Journal of Computational Physics 257 (2014) 830–862

Contents lists available at ScienceDirect

Journal of Computational Physics

www.elsevier.com/locate/jcp

High-order solution-adaptive central essentially non-oscillatory (CENO) method for viscous flows

Lucian Ivan *,[1], Clinton P.T. Groth

University of Toronto Institute for Aerospace Studies, 4925 Dufferin Street, Toronto, Ontario, M3H 5T6 Canada

ARTICLE INFO

Article history:
Received 13 December 2012
Received in revised form 7 September 2013
Accepted 24 September 2013
Available online 2 October 2013

Keywords:
High-order ENO scheme
Central stencil
Adaptive mesh refinement
Inviscid and laminar viscous flows

ABSTRACT

A high-order, central, essentially non-oscillatory (CENO), finite-volume scheme in combination with a block-based adaptive mesh refinement (AMR) algorithm is proposed for solution of the Navier–Stokes equations on body-fitted multi-block mesh. In contrast to other ENO schemes which require reconstruction on multiple stencils, the proposed CENO method uses a hybrid reconstruction approach based on a fixed central stencil. This feature is crucial to avoiding the complexities associated with multiple stencils of ENO schemes, providing high-order accuracy at relatively lower computational cost as well as being very well suited for extension to unstructured meshes. The spatial discretization of the inviscid (hyperbolic) fluxes combines an unlimited high-order k-exact least-squares reconstruction technique following from the optimal central stencil with a monotonicity-preserving, limited, linear, reconstruction algorithm. This hybrid reconstruction procedure retains the unlimited high-order k-exact reconstruction for cells in which the solution is fully resolved and reverts to the limited lower-order counterpart for cells with under-resolved/discontinuous solution content. Switching in the hybrid procedure is determined by a smoothness indicator. The high-order viscous (elliptic) fluxes are computed to the same order of accuracy as the hyperbolic fluxes based on a k-order accurate cell interface gradient derived from the unlimited, cell-centred, reconstruction. A somewhat novel h-refinement criterion based on the solution smoothness indicator is used to direct the steady and unsteady mesh adaptation. The proposed numerical procedure is thoroughly analyzed for advection–diffusion problems characterized by the full range of Péclet numbers, and its predictive capabilities are also demonstrated for several inviscid and laminar flows. The ability of the scheme to accurately represent solutions with smooth extrema and yet robustly handle under-resolved and/or non-smooth solution content (i.e., shocks and other discontinuities) is shown. Moreover, the ability to perform mesh refinement in regions of under-resolved and/or non-smooth solution content is also demonstrated.

1. Introduction and motivation

In spite of rapid advances in high performance computing, there are a number of physically-complex flows for which computational costs can make their numerical solution prohibitive and/or non-routine. Such flows would include but are certainly not limited to compressible turbulent and turbulent reactive flows. Effective numerical solution methods for such complex flows may require both high-order discretizations and adaptive mesh refinement (AMR). Moin and Mahesh [1]

* Corresponding author.
E-mail addresses: livan@math.uwaterloo.ca (L. Ivan), groth@utias.utoronto.ca (C.P.T. Groth).
[1] Currently at the Department of Applied Mathematics, University of Waterloo, Waterloo, Ontario, Canada N2L 3G1.

0021-9991/$ – see front matter
http://dx.doi.org/10.1016/j.jcp.2013.09.045

point out that even ENO schemes as high as sixth-order can be too dissipative for the prediction of shock/turbulence interactions without refinement of the mesh in the vicinity of the shock wave and Jameson [2] has shown that high-order schemes combined with AMR can be very effective for obtaining high solution accuracy and treating flows containing vortices, eddies, and turbulence.

In the past few decades, there have been numerous studies of high-order methods. For hyperbolic conservation laws and/or compressible flow simulations, the challenge has been to achieve accurate discretizations while coping with shocks in a reliable and robust fashion. The essentially non-oscillatory (ENO) finite-volume schemes first proposed by Harten et al. [3] provide a robust framework for high-order finite-volume discretizations of hyperbolic systems. The original ENO scheme is based on an adaptive-stencil strategy, in which the stencil leading to the "smoothest" reconstruction is selected and stencils containing discontinuities are thereby avoided. Although originally developed for structured regular mesh, Abgrall [4] and Sonar [5] have since extended the ENO concept for application to unstructured grids. In addition, weighted ENO (WENO) schemes have also been developed for both structured and unstructured meshes [6–8]. Nevertheless, the difficulty with these high-order approaches has been the extension of the method to multi-dimensional problems and large systems of coupled partial differential equations (PDEs). The computational challenges in these cases are primarily related to stencil selection, particularly for unstructured grids [4,5,9,10], and the poor conditioning of the linear systems that define the reconstructions for the stencils [9,10]. Moreover, the requirement of using a different stencil for each flow variable is also somewhat problematic, particularly for PDE systems with many dependent variables. Other perceived challenges associated with ENO and WENO schemes are their relatively large non-compact stencils which can give rise to difficulties with the application of boundary data and efficient parallel implementation of the algorithm. While successful implementations of this class of finite-volume scheme have been developed, in general the computational costs and complexity of the schemes have somewhat limited their widespread application. In spite of this, combinations of high-order ENO and WENO schemes with AMR for both structured and unstructured meshes have been developed and applied to a range of engineering problems. See, for example, the recent study of Wolf and Azevedo [11].

Other high-order finite-volume schemes have been considered, which may be more readily extended to multi-dimensional problems and unstructured meshes. For example, Barth and Fredrickson [12,13] previously proposed a high-order, finite-volume approach for unstructured mesh based on k-exact least-squares reconstruction. Following this work, Ollivier-Gooch [14,15] more recently proposed a data-dependent, weighted, least-squares, reconstruction procedure (DD-ENO), that uses a fixed stencil and seeks to enforce monotonicity by introducing data-dependent weights for each point in the stencil. To circumvent some deficiencies in the DD-ENO procedure, Ollivier-Gooch et al. [16,17] have also considered a high-order limiting strategy in combination with k-exact reconstruction. Capdeville also recently revisited the DD-ENO concept and formulated a compact Hermite least-square monotone (HLSM) reconstruction scheme for one space dimension [18]. Additionally, Colella et al. [19] have proposed a high-order finite-volume method for the solution of linear elliptic and hyperbolic PDEs in mapped coordinates. This approach was also extended to nonlinear systems of hyperbolic conservation laws on locally-refined grids by McCorquodale and Colella [20]. To suppress oscillations and ensure monotonicity, a rather elaborate combination of slope limiters, slope flattening, and artificial viscosity is applied in the latter. The proposed multi-dimensional limiting strategy in this case employs a modified version of the one-dimensional limiter formulated by Colella and Sekora [21] for preserving accuracy at smooth extrema.

High-order finite-element schemes have also been considered for problems involving discontinuities. The class of schemes now generally referred to as discontinuous Galerkin (DG) schemes have gained in popularity. Cockburn et al. [22–24] were the first to formulate a family of high-order, total-variation-bounded (TVB), schemes for nonlinear systems of conservation laws referred to as Runge–Kutta, discontinuous, Galerkin (RKDG) methods. By combining elements from both finite-element and finite-volume methods, a family of flexible high-order numerical schemes were obtained with improved data locality and treatment for discontinuous solutions. The reduced stencil is beneficial for both boundary-condition and parallel implementation strategies. A variety of approaches have been considered to achieve non-oscillatory solutions. Both TVB [24] and WENO-type [25] limiting strategies, as well as artificial viscosity techniques, such as the PDE-based artificial viscosity model formulated by Barter and Darmofal [26,27], have been considered to avoid spurious oscillations in the proximity of shocks. Xu et al. [28] have also recently proposed a hierarchical limited reconstruction strategy in combination with WENO-type linear reconstruction. Despite the considerable interest generated by DG methods over the last decade, the high-order versions of the schemes do suffer from: (i) a generally more restrictive stability limit for time integration schemes that is exacerbated as the order is increased [29]; (ii) a relatively high computational cost per cell associated with the solution of additional unknowns; and (iii) a higher storage cost, particularly when implicit time-marching schemes are considered. Partly in response to some of these issues, Dumbser et al. [30] have recently formulated a unified framework for both finite-volume and DG methods in terms of a family of reconstructed DG schemes denoted as $P_N P_M$ methods.

In other work, Wang et al. [31–33] and Liu et al. [34] have proposed variants of the DG formulation, referred to as spectral volume (SV) and spectral difference (SD) methods. The DG, SV and SD methods share many similarities, with the main differences being related to how the degrees of freedom of each scheme are selected and updated. A new perspective regarding these methods was introduced by Huynh [35], who was able to recover each method as a particular case of a more general formulation of high-order schemes for hyperbolic conservation laws based on a flux reconstruction (FR) approach. The SV and SD methods usually fall back on total variation diminishing (TVD) limiting strategies for dealing with undesirable solution oscillations; however, Yang and Wang [36,37] have recently combined a SD scheme for unstructured grids with a compact high-order hierarchical moment limiter. More recently, Wang and Gao [38] have generalized Huynh's

flux reconstruction approach to mixed grids and formulated a unifying lifting collocation penalty (LCP) method for the Euler equations and applied it to smooth flows. Additionally, the LCP method is shown to be more efficient in terms of both memory usage and computational cost as compared to standard DG methods. Alternative high-order schemes, such as residual distribution schemes [39–42] and finite-difference algorithms [43], have also been developed and are currently being pursued for application on unstructured and/or multi-block structured meshes.

High-order schemes for PDEs having an elliptic nature have also been considered. In these cases, it is desirable that the discretization of the elliptic operator remain accurate while satisfying a maximum principle, even on stretched/distorted meshes [44,45]. Even standard lower-order spatial discretization procedure may not possess both of these characteristics. Barad and Colella have proposed a fourth-order-accurate AMR finite-volume scheme for Poisson's equation [46] and Ollivier-Gooch and Van Altena [47] describe a general high-order finite-volume framework for the solution of the advection–diffusion equation on unstructured mesh. The DG method was originally extended for the treatment of the Navier–Stokes equations with elliptic operators by Bassi and Rebay [48]. In their approach, the Navier–Stokes equations are recast as a coupled first-order system and then discretized using a standard DG formulation with a symmetric flux function for the elliptic fluxes. A similar technique was also advocated in the so-called local discontinuous Galerkin (LDG) method of Cockburn and Shu [49] convection–diffusion equations; however, in this case a non-symmetric treatment of the elliptic fluxes was proposed. The symmetric elliptic flux function of Bassi and Rebay can be shown to be unstable and various penalty methods have been devised to stabilize and couple the scheme based on computed solution jumps at the cell interfaces. In other high-order DG work related to diffusion problems, van Leer et al. [50,51] have proposed a recovery-based approach that allows the DG method to be applied directly to the diffusion equation and eliminates the introduction of the somewhat *ad hoc* penalty or coupling terms of other DG schemes. An alternative and interesting treatment for diffusion problems has also been proposed by Gassner et al. [52,53] based on a local analytical approximations to the generalized solution of the Riemann initial value problem for the diffusion equation. The latter readily allows for the treatment of rapidly varying values of the coefficient of diffusion [53] and has a number of similarities with the so-called direct discontinuous Galerkin (DDG) method of Liu and Yan [54], in which the high-order approximations for the elliptic fluxes are evaluated based on an analytical solutions to the heat equation. Finally, Oliver and Darmofal [55] have applied a DG scheme with PDE-based artificial viscosity method to the prediction of turbulent aerodynamic flows.

In other high-order treatments for PDEs with elliptic terms, Sun et al. [56] and May and Jameson [57,58] have also considered the application of the SV and SD methods, respectively, to viscous flows. The latter also applied the SD scheme in combination with AMR. Additionally, Kannan and Wang [59–62] have analyzed and evaluated various elliptic flux functions, including those based on the LDG, penalty, and DDG methods, for use with the SV approach and shown the importance of the elliptic flux discretization procedure for high-order discretizations on unstructured grids. Haga et al. [63] have also extended the so-called correction procedure via reconstruction (CPR) formulation to solve the Navier–Stokes equations on three-dimensional mixed grids composed of tetrahedrons and triangular prisms. Finally, in other earlier work, De Rango and Zingg [64] considered the application of high-order finite-difference methods to the prediction of turbulent aerodynamic flows.

In spite of the advances in high-order accurate methods, there is still no consensus for a robust, efficient, and high-order accurate scheme that fully deals with all of the aforementioned issues and is applicable to arbitrary meshes and complex boundary-condition problems. Examples of the complexities that can arise in the specification of boundary values for solution variables would include boundary conditions for three-dimensional geometries with irregular and complex topologies, constant mass-flow boundary conditions for internal flow applications, and time-varying boundary conditions involving interfaces between stationary and rotating frames in turbomachinery simulations. In the current study, the high-order central ENO (CENO) cell-centred finite-volume scheme proposed by Ivan and Groth [65,66,122] for inviscid flows is extended to the solution of the Navier–Stokes equations governing two-dimensional, compressible, viscous flows on body-fitted multi-block mesh. The CENO discretization of the inviscid (hyperbolic) flux is based on a hybrid solution reconstruction procedure that combines the unlimited high-order k-exact least-squares reconstruction technique of Barth [12] based on a fixed central stencil with a monotonicity-preserving, limited, piecewise, linear least-squares reconstruction algorithm [12]. Switching in the hybrid procedure is determined by a solution smoothness indicator that specifies whether or not solution content is resolved on the mesh. Limited reconstruction is applied to computational cells with discontinuous and under-resolved solution content and the high-order, unlimited, k-exact reconstruction scheme is used for cells in which the solution is fully resolved. To guarantee a consistent k-order accurate numerical scheme on arbitrary meshes, the proposed discretization of the viscous (elliptic) flux is based on a k-order accurate, average, interface gradient that can be derived from the unlimited, k-exact, cell-centred reconstruction. In this way, the same solution reconstruction is used in computing both inviscid and elliptic fluxes. The high-order CENO scheme is combined with a flexible, block-based, hierarchical data structure to facilitate parallel implementation via domain decomposition and automatic solution-directed mesh adaptation on body-fitted multi-block quadrilateral mesh [67–69]. An h-refinement criterion based on the solution smoothness indicator is proposed to control the refinement of the multi-block AMR mesh.

The key advances in the proposed high-order CENO finite-volume method include:

- the development of a hybrid single-stencil reconstruction procedure as directed by a rather effective and robust smoothness indicator;

- the formulation of a consistent high-order treatment of the viscous fluxes based on a uniquely defined k-order accurate interface gradient;
- the high-order treatment for non-rectilinear boundaries; and
- the extension of the high-order approach for use in multi-block AMR mesh as directed by a new and novel h-refinement criterion.

By adopting a hybrid approach based on a single central stencil for all solution variables, the CENO scheme overcomes some of the practical drawbacks of other ENO and WENO schemes for application to multi-dimensional problems with large numbers of unknowns, especially for unstructured meshes, which are: (i) stencil selection; (ii) high computational cost associated with multiple reconstructions on different stencils for each dependent variable; (iii) issues with singular stencils [10]; and (iv) the occurrence of negative weights [70]. Other issues associated with loss of strict monotonicity [4,71] and tuning of WENO weights to recover the central stencil, recognized to be the most accurate stencil, in the smooth parts of the solution [72,73] are also dealt with. The proposed central treatment for the elliptic flux, based on averaging of the reconstructed solution and its gradient, is rather inexpensive to compute and provides a consistent approach in which the hyperbolic and elliptic fluxes are of equal order and yet it avoids issues with solution decoupling and severely restrictive stability limits.

In what follows, the proposed high-order CENO finite-volume method, block-based AMR procedure, and high-order treatment of boundary data are all described. Specifically, the next section, Section 2, reviews the governing equations of interest here. This is followed in Section 3 by a detailed description of the high-order CENO finite-volume scheme, including the high-order evaluation of the numerical residual containing both inviscid and viscous fluxes. The same section provides details of the hybrid CENO reconstruction procedure, smoothness indicator, high-order boundary conditions implementation. The combination of the high-order CENO algorithm with a block-based AMR scheme, previously developed for second-order accuracy, is then described in Section 4, with an emphasis on the algorithmic extensions required to maintain high-order solution accuracy on the AMR mesh. Finally, the properties of the proposed CENO scheme are investigated in Section 5 by comparing numerical solutions to analytical results for a range of problems pertaining to an advection–diffusion model equation. The accuracy of the hyperbolic flux discretization is demonstrated for solutions of the Euler equations governing inviscid flows. Moreover, the capabilities of the proposed approach for laminar viscous flows governed by the Navier–Stokes equations are demonstrated by comparing numerical predictions to other numerical and experimental results reported in the literature for flow past a NACA0012 airfoil as well as a circular cylinder. Results obtained with the CENO scheme with AMR are also described to illustrate the capabilities of the proposed combined approach. For additional computational details and discussions of the CENO scheme, the reader is referred to the dissertation of Ivan [66].

2. Conservation equations

The proposed high-order finite-volume method is applied herein to solutions of both a scalar advection–diffusion model equation and the Navier–Stokes equations governing laminar compressible flows. The advection–diffusion equation is a convenient mathematical model for the development of numerical methods for it contains many of the features of more complicated PDEs, yet retains a simplicity that permits direct analysis. The particular form considered here is

$$\frac{\partial u}{\partial t} + \vec{\nabla} \cdot \left(\vec{V}(\vec{r})u\right) = \vec{\nabla} \cdot \left(\kappa(\vec{r})\vec{\nabla}u\right) + \phi(\vec{r}, u), \tag{1}$$

where t is the time, u is the solution, $\vec{V}$ is the advection velocity vector, κ is the diffusion coefficient, and ϕ is a non-linear source term. Here, $\vec{V}$ and κ are most generally functions of the position vector, $\vec{r}$. Based on the relative magnitudes of the advective and diffusive fluxes, solutions of the equation can range from those having a more hyperbolic nature and governed by wave propagation phenomena to those having a more elliptic nature and governed by diffusive processes. It is desirable that numerical solution schemes do not introduce excessive artificial dissipation, large dispersive error, and spurious oscillations arising from the discretization of the hyperbolic flux, and provide accurate discretizations of the elliptic flux while satisfying a maximum principle [74,75].

The Navier–Stokes equations governing viscous compressible gaseous flows have the general weak conservation form

$$\frac{\partial \mathbf{U}}{\partial t} + \vec{\nabla} \cdot \vec{\mathbf{F}} = \frac{\partial \mathbf{U}}{\partial t} + \vec{\nabla} \cdot \vec{\mathbf{F}}_{\mathrm{H}}(\mathbf{U}) + \vec{\nabla} \cdot \vec{\mathbf{F}}_{\mathrm{E}}(\mathbf{U}, \vec{\nabla}\mathbf{U}) = \mathbf{S}, \tag{2}$$

where $\mathbf{U}$ is the vector of conserved solution variables, and $\vec{\mathbf{F}}$ is the solution flux dyad, and $\mathbf{S}$ is a source vector representing various physical phenomena and/or body forces. The solution flux, $\vec{\mathbf{F}}$, is the sum of a hyperbolic (inviscid) term, $\vec{\mathbf{F}}_{\mathrm{H}}$, which depends on the solution vector and accounts for transport by wave phenomena and an elliptic (viscous) term, $\vec{\mathbf{F}}_{\mathrm{E}}$, associated with diffusion processes and dependent on both the solution vector and its gradient. In the absence of body forces, such as gravity, $\mathbf{S} = 0$ and, for two-dimensional planar flows, the solution vector, $\mathbf{U}$, is given by

$$\mathbf{U} = [\rho, \quad \rho v_x, \quad \rho v_y, \quad \rho e]^{\mathrm{T}}, \tag{3}$$

where ρ is the gas density, v_x and v_y are the velocity components in the x- and y-coordinate directions, $e = p/(\rho(\gamma - 1)) + (v_x^2 + v_y^2)/2$ is the specific total energy, $p = \rho RT$ is the pressure, T is the gas temperature, R is the gas constant, γ

is the specific heat ratio. The components of the hyperbolic flux dyad, $\vec{\mathbf{F}}_{\mathrm{H}}(\mathbf{F},\mathbf{G})$, and elliptic flux dyad, $\vec{\mathbf{F}}_{\mathrm{E}}(\mathbf{F}_{\mathrm{v}},\mathbf{G}_{\mathrm{v}})$, are in this case given by

$$\mathbf{F}=\left[\rho v_x,\ \rho v_x^2+p,\ \rho v_x v_y,\ v_x(\rho e+p)\right]^{\mathrm{T}},\qquad \mathbf{G}=\left[\rho v_y,\ \rho v_x v_y,\ \rho v_y^2+p,\ v_y(\rho e+p)\right]^{\mathrm{T}}, \tag{4}$$

and

$$\mathbf{F}_{\mathrm{v}}=-[0,\ \tau_{\mathrm{xx}},\ \tau_{\mathrm{yx}},\ v_x\tau_{\mathrm{xx}}+v_y\tau_{\mathrm{xy}}-q_{\mathrm{x}}]^{\mathrm{T}},\qquad \mathbf{G}_{\mathrm{v}}=-[0,\ \tau_{\mathrm{xy}},\ \tau_{\mathrm{yy}},\ v_x\tau_{\mathrm{yx}}+v_y\tau_{\mathrm{yy}}-q_{\mathrm{y}}]^{\mathrm{T}}. \tag{5}$$

The terms τ_{xx}, τ_{yy} and τ_{xy} are the components of the fluid stress tensor, $\vec{\vec{\tau}}=2\mu(\vec{\vec{S}}-\vec{\vec{I}}\,\vec{\nabla}\cdot\vec{V}/3)$, where μ is the dynamic viscosity and $\vec{\vec{S}}=[\vec{\nabla}\vec{V}+(\vec{\nabla}\vec{V})^{\mathrm{T}}]/2$ is the strain rate tensor. The heat flux vector, $\vec{q}=[q_{\mathrm{x}},q_{\mathrm{y}}]$, follows from Fourier's law and is given by $\vec{q}=-\kappa\vec{\nabla}T$, where κ is the thermal conductivity. For a calorically perfect polytropic gas, the ratio of specific heats, γ, is a constant and the specific heats are $C_v=R/(\gamma-1)$ and $C_p=\gamma R/(\gamma-1)$. In this case, Eq. (2) can be then re-written as

$$\frac{\partial\mathbf{U}}{\partial t}+\frac{\partial}{\partial x}(\mathbf{F}+\mathbf{F}_{\mathrm{v}})+\frac{\partial}{\partial y}(\mathbf{G}+\mathbf{G}_{\mathrm{v}})=0. \tag{6}$$

Note that in two space dimensions, the advection–diffusion equation above can also be written in the form of Eq. (6). By defining the scalar fluxes $f=V_x u$, $g=V_y u$, $f_{\mathrm{d}}=-\kappa\,\partial u/\partial x$, and $g_{\mathrm{d}}=-\kappa\,\partial u/\partial y$, and source $s=\phi$, one can write

$$\frac{\partial u}{\partial t}+\frac{\partial}{\partial x}(f+f_{\mathrm{d}})+\frac{\partial}{\partial y}(g+g_{\mathrm{d}})=s. \tag{7}$$

3. High-order CENO finite-volume scheme

A high-order, cell-centred, CENO, finite-volume scheme is proposed for solving mixed type systems of conservation laws as given in Eqs. (6) and (7) on body-fitted multi-block quadrilateral mesh in conjunction with a block-based adaptive mesh refinement technique. Use of this particular grid structure and AMR strategy has been shown to allow highly efficient and scalable parallel implementations of finite-volume methods [67–69].

