

Dynamic fisheye grids for binary black hole simulations

This content has been downloaded from IOPscience. Please scroll down to see the full text.

2014 Class. Quantum Grav. 31 065013

(<http://iopscience.iop.org/0264-9381/31/6/065013>)

View [the table of contents for this issue](#), or go to the [journal homepage](#) for more

Download details:

This content was downloaded by: scottcnoble

IP Address: 129.21.187.2

This content was downloaded on 18/03/2014 at 19:56

Please note that [terms and conditions apply](#).

Dynamic fisheye grids for binary black hole simulations

Miguel Zilhão and Scott C Noble

Center for Computational Relativity and Gravitation and School of Mathematical Sciences, Rochester Institute of Technology, Rochester, NY 14623, USA

E-mail: mzilhao@astro.rit.edu and scn@astro.rit.edu

Received 17 September 2013, revised 4 February 2014

Accepted for publication 5 February 2014

Published 27 February 2014

Abstract

We present a new warped gridding scheme adapted to simulating gas dynamics in binary black hole spacetimes. The grid concentrates grid points in the vicinity of each black hole to resolve the smaller scale structures there, and rarefies grid points away from each black hole to keep the overall problem size at a practical level. In this respect, our system can be thought of as a ‘double’ version of the fisheye coordinate system, used before in numerical relativity codes for evolving binary black holes. The gridding scheme is constructed as a mapping between a uniform coordinate system—in which the equations of motion are solved—to the distorted system representing the spatial locations of our grid points. Since we are motivated to eventually use this system for circumbinary disc calculations, we demonstrate how the distorted system can be constructed to asymptote to the typical spherical polar coordinate system, amenable to efficiently simulating orbiting gas flows about central objects with little numerical diffusion. We discuss its implementation in the HARM3D code, tailored to evolve the magnetohydrodynamics equations in curved spacetimes. We evaluate the performance of the system’s implementation in HARM3D with a series of tests, such as the advected magnetic field loop test, magnetized Bondi accretion, and evolutions of hydrodynamic discs about a single black hole and about a binary black hole. Like we have done with HARM3D, this gridding scheme can be implemented in other unigrid codes as a (possibly) simpler alternative to adaptive mesh refinement.

Keywords: general relativity, magnetohydrodynamics, black holes, fisheye coordinates

PACS numbers: 04.25.dg, 95.30.Qd, 04.25.D–, 98.62.Js

(Some figures may appear in colour only in the online journal)

1. Introduction

When dealing with numerical simulations with a large range in length scales, it is often challenging to find the right balance between adequate resolution to accurately represent the relevant parts of the physical system and computational expense. Such a situation is typically encountered when evolving systems of compact objects (e.g. binaries involving neutron stars and black holes) in general relativity (GR). In these cases, the stellar object's extent must be resolved by at most tenths of radii while the global spatial extent often must extend to $O(10^3)$ stellar radii or further in order to evolve other important aspects of the problem, e.g., gravitational radiation, an extended circumbinary disc.

The large range in spatial scales often leads those in the field of general relativistic magnetohydrodynamics (GRMHD) to resort to adaptive mesh refinement (AMR) techniques, where the spacing between grid points changes depending on the desired accuracy of the solution in a specific region. For instance, block structured AMR [1, 2] is used to resolve compact objects with a hierarchy of nested blocks of cells of ever decreasing size, where the grid cell size covering the stellar objects can be orders of magnitude smaller than the extent of the simulation domain [3–6]. Other codes refine locally as needed [7], while others use multiple coordinate patches to cover a domain in a way beneficial for the computation [8–12]. Though not adapted to GR yet, unstructured moving mesh codes show promise for automatic, local mesh refinement that achieves relatively low advection errors by moving cells at their local fluid element's velocity [13–15].

Each of these methods has its advantages and drawbacks. For instance, block structured AMR codes often employ Cartesian grids that lead to excessive loss of angular momentum and diffusion for nearly spherical flows. They also do not scale very well due to challenges in balancing the computational load of managing the often large grid hierarchy data structure [5, 16]. Multi-patch systems preserve angular momentum well [9], yet still require one to track the moving compact objects with a block mesh hierarchy, thereby sharing the scaling drawbacks of the block AMR codes. On the other hand, moving mesh codes seem to run at rates of several factors slower than those of their uniform static grid counterparts [14].

The main motivation for this work comes from our ongoing efforts toward modeling accretion flows around binary black holes [17]. In order to provide the scientific community with adequate predictions of electromagnetic signatures of supermassive binary black holes, accurate dynamic simulations of their gaseous environment must be made. Since the circumbinary gas is expected to be hot enough to be mostly ionized, GRMHD methods are essential for describing the role MHD physics plays in the accretion process. The magnetorotational instability [18, 19] is expected to be a principal driver of angular momentum transport and it—in part—leads to the turbulence that ultimately dissipates the free energy of orbital motion into radiation [20, 21]. In the near field regime when the black holes are separated by less than $100 GM/c^2$ (M is the total mass of the binary), several calculations have been made in the past few years that have begun to provide early descriptions of these events [22–26]. One aspect that remains lacking in these calculations is a satisfactory description of the conditions of the binary's environment prior to merger. Our research focuses on arriving at more realistic initial conditions by evolving from larger separations and for longer periods of time. We must then employ a GRMHD code that can accurately describe the turbulent gas in the disc and the material that falls to the black holes.

Unfortunately, all the codes used by other groups to explore this problem have used the previously mentioned block-structured AMR techniques in Cartesian coordinates. These methods typically lead to poor conservation of fluid angular momentum and excessive

dissipation at refinement boundaries. These two effects alter the disc's angular momentum evolution and thermodynamics in non-trivial ways.

We present here an alternative solution, which uses a predetermined coordinate system to concentrate more grid points in the spatial regions that require them with the GRMHD code called HARM3D [27]. The methodology we employ is similar to that employed in other GRMHD codes for fixed spacetimes (e.g., [27–29]), but our system can be time-dependent and track moving features in the solution. For the case of following two orbiting black holes, our system closely resembles a set of two ‘fisheye’ transformations, one about each black hole. Fisheye coordinates employ a transition function (e.g., a numerical approximation of the Heaviside function) to connect a region of high resolution to a region of low resolution, so that computational effort is focused on the region-of-interest [30–32]. Single, static fisheye coordinates have been used extensively in numerical relativity¹. For instance, they have been used for the vacuum two-body problem [35–40], and a series of nested fisheye coordinates have been used for GRMHD simulations of core collapse [41]. Besides our system being dynamic, another key distinction is that we can focus points about each black hole rather than have one focal region encompassing both. This ability increases our computational efficiency by reducing the number of points in the region between the two black holes, a volume that does not require as much resolution for our problem. Further, the system asymptotes to spherical coordinates away from the black holes, which is helpful for conforming more closely to the geometry of the circumbinary disc thereby providing better resolution of the disc's angular momentum and MHD turbulence [42].

Even though HARM3D lacks AMR (and, as mentioned above, AMR may not be best route), it has always supported arbitrary, time-dependent coordinate distortions. The coordinate transformation we present here thus concentrates grid cells in the vicinity of the black holes and on the plane of the disc. It also rarefies the number of cells as we move away from the binary so that the overall number of cells remains at a manageable level. We again emphasize that, even though all our discussion will happen in the context of our implementation of this coordinate system in the HARM3D code, there is nothing in its construction that relies on HARM3D, or on MHD evolutions. The coordinate transformation we will present can in principle be successfully implemented in any (sufficiently modular) unigrid code and could be a simpler alternative to AMR.

The paper is presented in the following way. In section 2 we cover some theoretical background, including some details about the MHD evolution procedure. In section 3 we introduce our principal ideas by describing the construction of a distorted coordinate system adapted to two-dimensional (2D) periodic Cartesian grids; we also describe how well this coordinate system performs with the advected magnetic field loop test. In section 4 we state our warped spherical coordinate transformation, and show examples of possible grid configurations. Results obtained with the ultimate spherical construction are illustrated in section 5, where we demonstrate its performance in evolving magnetized Bondi accretion, and hydrodynamic discs about single and binary black hole systems. We end with some final remarks in section 6.

2. Theoretical background

Before we begin our description of the coordinate system itself, let us first discuss how coordinates are often used when solving equations of motion (EOM) in the theory of GR, to provide a context. Given two sets of coordinate systems, x^μ and $x^{\mu'}$, that are related by a

¹ Fisheye coordinates are also used for visualizing streams of textual information [33, 34].

differential map, i.e. $x^\mu = x^\mu(x^{\mu'})$ and $x^{\mu'} = x^{\mu'}(x^\mu)$ are both continuous and differentiable, general covariance of GR implies that

$$\nabla_\mu T^{\mu\nu} = 0 = \nabla_{\mu'} T^{\mu'\nu'}, \quad (1)$$

where ∇_μ and $\nabla_{\mu'}$ each represent the differential operator associated with the metric in their own system:

$$\nabla_\mu g_{\nu\lambda} = 0 = \nabla_{\mu'} g_{\nu'\lambda'}, \quad (2)$$

and $(T_{\mu\nu}, g_{\mu\nu})$ are related to $(T_{\mu'\nu'}, g_{\mu'\nu'})$ via the well-known tensor transformation equation, which we state here for completeness:

$$A_{\mu_1 \dots \mu_{n_1}}{}^{\nu_1 \dots \nu_{n_2}} = A_{\mu'_1 \dots \mu'_{n_1}}{}^{\nu'_1 \dots \nu'_{n_2}} \prod_{i=1}^{n_1} \prod_{j=1}^{n_2} \frac{\partial x^{\mu'_i}}{\partial x^{\mu_i}} \frac{\partial x^{\nu_j}}{\partial x^{\nu'_j}}. \quad (3)$$

Greek indices represent spacetime coordinate indices.

The EOM are often solved numerically by discretizing the continuum domain into points and solving the equations for the fields at these points. One therefore must have a map from the computer memory's space to coordinate space of one's problem. Even for the simple example of uniform Cartesian coordinates, a diffeomorphism is used to connect the memory's index space, $x^{\mu'} = i^{\mu'} \in \mathbb{Z}$, to x^μ : $x^\mu = x_0^\mu + x^{\mu'} dx^\mu$ with the grid spacing being $\Delta x^\mu = \frac{\partial x^\mu}{\partial x^{\mu'}} \Delta x^{\mu'} = dx^\mu$ since $x^{\mu'} \in \mathbb{Z}$.² Often this map is constructed in a nonlinear way so that more grid points cover areas of interest. Sometimes, one chooses to discretize w.r.t. the non-uniformly spaced x^μ and solve for $(g_{\mu\nu}, T_{\mu\nu})$ (e.g., [29]), or one chooses to discretize w.r.t. the uniform $x^{\mu'}$ system and solve for the transformed quantities of interest $(g_{\mu'\nu'}, T_{\mu'\nu'})$ [27, 32]. The complication introduced by the former method is that finite differences must be performed in a way valid for non-uniform grids, while the latter scheme often requires one to either transform the Christoffel symbols for the EOM's source terms (or extrinsic curvature in the case of Einstein's equations), or finite difference the transformed metric $g_{\mu'\nu'}$ for them. In passing, we note that the latter method is used in HARM3D and we choose to finite difference the metric for the Christoffel symbols, primarily because deriving closed form expressions for the Christoffel symbols (transformed or otherwise) from our approximate binary black hole metric is impractical [17, 43]; this step, therefore, costs no additional overhead for our scheme.

2.1. MHD evolution

All our tests will involve the evolution of magnetized or unmagnetized matter, so we must introduce the EOM used by HARM3D. The equations will also illustrate how the coordinates are relevant to the EOM we will be using.

We assume that the gas does not self-gravitate and alter the spacetime dynamics, so we need only solve the GRMHD equations on a specified background spacetime, $g_{\mu\nu}(x^\lambda)$. As mentioned earlier, HARM3D solves the EOM in uniform $x^{\lambda'}$ coordinates so please interpret the following tensor indices to be 'primed.'

The EOM originate from the local conservation of baryon number density, the local conservation of stress–energy, and the induction equations from Maxwell's equations (please see [27] for more details). They take the form of a set of conservation laws:

$$\partial_t \mathbf{U}(\mathbf{P}) = -\partial_i \mathbf{F}^i(\mathbf{P}) + \mathbf{S}(\mathbf{P}) \quad (4)$$

² In the code, we promote $\{x^{\mu'}\}$ to double precision floating point numbers because we need to represent intermediate coordinates, e.g., $x^{\mu'} + \Delta x^{\mu'}/2$.

where \mathbf{U} is a vector of ‘conserved’ variables, \mathbf{F}^i are the fluxes, and \mathbf{S} is a vector of source terms. Explicitly, these are

$$\mathbf{U}(\mathbf{P}) = \sqrt{-g} [\rho u^t, T^t_t + \rho u^t, T^t_j, B^k]^T \quad (5)$$

$$\mathbf{F}^i(\mathbf{P}) = \sqrt{-g} [\rho u^i, T^i_t + \rho u^i, T^i_j, (b^i u^k - b^k u^i)]^T \quad (6)$$

$$\mathbf{S}(\mathbf{P}) = \sqrt{-g} [0, T^\kappa_\lambda \Gamma^\lambda_{t\kappa}, T^\kappa_\lambda \Gamma^\lambda_{j\kappa}, 0]^T \quad (7)$$

where g is the determinant of the metric, $\Gamma^\lambda_{\mu\kappa}$ are the Christoffel symbols, $B^\mu = \tilde{F}^{\mu t} / \sqrt{4\pi}$ is our magnetic field (proportional to the field measured by observers traveling orthogonal to the spacelike hypersurface), $\tilde{F}^{\mu\nu}$ is the Maxwell tensor, u^μ is the fluid’s 4-velocity, $b^\mu = \frac{1}{u^t} (\delta^\mu_\nu + u^\mu u_\nu) B^\nu$ is the magnetic 4-vector or the magnetic field projected into the fluid’s co-moving frame, and $W = u^t / \sqrt{-g^{tt}}$ is the fluid’s Lorentz function. The MHD stress-energy tensor, $T_{\mu\nu}$, is defined as

$$T_{\mu\nu} = (\rho h + \|b\|^2) u_\mu u_\nu + (p + \|b\|^2/2) g_{\mu\nu} - b_\mu b_\nu \quad (8)$$

where $\|b\|^2 \equiv b^\mu b_\mu$ is the magnetic energy density, p is the gas pressure, ρ is the rest-mass density, $h = 1 + \epsilon + p/\rho$ is the specific enthalpy, and ϵ is the specific internal energy.

