



A High-Order Finite-Volume Method with Anisotropic AMR for Ideal MHD Flows

L. Freret* and L. Ivan†

Institute for Aerospace Studies, University of Toronto, Toronto, ON, M3H5T6, Canada

H. De Sterck‡

School of Mathematical Sciences, Monash University, Melbourne, Australia

C.P.T. Groth§

Institute for Aerospace Studies, University of Toronto, Toronto, ON, M3H5T6, Canada

A high-order central essentially non-oscillatory (CENO) finite-volume scheme combined with a block-based anisotropic adaptive mesh refinement (AMR) algorithm is proposed for the solution of the ideal magnetohydrodynamics (MHD) equations. A generalized Lagrange multiplier (GLM) divergence correction technique is applied to achieve numerically-divergent free magnetic fields while preserving high-order accuracy. The cell-centered CENO method uses a hybrid reconstruction approach based on a fixed central stencil. Smooth and fully resolved solution data is interpolated using an unlimited high-order k -exact reconstruction. In cells deemed to have under-resolved/discontinuous solution content based on a smoothness indicator, the high-order reconstruction reverts to a lower-order limited linear approach. The high-order CENO finite-volume scheme is implemented with fourth-order spatial accuracy within a flexible multi-block hexahedral grid framework. Parallel implementation and local anisotropic grid adaptivity are achieved by using a hierarchical block-based domain decomposition strategy in which the connectivity and refinement history of grid blocks are tracked using a binary tree data structure. Physics-based refinement criteria are used for directing the mesh refinement. Numerical results, including solution-driven anisotropic refinement of cubed-sphere grids are presented to demonstrate the accuracy and efficiency of the proposed high-order solution method.

I. Introduction and Motivation

The aim of this paper is to combine a high-order spatial discretization scheme with an adaptive mesh refinement technique to accurately capture disparate spatial and temporal scales as in space-physics problems.¹⁻³ The high-order central essentially non-oscillatory (CENO) finite-volume scheme described in Ivan *et al.*^{4,5} uses a hybrid reconstruction approach based on a fixed central stencil. An unlimited high-order k -exact reconstruction is performed for all computational cells and the scheme reverts to a low-order limited linear approach for cells with under-resolved/discontinuous solution content. Switching in the hybrid procedure is determined by a smoothness indicator. The high-order CENO scheme has been successfully applied to a broad range of flows on multi-block structured meshes including inviscid flows,⁴ viscous flows,⁵ large-eddy simulation (LES) for turbulent premixed flames⁶ and magnetohydrodynamics (MHD) problems.⁴ The efficiency of the CENO scheme has also been assessed on cubed-sphere meshes.⁴ Moreover the method has been extended to unstructured meshes for laminar viscous flows⁷ and turbulent reactive flows.⁸

Block-based adaptive mesh refinement (AMR) approaches⁹⁻¹¹ are very attractive since they allow to refine automatically the mesh for treating the disparate spatial scales while requiring an overall light data structure due to grid connectivity being computed on a block level. Even if they can sometimes locally over resolve

*Post-doctoral fellow, lfreret@utias.utoronto.ca

†Research Scientist, lucian.ivan@utoronto.ca

‡Professor, hans.desterck@monash.edu

§Professor, senior member AIAA, groth@utias.utoronto.ca

the domain in areas flagged for refinement, the domain decomposition of a block-based adaptively-refined mesh is naturally suitable for parallel implementation and leads to highly scalable solution methods.

The multi-block AMR scheme considered here is based on the previous work by Gao and Groth¹² for reacting flows with isotropic refinement. It has also been applied to the solution of complex flow problems such as non-premixed laminar¹³ and turbulent¹⁴ flames as well as turbulent multi-phase rocket core flows,¹⁵ MHD simulations,^{1,16} and micron-scale flows.¹⁷ The application of the isotropic AMR method to cubed-sphere grids was also examined by Ivan *et al.*¹⁸ Additionally, it was originally extended to anisotropic refinement by Williamschen and Groth¹⁹ for non-viscous flows with second-order accuracy. More recently, Freret and Groth²⁰ reformulated the second-order anisotropic AMR scheme using a non-uniform treatment of the cells, both interior and ghost or halo cells, within a given block. The anisotropic approach makes direct use of the neighboring cells as the ghost cells, even for those at different levels of refinement as found at grid resolution changes. The resulting AMR multi-block scheme is more tractable and well suited for high-order finite-volume schemes. In the present study, the anisotropic refinement criteria is restricted to physics-based approaches as it is a good compromise between implementation simplicity and accuracy. For the flow problems presented in Section V, gradient of the density or curl of the velocity are used to direct the mesh refinement.

The goal of the present study is therefore twofold, namely: (i) to extend the enhanced anisotropic AMR algorithm for use in conjunction with the CENO high-order finite-volume scheme by developing efficient high-order accurate solution transfer operators; and (ii) to apply the proposed procedure to the solution of the ideal MHD equations. It is also shown that the newly-formulated high-order, anisotropic AMR solution procedure is amenable for application to the prediction of space-plasma flows. The solenoidal constraint of the magnetic field is controlled using the generalized Lagrange multiplier (GLM) technique proposed by Dedner *et al.*²¹ This avoids the computational cost associated with an elliptic cleaning step^{22,23} and the possibly more complicated treatment of staggered fields required by constrained transport algorithms.^{24–26}

The paper is organized as follows. The ideal MHD equations and the GLM formulation are described in Section II. In Section III we provide a brief outline of the high-order CENO scheme. The proposed anisotropic AMR block-based method is reviewed in Section IV with the necessary extension for use with the high-order spatial discretization scheme. Finally, fully three-dimensional numerical results are presented in Section V to demonstrate the accuracy and robustness of the fourth-order CENO method for Euler and MHD flows on Cartesian and cubed-sphere grids. Results for steady-state non-magnetized and MHD plasma smooth flows are presented to assess the accuracy of the scheme on a cubed-sphere grid. Robustness of the fourth-order CENO method in the presence of 3D discontinuous flow features and the capability of the implementation to dynamically refine the mesh in a time-dependent simulation are demonstrated.

II. Ideal Magnetohydrodynamics Equations

The hyperbolic system of ideal MHD equations is solved herein using a high-order Godunov-type finite-volume scheme with the GLM formulation proposed by Dedner *et al.*²¹ This approach²¹ couples the divergence constraint, $\nabla \cdot \vec{B} = 0$, with the induction equation through the introduction of a new potential variable, ψ . Thus, the system of conservation laws for which the solution is sought here may be expressed in conservation form as

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

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 conserved variable solution vector, \mathbf{U} , has the form

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

where ρ is the plasma density, $\vec{V} = (V_x, V_y, V_z)$ is the velocity field, $\vec{B} = (B_x, B_y, B_z)$ is the magnetic field, ρe is the total energy and ψ is the so-called generalized Lagrange multiplier variable associated with the treatment of the solenoidal condition, $\vec{\nabla} \cdot \vec{B} = 0$ for the magnetic field. Note that we use a standard orthogonal x, y, z Cartesian coordinate system in the conservation law discretization. The flux dyad, $\vec{\mathbf{F}}$, is

given by

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

The specific total plasma energy is $e = p/(\rho(\gamma - 1)) + V^2/2 + B^2/(2\rho)$, where p is the molecular pressure, $V = \sqrt{V_x^2 + V_y^2 + V_z^2}$ is the magnitude of the fluid velocity, and $B = \sqrt{B_x^2 + B_y^2 + B_z^2}$ is the magnitude of the magnetic field. The term $p_B = \frac{\vec{B} \cdot \vec{B}}{2}$ is known as the magnetic pressure. The numerical source term, \mathbf{S} , is due to the GLM-MHD formulation and given by

$$\mathbf{S} = [0, \vec{0}, \vec{0}, 0, -\frac{c_h^2}{c_p^2} \psi]^T, \quad (4)$$

in which the coefficients c_p and c_h control the relative rates of dissipation and transport of ψ , as well as the corresponding advection speed of the $\nabla \cdot \vec{B}$ cleaning mechanism, respectively. 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 $C_p = \gamma/(\gamma - 1)$. The column vector, \mathbf{Q} , appearing in Eq. (1) generally represents different volumetric sources arising from the physical modelling of space-physics problems such as sources associated with gravitational fields.