Each block of the body-fitted multi-block grid is associated with a structured quadrilateral grid with boundary geometry defined by possibly up to four body-fitted curved boundaries. The latter are represented as piecewise Lagrange polynomial splines of an order consistent with that of the finite-volume scheme. Interior inter-block boundaries are always represented by straight line segments and only block boundaries associated with the boundaries of the physical domain are represented by the body-fitted splines. Schematic representations of geometric elements of the grid encountered at block boundaries of a body-fitted quadrilateral mesh are depicted in Fig. 1(a). The standard quadrilateral computational cells, used for all interior cells and non-curved boundaries, as well as the special curved-boundary cell, introduced here specifically for the high-order treatment of curved physical boundaries, are both shown in Figs. 1(b) and 1(c), respectively. The introduction of the latter is a distinctive feature of the present method compared to previous lower-order implementations on similar meshes [67–69,76,77]. These generalized computational cells are introduced to represent in an accurate fashion the curved physical boundaries of the mesh. Note that the connector between any two adjacent mesh nodes on the curved boundary is a piecewise-smooth curve defined by the boundary splines of the grid block.

3.1. Semi-discrete form

The proposed finite-volume solution method starts with the integral form of Eqs. (6) and (7) applied to the two-dimensional polygonal control volume representing the computational cell (i,j) as shown in Fig. 1 and then makes use of the divergence theorem to arrive at the following semi-discrete form:

$$\frac{\mathrm{d}}{\mathrm{d}t}\bar{\mathbf{U}}_{i,j}(t)=-\frac{1}{A_{i,j}}\sum_{l=1}^{N_f}\sum_{m=1}^{N_G}(\omega\vec{\mathbf{F}}\cdot\vec{n}\,\Delta\ell)_{i,j,l,m}+\bar{\mathbf{S}}_{i,j}=\mathbf{R}_{i,j}(\bar{\mathbf{U}}), \tag{8}$$

where

$$\bar{\mathbf{U}}_{i,j}=\frac{1}{A_{i,j}}\iint_{\mathcal{A}_{i,j}}\mathbf{U}\,\mathrm{d}x\,\mathrm{d}y, \tag{9}$$

is the average value of the conserved solution vector for cell (i,j),

$$\bar{\mathbf{S}}_{i,j}=\frac{1}{A_{i,j}}\iint_{\mathcal{A}_{i,j}}\mathbf{S}(\mathbf{U})\,\mathrm{d}x\,\mathrm{d}y, \tag{10}$$

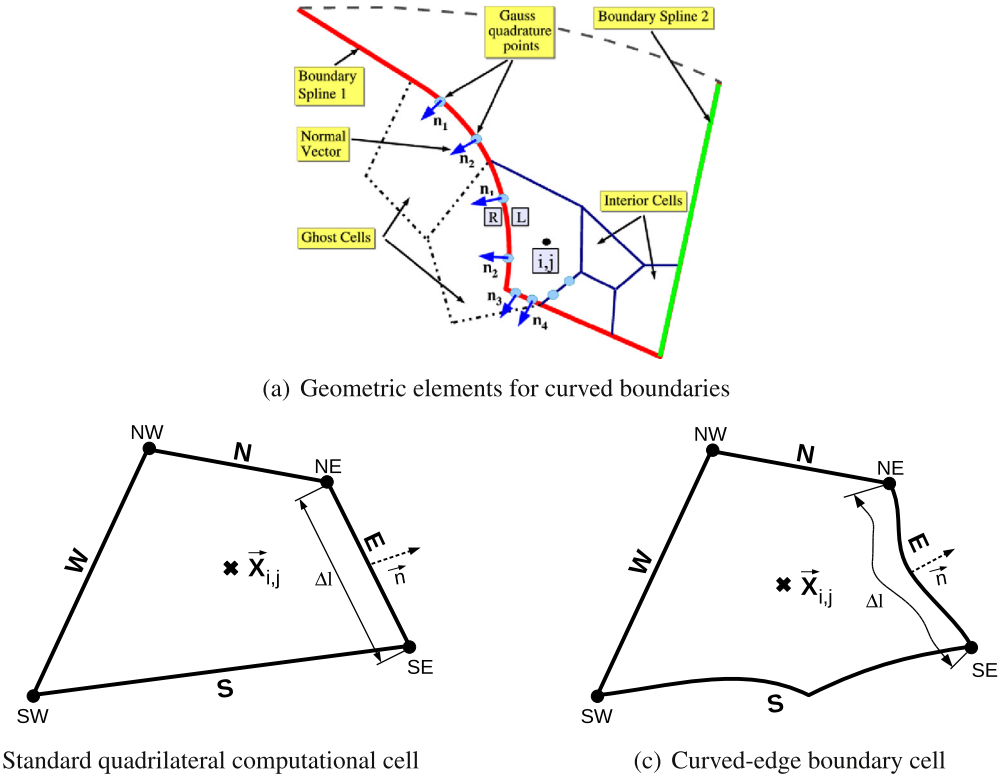

(a) Geometric elements for curved boundaries

(b) Standard quadrilateral computational cell

(c) Curved-edge boundary cell

Fig. 1. Features of multi-block, body-fitted, quadrilateral mesh showing: (a) schematic representation of geometric elements of the grid at the block boundaries of body-fitted mesh; (b) standard straight-edge quadrilateral computational cell; and (c) special curved-edge boundary cell for high-order treatment of curved physical boundaries considered in this study. Cardinal directions (i.e., N, S, E, W) relative to the element centroid, $\vec{X}_{i,j}$, are used to identify each face and node of the computational cells. Outward normal vectors, $\vec{n}$, and lengths, $\Delta\ell$, of the cell edges/faces are also shown.

is the corresponding average value of the source vector, $\mathbf{R}_{i,j}$ is the so-called residual vector, and $A_{i,j}$ is the cell area. An N_G-point Gaussian quadrature integration procedure is used to evaluate the solution flux along each of the N_f faces of the cell, where ω is the quadrature weighting coefficient and $\Delta\ell$ and $\vec{n}$ are the length of the cell face and unit vector normal to the cell face or edge, respectively.

The coupled non-linear system of ordinary differential equations (ODEs) given by Eq. (8) describe the time evolution of $\overline{\mathbf{U}}$. The proposed finite-volume scheme involves three main steps when integrating the ODEs forward in time: (i) reconstruction in which an approximation for $\mathbf{U}(\vec{X})$ is obtained within each computational cell; (ii) residual evaluation in which the residual vector, $\mathbf{R}$, is evaluated for each cell based on the reconstructed solution; and (iii) time integration, in which a time-marching scheme is used to advance the solution to the next time level based on the residual. Details about each of these three steps, with an emphasis on approaches used to ensure high-order spatial accuracy, now follow.

3.2. High-order residual evaluation

High-order solutions of Eq. (8) are sought here by evaluating the residual vector, $\mathbf{R}_{i,j}$, to high order, based on high-order reconstruction, an upwind discretization of the hyperbolic flux, and a centrally weighted discretization of the elliptic flux. A high-order evaluation of the net flux through the boundary of a computational cell first requires the selection of the number of Gauss quadrature points, N_G, required for the evaluation of the integrated cell face flux, $\vec{\mathbf{F}}\cdot\vec{n}$, to the desired level of accuracy [13]. The sum of the non-linear hyperbolic and elliptic fluxes, $\vec{\mathbf{F}}\cdot\vec{n}=\vec{\mathbf{F}}_{\mathrm{H}}(\mathbf{U})\cdot\vec{n}+\vec{\mathbf{F}}_{\mathrm{E}}(\mathbf{U},\vec{\nabla}\mathbf{U})\cdot\vec{n}$, must be estimated at each quadrature point, m, of a cell face, l, with a flux function that approaches the *true* flux at the rate imposed by the expected order of accuracy in the asymptotic limit of infinitely small mesh size. This requirement translates to a similar set of conditions on the accuracy of the solution vector, $\mathbf{U}$, and its gradient, $\vec{\nabla}\mathbf{U}$, used in calculating the hyperbolic and elliptic fluxes. In addition, the flux function should produce a dissipative and stable scheme.

The hyperbolic flux at each cell interface is determined here using an upwind flux formulation as originally pioneered by Godunov [78]. The flux is based on the solution of a local Riemann problem with discontinuous initial data defined by left and right solution quantities, $\mathbf{U}_l$ and $\mathbf{U}_r$ [79,80]. The hyperbolic numerical flux at each quadrature point is thus given by

$$\vec{\mathbf{F}}_{\mathrm{H}}\cdot\vec{n}=\mathcal{F}_{\mathrm{H}}(\mathbf{U}_l,\mathbf{U}_r,\vec{n}), \tag{11}$$

where the flux function $\mathcal{F}_H$ returns the solution to the Riemann problem in a direction defined by the cell face normal, $\vec{n}$, and with left and right solution values, $\mathbf{U}_l$ and $\mathbf{U}_r$. Both exact and approximate Riemann solvers can be used to solve the Riemann problem and evaluate the numerical flux [79,80]. Details of the flux functions considered herein are given in Sections 3.4 to follow. The left and right solution states, $\mathbf{U}_l$ and $\mathbf{U}_r$, are determined by performing cell-centred, piecewise, k-exact, polynomial reconstruction within each computational cell. As the truncation error for k-order exact reconstruction is $\mathcal{O}(\Delta x^{k+1})$, a $(k+1)$-order accurate spatial discretization can be achieved for smooth hyperbolic problems.

Numerical values for the elliptic fluxes, $\vec{\mathbf{F}}_E \cdot \vec{n}$, must also be evaluated at the quadrature points of each cell face. In general, the elliptic fluxes will have the form

$$\vec{\mathbf{F}}_E \cdot \vec{n} = \mathcal{F}_E(\mathbf{U}, \vec{\nabla}\mathbf{U}, \vec{n}), \tag{12}$$

where $\mathcal{F}_E$ is the corresponding elliptic flux function that returns the flux in direction $\vec{n}$ given the cell interface values of the solution, $\mathbf{U}$, and its gradient $\vec{\nabla}\mathbf{U}$. In the proposed scheme, the unlimited, cell-centred, k-exact, polynomial reconstruction used in the hyperbolic flux evaluation is re-used (to avoid considerable additional computational effort) to obtain a $(k+1)$-order accurate value of the interface solution, $\mathbf{U}$. Direct differentiation of this k-exact reconstructed solution is also then used to obtain a k-order accurate value of the interface gradient, $\vec{\nabla}\mathbf{U}$ (i.e., one order less accurate due to differentiation).

In this way, a consistent spatial discretization scheme is obtained for both hyperbolic and elliptic fluxes. The overall accuracy of the proposed scheme is k-order accurate on arbitrary and irregular meshes (including meshes resulting from the application of AMR). This is achieved through the use of k-exact solution reconstruction (i.e., reconstruction with $(k+1)$-order truncation error for the solution itself and, hence, the hyperbolic fluxes) without relying on regularity of the grid and error cancellation. Most finite-volume methods generally use $(k-1)$-exact solution reconstruction to achieve a k-order accurate scheme. The use of k-exact reconstruction is particularly helpful in ensuring uniform accuracy on irregular and solution adapted mesh. Appropriate choices for the evaluation of $\mathbf{U}$, $\vec{\nabla}\mathbf{U}$, and the elliptic flux function, $\mathcal{F}_E$, at the cell face are discussed below in Section 3.5.

The number of Gauss quadrature points used in the integration of the numerical flux on each cell face is selected to match the order of the reconstruction. Recall that an N_G-point Gaussian quadrature rule integrates polynomials of degree $2N_G - 1$ exactly, providing a $2N_G$-order accurate formula [81]. To obtain a consistent k-order accurate scheme for both hyperbolic and elliptic operators, with $(k+1)$-order solution reconstruction and k-order gradient representation following from k-exact reconstruction, requires a k-order accurate quadrature rule. Hence, the practice adopted here is to use one quadrature point ($N_G = 1$) for second-order schemes (piecewise quadratic, $k = 2$, reconstruction) and two quadrature points ($N_G = 2$) for third- and fourth-order schemes (piecewise cubic, $k = 3$, and quartic, $k = 4$, reconstruction). Note that piecewise linear, $k = 1$, reconstruction with one quadrature point is also appropriate for a second-order treatment of the hyperbolic fluxes, but $(k = 1)$-exact reconstruction is problematic for the evaluation of elliptic fluxes. For the latter, second-order accuracy is only formally achieved by relying on error cancellation on uniform regular grids.

Lastly, Gaussian quadrature is again used to evaluate the cell-average value of the source vector, $\bar{\mathbf{S}}_{i,j}$, and its contribution to the solution residual for each quadrilateral cell as defined by Eq. (10) above. As with the cell face flux integration, k-exact reconstruction requires a k-order accurate quadrature rule for the source vector, which in this case involves a two-dimensional integral. Thus, for piecewise quadratic, $k = 2$, reconstruction, a single quadrature point is used in evaluating $\bar{\mathbf{S}}_{i,j}$ whereas, for piecewise cubic, $k = 3$, and quartic, $k = 4$, reconstruction, four quadrature points (two quadrature points per dimension) are used in computing the cell-average source vector.

3.3. High-order CENO reconstruction

As with other finite-volume methods, a piecewise polynomial approach is used to perform the solution reconstruction within each computational cell. However, rather than selecting or weighting reconstructions from multiple stencils, a hybrid solution reconstruction procedure is used that combines the high-order k-exact least-squares reconstruction technique of Barth [12], based on a *fixed central stencil*, with a monotonicity preserving limited piecewise linear least-squares reconstruction algorithm [12]. The central stencil is optimal in the sense that it generally provides the most accurate reconstruction. In case of unstructured meshes, the central stencil be interpreted as the stencil that includes all nearest neighbour cells. The limited reconstruction procedure is applied only to computational cells with under-resolved solution content, thereby avoiding undesirable solution oscillations. The unlimited k-exact reconstruction scheme is used for all cells in which the solution is fully resolved. Switching in the hybrid procedure is determined by a solution smoothness indicator. Accuracy to any order is possible by simply expanding the support for the cell-centred reconstruction. Direct benefits and/or advantages of the proposed reconstruction procedure are two-fold: (i) reconstruction on multiple stencils, which in some cases can lead to poorly conditioned linear systems, is avoided; and (ii) the same fixed stencil is used for each variable and thus solution of the least-squares problem associated with reconstruction can be made very efficient. All of this also makes the hybrid CENO algorithm well suited for application to unstructured mesh. As will be shown, mesh adaptation can also be directed based on the ability of the scheme to differentiate between resolved and under-resolved or non-smooth solution content.

For hyperbolic conservation laws, the CENO reconstruction leads to a finite-volume scheme that is high-order accurate for smooth solutions, even at extrema, and avoids the appearance of $O(1)$ numerical oscillations in under-resolved regions

and for solutions containing strong discontinuities and/or shocks. However, the formal ENO property of uniform accuracy is lost for non-smooth solutions. Note that Harten and Chakravarthy [82] previously also proposed a technique to obtain an ENO-like reconstruction on a fixed central stencil by hybridizing a high-order reconstruction with a first-order formulation. Switching in their proposed hybrid scheme was based on undivided differences and the TVD property [83] and not directly on the smoothness of the reconstructions. More recently, Haselbacher [10] explored the use of fixed stencil central reconstructions in the formulation of WENO schemes for unstructured mesh. Haselbacher's approach, is quite different to the current approach and schemes of accuracy greater than second order (i.e., $k = 1$) were not considered.

The proposed reconstruction scheme depends intimately on the average value of the solution within each cell of the stencil. For non-linear systems of PDEs, the reconstruction can be theoretically applied to any of the dependent variables, regardless of their physical meaning and whether or not they are conserved, primitive, or characteristic quantities. Nevertheless, as the cell-averaged values of the conserved solution variables (e.g., $\mathbf{U} = [\rho, \rho v_x, \rho v_y, \rho e]^{\mathrm{T}}$ for the Navier–Stokes equations) are updated directly at each time level or iteration in a time-marching scheme, it would seem natural to perform the reconstruction using the conserved variables. While Zhang and Shu [84,85] have recently developed a rather elegant approach that allows the positivity-preserving high-order reconstruction of conserved variables and this approach has been applied to non-linear systems of conservation laws, including the equations governing inviscid compressible flows [85,86] and ideal magnetohydrodynamics [86], in many cases it is often preferable to perform the reconstruction in terms of other derived or so-called primitive variables (e.g., $\mathbf{W} = [\rho, v_x, v_y, p]^{\mathrm{T}}$ for the Navier–Stokes equations) in order to control more effectively monotonicity, positivity, and realizability of key physical solution values. This is the approach adopted here. To convert average conserved into average primitive solution values and/or vice-versa, a simple and straightforward point-wise mapping is used. A potential weakness of the latter is that the use of a point-wise mapping does not provide an exact conversion between the cell-averaged conservative and primitive variables when the relationship between variables is non-linear. Conversion errors are theoretically introduced. Recently, McCorquodale and Colella [20] proposed a mapping which provides a fourth-order accurate conversion between the cell-averaged values of conserved and primitive quantities. While a similar procedure is not considered here, numerical experiments, the results of which are in Section 5.6 to follow, have demonstrated that the errors introduced by the point-wise conversion do not detrimentally affect the order of accuracy of the proposed CENO scheme.

3.3.1. k-exact least-squares reconstruction

In piecewise k-exact polynomial reconstruction [12], it is assumed that a solution variable, u, at any location, $\vec{X}$, in computational cell (i, j) of a two-dimensional space has the general form

$$u^k_{i,j}(\vec{X}) = \sum_{\substack{p_1=0 \\ (p_1+p_2\leqslant k)}}^{k} \sum_{p_2=0}^{k} (x - \bar{x}_{i,j})^{p_1} (y - \bar{y}_{i,j})^{p_2} D^k_{p_1 p_2}, \tag{13}$$

where (x, y) are the Cartesian coordinates of the point of interest, $(\bar{x}_{i,j}, \bar{y}_{i,j})$ are the coordinates of the cell centroid, $\vec{X}_{i,j}$, k is the order of the piecewise polynomial interpolant, the summation indices p_1 and p_2 satisfy the condition that $p_1 + p_2 \leqslant k$, and $D^k_{p_1 p_2}$ are the coefficients of the k-exact polynomial approximation. The latter are in general functions of the mean or average solution, $\bar{u}_{i,j}$, within the cell and its neighbours. For a two-dimensional reconstruction, the number of coefficients, $\mathcal{N}_D$, for a particular order, k, is given by $\mathcal{N}_D = (k + 1)(k + 2)/2$. For example, cubic reconstruction, $k = 3$, requires the specification of $\mathcal{N}_D = 10$ polynomial coefficients for each solution variable in terms of cell-averaged solution values.

