



High-Order Finite-Volume Method with Block-Based AMR for Magnetohydrodynamics Flows

L. Freret¹ · L. Ivan^{1,2} · H. De Sterck³ · C. P. T. Groth¹

Received: 8 February 2017 / Revised: 27 September 2018 / Accepted: 28 September 2018 /
Published online: 12 October 2018
© Springer Science+Business Media, LLC, part of Springer Nature 2018

Abstract

A high-order central essentially non-oscillatory (CENO) finite volume scheme combined with a block-based adaptive mesh refinement (AMR) algorithm is proposed for the solution of the ideal magnetohydrodynamics equations. The high-order CENO finite-volume scheme is implemented with fourth-order spatial accuracy within a flexible multi-block, body-fitted, hexahedral grid framework. An important feature of the high-order adaptive approach is that it allows for anisotropic refinement, which can lead to large computational savings when anisotropic flow features such as isolated propagating fronts and/or waves, shocks, shear surfaces, and current sheets are present in the flow. This approach is designed to handle complex multi-block grid configurations, including cubed-sphere grids, where some grid blocks may have degenerate edges or corners characterized by missing neighboring blocks. A procedure for building valid high-order reconstruction stencils, even at these degenerate block edges and corners, is proposed, taking into account anisotropic resolution changes in a systematic and general way. Furthermore, a non-uniform or heterogeneous block structure is used where the ghost cells of a block containing the solution content of neighboring blocks are stored directly at the resolution of the neighbors. A generalized Lagrange multiplier divergence correction technique is applied to achieve numerically divergence-free magnetic fields while preserving high-order accuracy on the anisotropic AMR grids. Parallel implementation and local grid adaptivity are achieved by using a hierarchical block-based domain partitioning strategy in which the connectivity and refinement history of grid blocks are tracked using a flexible binary tree data structure. Physics-based refinement criteria as well as the CENO smoothness indicator are both 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 approach.

✉ L. Freret
lfreret@utias.utoronto.ca

¹ Institute for Aerospace Studies, University of Toronto, 4925 Dufferin Street, Toronto, ON M3H5T6, Canada

² Present Address: Canadian Nuclear Laboratories, 286 Plant Road, Chalk River, ON K0J 1J0, Canada

³ School of Mathematical Sciences, 9 Rainforest Walk, Monash University, Melbourne, VIC 3800, Australia

Keywords High-order finite volume scheme · Anisotropic block-based adaptive mesh refinement · Ideal magnetohydrodynamics (MHD) flows

1 Introduction

High-order numerical methods have been actively pursued in an effort to reduce the cost of large-scale scientific computing applications. High-order methods are often more efficient than low-order methods in terms of computational cost and CPU time for the same target accuracy [37,65]. The high-order central essentially non-oscillatory (CENO) finite-volume scheme described in Ivan et al. [37,38] uses a hybrid reconstruction approach that combines an unlimited high-order K -exact reconstruction [3] based on a fixed central stencil and a monotonicity preserving low-order limited linear reconstruction for cells with under-resolved or discontinuous solution content. Switching in the hybrid procedure is determined by a smoothness indicator. This hybrid approach presents two main advantages. First, it avoids the complexity associated with other ENO and WENO schemes that require reconstruction on multiple stencils to find the “smoothest” stencil among several candidates [33,40,43,57]. Second, mesh adaptation may be directed based on the ability of the scheme to differentiate between resolved and under-resolved or non-smooth solution content. The high-order CENO scheme has been successfully applied to a broad range of flows on multi-block structured meshes including inviscid flows [37], viscous flows [38], large-eddy simulation (LES) for turbulent premixed flames [59] and magnetohydrodynamics (MHD) problems [26,37], as well as on cubed-sphere meshes [37]. Furthermore, the relative advantages of the CENO scheme as compared to ENO and WENO methods make it readily applicable to general unstructured meshes as considered by Charest and Groth for laminar viscous flows [10] and turbulent reactive flows [11].

Adaptive mesh refinement (AMR) approaches [4,6] are very attractive since they allow automatic refinement of the mesh which can be very effective in the treatment of problems having disparate spatial scales. There are currently a number of existing AMR libraries available for use including patch-based [1,16], cell-based [8], and block-based [41,44,60] AMR methods. Block-based AMR approaches, such as the schemes developed by Groth and co-workers [21,28–30,32,55], are a very attractive option since the overall data structure for grid connectivity is relatively light, due to the feature that grid connectivity is determined at the block level, and the approach is well suited to parallel implementation via domain decomposition producing highly scalable and efficient methods. These block-based methods have proven to be very effective in the solution of both laminar [12] and turbulent [28–30] flames as well as turbulent multi-phase rocket core flows [55], MHD simulations [31,32,36], and micron-scale flows [47]. Nevertheless, virtually all of these implementations are restricted to refinement of the mesh in an isotropic fashion. Isotropic AMR methods can be prone to excessively large grid sizes when attempting to capture solutions with anisotropic features such as propagating fronts and/or waves, shocks, shear surfaces, and current sheets.

A far more efficient strategy for dealing with shocks and/or thin solution layers is to adopt an anisotropic procedure in which the mesh refinement is carried out in a directionally dependent manner. Multi-block anisotropic AMR schemes have been proposed in previous studies by Zhang and Groth [68] and Williamschen and Groth [66] in applications to inviscid flow prediction in two-dimensional (2D) and three-dimensional (3D) domains, respectively. More recently, a revised anisotropic block-based refinement approach has been developed by Freret and Groth [25] that adopts a non-uniform or heterogeneous block structure in the representa-

tion of each grid block. In the non-uniform block approach, the extended layers of “ghost” or “halo” cells which allow the update of the solution within each block in an independent manner (helpful in serial optimization and parallel implementation of the algorithm) are direct duplicates of those of the neighboring blocks, even if the block itself and its neighbors are not at the same level of refinement. While introducing so-called “hanging nodes” in the grid, this non-uniform representation of the neighboring blocks is naturally suited to application with high-order numerical schemes [25], as it does not require the use of high-order restriction and prolongation operators to evaluate the ghost-cell values representing neighboring blocks as can be found in the mapped multi-block grids approach [17,46]. Furthermore, it does not use any additional procedures to ensure conservation and stability properties of the scheme [5,13,52,53,56]. Instead, the spatial discretization is merely modified to take into account the grid resolution changes that can occur between interior and ghost cells [25].

The present study therefore considers the extension of the anisotropic AMR method of Freret and Groth [25], originally developed for a second-order finite-volume scheme and application to gas dynamics simulations, to use with the high-order CENO finite-volume method of Ivan et al. [37,38,58] for MHD computations. The application of the high-order CENO scheme with isotropic AMR was examined previously by Susanto et al. [58] for the solution of the ideal MHD equations on 2D domains. It is re-considered here for 3D flows with anisotropic AMR using the non-uniform or heterogeneous block structure proposed by Freret and Groth [25], in which the ghost cells of a block are stored directly at the resolution of the neighboring blocks, is exploited to enable a fourth-order CENO spatial discretization procedure on complex multi-block AMR grid, including cubed-sphere grid topologies, with unstructured root block topology. A binary tree data structure is used to track the connectivity and refinement history of the grid blocks and is sufficiently flexible to deal with some blocks that feature degenerate edges and corners, where the number of neighboring blocks and available ghost cells are smaller than would be expected for blocks in a regular, logically Cartesian, grid structure. A major technical difficulty in extending the anisotropic AMR method for use with high-order finite-volume spatial discretization schemes is to define valid reconstruction stencils taking into account the anisotropic resolution changes and degenerate block edges and corners. A systematic and general approach to build these high-order stencils is proposed. Physics-based and smoothness-based criteria are both used as refinement criteria to direct the anisotropic refinement of the mesh. The binary tree data structure is also readily modified to revert to isotropic refinement through the use of additional constraints on admissible refinement and coarsening of the grid blocks. For the solution of the ideal MHD equations considered herein, the solenoidal constraint on the magnetic field is controlled using the generalized Lagrange multiplier (GLM) technique proposed by Dedner et al. [22]. This avoids the computational costs associated with elliptic cleaning procedures [7, 39] and/or possibly more complicated treatments associated with constrained transport (CT) algorithms [2,19,24,49,50]. The latter usually necessitate staggered fields [2,19,24], although several non-staggered constrained transport techniques have been proposed recently [14,34, 35,49,50]. For example, Helzel et al. developed an original unstaggered CT method for the 3D MHD equations on Cartesian grids based on the wave-propagation method [35] and extended this work to mapped grids [34] by using a method of lines. More recently, Christlieb et al. [14] extended the latter constrained transport method to even high order discretization on non-staggered meshes. As will be shown, the proposed combination of the fourth-order CENO finite-volume and anisotropic AMR schemes yields a new powerful parallel and efficient simulation tool for the prediction of MHD flows on complex multi-block grids.

The paper is organized as follows. The ideal MHD equations and the GLM formulation are described in Sect. 2. In Sect. 3 a brief outline of the high-order CENO scheme is provided.

The proposed anisotropic AMR block-based method is discussed in Sect. 4 with the necessary extension for use with the high-order spatial discretization scheme on grids with degeneracies. Finally, fully three-dimensional numerical results are presented in Sect. 5 to demonstrate the accuracy and robustness of the adaptive fourth-order CENO method for Euler and MHD flows on Cartesian and cubed-sphere grids. Results for steady-state MHD plasma smooth flows are presented to assess the accuracy of the scheme on cubed-sphere grids. Furthermore, the 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 for time-dependent simulations are demonstrated.

2 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. [22]. This approach [22] couples the divergence constraint, $\nabla \cdot \mathbf{B} = 0$, with the induction equation through the introduction of a new potential variable, ψ . Thus, the modified system of conservation laws for which the solution is sought may be expressed in non-dimensional weak conservation form as

$$\frac{\partial \mathbf{U}}{\partial t} + \nabla \cdot \mathbf{F} = \mathbf{S} + \mathbf{Q}, \tag{1}$$

where \mathbf{U} is the conserved variable solution vector, \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

$$\mathbf{U} = [\rho, \rho \mathbf{V}, \mathbf{B}, \rho e, \psi]^T, \tag{2}$$

where ρ is the plasma density, $\mathbf{V} = (V_x, V_y, V_z)$ is the velocity field, $\mathbf{B} = (B_x, B_y, B_z)$ 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, $\nabla \cdot \mathbf{B} = 0$, for the magnetic field. Note that the standard orthogonal x, y, z Cartesian coordinate system is used in the conservation law discretization. The flux dyad, \mathbf{F} , is given by

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

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 $p_B = (\mathbf{B} \cdot \mathbf{B})/2$ is known as the magnetic pressure. The numerical source term, \mathbf{S} , is due to the GLM-MHD formulation and given by

$$\mathbf{S} = \left[0, \mathbf{0}, \mathbf{0}, 0, -\frac{c_h^2}{c_p^2} \psi \right]^T, \tag{4}$$

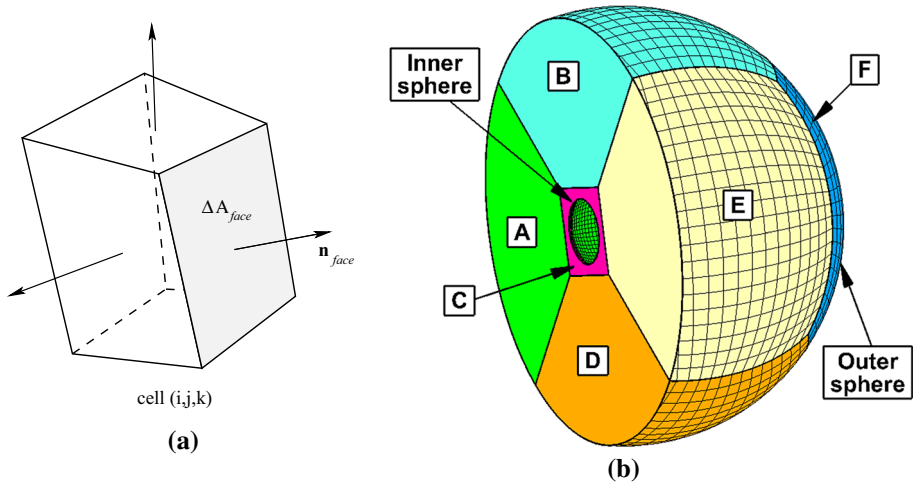


Fig. 1 **a** Hexahedral computational cell and **b** example of a cubed-sphere mesh composed of six blocks from **A** to **F** with each block containing of three-dimensional hexahedral cells. **a** Hexahedral cell showing face normals. **b** Cross-section of the cubed-sphere grid composed of six blocks, from **A** to **F**. Grid block boundaries are shown with bold lines. For clarity, mesh is shown only on the inner and outer spheres

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 $\nabla \cdot \mathbf{B}$ cleaning mechanism. 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 modeling of space-physics problems, such as sources associated with gravitational fields.