If not otherwise noted, we use the following numerical techniques. Piecewise parabolic reconstruction of the primitive variables is performed at each cell interface for calculating the local Lax–Friedrichs flux [28], and a three-dimensional (3D) version of the FluxCT algorithm is used to impose the solenoidal constraint, $\partial_i \sqrt{-g} B^i = 0$ [44]. The electromotive forces (EMFs) are calculated midway along each cell edge using piecewise parabolic interpolation of the fluxes from the induction equation. A second-order accurate Runge–Kutta method is used to integrate the EOM using the method of lines once the numerical fluxes are found. The primitive variables are found from the conserved variables using the ‘2D’ scheme of [45]. Please see [27] and the [appendix](#) for more details.

3. Warped Cartesian coordinates

As mentioned in section 1, our goal is to construct a grid tailored to evolving accretion discs in the background of binary black hole spacetimes. For this purpose, in order to better resolve the accretion disc itself, a spherical coordinate system is desired. As proof of principle and as a first test, however, we will start by presenting a 2D warped coordinate system adapted to a (2D) periodic Cartesian grid between two coordinate systems: $\{T, X, Y\}$ and $\{t, x, y\}$. This warped Cartesian coordinate system also serves a pedagogical purpose, since it is not as involved as the warped spherical construction that will follow.

3.1. Implementation

We fix the following notation. The physical (spatial) coordinates are labeled $\{X, Y\}$, they are assumed to span the physical range $[X_{\min}, X_{\max}]$, $[Y_{\min}, Y_{\max}]$ and can have (in general) a non-uniform grid spacing. The ‘numerical’ coordinates (i.e., the coordinates actually used to evolve the equations in the numerical code) are labeled $\{x, y\}$, are uniformly discretized, and span $[0, 1]$.

Our immediate goal is to develop a coordinate transformation between these two coordinate systems that can concentrate grid cells on certain (predetermined) physical regions. In general, we will be interested in having two ‘focal points,’ where the density of cells is largest.

In the following, we summarize the desired properties of the coordinate transformation:

- (i) $T = t$;
- (ii) $x, y \in [0, 1], X \in [X_{\min}, X_{\max}], Y \in [Y_{\min}, Y_{\max}]$;
- (iii) $\Delta x, \Delta y = \text{constant}$;
- (iv) (X_1, Y_1) and (X_2, Y_2) are the (X, Y) coordinates of the two focal points, assumed to correspond to (x_1, y_1) and (x_2, y_2) in (x, y) coordinates;
- (v) $X_1, X_2, Y_1, Y_2, x_1, x_2, y_1, y_2$ are (generically) functions of time;
- (vi) $\frac{\partial X}{\partial x}$ and $\frac{\partial Y}{\partial y}$ should have local minima along the lines $\{(X_1, Y), (X_2, Y)\}$ and $\{(X, Y_1), (X, Y_2)\}$, respectively;
- (vii) $\frac{\partial X}{\partial x}$ and $\frac{\partial Y}{\partial y}$ should be periodic in x and y , respectively: $\frac{\partial X}{\partial x}(x) = \frac{\partial X}{\partial x}(x + 1), \frac{\partial Y}{\partial y}(y) = \frac{\partial Y}{\partial y}(y + 1)$;
- (viii) $\frac{\partial X}{\partial x} > 0, \frac{\partial Y}{\partial y} > 0$.

Property (vi) guarantees that resolution is highest at the focal points. Property (viii) ensures there are no coordinate singularities. The periodicity condition, (vii), makes our tests easier and will be used in the spherical case later on since the azimuthal coordinate has a similar symmetry.

The coordinate transformation satisfying these properties is most easily described using the following idealized functionals:

$$\sigma(z) = \begin{cases} -1 & : z < 0 \\ 1 & : z \geq 0 \end{cases} \quad (9)$$

$$\Sigma(z) = \int \sigma(z) dz = \begin{cases} -z & : z < 0 \\ z & : z \geq 0 \end{cases} \quad (10)$$

$$\tau(z, z_1, \delta) = \frac{1}{2} [\sigma(z - z_1 + \delta) - \sigma(z - z_1 - \delta)] \quad (11)$$

$$\mathcal{T}(z, z_1, \delta) = \frac{1}{2} [\Sigma(z - z_1 + \delta) - \Sigma(z - z_1 - \delta)] \quad (12)$$

$$\tilde{\tau}(z, z_1, \delta) = \tau(z, \bar{z}_1, \delta) + \tau(z, \bar{z}_1 - 1, \delta) + \tau(z, \bar{z}_1 + 1, \delta) \quad (13)$$

$$\tilde{\mathcal{T}}(z, z_1, \delta) = \mathcal{T}(z, \bar{z}_1, \delta) + \mathcal{T}(z, \bar{z}_1 - 1, \delta) + \mathcal{T}(z, \bar{z}_1 + 1, \delta) \quad (14)$$

where $\bar{z} = z - \text{floor}(z)$. Expression (11) (the ‘boxcar’ function) will be used throughout, and can be seen schematically in figure 1. Tilded expressions are nothing but the periodic equivalents of their non-tilded counterparts.

We found that following system meets our requirements:

$$\frac{1}{X_{\max} - X_{\min}} \frac{\partial X}{\partial x} = 1 - a_{x1} \tilde{\tau}(y, y_1, \delta_{y3}) [\tilde{\tau}(x, x_1, \delta_{x1}) - 2\delta_{x1}] - a_{x2} \tilde{\tau}(y, y_2, \delta_{y4}) [\tilde{\tau}(x, x_2, \delta_{x2}) - 2\delta_{x2}] \quad (15)$$

$$\frac{1}{Y_{\max} - Y_{\min}} \frac{\partial Y}{\partial y} = 1 - a_{y1} \tilde{\tau}(x, x_1, \delta_{x3}) [\tilde{\tau}(y, y_1, \delta_{y1}) - 2\delta_{y1}] - a_{y2} \tilde{\tau}(x, x_1, \delta_{x4}) [\tilde{\tau}(y, y_2, \delta_{y2}) - 2\delta_{y2}] \quad (16)$$

where the $a_{x,y}$ and $\delta_{x,y}$ are tunable parameters that approximately represent the relative amplitude and width of distortion, respectively.

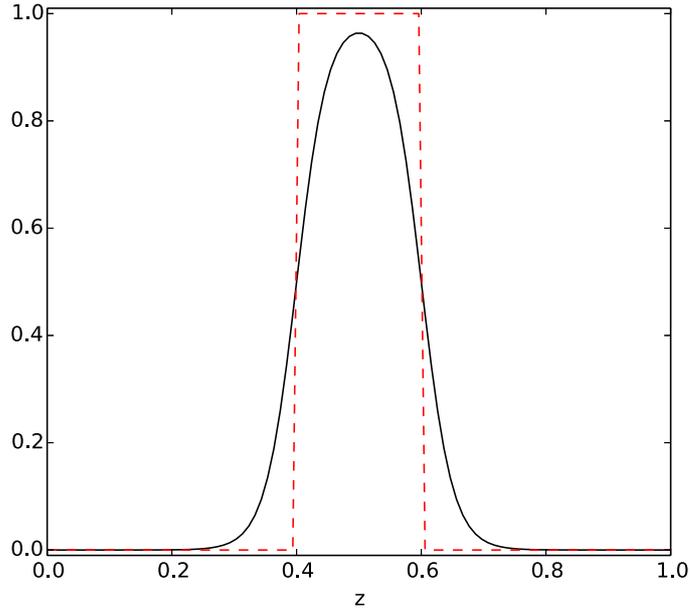


Figure 1. Ideal boxcar function (equation (11)) (dashed red line) and our approximate version (solid black line), with $z_1 = 0.5$, $\delta = 0.1$, $h = 20$.

We here note, though, that in practice we will not be using the idealized expressions (9) and (10). Instead, we will use numerical approximations to the step function, its integral, and its derivative (respectively):

$$\sigma(z, z_1, h) = \tanh(h(z - z_1)) \tag{17}$$

$$\Sigma(z, z_1, h) = \int \sigma(z, z_1, h) dz = \frac{1}{h} \ln \cosh(h(z - z_1)) \tag{18}$$

$$\sigma'(z, z_1, h) = \partial_z \sigma(z, z_1, h) = h \operatorname{sech}^2(h(z - z_1)) \tag{19}$$

where h is a steerable parameter controlling the ‘steepness’ of the function’s transition. The other functions are defined as before without any change. To accommodate the fact that the approximate boxcar function does not have the same integral as the ideal boxcar function, we must normalize it differently in equations (15) and (16) by making the following change: $2\delta_{x1} \rightarrow \tilde{T}(1, x_1, \delta_{x1}) - \tilde{T}(0, x_1, \delta_{x1})$, and analogously for $\delta_{x2}, \delta_{y1}, \delta_{y2}$.

This yields the final expressions for $X(x, y)$ and $Y(x, y)$:

$$\begin{aligned} \frac{X(x, y) - X_{\min}}{X_{\max} - X_{\min}} = & x - a_{x1} \tilde{\tau}(y, y_1, \delta_{y3}) \{ \tilde{T}(x, x_1, \delta_{x1}) - \tilde{T}(x_1, x_1, \delta_{x1}) - (x - x_1)[\tilde{T}(1, x_1, \delta_{x1}) \\ & - \tilde{T}(0, x_1, \delta_{x1})] \} - a_{x2} \tilde{\tau}(y, y_2, \delta_{y4}) \{ \tilde{T}(x, x_2, \delta_{x2}) - \tilde{T}(x_2, x_2, \delta_{x2}) \\ & - (x - x_2)[\tilde{T}(1, x_2, \delta_{x2}) - \tilde{T}(0, x_2, \delta_{x2})] \} \end{aligned} \tag{20}$$

$$\begin{aligned} \frac{Y(x, y) - Y_{\min}}{Y_{\max} - Y_{\min}} = & y - a_{y1} \tilde{\tau}(x, x_1, \delta_{x3}) \{ \tilde{T}(y, y_1, \delta_{y1}) - \tilde{T}(y_1, y_1, \delta_{y1}) - (y - y_1)[\tilde{T}(1, y_1, \delta_{y1}) \\ & - \tilde{T}(0, y_1, \delta_{y1})] \} - a_{y2} \tilde{\tau}(x, x_2, \delta_{x4}) \{ \tilde{T}(y, y_2, \delta_{y2}) - \tilde{T}(y_2, y_2, \delta_{y2}) \\ & - (y - y_2)[\tilde{T}(1, y_2, \delta_{y2}) - \tilde{T}(0, y_2, \delta_{y2})] \}. \end{aligned} \tag{21}$$

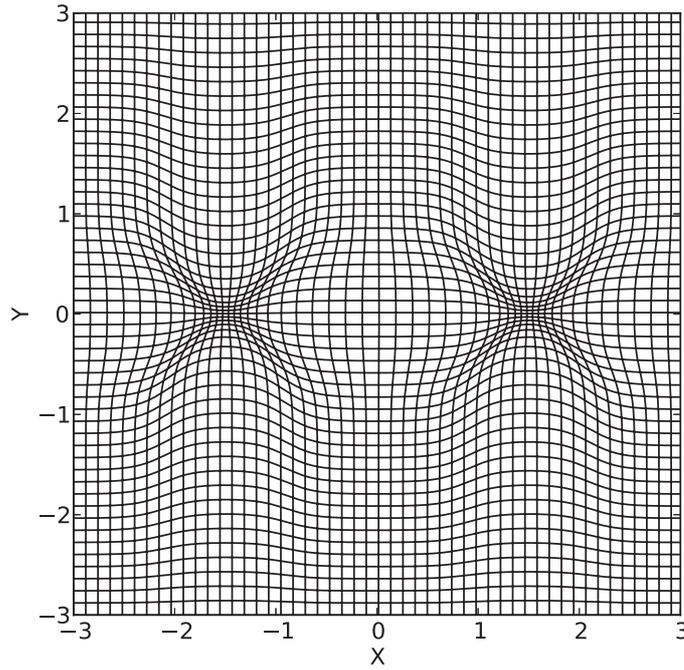


Figure 2. Example grid with $\delta_{x1} = \delta_{x2} = \delta_{x3} = \delta_{x4} = \delta_{y1} = \delta_{y2} = \delta_{y3} = \delta_{y4} = 0.1$, $h_{x1} = h_{x2} = h_{y1} = h_{y2} = 20$, $a_{x1} = a_{x2} = a_{y1} = a_{y2} = 1$, $x_1 = 0.25$, $x_2 = 0.75$, $y_1 = y_2 = 0.5$, $X_{\max} = -X_{\min} = 3$, $Y_{\max} = -Y_{\min} = 3$.

With this coordinate system we have in general two focal points, (x_1, y_1) and (x_2, y_2) , that have a larger density of grid cells in their vicinity. Parameters $a_{x,y}$ and $\delta_{x,y}$ control the transition to the outer region. As stated before, the location of these focal points can vary in any prescribed continuous fashion, to better accommodate the dynamics of the physical system. For the sake of clarity, we have dropped the inclusion of the h parameters in the τ functions in the above expressions. In general, there is an h parameter for every δ parameter, i.e. we have h_{x1} , h_{x2} , h_{y1} , and h_{y2} . The plethora of parameters illustrates the great control one has in fine-tuning a system of coordinates. We show in figure 2 an example with two focal points with symmetric distortions.

3.2. Advected field loop

As a first test of the performance of these coordinates—and validation of its implementation in HARM3D—we present results from the so-called ‘advected field loop’ test, wherein one starts with a circular loop of magnetic field at the center and constant gas density, pressure, and velocity throughout [4, 46–49]. Even though the exact solution is trivial (i.e. no evolution in the field except for the configuration’s translation in space), the test is difficult for constrained transport (CT) schemes to evolve without the field significantly diffusing and deforming. Since the focal point moves through the domain, the coordinates change in time. One could imagine that the moving distortion could lead to additional diffusion or deformation of the field loop. Our aim here is to compare our results from using the distorted coordinates to what results from using uniform coordinates, all with HARM3D.