In determining $D^k_{p_1 p_2}$, the following conditions should be satisfied by the k-exact reconstruction procedure:

- the reconstruction procedure must reproduce exactly polynomials of degree $N \leqslant k$;
- the reconstruction must preserve the mean or average value within the computational cell; and
- the reconstruction must have compact support.

The first condition is equivalent to

$$u^k_{i,j}(\vec{X}) - u_{\mathrm{exact}}(\vec{X}) = \mathcal{O}\left(\Delta x^{k+1}\right), \tag{14}$$

which is assumed to hold anywhere in the vicinity of cell (i, j). The second condition requires the integral of the piecewise polynomial approximation to recover the cell average

$$\bar{u}_{i,j} = \frac{1}{A_{i,j}} \iint\limits_{\mathcal{A}_{i,j}} u^k_{i,j}(\vec{X})\, \mathrm{d}x\, \mathrm{d}y. \tag{15}$$

Finally, the third condition involves the number and locality of neighbouring solution values used in the supporting stencil. The minimum size of the compact stencil is determined by the number of required unknown coefficients, but in practice, additional neighbours are included to make the reconstruction more robust in the presence of stretched meshes and solution

gradients not aligned with the mesh. For body-fitted quadrilateral mesh of interest here, fixed central stencils which include the 8 nearest neighbour cells for $k=1$ and 24 nearest neighbours for $k=2$, 3, and 4 are used.

The coefficients, $D^k_{p_1p_2}$, can be computed by assuming that the k-exact polynomial for $u^k_{i,j}(\vec{X})$ given in Eq. (13) above extends to all computational cells within the supporting stencil. This results in an overdetermined system of linear equations for $D^k_{p_1p_2}$ which can then be solved via a least-squares approach. The overdetermined linear system may be written as

$$\mathbf{A}\mathbf{x}-\mathbf{b}=\mathbf{e}, \tag{16}$$

where the left-hand-side coefficient matrix, $\mathbf{A}$, depends only on the mesh geometry and can be largely pre-computed and re-used. The right-hand-side vector, $\mathbf{b}$, contains the average values of the solution values in each cell and $\mathbf{x}$ is a vector containing the desired coefficients for the k-exact reconstruction. The vector, $\mathbf{e}$, represents the error associated with the least-squares solution, which is minimized in a least-squares sense. To localize and improve the accuracy of the reconstruction for stretched meshes with surface curvature, geometric weights of the form $w_{\gamma,\delta}=|\Delta\vec{X}_{\gamma,\delta}|^{-\theta}$ are applied to the equation for each neighbouring control volume, (γ,δ), where $\Delta\vec{X}_{\gamma,\delta}=\vec{X}_{\gamma,\delta}-\vec{X}_{i,j}$ and appropriate values for the exponent θ are either one or two [87]. The preservation of the average value, $\bar{u}_{i,j}$, within the reconstructed cell can be either explicitly enforced by expressing the coefficient, D^k_{00}, as a function of the other unknowns or by solving a linear equality constrained least-squares problem with the full set of equations. The former technique gives rise to a computationally more efficient implementation and is recommended here.

Both Householder QR factorization and singular value orthogonal decomposition (SVD) can be used to solve the weighted least-squares problem of Eq. (16) [88]. In general, the latter is particularly advantageous when used in the proposed CENO reconstruction scheme. The SVD approach permits the computation of a pseudo-inverse matrix, $\mathbf{A}^\dagger$, with which the solution to the least-squares problem is then given by the matrix–vector product $\mathbf{x}=\mathbf{A}^\dagger\mathbf{b}$ [89]. 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. The repeated evaluation of the pseudo inverse is avoided and the computation of the coefficients associated with the reconstruction polynomials then reduces to the evaluation of simple matrix–vector products, which can be performed rather efficiently. Storage and re-use of $\mathbf{A}^\dagger$ was found to reduce significantly the computational costs of performing the high-order reconstructions. Although requiring some additional storage, the added memory requirements are not that substantial and are generally readily available on most modern distributed memory architectures.

One concern with k-exact reconstruction is that the stencil size and hence the size of the least-squares problem grow with the order of the scheme. This can potentially lead to issues with the conditioning and/or invertibility of the least-squares problem. Aside from the stencil size and desired order of the polynomial reconstruction, the conditioning of the least-squares system will in general be very dependent on features of the mesh, such as the cell size or spacing, cell aspect ratio, and irregularity of the mesh topology. Recently, Jalali and Ollivier-Gooch [90] have shown that a rather simple column-scaling procedure can be applied to the least-squares problem to improve the conditioning and make it independent of the mesh size and aspect ratio. This column-scaling conditioning of the least-squares problem can be introduced by defining a diagonal column-scaling matrix, $\mathbf{D}_C$, with diagonal entries that are equal to the inverse of the maximum absolute value of the entries of each column of matrix $\mathbf{A}$. Post multiplication of $\mathbf{A}$ by $\mathbf{D}_C$ results in the product, $\mathbf{A}\mathbf{D}_C$, the entries of which are all in the range $[-1,1]$. As a result, the conditioning of this modified reconstruction matrix is significantly improved compared to the original. With the column scaling, the modified least-squares problem can be expressed as

$$\mathbf{A}\mathbf{D}_C\mathbf{D}_C^{-1}\mathbf{x}-\mathbf{b}=\mathbf{A}\mathbf{D}_C\mathbf{z}-\mathbf{b}=\mathbf{e}, \tag{17}$$

and the solution for the reconstruction coefficients, $\mathbf{x}$, is then obtained from the scaled solution of the least-squares problem, $\mathbf{z}$, by using $\mathbf{x}=\mathbf{D}_C\mathbf{z}$. In the case that the SVD approach is used to solve the least-squares problem, the solution is then given simply by $\mathbf{x}=\mathbf{D}_C(\mathbf{A}\mathbf{D}_C)^\dagger\mathbf{b}$ where the matrix $\mathbf{D}_C(\mathbf{A}\mathbf{D}_C)^\dagger$ can be pre-computed and stored for re-use.

To examine these conditioning issues, the condition number, $\kappa(\mathbf{A})$, of the coefficient matrix, $\mathbf{A}$, for least-squares problems corresponding to different reconstruction orders, k, on several representative quadrilateral meshes was analyzed. Table 1 provides the maximum values of the computed condition number as defined by the L_∞-norm of $\mathbf{A}$ and its pseudo-inverse $\mathbf{A}^\dagger$ given by

$$\kappa_\infty(\mathbf{A})=\|\mathbf{A}\|_\infty\left\|\mathbf{A}^\dagger\right\|_\infty, \tag{18}$$

for linear ($k=1$), quadratic ($k=2$), cubic ($k=3$) and quartic ($k=4$) solution reconstructions on Cartesian grids with aspect ratio, a_r, of 1, 240, and 14,012, respectively, as well as on two quadrilateral meshes: one a smooth regular mesh and the other an irregular mesh, having the geometry depicted in Fig. 4. Quartic reconstruction, which produces a 4th-order scheme for both hyperbolic and elliptic fluxes, was deemed to be the highest order of interest and thus reconstruction orders up to $k=4$ were examined. The effect of different geometric weighting formulations, $w_{\gamma,\delta}$, on the condition number was also considered, with both inverse ($\theta=1$) and inverse squared ($\theta=2$) weighting examined. In addition, the effectiveness of the column-scaling procedure outlined above is also assessed. Results for all of these cases are given in the table.

Table 1 reveals that, without the application of weighting or scaling, the condition number increases with an increase in the reconstruction order as expected. Additionally, larger values for the condition number are observed as the mesh deviates from a regular square Cartesian grid, either through an increase in the aspect ratio or irregularities in the mesh. The largest

Table 1
The condition numbers of the coefficient matrix, **A**, obtained based on L_∞-norm of the matrix for different geometric-weighting and column-scaling strategies and for several representative meshes: Cartesian grids with aspect ratio, a_r, of 1, 240 and 14,012; as well as both a regular smooth mesh and an irregular grid for the Ringleb flow geometry as defined in Section 5. The irregular grid is depicted in Fig. 5(a).

Mesh	Geometric weighting $w_J = \lvert\Delta\vec{X}_{IJ}\rvert^{-\theta}$	Column scaling applied	$\kappa_\infty(\mathbf{A})$			
			$k=1$	$k=2$	$k=3$	$k=4$
Cartesian grid	$\theta=0$	no	2.0	25.4	452.2	6748.4
$a_r=1$	$\theta=1$	no	1.7	22.0	273.2	3223.8
$\Delta X=0.0965$	$\theta=2$	no	1.3	32.6	335.3	4007.2
	$\theta=1$	yes	1.7	5.3	15.1	49.3
Cartesian grid	$\theta=0$	no	227.6	292,335.1	5.71×10^{8}	9.69×10^{11}
$a_r=240$	$\theta=1$	no	1.8	1116.6	1.72×10^{6}	2.01×10^{9}
$\Delta X=0.0965$	$\theta=2$	no	239.0	245,311.9	3.15×10^{8}	4.03×10^{11}
	$\theta=1$	yes	1.8	16.6	83.7	395.3
Cartesian grid	$\theta=0$	no	12,780.4	8.40×10^{8}	8.78×10^{13}	1.10×10^{14}
$a_r=14,012$	$\theta=1$	no	1.8	60,118.4	4.66×10^{9}	9.17×10^{13}
$\Delta X=0.0965$	$\theta=2$	no	15,047.1	8.32×10^{8}	5.40×10^{13}	2.78×10^{14}
	$\theta=1$	yes	1.8	17.4	1357.1	4246.6
Regular	$\theta=0$	no	4.1	172.9	9232.4	524,247.3
smooth grid	$\theta=1$	no	2.2	77.5	3815.2	171,418.2
	$\theta=2$	no	3.2	141.9	4928.2	192,885.6
	$\theta=1$	yes	3.2	18.2	105.0	872.9
Irregular grid	$\theta=0$	no	5.2	191.7	10,552.5	699,826.2
	$\theta=1$	no	2.4	84.1	4661.8	199,873.8
	$\theta=2$	no	4.6	203.2	6688.5	240,616.7
	$\theta=1$	yes	3.9	21.9	134.7	1096.5

values for the condition number were encountered on the Cartesian grid with an aspect ratio of $a_r = 14{,}012$. The results of Table 1 also reveal that some improvements in the conditioning of the systems can be obtained through the use of geometric weighting. For the meshes considered herein, simple inverse-distance geometric weighting (i.e., $w_{\gamma,\delta} = |\Delta\vec{X}_{\gamma,\delta}|^{-1}$) provided lower values of the condition number for all reconstruction orders, as compared to the inverse-distance-squared weighting. More importantly, it is also very clear that the column-scaling conditioning procedure effectively eliminates the poor conditioning of the least-squares problems and makes the condition numbers effectively independent of the mesh size and aspect ratio. Even for the Cartesian meshes with a high aspect ratio of $a_r = 14{,}012$, the conditioning number is reduced from 1.10×10^{14} to 4246.6 through a combination of inverse-distance weighting and column-scaling. Based on these observations, it is felt that the conditioning of the least-squares problem can be dealt with, at least for reconstruction orders up to four, but caution should still be exercised when considering either higher than quartic reconstruction and/or highly irregular grids. Finally, it should be noted that for all of the numerical results pertaining to the validations cases and flow problems to follow, column-scaling conditioning was not required to achieve the desired computational precision; however, it was used in a few of the computations, including those for the NACA0012 airfoil, to ensure that full accuracy was achieved in the high-order reconstructions.

3.3.2. Monotonicity enforcement via smoothness indicator

Solution monotonicity is preserved in the hybrid scheme by reverting the high-order k-exact reconstruction to a limited piecewise linear ($k=1$) reconstruction where necessary. In the present work, a standard MUSCL-type limited linear reconstruction procedure [91] is used along with the slope limiters of Barth–Jespersen [12,123] and Venkatakrishnan [92]. This piecewise limited linear scheme is only applied in regions of the solution which are deemed as being non-smooth and fully replaces the high-order k-exact reconstruction locally only in these appropriately identified regions. In order to detect regions where the order of the reconstruction should be reduced and limited linear reconstruction applied, a smoothness indicator is computed for every variable individually within each cell once the unlimited k-exact reconstruction has been performed. The smoothness indicator is then used in the manner described below to ensure that limited linear reconstruction is applied only to cells with under-resolved and/or non-smooth solution content and the unlimited k-exact reconstruction is retained elsewhere.

The form of the smoothness indicator proposed here was inspired by the definition of multiple-correlation coefficients that are often used in evaluating the accuracy of curve fits [88]. The idea is to assess how accurately the cell-centred polynomial expansion represents the solution within the reconstruction stencil. This is achieved by comparing reconstructed solutions in neighbouring cells. In particular, the smoothness indicator, $\mathcal{S}$, is calculated in terms of a smoothness parameter, α, as well as in terms of the number of unknowns (degrees of freedom), $\mathcal{N}_{\mathrm{DOF}}$, and stencil size, $\mathcal{N}_{\mathrm{SOS}}$, for the reconstruction. The smoothness indicator is taken to have the form

$$\mathcal{S} = \frac{\alpha}{\max((1-\alpha),\epsilon)}\frac{(\mathcal{N}_{\mathrm{SOS}}-\mathcal{N}_{\mathrm{DOF}})}{(\mathcal{N}_{\mathrm{DOF}}-1)}, \tag{19}$$

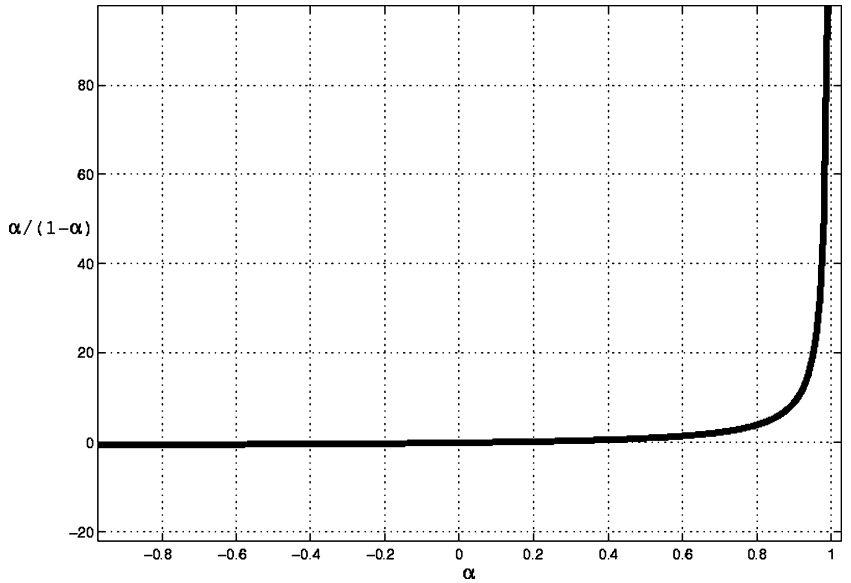

Fig. 2. The graph of $f(\alpha) = \frac{\alpha}{(1-\alpha)}$ illustrating the behaviour of the solution smoothness indicator as a function of the smoothness parameter, α.

where α is determined as follows

$$\alpha = 1 - \frac{\sum_{\gamma}\sum_{\delta}\left(u^k_{\gamma,\delta}(\vec{X}_{\gamma,\delta}) - u^k_{i,j}(\vec{X}_{\gamma,\delta})\right)^2}{\sum_{\gamma}\sum_{\delta}\left(u^k_{\gamma,\delta}(\vec{X}_{\gamma,\delta}) - \bar{u}_{i,j}\right)^2}, \tag{20}$$

and where the ranges of the indices, γ and δ, are taken to include all control volumes in the reconstruction stencil for cell (i, j), $\vec{X}_{\gamma,\delta}$ is the centroid of the cell (γ, δ), and the tolerance, ϵ, has been introduced in order to avoid division by zero. A suitable value for ϵ has been found to be 10^{-8}.

It should be evident that the parameter, α, is based on a comparison of the values of reconstructed solutions at the centroids of neighbouring cells used in the reconstruction procedure for cell (i, j). Additionally, the average value of the solution in the cell is used in normalizing the differences in the reconstructed solution values. The range for α is $-\infty < \alpha \leqslant 1$ and it rapidly approaches unity as the mesh is refined and/or the solution becomes smooth such that near perfect matching of the reconstructions in adjacent cells is achieved. To see this, first consider the numerator of the second term appearing in Eq. (20) for the smoothness parameter, α. 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 the numerator has the form

$$u^k_{\gamma,\delta}(\vec{X}_{\gamma,\delta}) - u^k_{i,j}(\vec{X}_{\gamma,\delta}) \approx \mathcal{O}\left(\Delta x^{k+1}\right). \tag{21}$$

Similarly, it can be shown that the denominator of the second term has the form

$$u^k_{\gamma,\delta}(\vec{X}_{\gamma,\delta}) - \bar{u}_{i,j} \approx \mathcal{O}(\Delta x). \tag{22}$$

From this it follows that

$$\alpha \approx 1 - \mathcal{O}\left(\Delta x^{2k}\right), \tag{23}$$

for smooth resolved solution content. In this case, it is evident that $\alpha \to 1$ 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 the reconstruction of non-smooth data, it is expected that

$$u^k_{\gamma,\delta}(\vec{X}_{\gamma,\delta}) - u^k_{i,j}(\vec{X}_{\gamma,\delta}) \approx \mathcal{O}(1), \qquad u^k_{\gamma,\delta}(\vec{X}_{\gamma,\delta}) - \bar{u}_{i,j} \approx \mathcal{O}(1), \tag{24}$$

and therefore α will be generally far from a value of one. The variation of $\alpha/(1-\alpha)$ is depicted in Fig. 2. The figure shows that $\mathcal{S}$ rapidly becomes large as α approaches unity. Also, by definition, $\mathcal{S}$ is invariant to scaling of the solution.

Once evaluated, the smoothness indicator, $\mathcal{S}$, is then compared to a pass/no-pass cutoff value, $\mathcal{S}_c$. An appropriate value for the cutoff has been determined from a wide range of numerical experiments. Values for $\mathcal{S}_c$ in the range 1000–5000 work well. For smooth solutions, the actual value of the smoothness indicator is typically orders of magnitude larger (e.g.,

$\sim$100,000). For non-smooth solutions (i.e., the interior of discontinuities), typical values for $\mathcal{S}$ are close to 10. For cells with $\mathcal{S} < \mathcal{S}_c$, the solution is deemed to be under-resolved and/or non-smooth and the high-order k-exact reconstruction is replaced by the limited linear reconstruction. In this case, the stencil is also reduced and includes just the 8 nearest neighbours so as to retain a compact stencil for the treatment of shocks and discontinuities. Note that this is only done for the hyperbolic flux evaluation. The k-exact reconstruction is retained in the evaluation of the elliptic fluxes. For $\mathcal{S} > \mathcal{S}_c$, the unlimited k-exact reconstruction is deemed acceptable and retained in the evaluation of both the hyperbolic and elliptic fluxes. Note that with this definition of the smoothness indicator, the effect of reconstructions containing discontinuous data is reflected in the smoothness indicator of just a few neighbouring cells and, consequently, solution discontinuities are typically contained within 7 to 10 cells. It is also worthwhile noting that while ENO and WENO schemes can suffer a loss in accuracy from the rapid switching from one-sided to centred stencils, the present hybrid scheme as directed by the smoothness indicator does not suffer from this effect as the high-order reconstruction is based only on the centred and most accurate stencil.

3.3.3. Reconstruction at boundaries and implementation of high-order boundary conditions

Consistent high-order treatment of boundary conditions is a crucial element for high-order methods. In order to retain high accuracy at boundaries, all geometric data (i.e., cell area, centroid, geometric moments, normals, edge lengths, locations of the Gauss quadrature integration points) must be evaluated to the same order of accuracy as that of the interior scheme and an appropriate high-order representation of the boundary geometry is required. In addition, a high-order prescription for solution values at the boundary based on the boundary equations in needed. One approach to imposing high-order boundary data is to make use of extra rows of ghost cells which are added beyond the geometric boundary of the computational domain. Solution values are then imposed in the ghost cells in such a way that the reconstructed solution values and/or fluxes at the boundary approximates the required boundary condition. An alternative approach is to enforce the boundary conditions by constraining the least-squares reconstruction in control volumes adjacent to the boundary as described by Ollivier-Gooch and Van Altena [47]. Constrained reconstruction is generally preferred here. Nevertheless, for some boundary data, this approach may prove difficult and therefore both procedures (i.e., ghost cell and constrained reconstruction treatments) have been implemented and are used. The present implementation of constrained reconstruction is summarized below.

