

Multi-Dimensional Finite-Volume Scheme for Hyperbolic Conservation Laws on Three-Dimensional Solution-Adaptive Cubed-Sphere Grids

Lucian Ivan^{*,a}, Hans De Sterck^a, Scott A. Northrup^b, Clinton P. T. Groth^b

^a*Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1*

^b*University of Toronto Institute for Aerospace Studies, Toronto, Ontario, Canada M3H 5T6*

Abstract

A scalable parallel and block-adaptive cubed-sphere grid simulation framework is described for solution of hyperbolic conservation laws in domains between two concentric spheres. In particular, the Euler and ideal magnetohydrodynamics (MHD) equations are considered. Compared to existing cubed-sphere grid algorithms, a novelty of the proposed approach involves the use of a fully multi-dimensional finite-volume method. This leads to important advantages when the treatment of boundaries and corners of the six sectors of the cubed-sphere grid is considered. Most existing finite-volume approaches use dimension-by-dimension differencing and require special interpolation or reconstruction procedures at ghost cells adjacent to sector boundaries in order to achieve an order of solution accuracy higher than unity. In contrast, in our multi-dimensional approach, solution blocks adjacent to sector boundaries can directly use physical cells from the adjacent sector as ghost cells while maintaining uniform second-order accuracy. This leads to important advantages in terms of simplicity of implementation for both parallelism and adaptivity at sector boundaries. Crucial elements of the proposed scheme are: unstructured connectivity of the six grid root blocks that correspond to the six sectors of the cubed-sphere grid, multi-dimensional k -exact reconstruction that automatically takes into account information from neighbouring cells isotropically and is able to automatically handle varying stencil size, and adaptive division of the solution blocks into smaller blocks of varying spatial resolution that are all treated exactly equally for inter-block communication, flux calculation, adaptivity and parallelization. The proposed approach is fully three-dimensional, whereas previous studies on cubed-sphere grids have been either restricted to two-dimensional geometries on the sphere or have grids and solution methods with limited capabilities in the third dimension in terms of adaptivity and parallelism. Numerical results for several problems, including systematic grid convergence studies, MHD bow-shock flows, and global modelling of solar wind flow are discussed to demonstrate the accuracy and efficiency of the proposed solution procedure, along with assessment of parallel computing scalability for up to thousands of computing cores.

Key words:

Magnetohydrodynamics (MHD), Cubed-sphere grid, Space-physics flows, Numerical approximation, Adaptive mesh refinement (AMR).

PACS: 47.11.Df, 47.40.-x, 52.65.Kj, 96.50.Ci

1. Introduction

High-performance computational methods for numerically solving conservation laws in domains between two concentric spheres are desirable in many fields of computational physics. Spherical shell domains are ubiquitously present in computations of global physical processes associated with geophysical and celestial bodies and the intervening space. Consequently, the discretization of spherical shell geometries is required, at least as a first approximation, in fields as diverse as space physics, astrophysics, climate and weather modelling, and geophysics. The

*Corresponding author

Email addresses: livan@uwaterloo.ca (Lucian Ivan), hdesterck@uwaterloo.ca (Hans De Sterck), northrup@utias.utoronto.ca (Scott A. Northrup), groth@utias.utoronto.ca (Clinton P. T. Groth)

choice of the grid is important in formulating an efficient and accurate numerical method for solving partial differential equations on such domains. For example, for space-physics problems, one of the main numerical challenges is provided by the presence of a wide variety of temporal and spatial scales on which interesting plasma physics phenomena occur throughout the vast domains associated with the large-scale space-weather environment. Numerical solutions of the equations arising in the modelling of these complex flows are computationally intensive and feasible only on massively parallel computers [1–4]. Therefore, scalable high-performance algorithms capable of efficiently resolving the solution features of these flows and of reducing the time required to obtain numerical solutions of these problems are actively researched.

In recent years, cubed-sphere grids have gained increasing popularity for simulating fluid flow in domains between concentric spheres, first in the area of climate and weather modelling [5–17], but more recently also in areas like astrophysics [18, 19]. Cubed-sphere grids (see Fig. 1) are attractive because they offer a nearly uniform covering of the spherical surface, while also providing logically Cartesian grids in each of the six sectors (or panels) of the grid, which can be exploited for efficient implementation. Cubed-sphere grids, however, also pose important numerical challenges due to the grid irregularity and non-trivial connectivity along the boundaries and at the corners of the six sectors of the grid. For these reasons, progress in terms of uniform second-order accuracy and higher, adaptivity, and parallelization has been slow.

This paper proposes a new parallel and adaptive cubed-sphere simulation framework for solving conservation laws. Compared to existing cubed-sphere grid algorithms, a novelty of the proposed approach involves the use of a fully multi-dimensional finite-volume method, which leads to important advantages when the treatment of boundaries and corners of the six identical sectors of the cubed-sphere grid is considered. This is in contrast to existing finite-volume schemes on cubed-sphere grids, most of which use dimension-by-dimension differencing and require special interpolation or reconstruction procedures at sector boundaries. Instead, the multi-dimensional finite-volume procedure applied in this work can directly use physical cells from the adjacent sector as ghost cells while maintaining full second-order accuracy and providing important additional advantages in terms of simplicity of implementation, parallelization and adaptivity. The multi-dimensional scheme is implemented into a fully-adaptive three-dimensional (3D) cubed-sphere grid framework, which is uniformly second-order accurate and has excellent parallel scalability on thousands of computing cores. These properties have been obtained by basing the framework on a consistent multi-block approach, with block-based adaptivity, unstructured root-block connectivity, and least-

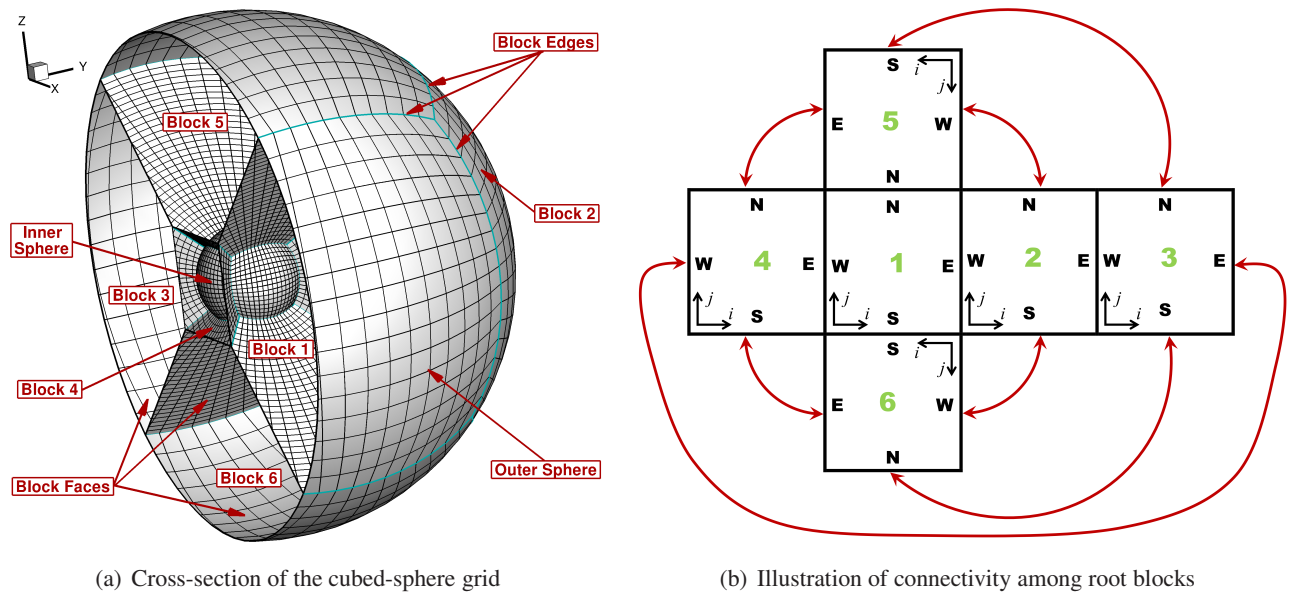


Figure 1: Three-dimensional cubed-sphere grid with six root blocks (corresponding to the six sectors of the grid) and depiction of inter-block connectivity. In our approach, the root blocks can be refined in a block-adaptive way. In panel (b), the block faces are denoted with the initials of the cardinal directions: North (N), East (E), South (S) and West (W).

squares based reconstruction that allows for variable stencil sizes in the upwind finite-volume formulation. We call our multi-block approach *consistent* to highlight that no ad-hoc procedures are employed between blocks at sector boundaries to handle the nontrivial grid connection there. Instead, in terms of the software engineering of the simulation framework, the built-in features of the blocks are made sufficiently rich to deal with sector boundaries automatically. This is achieved by adopting general mechanisms to handle the finite-volume spatial discretization procedure on general hexahedral elements, variable stencil size, unstructured root block connectivity and orientation of coordinate axes in the blocks. This approach results in uniform second-order accuracy, solution adaptivity and parallelization that are treated consistently at the boundaries between the six sectors of the cubed-sphere grid. The proposed scheme provides second-order accuracy on 3D cubed-sphere grids in a uniform manner in all three spatial directions, and naturally permits efficient leading-edge parallel and adaptive implementations on cubed-sphere grids. While the present study has been restricted to obtaining consistent second-order finite-volume discretizations on cubed-sphere grids, future follow-on work will consider high-order discretization up to fourth order. Preliminary results on ongoing work extending the proposed cubed-sphere grid framework to fourth-order accuracy are presented in the recent study by Ivan *et al.* [20].

The computational framework described significantly advances the state-of-the-art of numerical algorithms on cubed-sphere grids because it is fully 3D (with adaptivity and parallelism), whereas existing work is either 2D or has limited capabilities in the third dimension in terms of adaptivity and parallelism. Our scalable, parallel adaptive cubed-sphere grid simulation framework is applied to non-conducting compressible gaseous flow problems and to magnetohydrodynamics (MHD) flows from space physics. While preliminary results on the work presented here were published in an earlier conference paper [21], the present study provides a complete and comprehensive description of the approach.

Two-dimensional (2D) cubed-sphere grids are generated by projecting onto a sphere a Cartesian grid that is defined on the six sides of a cube. In the same fashion, a three-dimensional (3D) cubed-sphere grid can be obtained by overlaying a sequence of concentric 2D spherical shell grids in the radial direction and forming six three-dimensional blocks, each of which is enclosed by the union of four radial and two spherical faces (see Fig. 1). In contrast to the 2D counterpart, which requires the definition of a curved coordinate system for each of the six cubed-sphere sectors, the 3D cubed-sphere grid allows the use of a unique coordinate system (e.g., Cartesian) to discretize the governing conservation laws everywhere in the physical domain, which makes the use of a covariant transformation unnecessary [5, 6, 14, 17, 22]. Note also that, due to the large discrepancies between horizontal and vertical scales, 3D weather and climate models are formulated as layered models, and thus will still require the transformation of information between different coordinate systems of the cubed-sphere grid. However, the space physics flows we target are suitable for a fully-coupled 3D approach, which means that the complication of having to deal with multiple coordinate systems can be avoided: we use a single xyz Cartesian coordinate system throughout the entire simulation domain. The large-scale parallel domain partitioning of 3D cubed-sphere meshes is typically obtained by generating cuts in the radial direction, thereby increasing the number of partitioning blocks that can be farmed out to different processors to $6 \times N_c$, where N_c is the number of radial blocks in each sector. The cubed-sphere grids used in this paper are generated using the so-called cubic-gnomonic projection [5, 6, 10]. In particular, the angularly equidistant mapping described in [6] is used to generate the initial six blocks of the grids.

Our approach incorporates adaptive mesh refinement (AMR) in the multi-block framework. Multi-block grids (see Fig. 2) are grids composed of a number of adjacent logically Cartesian blocks that may differ in resolution. Although local enhancement of flow features has already been widely considered in the context of h -adaptation on 2D latitude-longitude spherical grids by, e.g., Jablonowski *et al.* [23], and more recently on 2D cubed-sphere grids by St-Cyr *et al.* [11], to our knowledge, a fully 3D adaptive mesh refinement procedure on cubed-sphere grids has not been described in the literature before, and the present study represents the first application of a block-based AMR algorithm to 3D cubed-sphere grids. Adaptive mesh refinement is an effective approach for coping with the computational cost of large-scale numerical simulations, such as those encountered in space-physics flows [2, 24–28]. Computational grids that automatically adapt to the solution of the governing equations are effective in treating problems with disparate length scales, providing the required spatial resolution while minimizing memory and storage requirements. The multi-block AMR scheme we use is based on the previous work of Gao and Groth [29, 30] for reacting flows, and has been extended and optimized for cubed-sphere grids in this work. As such,

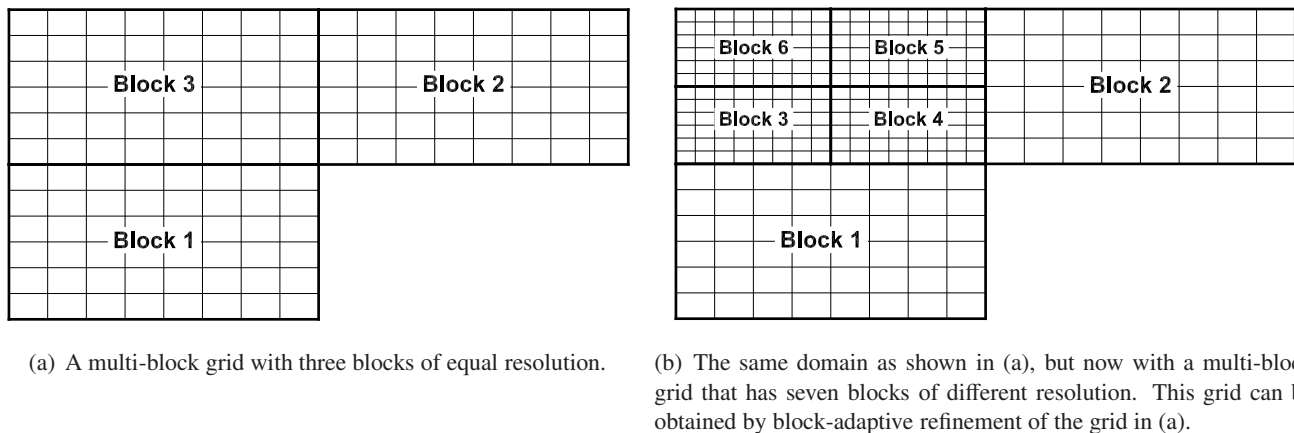


Figure 2: Examples of two-dimensional multi-block grids.

the multi-block AMR finite-volume framework is not restricted to cubed-sphere topologies and more general grid geometries can be treated with the proposed method. Multi-block meshes for cubed-sphere grids require unstructured root-block connectivity. This can be seen in Fig. 1(a): at the projection of the cube corners (i.e., where the root block edges intersect) three root blocks meet, and as a result, the root-block connectivity is as depicted in Fig. 1(b). The connectivity of the root blocks is termed unstructured because it cannot be mapped to a logically Cartesian mesh. The multi-block AMR scheme proposed earlier by Gao and Groth for reacting flows [29, 30] allows for unstructured root-block connectivity and is therefore well-suited for handling the cubed-sphere grid topology.

