

APPLICATION OF THE CUBED-SPHERE GRID TO TILTED BLACK HOLE ACCRETION DISKS

P. CHRIS FRAGILE¹, CHRISTOPHER C. LINDNER^{1,3}, PETER ANNINOS², AND JAY D. SALMONSON²

¹ Department of Physics and Astronomy, College of Charleston, Charleston, SC 29424, USA; fragilep@cofc.edu

² Lawrence Livermore National Laboratory, Livermore, CA 94550, USA

³ Astronomy Department, University of Texas, 1 University Station, C1400, Austin, TX 78712, USA

Received 2008 July 28; accepted 2008 September 23; published 2009 January 15

ABSTRACT

In recent work we presented the first results of global general relativistic magnetohydrodynamic (GRMHD) simulations of tilted (or misaligned) accretion disks around rotating black holes. The simulated tilted disks showed dramatic differences from comparable untilted disks, such as asymmetrical accretion onto the hole through opposing “plunging streams” and global precession of the disk powered by a torque provided by the black hole. However, those simulations used a traditional spherical–polar grid that was purposefully under-resolved along the pole, which prevented us from assessing the behavior of any jets that may have been associated with the tilted disks. To address this shortcoming we have added a block-structured “cubed-sphere” grid option to the Cosmos++ GRMHD code, which will allow us to simultaneously resolve the disk and polar regions. Here we present our implementation of this grid and the results of a small suite of validation tests intended to demonstrate that the new grid performs as expected. The most important test in this work is a comparison of identical tilted disks, one evolved using our spherical–polar grid and the other with the cubed-sphere grid. We also demonstrate an interesting dependence of the early-time evolution of our disks on their orientation with respect to the grid alignment. This dependence arises from the differing treatment of current sheets within the disks, especially whether or not they are aligned with symmetry planes of the grid.

Key words: accretion, accretion disks – black hole physics – methods: numerical – MHD – relativity

Online-only material: color figures

1. INTRODUCTION

We have recently undertaken a series of numerical studies of titled accretion disks around rapidly rotating black holes, first in the hydrodynamic (Fragile & Anninos 2005) and then in the magnetohydrodynamic (MHD; Fragile et al. 2007b) limits. All of these simulations have been fully general relativistic, using the Kerr–Schild metric to represent the spacetime of the black hole.

Tilted accretion disks are of particular interest because they are subject to differential warping due to the Lense–Thirring precession of the rotating black hole. For very thin disks, close to the black hole the competition between the differential twisting and “viscous” damping may cause the angular momenta of the disk and hole to align. Further out in the disk, beyond some warp radius, the disk maintains its misaligned state.

For moderately thin to thick disks, such as the ones we simulated previously, the situation is more complex and interesting. The primary difference is that warping is transported via bending waves rather than diffusively, as for thin disks. One consequence of this is that the midplane of a thick disk does not tend to align with the symmetry plane of the black hole at small radii, as in the thin disk case. In fact, the relative tilt between the black hole and disk angular momenta can *increase* at small radii. Having the tilted disk penetrate very close to the black hole has many interesting consequences. For instance, we found that accretion onto the hole occurs predominantly through two opposing “plunging streams” that start from high latitudes with respect to both the black hole and disk midplanes (Fragile et al. 2007a). There is also a strong epicyclic driving within the disk attributable to the gravitomagnetic torque of the misaligned (tilted) black hole (Fragile & Blaes 2008). The induced motion of the gas can be coherent over the scale of the entire

disk. The gas also experiences periodic (twice per orbit) compressions. The compressions occur as the gas orbits past the line of nodes between the black hole symmetry plane and disk midplane. Near the black hole these compressive motions can become supersonic and transform into a pair of quasistationary shocks. The shocks act to enhance angular momentum transport and dissipation near the hole, forcing some material to plunge toward the black hole from well outside the innermost stable circular orbit. Finally, because we are simulating disks with finite radial extents and fast sound-crossing times, the torque of the black hole causes the entire disk body to precess globally.

The main shortcoming of our work so far comes from limitations imposed on us by our use of a spherical–polar grid. First, construction of a uniform spherical–polar grid in three dimensions results in very small zones surrounding the two poles, where all of the lines of longitude converge. These very small zones constrain the Courant-limited time step to be exceedingly small, such that the required CPU cycle count becomes prohibitively large. To avoid this problem, researchers have either excised a small conical section around each pole (e.g., De Villiers & Hawley 2003) or used a lower grid resolution near the poles (Fragile et al. 2007b). Although these techniques are reasonable when one is primarily interested in studying the equatorial region (where a disk may form), these are not satisfactory when one is interested in what is happening in the polar regions (where jets may form). A second problem with the spherical–polar grid is that the poles themselves actually represent coordinate singularities, which present significant challenges for numerical advection and curvature coupling schemes (e.g., solving Riemann curvature source terms).

For these reasons we have added the cubed-sphere grid (Sadourmy 1972; Ronchi et al. 1996) as an option within our numerical code, Cosmos++. The advantage of this grid

construction is that its topological properties more closely resemble a Cartesian coordinate system than a spherical–polar system. The cubed-sphere grid uses a more uniform zone spacing than spherical polar, so the time step can remain reasonably large even in high-resolution simulations. Also important, the grid does not contain any coordinate singularities except at the origin, which is not a concern for our intended use since we truncate the grid just inside the event horizon of the black hole. Ours is not the first application of the cubed-sphere grid to problems in computational astrophysics; it has been used previously to study accretion onto rotating stars with inclined magnetic fields (Koldoba et al. 2002; Romanova et al. 2003) and a few problems in stellar evolution (Dearborn et al. 2005, 2006). However, our work is one of the first applications of this grid to the study of black hole accretion disks and their attendant jets (see also Zink et al. 2008).

The paper is organized as follows: Section 2 describes the cubed-sphere mesh in detail and our particular implementation. In Section 3 we discuss results of basic gradient tests on the cubed-sphere grid. In Section 4 we compare two sets of numerical simulations of black hole accretion disks. In the first set we compare simulations of disks accreting onto a Schwarzschild black hole. We compare different grids, resolutions, and orientations of the disk with respect to the grid. Because these simulations use a Schwarzschild black hole, the orientation should have no physical meaning. However, we show that there are, nevertheless, considerable differences in their evolution at early times. We present the case that the differences have to do with the differing treatments of the midplane current sheets in the disks, which forms from the differential winding of our initial poloidal field loops. Finally, we compare simulations of tilted accretion disks around Kerr black holes. We use a tilt of $\beta_0 = 15^\circ$ and a spin of $a/M_{\text{BH}} = J_{\text{BH}}/M_{\text{BH}}^2 = 0.5$, in geometrized units where $G = c = 1$, and M_{BH} and J_{BH} are the mass and angular momentum of the black hole, respectively. One of these simulations is run on a spherical–polar mesh, the others on the cubed-sphere grid. We demonstrate that, for the most part, the simulations agree very well.

2. THE CUBED SPHERE

The cubed-sphere grid gets its name from its construction—it is actually composed of six “blocks” that are morphed into segments of a sphere. Each block is constructed of segments of concentric radial shells. In the present work, these shells are spaced exponentially based upon their distance from the hole, similar to a logarithmic radial coordinate. The other two coordinates are constructed such that, on any given block, the grid lines trace out “great circles” on each radial shell segment. It is as if there are two longitude coordinates, ϕ_1 and ϕ_2 , on each block. The range of ϕ_1 and ϕ_2 on each block is $\pi/2$ so that the full 4π steradian is covered by the six blocks.