We specify the initial conditions as follows. Working in the magnetic loop's rest frame, we specify the vector potential to be $A_\mu = (0, 0, 0, A_z)$, where

$$A_z = \begin{cases} A_{\text{loop}} (R_{\text{loop}} - r); & r \leq R_{\text{loop}} \\ 0; & r > R_{\text{loop}} \end{cases}. \quad (22)$$

We next perform a coordinate transformation on A_μ to the boosted frame with coordinate velocity $-u^i/u^t$ and then transform to the numerical coordinates $A_{\mu'}$. We finally compute $B^{i'}$ from $A_{\mu'}$ through a finite difference procedure consistent with our CT scheme (equation (A.14)). We choose parameters $A_{\text{loop}} = 10^{-3}$, $\rho = 1.0$, $P = 3.0$ for the initial magnetic field and hydrodynamic configuration. The fluid is given uniform velocity throughout, $u^i/u^t = [1/12, 0, 0]$, and periodic boundary conditions are imposed at each edge.

For this physical scenario, we set up the coordinate system with only one focal point, fix the parameters in a way such that the cells with smaller aspect ratio lie well outside of the loop, and make the local grid spacing inside the loop resemble a uniform one, cf figure 3. We evolve the system for one period, until $t = 72$, and compare the results against the initial configuration at $t = 0$. These results should be compared with those obtained with no warping, which we present in figure 4. We also explore the differences between two sets of numerical schemes used for reconstructing the EMFs at the cell edges for the CT scheme, and primitive variables at the cell interfaces. One uses nearest neighbor (2-point) averaging for the EMFs and piecewise linear reconstruction (monotonized central limiter [50]) for the primitive variables (the so-called 'linear' method); the other uses a piecewise parabolic interpolation method [51] for both the EMFs and primitive variables (the so-called 'parabolic' method). Once reconstructed at the cell edges, the EMFs are used in the same way for the two FluxCT schemes, which makes both CT schemes second-order accurate even though the parabolic reconstruction often yields more accurate results; please see the appendix for more details regarding our FluxCT schemes.

As can be appreciated from these figures, some (inevitable) diffusion happens for all cases. The warped evolutions, however, while developing some asymmetry, seem to diffuse much less overall than their uniform counterparts. We also find that the evolutions using the parabolic method also diffuse less than those using the linear method. The parabolic method seems to result in a higher effective resolution with the same number of cells. On the other hand, the higher order interpolation stencils used by the parabolic method introduce short wavelength structure into the solution. The differences in rates of diffusion between the various methods can be quantitatively confirmed in figure 5, where we integrate the magnetic field's energy density over the initial field loop's area at each time step. Specifically, we show

$$\frac{1}{A_{\text{loop}}^2} \int_{r < R_{\text{loop}}} ||b||^2 \sqrt{-g'} dx^1 dx^2 \quad (23)$$

as a function of time, where the radius r used in the integral's limits is the 2D radius from the loop's instantaneous center (as predicted by its advection velocity). We note that the integral is performed w.r.t. numerical coordinates in both the uniform and warped cases. Diffusion of magnetic field leads to a loss of field within this area. Similar to the effect of the warped coordinates, the parabolic method effectively increases the resolution and diminishes the rate of diffusion outside the loop. The warped/linear combination, however, is even less diffusive than the uniform/parabolic combination. Obviously, the warped/parabolic combination is the least diffusive of all the configurations, resulting in only a 10% loss of magnetic field energy out of the loop after a period. In order to measure the effective resolution of the warped system, we performed a uniform Cartesian run with double the number of cells per dimension (i.e. 768×768 cells), as this is approximately the local resolution of the warped grid within the field loop. We found that the warped system leads to even less diffusion than the high-resolution

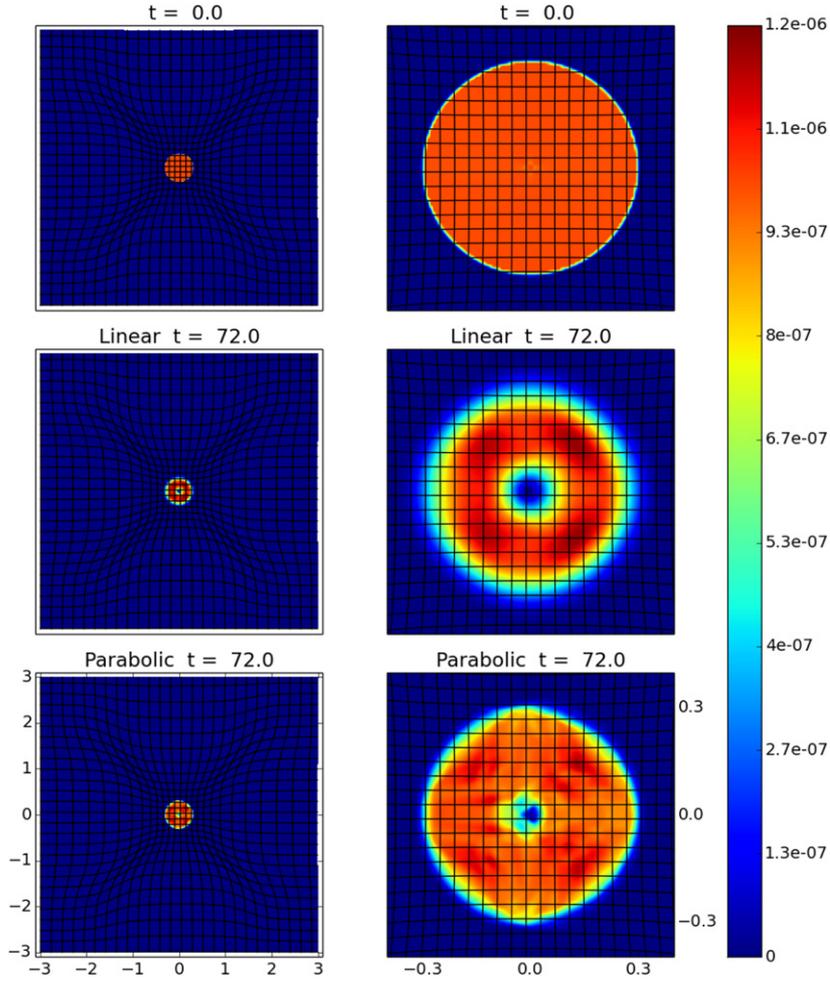


Figure 3. Two-dimensional distributions of magnetic field energy density ($\|b\|^2$) plotted at initial time (top row) and at the final time of the test (bottom two rows). We show results from runs using the linear reconstruction method (middle row), and the parabolic reconstruction method (bottom row). In all runs, we used 384×384 cells. The loop was advected across the grid with a constant velocity along the x -axis, and the ‘warp’ was transported at the same rate. We show both the entire domain with 15 cells per grid line shown (left column), and a view focusing on the field loop with 5 cells per grid line shown (middle column). Grid parameters used for this simulation were $\delta_{x1} = \delta_{x3} = \delta_{y1} = \delta_{y3} = 0.25$, $h_{x1} = h_{y1} = 12$, $a_{x1} = a_{y1} = 1$, $a_{x2} = a_{y2} = 0$, $X_{\max} = -X_{\min} = 3$, $Y_{\max} = -Y_{\min} = 3$. The center of the warp, (x_1, y_1) , is set to coincide with the center of the loop as predicted by the advection equation.

uniform scheme. Comparing the distributions of $\|b\|^2$ between the two least diffusive runs (figures 3 and 4), we see that less magnetic energy density has diffused from the center in the ‘warped parabolic’ run than in the ‘uniform high-resolution parabolic’ run. Even though we tuned the warped coordinates to be nearly uniform within the loop’s region, it is not perfectly uniform: there is $\sim 10\%$ variation in the grid spacing from the center of the loop to the loop’s edge, making the grid spacing at the loop’s center $\sim 5\%$ smaller in the warped run

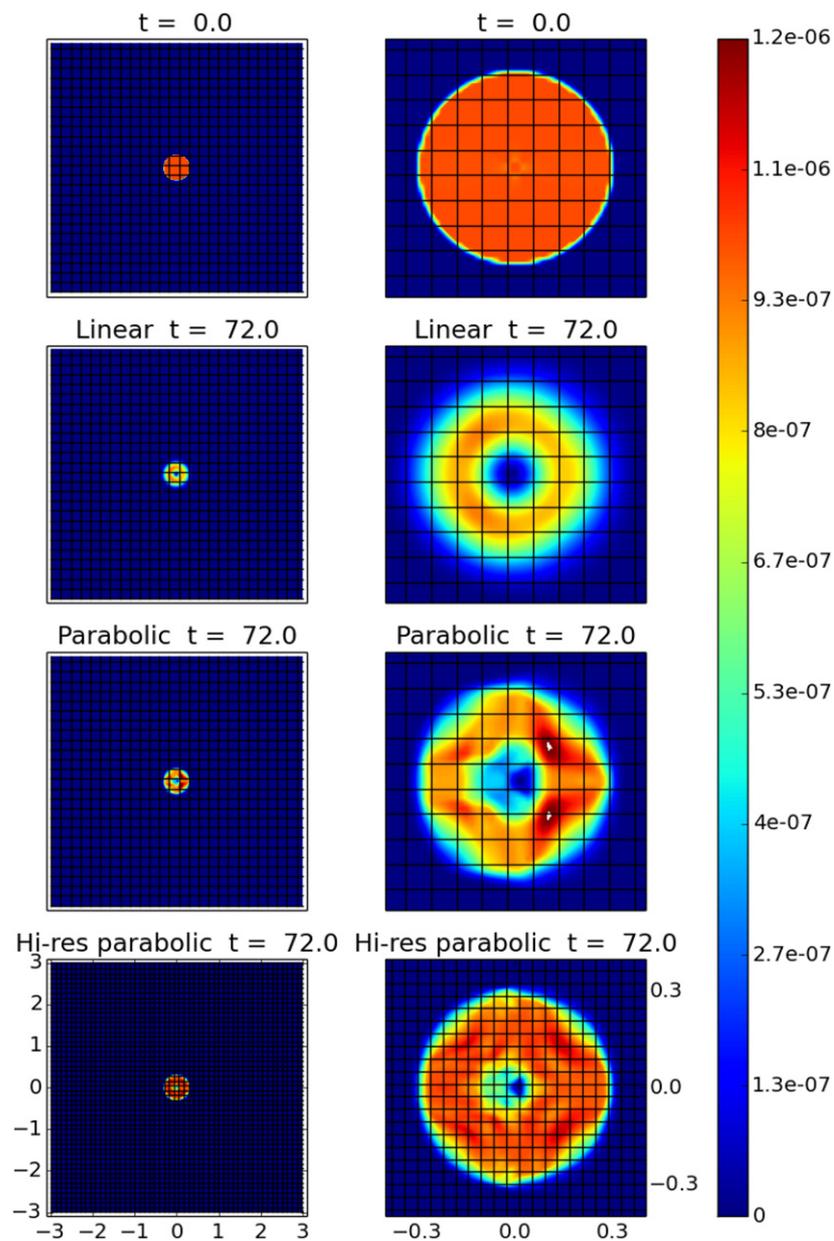


Figure 4. Equivalent runs to those illustrated in figure 3, but using a static uniform grid with 384×384 cells. We show both the entire domain with 15 cells per grid line shown (left column), and a view focusing on the field loop with 5 cells per grid line shown (right column). Bottom row shows a higher-resolution unigrid case (768×768 cells). The grid spacing here was chosen to match the smallest grid spacings found in the warped configurations.

than in the high-resolution uniform run. This slight difference in grid spacing likely explains the differences between the two runs shown in figure 5 since they are of the same relative magnitude (of a few per cent). The fact that the warped run performs at least as well as the

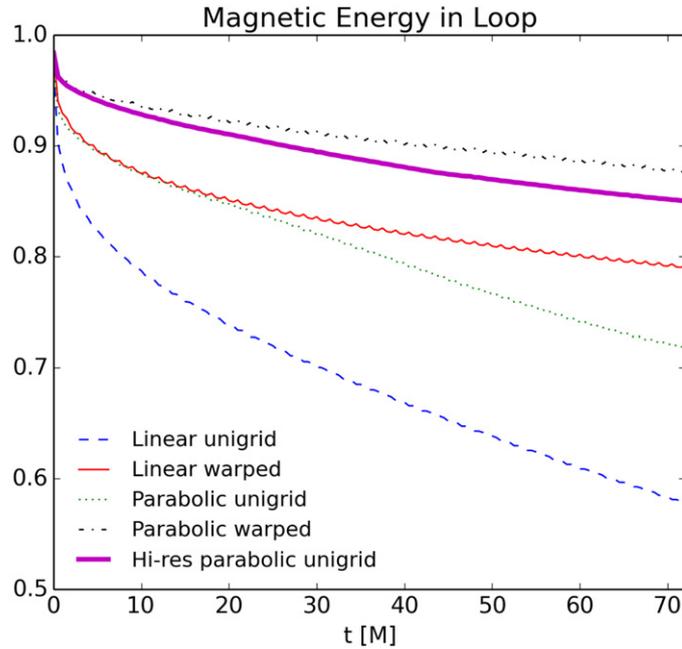


Figure 5. Integral of $\|b\|^2$ over the field loop as function of time (equation (23)) using a variety of run configurations. The standard linear method (solid and dashed) and the parabolic method (dotted and dash-dotted) were used, both with a uniform Cartesian grid (dashed and dotted) and a warped mesh (solid and dash-dotted). Another configuration included using the parabolic method with a higher-resolution uniform grid (768×768 cells), whose grid spacing was chosen to match the smallest grid spacings found in the warped configurations.

high-resolution uniform run verifies that the coordinate system was successfully implemented and is effective at moving a mesh refinement without introducing spurious behavior in the refined region.

We conclude with a note that the parabolic method is the default choice of methods in HARM3D, and it is used in the rest of the paper.

4. Warped spherical coordinates

Having obtained encouraging results with the warped Cartesian implementation, we have designed an analogous coordinate transformation adapted to spherical coordinates. Recall that we are ultimately interested in evolving accretion discs around binary black holes, so our goal is a transformation that accurately resolves the region around the black holes, resolves the small aspect ratio of the disc, and conforms to the near azimuthally-symmetric shape of the disc past the binary's orbit. The full coordinate transformation will be a 3D one, but let us describe the 2D version since the transformation along the third dimension is straightforward. We also outline first the (simpler) transformation where the two focal points occur at the same radial distance from the coordinate origin, and then deal with the case where the focal points can be at different radii (e.g., for non-equal mass binaries). In all cases, the center of mass of the binary will be fixed at the coordinate system's origin.

