



High-Order Finite-Volume Scheme with Anisotropic Adaptive Mesh Refinement: Efficient Inexact Newton Method for Steady Three-Dimensional Flows

L. Freret*, C. P. T. Groth[†]

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

T. B. Nguyen[‡] H. De Sterck[§]

School of Mathematical Science, 9 Rainforest Walk, Monash University, VIC 3800, Australia

A high-order finite-volume method with anisotropic adaptive mesh refinement (AMR) is combined with a parallel inexact Newton method integration scheme and described for the solution of compressible fluid flows governed by Euler and Navier-Stokes equations on three-dimensional multi-block body-fitted hexahedral meshes. The proposed approach combines a family of robust and accurate high-order central essentially non-oscillatory (CENO) spatial discretization schemes with a scalable and efficient Newton-Krylov-Schwarz (NKS) algorithm and a block-based anisotropic AMR. The CENO scheme is based on a hybrid solution reconstruction procedure that provides high-order accuracy in smooth regions (even for smooth extrema) and non-oscillatory transitions at discontinuities. The implicit time stepping scheme is based on Newton's method where the set of linear systems are solved using the generalized minimal residual (GMRES) algorithm preconditioned by a domain-based additive Schwarz technique. The latter uses the domain decomposition provided by the block-based AMR scheme leading to a fully parallel implicit approach with an efficient scalability of the overall scheme. The anisotropic AMR method is based on a binary tree and hierarchical data structure to permit local anisotropic refinement of the grid in preferred directions as directed by appropriately specified physics-based refinement criteria. Application and numerical results will be discussed for several steady inviscid and viscous problems and the computational performance of the overall scheme is demonstrated for a range of fluid flows.

I. Introduction

One highly successful approach for reducing the computational cost of computational fluid dynamic (CFD) solution is to make use of solution-directed mesh adaptation where the underlying computational mesh automatically adapts according to the solution, adding mesh resolution only where required. Methods based on adaptive mesh refinement (AMR) have proved to be particularly effective for the solution of conservation equations on structured cartesian and body-fitted meshes and have been developed for a wide range of problems.¹⁻⁶ The non-uniform block-based AMR of Freret and Groth⁷ has been used for the solution of inviscid and viscous flows with an anisotropic mesh and a low-order finite-volume scheme. This block-based AMR approach is also well suited for use with high-order spatial schemes^{8,9} and has been coupled for this purpose to the solution of the ideal magnetohydrodynamics (MHD) equations with the high-order CENO finite volume scheme of Ivan et al.^{10,11} The latter uses a hybrid reconstruction approach based on a fixed central stencil. An unlimited high-order k -exact reconstruction is performed in the cells where the solution is well resolved while the scheme reverts to a low-order limited linear approach for cells with under-resolved/discontinuous solution content. Switching in the hybrid procedure is determined by a smoothness

*Post-doctoral fellow

[†]Professor, Senior Member AIAA.

[‡]Post-doctoral fellow

[§]Professor

indicator. The CENO high-order scheme has been successfully applied to a broad range of flows on uniform cartesian meshes including non-viscous flows,¹⁰ viscous flows,¹¹ Large-Eddy Simulation (LES) for turbulent premixed flames¹² and MHD problems.^{10,13} The efficiency of the CENO scheme has also been assessed on cubed-sphere meshes¹⁰ and extended to unstructured meshes for laminar viscous flows¹⁴ and turbulent reactive flows.¹⁵ For many applications, the high-order CENO scheme has been shown to be more efficient than a second-order method in terms of execution time and number of computational cells required to achieve a target discretization error.^{10,12}

For steady flow problems, the inherent disparate temporal scales of the governing conservation equations and the varying mesh spacing arising from the AMR strategy give rise to the need for solution of a coupled system of non-linear algebraic equations with a high degree of numerical stiffness. One common approach for the efficient solution of non-linear algebraic equations is Newton's method. In this study, an inexact Newton method is considered for the solution of the stiff non-linear equations arising from the high-order spatial discretization of the partial differential equations. In this method, the non-linear system is solved by first applying a Newton linearization, then using a linear solver to solve the resulting linear systems for each Newton iteration.^{1,16} The linear system arising from the linearization is both large and sparse and is solved by a right preconditioned GMRES method.¹⁷ A domain-based additive Schwarz preconditioning technique is used as the global preconditioner and incomplete lower-upper factorization with fill (ILU(2)) is used as the local preconditioner to improve the convergence rate. The Schwarz preconditioner provides the advantage of using the same domain decomposition procedure used by the block-based AMR scheme, making parallel implementation rather straightforward. The combination of the strengths of block-based adaptive mesh refinement along with an implicit Newton-Krylov-Schwarz (NKS) time stepping scheme to achieve a scalable parallel framework has been considered previously for lower-order spatial discretization of two-dimensional (2D) and three-dimensional (3D) steady and unsteady as well as non-reactive and reactive flows.^{7,16,18,19} It is reconsidered here for high-order finite-volume discretizations of 3D steady flows.

I.A. Scope

The aim of this study is to combine the high-order CENO finite-volume scheme of Ivan et al.^{10,11} and the anisotropic block-based AMR technique with an efficient and scalable inexact Newton method^{7,16,18,19} for computing steady-state solutions of the Euler and Navier-Stokes equations governing inviscid and viscous compressible gaseous flows. Furthermore, the benefits and computational performance of this proposed combination are demonstrated by considering numerical results for a range of flow problems. The remainder of the paper is organized as follows. The high-order CENO finite-volume scheme with anisotropic block-based AMR is described in Section II. The inexact Newton's method for the solution methodology to steady-state problems is presented in Section III. Qualitative and quantitative comparisons are then in Sections IV and V to demonstrate the potential of this combined AMR high-order discretization and implicit time stepping scheme. For the numerical results presented in this manuscript, numerical solution of the Euler and Navier-Stokes equations are considered.

II. High-Order Finite-Volume Scheme with Anisotropic AMR

The details of the parallel adaptive mesh refinement finite-volume spatial discretization method are now described when the numerical solution of the Euler and Navier-Stokes equations are considered. The CENO scheme using a high-order polynomial reconstruction approach with monotonicity enforcement is presented and its implementation within the parallel block-based AMR framework is also described.