An important aspect in the development of numerical algorithms on cubed-sphere grids is related to the treatment of interfaces between adjacent grid sectors, across which grid lines have slope discontinuities [5, 6, 9, 14, 19, 22]. If special care is not taken, the accuracy of the numerical discretizations that are based on one-dimensional approaches may be significantly degraded at these interfaces. In traditional implementations of cubed-sphere grid codes the special treatment required at these interfaces for maintaining adequate accuracy poses difficulties for consistent implementation of AMR and parallelization across the computational grid. Several approaches have been proposed in the literature for coupling the solution blocks at sector boundaries, including direct information transfer from interior cells of neighbouring blocks to layers of overlapping “ghost” (or halo) cells [9], one-dimensional interpolation known as *cascade interpolation* which makes use of some particular features of the cubed-sphere grid [6], as well as a one-sided reconstruction followed by high-order Gaussian quadrature to determine the average solution data within each ghost element [14]. We propose a new approach to deal with this accuracy issue at sector boundaries. We use a fully multi-dimensional finite-volume spatial discretization method based on least-squares that directly employs physical cells from the adjacent sector as ghost cells and automatically handles varying stencil size, while maintaining full second-order accuracy.

The remainder of the paper is organized as follows. The next section discusses governing equations of interest. In Section 3, a detailed overview of the proposed block-adaptive cubed-sphere grid simulation framework for conservation laws is provided. In Section 4, the accuracy of the numerical procedure is demonstrated by comparing with analytical solutions and with highly-accurate one-dimensional (1D) transonic flow solutions. A new axi-symmetric MHD test problem is presented that has an exact solution in a domain between two concentric spheres, and can thus be used to test the order of convergence of MHD solvers in this type of domains. Additionally, numerical results are provided to illustrate the capabilities of the adaptive mesh refinement (AMR) algorithm for magnetized and non-magnetized space-plasma problems, such as spherically-symmetric transonic winds and supersonic bow-shock flows past a sphere. Finally, a more realistic space-physics problem is considered by developing a global MHD model similar to the one proposed by Groth *et al.* [2] for the time-averaged or quasi-steady solar wind at minimum solar activity.

2. Governing Equations

We consider the solution of hyperbolic systems of conservation laws of the form given by

$$\partial_t \mathbf{U} + \vec{\nabla} \cdot \vec{\mathbf{F}} = \mathbf{S} + \mathbf{Q}, \quad (1)$$

where \mathbf{U} is the conserved variable solution vector, $\vec{\mathbf{F}}$ is the system flux dyad, and \mathbf{S} and \mathbf{Q} are volumetric source terms. The numerical solution of two prototypical hyperbolic systems will be considered: the Euler and the ideal MHD equations for inviscid compressible non-conducting and conducting ideal gases, respectively. The conserved variable solution vector, \mathbf{U} , for these cases has the form

$$\mathbf{U} = \left[\rho, \rho \vec{V}, \vec{B}, \rho e \right]^T, \quad (2)$$

where ρ is the gas density, $\vec{V} = (V_x, V_y, V_z)$ is the velocity, $\vec{B} = (B_x, B_y, B_z)$ is the magnetic field and ρe is the total energy. The flux dyad, $\vec{\mathbf{F}}$, is given by

$$\vec{\mathbf{F}} = \begin{bmatrix} \rho \vec{V} \\ \rho \vec{V} \vec{V} + \left(p + \frac{\vec{B} \cdot \vec{B}}{2} \right) \vec{I} - \vec{B} \vec{B} \\ \vec{V} \vec{B} - \vec{B} \vec{V} \\ \left(\rho e + p + \frac{\vec{B} \cdot \vec{B}}{2} \right) \vec{V} - (\vec{V} \cdot \vec{B}) \vec{B} \end{bmatrix}. \quad (3)$$

In Eqs. (2) and (3) the specific total plasma energy is $e = p/(\rho(\gamma - 1)) + V^2/2 + B^2/(2\rho)$, where p is the molecular pressure, V is the magnitude of the fluid velocity, and B is the magnitude of the magnetic field. The term $p_B = \frac{\vec{B} \cdot \vec{B}}{2}$ in Eq. (3) is known as the magnetic pressure.


The so-called Powell source term, \mathbf{S} , is given by

$$\mathbf{S} = - \begin{bmatrix} 0 \\ \vec{B} \\ \vec{V} \\ \vec{V} \cdot \vec{B} \end{bmatrix} \nabla \cdot \vec{B} = \mathbf{S} \nabla \cdot \vec{B}, \quad (4)$$

the role of which is not only to make the MHD equations Galilean invariant and symmetrizable, as proved by Godunov [31], but also to provide a numerically stable way of controlling the errors in the divergence free condition, $\nabla \cdot \vec{B} = 0$, a technique first advocated by Powell and co-workers [24]. Note that our framework can also be used with other mechanisms to control $\nabla \cdot \vec{B}$, for example, the generalized Lagrange multiplier mechanism of [32]. However, these alternative approaches are not considered here.

Note that Eqs. (1)-(4) represent the non-dimensional scaled form of the MHD equations following from the non-dimensionalization described by, e.g., Powell *et al.* [24] and by Groth *et al.* [2] for the solar wind. The ideal gas equation of state $p = \rho RT$ is assumed, where T is the gas temperature and $R = 1/\gamma$ is the gas constant. For a polytropic gas (thermally and calorically perfect), the ratio of plasma specific heats, γ , is a constant, and the specific heats are given by $C_v = 1/(\gamma - 1)$ and $C_p = \gamma/(\gamma - 1)$. Unless specified otherwise, di-atomic gases are used throughout this paper, which corresponds to $\gamma = 7/5 = 1.4$. The aforementioned equations can be simplified by setting the magnetic field, \vec{B} , equal to zero, in which case the conservation laws of the non-magnetized plasma reduce to the Euler equations of gas dynamics in non-dimensional form.

The column vector, \mathbf{Q} , appearing in Eq. 1 is used here to represent different volumetric sources arising from the physical modelling of the space-physics problem studied. Thus, the volumetric source term $\mathbf{Q} = \mathbf{Q}_G + \mathbf{Q}_H + \mathbf{Q}_O$ accounts for sources associated with external gravitational fields, \mathbf{Q}_G , effects of coronal heating processes and heat and radiation transfer, \mathbf{Q}_H , as well as other effects, \mathbf{Q}_O . Note that the source term \mathbf{Q}_H is used only for modelling the time-averaged solar wind (see Sect. 4.2.3 for details) whereas the term \mathbf{Q}_O is used to construct an analytical solution to the MHD equations by the method of manufactured solutions [33], as described in Sect. 4.2.1.



(a) Cross-section of a cubed-sphere grid with three refinement levels.

(b) Depiction of octree data structure used to keep track of connectivity among solution blocks.

Figure 3: Illustration of three-dimensional block-based adapted cubed-sphere grids showing the blocks and block boundaries.

3. Parallel Adaptive Second-Order-Accurate Simulation Framework on Cubed-Sphere Grids

3.1. Block-Structured Adaptive Mesh Refinement Approach

A variety of adaptive mesh refinement techniques have emerged in the literature, which can broadly be classified into cell-based, patch-based, block-based and hybrid block-based techniques. A recent review and discussion of these strategies is provided by Gao and Groth in [30]. Among these AMR strategies, block-based AMR approaches lend themselves readily to efficient parallelization on distributed-memory clusters. This feature made the block-based AMR technique attractive for application to different configurations (e.g., on Cartesian, curvilinear and/or body-fitted meshes in conjunction with schemes of various orders of accuracy). AMR has been applied to a variety of engineering and scientific problems by Quirk [34], Berger and Colella [35–37], Gombosi and co-workers [2, 3, 24, 38, 39], Groth and co-workers [30, 40–43], Keppens *et al.* [44] and Jablonowski *et al.* [23], among many others.

The block-based AMR strategy considered in this work follows the approach of Gao and Groth [29, 30] and mesh adaptation is accomplished by dividing and coarsening appropriate solution blocks. To illustrate the general setting for our approach, an example of a 3D block-structured cubed-sphere mesh containing blocks of different resolution levels is depicted in Fig. 3 along with the corresponding hierarchical octree data structure. A refinement step divides a “parent” block into eight “children”, and the reverse happens for coarsening. In general, each block has an equal number of cells for algorithm simplicity and computational efficiency (load balancing). Furthermore, for accuracy reasons and reduced algorithm complexity, the maximum resolution change between adjacent blocks is limited to two. In the adaptation process a grid hierarchy gets generated from a set of initial (i.e., root) blocks in the form of a sequence of nested grids that can be conveniently tracked with a hierarchical block-tree data structure, as the one depicted in Fig. 3(b).

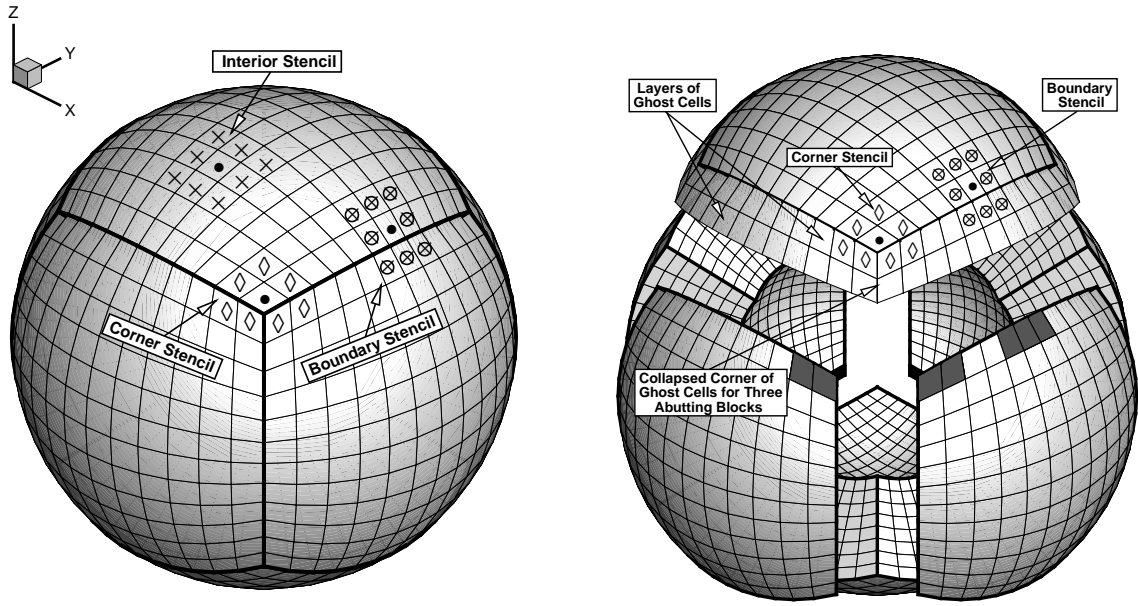
The octree data structure illustrated in Fig. 3(b) contains multiple “roots” and multiple “trees” that keep track of the connectivity and the mesh resolution change between solution blocks. The blocks of the initial mesh represent the roots of the forest. Traversal of the multi-tree structure by recursively visiting the parents and children of solution blocks can be used to determine block connectivity. However, in order to reduce overhead associated with accessing solution information from adjacent blocks, the neighbours of each block are pre-computed and stored, providing direct interconnects between blocks in the hierarchical data structure that are neighbours in the grid. In general, a given solution block in the interior of one of the grid sectors has a minimum of 26 neighbouring blocks (all neighbouring blocks in a $3 \times 3 \times 3$ block stencil), the actual number depending on the level of mesh resolution

change encountered on each of the 26 boundary elements (6 faces, 12 edges and 8 vertices), and having a maximum value of 56 neighbouring blocks. For blocks adjacent to domain boundaries or to cubed sphere corners the number of neighbouring blocks is smaller. In order to facilitate information exchange between adjacent blocks, additional layers of “ghost” cells that overlap with interior cells of adjacent blocks are added to each solution block. The numerical scheme used in this work (see Sect. 3.3), handles consistently the computational cells near block interfaces by making use of the overlapping ghost cells and accommodating variable stencil sizes. For our second-order accurate method we employ two ghost cell layers, which represents a small memory overhead relative to the benefits to be gained when the parallel aspects of the algorithm are considered, as described below. One of the advantages of the hierarchical octree data structure is that it readily permits local mesh refinement. Local modifications to the multi-block mesh can be performed without re-gridding the entire mesh and re-calculating all solution block connectivities.