Boundary condition prescription via constrained reconstruction is more accurate and straightforward for boundary geometries that are curved. In this case, the boundary data is enforced using only information from the interior of the computational domain and by the imposition of boundary conditions at the Gaussian quadrature integration points of cell faces lying along the boundary of the domain. This yields additional constraints on the cell-centred reconstruction procedure for interior cells adjacent to the boundary. By constraining the least-squares reconstruction in control volumes near boundaries, complex boundary data can be enforced. Two basic types of boundary constraints have been developed here: (i) general Robin boundary conditions (i.e., linear combinations of Dirichlet and Neumann boundary data) which can be applied individually to any variable; and (ii) linear boundary relations between variables which can be applied as coupled constraints to a set of reconstructed variables.

Robin boundary data for dependent variable, u, in terms of a k-degree reconstruction polynomial, $u^k(\vec{X})$, at a given quadrature point, $\vec{X}_g$, is imposed using a general constraint which can be expressed as

$$\begin{aligned} f(\vec{X}_g) &= \left[a(\vec{X})u(\vec{X}) + b(\vec{X})\frac{\partial u(\vec{X})}{\partial n} \right]_{\vec{X}_g} \\ &= \sum_{\substack{p_1=0 \\ (p_1+p_2\leqslant k)}}^{k} \sum_{p_2=0}^{k} \left\{ \Delta X_g^{(p_1-1)} \Delta Y_g^{(p_2-1)} \left[a\Delta X_g \Delta Y_g + bp_1 \Delta Y_g n_x^g + bp_2 \Delta X_g n_y^g \right] \right\} D^k_{p_1 p_2} \end{aligned} \tag{25}$$

where $a(\vec{X})$ and $b(\vec{X})$ are the coefficients defining the contributions of the Dirichlet and Neumann components to the boundary value, f, respectively, $\Delta X_g = x_g - \bar{x}_{i,j}$, $\Delta Y_g = y_g - \bar{y}_{i,j}$, and $\vec{n}^g(n_x^g, n_y^g)$ is the normal unit vector at the point $\vec{X}_g$. The constraint of Eq. (25) can be added to the system of linear equations making up the least-squares problem for a boundary cell and the resulting system can then be solved in a least-squares sense while ensuring that the constraint is exactly satisfied. The latter can be accomplished by applying Gaussian elimination with pivoting to eliminate the constraints and then obtaining a least-squares solution for the remaining equations in the overdetermined linear system.

To illustrate the imposition of boundary conditions via coupled linear constraints, consider the specification of boundary data for an inviscid solid wall or zero-shear slip boundary for which $\vec{V} \cdot \vec{n} = 0$, where $\vec{V}$ is the velocity vector. In two dimensions, this condition is imposed using coupled constrained reconstruction for velocity components, (u, v), and requiring that $\vec{V}$ is tangent to the boundary at every Gauss integration point. In this case, the coefficients of the reconstruction for the x- and the y-velocity components, $(D^k_{p_1 p_2})_u$ and $(D^k_{p_1 p_2})_v$ respectively, must satisfy the condition

$$\sum_{\substack{p_1=0 \\ (p_1+p_2\leqslant k)}}^{k} \sum_{p_2=0}^{k} \Delta X_g^{p_1} \Delta Y_g^{p_2} n_x^g \left(D^k_{p_1 p_2}\right)_u + \sum_{\substack{p_1=0 \\ (p_1+p_2\leqslant k)}}^{k} \sum_{p_2=0}^{k} \Delta X_g^{p_1} \Delta Y_g^{p_2} n_y^g \left(D^k_{p_1 p_2}\right)_v = 0. \tag{26}$$

Eq. (26) couples the linear systems defining the reconstructions for variables u and v within the boundary cell. Gaussian elimination can again be applied followed by solution of the coupled least-squares problem to determine the reconstructed velocity components at the slip boundary.

3.4. Inviscid (hyperbolic) flux evaluation

Given the cell-centred CENO reconstruction for each cell, the values of the solution flux can be calculated at each integration point. For the hyperbolic flux, upwinding is used, which for the advection–diffusion equation of Eq. (1), yields an expression for the numerical flux at each quadrature point given by

$$\vec{F}_{\mathrm{H}} \cdot \vec{n} = \mathcal{F}_{\mathrm{H}}(u_l, u_r, \vec{n}) = \begin{cases} u_l(\vec{V} \cdot \vec{n}) & \text{if } \vec{V} \cdot \vec{n} \geqslant 0, \\ u_r(\vec{V} \cdot \vec{n}) & \text{if } \vec{V} \cdot \vec{n} < 0, \end{cases} \tag{27}$$

where the left and right solution values, u_l and u_r, are the CENO reconstructed solution values to the left and right of the inter-cellular face, respectively. For the Navier–Stokes equations, upwinding of the hyperbolic flux, $\vec{\mathbf{F}}_{\mathrm{H}} \cdot \vec{n}$, at the quadrature points is achieved by using Eq. (11) where $\mathcal{F}_{\mathrm{H}}$ represents the solution of the Riemann problem for the inviscid form of the conservation equations in a direction, $\vec{n}$, normal to the face, and the left and right solution vectors, $\mathbf{U}_l$ and $\mathbf{U}_r$, are given by the CENO reconstructed solutions at the face of interest arising from the two cells sharing the interface. Both exact and approximate Riemann solvers can be used here to solve the Riemann problem and evaluate the numerical flux. In particular, the approximate Roe Riemann solver [93], HLLE and modified HLLE flux functions due to Linde [94–96], and exact Riemann solver of Gottlieb and Groth [79] have all been implemented and may be used.

3.5. Viscous (elliptic) flux evaluation

Numerical values of the elliptic fluxes are also required at each quadrature point. For the advection–diffusion equation, given left and right $(k+1)$-order accurate unlimited reconstructions, $u_l^k(\vec{X})$ and $u_r^k(\vec{X})$, for the scalar field, u, a k-order solution gradient at the cell face is obtained as the arithmetic mean of the left and right reconstructed gradients. A numerical approximation for the elliptic flux, $\vec{F}_{\mathrm{E}} \cdot \vec{n} = -\kappa\, \vec{\nabla} u \cdot \vec{n}$, at location $\vec{X}$ is then evaluated as

$$\vec{F}_{\mathrm{E}} \cdot \vec{n} = \mathcal{F}_{\mathrm{E}}(\mathbf{u}, \vec{\nabla}\mathbf{u}, \vec{n}) = -\kappa \left[\frac{1}{2} \left(\vec{\nabla} u_l^k(\vec{X}) + \vec{\nabla} u_r^k(\vec{X}) \right) \right] \cdot \vec{n}. \tag{28}$$

Similarly, elliptic fluxes for the Navier–Stokes system of Eq. (12) are evaluated as

$$\vec{\mathbf{F}}_{\mathrm{E}} \cdot \vec{n} = \mathcal{F}_{\mathrm{E}}(\mathbf{U}, \vec{\nabla}\mathbf{U}, \vec{n}) = \vec{\mathbf{F}}_{\mathrm{E}}(\mathbf{U}, \vec{\nabla}\mathbf{U}) \cdot \vec{n} = \vec{\mathbf{F}}_{\mathrm{E}} \left(\frac{1}{2} \left[\mathbf{U}_l^k(\vec{X}) + \mathbf{U}_r^k(\vec{X}) \right], \frac{1}{2} \left[\vec{\nabla} \mathbf{U}_l^k(\vec{X}) + \vec{\nabla} \mathbf{U}_r^k(\vec{X}) \right] \right) \cdot \vec{n}, \tag{29}$$

where the reconstructed gradient, $\vec{\nabla}\mathbf{U}^k$, at location, $\vec{X}$, is given by

$$\vec{\nabla} \mathbf{U}^k(\vec{X}) = \left. \frac{\partial U^k}{\partial x} \right|_{\vec{X}} \hat{\imath} + \left. \frac{\partial U^k}{\partial y} \right|_{\vec{X}} \hat{\jmath}, \tag{30}$$

and where $\hat{\imath}$ and $\hat{\jmath}$ are the Cartesian unit vectors for the two-dimensional space. The derivatives of the reconstructed polynomial in the x- and y-directions, are based on the polynomial coefficients of the unlimited k-exact reconstruction and given by

$$\left. \frac{\partial U^k}{\partial x} \right|_{\vec{X}} = \sum_{\substack{p_1=0 \\ (p_1+p_2 \neq 0)}}^{k} \sum_{p_2=0}^{k} p_1 (x - \bar{x}_{i,j})^{p_1 - 1} (y - \bar{y}_{i,j})^{p_2} D^k_{p_1 p_2}, \tag{31a}$$

$$\left. \frac{\partial U^k}{\partial y} \right|_{\vec{X}} = \sum_{\substack{p_1=0 \\ (p_1+p_2 \neq 0)}}^{k} \sum_{p_2=0}^{k} p_2 (x - \bar{x}_{i,j})^{p_1} (y - \bar{y}_{i,j})^{p_2 - 1} D^k_{p_1 p_2}. \tag{31b}$$

While other methods for the numerical evaluation of the elliptic fluxes would also be possible, such as the method based on the solution of the generalized Riemann problem for diffusion equations as proposed by Gassner et al. [52,53] or the DDG elliptic flux of Liu and Yan [54], the evaluation of the elliptic fluxes in terms of simple arithmetic means of the unlimited cell-centred k-exact reconstructed solutions and their gradients in the cells to the right and left of the interface is certainly straightforward and offers some important computational advantages. In particular, a k-order accurate flux is obtained by using just cell centred reconstructions. However, it is important to assess other properties of the proposed elliptic operator, such as positivity of the operator (related to local satisfaction of a discrete maximum principle) as well as the susceptibility of the operator to odd–even solution decoupling, as similar formulations for standard second-order discretizations can lead

to numerical difficulties related to these issues [44,97]. For these purposes, the proposed discretization scheme for the elliptic flux was assessed by considering its application to the Laplacian operator, $L(u) = \nabla^2 u$, as was done previously by Coirier [44]. In this case, sufficient conditions for the discrete Laplacian based on an N-point stencil given by

$$\tilde{L}(u) \approx \sum_{n=0}^{N} \alpha_n \bar{u}_n, \tag{32}$$

to satisfy a discrete version of the maximum principle are that the weights or influence coefficients, α_n, satisfy $\alpha_n \geqslant 0$ for $n \in [1, N]$ and $\alpha_0 < 0$. As proposed by Coirier [44], the positivity and stability of the scheme can be characterized in terms of α_0 and $\tilde{\alpha}_{\min}$ where

$$\tilde{\alpha}_{\min} = \frac{\min(\alpha_n, 0)}{\sqrt{\frac{1}{N}\sum_n \alpha_n^2}}, \quad n \in [1, N]. \tag{33}$$

Ideally, $\alpha_0 < 0$ for stability and $\tilde{\alpha}_{\min} = 0$ for positivity [45]. Odd–even decoupling properties of the scheme can be assessed by checking for vanishing values for the weights, α_n, for $n \in [1, N]$.

For a given elliptic discretization scheme and supporting stencil, the weights for the approximate Laplacian operator and consequently the positivity of the scheme depend only on local mesh geometry. In the current study, a range of two-dimensional mesh topologies were therefore considered, including Cartesian, stretched, and randomly disturbed quadrilateral grids. For each grid, the associated weights for the proposed schemes were computed numerically. Analysis of the high-order CENO discretization schemes in this manner showed that odd–even solution decoupling does not occur for schemes of all orders $2 \leqslant k \leqslant 4$, even for the second-order scheme with $k = 2$. Furthermore, it was found that $\alpha_0 < 0$ in all cases, implying that all of the schemes are stable. Unfortunately, it was also found that $\tilde{\alpha}_{\min} < 0$ for the proposed elliptic discretization scheme of all orders, implying that, while stable, none of the discretizations satisfy a discrete maximum principle. This result agrees with the general perception that, for finite-volume discretizations, accuracy and positivity are essentially conflicting properties [44]. However, it is important to note that the lack of strict positivity for the elliptic operator does not seem to present an obvious issue in practice, at least for the range of numerical problems that are examined in what follows.

Note that for square Cartesian meshes, values for $\tilde{\alpha}_{\min}$ were found to be -0.823 for $k = 2$, -0.362 for $k = 3$ and -0.854 for $k = 4$ when inverse-distance geometric weighting is used in the k-exact reconstruction. The positivity is improved by using an inverse-distance-squared geometric weighting, for which $\tilde{\alpha}_{\min}$ was found to be -0.051 for $k = 2$, -0.247 for $k = 3$ and -0.324 for $k = 4$. Nevertheless, for non-Cartesian meshes, large variations in the value of $\tilde{\alpha}_{\min}$ were observed ($-5 < \tilde{\alpha}_{\min} < 0$), depending on the regularity and topology of the mesh.

Aside from positivity and decoupling properties, the stability of the proposed high-order discretization scheme for the elliptic operator is also of concern. While either Fourier or eigensystem analysis can be applied to the discrete spatial operator of the CENO scheme and assessed for various time-marching schemes [98], such a systematic analysis has been reserved for follow-on studies in which optimal high-order time-marching schemes are considered for use in conjunction with the CENO finite-volume method. However, numerical experiments with the proposed elliptic flux function and CENO scheme, carried out using the selected explicit time-marching schemes described in the next section, suggest that the high-order CENO schemes are conditionally stable for these rather standard explicit schemes and that the stability limits are defined by Courant–Friedrichs–Lewy or CFL-like and Neumann-like conditions for the hyperbolic and elliptic operators, respectively, that appear to be only very slightly more restrictive than that those of standard second-order discretizations. Penalty terms for the elliptic fluxes do not appear to be required for stability but may enhance the stability and/or accuracy of the scheme and will be assessed more carefully in future follow-on studies, along with a formal stability analysis of the high-order scheme.

3.6. Explicit time-marching schemes

Efficiency and accuracy of the temporal discretization were not the focus here. While possibly more efficient explicit and implicit high-order time-marching schemes have been examined in several recent studies (see, for example, the studies by Balsara et al. [99] and Bijl et al. [100]), integration of Eq. (8) was carried out using standard explicit time-marching schemes for both steady and unsteady problems. For steady, time-invariant problems, for which $\mathbf{R}_{i,j}(\bar{\mathbf{U}}) = 0$, the explicit optimally-smoothing multi-stage schemes developed by van Leer et al. [101,102] were used. For time-accurate calculations, standard two- and four-stage Runge–Kutta schemes [98,103,104] were used, depending on the accuracy of the spatial reconstruction. When applied in conjunction with the AMR scheme to be described in the section to follow, the explicit time integration procedure for time accurate computations is applied to all cells on all grid levels with a single global time step and no effort is made to reduce the computational effort by advancing the solutions on each level of mesh resolution with its own optimal time step. The formulation and use of more efficient, high-order, time-marching schemes for use in conjunction with the proposed high-order finite-volume method is beyond the present scope but will certainly be considered in future follow-on research.

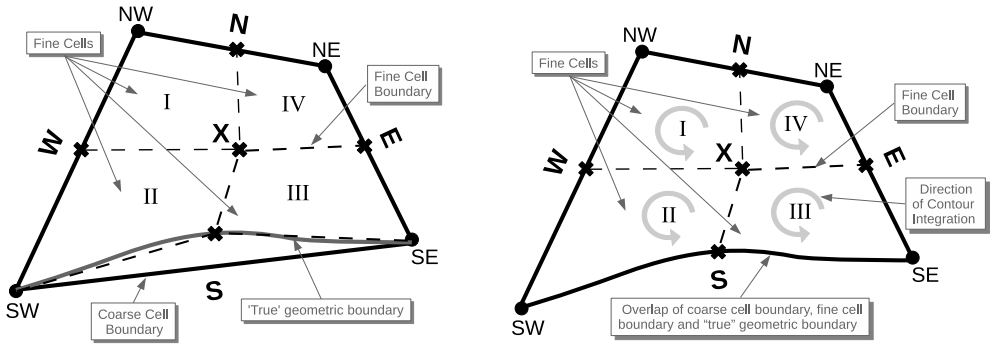

(a) Low-order boundary representation (b) High-order boundary representation

Fig. 3. Depiction of a coarse cell division into four fine cells for low- and high-order block boundary elements.

4. Adaptive mesh refinement

The development of the proposed high-order CENO finite-volume algorithm is considered here in conjunction with the block-based AMR scheme originally formulated by Sachdev et al. [67], Gao and Groth [68,105], and Gao et al. [69] for second-order, finite-volume schemes on body-fitted mesh. As in the lower-order approach, a flexible hierarchical data structure is used in combination with the CENO scheme to facilitate automatic solution-directed mesh adaptation on two-dimensional, body-fitted, multi-block, quadrilateral mesh. The AMR technique allows for the use of anisotropic mesh and is well suited to parallel implementation via domain decomposition. The high-order variant requires inter-block communication of high-order solution content, as well as high-order treatments of physical boundaries, boundary conditions, and solution transfer between AMR grids. While the combination of high-order CENO finite-volume and AMR strategies has been considered by Ivan and Groth for inviscid flows [65,66,122], the application to the advection–diffusion equation and fully-viscous flows is the current focus.

In the proposed high-order, block-based, AMR approach, mesh adaptation is accomplished by the dividing and coarsening of appropriate solution blocks. In regions requiring increased cell resolution, a "parent" block is refined by dividing itself into four "children" or "offspring". Each of the four quadrants or sectors of a parent block becomes a new block having the same number of cells as the parent and thereby doubling the cell resolution in the region of interest. This process can be reversed in regions that are deemed over-resolved and four children are coarsened into a single parent block. The mesh refinement is constrained such that the grid resolution changes by only a factor of two between adjacent blocks. A hierarchical tree-like data structure with multiple "roots", multiple "trees", and additional interconnects between the "leaves" of the trees is used to keep track of mesh refinement and the connectivity between solution blocks.

In order to apply the CENO finite-volume scheme to all grid blocks in an independent manner, solution information is shared between adjacent blocks having common interfaces. This information is contained in and exchanged using additional layers of overlapping "ghost" cells associated with each block. Additional inter-block communication is also required at interfaces with resolution changes to enforce strictly the flux conservation properties of the scheme [106,107]. In particular, the interface fluxes computed on more refined blocks are used to correct the interface fluxes computed on coarser neighbouring blocks and ensure the solution fluxes are conserved across block interfaces.

During refinement of grid blocks, the geometry of newly created cells is obtained by dividing each coarse interior cell into four fine cells denoted as *I*, *II*, *III* and *IV*, and illustrated in Fig. 3. The newly introduced nodes are merely the midpoints of each coarse cell face and weighted averages of the four cell nodes as defined in the curvilinear or computational coordinate frame of the grid. Second-order estimates of the grid metrics are used to define the transformation between the curvilinear coordinate frame and physical space so as to preserve the smoothness and stretching of the grid lines in the body-fitted mesh as described by Gao and Groth [105]. Nodes of the refined mesh inserted at curved physical boundaries are made to lie on the high-order splines defining each boundary. The utilization of high-order boundary elements (i.e., cells with non-straight, curved edges), as depicted in Fig. 3(b), ensures naturally the equality of the coarse element area, A_Ω, with the summation of the areas of the offspring, $A_I + A_{II} + A_{III} + A_{IV}$. It then follows that the conservation property for the transfer of solution content between coarse and fine mesh can be expressed simply as

$$\bar{u}_\Omega A_\Omega = \bar{u}_I A_I + \bar{u}_{II} A_{II} + \bar{u}_{III} A_{III} + \bar{u}_{IV} A_{IV}. \tag{34}$$

The hybrid CENO solution reconstruction procedure is used in conjunction with standard multi-grid-type restriction and prolongation operators to evaluate the solution on all blocks created by coarsening and division processes, respectively. High-order operators for restriction and prolongation of the solution between coarse and fine meshes are used here based on the coarse-to-fine cell conservation property for conserved variables as expressed by Eq. (34). High-order restriction from the fine to coarse cells is straightforward and can be achieved by simple application of the conservation property. The high-order prolongation from coarse to fine cells can be accomplished by direct integration of the appropriate coarse-cell

polynomial reconstructions for the conserved variables over the domain of each newly created fine cell. While satisfying the conservation property above, this prolongation operator does not explicitly enforce solution positivity and, in some instances, negative pressures and densities may result when solving the Navier–Stokes equations. Although positivity-preserving high-order reconstructions of conserved variables are certainly possible as discussed by Zhang and Shu [84,85] and Balsara [86], a more pragmatic approach is currently adopted here. In situations where solution positivity is violated by the high-order prolongation operator, simple direct injection of the solution is used with $\bar{u}_I = \bar{u}_{II} = \bar{u}_{III} = \bar{u}_{IV} = \bar{u}_\Omega$.