II.A. Three-Dimensional Governing Equations

The conservation form of the governing equations for a compressible polytropic gas, either inviscid or viscous, can be written as

$$\frac{\partial \mathbf{U}}{\partial t} + \vec{\nabla} \cdot \vec{\mathbf{F}} = 0, \quad (1)$$

where \mathbf{U} is the vector of conserved solution variables. For a three dimensional Cartesian coordinate system, the Euler equations governing inviscid gaseous flows can be written as

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial \mathbf{F}}{\partial x} + \frac{\partial \mathbf{G}}{\partial y} + \frac{\partial \mathbf{H}}{\partial z} = 0, \quad (2)$$

with

$$\mathbf{U} = \begin{bmatrix} \rho \\ \rho u \\ \rho v \\ \rho w \\ \rho e \end{bmatrix}, \mathbf{F} = \begin{bmatrix} \rho u \\ \rho u^2 + p \\ \rho uv \\ \rho uw \\ \rho uh \end{bmatrix}, \mathbf{G} = \begin{bmatrix} \rho v \\ \rho v^2 + p \\ \rho vw \\ \rho v^2 + p \\ \rho vh \end{bmatrix}, \mathbf{H} = \begin{bmatrix} \rho w \\ \rho w^2 + p \\ \rho vw \\ \rho w^2 + p \\ \rho wh \end{bmatrix}, \quad (3)$$

where ρ is the gas density, u , v , w are the components of the velocity vector in the x , y , z directions, $e = p/(\rho(\gamma - 1)) + u^2/2$ is the specific total energy, p is the gas pressure and γ is the ratio of specific heats. The vectors \mathbf{F} , \mathbf{G} , \mathbf{H} , are the inviscid flux in the x , y , z directions, respectively such that $\vec{\mathbf{F}} = \vec{\mathbf{F}}_H = [\mathbf{F}, \mathbf{G}, \mathbf{H}]$ is the solution flux. The ideal gas equation of state $p = \rho RT$ is used to close the system, where T is the gas temperature and R is the ideal gas constant.

In the case of the Navier-Stokes equations governing viscous gaseous flows, the conservation equations can be expressed as

$$\frac{\partial \mathbf{U}}{\partial t} + \frac{\partial (\mathbf{F} - \mathbf{F}_v)}{\partial x} + \frac{\partial (\mathbf{G} - \mathbf{G}_v)}{\partial y} + \frac{\partial (\mathbf{H} - \mathbf{H}_v)}{\partial z} = 0, \quad (4)$$

where the viscous flux vector $\vec{\mathbf{F}}_E = [\mathbf{F}_v, \mathbf{G}_v, \mathbf{H}_v]$ in the x , y , and z directions respectively, is given by:

$$\mathbf{F}_v = \begin{bmatrix} 0 \\ \tau_{xx} \\ \tau_{xy} \\ \tau_{xz} \\ -q_x + u\tau_{xx} + v\tau_{xy} + w\tau_{xz} \end{bmatrix}, \mathbf{G}_v = \begin{bmatrix} 0 \\ \tau_{xy} \\ \tau_{yy} \\ \tau_{zy} \\ -q_y + u\tau_{xy} + v\tau_{yy} + w\tau_{yz} \end{bmatrix}, \mathbf{H}_v = \begin{bmatrix} 0 \\ \tau_{xz} \\ \tau_{yz} \\ \tau_{zz} \\ -q_z + u\tau_{xz} + v\tau_{yz} + w\tau_{zz} \end{bmatrix}. \quad (5)$$

In the preceding equation, q_x , q_y and q_z are the heat flux components and τ_{xx} , τ_{xy} , τ_{xz} , τ_{yy} , τ_{yz} , τ_{zz} are the elements of the fluid stress tensor.

II.B. Finite-Volume Method and Semi-Discrete Form

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

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

where N_g is the number of Gauss quadrature points and \vec{n} is the local normal of the face f at each of the N_g Gauss quadrature points. The hexahedral cells are contained within logically cartesian blocks that form a multi-block body-fitted mesh with general unstructured connectivity between blocks as shown in Figure 1. The total number of Gauss integration points, N_g , at which the numerical flux is evaluated is chosen as the minimum required to preserve the targeted rate of convergence for solution accuracy. In this work, standard tensor-product quadrature consisting of four Gauss quadrature points are used for the cell faces, providing a fourth order accurate spatial discretization. The latter is the target accuracy for the high-order scheme considered here.

The numerical hyperbolic fluxes, $\vec{\mathbf{F}}_H \cdot \vec{n}$, at each Gauss quadrature points on each face of a cell (i, j, k) are determined from the solution of a Riemann problem. Given the left and right interface solution values, \mathbf{U}_l and \mathbf{U}_r , an upwind numerical flux is evaluated by solving a HLLC approximate Riemann problem²⁰ in the direction defined by the normal to the face. The values of \mathbf{U}_l and \mathbf{U}_r are determined by performing a k -exact polynomial CENO reconstruction as detailed in the next section. Numerical values for the elliptic fluxes $\vec{\mathbf{F}}_E \cdot \vec{n} = \mathcal{F}_E(\mathbf{U}, \vec{\nabla} \mathbf{U}, \vec{n})$ are determined given the cell interface values of the solution \mathbf{U} , its gradient $\vec{\nabla} \mathbf{U}$ and the vector normal \vec{n} . A $(k+1)$ -exact polynomial reconstruction is used to obtain a $(k+2)$ -order accurate value of the interface solution \mathbf{U} . Direct differentiation of this $(k+1)$ -exact reconstructed solution is then used to obtain a $(k+1)$ -order accurate value of the interface gradient $\vec{\nabla} \mathbf{U}$. The resulting scheme for the Navier-Stokes equations is $(k+1)$ -order accurate. Practically, only one $(k+1)$ -exact polynomial reconstruction is performed for evaluating both the hyperbolic and elliptic fluxes to avoid additional computational effort.