3 High-Order CENO Finite-Volume Scheme

3.1 Semi-discrete Finite-Volume Formulation

The semi-discrete form of the finite-volume formulation applied to Eq. (1) for a hexahedral computational cell (ijk) of a three-dimensional grid as shown in Fig. 1a can be expressed as

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

where N_g is the number of Gauss quadrature points and \mathbf{n} is the local normal of the face f at each of the N_g Gauss quadrature points. The integer N_f is the number of faces of the cell. The six faces (west, east, south, north, bottom and top) of the cell are indicated by $\{W, E, S, N, B, T\}$. The hexahedral cells are contained within grid blocks of a multi-block body-fitted mesh as depicted in Fig. 1b. To accommodate complex meshes, such as the cubed-sphere grid of Fig. 1b, the connectivity between blocks in the multi-block mesh is defined using an unstructured data structure [29].

Standard tensor-product Gaussian quadrature is used in the evaluation of the integrals of numerical flux on each cell face as described by Ivan et al. [37] and the number of quadrature points, N_g , is selected to be the minimum required to preserve the desired order of solution accuracy. For instance, four Gauss quadrature points per face are required for the fourth-order accurate spatial discretization considered here. Upwind values of the numerical fluxes, $\mathbf{F} \cdot \mathbf{n}$, at each Gauss quadrature point on each face of a cell (ijk) are determined from the solution of a Riemann problem. Given the left and right interface solution values, \mathbf{U}_l and \mathbf{U}_r , the upwind numerical flux is evaluated by solving a Riemann problem in a direction normal to the face. The values of \mathbf{U}_l and \mathbf{U}_r are determined via a high-order CENO polynomial reconstruction, as detailed in the next section. In the present computational studies, the HLLC and local Lax-Friedrichs numerical flux functions are used in obtaining solutions for the Euler equations governing non-magnetized flows and the MHD equations for magnetized flows, respectively [23,42]. 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.

3.2 High-Order CENO Reconstruction

As indicated previously, the hybrid CENO scheme of Ivan and Groth [37,38] is used to perform the piecewise cell-centered high-order reconstruction of the solution within each hexahedral cell of the computational grid. The hybrid CENO procedure uses a multidimensional unlimited K -exact reconstruction scheme as originally developed by Barth [3] in regions of smooth solution content and reverts to a low-order limited piecewise linear reconstruction in regions deemed as non-smooth or under-resolved by a solution smoothness indicator, thus ensuring monotone solutions near discontinuities and shocks. The K -order Taylor polynomial for the high-order K -exact reconstruction of a scalar solution quantity, u_{ijk} , within a cell (ijk) with cell-centroid $(x_{ijk}, y_{ijk}, z_{ijk})$ can be expressed as


$$u_{ijk}^K(x, y, z) = \sum_{\substack{p_1=0 \\ p_1+p_2+p_3 \leq K}}^K \sum_{p_2=0}^K \sum_{p_3=0}^K (x - x_{ijk})^{p_1} (y - y_{ijk})^{p_2} (z - z_{ijk})^{p_3} D_{p_1 p_2 p_3}, \quad (6)$$

where the coefficients, $D_{p_1 p_2 p_3}$, are the unknown derivatives of the Taylor polynomial expansion. A total of 20 unknown derivatives must be evaluated for piecewise cubic ($K = 3$) reconstruction, which is required to achieve a spatially fourth-order accurate scheme. The coefficients are obtained via the solution of a constrained least-squares problem as described by Ivan et al. [37]. To obtain an exactly determined (as a very minimum) or more preferably an overdetermined set of linear equations for the unknown derivatives, the number of neighboring cells considered in the reconstruction stencil must be at least equal to the number of unknown coefficients. In this study, a stencil that includes current cell of interest, the first ring of nearest 26 neighboring cells, plus 6 extra next-to-nearest cells in each coordinate direction is used providing a total of 33 cells and conditions for the least-squares reconstruction when a uniform logically Cartesian mesh is used. That is the stencil, S_t , is given by

$$S_t = \{(d_x, d_y, d_z) | d_x, d_y, d_z \in \{-1, 0, 1\}\} \cup \{(0, 0, \pm 2), (0, \pm 2, 0), (\pm 2, 0, 0)\}, \quad (7)$$

where the triplet (d_x, d_y, d_z) represents a search direction for building the stencil. Depiction of the reconstruction stencil in the rather general case of a Cartesian uniform grid is provided

Fig. 2 Typical 33-cell reconstruction stencil in a Cartesian uniform mesh for a cell (ijk) 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 used for the ghost cells. The stencil is composed of the central cell, 26 first-ring cells, plus 6 extra second-ring cells



in Fig. 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 the reconstruction of a solution variable in cell (ijk) is non-smooth, is computed as follows [36]:

$$\left. \begin{aligned} \mathcal{S} &= \frac{\alpha}{\max(1 - \alpha, \epsilon)} \frac{\mathcal{N}_{\text{SOS}} - \mathcal{N}_D}{\mathcal{N}_D - 1}, \text{ with} \\ \alpha &= 1 - \frac{\sum_{\gamma} \sum_{\delta} \sum_{\xi} (u_{\gamma\delta\xi}^K(\mathbf{X}_{\gamma\delta\xi}) - u_{ijk}^K(\mathbf{X}_{\gamma\delta\xi}))^2}{\sum_{\gamma} \sum_{\delta} \sum_{\xi} (u_{\gamma\delta\xi}^K(\mathbf{X}_{\gamma\delta\xi}) - \bar{u}_{ijk})^2}, \end{aligned} \right\} \quad (8)$$

where \mathcal{N}_{SOS} is an integer value representing the size of stencil (i.e., number of cells in stencil, 33 here), \mathcal{N}_D is an integer value representing the degrees of freedom in the K -exact polynomial reconstruction (i.e., the number of unknowns, 20 here), and the ranges of the indices $(\gamma\delta\xi)$ are taken to include either the whole or a subset of the supporting reconstruction stencil for the cell (ijk) . 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 this latter case, a linear reconstruction is performed ($K = 1$ in Eq. (6)) and Venkatakrishnan limiter [64] is used. For all of the simulations presented in Sect. 5, the smoothness indicator cut-off value is $\mathcal{S}_C = 1500$.

3.3 Choice of Variables for Cell-Centered Reconstruction

Like other Godunov-type finite volume methods, the proposed CENO scheme has the flexibility of applying the cell-centered solution reconstruction to either the set of conserved variables, $\mathbf{U} = [\rho, \rho\mathbf{V}, \mathbf{B}, \rho e, \psi]^T$, or primitive variable set, $\mathbf{W} = [\rho, \mathbf{V}, \mathbf{B}, p, \psi]^T$. 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, the conservation properties of the scheme are assured regardless of the choice of reconstruction variables, due to the nature of the finite-volume scheme and use of Riemann-solver-based flux functions. Furthermore, it is known that cell-centered piecewise reconstructions in terms of this set provides the lower-quality shock-capturing fidelity when compared to the reconstruction performed either in terms of the primitive or in characteristic variables [67]. As a result, the use of the primitive variable set has 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 key flow variables such as density and pressure. Unfortu-

nately, as recently demonstrated by Charest [9], 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. While Charest [9] and Ivan and Groth [38] have shown that this error is not introduced for steady problems and the expected formal accuracy is always recovered in this case, the loss of accuracy for unsteady flows is undesirable and to be avoided. Note that McCorquodale and Colella [45] have proposed a nonlinear transformation which provides a fourth-order accurate conversion between the aforementioned solution variables.

Fortunately, the hybrid nature of the CENO reconstruction scheme offers a rather simple solution to the loss of accuracy for unsteady problems: the unlimited high-order reconstruction can be performed in terms of the conserved variables and, where the scheme reverts to a limited linear reconstruction for the treatment of non-smooth solution content, the primitive variables set can be adopted. In this way, tight control is maintained on the monotonicity and positivity of key flow variables while avoiding the introduction of the temporal errors and incurring virtually no additional computational overhead. The conservation properties of the scheme are also unaffected. This new set of combined reconstruction variables is used in the CENO scheme implementation in this work and, in Sect. 5.1, it is evaluated and compared to the original CENO scheme in which the primitive variables were reconstructed in both the high-order and low-order representations of the solution. It is shown via a convergence study for an unsteady flow problem that the expected high-order accuracy is recovered in the asymptotic limit only for the set of combined variables, while still affording good control of solution monotonicity.

3.4 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 accuracy such as flux integration and calculation of geometric properties (e.g., area, volume, centroid, moments, etc.). The solution adopted by Ivan et al. [37] in which the faces of general hexahedral cells are defined in terms of a trilinear mapping 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.

4 Parallel Anisotropic Block-Based AMR

A flexible block-based hierarchical binary tree data structure is used in conjunction with the high-order spatial discretization procedure described in Sect. 2 to enable automatic solution-directed mesh adaptation on body-fitted multi-block mesh. The general AMR framework of Freret and Groth [25], originally implemented for a second-order accurate discretization based on extensions to the previous work by Williamschen and Groth [66], is adapted herein for use with high-order CENO finite-volume scheme using a non-uniform treatment for the grid block neighbors. In this way, anisotropic mesh refinement and efficient and highly scalable parallel implementation are achieved via domain partitioning.

4.1 Hierarchical Binary Tree Data Structure

In the proposed anisotropic AMR scheme, mesh adaptation is accomplished by refining and coarsening grid blocks. Figure 3a shows the resulting multi-block domain after several refinements of an initial mesh consisting of a single block. Each level of refinement introduces new blocks called “children” from a given “parent” block. Each refined parent block will introduce either two, four, or eight children, depending of the mesh refinement criteria (please refer to the section to follow). Each child block contains the same number of cells as the parent and this effectively doubles the mesh resolution in the desired directions. The children can of course be further refined and the refinement process can be reversed in regions that are deemed over-resolved and two, four or eight children, depending on the local solution, are coarsened or merged into a single parent block. A flexible hierarchical tree data structure is used for tracking the grid 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. The binary tree associated with the multi-block domain in Fig. 3a is shown in Fig. 3b. The initial single block corresponds to the single root node. The “leaves” are the blocks of the multi-block structure. The solid (respectively dashed and dotted) arrows represent refinement in the x -, y - and z -directions, respectively, for this Cartesian mesh. Level and sector triplets are used for connectivity purposes. Local refinement and coarsening of the mesh can be directed according to adjoint error estimation strategies [18,51,62,63] or so-called physics-based refinement criteria [25,66]. While the former are expected to be more accurate, the latter have been used extensively because of their relative simplicity.

4.2 Evaluation of Mesh Refinement Criteria

Refinement criteria are required for the anisotropic AMR scheme, with separate indicators needed for refinement in each direction. These measures are scaled by the area of the cell and normalized by the magnitude of the solution variables involved and then used to direct the mesh refinement. Partial derivatives of the primitive solution variables with respect to each coordinate direction are an intuitive choice and are adopted for the type of physics-based refinement criteria considered here. Expressions for the directional refinement indicators r_γ , r_η , and r_ζ are as follows:

$$r_\gamma = \frac{1}{u} (\nabla u \cdot \Delta \tilde{\gamma}), \quad r_\eta = \frac{1}{u} (\nabla u \cdot \Delta \tilde{\eta}), \quad r_\zeta = \frac{1}{u} (\nabla u \cdot \Delta \tilde{\zeta}), \quad (9)$$

where $\Delta \tilde{\gamma}$, $\Delta \tilde{\eta}$ and $\Delta \tilde{\zeta}$ are the vector differences between the midpoints of the faces in the γ , η , and ζ directions, respectively. In equation (9), u represents any quantity of interest. In the present study, the flow density is largely used as the relevant variable for directing refinement, i.e., $u = \rho$.