II.A. Semi-Discrete Finite-Volume Formulation


The semi-discrete form of the finite-volume formulation applied to Eq. (1) for a hexahedral computational cell (i, j, k) of a three-dimensional grid is given by

$$\frac{d\bar{\mathbf{U}}_{ijk}}{dt} = -\frac{1}{V_{ijk}} \sum_{f=1}^6 \sum_{m=1}^{N_g} (\tilde{\omega} \vec{\mathbf{F}} \cdot \vec{n})_{i,j,k,f,m} + (\bar{\mathbf{S}})_{ijk} + (\bar{\mathbf{Q}})_{ijk} = (\bar{\mathbf{R}})_{ijk}(\bar{\mathbf{U}}), \quad (5)$$

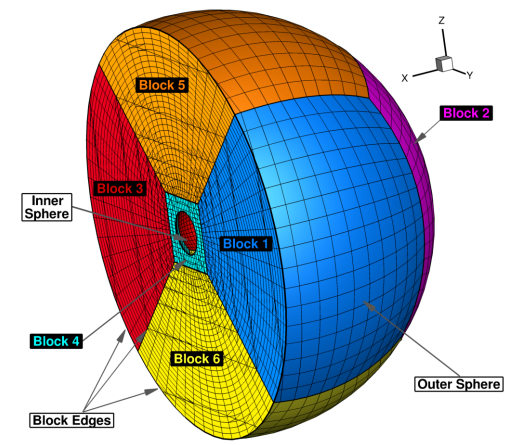
where N_g is the number of Gauss quadrature points and \vec{n} is the local normal of the face f at each of the N_g Gauss quadrature points. The hexahedral cells are contained inside a single block with that block being a part of a multi-block body-fitted mesh as shown in Figure 1. To accommodate complex meshes, as the cubed-sphere grid in Figure 1-right, the connectivity between blocks in the multi-block mesh is defined using an unstructured data structure.¹² The total number of Gauss integration points, N_g , at which the numerical flux is evaluated is chosen as the minimum required to preserve the desired order of solution accuracy. In this work, standard tensor-product quadratures are used for each cell face, as described in Ivan *et al.*⁴ For instance, four Gauss quadrature points per face are required for a fourth order accurate spatial discretization, which is the targeted accuracy order for the scheme considered here. The numerical fluxes, $\vec{\mathbf{F}} \cdot \vec{n}$, at each Gauss quadrature point on each face of a cell (i, j, k) are determined from the solution of a Riemann problem. Given the left and right interface solution values, \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. The values of \mathbf{U}_l and \mathbf{U}_r are determined by performing a CENO polynomial reconstruction, as detailed in the next section. The contributions of the volumetric sources $\bar{\mathbf{S}}_{ijk}$ and $\bar{\mathbf{Q}}_{ijk}$ are evaluated to fourth-order accuracy by again using a standard tensor-product Gauss quadrature with twenty-seven points for the volumetric integration.

III. High-Order CENO Finite-Volume Scheme

The hybrid CENO finite-volume method for conservation laws originally proposed by Ivan and Groth^{4,5} is used to discretize the governing equations on a hexahedral computational grid. The hybrid CENO procedure uses the multidimensional unlimited k -exact reconstruction of Barth²⁷ in smooth regions and reverts to a



Hexahedral cell (i, j, k) showing face normal.



Cross-section of the cubed-sphere grid. Grid block boundaries are shown with bold lines

Figure 1. Three-dimensional hexahedral cell (left) and example of a cubed-sphere mesh composed of 6 blocks (right).

limited piecewise-linear reconstruction algorithm in regions deemed as non-smooth or under-resolved by a solution smoothness indicator, thus providing monotone solutions near discontinuities. The k -order Taylor series for the spatial distribution of a scalar solution quantity, U_{ijk} , within a cell with index ijk about the cell-centroid $(x_{ijk}, y_{ijk}, z_{ijk})$ can be expressed as:

$$U_{ijk}^k(x, y, z) = \sum_{p_1=0}^k \sum_{p_2=0}^k \sum_{p_3=0}^k (x - x_{ijk})^{p_1} (y - y_{ijk})^{p_2} (z - z_{ijk})^{p_3} D_{p_1 p_2 p_3}. \quad (6)$$

The coefficients, $D_{p_1 p_2 p_3}$, are the unknown derivatives for the Taylor polynomial. A total of 20 unknown derivatives needs to be computed for a piecewise cubic ($k = 3$) reconstruction, which is required to achieve a fourth-order accurate solution procedure. They are obtained by solving a constrained least-squares problem as detailed in Ivan *et al.*⁴ To obtain an exact determined or overdetermined set of equations to solve for the unknown derivatives, the number of neighbors considered in the stencil must be at least equal to the number of unknown derivatives, but it is not desirable to make the overdetermined system larger than necessary. In this study a stencil including the first ring plus 6 extra cells is used giving a total of 33 cells when a uniform mesh is used. Depiction of the stencil in the rather general case of a Cartesian uniform grid is provided in Figure 2. In the stencil representation, cells colored yellow represent the interior domain of a block and the blue colored cells represent ghost cells for that block.

The smoothness indicator, \mathcal{S} , used to determine whether a flow variable in cell (i, j, k) is deemed non-




Figure 2. Example of cell stencil in a Cartesian uniform mesh for a cell (i, j, k) located at the block boundary such that its stencil includes ghost cells from neighboring blocks. The interior domain is represented in yellow and blue color is for the ghost cells. The stencil is composed of 27 first-ring cells plus 6 extra second-ring cells.

smooth, is computed as following:¹⁸

$$\mathcal{S} = \frac{\alpha}{\max(1 - \alpha, \epsilon)} \frac{\mathcal{N}_{\text{SOS}} - \mathcal{N}_D}{\mathcal{N}_D - 1}, \quad \alpha = 1 - \frac{\sum_{\gamma} \sum_{\delta} \sum_{\zeta} (u_{\gamma\delta\zeta}^K(\vec{X}_{\gamma\delta\zeta}) - u_{ijk}^K(\vec{X}_{\gamma\delta\zeta}))^2}{\sum_{\gamma} \sum_{\delta} \sum_{\zeta} (u_{\gamma\delta\zeta}^K(\vec{X}_{\gamma\delta\zeta}) - \bar{u}_{ijk})^2}, \quad (7)$$

where \mathcal{N}_{SOS} stands for the size of stencil, \mathcal{N}_D degrees of freedom and the ranges of the indices ($\gamma\delta\zeta$) are taken to include either the whole or a subset of the supporting reconstruction stencil for the cell (i, j, k) . Based on the magnitude relative to a chosen cut-off value \mathcal{S}_C , a larger value of \mathcal{S} indicates smooth variations while a small value of \mathcal{S} indicates non-smooth or under-resolved solution content. For all simulations presented in Section V, the smoothness indicator cut-off value is $\mathcal{S}_C = 1500$.

III.A. Choice of Variables for Reconstruction

Like other Godunov-type finite volume methods, the proposed CENO scheme has the flexibility of carrying out the solution reconstruction on either the conserved variable vector, $\mathbf{U} = [\rho, \rho\vec{V}, \vec{B}, \rho e, \psi]^T$, or the primitive variable vector, $\mathbf{W} = [\rho, \vec{V}, \vec{B}, p, \psi]^T$. Other choices, such as the entropy variables are also possible. While the use of conserved variables in the reconstruction would seem to be the natural choice, it is known that the reconstruction in terms of this set provides the worst shock capturing fidelity when compared to the reconstruction performed either in primitive or in characteristic variables.²⁸ As a result, the use of primitive variables have been preferred to date in the application of the CENO reconstruction and limiting procedure (and indeed in many Godunov-type methods) so as to maintain tighter control on the positivity of certain flow variables (density and pressure). Unfortunately, as demonstrated by Charest²⁹, the point-wise mapping to convert average conserved solution variables into average primitive solution values introduces a temporal error that is at best of second-order, leading to a loss of accuracy for unsteady flow problems. Charest²⁹ and Ivan and Groth⁵ have shown that this error is not introduced for steady problems and the expected formal accuracy is always recovered in this case.

To maintain a tight control on the flow variables and avoid the introduction of a temporal error, the CENO scheme implementation has been reformulated in this work. An unlimited high-order reconstruction of conserved variables is performed in regions of smooth flow and the scheme reverts to a limited second-order reconstruction of primitive variables where the solution is non-smooth. In Section V.A this new set of combined variables is evaluated and compared to the original 3D CENO scheme where primitive variables are reconstructed with high-order and low-order accuracy. The convergence study provided for an unsteady flow problem shows that the high-order accuracy is recovered in the asymptotic limit only for the set of combined variables, while still allowing good control of monotonicity.

III.B. Hexahedral Elements and Trilinear Transformation

In order to obtain fourth-order accuracy on generic hexahedral volumes, the non-planar faces must be carefully treated with regard to all geometric operations affecting the numerical procedure such as flux integration and calculation of geometric properties (e.g., area, volume, centroid, moments, etc.). The solution adopted by Ivan *et al.*⁴ in which general hexahedral cells are defined in terms of a trilinear mapping of the faces is used in the present work. This allows geometric computations to be performed with high-order accuracy by transforming the element to a reference cubic cell.