II.C. High-Order CENO Spatial Discretization Scheme

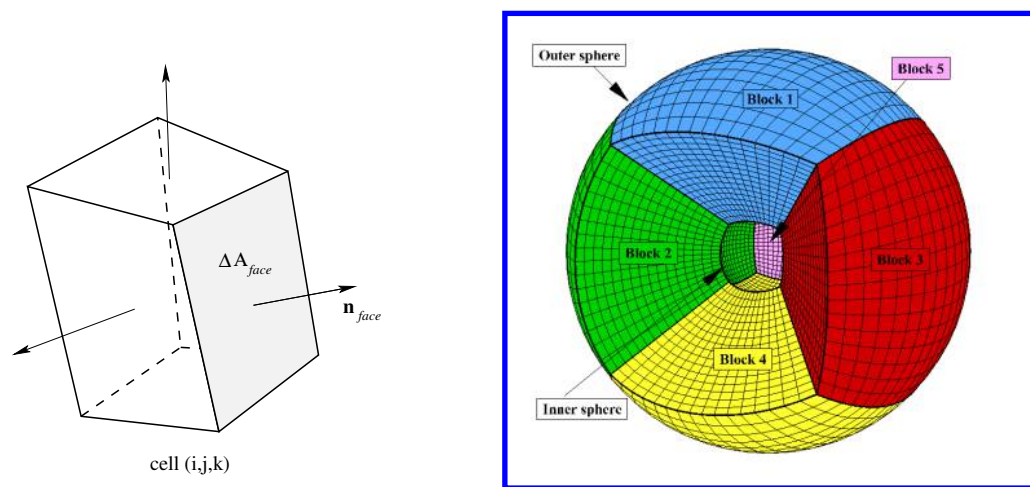
The hybrid CENO finite-volume method for conservation laws originally proposed by Ivan and Groth^{10,11} is used to discretize the governing equations on a hexahedral computational grid. The hybrid CENO procedure uses the multidimensional unlimited k -exact reconstruction of Barth²¹ in smooth regions and reverts to a limited piecewise-linear reconstruction algorithm in regions deemed as non-smooth or under-resolved by a solution smoothness indicator, \mathcal{S} , thus providing monotone solutions near discontinuities.

The k^{th} -order Taylor series representing a k -exact reconstruction of a scalar solution quantity, U_{ijk} , within a cell with index ijk about the cell-centroid $(x_{ijk}, y_{ijk}, z_{ijk})$ can be expressed as:

$$U_{ijk}^k(x, y, z) = \sum_{p_1=0}^k \sum_{p_2=0}^k \sum_{p_3=0}^k (x - x_{ijk})^{p_1} (y - y_{ijk})^{p_2} (z - z_{ijk})^{p_3} D_{p_1 p_2 p_3}. \quad (7)$$

The constant coefficients, $D_{p_1 p_2 p_3}$, are referred to as the unknown derivatives and their number, for the target fourth-order accurate is equal to 20 for the Euler equations ($k = 3$ piecewise cubic reconstruction) and 35 for the Navier-Stokes equations ($k = 4$ piecewise quartic reconstruction). Both Householder QR factorization and singular value orthogonal decomposition (SVD) can be used to solve the weighted least-squares problem associated with the CENO reconstruction.¹⁰ The latter is exploited here. The SVD approach permits the computation of a pseudo-inverse matrix after which the solution of the least-squares problem is then given by a simple matrix-vector product. The use of a single fixed stencil, the same for all dependent variables, allows the pseudo-inverse matrix to be stored and re-used in the reconstruction of all variables, thereby avoiding the repeated evaluation of the pseudo inverse. This was found to reduce significantly the computational costs of performing the CENO reconstruction without requiring substantial additional storage.¹⁰ Additionally, there are conventionally issues with high-order polynomial reconstruction related to conditioning and/or invertibility that generally increase with the order of the scheme as well as can be very dependent on mesh features, such as cell size, aspect ratio, and topology. However, a rather simple column-scaling procedure is applied here to the least-squares problem which significantly improves the conditioning, makes it virtually independent of the mesh, and affords robust and reliable solutions to the least-squares problem.¹⁰

In order to obtain an exactly determined or overdetermined set of equations to solve for the unknown derivatives, a stencil including the two nearest rings of neighbours is used whatever is the mesh discretization size in the neighbouring cells. In particular, for fourth-order spatial accuracy ($k = 3$ or cubic reconstruction for the Euler equations, $k = 4$ or quartic reconstruction for the Navier-Stokes equations) a central stencil containing $5 \times 5 \times 5$ cells is used in a region with uniform resolution, and near resolution changes or where the grid connectivity is irregular (such as at cubed-sphere sector edges). Depiction of the reconstruction stencil in the rather general case of a cubed-sphere grid is provided in Figure 2. In the stencil representation,



Hexahedral cell i showing face normal.

Example of cubed-sphere mesh composed of 6 blocks. One block has been removed for clarity

Figure 1. Three-dimensional hexahedral cell (left) and example of a cubed-sphere mesh with one of the six blocks removed for clarity (right).




Figure 2. Typical 125-cell reconstruction stencil in a cubed-sphere uniform mesh for a cell (ijk) located at the block boundary such that its stencil includes ghost cells from neighboring blocks. The interior domain is represented in blue and gray color is used for the ghost cells. (a) interior and ghost cells have the same mesh refinement levels, (b) interior cells have a coarser mesh refinement, (c) ghost cells have a coarser mesh refinement.

cells colored blue represent the interior domain of a block and the gray colored cells represent ghost cells for that block. The resulting overdetermined system provides the values of the Taylor derivatives defining unlimited high-order reconstruction polynomial valid for regions with smooth solution content. For the limited linear reconstruction used in regions deemed to be non-smooth, a central stencil of $3 \times 3 \times 3$ cells is used, thus providing an overdetermined set of equations to solve for just the three first-order unknown spatial derivatives. The mechanism used to switch between the unlimited high-order reconstruction and the piecewise limited linear reconstruction, as well as the slope limiting procedure, are fully described in the 3D context by Ivan et al.¹⁰ A k -order solution gradient at the cell face is obtained as the arithmetic mean of the left and right reconstructed gradients as done by Ivan et al.¹⁰

II.D. Block-Based Anisotropic Adaptive Mesh Refinement

A flexible block-based hierarchical binary tree data structure is used in conjunction with the high-order spatial discretization procedure described in Section II.C to facilitate automatic solution-directed anisotropic mesh adaptation on body-fitted multi-block mesh. The general AMR framework of Freret and Groth,⁷ based on extensions to the previous work by Williamschen and Groth,²² is well suited and readily allows the use of high-order spatial discretization by adopting a non-uniform representation of the cells within each block. An example of a non-uniform block obtained from a multi-block structure is shown in Figure 3. It uses directly the neighboring cells as the ghost cells, even those at different levels of refinement as found at grid resolution changes.

Mesh adaptation is accomplished by refining and coarsening grid blocks. Each refinement produces new blocks called “children” from a “parent” block and the children can be refined further. This refinement




Figure 3. Example of a non-uniform structured mesh block (right) obtained from a block-based anisotropic AMR grid mesh (left). This block is called non-uniform because its ghost cells may have different resolution from the interior cells.

process can be reversed in regions that are deemed over-resolved and two, four or eight children can coarsen or merge into a single parent block. Figure 4 shows the resulting binary tree after several refinements of an initial mesh consisting of a single block. In the present work, heuristic or physics-based refinement criteria based on the gradients of the density and/or velocity field are used to direct the mesh refinement. Directional dependent refinement criteria are developed based on the three components of the gradient vectors such that the mesh can be refined in an anisotropic way.

