Journal of Computational Physics 282 (2015) 157–182

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

High-order central ENO finite-volume scheme for hyperbolic conservation laws on three-dimensional cubed-sphere grids

L. Ivan [a,*], H. De Sterck [a], A. Susanto [a], C.P.T. Groth [b]

[a] Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, N2L 3G1, Canada
[b] University of Toronto Institute for Aerospace Studies, Toronto, Ontario, M3H 5T6, Canada

ARTICLE INFO

Article history:
Received 4 May 2014
Received in revised form 16 October 2014
Accepted 1 November 2014
Available online 10 November 2014

Keywords:
High-order schemes
CENO
Finite-volume methods
Conservation laws
3D cubed-sphere grids
Compressible gas dynamics
Compressible magnetohydrodynamics

ABSTRACT

A fourth-order accurate finite-volume scheme for hyperbolic conservation laws on three-dimensional (3D) cubed-sphere grids is described. The approach is based on a central essentially non-oscillatory (CENO) finite-volume method that was recently introduced for two-dimensional compressible flows and is extended to 3D geometries with structured hexahedral grids. Cubed-sphere grids feature hexahedral cells with nonplanar cell surfaces, which are handled with high-order accuracy using trilinear geometry representations in the proposed approach. Varying stencil sizes and slope discontinuities in grid lines occur at the boundaries and corners of the six sectors of the cubed-sphere grid where the grid topology is unstructured, and these difficulties are handled naturally with high-order accuracy by the multidimensional least-squares based 3D CENO reconstruction with overdetermined stencils. A rotation-based mechanism is introduced to automatically select appropriate smaller stencils at degenerate block boundaries, where fewer ghost cells are available and the grid topology changes, requiring stencils to be modified. Combining these building blocks results in a finite-volume discretization for conservation laws on 3D cubed-sphere grids that is uniformly high-order accurate in all three grid directions. While solution-adaptivity is natural in the multi-block setting of our code, high-order accurate adaptive refinement on cubed-sphere grids is not pursued in this paper. The 3D CENO scheme is an accurate and robust solution method for hyperbolic conservation laws on general hexahedral grids that is attractive because it is inherently multidimensional by employing a K-exact overdetermined reconstruction scheme, and it avoids the complexity of considering multiple non-central stencil configurations that characterizes traditional ENO schemes. Extensive numerical tests demonstrate fourth-order convergence for stationary and time-dependent Euler and magnetohydrodynamic flows on cubed-sphere grids, and robustness against spurious oscillations at 3D shocks. Performance tests illustrate efficiency gains that can be potentially achieved using fourth-order schemes as compared to second-order methods for the same error level. Applications on extended cubed-sphere grids incorporating a seventh root block that discretizes the interior of the inner sphere demonstrate the versatility of the spatial discretization method.

* Corresponding author.
E-mail addresses: livan@uwaterloo.ca (L. Ivan), hdesterck@uwaterloo.ca (H. De Sterck), asusanto@uwaterloo.ca (A. Susanto), groth@utias.utoronto.ca (C.P.T. Groth).

http://dx.doi.org/10.1016/j.jcp.2014.11.002
0021-9991/

1. Introduction

Over the past decade, cubed-sphere grids [1,2] (see Fig. 1) have rapidly gained popularity for simulations with spherical geometry in a variety of application domains. They are often employed in the modeling of weather, climate, and the oceans [1–16], and have been used in space physics and astrophysics [17,18]. This paper presents a fourth-order accurate finite-volume scheme for conservation laws on three-dimensional (3D) cubed-sphere grids. It extends the central essentially non-oscillatory (CENO) finite-volume schemes that were presented in [19–24] for two-dimensional (2D) compressible flow simulation to 3D geometries with structured hexahedral grids, with particular emphasis on cubed-sphere mesh topologies. The resulting 3D CENO scheme on general hexahedral cells is used to formulate a finite-volume conservation law scheme on 3D cubed-sphere grids that is uniformly fourth-order accurate in all three grid directions and robust against spurious oscillations at shocks. Preliminary results of our approach to extend CENO to 3D hexahedral structured grids were reported in [25]. The CENO finite-volume method can also be applied to unstructured meshes, as is discussed for laminar viscous flows and turbulent reactive flows by Charest et al. in [26–29].

Our goal in this paper is to develop high-order finite-volume schemes on multi-block 3D cubed-sphere grids [30,31]. We employ a multi-block architecture where the six root blocks from Fig. 1 can be refined uniformly multiple times by recursively dividing each block into eight smaller blocks with the same number of cells as the original block [31]. Note that high-order accurate adaptive mesh refinement (AMR) on the cubed sphere is not pursued in this paper. The multi-block scheme we use is based on previous work of Gao and Groth [32,33] for reacting flows, which enables unstructured root block connectivity. Groth and co-researchers [34–38,33,39] have developed block-based AMR finite-volume schemes for the solution of a range of physically complex flows on multi-block body-fitted meshes with second-order accuracy. The application of these methods has included laminar flames [36,39] and high-pressure soot prediction [40–42], turbulent non-premixed flames [37,33,39] as well as turbulent multi-phase rocket core flows [35,43], micro-scale flows [44,38], and radiation transport [45]. Extensions of the multi-block methods for embedded boundaries not aligned with the mesh [46] and for use with an anisotropic refinement strategy [47,48] are also possible and have been developed. This adaptive multi-block approach was extended to cubed-sphere grids (with second-order accuracy) and magnetohydrodynamics (MHD) flow problems by Ivan, De Sterck and co-workers in [49,30,31,24]. Note that this manuscript focuses on describing the new 3D CENO method on non-adaptive cubed-sphere grids and does not pursue adaptive refinement with efficient high-order accurate solution transfer operators, which is subject of ongoing work.

Hexahedral cubed-sphere grid cells are challenging for high-order accurate finite-volume discretizations because they have nonplanar cell interfaces. A first specific challenge in extending the 2D CENO method from [19–23] to 3D grids with general hexahedral cells is, thus, the high-order treatment of hexahedral cells with nonplanar cell faces. We address this challenge by employing a trilinear representation of the finite-volume cells that facilitates high-order accurate numerical integration in a transformed domain [50–52]. Preliminary results of this approach to extend CENO to 3D hexahedral structured grids were reported in [25].

With the high-order CENO formulation on general hexahedral cells in place, a second specific challenge in employing this scheme on 3D cubed-sphere grids is how to build consistent stencils for high-order discretization at degenerate block edges and corners, i.e., block edges and corners where the number of neighboring blocks is smaller than in a grid with Cartesian topology, and where fewer ghost cells are available. This challenge is addressed in a way that is general for a variety of stencils and block configurations by requiring consistency in a double sense: first, the number of cells dropped in the degenerate stencil needs to be as small as possible compared to the base stencil while maintaining the same stencil structure in terms of connectivity to neighboring cells; and second, degenerate stencils generated for ghost cells need to be identical to the degenerate stencils generated for the corresponding physical cells in neighboring blocks. It is important to formulate a general strategy that meets these requirements in the complex degenerate 3D topology with missing ghost cells. We present a general rotation mechanism to automatically build consistent high-order stencils at these degenerate locations on the cubed-sphere grid. The mechanism we propose is sufficiently general to also handle the case of a cubed-sphere grid with an additional grid block filling the interior [53,54], where blocks may have corners where three degenerate edges meet. The resulting high-order CENO scheme extends the second-order cubed-sphere grid method from [30,31] to fourth-order accuracy. In fact, our approach is more general and the principle extends to orders of accuracy higher than four (by increasing the stencil size), but this paper focuses on a fourth-order accurate implementation.

Two-dimensional cubed-sphere grids can be obtained from a regular Cartesian grid that is defined on the six faces of a cube by projecting the cube faces onto a sphere, resulting in a surface grid for the sphere that is composed of six panels with quasi-uniform spacing and without polar singularities. A 3D cubed-sphere grid (see Fig. 1) is formed by stacking a sequence of concentric 2D cubed-sphere grids in the radial direction, resulting in a grid for the domain between the inner and outer concentric spheres that is composed of six structured sector blocks. Note that the mesh topology within these six blocks is Cartesian, but the connectivity between the six blocks (which form root blocks in our grid data structure) is unstructured (see Fig. 1). The radial edges of the six root blocks (the edges corresponding to the eight corners of the 2D cubed-sphere grid) are degenerate, in the sense that these edges have only three incident root blocks, whereas block edges in a structured, Cartesian topology have four incident blocks. Note that we will also consider the case of a cubed-sphere grid where the interior is filled by a seventh root block with Cartesian topology, which may, for example, model the interior of a planet. The cubed-sphere grids used in this paper are generated using the so-called cubic-gnomonic projection [1,2,4]. In particular, the angularly equidistant mapping described in [2] is used to generate the initial six blocks of the grids.

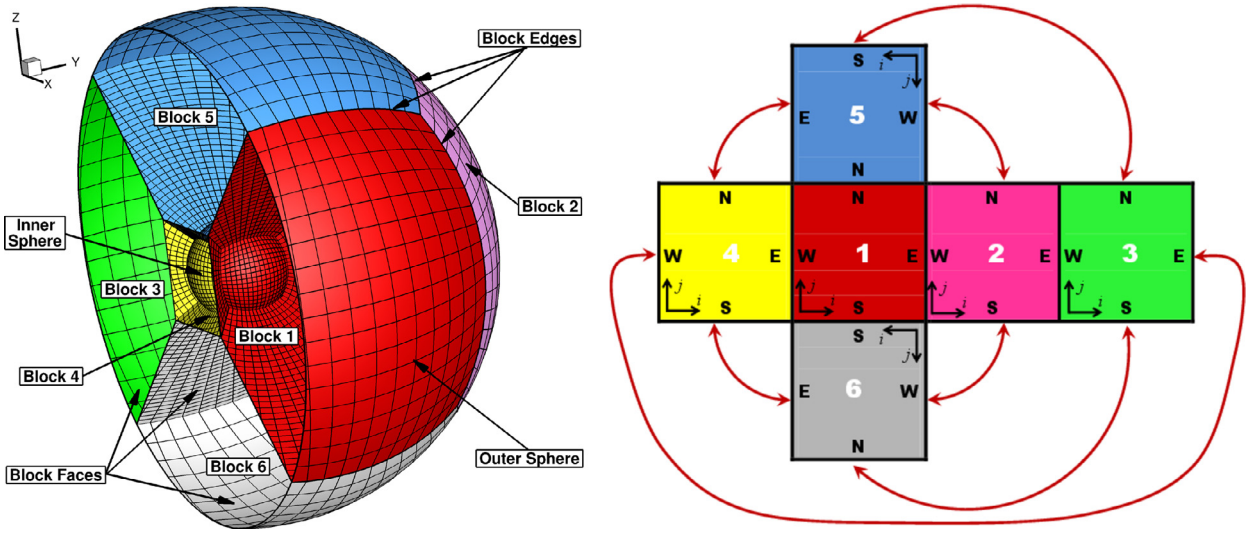

(a) Cross-section of the cubed-sphere grid (b) Illustration of unstructured connectivity among root blocks

Fig. 1. Three-dimensional cubed-sphere grid with six root blocks (corresponding to the six sectors of the grid) and depiction of inter-block connectivity. In our multi-block approach, the root blocks can be refined by recursively dividing each block into eight smaller blocks with the same number of cells as the original block. In panel (b), the block faces are denoted with the initials of the cardinal directions: North (N), East (E), South (S) and West (W). (Figures adapted from [31].)

Significant advances have been made for second-order and high-order discretizations on cubed-sphere grids in recent years [6,7,53,8–16]. Most of these works target 2D cubed-sphere grids with curvilinear coordinates, and only few address high-order discretizations on 3D cubed-sphere grids. An example of the latter is [9], which presents a 3D dynamical core on a cubed-sphere grid that achieves fourth-order accuracy in the horizontal direction and second-order in the vertical direction. To our knowledge, our method is the first to offer full fourth-order discretization accuracy on 3D cubed-sphere grids in all grid directions, and it does so uniformly using the same scheme in the three grid directions. Note that we use a standard orthogonal x, y, z (Cartesian) coordinate system in our conservation law discretization on domains between two concentric spheres, and for our purposes there is no need to define separate curvilinear coordinate systems on each of the cubed-sphere grid sectors as is usually done for layered models in atmospheric flow simulation.

Obtaining high-order accurate finite-volume discretizations on cubed-sphere grids is a challenge: grid lines have slope discontinuities at the interfaces between the six root blocks, and the connectivity is unstructured at the radial edges of the root blocks. This means that special reconstruction procedures are required for obtaining an order of accuracy higher than one when finite-volume schemes are employed that use the direction-by-direction approach that is standard for grids with Cartesian topology [1,2,55,53,6,18,56]. In contrast, the proposed CENO method is multidimensional (as opposed to direction-by-direction) and uses a least-squares reconstruction mechanism with an overdetermined stencil. As a consequence, our approach can directly employ physical cells from adjacent sectors as ghost cells when required (so no complicating special reconstruction procedures are required at interfaces between the root blocks), and it can automatically handle varying stencil size at the degenerate block edges. This means that uniform fourth-order accuracy is naturally obtained in all three grid directions and at all locations of the grid.

More generally, the development of high-order accurate and robust discretizations for hyperbolic conservation laws, especially in 3D, also remains a challenging and active area of research. In recent years, significant advances have been made in proposing new discretization schemes of various types for 3D hyperbolic conservation laws, including finite difference methods (e.g., [57]), discontinuous Galerkin methods (e.g., [58–60]), finite-volume methods (e.g. [61–65]), and combinations of these approaches (e.g., [66]). Our 3D CENO method proposes a different approach on general hexahedral grids that combines high-order accuracy with robustness at discontinuities, and it is attractive because it is inherently multidimensional by employing a K-exact overdetermined reconstruction scheme, and it avoids the complexity of considering multiple non-central stencil configurations that characterizes traditional ENO schemes. Our approach is general and can be extended to unstructured grids as discussed by Charest et al. in [27–29] for laminar viscous flows and turbulent reactive flows. For 3D MHD flows, we extend the 2D fourth-order accurate CENO MHD solver from [24] to 3D, using the generalized Lagrange multiplier (GLM) approach from [67] to control errors in the divergence of the magnetic field for the high-order accurate simulation (as in [57,24]). While many space physics simulation codes were only second-order accurate initially [68–70, 31], there have been significant developments extending these approaches toward high-order accurate discretizations [57, 71–76], but these developments are often specific to structured grids, and our parallel 3D fourth-order MHD method further extends these developments to general hexahedral grids and significantly advances the state-of-the-art of large-scale simulation frameworks for cubed-sphere grids.

The remainder of this paper is structured as follows. Section 2 describes our fourth-order accurate finite-volume scheme for conservation laws on 3D cubed-sphere grids, using a high-order polynomial reconstruction approach with CENO monotonicity enforcement combined with trilinear representation of the cubed-sphere grid cells to obtain high-order accurate flux calculations. The specific difficulty of generating consistent stencils near degenerate block edges is addressed using a general rotation mechanism. Section 3 describes extensive numerical tests demonstrating the accuracy and robustness of the fourth-order CENO method for Euler and MHD flows on Cartesian and cubed-sphere grids. Section 4 concludes.

2. Fourth-order CENO method for hyperbolic conservation laws on 3D cubed-sphere grids

2.1. Finite-volume formulation

We consider nonlinear conservation laws of the form

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

where $\mathbf{U}$ is the vector of conserved variables, $\vec{\mathbf{F}}$ is the flux dyad, and $\mathbf{S}$ and $\mathbf{Q}$ are numerical and physical source terms that may arise for certain equation sets and application problems. While our fourth-order cubed-sphere grid method is developed for general conservation laws, we consider in this paper the particular cases of the MHD and Euler equations. For MHD (with GLM approach to control divergence errors as in [67,24]) $\mathbf{U}$ is given by

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

where ρ is the gas density, $\vec{V} = (V_x, V_y, V_z)$ is the velocity, $\vec{B} = (B_x, B_y, B_z)$ is the magnetic field, ρe is the total energy density, and ψ is the generalized Lagrange multiplier employed to control errors in the divergence of the magnetic field. Here, the total energy is given by $\rho e = p/(\gamma - 1) + \rho V^2/2 + B^2/2$, where V and B are the magnitudes of the velocity and magnetic field vectors, respectively, and γ is the ratio of specific heats. The flux dyad, $\vec{\mathbf{F}}$, is given by

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

where $\mathbb{I}$ is the 3×3 identity matrix. The numerical source term employed to control magnetic field divergence [67,24], $\mathbf{S}$, is given by

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

The parameters c_h and c_p are chosen as in [24], following [67]. When $\vec{B}$ and ψ are set to zero, the equations reduce to the Euler equations of gas dynamics.

We formulate the finite-volume method in terms of hexahedral cells in a 3D structured grid block with indices ijk. The semi-discretization for the temporal evolution of the cell average $\bar{\mathbf{U}}_{ijk}$ in cell ijk is derived from

$$\frac{d\bar{\mathbf{U}}_{ijk}}{dt} = \frac{1}{V_{ijk}}\left[-\oiint_{\partial\mathcal{V}_{ijk}} \vec{\mathbf{F}}\cdot\vec{n}\,\mathrm{da} + \iiint_{\mathcal{V}_{ijk}} (\mathbf{S} + \mathbf{Q})\,\mathrm{dv}\right], \tag{5}$$

where V_{ijk} is the volume of cell $\mathcal{V}_{ijk}$, $\vec{n}$ is the unit outward normal of the cell surface $\partial\mathcal{V}_{ijk}$, and da and dv are surface and volume elements, respectively. Approximating the surface and volume integrals numerically we obtain

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

where f is used to indicate the six interfaces of the hexahedral cell, and m indicates the Gauss quadrature points over each interface. Here, $\vec{\mathbf{F}}_{\mathrm{num}}$ is the numerical flux function, $\bar{\mathbf{S}}_{ijk}$ and $\bar{\mathbf{Q}}_{ijk}$ are cell-averaged source terms (obtained using Gaussian quadrature), and $\mathbf{R}_{ijk}$ denotes the residual vector in cell ijk, which depends on the set of cell averages $\bar{\mathbf{U}}$.