III.B.1. Trilinear Transformation

The trilinear transformation maps a hexahedron from physical space into a reference unitary cube in the canonical space. Figure 3 illustrates the procedure to map a unitary cube to a generic hexahedron (with non-planar faces) by the following trilinear transformation

$$\vec{r}(p, q, r) = \vec{A} + \vec{B}p + \vec{C}q + \vec{D}r + \vec{E}pq + \vec{F}pr + \vec{G}qr + \vec{H}pqr, \quad (8)$$

where p, q and r are Cartesian coordinates in the canonical space of the reference cube, and $\vec{A}, \vec{B}, \vec{C}, \vec{D}, \vec{E}, \vec{F}, \vec{G}$ and \vec{H} are the transformation vector coefficients that are computed by imposing the one-to-one correspondence between the vertices of the hexahedron and those of the reference cube. Once the




Figure 3. A general hexahedral cell in physical space having faces with non-planar vertices (left) to which a reference unit cube (right) is mapped by applying a trilinear transformation $\vec{r}(p, q, r)$. Four local unit normals are also shown for one face of the general hexahedron.

transformation coefficients are determined, any local or integrated quantities involving the geometry of the hexahedron (e.g., centroid, volume, face normal, etc.) can be computed by carrying out the required evaluations within the reference cube.

IV. Parallel Anisotropic Block-Based AMR

A flexible block-based hierarchical binary tree data structure is used in conjunction with the spatial discretization procedure described in Section II to facilitate automatic solution-directed mesh adaptation on body-fitted multi-block mesh. The general AMR framework of Freret and Groth,²⁰ based on extensions to the previous work by Williamschen and Groth,¹⁹ is well suited and readily allows the use of high-order spatial discretization by adopting a non-uniform representation of the cells within each block. Anisotropic mesh refinement and efficient and highly scalable parallel implementation are achieved via domain partitioning.

In the proposed anisotropic AMR scheme, mesh adaptation is accomplished by refining and coarsening grid blocks. Figure 4 shows the resulting binary tree after several refinements of an initial mesh consisting of a single block. Each refinement produces new blocks called “children” from a “parent” block and the children can be refined further. This refinement process can be reversed in regions that are deemed over-resolved and two, four or eight children are coarsened or merged into a single parent block. A flexible hierarchical tree data structure is used for tracking block connectivity and mesh refinement history. Within this structure, each node of the binary tree stores references to two child nodes, corresponding to a grid block refinement in one of three coordinate directions. Local refinement and coarsening of the mesh can be directed according to adjoint error estimation strategies^{30–32} or so-called physics-based refinement criteria.^{19,20} While the former are expected to be more optimal, the latter have been used extensively because of their relative simplicity.

The anisotropic block-based AMR scheme implemented in this study employs six steps. They are:




Figure 4. 3D binary tree and the corresponding blocks after several anisotropic refinements.

1. **Refinement Flags Evaluation:** blocks are flagged to refine, coarsen or keep their refinement level for each computational direction based on refinement criteria calculated for each block,
2. **Conflict Checking:** refinement flags are modified to ensure that there is no violation in refinement or coarsening rules,
3. **New Block Computation:** blocks are refined and/or coarsened in the flagged directions. The binary tree is updated to reflect the new grid topology,
4. **Connectivity Rearrangement:** connectivity rearrangement is performed to ensure that parents of all blocks flagged to coarsen are split in the appropriate direction,
5. **Neighbor Information Update:** neighbor information or computation of solution block connectivity is re-computed and stored for each block,
6. **Message Passing:** solution and geometry information is shared between adjacent blocks through message passing procedure.

As outlined by Williamschen and Groth,¹⁹ restrictions on the assignment of the refinements flags are used to eliminate and prevent undesirable mesh topologies including: 1) the maximum level difference in any direction being greater than one, i.e., any neighbor can be only one level coarser or one level finer; 2) a block being flagged to coarsen that does not have a neighbor/sibling flagged to coarsen or after refinement the neighbor/sibling is not at the same refinement levels; 3) an invalid binary tree connectivity. In order to get an optimal anisotropic mesh, faces associated with a given direction are allowed not to have the same level in both tangential directions; an example of this configuration is given in Figure 5-a. A face is also allowed not to span its neighbors; this configuration is represented in Figure 5-b. Being free of these constraints requires extra care during the coarsening procedure as some connectivity rearrangements of the binary tree may be required,¹⁹ and calculation of the numerical fluxes may occur through faces that are not a priori defined, as shown in Figure 5-b. In terms of computation time, allowing these configurations adds very little to the actual overall computational cost of the problem while greatly increasing the effectiveness of the anisotropic AMR algorithm.

IV.A. High-Order Solution Transfer

In the refinement process of a grid block, the geometry of the newly created cells belonging to the offspring is obtained by dividing the domain of each coarse interior cell into two, four or eight fine cells. In order to maintain the high-order solution accuracy for combination of the CENO algorithm with the aforementioned block-based AMR approach, accurate solution transfer operators must be formulated. One advantage of using non-uniform blocks is that restriction and prolongation operators are only required in the interior domain to transfer the solution content from coarse to fine cells, but not in the ghost cells for overlapping cell domains. In particular, to distribute the average solution quantity among offspring with high-order accuracy the high-order polynomial reconstructions of all solution variables in the coarse cell are integrated over the domain of each new fine cell having a volume, V_{fine} ,

$$\bar{u}_{fine} = \frac{1}{V_{fine}} \iiint_{V_{fine}} u_{coarse}^k(\vec{X}) dV, \quad (9)$$

where the volume integral is computed exactly for the reconstruction polynomial with an appropriate-order tensor-product Gauss quadrature volumetric integration technique (twenty-seven quadrature points are used for fourth-order spatial accuracy).



Figure 5. Examples of refined blocks that do not violate the 3D rules.

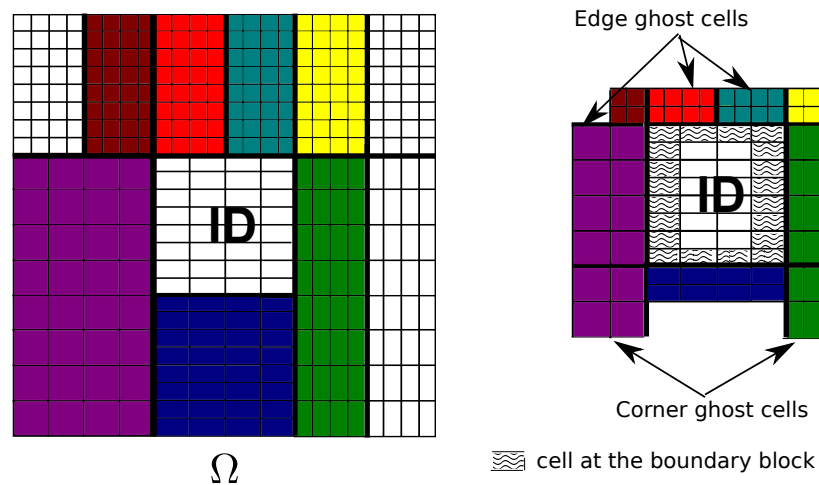


Figure 6. Example of a domain Ω divided into 11 blocks of 4×8 cells (left) and illustration of a non-uniform structured mesh block composed by the interior domain (ID) extended by a layer of 2 ghost cells provided by the neighboring blocks (right). The resulting non-uniform block has 5 edge blocks and 4 corner blocks.

IV.B. Non-Uniform Blocks

In order to apply the finite-volume scheme to all blocks in a more independent manner, solution information is shared between adjacent blocks having common interfaces. This information is stored in additional layers of overlapping ghost cells associated with each block. The non-uniform representation of the blocks adopted here within the multi-block structure uses directly the neighboring cells as the ghost cells, even for those at different levels of refinement as found at grid resolution changes.²⁰ An example of a non-uniform block obtained from a multi-block structure is shown in Figure 6.