As the CENO scheme has its own measure of solution smoothness through the smoothness indicator, \mathcal{S} , and this smoothness measure was proposed previously by Ivan and Groth [38] for directing mesh refinement within an isotropic AMR approach. Alternate smoothness-based anisotropic refinement criteria are proposed here which can be defined in terms of the following directional indicators:

$$r_\gamma = \exp(-\mathcal{S}_\gamma), \quad r_\eta = \exp(-\mathcal{S}_\eta), \quad r_\zeta = \exp(-\mathcal{S}_\zeta), \quad (10)$$

with, for direction γ ,

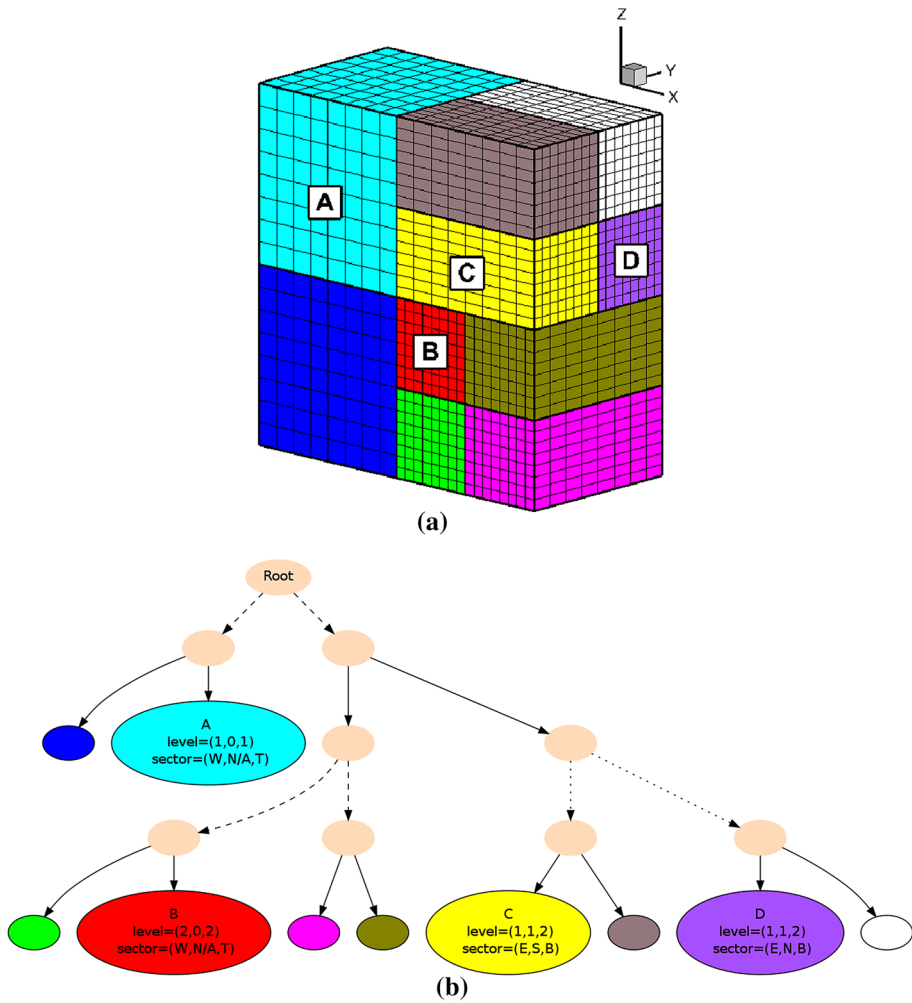


Fig. 3 Example of a multi-block domain after several anisotropic refinements initially composed of an initial single block and its corresponding binary tree. **a** Example of a multi-block domain obtained after several refinements of an initial single block. **b** Depiction of the binary tree associated with the above multi-block structure. The peach colored ovals represent nodes of the binary tree. The root node corresponds to the initial block. The leaves are the blocks of the multi-block structure. The solid (respectively dashed and dotted) arrows represent refinement in x (respectively y and z -direction. Level and sector triplets are used for connectivity purposes. The initial level is $(0,0,0)$ and its components increase with the number of refinements associated with each direction. The sectors indicate the position relative to neighbors in each direction, and are given by $W, E, S, N, B, T, N/A$ for west, east, south, north, bottom, top and undefined respectively. A sector component can be undefined if it has not been yet refined in the corresponding direction

$$\left. \begin{aligned}
 \alpha_\gamma &= \frac{\alpha_\gamma}{\max(1 - \alpha_\gamma, \epsilon)}, \text{ and} \\
 \alpha_\gamma &= 1 - \frac{\sum_\delta \sum_\zeta (u_{\gamma\delta\zeta}^K(\mathbf{X}_{\gamma\delta\zeta}) - u_{ijk}^K(\mathbf{X}_{\gamma\delta\zeta}))^2}{\sum_\delta \sum_\zeta (u_{\gamma\delta\zeta}^K(\mathbf{X}_{\gamma\delta\zeta}) - \bar{u}_{ijk})^2}
 \end{aligned} \right\} \quad (11)$$

While the solution smoothness, \mathcal{S} , is evaluated within the complete 33-cell stencil, the value \mathcal{S}_γ is restricted through α_γ to the neighboring cells aligned with the cell (ijk) in the γ direction, thus providing a smoothness indicator for that particular direction only. Corresponding expressions for \mathcal{S}_η and \mathcal{S}_ζ readily follow from Eq.(8).

It should be noted that, for both sets of refinement indicators defined by Eqs. (9)–(10) above, the indicators are first evaluated for every cell in each grid block. The maximum value of the indicator for each direction over all cells within the block is then taken to be the value for the directional indicator for refinement of the entire block in each of the logical coordinate directions. While the refinement of the entire block is recognized to be non-optimal in terms of refinement efficiency (typically, larger numbers of new computational cells are created during the refinement process as compared to cell-based AMR approaches), such block-based AMR methods provide effective treatments for disparate scales while readily lending themselves to efficient and scalable parallel implementations via domain decomposition on distributed memory parallel computers [15,21,28,29,32,55]. In the present approach, the block-based AMR strategy and binary tree data structure, combined with domain decomposition, results in a highly efficient and scalable parallel implementation.

4.3 Flagging of Grid Blocks for Refinement

Having arrived at appropriate criteria for directing the mesh refinement, the anisotropic block-based AMR scheme then proceeds according to the following six steps:

1. **Assignment of Refinement Based on Refinement Criteria:** blocks are flagged to refine, coarsen or maintain their refinement level for each computational direction based on the directional-dependent refinement criteria calculated for each block.
2. **Conflict Checking:** refinement flags are modified to ensure that there is no violation of predefined rules for block refinement or coarsening.
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 directions.
5. **Neighbor Information Update:** neighbor information or solution block connectivity is re-computed and stored for each block.
6. **Message Passing:** solution and geometry information is shared between adjacent blocks via a message passing procedure.

As outlined by Williamschen and Groth [66], restrictions or predefined rules on the assignment of the refinements flags are used to eliminate and prevent undesirable mesh topologies including: (i) 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; (ii) 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; and (iii) an invalid binary tree connectivity. However, several mesh block topologies are allowed here in order to arrive at a more optimal anisotropic mesh. Firstly, block faces associated with a given direction are not restricted to be at the same level in both tangential directions; an example of this configuration is given in Fig. 4a. Additionally, a face is not required to fully span its neighbors; this configuration is represented in Fig. 4b. Being free of these constraints requires extra care during the coarsening procedure as some connectivity rearrangements of the binary tree may be required [66], and calculation of the numerical fluxes may occur through faces that are not defined a priori, as shown in Fig. 4b. However, in terms of computational time, allowing these configurations




Fig. 4 Examples of refined blocks that do not violate the 3D rules

incurs very little overhead to the overall computational cost of the problem while greatly increasing the effectiveness of the anisotropic AMR algorithm.

It is worth mentioning that an isotropic AMR procedure can be readily obtained using the binary tree data structure. While more classical octree data structures [29] are usually employed with isotropic AMR, with minimal effort, the anisotropic AMR scheme as described above can degenerate to an isotropic approach by making only two modifications. Firstly, the refinement criteria must be appropriately modified to no longer be directionally based. For example, in the numerical results presented in Sect. 5.4, an isotropic refinement indicator based on the curl of the velocity field is used. The second modification to arrive at an isotropic AMR scheme occurs in the conflict checking procedure. Steps must be taken to enforce that any block flagged for coarsening must have its seven siblings also flagged for coarsening. If at least one of the siblings is flagged for refinement or maintaining its current level of refinement, all children are prevented from undergoing the grid coarsening. The conflict checking procedure is finally re-applied to remove any refinement conflicts that may have been introduced by the modified flags.

4.4 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 sub-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 are required. One advantage of using the non-uniform block representation (please refer to the subsection to follow) 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 of the overlapping cell domains. In particular, to distribute the average solution quantity among offspring at high-order accuracy the high-order polynomial reconstructions of all solution variables in the coarse cell are integrated over the subdomains of each new fine cell having a volume, V_{fine} ,

$$\bar{u}_{fine} = \frac{1}{V_{fine}} \iiint_{V_{fine}} u_{coarse}^k(\mathbf{X}) dV, \quad (12)$$

where the volume integral is computed exactly for the given reconstruction polynomial with an appropriate-order, volumetric, tensor-product, Gaussian quadrature integration technique.

4.5 Non-uniform Blocks

In order to apply the proposed high-order finite-volume scheme to all of the grid 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 [25]. Within the data structure for the non-uniform block representation, the regions adjacent to block faces, edges and corners are referred to as boundary elements. The topological arrangement of these boundary elements is described in terms of the triplet (be_i, be_j, be_k) , $be_\xi \in \{-1, 0, 1\}$ for $\xi = \{i, j, k\}$. There are 26 boundary elements corresponding to the regions opposing the 6 faces, 12 edges and 8 corners of a given block, in addition to the interior of the block itself. Using this denomination, the triplet $(-1, -1, -1)$ represents the west-south-bottom corner, $(0, 0, 1)$ identifies the top face and $(0, 0, 0)$ is the interior domain. In particular, a maximum of four, two, or one ghost-cell block(s) can span a face, an edge, or a corner, respectively.

Examples of non-uniform blocks obtained from representative multi-block structures are shown in Fig. 5. In Fig. 5a, block **B** has been extracted from the multi-block structure shown previously in Fig. 3b. It is composed of an interior domain (in red) containing the computational cells and extended in all directions by ghost cell blocks. For clarity, the ghost cell blocks are shown only in the x and z directions. In Fig. 5b, block **E** has been extracted from the cubed-sphere grid previously depicted in Fig. 1b. It is composed of the interior domain (shown in light yellow) and four faces; ghost blocks along the radial direction have been removed for clarity purposes. As discussed by Ivan et al. [37], the cubed-sphere grids have degenerate edges and corners characterized by missing ghost cells. These degeneracies are naturally taken into account merely by the absence of corner and edge ghost cell blocks.

The non-uniform block treatment adopted here presents many advantages as outlined by Freret and Groth [25]. Firstly, high-order restriction and prolongation operators are not required to evaluate the ghost cell solution values at resolution changes. Moreover, there is no need for an additional correction to enforce the flux conservation properties of the finite-volume scheme. This property is automatically satisfied by the non-uniform treatment. Finally, as considered in the present work, the non-uniform approach allows for an efficient way to treat the ghost cells. 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. Whereas in the original implementation of Ivan et al. [36,37] 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, the current implementation achieves the same accuracy with only two layers of ghost cells as follows. Instead of reconstructing the solution in the 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. In this way, the number of ghost cells is greatly minimized at the expense of increasing somewhat the size and number of send/receive messages containing information between the neighboring blocks. For example, in the simulation discussed in Sect. 5.4, the use of this reduced storage treatment for the ghost cells represents a memory savings of about 60% compared to the original rather non-optimal approach.

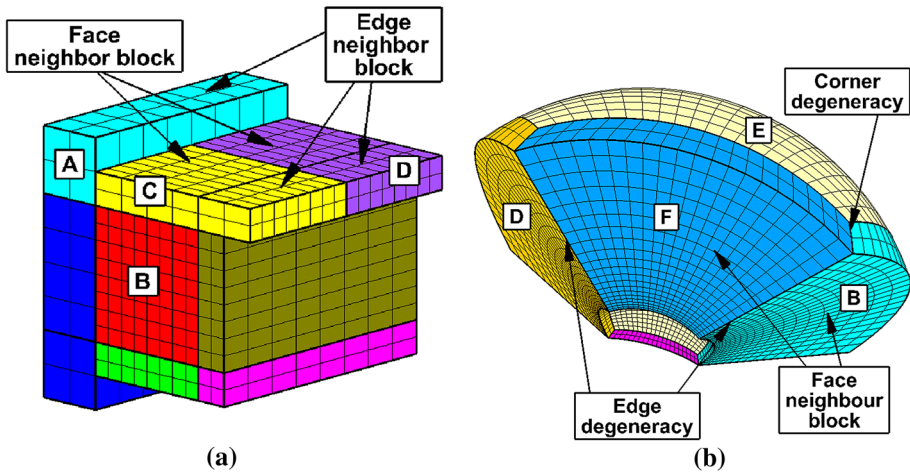


Fig. 5 Depiction of non-uniform structured mesh blocks in case of **a** a Cartesian mesh and **b** a cubed-sphere grid. **a** Depiction of non-uniform block **B** extracted from the multi-block domain presented in Fig. 3a. It is composed of the interior domain (in red) extended (only in the x and z direction for clarity) by a layer of 2 ghost cells provided by the neighboring blocks. Blocks **C** and **D** from Fig. 3a share the top face as well as the east-top edge of block **B**. **b** Depiction of non-uniform block **E** extracted from the multi-block domain presented in Fig. 1a. It is composed of the interior domain (in light yellow) extended (only in the azimuthal and zenith directions for clarity) by a layer of 2 ghost cells provided by the neighboring blocks. This non-uniform block presents edge and corner degeneracies where faces meet

The non-uniform block strategy is also very convenient when cubed-sphere meshes are considered. As the reconstruction stencils must be consistent with each other for cells located in the interior domain and for the same cell occurring as a ghost cell in a neighboring block, Ivan et al. [37] proposed a rotation mechanism for generating consistent stencils near degenerate block edges. With the new non-uniform block treatment of the ghost cells, this somewhat complicated though effective technique is no longer required.