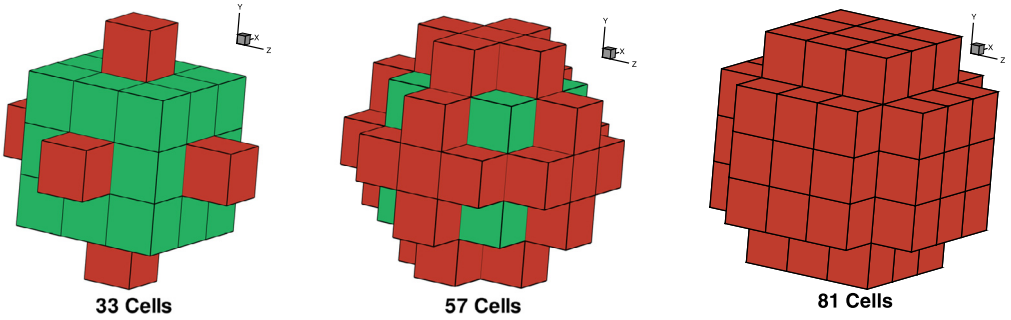

(a) 27 first-ring cells plus 6 extra second-ring cells for a total of 33 cells

(b) 27 first-ring cells plus 6 × 5 extra second-ring cells for a total of 57 cells

(c) 27 first-ring cells plus 6 × 9 extra second-ring cells for a total of 81 cells

Fig. 2. Examples of several central reconstruction stencils for a cell ijk ranging in size from 33 to 81 cells that can be used to determine the 20 polynomial coefficients of a cubic reconstruction ($K = 3$). The first-degree neighbors of the cell ijk are shown in green whereas dark red is used for the second-degree neighbors.

2.2. K-exact reconstruction and stencil choice

The 2D high-order CENO method of [19–24] achieves high-order accuracy by employing K-exact polynomial reconstruction [77] of the solution field in each cell using the cell average and the cell averages of neighboring cells in a reconstruction stencil. In our 3D extension, we compute a K-exact polynomial reconstruction for each solution variable in cell ijk of the form

$$u^K_{ijk}(\vec{X}) = \sum_{\substack{p_1=0 \\ (p_1+p_2+p_3\le K)}}^{K} \sum_{p_2=0}^{K} \sum_{p_3=0}^{K} (x-\bar{x}_{ijk})^{p_1} (y-\bar{y}_{ijk})^{p_2} (z-\bar{z}_{ijk})^{p_3} D^K_{p_1 p_2 p_3}, \tag{7}$$

where K is the degree of the polynomial, $\vec{X} = (x, y, z)$ is the coordinate vector, $(\bar{x}_{ijk}, \bar{y}_{ijk}, \bar{z}_{ijk})$ are the coordinates of the centroid of cell ijk, and the $D^K_{p_1 p_2 p_3}$ are the polynomial coefficients that are determined such that $u^K_{ijk}(\vec{X})$ matches cell averages in the reconstruction stencil with high accuracy. There are

$$\mathcal{N}_D = \frac{(K+1)(K+2)(K+3)}{6} \tag{8}$$

polynomial coefficients $D^K_{p_1 p_2 p_3}$, and they can be computed such that a function $u_{\text{exact}}(\vec{X})$ is represented with accuracy

$$u^K_{ijk}(\vec{X}) - u_{\text{exact}}(\vec{X}) = \mathcal{O}\left(\Delta x^{K+1}\right), \tag{9}$$

while a polynomial with degree K can be reconstructed exactly [77]. In this paper, we use cubic reconstruction ($K = 3$) to obtain a fourth-order accurate numerical scheme. In this case, $\mathcal{N}_D = 20$. We choose the coefficients such that the cell average of the reconstruction exactly equals the cell average of variable u in cell ijk,

$$\bar{u}_{ijk} = \frac{1}{V_{ijk}} \iiint_{\mathcal{V}_{ijk}} u^K_{ijk}(\vec{X})\, \mathrm{dv}, \tag{10}$$

and for all cells $\gamma\delta\zeta$ in the reconstruction stencil of cell ijk we also desire that the cell averages of the reconstruction equal the cell averages $\bar{u}_{\gamma\delta\zeta}$:

$$\left(\frac{1}{V_{\gamma\delta\zeta}} \iiint_{\mathcal{V}_{\gamma\delta\zeta}} u^K_{ijk}(\vec{X})\, \mathrm{dv} \right) - \bar{u}_{\gamma\delta\zeta} = 0. \tag{11}$$

As in [77,19–24], we use overdetermined reconstruction stencils, and impose Eq. (11) in the least-squares sense, while Eq. (10) is imposed exactly. Note that the use of overdetermined stencils is especially attractive in the case of cubed-sphere grids, where the grid topology is unstructured near root-block edges and smaller numbers of neighbor cells are available than in regular Cartesian topology. Our approach naturally allows for reconstruction stencils of reduced size, without reducing the order of accuracy. The reconstruction stencils we consider in this paper have 33, 57 and 81 cells, and are depicted in Fig. 2. We will compare these stencils in numerical tests which show that the 33-cell stencil is optimal for our fourth-order accurate simulations (see Section 3 below).

Eq. (10) can be enforced analytically by replacing u^K_{ijk} with Eq. (7) and expressing the first coefficient, D^K_{000}, as a function of the other $M = \mathcal{N}_D - 1$ polynomial unknowns as

$$D^K_{000} = \bar{u}_{ijk} - \sum_{\substack{p_1=0 \\ (1\le p_1+p_2+p_3\le K)}}^{K} \sum_{p_2=0}^{K} \sum_{p_3=0}^{K} D^K_{p_1p_2p_3} \left(\overline{x^{p_1} y^{p_2} z^{p_3}}\right)_{ijk}, \tag{12}$$

where the geometric moment $(\overline{x^{p_1} y^{p_2} z^{p_3}})_{ijk}$ of powers (p_1, p_2, p_3) is given by

$$\left(\overline{x^{p_1} y^{p_2} z^{p_3}}\right)_{ijk} = \frac{1}{V_{ijk}} \iiint_{\mathcal{V}_{ijk}} (x - \bar{x}_{ijk})^{p_1} (y - \bar{y}_{ijk})^{p_2} (z - \bar{z}_{ijk})^{p_3} \, \mathrm{dv}. \tag{13}$$

Substituting u^K_{ijk} from Eq. (7) in Eq. (11) and using Eq. (12) for D^K_{000} the following overdetermined linear system for the M unknowns is obtained:

$$\mathbb{L}\mathbf{D} - \mathbf{B} = \begin{bmatrix} L_1 \\ L_2 \\ \vdots \\ L_J \\ \vdots \\ L_{N_n} \end{bmatrix}_{N_n \times M} \begin{pmatrix} D^K_{001} \\ D^K_{002} \\ \vdots \\ D^K_{p_1p_2p_3} \\ \vdots \\ D^K_{K00} \end{pmatrix}_{M\times 1} - \begin{pmatrix} w_1(\bar{u}_1 - \bar{u}_I) \\ w_2(\bar{u}_2 - \bar{u}_I) \\ \vdots \\ w_J(\bar{u}_J - \bar{u}_I) \\ \vdots \\ w_{N_n}(\bar{u}_{N_n} - \bar{u}_I) \end{pmatrix}_{N_n \times 1} = \begin{pmatrix} 0 \\ 0 \\ \vdots \\ 0 \\ \vdots \\ 0 \end{pmatrix}_{N_n \times 1}, \tag{14}$$

where N_n is the number of neighbors in the reconstruction stencil. Here we have used the shorthand notation I to refer to cell ijk about which we do the reconstruction, and we have used indices $1, 2, \ldots, J, \ldots, N_n$ to refer to the other cells in the stencil of cell ijk. The w_J are geometric weights for each control volume J which serve the purpose of improving the locality of the reconstruction, becoming especially important for stretched meshes with surface curvature [78]. We use

$$w_J = \frac{1}{\|\Delta\vec{X}_{IJ}\|^\theta}, \tag{15}$$

with $\Delta\vec{X}_{IJ} = (\Delta x_{IJ}, \Delta y_{IJ}, \Delta z_{IJ}) = \vec{X}_J - \vec{X}_I$ and $\theta = 1, 2$. We use $\theta = 2$ except where noted otherwise. Row L_J of matrix $\mathbb{L}$ for a neighboring cell J is given by

$$L_J = \left(w_J \widehat{(x^0 y^0 z^1)}_{IJ} \quad w_J \widehat{(x^0 y^0 z^2)}_{IJ} \quad \ldots \quad w_J \widehat{(x^{p_1} y^{p_2} z^{p_3})}_{IJ} \quad \ldots \quad w_J \widehat{(x^K y^0 z^0)}_{IJ}\right). \tag{16}$$

The matrix coefficients $\widehat{(x^{p_1} y^{p_2} z^{p_3})}_{IJ}$ can be computed efficiently using only the $(\overline{x^{p_1} y^{p_2} z^{p_3}})$ moments:

$$\begin{aligned} \widehat{\left(x^{p_1} y^{p_2} z^{p_3}\right)}_{IJ} &= \left(\frac{1}{V_J} \iiint_{\mathcal{V}_J} (x - x_I)^{p_1} (y - y_I)^{p_2} (z - z_I)^{p_3} \, \mathrm{dv}\right) - \left(\overline{x^{p_1} y^{p_2} z^{p_3}}\right)_I \\ &= \sum_{\ell=0}^{p_1} \sum_{q=0}^{p_2} \sum_{m=0}^{p_3} \left[C^\ell_{p_1} C^q_{p_2} C^m_{p_3} \Delta x^\ell_{IJ} \Delta y^q_{IJ} \Delta z^m_{IJ} \left(\overline{x^{(p_1-\ell)} y^{(p_2-q)} z^{(p_3-m)}}\right)_J\right] - \left(\overline{x^{p_1} y^{p_2} z^{p_3}}\right)_I, \end{aligned} \tag{17}$$

where the binomial coefficients C^β_α can be computed efficiently by recursion as

$$C^\beta_\alpha = \frac{\alpha - \beta + 1}{\beta} C^{\beta-1}_\alpha, \qquad C^0_\alpha = 1. \tag{18}$$

See [20] for the equivalent expressions in 2D.

The solution of the overdetermined linear system equation (14) can be obtained via the normal equations, or, with better numerical stability, using QR factorization or by multiplication with the pseudo-inverse of $\mathbb{L}$ [79,20]. In each time step, the constrained least-squares reconstruction problem is solved for each cell and for each primitive variable. Matrix $\mathbb{L}$ depends completely on geometry and is the same for all least-squares problems in a given cell ijk and for all time steps, so its inverse can be precomputed and reused to provide computational speedup. Column scaling can be applied to improve conditioning of the system. (See [19,20,22] for details in the 2D case.)

A specific technical difficulty in extending the 2D high-order CENO finite-volume method to 3D grids composed of general hexahedral cells with nonplanar surfaces, namely, the high-order accurate computation of the surface and volume integrals that were introduced in this and the previous subsection, is discussed in Section 2.4.

2.3. CENO monotonicity enforcement

In order to control spurious oscillations at discontinuities, we use the CENO monotonicity procedure that was introduced by Ivan and Groth [19,23] for the 2D Euler equations, and has since been extended to the Navier–Stokes equations [20–22] and MHD [24] in 2D. The CENO procedure switches between an unlimited high-order accurate reconstruction (piecewise cubic in this paper, leading to a fourth-order accurate scheme) and a limited piecewise-linear reconstruction (second-order accurate), with the switching based on the smoothness indicator introduced in [19]. The smoothness indicator is computed in each cell for every reconstructed variable to determine whether the flow is locally smooth and well-resolved. For cells containing non-smooth or under-resolved solution content, the unlimited K-exact reconstruction is switched to limited piecewise linear reconstruction. It should be emphasized that local switching to lower order is only performed for those reconstructed variables that are deemed non-smooth or under-resolved. Note that the CENO scheme is called central because both the high-order and the low-order stencils are central with respect to the cell. The method is called an ENO method because it satisfies the ENO property of [80], which allows the presence of small spurious oscillations that have a magnitude on the order of the truncation error, but it does not allow $O(1)$ Gibbs-like oscillations at discontinuities. The fixed stencil used during the CENO reconstruction procedure avoids the complexity of considering multiple non-central stencil configurations that characterizes traditional ENO schemes.

Limiting for the second-order method is performed using the procedure developed by Park et al. [81] specifically for multiple dimensions in conjunction with the slope limiter function of [82]. The use of the multidimensional limiter formulation improved the representation of linear functions without unnecessary limiting or clipping of solution content. The limiting is applied at the corners of the hexahedral cells and not at the quadrature points for the faces. Apart from monotonicity, this limiting strategy offers significant computational advantages by reducing, at least by a factor of three, the number of reconstruction evaluations required for the calculation of the limiter.

Full details on the CENO smoothness indicator and switching mechanism are described in the 2D context in [19,23, 20–22], and here we briefly present the generalization of the relevant formulas to 3D. For every cell ijk and primitive solution variable we compute a variable $\mathcal{S}$, the smoothness indicator, by

$$\mathcal{S} = \frac{\alpha\, c_s}{\max(1-\alpha, \epsilon)}, \tag{19}$$

where α is given by

$$\alpha = 1 - \frac{\sum_\gamma \sum_\delta \sum_\zeta (u^K_{\gamma\delta\zeta}(\vec{X}_{\gamma\delta\zeta}) - u^K_{ijk}(\vec{X}_{\gamma\delta\zeta}))^2}{\sum_\gamma \sum_\delta \sum_\zeta (u^K_{\gamma\delta\zeta}(\vec{X}_{\gamma\delta\zeta}) - \bar{u}_{ijk})^2}, \tag{20}$$

and c_s is given by

$$c_s = \frac{\mathcal{N}_{\mathrm{SOS}} - \mathcal{N}_{\mathrm{D}}}{\mathcal{N}_{\mathrm{D}} - 1}. \tag{21}$$

Here, $\mathcal{N}_{\mathrm{SOS}}$ indicates the size of the stencil used for reconstruction, $\mathcal{N}_{\mathrm{D}}$ (number of degrees of freedom) denotes the number of unknown polynomial coefficients (e.g., 20 for $K=3$), $\epsilon = 10^{-8}$ is introduced to avoid division by zero, and $\vec{X}_{\gamma\delta\zeta}$ is the centroid of cell $\gamma\delta\zeta$ in the stencil of cell ijk. (The stencil used for computing the smoothness indicator can be chosen smaller than the reconstruction stencil. In particular, we compute the smoothness indicator using a 27-cell stencil (with 26 first-degree neighbors).)

The parameter α, with range $-\infty < \alpha \leq 1$, measures how accurately centroidal solution values of neighboring cells can be reproduced using the reconstruction for cell ijk: in smooth flow α is close to one, while at cells with a discontinuity or an under-resolved feature, α tends away from one. The smoothness indicator $\mathcal{S}$ has range $-c_s < \mathcal{S} < c_s/\epsilon$: for smooth variation $\mathcal{S}$ is large, and for non-smooth or under-resolved features $\mathcal{S}$ is small. We then choose a cut-off value, $\mathcal{S}_\mathcal{C}$, and when $\mathcal{S} > \mathcal{S}_\mathcal{C}$ the high-order reconstruction is used, while the limited low-order reconstruction is used otherwise. As in previous studies [22,24], values for $\mathcal{S}_\mathcal{C}$ are typically selected in the range of 1000 to 5000. For the simulations presented in this paper we choose $\mathcal{S}_\mathcal{C} = 1500$, except where noted. The form of the smoothness indicator is inspired by the definition of multiple-correlation coefficients and least-squares goodness-of-fit testing; see [20] for a more detailed discussion with further motivation for the approach.