The non-uniform block treatment presents many advantages as outlined by Freret and Groth:²⁰ First no restriction or prolongation operators at resolution changes are required to evaluate the ghost cell solution values. Next, there is no need for an additional correction to enforce the flux conservation properties of the finite-volume scheme. This property is automatically satisfied. Finally, as considered in the present work, an efficient way to treat the ghost cells is proposed. In general, the number of ghost cell layers depends on the size of the reconstruction stencil and the size of the stencil used in evaluating the smoothness indicator. Assuming a two-layer stencil is used to compute the high-order reconstruction and a one-layer stencil for evaluating the smoothness indicator, the total number of ghost cell layers required at each block boundary is four. While in the original implementation of Ivan *et al.*^{4,18} four layers of ghost cells were used to get a fourth-order accurate solution reconstruction for the cells in the interior domain and the first two layers of ghost cells, just two layers of ghost cells are used in the current implementation. Instead of reconstructing the solution in ghost cells, the new formulation relies on communicating any additional needed information (i.e., solution data, geometry, derivatives) from the corresponding interior cells of the neighboring block to the ghost cells, which are used only as a storage space. In this way, the number of ghost cells is minimized at the expense of increasing somewhat the size and number of the send/receive information between neighboring blocks. For example, in the simulation presented in Section V.E, this implementation represents a memory saving of 60% compared to the original approach.

The non-uniform block strategy is also very convenient when cubed-sphere meshes are considered. As discussed by Ivan *et al.*,⁴ the cubed-sphere grids have degenerate corners characterized by missing ghost cells. Since stencils need to be consistent between a cell located in the interior domain and the same cell as a ghost cell in a neighboring block, Ivan *et al.*⁴ proposed a rotation mechanism for generating consistent stencils near degenerate block edges. With the new non-uniform block treatment of the ghost cells, this relatively complicated technique is no longer required.

One challenge with the present formulation is that the different refinement configurations that can occur with anisotropic AMR have to be taken into account during the procedure for the construction of the reconstruction stencil. As the current anisotropic AMR implementation is extremely flexible, there are up to 16 different types of resolution changes at block boundaries that can be summarized in four groups as shown in Figure 7. In the stencil representation, cells colored yellow represent the interior domain of a block and the

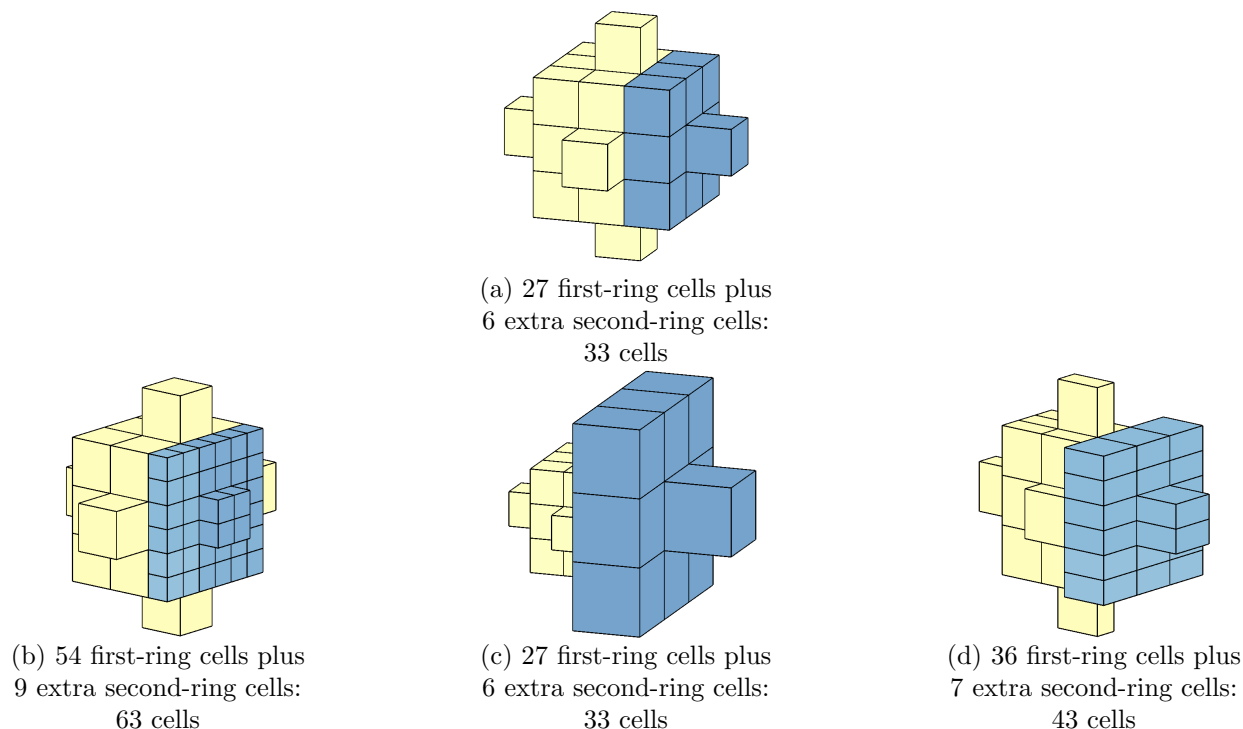


Figure 7. Examples of cell stencil construction in a Cartesian mesh for a cell (i, j, k) located at the block boundary such that its stencil includes ghost cells from neighboring blocks. The interior domain is represented in yellow and blue is for the ghost cells. (a) Neighbor has the same mesh refinement. Stencil is identical as an interior cell. (b) Two neighbors with finer mesh are facing the block. Neighbor search direction includes more neighbors. (c) Neighbor has a coarser mesh. Number of cells identical to uniform configuration. (d) Neighbor has a finer mesh in a direction and a coarser mesh in the other direction. A total of 43 ghost cells are included in the stencil.

blue colored cells represent ghost cells for that block. All four different stencils are based on a 33-cells stencil which has been found optimal for the CENO scheme.⁴ Figure 7(a) is a stencil obtained for a cell (i, j, k) located at the boundary block facing a neighbor with the same mesh refinement. The shape of the stencil is exactly the same as an interior cell stencil and counts 33 cells. Figure 7(b) represents a stencil built with a neighbor having a finer mesh. Because the mesh is finer in both tangential directions, the number of cells included in the stencil is equal to 63. With a finer mesh in one direction only, the number of cells included in the stencil would be 43. Figure 7(c) depicts a stencil of a cell having a neighbor with a coarser mesh. The shape of the stencil is different from the uniform case but the number of cells included in the stencil remains 33 since the neighboring mesh is coarser in the two tangential directions. With a mesh coarser in just one direction, the number of cells would be 43. Figure 7(d) is specific to anisotropic AMR since the neighbor represented in blue has a finer mesh in a direction and a coarser mesh in the other direction, a situation which cannot occur with isotropic refinement. Again the shape of the stencil is very close to the uniform reference stencil of Figure 7(a) but includes a total of 43 cells. Thus for all mesh refinement configurations of a Cartesian grid, at least 33 neighbors are found in the cell stencil ensuring that Eq (6) is overdetermined.

As explained in Freret and Groth,²⁰ the numerical flux evaluation also becomes more challenging when a non-uniform mesh is used, more particularly at block boundaries. One must now account for the 16 different types of resolution changes at block boundaries that can be summarized in three groups as shown in Figure 8. In Figure 8, case (a) involves a cell which has an equal or finer mesh in one or several directions than its neighbor ghost cell. Case (b) shows an example of a cell having a neighbor with a coarser mesh discretization and case (c) illustrates two cells with 2 opposite resolution changes in two different directions, as also shown in Figure 5-b. Numerical evaluations of the fluxes for case (a) and (b) are rather straightforward since they represent a situation where a coarse face spans its finer neighbor. For case (a) all needed quantities for the flux computation are based on the interior cell, as shown in Figure 8 (top). For case (b) the evaluation of the fluxes is done through the faces of the ghost cells, taking the opposite of the normal vector, as depicted in Figure 8 (middle). The high-order flux evaluation is done by summing the flux through the finer neighboring