4.6 Determination of Cell-Centered Reconstruction Stencils with AMR

One challenge with the present formulation is that the different refinement configurations that can occur as part of the anisotropic AMR procedure, including grid blocks with degenerate edges and corners, must be accounted for when determining the reconstruction stencils of the cells. A new neighbor search algorithm for constructing the stencils has been developed here, which exploits the data structure of the non-uniform block treatment. In the search algorithm, stencils are built for all cells lying within the interior of a block. This includes interior cells that are remote from block boundary elements as well as interior cells that are close to boundary elements. The ghost cells of a block are only used for storage space of solution information and stencil determination is not required for these cells. The neighbor search algorithm is applied in each of the 32 search directions associated with an interior cell as defined by Eq (7). For each of these search direction, the number of cells to be added to the stencil in that direction is determined: one cell neighbor in the case of a uniform or a coarser mesh, two or four cells for a neighbor block having a finer mesh resolution in one or two of directions tangential to the search direction, respectively, and no cell in the case of an edge or a corner degeneracy. The full procedure for computation of the stencil employs four steps and can be described as follows:

1. **Boundary element destination:** First, the boundary element indicator triplet, (be_i, be_j, be_k) , is evaluated corresponding to either the interior of the block itself or one of the 6 faces, 12 edges or 8 corners of the block. This first step is simple and based on the search directions, (d_i, d_j, d_k) and the cell index itself, $(i, j, k) \in [0, N_i] \times [0, N_j] \times [0, N_k]$, with N_ξ being the number of cells in the ξ direction. The value for be_ξ is defined for $\xi \in \{i, j, k\}$ such that

if $\xi + d_\xi < 0$ **then**

$$be_\xi = -1$$

$$\Delta_\xi = \xi + d_\xi$$

end if

if $\xi + d_\xi > N_\xi$ **then**

$$be_\xi = 1$$

$$\Delta_\xi = \xi + d_\xi - N_\xi$$

else

$$be_\xi = 0$$

end if

If a value of $(0, 0, 0)$ is found for the boundary element indicator, the boundary element destination is the interior domain of the block itself. In this case, the neighbor search algorithm results in a single cell lying in the interior block which can be identified as $C = \{(i + d_i, j + d_j, k + d_k)\}$ and the neighbor search for the given direction is complete. Steps 2 through 4 can be skipped. Figure 6 provides an example of the stencil computation for a cell lying near a boundary element. Cells colored yellow represent neighbouring cells that were found to lie within the interior domain of a block. In the case that the value of the boundary element triplet is not $(0, 0, 0)$ but the corresponding boundary element is degenerate, no neighbor is found, $C = \{\emptyset\}$, and again the search for the given direction is complete. Steps 2 to 4 are again skipped. In the example of Fig. 6, the boundary element is highlighted in light gray and corresponds to the east face.

2. **Block destination:** If the boundary element identified by the indicator, (be_i, be_j, be_k) , contains just one ghost block, step 2 is not required. If several blocks span the boundary element, the block destination is identified using its relative position to the interior block, as encoded in the sector vector that is stored in the binary tree data structure. Using the index, (i, j, k) , for the cell lying in the interior of the block, a simple test $\xi + d_\xi < N_\xi/2$ is used to discriminate between either the east or west (i), south or north (j) and bottom or top (k) of the boundary element. In Fig. 6, the block destination shown in green has a south and bottom sector.
3. **Cell index destination:** For the current search direction, a unique target cell is now selected for inclusion in the stencil. The index of this cell lying within the ghost block is provided using the following pseudo-code:

switch be_ξ **do**

case -1:

$$\xi^* = N_\xi(B_d) + \Delta_\xi + 1$$

case 1 :

$$\xi^* = \Delta_\xi - 1$$

case 0:

if $l_{\text{source}}(\xi) = l_{\text{dest}}(\xi)$ **then**

$$\xi^* = \xi + d_\xi$$

end if

if $l_{\text{source}}(\xi) < l_{\text{dest}}(\xi)$ **then** % neighbor has a finer mesh

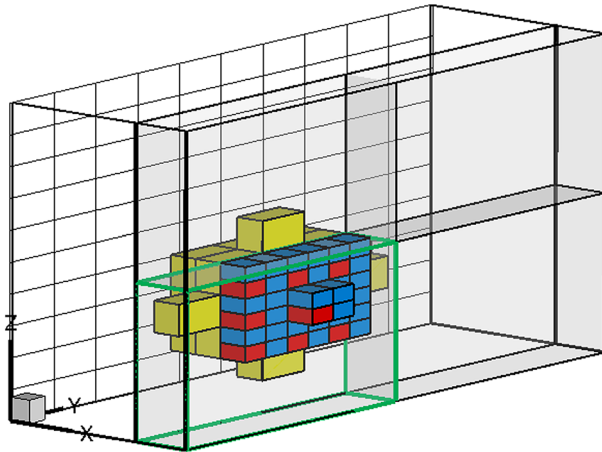


Fig. 6 Depiction of a non-uniform block to highlight the four steps of the neighbor search algorithm for a cell close to the block boundary, having ghost cells in its reconstruction stencil. For each search direction, one has to: 1—identify the boundary element destination (a face here) highlighted in light gray, 2—identify the block destination within the boundary element (one of the four blocks spanning the face, here the one shown in green), 3—identify a unique neighbor for each search direction cell (in red), 4—in case the neighboring mesh is finer, extend the stencil by filling the gaps in the tangential finer cells (in blue). The yellow cells represent interior cells for which the neighbor search algorithm is greatly simplified with the use of step 1 only

```

    ξ* = modulo ((ξ + dξ) × 2, Nξ(Bd))
end if
if lsource(ξ) > ldest(ξ) then % neighbor has a coarser mesh
    ξ* = ξ/2 + dξ
end if

```

where $l_{source}(\xi)$ and $l_{dest}(\xi)$ are the refinement levels of the source (interior domain) and destination blocks in the ξ direction, respectively. The value $N_{\xi}(B_d)$ corresponds to the number of cells of the block destination B_d in the ξ direction. An index ensemble is defined such that $C_{\xi} = \{\xi^*\}$ and the unique cell index destination is given by $C = (C_i, C_j, C_k)$. In Fig. 6, cells identified by step 3 are shown in red.

- Cell index extension:** In the cases that the neighboring mesh is finer in the tangential ξ -direction, the index ensemble C_{ξ} is enriched to ensure that there is no gap between the cells as follows:

```

if lsource(ξ) < ldest(ξ) & beξ = 0 then
    Cξ = {ξ*, ξ* + 1}
end if

```

Finally, the neighbor cell indices are $C = (C_i, C_j, C_k)$ for a given cell (i, j, k) and a search direction (d_i, d_j, d_k) . As depicted in Fig. 6, the blue colored cells are added to complete the stencil.

The application of the preceding neighbor search algorithm applied on a Cartesian mesh to a cell located close enough to the block boundaries such that its stencil includes cells in the ghost cell area is shown in Fig. 7. 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 Fig. 7. In the stencil representation, cells colored yellow represent the interior domain of a block and the blue colored cells represent

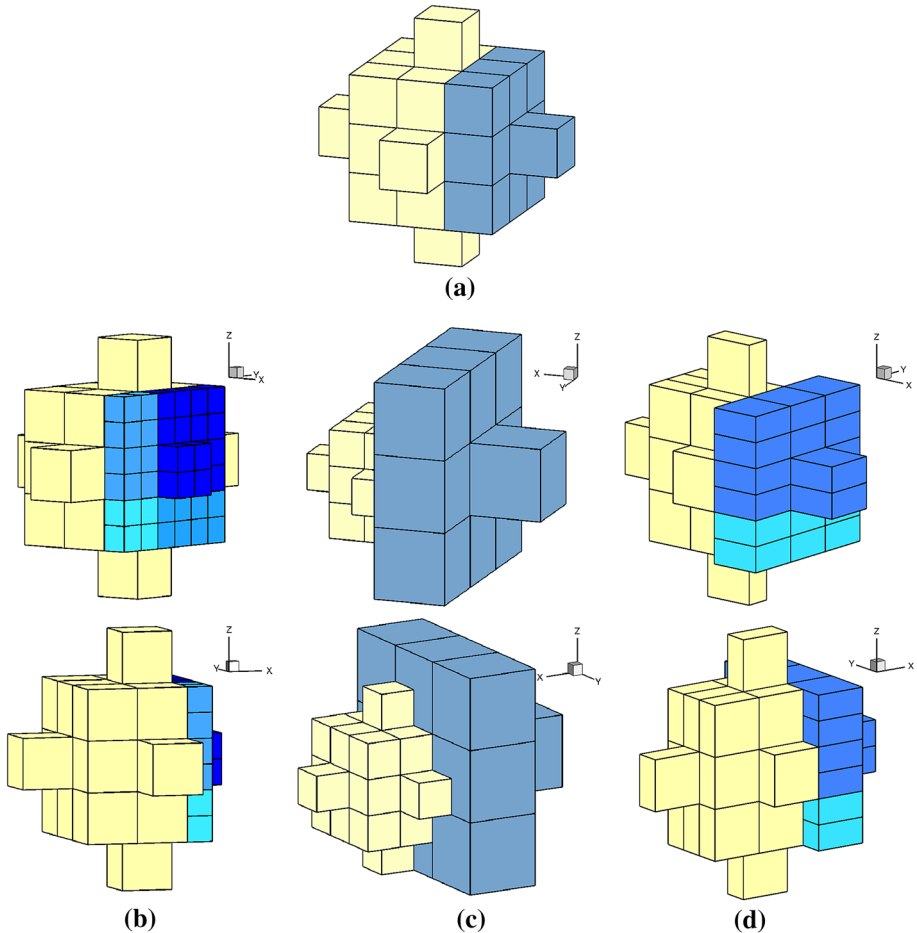


Fig. 7 Examples of cell stencil construction in a Cartesian mesh for a cell (ijk) located at the block boundary such that its stencil includes ghost cells from neighboring blocks. The interior domain is represented in yellow and blue is for the ghost cells. The different blue color shades are for different ghost blocks spanning an interior block face. **a** Neighbor has the same mesh refinement. Stencil is identical to that of an interior cell. **b** Four neighbor blocks with finer mesh are facing the block. Neighbor search directions produce more than one cell. **c** Neighbor has a coarser mesh. Number of cells identical to uniform configuration. **d** Two neighbor blocks having a finer mesh in a direction and a coarser mesh in the other direction. A total of 43 cells are included in the stencil. **a** 27 first-ring cells plus 6 extra second-ring cells: 33 cells. **b** Two views of a stencil built with four neighbors having a finer mesh, 54 first-ring cells plus 9 extra second-ring cells: 63 cells. **c** Two views of a stencil built with a neighbor having a coarser mesh, 27 first-ring cells plus 6 extra second-ring cells: 33 cells. **d** Two views of a stencil built with two neighbors having both a finer and a coarser mesh, 36 first-ring cells plus 7 extra second-ring cells: 43 cells

ghost cells for that block. All four different stencils are based on the baseline 33-cell uniform Cartesian stencil defined by Eq (7) which has been found optimal for the fourth-order CENO scheme [37]. Stencil construction using the neighbor search algorithm would also be valid for other more extensive stencils discussed by Ivan et al. [37]. Figure 7a is a stencil obtained for a cell ijk located at the block boundary facing a neighbor with the same mesh refinement. The shape of the stencil is exactly the same as an interior cell stencil and the stencil contains

33 cells. Figure 7b represents a stencil built with four neighbor ghost blocks having a finer mesh. Because the mesh is finer in both tangential directions, the number of cells included in the stencil equals 63. With a finer mesh in one direction only (2 neighbor ghost blocks spanning the face), the number of cells included in the stencil would be 43. Figure 7c 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 7d is specific to anisotropic AMR since the neighbors represented in blue have 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 Fig. 7a 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. Figure 8 shows three stencils obtained with the above neighbor search algorithm applied to a cell of a cubed-sphere grid having an edge degeneracy. Figure 8a (respectively Figs. 8b and 8c) depicts a stencil with finer (respectively coarser and both finer and coarser) neighboring ghost blocks. For this mesh topology, the minimum size of the various stencils is 29, again ensuring an overdetermined system for solution of the least squares problem associated with the evaluation of the high-order K -exact reconstruction polynomial.