4.1. Circular azimuthal focusing

In this 2D version of the warped coordinates, constructed to accommodate an equal-mass binary black hole spacetime, we consider only the transformation between $\{r, \phi\}$ and $\{x, y\}$, where $\{r, \phi\}$ are our physical coordinates, with ϕ being the periodic azimuthal coordinate and r the radial coordinate. In the numerical $\{x, y\}$ coordinate system, y will be the radial-like coordinate and x will be the azimuthal-like one; we are therefore interested in keeping the periodicity assumption (Property (vii) from the previous section) along x , but not along y . Thus, for the ϕ coordinate, we can use the same transformation as before (equation (20)) (but without periodicity in y),

$$\begin{aligned} \phi(x, y) = & 2\pi x - 2\pi a_{x1} \tau_{y3}(y, t) [\tilde{T}_{x1}(x, t) - \tilde{T}_{x1}(x_1(t), t) - (x - x_1(t))(\tilde{T}_{x1}(1, t) - \tilde{T}_{x1}(0, t))] \\ & - 2\pi a_{x2} \tau_{y4}(y, t) [\tilde{T}_{x2}(x, t) - \tilde{T}_{x2}(x_2(t), t) - (x - x_2(t))(\tilde{T}_{x2}(1, t) - \tilde{T}_{x2}(0, t))] \end{aligned} \quad (24)$$

where we have introduced the shorthand notations

$$\begin{aligned} \tilde{T}_{xi}(x, t) & \equiv \tilde{T}(x, x_i(t), \delta_{xi}, h_{xi}), \\ \tau_{y3}(y, t) & \equiv \tau(y, y_1(t), \delta_{y3}, h_{y3}), \\ \tau_{y4}(y, t) & \equiv \tau(y, y_2(t), \delta_{y4}, h_{y4}). \end{aligned}$$

For the r coordinate, though, we are interested in a different transformation. We want to focus cells in the vicinity of the black holes, but rarefy them at an exponential rate as one proceeds away from them—similar to what is done in single black hole accretion disc simulations (e.g., [27]). Rarefaction near the origin is motivated by the fact that the origin is not a point of special interest for our problem³, and because there will already be many cells in its vicinity due to the focusing of azimuthal spacing that naturally arises there in spherical coordinates. Radial grid spacing is increased with radius in order to cover a larger radial extent and is justified because characteristic radial scales of the disc's turbulence tend to scale with their radius (at least for discs of constant scale height). We therefore consider the following relation

$$r(y) = R_{\text{in}} + (b_r - sa_r)y + a_r [\sinh(s(y - y_b)) + \sinh(sy_b)], \quad (25)$$

$$a_r \equiv \frac{R_{\text{out}} - R_{\text{in}} - b_r}{\sinh(s(1 - y_b)) + \sinh(sy_b) - s}, \quad (26)$$

where s controls the strength of the transition in resolution near y_b . The meaning of the parameters is more readily gleaned when looking at $\partial r / \partial y$:

$$\frac{\partial r}{\partial y} = b_r + sa_r [\cosh(s(y - y_b)) - 1]. \quad (27)$$

We see immediately that b_r is the minimum of $\frac{\partial r}{\partial y}$ which occurs at $y = y_b$, a_r is the amplitude of the nonlinear term, and that the radial grid spacing increases away from $y = y_b$. The radii R_{in} and R_{out} are defined as the radius of the inner edge of the domain and the radius of the outer edge of the domain, respectively. It is implicit in this transformation that $y_{\text{in}} = 0$ and $y_{\text{out}} = 1$, because of Property (ii) of section 3.1. The parameter y_b is determined by inverting $r(y_b) = r_b$ numerically using (25)⁴ where r_b is the radius of the black holes' orbit. The grid is therefore parameterized by $\{b_r, R_{\text{in}}, r_b, R_{\text{out}}, s\}$. In figure 6 we show an example.

³ As far as the authors can tell from previous work, no small scale phenomena develops at the center of mass of the system [22–25]. This is to be expected as the center of mass is a point of unstable equilibrium: gas elements near it will adiabatically expand and fall toward the nearest black hole if left alone.

⁴ The inversion procedure is performed via a Newton-Raphson scheme.

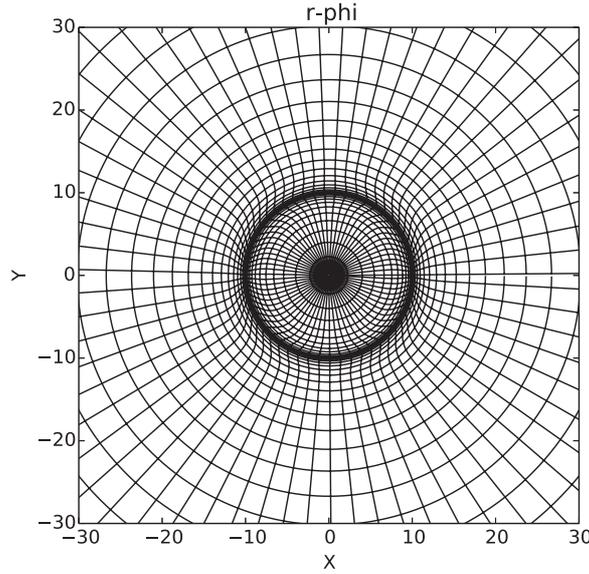


Figure 6. Example of a warped spherical grid, using circular azimuthal warping only. The parameters are $\delta_{x1} = \delta_{x2} = \delta_{y3} = \delta_{y4} = 0.1$, $a_{x1} = a_{x2} = 1$, $h_{x1} = h_{x2} = h_{y3} = h_{y4} = 20$, $s = 10^{-2}$, $b_r = 6.5$, $R_{in} = 0$, $R_{out} = 300$, $x_1 = 0$, $x_2 = 0.5$, $r_b = 10$.

4.2. General warped coordinates

As defined in the previous section, $r(y)$ will have one focal point at $r = r_b$ for all ϕ . When evolving black holes with unequal mass ratios in the center of mass frame, however, the radial coordinates of the black holes will be different. In this section we will generalize the transformation from the previous section to accommodate focusing at different radii as well as introduce the dependence along the third dimension.

In order to focus resolution at different radii, we will smoothly interpolate between different $r(y)$ that use three different sets of $\{b_r, r_b\}$, two for the black holes and one for the intermediate region:

$$r(x, y) = r_3 + (r_1 - r_3) \tilde{\tau}(x, x_1, \delta_{x1}) + (r_2 - r_3) \tilde{\tau}(x, x_2, \delta_{x2}), \quad (28)$$

where $r_{1,2,3} = r_{1,2,3}(y)$ (given by equation (25)), $r_{1,2}$ have values tailored to the black hole at—respectively— $x_{1,2}$. The intermediate profile, $r_3(y)$, is the value of $r(y)$ at the azimuthal points intermediate to the two black holes. Hence, the numerical inversion of $r(y_i) = r_{bi}$ ($i = 1, 2, 3$) needs to be calculated for all x and whenever the black hole positions change. We set $b_{1,2,3}$ (we omit the subscript r in the parameters b from now on to simplify the notation) depending on the resolution requirements. We must set $b_3 > b_{1,2}$ to focus more radial zones at $\phi_{1,2}$. In practice, we see no reason to have R_{in} and R_{out} different between the regions, as this ensures that the inner and outer edges of the domain are at constant radii which better accommodates the outflow boundary conditions imposed there. The circular azimuthal focusing coordinate system described in section 4.1 is a special case of the general warped system when $r_1 = r_2 = r_3$.

The previous set of coordinates lacks dependence on poloidal angle (θ). We introduce this dependence here. Like x and y , we will assume that $z \in [0, 1]$.

Like that used in previous calculations, the poloidal coordinate is refined near the equator. We need not distort it further about each black hole as the resolution in the midplane is already sufficiently fine. Using a construction similar to that already presented

$$\theta(z) = \pi \{z - a_z [\tilde{\mathcal{T}}_{z1}(z) - \tilde{\mathcal{T}}_{z1}(z_1) - (z - z_1)(\tilde{\mathcal{T}}_{z1}(1) - \tilde{\mathcal{T}}_{z1}(0))]\} \quad (29)$$

where we will fix the z coordinate of the focus to be $z_1 = 1/2$ and not a function of time like x_1, x_2, y_1, y_2 (though in principle it could be), since the black holes are always located in the equatorial plane.

Far from the black holes, we want the coordinates to approach uniform spherical coordinates. This is done by multiplying the distortion terms by z -dependent boxcar functions:

$$r(x, y, z, t) = r_3(y, t) + \tilde{\tau}_{z1}(z) \{ [r_1(y, t) - r_3(y, t)] \tilde{\tau}_{x3}(x, t) + [r_2(y, t) - r_3(y, t)] \tilde{\tau}_{x4}(x, t) \}, \quad (30)$$

$$\begin{aligned} \phi(x, y, z, t) = 2\pi \{ &x - \tilde{\tau}_{z1}(z) [a_{x1} \tau_{y3}(y, t) (\tilde{\mathcal{T}}_{x1}(x, t) - \tilde{\mathcal{T}}_{x1}(x_1(t), t) \\ &- (x - x_1(t)) (\tilde{\mathcal{T}}_{x1}(1, t) - \tilde{\mathcal{T}}_{x1}(0, t))) + a_{x2} \tau_{y4}(y, t) (\tilde{\mathcal{T}}_{x2}(x, t) - \tilde{\mathcal{T}}_{x2}(x_2(t), t) \\ &- (x - x_2(t)) (\tilde{\mathcal{T}}_{x2}(1, t) - \tilde{\mathcal{T}}_{x2}(0, t)))] \} \end{aligned} \quad (31)$$

where

$$r_i(y, t) = R_{\text{in}} + b_i y + (R_{\text{out}} - R_{\text{in}} - b_i) \left[\frac{\sinh [s_i(y - y_i(t))] + \sinh [s_i y_i(t)] - s_i y}{\sinh [s_i(1 - y_i(t))] + \sinh [s_i y_i(t)] - s_i} \right] \quad (32)$$

and, as in the previous section,

$$\begin{aligned} \tau_{y3}(y, t) &\equiv \tau(y, y_1(t), \delta_{y3}, h_{y3}) \\ \tau_{y4}(y, t) &\equiv \tau(y, y_2(t), \delta_{y4}, h_{y4}) \\ \tilde{\tau}_{x3}(x, t) &\equiv \tilde{\tau}(x, x_1(t), \delta_{x3}, h_{x3}) \\ \tilde{\tau}_{x4}(x, t) &\equiv \tilde{\tau}(x, x_2(t), \delta_{x4}, h_{x4}) \\ \tilde{\tau}_{z1}(x, t) &\equiv \tilde{\tau}(z, z_1(t), \delta_z, h_z) \\ \tilde{\mathcal{T}}_{x1}(x, t) &\equiv \tilde{\mathcal{T}}(x, x_1(t), \delta_{x1}, h_{x1}) \\ \tilde{\mathcal{T}}_{x2}(x, t) &\equiv \tilde{\mathcal{T}}(x, x_2(t), \delta_{x2}, h_{x2}). \end{aligned}$$

We thus arrive at the final form of our coordinate transformation, comprised of equations (29)–(31). Figures 7 and 8 showcase possible grid configurations. The final grid has a number of free parameters which need to be manually set on a case-by-case basis depending on the physical situation at play. In the following we try to summarize the meaning of these parameters.

- R_{in} and R_{out} are the radial coordinates of the inner and outer edges of the physical domain;
- x_1, x_2, y_1, y_2 are the (x, y) coordinates of the two focal points (black holes). y_3 is the y coordinate of the intermediate region focal point. In general these have to be set at each time-step, and for our specific implementation this is done as follows: let r_1, r_2, ϕ_1, ϕ_2 be the (r, ϕ) coordinates of each black hole at every time-step (which are known); $r_3 = \max(r_1, r_2)$, we set $x_{1,2}$ to be

$$x_{1,2} = \frac{\phi_{1,2} \bmod 2\pi}{2\pi} \quad (33)$$

and we compute $y_{1,2,3}$ via a Newton–Raphson root-finding procedure⁵

$$r_{1,2,3}(y_{1,2,3}) = r_{1,2,3}. \quad (34)$$

z_1 , the z coordinate of the focal points, is kept fixed at $z_1 = 1/2$ (the equatorial plane).

⁵ We note that one would in principle have to invert relations (30) and (31), but we find that (33) and (34) give sufficient accuracy for our purposes.

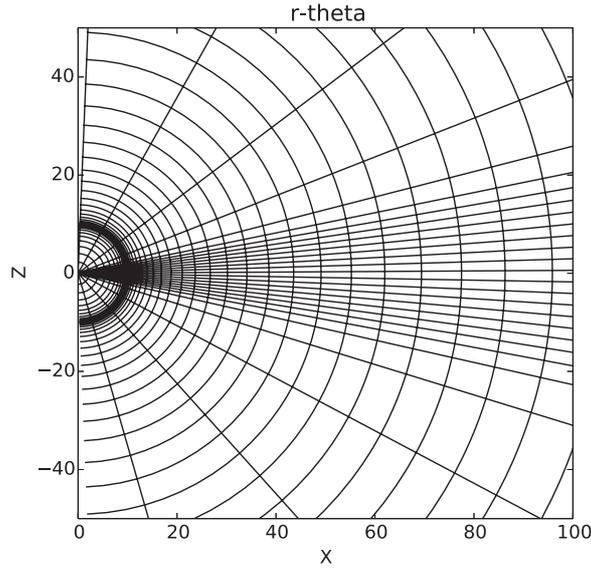


Figure 7. Poloidal slice of a warped spherical grid, with parameters $\delta_{x1} = \delta_{x2} = \delta_{y3} = \delta_{y4} = 0.1$, $\delta_z = 0.4$, $a_{x1} = a_{x2} = 1$, $a_z = 4$, $h_{x1} = h_{x2} = h_{y3} = h_{y4} = h_z = 20$, $s_i = 10^{-2}$, $b_i = 6.5$, $R_{in} = 0$, $R_{out} = 300$, $x_1 = 0$, $x_2 = 0.5$, $r_{b1} = r_{b2} = r_{b3} = 10$. The equatorial slice of these coordinates is identical to that shown in figure 6.