A high-order accurate solution transfer from the coarse cell to the fine cells is required to distribute the average solution quantity among offspring with high-order accuracy. The high-order polynomial reconstructions of all solutions variables in the coarse cell are integrated over the domain of each new fine cell having a volume, V_{fine} ,

$$\bar{u}_{fine} = \frac{1}{V_{fine}} \iiint_{V_{fine}} u_{coarse}^k(\vec{X}) dV = \frac{1}{V_{fine}} \sum_{m=1}^{N_g} \omega_m u_{coarse}^k(\vec{X}_m) \quad (8)$$

where the volume integral is computed exactly for the reconstruction polynomial with an appropriate-order tensor-product Gauss quadrature volumetric integration technique ($N_g = 27$ quadrature points are used for fourth-order spatial accuracy¹⁰). Here, ω_m and \vec{X}_m are the Gauss weights and Gauss quadrature points in the fine cell.

III. Inexact Newton Method

The Jacobian-free inexact Newton method used in this work follows the algorithm developed by Groth and co-workers.^{16, 18, 23} The latter is well adapted for computations on large multi-processor distributed-memory parallel clusters. For steady-state problems, the semi-discreted form of the governing equations given above reduces to

$$\frac{d\bar{\mathbf{U}}}{dt} = \mathbf{R}(\bar{\mathbf{U}}) = 0. \quad (9)$$

The solution of the coupled system of non-linear algebraic equations defined by Eq. (9) is obtained via Newton's method by iteratively solving a sequence of linear systems given an initial estimate \mathbf{U}^0 . Improved estimates are successively obtained by solving

$$\left(\frac{\partial \mathbf{R}}{\partial \bar{\mathbf{U}}} \right)^k \Delta \bar{\mathbf{U}}^k = \mathbf{J}(\bar{\mathbf{U}}^k) \Delta \bar{\mathbf{U}}^k = -\mathbf{R}(\bar{\mathbf{U}}^k), \quad (10)$$

where $\mathbf{J} = \partial \mathbf{R} / \partial \bar{\mathbf{U}}$ is the residual Jacobian. Hence, the improved solution at the k^{th} Newton step iteration is updated with

$$\bar{\mathbf{U}}^{k+1} = \bar{\mathbf{U}}^k + \Delta \bar{\mathbf{U}}^k.$$




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

Newton iterations are performed until a desired reduction of the residual norm is achieved, that is

$$\|\mathbf{R}(\bar{\mathbf{U}}^k)\| < \epsilon \|\mathbf{R}(\bar{\mathbf{U}}^0)\|.$$

The tolerance ϵ used in this work for steady flows is taken in the range of $[10^{-10}, 10^{-7}]$.

As noted above, each step of Newton's method requires the solution of a linear problem of the form

$$\mathbf{J}\mathbf{x} = \mathbf{b} \quad (11)$$

where $\mathbf{x} = \Delta\bar{\mathbf{U}}$ and $\mathbf{b} = -\mathbf{R}(\bar{\mathbf{U}})$. This linear system is large, sparse, and non-symmetric and is solved herein using the iterative generalized minimal residual (GMRES) algorithm of Saad and co-workers.¹⁷ A combination of an additive Schwarz global preconditioner and a block incomplete lower-upper (BILU(2)) local preconditioner is used provide a parallel implementation of the Newton method and to speed-up the convergence of the GMRES algorithm. As discussed by Dembo et al.,²⁴ determining the exact solution of Eq. (11) is not necessary for rapid convergence of Newton's method and partial convergence of the linear problem can prove to be computationally more efficient. A relatively large convergence tolerance of 0.1 for the iterative solution of the linear problem is used in the resulting inexact Newton's method of this work.

III.A. Implicit-Euler Startup For Steady-State Problems

In order to ensure that the proposed inexact Newton's method is globally convergent for the solution of steady-state problems of interest here, a so-called switched-evolution-relaxation (SER) startup strategy based on the implicit Euler time marching scheme as proposed by Mulder and Van Leer,²⁵ which provides a smooth switching between time-integration of the governing equations and Newton's method, is used and combined with the Weiss-Smith low-Mach-number preconditioner.²⁶ The latter provides improved convergence for low-Mach-number flows. The application of this startup procedure leads to the solution of a linear system of the form

$$\left[\frac{\mathbf{\Gamma}}{\Delta\tau^k} + \left(\frac{\partial\mathbf{R}}{\partial\bar{\mathbf{U}}} \right)^k \right] \Delta\bar{\mathbf{U}}^k = -\mathbf{R}^k, \quad (12)$$

where $\mathbf{\Gamma}$ is the Weiss-Smith²⁶ preconditioning matrix for the conserved variables and $\Delta\tau^k$ represents an artificial time-step. As $\Delta\tau^k \rightarrow \infty$, Newton's method of Eq. (10) is recovered. In the SER startup procedure, the time-step $\Delta\tau$ is varied, starting from a relatively small and finite value and is gradually increased to become very large as the desired steady solution is approached. As the time step becomes large, the quadratic convergence of Newton's method is recovered. The time-step size is determined by considering the inviscid Courant-Friedrichs-Lewy (CFL) condition based on the pseudo-compressible system. The maximum permissible time-step for each local cell is determined by

$$\Delta\tau^k \leq \text{CFL}^k \left(\frac{\Delta x}{\max(|\mathbf{u}| + a)} \right),$$

where $\Delta x = V^{1/3}$ is a measure of the grid size and a is the sound speed. The CFL number for the k^{th} iteration is then computed using the following relation

$$\text{CFL}^k = \text{CFL}^0 \frac{\|\mathbf{R}(\bar{\mathbf{U}}^0)\|}{\|\mathbf{R}(\bar{\mathbf{U}}^k)\|}.$$

In this work, a value of $\text{CFL}^0 = 1$ is used.