The difficulty with the cubed-sphere grid is that the “great circles” cannot be made continuous across all six blocks, and hence the block-structured nature of the mesh. Stated differently, the coordinates ϕ_1 and ϕ_2 cannot maintain a consistent orientation across all blocks. At each block boundary, the coordinate system has a discontinuous jump. Fortunately, this can be handled with the proper application of boundary conditions and communication between blocks, as we shall describe.

Another problem with the cubed-sphere grid is that the ϕ_1 and ϕ_2 coordinates are not orthogonal. There are techniques available to try to improve the orthogonality of the cubed-sphere grid at the cost of reducing its uniformity. However, such

techniques have been shown not to perform significantly better than the standard cubed-sphere grid implemented here (Putman & Lin 2007). Furthermore, such techniques are not necessary in our Cosmos++ code, which is designed with tremendous mesh flexibility to handle a variety of grids including fully unstructured and nonorthogonal ones.

2.1. Implementation Within Cosmos++

When working with more traditional spherical–polar meshes, the Cosmos++ code actually evolves the MHD equations in a generalized coordinate system $\{x_0, x_1, x_2, x_3\}$, with the curvature implemented through metric terms. This is done even for the Newtonian formulation. The corresponding physical coordinates in general relativity are the Kerr–Schild polar coordinates $\{t, r, \theta, \phi\}$. For the cubed sphere, instead, we construct the grid in physical space using the Kerr–Schild Cartesian coordinate system $\{t, x, y, z\}$. The two Kerr–Schild coordinate systems are related through the following transformations:

$$\begin{aligned} x &= r \sin \theta \cos \phi - a \sin \theta \sin \phi, \\ y &= r \sin \theta \sin \phi + a \sin \theta \cos \phi, \\ z &= r \cos \theta, \end{aligned} \quad (1)$$

or

$$\begin{aligned} r^2 &= \frac{(x^2 + y^2 + z^2 - a^2) + \sqrt{(x^2 + y^2 + z^2 - a^2)^2 + 4a^2z^2}}{2}, \\ \sin \theta &= \left(\frac{x^2 + y^2}{r^2 + a^2} \right)^{1/2}, \\ \cos \theta &= \frac{z}{r}, \\ \sin \phi &= \frac{ry - ax}{\sqrt{(r^2 + a^2)(x^2 + y^2)}}, \\ \cos \phi &= \frac{rx + ay}{\sqrt{(r^2 + a^2)(x^2 + y^2)}}. \end{aligned} \quad (2)$$

Ultimately the Cosmos++ code just needs to know the coordinate locations of all the zone vertices. From those it is able to fully reconstruct all of the necessary zone properties such as volumes and face areas. We find it easiest for the cubed-sphere grid to start from the cubed-sphere coordinates $\{r, \phi_1, \phi_2\}$ of each vertex, then use the transformations given in the Appendix to convert to the Kerr–Schild polar coordinates $\{r, \theta, \phi\}$, and finally use Equation (1) to obtain the correct Kerr–Schild Cartesian coordinates $\{x, y, z\}$ from the polar ones. For convenience we label the six blocks 0–5, with their orientations described in the Appendix. Samples of blocks 0, 1, and 2 are illustrated in Figure 1.

One consequence of using Kerr–Schild Cartesian coordinates is that, whereas in the spherical–polar case we were able to tilt the black hole with respect to the grid (Fragile & Anninos 2005; Fragile et al. 2007a, 2007b), in the cubed-sphere case the black hole must remain aligned with the grid for a rotating black hole ($a \neq 0$). This is because in Kerr–Schild Cartesian coordinates, the z -axis is chosen to be the spin axis of the black hole and the event horizon is only symmetric about this axis. Thus, in order to get the inner boundary of the cubed-sphere grid to align with the black hole event horizon, the black hole spin axis must align with the grid z -axis.

2.2. Boundary Conditions

As we said, one of the difficulties with the cubed sphere is that the coordinates are not continuous across block boundaries. This




Figure 1. Examples of blocks 0 (right), 1 (center), and 2 (left) that might make up a cubed-sphere grid. Note that in this illustration we use a very low resolution for clarity.

requires some care when setting up communications between blocks. Even once the communication pattern between blocks is established, there are certain subtleties about the grid that must be dealt with. For instance, as we show in Section 3, the gradient operator can only be made to converge properly (second order in our case) if a set of ghost zones are constructed that are an extension of the coordinates on the current block. However, such ghost zones then do not correspond directly to any of the zones on the neighboring block; instead, they tend to straddle more than one zone, and a simple domain exchange is not exactly valid. Fortunately, at any inter-block boundary it is only one of the ϕ coordinates that is discontinuous; the radial coordinate and the other ϕ coordinate are consistent across any given inter-block boundary (Ronchi et al. 1996). Therefore, for a single layer grid with uniform zone spacing, the ghost zones of one block will never overlap more than two zones on the neighboring block. In such a case, we can get away with applying a boundary condition that simply fills the ghost zone with a field F_W that is a weighted average of the fields in the two real zones it overlaps, F_0 and F_1 . The weighted average we use is

$$F_W = \frac{F_0(L - |\mathbf{x}_0 - \mathbf{x}|) + F_1(L - |\mathbf{x}_1 - \mathbf{x}|)}{L}, \quad (3)$$

where

$$L = |\mathbf{x}_1 - \mathbf{x}| + |\mathbf{x}_0 - \mathbf{x}|, \quad (4)$$

and \mathbf{x} , \mathbf{x}_0 , and \mathbf{x}_1 are the coordinate centers of the ghost zone and the two real zones it overlaps on the neighboring block, respectively. This weighting scheme is applied any time a normal domain exchange would be needed between neighboring processors, such as after fields are updated, but before any gradients are taken. Something slightly different must be done for advection as explained in the following section.

2.3. Advection

Because Cosmos++ was written using finite-volume methods, and designed for arbitrary mesh topologies, few changes were needed to apply the code to the cubed-sphere mesh. The one thing (in addition to the ghost zone construction) that was modified, if only slightly, was the algorithm for advection. A number of different options for advection are available in Cosmos++ including upwind, subzonal polyhedral reconstruction and global

monotonic flux-limiter methods, described by Anninos et al. (2005). These methods are designed to operate on multidimensional vector quantities (e.g., gradients) constructed from the convex attributes of arbitrary covariant cell geometries and connectivities. However, for the cubed-sphere mesh we found that flux estimates performed with local one-dimensional-limited projections (or differences) across individual cell faces are generally more robust than computing vector fluxes across the entire cell structure, even with appropriate multidimensional flux limiters.