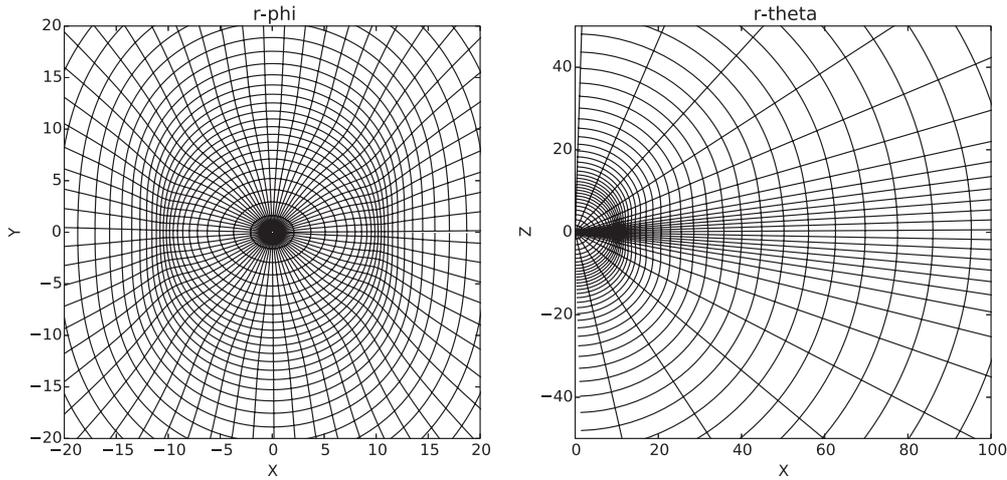


Figure 8. Poloidal and equatorial slices of a different warped spherical grid configuration, with parameters $\delta = 0.1$, $\delta_z = 0.4$, $a_{x1} = a_{x2} = 1$, $a_z = 4$, $h_{x1} = h_{x2} = h_{y3} = h_{y4} = h_z = 10$, $s_{1,2} = 10^{-2}$, $s_3 = 4$, $b_{1,2} = 6.5$, $b_3 = 40$, $R_{in} = 0$, $R_{out} = 300$, $x_1 = 0$, $x_2 = 0.5$, $r_{b1} = 10$, $r_{b2} = 15$, $r_{b3} = 12$.

- δ and h parameters are related to the width and steepness of the approximate (non-ideal) step-functions, which control the extent of and transition rate to the warped regions.
- b_i is the minimum of $\frac{\partial r_i}{\partial y}$ which occurs at $y = y_i$, and the transition steepness is controlled by s_i .
- a_{x1} , a_{x2} , a_z are coupling parameters.

Figure 9 depicts the grid configuration we devised for our evolutions of a binary black hole system. The grid is dynamic, with the location of the focal points tied to the location of the binary.

4.3. Implementation details and computational effort

In terms of computational cost, evaluating $x^\nu(x'^\mu)$ and their derivatives as they are written here for each cell would be computationally intensive and account for a large fraction of the runtime, largely because they involve expensive transcendental functions (e.g., \tanh). One way we have dramatically reduced the computational cost is by separately computing functions by their dependent variables—equations (29)–(31) can each be expressed as a sum of products of functions that each depend only on time and one of the x^i variables. These time-dependent one-dimensional (1D) functions (and their derivatives) are evaluated once per time step along each dimension and stored in 1D arrays, which are used in all the necessary coordinate calculations thereafter. Since our simulations are 3D, several 1D function evaluations amount to a trivial fraction of runtime.

That is not to say that the overhead of the coordinate transformation is insignificant; the effort it takes transforming all tensors to and from our dynamic coordinates—i.e. in the manner of equation (3)—is larger than any other operation related to the dynamic coordinate system and amounts to be about $\simeq 6\text{--}11\%$ of runtime for a typical circumbinary disc evolution. However, these operations are necessary in our simulations anyway because our spacetime metric is defined in Cartesian coordinates and we choose to simulate the gas in spherical coordinates [17]. As a result, our circumbinary disc runs suffer very little additional effort ($\lesssim 1\text{--}2\%$ of runtime) from using dynamic coordinates instead of static coordinates, at least for the binary black hole runs that we will describe later in the paper. Since the computations made for the dynamic coordinates are performed locally on a processor per cell per time step, their contribution to the overall effort does not depend on the number of processors used in the run and should not depend strongly on the computer’s architecture. Our timing measurements were made using the Stampede cluster at the Texas Advanced Computing Center and the BlueSky cluster at RIT.

5. Results

To test our implementation and measure the effect of the warped coordinate system, we have performed a number of tests and, when possible, matched the results against those obtained using ‘standard’ well-tested coordinate systems present in the HARM3D code. We present below results obtained for the evolution of: the magnetized Bondi solution, section 5.1; an accretion disc in the background of a single black hole, section 5.2; a circumbinary accretion disc, section 5.3.

5.1. Bondi flow

The Bondi solution is a solution to the equations describing spherically symmetric, time-independent accretion of non-magnetized gas onto a central, perfectly absorbing object (e.g., a black hole) [52]. The particular solution we use here has $\dot{M} = 4\pi r^2 \rho u^r = 4\pi$, the adiabatic index of the equation of state is $\gamma = 4/3$, and the radial coordinate of the sonic point is $r_s = 8M$. Our line element is the Schwarzschild solution written in Kerr–Schild coordinates. We add a weak, radial magnetic field to the initial conditions in order to validate the magnetic field evolution routines, too. Adding a radial magnetic field to the system does not alter the

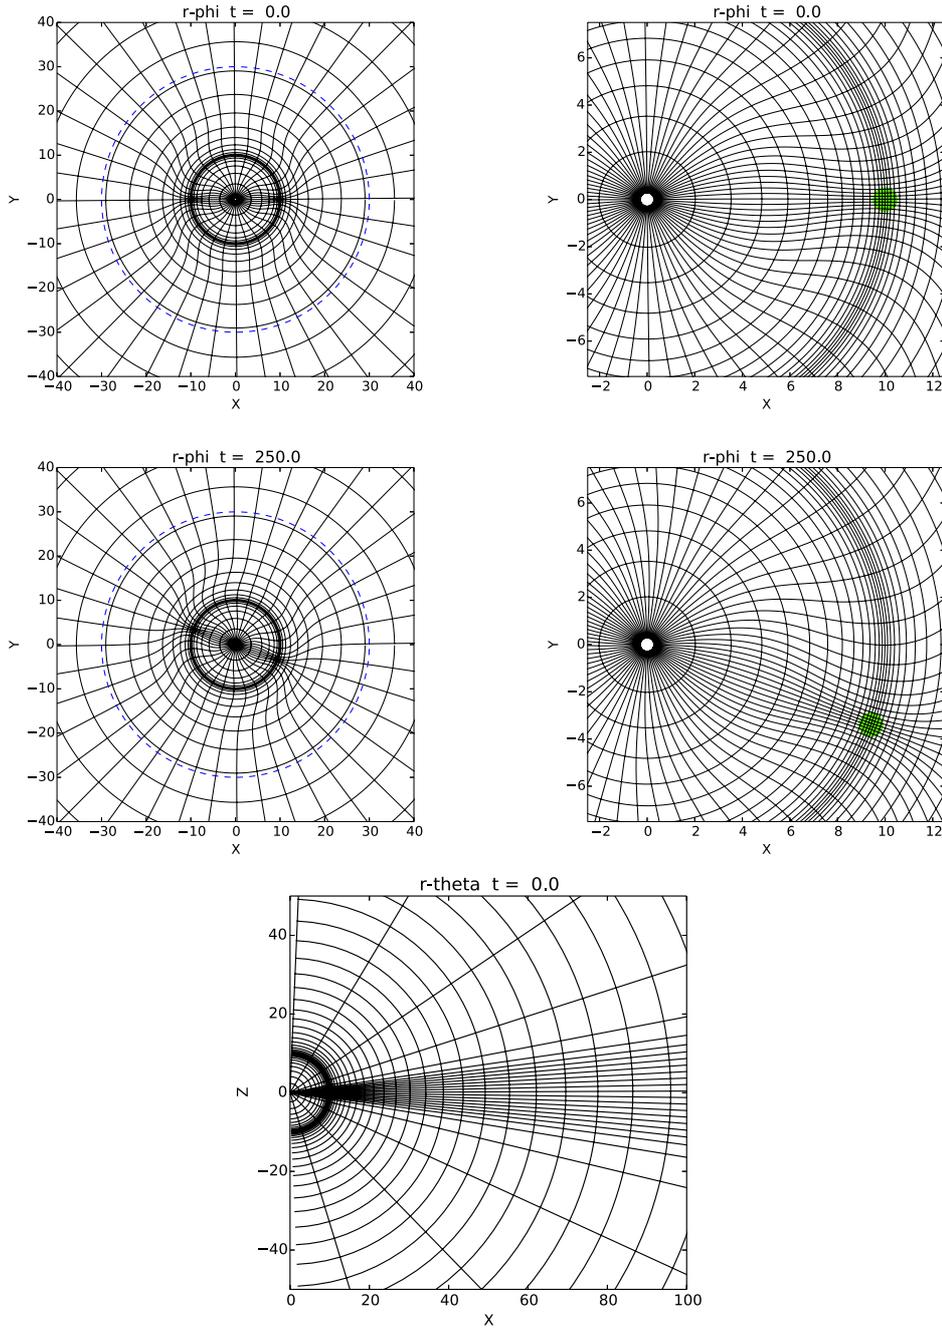


Figure 9. Example grid used for a binary black hole system. The green-shaded area marks the black holes' event horizons, and the dashed blue line marks the location of the accretion disc's inner radius. Grid parameters are as follows: $\delta_{x1} = \delta_{x2} = 0.1$, $\delta_{x3} = \delta_{x4} = \delta_{y3} = \delta_{y4} = 0.18$, $a_{x1} = a_{x2} = 1.5$, $a_z = 4.3$, $h_{x1} = h_{x2} = h_{x3} = h_{x4} = h_{y3} = h_{y4} = h_z = 20$, $s_i = 10^{-2}$, $b_i = 6.5$, $R_{in} = 0$, $R_{out} = 300$, $r_b = 10$; x_1 and x_2 are set at each time-step to coincide with the location of the black holes. Parameters are chosen so that there are approximately 32 cells spanning each black hole horizon in each dimension (black hole horizons span $M \equiv M_{BH1} + M_{BH2} = 1$ in this coordinate system).

solution analytically, but it does test how a code handles terms in the EOM at the truncation error level that do not satisfy the Bondi equations. Relative to evolutions in standard spherical coordinates, we expect these terms to result in a larger effect for our warped coordinates as the system no longer conforms to the symmetry of the problem.

For this test, we use the warped spherical system described in section 4.2 using a similar setup to that illustrated in figure 6. We stress that this grid was *not* tailored for this example; our goal here is not to evolve the Bondi solution as accurately as possible, but rather to measure possible artifacts introduced by the warped and dynamic grid during the evolution. We compare the warped evolution to a run with a coordinate configuration we call ‘unwarped,’ which is the same as the warped coordinate setup except we move the focal points to the origin by setting $r_{b1} = r_{b2} = r_{b3} = 0$ and $a_{x1} = a_{x2} = 0$. The unwarped setup is static, azimuthally symmetric, and is similar to the grids used in single black hole accretion disc studies where $\Delta r/r$ is constant [27].

We follow the procedure outlined in [28] to construct the initial data. We have performed pure hydrodynamic as well as magnetized evolutions. For the magnetized cases, we prescribe a vector potential with the form

$$A_\mu = (0, 0, 0, -A_0 \cos \theta), \quad (35)$$

which guarantees a radial divergence-less magnetic field with

$$||b||^2 = \frac{A_0^2}{r^4}. \quad (36)$$

The problem was integrated until $t = 50M$ in a 2D equatorial domain, with $r \in [1.7, 100]M$, 256×256 cells and a time step of $dt = 0.01M$. We used a warped spherical grid with all the parameters set the same as in figure 6, except here we use $R_{\text{in}} = 1.7$ and $R_{\text{out}} = 100$ instead. We fixed the orbital angular velocity of the (fictitious) black holes—or the grid’s focal points—to be $\omega = 2\pi/50$. The orbital frequency of the warp used here is more than ten times as large as what would be used for a binary evolution since the binary’s orbital period is $P_{\text{orb}} \simeq 561M$ at this separation (20M). Since Bondi is a steady-state solution, we (ideally) expect to see no evolution in our primitive variables. Since the grid *does* evolve, however, in order to quantitatively measure deviations from a variable’s initial state *at the same physical points*, we measure such deviations after exactly one orbital period:

$$\frac{\Delta \rho}{\rho} = \left| \frac{\rho(t = 50M) - \rho(t = 0)}{\rho(t = 0)} \right|. \quad (37)$$

For our specified angular velocity of $\omega = 2\pi/50$, this guarantees that we are performing the comparison at the same physical points. Figure 10 shows a contour plot of this quantity; inspecting the figure, we see relative differences of the order of 10% close to the black hole horizon for the pure hydrodynamic case, and values close to 50% for the magnetized case, which may at first seem worrisome. These rather large values can be explained by the fact that, as we have already stressed, the grid we used for these evolutions was tailored for binary black hole evolutions, *not* for single black hole ones. Thus, by construction, the radial grid spacing increases as one approaches the black hole horizon (as can be observed in figure 6, for example, where the black holes would be located at $r = 10M$ with horizon radii of $0.5M$), resulting in fewer cells there than in the unwarped setup shown here. This is particularly critical in the magnetized case, where the larger values of the magnetic field close to the black hole horizon would require a larger number of cells to keep the truncation error constant. We also note that the relative solution errors seem to equilibrate quickly, with very little evolution happening after $t \sim 7M$. By increasing the overall number of cells we have further verified that these relative changes do decrease, and thus obtained further evidence that the relatively large values

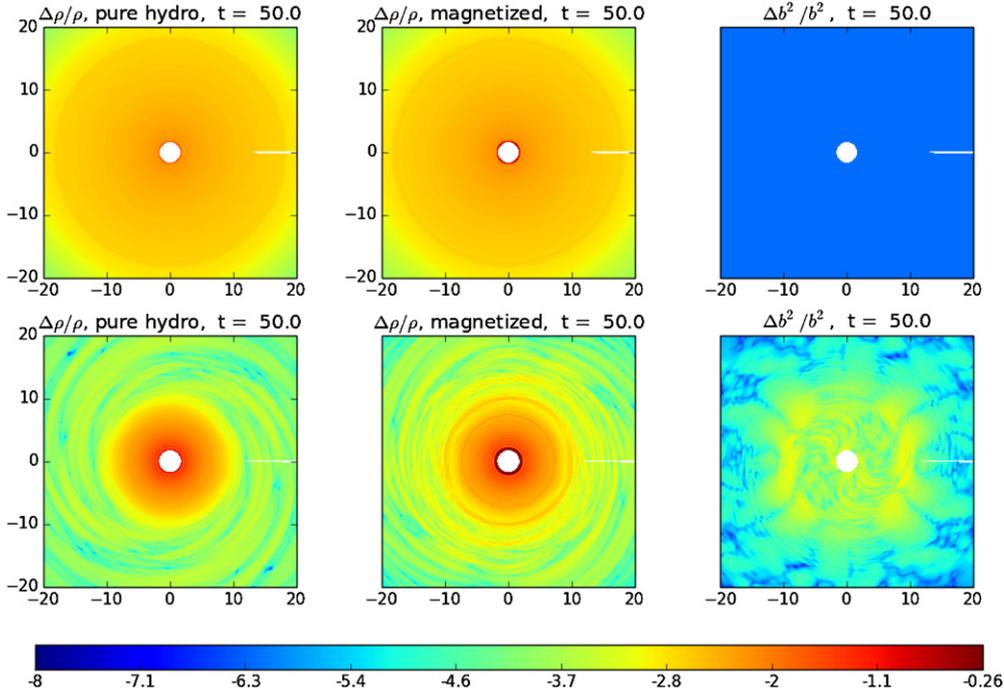


Figure 10. Contour plots (in logarithmic units) of the relative change in density (ρ) (equation (37), left and middle columns) as well as the relative error in magnetic energy density ($\|b\|^2$) (equation (38), right column) at $t = 50M$ for a 2D equatorial Bondi solution test. The pure hydrodynamic Bondi test (left column) is compared with the magnetized case (middle and right columns). Evolutions were made with the warped coordinates (bottom row), and for an unwarped coordinate system with the distortions effectively moved to the origin (top row); see text for the parameters we used. The narrow white wedge appearing in all plots is a visualization artifact that marks the edge between the beginning and end of x^3 .

observed are indeed artifacts of insufficient resolution in the corresponding region, rather than an error or problem with our implementation.