4.7 Numerical Flux Evaluation with AMR

As discussed by Freret and Groth [25], the numerical flux evaluation also becomes more challenging when a non-uniform mesh is used, more particularly at block boundaries. The possibility of 16 different types of resolution changes at the block boundaries must be accounted for, which can be summarized in three groups as shown in Fig. 9. In Fig. 9, case (a) involves a cell which has an equal or a finer mesh in one or several directions compared to its neighboring 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 opposite resolution changes in two different directions, as also shown in Fig. 4b. 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 Fig. 9a. For case (b) the evaluation of the fluxes is performed through the faces of the ghost cells, taking the opposite of the normal vector, as depicted in Fig. 9b. The high-order flux evaluation is done by summing the flux through the finer neighboring cells to get the total flux of a coarser cell. Case (c) of Fig. 9c 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.

5 Numerical Results

A number of numerical results are now presented to demonstrate the accuracy and efficiency of the overall fourth-order CENO finite-volume scheme with adaptive mesh refinement. Results for uniform, isotropic, and anisotropic AMR meshes are included, as well as for cubed-sphere grids, so to provide a broad evaluation of the AMR strategy. In Sect. 5.1, numerical results for two different choices of reconstruction variables within the CENO

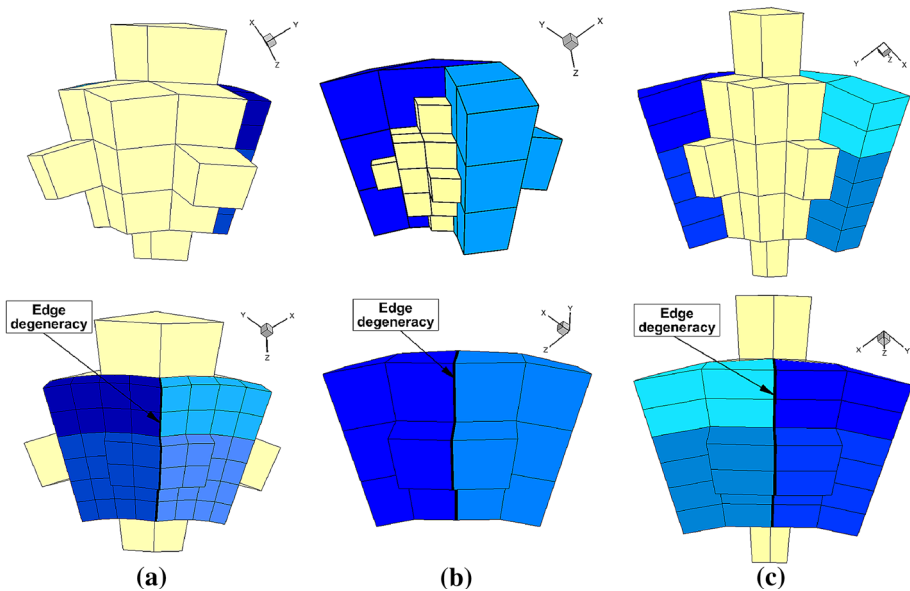


Fig. 8 Examples of cell stencil construction in a cubed-sphere mesh for a cell (i, j, k) located at the block boundary adjacent to a degenerate edge, such that its stencil includes ghost cells from neighboring face blocks. The interior domain is represented in yellow and blue is used for the ghost cells. The different blue color shades are for different ghost blocks spanning interior block faces. **a** Neighbors with finer mesh are facing the block. Neighbor search directions produce more than one cell. **b** Neighbors have a coarser mesh. Since the edge block cells are missing, the number of cells is less than 33. **c** Neighbors having a finer mesh in a direction and a coarser mesh in the other direction. A total of 43 cells are included in the stencil. **a** Two views of a stencil built with neighbors having a finer mesh, 59 first-ring cells plus 12 extra second-ring cells: 71 cells. **b** Two views of a stencil built with neighbors having a coarser mesh, 23 first-ring cells plus 6 extra second-ring cells: 29 cells. **c** Two views of a stencil built with neighbors having both a finer and a coarser mesh, 35 first-ring cells plus 8 extra second-ring cells: 43 cells

scheme are first considered for a transient problem involving the transport of an iso-density vortex. Numerical results for both non-magnetized and magnetized MHD results are next considered to evaluate grid convergence for smooth steady flows (Sect. 5.2) and robustness of the high-order method at shocks (Sect. 5.3). In Sect. 5.4, the proposed high-order AMR procedure is applied to the solution of an unsteady, 3D, iso-density, magnetized, vortex problem in which the vortex is transported across a Cartesian mesh corresponding to a periodic domain in order to show the ability of the combined AMR scheme to capture features or quantities of interest for smooth flows. Sects. 5.5 and 5.6 provide additional numerical results for a magnetically dominated bow shock flow with intermediate shocks as well as for a solar wind-magnetosphere interaction problem with a northward interplanetary magnetic field (IMF), respectively. For the cases considered, high-order reconstruction stencils based on the 33 cells as depicted in Fig. 7 were used to obtain the high-order results. In the numerical solution of all of the steady flow problems, a multi-stage optimal smoothing scheme [61] with local time stepping was used. For the time-accurate calculations, the standard, four-stage, fourth-order, Runge-Kutta time-marching scheme was employed.

It should be noted that the time integration procedure for time-accurate computations is applied globally here to all computational cells within the AMR grid using a single global time step. While the implementation of time-marching schemes using a hierarchical approach

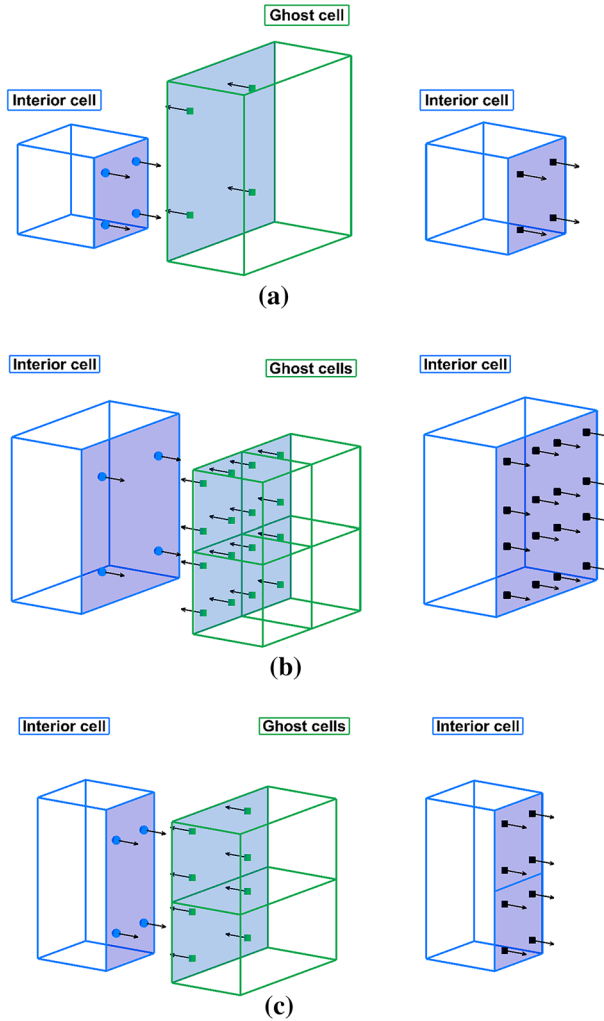


Fig. 9 Three scenarios to evaluate the numerical flux through a face of an interior cell depending on its mesh resolution and its neighbor cell mesh resolution. The blue (respectively green) colored points represent the Gauss quadrature points of the interior (respectively ghost) cell face. The black colored points are the Gauss quadrature points where the numerical fluxes are evaluated. **a** (left) Exploded view of an interior cell having a finer mesh than its neighbor ghost cell. The intersection of their face is the interior cell face. (right) The numerical fluxes are evaluated at the Gauss quadrature points of the interior cell face, represented by the black points. **b** (left) Exploded view of an interior coarse cell spanned with 4 ghost cells. (right) The intersection of the interior cell and the ghost cells faces are the ghost cells faces. The numerical fluxes are evaluated at each Gauss quadrature of the ghost cells face using the opposite of the original normal vectors. **c** (left) Exploded view of an interior cell extruded from a mesh that is coarser in one direction and finer in the other direction compared to its neighbor mesh. (right) The intersection of their faces are new cell faces, both of them having 4 Gauss quadrature points (colored in black). The numerical fluxes are computed at each of these quadrature points

in which the time step is dictated by the local grid resolution (e.g., as described by Berger and Olinger [4,6]) is possible, it has been the experience of the authors that the resulting potential computational efficiencies are not significant for AMR meshes in which approximately half

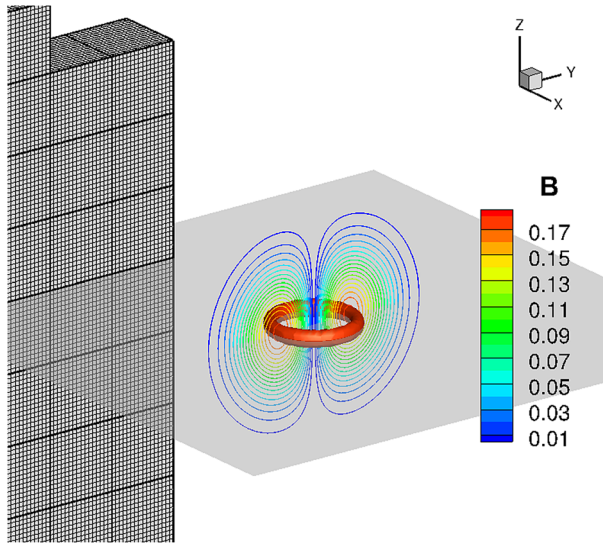


Fig. 10 Iso-density MHD vortex flow showing the iso-surface of the magnetic field ($B = 0.18$) contours of the initial flow solution on a Cartesian grid

of the computational cells reside at the finest level of resolution as commonly occurs in the MHD problems of interest here. For this reason, a global time-stepping approach is deemed sufficient and was used here. It should also be noted that the focus of the present study is on combining the high-order CENO finite-volume spatial discretization scheme with the proposed anisotropic AMR approach. More effective temporal discretization strategies will be considered in other follow-on research.

5.1 Effect of Choice of Reconstruction Variables

As a first evaluation of the fourth-order 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. [48]. This is a smooth time-dependent 3D test problem with an exact solution. The underlying stationary solution for the 3D iso-density MHD vortex flow is given by

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

where $\mu = \kappa = 1/(2\pi)$, and $q = 1$. This stationary solution is then 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 10 depicts this flow solution on the Cartesian grid. Initially the vortex center was positioned at $(0, 0, 0)$ and the simulations were carried out until time $t = 10$ at which time




Fig. 11 Effects of the choice of reconstruction variables on prediction accuracy showing comparisons of the predicted L_1 , L_2 and L_∞ error norms for both high- and low-order reconstruction in terms of primitive variables (solid lines with symbols) and mixed reconstruction with high-order reconstruction using conserved variables and low-order reconstruction using primitive variables (dashed lines). **a** Error norms in the velocity (solid lines with symbols) and momentum (dashed lines) versus grid size. **b** Error norms in the magnetic field as a primitive variables (solid lines with symbols) and conserved variables (dashed lines) versus grid size. **c** Error norms in the pressure (solid lines with symbols) and energy (dashed lines) versus grid size

the vortex has transported one period through the grid and is back to its original position. A grid convergence study was performed to assess the accuracy of the CENO scheme for this unsteady flow problem. Uniform mesh refinement with no AMR was considered for this first case. For the study, each grid block contained $8 \times 8 \times 8 = 512$ cells and the initial grid consisted of 64 blocks with a total of 32,768 cells. Uniform refinement was then applied to each subsequently finer mesh up to a finest mesh consisting of 262,144 blocks and 134,217,728 cells.