The method is only slightly modified from that presented by Anninos et al. (2005), so we present only an abbreviated discussion. The advection terms are solved for each evolved field quantity using an upwind time-explicit, first-order forward Euler scheme with appropriately time-centered fluxes. Letting \mathbf{F} represent any of the evolved fields (or their consistent transport counterparts with $\mathbf{F} \rightarrow \mathbf{F}/D$, where D is the mass density), the discrete finite-volume representation of the advection source term can be written as

$$\partial_i(\mathbf{F}V^i) = -\frac{1}{V_z} \sum_f^{\text{faces}} (\mathbf{F}^* V^i A_i)_f, \quad (5)$$

where V_z is the local donor cell volume of zone z , $(A_i)_f$ is the inward pointing area normal vector associated with face f of the donor cell, and $(V^i)_f$ is the face-centered velocity defined as a weighted average across neighboring cells. In Anninos et al. (2005), the quantity $(\mathbf{F}^*)_f$ represents piecewise linearly reconstructed zone-centered fields extrapolated to each cell face by a monotonic Taylor's series expansion, $\mathbf{F}^* = \mathbf{F}_z + (\partial_i \mathbf{F})_z^L (r^i - r_z^i)$, projected from the donor cell center r_z^i to either the face center $r^i = r_f^i$ or the advection control volume center $r^i = r_f^i - (\Delta t/2)(V^i)_f$, over a time-step interval Δt . The zone-centered-limited gradient $(\partial_i \mathbf{F})_z^L$ forces monotonicity in the extrapolated fields using polyhedral subzonal interpolations and control volume integrals to construct upwind, downwind, and centered variations. The difference here, for the cubed sphere, is that the monotonic multidimensional gradient is replaced by a local one-dimensional calculation separately across each donor cell face, and along the direction of the cell face normal (perpendicular to the cell face) using the generalized minmod

limiter in the form

$$\nabla \mathbf{F}_\perp = \left[\frac{1}{2} \left(\frac{a}{|a|} + \frac{b}{|b|} \right) \times \frac{1}{2} \left(\frac{a}{|a|} + \frac{c}{|c|} \right) \times \frac{1}{2} \left(\frac{b}{|b|} + \frac{c}{|c|} \right) \right] \times \min(|a|, |b|, |c|), \quad (6)$$

where $a = (1+\lambda)\nabla \mathbf{F}_D$, $b = \nabla \mathbf{F}_C$, $c = (1+\lambda)\nabla \mathbf{F}_U$, λ is an order parameter between zero and unity specifying the steepness of the applied limiter, and $\nabla \mathbf{F}_U$, $\nabla \mathbf{F}_D$, and $\nabla \mathbf{F}_C$ are the upwind, downwind, and center-difference gradients, respectively. The upwind and downwind gradients are defined as $\nabla \mathbf{F}_{U(D)} = k\delta \mathbf{F}/\delta s$, where $k = \pm 1$ depending on the upwind direction with respect to the coordinate orientation, $\delta \mathbf{F} = \mathbf{F}_z - \mathbf{F}_{\text{opp}}$ is the difference between donor and opposite cell field values, and δs is the magnitude of the distance between donor and neighbor cell centers. The center-difference representation of the gradient is approximated as $\nabla \mathbf{F}_C = \sum_{\text{faces}} (k/2)(\delta \mathbf{F}/\delta s)$, where the sum is over opposite cell face pairs. A projected estimate for the advected fields contributing to the flux in Equation (5) at each cell face is provided by the donor cell as $\mathbf{F}^* = \mathbf{F}_z + \delta \mathbf{F} = \mathbf{F}_z - k\nabla \mathbf{F}_\perp \times \delta r$, where $\delta r = |\vec{x}_f - \vec{x}_z - 0.5\Delta t(V^i A_i)\vec{A}/(A^j A_j)|$ is the distance to the advection control volume center along the direction aligned parallel to the cell face normal vector \vec{A} (between neighbor zone centers).

For advection from one block to another, in order to conserve mass, energy, and momentum to round-off instead of truncation, it is important *not* to interpolate values between ghost zones as was done for the extrapolated field gradients in the previous section. Instead, for advection we use the ghost zones as “buckets” to capture material advecting off of the host block. The mass, energy, and momentum collected in this bucket are then deposited into the corresponding real zone on the neighboring block that shares a face with the originating real zone as part of a final loop in the advection routine. This is appropriate since zones along inter-block boundaries share faces with only a single neighbor.

3. GRADIENT TEST

Because the cubed-sphere grid uses the same basic gradient operators that were already tested in Cosmos++ (Anninos et al. 2005), we fully expect the same second-order convergence for smooth fields, at least in the interior zones. Nevertheless, it is worthwhile to conduct a simple gradient test for a variety of fields to verify second-order convergence over the entire domain, including at the inter-block boundaries where we have introduced a new procedure for interpolation of fields beyond local grid domains.

In our first attempt at implementing the cubed-sphere grid, we actually did not achieve uniform second-order convergence. In that attempt, instead of constructing the ghost zones as extensions of each block as described in Section 2, we constructed the ghost zones to be exact replicas of the nearest zone on the neighboring block and to mimic the behavior of periodic boundaries on spherical–polar grids. However, this introduces a discontinuity into the gradient operator and actually prevents the convergence of gradients at the inter-block boundaries. For interior zones *not touching an inter-block boundary*, we found the L1-normalized error for the gradient of a simple scalar field to converge at second order as expected (the L1-normalized error is defined as $E_1 = \sum_{i,j,k} |a_{i,j,k} - A_{i,j,k}|/(n_i n_j n_k)$, where $a_{i,j,k}$ and $A_{i,j,k}$ are the numerical and exact solutions, respectively, in each zone n_i , n_j , and n_k are the number of zones in each of the

three directions). However, for the interior zones *touching the inter-block boundaries* (not the ghost zones themselves, but the zones that touch them), the L1-normalized error *did not converge*. To explain where this failure arises we first note that the gradient of a generic field F in Cosmos++ is calculated as (akin to Equation (5))

$$G_i = \partial_i F = -\frac{1}{V_z} \sum_f^{\text{faces}} (F^* A_i)_f, \quad (7)$$

where the summation is performed over all cell faces. The problem arises in calculating F^* , the face-centered value of the field; Cosmos++ uses a simple average of the zone-centered values F_z in the two zones adjoining at face f . However, when the line connecting the two zone centers does not pass through the center of the zone face, as is the case for nearest neighbor cells across an inter-block boundary, this simple averaging does not give the correct face-centered value F^* . In fact, it is relatively easy to show in this case that the absolute error ($|a_{i,j,k} - A_{i,j,k}|$) in each zone along the inter-block boundary remains essentially constant, regardless of the resolution (it only depends weakly on the location of the zone along the boundary), thus explaining the nonconvergence in these zones.

The ghost-zone construction described in Section 2, on the other hand, which is the only one used for the remainder of this work, restores second-order convergence in all interior zones by giving a properly extrapolated value for F^* . Here, $F^* = 0.5(F_z + F_W)$ is a simple average of the zone-centered value F_z in the interior zone and the ghost-zone-weighted average F_W from Equation (3). We have confirmed that all interior zones (including those touching the inter-block boundaries) give errors at the level of round-off for flat fields and second-order convergence for all linear and higher order fields.

4. TILTED ACCRETION DISKS

Having demonstrated that our implementation of the cubed-sphere grid preserves the correct convergence order for our code, we can confidently move on to testing our primary application of interest: black hole accretion disks. We begin with a review of how the simulations are initialized and then consider two sets of test cases: in Section 4.1 we study disks of differing alignments relative to a Schwarzschild black hole; in Section 4.2, we compare tilted disk simulations around a Kerr black hole, one carried out on a spherical–polar mesh and the others on the cubed sphere.