For the magnetized case, a simpler quantitative test is achieved by simply evaluating the following quantity throughout the evolution

$$\frac{\Delta\|b\|^2}{\|b\|^2} = \left| \frac{\|b\|^2 r^4 - A_0^2}{A_0^2} \right|. \quad (38)$$

We plot this quantity at $t = 50M$ in figure 10, which demonstrates much less deviation in $\|b\|^2$ from its initial condition than what was seen in ρ . Apparently, our evolution procedure (including, e.g., our parabolic FluxCT scheme) introduces an insignificant level of numerical error into the magnetic field when using warped coordinates.

As a point of reference, we have measured the relative errors in ρ and $\|b\|^2$ for the magnetized Bondi test when using unwarped coordinates, which were described previously. The differences we find between the two sets of figures then highlight the effects from the truncation errors introduced by the warped coordinates. As expected for the unwarped case, we see no variation in relative errors in the azimuthal dimension. The significantly lower magnetic field errors in the unwarped system are likely because only one of the induction equations is

nontrivial (e.g., $\partial_t B^1 \dots$), whereas—in the warped case—two induction equations are nontrivial which allows the magnetic field components to dynamically interact with each other. Also, the largest values observed with the warped coordinates occur in the immediate vicinity of the black hole horizon, where resolution is coarser than in the unwarped case; the unwarped system’s radial grid spacing shrinks with diminishing radius, unlike the warped system whose radial grid spacing continuously grows from $r = 10 M$ to the horizon. Specifically, the radial step size in the vicinity of the black hole horizon is roughly $dr \simeq 0.056 M$ in the unwarped case, as opposed to $dr \simeq 0.27 M$ for the warped case. Even so, the contrast between the density errors is not large (i.e. it is within a factor of 2) near the horizon, and the warped system becomes more accurate at larger radii since its radial resolution there is higher than that of the unwarped system.

5.2. Disc with single black hole

In order to test the effect of the coordinates’ distortion on an accretion disc, we have evolved an accretion disc in the background of a single (non-spinning) black hole with a numerical grid adapted to the binary configuration of figure 9. For details on the initial data and evolution procedure we refer the reader to [17].

We performed a 2D equatorial evolution, with 400×400 cells; the disc was initialized with an aspect ratio of $H/r = 0.1$ ⁶, inner edge at $r_{\text{in}} = 40 M$, and pressure maximum at $r_p = 78 M$. Our motivation behind our choice of the disc’s location and extent was to approximate the quasi-steady state of the circumbinary disc we have seen develop about an equal mass black hole binary [17, 53]. We evolved the configuration until $t = 16000 M$, corresponding to roughly 10 orbits at the inner edge of the disc, 3.7 orbits at its pressure maximum, and about 28 orbits of the warped grid’s focal points.

Snapshots of the fluid density ρ for two different time steps can be seen in figure 11. Again, we would like to emphasize that the black hole at the origin is not as well resolved as it typically is in single BH accretion disc simulations. As a result, there are fluctuations in the fluid—that begin with density and pressure amplitudes close to that of the atmosphere—that propagate from near the horizon to the inner edge of the disc. The disc is further perturbed by the time-varying azimuthal spacing of the warped coordinate system. By $t \sim 12000 M$, we can see perturbations have grown near the inner edge of the disc to amplitudes of a few orders of magnitude above the atmosphere level, which causes matter to start inflowing. To give a sense of scale, it takes an equal mass binary at the separation of the warps (20 M) approximately 13000 M of time to inspiral. In addition, the outer edge of the disc appears to remain unperturbed after its initial relaxation to a shallower gradient. We expect that magnetic stresses will dramatically alter this picture and give rise to a much more rapid development of accretion onto the black hole(s). Our neglect of magnetic fields here is meant to emphasize the efficacy of our warped coordinates in maintaining stable circularly orbiting distributions of gas.

We have further measured the diffusion of mass from the disc by computing the relative change in rest mass enclosed in the initial volume of the disc. The result obtained can be seen in figure 12. We compare the warped system’s performance to that of an unwarped system, which has been well tested and used in many production runs in HARM3D. The unwarped coordinates used in this test are defined in the following way: $x^3 = \phi$, $\Delta\phi = 2\pi/400$, $x^1 = \ln(r) = x^1_0 + i\Delta x^1$, $x^1_0 = \ln(R_{\text{in}})$, $\Delta x^1 = \ln(R_{\text{out}}/R_{\text{in}})/400$, $R_{\text{in}} = 1.845 M$, and $R_{\text{out}} = 300 M$.

⁶ Since this setup was assumed to be independent of polar angle, the aspect ratio serves as a parameter to control the pressure and angular velocity distribution in the disc and not the geometrical shape of the disc.

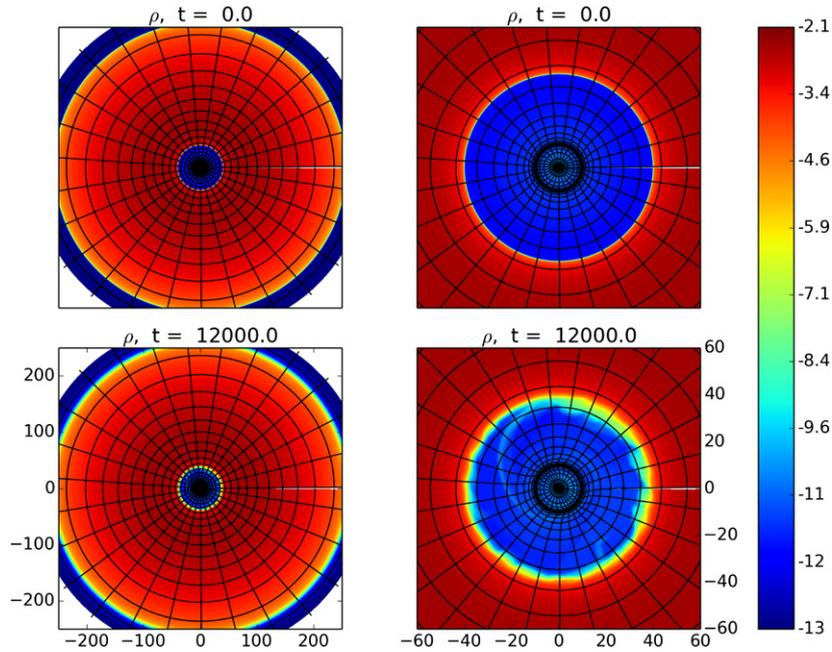


Figure 11. Snapshots of the fluid density ρ (logarithmic scale) at $t = 0$ and $t = 12000 M$ for an accretion disc evolution in the background of a non-spinning single black hole. We have used the warped spherical grid of figure 9. The narrow white wedge appearing in all plots is a visualization artifact that marks the edge between the beginning and end of x^3 .

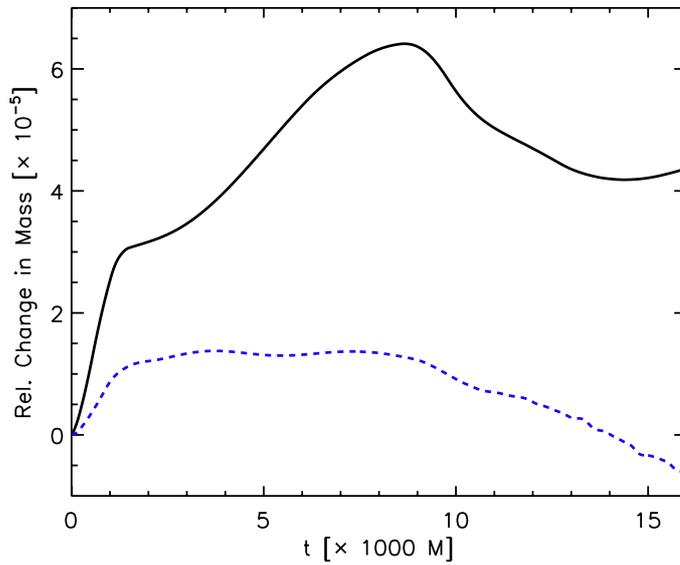


Figure 12. Relative change in rest mass enclosed in the initial volume of the disc. We plot the result obtained with the warped grid configuration of figure 9 (dashed blue line) and with the standard ϕ -independent coordinate system of HARM3D (solid black line).

After about $10^4 M$ or 6.3 orbits at $r = R_{\text{in}}$, the dynamic coordinates have perturbed the inner edge enough so that a steady, though relatively weak, flow of material develops from the disc's edge. This rate of inflow is smaller than the rate of mass increase the two simulations experienced at the very beginning of their runs; the secular late time mass loss of the warped run is 1/4 times that of the warped run's initial mass increase and 1/12 times that of the unwarped run's. The absorption of mass in the two runs is due to the initial data settling about a slightly different equilibrium state after evolution begins that is set by the truncation level of the EOM. We also see larger swings in mass in the unwarped case than the rate of mass loss in the warped case. The explanation for the poorer performance of the unwarped run is that more of its resolution is focused near the horizon and not through the disc in comparison. For instance, the unwarped radial grid spacing is about twice that of the warped case at the outer edge of the disc.

5.3. Disc with black hole binary

Finally, we present preliminary results for an evolution of an actual circumbinary accretion disc, which was the main motivation for this work. For our analytic binary black hole spacetime model, we use a patchwork system of various approximate spacetime metrics (i.e., boosted black hole, post-Newtonian, and post-Minkowski metrics) to cover the entire physical domain, and neglect the effect of the disc on the binary's orbit. The construction of the metric, initial data, and evolution procedure for this configuration is detailed in [17, 43], to where we refer the interested reader. The only difference here is the disc is not magnetized and the disc lies only in the equator (i.e. the disc is assumed to be unstratified and is resolved by only one cell in the vertical dimension). Our test starts with two equal mass black holes at a fixed separation of $20 M$, $M \equiv M_{\text{BH1}} + M_{\text{BH2}}$. The grid is one we have used before, namely that described in figure 9. This grid is constructed so that there are approximately 32 cells spanning each black hole horizon in each dimension. The distortions track the black holes on their orbits.

Since the black holes now reside on the numerical domain, we 'excise' a small volume of cells about each black hole's singularity from the update procedure to maintain a stable hydrodynamic evolution. The innermost metrics used to cover the spacetime about each black hole are written in horizon-penetrating coordinates, which allows us to excise cells well within the horizon. Our excision method involves treating cells in three different ways depending on their proximity to the black holes' centers. So-called 'excised' cells are not evolved, and primitive variables and fluxes are not calculated at their faces. There are also regular or 'evolved' cells that are updated in the usual way and are shielded from the excised cells by so-called 'buffer' cells. The buffer cells ensure that evolved cells never use any data from excised cells, e.g., for the primitive variable reconstruction, the flux calculation, or the FluxCT procedure. The buffer cells are not updated, but primitive variables are reconstructed and fluxes calculated at their faces. In the test presented here, a cell is excised if any part of it lies within half a black hole's horizon radius from the black hole's center. Since the stencil for our finite volume and FluxCT procedures span three cells on each side of a cell, buffer cells extend three cells in each dimension from the surface of the excised region. The remaining cells are evolved cells.