3.2. Multi-Block Cubed-Sphere Grid with Unstructured Root-Block Connectivity

The main technical difficulty in applying the block-structured adaptive multi-block concept to cubed-sphere grids is to deal consistently with the unstructured connectivity between adjacent blocks that occurs at sector boundaries and sector corners. The first issue arises at sector corners (see Fig. 4): grid cells adjacent to one of the eight sector corners have only seven neighbouring cells (in 2D), while all other cells have 8 neighbours (these neighbours are used in stencils for solution reconstruction and flux calculation, see next section). This issue is dealt with in our approach by automatically detecting blocks with such corner cells, and by assigning “collapsed” corner ghost cells to those blocks sharing the relevant corner (as in [30]). In practice, this is implemented by tagging those collapsed corner ghost cells as inactive and setting default values in the memory assigned to them so as to maintain consistency of the memory layout for solution blocks. Note that in one sector corner each of the three blocks has a number of collapsed cells equal to four times the number of cells in the radial direction (i.e., interior cells plus four ghost cells). This shows that the memory overhead of “storing” collapsed cells is only dependent on the radial resolution along sector corners and remains constant as more cells are added in the latitudinal and longitudinal directions. Furthermore, the memory assigned to collapsed cells becomes even less significant relative to the total memory requirements as more blocks are added in the angular directions by the AMR process, as blocks interior to cubed-sphere grid sectors do not require this type of ghost cell treatment. These collapsed ghost cells are not used in the stencils for reconstruction computation, so grid cells adjacent to sector corners employ reduced stencil sizes. The flexible finite-volume reconstruction mechanism of the proposed solution method (see the next subsection) handles this consistently across block cells without reducing the local order of accuracy. In our implementation, blocks with collapsed ghost cells are detected from the block connectivity data structure: if no neighbouring block is found in the direction of a block corner the corresponding corner ghost cells are taken to be collapsed.

The second issue is the connectivity between adjacent blocks (recall Fig. 1(b)): on cubed-sphere grids, it is no longer true that, for example, eastern block boundaries are always adjacent to western block boundaries, and it is necessary to carefully keep track of the block interface adjacency, and of the orientation of index axes in adjacent blocks (i.e., the i , j , and k indices in the logically Cartesian data structure of a block may correspond to a different orientation of the i , j , and k indices in a neighbouring block). We use the unstructured root-block approach of [30] to handle this issue in a systematic way. The logic employed follows the methodology proposed in the Computational Fluid Dynamics General Notation System (CGNS) [45]. A typical refined block stores the relative orientation of the neighbouring blocks in the directions of all 26 boundary elements (6 faces, 12 edges and 8 vertices), and for each neighbour the orientation of the i , j , and k axes relative to the orientation of the i , j and k axes in the current block is stored in compact form as a three-component transformation array. For example, consider the index system of block 1 in Fig. 1(b) in relationship to its south neighbour, block 6. The i and j indices of block 6 run in directions opposite to the i and j indices in block 1. This is indicated by negative signs in the block 1-to-block 6 transformation array, which is given by $[-1, -2, +3]$. The components 1, 2 and 3 in the array signify the i , j and k indices in block 6. The value 1 in the first component of the array means that the i index from block 6 is associated with the first index of block 1 (its i index), and the negative sign indicates that they run in opposite directions. The value -2 in the second position means that the j indices of the two blocks are also associated and running in opposite directions, and the value $+3$ in the third direction indicates that the k indices have the same orientation. Similarly, considering



(a) Compact view of a cubed-sphere grid and depiction of cells participating in reconstruction stencils in different regions. The data from the cells residing on different blocks than the reconstructed cell is locally reproduced by using overlapping layers of ghost cells, as shown in (b).

(b) Exploded view of the cubed-sphere grid shown in (a) illustrating the ghost cell layers for one block and the actual cells forming the reconstruction stencils. The interior cells marked with shaded gray are the cells from adjacent blocks which have their data duplicated in the ghost cells.

Figure 4: Examples of interior, boundary and corner reconstruction stencils for different cells of a cubed-sphere mesh with six $12 \times 12 \times 8$ blocks and a total of 6,912 cells. The cell of which the solution is reconstructed is marked with \bullet symbol and the neighbouring cells that are part of the reconstruction stencil are marked with \times , \diamond and \otimes symbols for the interior, corner and boundary reconstruction stencil, respectively.

block 2, the block 2-to-block 5 transformation array is given by $[-2, +1, +3]$, because the second index in block 5 (j) corresponds to the first index in block 2 (i) (running in the opposite direction), the first index in block 5 (i) corresponds to the second index in block 2 (j) (running in the same direction), and the third index in block 5 (k) corresponds to the third index in block 2 (k) (running in the same direction). These transformation arrays provide a convenient short-hand notation for the transformation matrices [29, 30, 45] describing the relation between indices of two adjacent blocks, which can be used to exchange solution information between blocks having common interfaces in a general and consistent way. For more details, see [30].

The connectivity information is propagated from the root blocks to refined blocks via the octree data structure, in such a way that each refined block stores a transformation array describing index axis alignment with all of its (typically 26) neighbour blocks. These block-to-block transformation arrays are used in the solution procedure to properly compute numerical fluxes through the block boundaries (via the ghost cell and reconstruction mechanisms, see below). Note that the transformation array mechanism is implemented uniformly for all blocks to maintain algorithmic consistency, but it only results in nontrivial action at block-block interfaces along sector boundaries.

3.3. Multi-Dimensional Second-Order Finite-Volume Scheme with Uniform Treatment of Sector Boundaries and Corners

Numerical solutions of Eq. 1 on the cubed-sphere grids are sought by applying a Godunov-type finite-volume spatial discretization procedure [46] in conjunction with second-order polynomial reconstruction and Riemann-solver based flux functions. Thus, the set of coupled ordinary differential equations (ODEs) resulting from the application of this finite-volume discretization to Eq. 1 for cell (i, j, k) of a multi-block mesh composed of hexahedral

computational cells is given by

$$\frac{d\mathbf{U}_{i,j,k}}{dt} = -\frac{1}{V_{i,j,k}} \sum_{m=1}^{N_f} (\vec{\mathbf{F}} \cdot \vec{\mathbf{n}} \Delta A)_{i,j,k,m} + \bar{\mathbf{S}}_{i,j,k} + \bar{\mathbf{Q}}_{i,j,k} = \mathbf{R}_{i,j,k}(\mathbf{U}), \quad (5)$$

where $V_{i,j,k}$ is the cell volume, N_f is the number of cell faces, and $\vec{\mathbf{n}}$ and ΔA are the unit outward normal vector and the area of cell face m , respectively. The numerical fluxes, $\vec{\mathbf{F}} \cdot \vec{\mathbf{n}}$, at the midpoint of each face of cell (i, j, k) are determined from the solution of a Riemann problem. Given the left and right interface solution states, \mathbf{U}_l and \mathbf{U}_r , an upwind numerical flux is evaluated by solving a Riemann problem in the direction defined by the normal to the face. In our computational studies, the Lax-Friedrichs [47], HLLC [48], and modified HLLC [49] approximate Riemann solvers are used in evaluating the numerical flux. The left and right initial states for the Riemann problem, \mathbf{U}_l and \mathbf{U}_r , are determined based on the associated primitive variables, $\mathbf{W}_{l,r} = [\rho, \vec{V}, \vec{B}, p]_{l,r}$, at the interface. These primitive variable states \mathbf{W}_l and \mathbf{W}_r at the cell interfaces are obtained by performing the multi-dimensional limited piecewise linear solution reconstruction proposed by Barth [50] in conjunction with the Venkatakrishnan [51] slope limiter, so as to avoid the occurrence of spurious oscillations and non-physical solutions (i.e., negative density and/or pressure). Thus, the limited linear representation of the primitive solution for cell (i, j, k) is written as a polynomial expansion around the cell centroid, $\vec{x}_{i,j,k}$, as

$$\mathbf{W}_{i,j,k}(\vec{x}) = \bar{\mathbf{W}}_{i,j,k} + \Phi_{i,j,k} \vec{\nabla} \mathbf{W} \cdot (\vec{x} - \vec{x}_{i,j,k}), \quad (6)$$

where $\Phi_{i,j,k}$ is the vector of limiter values that are determined individually for each solution variable, and $\bar{\mathbf{W}}_{i,j,k}$ is the cell average solution vector. The gradients of the primitive variables, $\vec{\nabla} \mathbf{W}$, must be determined for each cell using the local solution data and follow from the least-squares solution of the unlimited (i.e., $\Phi_{i,j,k} = \vec{1}$) linear reconstruction based on a three-dimensional supporting stencil which generally includes 27 cells in total (i.e., the $3 \times 3 \times 3$ block composed of the reconstructed cell plus the 26 nearest neighbouring cells). For each solution variable reconstruction, $W_{i,j,k}(\vec{x})$, three unknowns (i.e., the x -, y - and z -direction components of $\vec{\nabla} \mathbf{W}$) are calculated by imposing 26 equations of the form

$$\bar{W}_{\gamma,\delta,\zeta} = \frac{1}{V_{\gamma,\delta,\zeta}} \iiint_{\mathcal{V}_{\gamma,\delta,\zeta}} W_{i,j,k}(\vec{x}) dV, \quad (7)$$

which expresses the condition that the average of the reconstructed polynomial for cell (i, j, k) over each neighbouring cell (γ, δ, ζ) in the reconstruction stencil is required to recover the existing volume-averaged solution, $\bar{W}_{\gamma,\delta,\zeta}$, in that cell. The resulting overdetermined system of linear equations is solved in the least-squares sense [50]. The limiter values are then determined based on the reconstructed solution at the midpoint of each face (i.e., at the flux integration points) by applying the selected slope limiter. The reconstruction procedure could in principle be carried out with smaller reconstruction stencils (for example, a seven-point reconstruction stencil would lead to six equations), but it would then lose much of its multi-dimensional character, and we have indeed found that the 27-point reconstruction stencil provides better robustness against non-uniformity in the grid and solution gradients that are not aligned with the grid. The least-squares procedure allows for flexible stencil sizes, without reducing the local order of accuracy. In particular, for grid cells adjacent to the radial lines that pass through cubed-sphere sector corners, the reconstruction stencil size is reduced from 27 to 24, which is handled naturally and consistently by the least-squares reconstruction procedure.

Figure 4 shows the computational cells belonging to a spherical face that are included in the reconstruction stencil of representative cells that lie either in the interior, at a boundary, or at a corner of the cubed-sphere sectors. The complete 3D reconstruction stencils are obtained by including the radial neighbours (i.e., above and below) of the cells depicted in Fig. 4(a) and (b). The overlapping layers of ghost cells depicted in these figures are used to mirror the information from the neighbouring block and compute updates in cells adjacent to the block boundary. Note that the stencils shown in the figures are reconstruction stencils, i.e., the stencils used for computing the linear reconstruction in a single cell. These reconstruction stencils contain one layer of neighbours. The actual computational stencils of our method (the cells on which a cell depends for its update) contain cells from two layers of neighbours. In our approach two layers of ghost cells are used. Solution reconstruction is performed in interior

cells but also in the first layer of ghost cells, using the solution values in the second layer of ghost cells. This means that only one communication step is required per time integration stage. Without the second layer of ghost cells, a second communication step would be required in each stage of the time integration, in between the solution reconstruction computation and the flux calculation, since reconstructed values are required at both sides of the block boundaries to compute the fluxes. The multi-dimensional reconstruction procedure is observed to perform accurately and robustly on cubed-sphere grids (see the convergence results below), avoiding any special interpolation treatment of cells near grid sector boundaries. Note that this is in contrast to other previous finite-volume algorithms on the cubed sphere, which require either one-dimensional interpolation [6, 19] or one-sided reconstruction [5, 14] near sector boundaries to obtain order of accuracy greater than one.

The numerical procedure outlined above results in a second-order accurate upwind finite-volume scheme. The average value of the Powell source term, $\bar{\mathbf{S}}_{i,j,k}$, of cell (i, j, k) is computed with second-order accuracy by estimating the quantity \mathcal{S} based on the appropriate cell-averaged solution variables (i.e., the value of the reconstructed polynomial at the cell centroid) and multiplying it with the following discretization for $\nabla \cdot \vec{\mathbf{B}}$:

$$(\nabla \cdot \vec{\mathbf{B}})_{i,j,k} = \frac{1}{V_{i,j,k}} \sum_{m=1}^{N_f} (\vec{\mathbf{B}}_f \cdot \vec{n} \Delta A)_{i,j,k,m}, \quad (8)$$

where $\vec{\mathbf{B}}_f$ is an interface magnetic field computed as the arithmetic mean of the left and right reconstructed values, that is, $\vec{\mathbf{B}}_f = (\vec{\mathbf{B}}_l + \vec{\mathbf{B}}_r)/2$. The contribution of all other volumetric sources to the solution residual, $\mathbf{R}_{i,j,k}$, is also evaluated with second-order accuracy by computing the average source term $\bar{\mathbf{Q}}_{i,j,k}$ at the centroid of cell (i, j, k) . The test problems considered in this paper are all steady-state problems. To compute the steady-state solutions, the coupled system of nonlinear ODEs given by Eq. 5 is solved using the explicit five-stage scheme from [52] that was derived to provide optimum damping for second-order upwinding. This scheme achieves fast convergence for steady-state simulations. Since we compute steady-state solutions, we can use local time-stepping, in which every cell employs the maximal time-step that obeys the Courant-Friedrichs-Lewy (CFL) stability condition. The steady-state nature of the solution was determined based on the reduction of the L_2 -norm of the solution residual over all grid cells.