Most of the accretion disk simulations presented in this work using the cubed-sphere grid are carried out at a resolution of $128 \times 64 \times 64 \times 6$, where there are 128 radial shells and each of the blocks are resolved with 64×64 angular zones. Along its symmetry planes, such a grid looks like a spherical–polar grid of resolution $128 \times 128 \times 256$. However, the more uniform distribution of zones in the cubed-sphere grid means we are able to achieve such resolution with a smaller number of zones overall (by a factor of 3/4). Also, because of the more uniform zone sizing, we are able to run with a Courant time step that is almost 30 times larger than could be used with a spherical–polar grid of that resolution, which means the required CPU cycle-count is smaller by the same factor. An image of the actual grid used in these simulations is shown in the left panel of Figure 2. This can be compared with the spherical–polar grid used in Fragile et al. (2007b), including the underresolved polar regions, which is shown in the right panel of Figure 2. The time




Figure 2. Left: plot of the cubed-sphere grid geometry used for the disk simulations presented in this work. Right: plot of the spherical–polar grid used in Fragile et al. (2007b), including an under-resolved polar region.

step for the cubed-sphere grid is even 25% larger than for that special grid, where the pole was under-resolved precisely to keep the time step reasonable.

The inner and outer radial boundaries are set at $0.98r_{\text{BH}}$ and $120r_G$, respectively, where r_{BH} is the radius of the black hole horizon and $r_G = GM_{\text{BH}}/c^2$ is the gravitational radius. Note that, because we use the Kerr–Schild form of the Kerr metric, we are able to place the inner radial boundary inside the black hole horizon. In principle, this should keep the inner boundary causally disconnected from the flow, although numerically there is still some communication. At both the inner and outer radial boundaries we apply “outflow” conditions: Fluid variables are set the same in the external boundary zone as in the neighboring internal zone, except for velocity, which is chosen to satisfy

$$V_{\text{ext}}^r = \begin{cases} V_{\text{int}}^r & V^r \text{ points off the grid,} \\ -V_{\text{int}}^r & V^r \text{ points onto the grid.} \end{cases} \quad (8)$$

For the initial conditions of the simulations we start from the commonly used analytic solution for a hydrostatic fluid torus orbiting the black hole. In this case, we choose the torus parameters to be the black hole spin (a/M_{BH}), the inner radius of the torus ($r_{\text{in}} = 15r_G$), the radius of the pressure maximum of the torus ($r_{\text{center}} = 25r_G$), and the power-law exponent ($q = 1.68$) used in defining the specific angular momentum distribution,

$$\ell = -u_\phi/u_t = k\Lambda^{2-q}, \quad (9)$$

where $u_\mu = g_{\mu\nu}u^\nu$, $g_{\mu\nu}$ is the 4-metric, and u^μ is the fluid 4-velocity. We then follow the procedure in Chakrabarti (1985) to solve for the initial state of the torus. Knowledge of r_{center} leads directly to a determination of ℓ_{center} by setting it equal to the geodesic value at that radius. The numerical value of k comes directly from the choice of q and the determination of Λ_{center} , where

$$\frac{1}{\Lambda^2} = -\frac{g_{t\phi} + \ell g_{tt}}{\ell g_{\phi\phi} + \ell^2 g_{t\phi}}. \quad (10)$$

Finally, having chosen r_{in} we can obtain $u_{\text{in}} = u_t(r_{\text{in}})$, the surface binding energy of the torus, from $u_t^{-2} = g^{tt} - 2\ell g^{t\phi} + \ell^2 g^{\phi\phi}$.

The solution of the torus variables can now be specified. The internal energy of the torus is (De Villiers et al. 2003)

$$\epsilon(r, \theta) = \frac{1}{\Gamma} \left[\frac{u_{\text{in}} f(\ell_{\text{in}})}{u_t(r, \theta) f(\ell(r, \theta))} \right], \quad (11)$$

where $\ell_{\text{in}} = \ell(r_{\text{in}})$ is the specific angular momentum of the fluid at the surface and

$$f(\ell) = |1 - k^{2/n} \ell^\alpha|^{1/\alpha}, \quad (12)$$

where $n = 2 - q$ and $\alpha = (2n - 2)/n$. Assuming an isentropic equation of state for the initialization only, the gas pressure and density must be related by the expression $P = \rho\epsilon(\Gamma - 1) = \kappa\rho^\Gamma$, and so the density is given by $\rho = [\epsilon(\Gamma - 1)/\kappa]^{1/(\Gamma - 1)}$. We take $\Gamma = 5/3$ and $\kappa = 0.01$ (arbitrary units). Finally, the angular velocity of the fluid is specified by

$$\Omega = V^\phi = -\frac{g_{t\phi} + \ell g_{tt}}{g_{\phi\phi} + \ell g_{t\phi}}. \quad (13)$$

The torus is then seeded with weak poloidal magnetic field loops with non-zero spatial components $B^r = -\partial_\theta A_\phi$ and $B^\theta = \partial_r A_\phi$, where

$$A_\phi = \begin{cases} b(\rho - \rho_{\text{cut}}) & \text{for } \rho \geq \rho_{\text{cut}}, \\ 0 & \text{for } \rho < \rho_{\text{cut}}. \end{cases} \quad (14)$$

The parameter $\rho_{\text{cut}} = 0.5\rho_{\text{max},0}$ is used to initially keep the field a suitable distance inside the surface of the torus, where $\rho_{\text{max},0}$ is the initial density maximum within the torus. Using the constant b in Equation (14), the field is normalized such that initially $\beta_{\text{mag}} = P/P_B \geq \beta_{\text{mag},0} = 10$ throughout the torus, where P_B is the magnetic pressure. The magnetic field is added in order to seed the magnetorotational instability (MRI; Balbus & Hawley 1991), which is now commonly believed to be the source of angular momentum transport within black hole accretion disks (Balbus & Hawley 1998).

As mentioned in Section 2, our implementation of the cubed sphere requires that the black hole be aligned with the grid. Therefore, unlike our previous work where we tilted the black hole, if we want a tilted configuration now we must tilt the disk. By itself, tilting the disk is a rather trivial operation, simply requiring the following coordinate transformation be applied prior to constructing the torus:

$$\begin{aligned} x' &= x \cos \beta_0 - z \sin \beta_0, \\ y' &= y, \\ z' &= x \sin \beta_0 + z \cos \beta_0, \end{aligned} \quad (15)$$

Table 1
Schwarzschild Simulation Parameters

Simulation	a/M	Tilt Angle	Resolution ^a	Start ^b Time	End ^b Time	\dot{M}^c
0L ^d	0	0	$64 \times 32 \times 32 \times 6$	0	4	-0.0129
0H ^d	0	0	$128 \times 64 \times 64 \times 6$	0	4	-0.0090
015H ^d	0	15°	$128 \times 64 \times 64 \times 6$	0	4	-0.0085
0SP ^e	0	0	128^3	0	4	-0.0141

Notes.

^a In the case of the spherical–polar grid this represents the equivalent peak resolution of an unrefined grid.

^b In units of $t_{\text{orb}} = 785GM/c^3$, the geodesic orbital period at the initial pressure maximum r_{center} .

^c Calculated from the slopes of \dot{M} vs. t over the interval $3 \leq t/t_{\text{orb}} \leq 4$.

^d Cubed-sphere grid.

^e Multiresolution-layer spherical–polar grid.

where β_0 is the initial tilt of the disk. However, as we describe in the following section, we were surprised to discover that our tilted disks evolved differently than our untilted disk, at least at early times, even for nonrotating, Schwarzschild black holes, for which a tilt should have no physical meaning or significance.