To investigate the accuracy issues associated with the choice of reconstruction variables as raised in Sect. 3.3, the series of simulations for the iso-density vortex problem on the sequence of uniformly refined meshes was first performed 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 this

first set of cases are presented with solid lines with symbols in Fig. 11a–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 $S_{prim} = \{u, B_x, p\}$. For the initial coarse 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 Fig. 11a–c. For the next refined meshes, all variables in each cell are flagged as smooth and a fourth-order unlimited polynomial reconstruction is performed for all variables 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 Charest [9], which indicate 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 [9,38].

A second series of simulations was then performed on the uniformly refined meshes in which the high-order reconstruction was carried out for the conserved variables and the low-order reconstruction was performed in terms of primitive variables, only for regions of the flow which are deemed as non-smooth. The corresponding numerical results for this second set of cases are presented with dashed lines in Fig. 11a–c. Figure 11 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 non-smooth for the two coarsest meshes, a low-order reconstruction is performed in these cells and global second-order accuracy is obtained. 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 [9]. 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 positivity of 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.


5.2 Manufactured Solution for Superfast MHD Outflow

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. [36]. The exact solution is given by

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

and corresponds to the solution having an additional source term \mathbf{Q} in the governing equations as detailed by Ivan et al. [36]. In this flow, the magnetic field is irrotational and aligned everywhere with the velocity. The computational domain used for this convergence study is

Fig. 12 Manufactured solution showing predicted L_1, L_2, L_∞ error norms for the first component of the magnetic field B_x versus grid size for uniform refinement (dashed dotted lines), isotropic AMR (solid lines), and anisotropic AMR (dashed lines) of the mesh. The L_1, L_2, L_∞ error norms are represented respectively by square, triangle and circle symbols



defined by inner and outer spheres of radius $R_i = 2$ and $R_o = 4$. To achieve fourth-order accuracy for this smooth problem, it is necessary to provide a high-order approximation to the average source term, \overline{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: both the inflow and outflow boundary conditions are specified using high-order accurate integration of the exact solution in ghost cells to determine the ghost cell averages.

Figure 12 depicts the computed error norms in the first component of the magnetic field B_x as a function of grid size. The theoretical fourth-order convergence is obtained with the CENO scheme on the cubed sphere for all error norms with successive uniformly refined grids (dashed dotted lines with symbols). Physics-based refinement criteria in terms of the plasma density were used in the computations with AMR. With isotropic AMR, after the first refinement, the convergence rate is around 7 for all error norms (solid lines). The plateau observed for the first isotropic refinement illustrates the limitations of the physics- or gradient-based refinement criteria. There is indeed no guarantee of a reduction in the solution error as the mesh is refined. For anisotropic AMR, the error decreases monotonically and a convergence rate reaches a value of 13 for all error norms (dashed lines), reflecting the important mesh savings that can be realized compared to both the uniform and isotropic refinement approaches. This is due to the refinement occurring only in the radial direction, thus allowing a large reduction in the number of cells to solve the problem for a given target accuracy. For example, to achieve an error of $L_1 = 5 \times 10^{-8}$, the anisotropic AMR approach requires only 6.25% of the total number of cells of the uniformly refined approach and exhibits an overall refinement efficiency of 98% as defined by the number of actual cells in the computational mesh to those in the equivalent uniformly refined mesh having the same finest mesh level [55]. Note that these accuracy results for the MHD outflow with AMR provides a partial validation of the proposed approach to obtain high-order accuracy using non-uniform stencils at different mesh refinement levels with a minimal ghost cells strategy that ensures consistent stencils near boundaries and corners.

5.3 Non-magnetized Bow Shock Flow

The application of the proposed AMR algorithm to a 3D non-magnetized bow-shock flow around a solid sphere body with reflecting boundary conditions is considered next. The

Fig. 13 Non-magnetized bow shock showing initial cubed-sphere mesh with five sectors highlighted in different colors and composed of 320 blocks. The grid blocks are shown and the cells within the grid blocks are shown in the dark-purple sector. Inflow boundary conditions are imposed on the outer sphere, outflow supersonic boundary conditions are applied on the back panels and reflection is used at the inner sphere

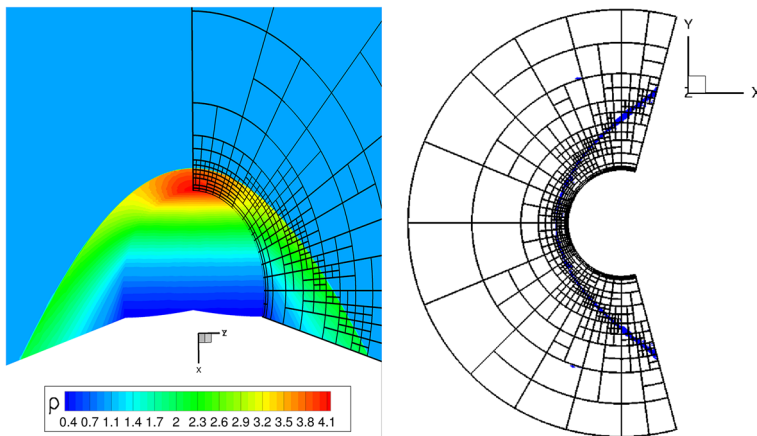
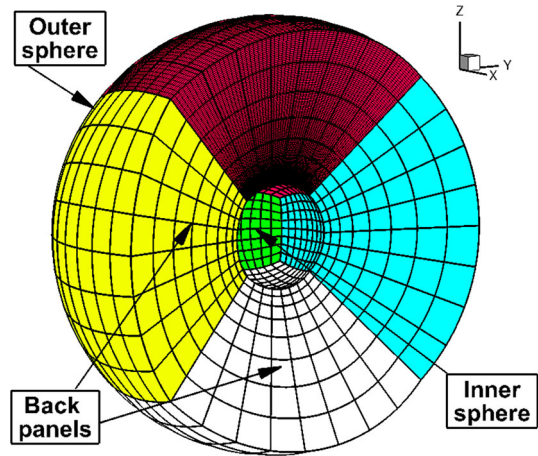


Fig. 14 Non-magnetized bow shock showing slice $z = 0$ of the anisotropic mesh and predicted density field obtained on the fourth level of mesh (left) and solution smoothness indicator (in blue) indicating cells deemed as non-smooth (right)

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 with a cubed-sphere grid to reduce the number of computational cells and focus on the upstream side of the sphere, as shown in Fig. 13. The supersonic inflow solution parameters used for this problem are $\rho = 1$, $v_x = 1.4943$, $v_y = v_z = 0$ and $p = 0.2$ such that the flow Mach number is $M = 2.825$. Inflow boundary condition is applied on the outer sphere R_o , a reflection boundary condition is imposed on the inner sphere R_i , and an outflow boundary condition is used for the back panels of the outer sphere.

The initial grid for this steady bow-shock problem consisted of a total of 320 blocks of $8 \times 8 \times 16$ cells, for a total of 32,768 hexahedral cells, with 8 blocks in the radial direction. The mesh was then subjected to 4 levels of adaptive mesh refinements. The anisotropic refinement criteria used for this simulation were based on the smoothness indicator as detailed in Eq. (10).

The predicted solution density and the smoothness indicator in the Cartesian plane $\{(x, y), z = 0\}$ obtained on the final mesh are both shown in Fig. 14. Similar 3D views

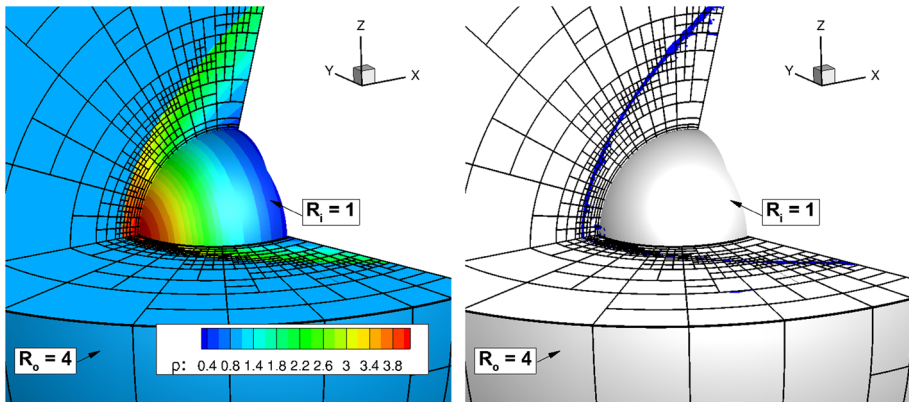


Fig. 15 Non-magnetized bow shock: 3D view of the anisotropic mesh and density field obtained on the fourth level of mesh (left) and solution smoothness indicator (in blue) indicating cells deemed as non smooth (right)

are given in Fig. 15. The final anisotropic mesh consists of 24,884 blocks and 80,116,416 computational cells. In this case the overall refinement efficiency is 84% compared to a uniform mesh. As shown in the two figures, refinement of the mesh occurs all along the shock profile and near the stagnation point. The number of cells deemed as non-smooth is around 10% for the initial mesh and up to 25% for the final mesh. In Figs. 14-(right) and 15-(right), the smoothness indicator is shown in blue on the final mesh. As can be seen, cells located near the shock are detected as non-smooth. The smoothness indicator thus provides a very good way to identify the shock. It is indeed worth mentioning that using the smoothness indicator as a refinement criteria allows the refinement all along the shock even in the outer regions of the bow shock that can be missed using physics-based refinement criteria in terms of the gradient of density [27].

5.4 3D Iso-Density MHD Vortex in a Periodic Cartesian Box

The ability of the proposed AMR procedure to track the advection of a 3D iso-density vortex in a periodic box is now evaluated. The problem setup is identical to that described in Sect. 5.1. The number of root blocks in the initial mesh was modified so as to contain $3 \times 3 \times 3$ blocks. At the start of the simulation, this initial mesh with 27 root blocks containing $8 \times 8 \times 8$ cells was refined eight times based on the initial conditions and a maximum refinement level of six was achieved. As the vortex problem is mostly isotropic, physics-based refinement criteria in terms of the curl of the velocity field was used to direct an isotropic refinement procedure in which the binary tree was forced here to behave as an octree, as described in Sect. 4.

Figure 16 shows 6 slices of the predicted magnitude of the magnetic field obtained for the unsteady MHD vortex problem on the isotropically refined meshes at (a) $t = 0$, (b) $t = 1.25$, (c) $t = 2.84$ from top-left to top-right and (d) $t = 6$, (e) $t = 7.59$, and (f) $t = 10$ from bottom-left to bottom-right. Following the initial mesh refinement, the unsteady AMR refinement and coarsening procedure was applied every 10 time-steps until a maximum time of $t = 10$ was reached. From the figure it is apparent that, throughout the simulation, the AMR tracks the vortex very well, clustering the blocks in the vortex area and with a minimal variation in the number of blocks. An overall reduction of 99% in the number of cells is

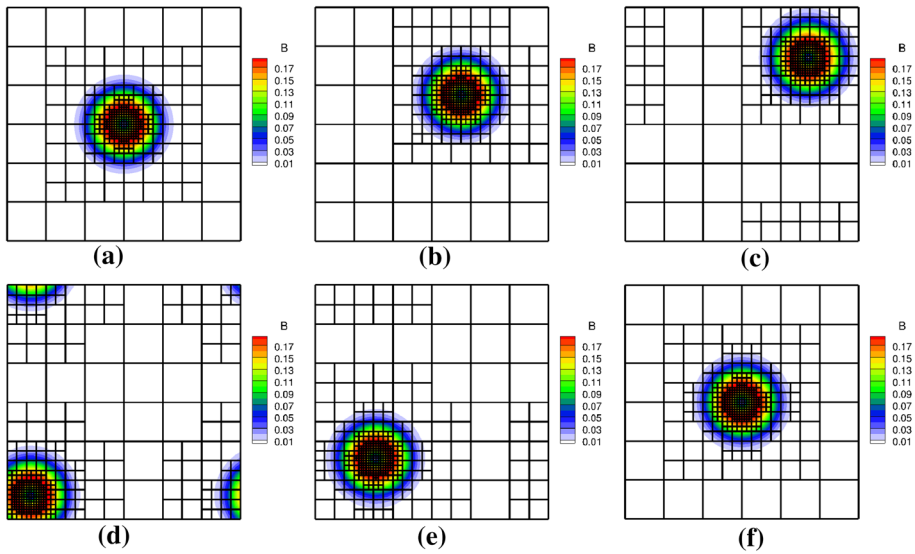


Fig. 16 Iso-density vortex in a periodic box showing isotropic AMR mesh and predicted magnitude of the magnetic field obtained at **a** $t = 0$, **b** $t = 1.25$, **c** $t = 2.84$, **d** $t = 6$, **e** $t = 7.59$ and **f** $t = 10$ in a XY plane of a periodic 3D box. **a** slice $z = 0$, 4024 blocks, $t = 0$, **b** slice $z = 2.5$, 5641 blocks, $t = 1.25$, **c** slice $z = -4.2$, 5207 blocks, $t = 2.84$, **d** slice $z = 2$, 5508 blocks, $t = 6$, **e** slice $z = -4.8$, 5606 blocks, $t = 7.59$, **f** slice $z = 0$, 5592 blocks, $t = 10$

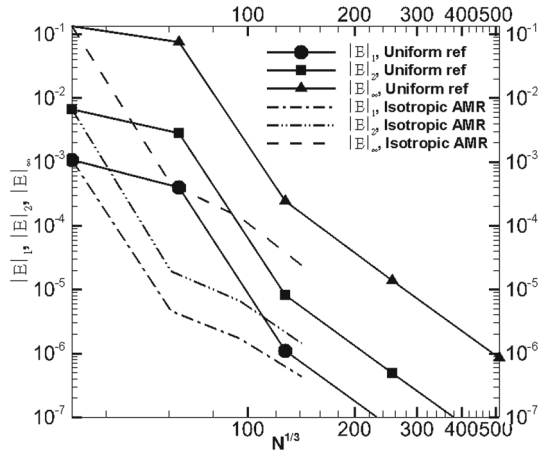
achieved compared to that which would be required for a uniformly refined mesh with a resolution equal to that on the finest level.