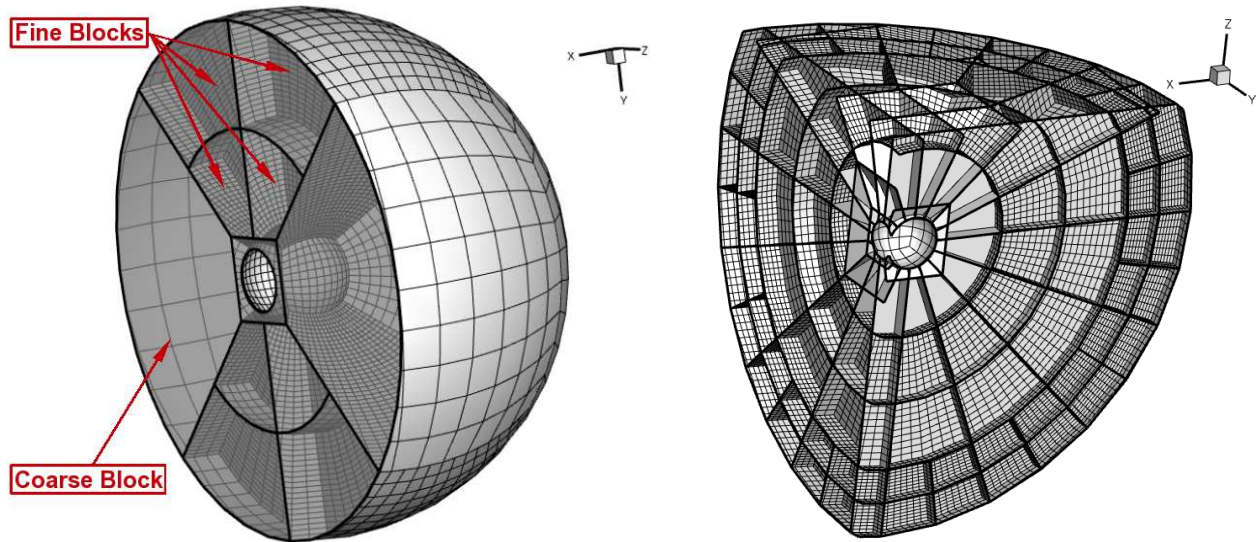
While the explicit time-marching scheme applied here is not optimal, it was sufficient for the steady-state problems of interest herein. Future follow-on research will consider improving the computational efficiency of the cubed-sphere framework by combining the proposed spatial discretization procedure with more effective time-marching approaches such as the parallel implicit Newton-Krylov-Schwarz algorithm of Northrup and Groth [53, 54].

3.4. Block-Based Adaptive Refinement and Coarsening Strategy

The flexible block-based hierarchical data structure discussed in Sect. 3.1 is used in conjunction with the spatial discretization procedure outlined above to facilitate automatic solution-directed mesh adaptation on 3D cubed-sphere grids. Mesh adaptation is accomplished by refining and coarsening appropriate solution blocks. In the case of refinement, each of the eight octants of a parent block becomes a new block having the same number of cells as the parent, thereby doubling the cell resolution in the region of interest. The additional mesh nodes in the new fine blocks are inserted in the middle between the nodes inherited from the parent. This process can be reversed in regions that are deemed over-resolved and eight children are coarsened into a single parent block, the mesh nodes of which are obtained by eliminating every other interior node from the children blocks. Mesh refinement criteria are used to decide on adaptive resolution changes in a procedure that can be summarized as follows (see [30, 53] for a more extensive discussion).

Local refinement and coarsening of the mesh is directed according to so-called physics-based refinement criteria. In particular, for the test cases considered in this paper, density gradient and divergence and curl of velocity are used in the decision to refine or coarsen a solution block. For definiteness, it is discussed here in the case where the density gradient is used as the basis for refinement.

When refinement and coarsening are desired in the course of a simulation, the following procedure is carried out. First, the maximum per-block value of the density gradient in each cell is computed, and then a normalized scale for the variation of the refinement measure per block is built based on the global minimal and maximal values



(a) Cross-section of a cubed-sphere grid with one refinement level (b) Complex cut into an adapted cubed-sphere grid with a few refinement levels

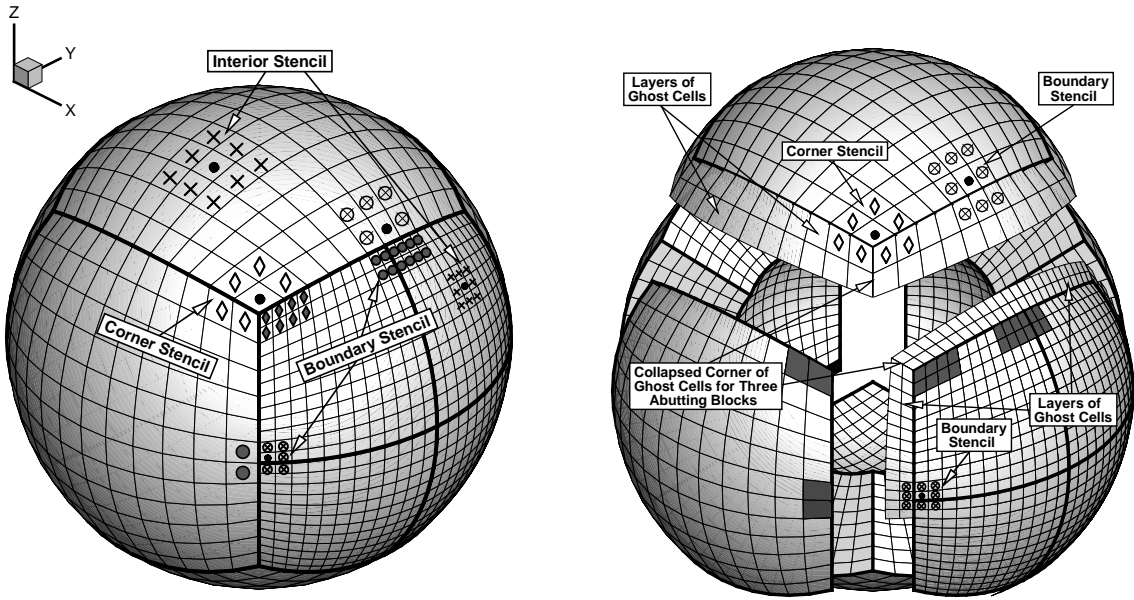
Figure 5: Illustration of three-dimensional block-based adapted cubed-sphere grids showing the block boundaries and the associated meshes.

for all solution blocks. Based on user-specified threshold values for refinement and coarsening on a scale between zero to one, blocks with maximum density gradient values above the refinement and below the coarsening thresholds are flagged for refinement and coarsening, respectively. All the other blocks are flagged to keep their current mesh resolution. The final decision about which blocks get refined/coarsened or are left unchanged employs a check for possible violations of the condition that adjacent blocks may not differ by more than a factor of two in resolution after refinement/coarsening. The conflicts are resolved according to the algorithm in [29, 30], in which precedence is given to block refinement. After these conflicts are resolved, the appropriately flagged blocks are refined and coarsened as scheduled. The newly created blocks during refinement are distributed evenly among computing cores (for load balancing) in accordance to a resource list mirrored on all computing cores which stores information about the distribution of solution blocks and the available computational resources.

For the steady-state computations to be presented in the numerical results section of this paper, this refinement/coarsening mechanism is invoked at regular intervals during the convergence history. Our implementation also allows for time-dependent, dynamical refinement and coarsening on the cubed-sphere grid, simply by applying the procedure at regular intervals during time-dependent computations. Note that blocks at sector interfaces and sector corners of the cubed-sphere grid are handled by the refinement/coarsening mechanism in a fully consistent way. For example, when refined blocks are created, their neighbours are determined and blocks adjacent to sector corners receive less than the full amount of neighbouring blocks (due to “missing” corner blocks), and will automatically be assigned collapsed corner ghost cells. The transformation arrays for these blocks take into account the unstructured connectivity automatically.

Two illustrative examples of the application of the AMR algorithm described above to cubed-sphere grids are shown in Fig. 5. Note also that a body-fitted mesh is readily obtained for the cubed-sphere approach by simply ensuring that newly inserted mesh nodes at the inner and outer spherical boundaries belong exactly to the appropriate spherical shell.

Standard multi-grid-type restriction and prolongation operators [55] are used to evaluate the solution on all blocks created by the coarsening and division processes, respectively. A coarse cell has mean solution data calculated from the corresponding eight fine cells it overlaps with in a way that satisfies conservation [37, 55]. In particular, each mean solution variable of the coarse (parent) grid cell is determined by summing the corresponding eight mean solution variables of the fine grid cells weighted by the fraction of the fine grid cell volume to the coarse grid cell volume. The volume-averaged solution of a fine cell is derived from the solution of a parent coarse cell by simply



(a) Compact view of a refined cubed-sphere grid and depiction of cells participating in reconstruction stencils in different regions. The data from the cells residing on different blocks than the reconstructed cell is reproduced by using overlapping layers of ghost cells, as shown in (b).

(b) Exploded view of the refined cubed-sphere grid shown in (a) illustrating the ghost cell layers for one coarse and four refined blocks. Additionally, the actual cells forming the reconstruction stencils are shown for both low- and high-resolution blocks. The interior cells marked with shaded gray are the cells from adjacent blocks which provide data to the ghost cells at different transitions of mesh resolution (by duplication for no resolution change, by prolongation for coarse-to-fine change, and by restriction for fine-to-coarse change).

Figure 6: Examples of interior, boundary and corner reconstruction stencils on a refined cubed-sphere mesh with five coarse sectors and one refined sector (all 13 blocks in the grid are $12 \times 12 \times 8$, and there are 14,976 total cells). The cell of which the solution is reconstructed is marked with \bullet symbol and the neighbouring cells that are part of the stencil are marked with \times , \diamond and \otimes symbols for the interior, corner and boundary reconstruction stencil, respectively.

evaluating the limited linear reconstruction of the parent at the centroid of the fine cell. Owing to the linearity of the reconstruction, mass is conserved in this process, as can be shown using the mean value theorem. The same restriction and prolongation operators are also used for ghost cells at block interfaces where the resolution changes, in all time steps (see Fig. 6). This procedure directly influences the domain of dependence of solution reconstructions for cells that include ghost cells in their supporting stencils. Thus, even though reconstruction stencils are always locally constructed in the fashion described in Sect. 3.3, the actual domain of dependence depends on the mesh resolution transition at block interfaces. As can be seen from Fig. 6(b), the stencil construction algorithm outlined in Sect. 3.3 can be applied in the same way to each of the cubed-sphere blocks regardless of their spatial resolution.

Within the AMR approach, additional inter-block communication is also required at interfaces with resolution changes to strictly enforce the flux conservation properties of the finite-volume scheme [36, 37]. In particular, the interface fluxes computed on more refined blocks are used to correct the interface fluxes computed on coarser neighbouring blocks and ensure that the solution fluxes are conserved across block interfaces.

3.5. Parallelization with Uniform Treatment of Sector Boundaries and Corners

With the mechanisms of unstructured root-block connectivity and multi-dimensional reconstruction with flexible stencil sizes in place, parallelization can be performed in a way that is fully consistent at sector boundaries and corners. As in the other previous work of Groth *et al.*, e.g., [30, 40–43], an efficient domain partitioning is achieved in our implementation by distributing the active solution blocks equally among available processor cores, with more than one block permitted per processor core. This approach efficiently exploits the self-similar nature of the solution

blocks and readily produces an effective load balancing. Inter-processor communication is mainly associated with block interfaces and involves the exchange of ghost-cell solution values and conservative flux corrections at every stage of the multi-stage time integration procedure. To improve the efficiency of the parallel communication, message passing of the ghost-cell values and flux corrections is performed by consolidating inter-processor messages. To further reduce the overall communication cost among blocks it may be advantageous to place nearest-neighbour blocks on the same or nearest processor, which is usually realized by utilizing space-filling curves [56]. In the proposed solution method, Morton ordering [57] has been implemented but it is not used here based on considerations that, for the steady-state flow simulations and explicit time-marching schemes considered, the communication cost is relatively low and the benefits of using a space-filling curve would amount to a maximum of 5-10% reduction in execution time. Note however that for simulation of unsteady flows in which block refinement and coarsening occur more frequently, Morton ordering has the potential to generate more significant computational savings. The number of CPU cores is fixed in advance, but it is worth noting that our flexible parallelization model allows for situations in which the number of CPU cores is larger than the number of solution blocks. This is typically encountered during initial stages of AMR simulations where a small number of coarse blocks are adapted until enough new blocks are created to fully populate the available number of computing cores as evenly distributed as possible, providing overall a well-load-balanced simulation.

Blocks adjacent to grid sector corners feature “collapsed” corner ghost cells, and since there is no neighbour block associated with the collapsed ghost cells in the data structure, messages are not sent for collapsed ghost cells. Blocks adjacent to grid sector boundaries may have local i, j, k index systems that are oriented differently from their neighbouring blocks across the grid sector boundary, but this is taken into account naturally and consistently when filling in ghost cells from messages sent by the neighbouring blocks, using the transformation array mechanism described above. In this way, all solution blocks are treated equally in a way that facilitates consistent parallelization across grid sector boundaries and corners. The combination of this scalable domain partitioning and the effective AMR-based block-multiplication procedure have allowed us to perform efficient parallel calculations on 3D cubed-sphere grids with in excess of 3,000 computing processor cores. A quantitative demonstration of the parallel performance is provided for a representative space-physics problem in Sect. 4.2.3.

4. Numerical Results

A set of numerical results demonstrating the accuracy and capabilities of the proposed cubed-sphere framework is now described for a range of flow problems with magnetized and non-magnetized plasma. Results for both fixed and AMR meshes are included. All computations were performed on a high-performance parallel cluster consisting of 3,780 Intel Xeon E5540 (2.53GHz) nodes with 16GB RAM per node. The cluster is connected with a high-speed, low-latency, non-blocking 4x-DDR InfiniBand switched fabric communications link.