4.1. Schwarzschild Black Hole

Here we compare simulations of black hole accretion disks carried out for a Schwarzschild black hole ($a/M_{\text{BH}} = 0$). Table 1 summarizes the parameters for these runs. In our naming convention, the first number indicates the dimensionless spin of the black hole (a/M) without the decimal; the second number, if present, gives the tilt angle of the disk in degrees; the final letter is used to distinguish what resolution the simulation is carried out at, “H” being our high resolution ($128 \times 64 \times 64 \times 6$) and “L” being low ($64 \times 32 \times 32 \times 6$). The two main simulations, 0H and 015H, begin from identical initial conditions except for the tilt of the disk with respect to the grid, which are 0° and 15°, respectively. We also include results of a simulation that uses the spherical–polar grid from Fragile et al. (2007b); this simulation is denoted by the suffix “SP” and has an equivalent peak resolution of 128^3 . We showed in Fragile et al. (2007b) that this was roughly the minimum resolution needed to get a relatively well converged result for this type of problem. Therefore, in the present work, we do not expect our low-resolution simulation (0L) to be converged; it is instead included for the purpose of estimating the rate of convergence when using the cubed-sphere grid.

Our first concern with the cubed-sphere grid is that the angular momentum conservation may not be sufficient for the purpose of following the long-term evolution of an accretion disk, particularly as the flow crosses the coordinate discontinuities at block boundaries. At a minimum we want to quantify our angular momentum conservation error, which we do graphically in Figure 3, where we plot the total angular momentum in each simulation as a function of time for runs 0L, 0H, 015H, and 0SP. The total angular momentum is defined as

$$\int_V T_{\phi}^0 \sqrt{-g} dV, \quad (16)$$

where $T_{\phi}^0 = (\rho h + 2P_B)u^0 u_{\phi} - (B^0 B_{\phi})/(4\pi)$, g is the determinant of the 4-metric, B^{μ} is the magnetic field 4-vector, and

$$h = 1 + \epsilon + \frac{P}{\rho} \quad (17)$$




Figure 3. Plot of the total angular momentum as a function of time for simulations 0H (solid), 015H (dashed), 0L (dot-dashed), and 0SP (dotted). All plots have been normalized to their initial angular momenta. Simulation 0SP uses the spherical–polar grid described in Fragile et al. (2007b).

is the relativistic enthalpy. We only plot part of the first orbital period ($t \leq 0.8t_{\text{orb}}$) of data because after this time significant amounts of angular momentum begin to advect into the black hole and leave the outer boundary of the grid in jets and winds, so that it is much more difficult to track the global conservation. Ideally all the lines in Figure 3 would be perfectly horizontal, indicating exact angular momentum conservation, but we do not really expect this (slight imbalances in the momentum “source” terms and imperfect boundary conditions, for instance, can prevent exact conservation). The angular momentum conservation in our “low” resolution simulation 0L is 0.18% (extrapolated to a full orbital period); this drops down to 0.048% per orbital period at our normal resolution, about the level of convergence (second order) we expect. Furthermore, it appears that this error is not strongly dependent on the orientation of the disk with respect to the grid, based on a comparison of simulations 0H and 015H. Simulation 0SP is included to give some indication of our typical angular momentum conservation error on the multiresolution-layer spherical–polar grid used in our previous work. The error in this case is 0.019% per orbital period, somewhat better but still comparable to simulation 0H, suggesting we suffer only a small degradation in angular momentum conservation in going from our spherical–polar grid to the cubed-sphere grid.

Because we are simulating a nonrotating black hole in this section, any tilt we assign the disk has no physical meaning; it can only be defined relative to the grid. We would expect, therefore, that this tilt would not have any physical effect on the evolution. Interestingly, that is not what we find at early times. The difference is perhaps shown most graphically in Figure 4, which shows the gas density of the disk for simulations 0H and 015H along one azimuthal slice after one orbital period (t_{orb}) at the initial pressure maximum (r_{center}). In simulation 0H (left panel of Figure 4), the disk has spread radially to such an extent that it reaches all the way to the event horizon of the black hole (inner boundary of the computational grid). In simulation 015H, on the other hand (right panel), the disk has hardly spread radially at all, having started at $r_{\text{in}} = 15r_G$ and only penetrated to $12r_G$.

The primary mechanism responsible for the radial spreading of the disk over the first orbital period is not solely the MRI,




Figure 4. Plot of logarithm of density (normalized to $\rho_{0,\max}$) along an azimuthal slice at $\phi = 0$ at $t = 1t_{\text{orb}}$ for simulations 0H (left panel) and 015H (right panel). Because the black hole is not rotating in this simulation, the tilt should have no physical effect and we would expect the two simulations to evolve nearly identically. The observed differences are due to the numerical treatment of the current sheet that forms in the midplane of the disk, as described in the text. (A color version of this figure is available in the online journal.)

but also the differential winding of the initial radial component of the poloidal field loops, the so-called Ω -dynamo. The amplified toroidal and radial field components allow for efficient angular momentum transport essentially from the beginning of the simulation. This is, of course, peculiar to an initial field configuration such as ours which includes a radial field component. If, instead, we started from a purely toroidal field, differential winding would not play a role initially and angular momentum transport would have to await a more complete development of the MRI, which occurs on roughly an orbital timescale.

Something in simulation 015H appears to be shorting out the shear amplification of the field as compared to simulation 0H. Growth of the MRI also appears to be delayed, as evidenced by the less turbulent appearance of simulation 015H in Figure 4. This may be related to the lack of an Ω -dynamo since the MRI has less field to grow on whenever this is inactive (Hawley & Krolik 2002). Furthermore, we can see for certain in Figure 5 that the total magnetic energy is growing more slowly in simulation 015H than in 0H (and OSP). Here we define the magnetic and kinetic energies as

$$\sqrt{-g} \left[(g^{00} + 2u^0 u^0) P_B - \frac{B^0 B^0}{4\pi} \right] \quad (18)$$

and $Dh(u^0 - 1)$, respectively, where $D = W\rho$ is the generalized fluid density with boost $W = \sqrt{-g}u^0$. Both energies are summed over the entire simulation domain. All three simulations show a very rapid initial growth of the magnetic energy due to the combination of shear amplification and the MRI. They also show a gradual increase in kinetic energy over the first orbit as gravitational potential energy is converted into kinetic. After approximately $1t_{\text{orb}}$ the growth of the magnetic fields saturates. At about the same time in simulations 0H and OSP, kinetic energy begins accreting into the black hole in significant amounts, accounting for the sudden change in slope. This happens about an orbit and a half later in simulation 015H.

The culprit for the retarded field growth in simulation 015H appears to be the numerical treatment of the current sheet that forms in the midplane of the disk as a result of the




Figure 5. Plot of the total magnetic (top) and kinetic (bottom) energies as functions of time for simulations 0H (solid), 015H (dashed), and OSP (dotted). All plots have been normalized by the initial kinetic energy of simulation 0H.

differential winding. For an untilted simulation, such as 0H, this current sheet resides almost exactly along an interfacial boundary, right along one of the symmetry planes of the grid (see the left panel of Figure 6). Furthermore, because this is a nearly perfect symmetry plane for the flow, there is very little advection of fluid across this boundary, and so the current sheet remains relatively stationary. In effect, the current sheet remains unresolved, because it spans less than a full zone's width in the vertical direction. Note how narrow the current sheet is in the left panel of Figure 6. This is not the case for a tilted-disk simulation. By necessity, the disk midplane in a tilted disk is no longer aligned with any symmetry plane of the grid (right panel of Figure 6). This means that the disk midplane, and more importantly the midplane current sheet, passes through the interiors of some zones rather than always along their boundary. Numerically this is a critical distinction. For a zone-centered code such as Cosmos++ whenever a current sheet is aligned along an interfacial boundary that experiences no advection, as is approximately the case in our untilted simulations (0H and




Figure 6. Pseudocolor plot representing the value of B^ϕ (in code units) along an azimuthal slice at $\phi = 0$ at $t = 1t_{\text{orb}}$ for simulations 0H (left panel) and 015H (right panel). The midplane current sheet (represented by the line where the color changes from red to blue) remains essentially subzonal in simulation 0H, whereas it is spread across approximately three zones in simulation 015H (greenish yellow zones between red and blue).