The preceding approach for the prolongation of solution content is rather straightforward when the coarse-cell reconstruction is carried out in terms of the conserved variables. When other non-conservative variables are used in the reconstruction as also proposed here, conservative prolongation is achieved in terms of the derived and/or primitive solution quantities, $\bar{W}$, by first determining the average integrated value of the primitive variable, $\bar{W}'_f$, in each newly created fine cell, f, then by evaluating an estimate for the corresponding cell average of the conservative variables using the expression $\bar{U}'_f = f(\bar{W}'_f)$, where $f(W)$ represents the conversion from primitive to conserved solution quantities, and finally by determining the desired values of conserved solution in each fine cell, $\bar{U}_f$, by correcting the preceding estimate with a term related to the total integrated error in the conserved quantity. The corrected conserved solution for fine cell, f, is then given by

$$\bar{U}_f = \bar{U}'_f + \Delta\bar{U}_f = f\left(\bar{W}'_f\right) + \frac{1}{A_\Omega}\left[\bar{U}_\Omega A_\Omega - \sum_{i=I}^{IV} f\left(\bar{W}'_i\right) A_i\right], \tag{35}$$

where f corresponds to the fine cells associated with sectors I, II, III and IV of the coarse cell Ω. As discussed above, direct injection is again used in cases for which non-physical prolonged solutions are encountered.

While heuristic, physics-based, refinement criteria have been used in previous work to direct the coarsening and division of grid blocks [67–69,105], a novel h-refinement criterion based on the CENO solution smoothness indicator is instead used here to control refinement of the body-fitted, multi-block, AMR mesh. The proposed h-refinement criterion is based solely on values of a refinement parameter, $\mathcal{R}_s$, for solution variable, s, given by

$$\mathcal{R}_s = \exp\left(-\frac{\max(0, \mathcal{S}_s)}{\mathcal{U}_s \mathcal{S}_c}\right), \tag{36}$$

where $\mathcal{S}_s$ is the smoothness indicator for variable s, $\mathcal{S}_c$ is the corresponding cut-off value, and $\mathcal{U}_s$ is an appropriate scaling coefficient. Grid blocks having values of $\mathcal{R}_s$ greater than a specified maximum are flagged for division and those having values less than a specified minimum are tagged for coarsening. By definition, the proposed refinement parameters are non-dimensional and have values in the range $[0, 1]$. Consequently, the importance of different solution quantities in directing the AMR is based only on their relative smoothness.

Although the block-based AMR approach described above is somewhat less flexible and incurs some inefficiencies in solution resolution as compared to cell-based approaches [108–111] (i.e., for the same solution accuracy, generally more computational cells are introduced in the adapted grid), the method offers many advantages over cell-based techniques when parallel implementation and performance of the solution algorithm are taken into account. In particular, the multi-block quadrilateral mesh and quadtree data structure lend themselves naturally to geometric domain decomposition based on the grid block. The solution blocks can be easily evenly distributed to the available processor cores, with more than one block permitted on each core. In this way, very efficient and scalable implementations of the CENO finite-volume scheme on distributed-memory parallel architectures can be achieved. The parallel efficiency of the proposed scheme is examined in Section 5.7.2.

5. Numerical results

To demonstrate the capabilities of the proposed high-order CENO scheme, numerical results are now presented for solution reconstruction in two-space dimensions, as well as for various two-dimensional flow problems governed by the advection–diffusion and Navier–Stokes equations. Moreover, the influence of the choice of reconstruction variables on the accuracy of the CENO scheme is assessed for Ringleb's flow, an inviscid flow having an analytical solution. Numerical results are considered for both fixed and AMR meshes.

5.1. Reconstruction of two-dimensional smooth trigonometric function

The properties of the CENO reconstruction for two-dimensional smooth solutions are illustrated by considering reconstruction of the trigonometric function, $u(x, y)$, given by

$$u(x, y) = 1.1 + \cos\left(\pi x^2 + 4\pi y\right), \tag{37}$$

on body-fitted computational domains defined between streamlines corresponding to $k = 0.75$ and $k = 1.5$ and the iso-velocity contour $q = 0.5$ for Ringleb's flow of an inviscid compressible gas [112]. Refer to Section 5.6 for further details of the Ringleb's flow. The exact distribution of the function, u, and the corresponding cell-centred, quartic ($k = 4$), CENO, reconstructed solution for a regular body-fitted mesh with 40×40 cells are compared in Fig. 4. From the figure, it is evident

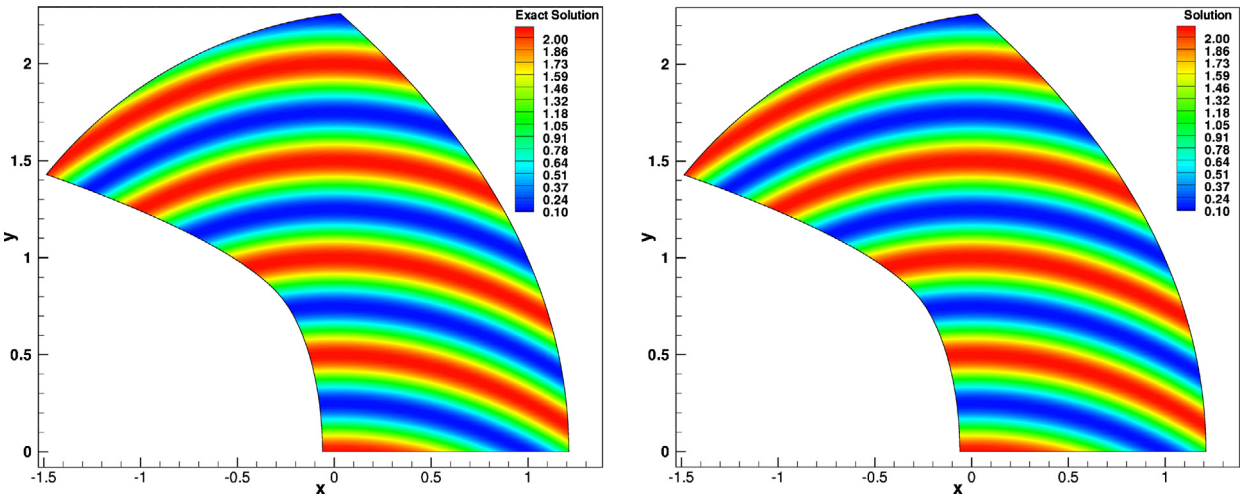

(a) Exact function (b) Quartic ($k=4$) CENO reconstruction

Fig. 4. Exact distribution of smooth function, $u(x, y) = 1.1 + \cos(\pi x^2 + 4\pi y)$, and corresponding quartic CENO reconstruction on a regular body-fitted mesh having 40×40 computational cells.

the quartic CENO reconstruction is able to capture all of the smooth extrema associated with the function, even on this relatively coarse mesh, providing a solution profile that is visually indistinguishable from the original exact function. Note that the gradients of the function are not at all aligned with the computational mesh.

To quantify the CENO performance and reconstruction errors for this smooth function, convergence studies were carried out using two mesh types with mesh densities ranging from 10×10 to 160×160 computational cells. A sequence of regular body-fitted grids with smooth grid lines was considered along with a second sequence of irregular, non-smooth, grids produced by perturbing randomly the nodes of the regular meshes. An irregular perturbed mesh from the second sequence with 40×40 cells is depicted in Fig. 5(a). The accuracy of the CENO reconstruction procedure for cubic ($k = 3$) and quartic ($k = 4$) representations was evaluated by considering reconstructed solution accuracies for the following cases: (i) k-exact reconstruction; (ii) CENO reconstruction with a value for the cutoff parameter of $\mathcal{S}_c = 1000$; (iii) CENO reconstruction with $\mathcal{S}_c = 5000$. The L_1-, L_2-, and L_∞-norms of the reconstruction error defined in terms of the integrated difference between the smooth function and reconstructed solutions in each computation cell were computed in each case. For each cell, adaptive quadrature was used in integrating the difference between the functional and reconstructed solution over the cell area.

Fig. 5 shows the convergence of the L_1, L_2, and L_∞-norms of the reconstruction error with increasing mesh resolution for the k-exact method and CENO schemes with $\mathcal{S}_c = 1000$ and $\mathcal{S}_c = 5000$ on the regular and perturbed meshes. Firstly, the results of the figure verify that the expected theoretical order of accuracy or convergence rate is achieved in each case (four for the cubic and five for the quartic reconstruction schemes), at least in the asymptotic limit. Additionally, the reconstruction error norms obtained on the perturbed meshes are very similar to those corresponding to the regular grids, although the absolute error is slightly larger. The influence of the cutoff value, $\mathcal{S}_c$, can also be inferred from the figure. On coarse meshes, the smooth function is deemed to be under-resolved and the CENO reconstruction procedure reverts to piecewise limited linear reconstruction with similar accuracies in each case. As the mesh density increases, the function is eventually resolved, the CENO reconstructions deemed smooth, and k-exact reconstruction is recovered. While the extent of the transition from limited piecewise linear to k-exact reconstruction is somewhat sensitive to the cutoff value, $\mathcal{S}_c$, in the asymptotic limit the results are insensitive to values of $\mathcal{S}_c$ in the recommended range (i.e., 1000–5000). It seems that for complex but smooth functional variations of the type considered here, a minimum of about 60 to 100 cells per space dimension is needed before the cubic CENO reconstruction scheme identifies the solution as being smooth and the unlimited k-exact reconstruction based on the central stencil is used. This minimum cell resolution is reduced to about 40 to 60 cells per direction when quartic CENO reconstruction is used. This minimum mesh resolution requirement for feature detection and recovery of the asymptotic solution may be slightly larger than that of other ENO and WENO schemes (see, for example, Jiang and Shu [6] and Balsara and Shu [113]), but not by a significant margin.

5.2. Residual accuracy assessment for Poisson's equation

The accuracy of the proposed, high-order, finite-volume, discretization scheme for elliptic operators is confirmed by considering the convergence of the discretized residual error for a Poisson equation of the form

$$\vec{\nabla} \cdot (\vec{\nabla} u) = a e^{\beta u}, \tag{38}$$

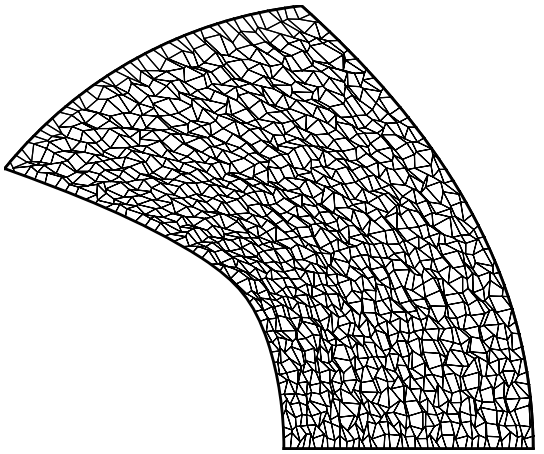

(a) Irregular perturbed grid with 40×40 cells

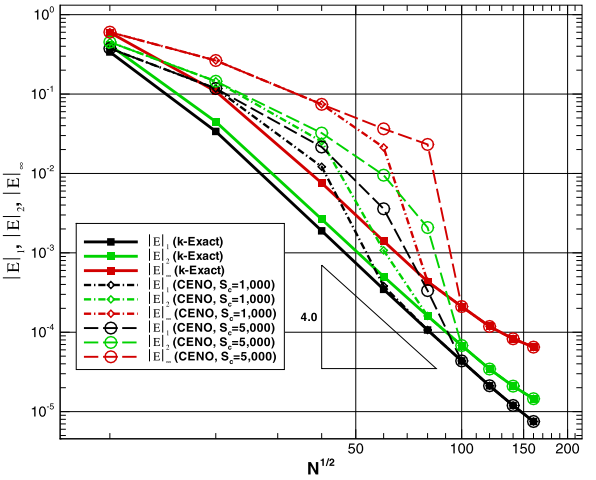

(b) Cubic ($k = 3$) reconstruction error on regular mesh

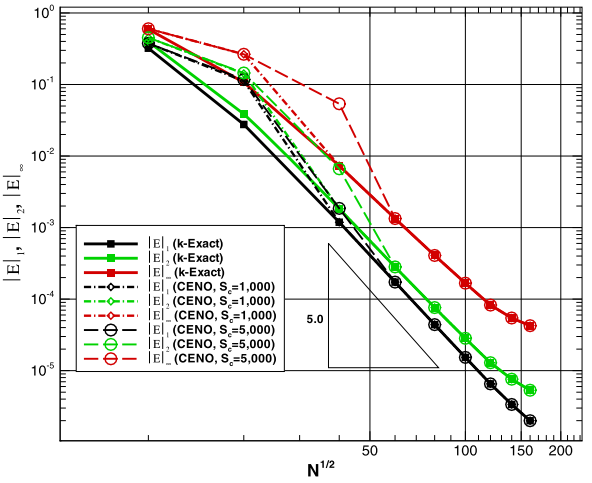

(c) Quartic ($k = 4$) reconstruction error on regular mesh

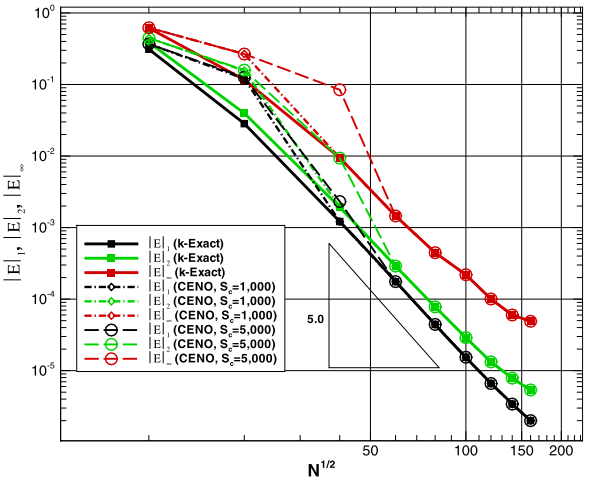

(d) Quartic ($k = 4$) reconstruction error on perturbed mesh

Fig. 5. Example of irregular perturbed grid with 40×40 cells and L_1, L_2, and L_∞-norms of the cubic and quartic reconstruction error for smooth function, $u(x, y) = 1.1 + \cos(\pi x^2 + 4\pi y)$, obtained using k-exact and CENO reconstruction schemes, the latter with $\mathcal{S}_c = 1000$ and $\mathcal{S}_c = 5000$, as a function of mesh resolution for regular and perturbed meshes.

on a rectangular domain defined by $0.5 \leqslant x, y \leqslant 4.5$ and having the exact solution

$$u(x, y) = \frac{1}{\beta}\left[\ln\left(\frac{8C}{a\beta}\right) - 2\ln\left|(x + A)^2 + (y + B)^2 - C\right|\right], \tag{39}$$

where $A = 2.0$, $B = 1.0$, $C = 2.0$, $a = 2.5$ and $\beta = 0.001$. The L_1-, L_2-, and L_∞-norms of the errors in the discrete residual (i.e., the flux integral for the elliptic operator minus the area integral of the source term) are shown in Figs. 6(a) and 6(b), as a function of the mesh density for sequences of regular Cartesian grids with uniform mesh spacing as well as randomly disturbed variants of these meshes. The numerical results for quadratic ($k = 2$), cubic ($k = 3$), and quartic ($k = 4$) reconstruction were obtained by applying the unlimited k-exact reconstruction to the exact solution and then evaluating the flux integral and the area integral of the source term for each computational cell (the difference between the values of the two latter integrals is expected to be identically zero on all meshes). The corresponding slopes of the L_1-, L_2- and L_∞-norm error curves as determined from the computed errors are also shown in Table 2 for both Cartesian and disturbed meshes and compared to the expected theoretical residual accuracies (convergence rates) of the schemes.

The results of Fig. 6 and Table 2 show that, in agreement with theoretical expectations, both cubic and quartic interpolants produce 4th-order schemes on uniform Cartesian meshes in all error norms, whereas the quadratic scheme yields a second-order accurate discrete elliptic operator. For uniform Cartesian meshes, the absolute error of the quartic reconstruction scheme is only marginally lower than those of the cubic reconstruction scheme. This is because error cancellation

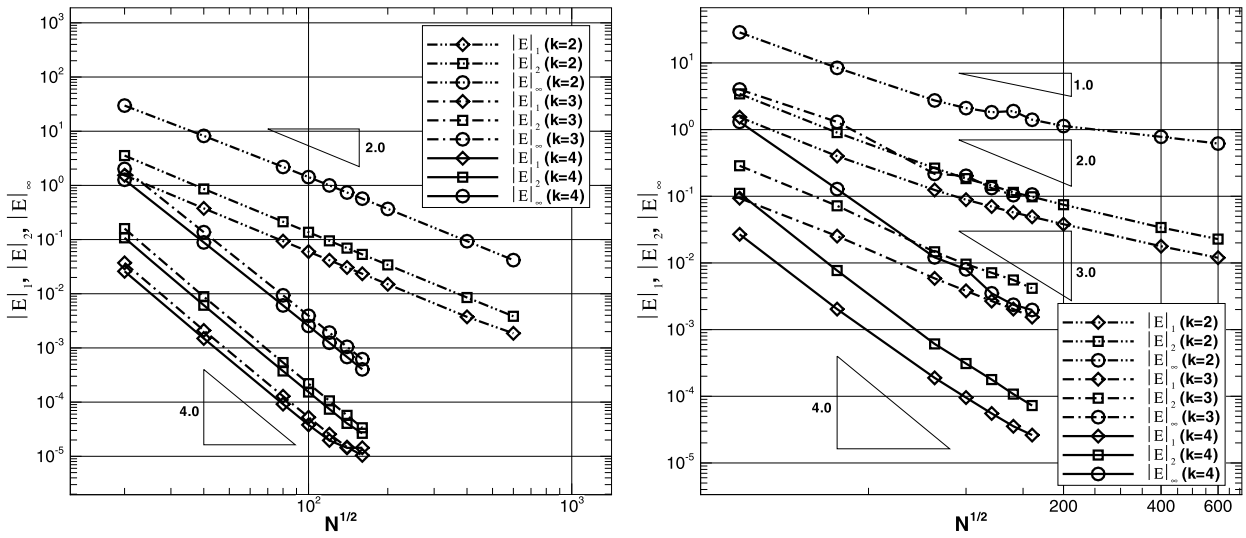

(a) Error norms for regular Cartesian meshes (b) Error norms for distorted meshes

Fig. 6. L_1-, L_2-, and L_∞-norms of the discretized residual error for Poisson's equation, $\vec{\nabla}\cdot(\vec{\nabla}u) = ae^{\beta u}$, obtained using CENO, finite-volume discretization of the elliptic operator based on unlimited, piecewise, quadratic ($k=2$), cubic ($k=3$), and quartic ($k=4$) k-exact reconstructions as a function of mesh density cells for regular Cartesian meshes and randomly distorted meshes.

Table 2
Estimated slopes of the curves of Fig. 6 for the L_1-, L_2- and L_∞-norm error in the solution residual compared to expected theoretical residual accuracies (convergence rates).

Mesh	Reconstruction order	L_1	L_2	L_∞
Cartesian grid	$k=2$	1.99	2.01	1.93
	$k=3$	3.98	4.07	3.89
	$k=4$	3.77	4.02	3.87
Distorted grid	$k=2$	1.41	1.47	1.10
	$k=3$	1.99	2.05	1.86
	$k=4$	3.35	3.53	3.18

on the Cartesian mesh yields a 4th-order discretization scheme for the Poisson operator, in the case of $k=3$ (cubic) reconstruction. However, error cancellation does not occur on the irregular randomly-disturbed meshes and the convergence rates (slopes) of the error norms for cubic and quartic reconstructions differ by more than one. The quartic ($k=4$) scheme provides a significantly more accurate residual estimation that approaches theoretical expectations, even for the sequence of irregular randomly-disturbed meshes which lack the usual nesting property of mesh sequences used in rigorous convergence studies. Due to the lack of error cancellation and irregular nature of the meshes, cubic reconstruction yields a scheme that is only about second-order accurate on the perturbed grids.

5.3. Solution of Laplace's equation on domains with curved boundaries

Numerical solution of Laplace's equation is explored next for Dirichlet boundary-value problems on domains having curved boundaries. The circular-shaped annular domain shown in Fig. 7(a) was considered and Dirichlet boundary conditions were implemented via constrained reconstruction based on the exact solution for this problem given by

$$u(x,y) = e^{\mu x}\big(A\cos(\mu y) + B\sin(\mu y)\big), \tag{40}$$

where $A=1$, $B=2$ and $\mu=1.5$. A predicted solution obtained using the proposed 4th-order ($k=4$), CENO, finite-volume scheme on a body-fitted curvilinear mesh with 40×40 cells is depicted in Fig. 7(a). The L_1-, L_2-, and L_∞-norms of the error in the predicted solutions compared to the exact result for both cubic ($k=3$) and quartic ($k=4$) interpolants are given in Fig. 7(b) for this problem. The slopes of the L_1- and L_2-norms achieve asymptotic values of -3.86 and -3.85 for $k=3$ and -3.86 and -3.81 for $k=4$, respectively, confirming the expected 4th-order accuracy on curvilinear mesh. While the cubic and quartic reconstruction schemes have virtually identical accuracy, there is about one order difference between the absolute magnitudes of the solution errors, demonstrating the benefits of using a quartic interpolant.

