diagram, Figure 2.3b) by taking the period measured from the graph, dividing by 2π , and then subtracting m . Note that ϕ_m must be plotted in the inertial frame. Here we usually will present results only at the radius of pressure maximum, R_0 , and in the equatorial plane. In order to gain a better understanding of the structure of these fluctuations at different radii, however, we can perform the same Fourier transform at many different radii along the equatorial plane. We can then obtain a more complete picture of the structure of the unstable eigenvector by plotting, for one particular point in time, the amplitude of a single Fourier mode as a function of radius (a “ $D_m - r$ ” diagram, Figure 2.3c), and the phase angle of a single Fourier mode as a function of radius (a “ $\phi_m - r$ ” diagram, Figure 2.3d).

2.5 Quantifying Results

We have devised a formulaic method for quantifying both the growth rate and the quality of the measurement. The idea is that an analysis of simulations that display long, relatively quiet periods of exponential growth should produce much smaller error bars than the analysis of simulations that display only short, noisy segments of exponential growth. A formulaic method of this type minimizes human judgment and helps ensure that the measurements are reproducible.

As is illustrated in Figures 2.4a and 2.5a we begin the analysis by generating a $D_m - t$ plot from the results of each simulation. (A full suite of $D_m - t$ plots from our 5 LOR evolutions is shown in Figure 3.1.) This plotted $D_m - t$ curve contains 100 individual data points for each orbit. Starting at $t = t_{\text{orb}}/8$, with a window of width $t_{\text{orb}}/4$ centered on that time, a linear regression is used to determine the slope, S . This window is moved continuously across the $D_m - t$ plot to generate a set of data, $S(t)$, as shown in Figures 2.4b and 2.5b.

The relatively flat portion of the Figure 2.4b $S(t)$ curve – bounded by the two vertical solid lines at $t_{\text{start}} = 1.7 t_{\text{orb}}$ and $t_{\text{end}} = 3.76 t_{\text{orb}}$ – corresponds to the period of relatively clean exponential growth in the Figure 2.4a $D_m - t$ plot. In order to determine the best portion of the simulation over which to measure a slope of the $D_m - t$ graph, that is, in order to determine the best time boundaries, t_{start} and t_{end} , we first define the error bar size as,

$$\delta_{y_2} = \frac{\sqrt{\frac{1}{i_{\text{end}} - i_{\text{start}}} \sum_{i=i_{\text{start}}}^{i_{\text{end}}} (S_i - \bar{S})^2}}{t_{\text{end}} - t_{\text{start}}}, \quad (2-22)$$

where i_{start} and i_{end} are the starting and ending points of the selected section respectively, and S_i is the value of $S(t)$ at each time t_i . Every allowable combination of i_{start} and i_{end} is

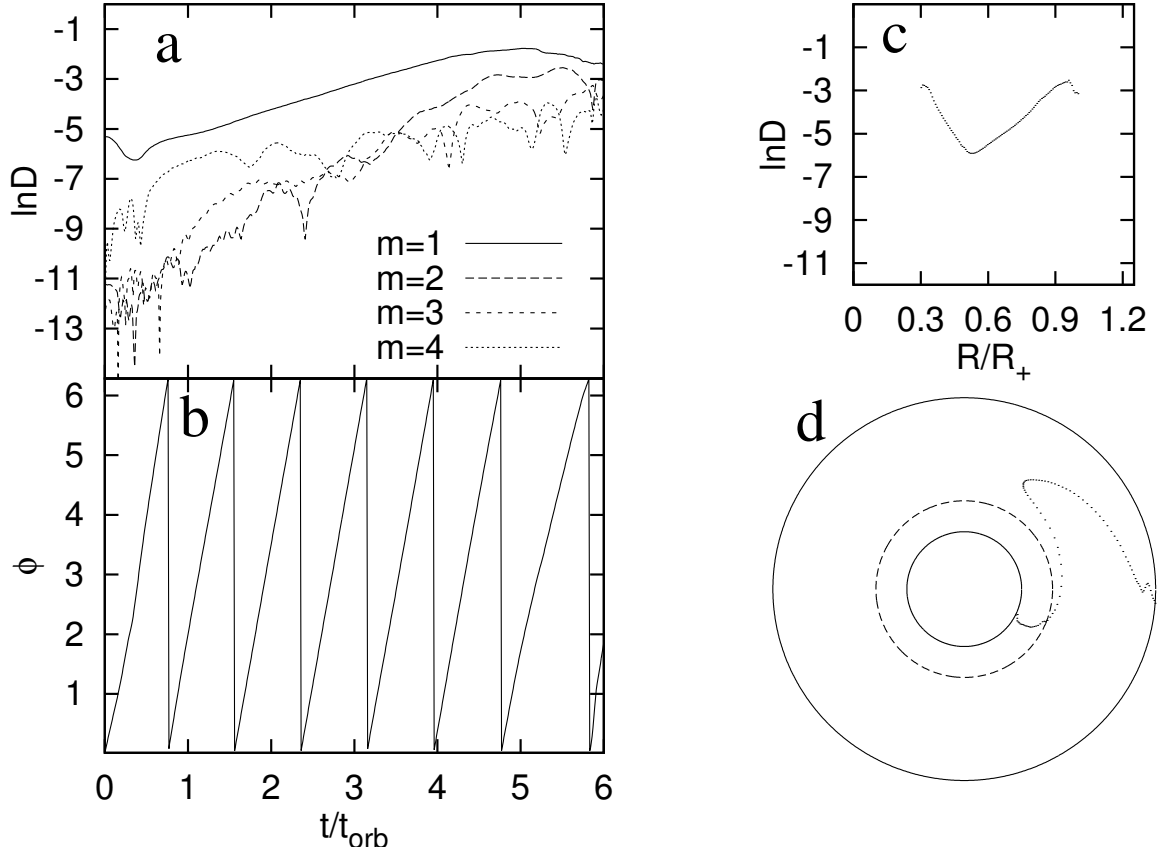


Fig. 2.3.— Data from the Cartesian advection simulation of Model 3 ($R_-/R_+ = 0.3$) is used to demonstrate four diagnostic diagrams, providing a direct comparison to Figure 2 in Woodward et al. (1994). (a) A “ $D_m - t$ ” diagram showing the Fourier amplitude of modes $m = 1, 2, 3, 4$ at the radius of pressure maximum in the equatorial plane. (b) A “ $\phi_m - t$ ” diagram showing the phase angle of the $m = 1$ mode, again at the radius of pressure maximum in the equatorial plane. (c) A “ $D_m - r$ ” diagram showing the amplitude of the $m = 1$ mode as a function of radius in the equatorial plane at time $t = 2.5 t_{\text{orb}}$. (d) A “ $\phi_m - r$ ” diagram showing the azimuthal location of the density maximum ($m = 1$) as a function of radius in the equatorial plane at time $t = 2.5 t_{\text{orb}}$.




Fig. 2.4.— Data from a simulation of Model 3 ($R_-/R_+ = 0.3$) using the hybrid scheme and 5 LOR. (a) Shows a $D_m - t$ plot, with vertical lines marking the starting and ending points of the region used to determine \bar{S} ; the portion of the curve used to measure this slope is shown in bold. A dashed line with a slope equal to the \bar{S} is also shown. (b) Shows the windowed slope measurement, $S(t)$, with vertical lines marking the starting and ending points of the region used to measure \bar{S} . The horizontal dashed line identifies the measured slope.




Fig. 2.5.— Same as Figure 2.4, but from a simulation of Model 7 ($R_-/R_+ = 0.7$) using the Cartesian momentum advection scheme and 5 LOR.

considered (as long as $t_{\text{end}} - t_{\text{start}} \geq 1.0 t_{\text{orb}}$), and the best time boundaries are defined by the interval which generates the smallest error bar, $\delta_{y_2}|_{\text{min}}$. We then define the measured

slope to be the mean of $S(t)$ in the selected region,

$$2\pi y_2 = \bar{S} = \frac{1}{i_{\text{end}} - i_{\text{start}}} \sum_{i=i_{\text{start}}}^{i_{\text{end}}} S_i. \quad (2-23)$$

The magnitude of $\delta_{y_2}|_{\text{min}}$ is not meaningful by itself, but it can be used to compare the relative quality between two separately measured slopes.