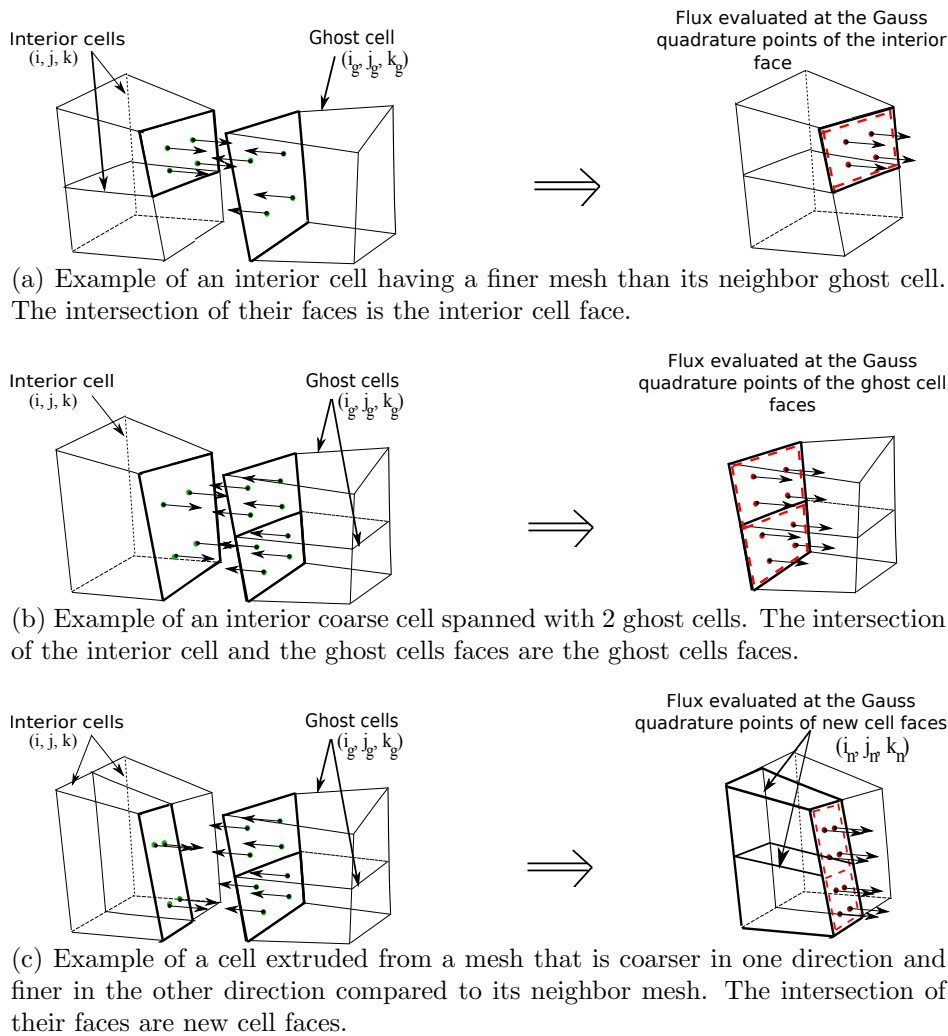


Figure 8. Three scenarios to evaluate the numerical flux through a face of an interior cell depending on its mesh resolution and its neighbor cell mesh resolution.

cells to get the total flux of a coarser cell. In Figure 8 (bottom), case (c) requires some extra considerations since new Gauss quadrature points and associated weights have to be evaluated on the face that intersects both cells with opposite resolution changes. The total high-order numerical flux is obtained by summing the flux computed on each sub-face.

V. Numerical Results

A number of numerical results are now presented to demonstrate the accuracy of the overall high-order CENO finite-volume scheme with adaptive mesh refinement. Results for uniform, isotropic and anisotropic AMR meshes are included as well as cubed-sphere grids to provide a broad range of evaluation of the AMR scheme. In Section V.A two sets of reconstruction variables within the CENO scheme are first studied for the 3D iso-density vortex transient problem. Euler and MHD results are next considered to evaluate grid convergence for smooth steady flows (Sections V.B and V.C) and robustness of the high-order method at shocks (Section V.D). In Section V.E, the AMR procedure is applied for the 3D iso-density magnetized vortex unsteady problem to show the ability of the AMR scheme to capture some features or quantities of interest even for a smooth flow. For steady time-invariant problems the multi-stage optimal smoothing scheme³³ with local time stepping is used. For time-accurate calculations, the standard four-stage fourth-order Runge-Kutta time-marching scheme is employed. The HLLE and Lax-Friedrichs numerical flux function are used

for the Euler and MHD equations, respectively. The stencils based on 33 cells from Figure 7 are used to obtain the high-order results.

V.A. Choice of Reconstruction Variables

As a first evaluation of the CENO scheme, the numerical convergence of the local and global error norms of different flow quantities for an unsteady flow problem is examined. The MHD flow problem of interest consists of a magnetized vortex structure in force equilibrium advected by a uniform flow field from Mignone *et al.*³⁴ This is a smooth time-dependent 3D problem with an exact solution. The stationary 3D iso-density MHD vortex flow is given by

$$\begin{aligned} \rho &= 1, \\ \vec{V} &= (-y, x, 0)\kappa \exp(q(1 - R^2)), \\ \vec{B} &= (-y, x, 0)\mu \exp(q(1 - R^2)), \\ p &= 1 + \frac{1}{4q}(\mu^2(1 - 2q(R^2 - z^2)) - \kappa^2\rho) \exp(2q(1 - R^2)), \end{aligned}$$

where $\mu = \kappa = 1/(2\pi)$, and $q = 1$. This solution is translated with a background velocity that equals $(1, 1, 2)$. This flow problem is simulated on a periodic Cartesian box with x, y and z varying between $[-5, 5]$. Figure 9 depicts the flow solution on the Cartesian grid. Initially the vortex center is positioned at $(0, 0, 0)$ and the simulation runs until time $t = 10$ (the vortex has done one period and is back to its original position). A convergence study is performed to assess the accuracy of the CENO scheme for unsteady flow problems. Each block contains $8 \times 8 \times 8$ cells and the initial grid of 64 blocks is uniformly refined to reach 262,144 blocks for the largest simulation.

The first series of simulations for the iso-density vortex problem was performed herein in which reconstruction of the primitive variables was considered for both the high-order unlimited reconstruction and the low-order limited reconstruction of the hybrid CENO method. The numerical results for these cases are presented with dashed lines in Figures 10(a)-(c). The L_1, L_2 and L_∞ error norms for the solution obtained on the sequence of uniformly refined grids are shown for a subset of the primitive variables vector $\mathcal{S}_{prim} = \{u, B_x, p\}$. For the initial mesh (32^3 cells), 56% of the cells are deemed as non-smooth and a limited linear reconstruction is performed for all variables in these cells. For the second mesh (64^3 cells), only 17% of the cells are detected with a non-smooth content and reverted to a second-order limited reconstruction. Hence a global second order accuracy is observed for all variables for the two coarsest meshes in Figure 10(a)-(c). For the next refined meshes, all cells are flagged as smooth and a fourth-order unlimited polynomial reconstruction is performed in all cells. However, the theoretical fourth-order accuracy is not achieved. Instead, the effective order of convergence is close to 2 for the three variables and for all error norms. This result agrees with the findings of Charest²⁹ which indicates a loss of accuracy when high-order reconstruction of primitive variables is used in the solution of unsteady problems, even though formal accuracy is retained for steady problems.^{5,29}

A second series of simulations was performed in which a high-order reconstruction of the conserved variables was performed and, where the flow was deemed as non-smooth, low-order reconstruction of the primitive variables was carried out. The corresponding numerical results are presented with solid lines in

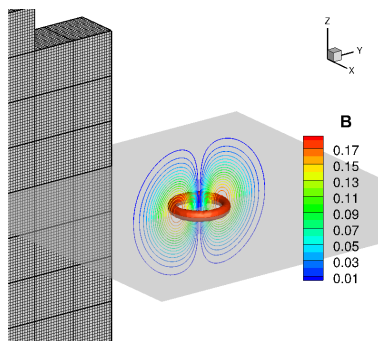





Figure 9. Magnetic field iso-surface ($B = 0.18$) and contours of the initial flow solution obtained on a Cartesian grid.



(a) Error norms in the velocity (dashed lines) and momentum (solid lines) versus grid size.



(b) Error norms in the magnetic field as a primitive variables (dashed lines) and conserved variables (solid lines) versus grid size.



(c) Error norms in the pressure (dashed lines) and energy (solid lines) versus grid size.

Figure 10. Comparison of L_1 , L_2 and L_∞ error norms with high- and low-order reconstruction with primitive variables (dashed lines) and a mixed polynomial reconstruction with high-order reconstruction using conserved variables and low-order reconstruction using primitive variables (solid lines).

Figures 10(a)-(c). Figure 10 shows the L_1 , L_2 and L_∞ error norms for a subset of the conserved variables $S_{cons} = \{\rho u, B_x, \rho e\}$ obtained on these series of grids. As the solution content is deemed as non-smooth for the two coarsest meshes, a low-order reconstruction is performed in these cells and a global second-order accuracy is provided. However, for further uniformly refined meshes where all solution content is detected as smooth, the fourth-order theoretical accuracy is recovered for all solution variables and for all error norms. Hence, for unsteady flow problems, the CENO scheme does not suffer from the temporal error introduced by the transformation from primitive to conserved variables.²⁹ Instead, a high-order accuracy is provided in case of smooth unsteady flow problem by considering conserved variables reconstruction and, the scheme reverts automatically to a low-order accuracy with a tight control on the flow quantities by reconstructing primitive variables. The possibility of switching reconstruction variables within the hybrid CENO scheme is certainly an attractive feature of the method when monotonicity of the solution is a concern.