(A color version of this figure is available in the online journal.)

OSP), there can be no numerical reconnection and magnetic fields are preserved. If, on the other hand, the current sheet passes through a zone center, as it does in our tilted simulation (015H), numerical reconnection is greatly enhanced. The effect is to drain energy from the magnetic field. In the present work, which uses the internal energy-conserving form of Cosmos++ this energy is simply lost from the simulations (see Fragile & Meier 2008 for a discussion of the implications of the different forms of energy conservation in numerical simulations of black hole accretion disks).

This is a somewhat worrisome discovery; however, we emphasize that it is restricted to the particular field geometry we start with, as no strong midplane current sheet forms if one starts from a purely toroidal field. Furthermore, as the disk becomes more turbulent with the action of the MRI, we find that the discrepancies between the tilted and untilted simulations are dramatically reduced to the point that, at late times, they are nearly indistinguishable. For instance, in Figure 7, we show plots equivalent to Figure 4, except at $t = 4t_{\text{orb}}$ as opposed to $1t_{\text{orb}}$, which show the two disks to be nearly identical. The late-time mass accretion rates are also quite similar (see Table 1).

For a more rigorous comparison, in Figure 8, we present time- and shell-averaged values of density (ρ), gas pressure (P), dimensionless stress (α), plasma magnetization parameter (β_{mag}), specific angular momentum (ℓ), and radial inflow velocity (\bar{V}^r) as functions of radius for simulations 0H, 015H, and OSP at late time, where

$$\alpha = \left\langle \frac{|u^r u^\phi| |B|^2 - B^r B^\phi}{4\pi P} \right\rangle \quad (19)$$

and $\bar{V}^r = \langle \rho V^r \rangle / \langle \rho \rangle$. Angle brackets indicate that a radial shell average has been taken, where

$$\langle \mathcal{Q} \rangle(r, t) = \frac{1}{A} \int_0^{2\pi} \int_0^\pi \mathcal{Q} \sqrt{-g} d\theta d\phi, \quad (20)$$

and $A = \int_0^{2\pi} \int_0^\pi \sqrt{-g} d\theta d\phi$ is the surface area of a given radial shell. The time averaging is done over the interval $3 \leq t/t_{\text{orb}} \leq 4$. The shell averages for P , α , β_{mag} , ℓ , and \bar{V}^r are

mass weighted. Measurements of ρ , P , and ℓ show very good agreement between all three simulations, with errors everywhere $\lesssim 20\%$ and generally much less. The discrepancies in α , β_{mag} , and \bar{V}^r are similarly small for simulations 0H and OSP, but considerably larger for simulation 015H. This is not unexpected as these quantities depend sensitively on the distribution of magnetic field, meaning they are more affected by the delayed growth of the MRI.

4.2. Kerr Black Hole

Having shown that the late-time evolution of simulated black hole accretion disks on our cubed-sphere grid is relatively independent of the orientation of the disk with respect to the grid by analyzing a few Schwarzschild test cases, we can now evaluate the treatment of tilted accretion disks around modestly rotating ($a/M_{\text{BH}} = 0.5$) Kerr black holes. Here our test simulations (515L and 515H), which use the new cubed-sphere grid, are compared to a reference simulation (515SP), which uses the multiresolution-layer spherical-polar grid from Fragile et al. (2007b; shown in the right panel of Figure 2). All simulations have an initial tilt angle $\beta_0 = 15^\circ$. Again we do not expect the low-resolution simulation (515L) to be converged; instead, it is included to provide some indication of the rate of convergence. The parameters for each run are described in Table 2.

First we show in Figure 9 that the general disk properties of simulations 515H and 515SP are quite similar. Again, the largest discrepancies are in the dimensionless stress α and the plasma magnetization parameter $\beta_{\text{mag}} = P/P_B$. This is not surprising since both of these properties have been shown in previous studies to be quite sensitive to resolution (Hawley et al. 1996; Fromang & Papaloizou 2007), and although the total number of zones in these two simulations is comparable, the distribution of those zones is considerably different. The level of agreement in the other parameters is really quite remarkable given the very different structures of the grids. Of course, this was exactly what we were hoping to see.

Now, because the black hole is rotating, the tilt of the disk has some physical meaning and consequently causes changes in




Figure 7. Same as Figure 4, except at time $t = 4t_{\text{orb}}$ instead of $1t_{\text{orb}}$. Here the discrepancies between the untilted (left panel) and tilted (right panel) simulations are greatly diminished.

(A color version of this figure is available in the online journal.)




Figure 8. Main disk properties plotted as a function of radius for simulations 0H, 015H, and 0SP. The data have been time averaged over the final orbital period of each simulation ($3 \leq t/t_{\text{orb}} \leq 4$). P , α , β , l , and V^r are mass-weighted averages of the pressure, dimensionless stress, plasma equipartition parameter, specific angular momentum, and radial inflow velocity, respectively.

(A color version of this figure is available in the online journal.)




Figure 9. Main disk properties plotted as a function of radius for simulations 515H and 515SP. The data have been time averaged over the final two orbital periods of each simulation ($8 \leq t/t_{\text{orb}} \leq 10$). P , α , β , l , and v^r are mass-weighted averages of the pressure, dimensionless stress, plasma equipartition parameter, specific angular momentum, and radial inflow velocity, respectively.

(A color version of this figure is available in the online journal.)

Table 2
Kerr Simulation Parameters

Simulation	a/M	Tilt Angle	Resolution ^a	End ^b Time	\dot{M}^c
515L ^d	0.5	15°	$64 \times 32 \times 32 \times 6$	10	-0.0032
515H ^d	0.5	15°	$128 \times 64 \times 64 \times 6$	10	-0.0114
515SP ^e	0.5	15°	128^3	10	-0.0122

Notes.

^a In the case of the spherical–polar grid this represents the equivalent peak resolution of an unrefined grid.

^b In units of $t_{\text{orb}} = 789GM/c^3$, the geodesic orbital period at the initial pressure maximum r_{center} .

^c Calculated from the slopes of \dot{M} vs. t over the interval $3 \leq t/t_{\text{orb}} \leq 4$.

^d Cubed-sphere grid.

^e Multiresolution-layer spherical–polar grid.

its evolution relative to an untilted disk, as described in Fragile et al. (2007b) and Fragile & Blaes (2008). For instance, although the disk begins with a uniform tilt of $\beta_0 = 15^\circ$, we expect a warp caused by the gravitomagnetic torque of the black hole

to propagate through the disk as a bending wave. This will cause the tilt β to become an oscillating function of radius (Ivanov & Illarionov 1997; Lubow et al. 2002). In Figure 10, we plot $\beta(r)$, time averaged over the interval $8t_{\text{orb}} \leq t \leq 10t_{\text{orb}}$, for simulations 515L, 515H, and 515SP. As in previous work (Fragile & Anninos 2005; Fragile et al. 2007b), we recover the tilt using the definition

$$\beta(r) = \arccos \left[\frac{\mathbf{J}_{\text{BH}} \cdot \mathbf{J}_{\text{Disk}}(r)}{|\mathbf{J}_{\text{BH}}| |\mathbf{J}_{\text{Disk}}(r)|} \right], \quad (21)$$

where \mathbf{J}_{BH} is the angular momentum vector of the black hole and $\mathbf{J}_{\text{Disk}}(r)$ is the angular momentum vector of each radial shell of the simulation domain (dominated by the disk). Again simulations 515H and 515SP produce remarkably similar results, with discrepancies no larger than $\sim 10\%$ and generally much smaller. The discrepancies likely have their root in the small differences in conditions at the inner edge of the disk (see Figure 9) where the bending waves are launched. The simulation 515L exhibits considerably larger discrepancies over most of the disk as expected.