Figures 2.4 and 2.5 show examples of $D_m - t$ plots (top) with their corresponding $S(t)$ plots (bottom). The dashed lines shown in both the $D_m - t$ and $S(t)$ plots represent the measured slope, \bar{S} , for the relevant model simulation. Figure 2.4 shows data from a simulation that generated a relatively small $\delta_{y_2}|_{\text{min}}$. In contrast, Figure 2.5 shows data that generated a larger $\delta_{y_2}|_{\text{min}}$. Notice that, when compared to Figure 2.4, the region of exponential growth is shorter and $S(t)$ is much less constant in the selected region identified in Figure 2.5.

Additionally, we sought to measure the y_1 parameter in a way that minimized human interaction. The portion of the $\phi_m - t$ graph that corresponds with the region selected above (t_{start} to t_{end}) was fitted to a line using a least-squares method. The slope of this line is used to compute y_1 , and the residuals of this fit are then used to compute δ_{y_1} .

3 Results

3.1 Uniform Specific Angular Momentum Tori

Figure 3.1 shows $D_m - t$ plots resulting from 14 separate simulations, all run at 5 LOR, with the selected linear portion of the plot in bold, and a dashed line representing the linear fit. Table 3.1 contains a separate row of data corresponding to each of these 14 simulations. Column 1 identifies the initial model configuration as specified in Table 2.1, column 2 identifies the azimuthal mode perturbation that we introduced at the start of each run, and column 3 indicates whether the simulation was carried out using a standard Cartesian (C) advection scheme or our new hybrid (H) scheme. The measured slope, \bar{S} , associated value of y_2 , and error bar, $\delta_{y_2}|_{\min}$, corresponding to the Figure 3.1 plots are tabulated in columns 4, 5, and 6, respectively, of Table 3.1. From the $\phi_m - t$ plot of each simulation (not shown), we measured $d\phi/dt$, y_1 , and δ_{y_1} and recorded the results in columns 7, 8, and 9, respectively, of Table 3.1.

When we apply the post-processing algorithm to the results of our simulations, we see a good match with both the results of Kojima and WTH, although error bars were not provided in either of these earlier works. In Figures 3.2 and 3.3 we plot, respectively,

Table 3.1. Torus Mode Characteristics (5 LOR)

Model (1)	m_{kick} (2)	Scheme ^a (3)	\bar{S} (4)	y_2 (5)	δ_{y_2} (6)	$\frac{d\phi}{dt}$ (7)	y_1 (8)	δ_{y_1} (9)
1	1	C	0.760	0.121	0.007	2.294	0.365	0.0006
1	1	H	0.752	0.120	0.009	2.103	0.335	0.0005
2	1	C	0.792	0.126	0.005	1.885	0.300	0.0002
2	1	H	0.780	0.124	0.005	1.834	0.292	0.0004
3	1	C	1.031	0.164	0.006	1.578	0.251	0.0004
3	1	H	0.994	0.158	0.005	1.579	0.251	0.0003
4	1	C	1.047	0.167	0.008	1.064	0.169	0.0005
4	1	H	1.064	0.169	0.011	1.064	0.169	0.0002
5	2	C	1.253	0.199	0.020	1.251	0.199	0.0017
5	2	H	1.093	0.174	0.016	1.256	0.200	0.0027
6	2	C	1.313	0.209	0.027	0.594	0.095	0.0011
6	2	H	1.320	0.210	0.011	0.792	0.126	0.0007
7	3	C	1.220	0.194	0.037	0.468	0.074	0.0017
7	3	H	1.208	0.192	0.014	0.553	0.088	0.0005

^aH denotes hybrid advection scheme, C denotes Cartesian momentum advection.




Fig. 3.1.— A series of $D_m - t$ plots from every initial model listed in Table 2.1 labeled with either “H” or “C” for simulations using the hybrid scheme or the Cartesian momentum advection scheme, respectively, followed by the model number. All simulations were performed using 5 LOR. In each case the portion of the plot used to measure \bar{S} is highlighted in bold, and a dashed line with the measured slope is also plotted.

the imaginary (y_2) and real (y_1) components of the eigenfrequency as determined from our simulations, from WTH, and from the linear analysis of Kojima. The three dashed curves connect discrete points from the Kojima linear analysis. Because the modes are completely uncoupled in his simulations, Kojima is able to measure growth rates for all modes over most of the range of R_+/R_- that we are interested in. Results from WTH are marked by solid diamonds. Our current results, shown as symbols with error bars, represent the $m = 1$ growth rate for the four fattest tori (the points marked by red boxes), the $m = 2$ growth rate for the next 2 models (points marked by green circles), and the $m = 3$ growth rate for the slimmest torus (points marked by blue triangles). Results from simulations performed with the traditional Cartesian advection scheme are indicated by open symbols while results using our hybrid scheme are indicated by filled symbols. For each model these symbols are separated horizontally on the graph purely for visibility; both runs were started from identical initial states.

While comparison to Kojima’s results provides a good sanity check for our results, they do not really provide an ideal solution to our problem. Kojima’s linear analysis problem has different boundary conditions than the nonlinear simulations performed here and by WTH. Kojima assumes that the surface of the torus remains fixed in space, whereas the hydrocodes allow the torus surface to distort. In the linear analysis, the individual modes do not couple,




Fig. 3.2.— Comparison of the imaginary component (y_2) of the eigenfrequency of various unstable modes between this work and Woodward et al. (1994) and Kojima (1986). The red, green, and blue dashed curves connect discrete points from Kojima’s linear analysis for $m = 1, 2$, and 3 respectively. The points marked with red boxes, green circles, and blue triangles show values measured in this work for $m = 1, 2$, and 3 respectively; open symbols represent the Cartesian momentum advection scheme and filled symbols represent the hybrid scheme. Red and green diamonds represent $m = 1$ and 2 results published in WTH, respectively. Our measured growth rates show good agreement with both previous studies. As described in the text, error bars on data points from this work represent the relative quality of measurements.

whereas in the nonlinear hydrocode they do. Thus we do not expect that our results will converge to Kojima’s, no matter how much resolution we use.

In addition to the information displayed in Figure 3.2 highlighting the fastest growing modes, we observe in all simulations an unphysical development of $m = 4$ distortions. This is undoubtedly due to the 4-fold symmetry of the underlying Cartesian grid. As is illustrated more fully below, in all simulations the time-evolutionary development of $m = 4$ distortions is strongly resolution dependent. The simulations using the hybrid scheme show less unphysical $m = 4$ growth than do the simulations advecting Cartesian momentum.

Figure 3.4 shows the global conservation of angular momentum for a $R_-/R_+ = 0.7$ torus at 3 different LOR. These curves show the difference between the initial total angular momentum, L_0 , and the angular momentum, L , at time t , divided by the initial total angular momentum. The red, green, and blue dashed curves show data from the Cartesian momentum advection scheme at 5, 4, and 3 LOR respectively; the red, green, and blue solid curves show data from the hybrid scheme at 5, 4, and 3 LOR respectively. The simulations




Fig. 3.3.— As in Figure 3.2, but showing the real component (y_1) of the eigenfrequency of various unstable modes. Our measured frequencies show good agreement with both previous studies.

run with the hybrid scheme (solid curves) show non-conservation at levels $\Delta L/L_0 \lesssim 10^{-13}$ due only to machine truncation error. The hybrid scheme simulation at 3 LOR (solid blue curve) seems to display better conservation than the 4 and 5 LOR simulations (overlapping solid green and red curves), however this is only because the 3 LOR simulation takes far fewer timesteps than the higher resolution simulations and thus accumulates less error due to truncation. The simulations run with the Cartesian momentum advection scheme display resolution dependent global angular momentum conservation, at levels of $\Delta L/L_0 \approx 10^{-5}$, 10^{-4} , and 10^{-3} , for 5, 4, and 3 LOR, respectively.