III.B. Parallel implementation

The computational domain of the proposed finite-volume scheme is represented by a multi-block mesh. The multi-block structure of the mesh is used to prescribe the block partitions of the global additive Schwarz preconditioner used in the GMRES algorithm of the inexact Newton's method. The multi-block nature of the mesh and the Schwarz preconditioning lend themselves naturally to an efficient and scalable parallel implementation of the Newton method algorithm via domain decomposition.^{16, 18, 19} The domain decomposition




Figure 5. Repartition of solution blocks in a multi-block hexahedral mesh (left) before and (right) after the space filling curve for efficient load balancing of blocks on 8 processors. All blocks of a same processor are represented by the same color. The black line represents the space filling curve passing through each of the solution blocks in a multi-block hexahedral mesh.

is carried out by distributing the mesh blocks out to the processors, with several blocks permitted on each processor. One of the key issues related to AMR is dynamic load balancing (DLB), which allows large-scale adaptive applications to be run efficiently on parallel architectures. Our DLB scheme belongs to the class of diffusion-based schemes, as opposed to the scratch-and-remap schemes. The formers are scalable since they only require local information. During AMR, the load balance is modified during both the refinement and the coarsening procedures. During the refinement process, when newly child blocks are created, blocks are taken from a pool of available resources and distributed among the processors such that the load balancing remains optimal. When blocks are coarsened, one of the blocks being coarsened gives its computational resources to the new coarse block, and the remaining blocks are given back to the pool of available resources, to be used later for any newly refined grid blocks. This scalable parallel AMR algorithm is generally sufficient to obtain and maintain an imbalance of less than 10%

When the imbalance is greater than this threshold, a scratch-and-remap strategy is applied. It consists in a space-filling curve²⁷ which simplify the partition problem by mapping a n-dimensional space to one dimension, at a very low computational cost. As proposed in Gao et al.²⁸ following the ideas of Aftosis et al,²⁹ a recursive post order traversal of the tree data structure is used as the curve generate function. Figure 5 shows the impact of the space filling curve (black line) passing through each of the solution blocks (colored blocks) in a channel with a non-smoothed bump.

IV. Numerical Results for the Euler Equations

Numerical results associated with the application of the proposed Newton method for two inviscid flow problems governed by the Euler equations will now be discussed. Numerical results are described for a steady-state supersonic radial outflow and a steady-state supersonic flow past a sphere, both of them using a cubed-sphere mesh. The flow problems have been chosen to show the potential of the proposed 3D scheme for significantly reducing computational complexity as well as the validity of the algorithm for solving both smooth and non-smooth flows. For the first flow problem, comparing the results with 3D uniform AMR illustrates the potential of the AMR scheme for reducing computational costs.

IV.A. Steady Supersonic Spherical Outflow

A spherical computational domain is created using the cubed-sphere mesh proposed by Ronchi et al.³⁰ and implemented within the block-based AMR framework by Ivan et al.^{10,31} The domain consists of six sectors (or blocks) connected with edges degeneracies, forming an inner hollow sphere and an outer spherical shell. As a first investigation into the use of the high-order CENO spatial and NKS implicit time stepping schemes within the non-uniform block-based anisotropic AMR framework on a cubed-sphere mesh, a steady-state flow problem with an analytical solution is examined. In this test-case, we consider a spherical inflow with radius $R_i = 1$ m and a spherical outflow radius of $R_o = 4$ m. Supersonic air enters through the inner sphere with a velocity only in the radial direction (V_r) subsequently expands and exits supersonically through the outflow sphere. The inflow is fixed with flow density $\rho_i = 10$ kg/m³, radial velocity $V_{r,i} = 4.5$ m/s, and pressure $p_i = 26$ Pa. A ratio of specific heats $\gamma = 1.4$ is assumed. At the outflow boundary, the gas is assumed to exit at supersonic speed. The analytical solution to this flow problem which is described by Ivan

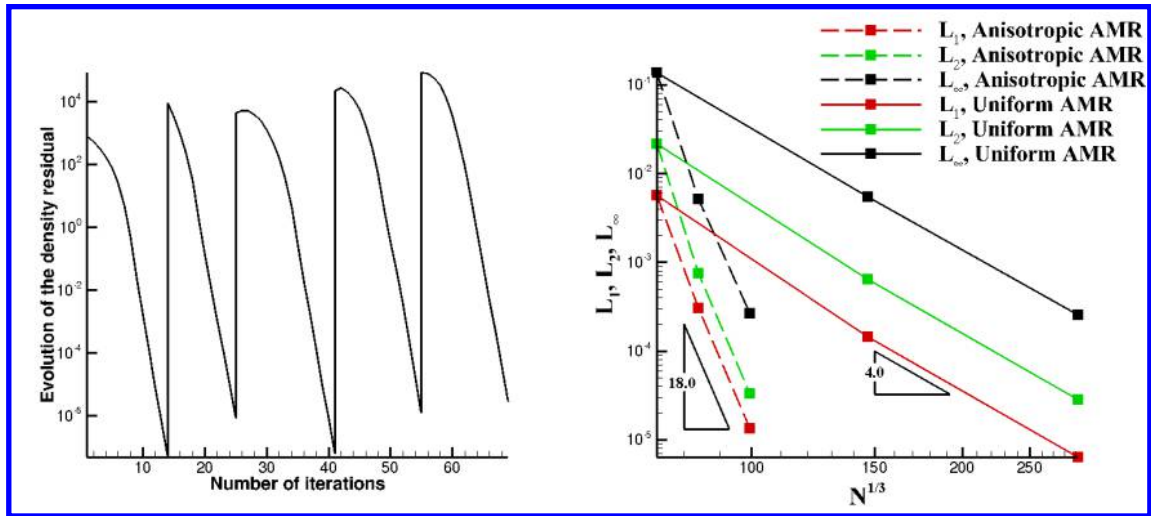


Figure 6. (left) Convergence history of the Newton scheme for the numerical prediction of steady three-dimensional supersonic outflow through 5 anisotropic AMR meshes, (right) L_1, L_2, L_∞ error norms for the fourth-order CENO scheme with anisotropic AMR (dashed lines) compared to successive uniform refinements (solid lines) as a function of the number of computational cells.

et al.³¹ is taken as the numerical solution at any radial location (r) to the equation

$$C_3 - \frac{1}{r^2 V_r [(C_2 - V_r^2)^{\frac{1}{\gamma-1}}]} = 0, \quad (13)$$

where

$$C_3 = \frac{1}{\left(\frac{2\gamma}{\gamma-1} \frac{p_i}{\rho_i}\right)^{\frac{1}{\gamma-1}} R_i^2 V_{r,i}}, \quad C_2 = \frac{2\gamma}{\gamma-1} \frac{p_i}{\rho_i} + V_{r,i}^2, \quad (14)$$

are constants depending on the inflow conditions. The initial mesh consists in a cubed-sphere grid with 1 block in the radial direction, that is a total of 6 blocks, each containing $64 \times 64 \times 16$ cells.