V.B. Steady Supersonic Outflow of Non-Magnetized Plasma

To assess the accuracy of the finite-volume scheme on cubed-sphere grids, numerical convergence studies for a spherically symmetric expanding supersonic non-magnetized plasma flow have been performed and are now

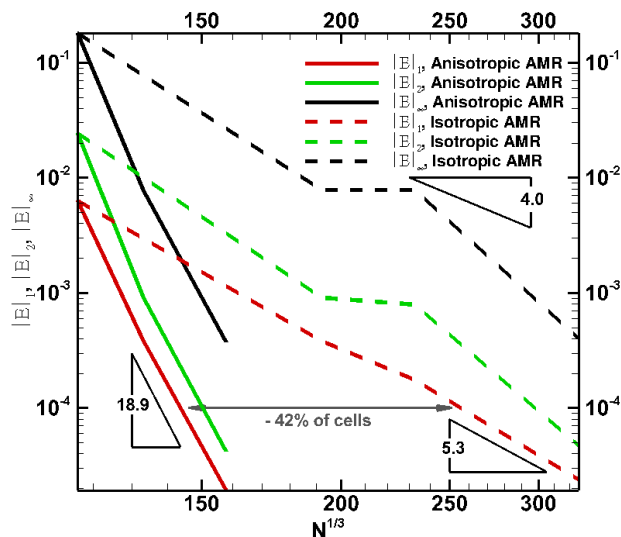


Figure 11. L_1, L_2, L_∞ error norms of the solution density obtained with the fourth-order CENO scheme with anisotropic AMR (solid lines) compared to isotropic AMR (dashed lines).

considered. The accuracy of the fourth-order CENO scheme for a series of isotropic and anisotropic refined AMR meshes was determined and is compared here. The computational domain of interest is defined by inner and outer spheres of radius $R_i = 1$ and $R_o = 4$, respectively. The initial mesh is composed of $128 \times 128 \times 16$ cells for each of the 6 sectors forming the cubed-sphere grid. The mesh undergoes three adaptive mesh refinements and is converged to machine accuracy on each level of grid. The refinement criteria is based on the gradient of density for both anisotropic and isotropic refinements. For boundary data, the exact solution is imposed on the inner sphere: $\rho_i = 10$, $V_{r,i} = 4.5$, $\vec{V}_{||,i} = \vec{0}$ and $p_i = 26$. An outflow supersonic boundary condition is imposed at R_o . As described by Ivan *et al.*,¹⁸ the analytical solution of this flow problem can be obtained in spherical coordinates as the solution of the following equation

$$C_3 - \frac{1}{r^2 V_r \left[(C_2 - V_R^2)^{\frac{1}{\gamma-1}} \right]} = 0, \quad (10)$$

where C_2 and C_3 are constants depending on the inflow conditions. The resulting L_1, L_2 and L_∞ norms of the error in the predicted solution density obtained on a series of grids are given in Figure 11.

The convergence results of Figure 11 show that both isotropic AMR (dashed lines) and anisotropic AMR (solid lines) provide a convergence rate higher than the theoretical fourth-order accuracy. For anisotropic refinement of the mesh, the effective convergence rate approaches 19 and for isotropic AMR it is close to 5. As noted previously^{19,20}, the solution varies only along the radial direction and the anisotropic AMR exploits this feature by refining only in the radial direction, thus avoiding the introduction of an unnecessary large number of computational cells. For an error target $L_1 = 10^{-4}$, the memory requirement of the anisotropic AMR is only 58% of the memory requirement of the isotropic AMR. A theoretical fourth-order accuracy of the uniform refinements leads to a mesh saving of 88% for the anisotropic AMR compared to uniform AMR. Finally it is worth mentioning that the physics-based refinement criteria does not control any reduction in the error norms. For that reason a plateau in L_2 and L_∞ error norms is observed with the isotropic refinement for the second level of refinement. To improve this behavior an error-based refinement strategy can be used as the one developed by Venditti and Darmofal^{30,31} and Narechania *et al.*³²

V.C. Steady Supersonic Outflow of Magnetized Plasma

The 3D steady-state axi-symmetric solution of a MHD plasma is now considered on a spherical shell domain flowing outward at superfast speeds as proposed by Ivan *et al.*¹⁸ The exact solution is given by

$$\begin{aligned} \rho &= R^{-5/2}, \\ \vec{V} &= \left(\frac{x}{\sqrt{R}}, \frac{y}{\sqrt{R}}, \frac{z}{\sqrt{R}} \right) + \kappa R^{-5/2}, \end{aligned}$$

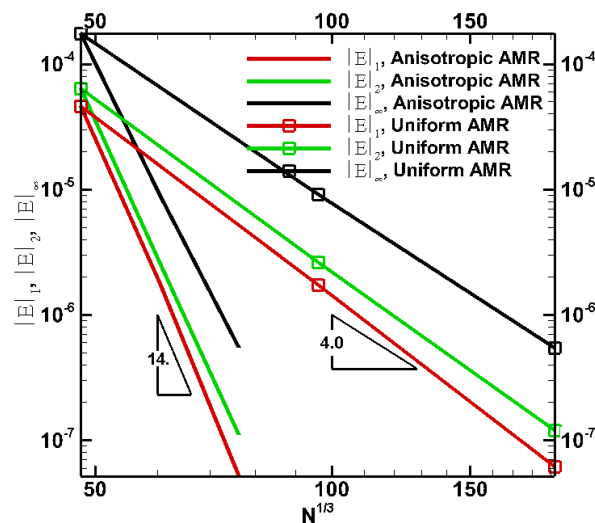


Figure 12. Error norms versus grid size for the first component of the magnetic field B_x .

$$\begin{aligned}\vec{B} &= \left(\frac{x}{R^3}, \frac{y}{R^3}, \frac{z}{R^3} \right) + \kappa, \\ p &= R^{-5/2},\end{aligned}$$

and corresponds to the solution having the additional source term in the governing equations given by

$$\mathbf{Q} = \begin{bmatrix} 0 \\ \frac{1}{2}xR^{-5/2}(R^{-1} - 5R^{-2} - \kappa z) \\ \frac{1}{2}yR^{-5/2}(R^{-1} - 5R^{-2} - \kappa z) \\ \frac{1}{2}zR^{-5/2}(R^{-1} - 5R^{-2} - \kappa z) + \frac{5}{2}R^{-1/2}\kappa(1 + \kappa Rz) + \kappa R^{-1/2} \\ 0 \\ \frac{1}{2}R^{-2} + \kappa z(3.5R^{-1} + 2\kappa z) + \frac{(\kappa R)^2}{2}(7 + 5\kappa Rz) \end{bmatrix}.$$

As suggested by Ivan *et al.*,^{4,18} the perturbation parameter is taken as $\kappa = 0.017$ such that the solution has significant latitudinal variation. In this flow, the magnetic field is irrotational and aligned everywhere with the velocity. The computational domain used for this convergence study is defined by inner and outer spheres of radius $R_i = 2$ and $R_o = 4$. To achieve high-order accuracy for this problem, it is necessary to provide a high-order approximation to the average source term, $(\mathbf{Q})_{ijk}$, in the numerical residual by integrating the analytical expression of the source term with high-order accuracy. Moreover, high-order boundary conditions must be imposed, which we achieved by specifying both the inflow and outflow boundary conditions using high-order accurate integration of the exact solution in ghost cells to determine the ghost cell averages.

Figure 12 depicts the error norms in density as a function of grid size for the supersonic outflow problem. The theoretical fourth-order of convergence is obtained with the CENO scheme on the cubed-sphere for all error norms with successive uniformly refined grid. With anisotropic AMR, the convergence rate reaches the value of 14. In terms of computational cells, to reach $L_1 = 5e^{-08}$, the anisotropic AMR simulation uses only 6.25% of the total number of cells of the uniform AMR. This validates the approach to obtain high-order accuracy using non-uniform stencils at different mesh refinement levels and a minimal ghost cells strategy that ensures consistent stencils near boundaries and corners.