As discussed in [22], using standard Taylor approximation theory and assuming u is a continuous and differentiable function and the mesh is Cartesian with uniform spacing Δx, it is rather straightforward to show that

$$u^K_{\gamma\delta\zeta}(\vec{X}_{\gamma\delta\zeta}) - u^K_{ijk}(\vec{X}_{\gamma\delta\zeta}) \approx \mathcal{O}\left(\Delta x^{K+1}\right). \tag{22}$$

Similarly, it can be shown that

$$u^K_{\gamma\delta\zeta}(\vec{X}_{\gamma\delta\zeta}) - \bar{u}_{ijk} \approx \mathcal{O}(\Delta x). \tag{23}$$

This implies that

$$\alpha \approx 1 - \mathcal{O}\left(\Delta x^{2K}\right), \tag{24}$$

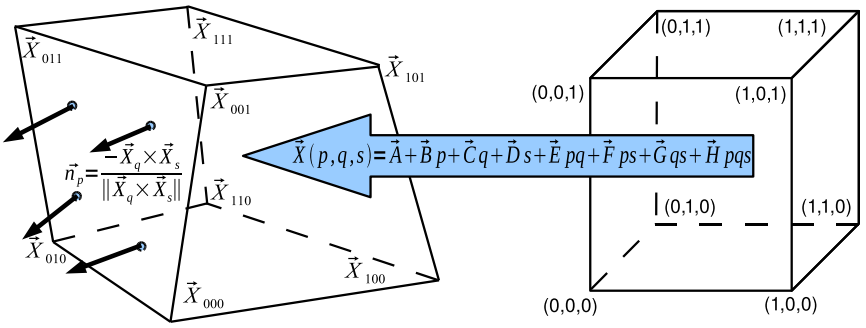

Fig. 3. A general hexahedral cell in physical space having faces with nonplanar vertices (left) to which a reference unit cube (right) is mapped by defining a trilinear transformation $\vec{X}(p,q,s)$. The position vectors for the hexahedral cell vertices, $\vec{X}_{000}$, $\vec{X}_{100}$, $\vec{X}_{010}$, $\vec{X}_{001}$, $\vec{X}_{110}$, $\vec{X}_{101}$, $\vec{X}_{011}$ and $\vec{X}_{111}$, are mapped one-to-one to the vertices of the reference cube, where the Cartesian coordinates in the unit cube, (p,q,s), take binary values. Additionally, the unit normals at four points located on the hexahedral face defined by constant coordinate $p=0$ are shown.

for smooth resolved solution content. In this case, it is evident that $\alpha \to 1$ and $\mathcal{S} \to \infty$ as $\Delta x \to 0$ and this occurs at a rate that is significantly more rapid than the formal order of accuracy of the scheme. In this way, the asymptotic accuracy of the K-exact reconstruction is recovered as $\Delta x \to 0$. Conversely, for non-smooth solutions it is expected that

$$u^K_{\gamma\delta\zeta}(\vec{X}_{\gamma\delta\zeta}) - u^K_{ijk}(\vec{X}_{\gamma\delta\zeta}) \approx \mathcal{O}(1), \qquad u^K_{\gamma\delta\zeta}(\vec{X}_{\gamma\delta\zeta}) - \bar{u}_{ijk} \approx \mathcal{O}(1), \tag{25}$$

and therefore α will generally not be close to one.

As explained in [24], it is beneficial for stability to detect nearly-uniform regions by considering

$$\xi_{ijk} = \sqrt{\frac{1}{\mathcal{N}_D - 1} \sum_{\substack{p_1=0 \\ (1\le p_1+p_2+p_3\le K)}}^{K} \sum_{p_2=0}^{K} \sum_{p_3=0}^{K} \left(D^K_{p_1p_2p_3}\right)^2 \xi_V(p_1,p_2,p_3)}, \tag{26a}$$

$$\xi_V(p_1,p_2,p_3) = \left(V_{ijk}^{\frac{2}{3}}\right)^{p_1+p_2+p_3}, \tag{26b}$$

which measures the variability of solution variable u in cell ijk. (It takes into account derivatives at the centroid of cell ijk.) When ξ_{ijk} is smaller than a threshold value (low variability), high-order reconstruction is always used, and only when ξ_{ijk} is greater than the threshold the smoothness indicator is computed and the CENO switching mechanism is activated. In particular, the smoothness indicator for the solution variable u is evaluated in cell ijk when

$$\xi_{ijk} > \epsilon_A + \epsilon_R \bar{u}_{ijk}, \tag{27}$$

where ϵ_A and ϵ_R represent absolute and relative variability thresholds, chosen to be 10^{-5} for the simulations performed in this paper, except where noted otherwise. Notice that the ξ_V term provides an appropriate length scaling for each derivative based on the cell volume, and it has only K distinct values due to the common variation range for p_1, p_2 and p_3. Moreover, using only the first order derivatives to measure the variability of solution variable u is less expensive and has been found to perform equally well in our numerical studies. This simplification is obtained in Eq. (26a) by taking $\mathcal{N}_D = 4$ and computing the summation terms with $K = 1$. This simplification is used in our numerical tests.

2.4. High-order integration for hexahedral cells with nonplanar cell faces

A specific technical difficulty in obtaining high-order accuracy on 3D cubed-sphere grids is to properly handle the nonplanar cell faces of the cubed-sphere grid cells. We use a trilinear representation for the nonplanar cell faces which can be exploited to obtain sufficiently accurate volume and surface integrals [50–52], as we now briefly explain (see Fig. 3).

Given a hexahedral cell in physical space with vertices defined by position vectors $\vec{X}_{000}$, $\vec{X}_{100}$, $\vec{X}_{010}$, $\vec{X}_{001}$, $\vec{X}_{110}$, $\vec{X}_{101}$, $\vec{X}_{011}$, $\vec{X}_{111}$, a trilinear mapping from a reference unit cube can be defined as

$$\vec{X}(p,q,s) = \vec{A} + \vec{B}p + \vec{C}q + \vec{D}s + \vec{E}pq + \vec{F}ps + \vec{G}qs + \vec{H}pqs, \tag{28}$$

where, p,q,s are the coordinates in the reference domain, and $\vec{A} = \vec{X}_{000}$, $\vec{B} = \vec{X}_{100} - \vec{A}$, $\vec{C} = \vec{X}_{010} - \vec{A}$, $\vec{D} = \vec{X}_{001} - \vec{A}$, $\vec{E} = \vec{X}_{110} - \vec{A} - \vec{B} - \vec{C}$, $\vec{F} = \vec{X}_{101} - \vec{A} - \vec{B} - \vec{D}$, $\vec{G} = \vec{X}_{011} - \vec{A} - \vec{C} - \vec{D}$, $\vec{H} = \vec{X}_{111} - \vec{A} - \vec{B} - \vec{C} - \vec{D} - \vec{E} - \vec{F} - \vec{G}$ (see, e.g., [50] for details). Note that the vertices of the hexahedral cell in physical space are obtained for $(p,q,s) \in \{0,1\} \times \{0,1\} \times \{0,1\}$. Thus, $\vec{X}_{000} = \vec{X}(0,0,0)$, $\vec{X}_{100} = \vec{X}(1,0,0)$, $\vec{X}_{010} = \vec{X}(0,1,0)$, $\vec{X}_{001} = \vec{X}(0,0,1)$, $\vec{X}_{110} = \vec{X}(1,1,0)$, $\vec{X}_{101} = \vec{X}(1,0,1)$, $\vec{X}_{011} = \vec{X}(0,1,1)$, $\vec{X}_{111} = \vec{X}(1,1,1)$, and the six trilinear surfaces are described by fixing one of p, q and s to 0 or 1. The tangent vectors to the coordinate lines are defined by

$$\frac{\partial \vec{X}}{\partial p} \equiv \vec{X}_p(q,s) = \vec{B} + \vec{E}q + \vec{F}s + \vec{H}qs, \tag{29a}$$

$$\frac{\partial \vec{X}}{\partial q} \equiv \vec{X}_q(p,s) = \vec{C} + \vec{E}p + \vec{G}s + \vec{H}ps, \tag{29b}$$

$$\frac{\partial \vec{X}}{\partial s} \equiv \vec{X}_s(p,q) = \vec{D} + \vec{F}p + \vec{G}q + \vec{H}pq. \tag{29c}$$

The determinants of the Jacobians for volume and surface integration are given by

$$\det \mathbf{J}(p,q,s) \equiv \left|\frac{\partial(x,y,z)}{\partial(p,q,s)}\right| = \vec{X}_p \cdot (\vec{X}_q \times \vec{X}_s), \tag{30a}$$

$$\det \mathbf{J}_p(q,s) \equiv \|\vec{X}_q \times \vec{X}_s\|\big|_{p=\mathrm{ct}}, \tag{30b}$$

$$\det \mathbf{J}_q(p,s) \equiv \|\vec{X}_s \times \vec{X}_p\|\big|_{q=\mathrm{ct}}, \tag{30c}$$

$$\det \mathbf{J}_s(p,q) \equiv \|\vec{X}_p \times \vec{X}_q\|\big|_{s=\mathrm{ct}}, \tag{30d}$$

and the normal vectors to the coordinate planes by

$$\vec{n}_p(q,s) \equiv -\left.\frac{\vec{X}_q \times \vec{X}_s}{\det \mathbf{J}_p}\right|_{p=\mathrm{ct}}, \tag{31a}$$

$$\vec{n}_q(p,s) \equiv -\left.\frac{\vec{X}_s \times \vec{X}_p}{\det \mathbf{J}_q}\right|_{q=\mathrm{ct}}, \tag{31b}$$

$$\vec{n}_s(p,q) \equiv -\left.\frac{\vec{X}_p \times \vec{X}_q}{\det \mathbf{J}_s}\right|_{s=\mathrm{ct}}. \tag{31c}$$

2.4.1. Volumetric integrals

The volume of hexahedral cell $I = ijk$ is defined by

$$V_I = \iiint_{\mathcal{V}_{ijk}} \mathrm{dv}, \tag{32}$$

where $\mathrm{dv} = \mathrm{d}x\,\mathrm{d}y\,\mathrm{d}z$ is the volume element. The cell centroid is

$$\vec{X}_I = \frac{1}{V_I} \iiint_{\mathcal{V}_{ijk}} \vec{g}(\vec{X})\,\mathrm{dv}, \tag{33}$$

where the vector-valued function $\vec{g}(\vec{X}) = [x, y, z]^{\mathrm{T}}$, which is the position vector in the physical space. Thus, the more general procedure of evaluating a volumetric integral of a continuous smooth function, $g(\vec{X})$, over a control volume $\mathcal{V}_{ijk}$ is discussed here, which recovers the volume calculation for $g(\vec{X}) = 1$ and, for example, the x-coordinate of the cell centroid for $g(\vec{X}) = x$.

To evaluate the volumetric integral,

$$\mathcal{I} = \iiint_{\mathcal{V}_{ijk}} g(\vec{X})\,\mathrm{dv}, \tag{34}$$

the variables and integration domain are changed to those of the reference unit cube by making use of the trilinear transformation, $\vec{X} = \vec{X}(p,q,s)$, and its transformation Jacobian determinant, $\det \mathbf{J}$, [50]. Thus, the volumetric integral, $\mathcal{I}$, is calculated in the canonical space (p,q,s) as

$$\mathcal{I} = \int_0^1\int_0^1\int_0^1 g\big(\vec{X}(p,q,s)\big) \det \mathbf{J}\,\mathrm{d}p\,\mathrm{d}q\,\mathrm{d}s \tag{35a}$$

$$\simeq \sum_{m=1}^{N_v} g\big(\vec{X}(p_m,q_m,s_m)\big)(\det \mathbf{J})_m \omega_m = \sum_{m=1}^{N_v} g(\vec{X}_m)\tilde{\omega}_m, \tag{35b}$$

where N_v is the number of volumetric Gauss points (see Fig. 4). For our fourth-order accurate scheme, we use 27 Gauss points in the volume of the hexahedral cell, and 4 Gauss points on each cell face. All volumetric integrals are computed with

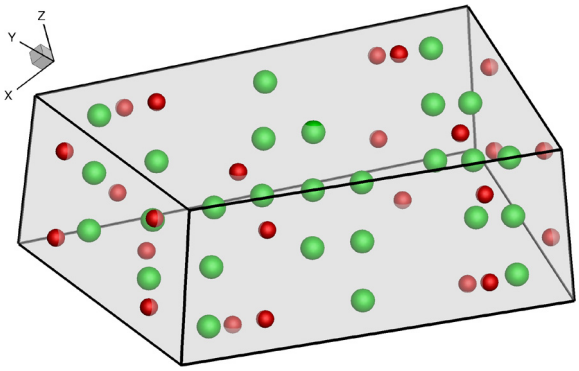

Fig. 4. Volumetric (green) and surface (red) Gauss quadrature points used in the evaluation of volumetric and surface integrals over a hexahedral control volume to achieve fourth-order accuracy. Note that there are 27 volumetric points, and 4 surface points for each hexahedral face. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Table 1
Summary of volumetric integrals used in the formulation of the CENO finite-volume scheme.

Symbol	Description	Definition	Section
V_I	volume of control volume I	Eq. (32)	2.4.1
$\vec{X}_I$	centroid of control volume I	Eq. (33)	2.4.1
$(\overline{x^{p_1} y^{p_2} z^{p_3}})_I$	geometric moment of powers (p_1, p_2, p_3) of control volume I about its own centroid	Eq. (13)	2.2
$\bar{u}_I$	mean value of u in control volume I	Eq. (10)	2.2
$(\bar{\mathbf{Q}})_I$	average source term in control volume I	Eq. (5)	2.1, 3.3.1

Table 2
Surface integrals used in the formulation of the CENO finite-volume scheme.

Symbol	Description	Definition	Section
$\oiint_{\partial\mathcal{V}_I} \vec{\mathbf{F}} \cdot \vec{n}\,\mathrm{da}$	flux through the boundaries of control volume I	Eq. (5)	2.1

a fifth-order Gaussian integration rule derived with three points in each canonical direction (i.e., a $3 \times 3 \times 3$ product rule [52]) and $N_v = 27$ volumetric Gauss points. In this rule, the basic abscissae and weights generated for a $[0, 1]$ domain [83] are $[0.5, 0.5 \pm \sqrt{3/5}]$ and $[8/18, 5/18, 5/18]$, respectively. Note that the Gaussian abscissas $\vec{X}_m$ and weights $\tilde{\omega}_m = (\det \mathbf{J})_m \omega_m$ are needed for initial computation of volumetric integrals of various quantities (see Table 1) but need not be stored during the actual simulation, unless required for integration of time-dependent volumetric source terms. Table 1 summarizes the different volumetric integrals that need to be computed in the 3D CENO scheme.

2.4.2. Surface integrals

Surface integrals are evaluated in a similar fashion. For example, to evaluate a surface integral along a surface with constant p-coordinate, $\mathcal{A}_p$, the following expression is used:

$$\mathcal{I}_p = \iint_{\mathcal{A}_p} g(\vec{X})\,\mathrm{da} = \int_0^1 \int_0^1 g\big(\vec{X}(p,q,s)\big) \det \mathbf{J}_p \,\mathrm{d}q\,\mathrm{d}s \tag{36a}$$

$$\simeq \sum_{m=1}^{N_g} g\big(\vec{X}(p_m, q_m, s_m)\big)\,(\det \mathbf{J}_p)_m \omega_m = \sum_{m=1}^{N_g} g(\vec{X}_m)\,\tilde{\omega}_m, \tag{36b}$$

and the expressions for integrals along surfaces $\mathcal{A}_q$ and $\mathcal{A}_s$ follow by cyclic permutation. Here, N_g is the number of Gauss points on the surfaces (see Fig. 4). The $\vec{X}_m$ and $\tilde{\omega}_m = (\det \mathbf{J}_p)_m\,\omega_m$ are stored in our framework to increase the computational performance. Table 2 indicates that surface integrals are used to compute the fluxes through cell faces in the 3D CENO scheme.

2.5. Multi-block cubed-sphere mesh

As described in detail in [30,31], the cubed-sphere mesh is implemented using an MPI-parallel multi-block method with octree data structures, where the six root blocks can be refined recursively by dividing each block into eight smaller blocks with the same number of cells as the original block. Boundary conditions are imposed using ghost cells. All blocks have

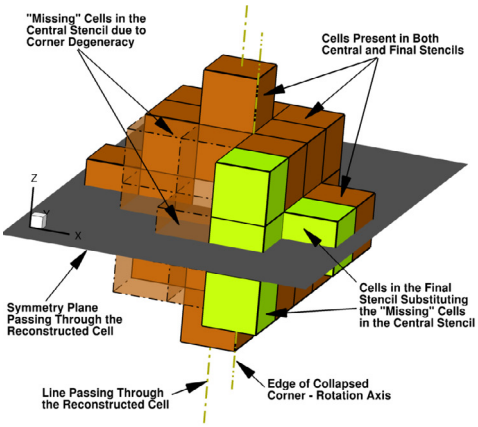

(a) 3D diagram of the stencil.

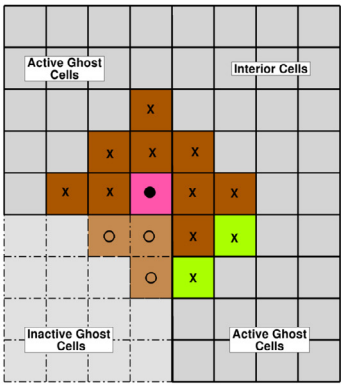

(b) 2D diagram showing the stencil in the symmetry plane and active and inactive ghost cells.

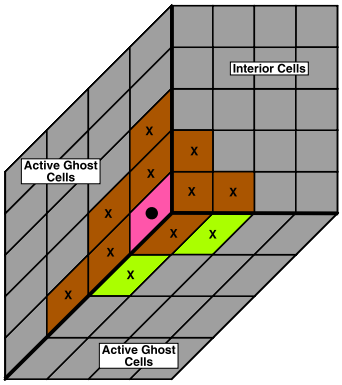

(c) Compact representation of the stencil in the symmetry plane.

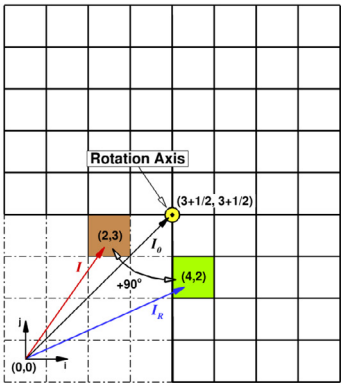

(d) Illustration of the index vectors used in Eq.(37) to map cell (2,3) into (4,2).

Fig. 5. Example of a degenerate stencil with 33 cells for a reconstructed ghost cell touching a degenerate edge. The cells of the original 33-cell central stencil remaining active in the final stencil are colored dark brown. The cells colored in green are replacing the cells identified as "missing" in the original stencil, which are colored light brown/transparent. The final stencil has 30 cells, with three cells being "lost" due to the degeneracy of the grid block edge and the collapsing of the ghost cells in the corner opposite to the interior cells of the grid block. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

the same number of cells, and load balancing is achieved by distributing the blocks over the processor cores as evenly as possible, where multiple blocks are normally allocated per core.