Figures 3.5 - 3.10 present data from two initial models, each evolved in 6 separate simulations (2 schemes \times 3 LORs). Figures 3.5 - 3.7 present data from Model 3 ($R_-/R_+ = 0.3$); Figures 3.8 - 3.10 present data from Model 7 ($R_-/R_+ = 0.7$). In Figures 3.5 and 3.8, each of the six panels shows the $D_m - t$ plot for the fastest growing unstable mode (either $m = 1$ or $m = 3$, solid curves) and for $m = 4$ (dashed curves) from that initial model evolved in a separate simulation. The plots shown in the left column were produced from simulations evolved with the hybrid scheme; plots in the right column were produced from simulations evolved advecting Cartesian momentum. Figures 3.6 and 3.9 show the same data, but each frame combines all three LOR and separates the plots of the two modes ($m = 1$ and $m = 4$ for Figure 3.6, $m = 3$ and $m = 4$ for Figure 3.9). Figures 3.7 and 3.10 combine the hybrid scheme and Cartesian plots while separating out the different LOR.

In Figure 3.5 we see that all six evolutions show virtually identical behavior of the growth of the $m = 1$ mode with time. Hence the measured growth rate is quite independent




Fig. 3.4.— The accumulated change in Model 7’s total angular momentum ($L - L_0$), measured relative to its initial value, L_0 , is shown as a function of time, t/t_{orb} , for 6 different simulations – 3 different levels of refinement and using Cartesian momentum advection (dashed curves) or the hybrid scheme (solid curves). In the case of both schemes, the red, green, and blue curves show data from simulations conducted with 5, 4, and 3 LOR respectively. When using the Cartesian momentum advection scheme, the level of angular momentum conservation shows clear resolution dependence. The hybrid scheme conserves angular momentum at a level set by machine truncation error.

of the chosen advection scheme or selected LOR. However, from Figure 3.5 we can also see that the time-dependent behavior of the amplitude of the $m = 4$ distortion seems to be strongly resolution dependent. As stated earlier, this likely reflects the 4-fold symmetry of the Cartesian grid structure. In Figure 3.6 the two upper panels ($m = 1$) show more clearly that results from both the hybrid scheme and Cartesian scheme are nearly identical, with the lower resolution simulations displaying only slightly noisier $D_m - t$ plots. The lower two panels show how the amplitude of $m = 4$ fluctuations decreases dramatically with increasing resolution. We can see from the left column of Figure 3.7, showing $m = 1$, the difference between the Cartesian and hybrid scheme plots differs less with increasing resolution, suggesting that




Fig. 3.5.— Data from Model 3 ($R_-/R_+ = 0.3$) simulations. Each of the six panels shows the $D_m - t$ plot for $m = 1$ (solid curve) and $m = 4$ (dashed curve) obtained from, as labeled, either the hybrid scheme or the Cartesian advection scheme for 4, 5, or 6 LOR. While the time-dependent growth of the unstable $m = 1$ mode is very similar in all cases, the amplitude of $m = 4$ appears to be strongly resolution dependent, reflecting the 4-fold symmetry of the Cartesian grid structure.




Fig. 3.6.— The same information as shown in Figure 3.5, but grouped differently. Each frame shows all three levels of refinement for either $m = 1$ (top row) or $m = 4$ (bottom row), using either the hybrid scheme (left column) or Cartesian momentum advection scheme (right column). The two upper panels ($m = 1$) show that both the hybrid and Cartesian seem to converge, as the lower resolution simulations display noisier $D_m - t$ plots. The lower two panels show how the amplitude of $m = 4$ fluctuations decreases dramatically with increasing grid resolution.




Fig. 3.7.— The same information as shown in Figure 3.5, but grouped differently. Each frame compares hybrid (solid curve) to Cartesian (dashed curve) momentum advection schemes, for either $m = 1$ (left column) or $m = 4$ (right column), and at 4, 5, and 6 levels of refinement (bottom, middle, and top rows, respectively). The left column ($m = 1$) illustrates how the difference between the Cartesian and hybrid scheme plots differs less with increasing resolution, suggesting that they are both converging to the same answer. The right column ($m = 4$) shows that, at each specified resolution, the hybrid scheme generally produces slightly lower amplitude $m = 4$ distortions than the Cartesian advection scheme.

they are both converging to the same answer. In the right column, showing $m = 4$, we see that at each resolution the hybrid scheme has slightly lower amplitude.

In contrast to the previous set of figures, Figures 3.8-3.10 (the slimmest torus evolutions) show a dramatic difference between the hybrid scheme and Cartesian momentum advection scheme. In Figure 3.8, the differences are quite apparent not only between the left and right columns (hybrid scheme and Cartesian, respectively) but also between each row (different LOR). This is a model in which the most rapidly growing mode should be $m = 3$, but development of unphysical $m = 4$ distortions can dominate. Note that for the hybrid scheme, $m = 4$ dominates only at 4 LOR. For the Cartesian, $m = 4$ dominates at 4 and 5 LOR, and is only matched by $m = 3$ at 6 LOR. The increasing amplitude of the $m = 3$ mode compared to the $m = 4$ mode with resolution is indicative that a dominant $m = 3$ is the true character of the most unstable eigenmode for this model. This is also consistent with the relative amplitudes measured by Kojima (1986), as shown in Figure 3.2. Focusing on the top row of Figure 3.9, we see that the $m = 3$ mode amplitude tracks the amplitude of the next higher level of resolution until a certain point where it turns off (for hybrid 4 LOR, Cartesian 4 and 5 LOR). Looking below to the $m = 4$ plots, we can see that these turn-off points correspond to a time where the $m = 4$ amplitude has risen to $\ln D_4 \approx -3$. Apparently, once a mode reaches this amplitude, the modes are no longer sufficiently decoupled, and the $m = 3$ cannot continue to grow exponentially at the rate predicted by simulations that do not allow different modes to couple. We also observe that the amplitude of the $m = 4$ mode decreases with increasing resolution. In Figure 3.10 we again see dramatic differences between the hybrid scheme and Cartesian momentum scheme, especially in $m = 4$.

For the $R_-/R_+ = 0.7$ model, simulations that relied on Cartesian momentum advection failed to achieve qualitative convergence even at 6 LOR. So, in order to show that both schemes (the Cartesian momentum advection and the hybrid scheme) ultimately converge to the same answer, the model was re-run advecting Cartesian momentum at 7 LOR. This is impossible using the OpenMP version of the code, so the simulation was run using Octopus, which employs computational fluid algorithms identical to those used in Scorpio, but built within the HPX parallel programming framework¹. The Octopus code exposes a greater degree of parallelism than the OpenMP code, making the simulation practical. Octopus parallelizes the invocation of various “kernels” to every subgrid in the octree-based grid structure. In contrast, the OpenMP code only parallelizes local loops inside these kernels while the invocation of the kernels is done serially. For example, in Octopus, the computation of fluxes for all subgrids is done in parallel; in the OpenMP code, the computation of the fluxes contained within each subgrid are computed in parallel, but only one subgrid’s fluxes are computed at a time.

The $D_m - t$ plots presented in Figure 3.11 show the development of the $m = 3$ mode from seven separate simulations of Model 7 ($R_-/R_+ = 0.7$). Combining the top two panels of Figure 3.9, three different resolutions, 4, 5, and 6 LOR, are shown for both the hybrid scheme

¹<http://stellar.cct.lsu.edu/>




Fig. 3.8.— Data from Model 7 ($R_-/R_+ = 0.7$) simulations. Each of the six panels shows the $D_m - t$ plot for $m = 3$ (solid curve) and $m = 4$ (dashed curve) obtained from, as labeled, either the hybrid scheme or the Cartesian advection scheme for 4, 5, or 6 LOR.