Figure 10. Plot of the tilt β as a function of radius through the disk for simulations 515L, 515H, and 515SP. The data for this plot has been time averaged over the final two orbital periods of each simulation ($8 \leq t/t_{\text{orb}} \leq 10$). The initial tilt was $\beta_0 = 15^\circ$.

Along with warping the disk, the gravitomagnetic torque of the black hole also causes it to precess, particularly in disks such as the ones in our simulations where the fast sound-crossing time causes the disk material to be tightly coupled and respond globally to the torque of the black hole. Global precession of this nature has been noted before in low Mach number hydrodynamic (Fragile & Anninos 2005) and MHD (Fragile et al. 2007b) disks. We track the overall precession (twist), defined as

$$\gamma = \arccos \left[\frac{\mathbf{J}_{\text{BH}} \times \mathbf{J}_{\text{Disk}}}{|\mathbf{J}_{\text{BH}} \times \mathbf{J}_{\text{Disk}}|} \cdot \hat{\mathbf{y}} \right], \quad (22)$$

where \mathbf{J}_{Disk} is the total angular momentum vector of the disk and $\hat{\mathbf{y}}$ is the unit vector that points along the initial line of nodes between the black hole symmetry plane and disk midplane. In order to capture twists larger than 180° , we also track the projection of $\mathbf{J}_{\text{BH}} \times \mathbf{J}_{\text{Disk}}$ onto $\hat{\mathbf{x}}$, allowing us to break the degeneracy in arccos. By plotting the cumulative precession as a function of time as we have done in Figure 11, we make it easy to calculate the precession period of the disk—in this case $0.7(M/M_\odot)$ s, which agrees nicely with our predictions for a black hole of this spin (Fragile et al. 2007b).

5. CONCLUSION

In this paper we have presented our implementation of the cubed-sphere grid within Cosmos++. The cubed-sphere grid has at least three significant advantages over more-traditional grid options: (1) it has topological properties similar to a Cartesian grid, but generally conserves angular momentum much better (and nearly as well as a spherical–polar mesh); (2) it can run at a larger Courant-limited time step than a spherical–polar mesh at comparable resolution (almost a factor of 30 at the resolution used in this work); and (3) it distributes zones more evenly than a spherical–polar mesh, which is desirable for problems where the symmetry is imperfect, such as in tilted accretion disks around rotating black holes, a problem of particular interest to us.

Section 2 and the Appendix give detailed prescriptions for the construction of the cubed-sphere grid and convey a few “lessons learned” in regard to extrapolating fields at inter-block boundaries and applying limiters to the field gradients




Figure 11. Plot of the precession (twist) γ as a function of time for simulations 515L, 515H, and 515SP. The slope of this plot can be used to estimate the precession period of the disk as a whole, which is $0.7(M/M_\odot)$ s.

in advection. After implementing these lessons ourselves, we found we could recover second-order convergence over the entire grid, including along the inter-block boundaries.

To specifically demonstrate that the cubed-sphere grid is a viable option for the black hole accretion disk work we have in mind, we have carried out a series of such simulations on our new cubed-sphere grid, using results from our spherical–polar grid as a reference standard. From these tests we conclude the following.

1. The cubed-sphere grid conserves angular momentum nearly as well as a spherical–polar grid at comparable resolution.
2. The angular momentum conservation error on the cubed-sphere grid is only weakly dependent on the tilt of the disk.
3. Results on the cubed-sphere grid converge to the same solutions obtained on a spherical–polar grid when the two grids approach comparable resolutions. This is true for both untilted and tilted disks.
4. Important disk properties, such as density, pressure, specific angular momentum, inflow velocity, tilt, and twist, agree to better than 10%–20% for simulations carried out on cubed-sphere and spherical–polar grids with roughly $(2\text{--}3) \times 10^6$ zones.

During our testing, we made one surprise discovery—that the early-time evolution was considerably different between our untilted and tilted simulations on the cubed-sphere grid. We found this to be true even for nonrotating Schwarzschild black holes, for which a tilt should have no physical meaning or significance. This is something we had not seen on the spherical–polar grid, but there we had tilted the black hole, not the disk as we do now. We did not anticipate how important this difference would be for the initial growth and development of the Ω -dynamo and MRI.

We attribute the disparate early-time behavior to the differing ways in which the strong initial current sheet in our disk is handled numerically when it is tilted with respect to the grid. This is another reminder of the important role that numerical reconnection plays in the evolution of numerically simulated magnetized flows even though this topic is perhaps not given enough emphasis in the literature. The appearance of current sheets is virtually unavoidable in strongly sheared MHD flows such as accretion disks. One possible technique for treating the current sheets more consistently throughout the simulation

may be to use an artificial resistivity. This would ensure that the current sheets are always resolved in a similar fashion regardless of their orientation with respect to the grid. However, this technique has only been implemented very recently in relativistic MHD (Komissarov 2007). An alternative, although only partial, solution might be to use a total-energy conserving scheme instead of the internal-energy conserving one used here. This, at least, guarantees that the energetics of the flow remain consistent by recapturing in the form of thermal energy any energy lost through magnetic reconnection. When coupled with a radiative cooling package, this can give a much more physical description of the evolution of the flow (Fragile & Meier 2008).

Although the numerical treatments of current sheets and reconnection are important to understand and appreciate, it is equally important in the context of this paper to point out that we demonstrated by numerical example that the long-term evolution of our disks is relatively unaffected by whether or not they are tilted with respect to the grid. As expected, only when the tilt is relative to a rotating black hole are there long-term implications within the disk.

We are not surprised to find significant discrepancies between our “low” and “high” resolution simulations, as previous experience had shown us that 128^3 was roughly the minimum resolution necessary to follow the evolution of black hole accretion disks in global general-relativistic MHD simulations such as these. Below that resolution the characteristic MRI wavelength ($\lambda_{\text{MRI}} \equiv 2\pi v_A/\Omega$, where v_A is the Alfvén speed) is not covered by a sufficient number of zones over much of the disk volume. This has nothing to do with the cubed-sphere grid itself, but is rather a universal constraint for these types of problems.

Overall we consider our experimentation with the cubed-sphere grid to be a success. In future work we will present further analysis of tilted disks (and their associated jets) evolved using this new grid option.

We would like to recognize Joseph Niehaus for his help testing the cubed-sphere grid. We gratefully acknowledge the support of Faculty Research and Development grants from the College of Charleston, SURF, and RPG grants from the College of Charleston 4th Century Initiative Program, and a REAP grant from the South Carolina Space Grant Consortium. A portion of this work was performed under the auspices of the U.S. Department of Energy by Lawrence Livermore National Laboratory under Contract DE-AC52-07NA27344. This work was supported by the National Science Foundation through TeraGrid resources provided by the Texas Advanced Computing Center (TACC). This work also made use of computing resources provided by the Barcelona Supercomputing Center under activity AECT-2007-3-0002.