2.6. Rotation mechanism for generating consistent stencils near degenerate block edges

In a multi-block grid with regular Cartesian topology throughout the grid, four grid blocks are incident to any edge of a grid block. In the cubed-sphere grid of Fig. 1, however, the connectivity of the six root blocks is unstructured (i.e., the topology is non-Cartesian), and any grid block edge along one of the eight radial lines that go through the eight corners of the cubed-sphere grid have only three incident grid blocks (as depicted in a 2D projection in Fig. 5(c)). We call grid block edges along which the number of incident grid blocks is less than four *degenerate* edges. In our implementation, blocks along these edges have missing ghost cells in the opposite corner along the degenerate edge (see Figs. 5(a) and 5(b)) (we say that the ghost cells in the corner are *collapsed* or *inactive*, see also [31]). This means that the CENO reconstruction stencils we use (Fig. 2) need to be adjusted since the topology is different and not all cells used in the original stencils are available. In this section we propose a general rotation-based mechanism to automatically derive modified stencils near degenerate edges. The mechanism is sufficiently general that it can also be used to derive modified stencils for grid blocks that have three degenerate edges intersecting in a corner. This occurs, for example, when considering a cubed-sphere grid in which the interior of the inner sphere is filled by a seventh root block with Cartesian topology, see Fig. 6. This seventh

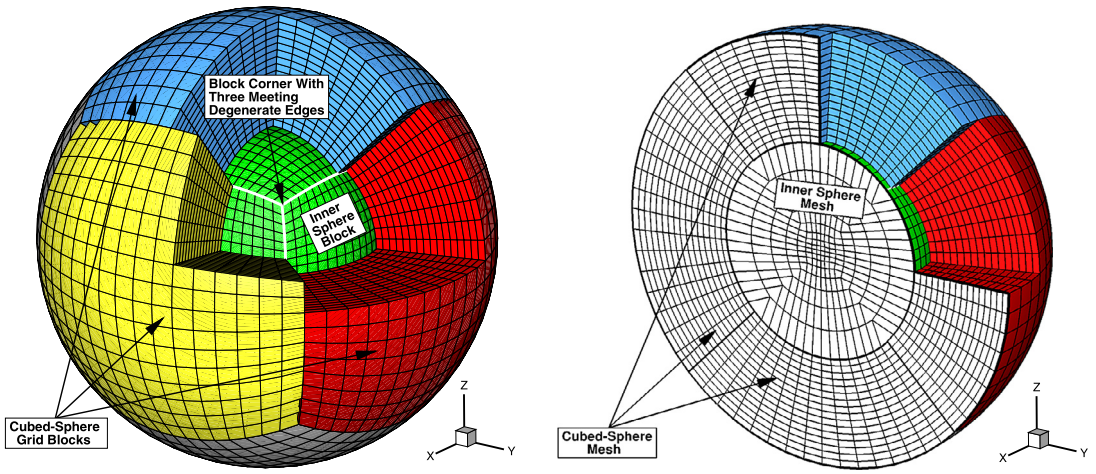

(a) 3D Representation. (b) Cross Section.

Fig. 6. Nested spherical grid comprising one spherical root block filling the interior of a cubed-sphere grid with six root blocks.

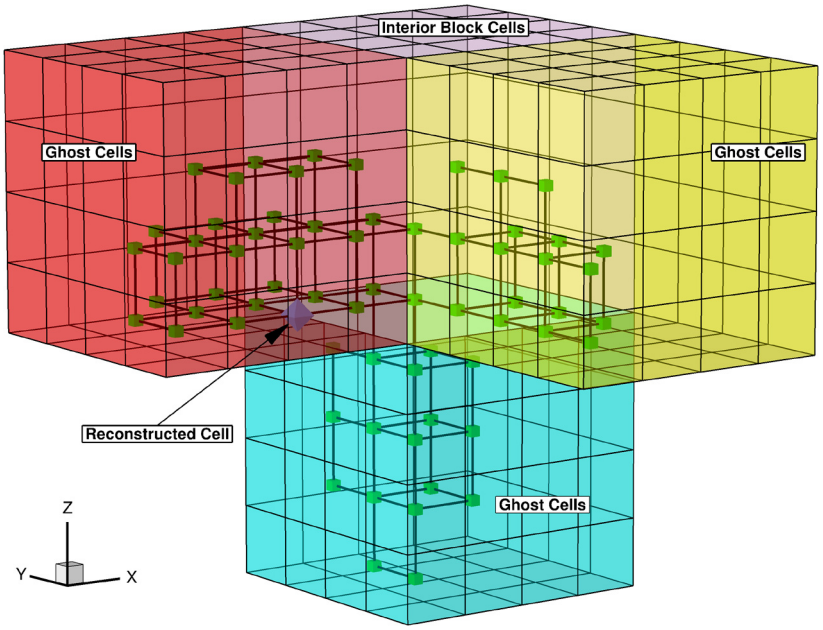

Fig. 7. Grid block with three degenerate edges. The green dots indicate the degenerate stencil obtained using the proposed rotation mechanism for a base stencil with 81 cells centred about the reconstruction cell (purple). All ghost cells lying in the non-colored blocks are collapsed. The resulting degenerate stencil has 66 cells. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

block may, for example, represent the interior of a lunar or planetary object, in which a magnetic diffusion equation may be solved when modeling interaction with the solar wind [84]. The grid inside the interior sphere is generated as in [54]. In this type of grid, grid blocks that contain one of the eight corners of the interior root block, have three degenerate edges emanating from that corner, see Figs. 6 and 7. (In fact, four degenerate block edges emanate from the corner (including one in the radial direction).) We will apply the proposed rotation mechanism about each of the three degenerate edges to obtain modified stencils in three-way degenerate blocks in a systematic way.

The rotation mechanism we propose functions as follows. We explain the mechanism for the case depicted in Fig. 5, where we want to use the 33-cell base reconstruction stencil of Fig. 2 for computing the reconstruction in a ghost cell of a grid block that lies next to the corner of the block. The grid block is one-way degenerate with the degenerate edge in the z direction (indicated by the dash-dotted line 'Edge of Collapsed Corner' in Fig. 5(a)). Since the corner ghost cells are missing (light brown/transparent cells), the stencil needs to be adjusted. Consider, for simplicity, the central xy plane of the stencil. By folding the two remaining layers of active ghost cells toward each other (as in Fig. 5(c)), it can be seen that the two green cells now become distance-two neighbors of the pink cell (where the reconstruction is computed), and can be included

in the modified stencil (since distance-two cells are included in the original stencil). More generally, the non-Cartesian topology makes the collapsed (inactive) ghost cells unavailable to the stencil, but equivalent cells (in terms of distance) in the remaining ghost cell layers are available to complete the stencil. A general mechanism that works for all stencils of Fig. 2 can be formulated in terms of a rotation mechanism: in Figs. 5(b) and 5(d), we rotate the collapsed ghost cells that were part of the original stencil (light brown) by 90 degrees about the degenerate edge away from the reconstruction cell (pink). A rotated cell may land on a cell that does not form part of the existing stencil cells, in which case those new cells (the green cells in Figs. 5(b) and 5(c)) are added to the modified stencil. In case a rotated cell lands on an existing stencil cell (e.g., the brown cell next to the two green cells), no new cell is added to the stencil. The same mechanism is applied to the stencil planes above and below the central plane in Fig. 5, resulting in two more green cells added to the modified stencil. As a result, the 33-cell stencil, which was missing seven collapsed ghost cells, is modified by including four new rotated collapsed ghost cells, resulting in a final 30-cell stencil. This rotation mechanism provides modified stencils that are consistent with the original stencils from Fig. 2 in a double sense: first, the number of cells dropped in the degenerate stencil is minimized by incorporating new cells taking into account the new neighbor configuration and distances in the non-Cartesian topology; and second, degenerate stencils generated for ghost cells are identical to the degenerate stencils generated for the corresponding physical cells in neighboring blocks (which is required for consistency of flux calculation, which we rely on for conservation in the implementation of our finite-volume methods).

In terms of implementation, we obtain the coordinates of rotated collapsed ghost cells by the following matrix rotation mechanism. For simplicity, we formulate the equations in the 2D setting corresponding to Fig. 5(d). We choose the origin of the ij coordinate system for the block in the center of the lower left cell of Fig. 5(d) (which is located in the collapsed ghost cell region). Since there are four layers of ghost cells for the block depicted in the figure, the rotation axis (the degenerate block edge) is located at coordinate $\vec{I}_0 = (3 + 1/2, 3 + 1/2)$. Let $\vec{I}$ be the coordinate of a collapsed ghost cell that is rotated towards the active ghost cell layer about the degenerate block edge (e.g., $\vec{I} = (2, 3)$ for the brown cell in Fig. 5(d)), and let $\vec{I}_R$ be the coordinate of its image under the rotation ($\vec{I}_R = (4, 2)$ for the green cell). The cell coordinate after rotation, $\vec{I}_R$, is obtained by

$$\vec{I}_R = \vec{I}_0 + \mathbb{R}(\vec{I} - \vec{I}_0), \tag{37}$$

where the rotation matrix for the case of Fig. 5(d) is given by

$$\mathbb{R} = \begin{bmatrix} 0 & -1 \\ 1 & 0 \end{bmatrix}. \tag{38}$$

For three-way degenerate blocks (see Fig. 7), the collapsed ghost cells are rotated about each of the three degenerate edges to identify cells that can be added to the modified stencil. Note that the 27-cell stencil we use for our second-order scheme is also reduced at degenerate edges, and we use the same rotation mechanism to determine the degenerate stencil. On the other hand, the 27-cell stencil cannot be used for cubic reconstruction (fourth-order scheme), since distance-two neighbors need to be added at least in the directions of the grid lines to have a well-posed reconstruction problem.

We emphasize the value of a systematic automatic procedure for determining these degenerate stencils. While in principle these special stencils could be hand-derived and hard-coded, this would be a formidable task and would be error-prone due to the multitude of cases that occur for the different stencils of Fig. 2 and for the many different geometrical configurations of degenerate blocks in 3D space and their ghost cells and physical cells located at different positions in the block. Hard-coding these special cases would also lead to unmanageable additional complexity of the block-parallel 3D computer code. For this reason, our automatic procedure to select appropriate smaller stencils using a general rotation-based mechanism is an important ingredient of our approach. Furthermore, the mechanism is attractive because it is sufficiently flexible to handle three-way degenerate blocks that occur, for example, when a seventh root block is incorporated into the cubed-sphere grid to discretize the interior of the inner sphere of the grid.

2.7. Time integration

We use standard explicit second-order and fourth-order Runge–Kutta time integration methods for the second-order and fourth-order accurate spatial discretizations to be compared in our time-dependent numerical test problems (see [24,20] and references therein). For steady-state simulations, we use a five-stage optimally smoothing method regardless of the solution accuracy.

3. Numerical results

To validate the proposed 3D CENO method, we perform extensive numerical tests, including CENO reconstruction tests of known functions, 3D simulation problems on rectangular box domains, and problems on cubed-sphere grids in various configurations (with five, six and seven root blocks). We include Euler and MHD problems to evaluate fourth-order accurate grid convergence and robustness of the high-order method at shocks. We also compare the efficiency of the fourth-order method compared to a second-order method [24]. Except when noted otherwise, we employ the local Lax–Friedrichs numerical flux function in all tests, as in [24]. We compare fourth-order CENO with the second-order limited method described in [31]. We

use the 33-cell stencil from Fig. 2(a) for high-order results, except where noted. For the second-order comparison tests we use a standard 27-cell stencil. Large parallel simulation runs were performed on the GPC system at SciNet, which consists of 3780 Intel Xeon E5540 (2.53 GHz) nodes with 16 GB RAM per node and a total of 30,240 cores. The cluster is connected with a high-speed InfiniBand switched fabric communications link.

In grid convergence studies based on an exact solution, the L_1, L_2, and L_∞ norms of the numerical solution error are computed as follows:

$$L_1 = |E|_1 = \frac{1}{V_T}\sum_{i,j,k}\iiint_{\mathcal{V}_{ijk}} \left|u^K_{ijk}(\vec{X}) - f(\vec{X})\right| \mathrm{dv}, \tag{39}$$

$$L_2 = |E|_2 = \sqrt{\frac{1}{V_T}\sum_{i,j,k}\iiint_{\mathcal{V}_{ijk}} \left[u^K_{ijk}(\vec{X}) - f(\vec{X})\right]^2 \mathrm{dv}}, \tag{40}$$

$$L_\infty = |E|_\infty = \max_{i,j,k}\left(\frac{1}{V_{ijk}}\iiint_{\mathcal{V}_{ijk}} \left|u^K_{ijk}(\vec{X}) - f(\vec{X})\right| \mathrm{dv}\right), \tag{41}$$

where V_T is the total volume of the computational domain, $f(\vec{X})$ is the exact solution evaluated at point $\vec{X}$, and the summation is taken over all the grid cells.

3.1. Three-dimensional CENO reconstructions

We describe three examples of smooth function reconstructions to demonstrate the accuracy of the CENO reconstruction that lies at the core of the proposed high-order finite-volume scheme. These reconstruction tests proceed by first computing highly accurate cell averages for a given function, then using these cell averages to compute high-order polynomial reconstructions in the cells, and finally computing the error between the original function and the polynomial reconstruction by high-accuracy numerical integration over each cell. The order of convergence of this error as grids are refined measures the order of accuracy of the CENO reconstruction, which determines the theoretical order of accuracy of the numerical simulation method. The discrete initial data in these cases is generated by accurately integrating the exact solution over cells to obtain the required cell averages using the integration procedure outlined in Section 2.4 in combination with an adaptive approach [85,20], and this adaptive approach is also used to compute the error with high accuracy.

3.1.1. Problem 1: Reconstruction of a smooth function in a box with distorted grid

To assess the accuracy of the 3D high-order CENO procedure on meshes containing hexahedral cells with nonplanar faces, reconstructions of the smooth function $f(x, y, z) = (\cos(\pi(y+1)) - \cos(\pi z))e^{-\pi(x+1)}$ are compared to the exact solution on distorted meshes for a rectangular box domain. The meshes are distorted in such a way that cell faces are nonplanar. The computational domain is the rectangular box defined by $0 < x < 1$ and $-1 < y, z < 1$. Fig. 8(a) illustrates how the interior nodal points have been perturbed randomly, and depicts the high-order solution reconstruction obtained on a structured 3D mesh with eight blocks of $4 \times 8 \times 8$ cells and 2048 total computational cells.

Fig. 8(b) shows grid convergence reconstruction studies with the fourth-order CENO method on a series of meshes. We use $\mathcal{S}_\mathcal{C} = 1500$ for the smoothness indicator cut-off value. The initial mesh has one block with $8 \times 16 \times 16$ cells and 2048 total cells and the final mesh has 4096 blocks and $\mathbf{N} = 8{,}388{,}608$ total computational cells. The expected theoretical asymptotic convergence rate of the fourth-order accurate method is achieved in all error norms. As the mesh is refined, the slopes of the L_1-, L_2-, and L_∞-error norms approach -4.087, -4.075 and -3.836, respectively, thereby providing validation for the proposed trilinear-based CENO reconstruction procedure to general hexahedral cells with nonplanar faces.

3.1.2. Problem 2: Reconstruction of a modulated exponential on the cubed sphere

To demonstrate the spatial accuracy of the high-order hybrid CENO scheme on 3D cubed-sphere grids, we consider the solution reconstruction of $f(x, y, z) = (1 - R + R^2)e^{x+y+z}$ on the computational domain defined by two concentric spheres with inner and outer radius $R_i = 1$ and $R_o = 3$, where R denotes the radius. As depicted in Fig. 9(a), this function exhibits a large smooth variation spanning several orders of magnitude that is oriented along the line connecting two diametrically-opposed cubed-sphere corners, where the function maximum and minimum occur. Therefore, this is a good test case for studying how the locally-reduced reconstruction stencil employed by the high-order reconstruction approach handles solution extrema occurring at the sector corners of the cubed sphere (where the stencil size is reduced). In Fig. 9(b), we compare error norms of the fourth-order CENO reconstruction with a second-order limited K-exact procedure ($K = 1$) which is based on a linear reconstruction with a stencil of 26 first-degree neighbors. The computational meshes used in this grid convergence study range in size from 3072 to 100,663,296 cells, using from six to 3072 solution blocks.