Fig. 3.9.— The same information as shown in Figure 3.8, but grouped differently. Each frame shows all three levels of refinement for either $m = 3$ (top row) or $m = 4$ (bottom row), using either the hybrid scheme (left column) or Cartesian momentum advection scheme (right column).

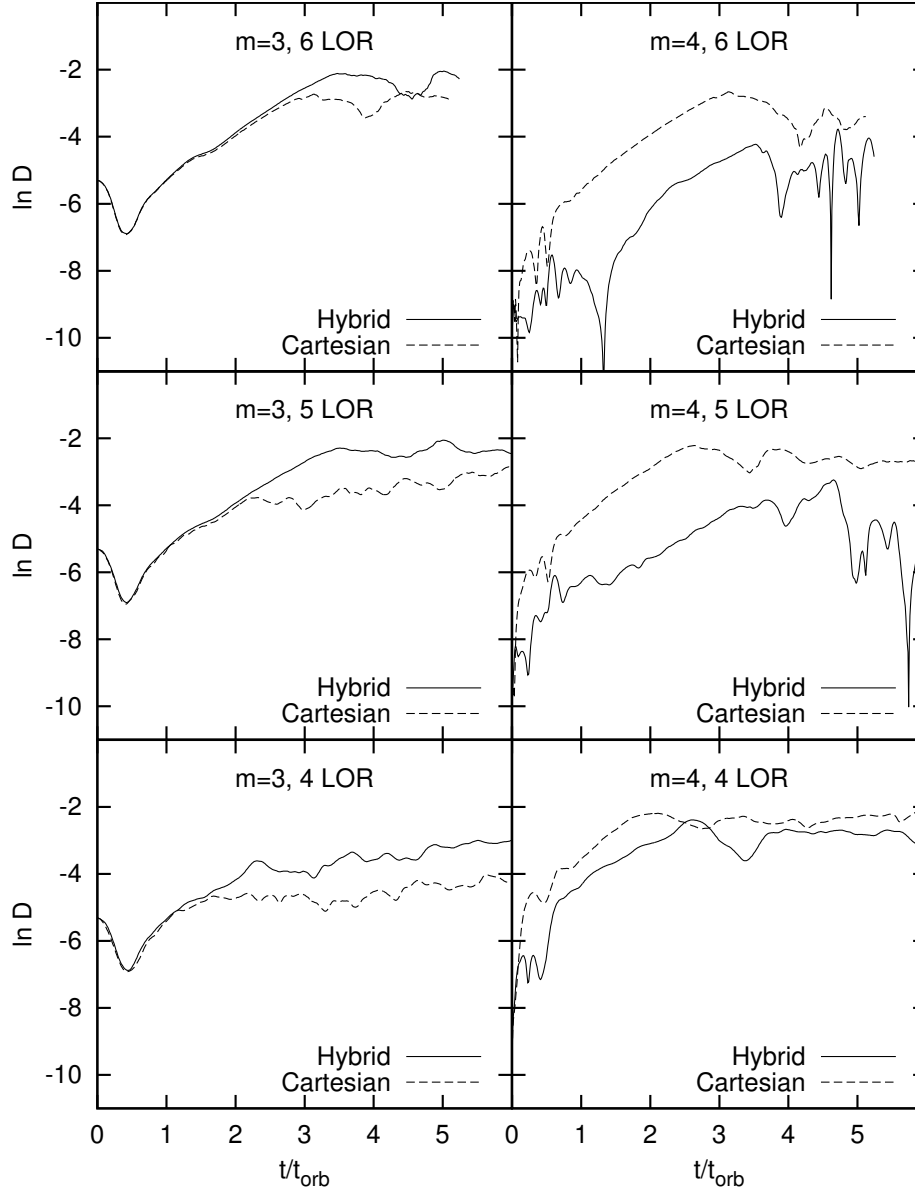


Fig. 3.10.— The same information as shown in Figure 3.8, but grouped differently. Each frame compares hybrid (solid curve) to Cartesian (dashed curve) momentum advection schemes, for either $m = 3$ (left column) or $m = 4$ (right column), and at 4, 5, and 6 levels of refinement (bottom, middle, and top rows, respectively).




Fig. 3.11.— $D_m - t$ plots showing the development of the $m = 3$ mode from seven different simulations of Model 7 ($R_-/R_+ = 0.7$). The hybrid scheme simulations (solid curves) show qualitative convergence at 5 LOR, while the Cartesian momentum advection scheme (dashed curves) does not converge until 7 LOR, requiring a factor of ~ 30 more computational zones.

and the Cartesian momentum advection; additionally, 7 LOR is shown for the Cartesian simulation (red dashed curve). While the 4, 5, and 6 LOR curves for the Cartesian simulations don't agree with the results of the converged hybrid scheme curves (at 5 and 6 LOR), the 7 LOR Cartesian does lie almost directly on top of those curves. This demonstrates that both schemes do ultimately converge to the same answer, but that, in this case, the hybrid scheme achieves qualitative convergence a full 2 levels of refinement sooner, and with a factor of ~ 30 fewer fine zones. Figure 3.12 shows the behavior of the $m = 4$ distortion for the same set of simulations. We see that the amplitude of the $m = 4$ fluctuation is strongly resolution dependent and, furthermore, that the hybrid scheme demonstrates much lower levels of the unphysical $m = 4$ distortion than the Cartesian advection scheme.

Figure 3.13 illustrates what the nonlinear amplitude structure of these modes looks like physically. These equatorial ($z = 0$) plane mass density plots drawn from the Cartesian momentum advection simulations confirm what the Figure 3.11 and 3.12 $D_m - t$ plots tell us about the relative amplitudes of the $m = 3$ mode and the $m = 4$ distortion. Specifically,




Fig. 3.12.— $D_m - t$ plots showing the development of the $m = 4$ distortion in Model 7 ($R_-/R_+ = 0.7$). At each specified LOR, the hybrid scheme (solid curves) shows much lower levels of development of this unphysical distortion than the Cartesian advection scheme (dashed curves).

at 4 and 5 LOR, the $m = 4$ distortion is clearly dominant. At 6 LOR, the $m = 3$ mode and $m = 4$ distortion are approximately the same amplitude, and the mass density distribution shows this. At 7 LOR, the mass density distribution shows a dominant $m = 3$ mode.

3.2 Near-Keplerian Disk Simulation

As an additional test of the method, we simulated a disk with a near-Keplerian rotation profile. As with the constant specific angular momentum tori, we assume a polytropic equation of state (equation 2-1) and $n = 3$. Instead of $q = 2$, we use $q = 1.52$ to determine the rotation profile in equation (2-2). Figure 3.14 shows the meridional-plane cross section of




Fig. 3.13.— Mass density plots from the Cartesian momentum advection simulations of the Model 7 ($R_-/R_+ = 0.7$) torus, showing the progression from $m = 4$ dominated evolutions at lower resolutions to $m = 3$ dominated evolutions at the highest resolution. Each slice shows the mass density in the $z = 0$ plane, and is labeled with the number of LOR used in the simulation. Data for 4, 5, 6, and 7 LOR are taken from $t = 1.8t_{\text{orb}}$, $2.1t_{\text{orb}}$, $3.0t_{\text{orb}}$, and $2.5t_{\text{orb}}$ respectively.