Within this smooth flow, a high-order k -exact polynomial reconstruction is performed for all cells. The HLLE approximate Riemann solver²⁰ is used with the Newton-Krylov-Schwarz algorithm with a GMRES tolerance of 0.1 and ILU(2) to achieve a steady-state solution. Newton iterations are stopped when the non-linear residual is reduced by 10 orders of magnitude ($\epsilon = 10^{-10}$). Four levels of anisotropic adaptive refinements are applied and the solution is converged by the inexact Newton method at each refinement




Figure 7. Slices of the predicted solution showing density in kg/m^3 and block outlines for anisotropic (left) and uniform (right) AMR after 2 mesh refinements. The anisotropic mesh exhibits a total cell count of 983,040 and the uniform mesh is composed of 25,165,824 cells.

level. The mesh refinement criteria is based on the gradient of the density. The density residual convergence history as a function of the number of Newton steps obtained on the initial coarse mesh and the subsequent sequence of adaptively refined anisotropic meshes is presented in Figure 6-left. The convergence results shows the rapid convergence of the solution to the steady state solution on each mesh level within less than 20 Newton steps for each mesh.

The predicted flow density obtained using anisotropic AMR after two levels of refinement is depicted on a slice ($x = 0$) of the cubed-sphere in Figure 7-left. It can be observed that the solution varies only along the radial direction and that anisotropic AMR exploits this feature by refining only in the radial direction. For comparison, the corresponding results obtained from two successive uniform refinements is depicted in Figure 7-right. A large reduction of 96% in the total cell count is obtained for the solution on the anisotropic mesh compared to the uniform refinements.

To assess the accuracy of the overall numerical scheme, the converged solution is compared to the exact solution given by Eq. (13-14), and the L_1 , L_2 and L_∞ error norms in the density, ρ , are computed on a successively uniformly refined and anisotropic AMR meshes. The predicted solution errors as a function of the mesh density are shown in Figure 6-right for the initially coarse and first two refined meshes, where the anisotropic AMR results are represented by the dashed lines and the uniform AMR results by the solid lines. The fourth-order theoretical accuracy for this smooth flow is achieved by the numerical scheme in all error norms as the mesh is uniformly refined. When considering anisotropic AMR, the effective convergence rate reaches the value of 18. This is reflected in a important mesh saving (96%) on the final anisotropic mesh shown.


IV.B. Steady Supersonic Flow Past a Sphere

The numerical solution of a supersonic flow past a sphere is considered next. The geometry consists of a half sphere of radius 1 m with a far-field boundary at 4 m. The cubed-sphere mesh for this test-case has only 5 sectors to reduce the problem size, each block containing $16 \times 16 \times 16$ cells. Two initial uniform refinements are applied leading to a total of 320 blocks.


Free-stream air with $p = 101.325$ kPa, $\rho = 1.225$ kg/m³, $M = 2.0$ and $\gamma = 1.4$ enters through the far-field in the positive x direction and exits through the outflow boundary. The far-field boundary condition is fixed at the inflow conditions and the outflow boundary condition is imposed by linear extrapolation since the air exits supersonically. A reflection boundary condition is imposed on the surface of the inner sphere.

In the simulations, the HLLC flux function with the Venkatakrishnan limiter was used together with the Newton-Krylov-Schwarz algorithm with a GMRES tolerance of 0.1 and ILU(2) to achieve a steady-state solution. Newton iterations are stopped when the non-linear residual is reduced by 9 orders of magnitude ($\epsilon = 10^{-9}$). The Euler startup algorithm described in Section III.A is used here. The cut-off value of the smoothness indicator within the CENO scheme $\mathcal{S}_C = 1500$ is used to switch between the fourth-order reconstruction and the limited second-order reconstruction. Four anisotropic AMR procedure are completed and the solution is converged at each refinement level.


Figure 8 (a) depicts the predicted density contours obtained on the final mesh and Figure 8 (b) shows the cells (in blue) where a second-order limited linear reconstruction is applied. It is quite apparent that the anisotropic adapted mesh is well aligned with the bow shock near the stagnation point and the solution smoothness indicator is well designed to detect the flow shock. After 4 anisotropic AMR procedures, the mesh saving for the anisotropic AMR is 99% compared to an equivalent sequence of uniform refinements. The convergence history of the density residual on all meshes is presented as a function of the number of iterations in Figure 8 (c). The convergence history illustrates that a larger number of iterations are needed on the initial mesh which is expected since the initial flow is taken to be a uniform flow. After this SER startup phase, a rapid convergence of the solution is obtained on all meshes in less than 75 iterations on all adaptively refined meshes. Figure 8 (d) depicts the evolution of the residual as a function of the total computational time. Results of this figure indicates that, while the time for solution increases on each mesh, the convergence is obtained very quickly on the three first meshes. The last AMR procedure requires more computational time as 56,660 blocks are created during the third mesh refinement procedure, this represents an increase of 275% in the number of blocks. Note that for all mesh levels, the cost of performing the mesh adaptation remains a rather small function of the cost of determining the solution. Due to the very large number of blocks on the final mesh, the residual convergence requires a large computational time.



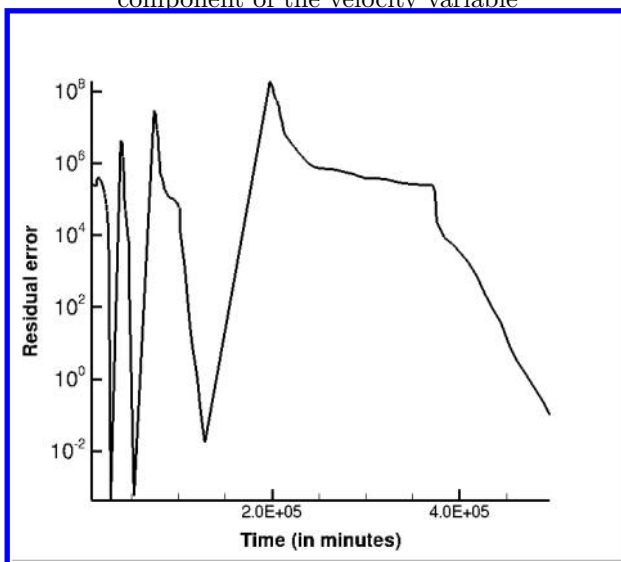
(a) Density field on an anisotropic mesh of 77,235 blocks (39,544,320 cells)



(b) Smoothness indicator of the first component of the velocity variable



(c) Density residual convergence history versus the number of iterations through 4 anisotropic AMR meshes