The error norms obtained for the two reconstruction procedures show that both schemes achieve the theoretical convergence accuracy in the asymptotic limit. Moreover, it can be observed by inspecting the convergence plot of the fourth-order reconstruction that the CENO procedure with a cut-off value of 1500 is able to handle the exponential solution variation both accurately and robustly: on the first mesh containing only $8 \times 8 \times 8$ cells per grid sector, the function is under-resolved

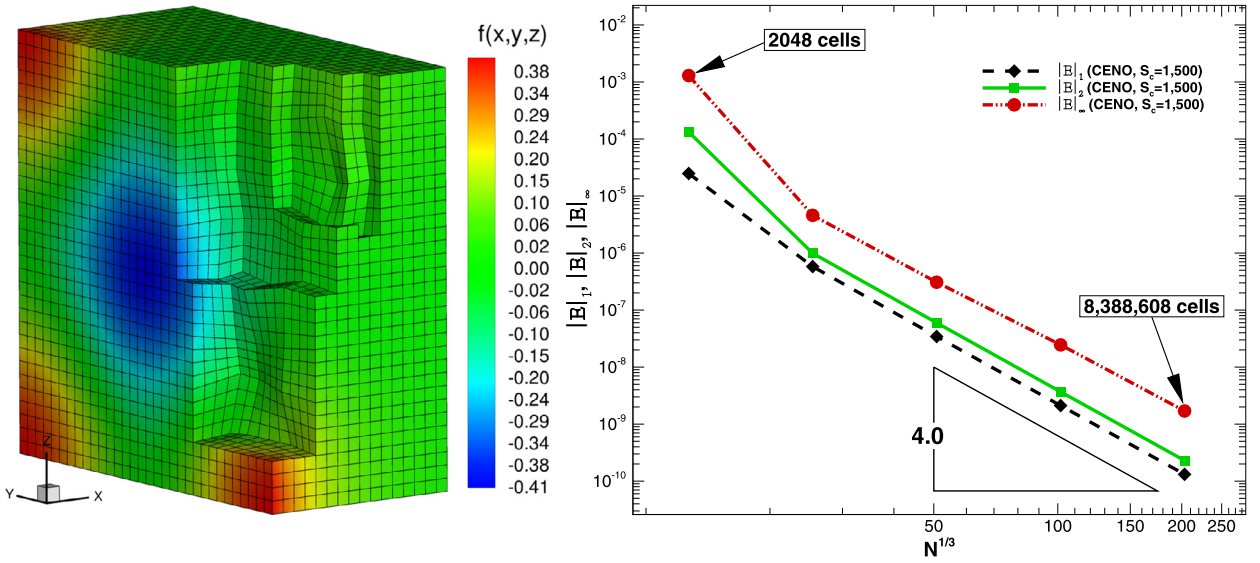

(a) Reconstructed solution. (b) Error norms.

Fig. 8. Problem 1: (a) Solution reconstruction obtained using the fourth-order CENO scheme on a distorted mesh with 8 blocks of size $4 \times 8 \times 8$ and 2048 cells; and (b) L_1-, L_2-, and L_∞-error norms for the fourth-order ($K = 3$) CENO scheme versus the equivalent number of grid nodes in one direction. Fourth-order accurate convergence is obtained for all three error norms.

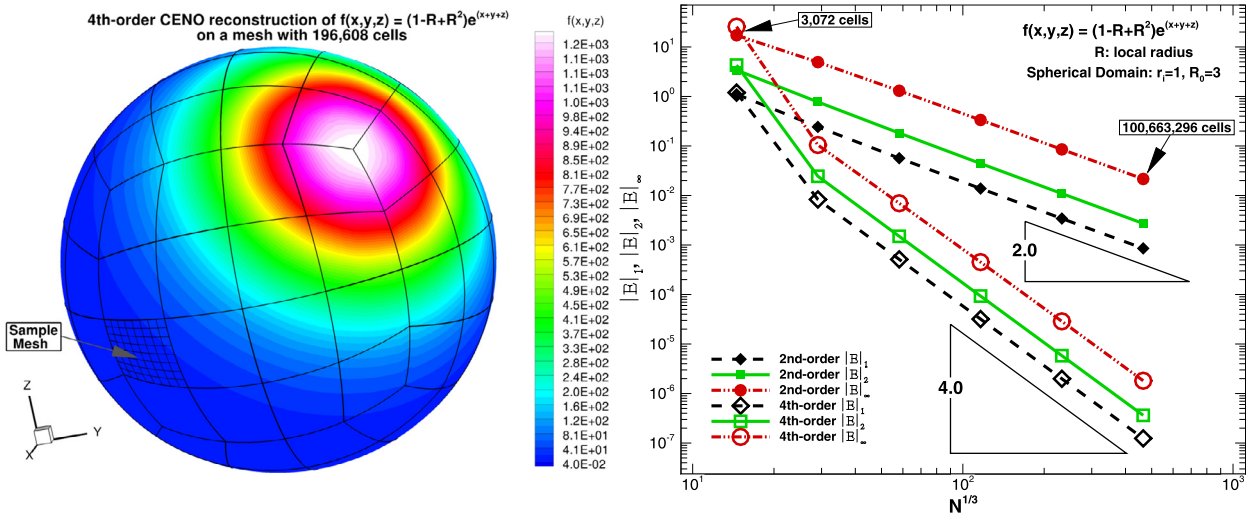

(a) Reconstructed solution. (b) Error norms.

Fig. 9. Problem 2: (a) Solution reconstruction of $f(x, y, z) = (1 - R + R^2)e^{x+y+z}$ on a cubed-sphere grid with 384 $8 \times 8 \times 8$ blocks and 196,608 cells obtained with the fourth-order ($K = 3$) CENO method; and (b) L_1-, L_2-, and L_∞-error norms for a second-order ($K = 1$) limited reconstruction and a fourth-order ($K = 3$) CENO approach.

and the high-order CENO procedure avoids creating oscillatory reconstructions by switching to monotonic piecewise linear approximations with larger errors but maintaining monotonicity, whereas on more refined meshes the K-exact reconstruction is retained everywhere and the reconstruction procedure achieves its theoretical accuracy. Thus, the L_1, L_2, and L_∞ norms of the solution error in the asymptotic limit are -2.068, -2.121 and -2.101 for the second-order reconstruction method and -4.034, -4.055 and -4.081 for the fourth-order CENO reconstruction. The improved accuracy exhibited by the high-order algorithm translates into significant savings in terms of number of computational cells for a targeted solution error. For example, even to achieve a modest $L_1 = 10^{-3}$ solution error, the second-order method requires about 100,663,296 cells, which is more than 500 times the mesh requirements of the high-order method at 196,608 cells, and this factor increases as the error level decreases.

3.1.3. Problem 3: Reconstruction of a spherically symmetric function on the cubed sphere

The distribution of errors on cubed-sphere grids is not expected to be uniform at any given radius, for several reasons: grid lines have kinks at the boundaries of the six cubed-sphere grid sectors, cells are smaller and more deformed near sector

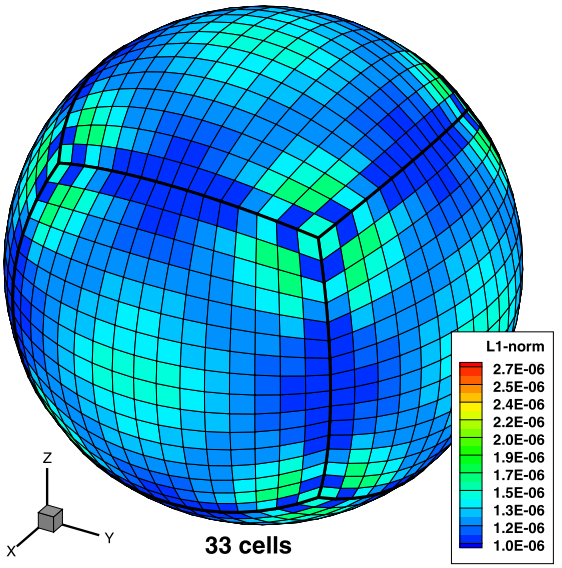

(a) L_1 error norm for the stencil with 33 cells.

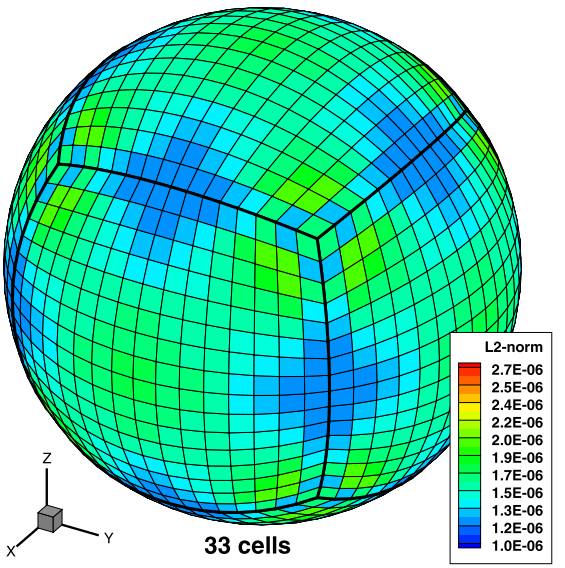

(b) L_2 error norm for the stencil with 33 cells.

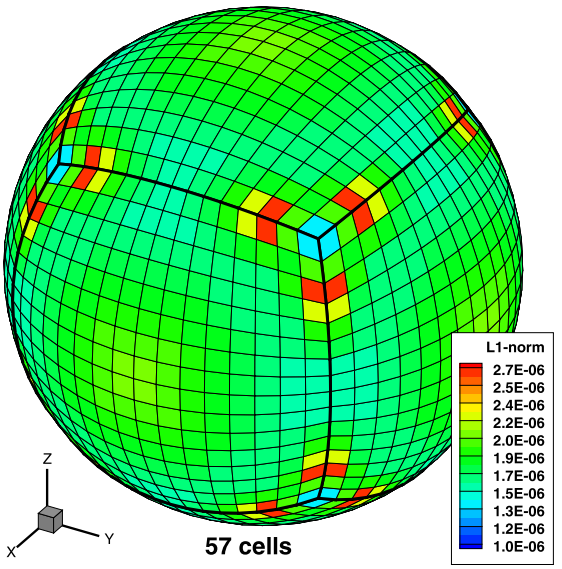

(c) L_1 error norm for the stencil with 57 cells.

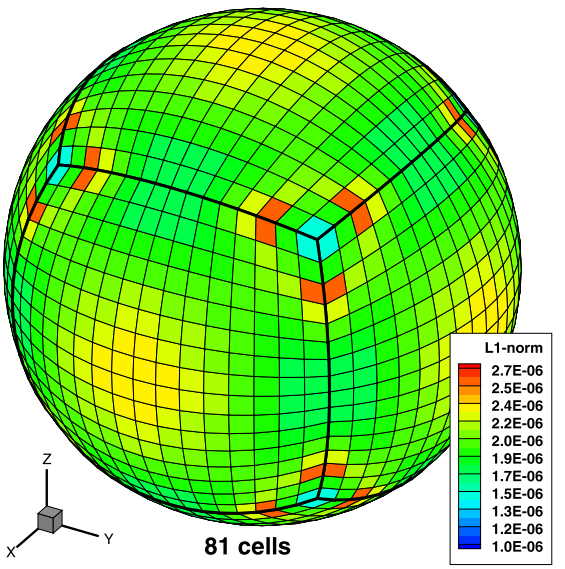

(d) L_1 error norm for the stencil with 81 cells.

Fig. 10. Problem 3: Error norms of cubic reconstruction at $R = 2.6$ for $f(R) = R^{-2.5}$.

boundaries and corners, and the grid topology is unstructured (with varying stencil size) at the sector corners. It is therefore of interest to investigate the error distribution in reconstructions obtained by the proposed fourth-order CENO procedure. In the previous two test problems we have established that fourth-order convergence is obtained overall (including near sector corners where CENO stencil sizes are reduced), but it is also desirable that the magnitude of the error for any grid resolution is not excessively larger near sector boundaries and corners than in other grid regions, which we investigate in this section. We do this for the function $f(R) = R^{-2.5}$, which is spherically symmetric, such that any non-uniformity in the reconstruction error can only be attributed to non-symmetric variations in the grid and to the differences in the selection of the reconstruction stencil at sector corners of the cubed-sphere grid. We reconstruct the function $f(R) = R^{-2.5}$ on a spherical shell with inner and outer radii $R_i = 2$ and $R_o = 3.5$. The cubed-sphere grid used consists of six $16 \times 16 \times 16$ blocks and the error distribution is plotted in Fig. 10 on a sphere of radius $R = 2.6$ in the interior of the domain to avoid effects from the boundary condition implementation. Note that local L_1 and L_2 errors are plotted according to Eqs. (39) and (40), where the integrals are taken over one cell and are normalized by the cell volume. We find a fairly symmetric error distribution, indicating that the least-squares based K-exact reconstruction succeeds in providing almost uniform accuracy, including in irregular cells near sector boundaries and corners.

It is interesting to note in Fig. 10(a) that, on the gnomonic cubed-sphere grid we use, the error is actually the smallest near the sector boundaries (away from the corners), and this is likely due to the smaller size of the cells there (and the slope discontinuities in the grid lines are small at those locations). The cells near the centre of the sectors are somewhat larger and the error is somewhat larger there. Not unexpectedly the error is largest near the sector corners, where grid lines have significant slope discontinuities and cells are deformed, and reconstruction stencils have reduced size. Nevertheless, the overall difference in error is within a factor of about two.

Comparing stencils with 33, 57 and 81 cells, it is interesting to note that the smaller stencil (33 cells) features a maximum L_1 error that is about 50% smaller than for the larger stencils. This may seem counter-intuitive at first, but can be explained by observing that we use each stencil with a fourth-order accurate reconstruction, so the accuracy of the numerical results depends on the accuracy of the polynomial coefficients that are obtained by the least-squares fitting process. Fig. 10 indicates that polynomial reconstruction is more accurate for the smaller-size stencils, which can be expected since they use information that is more local to the cell for which the reconstruction is computed, whereas the larger stencils also use more distant information. We conclude that the size-33 stencil is the best choice when using fourth-order reconstruction, while larger stencils may be used when polynomials of degree higher than four are used for reconstruction (in which case 33 cells do not provide enough conditions to constrain the polynomial coefficients).

3.2. Three-dimensional flow problems in a rectangular box

We now demonstrate the performance of the fourth-order CENO scheme on two standard test problems for MHD flows in rectangular boxes.

3.2.1. Problem 4: Magnetostatic MHD problem

We consider the 3D magnetohydrostatic test case proposed by Warburton and Karniadakis [59]. In this problem the fluid is static (i.e., has zero velocity) and the magnetic field is irrotational. The domain of the problem is the rectangular box defined by $0 < x < 1$ and $-1 < y, z < 1$. The 3D analytical solution is given by

$$\begin{aligned}
&\rho = 1,\\
&\vec{V} = \vec{0},\\
&B_x = \big(\cos\big(\pi(y+1)\big) - \cos(\pi z)\big) f(x),\\
&B_y = \cos(\pi z) f(y) + \sin\big(\pi(y+1)\big) f(x),\\
&B_z = \sin(\pi z)\big(f(y) - f(x)\big),\\
&p = 5(\gamma - 1),\\
&\psi = 0,
\end{aligned} \tag{42}$$

where $f(u) = e^{-\pi(u+1)}$ and $\gamma = 5/3$. Following [59], this test has been performed as an initial value problem with the exact solution used to provide the initial condition and the values for the Dirichlet boundary conditions for all domain boundaries. Numerical simulations for this problem have been performed with the fourth-order CENO scheme on a series of Cartesian meshes ranging in size from $4 \times 8 \times 8 = 256$ to $128 \times 256 \times 256 = 8{,}388{,}608$ cells, until the solution reached the steady state. We use the local Lax–Friedrichs (LF) flux function.

Fig. 11(a) shows the magnitude of the magnetic field vector obtained using the fourth-order ($K = 3$) CENO-GLM scheme on a Cartesian mesh with one $8 \times 16 \times 16$ block. The L_1, L_2, and L_∞ norms of the error in the x-component of the magnetic field, B_x, are given in Fig. 11(b). We use $\mathcal{S}_\mathcal{C} = 1500$ for the smoothness indicator cut-off value for the CENO results, and also include error plots for K-exact reconstruction without switching to the limited second-order scheme. Error measurements in the other components of the magnetic field, B_y and B_z, behave similarly. The results of the convergence study in Fig. 11(b) clearly show that the fourth-order theoretical accuracy is achieved by the numerical scheme in all error norms in the asymptotic limit, thereby demonstrating the high-order accuracy of the proposed CENO-GLM formulation for ideal MHD simulations.

3.2.2. Problem 5: 3D rotated MHD Shu–Osher problem

To demonstrate the overall robustness and advantages of our fourth-order scheme in regions of rapid smooth variation near a discontinuity, we consider a 3D rotated version of the shock tube problem which was proposed by Shu and Osher [86], and extended to MHD in [24]. As depicted in Fig. 12(a), a sinusoidal density perturbation is added downstream of a purely advecting superfast shock wave (with $\gamma = 5/3$). The interaction of the shock wave with the sinusoidal part of the density field gives rise to fast oscillations and complex flow features downstream of the shock. The Shu–Osher shock tube problem provides an excellent testbed to highlight the benefits of the improved accuracy of high-order numerical schemes, while at the same time the presence of the shock puts the robustness and stability of the schemes to test.

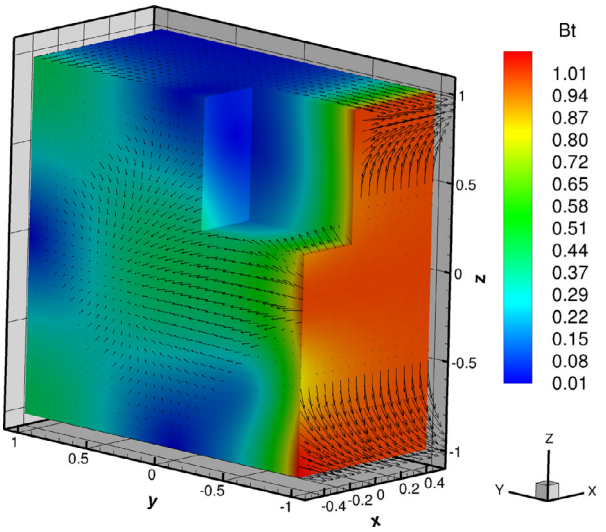

(a) Magnetic vector field and magnitude. The top-front corner has been removed to allow the visualization of the interior solution.