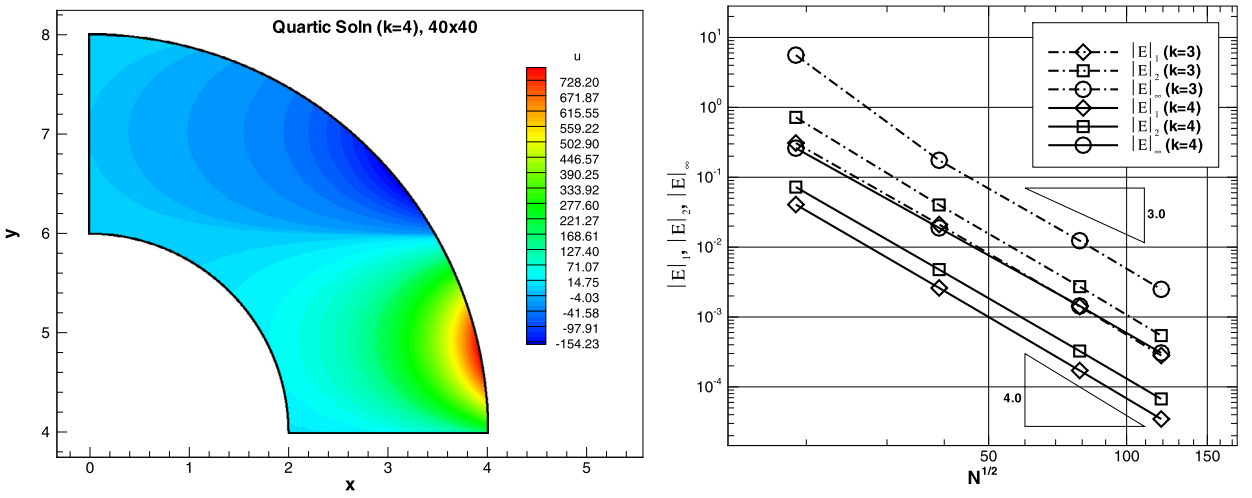

Fig. 7. (a) Fourth-order solution to the Laplace equation, u, obtained using the 4th-order, CENO, finite-volume discretization of the elliptic operator based on unlimited quartic ($k=4$) reconstruction on a mesh with 40×40 computational cells; and (b) computed L_1-, L_2-, and L_∞-norms of the solution error as a function of mesh density obtained using CENO, finite-volume discretization of the elliptic operator based on unlimited cubic ($k=3$) and quartic ($k=4$) k-exact interpolants.

5.4. Solution of advection–diffusion equation for a range of Péclet numbers

The application of the proposed CENO scheme is now considered for problems involving the advection–diffusion equation for a range of Péclet numbers. In particular, solution of Eq. (1) with constant advection velocity, $\vec{V}=[v_0,0]$, and constant diffusion coefficient, $\kappa(x,y)=\kappa_0=0.01$, on a rectangular domain of length $L=3$ and unit width was considered for three different Péclet numbers as dictated by the value of v_0. Péclet numbers of $\mathrm{Pe}=0.1$, 1, and 10 were examined corresponding to diffusion-dominated, advection–diffusion balanced, and advection-dominated regimes. The boundary data for this channel-flow problem, as illustrated in Fig. 8, were

$$u(x,0)=u(x,1)=0,\qquad u(0,y)=\sin(\pi y),\qquad \left.\frac{\partial u}{\partial x}\right|_{x=L}=0. \tag{41}$$

An exact analytical solution to this boundary-value problem can be determined by the method of separation of variables and has the form

$$u(x,y)=\frac{\sin(\pi y)}{(R[e^{L(R-1)}]^{r_1}-1)}\left(R\left[e^{(RL+x-L)}\right]^{r_1}-e^{r_2 x}\right), \tag{42}$$

where $r_{1,2}=(v_0/2\kappa_0)\pm\sqrt{(v_0^2/4\kappa_0^2)+\pi^2}$ and $R=r2/r1$. An example numerical solution obtained for $\mathrm{Pe}=10$ on an 80×40 Cartesian mesh is shown in Fig. 8(a). Similar problems were considered previously by Ollivier-Gooch and Van Altena [47] for evaluating high-order schemes.

The various norms of the solution error for the advection–diffusion problem with $\mathrm{Pe}=0.1$, 1, and 10 are shown in Figs. 8(b)–8(d). It is evident that the errors for the high-order, CENO, finite-volume scheme with quartic reconstruction are consistently lower than those of the scheme employing a cubic interpolant by at least one order of magnitude for all Péclet numbers. Furthermore, the computed convergence rates for L_1 and L_2 error norms with $k=4$ are -4.02 and -4.08 for $\mathrm{Pe}=0.1$, -4.30 and -4.46 for $\mathrm{Pe}=1.0$, and -3.92 and -3.95 for $\mathrm{Pe}=10.0$, respectively, indicating that a uniformly accurate fourth-order scheme is recovered in all cases, in agreement with theoretical expectations. For $k=3$, the convergence rates of the error norms are -3.92 and -3.85 for $\mathrm{Pe}=0.1$, -3.88 and -3.81 for $\mathrm{Pe}=1.0$ and -3.53 and -3.62 for $\mathrm{Pe}=10.0$, respectively, indicating a slight loss in accuracy for high Péclet numbers. For the same accuracy level, the CENO scheme with cubic reconstruction requires almost twice as many computational cells as the scheme based on quartic reconstruction. The increased accuracy of the quartic-based scheme comes at a relatively small additional computational cost: the quartic and cubic use the same reconstruction stencil, the only additional costs of the quartic reconstruction scheme are associated with storage, computation, and use of five additional coefficients in the least-squares reconstruction procedure.

5.5. Solution of advection equation with AMR

The application of the proposed high-order, CENO, finite-volume, scheme to the solution of the pure advection equation ($\mathrm{Pe}\to\infty$) is now examined. The particular case of circular advection with a constant angular velocity was considered on

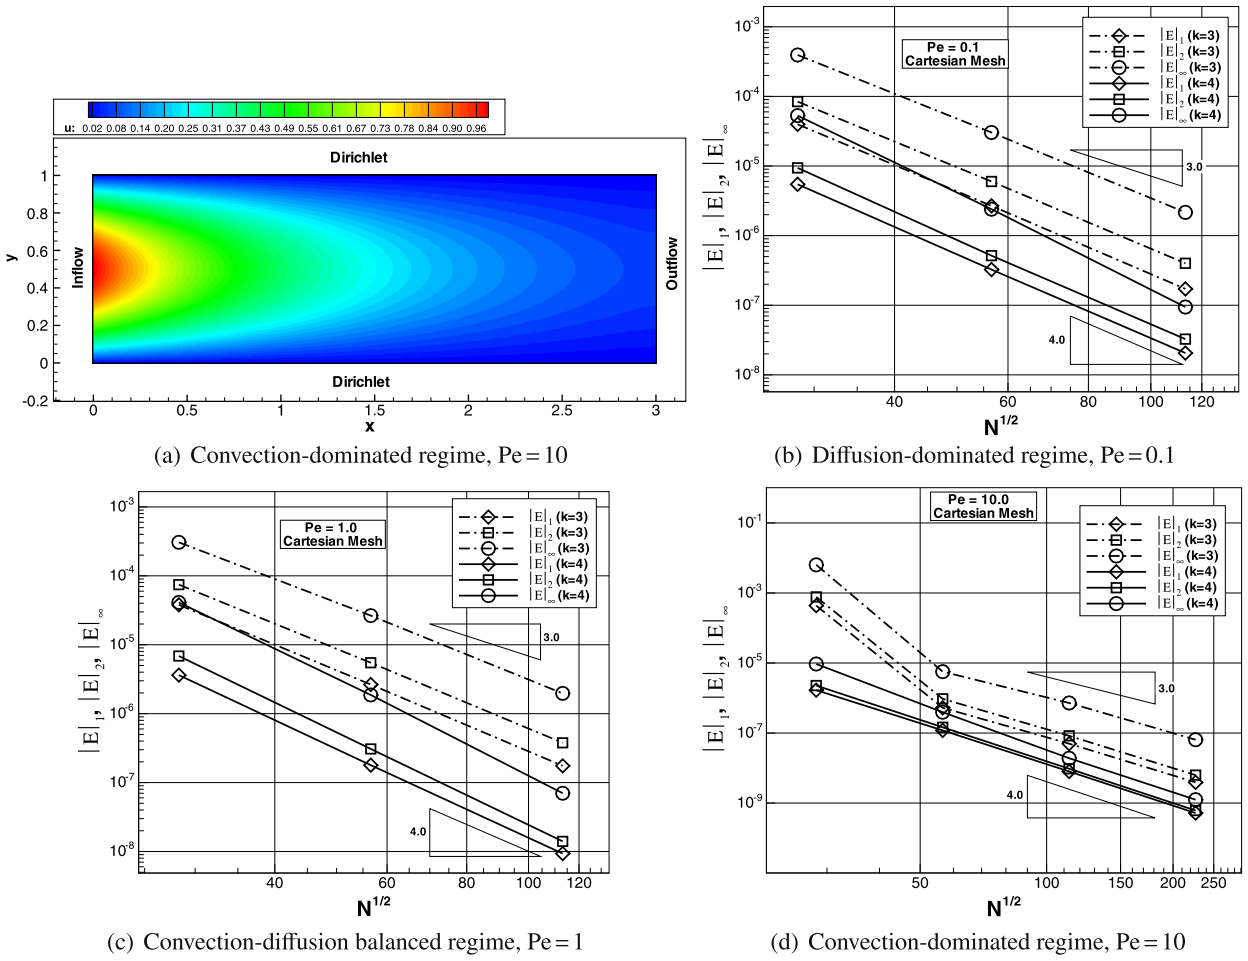

(a) Convection-dominated regime, Pe = 10

(b) Diffusion-dominated regime, Pe = 0.1

(c) Convection-diffusion balanced regime, Pe = 1

(d) Convection-dominated regime, Pe = 10

Fig. 8. (a) Numerical prediction of the solution, u, for the advection–diffusion channel flow problem with Pe = 10 obtained using 80×40 Cartesian mesh; and (b)–(d) L_1-, L_2-, and L_∞-norms of the solution error for the advection–diffusion channel flow problem as a function of mesh density obtained with cubic ($k = 3$) and quartic ($k = 4$) CENO reconstruction for Pe = 0.1, 1, and 10.

a square solution domain with a dimension of 1.5. A prescribed inflow solution was specified along the bottom of the domain ($y = 0$) having the form

$$u(x,0) = \begin{cases} e^{2(x-0.4)} \sin^6[2\pi(x-0.4)], & \text{for } 0.4 \leqslant x \leqslant 1.2, \\ 0, & \text{for } x < 0.4 \text{ or } x > 1.2, \end{cases} \tag{43}$$

and an outflow boundary condition was applied along the left boundary ($x = 0$). The discontinuous inflow function was chosen to test both the accuracy and robustness of the scheme as well as the capabilities of the proposed AMR algorithm.

Numerical results for circular advection problem with AMR are given in Fig. 9. Predicted steady-state solutions were obtained on a sequence of adaptively refined meshes starting with an initial mesh consisting of 16 10×10 solution blocks as shown in Fig. 9(a). An intermediate mesh resulting from three levels of adaptive refinement is depicted in Fig. 9(b) and the final mesh consisting of 2911 blocks and 291,100 computational cells is shown in Fig. 9(c). The predicted solution obtained using the 4th-order CENO scheme ($k = 3$) on the final mesh is given in Fig. 9(d) and a comparison between the predicted profile along the cross-section A–A to the exact solution is provided in Fig. 9(e). It should be evident from the quality of the predicted solution that the proposed AMR scheme in combination with the h-refinement criteria based on the smoothness indicator is capable of refining both under-resolved and non-smooth regions of the solution, yet will not unnecessarily refine resolved solution content. Smooth extrema are well captured by the high-order scheme whereas discontinuities are readily identified by the smoothness indicator and subsequently resolved by the AMR procedure.

5.6. Solution of Euler equations for Ringleb's flow

The accuracy of the CENO approach is now demonstrated for the discretization of the hyperbolic operator representing the Navier–Stokes equations. The hyperbolic operator in this case corresponds to the well-known Euler equations of

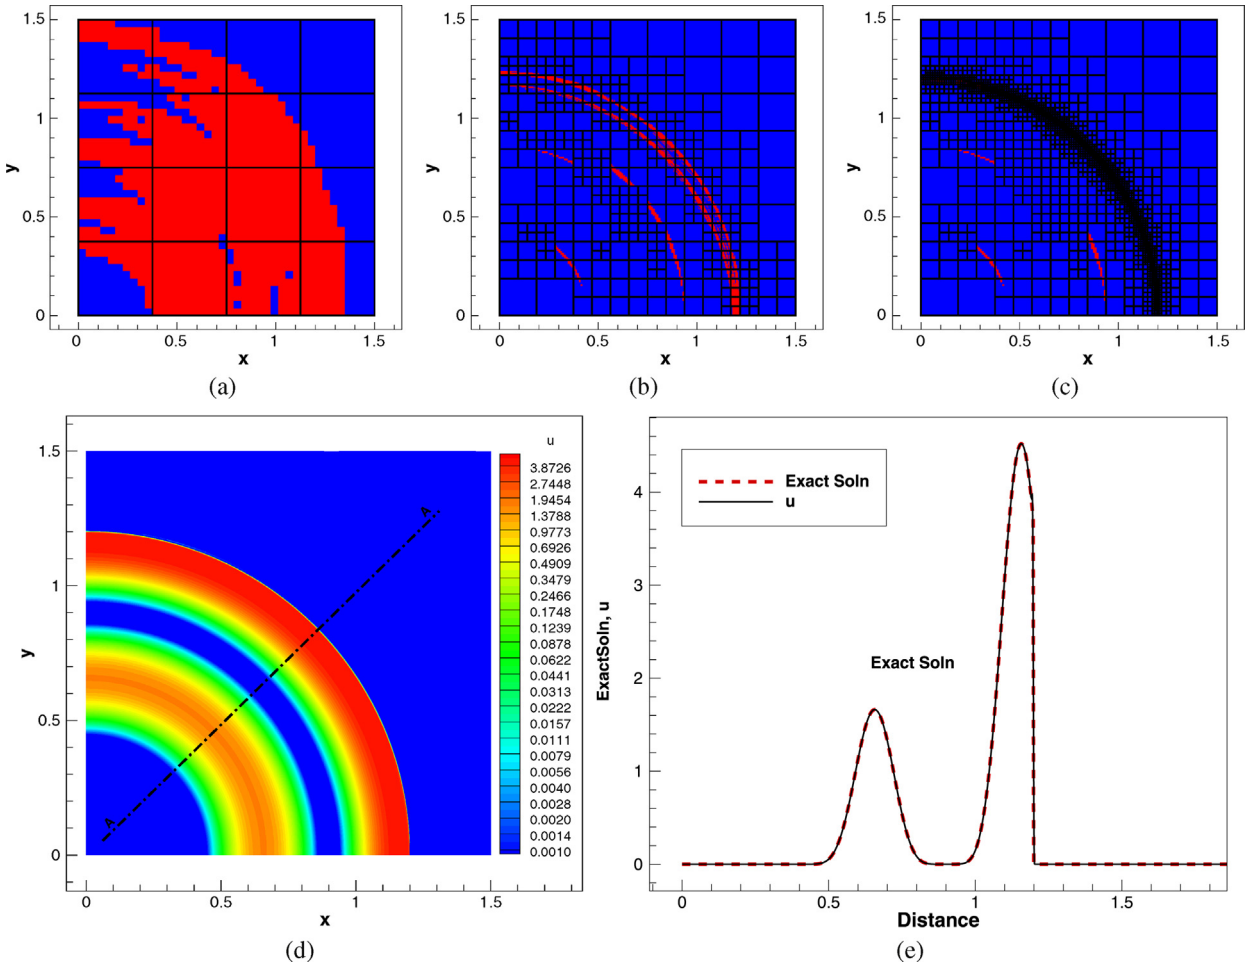

Fig. 9. Numerical solution of circular advection problem with AMR showing: (a) initial mesh with 16 10×10 blocks with regions in which limited linear reconstruction is used shown in red; (b) refined mesh after 3 levels of refinement; (c) final refined mesh after 6 refinement levels with 2911 10×10 blocks, 291,100 computational cells, and a refinement efficiency of $\eta = 0.955$; (d) prediction solution, u, obtained using the fourth-order ($k = 3$), CENO, finite-volume scheme on finest mesh; and (e) comparison of exact analytical and predicted solutions for u along cross section A–A. (For interpretation of the references to colour in this figure, the reader is referred to the web version of this article.)

compressible gas-dynamics which follow from the Navier–Stokes equations under the inviscid assumption (i.e., $\mu = \nu = 0$). Ringleb's flow involves compressible, inviscid, isentropic, and irrotational flow in a stream tube as described by the Euler equations and exact solutions for this smooth continuous flow field can be determined by analytical means [109,112]. A transonic variant of Ringleb's flow defined by $k_{\min} = 0.5$, $k_{\max} = 1.2$, and $q = 0.3$ was considered here. In this case, the inflow to the converging stream tube is subsonic with a Mach number of about $M \approx 0.34$ and, following convergence and turning of the flow, a portion of the outflow is supersonic with $M \approx 1.30$.

Numerical solutions of the Ringleb flow problem on body-fitted multi-block quadrilateral mesh were determined and compared to the exact analytical result as part of a systematic grid convergence study. In addition, a head-to-head comparison of the CENO scheme predictions for cell-centred reconstruction based on both primitive and conserved dependent variable sets was also carried out to assess the effects of this choice on solution accuracy and errors that may be introduced by working with a primitive set. Note that it is important to assess any potential sources of errors that may be introduced by the conversion from conserved to primitive variables as discussed previously in Section 3.3. Reflection boundary conditions were applied along the streamline boundaries by enforcing the inviscid (slip) condition, $\vec{V} \cdot \vec{n} = 0$, at all Gauss integration points via the constrained least-squares reconstruction procedure outlined in Section 3.3.3 above.

The predicted Mach number distribution for the Ringleb flow problem obtained using the 4th-order, CENO, finite-volume scheme on an 80×80 body-fitted mesh is given in Fig. 10(a). The computed L_1-, L_2-, and L_∞-norms of the solution error in both predicted density and pressure obtained using the 3rd- and 4th-order versions of the CENO scheme with $k = 2$ and $k = 3$, respectively, are shown in Figs. 10(b) and 10(c). The slopes of the L_1- and L_2-norm curves of Fig. 10(b) for the 3rd- and 4th-order CENO schemes based on reconstruction of the primitive variables are -3.06, -3.00, -3.93, and -4.02,

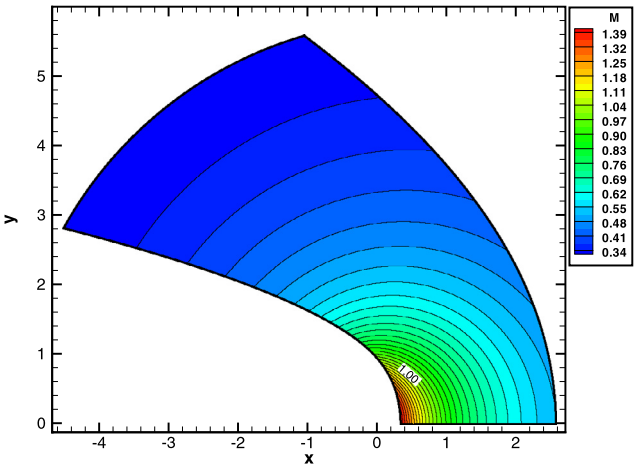

(a) Solution of 4th-order ($k=3$) CENO scheme

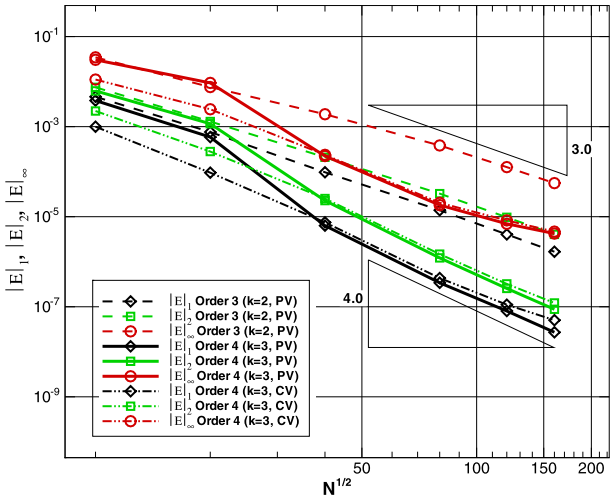

(b) Density error norms

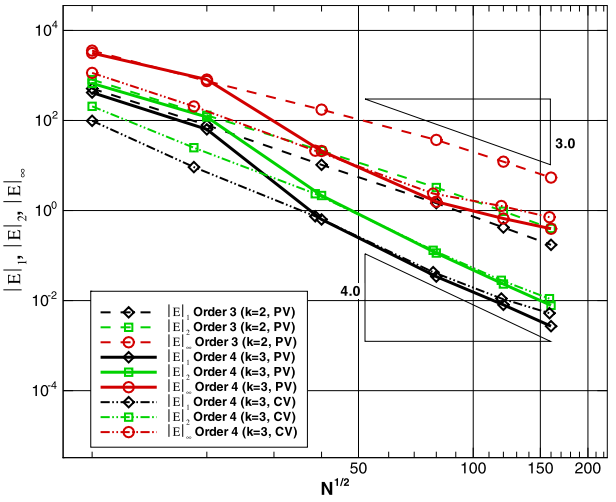

(c) Pressure error norms

Fig. 10. Numerical solution of transonic Ringleb's flow showing: (a) predicted distribution for Mach number, M, obtained using 4th-order, CENO, finite-volume scheme; (b) L_1-, L_2-, and L_∞-norms of the solution error in density, ρ, for 3rd- and 4th-order CENO schemes as a function of mesh resolution; and (c) L_1-, L_2-, and L_∞-norms of the solution error in pressure, p, for 3rd- and 4th-order CENO schemes as a function of mesh resolution. The predicted error norms for the 4th-order CENO scheme are given for reconstruction in terms of both primitive (PV) and conservative (CV) solution variables.

respectively, illustrating that the expected formal accuracy of the schemes is maintained on body-fitted mesh with strongly curved boundaries by using the constrained least-squares reconstruction procedure.