Figure 17 shows the error norms of the first component of the magnetic field B_x when the isotropic AMR was applied, along with comparisons to those results obtained for uniform refinement. Maximum levels for the refinement of 3, 4 and 5 were imposed, respectively. For each case, the simulation was carried out up to $t = 10$, with the isotropic AMR being applied every 20 time steps, and the error norms were then evaluated. The error norms for the uniform refined meshes are shown with solid lines and symbols as a reference and the results for the isotropic AMR are depicted with dashed lines. The mesh savings offered by the isotropic AMR approach compared to a standard uniform approach when comparing the predicted L_∞ norms of the error for the same accuracy is rather substantial: the mesh is reduced by a factor of 5–12 depending on the desired accuracy. It should be noted that, for this test case, the vortex represents a significant portion of the domain. As such it is therefore expected that the convergence or slopes of the error norms for both uniform and isotropic adaptive refinement should be similar.

5.5 Magnetically Dominated Bow Shock Flow

As a next case, numerical predictions of a stationary MHD bow shock flow around a perfectly conducting sphere are considered. For this case, a uniform superfast plasma flow is incident on the sphere and a stationary MHD bow shock is formed. The upstream magnetic field lines are aligned with the x -direction and form an angle $\theta_{vB} = 3.8^\circ$ with the upstream velocity field lines. The Alfvénic Mach number is such that $M_{Ax} = 1.5$ and the upstream plasma $\beta = 2p/B^2$. In such conditions, the flow is said to be magnetically dominated as thermal

Fig. 17 Iso-density vortex in a periodic box showing predicted error norms of B_x for both uniform refinement (solid lines with symbols) and isotropic AMR (dashed lines) of the mesh



and dynamical pressure effects are dominated by magnetic effects [20]. It has been shown that an upstream flow with such a strong magnetic field leads to a bow shock front followed by a secondary shock front [20].

The computational domain used for this problem was again based on the cubed-sphere grid with only five sectors, as described previously in Sect. 5.3. The initial mesh consisted of 640 blocks of $8 \times 8 \times 16$ cells with a total of 65,536 hexahedral cells. Converged steady-state solutions were obtained on sequence of successively refined meshes. Four successively refined anisotropic AMR procedure were applied in which physics-based refinement criteria based on the gradient of the velocity field amplitude were used. On the final mesh, 19,083 blocks were used for a total of 19,540,992 cells and the refinement efficiency was 99.27%.

Figure 18a depicts the predicted density distribution in the (x, y) plane obtained on the finest anisotropic AMR mesh using the fourth-order CENO finite-volume scheme, which is in very good agreement with those of the reference work by De Sterck and Poedts [20]. Figures 18b, c show the predicted magnetic field amplitude and the blocks of the computational grid in the (x, y) plane for the final anisotropically refined mesh, respectively. Lastly, the regions where the smoothness indicator is active is shown in blue in Fig. 18d. The latter corresponds to the computational cells in which a limited second-order reconstruction has been applied. On the final mesh, just 15% of cells are detected as non-smooth, most lying within the shocks, again illustrating the effectiveness of the smoothness indicator in detecting non-smooth solution content.

5.6 Earth’s Magnetosphere with Northward Interplanetary Magnetic Field

As a last example, numerical results are now presented for the interaction of a solar wind with a magnetized sphere representative of Earth’s magnetosphere. For the case considered, an incoming high-speed solar wind flow with a northward background magnetic field interacts with an assumed idealized magnetic dipole representing the magnetosphere. In order to compare the predicted results to those presented previously by Powell et al. [54], similar free stream and boundary conditions were applied; that is, the solar wind velocity was taken to be $u_x = 400$ km/s, the strength of the magnetic dipole was taken to be 3×10^{-5} Tesla $\times R_E^3$, where R_E the radius of Earth, the free-stream conditions imposed on the outer sphere boundary located at $R_o = 384R_E$, and constant plasma conditions of $\rho = 1$, $\mathbf{U} = 0$, $B_r = 0$, $p = 8$




Fig. 18 Magnetically dominated bow shock flow over a sphere showing **a** the predicted density distribution in the (x, y) plane; **b** the predicted distribution of the magnetic field amplitude in the (x, y) plane; **c** the blocks of the anisotropically refined AMR mesh in the (x, y) plane; and **d** smoothness indicator in the (x, y) plane, for a final mesh consisting of 19,083 blocks and 19,540,992 cells. **a** Density contours. **b** Magnetic field amplitude. **c** Final mesh of 19,083 blocks. **d** Smoothness indicator

at an inner spherical boundary representing the ionosphere located at $R_i = 3R_E$. The computational domain considered here is a standard cubed-sphere grid with 6 sectors. The initial mesh consists of 288 blocks of $12 \times 12 \times 12$ cells for a total of 497,664 computational cells. Steady solutions for the magnetosphere problem were obtained on a sequence of successively anisotropically refined meshes. Five levels of anisotropic AMR were applied using a refinement criteria based on the pressure gradient. The resulting final mesh after 5 levels of refinement was composed of 2306 blocks and 3,984,768 hexahedral cells. The refinement efficiency on the finest mesh was 99.80%.

Figure 19a shows the predicted pressure distribution and magnetic field lines of the steady-state solution for the magnetosphere/solar wind interaction with northward IMF on the finest mesh following the 5 levels of anisotropic refinement. It can be observed that the magnetic field lines of the solar wind are parallel to the terrestrial magnetic dipole field lines and this situation results in a so-called “closed magnetosphere” with very little reconnection between

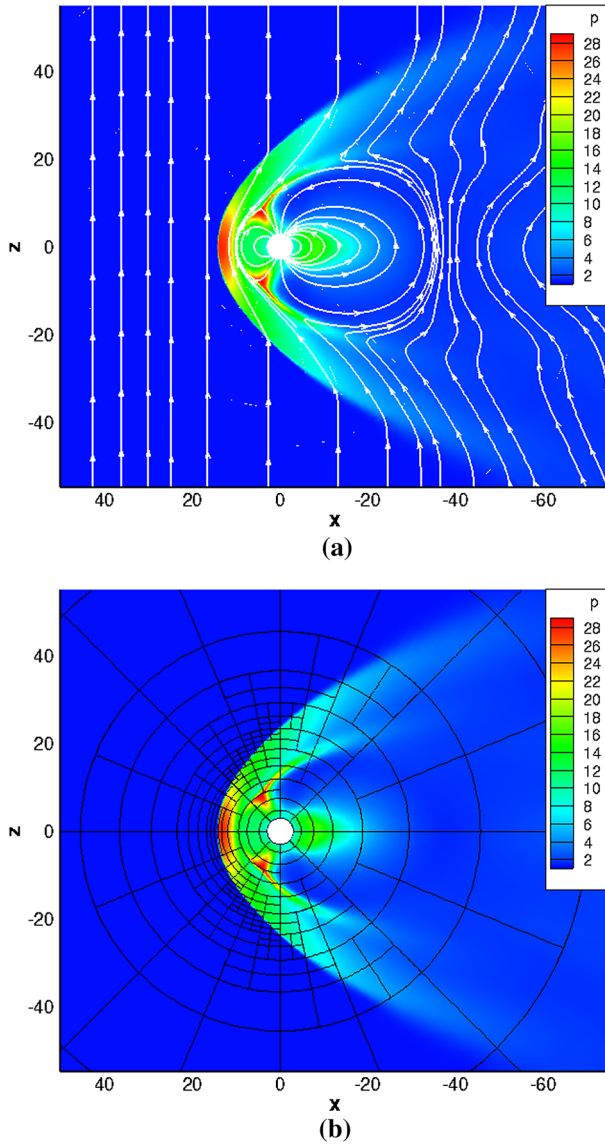


Fig. 19 Magnetosphere interaction of solar with with Northward IMF showing **a** predicted pressure contours and magnetic field lines in the (x, z) plane and **b** anisotropic block-based AMR mesh after 5 levels of refinement. **a** Pressure contours with magnetic field lines. **b** Final mesh with 2306 blocks and 3,984,768 hexahedral cells

terrestrial and IMF field lines. The length of tail of the magnetosphere as represented by the last closed magnetic field lines on the night side (downstream of the Earth) is evaluated to be $x = -35R_E$. This result, as well as the overall predicted structure of the magnetosphere, are in very close agreement with the predictions presented in the previous study by Powell et al. [54]. Figure 19b provides a plot of the grid blocks contained in the final anisotropically refined AMR mesh. It is evident that the blocks are clustered along the bow shock on the day

side. The adapted mesh is clearly well aligned with the bow shock near the stagnation point (refinement is highly anisotropic) but gradually loses alignment moving toward the outer regions of the bow shock where a more isotropic refinement of the mesh is recovered.

6 Conclusions

A fourth-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 fourth-order AMR scheme is 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.

Acknowledgements This work was supported by the Canadian Space Agency and by the Natural Sciences and Engineering Research Council (NSERC) of Canada. In particular, the authors would like to acknowledge the financial support received from the Canadian Space Agency through the Geospace Observatory Canada program. 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.