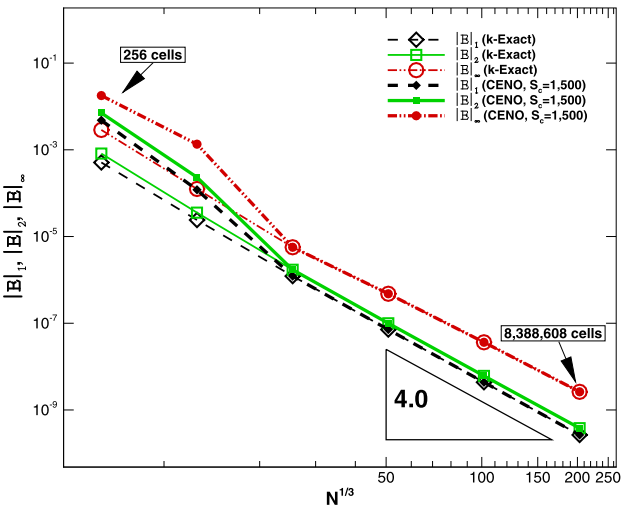

(b) Error norms.

Fig. 11. Problem 4: (a) Fourth-order ($K = 3$) CENO-GLM results for the magnetic field magnitude on a Cartesian grid with 2048 cells; and (b) L_1, L_2, and L_∞ norms of the error in the x-component of the magnetic field, B_x, as a function of mesh resolution.

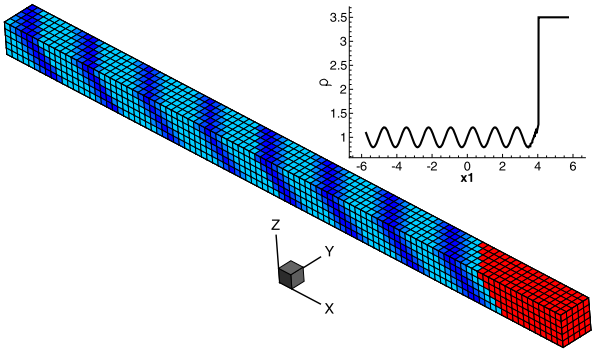

(a) Initial density profile as a function of the non-rotated coordinate x_1 and as a three-dimensional representation in the rotated frame.

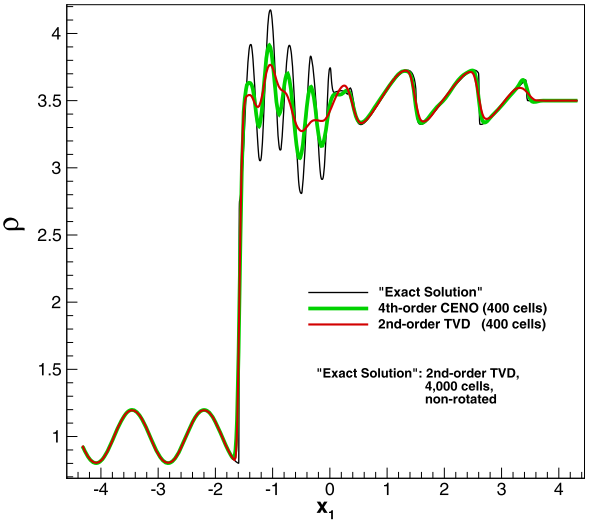

(b) Final density along the x_1 axis.

Fig. 12. Problem 5: (a) Initial density profile in the MHD Shu–Osher problem, rotated with $\varphi = 45°$ in the xy plane and $\beta = 35.2644°$ in the xz plane; and (b) comparison of the density solution at $t = 0.6906$. As can be observed from the figure, the fourth-order method produces results that are much closer to the non-rotated reference result in the highly oscillatory region, illustrating the benefits of high-order accuracy.

The unperturbed initial conditions (in terms of vector field components along and perpendicular to the x_1 direction of the unrotated shock tube problem) are given in [24]:

$$(\rho, v_\perp, v_\parallel, v_z, B_\perp, B_\parallel, B_z, p, \psi) = \begin{cases} (1, 0, 0, 0, 1, 1, 0, 1, 0) & \text{for } x_1 < 4, \\ (3.5, 5.8846, 1.1198, 0, 1, 3.6359, 0, 42.0267, 0) & \text{for } x_1 > 4, \end{cases} \tag{43}$$

and a sinusoidal density perturbation is added to the downstream part of the density field:

$$\rho_l = 1 + 0.2\sin(5x_1), \qquad \rho_r = 3.5. \tag{44}$$

The one-dimensional shock tube problem (along the x_1-axis) is rotated to the $(1, 1, 1)$ diagonal direction in the x, y, z coordinate system, resulting in angles of $\varphi = 45°$ in the xy plane and $\beta = 35.2644°$ in the xz plane. The comparison of the

density profiles between the fourth-order and second-order methods is shown in Fig. 12(b), where 400 cells are used in the x direction with a box length of 15. The grid size is the same in all three spatial directions. We use $\mathcal{S}_C = 180$ for the smoothness indicator cut-off value (because that produces a result with low diffusion), use the local Lax-Friedrichs (LF) flux function, and choose $\epsilon_A = \epsilon_R = 2 \times 10^{-5}$. It is clear that, for the same number of points, the fourth-order CENO method captures the small-scale flow features much better than the second-order method. At the same time, the high-order CENO approach is stable at the shock and does not produce spurious oscillations.

3.3. Three-dimensional flow problems on cubed-sphere grids

We finally describe extensive test results for three model problems on cubed-sphere grids, demonstrating high-order accuracy on the cubed-sphere grid with nonplanar cell surfaces including uniform fourth-order convergence at sector boundaries and corners of the cubed-sphere grid. We consider cubed-sphere grids with five, six and seven root blocks, and discuss performance gains obtained with the high-order method as compared to a second-order scheme.

3.3.1. Problem 6: MHD manufactured solution

To assess the accuracy of the high-order CENO finite-volume scheme on cubed-sphere grids, we present grid convergence studies for a 3D steady-state axi-symmetric solution of a MHD plasma on a spherical shell domain flowing outward at superfast speeds. This problem was proposed in [31] and represents an adequate 3D test case for solution accuracy on the cubed sphere because the 3D cubed-sphere grid is not axi-symmetric. In this test problem, the exact solution is specified using the primitive variables as

$$\mathbf{W}(x, y, z) = \left[R^{-\frac{5}{2}}, \quad \frac{x}{\sqrt{R}}, \quad \frac{y}{\sqrt{R}}, \quad \frac{z}{\sqrt{R}} + \kappa R^{\frac{5}{2}}, \quad \frac{x}{R^3}, \quad \frac{y}{R^3}, \quad \frac{z}{R^3} + \kappa, \quad R^{-\frac{5}{2}}\right]^{\mathrm{T}}. \tag{45}$$

This is an exact solution of ideal MHD if a source term is added to the equations given by

$$\mathbf{Q} = \begin{bmatrix} 0, \\ \frac{1}{2} x R^{-\frac{5}{2}}(R^{-1} - 5R^{-2} - \kappa z), \\ \frac{1}{2} y R^{-\frac{5}{2}}(R^{-1} - 5R^{-2} - \kappa z), \\ \frac{1}{2} z R^{-\frac{5}{2}}(R^{-1} - 5R^{-2} - \kappa z) + \frac{5}{2} R^{-\frac{1}{2}} \kappa (1 + \kappa R z) + \kappa R^{-\frac{1}{2}}, \\ \vec{0}, \\ \frac{1}{2} R^{-2} + \kappa z (3.5 R^{-1} + 2\kappa z) + \frac{(\kappa R)^2}{2}(7 + 5\kappa R z) \end{bmatrix}. \tag{46}$$

As suggested in [31], the perturbation parameter is taken as $\kappa = 0.017$ such that the solution has significant latitudinal variation yet the flow remains supersonic in the whole domain. Note that in this flow the magnetic field is irrotational and aligned everywhere with the velocity. The computational domain used for this convergence study is defined by inner and outer spheres of radius $R_i = 2$ and $R_o = 3.5$, respectively. To achieve high-order accuracy for this problem, it is necessary to provide a high-order approximation to the average source term, $(\overline{\mathbf{Q}})_{ijk}$, in the numerical residual (Eq. (6)) by integrating the analytical expression of the source term (Eq. (46)) with high-order accuracy.

Moreover, high-order boundary conditions must be imposed, which we achieved by specifying both the inflow and outflow boundary conditions using high-order accurate integration of the exact solution in ghost cells to determine the ghost cell averages. These ghost cell averages are used to reconstruct the solution in the ghost cells, and the reconstructed values at the domain boundaries are used together with an interior solution state to solve a local Riemann problem at domain boundaries, in the same way as at interior cell boundaries. We use $\mathcal{S}_C = 1500$ for the smoothness indicator cut-off value, use the local Lax–Friedrichs (LF) flux function, and use $\gamma = 7/5$.

Fig. 13 depicts the error norms in density as a function of grid size. It can be seen that uniform fourth-order convergence is obtained with the fourth-order scheme on the cubed-sphere, including near sector corners and boundaries, despite the grid irregularities and reduced stencil sizes there. This validates our approach to obtain high-order accuracy using overdetermined K-exact polynomial reconstruction, trilinear representations of nonplanar cell surfaces, and a rotation mechanism to build consistent stencils near sector boundaries and corners. Fig. 13(a) shows that our fourth-order method (with stencil size 33) converges with order four for all error norms, and, when compared with the second-order scheme of [30,31], much higher accuracy is obtained for a given number of grid cells (as much as four orders of magnitude lower). The number of computational cells required to achieve a target discretization error is up to 500 smaller for the fourth-order method. These results were obtained on a series of grids ranging in size from six $8 \times 8 \times 10$ initial blocks to 3072 $32 \times 32 \times 40$ cubed-sphere blocks, which corresponds to 3840 and 125,829,120 total cells, respectively, using up to 1536 CPU cores. As the mesh is refined, the slopes of the L_1, L_2, and L_∞ norms approach in the asymptotic limit -2.06, -2.10 and -2.05, respectively, for the second-order scheme, and -4.04, -4.05 and -4.019 for the fourth-order ($K = 3$) scheme. Fig. 13(b) compares error norms using base stencils with 33, 57 and 81 cells. As expected, all three choices lead to convergence with fourth-order accuracy. Consistent with what we observed in Fig. 10, the larger stencils lead to somewhat larger errors.

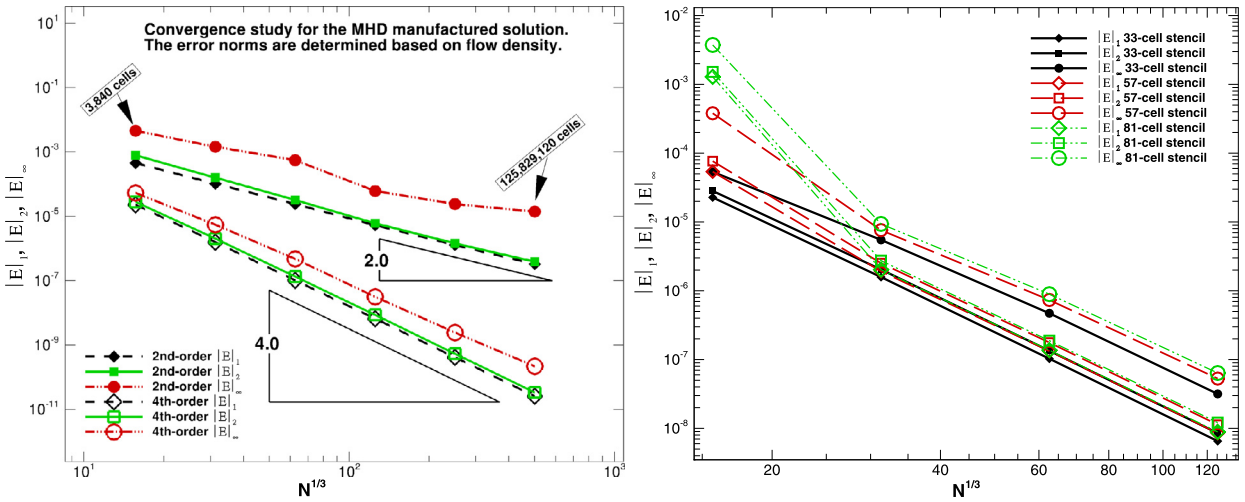

(a) Error norms for second-order scheme versus fourth-order scheme (stencil of size 33).

(b) Error norms for fourth-order scheme with various stencil sizes.

Fig. 13. Problem 6: Comparisons of L_1, L_2, and L_∞ error norms in solution density for the MHD manufactured solution problem, for the fourth-order CENO method and a limited second-order scheme, and for various stencil sizes.

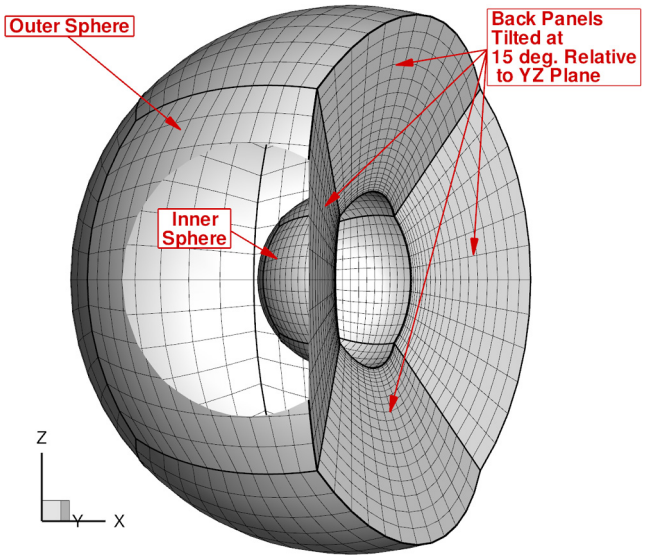

Fig. 14. Problem 7: Computational domain, discretized with a cubed-sphere mesh consisting of five root blocks.

3.3.2. *Problem 7: Euler bow shock*

Robustness of the fourth-order CENO method in the presence of 3D discontinuous flow features is next demonstrated using a 3D Euler bow shock flow around a sphere. We are mainly interested in testing robustness at the bow shock, so we simulate only the domain in front of the sphere by considering a cubed-sphere grid with five root blocks as indicated in Fig. 14, which illustrates the versatility of our multi-block cubed-sphere approach. The inflow conditions for the bow shock are $\rho = 1$, $p = 1.49492$, and $v_x = 2$, for a Mach number of approximately 2.825. The simulation domain is defined by two concentric spheres with inner and outer radius $R_i = 1$ and $R_o = 8$, and the four back panels of the 'missing sector' are tilted at 15° relative to the Cartesian yz plane. We use $\mathcal{S}_\mathcal{C} = 1500$ for the smoothness indicator cut-off value, use the HLLE flux function, and choose $\gamma = 7/5$, $\theta = 1$ (see Eq. (15)) and $\epsilon_A = \epsilon_R = 10^{-6}$. A cubed-sphere grid is used with 320 $16 \times 16 \times 32$ self-similar solution blocks with a total of 2,621,440 computational cells, using 320 CPU cores. Fig. 15 shows Mach number plots indicating that the fourth-order CENO scheme succeeds in producing flow solutions that preserve monotonicity at shocks on cubed-sphere grids with general hexahedral elements. Statistics collected at the end of the simulation revealed that about 7% of all cells contain non-smooth data in all solution variables (i.e., requiring limited linear reconstruction for all flow unknowns). An additional 4% of all cells used linear reconstruction for at least one but not all solution variables.

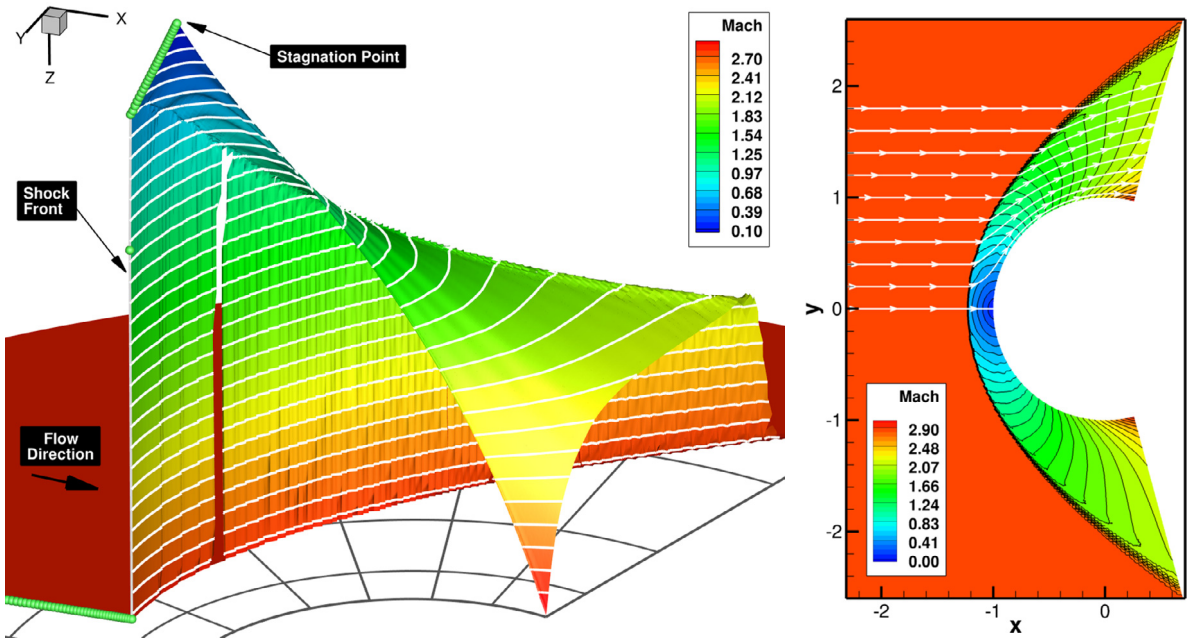

Fig. 15. Problem 7: Hydrodynamic bow shock over a sphere. (a) Three-dimensional representation of the Mach number distribution in the *xy* plane (the green dots mark the solution at cell centroids along the stagnation line passing through the tip of the sphere); and (b) planar representation of the Mach number distribution in the *xy* plane. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Thus, the limited linear reconstruction procedure was applied in just 11% of the cells in the computational domain. This includes cells at the 3D shock surface, and cells in the vicinity of the 3D curved domain boundary.