The error norms of the 4th-order CENO schemes based on reconstruction of primitive (PV) and conserved (CV) dependent variables, respectively, are also compared for this transonic flow in Figs. 10(b) and 10(c). The results show that the convergence rates of the error norms in density for the reconstruction methods based on primitive and conserved variables are very similar. Both methods recover the expected order of accuracy in the asymptotic limit. Furthermore, inspection of Figs. 10(b) and 10(c) reveals that the error norms in predicted pressure exhibit almost identical trends to those of density. It would seem that the same order of accuracy is recovered for all primitive variables. While formal analysis suggests that the point-wise mapping of average conserved to average primitive variables has the potential to introduce unwanted low-order errors, the numerical experiments for Ringleb's flow indicate that these errors are relatively small and do not noticeably affect the accuracy of the scheme.

5.7. Solution of Navier–Stokes equations

Finally, the application of the proposed high-order, CENO, finite-volume scheme to the numerical solution of the full Navier–Stokes equations is considered to illustrate the predictive capabilities of the solution method for laminar viscous flows. Numerical results are presented for a subsonic laminar flow past a NACA0012 airfoil at zero angle of attack, with

a free-stream Mach number of $M_\infty = 0.5$ and a Reynolds number of $Re = 5000$. This is a useful case for verification and validation purposes for it simultaneously tests the viscous flux discretization, inviscid flux formulation, effect of curved boundary treatment, and the ability to predict the underlying physics like flow separation and re-circulation. Moreover, this laminar airfoil case has been considered quite extensively in other previous studies [59,61,114,115], and reported results for low- and high-order schemes on both structured and unstructured meshes are available for comparison. Additionally, numerical predictions are presented for both steady and unsteady, subsonic, laminar flow past a circular cylinder with $M_\infty = 0.1$. The simulations were carried out for two Reynolds numbers, $Re = 30$ and $Re = 110$. There are no three-dimensional flow effects in either case [116]. For $Re = 30$, the flow is steady and possesses a steady wake with symmetric separated flow behind the cylinder. For $Re = 110$, the flow is unsteady and characterized by the periodic shedding of vortices and the formation of a von Kármán vortex street behind the cylinder. Fixed body-fitted meshes were used to obtain the steady flow solutions and the AMR scheme was used in obtaining high-order unsteady flow results.

5.7.1. Subsonic laminar flow past a NACA0012 symmetric airfoil

As mentioned above, this laminar airfoil case has been examined in previous studies [59,61,114,115]. The free-stream Mach and Reynolds numbers for this flow are $M_\infty = 0.5$ and $Re = 5000$, respectively, and for zero angle of attack the flow is both steady and symmetric with respect to the stagnation stream line and symmetric airfoil. However, flow separation occurs near the trailing edge, resulting in the formation of small re-circulation bubbles on the upper and lower surfaces of the airfoil that extend into the near-wake region and present numerical challenges for the accurate numerical prediction of the airfoil drag.

The computational domain for this case was constructed to be similar to those of previous studies [115]. The outer boundary of the domain was positioned at about 24 chord lengths from the airfoil surface. An O-type meshing strategy was used and a sequence of five uniformly refined multi-block body-fitted meshes arising from an initial coarse mesh were considered. The initial coarse mesh consisted of 96 cells along the upper and lower surfaces of the airfoil (48 cells for each side) and 32 cells in the direction normal to the surface. The 3072 cells were grouped into two initial grid blocks, each consisting of 48×32 cells. The grid lines in the tangential direction to the airfoil were clustered such that cell aspect ratios near the airfoil surface were in the range from 10 to 20. The higher-resolution meshes in the sequence were then generated by uniformly and successively refining each grid block to arrive at five meshes, having 8, 32, 128, 512 and 2048 solution blocks and a total number of $N_1 = 12{,}288$, $N_2 = 49{,}152$, $N_3 = 49{,}152$, $N_4 = 786{,}432$ and $N_5 = 3{,}145{,}728$ computational cells, respectively. Fig. 11(a) depicts the second computational grid in the sequence having 49,152 cells, showing the grid structure in the vicinity of the airfoil.

Computations were performed on each of the first four meshes using both high- and low-order schemes. Simulations on the finest mesh with $N_5 = 3{,}145{,}728$ computational cells were performed only for the low-order scheme to provide reference values and confirm the convergence of the results. Results were obtained for the proposed 4th-order, CENO, finite-volume, scheme with quartic reconstruction as well as for a standard, second-order, finite-volume method that used a combination of piecewise linear least-squares and diamond path reconstructions in the discretization of the hyperbolic and elliptic operators, respectively. The latter provides a reference for the high-order results. Details of the second-order scheme are given by Sachdev et al. [67]. The CFL number in these simulation was taken 0.4 for both the low- and high-order schemes. No-slip boundary conditions for velocity and adiabatic boundary conditions for temperature were imposed on the airfoil surface. The constrained reconstruction algorithm proposed here was used in the enforcement of these conditions in the high-order CENO computations.

Fig. 11(b) depicts the spatial distribution of the flow Mach number determined using the 4th-order ($k = 4$), CENO, finite-volume scheme on the second body-fitted mesh with $N_2 = 49{,}152$ computational cells. A close-up view of the predicted solution for the Mach number in the trailing-edge region of the flow is shown in Fig. 11(d). The re-circulation bubbles appear as regions of low Mach number and can be clearly identified by the symmetric set of flow streamlines. When compared to similar previous results [59,61,114,115], this predicted solution would seem to agree very well with those solutions obtained in previous studies, at least qualitatively.

The solution accuracy of the 4th-order CENO scheme is assessed here on the sequence of five grids based on the prediction of the extent of the re-circulation bubbles as measured by the locations of the separation point as well as the computed values of the pressure, viscous, and total drag. Similar metrics have been used and reported elsewhere [59,61,114,115]. The predictions of the 4th-order scheme for the separation point and coefficients of pressure drag, $C_{D,p}$, viscous drag, $C_{D,f}$, and total drag, C_D, are given and compared to those of the 2nd-order method in Fig. 11 for each of the five grids. The variation with mesh resolution of total drag is given in Fig. 11(e). Fig. 11(f) depicts the predicted values of the two contributions to the total drag coefficient, $C_{D,p}$ and $C_{D,f}$, as a function of mesh density. Finally, the predicted chord locations of the flow separation point, x_s, as identified by the zero of the derivative of tangential velocity component in the wall normal direction, are shown in Fig. 11(c). The separation location is expressed as a percentage of the chord length.

From the results of Fig. 11 is evident that, as the mesh is successively refined, the 2nd- and 4th-order schemes appear to be converging asymptotically to identical results for the various drag coefficients and separation location. The value of the total drag coefficient, C_D, predicted by the 2nd-order scheme on the 3,145,728-cell mesh is 0.05555 and the 4th-order CENO scheme yields a nearly equivalent value of 0.05553 on the mesh with 786,432 control volumes. Similarly, the predicted locations of the separation point, x_s, for the 2nd- and 4th-order schemes on the N_4 and N_5 meshes are 81.0% and 80.9%,

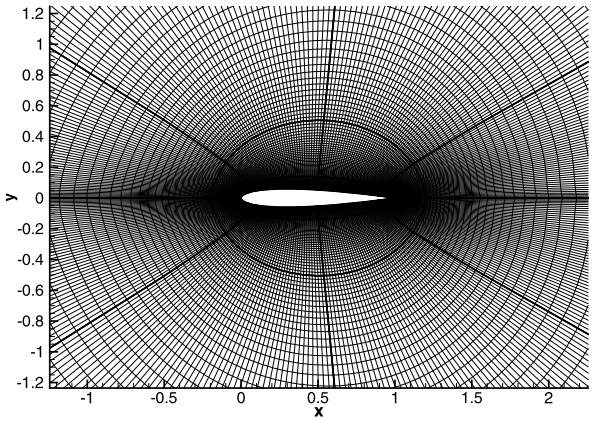

(a) Computational grid with 32 blocks of 48×32 cells and 49,152 cells in total. The thicker lines are block boundaries.

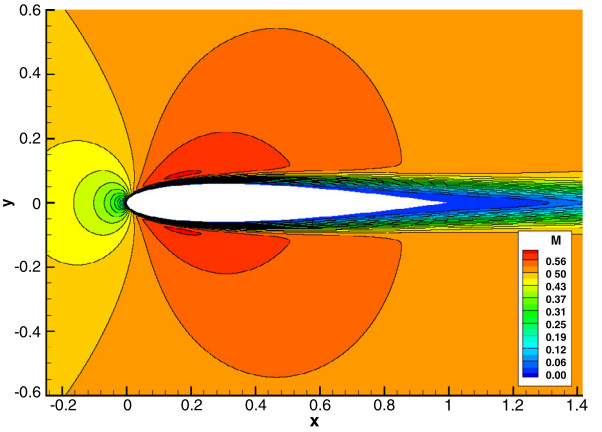

(b) Computed Mach contours (20 equally-spaced levels) by the 4th-order CENO scheme on the grid shown in (a).

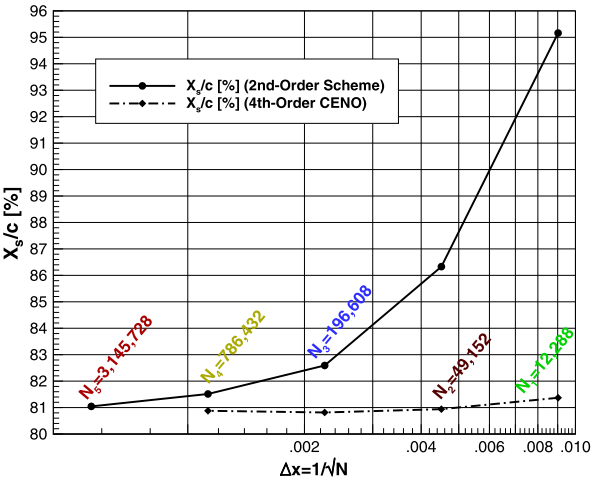

(c) Location of the separation point in percentages of chord length as a function of mesh resolution.

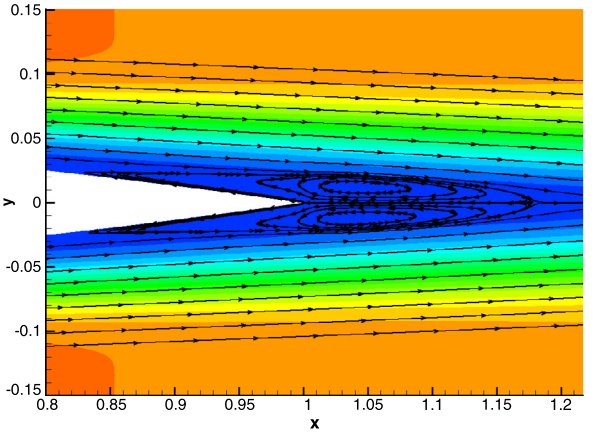

(d) Solution close-up in the trailing edge region, and illustration of the laminar separation bubble by velocity streamlines.

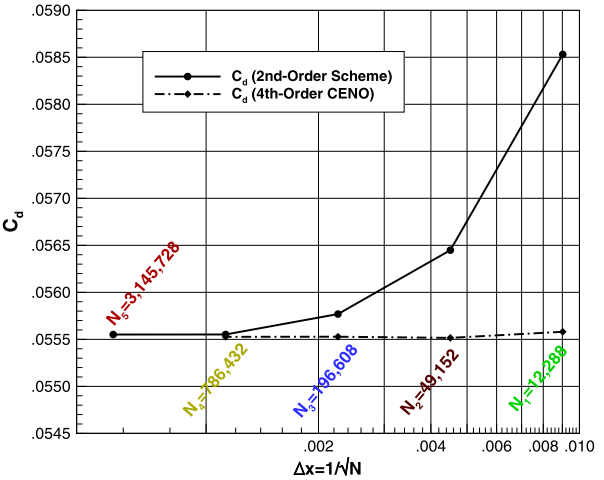

(e) Total drag coefficient as a function of mesh resolution, Δx.

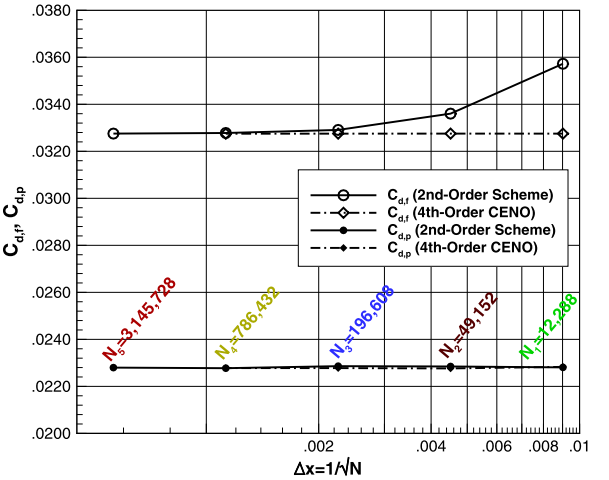

(f) Variation of drag coefficients with Δx.

Fig. 11. Computed Mach number distribution by the 4th-order CENO scheme and comparison of location of the separation point and the drag coefficients determined by the 2nd- and 4th-order schemes as a function of mesh resolution.

respectively. These values compare well and are within engineering accuracy to the values reported in previous studies [59,61,114,115]. It is also obvious from Fig. 11(c) that the pressure drag can be captured relatively accurately by both low- and high-order schemes, even on relatively coarse meshes. However, in contrast, the viscous drag is only predicted accurately on coarser meshes by the 4th-order scheme, whereas the 2nd-order method requires a significantly finer mesh to achieve a similar level of accuracy. A similar behaviour, but with greater differences, can be observed for the prediction of the flow separation point.

The results of Fig. 11 also demonstrate the potential benefits of the proposed high-order finite-volume scheme. Assuming that the values of the total drag and separation point as determined by the reference 2nd-order solution method are the correct fully-converged results, then the estimated relative errors in the predictions obtained by the 4th-order scheme on the coarsest mesh, N_1, with 12,288 cells is about 0.05% for C_D and 0.4% for x_s, whereas the corresponding errors for the 2nd-order scheme on the 12,288-cell mesh are 5.36% and 17.4%, respectively. Both errors for the 4th-order scheme on the coarse mesh are considerable less than 1% even though the coarse mesh is 256 times smaller than the finest mesh and has a spatial resolution that is coarser by a factor of 16. Moreover, to obtain an equivalent error in the prediction of the separation point (i.e., a 0.4% error in x_s), it is estimated from Fig. 11(c) that the second-order method would require a two-dimensional computational mesh that is about 123 times larger than the coarsest mesh, N_1, and has a finer spatial resolution by a factor of about 11. In terms of the accuracy as a function of the number of unknowns or degrees of freedom, the performance of the 4th-order CENO scheme is very similar if not slightly superior to that offered by the SD scheme of Sun et al. for this case [115], for which it can be estimated from their results that about 15,000 degrees of freedom are required to obtain an error in the separation point of less than 0.4%. The latter comparison is obviously affected by possible differences in the actual meshes used to obtain the solutions, but shows a rather good performance of the proposed CENO scheme relative to other recently proposed spatial discretization schemes.

The preceding observations suggest that there is a significant performance gain, in terms of accuracy, in going from the standard second-order approach to the proposed 4th-order CENO scheme. While no attempt has been made here to quantify directly the overall computational savings offered by the high-order method (such an evaluation was deemed beyond the scope of the present study due to several inefficiencies and inequities in the implementations of the reference second-order scheme, high-order approach, and time-marching schemes which could potentially lead to unfair comparisons), estimates of the relative cost of performing one residual evaluation per computational cell for the reference second-order and 4th-order CENO schemes are also quite favourable. While representing about 75% of the total computing time for the simulations with explicit time marching, the cost of evaluating the high-order spatial residual, $\mathbf{R}_{i,j}$, for cell (i, j) was found to be only about 2.9–3.4 times that of the second-order operator. The latter would imply that the proposed high-order CENO scheme has the potential to provide large computational savings when accurate predictions of separation and drag are required.

5.7.2. Steady subsonic laminar flow past a circular cylinder

The flow-field geometry considered in the simulations of the steady, subsonic, laminar flow past a cylinder consists of a domain between two concentric cylinders. The inner cylinder has a diameter $d_i = 0.0001$ m and the ratio of the outer and inner cylinder diameters is 40. Results for five non-uniformly-spaced, body-fitted, O-grid meshes were considered with clustering of the grid lines toward the inner cylinder. The five grids contained $M_1 = 3200$, $M_2 = 12{,}800$, $M_3 = 51{,}200$, $M_4 = 204{,}800$, and $M_5 = 3{,}276{,}800$ computational cells, respectively. As for the airfoil case, the computations for the cylinder were again performed on each of the meshes using both high- and low-order schemes. Results were obtained for the proposed 4th-order, CENO, finite-volume, scheme with quartic reconstruction as well as for the standard, second-order, finite-volume method described above. No-slip boundary conditions for velocity and adiabatic boundary conditions for temperature were imposed on the inner cylinder wall. The constrained reconstruction algorithm was again used in the enforcement of these conditions in the high-order CENO computations.

The predicted Mach number distribution for steady laminar flow past the cylinder at $\mathrm{Re} = 30$ obtained using the 4th-order ($k = 4$), CENO, finite-volume scheme on body-fitted mesh M_3 is shown in Fig. 12, along with the predicted cylinder drag coefficients obtained using both the second- and 4th-order schemes on the sequence of five computational meshes, M_1, M_2, M_3, M_4, through M_5. The re-circulating flow regions behind the cylinder following separation can be clearly identified by the streamlines of Fig. 12(a). As indicated in Fig. 12(b), the drag coefficients, C_D, predicted by the high-order scheme were $C_D = 1.7498$ for mesh M_1, $C_D = 1.7512$ for mesh M_2, $C_D = 1.7522$ for mesh M_3, $C_D = 1.7528$ for mesh M_4, and $C_D = 1.7541$ for mesh M_5. These values are generally in good agreement with the curve fits proposed by Henderson [117], which yield a value of $C_D = 1.737$ for this case, as well as available experimental data from the open literature [118,119]. Note that Henderson's fits are based on thorough numerical studies near the onset of vortex shedding and have been verified against experimental data [117].