V.D. Non-Magnetized Bow Shock

The application of the proposed AMR algorithm to a 3D non-magnetized bow-shock flow around a perfectly conducting sphere is considered next. The inflow solution parameters used for this problem are $\rho = 1$, $v_x = 1.4943$, $v_y = v_z = 0$ and $p = 0.2$. The computational domain is defined by inner and outer spheres of radius $R_i = 1$ and $R_o = 4$. The mesh contains only five sectors instead of the typical six associated

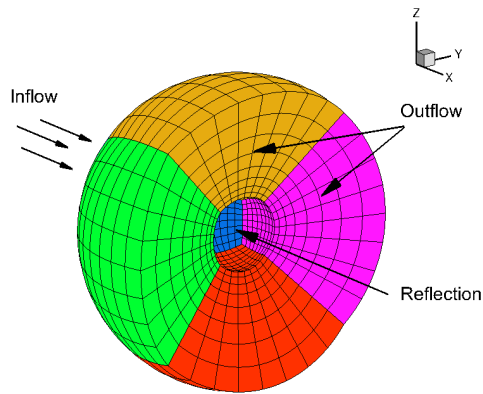


Figure 13. Initial mesh with five sectors highlighting the boundary conditions

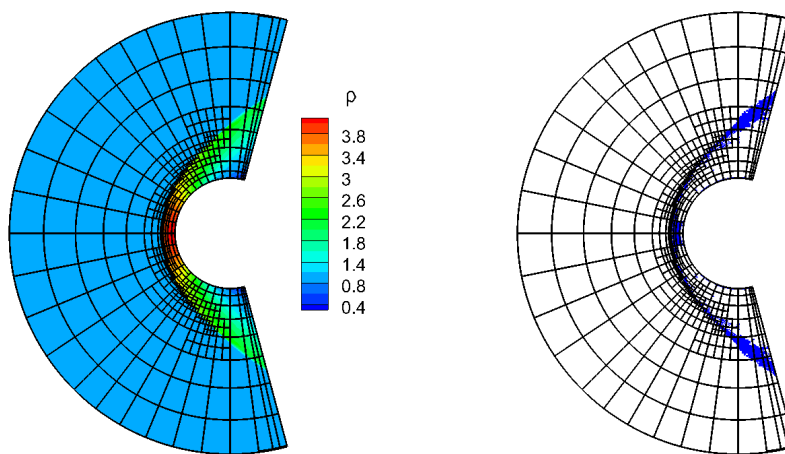


Figure 14. Slice $z = 0$ of the anisotropic mesh and density field obtained on the fifth level of mesh (left) and solution smoothness indicator (in blue) indicating cells deemed as non smooth (right).

with a cubed-sphere grid to reduce the number of computational cells and focus on the upstream side of the sphere, as shown in Figure 13. Reflection boundary condition is imposed on the inner sphere, inflow boundary condition is applied on the outer sphere and an outflow boundary condition is used for the back panels of the outer spheres. The initial grid consists of a total of 320 blocks of $8 \times 8 \times 16$ cells with 8 blocks in the radial direction.

The predicted density distribution and the smoothness indicator in the Cartesian plane $\{(x, y), z = 0\}$ obtained after 5 adaptive refinement levels are shown in Figure 14. This final mesh consists in 7986 blocks and 8,177,664 computational cells. This steady-state solution is obtained by iteratively getting a converged solution on a given mesh and then adapt the grid according to the refinement criteria based on the gradient of the density. This procedure is repeated five times. As shown in Figure 14-(left), refinement of the mesh occurs along the shock profile and near the stagnation point. The refinements is mostly anisotropic and aligned with the bow shock. The outer regions of the bow shock are isotropically refined. The number of cells deemed as non-smooth is around 10% for the initial mesh and up to 16% for the final mesh. In Figure 14-(right), the smoothness indicator is shown in blue for the final mesh. As can be seen, only cells located near the shock are detected as non-smooth. The smoothness indicator thus provides a very good way to identify the shock. This makes it a very good candidate for directing the mesh and it will be considered for future work.

V.E. 3D Iso-density MHD Vortex in a Cartesian Box

As a final example, the ability of the AMR procedure to follow the advection of a 3D iso-density vortex in a velocity background is evaluated. The problem setup is identical to that described in Section V.A. The number of root blocks in the initial mesh was modified so as to contain $3 \times 3 \times 3$ blocks and the overall size of the domain was $[-20, 20]^3$ such that the vortex represents a smaller volume of the domain Ω . This provides a domain of sufficient size to evaluate properly the performance of the AMR scheme. The final time was selected to be $t = 40$ such that the vortex has done one period and is back to its initial position. All other parameters are unchanged. Before the simulation begins, the initial mesh with 27 root blocks containing $8 \times 8 \times 8$ cells was refined five times based on the initial conditions and a maximum refinement level of five is assigned. Since this problem is mostly isotropic, a refinement based on the curl of the velocity field is used as an isotropic refinement criteria.

Figure 15 shows 6 slices of the magnitude of the magnetic field with the mesh refinement plotted at $t = 0$ (a), $t = 8$, (b) and $t = 16$, (c) from top-left to top-right and $t = 24$ (d), $t = 33$, (e) and $t = 40$, (f) from bottom-left to bottom-right. Following the initial refinements, the AMR refinement and coarsening procedures are repeated every 10 time-steps until the maximum time $t = 40$ is reached. Throughout the simulation, the AMR tracks the vortex well, clustering the blocks in the vortex area. It is worth mentioning that the final mesh has the same number of blocks with the same repartition of the blocks as the initial mesh. There is no deterioration of the quality of the mesh. An overall reduction of 99% in the cells number is obtained compared to that of the usual uniform mesh with a resolution equal to that on the finest level.

VI. Conclusions

A high-order CENO finite-volume method combined with an efficient anisotropic block-based AMR scheme has been proposed and described for the solution of the compressible form of the ideal MHD equations on three-dimensional multi-block body-fitted hexahedral meshes. The proposed high-order AMR scheme is

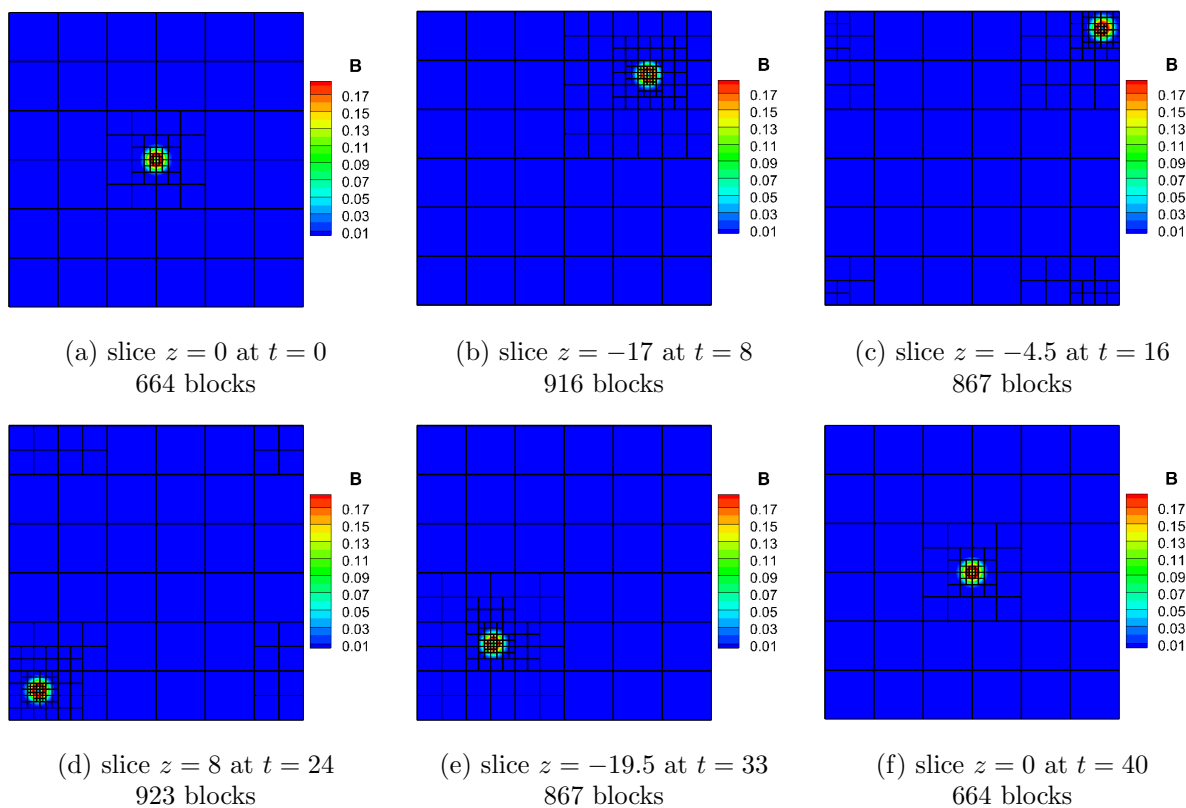


Figure 15. Isotropic mesh and magnitude of the magnetic field obtained at (a) $t = 0$, (b) $t = 8$, (c) $t = 16$, (d) $t = 24$, (e) $t = 33$ and (f) $t = 40$ in a XY plane of a periodic 3D box.

readily applicable to cubed-sphere meshes and has considerable potential for the global MHD solution of space plasmas. Numerical results have been presented and compared to analytical solutions for several benchmark problems. For these cases, the formal accuracy of the high-order CENO method has been established. Furthermore, the ability of the anisotropic AMR technique to provide efficient local refinement of the multi-block grid has been demonstrated for smooth unsteady flow problems as well as for steady flows with strong shocks.