3.3.3. *Problem 8: MHD iso-density vortex: assessment of computational performance*

As our final test problem, we discuss the 3D iso-density MHD vortex advection problem from [57]. This is a smooth time-dependent 3D test problem with an exact solution, consisting of a magnetized vortex structure in force equilibrium that is advected by a uniform flow field.

The stationary 3D iso-density MHD vortex flow from [57] is given by

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

where $\mu = \kappa = 1/(2\pi)$, and $q = 1$. This solution is translated with a background velocity that equals $(1, 1, 2)$.

We first simulate this flow problem on a periodic Cartesian box with x, y, and z varying between $[-5, 5]$. Initially the vortex centre is positioned at $(0, 0, 0)$ and the simulation runs until time $t = 10$. We use $\mathcal{S}_C = 1500$ for the smoothness indicator cut-off value, use the HLLE flux function, and use $\gamma = 5/3$. Fig. 16(a) depicts the flow solution on the Cartesian grid, and Fig. 16(b) demonstrates fourth-order convergence. The initial grid consists of 64 blocks with $8 \times 8 \times 8$ cells, corresponding to a total of 32,768 cells. The largest high-order simulation used 512 blocks with $8 \times 8 \times 8$ cells, corresponding to a total of 16,777,216 cells, while the largest low-order simulation used 4096 blocks with $64 \times 64 \times 64$ cells corresponding to 1,073,741,824 cells, on 4096 CPU cores. Fig. 16(b) also illustrates the computational gains that can be made by using a fourth-order accurate numerical method as compared to a second-order method: for the L_1 error level indicated in the figure by the horizontal line, the second-order method requires 1,073,741,824 computational cells, while the fourth-order method would require 1,507,188 cells (estimated from the graph). (The fourth-order result nearest to the intersection of the fourth-order convergence curve and the horizontal line required 2,097,152 cells.) In terms of the total execution time, the second-order method requires a total execution time (in core hours) that is greater by a factor of about 700 than the fourth-order method to achieve the same error. This is confirmed in Fig. 16(c): in a plot of error versus total compute time, the fourth-order curves drop much faster than the second-order curves. The cross-over points between the second-order

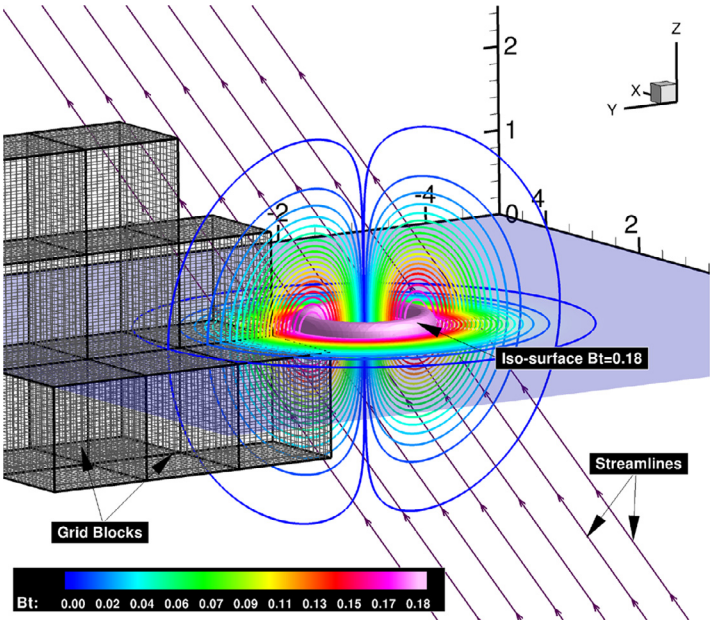

(a) Flow solution on Cartesian grid (selected grid blocks shown). (Total magnetic field contours and streamlines.)

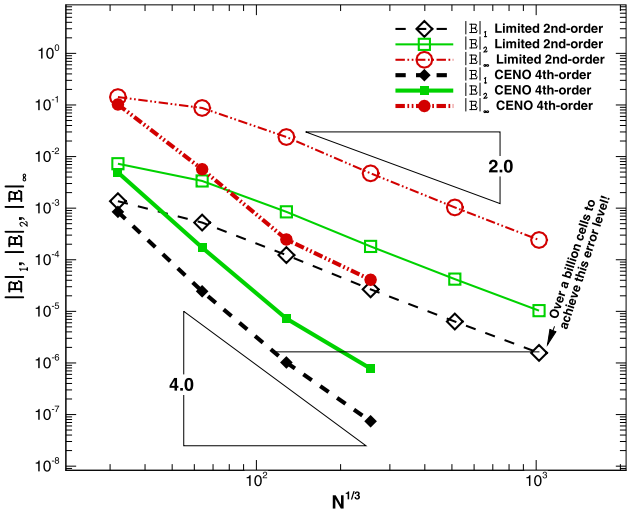

(b) Error norms versus grid size.

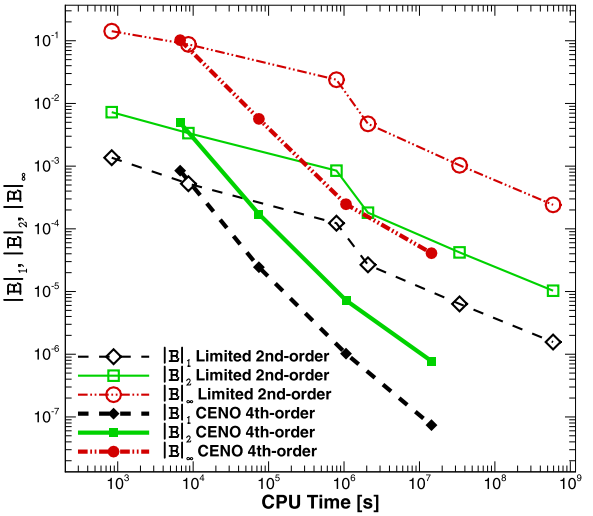

(c) Error norms versus total compute time.

Fig. 16. Problem 8: Iso-density vortex problem solved on a Cartesian grid in a periodic box.

and fourth-order curves occur early, and, beyond the cross-over points, decreasing target error levels correspond to increasingly higher efficiency of the fourth-order method compared to the second-order method. In other words, the fourth-order method starts to be more efficient than second-order already for modest target errors, and becomes more efficient by very large factors (easily more than three orders of magnitude in this test) as the required accuracy increases. While in general efficiency gains will depend on the problem and the problem size, this example clearly demonstrates the potential efficiency gains that can be offered by our high-order discretization scheme.

We have assessed the computational cost of the main steps of the high-order algorithm based on two test cases, the hydrodynamic bow shock flow and the convection of the iso-density magnetized vortex, both on the cubed-sphere grid. This analysis showed that the high-order explicit CENO scheme spends about 94% of the total CPU time in computing the high-order spatial residual; the remaining 6% is spent on the solution update and other operations like exchanging ghost cells and synchronizing the global time step. For computing the high-order spatial residual, about 69% of the total time is dedicated to performing the high-order CENO solution reconstruction. The remaining 25% consists of polynomial evaluations at the interfaces, evaluation of the Riemann fluxes there, and geometry evaluation. The high-order CENO solution reconstruction procedure (69% of the total time) consists in performing the least-squares reconstruction (44% of the total

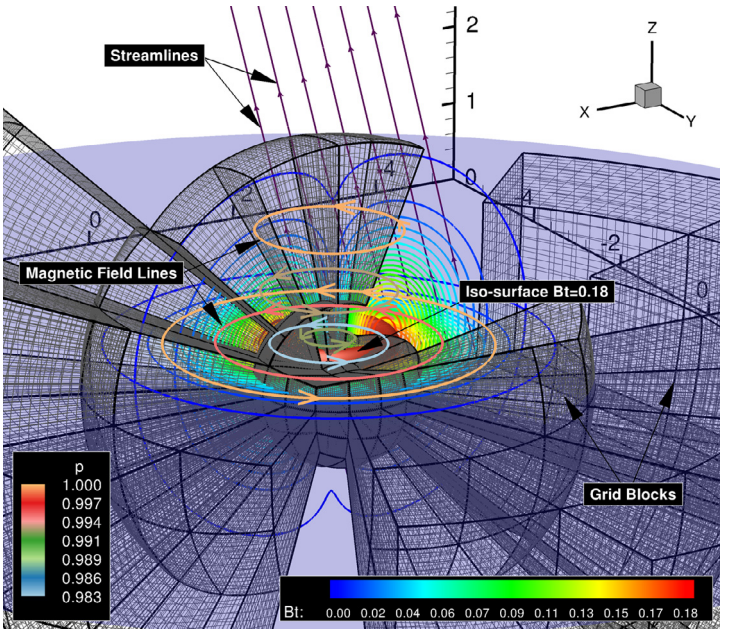

(a) Flow solution on cubed-sphere grid (selected grid blocks shown) at simulation time t=1.57812. (Total magnetic field contours, streamlines, and magnetic field lines coloured by pressure values.)

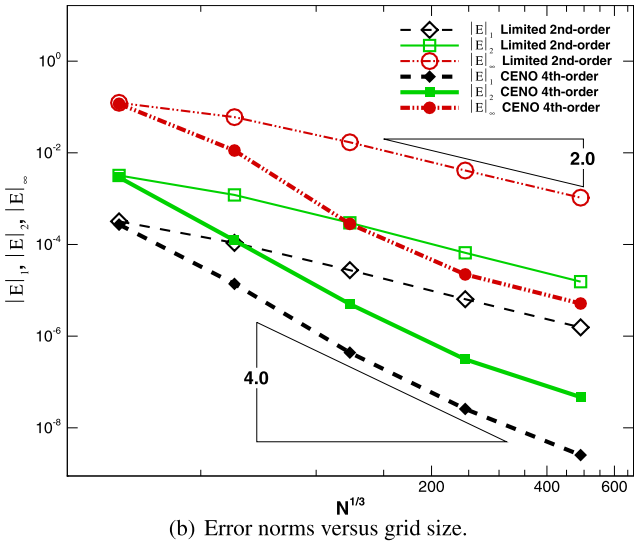

(b) Error norms versus grid size.

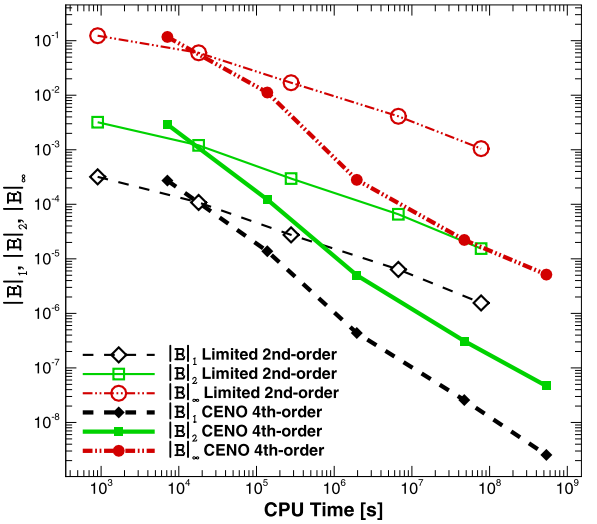

(c) Error norms versus total compute time.

Fig. 17. Problem 8: Iso-density vortex problem solved on a cubed-sphere grid with an additional seventh root block filling the interior sphere.

simulation cost), and the computation and analysis of smoothness indicators and the enforcement of solution monotonicity in cells deemed as non-smooth (25% of the total simulation cost). The matrix-vector multiplications in the least-squares reconstruction represent about 36% of the total simulation cost. It has also been determined that forming the reconstruction stencil amounts to less than 1% of the total simulation cost. Note, however, that these values can vary depending on the particular flow problem under study.

As a final demonstration of the capabilities of our fourth-order CENO method on the cubed-sphere grid, we report on a simulation of 3D iso-density MHD vortex advection on a cubed-sphere grid with an additional root block filling the interior of the inner sphere (as in Fig. 6(b)). The radius of the outer and inner spheres are $R = 9$ and $R = 1$, respectively, and the inner sphere is filled by a seventh root block with Cartesian topology. In this problem the 3D vortex is initially located outside the inner sphere (centred at $(-2, -2.5, -3)$), and is then translated with velocity $(1, 1, 2)$ over the cubed-sphere grid, passing through the inner sphere. The simulation is performed until $t = 3$, when the centre of the vortex is located at $(1, 0.5, 3)$. Fig. 17(a) shows the flow solution on the cubed-sphere grid at $t = 1.57812$, when the vortex passes through a corner of the seventh root block. The initial grid consists of 7 blocks with $16 \times 16 \times 16$ cells, corresponding to a total

of 28,672 cells The final grid uses 3584 blocks of $32 \times 32 \times 32$ cells corresponding to a total of 117,440,512 cells on 1792 CPU cores. Fig. 17(b) demonstrates that this simulation achieves fourth-order accuracy for a time-dependent flow on the cubed-sphere grid, and as such confirms the validity of our approaches for achieving high-order accuracy on grids with nonplanar cell surfaces and using a rotation mechanism to determine degenerate stencils at block boundaries and corners with unstructured grid topology. Fig. 17(c) confirms for the cubed-sphere grid the potential efficiency gains of the fourth-order method that were observed in Fig. 16(c) for a Cartesian grid. The large performance gains of the fourth-order method compared to the second-order scheme indicate that our approach of trilinear geometry representations and general rotation-based stencil modification mechanisms at cubed-sphere grid sector edges results in an efficient high-order method on the cubed-sphere grid, providing high-order accuracy as demonstrated in Fig. 17(b).

4. Conclusions

We have presented a new fourth-order accurate finite-volume scheme for conservation laws on 3D cubed-sphere grids. Our approach is based on the CENO finite-volume method for compressible fluid flows that was introduced for 2D gas dynamics and MHD in [19–24]. In this paper we have extended the CENO approach to 3D on general hexahedral grids, with particular attention to handling hexahedral cells with nonplanar surfaces using a trilinear surface representation. The K-exact multidimensional polynomial reconstruction used in CENO employs overdetermined stencils with a constrained least-squares solution procedure, and as such is ideally suited to handle with high-order accuracy the unstructured grid topology that occurs at boundaries and corners of the six root blocks of the cubed sphere, where stencil sizes vary and grid lines have slope discontinuities. However, an automatic procedure is required to select appropriate smaller stencils at those degenerate block boundaries, and we have proposed a general rotation-based mechanism to achieve this goal. Note that our approach facilitates high-order CENO simulations on other complex multi-block geometries that may require unstructured block connectivity, e.g., complex cylindrical domains for internal flow simulation.

There has been extensive work in recent years on high-order accurate discretizations for cubed-sphere grids, in particular in the context of high-order methods for flows on 2D spherical surfaces with curvilinear coordinates that are normally combined with separate discretizations in the vertical direction for applications in weather and climate simulation [6–16]. However, we describe the first 3D cubed-sphere grid finite volume discretization for conservation laws that is uniformly high-order accurate in all three grid directions. Our approach also stands out for its flexibility in allowing additional root blocks in the grid, for example, to mesh the interior of the inner sphere with an additional root block. More generally, we have successfully extended the CENO methodology to 3D conservation law flows on general hexahedral grids, demonstrating the accuracy, robustness and versatility of the high-order CENO approach in 3D. In terms of MHD flows, we extend the 2D GLM-based fourth-order MHD solver from [24] to 3D flows on complex cubed-sphere grid domains. In extensive numerical tests we have demonstrated fourth-order convergence of the 3D CENO method and robustness against spurious oscillations at 3D shocks. We have shown that high-order simulations can be much more efficient than second-order methods for the same error level, in terms of the number of computational cells and the execution time.

Future work will include efficient high-order accurate interpolation and restriction for dynamic solution-adaptive refinement and coarsening using the 3D CENO method, and the study of efficient linear and nonlinear solvers for implicit time integration using our high-order discretizations. Applications of the 3D CENO method on cubed-sphere grids to space physics problems are currently being pursued. For example, the interaction of the solar wind with magnetic anomalies at the surfaces of the Moon and Mars is a problem of interest.

Acknowledgements

This work was supported by the Canadian Space Agency and by NSERC of Canada. Computations were performed on the GPC supercomputer at the SciNet HPC Consortium and on the facilities of the Shared Hierarchical Academic Research Computing Network (SHARCNET: www.sharcnet.ca). The SciNet and SHARCNET consortia are funded by: the Canada Foundation for Innovation under the auspices of Compute/Calcul Canada; the Government of Ontario; Ontario Research Fund – Research Excellence; and the member institutions of the consortia.