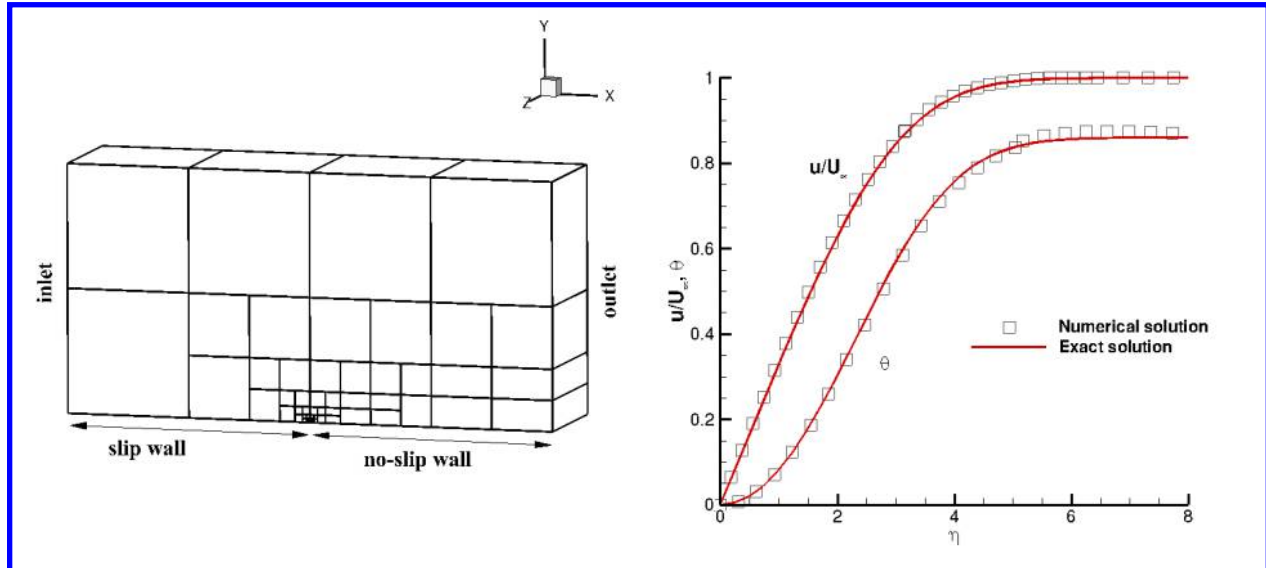
(d) Density residual convergence history versus time (in minutes) through 4 anisotropic AMR meshes

Figure 8. Numerical results of a steady flow past a sphere obtained on successive anisotropic AMR meshes. (a) The density field is presented on the final mesh which is composed of 77,235 blocks. Compared to a uniform refinement, a mesh saving of 99% is obtained for the final anisotropic refinement. (b) The smoothness indicator is indicated in blue for the first component of the velocity field. A low-order reconstruction is performed in these cells that represent about 30% of the total number of cells. Evolution of the density residual versus (c) the number of iterations and (d) time through 4 anisotropic AMR meshes.

V. Numerical Results for the Navier-Stokes Equations

V.A. Subsonic Laminar Boundary-Layer Flow Past a Flat Plate

Numerical results for the steady laminar flow over a flat plate are computed. A rectangular domain




(a) Final computational mesh used to study steady laminar flow over a flat plate composed by 71 blocks (b) non-dimensional velocity profiles compared with Blasius solution

Figure 9. (a) Final computational mesh for the subsonic laminar boundary-layer flow past a flat plate and (b) Comparison between non-dimensional numerical and analytical velocity profiles.

$[-2L, 2L] \times [0, 2L] \times [0, L]$ is used. The far-field and upstream of the plate are modeled by a reflection/slip condition while the presence of the plate, on $x = [0, 2L], y = 0$ is simulated using a no-slip condition. A uniform velocity profile is specified at the inlet ($x = -2L$) and a zero-gradient condition on the fluid velocity is applied at the outlet ($x = 2L$). Pressure is held fixed and equal to 0 at the outlet. The free-stream Mach number and Reynolds number based on the plate length ($L = 0.002$ m) are $M_\infty = 0.2$ and $Re_L = 9318.92$, respectively. The initial condition is a uniform flow of air. The initial mesh consists of 2 blocks of $16 \times 16 \times 8$ cells. Five anisotropic refinements are performed during the simulation and the solution is converged at each level of refinement.

In Figure 9 (a), the mesh geometry is depicted at the final level of refinement, where 71 blocks are clustered near the leading edge of the plate. No refinement is performed in the z -direction. The results are obtained using the inexact Newton's method described in Section III for steady flows. Newton iterations are stopped when the non-linear residual is reduced by 10 orders of magnitude ($\epsilon = 10^{-10}$). At the exit plane of the flat plate at $x = 0.004$ m, velocity data are extracted across the boundary layer and compared with the Blasius analytical solution. This is shown in Figure 9 (b) where the velocity profile is plotted using similar variables from the Blasius solution. The predicted non-dimensional boundary layer profiles, u/U_∞ and $\theta = v\sqrt{Re_x}/U_\infty$, are represented at $Re_x = 8,900$ where U_∞ is the free-stream velocity. Here, the non-dimensional distance from the plate (η) is defined as $y\sqrt{Re_x}/x$.

The convergence history for the L_1 norm of the solution residual is provided in Figure 10 (a). Excellent convergence characteristics are displayed by the parallel Newton-Krylov-Schwarz solution algorithm, requiring a maximum of only 50 Newton iterations and 3163 evaluations of Eq.(11) to reduce the equation residuals by ten orders of magnitude. Figure 10 (b) presents the evolution of the residual versus total computational time. From the convergence history as a function of time, it can be seen that the increases in the computational cost to obtain the solution on each successively refined mesh are not that substantial for this rather simple viscous flow case. This is because for all meshes there is at most one grid block per processor as the mesh is refined (71 blocks are present on the final mesh) and the proposed solution algorithm exhibits very good parallel efficiency.



(a) Evolution of the residual versus the number of iterations through six successively refined meshes.

(b) Evolution of the residual versus time through six successively refined meshes.

Figure 10. L_1 norm of the solution residual history obtained for the subsonic laminar boundary-layer flow past a flat plate on six successive anisotropic mesh refinement levels versus (a) the number of iterations and (b) time (in minutes).

VI. Conclusions

A fourth-order CENO finite-volume scheme has been developed and coupled to an inexact Newton method and anisotropic block-based AMR for the prediction of three-dimensional compressible flows. Three-dimensional anisotropic AMR results have been obtained for two inviscid steady flow problems on cubed-sphere grids, one with smooth solution content and one with discontinuities (shocks) and a viscous flow with boundary layers. Efficient and robust convergence of the Newton method is demonstrated for both smooth and non-smooth solutions. Moreover, comparisons of the predicted solutions to those obtained on uniformly refined meshes and the corresponding accuracy assessments have demonstrated the high computational efficiency of the proposed approach. High accurate solutions have been obtained with a reduced computational effort and significant reductions in mesh size.