References

- Adams, M., Colella, P., Graves, D.T., Johnson, J.N., Keen, N.D., Ligoocki, T.J., Martin, D.F., McCorquodale, P.W., Modiano, D., Schwartz, P.O., Sternberg T. D. Van Straalen, B.: Chombo: Software package for AMR applications—design document. Lawrence Berkeley National Technical Report LBNL-6616E
- Balsara, D.: Divergence-free reconstruction of magnetic fields and WENO schemes for magnetohydrodynamics. *J. Comput. Phys.* **228**(14), 5040–5056 (2009)
- Barth, T.: Recent developments in high order k-exact reconstruction on unstructured meshes. In: 31st Aerospace Sciences Meeting (1993)
- Bell, J., Berger, M., Saltzman, J., Welcome, M.: Three-dimensional adaptive mesh refinement for hyperbolic conservation laws. *SIAM J. Sci. Comput.* **15**(1), 127–138 (1994)
- Berger, M.: On conservation at grid interfaces. *SIAM J. Numer. Anal.* **24**, 967–984 (1987)
- Berger, M., Olinger, J.: Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.* **53**(3), 484–512 (1984)
- Brackbill, J., Barnes, D.: The effect of nonzero $\nabla \cdot \mathbf{B}$ on the numerical solution of the magnetohydrodynamics equations. *J. Comput. Phys.* **35**(3), 426–430 (1980)
- Burstedde, C., Wilcox, L.C., Ghattas, O.: *p4est*: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM J. Sci. Comput.* **33**(3), 1103–1133 (2011)
- Charest, M.R.J.: Effect of variables choices on Godunov-type high-order finite-volume methods (to be submitted)
- Charest, M.R.J., Groth, C.P.T.: A high-order central ENO finite-volume scheme for three-dimensional low-speed viscous flows on unstructured mesh. *Commun. Comput. Phys.* **17**, 615–656 (2015)
- Charest, M.R.J., Groth, C.P.T.: A high-order central ENO finite-volume scheme for three-dimensional turbulent flows on unstructured mesh. AIAA Paper (June 2013)
- Charest, M.R.J., Groth, C.P.T., Gülder, Ö.L.: A computational framework for predicting laminar reactive flows with soot formation. *Combust. Theory Model.* **14**(6), 793–825 (2010)
- Chen, Y., Toth, G., Gombosi, T.: A fifth-order finite difference scheme for hyperbolic equations on block-adaptive curvilinear grids. *J. Comput. Phys.* **305**, 604–621 (2016)

14. Christlieb, A.J., Rossmanith, J.A., Tang, Q.: Finite difference weighted essentially non-oscillatory schemes with constrained transport for ideal magnetohydrodynamics. *J. Comput. Phys.* **268**, 302–325 (2014)
15. Clauer, C.R., Gombosi, T.I., De Zeeuw, D.L., Ridley, A.J., Powell, K.G., Van Leer, B., Stout, Q.F., Groth, C.P.T.: High performance computer methods applied to predictive space weather simulations. *IEEE Trans. Plasma Sci.* **28**, 1931–1937 (2000)
16. Clawpack Development Team: Clawpack software (2017). <https://doi.org/10.5281/zenodo.262111>. <http://www.clawpack.org>. Version 5.4.0
17. Colella, P., Dorr, M., Hittinger, J.A., Martin, D.: High-order finite-volume methods in mapped coordinates. *J. Comput. Phys.* **230**, 2952–2976 (2011)
18. Davis, B.N., LeVeque, R.J.: Adjoint methods for guiding adaptive mesh refinement in tsunami modeling. *Pure Appl. Geophys.* **173**, 4055–4074 (2016)
19. De Sterck, H.: Multi-dimensional upwind constrained transport on unstructured grid for shallow water magnetohydrodynamics. *AIAA* (2001)
20. De Sterck, H., Poedts, S.: Intermediate shocks in three-dimensional magnetohydrodynamic bow-shock flows with multiple interacting shock fronts. *Phys. Rev. Lett.* **84**(24), 5524–5527 (2000)
21. De Zeeuw, D., Gombosi, T., Groth, C.P.T., Powell, K., Stout, Q.: An adaptive MHD method for global space weather simulations. *IEEE Trans. Plasma Sci.* **105**, 1956–1965 (2000)
22. Dedner, A., Kemm, F., Kroner, D., Munz, C., Schnitzer, T., Wesenberg, M.: Hyperbolic divergence cleaning for the MHD equations. *J. Comput. Phys.* **175**(2), 645–673 (2002)
23. Einfeldt, B.: On Godunov-type methods for gas dynamics. *SIAM J. Numer. Anal.* **25**(2), 294–318 (1988)
24. Evans, C.R., Hawley, J.F.: Simulation of magnetohydrodynamics flows: a constrained transport method. *Astrophys. J.* **332**, 659–677 (1988)
25. Freret, L., Groth, C.P.T.: Anisotropic non-uniform block-based adaptive mesh refinement for three-dimensional inviscid and viscous flows. In: 22nd AIAA Computational Fluid Dynamics Conference (2015)
26. Freret, L., Groth, C.P.T.: A parallel high-order CENO finite-volume scheme with AMR for three-dimensional ideal MHD flows. In: International Conference On Spectral and High-Order Methods (2016)
27. Freret, L., Groth, C.P.T.: A high-order finite-volume method with anisotropic AMR for ideal MHD flows. In: 55th AIAA Aerospace Science Meeting (2017)
28. Gao, X., Groth, C.P.T.: A parallel adaptive mesh refinement algorithm for predicting turbulent non-premixed combustions flows. *Int. J. Comput. Fluid Dyn.* **20**(5), 349–357 (2006)
29. Gao, X., Groth, C.P.T.: A parallel solution adaptive method for three-dimensional turbulent non-premixed combustions flows. *J. Comput. Phys.* **229**(9), 3250–3275 (2010)
30. Gao, X., Northrup, S.A., Groth, C.P.T.: Parallel solution-adaptive method for two-dimensional non-premixed combustions flows. *Int. J. Prog. Comput. Fluid Dyn.* **11**(2), 76–95 (2011)
31. Groth, C.P.T., De Zeeuw, D., Powell, K., Gombosi, T., Stout, Q.: A parallel adaptive 3D MHD scheme for modeling coronal and solar wind plasma flows, pp. 193–198 (1999)
32. Groth, C.P.T., De Zeeuw, D.L., Gombosi, T.I., Powell, K.G.: Global three-dimensional MHD simulation of a space weather event: CME formation, interplanetary propagation, and interaction with the magnetosphere. *J. Geophys. Res.* **105**(A11), 25053–25078 (2000)
33. Harten, A., Engquist, B., Osher, S., Chakravarthy, S.: Uniformly high-order accurate essentially non-oscillatory scheme III. *J. Comput. Phys.* **131**(1), 3–47 (1997)
34. Helzel, C., Rossmanith, J.A., Taetz, B.: A high-order unstaggered constrained transport method for the three-dimensional ideal magnetohydrodynamics equations based on the method of lines. *J. Sci. Comput.* **35**(2), 623–651 (2013)
35. Helzel, C., Rossmanith, J.A., Taetz, B.: An unstaggered constrained transport method for the 3d ideal magnetohydrodynamic equations. *J. Comput. Phys.* **230**, 3803–3829 (2011)
36. Ivan, L., De Sterck, H., Northrup, S.A., Groth, C.P.T.: Multi-dimensional finite-volume scheme for hyperbolic conservation laws on three-dimensional solution-adaptive cubed-sphere grids. *J. Comput. Phys.* **255**, 205–227 (2013)
37. Ivan, L., De Sterck, H., Susanto, A., Groth, C.P.T.: High-order central ENO finite-volume scheme for hyperbolic conservation laws on three-dimensional cubed-sphere grids. *J. Comput. Phys.* **282**, 157–182 (2015)
38. Ivan, L., Groth, C.P.T.: High-order solution-adaptive central essentially non-oscillatory CENO method for viscous flows. *J. Comput. Phys.* **257**, 830–862 (2013)
39. Jiang, B., Lin, T., Pavinelli, L.: Large-scale computation of incompressible viscous flow by least-squares finite element method. *Comput. Methods Appl. Mech. Eng.* **144**, 213–231 (1994)
40. Jiang, G., Shu, C.: Efficient implementation of weighted ENO schemes. *J. Comput. Phys.* **126**, 202–228 (1996)

41. Keppens, R., Maliani, Z., Van Marle, A.J., Delmont, P., Vlasis, A., van der Holst, B.: Parallel, grid-adaptive approaches for relativistic hydro and magnetohydrodynamics. *J. Comput. Phys.* **231**(1), 718–744 (2012)
42. LeVeque, R.J.: *Numerical Methods for Conservation Laws*. Birkhäuser, Basel (1992)
43. Liu, X.D., Osher, S., Chan, T.: Weighted essentially non-oscillatory schemes. *J. Comput. Phys.* **115**(1), 200–212 (1994)
44. MacNeice, P., Olson, K., Mobarrry, C., de Fainchtein, R., Packer, C.: Paramesh: a parallel adaptive mesh refinement community toolkit. *Comput. Phys. Commun.* **126**, 330–354 (2000)
45. McCorquodale, P., Colella, P.: A high-order finite volume method for conservation laws on locally refined grids. *Commun. Appl. Math. Comput. Sci.* **6**(1), 1–25 (2011)
46. McCorquodale, P., Dorr, M., Hittinger, J., Colella, P.: High-order finite-volume methods for hyperbolic conservation laws on mapped multiblock grids. *J. Comput. Phys.* **288**, 181–195 (2015)
47. McDonald, J.G., Sachdev, J.S., Groth, C.P.T.: Application of gaussian moment closure to micron-scale flows with moving embedded boundaries. *AIAA J.* **52**(9), 1839–1857 (2014)
48. Mignone, A., Tzeferacos, P., Bodo, G.: High-order conservative finite difference GLM-MHD schemes for cell-centered MHD. *J. Comput. Phys.* **229**, 5896–5920 (2010)
49. Mocz, P., Pakmor, R., Springel, V., Vogelsberger, M., Marinacci, F., Hernquist, L.: A moving mesh unstaggered constrained transport scheme for magnetohydrodynamics. *Mon. Not. R. Astron. Soc.* **463**(1), 477–488 (2016)
50. Mocz, P., Vogelsberger, M., Hernquist, L.: A constrained transport scheme for MHD on unstructured static and moving meshes. *Mon. Not. R. Astron. Soc.* **442**(1), 43–55 (2014)
51. Narechania, N., Freret, L., Groth, C.P.T.: Block-based anisotropic AMR with *A Posteriori* adjoint-based error estimation for three-dimensional inviscid and viscous flows. In: *23rd AIAA Computational Fluid Dynamics* (2017)
52. Olsson, F., Petersson, N.: Stability of interpolation on overlapping grids. *Comput. Fluids* **25**, 583–605 (1996)
53. Pärt-Enander, E., Sjögreen, B.: Conservative and non-conservative interpolation between overlapping grids for finite volume solutions of hyperbolic problems. *Comput. Fluids* **23**, 551–574 (1994)
54. Powell, K.G., Roe, P.L., Linde, T.J., Gombosi, T.I., De Zeeuw, D.L.: A solution-adaptive upwind scheme for ideal magnetohydrodynamics. *J. Comput. Phys.* **154**, 284–309 (1999)
55. Sachdev, J.S., Groth, C.P.T., Gottlieb, J.J.: A parallel solution-adaptive scheme for multi-phase core flows in solid propellant rocket motors. *Int. J. Comput. Fluid Dyn.* **19**(2), 159–177 (2005)
56. Shen, C., Qiu, J., Christlieb, A.: Adaptive mesh refinement based on high order finite difference WENO scheme for multi-scale simulations. *J. Comput. Phys.* **230**, 3780–3802 (2011)
57. Shu, C.W.: High-order weighted non-oscillatory schemes for convection dominated problems. *SIAM Rev.* **51**(1), 82–126 (2009)
58. Susanto, A., Ivan, L., De Sterck, H., Groth, C.P.T.: High-order central ENO finite-volume scheme for ideal MHD. *J. Comput. Phys.* **250**, 141–164 (2013)
59. Tobaldini Neto, L., Groth, C.P.T.: A high-order finite-volume scheme for large-eddy simulation of turbulent premixed flames. *AIAA Paper* (January 2014)
60. Toth, G., van der Holst, B., Sokolov, I., De Zeeuw, D., Gombosi, T., Fand, F., Manchester, W., Meng, X., Najib, D., Powell, K., Stout, Q., Glocer, A., Ma, Y., Opher, M.: Adaptive numerical algorithms in space weather modeling. *J. Comput. Phys.* **231**(1), 870–903 (2012)
61. Van Leer, B., Tai, C.H., Powell, K.G.: Design of optimally-smoothing multi-stage schemes for the euler equations. *Tech. Rep. 89-1933-CP*, AIAA (1989)
62. Venditti, D., Darmofal, D.: Anisotropic grid adaptation for functional outputs: application to two-dimensional viscous flows. *J. Comput. Phys.* **187**, 22–46 (2003)
63. Venditti, D., Darmofal, D.: Anisotropic adaptation for functional outputs of viscous flow simulations. *AIAA Paper* (June 2003)
64. Venkatakrishnan, V.: On the accuracy of limiters and convergence to steady state solutions. In: *31st Aerospace Sciences* (1993)
65. Wang, Z.J., Fidkowski, K., Abgrall, R., Bassi, F., Caraeni, D., Cary, A., Deconinck, H., Hartmann, F., Hillewaert, K., Huynh, H.T., Kroll, N., May, G., Persson, P., Leer, B.V., Visbal, M.: High-order CFD methods: current status and perspective. *Int. J. Numer. Methods Fluids* **72**, 1–42 (2012)
66. Williamschen, M.J., Groth, C.P.T.: Parallel anisotropic block-based adaptive mesh refinement algorithm for three-dimensional flow. In: *21st AIAA Computational Fluid Dynamics Conference* (2013)
67. Zanotti, O., Dumbser, M.: Efficient conservative ADER schemes based on WENO reconstruction and space-time predictor in primitive variables. *Comput. Astrophys. Cosmol.* **3**, 1 (2016)
68. Zhang, Z.J., Groth, C.P.T.: Parallel high-order anisotropic block-based adaptive mesh refinement finite-volume scheme. *Paper 2011-3695*, AIAA (2011)