References

[1] R. Sadourny, Conservative finite-difference approximations of the primitive equations on quasi-uniform spherical grids, Mon. Weather Rev. 100 (2) (1972) 136–144.
[2] C. Ronchi, R. Iacono, P.S. Paolucci, The “Cubed Sphere”: a new method for the solution of partial differential equations in spherical geometry, J. Comput. Phys. 124 (1996) 93–114.
[3] A. Adcroft, J.-M. Campin, C. Hill, J. Marhall, Implementation of an atmosphere–ocean general circulation model on the expanded spherical cube, Mon. Weather Rev. 132 (2004) 2845–2863.
[4] W.M. Putman, S.-J. Lin, Finite-volume transport on various cubed-sphere grids, J. Comput. Phys. 227 (2007) 55–78.
[5] A. St-Cyr, C. Jablonowski, J.M. Dennis, H.M. Tufo, S.J. Thomas, A comparison of two shallow-water models with nonconforming adaptive grids, Mon. Weather Rev. 136 (2008) 1898–1922.
[6] P.A. Ullrich, C. Jablonowski, B. van Leer, High-order finite-volume methods for the shallow-water equations on the sphere, J. Comput. Phys. 229 (17) (2010) 6104–6134.
[7] C. Chen, F. Xiao, Shallow water model on cubed-sphere by multi-moment finite volume method, J. Comput. Phys. 227 (2008) 5019–5044.

[8] C.G. Chen, F. Xiao, X.L. Li, Y. Yang, A multi-moment transport model on cubed-sphere grid, Int. J. Numer. Methods Fluids 67 (2011) 1993–2014.
[9] P.A. Ullrich, C. Jablonowski, MCore: a non-hydrostatic atmospheric dynamical core utilizing high-order finite-volume methods, J. Comput. Phys. 231 (15) (2012) 5078–5108.
[10] L. Bao, R.D. Nair, H.M. Tufo, A mass and momentum flux-form high-order discontinuous Galerkin shallow water model on the cubed-sphere, J. Comput. Phys. 271 (2014) 224–243.
[11] L.M. Harris, S.-J. Lin, A two-way nested global-regional dynamical core on the cubed-sphere grid, Mon. Weather Rev. 141 (1) (2013) 283–306.
[12] J.-P. Croisille, Hermitian compact interpolation on the cubed-sphere grid, J. Sci. Comput. 57 (1) (2013) 193–212.
[13] C. Erath, P.H. Lauritzen, H.M. Tufo, On mass conservation in high-order high-resolution rigorous remapping schemes on the sphere, Mon. Weather Rev. 141 (6) (2013) 2128–2133.
[14] W. Guo, R.D. Nair, J.-M. Qiu, A conservative semi-Lagrangian discontinuous Galerkin scheme on the cubed sphere, Mon. Weather Rev. 142 (1) (2014) 457–475.
[15] H. Yang, C. Yang, X.-C. Cai, Parallel domain decomposition methods with mixed order discretization for fully implicit solution of tracer transport problems on the cubed-sphere, J. Sci. Comput. (2014) 1–23.
[16] C. Erath, R.D. Nair, A conservative multi-tracer transport scheme for spectral-element spherical grids, J. Comput. Phys. 256 (2014) 118–134.
[17] A.V. Koldoba, M.M. Romanova, G.V. Ustyugova, R.V.E. Lovelace, Three-dimensional magnetohydrodynamic simulations of accretion to an inclined rotator: the "Cubed Sphere" method, Astrophys. J. 576 (2002) L53–L56.
[18] P.C. Fragile, C.C. Lindner, P. Anninos, J.D. Salmonson, Application of the cubed-sphere grid to tilted black hole accretion disks, Astrophys. J. 691 (2009) 482–494.
[19] L. Ivan, C.P.T. Groth, High-order central ENO finite-volume scheme with adaptive mesh refinement, AIAA, 2007, Paper 2007-4323.
[20] L. Ivan, Development of high-order CENO finite-volume schemes with block-based adaptive mesh refinement, Ph.D. thesis, University of Toronto, 2011, http://hdl.handle.net/1807/29759.
[21] L. Ivan, C.P.T. Groth, High-order solution-adaptive central essentially non-oscillatory (CENO) method for viscous flows, AIAA, 2011, Paper 2011-0367.
[22] L. Ivan, C.P.T. Groth, High-order solution adaptive central essentially non-oscillatory (CENO) method for viscous flows, J. Comput. Phys. 257 (2013) 830–862.
[23] L. Ivan, C.P.T. Groth, High-order central ENO scheme with adaptive mesh refinement for hyperbolic conservation laws, Commun. Comput. Phys. (2013), submitted for publication.
[24] A. Susanto, L. Ivan, H. De Sterck, C. Groth, High-order central ENO finite-volume scheme for ideal MHD, J. Comput. Phys. 250 (2013) 141–164.
[25] L. Ivan, A. Susanto, H. De Sterck, C. Groth, High-order central ENO finite-volume scheme for MHD on three-dimensional cubed-sphere grids, in: Proceedings of the 7th International Conference on Computational Fluid Dynamics, Hawaii, USA, July 9–13, 2012, paper ICCFD7-3605.
[26] S.D. McDonald, M.R.J. Charest, C.P.T. Groth, High-order CENO finite-volume schemes for multi-block unstructured mesh, AIAA, 2011, paper 2011-3854.
[27] M.R.J. Charest, C.P.T. Groth, P.Q. Gauthier, High-order CENO finite-volume scheme for low-speed viscous flows on three-dimensional unstructured mesh, in: Proceedings of the 7th International Conference on Computational Fluid Dynamics, Hawaii, USA, July 9–13, 2012, paper ICCFD7-1002.
[28] M.R.J. Charest, C.P.T. Groth, A high-order central ENO finite-volume scheme for three-dimensional turbulent reactive flows on unstructured mesh, AIAA, 2013, paper 2013-2567.
[29] M.R.J. Charest, C.P.T. Groth, P.Q. Gauthier, A high-order central ENO finite-volume scheme for three-dimensional low-speed viscous flows on unstructured mesh, Commun. Comput. Phys. (2014), submitted for publication.
[30] L. Ivan, H. De Sterck, S.A. Northrup, C.P.T. Groth, Three-dimensional MHD on cubed-sphere grids: parallel solution-adaptive simulation framework, AIAA, 2011, paper 2011–3382.
[31] L. Ivan, H. De Sterck, S.A. Northrup, C.P.T. Groth, Multi-dimensional finite-volume scheme for hyperbolic conservation laws on three-dimensional solution-adaptive cubed-sphere grids, J. Comput. Phys. 255 (2013) 205–227.
[32] X. Gao, A parallel solution-adaptive method for turbulent non-premixed combusting flows, Ph.D. thesis, University of Toronto, 2008.
[33] X. Gao, C.P.T. Groth, A parallel solution-adaptive method for three-dimensional turbulent non-premixed combusting flows, J. Comput. Phys. 229 (2010) 3250–3275.
[34] C.P.T. Groth, S.A. Northrup, Parallel implicit adaptive mesh refinement scheme for body-fitted multi-block mesh, AIAA, 2005, paper 2005-5333.
[35] J.S. Sachdev, C.P.T. Groth, J.J. Gottlieb, A parallel solution-adaptive scheme for predicting multi-phase core flows in solid propellant rocket motors, Int. J. Comput. Fluid Dyn. 19 (2) (2005) 159–177.
[36] S.A. Northrup, C.P.T. Groth, Solution of laminar diffusion flames using a parallel adaptive mesh refinement algorithm, AIAA, 2005, paper 2005-0547.
[37] X. Gao, C.P.T. Groth, A parallel adaptive mesh refinement algorithm for predicting turbulent non-premixed combusting flows, Int. J. Comput. Fluid Dyn. 20 (5) (2006) 349–357.
[38] C.P.T. Groth, J.G. McDonald, Towards physically-realizable and hyperbolic moment closures for kinetic theory, Contin. Mech. Thermodyn. 21 (6) (2009) 467–493.
[39] X. Gao, S.A. Northrup, C.P.T. Groth, Parallel solution-adaptive method for two-dimensional non-premixed combusting flows, Prog. Comput. Fluid Dyn. 11 (2) (2011) 76–95.
[40] M.R.J. Charest, C.P.T. Groth, O.L. Gülder, A computational framework for predicting laminar reactive flows with soot formation, Combust. Theory Model. 14 (6) (2010) 793–825.
[41] M.R.J. Charest, H.I. Joo, C.P.T. Groth, O.L. Gülder, Experimental and numerical study of soot formation in laminar ethylene diffusion flames at elevated pressures from 10 to 35 atm, Proc. Combust. Inst. 33 (2011) 549–557.
[42] M.R.J. Charest, C.P.T. Groth, O.L. Gülder, Effects of gravity and pressure on laminar co-flow methane–air diffusion flames at pressures from 1 to 60 atmospheres, Combust. Flame 158 (5) (2011) 860–875.
[43] J.S. Sachdev, C.P.T. Groth, J.J. Gottlieb, Parallel AMR scheme for turbulent multi-phase rocket motor core flows, AIAA, 2005, paper 2005-5334.
[44] J. McDonald, C.P.T. Groth, Numerical modeling of micron-scale flows using the gaussian moment closure, AIAA, 2005, paper 2005-5035.
[45] M.R.J. Charest, C.P.T. Groth, O.L. Gülder, Solution of the equation of radiative transfer using a Newton–Krylov approach and adaptive mesh refinement, J. Comput. Phys. 231 (2012) 3023–3040.
[46] J.S. Sachdev, C.P.T. Groth, A mesh adjustment scheme for embedded boundaries, Commun. Comput. Phys. 2 (6) (2007) 1095–1124.
[47] Z.J. Zhang, C.P.T. Groth, Parallel high-order anisotropic block-based adaptive mesh refinement finite-volume scheme, AIAA, 2011, paper 2011-3695.
[48] M.J. Williamschen, C.P.T. Groth, Parallel anisotropic block-based adaptive mesh refinement algorithm for three-dimensional flows, AIAA, 2013, paper 2013-2442.
[49] L. Ivan, S.A. Northrup, C.P.T. Groth, H. De Sterck, Adaptive cubed-sphere simulation framework for space physics flows, in: Proceedings of the 18th Annual Conference of the CFD Society of Canada, London, Canada, May 18–19, 2010, 2010.
[50] O. Zienkiewicz, R. Taylor, The Finite Element Method, vol. 1, fifth edn., Butterworth–Heinemann, 2000.
[51] J.F. Thompson, Z.U.A. Warsi, C.W. Mastin, Numerical Grid Generation—Foundations and Applications, North-Holland, New York, 1985.
[52] C.A. Felippa, A compendium of FEM integration formulas for symbolic work, Eng. Comput. 21 (8) (2004) 867–890.
[53] P. Colella, M. Dorr, J. Hittinger, P. McCorquodale, D.F. Martin, High-order finite-volume methods on locally-structured grids, in: N.V. Pogorelov, et al. (Eds.), Numerical Modeling of Space Plasma Flows: ASTRONUM-2008, in: ASP Conf. Proc., vol. 406, ASP, San Francisco, 2009, pp. 207–216.

[54] D.A. Calhoun, C. Helzel, R.J. LeVeque, Logically rectangular finite volume grids and methods for circular and spherical domains, SIAM Rev. 50 (2008) 723–752.
[55] J.A. Rossmanith, A wave propagation method for hyperbolic systems on the sphere, J. Comput. Phys. 213 (2) (2006) 629–658.
[56] P.A. Ullrich, Atmospheric modeling with high-order finite-volume methods, Ph.D. thesis, University of Michigan, 2011.
[57] A. Mignone, P. Tzeferacos, G. Bodo, High-order conservative finite difference GLM-MHD schemes for cell-centered MHD, J. Comput. Phys. 229 (2010) 5896–5920.
[58] B. Cockburn, C.-W. Shu, The Runge–Kutta discontinuous Galerkin method for conservation laws V: multidimensional systems, J. Comput. Phys. 141 (2) (1998) 199–224.
[59] T.C. Warburton, G.E. Karniadakis, A discontinuous Galerkin method for the viscous MHD equations, J. Comput. Phys. 152 (1999) 608–641.
[60] H. Luo, Y. Xia, S. Li, R. Nourgaliev, C. Cai, A Hermite WENO reconstruction-based discontinuous Galerkin method for the Euler equations on tetrahedral grids, J. Comput. Phys. 231 (16) (2012) 5489–5503.
[61] J. Kleimann, A. Kopp, H. Fichtner, R. Grauer, K. Germaschewski, Three-dimensional MHD high-resolution computations with CWENO employing adaptive mesh refinement, Comput. Phys. Commun. 158 (2004) 47–56.
[62] A. Nejat, C.F. Ollivier-Gooch, A high-order accurate unstructured finite volume Newton–Krylov algorithm for inviscid compressible flows, J. Comput. Phys. 227 (2008) 2582–2609.
[63] D.S. Balsara, T. Rumpf, M. Dumbser, C.-D. Munz, Efficient, high accuracy ADER-WENO schemes for hydrodynamics and divergence-free magnetohydrodynamics, J. Comput. Phys. 228 (2009) 2480–2516.
[64] P. Colella, M. Dorr, J. Hittinger, D. Martin, High-order, finite-volume methods in mapped coordinates, J. Comput. Phys. (ISSN 0021-9991) 230 (8) (2011) 2952–2976, http://dx.doi.org/10.1016/j.jcp.2010.12.044, http://www.sciencedirect.com/science/article/pii/S0021999111000064.
[65] D.S. Balsara, Self-adjusting, positivity preserving high order schemes for hydrodynamics and magnetohydrodynamics, J. Comput. Phys. (ISSN 0021-9991) 231 (22) (2012) 7504–7517, http://dx.doi.org/10.1016/j.jcp.2012.01.032, http://www.sciencedirect.com/science/article/pii/S0021999112000629.
[66] M. Dumbser, D.S. Balsara, E.F. Toro, C.-D. Munz, A unified framework for the construction of one-step finite volume and discontinuous Galerkin schemes on unstructured meshes, J. Comput. Phys. 227 (18) (2008) 8209–8253.
[67] A. Dedner, F. Kemm, D. Kröner, C.-D. Munz, T. Schnitzer, M. Wesenberg, Hyperbolic divergence cleaning for the MHD equations, J. Comput. Phys. 175 (2) (2002) 645–673, http://dx.doi.org/10.1006/jcph.2001.6961.
[68] G. Tóth, B. van der Holst, I.V. Sokolov, D.L.D. Zeeuw, T.I. Gombosi, F. Fang, W.B. Manchester, X. Meng, D. Najib, K.G. Powell, Q.F. Stout, A. Glocer, Y.-J. Ma, M. Opher, Adaptive numerical algorithms in space weather modeling, J. Comput. Phys. 231 (3) (2012) 870–903, http://dx.doi.org/10.1016/j.jcp.2011.02.006, http://www.sciencedirect.com/science/article/pii/S002199911100088X.
[69] B. van der Holst, R. Keppens, Hybrid block-AMR in Cartesian and curvilinear coordinates: MHD applications, J. Comput. Phys. 226 (1) (2007) 925–946, http://dx.doi.org/10.1016/j.jcp.2007.05.007.
[70] R. Keppens, Z. Meliani, A.J. van Marle, P. Delmont, A. Vlasis, B. van der Holst, Parallel, grid-adaptive approaches for relativistic hydro and magnetohydrodynamics, J. Comput. Phys. 231 (3) (2012) 718–744.
[71] J. Raeder, D. Larson, W. Li, E.L. Kepko, T. Fuller-Rowell, Open GGCM simulations for the THEMIS mission, in: The THEMIS Mission, Springer, 2009, pp. 535–555.
[72] J. Lyon, J. Fedder, C. Mobarry, The Lyon–Fedder–Mobarry (LFM) global MHD magnetospheric simulation code, J. Atmos. Sol.-Terr. Phys. 66 (15) (2004) 1333–1350.
[73] B. van der Holst, I. Sokolov, X. Meng, M. Jin, W. Manchester IV, G. Tóth, T. Gombosi, Alfvén wave solar model (AWSoM): coronal heating, Astrophys. J. 782 (2) (2014) 81.
[74] R. Keppens, O. Porth, Scalar hyperbolic PDE simulations and coupling strategies, J. Comput. Appl. Math. 266 (2014) 87–101.
[75] O. Porth, C. Xia, T. Hendrix, S. Moschou, R. Keppens, MPI-AMRVAC for solar and astrophysics, Astrophys. J. Suppl. Ser. 214 (1) (2014) 4.
[76] A. Mignone, High-order conservative reconstruction schemes for finite volume methods in cylindrical and spherical coordinates, J. Comput. Phys. 270 (2014) 784–814.
[77] T.J. Barth, Recent developments in high order K-exact reconstruction on unstructured meshes, AIAA, 1993, paper 93-0668.
[78] D.J. Mavriplis, Revisiting the least-squares procedure for gradient reconstruction on unstructured meshes, AIAA, 2003, paper 2003-3986.
[79] C. Lawson, R. Hanson, Solving Least Squares Problems, Prentice-Hall, Inc., 1974.
[80] A. Harten, High resolution schemes for hyperbolic conservation laws, J. Comput. Phys. 49 (1983) 357–393.
[81] J.S. Park, S.-H. Yoon, C. Kim, Multi-dimensional limiting process for hyperbolic conservation laws on unstructured grids, J. Comput. Phys. 229 (2010) 788–812.
[82] V. Venkatakrishnan, On the accuracy of limiters and convergence to steady state solutions, AIAA, 1993, paper 93-0880.
[83] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes. The Art of Scientific Computing, third-edn., Cambridge University Press, New York, 2007.
[84] E.M. Harnett, R. Winglee, Two-dimensional MHD simulation of the solar wind interaction with magnetic field anomalies on the surface of the Moon, J. Geophys. Res. Space Phys. (1978–2012) 105 (A11) (2000) 24997–25007.
[85] W. Gander, W. Gautschi, Adaptive quadrature – revisited, BIT 40 (2000) 84–101.
[86] C.-W. Shu, S. Osher, Efficient implementation of essentially non-oscillatory shock-capturing schemes: II, J. Comput. Phys. 83 (1989) 32–78.