4.1. Results for Non-Conducting Fluid

4.1.1. Fully Supersonic Flow

To assess the accuracy of the finite-volume scheme on cubed-sphere grids, uniform convergence studies with a spherically symmetric expansion of a supersonic non-magnetized gas flow (without gravitational field) are performed. Since the cubed-sphere grid is not spherically symmetric, this problem provides a relevant test case for our implementation (in contrast, this problem is not interesting for a latitude-longitude (or spherical-polar) grid). The computational domain of this test case is defined by inner and outer spheres of radius $R_i = 1$ and $R_o = 4$, respectively. With V_r the radial velocity and \vec{V}_{\parallel} the velocity parallel to the inflow boundary, the boundary conditions at R_i are $\rho_i = 10$, $V_{r,i} = 4.5$, $\vec{V}_{\parallel,i} = 0$, and $p_i = 26$. The analytical solution of this flow problem can be obtained in spherical coordinates by integrating the Euler equations or by manipulating algebraically the flow invariants (i.e. mass flow rate, entropy and total enthalpy). Thus, the radial velocity, V_r , at any spherical radius, r , is obtained numerically by solving the non-linear expression

$$C_3 - \frac{1}{r^2 V_r \left[(C_2 - V_r^2)^{\frac{1}{\gamma-1}} \right]} = 0, \quad (9)$$




Figure 7: L_1 , L_2 , and L_∞ error norms in the predicted solution density for a supersonic expansion.

for V_r , where C_2 and C_3 are flow constants that can be determined based on the inflow parameters ρ_i , p_i and $V_{r,i}$:

$$C_3 = \frac{1}{\left(\frac{2\gamma}{\gamma-1} \frac{p_i}{\rho_i} \frac{1}{\gamma-1}\right) R_i^2 V_{r,i}}$$

$$C_2 = \frac{2\gamma}{\gamma-1} \frac{p_i}{\rho_i} + V_{r,i}^2. \quad (10)$$

Density and pressure as a function of radius can then be obtained from Eq. 10.

The L_1 , L_2 , and L_∞ norms of the error in the predicted solution density obtained on a series of grids ranging in size from $8 \times 8 \times 16$ (or an angular resolution of 11.25°) to $128 \times 128 \times 256$ (0.70°) for each of the six cubed-sphere blocks, which corresponds to 6,144 and 25,165,824 total cells, respectively, are given in Fig. 7 for this supersonic flow. Note that the error norms are evaluated in this work based on the point-wise differences between the exact solution computed at the centroid of each computational cell and the corresponding average solution (i.e., the value of the reconstruction at the cell centroid). Each L_p -error norm is then computed by applying the appropriate p -norm to the vector of cell average errors, determined as the product of the point-wise errors and the corresponding cell volumes, and by normalizing with the volume of the whole domain for L_1 and L_2 norms, and that of the appropriate computational cell for L_∞ norm.

The results in Fig. 7 show that the second-order theoretical accuracy for this smooth flow is achieved in all error norms as the mesh is refined, indicating that the finite-volume algorithm handles accurately the non-orthogonal grid near the projection of the cube corners and at sector boundaries. Figure 8 provides some more insight into the angular distribution of the error. For a radius of $R = 2.5$ and two different grid resolutions (corresponding to the first two resolutions of Fig. 7), the pointwise difference is plotted in each cell between the exact entropy and the numerically determined entropy at the cell center, divided by the exact entropy (relative error). The flow solution is spherically symmetric but the grid is not, and a non-uniform error distribution is expected. The relative error in entropy is large for the first very coarse grid (up to 85%), but it is confirmed that the error is reduced everywhere (including at sector boundaries and corners) as the grid is refined by a factor of two in each direction, with a ratio that is close to the expected ratio of four, consistent with the findings on near-uniform second-order convergence of Figure 7. On the refined grid of Figure 8(b), the variation in the error is small (ranging between 11.53% and 15.80%). The error at the sector boundaries is somewhat larger than the error in the interior cells of the grid sectors ($\approx 15.5\%$ versus 13.5%),




Figure 8: Distribution of the relative error in the entropy for the fully supersonic Euler flow at a radius of $R = 2.5$ (middle of the domain) and for two different grid resolutions. The plots show in each cell the pointwise difference between the exact entropy at the cell center and the numerically determined entropy at the cell center, divided by the exact entropy. (a) $8 \times 8 \times 16$ root blocks; and (b) $16 \times 16 \times 32$ root blocks.

but the differences are small.

4.1.2. Spherically Symmetric Transonic Wind

To illustrate the accuracy of the algorithm for solar wind-like solutions, the expansion of an inviscid compressible non-conducting gas under the influence of a gravitational field is now considered. Similar validation cases have been studied in [58] and the Newton Critical Point (NCP) method outlined there for one-dimensional problems has been used here to provide a highly accurate adaptively refined reference solution with 918 points.

A spherically symmetric external gravitational field $\vec{g} = -g_*/r^2 \hat{e}_r$ is considered. Here, r is the normalized distance to the solar or planet center, \hat{e}_r is the radial unit vector, and the non-dimensional constant $g_* = GM_*/(l_o a_o^2)$ is computed based on the gravitational constant, G , the solar or planetary mass, M_* , the reference length scale, l_o , and the ion-acoustic wave speed, a_o , of a suitable reference solution state. Thus, the volumetric source term \mathbf{Q}_G , which accounts for the effect of gravitational acceleration, is expressed as a function of the radial position vector, $\vec{r} = r\hat{e}_r$, in the following vector form

$$\mathbf{Q}_G = -\frac{\rho g_*}{r^3} \left[0, \vec{r}, \vec{0}, \vec{r} \cdot \vec{V} \right]^T. \quad (11)$$

As in [58], normalization factors have been used to provide non-dimensional variables. The spherically symmetric subsonic inflow at the inner sphere $R_i = 1$ is defined by the dimensionless density $\rho = 5$ and pressure $p = 23$. We choose $GM_* = 14$ which allows for a stationary transonic outflow solution in which the radial flow velocity is initially subsonic, then passes through a critical point where the local Mach number is one, and subsequently takes on supersonic values beyond the critical radius [58]. Although a crude approximation of a real solar wind, the accurate capturing of the transonic behaviour outlined above is a good test case for our code.

The predicted Mach number distribution for this spherically symmetric transonic flow obtained on a uniform cubed-sphere mesh M3 with 1,228,800 cells (2.25° angular resolution) is shown in Fig. 9(a). In this case, 128 cells are used to resolve the radial direction (i.e., 128 uniform radial points in the 3D simulation have to be compared to 918 highly adaptive points in the 1D reference solution). At the inner boundary, V_r is extrapolated from the computational domain, and $\vec{V}_\parallel = 0$. At the outer boundary, all state variables are extrapolated to the ghost cells. Solving Riemann problems at the boundaries ensures the correct propagation of the flow characteristics, thereby allowing the solution residual to converge to machine accuracy. The result shows that the steady-state transonic


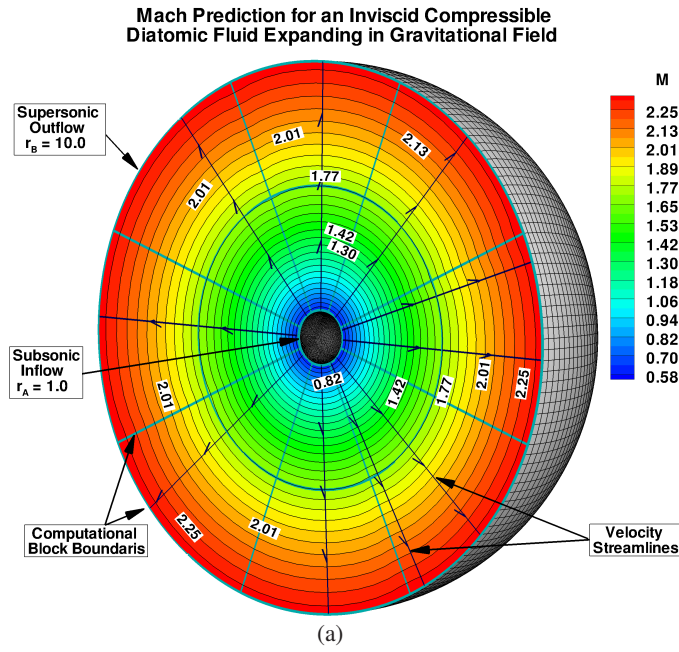


Figure 9: (a) Predicted Mach number distribution for spherically symmetric transonic wind obtained on the uniform mesh M3 with 1,228,800 cells; and (b) variation of flow properties in the x -axis direction predicted on M3 are compared against those obtained with the NCP method in one-dimension. The velocity prediction on meshes M1 and M2 (see text) is also shown.

solution has been correctly captured and the Mach number contour lines are close to circular indicating the good homogeneity of the flow solution at this grid resolution.

In Fig. 9(b), the flow properties of the predicted solution sampled along the x -axis are compared against those obtained with the NCP method on a non-uniform grid so as to capture the sharp density and pressure gradients near the inflow boundary. Additionally, the prediction of the radial velocity in the x -direction obtained with meshes M1 and M2 with 19,200 (9°) and 153,600 (4.5°) cells, respectively, has been also plotted. Note that these two meshes have 32 and 64 cells, respectively, in the radial direction. It is quite noticeable that the solution significantly improves as the mesh is refined, and it approaches in the convergence limit the 1D transonic solution predicted by the NCP method.

To assess the capability of the computational framework to improve the prediction of flow properties in regions of large solution gradients and the accuracy of the critical point location, the same transonic wind problem has

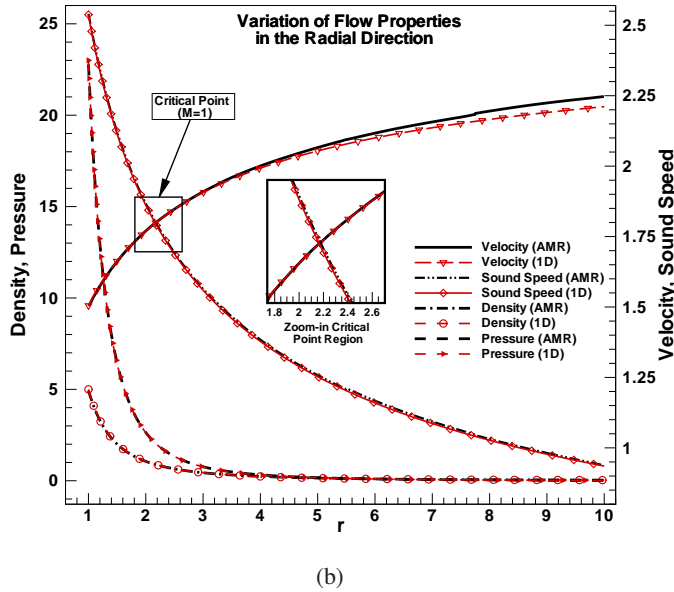
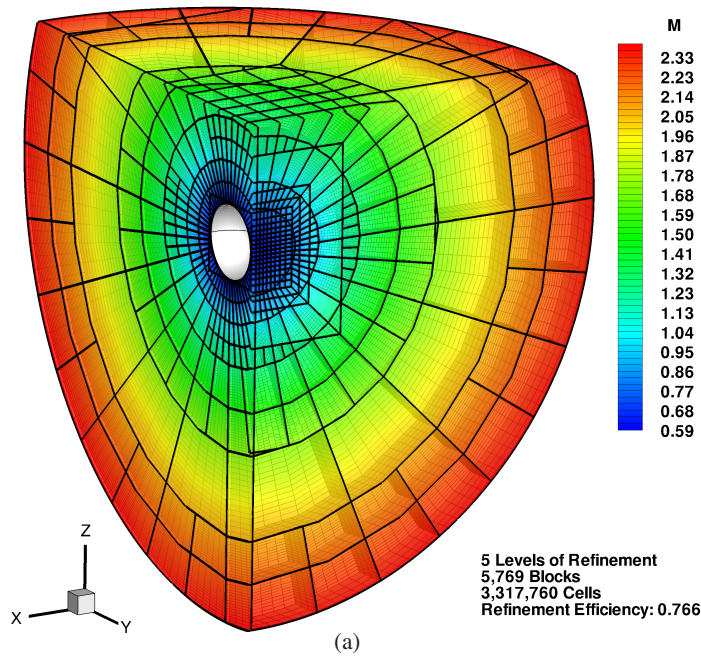


Figure 10: (a) Predicted Mach number distribution for spherically symmetric transonic wind obtained on the adapted cubed-sphere mesh after 5 refinement levels and with 3,317,760 cells; and (b) the flow properties in the x -axis direction are compared against those predicted with the NCP method.

been reconsidered with automatic solution-directed mesh adaptation. The initial mesh in this case consisted of 48 blocks, each with $6 \times 6 \times 16$ (15°) cells. The AMR procedure described in Sect. 3.4 has been applied at sequential intervals using refinement and coarsening threshold values of 0.7 and 0.3, respectively, in conjunction with density gradient and divergence of velocity refinement criteria. On each intermediate AMR mesh the solution residual is reduced several orders of magnitude, whereas the solution on the final mesh has been fully converged so as to obtain, for example, a constant mass flow rate through spherical surfaces in a discrete sense. The contour plot of the Mach number obtained on the final refined mesh after five levels of refinement applied to the 3D cubed-sphere grid is shown in Fig. 10(a). The block boundaries shown in the same figure indicate that a fairly symmetric refinement is automatically generated so as to improve the prediction of solution gradients near the inflow boundary, and demonstrate that the proposed AMR approach can be successfully applied to three-dimensional cubed-sphere

grids. The comparison of flow properties predicted by the 1D NCP method and the 3D AMR algorithm described in this work along the x -axis are shown in Fig. 10(b). It is quite obvious that the two solutions show excellent agreement in the transonic region and only slight deviations are encountered close to the supersonic outflow due to reduced resolution in the adaptive grid. Note that the 3D adaptive solution has approximately 136 points in the radial direction, while the 1D reference solution has 918 adaptive points. Notice, however, that for a similar number of radial points the AMR grid captures the velocity variation significantly better than the finest uniform mesh, M3, for almost the whole domain. The highest resolution obtained in the AMR simulation is at the inner boundary and consists of 0.94° angular resolution and $0.035R_i$ radial resolution. The somewhat under-resolved supersonic outflow in the AMR case is only a consequence of the parameters used for the gradient-based refinement criteria and of available computational resources, and does not represent a limitation of the AMR framework.