the disk (dashed blue curve) compared to the thickest (Model 7, solid red curve) and thinnest (Model 1, dashed green line) uniform specific angular momentum tori described in §2.1. We selected a ratio $R_-/R_+ = 0.1$ and a rotational profile index $q = 1.52$ to generate a disk with a height approximately equal to that of the Model 7 torus and a width approximately equal to that of the Model 1 torus. The initial axisymmetric density distribution is given by,

$$\rho = \left[\frac{H}{4K} \right]^3, \quad (3-1)$$

where,

$$H = \frac{GM_{\text{pt}}}{\sqrt{R^2 + Z^2}} - \Psi(R) + C, \quad (3-2)$$

and,

$$C = -\frac{GM_{\text{pt}}}{R_0} \left(1 + \frac{1}{2(1-q)} \right), \quad (3-3)$$

and,

$$\Psi(R) = -\frac{1}{2(1-q)} \left[\frac{GM_{\text{pt}}}{R_0} \left(\frac{R}{R_0} \right)^{2(1-q)} \right], \quad (3-4)$$

and finally,

$$R_0 = \left[\frac{R_-^{2(1-q)} - R_+^{2(1-q)}}{-2(1-q)(1/R_- - 1/R_+)} \right]^{\frac{1}{3-2q}}. \quad (3-5)$$

This simulation was performed with Octopus using the hybrid scheme. It was run on 256 Intel Sandy Bridge processors on SuperMike II, and took approximately 89 hours to complete. The simulation used a grid rotating at the same rate as the disk at pressure maximum, $\Omega_0 = \Omega(R_0)$. The simulation ran for $9.87t_{\text{orb}}$, where $t_{\text{orb}} = 2\pi/\Omega_0$. During the time it took for the fluid at R_0 to complete 9.87 orbits, the fluid at the inner edge of the disk completed approximately 30 orbits, and the fluid at the outer edge of the disk completed approximately 0.92 orbits.

$D_m - t$ plots for the disk evolution (Figure 3.15) show that the $m = 4$ distortion begins at a much higher level than the other modes. By the end of the simulation, the other modes had risen to approximately the same level as the $m = 4$ distortion. None of the modes seem to have entered any period of exponential growth, although the $m = 4$ distortion did display a brief segment of exponential change in amplitude.

Figures 3.16 and 3.17 present the $\phi_m - t$ plots for the disk evolution in the lab and the rotating frames, respectively. The lack of a clean saw-tooth shape in any of the plots points to a lack of global organization. The $m = 4$ plot in Figure 3.17 shows that the $m = 4$ disturbance moves very little in relation to the computational grid.

Lastly, Figure 3.18 presents $\phi_m - r$ plots for all four modes ($m = 1, 2, 3$ and 4), taken from the end of the simulation ($t = 9.87t_{\text{orb}}$). The lack of any distinctive shape in these plots shows that there is no global organization in any of the modes.




Fig. 3.14.— Meridional-plane cross sections of the $q = 1.52$ near-Keplerian disk (dashed blue curve) and the Model 1 and Model 7 $q = 2$ uniform specific angular momentum tori (red solid curve and green dashed curve, respectively). Lengths are normalized to the radius of the outer edge, R_+ .




Fig. 3.15.— $D_m - t$ plots for the near-Keplerian disk evolution, taken from the radius at pressure maximum, $R_0 \approx 0.209R_+$, in the equatorial plane. $m = 1, 2, 3,$ and $4,$ are represented by solid red, dashed green, dashed blue, and dotted magenta curves, respectively.




Fig. 3.16.— $\phi_m - t$ plots for the near-Keplerian disk evolution, taken from the radius at pressure maximum, $R_0 \approx 0.209R_+$, in the equatorial plane. $m = 1, 2, 3,$ and $4,$ are represented by solid red, dashed green, dashed blue, and dotted magenta curves, respectively.




Fig. 3.17.— Same as Figure 3.16, but plotted in the rotating frame.




Fig. 3.18.— $\phi_m - r$ plots for the near-Keplerian disk evolution, taken from the equatorial plane at time $t = 9.87t_{orb}$. The inner and outer edges of the disk are marked by solid lines and the radius of pressure maximum, $R_0 \approx 0.209R_+$, is marked by a dashed line. None of the four modes show any global organization₃₆

4 Conclusion

4.1 This Work

Our ultimate goal is to model in as realistic a manner as possible the dynamical evolution of mass-transferring binary systems. This can only be accomplished if the hydrodynamic code that is used to perform each simulation conserves angular momentum extremely well. We also need to have the flexibility of AMR to adequately resolve spatial features across many orders of magnitude in length scales simultaneously. The hybrid scheme described here allows us to conserve angular momentum to high accuracy on a refined Cartesian mesh, facilitating the use of AMR.

Our hybrid scheme is an implementation of the theoretical formulation developed by Call et al. (2010), which shows that we have the freedom to choose different coordinate bases for the transport velocity relative to the grid and the advected momentum quantities. In the past, these chosen basis sets typically have been the same – resulting in the advection of Cartesian momentum components on a Cartesian mesh, or cylindrical momentum components on a cylindrical mesh. In the hybrid scheme implemented here, we have chosen to advect cylindrical momentum components across a rotating Cartesian mesh. This allows us to conserve angular momentum to machine precision while capitalizing on the advantages of a Cartesian mesh, such as mesh refinement.

In order to test this method, we followed the development of nonaxisymmetric instabilities in massless PP tori having $n = 3$ and $q = 2$ (uniform specific angular momentum). This is a well-defined, fully three-dimensional problem with a reproducible solution. We evolved seven different initial tori with aspect ratios ranging from $R_-/R_+ = 0.1$ to 0.7. We chose to evolve two particular models, $R_-/R_+ = 0.3$ and 0.7, using several different grid resolutions. We compared our results to the linear stability analysis of Kojima (1986) and to the nonlinear hydrodynamics results of Woodward et al. (1994). Our code achieved good agreement with results from these previous studies.

We also introduced a prescriptive method for measuring the real and imaginary parts of the eigenfrequency of unstable modes, attaching an uncertainty to those measurements. This was done in an effort to increase transparency, reduce the influence of human judgment, and facilitate the reproducibility of these simulations. Through this work we have illustrated the utility of the PP tori as a new test problem, to be added to the standard suite of hydrodynamic test problems, that provides a means for measuring the ability of a particular code to correctly transport and conserve angular momentum.

A comparison of the resolution dependence of the hybrid scheme compared to the Cartesian momentum advection scheme shows that the hybrid scheme achieves qualitative convergence at grid resolutions that are equal to or lower than the Cartesian scheme. Specifically, we observe that in the $R_-/R_+ = 0.7$ torus, the hybrid scheme achieves qualitative convergence at only 5 LOR, whereas the Cartesian advection scheme required 7 LOR to achieve the same convergence – requiring a factor of ~ 30 more computation zones. The hybrid

scheme also reduces the level of unphysical $m = 4$ distortions that characteristically appear in simulations involving angular motion across a Cartesian grid.

Here we have demonstrated the utility of the hybrid scheme, which is only one very specific implementation of the formalism presented by Call et al. (2010), which can be applied in a fully relativistic generalized coordinate system.

4.2 Future Goals

We strive to eventually achieve the ability to simulate binary systems with enough accuracy to make meaningful comparisons to observations. In order to do this, these simulations must be as realistic as possible. Much work has been done by groups worldwide to incorporate important physics (i.e. gravity, magnetohydrodynamics, realistic equations of state, radiation transfer, nuclear burning) into current codes. However, the more critical problem is the inability of current codes to simultaneously resolve all of the important time and length scales. Lack of computational resources is not an issue – we have very large computers – but current algorithms are unable to utilize them efficiently. Scaling is the problem – past a certain point, using more computational resources will not speed up a simulation. As we move forward into the future, supercomputers will continue to become larger, increasing the gap between the maximum available computational power and the amount of this power we are able to effectively utilize.

Supercomputers are not only becoming larger, but more heterogeneous as well. Attaching accelerators, either graphics-processing units (GPUs) or the new Intel Xeon Phi coprocessor is becoming more common. For example, China’s Tianhe-2 supercomputer (#1 on the TOP500 as of November 2013) has three Xeon Phi coprocessors attached to each compute node. Current algorithms not only scale poorly, but can’t utilize coprocessors without complicated application-specific code.