Acknowledgments

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

References

- ¹Northrup, S. and Groth, C., "Solution of laminar diffusion flames using a parallel adaptive mesh refinement algorithm," No. 2005-0547, AIAA, 2005.
- ²Gao, X., *A Parallel Solution-Adaptive Method for Turbulent Non-Premixed Combusting Flows*, Ph.D. thesis, University of Toronto, 2008.
- ³Sachdev, J., Groth, C., and Gottlieb, J., "A parallel solution-adaptive scheme for multi-phase core flows in solid propellant rocket motors," *International Journal of Computational Fluid Dynamics*, Vol. 19, No. 2, 2005, pp. 159–177.
- ⁴Groth, C., Zeeuw, D. D., Powell, K., Gombosi, T., and Stout", Q.
- ⁵Groth, C., Zeeuw, D. D., Gombosi, T., and Powell, K., "Global Three-Dimensional MHD Simulation of a Space Weather Event: CME Formation, Interplanetary Propagation, and Interaction with the Magnetosphere," *Journal of Geophysical Research*, Vol. 105, No. A11, 2000, pp. 25,053–25,078.

⁶McDonald, J. and Groth, C., *Numerical modeling of micron-scale flows using the Gaussian moment closure*, University of Toronto, 2005.

⁷Freret, L. and Groth, C., “Anisotropic Non-Uniform Block-Based Adaptive Mesh Refinement for Three-Dimensional Inviscid and Viscous Flows,” 2015.

⁸Freret, L., Groth, C., and De Sterck, H., “A Parallel High-Order CENO Finite-Volume Scheme With AMR For Three-Dimensional Steady Ideal MHD Flows,” 2016.

⁹Freret, L., Ivan, L., De Sterck, H., and Groth, C., “A High-Order Finite-Volume Method with Anisotropic AMR for Ideal MHD Flows,” 2017.

¹⁰Ivan, L., De Sterck, H., Susanto, A., and Groth, C., “High-order central ENO finite-volume scheme for hyperbolic conservation laws on three-dimensional cubed-sphere grids,” *Journal of Computational Physics*, 2015.

¹¹Ivan, L. and Groth, C., “High-order solution-adaptive central essentially non-oscillatory (CENO) method for viscous flows,” *Journal of Computational Physics*, Vol. 257, 2013, pp. 830–862.

¹²Tobaldini Neto, L. and Groth, C., “A High-Order Finite-Volume Scheme for Large-Eddy Simulation of Turbulent Premixed Flames,” AIAA paper, January 2014.

¹³Susanto, A., Ivan, L., De Sterck, H., and Groth, C., “High-Order Central ENO Finite-Volume Scheme for Ideal MHD,” *Journal of Computational Physics*, Vol. 250, 2013, pp. 141–164.

¹⁴Charest, M. and Groth, C., “A High-Order Central ENO Finite-Volume Scheme for Three-Dimensional Low-Speed Viscous Flows on Unstructured Mesh,” *Communications in Computational Physics*, Vol. 17, 2015, pp. 615–656.

¹⁵Charest, M. and Groth, C., “A High-Order Central ENO Finite-Volume Scheme for Three-Dimensional Turbulent Flows on Unstructured Mesh,” AIAA paper.

¹⁶Groth, C. and Northrup, S., “Parallel implicit adaptive mesh refinement scheme for body-fitted multi-block mesh,” *17th AIAA Computational Fluid Dynamics Conference, Toronto, Ontario, Canada, AIAA paper*, Vol. 5333, 2005.

¹⁷Saad, Y. and Schultz, M., “GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems,” *SIAM J. Sci. Stat. Comput.*, Vol. 7, No. 3, 1986.

¹⁸Northrup, S. and Groth, C., “Parallel implicit adaptive mesh refinement scheme for unsteady fully-compressible reactive flows,” *21st AIAA Computational Fluid Dynamics Conference, San Diego, California, AIAA paper*, Vol. 2433, 2013.

¹⁹Charest, M., Groth, C., and Gülder, Ö., “A computational framework for predicting laminar reactive flows with soot formation,” *Combustion Theory and Modelling*, Vol. 14, No. 6, 2010, pp. 793–825.

²⁰Einfeldt, B., “On Godunov-Type Methods for Gas Dynamics,” Vol. 25, 1988, pp. 294–318.

²¹Barth, T., “Recent developments in high order k-exact reconstruction on unstructured meshes,” 1993.

²²Williamschen, M. J. and Groth, C. P. T., “Parallel anisotropic block-based adaptive mesh refinement algorithm for three-dimensional flows,” *21th AIAA Computational Fluid Dynamics Conference*, June 2013.

²³Charest, M., Groth, C., and Gülder, Ö., “Solution of the equation of radiative transfer using a Newton-Krylov approach and adaptive mesh refinement,” *J. Comput. Phys.*, Vol. 231, No. 8, 2012, pp. 3023–3040.

²⁴Dembo, R., Eisenstat, S., and Steihaug, T., “Inexact Newton methods,” *SIAM, Journal on Numerical Analysis*, Vol. 19, 1982.

²⁵Mulder, W. and Van Leer, B., “Experiments with implicit upwind methods for the Euler equations,” *Journal of Computational Physics*, Vol. 59, 1985, pp. 232–246.

²⁶Weiss, J.M., M. J. S. W., “Implicit Solution of Preconditioned Navier-Stokes Equations Using Algebraic Multigrid,” *AIAA Journal*, Vol. 37.

²⁷Sagan, H., *Space-Filling Curves*.

²⁸Gao, X., Northrup, S., and Groth, C., “Parallel solution-adaptive method for two-dimensional non-premixed combusting flows,” *Progress in Computational Fluid Dynamics*, Vol. 11, No. 2, 2011, pp. 76–95.

²⁹Aftosmis, M., Berger, M., and Melton, J., “Robust and Efficient Cartesian Mesh Generation for Component-Based Geometry,” *AIAA Journal*, 1998.

³⁰C. Ronchi, R. I. and Paolucci, P., “The cubed sphere: a new method for the solution of partial differential equations in spherical geometry,” *Journal of computational physics*, Vol. 124, No. 1, 1996, pp. 93–114.

³¹Ivan, L., De Sterck, H., S., N., and Groth, C., “Multi-dimensional finite-volume scheme for hyperbolic conservation laws on three-dimensional solution-adaptive cubed-sphere grids,” *Journal of Computational Physics*, Vol. 255, 2013, pp. 205–227.