In figure 13 we can see the density of the fluid for three different time steps, $t = 0 M, 2200 M, 4400 M$; these times occur after approximately 0, 4, 8 orbits, respectively. Overdense regions of densities just above the atmosphere level quickly arise, once the simulation begins, trailing each black hole. The overdensities are due to the fact that the numerically imposed floor or atmosphere is not in equilibrium with the binary's potential, allowing matter—which is artificially fed into the domain from the floor—to condense about

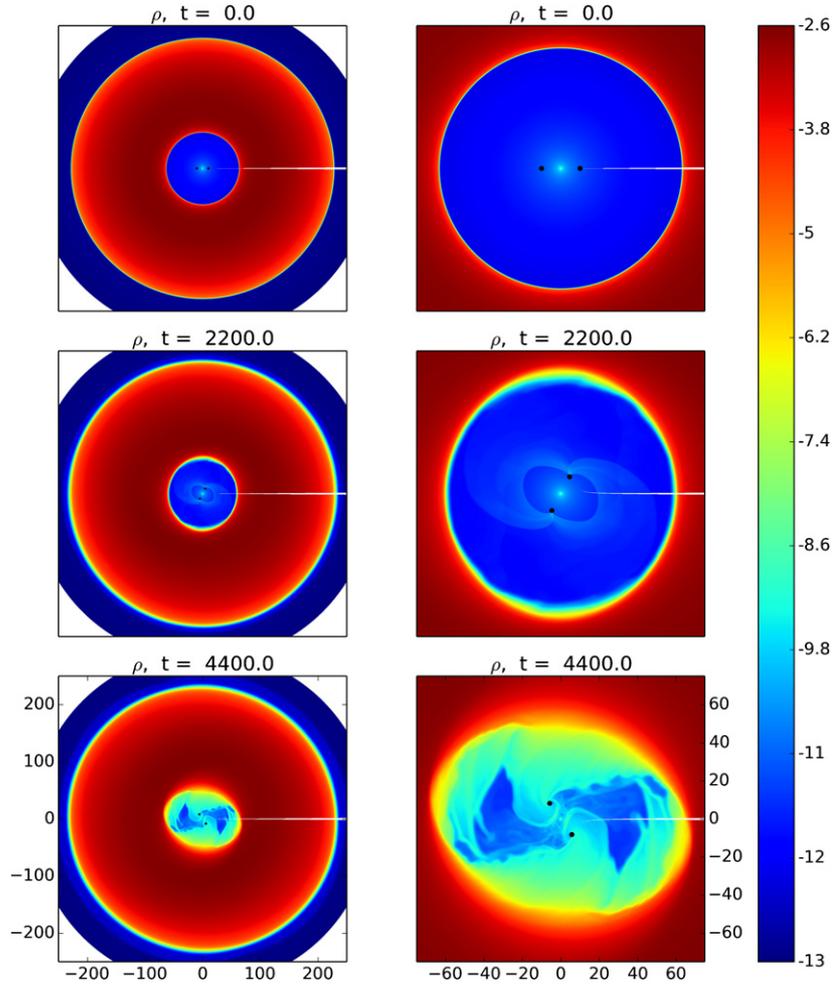


Figure 13. Plotted is the density of the fluid, ρ , (in logarithmic scale) for a 2D equatorial evolution of an accretion disc in the background of a black hole binary at $t = 0$ M, $t = 2200$ M and $t = 4400$ M. The black circles depict the black holes' horizons. We are here omitting the grid lines so as to not clutter the figure, but we note that the grid has the same configuration as the one presented in figure 9. The narrow white wedge appearing in all plots is a visualization artifact that marks the edge between the beginning and end of x^3 .

each black hole. This phenomenon is similar to what is seen with binaries in a similar environment—that of a non-rotating cloud of gas [54]. By $t = 2200$ M, the inner edge of the disc has begun to noticeably respond to the time-varying quadrupolar gravitational potential and distort from its original circular configuration.

Very little accretion from the disc occurs until $t \simeq 4000$ M, after which gravitational torques begin to more efficiently draw material in toward the binary and dense accretion streams form from the disc to each black hole. The presence of these accretion streams is characteristic of circumbinary disc evolutions [23–25, 55, 56]. One key difference we find is that the accretion streams seem to exhibit the Kelvin–Helmholtz instability, recognizable by the turbulent eddies and waves that form along the shear layer at the edge of the streams.

As far as the authors know, the Kelvin–Helmholtz instability has not been observed before in the context of circumbinary accretion onto black holes. We suppose that we may be seeing it because of the higher effective accuracy of our hydrodynamics methods. We intend to further investigate the accretion streams and verify that this is in fact the Kelvin–Helmholtz instability in future work.

We find that the warped coordinates sufficiently resolve material flowing near each black hole and that our excision procedure works as desired. Material flows into each black hole without reflection, and no hydrodynamic waves appear to emanate from them. The warped grid also conforms to the symmetry of the inner edge of the disc well enough so that accretion occurs later in time, on the timescale set by the binary’s gravitational torques, and not by the effective grid scale viscosity present in—for instance—Cartesian coordinates. Future efforts include exploring the role the Kelvin–Helmholtz instability plays in the evolution of the streams, and incorporating magnetic fields in 3D simulations using these coordinates.

6. Final remarks

In this paper, we have constructed a warped gridding scheme adapted to binary black hole simulations and discussed its implementation in the GRMHD `HARM3D` code.

In order to demonstrate the coordinate system’s efficacy and to validate its implementation a number of tests were performed. We have shown results obtained for the advection of a magnetic field loop, the evolution of the Bondi solution, and the evolution of a stationary gaseous disc in the background of a single black hole. We lastly presented preliminary results for the evolution of a circumbinary disc, the main motivation for this work. The warped coordinate evolution in the advected field loop test performed even better than the uniform Cartesian system, demonstrating that the distortions do not lead to extra diffusion in the magnetic field evolution. The dynamic and asymmetric aspect of the coordinates does introduce additional errors at the truncation level when compared to symmetric grids and solutions (e.g., the Bondi solution, stationary disc), but these errors do diminish with increasing cell count and are much smaller than what we would find if we were to use Cartesian grids with the same cell count. For instance, the stationary disc solution accreted no mass until after tens of binary orbits and several orbits at the disc’s inner edge—whereas disc evolutions on Cartesian grids immediately start accreting gas. Also, our results demonstrated that the distortions in the grid were able to resolve small scale features as expected and produced little—if any—anomalous features in the solutions.

We note that, while our results were very good for the warped grid configurations we tested, they do not necessarily extrapolate to different grid configurations. Indeed, we observed poorer results when the grid chosen was significantly distorted and not at all adapted to the symmetries of the problem. Care must therefore be taken when choosing the grid parameters, and some degree of testing and experimentation should be expected when using this construction.

Another consideration to make is one of computational cost. While the evaluation of all the functions related to the coordinate transformation carries some overhead in the code, we mitigated this impact by eliminating redundancies by storing auxiliary terms in memory rather than recalculating them. In our case, our spacetime metric is specified in Cartesian coordinates, so a coordinate transformation to spherical coordinates is necessary regardless of whether we use the coordinates described here.

We stress that our construction and resulting grid are sufficiently general that they can be applied to very different problems and codes. As are many other GRMHD codes (e.g., [4]), `HARM3D` is written in a covariant way making the adoption of this coordinate system

straightforward. For instance, it would be interesting to see if dynamic gauge conditions used to evolve punctures, would be stable under these dynamic coordinate changes. Our approach could be especially interesting for modular unigrid codes as a simpler alternative to AMR, particular for problems tracking compact features. The satisfactory results presented here bring us high hopes for ongoing efforts toward modeling circumbinary accretion flows.

Acknowledgments

We appreciate the helpful input from our collaborators M Campanelli, Y Zlochower, J Krolik, C Lousto, B Mundim, H Nakano, and K Sorathia. MZ is supported by NSF grants nos OCI-0832606, PHY-0969855, AST-1028087, and PHY-1229173. SCN is supported by NSF grant nos OCI-0725070, OCI-0832606, AST-1028087, PHY-1125915. Computational resources were provided by XSEDE allocation TG-PHY060027N, and by NewHorizons and BlueSky Clusters at Rochester Institute of Technology, which were supported by NSF grant nos PHY-0722703, DMS-0820923, AST-1028087, and PHY-1229173.

Appendix. Magnetic fields on generally warped coordinates

It is often the case that initial data routines assume properties of the numerical coordinates on which they reside. For instance, many initial data distributions are constant with respect to the azimuthal coordinate—which is often assumed to be the same as one of the numerical coordinates—so they tend to copy the data along this azimuthal array dimension, or—for instance—use this independence to derive the magnetic field in a simple way. When using warped coordinates, however, we can no longer make such simplifying assumptions and must calculate the initial data in the most general way. We have made these changes to the relevant initial data routines used by HARM3D so that assumptions on the coordinate system no longer remain. In the following, we describe the new method for initializing the magnetic field.

First, let us present a short review of classical electrodynamics (please see [57, 58] for more details). Using the Faraday tensor, $F^{\mu\nu}$, and the Maxwell tensor, $\tilde{F}^{\mu\nu}$, we can succinctly express the Maxwell equations as

$$\nabla_{\mu} F^{\nu\mu} = J^{\nu}, \quad (\text{A.1})$$

$$\nabla_{\mu} \tilde{F}^{\mu\nu} = 0. \quad (\text{A.2})$$

We have absorbed factors of $(1/\sqrt{4\pi})$ into the tensors for the sake of simplicity. Note that $F^{\mu\nu}$ and $\tilde{F}^{\mu\nu}$ are both antisymmetric, e.g. $F^{\mu\nu} = -F^{\nu\mu}$ and $\tilde{F}^{\mu\nu} = -\tilde{F}^{\nu\mu}$. The first set yields the evolution equations for the electric field, while the second set gives rise to the induction equations and the solenoidal constraint. The electromagnetic part of the stress–energy tensor is

$$T_{\text{EM}}^{\mu\nu} = F^{\mu\lambda} F^{\nu}_{\lambda} - \frac{1}{4} g^{\mu\nu} F^{\lambda\kappa} F_{\lambda\kappa} = ||b||^2 u^{\mu} u^{\nu} + \frac{1}{2} ||b||^2 g^{\mu\nu} - b^{\mu} b^{\nu}, \quad (\text{A.3})$$

where the magnetic 4-vector is $b^{\mu} = \tilde{F}^{\nu\mu} u_{\nu}$, and $\tilde{F}^{\mu\nu} = b^{\mu} u^{\nu} - b^{\nu} u^{\mu}$. We typically use the spatial magnetic field vector,

$$B^i = \frac{1}{\alpha} \mathcal{B}^i = \tilde{F}^{it} \quad (\text{A.4})$$

where

$$\mathcal{B}^{\mu} = n_{\nu} \tilde{F}^{\nu\mu} \quad (\text{A.5})$$

and $n_\mu = [-\alpha, 0, 0, 0]$, $n^\mu = \frac{1}{\alpha}[1, -\beta^i]$. One can easily show that the i -component of equation (A.2) gives rise to the induction equation:

$$\partial_t \sqrt{-g} B^i + \partial_j \sqrt{-g} (b^j u^i - b^i u^j) = 0 \quad (\text{A.6})$$

and the solenoidal constraint (also called as ‘divergence-less constraint’, also called as ‘div.b=0’ condition, also called as ‘no magnetic monopoles’ condition, etc):

$$\partial_i \sqrt{-g} B^i = 0. \quad (\text{A.7})$$

Note that this comes from the time component of equation (A.2):

$$0 = \nabla_\mu {}^* F^{\mu t} = \frac{1}{\sqrt{-g}} [\partial_\mu (\sqrt{-g} {}^* F^{\mu t}) + \sqrt{-g} \Gamma^t{}_{\mu\kappa} {}^* F^{\mu\kappa}] \quad (\text{A.8})$$

$$= \frac{1}{\sqrt{-g}} [\partial_i \sqrt{-g} B^i] \quad (\text{A.9})$$

where equation (A.9) follows equation (A.8) because ${}^* F^{\mu\nu}$ is antisymmetric and the Christoffel symbol is symmetric in the lower indices. Note that this is a gauge-independent constraint, i.e. one can perform an arbitrary coordinate transformation at the beginning and end up with a divergence-less quantity in that new coordinate system. All that is required is that both the derivative operator and the magnetic field vector be in the *same* coordinate system. Since our finite difference operators are w.r.t. the uniform numerical coordinates, then B^i must be in numerical coordinates as well.

Finite difference solutions are only accurate to truncation error. Different stencils—or finite difference approximations—for the divergence operator of the solenoidal constraint give rise to different truncation errors. The idea behind the CT method is that one special stencil for the difference operation is chosen so that the solenoidal constraint is satisfied to machine precision. This special stencil or finite difference operator is called ‘compatible’ with the CT scheme. The finite difference operator for first-order spatial derivatives compatible with our FluxCT scheme is:

$$\hat{\Delta}_{x^1} B_{[l_1-1/2, l_2-1/2, l_3-1/2]} = \frac{1}{4\Delta_{x^1}} \sum_{p=-1}^0 \sum_{q=-1}^0 (B_{[l_1, l_2+p, l_3+q]} - B_{[l_1-1, l_2+p, l_3+q]}) \quad (\text{A.10})$$

with similar equations for the $\hat{\Delta}_{x^2}$ and $\hat{\Delta}_{x^3}$ operators, and where $B_{[i, j, k]}$ is the value of a magnetic field component at the numerical coordinates of the center of the cell with spatial indices $\{i, j, k\}$. Our compatible stencil is centered at the cell corners while our magnetic field vectors are cell-centered quantities. One can describe the compatible stencil as a corner-centered difference of edge-centered quantities, which are found by averaging the four nearest cell-centered values to the edge centers. To be explicit, the following divergence operation is preserved to round-off errors in our FluxCT scheme:

$$\nabla_{\text{CT}} \cdot \mathbf{B} = [\hat{\Delta}_{x^1} B^1 + \hat{\Delta}_{x^2} B^2 + \hat{\Delta}_{x^3} B^3]_{[l_1-1/2, l_2-1/2, l_3-1/2]}. \quad (\text{A.11})$$

In order to initialize the magnetic field such that it is divergence-less w.r.t. this stencil, we set the vector potential at the cell corners and take the curl of it using the same derivative operators as were used for ∇_{CT} . We remind the reader that the covariant 4-vector potential is $\mathcal{A}_\mu = \Phi n_\mu + A_\mu$, where Φ is the electrostatic scalar potential, and A_μ is the spatial vector potential, where $A_\mu n^\mu = 0$, $A_t = \beta^i A_i$, $A^t = 0$, and $\mathcal{A}_i = A_i$ [57, 58]. The magnetic field is calculated from the curl of the vector potential:

$$B^i = \frac{1}{\sqrt{-g}} \tilde{\epsilon}^{ijk} \partial_j A_k, \quad (\text{A.12})$$

where

$$\tilde{\epsilon}^{ijk} = \sqrt{\gamma}\epsilon^{ijk} = \sqrt{\gamma}n_\mu\epsilon^{\mu ij k} = -\sqrt{-g}\epsilon^{tijk}. \quad (\text{A.13})$$

and $\tilde{\epsilon}^{ijk}$ is the anti-symmetric permutation tensor that equals 1 for even permutations of $(1, 2, 3)$, -1 for odd permutations, and zero for repeating indices. Also, here, $\epsilon^{\mu\nu\kappa\lambda}$ and ϵ^{ijk} are, respectively, the standard four-dimensional and three-dimensional (3D) Levi-Civita symbols. It is tedious, though straightforward, to show that

$$\mathbf{B}_{[l_1, l_2, l_3]}^i = \left[\frac{1}{\sqrt{-g}} \tilde{\epsilon}^{ijk} \hat{\Delta}_{x^j} A_k \right]_{[l_1, l_2, l_3]} \Rightarrow [\nabla_{\text{CT}} \cdot \mathbf{B}]_{[l_1, l_2, l_3]} = 0. \quad (\text{A.14})$$

We typically set A_ϕ , transform A_i to numerical coordinates, and evaluate the curl in numerical coordinates using $\hat{\Delta}_{x^j}$ to approximate the derivatives in the curl expression. We need not calculate the \mathcal{A}_i component before transforming since $dt/dx^i = 0$ for the coordinates of interest here.