We have to turn our algorithms “upside-down” – this is what Octopus has done in order to take advantage of HPX’s message-driven computational model. HPX naturally utilizes the inhomogeneous nature of today’s supercomputers and easily handles the load balancing problem inherent in AMR algorithms. HPX will eventually provide the ability to scale up to millions of computing cores, enabling us to perform astrophysics simulations of unprecedented size.

Our hybrid scheme provides the ability to conserve angular momentum on a refined Cartesian mesh to the degree needed by an astrophysics code to perform binary orbits. We believe the combination of our hybrid scheme, Cartesian AMR, and HPX – implemented in Octopus and demonstrated here – can be the key to more fully utilizing the computational resources available today and enabling transformative astrophysics simulations. If combined with the physics already implemented by other groups, such a tool would bring the goal of simulating binary systems with enough physical realism to compare to observational data closer than ever before.

References

- Call, J. M., Tohline, J. E., & Lehner, L. 2010, *Classical and Quantum Gravity*, 27, 175002
- Clayton, G. C. 1996, *PASP*, 108, 225
- Colella, P., & Woodward, P. R. 1984, *Journal of Computational Physics*, 54, 174
- Diemer, B., Kessler, R., Graziani, C., et al. 2013, *Astrophysical Journal*, 773, 119
- D'Souza, M. C. R., Motl, P. M., Tohline, J. E., & Frank, J. 2006, *Astrophysical Journal*, 643, 381
- Even, W., & Tohline, J. E. 2009, *Astrophysical Journal Supplement*, 184, 248
- Fragile, P. C., Lindner, C. C., Anninos, P., & Salmonson, J. D. 2009, *Astrophysical Journal*, 691, 482
- Frank, J., King, A., & Raine, D. J. 2002, *Accretion Power in Astrophysics: Third Edition*
- Frank, J., & Robertson, J. A. 1988, *Monthly Notices of the Royal Astronomical Society*, 232, 1
- Fryer, C. L., Ruitter, A. J., Belczynski, K., et al. 2010, *Astrophysical Journal*, 725, 296
- Fryxell, B., Olson, K., Ricker, P., et al. 2000, *Astrophysical Journal Supplement*, 131, 273
- Geier, S., Nesslinger, S., Heber, U., et al. 2007, *Astronomy and Astrophysics*, 464, 299
- Hawley, J. F. 1987, *Monthly Notices of the Royal Astronomical Society*, 225, 677
- . 1990, *Astrophysical Journal*, 356, 580
- Hawley, J. F., Richers, S. A., Guan, X., & Krolik, J. H. 2013, *Astrophysical Journal*, 772, 102
- Hawley, W. P., Athanassiadou, T., & Timmes, F. X. 2012, *Astrophysical Journal*, 759, 39
- Hayes, J. C., Norman, M. L., Fiedler, R. A., et al. 2006, *Astrophysical Journal Supplement*, 165, 188
- Hillebrandt, W., & Niemeyer, J. C. 2000, *Annual Review of Astronomy and Astrophysics*, 38, 191
- Kojima, Y. 1986, *Progress of Theoretical Physics*, 75, 251
- Kurganov, A., & Tadmor, E. 2000, *Journal of Computational Physics*, 160, 241

- MacNeice, P., Olson, K. M., Mobarrry, C., de Fainchtein, R., & Packer, C. 2000, *Computer Physics Communications*, 126, 330
- Marcello, D. C., & Tohline, J. E. 2012, *Astrophysical Journal Supplement*, 199, 35
- Mignone, A., Bodo, G., Massaglia, S., et al. 2007, *Astrophysical Journal Supplement*, 170, 228
- Motl, P. M., Frank, J., Tohline, J. E., & D'Souza, M. C. R. 2007, *Astrophysical Journal*, 670, 1314
- Motl, P. M., Tohline, J. E., & Frank, J. 2002, *Astrophysical Journal Supplement*, 138, 121
- Nelemans, G., Yungelson, L. R., Portegies Zwart, S. F., & Verbunt, F. 2001, *Astronomy and Astrophysics*, 365, 491
- New, K. C. B., & Tohline, J. E. 1997, *Astrophysical Journal*, 490, 311
- Norman, M. L., & Wilson, J. R. 1978, *Astrophysical Journal*, 224, 497
- Norman, M. L., Wilson, J. R., & Barton, R. T. 1980, *Astrophysical Journal*, 239, 968
- Papaloizou, J. C. B., & Pringle, J. E. 1984, *Monthly Notices of the Royal Astronomical Society*, 208, 721
- Perlmutter, S., Aldering, G., Goldhaber, G., et al. 1999, *Astrophysical Journal*, 517, 565
- Phillips, M. M. 1993, *Astrophysical Journal Letters*, 413, L105
- Riess, A. G., Filippenko, A. V., Challis, P., et al. 1998, *AJ*, 116, 1009
- Rosswog, S., Kasen, D., Guillochon, J., & Ramirez-Ruiz, E. 2009, *Astrophysical Journal Letters*, 705, L128
- Schaefer, B. E., & Pagnotta, A. 2012, *Nature*, 481, 164
- Shu, C.-W., & Osher, S. 1988, *Journal of Computational Physics*, 77, 439
- Staff, J. E., Menon, A., Herwig, F., et al. 2012, *Astrophysical Journal*, 757, 76
- Webbink, R. F. 1984, *Astrophysical Journal*, 277, 355
- Whelan, J., & Iben, Jr., I. 1973, *Astrophysical Journal*, 186, 1007
- Woodward, J. W., Tohline, J. E., & Hachisu, I. 1994, *Astrophysical Journal*, 420, 247
- Yungelson, L., Nelemans, G., & Portegies Zwart, S. F. 2004, in *Revista Mexicana de Astronomia y Astrofisica*, vol. 27, Vol. 20, *Revista Mexicana de Astronomia y Astrofisica Conference Series*, ed. G. Tovmassian & E. Sion, 117–119

Zink, B., Schnetter, E., & Tiglio, M. 2008, *Physical Review D*, 77, 103015

Zurek, W. H., & Benz, W. 1986, *Astrophysical Journal*, 308, 123

Appendix: Momentum Equation Formulation

This appendix is meant to more completely explain the origin of the momentum equations used in this work, and their place in the larger context of astrophysical hydrodynamics codes. The theoretical justification for the equations we use comes from Call, Tohline, & Lehner (2010), which describes a generalized version of the hybrid advection scheme used in this work, expanded to general relativity and to any curvilinear coordinate system. The equations used in this work differ from those typically seen in the community because they exploit two key advantages of the formalism described by Call, Tohline, & Lehner (2010):

- We are allowed to advect inertial-frame quantities on a rotating grid.
- We are allowed to choose different coordinate bases for the advected momentum quantities and for the grid on which we choose to advect these quantities.

Specifically, in the hybrid scheme as implemented here, we choose to advect cylindrical momentum quantities (that is, radial momentum, angular momentum, and z -direction momentum) measured in the inertial frame on a Cartesian mesh rotating with a fixed angular velocity. Below we will show how this simplifies and in certain cases eliminates the source terms from individual components of the momentum equation.

A.1 Statements of Conservation

We typically encounter hyperbolic PDEs of the following form:

$$\frac{d}{dt}\Psi + \Psi\nabla \cdot \mathbf{u} = S, \quad (\text{A1})$$

where \mathbf{u} is the velocity of the fluid as viewed from an inertial frame of reference, and the total time derivative is the Lagrangian derivative, following an individual fluid element as it moves through space. When the source term S is zero, then Ψ represents the volume density of a conserved quantity. Mass density, ρ , for example, is a conserved quantity, and the continuity equation does in fact have the form:

$$\frac{d}{dt}\rho + \rho\nabla \cdot \mathbf{u} = 0. \quad (\text{A2})$$

We should then expect that in the case of an axisymmetric distribution of fluid moving in an axisymmetric potential, when the azimuthal component of the angular momentum is conserved, that we will encounter an equation of the form,

$$\frac{d}{dt}(\rho R u_\varphi) + (\rho R u_\varphi)\nabla \cdot \mathbf{u} = 0, \quad (\text{A3})$$

where R is the cylindrical radius, and $u_\varphi = R\dot{\varphi}$ is the azimuthal component of the inertial velocity field \mathbf{u} .