4.2. Results for Magnetized Plasma

4.2.1. Systematic Grid Convergence Studies Based on MHD Manufactured Solution

To assess the accuracy of the finite-volume scheme on cubed-sphere grids, we have performed convergence studies for a steady-state axi-symmetric exact solution of magnetized gas flowing outward at supersonic speeds. (Note that the cubed-sphere grid is not axi-symmetric, so this problem is adequate for testing the three-dimensionality of our framework.)

As far as we know, there are no non-trivial MHD test problems with an exact solution in a domain between two concentric spheres (except for problems with rather trivial radial 1D solutions similar to the problem of Section 4.1.1), which can be used for quantitative grid convergence studies. To fill this void, we propose a new fully supersonic axi-symmetric MHD test problem that has an exact solution in a domain between two concentric spheres, and can thus be used to test the order of convergence of MHD solvers in this type of domains.

In this test problem, the exact solution is specified as

$$\mathbf{U}(x, y, z) = \left[r^{-\frac{5}{2}}, \frac{x}{\sqrt{r}}, \frac{y}{\sqrt{r}}, \frac{z}{\sqrt{r}} + \kappa r^{\frac{5}{2}}, \frac{x}{r^3}, \frac{y}{r^3}, \frac{z}{r^3} + \kappa, r^{-\frac{5}{2}} \right]^T, \quad (12)$$

and the volumetric source term \mathbf{Q} in Eq.(1) includes the term \mathbf{Q}_O , which represents the residual obtained for the exact solution (Eq. 12). (Note that an exact solution of a modified equation system obtained in this way is sometimes called a ‘manufactured solution’ [33]). Here $\kappa = 0.017$ is a perturbation parameter chosen such that the solution has significant latitudinal variation yet the flow remains supersonic in the whole domain. As can be observed easily from Eq. (12), the velocity and the magnetic field have been chosen aligned everywhere such that $\vec{V} \times \vec{B} = 0$ and therefore, no source terms arise in the induction equation. Moreover, the magnetic field has been constructed by considering a background inverse squared radial function, $\vec{B}_0 = r^{-2}\hat{e}_r$, plus a perturbation field, $\vec{B}_1 = \nabla f$, derived from the spherical harmonic function of degree one and order zero $f(x, y, z) = \kappa z$, and it thus automatically satisfies the divergence-free condition. This magnetic field is also irrotational.

Thus, the final form of \mathbf{Q}_O for a di-atomic gas is a function only of the space coordinates and can be written as

$$\mathbf{Q}_O = \begin{bmatrix} 0, \\ \frac{1}{2}xr^{-\frac{5}{2}}(r^{-1} - 5r^{-2} - \kappa z), \\ \frac{1}{2}yr^{-\frac{5}{2}}(r^{-1} - 5r^{-2} - \kappa z), \\ \frac{1}{2}zr^{-\frac{5}{2}}(r^{-1} - 5r^{-2} - \kappa z) + \frac{5}{2}r^{-\frac{1}{2}}\kappa(1 + \kappa rz) + \kappa r^{-\frac{1}{2}}, \\ \vec{0}, \\ \frac{1}{2}r^{-2} + \kappa z(3.5r^{-1} + 2\kappa z) + \frac{(\kappa r)^2}{2}(7 + 5\kappa rz) \end{bmatrix}. \quad (13)$$

The computational domain used for this convergence study is defined by inner and outer spheres of radius $R_i = 2$ and $R_o = 3.5$, respectively. The inflow boundary condition is specified based on the exact solution and the outflow uses linear extrapolation.

The L_1 , L_2 , and L_∞ norms of the error in the predicted solution density at cell centroids obtained on a series of grids ranging in size from $8 \times 8 \times 10$ (11.25°) to $256 \times 256 \times 320$ (0.352°) for each of the six cubed-sphere blocks, which




Figure 11: L_1 , L_2 , and L_∞ error norms in the predicted solution density for the manufactured MHD solution described in the text.

corresponds to 3,840 and 125,829,120 total cells, respectively, are given in Fig. 11 for this supersonic flow. As the mesh is refined, the slopes of the L_1 -, L_2 - and L_∞ -norms approach in the asymptotic limit -2.06, -2.10 and -2.05, respectively. The results clearly show that the second-order theoretical accuracy of the scheme can also be achieved in all error norms for smooth but non-radial flows with a magnetic field, thereby providing support for the validity of the applied spatial discretization procedure on cubed-sphere grids.

4.2.2. Magnetically Dominated MHD Bow Shock

We now consider the application of the proposed AMR algorithm to 3D MHD bow-shock flows around a perfectly conducting sphere. In particular, we consider the magnetically dominated upstream conditions from [59] in conjunction with solution-adaptive cubed-sphere grids. The inflow parameters used for this problem have been chosen as in [59] and are $\rho = 1$, $p = 0.2$, $B_x = 1$, $B_y = 0$, $v_x = 1.4943$ and $v_y = 0.1307$, which correspond to an upstream plasma characterized by $\beta = 2p/B^2 = 0.4$, an Alfvénic Mach number $M_{Ax} = 1.49$ along the upstream magnetic field lines and an angle $\theta_{v,B} = 5^\circ$ between the upstream velocity and magnetic vector fields. As shown in [59], this particular upstream configuration gives rise to 3D intermediate shocks and multiple interacting shock fronts, that are sought here to be well resolved by the adaptive mesh refinement procedure.

The computational domain for solving this problem is like the one depicted in Fig. 12 and is defined by inner and outer spheres of radius $R_i = 1$ and $R_o = 8$. Note that this mesh contains only five blocks instead of the typical six associated with a complete cubed-sphere mesh. This extra flexibility in our implementation allows for better resolution of the interesting flow features occurring on the upstream side of the sphere with a smaller domain and fewer computational cells as compared to the case in which the domain at the back of the sphere would have been included. In this grid, the back panel domain boundaries are tilted at 15° relative to the Cartesian (y, z) plane. Reflection boundary conditions (BCs) are imposed along the inner sphere. For the velocity and magnetic fields, the normal components are copied from the physical cells to the ghost cells with inverted sign, whereas the tangential components are copied without sign change, resulting in velocity and magnetic fields that are parallel to the inner sphere at the domain boundary. A free-stream boundary condition is applied to the outer sphere. Linear extrapolation BCs are implemented for the back panels of the outer boundary based on the fact that the flow normal to the panels is superfast everywhere between the inner and outer spheres and all eight characteristic waves leave the domain. Note that this flow feature and the use of Riemann-based flux functions avoid any requirements for a special treatment of the region where the bow shock intersects the domain boundaries. The initial multi-block grid for this problem consists of a total of $20 \times 8 \times 8 \times 10$ blocks, with four blocks in the radial direction for each of the five sectors.

The predicted Mach number and density distribution in the Cartesian (x, y) plane obtained after 7 adaptive re-




Figure 12: Cubed-sphere grid formed by only five root blocks, for simulation on one side of the sphere. A cut in the outer spherical geometry allows a better view of the inner spherical cap.

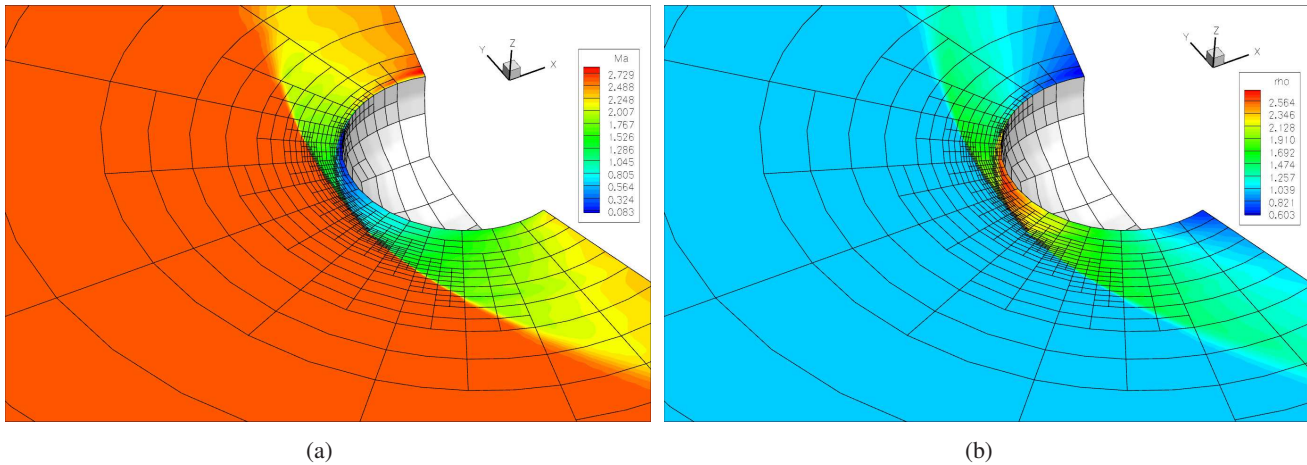


Figure 13: Predicted (a) acoustic Mach number and (b) density distribution in the Cartesian (x, y) plane for a magnetically dominated supersonic flow past a sphere obtained on the adapted cubed-sphere mesh after 7 refinement levels and with 22,693 blocks and 14,523,520 computational cells. Additionally, the block boundaries are depicted with solid black lines.

finement levels on the final refined mesh consisting of 22,693 blocks and 14,523,520 computational cells is shown in Fig. 13. This AMR steady-state solution is obtained in an automatic process starting from simulation on a coarse grid until a steady state is reached approximately, after which the grid is adaptively refined and the computation is resumed until a steady state is approximately reached again. This is repeated for each level of refinement. More detailed investigation of the numerical results (not shown) confirms that they are consistent with the results of [59]; all shocks arising in front of the sphere are captured by the refinement criteria and are well resolved by the 3D adaptation procedure, especially also in the region corresponding to the bifurcation of the shocks. (The bifurcation of the shocks is the most resolved region with a minimum angular resolution of 0.17° and minimum spatial distance of 0.00094 or approximately $1/1000 R_i$.)

4.2.3. Time-Invariant Solar Wind Solution

Finally, the application of the cubed-sphere simulation framework to a more realistic space-physics problem is considered. This test problem represents an initial step towards developing more sophisticated models for the solar wind and carrying out more advanced and complete space physics studies with the algorithm described here.

A numerical solution of the “steady” solar wind representative of solar minimum conditions (i.e., for the quiet Sun) is obtained based on the global 3D MHD model of Groth *et al.* [2]. The steady-state solar wind from $1 R_s$ (i.e., solar radius) to 1 AU is modelled by assuming that, at $1 R_s$, the inner solar corona is a large reservoir of hot plasma with an embedded magnetic multipole field described by a multipole expansion that includes terms up to the octupole moment (see [2, 60] for details). The plasma temperature (the sum of the ion and electron temperatures) of the reservoir is assumed to be $T_s = 2.85 \times 10^6$ K, and the plasma density is taken to be $1.5 \times 10^{14} \text{ m}^{-3}$. The solar magnetic field is azimuthally symmetric about the magnetic axis which is aligned with the Cartesian z -coordinate direction. The multipole expansion has been chosen such as to obtain a maximum field strength of 8.4 G at the magnetic poles and a strength of 2.2 G at the solar magnetic equator. For these solar wind calculations, the MHD equations are solved for the perturbation of the magnetic field from the embedded (“intrinsic”) magnetic multipole field, similarly to the approach used in Groth *et al.* [2] which follows the pioneering work of Tanaka [61]. The single-fluid plasma is assumed to have a particle mass equal to the proton mass, m_p , and a polytropic index $\gamma = 5/3$ is used everywhere in the domain. Note that more realistic solar-wind models have been previously considered which account for the rotation of the Sun and the tilting of the solar magnetic axis with respect to the ecliptic plane (see e.g., [2, 60]), and our approach can be extended to include these effects as well.

To reproduce a realistic solar wind using the ideal MHD description, which neglects the thermal conduction that is important in the vicinity of the Sun [62], and to allow the use of an ideal plasma gas with a polytropic index $\gamma = 5/3$ throughout the computational domain such that adiabatic cooling at larger heliospheric distances is correctly modelled, Groth *et al.* [2] have included an extra source term in the energy equation that mimics the effects of heat conduction and energy dissipation above the transition region such as to reproduce solar wind characteristics determined by *in situ* observations. The global solar wind model with additional heating source term proposed by Groth *et al.* [2] has proved its use in several numerical simulations of coronal mass ejection (CME) evolution [2, 60, 63, 64] and the same heating term is used here as well. Thus, the heating source term \mathbf{Q}_H is given by

$$\mathbf{Q}_H = -\rho \left[0, \vec{0}, \vec{0}, \rho q \left(T_o - \gamma \frac{p}{\rho} \right) \right]^T, \quad (14)$$

where $q(x, y, z)$ is an exponentially-decaying specific heat capacity function of the radial distance from the Sun and $T_o = T_o(x, y, z)$ is a prespecified “target” temperature which has a particular spatial distribution [60].