In order to enforce the divergence-less constraint, we use a 3D version of the FluxCT algorithm described in [28, 59]. The principal idea behind most CT schemes is that a consistent set of EMFs are used to reconstruct the fluxes in the induction equation so that, when the discretized time derivative and the CT-consistent divergence operator both act on the magnetic field, the result is zero algebraically. Please see [60] for a thorough description of how CT schemes are constructed, as we will only describe here what is necessary for reproducing our scheme. The first step in our procedure is to reconstruct the primitive variables and calculate the numerical fluxes at each cell face. It is easy to see that the spatial fluxes of the induction equation are components of the Maxwell tensor:

$$\partial_t \sqrt{-g} B^i + \partial_j (\sqrt{-g} F^{ij}) = 0. \quad (\text{A.15})$$

If a cell's center is located at the point with indices $[l_1, l_2, l_3]$, let that cell's faces in the x^i direction be offset by $\pm 1/2$ in that dimension, e.g., the faces orthogonal to x^1 have indices $[l_1 \pm 1/2, l_2, l_3]$. The cell's edges along that face are then half intervals in the other dimensions, e.g., for the top x^1 face the indices are $\{[l_1 + 1/2, l_2 \pm 1/2, l_3], [l_1 + 1/2, l_2, l_3 \pm 1/2]\}$. We next use the numerical fluxes from the first step to interpolate for a common set of edge-centered EMFs shared by all cells. This procedure exploits the analytic identity that $\mathcal{E}_k = -\tilde{\epsilon}_{ijk} b^j u^i$ and $F^{ij} = -\tilde{\epsilon}^{ijk} \mathcal{E}_k$. For each EMF component, an average is taken between two interpolated quantities, with each interpolation performed along a different dimension. Specifically, the EMFs are:

$$\mathcal{E}_{1[l_1, l_2-1/2, l_3-1/2]} = \frac{1}{2} [\langle \hat{F}^{32} \rangle_{[l_1, l_2-1/2, l_3-1/2]} - \langle \hat{F}^{23} \rangle_{[l_1, l_2-1/2, l_3-1/2]}], \quad (\text{A.16})$$

$$\mathcal{E}_{2[l_1-1/2, l_2, l_3-1/2]} = \frac{1}{2} [\langle \hat{F}^{13} \rangle_{[l_1-1/2, l_2, l_3-1/2]} - \langle \hat{F}^{31} \rangle_{[l_1-1/2, l_2, l_3-1/2]}], \quad (\text{A.17})$$

$$\mathcal{E}_{3[l_1-1/2, l_2-1/2, l_3]} = \frac{1}{2} [\langle \hat{F}^{21} \rangle_{[l_1-1/2, l_2-1/2, l_3]} - \langle \hat{F}^{12} \rangle_{[l_1-1/2, l_2-1/2, l_3]}], \quad (\text{A.18})$$

where \hat{F}^{ij} are the numerical fluxes calculated in the first step, and the angle brackets denote that the interior quantity has been interpolated to the indicated position. For the 'linear' FluxCT procedure, we use centered first-order interpolation, equivalent to nearest neighbor averaging:

$$\langle \hat{F}^{12} \rangle_{[l_1-1/2, l_2-1/2, l_3]} = \frac{1}{2} [\hat{F}_{[l_1, l_2-1/2, l_3]}^{12} + \hat{F}_{[l_1-1, l_2-1/2, l_3]}^{12}], \quad (\text{A.19})$$

$$\langle \hat{F}^{21} \rangle_{[l_1-1/2, l_2-1/2, l_3]} = \frac{1}{2} [\hat{F}_{[l_1-1/2, l_2, l_3]}^{21} + \hat{F}_{[l_1-1/2, l_2-1, l_3]}^{21}], \quad (\text{A.20})$$

$$\langle \hat{F}^{13} \rangle_{[l_1-1/2, l_2, l_3-1/2]} = \frac{1}{2} [\hat{F}_{[l_1, l_2, l_3-1/2]}^{13} + \hat{F}_{[l_1-1, l_2, l_3-1/2]}^{13}], \quad (\text{A.21})$$

$$\langle \hat{F}^{31} \rangle_{[l_1-1/2, l_2, l_3-1/2]} = \frac{1}{2} [\hat{F}^{31}_{[l_1-1/2, l_2, l_3]} + \hat{F}^{31}_{[l_1-1/2, l_2, l_3-1]}], \quad (\text{A.22})$$

$$\langle \hat{F}^{23} \rangle_{[l_1, l_2-1/2, l_3-1/2]} = \frac{1}{2} [\hat{F}^{23}_{[l_1, l_2, l_3-1/2]} + \hat{F}^{23}_{[l_1, l_2-1, l_3-1/2]}], \quad (\text{A.23})$$

$$\langle \hat{F}^{32} \rangle_{[l_1, l_2-1/2, l_3-1/2]} = \frac{1}{2} [\hat{F}^{32}_{[l_1, l_2-1/2, l_3]} + \hat{F}^{32}_{[l_1, l_2-1/2, l_3-1]}]. \quad (\text{A.24})$$

For the ‘parabolic’ FluxCT scheme, the interpolation is performed using the same piecewise parabolic interpolation method we use for reconstructing face-centered quantities when computing the numerical fluxes. Lastly, the ultimate divergence-less-preserving induction equation fluxes (\tilde{F}^{ij}) are calculated by linearly interpolating a pair of edge EMFs to the face’s center:

$$\tilde{F}^{12}_{[l_1, l_2-1/2, l_3]} = -\frac{1}{2} [\mathcal{E}_{3[l_1-1/2, l_2-1/2, l_3]} + \mathcal{E}_{3[l_1+1/2, l_2-1/2, l_3]}], \quad (\text{A.25})$$

$$\tilde{F}^{21}_{[l_1-1/2, l_2, l_3]} = \frac{1}{2} [\mathcal{E}_{3[l_1-1/2, l_2-1/2, l_3]} + \mathcal{E}_{3[l_1-1/2, l_2+1/2, l_3]}], \quad (\text{A.26})$$

$$\tilde{F}^{13}_{[l_1, l_2, l_3-1/2]} = \frac{1}{2} [\mathcal{E}_{2[l_1-1/2, l_2, l_3-1/2]} + \mathcal{E}_{2[l_1+1/2, l_2, l_3-1/2]}], \quad (\text{A.27})$$

$$\tilde{F}^{31}_{[l_1-1/2, l_2, l_3]} = -\frac{1}{2} [\mathcal{E}_{2[l_1-1/2, l_2, l_3-1/2]} + \mathcal{E}_{2[l_1-1/2, l_2, l_3+1/2]}], \quad (\text{A.28})$$

$$\tilde{F}^{23}_{[l_1, l_2, l_3-1/2]} = -\frac{1}{2} [\mathcal{E}_{1[l_1, l_2-1/2, l_3-1/2]} + \mathcal{E}_{1[l_1, l_2+1/2, l_3-1/2]}], \quad (\text{A.29})$$

$$\tilde{F}^{32}_{[l_1, l_2-1/2, l_3]} = \frac{1}{2} [\mathcal{E}_{1[l_1, l_2-1/2, l_3-1/2]} + \mathcal{E}_{1[l_1, l_2-1/2, l_3+1/2]}]. \quad (\text{A.30})$$

These fluxes are those used to update the magnetic field components.

References

- [1] Berger M J and Oliger J 1984 *J. Comput. Phys.* **53** 484–512
- [2] Berger M J and Colella P 1989 *J. Comput. Phys.* **82** 64–84
- [3] Fixed Mesh Refinement with Carpet www.tat.physik.uni-tuebingen.de/~schnette/carpet/
- [4] Moesta P *et al* 2013 arXiv:1304.5544
- [5] Anderson M, Hirschmann E, Liebling S L and Neilsen D 2006 *Class. Quantum Grav.* **23** 6503–24
- [6] East W E, Pretorius F and Stephens B C 2012 *Phys. Rev. D* **85** 124010
- [7] Anninos P, Fragile P C and Salmonson J D 2005 *Astrophys. J.* **635** 723–40
- [8] Calabrese G and Neilsen D 2004 *Phys. Rev. D* **69** 044020
- [9] Zink B, Schnetter E and Tiglio M 2008 *Phys. Rev. D* **77** 103015
- [10] Duez M D *et al* 2008 *Phys. Rev. D* **78** 104015
- [11] Fragile P C, Lindner C C, Anninos P and Salmonson J D 2009 *Astrophys. J.* **691** 482–94
- [12] Reisswig C, Haas R, Ott C D, Abdikamalov E, Mösta P, Pollney D and Schnetter E 2013 *Phys. Rev. D* **87** 064023
- [13] Springel V 2010 *Mon. Not. R. Astron. Soc.* **401** 791–851
- [14] Pakmor R, Bauer A and Springel V 2011 *Mon. Not. R. Astron. Soc.* **418** 1392–401
- [15] Duffell P C and MacFadyen A I 2011 *Astrophys. J. Suppl.* **197** 15
- [16] Palenzuela C 2012 Private communication
- [17] Noble S C, Mundim B C, Nakano H, Krolik J H, Campanelli M, Zlochower Y and Yunes N 2012 *Astrophys. J.* **755** 51
- [18] Balbus S A and Hawley J F 1991 *Astrophys. J.* **376** 214–33
- [19] Balbus S A and Hawley J F 1998 *Rev. Mod. Phys.* **70** 1–53
- [20] Noble S C and Krolik J H 2009 *Astrophys. J.* **703** 964–75
- [21] Schnittman J D, Krolik J H and Noble S C 2013 *Astrophys. J.* **769** 156
- [22] Bode T, Haas R, Bogdanović T, Laguna P and Shoemaker D 2010 *Astrophys. J.* **715** 1117–31
- [23] Farris B D, Liu Y T and Shapiro S L 2011 *Phys. Rev. D* **84** 024024
- [24] Bode T, Bogdanović T, Haas R, Healy J, Laguna P and Shoemaker D 2012 *Astrophys. J.* **744** 45
- [25] Farris B D, Gold R, Paschalidis V, Etienne Z B and Shapiro S L 2012 *Phys. Rev. Lett.* **109** 221102

- [26] Giacomazzo B, Baker J G, Miller M C, Reynolds C S and van Meter J R 2012 *Astrophys. J. Lett.* **752** L15
- [27] Noble S C, Krolik J H and Hawley J F 2009 *Astrophys. J.* **692** 411–21
- [28] Gammie C F, McKinney J C and Tóth G 2003 *Astrophys. J.* **589** 444–57
- [29] De Villiers J P and Hawley J F 2003 *Astrophys. J.* **592** 1060–77
- [30] Baker J, Brüggmann B, Campanelli M and Lousto C O 2000 *Class. Quantum Grav.* **17** L149–56
- [31] Baker J, Brüggmann B, Campanelli M, Lousto C O and Takahashi R 2001 *Phys. Rev. Lett.* **87** 121103
- [32] Baker J, Campanelli M and Lousto C O 2002 *Phys. Rev. D* **65** 044001
- [33] Gansner E, Hu Y and North S 2012 arXiv:1206.3980
- [34] Yasseri T, Spoerri A, Graham M and Kertész J 2013 arXiv:1305.5566
- [35] Zlochower Y, Baker J G, Campanelli M and Lousto C O 2005 *Phys. Rev. D* **72** 024021
- [36] Campanelli M, Lousto C O, Marronetti P and Zlochower Y 2006 *Phys. Rev. Lett.* **96** 111101
- [37] Campanelli M, Lousto C O and Zlochower Y 2006 *Phys. Rev. D* **73** 061501
- [38] Campanelli M, Lousto C O and Zlochower Y 2006 *Phys. Rev. D* **74** 041501
- [39] Campanelli M, Lousto C O and Zlochower Y 2006 *Phys. Rev. D* **74** 084023
- [40] Campanelli M, Lousto C O, Zlochower Y, Krishnan B and Merritt D 2007 *Phys. Rev. D* **75** 064030
- [41] Shibata M, Liu Y T, Shapiro S L and Stephens B C 2006 *Phys. Rev. D* **74** 104026
- [42] Hawley J F, Richers S A, Guan X and Krolik J H 2013 *Astrophys. J.* **772** 102
- [43] Mundim B C, Nakano H, Campanelli M, Noble S C, Yunes N and Zlochower Y 2013 arXiv:1312.6731
- [44] Tóth G 2000 *J. Comput. Phys.* **161** 605–52
- [45] Noble S C, Gammie C F, McKinney J C and Del Zanna L 2006 *Astrophys. J.* **641** 626–37
- [46] Devore C R 1991 *J. Comput. Phys.* **92** 142
- [47] Tóth G and Odstrcil D 1996 *J. Comput. Phys.* **128** 82–100
- [48] Gardiner T A and Stone J M 2005 *J. Comput. Phys.* **205** 509–39
- [49] Beckwith K and Stone J M 2011 *Astrophys. J. Suppl.* **193** 6
- [50] van Leer B 1977 *J. Comput. Phys.* **23** 276
- [51] Colella P and Woodward P R 1984 *J. Comput. Phys.* **54** 174–201
- [52] Bondi H 1952 *Mon. Not. R. Astron. Soc.* **112** 195
- [53] Shi J M, Krolik J H, Lubow S H and Hawley J F 2012 *Astrophys. J.* **749** 118
- [54] Bode T, Haas R, Bogdanovic T, Laguna P and Shoemaker D 2010 *Astrophys. J.* **715** 1117–31
- [55] Hayasaki K, Mineshige S and Sudou H 2007 *Publ. Astron. Soc. Japan* **59** 427–41
- [56] Cuadra J, Armitage P J, Alexander R D and Begelman M C 2009 *Mon. Not. R. Astron. Soc.* **393** 1423–32
- [57] Jackson J D 1998 *Classical Electrodynamics* 3rd edn (New York: Wiley)
- [58] Baumgarte T W and Shapiro S L 2003 *Astrophys. J.* **585** 921–9
- [59] Tóth G 2000 *J. Comput. Phys.* **161** 605–52
- [60] Evans C R and Hawley J F 1988 *Astrophys. J.* **332** 659–77