Acknowledgments

This work was supported by the Canadian Space Agency and by NSERC of Canada. Computational resources for performing all of the calculations reported herein were provided by the SciNet High Performance Computing Consortium at the University of Toronto and Compute/Calcul Canada through funding from the Canada Foundation for Innovation (CFI) and the Province of Ontario, Canada. Additionally, the authors would like to acknowledge the financial support received from the Canadian Space Agency through the Geospace Observatory Canada.

References

- ¹Groth, C. P. T., De Zeeuw, D. L., Gombosi, T. I., and Powell, K. G., “Global Three-Dimensional MHD Simulation of a Space Weather Event: CME Formation, Interplanetary Propagation, and Interaction with the Magnetosphere,” *Journal of Geophysical Research*, Vol. 105, No. A11, 2000, pp. 25,053–25,078.
- ²De Zeeuw, D. L., Gombosi, T. I., Groth, C. P. T., Powell, K. G., and Stout, Q. F., “An Adaptive MHD Method for Global Space Weather Simulations,” *IEEE Transactions on Plasma Science*, Vol. 105, pp. 1956–1965.
- ³Clauer, C. R., Gombosi, T. I., De Zeeuw, D. L., Ridley, A. J., Powell, K. G., van Leer, B., Stout, Q. F., and Groth, C. P. T., “HighPerformance Computer Methods Applied to Predictive Space Weather Simulations,” *IEEE Transactions on Plasma Science*, Vol. 28, 2000, pp. 1931–1937.
- ⁴Ivan, L., De Sterck, H., Susanto, A., and Groth, C. P. T., “High-order central ENO finite-volume scheme for hyperbolic conservation laws on three-dimensional cubed-sphere grids,” *Journal of Computational Physics*, 2015.
- ⁵Ivan, L. and Groth, C. P. T., “High-order solution-adaptive central essentially non-oscillatory (CENO) method for viscous flows,” *Journal of Computational Physics*, Vol. 257, 2013, pp. 830–862.
- ⁶Tobaldini Neto, L. and Groth, C. P. T., “A High-Order Finite-Volume Scheme for Large-Eddy Simulation of Turbulent Premixed Flames,” *AIAA paper*, January 2014.
- ⁷Charest, M. R. J. and Groth, C. P. T., “A High-Order Central ENO Finite-Volume Scheme for Three-Dimensional Low-Speed Viscous Flows on Unstructured Mesh,” *Communications in Computational Physics*, Vol. 17, 2015, pp. 615–656.
- ⁸Charest, M. R. J. and Groth, C. P. T., “A High-Order Central ENO Finite-Volume Scheme for Three-Dimensional Turbulent Flows on Unstructured Mesh,” *AIAA paper*, June 2013.
- ⁹Berger, M. J. and Olinger, J., “Adaptive Mesh Refinement for Hyperbolic Partial Differential Equations,” *Journal of Computational Physics*, 1983.
- ¹⁰Quirk, J. J., *An Adaptive Grid Algorithm for Computational Shock Hydrodynamics*, Ph.D. thesis, Cranfield Institute of Technology, January 1991.
- ¹¹Bell, J., Berger, M., Saltzman, J., and Welcome, M., “Three-Dimensional Adaptive Mesh Refinement for Hyperbolic Conservation Laws,” *SIAM Journal on Scientific Computing*, Vol. 15, No. 1, January 1994, pp. 127–138.
- ¹²Gao, X. and Groth, C., “A parallel solution-adaptive method for three-dimensional turbulent non-premixed combustions flows,” *Journal of Computational Physics*, Vol. 229, No. 9, 2010, pp. 3250–3275.
- ¹³Northrup, S. A. and Groth, C. P. T., “Solution of laminar diffusion flames using a parallel adaptive mesh refinement algorithm,” No. 2005-0547, AIAA, 2005.
- ¹⁴Gao, X., *A Parallel Solution-Adaptive Method for Turbulent Non-Premixed Combusting Flows*, Ph.D. thesis, University of Toronto, 2008.
- ¹⁵Sachdev, J. S., Groth, C. P. T., and Gottlieb, J. J., “A parallel solution-adaptive scheme for multi-phase core flows in solid propellant rocket motors,” *International Journal of Computational Fluid Dynamics*, Vol. 19, No. 2, 2005, pp. 159–177.
- ¹⁶Groth, C. P. T., De Zeeuw, D. L., Powell, K. G., Gombosi, T. I., and Stout, Q. F., “A Parallel Solution-Adaptive Scheme for Ideal Magnetohydrodynamics,” 1999.
- ¹⁷McDonald, J. G. and Groth, C. P. T., *Numerical modeling of micron-scale flows using the Gaussian moment closure*, University of Toronto, 2005.
- ¹⁸Ivan, L., De Sterck, H., Northrup, S. A., and Groth, C. P. T., “Multi-dimensional finite-volume scheme for hyperbolic conservation laws on three-dimensional solution-adaptive cubed-sphere grids,” *Journal of Computational Physics*, Vol. 255, 2013, pp. 205–227.
- ¹⁹Williamschen, M. J. and Groth, C. P. T., “Parallel anisotropic block-based adaptive mesh refinement algorithm for three-dimensional flows,” *21th AIAA Computational Fluid Dynamics Conference*, June 2013.
- ²⁰Freret, L. and Groth, C. P. T., “Anisotropic Non-Uniform Block-Based Adaptive Mesh Refinement for Three-Dimensional Inviscid and Viscous Flows,” *22nd AIAA Computational Fluid Dynamics Conference*, June 2015.
- ²¹Dedner, A., Kemm, F., Kroner, D., Munz, C., Schitzer, T., and Wesenberg, M., “Hyperbolic divergence cleaning for the MHD equations,” *Journal of Computational Physics* 2002, 2002, pp. 645–673.

- ²²Jiang, B. Lin, T. L. and Povinelli, L. A., "Large-scale computation of incompressible viscous flow by least-squares finite element method," *Computer Methods in Applied Mechanics Engineering*, Vol. 144, 1994, pp. 213–231.
- ²³Brackbill, J. and Barnes, D., "The effect of nonzero div B on the numerical solution of the magnetohydrodynamics equations," *Journal of Computational Physics*, Vol. 35, 1980.
- ²⁴Evans, C. R. and Hawley, J. F., "Simulation of magnetohydrodynamics flows: A constrained transport method," *Astrophysical Journal*, Vol. 332, 1988.
- ²⁵De Sterck, H., "Multi-dimensional upwind constrained transport on unstructured grid for shallow water magnetohydrodynamics," *AIAA paper*, 2001.
- ²⁶Balsara, D., "Divergence-free reconstruction of magnetic fields and WENO schemes for magnetohydrodynamics," *Journal of Computational Physics*, Vol. 228, 2009, pp. 5040–5056.
- ²⁷Barth, T., "Recent developments in high order k-exact reconstruction on unstructured meshes," 1993.
- ²⁸Zanotti, O. and Dumbser, M., "Efficient conservative ADER schemes based on WENO reconstruction and space-time predictor in primitive variables," *Computational Astrophysics and Cosmology Simulation, Data Analysis and Algorithms*, Vol. 3, 2016.
- ²⁹Charest, M. R. J., "Effect of variables choices on Godunov-type high-order finite-volume methods," *to be submitted*.
- ³⁰Venditti, D. and Darmofal, D., "Anisotropic Adaptation for Functionnal Outputs of Viscous Flow Simulations," *AIAA paper*, June 2003.
- ³¹Venditti, D. and Darmofal, D., "Anisotropic Grid Adaptation for Functionnal Outputs: Application to Two=Dimensional Viscous Flows," *Journal of Computational Physics*, Vol. 187, pp. 22–46.
- ³²Narechania, N., Freret, L., and Groth, C. P. T., "Block-Based Anisotropic AMR with *A Posteriori* Adjoint-Based Error Estimation for Three-Dimensional Inviscid and Viscous Flows," *to be submitted*.
- ³³van Leer, B., Tai, C. H., and Powell, K. G., "Design of Optimally-Smoothing Multi-Stage Schemes for the Euler Equations," Paper 89-1933-CP, AIAA, June 1989.
- ³⁴Mignone, A., Tzeferacos, P., and Bodo, G., "High-order conservative finite difference GLM-MHD schemes for cell-centered MHD," *Journal of Computational Physics*, Vol. 229, pp. 5896–5920.