For this numerical simulation the computational domain is the spherical shell defined by inner and outer spheres of radius $R_i = 1$ and $R_o = 215$, which represent distances normalized by the solar radius, R_s . The cubed-sphere computational grid is stretched with a higher density of grid cells towards the inner sphere and consists of 768 self-similar $16 \times 16 \times 20$ blocks and 3,932,160 cells, providing an angular resolution of 1.4° and a minimum cell size at the solar surface of $1/4 R_s$. Boundary conditions are implemented at the inner sphere, where the flow is subsonic, according to the propagation of the characteristics, to conservation principles for radial mass and magnetic fluxes, and to physical conditions for the directions of velocity and magnetic field vectors in perfectly conducting fluids (see, e.g., [4] for details). Moreover, at the inner boundary linear extrapolation of density gradient is applied and a fixed solar corona temperature is enforced whereas the plasma is permitted to freely leave the reservoir, but no “backflow” is allowed. At the outer boundary, the flow is supersonic and consequently all variables are simply extrapolated. The initial condition for the simulation is provided by Parker’s isothermal solar wind [65] and the initial magnetic field is obtained from the aforementioned multipole expansion [2].

The steady-state solar wind was obtained by using local time-stepping to speed up convergence of the numerical solution, which was reached after $\sim 22,000$ iterations and a solution residual reduction on the level of $\sim 10^{-7}$. The wall-clock time required for the simulation was approximately 6.5 hours using 384 computing cores of the type mentioned at the beginning of Sect. 4. The CPU time is mostly spent in performing operations local to solution blocks while the inter-processor communication, mostly related to the update of block interfaces, has a relatively minor cost in comparison with the arithmetic intensity of the solar-wind model and the solution procedure. This fact




Figure 14: Representation of the steady-state solar wind solution in the (y, z) meridional plane. The colour shading represents the magnitude of the velocity vector, $V = \|\vec{V}\|_2$, in km/s and the solid lines are magnetic field lines coloured based on the magnitude of $\log_{10}(B)$. The arrows of the magnetic lines indicate the magnetic polarity chosen for this problem (lines leave from the North pole of the Sun and enter into the South pole). The boundaries of the mesh blocks at the solar surface are also shown in the figure.

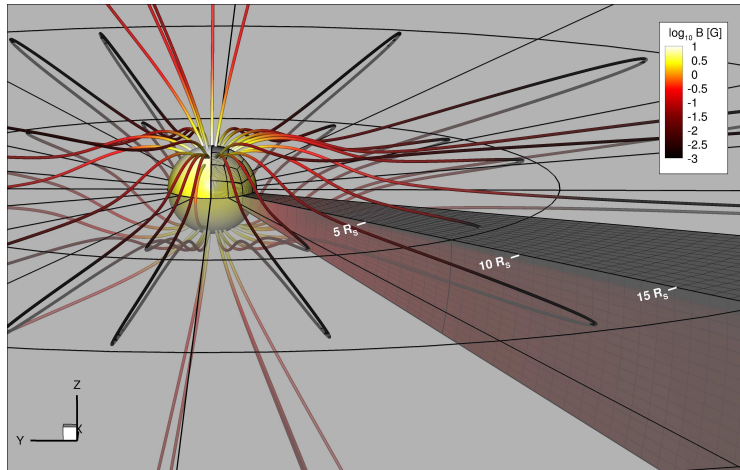


Figure 15: Three-dimensional representation of open and closed magnetic field lines in the corona region. Two blocks of the cubed-sphere grid and the associated meshes are completely shown together with the edges of the blocks in the equatorial plane and the edges and face nodes of several blocks at the inner sphere.

is demonstrated later in this section based on quantitative measurements of parallel performance.

Figure 14 shows a meridional cut through the numerical solution of the steady-state solar wind obtained using the algorithm described above. The shading represents the magnitude of the velocity field, and the solid lines correspond to the predicted field lines which have been coloured based on the logarithm of the magnitude of the magnetic field, $\log_{10}(B)$. A three-dimensional representation of several magnetic field lines in the solar corona is shown in Fig. 15, in which the same colouring scheme as the one just mentioned was used. Additionally, this last figure depicts elements of the computational grid. Inspection of Fig. 14 reveals a bimodal outflow pattern with slow wind leaving the Sun near the equator, and high-speed solar wind of around 700 km/s in the regions of open magnetic field lines emanating from the coronal holes. This is an important feature of the solar wind which has been revealed by *in situ* observations [2, 66]. Other noticeable features present in the solution are the formation of a “helmet streamer” configuration and a heliospheric current sheet, the latter representing the relatively-thin equatorial region across which the magnetic field changes the polarity rapidly. It can also be observed that in the regions of closed field lines near the solar surface the predicted plasma velocity is very small and subsonic, in the range of 10-100 km/s.

The current problem also represents a practical test case for assessing the parallel performance of the proposed

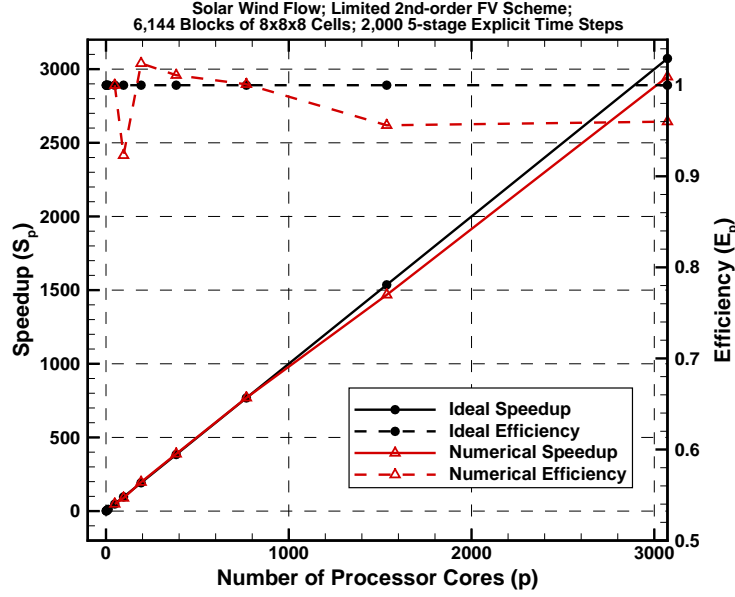


Figure 16: The parallel speedup and efficiency achieved for the time-invariant solar wind problem by the 2nd-order finite-volume algorithm on cubed-sphere grids as a function of the number of computing cores.

formulation on cubed-sphere grids by measuring/evaluating the parallel speedup and parallel efficiency relative to 48 processor cores defined as

$$S_p = \frac{t_{48}}{t_p}; \quad E_p = \frac{S_p}{p/48}, \quad (15)$$

respectively, where t_{48} and t_p are the execution times required to solve the problem by 48 and by p processor cores, respectively. Here, the parallel speedup, also known as strong scaling, is measured by considering a fixed problem size of 6,144 $8 \times 8 \times 8$ blocks and a fixed number of 2,000 multi-stage explicit time steps, and performing the simulation on an increasing number of processor cores, p , up to $p = 3,072$. Note that a perfect speedup corresponds to the ideal situation in which the execution time of the simulation on p computing cores is $p/48$ times smaller than the reference time, t_{48} . In Fig. 16, the parallel speedup and efficiency achieved by the method as a function of the number of computing cores is plotted and compared with the corresponding ideal variations of the parallel-performance parameters. The results show that the multi-dimensional second-order finite-volume scheme applied to cubed-sphere grids and MHD flows is capable to provide very good parallel scalability up to 3,072 computing cores even though the computational blocks contain a relatively small number of cells, namely 512, and the amount of computational work per computing core decreases from 128 to a minimum of 2 blocks/core, or 1.56% of the initial value. Hence, the parallel efficiencies achieved by the numerical scheme are 0.923, 1.024, 1.011, 1.001, 0.956 and 0.960 on 96, 192, 384, 768, 1,536 and 3,072 computing cores, respectively. For the sake of reference, $t_{3,072} = 2$ min on the aforementioned high-performance parallel cluster (see Sect. 4). Note that the slightly higher values than one of several parallel efficiencies can be interpreted as being due to a cache effect, in which the increases in available cache for the program can overcome the computational cost of the additional inter-block parallel communication incurred. It is also important to note that the results depicted in Fig. 16 are for strong scaling tests, which are more challenging than weak scaling tests, in which the number of grid cells per processor is kept constant, and for which even better parallel efficiency can thus be expected.

5. Discussion and Concluding Remarks

A parallel block-based solution-adaptive simulation framework has been described for three-dimensional cubed-sphere grids and flows governed by hyperbolic conservation laws. In particular, the method has been developed for the Euler and MHD equations, targeting application to space-physics problems.

The main novelty compared to existing cubed-sphere grid algorithms is that we use a fully multi-dimensional finite-volume approach. A major advantage of our multi-dimensional approach is that no special interpolation or reconstruction is needed at ghost cells adjacent to the boundaries and corners of the six sectors of the cubed-sphere grid, whereas existing finite-volume schemes, most of which use dimension-by-dimension differencing, need special interpolation or reconstruction at ghost cells adjacent to sector boundaries. In contrast, in our fully multi-dimensional approach solution blocks adjacent to sector boundaries can directly use physical cells from the adjacent sector as ghost cells while maintaining full accuracy, and no special interpolation or reconstruction is required. The fully multi-dimensional approach naturally maintains uniform second-order accuracy at the sector boundaries and corners of the cubed-sphere grid, and has important additional advantages in terms of simplicity of implementation, parallelization and adaptivity. The proposed scheme constitutes a significant advance in 3D cubed-sphere grid simulation capabilities: it is the first 3D cubed-sphere grid framework with fully dynamical solution-adaptive grid refinement and coarsening, it is scalable up to several thousands of processor cores, and uniform second-order accuracy has been demonstrated for relevant PDEs discretized uniformly in all three dimensions (i.e., without using a layered approach as is common in atmospheric applications), in a truly multi-dimensional way. Compared to existing cubed-sphere grid algorithms, a novelty of the proposed approach involves the use of a fully multi-dimensional finite-volume approach, which leads to important advantages when the treatment of boundaries and corners of the six identical sectors of the cubed-sphere grid is considered. We have achieved this by employing a consistent multi-block approach with unstructured root connectivity (following [30]), and by utilizing a solution reconstruction method that is multi-dimensional with flexible stencil size [50], which allowed us to obtain a framework with uniform second-order accuracy, solution-adaptive refinement, and efficient parallelization that are fully consistent at the sector boundaries and corners of the cubed-sphere grid.

The accuracy of the algorithm has been assessed based on exact solutions, and the theoretically expected second-order accuracy has been achieved. A new fully supersonic axi-symmetric MHD test problem has been presented that has an exact solution in a domain between two concentric spheres, and can thus be used to test the order of convergence of MHD solvers in this type of domains. Furthermore, the predictive capabilities of the framework have been illustrated for other test problems from the literature, including a radial transonic wind flow, a standard solar-wind model and bow shock flows. We have also demonstrated the ability of the scheme to adequately handle and resolve strong discontinuities/shocks with a number of computational elements that is reduced by adaptivity compared to uniformly refined grids. While this features have been demonstrated for the Euler and ideal MHD systems, it is believed that they will also carry over to other systems of conservation laws.

The parallel performance of the proposed 3D adaptive numerical procedure has been shown to be very high for problems with large numbers of solution blocks, computed on several thousands of processor cores. Future research will involve application of the algorithm to complex space-physics problems, the coupling with an effective parallel implicit algorithm (see, for example, the methods of Northrup and Groth [53, 54]), applying the AMR framework to unsteady flows (see, for example, our paper on high-order 2D MHD [67]), and the extension of the numerical scheme to high-order accuracy (i.e., accuracy higher than second-order) as considered in [42, 43, 67, 68]. Using the high-order finite-volume approach of [42, 43, 67, 68], the main technical challenge in extending our approach on the cubed-sphere to fourth-order accuracy is to properly deal with hexahedral cells that have non-planar faces. This can be handled by employing trilinear representations of these non-planar faces. Preliminary results on ongoing work extending the proposed cubed-sphere grid framework to fourth-order accuracy are presented in the recent study by Ivan *et al.* [20].

Acknowledgements

This work was supported by CSA (Canadian Space Agency) CGSM Contract No. 9F007-080157/001/ST. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium. SciNet is funded by: the Canada Foundation for Innovation under the auspices of Compute Canada; the Government of Ontario; Ontario Research Fund - Research Excellence; and the University of Toronto.