APPENDIX A

CUBED-SPHERE TRANSFORMATIONS

Included in this Appendix are the transformations necessary to go from the cubed-sphere coordinates $\{r, \phi_1, \phi_2\}$ to the corresponding spherical-polar ones $\{r, \theta, \phi\}$ on each block.

- (1) Block 0 (centered about the $+x$ -axis; $\pi/4 \leq \phi_1 \leq 3\pi/4$; $-\pi/4 \leq \phi_2 \leq \pi/4$):

$$\cos \theta = \frac{\cos \phi_1 \cos \phi_2}{\sqrt{1 - (\cos \phi_1 \sin \phi_2)^2}},$$

$$\begin{aligned} \sin \theta &= \sqrt{1 - \cos^2 \theta}, \\ \sin \phi &= \sin \phi_2, \\ \cos \phi &= \cos \phi_2. \end{aligned} \quad (\text{A1})$$

- (2) Block 1 (centered about the $+y$ -axis; $\pi/4 \leq \phi_1 \leq 3\pi/4$; $\pi/4 \leq \phi_2 \leq 3\pi/4$):

$$\begin{aligned} \cos \theta &= \frac{\cos \phi_1 \sin \phi_2}{\sqrt{1 - (\cos \phi_1 \cos \phi_2)^2}}, \\ \sin \theta &= \sqrt{1 - \cos^2 \theta}, \\ \sin \phi &= \sin \phi_2, \\ \cos \phi &= \cos \phi_2. \end{aligned} \quad (\text{A2})$$

- (3) Block 2 (centered about the $-x$ -axis; $\pi/4 \leq \phi_1 \leq 3\pi/4$; $3\pi/4 \leq \phi_2 \leq 5\pi/4$):

$$\begin{aligned} \cos \theta &= \frac{-\cos \phi_1 \cos \phi_2}{\sqrt{1 - (\cos \phi_1 \sin \phi_2)^2}}, \\ \sin \theta &= \sqrt{1 - \cos^2 \theta}, \\ \sin \phi &= \sin \phi_2, \\ \cos \phi &= \cos \phi_2. \end{aligned} \quad (\text{A3})$$

- (4) Block 3 (centered about the $-y$ -axis; $\pi/4 \leq \phi_1 \leq 3\pi/4$; $5\pi/4 \leq \phi_2 \leq 7\pi/4$):

$$\begin{aligned} \cos \theta &= \frac{-\cos \phi_1 \sin \phi_2}{\sqrt{1 - (\cos \phi_1 \cos \phi_2)^2}}, \\ \sin \theta &= \sqrt{1 - \cos^2 \theta}, \\ \sin \phi &= \sin \phi_2, \\ \cos \phi &= \cos \phi_2. \end{aligned} \quad (\text{A4})$$

- (5) Block 4 (centered about the $+z$ -axis; $-\pi/4 \leq \phi_1 \leq \pi/4$; $-\pi/4 \leq \phi_2 \leq \pi/4$):

$$\begin{aligned} \cos \theta &= \frac{\cos \phi_1 \cos \phi_2}{\sqrt{1 - (\sin \phi_1 \sin \phi_2)^2}}, \\ \sin \theta &= \sqrt{1 - \cos^2 \theta}, \\ \sin \phi &= \frac{\cos \phi_1 \sin \phi_2}{\sin \theta \sqrt{1 - (\sin \phi_1 \sin \phi_2)^2}}, \\ \cos \phi &= \frac{\sin \phi_1 \cos \phi_2}{\sin \theta \sqrt{1 - (\sin \phi_1 \sin \phi_2)^2}}. \end{aligned} \quad (\text{A5})$$

- (6) Block 5 (centered about the $-z$ -axis; $-\pi/4 \leq \phi_1 \leq \pi/4$; $-\pi/4 \leq \phi_2 \leq \pi/4$):

$$\begin{aligned} \cos \theta &= \frac{-\cos \phi_1 \cos \phi_2}{\sqrt{1 - (\sin \phi_1 \sin \phi_2)^2}}, \\ \sin \theta &= \sqrt{1 - \cos^2 \theta}, \\ \sin \phi &= \frac{\cos \phi_1 \sin \phi_2}{\sin \theta \sqrt{1 - (\sin \phi_1 \sin \phi_2)^2}}, \\ \cos \phi &= \frac{-\sin \phi_1 \cos \phi_2}{\sin \theta \sqrt{1 - (\sin \phi_1 \sin \phi_2)^2}}. \end{aligned} \quad (\text{A6})$$

REFERENCES

- Anninos, P., Fragile, P. C., & Salmonson, J. D. 2005, *ApJ*, **635**, 723
- Balbus, S. A., & Hawley, J. F. 1991, *ApJ*, **376**, 214
- Balbus, S. A., & Hawley, J. F. 1998, *Rev. Mod. Phys.*, **70**, 1
- Chakrabarti, S. K. 1985, *ApJ*, **288**, 1
- De Villiers, J.-P., & Hawley, J. F. 2003, *ApJ*, **592**, 1060
- De Villiers, J., Hawley, J. F., & Krolik, J. H. 2003, *ApJ*, **599**, 1238
- Dearborn, D. S. P., Lattanzio, J. C., & Eggleton, P. P. 2006, *ApJ*, **639**, 405
- Dearborn, D. S. P., Wilson, J. R., & Mathews, G. J. 2005, *ApJ*, **630**, 309
- Fragile, P. C., & Anninos, P. 2005, *ApJ*, **623**, 347
- Fragile, P. C., Anninos, P., Blaes, O. M., & Salmonson, J. D. 2007a, arXiv:astro-ph/0701272
- Fragile, P. C., & Blaes, O. M. 2008, *ApJ*, **687**, 757
- Fragile, P. C., Blaes, O. M., Anninos, P., & Salmonson, J. D. 2007b, *ApJ*, **668**, 417
- Fragile, P. C., & Meier, D. L. 2008, *ApJ*, submitted
- Fromang, S., & Papaloizou, J. 2007, *A&A*, **476**, 1113
- Hawley, J. F., Gammie, C. F., & Balbus, S. A. 1996, *ApJ*, **464**, 690
- Hawley, J. F., & Krolik, J. H. 2002, *ApJ*, **566**, 164
- Ivanov, P. B., & Illarionov, A. F. 1997, *MNRAS*, **285**, 394
- Koldoba, A. V., Romanova, M. M., Ustyugova, G. V., & Lovelace, R. V. E. 2002, *ApJ*, **576**, L53
- Komissarov, S. S. 2007, *MNRAS*, **382**, 995
- Lubow, S. H., Ogilvie, G. I., & Pringle, J. E. 2002, *MNRAS*, **337**, 706
- Putman, W. M., & Lin, S.-J. 2007, *J. Comput. Phys.*, **227**, 55
- Romanova, M. M., Ustyugova, G. V., Koldoba, A. V., Wick, J. V., & Lovelace, R. V. E. 2003, *ApJ*, **595**, 1009
- Ronchi, C., Iacono, R., & Paolucci, P. S. 1996, *J. Comput. Phys.*, **124**, 93
- Sadourny, R. 1972, *Mon. Weather Rev.*, **144**, 136
- Zink, B., Schnetter, E., & Tiglio, M. 2008, *Phys. Rev. D*, **77**, 103015