A.2 Rotating Frame of Reference

In order to reduce the motion of the fluid through the computational grid (thereby reducing the effects of numerical diffusion and artificial viscosity), we often wish to view the fluid from a rotating reference frame. Mathematically, we will accomplish this by changing the velocity in the divergence term to account for the frame velocity field, that is, we will replace \mathbf{u} with,

$$\mathbf{u}' = \mathbf{u} - \mathbf{u}_{\text{frame}}. \quad (\text{A4})$$

If the velocity field, $\mathbf{u}_{\text{frame}}$, is divergence-free, then the transformation is trivial. For a frame rotating with angular velocity Ω_0 ,

$$\mathbf{u}_{\text{frame}} = R\Omega_0\hat{\mathbf{e}}_\varphi, \quad (\text{A5})$$

and, utilizing cylindrical coordinates,

$$\nabla \cdot \mathbf{u}_{\text{frame}} = \frac{\partial}{\partial R}(0) + \frac{1}{R} \frac{\partial}{\partial \varphi}(R\Omega_0) + \frac{\partial}{\partial z}(0) = 0. \quad (\text{A6})$$

Hence,

$$\frac{d}{dt}\Psi + \Psi\nabla \cdot \mathbf{u}' = \frac{d}{dt}\Psi + \Psi\nabla \cdot [\mathbf{u} - \mathbf{u}_{\text{frame}}] = \frac{d}{dt}\Psi + \Psi\nabla \cdot \mathbf{u} \quad (\text{A7})$$

so the new hyperbolic PDE becomes,

$$\frac{d}{dt}\Psi + \Psi\nabla \cdot \mathbf{u}' = S, \quad (\text{A8})$$

and we are confident that this new PDE represents the physics of the system as well as the original PDE.

A.3 Eulerian Representation

In order to follow the time-rate of change of a quantity with respect to a point in space fixed with respect to the chosen frame of reference, we must use the following transformation from the Lagrangian to the Eulerian representation:

$$\frac{d}{dt}\Psi \rightarrow \frac{\partial}{\partial t}\Psi + \mathbf{u}' \cdot \nabla\Psi. \quad (\text{A9})$$

We can then rewrite the hyperbolic PDE as,

$$\frac{\partial}{\partial t}\Psi + \mathbf{u}' \cdot \nabla\Psi + \Psi\nabla \cdot \mathbf{u}' = S \quad (\text{A10})$$

or, more succinctly,

$$\frac{\partial}{\partial t}\Psi + \nabla \cdot (\Psi\mathbf{u}') = S. \quad (\text{A11})$$

We can recover the inertial-frame version of the equation simply by setting $\Omega_0 = 0$, which is equivalent to setting $\mathbf{u}' = \mathbf{u}$,

$$\frac{\partial}{\partial t}\Psi + \nabla \cdot (\Psi\mathbf{u}) = S. \quad (\text{A12})$$

While the underlying physics is identical, a distinction must be made regarding how the two equations are interpreted. Equation (A12) represents the time-rate of change of Ψ at a fixed point in inertial space, while equation (A11) provides the time-rate of change of Ψ at a fixed point in the rotating coordinate frame. Notice that this is totally independent of what quantity, Ψ , we choose to advect.

A.4 Angular Momentum Conservation

When the three vector components of the Euler equation of motion are projected onto a non-rotating cylindrical coordinate grid, the azimuthal component may be written as,

$$\frac{d}{dt}(\rho R u_\varphi) + (\rho R u_\varphi)\nabla \cdot \mathbf{u} = -\frac{\partial}{\partial \varphi}p - \rho\frac{\partial}{\partial \varphi}\Phi. \quad (\text{A13})$$

For this equation, the source term is,

$$S = -\frac{\partial}{\partial \varphi}p - \rho\frac{\partial}{\partial \varphi}\Phi, \quad (\text{A14})$$

and $\Psi = (\rho R u_\varphi)$ is the inertial-frame angular momentum density with respect to the z -coordinate axis. This corresponds to “Case B ($\eta = 3$)” in Call, Tohline, & Lehner (2010). Angular momentum will be conserved locally if the source term, $S = 0$. This will happen if the azimuthal derivative of the gravitational potential and the azimuthal derivative in the pressure are both zero, or if these two terms balance one another (i.e., $\partial p/\partial \varphi = -\rho\partial\Phi/\partial \varphi$). Based on the discussion above, it is perfectly valid to view the flow from a rotating frame of reference, in which case the equation is simply,

$$\frac{d}{dt}(\rho R u_\varphi) + (\rho R u_\varphi)\nabla \cdot \mathbf{u}' = -\frac{\partial}{\partial \varphi}p - \rho\frac{\partial}{\partial \varphi}\Phi. \quad (\text{A15})$$

We can also rewrite these two equations in their Eulerian form,

$$\frac{\partial}{\partial t}(\rho R u_\varphi) + \nabla \cdot [(\rho R u_\varphi)\mathbf{u}] = S, \quad (\text{A16})$$

and, when we want to follow the fluid on the rotating coordinate grid,

$$\frac{\partial}{\partial t}(\rho R u_\varphi) + \nabla \cdot [(\rho R u_\varphi)\mathbf{u}'] = S. \quad (\text{A17})$$

When comparing equations (A16) and (A17), notice that the conserved quantity is the same – the z -component of the angular momentum measured in the *inertial* frame. The only

difference in the two equations is the “transport” velocity (\mathbf{u} for the nonrotating frame, \mathbf{u}' for the rotating reference frame).

Equation (A17) is different from the more familiar formulation, where the angular momentum density as well as the transport velocity is measured with respect to the rotating frame, i.e., where the angular momentum density is expressed in terms of the azimuthal component of the transport velocity, u'_φ . But, as a consequence, the source term in the more familiar formulation is more complicated. We can derive the more familiar formulation from equation (A17) by recognizing that,

$$u_\varphi = u'_\varphi + R\Omega_0. \quad (\text{A18})$$

So we can write,

$$\frac{\partial}{\partial t}[\rho R(u'_\varphi + R\Omega_0)] + \nabla \cdot \{[\rho R(u'_\varphi + R\Omega_0)]\mathbf{u}'\} = S_{\varphi i}, \quad (\text{A19})$$

where, as shorthand, we have used,

$$S_{\varphi i} \equiv -\frac{\partial}{\partial \varphi} p - \rho \frac{\partial}{\partial \varphi} \Phi. \quad (\text{A20})$$

This implies,

$$\frac{\partial}{\partial t}(\rho R u'_\varphi) + \nabla \cdot [(\rho R u'_\varphi)\mathbf{u}'] = S_{\varphi i} - \frac{\partial}{\partial t}[\rho R(R\Omega_0)] - \nabla \cdot \{[\rho R(R\Omega_0)]\mathbf{u}'\} \quad (\text{A21})$$

$$= S_{\varphi i} - R^2\Omega_0 \left\{ \frac{\partial}{\partial t}\rho + \nabla \cdot (\rho\mathbf{u}') \right\} - \rho\mathbf{u}' \cdot \nabla(R^2\Omega_0) \quad (\text{A22})$$

$$= S_{\varphi i} - 2\rho R u'_R \Omega_0, \quad (\text{A23})$$

where the last step is accomplished by making use of the continuity relation, $\partial\rho/\partial t + \nabla \cdot (\rho\mathbf{u}') = 0$. Notice that all velocities now refer to \mathbf{u}' , the velocity as measured in the rotating frame, which is the more familiar formulation. The appearance of a Coriolis term is the result of choosing to measure angular momentum in the rotating frame rather than in the inertial frame. This corresponds to “Case B ($\eta = 3'$)” in Call, Tohline, & Lehner (2010). In our hybrid scheme we have chosen to use equation (A17) instead of (A23) primarily because equation (A17) presents a simpler source term.