References

- [1] W. B. Manchester IV, A. Vourlidas, G. Tóth, N. Lugaz, I. I. Roussev, I. V. Sokolov, T. I. Gombosi, D. L. D. Zeeuw, M. Opher, Three-dimensional MHD Simulation of the 2003 October 28 coronal mass ejection: comparison with LASCO coronagraph observations, *Astrophys. J.* 684 (2008) 1448–1460.
- [2] C. P. T. Groth, D. L. De Zeeuw, T. I. Gombosi, K. G. Powell, Global Three-Dimensional MHD Simulation of a Space Weather Event: CME Formation, Interplanetary Propagation, and Interaction with the Magnetosphere, *J. Geophys. Res.* 105 (A11) (2000) 25,053–25,078.
- [3] T. I. Gombosi, K. G. Powell, D. L. D. Zeeuw, C. R. Clauer, K. C. Hansen, W. B. Manchester, A. J. Ridley, I. I. Roussev, I. V. Sokolov, Q. F. Stout, G. Tóth, Solution-Adaptive Magnetohydrodynamics for Space Plasmas: Sun-to-Earth Simulations, *Comp. Sci. & Eng.* 6 (2) (2004) 14–35.
- [4] C. Jacobs, Magnetohydrodynamic modelling of the solar wind and coronal mass ejections, Ph.D. thesis, University of Leuven, 2007.
- [5] R. Sadourny, Conservative Finite-Difference Approximations of the Primitive Equations on Quasi-Uniform Spherical Grids, *Mon. Weather Rev.* 100 (2) (1972) 136–144.
- [6] C. Ronchi, R. Iacono, P. S. Paolucci, The "Cubed Sphere": A New Method for the Solution of Partial Differential Equations in Spherical Geometry, *J. Comput. Phys.* 124 (1996) 93–114.
- [7] A. Adcroft, J.-M. Campin, C. Hill, J. Marshall, Implementation of an Atmosphere-Ocean General Circulation Model on the Expanded Spherical Cube, *Mon. Weather Rev.* 132 (2004) 2845–2863.
- [8] R. D. Nair, S. J. Thomas, R. D. Loft, A Discontinuous Galerkin Global Shallow Water Model, *Monthly Weather Review* 133 (4) (2005) 876–888.
- [9] J. A. Rossmanith, A wave propagation method for hyperbolic systems on the sphere, *Journal of Computational Physics* 213 (2) (2006) 629 – 658.
- [10] W. M. Putman, S.-J. Lin, Finite-volume transport on various cubed-sphere grids, *J. Comput. Phys.* 227 (2007) 55–78.
- [11] A. St-Cyr, C. Jablonowski, J. M. Dennis, H. M. Tufo, S. J. Thomas, A Comparison of Two Shallow-Water Models with Nonconforming Adaptive Grids, *Mon. Weather Rev.* 136 (2008) 1898–1922.
- [12] C. Chen, F. Xiao, Shallow water model on cubed-sphere by multi-moment finite volume method, *Journal of Computational Physics* 227 (10) (2008) 5019 – 5044.
- [13] C. Yang, J. Cao, X.-C. Cai, A fully implicit domain decomposition algorithm for shallow water equations on the cubed-sphere, *SIAM J. Sci. Comput.* 32 (2010) 418–438.
- [14] P. A. Ullrich, C. Jablonowski, B. van Leer, High-order finite-volume methods for the shallow-water equations on the sphere, *J. Comput. Phys.* 229 (17) (2010) 6104 – 6134, doi: 10.1016/j.jcp.2010.04.044.
- [15] C. Chen, F. Xiao, X. Li, An Adaptive Multimoment Global Model on a Cubed Sphere, *Monthly Weather Review* 139 (2) (2011) 523–548.
- [16] C. G. Chen, F. Xiao, X. L. Li, Y. Yang, A multi-moment transport model on cubed-sphere grid, *International Journal for Numerical Methods in Fluids* 67 (12) (2011) 1993–2014.
- [17] P. A. Ullrich, C. Jablonowski, MCore: A Non-hydrostatic Atmospheric Dynamical Core Utilizing High-Order Finite-Volume Methods, *J. Comput. Phys.* 231 (2012) 5078–5108.

- [18] A. V. Koldoba, M. M. Romanova, G. V. Ustyugova, R. V. E. Lovelace, Three-Dimensional Magnetohydrodynamic Simulations of Accretion to an Inclined Rotator: the "Cubed Sphere" Method, *Astrophys. J.* 576 (2002) L53–L56.
- [19] P. C. Fragile, C. C. Lindner, P. Anninos, J. D. Salmonson, Application of the cubed-sphere grid to tilted black hole accretion disks, *Astrophys. J.* 691 (2009) 482–494.
- [20] L. Ivan, A. Susanto, H. De Sterck, C. P. T. Groth, High-Order Central ENO Finite-Volume Scheme for MHD on Three-Dimensional Cubed-Sphere Grids, in: *Proceedings of the Seventh International Conference on Computational Fluid Dynamics (ICCFD7)*, Hawaii, 2012.
- [21] L. Ivan, H. D. Sterck, S. Northrup, C. P. T. Groth, Three-Dimensional MHD on Cubed-Sphere Grids: Parallel Solution-Adaptive Simulation Framework, Paper 2011-3382, AIAA, 2011.
- [22] P. A. Ullrich, Atmospheric Modeling with High-Order Finite-Volume Methods, Ph.D. thesis, University of Michigan, 2011.
- [23] C. Jablonowski, M. Herzog, J. E. Penner, R. C. Oehmke, Q. F. Stout, B. van Leer, K. G. Powell, Block-Structured Adaptive Grids on the Sphere: Advection Experiments, *Mon. Wea. Rev.* 134 (2006) 3691–3713.
- [24] K. G. Powell, P. L. Roe, T. J. Linde, T. I. Gombosi, D. L. De Zeeuw, A Solution-Adaptive Upwind Scheme for Ideal Magnetohydrodynamics, *J. Comput. Phys.* 154 (1999) 284–309.
- [25] C. P. T. Groth, D. L. D. Zeeuw, K. G. Powell, T. I. Gombosi, Q. F. Stout, A Parallel Solution-Adaptive Scheme for Ideal Magnetohydrodynamics, Paper 99-3273, AIAA, 1999.
- [26] B. van der Holst, R. Keppens, Hybrid block-AMR in cartesian and curvilinear coordinates: MHD applications, *J. Comput. Phys.* 226 (1) (2007) 925 – 946, doi:10.1016/j.jcp.2007.05.007.
- [27] B. van der Holst, W. M. IV, I. Sokolov, G. Tth, T. Gombosi, D. DeZeeuw, O. Cohen, Breakout Coronal Mass Ejection or Streamer Blowout: The Bugle Effect, *APJ* 693 (2) (2009) 1178.
- [28] G. Tóth, B. van der Holst, I. V. Sokolov, D. L. D. Zeeuw, T. I. Gombosi, F. Fang, W. B. Manchester, X. Meng, D. Najib, K. G. Powell, Q. F. Stout, A. Glocer, Y.-J. Ma, M. Opher, Adaptive numerical algorithms in space weather modeling, *J. Comput. Phys.* 231 (3) (2012) 870 – 903, URL <http://www.sciencedirect.com/science/article/pii/S002199911100088X>, doi:10.1016/j.jcp.2011.02.006.
- [29] X. Gao, A Parallel Solution-Adaptive Method for Turbulent Non-Premixed Combusting Flows, Ph.D. thesis, University of Toronto, 2008.
- [30] X. Gao, C. P. T. Groth, A parallel solution-adaptive method for three-dimensional turbulent non-premixed combusting flows, *J. Comput. Phys.* 229 (2010) 3250–3275.
- [31] S. K. Godunov, Symmetric Form of the Equations of Magnetohydrodynamics, *Numerical Methods for Mechanics of Continuum Medium, Siberian Branch of USSR Academy of Sciences* 1 (1972) 26–34.
- [32] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, M. Wesenberg, Hyperbolic Divergence Cleaning for the MHD Equations, *J. Comput. Phys.* 175 (2) (2002) 645 – 673, doi:10.1006/jcph.2001.6961.
- [33] P. J. Roache, *Verification and Validation in Computational Science and Engineering*, Hermosa Publisher, New Mexico, 1998.
- [34] J. J. Quirk, U. R. Hanebutte, A Parallel Adaptive Mesh Refinement Algorithm, Report 93-63, ICASE, 1993.
- [35] M. J. Berger, J. S. Saltzman, AMR on the CM-2, *Appl. Numer. Math.* 14 (1994) 239–253.

- [36] M. J. Berger, J. Olinger, Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations, *J. Comput. Phys.* 53 (1984) 484–512.
- [37] M. J. Berger, P. Colella, Local Adaptive Mesh Refinement for Shock Hydrodynamics, *J. Comput. Phys.* 82 (1989) 67–84.
- [38] Q. F. Stout, D. L. De Zeeuw, T. I. Gombosi, C. P. T. Groth, H. G. Marshall, K. G. Powell, Adaptive Blocks: A High-Performance Data Structure, in: *Proceedings of SC97*, San Jose, California, U.S.A., November 12–15, 1997.
- [39] C. P. T. Groth, D. L. De Zeeuw, T. I. Gombosi, K. G. Powell, A Parallel Adaptive 3D MHD Scheme for Modeling Coronal and Solar Wind Plasma Flows, *Space Sci. Rev.* 87 (1999) 193–198.
- [40] J. S. Sachdev, C. P. T. Groth, J. J. Gottlieb, A Parallel Solution-Adaptive Scheme for Predicting Multi-Phase Core Flows in Solid Propellant Rocket Motors, *Int. J. Comput. Fluid Dyn.* 19 (2).
- [41] X. Gao, C. P. T. Groth, A parallel adaptive mesh refinement algorithm for predicting turbulent non-premixed combusting flows, *Int. J. Comput. Fluid Dyn.* 20 (5) (2006) 349–357.
- [42] L. Ivan, C. P. T. Groth, High-Order Central ENO Finite-Volume Scheme with Adaptive Mesh Refinement, Paper 2007-4323, AIAA, 2007.
- [43] L. Ivan, C. P. T. Groth, High-Order Solution-Adaptive Central Essentially Non-Oscillatory (CENO) Method for Viscous Flows, Paper 2011-0367, AIAA, 2011.
- [44] R. Keppens, Z. Meliani, A. J. van Marle, P. Delmont, A. Vlasis, B. van der Holst, Parallel, grid-adaptive approaches for relativistic hydro and magnetohydrodynamics, *J. Comput. Phys.* 231 (3) (2012) 718–744.
- [45] NASA, CGNS, CFD General Notation System, retrieved from <<http://www.grc.nasa.gov/www/cgns/>> (Accessed Mar. 26, 2012).
- [46] S. K. Godunov, Finite-Difference Method for Numerical Computations of Discontinuous Solutions of the Equations of Fluid Dynamics, *Mat. Sb.* 47 (1959) 271–306.
- [47] G. Tóth, D. Odstrčil, Comparison of Some Flux Corrected Transport and Total Variation Diminishing Numerical Schemes for Hydrodynamic and Magnetohydrodynamic Problems, *J. Comput. Phys.* 128 (1996) 82–100.
- [48] A. Harten, P. D. Lax, B. van Leer, On Upstream Differencing and Godunov-Type Schemes for Hyperbolic Conservation Laws, *SIAM Rev.* 25 (1) (1983) 35–61.
- [49] T. Linde, A practical, general-purpose, two-state HLL Riemann solver for hyperbolic conservation laws, *Int. J. Numer. Meth. Fluids* 40 (2002) 391–402.
- [50] T. J. Barth, D. C. Jespersen, The Design and Application of Upwind Schemes on Unstructured Meshes, Paper 89-0366, AIAA, 1989.
- [51] V. Venkatakrisnan, On the Accuracy of Limiters and Convergence to Steady State Solutions, Paper 93-0880, AIAA, 1993.
- [52] B. van Leer, C. H. Tai, K. G. Powell, Design of Optimally-Smoothing Multi-Stage Schemes for the Euler Equations, Paper 89-1933-CP, AIAA, 1989.
- [53] S. A. Northrup, C. P. T. Groth, Prediction of Unsteady Laminar Flames Using a Parallel Implicit Adaptive Mesh Refinement Algorithm, in: *Proceedings of the 6th U.S. National Combustion Meeting*, 2009.
- [54] S. A. Northrup, C. P. T. Groth, Parallel Implicit AMR Scheme for Unsteady Reactive Flows, in: *Proceedings of the 18th Annual Conference of the CFD Society of Canada*, London, Canada, 2010.

- [55] D. J. Mavriplis, Multigrid Techniques for Unstructured Meshes, Report 95-27, ICASE, 1995.
- [56] M. J. Aftomis, M. J. Berger, S. M. Murman, Applications of space-filling curves to Cartesian methods for CFD, Paper AIAA-2004-1232, AIAA, 2004.
- [57] H. Samet, The design and analysis of spatial data structures, Addison-Wesley Series on Computer Science and Information Processing, Addison-Wesley, 1990.
- [58] H. De Sterck, S. Rostrup, F. Tian, A fast and accurate algorithm for computing radial transonic flows, *J. Comput. Appl. Math.* 223 (2) (2009) 916–928.
- [59] H. De Sterck, S. Poedts, Intermediate Shocks in Three-Dimensional Magnetohydrodynamic Bow-Shock Flows with Multiple Interacting Shock Fronts, *Phys. Rev. Lett.* 84 (24) (2000) 5524–5527.
- [60] W. B. Manchester IV, T. I. Gombosi, I. I. Roussev, A. Ridley, D. L. D. Zeeuw, I. V. Sokolov, K. G. Powell, Modeling a space weather event from the Sun to the Earth: CME generation and interplanetary propagation, *J. Geophys. Res.* 109 (A0) (2004) 2107–2121.
- [61] T. Tanaka, Finite Volume TVD Scheme on an Unstructured Grid System for Three-Dimensional MHD Simulation of Inhomogeneous Systems Including Strong Background Potential Fields, *J. Comput. Phys.* 111 (2) (1994) 381 – 389.
- [62] N. Meyer-Vernet, Basics of the Solar Wind, Cambridge University Press, URL <http://dx.doi.org/10.1017/CB09780511535765>, 2007.
- [63] N. Lugaz, W. B. Manchester IV, T. I. Gombosi, The evolution of coronal mass ejection density structures, *Astrophys. J.* 627 (2005) 1019–1030.
- [64] C. Jacobs, B. van der Holst, S. Poedts, Comparison between 2.5D and 3D simulations of coronal mass ejections, *Astronomy and Astrophysics* 470 (2007) 359–365.
- [65] E. N. Parker, Dynamics of the Interplanetary Gas and Magnetic Fields, *Astrophys. J.* 128 (3) (1958) 664–676.
- [66] T. E. Cravens, Physics of Solar System Plasmas, Cambridge University Press, New York, 1997.
- [67] A. Susanto, L. Ivan, H. De Sterck, C. P. T. Groth, High-Order Central ENO Finite-Volume Scheme for Ideal MHD, *J. Comput. Phys.* 250 (2013) 141–164.
- [68] L. Ivan, Development of High-order CENO Finite-volume Schemes with Block-based Adaptive Mesh Refinement, Ph.D. thesis, University of Toronto, URL <http://hdl.handle.net/1807/29759>, 2011.