A.5 Summary of Equations

The z -component of the momentum will always be advected with the following equation,

$$\frac{\partial}{\partial t} s_z + \nabla \cdot (s_z \mathbf{u}) = -\frac{\partial}{\partial z} p - \rho \frac{\partial}{\partial z} \Phi, \quad (\text{A24})$$

where,

$$\Phi = -\frac{GM_{\text{pt}}}{(R^2 + Z^2)^{1/2}}. \quad (\text{A25})$$

For momentum quantities in the plane normal to the z -axis, the equations will depend on the coordinate bases chosen to describe the advected quantities and whether the quantities are measured in the rotating frame or inertial frame.

A.5.1 Non-rotating frame

When we advect Cartesian momentum components, we will evolve the following equations,

$$\frac{\partial}{\partial t}s_x + \nabla \cdot (s_x \mathbf{u}) = -\frac{\partial}{\partial x}p - \rho \frac{\partial}{\partial x}\Phi, \quad (\text{A26})$$

$$\frac{\partial}{\partial t}s_y + \nabla \cdot (s_y \mathbf{u}) = -\frac{\partial}{\partial y}p - \rho \frac{\partial}{\partial y}\Phi. \quad (\text{A27})$$

When we advect cylindrical momentum components, we evolve,

$$\frac{\partial}{\partial t}s_R + \nabla \cdot (s_R \mathbf{u}) = -\frac{\partial}{\partial R}p - \rho \frac{\partial}{\partial R}\Phi + \frac{\ell_z^2}{\rho R^3}, \quad (\text{A28})$$

$$\frac{\partial}{\partial t}\ell_z + \nabla \cdot (\ell_z \mathbf{u}) = -\frac{\partial}{\partial \varphi}p - \rho \frac{\partial}{\partial \varphi}\Phi. \quad (\text{A29})$$

Notice that in the Cartesian equations, A26 and A27, there are no source terms other than pressure and gravity. The $\ell_z^2/\rho R^3$ term in equation A28 comes from the transformation from the Lagrangian form of the equation to the Eulerian form in cylindrical coordinates. This corresponds to Case B, $\eta = 2, 3$ from Table 2 in Call, Tohline, & Lehner (2010).

A.5.2 Rotating Frame, Rotating Frame Quantities

This is the more familiar formalism for hydrodynamics codes that use a rotating frame of reference. Advecting momentum quantities in a frame rotating with constant $\Omega = \Omega_0$, where the momentum quantities are measured with respect to the rotating frame. These equations are found in Norman & Wilson (1978) and New & Tohline (1997), and correspond to Case C, $\eta = 2, 3'$ from Table 2 in Call, Tohline, & Lehner (2010).

In terms of Cartesian momentum components and derivatives,

$$\frac{\partial}{\partial t}s'_x + \nabla \cdot (s'_x \mathbf{u}') = -\frac{\partial}{\partial x}p - \rho \frac{\partial}{\partial x}\Phi + 2\Omega_0 s'_y + \rho \Omega_0^2 x, \quad (\text{A30})$$

$$\frac{\partial}{\partial t}s'_y + \nabla \cdot (s'_y \mathbf{u}') = -\frac{\partial}{\partial y}p - \rho \frac{\partial}{\partial y}\Phi - 2\Omega_0 s'_x + \rho \Omega_0^2 y. \quad (\text{A31})$$

Rewritten in terms of cylindrical momentum components and derivatives,

$$\frac{\partial}{\partial t}s_R + \nabla \cdot (s_R \mathbf{u}') = -\frac{\partial}{\partial R}p - \rho \frac{\partial}{\partial R}\Phi + \frac{\ell_z'^2}{\rho R^3} + \rho \Omega_0^2 R + \frac{2\Omega_0 \ell_z'}{R}, \quad (\text{A32})$$

$$\frac{\partial}{\partial t}\ell'_z + \nabla \cdot (\ell'_z \mathbf{u}') = -\frac{\partial}{\partial \varphi}p - \rho \frac{\partial}{\partial \varphi}\Phi - 2\Omega_0 R s_R, \quad (\text{A33})$$

where $\mathbf{u}' = \mathbf{u} - \mathbf{u}_{\text{frame}}$, $\mathbf{u}_{\text{frame}} = \hat{\mathbf{e}}_\varphi(R\Omega_0)$, and $\ell'_z = \ell_z - \rho R^2\Omega_0$.

A.5.3 Rotating Frame, Inertial Frame Quantities

When advecting inertial frame quantities on a rotating grid, we see that the source terms are simplified from the previous case (advecting rotating frame quantities on a rotating grid). These correspond to Case C, $\eta = 2, 3$ from Table 2 in Call, Tohline, & Lehner (2010).

In terms of Cartesian momentum components and derivatives,

$$\frac{\partial}{\partial t}s_x + \nabla \cdot (s_x \mathbf{u}') = -\frac{\partial}{\partial x}p - \rho \frac{\partial}{\partial x}\Phi + \Omega_0 s_y, \quad (\text{A34})$$

$$\frac{\partial}{\partial t}s_y + \nabla \cdot (s_y \mathbf{u}') = -\frac{\partial}{\partial y}p - \rho \frac{\partial}{\partial y}\Phi - \Omega_0 s_x. \quad (\text{A35})$$

Notice that the source terms in equations A34 and A35 are 1/2 of the more familiar Coriolis source terms. This source term is a consequence of the fact that the coordinate bases that the momenta are measured in change in time compared to the bases that define the coordinate grid. Rewritten in terms of cylindrical momentum components and derivatives,

$$\frac{\partial}{\partial t}s_R + \nabla \cdot (s_R \mathbf{u}') = -\frac{\partial}{\partial R}p - \rho \frac{\partial}{\partial R}\Phi + \frac{\ell_z^2}{\rho R^3}, \quad (\text{A36})$$

$$\frac{\partial}{\partial t}\ell_z + \nabla \cdot (\ell_z \mathbf{u}') = -\frac{\partial}{\partial \varphi}p - \rho \frac{\partial}{\partial \varphi}\Phi. \quad (\text{A37})$$

We also choose to rewrite the previous two equations in terms of Cartesian derivatives, since in practice the code looks at these differences on a Cartesian grid,

$$\frac{\partial}{\partial t}s_R + \nabla \cdot (s_R \mathbf{u}') = -\frac{1}{R} \left(x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \right) p - \frac{\rho}{R} \left(x \frac{\partial}{\partial x} + y \frac{\partial}{\partial y} \right) \Phi + \frac{\ell_z^2}{\rho R^3}, \quad (\text{A38})$$

$$\frac{\partial}{\partial t}\ell_z + \nabla \cdot (\ell_z \mathbf{u}') = \left(y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y} \right) p + \rho \left(y \frac{\partial}{\partial x} - x \frac{\partial}{\partial y} \right) \Phi. \quad (\text{A39})$$

In this work, we use equations A34 and A35 when performing simulations advecting Cartesian momentum, and equations A38 and A39 when performing simulations advecting cylindrical momentum.

Vita

Zachary Duncan Byerly was born in Cape Town, South Africa, to Maud Walsh and Gary Byerly, in 1984 and has lived in Baton Rouge most of his life. He attended East Baton Rouge Parish public schools, graduating from McKinley High School in 2002. He completed his Bachelor of Science in physics at Louisiana State University and Agricultural and Mechanical College in May of 2007, and his Masters of Science in physics in August of 2011.