Several conclusions can be drawn from the predictions of the drag coefficient given in Fig. 12(b). First of all, it is quite apparent that as the mesh is refined, the predictions of the proposed 4th-order CENO scheme and the reference second-order method converge to the same value of the drag coefficient. On the finest mesh, M_5, with more than three million cells, the difference between the predictions of the two schemes is just 6.5 drag counts. Secondly, assuming that the value of the drag obtained using the high-order scheme on the finest mesh, M_5, is the exact converged result, the error in the drag coefficient for the 4th-order CENO scheme on the coarsest mesh, M_1, relative to this best estimate is only 0.245%. This is a rather small error in drag considering that mesh M_1 is 1024 times smaller than mesh M_5 and has a spatial resolution that is coarser by a factor of 32. In contrast, the estimated error in drag for the reference second-order scheme on mesh

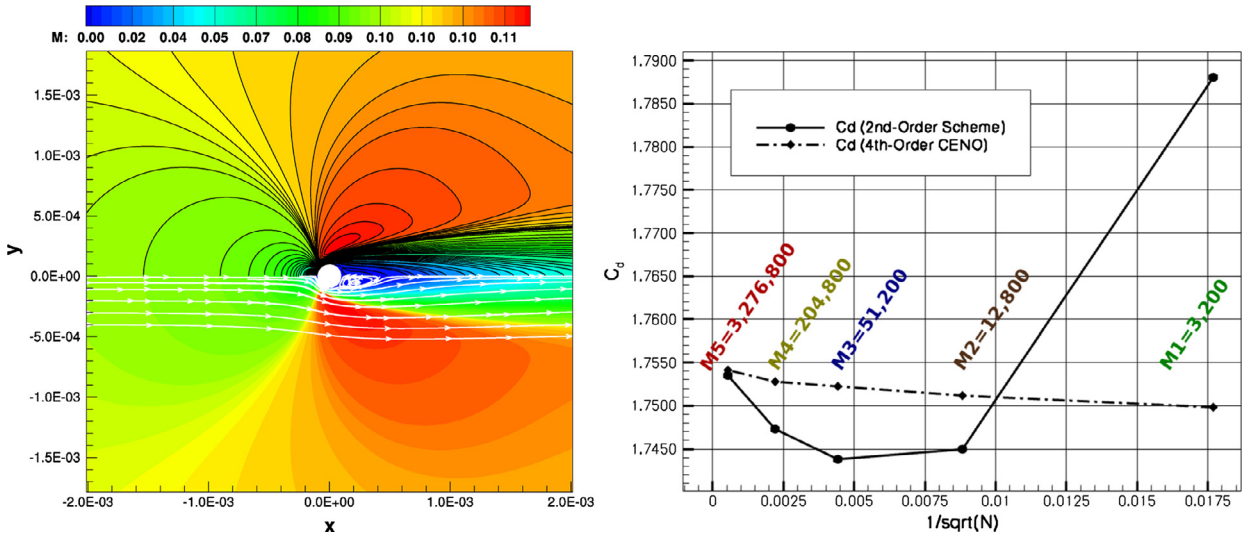

(a) Mach number prediction (b) Predicted drag coefficients of both schemes

Fig. 12. Predicted numerical results for steady, subsonic, laminar flow past a circular cylinder with a free-stream Mach number of $M_\infty = 0.1$ and Reynolds number of $Re = 30$ showing: (a) predicted distribution of Mach number, M, and streamlines obtained using the 4th-order, CENO, finite-volume scheme on mesh M_3 with 51,200 cells; and (b) comparison of predicted drag coefficients, C_D, for the 4th-order CENO and standard second-order finite-volume methods as a function of mesh density ($\Delta x = \frac{1}{\sqrt{N}}$) obtained using meshes M_1, M_2, M_3, M_4, and M_5.

M_1 is about 8 times larger with a 1.966% relative error compared to the 4th-order solution on the finest mesh. Finally, the results of Fig. 12(b) would suggest that to obtain the identical error in drag as that of the high-order CENO scheme for mesh M_1 (i.e., a 0.245% error in drag), the second-order method would require a two-dimensional computational mesh that is about 132 times larger than M_1 and has a finer spatial resolution by a factor of 11–12. Somewhat surprisingly, these values for the reductions in the mesh or degrees of freedom offered by the high-order approach are very similar to those found for the NACA0012 airfoil case described previously. As for the airfoil case, the possible performance gains, in terms of accuracy, in going from a standard second-order approach to the proposed 4th-order CENO scheme are very evident. When the CENO method is eventually combined with an efficient convergence acceleration technique and/or time-marching scheme, the potential for appreciable computations savings would seem rather significant.

5.7.3. Parallel performance

The parallel performance of the proposed high-order CENO algorithm has also been assessed for this steady, subsonic, laminar cylinder flow of the previous section. The parallel speedup and efficiency, defined as

$$S_p = \frac{t_1}{t_p}, \qquad E_p = \frac{S_p}{p}, \tag{44}$$

respectively, were both evaluated for this case, where t_1 and t_p are the total computational times required to solve the problem using a single processor core and using p cores, respectively. A strong scaling study was considered in which these parallel performance measures were estimated for a fixed size problem consisting of 2048, 20×20-cell, grid blocks with 819,200 computational cells. The scalability was computed using up to 1024 cores. The computations were performed on a high performance parallel cluster consisting 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.

Fig. 13 provides the actual computed parallel speedup and efficiency achieved for the 4th-order CENO method as a function of the number of computing cores. The results are compared directly to the corresponding ideal values assuming the computation for p cores is p times smaller than t_1. The results show that, despite the expected increases in inter-processor communication and decreases in computational work per core with increasing core count, rather impressive parallel efficiencies in excess of 70% are achieved in all cases up to 1024 cores for this relatively small two-dimensional strong-scaling problem.

5.7.4. Unsteady subsonic laminar flow past a circular cylinder with AMR

As a final assessment of the predictive capabilities of the proposed high-order, CENO, finite-volume schemes when also used in conjunction with the block-based AMR method, the prediction of unsteady vortex shedding for subsonic flow past a circular cylinder is considered. In this case, a cylinder flow with a Reynolds number of $Re = 110$ was studied. A computational domain with $d_i = 0.0001$ m and an outer boundary positioned at 80 inner diameters was used. A quasi-steady

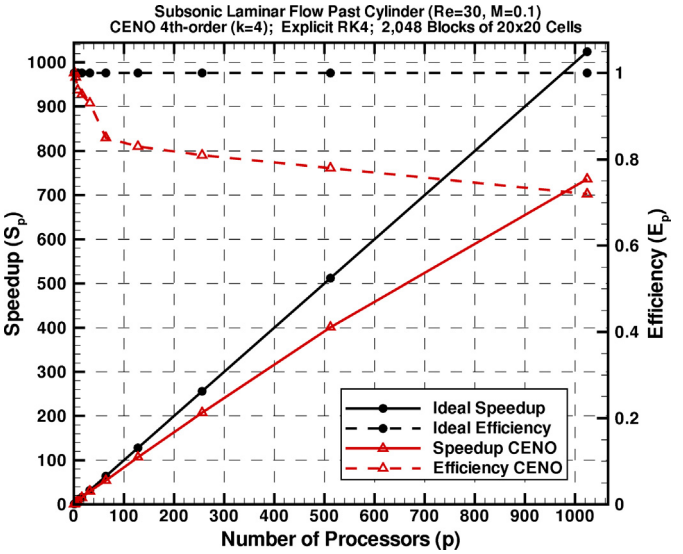

Fig. 13. Parallel speedup, S_p, and efficiency, E_p, of the 4th-order, CENO, finite-volume scheme as a function of the number of processor cores, p, for steady, subsonic, laminar flow past a circular cylinder with a free-stream Mach number of $M_\infty = 0.1$ and Reynolds number of $Re = 30$.

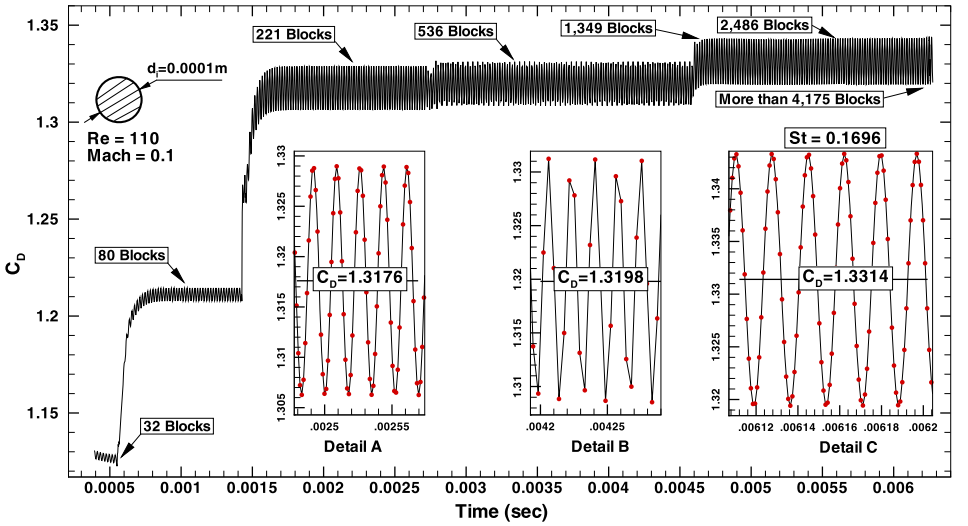

Fig. 14. Predicted drag coefficient, C_D, as a function of physical time, t, for unsteady, subsonic, laminar flow past a circular cylinder with a free-stream Mach number of $M_\infty = 0.1$ and Reynolds number of $Re = 110$ obtained using the 4th-order, CENO, finite-volume scheme with block-based AMR. The three inserts provide details of the instantaneous and mean drag for the computational meshes with 221 (Detail A), 536 (Detail B), and 2486 blocks (Detail C), respectively.

periodic solution was first obtained on an initial O-grid mesh of 32 self-similar 8×8 grid blocks containing 2048 computational cells. Mesh adaptation as directed by the proposed h-refinement criterion based on the CENO smoothness indicator was then applied successively to the coarse-mesh solution to improve the accuracy. Over a period of simulated physical time from $t = 0.00055$ s to $t = 0.006225$ s, a sequence of five new solution-dependent meshes were then created by applying the AMR procedure. In each case, a quasi-steady periodic solution was again obtained on the refined mesh prior to performing another level of the mesh adaptation. The five new meshes generated by the block-based refinement algorithm consisted of 80, 221, 536, 1349 and 2486 8×8 grid blocks and 5120, 14,144, 34,304, 86,336 and 159,104 cells, respectively. After $t = 0.006225$ s, further dynamic mesh adaptation was performed at regular time intervals. Following an initial increase to 4175 grid blocks, the AMR mesh resolution gradually increased during this final period until a quasi-steady periodic solution was obtained on meshes ranging in size from 10,000 to 11,000 grid blocks with 600,000–700,000 computational cells and 10–12 levels of mesh refinement.

The predicted unsteady drag coefficient, C_D, for the $Re = 110$ circular cylinder flow is depicted in Fig. 14. The simulated temporal variation of C_D is shown for the entire sequence of AMR meshes starting with the initial under-resolved mesh having just 32 grid blocks and 2048 computational cells, on through five additional grids for which quasi-steady periodic

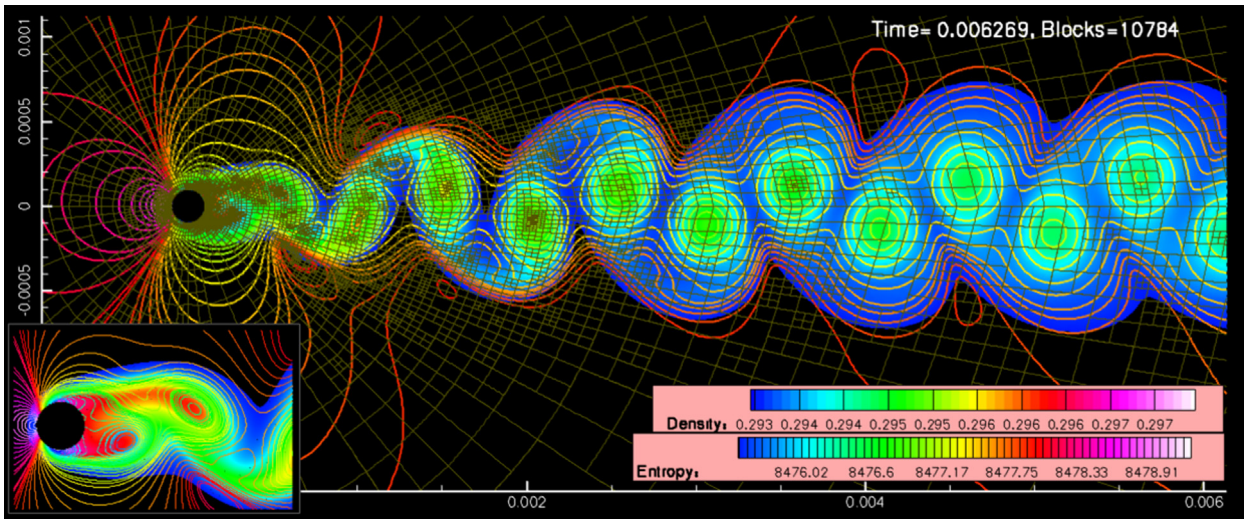

Fig. 15. Predicted instantaneous distributions of density, ρ, and entropy, s, as well as dynamically refined mesh, at $t = 0.006269$ s for unsteady, subsonic, laminar flow past a circular cylinder with a free-stream Mach number of $M_\infty = 0.1$ and Reynolds number of $Re = 110$ obtained using the 4th-order, CENO, finite-volume scheme. The instantaneous AMR mesh contains 10,784 grid blocks and 690,176 computational cells.

solutions were obtained without dynamic mesh adaptation, to the final dynamically adapted meshes having 10,000–11,000 grid blocks and 600,000–700,000 computational cells. The three inserts show the establishment of quasi-periodic behaviour as well as details of the instantaneous and mean drag for the computational meshes with 221 (Detail A), 536 (Detail B), and 2486 blocks (Detail C), respectively. While it is apparent that the mean drag is under-predicted on the initial mesh, the predicted solutions exhibit convergence to a grid-independent result as the mesh is successively refined by the AMR algorithm as described above. On the finest dynamically adapted meshes, the best predictions of the mean drag coefficient and Strouhal number, St, are $C_D = 1.3314$ and $St = 0.1696$, respectively, which both agree very well with established values from the literature. The Strouhal number for this case based on the relationship proposed by Roshko [120] to describe a best fit to experimental data is $St = 0.1711$. A value of $C_D = 1.34$ is reported in the computational study by Henderson [117].

The capabilities of the proposed AMR scheme for this case can be seen by considering Fig. 15, which depicts the predicted instantaneous distributions of density, ρ, and entropy, s, as well as dynamically refined mesh, at $t = 0.006269$ s for the unsteady, subsonic, laminar, cylinder flow. The instantaneous AMR mesh contains 10,784 grid blocks and 690,176 computational cells at $t = 0.006269$ s. The figure shows the distribution of the AMR blocks in the vicinity of the cylinder and the wake behind it containing the shed vortices. The h-refinement criterion based on the smoothness indicator is clearly able to identify the boundary layer, regions of flow detachment, shed vortices, and other wake-flow structures and these areas are targeted for mesh enrichment. When applied at a regular frequency throughout the calculation, the AMR mesh is then able to "track" and follow the unsteady flow features and vortices transported downstream, as shown in the figure. Note that the portion of the computational domain shown in Fig. 15 spans more than 60 cylinder diameters, yet the wake and shed vortices are well resolved, even at relatively large distances from the trailing edge of the cylinder.

5.7.5. Predicted drag for subsonic laminar flow past a circular cylinder

As a final summary of the subsonic, laminar, cylinder flow computations, a comparison of the predicted drag coefficients, C_D, for the $Re = 30$ and $Re = 110$ cases obtained using the proposed 4th-order, CENO, finite-volume scheme to the experimental measurements of Wieselsberger [119], two- and three-dimensional computations by Henderson [117,121], and other more recent numerical predictions by Sheard et al. [121] was made, the results of which are given in Fig. 16. It is quite evident from the figure that the drag predicted by the CENO scheme for both the steady ($Re = 30$) and unsteady ($Re = 110$) cases agree extremely well with the previously reported values. Despite only considering cylinder flows at two Reynolds numbers, it is felt that these numerical results demonstrate the validity of the proposed finite-volume approach and indicate that the high-order CENO scheme in combination with the block-based AMR algorithm has the potential to provide very accurate and reliable predictions of viscous laminar flows while significantly reducing the number of required computational elements.

6. Discussion and concluding remarks

A new high-order, CENO, finite-volume scheme with AMR has been proposed for solving the two-dimensional form of the Navier–Stokes equations governing compressible, laminar, viscous flows on body-fitted, multi-block, quadrilateral mesh. The procedure represents an extension of the CENO approach previously developed for inviscid flow simulations to viscous

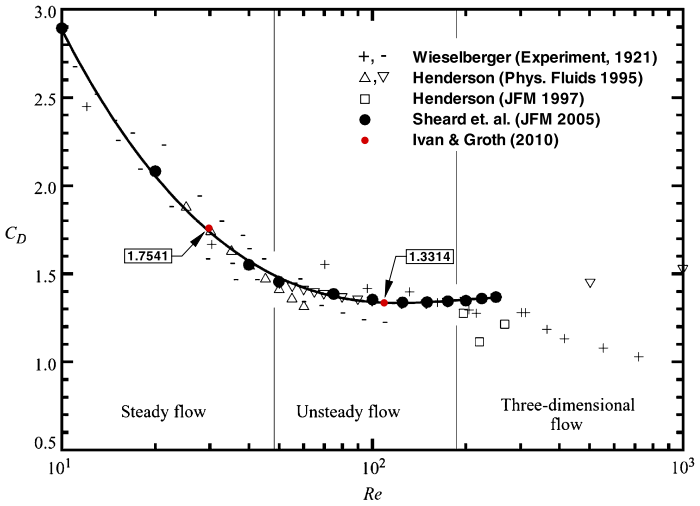

Fig. 16. Variation of the drag coefficient, C_D, as a function of the Reynolds number, Re, for subsonic laminar flow past a circular cylinder showing comparisons of the predictions of the 4th-order, CENO, finite-volume scheme for both $Re = 30$ and $Re = 110$ to experimental measurements of Wieselsberger [119], two- and three-dimensional computations by Henderson [117,121], and other more recent numerical predictions by Sheard et al. [121]. The vertical lines indicate the Reynolds numbers for transition from steady to unsteady flow ($Re \approx 47$) and from two- to three-dimensional flow ($Re \approx 189$).

flows of thermally and calorically perfect gases. The verification and validation of the proposed high-order AMR algorithm has been accomplished by comparing predicted solutions to a variety of available analytical results, previously reported computations, and experimental data. The numerical properties and validity of the proposed discretization scheme for the elliptic (viscous) operator were assessed and verified. The ability of the scheme to accurately represent solutions with smooth extrema and yet robustly handle under-resolved and/or non-smooth solution content (i.e., solutions with discontinuities) was demonstrated. The usefulness of the *h*-refinement criterion based on the smoothness indicator was evaluated and shown to provide a robust and reliable mesh adaptation algorithm that is capable of refining both under-resolved and non-smooth regions of the solution, while not unnecessarily over-refining resolved content. Finally, the extensive numerical experiments have well demonstrated the potential of the high-order and AMR schemes to reduce the overall computational cost of large-scale numerical simulations by reducing the grid resolution requirements for a given solution accuracy. Future research will involve further investigation of the CENO approach, the application of the method to more complex flows, its extension to three-dimensional problems and unstructured meshes, and the development of the scheme for use in conjunction with more efficient, high-order, time-marching schemes for both steady and unsteady applications.

Acknowledgements

The authors are very grateful to Dr. Gregory Sheard of the Mechanical and Aerospace Engineering Department at Monash University for providing the cylinder drag coefficient data. The authors would also like to thank the reviewers for their many useful comments and suggestions which have greatly helped to improve the manuscript and the discussion of results. 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

[1] P. Moin, K. Mahesh, Direct numerical simulation: a tool in turbulence research, Annu. Rev. Fluid Mech. 30 (1998) 539–578.
[2] L. Jameson, AMR vs. high order schemes, J. Sci. Comput. 18 (2003) 1–24.
[3] A. Harten, B. Enquist, S. Osher, S.R. Chakravarthy, Uniformly high order accurate essentially non-oscillatory schemes, III, J. Comput. Phys. 71 (1987) 231–303.
[4] R. Abgrall, On essentially non-oscillatory schemes on unstructured meshes: analysis and implementation, J. Comput. Phys. 114 (1994) 45–58.
[5] T. Sonar, On the construction of essentially non-oscillatory finite volume approximations to hyperbolic conservation laws on general triangulations: polynomial recovery, accuracy and stencil selection, Comput. Methods Appl. Mech. Eng. 140 (1997) 157.
[6] G.-S. Jiang, C.-W. Shu, Efficient implementation of weighted ENO schemes, J. Comput. Phys. 126 (1996) 202–228.
[7] O. Friedrich, Weighted essentially non-oscillatory schemes for the interpolation of mean values on unstructured grids, J. Comput. Phys. 144 (1998) 194–212.
[8] C. Hu, C.-W. Shu, Weighted essentially non-oscillatory schemes on triangular meshes, J. Comput. Phys. 150 (1999) 97–127.
[9] D. Stanescu, W. Habashi, Essentially nonoscillatory Euler solutions on unstructured meshes using extrapolation, AIAA J. 36 (1998) 1413–1416.
[10] A. Haselbacher, A WENO reconstruction algorithm for unstructured grids based on explicit stencil construction, Paper 2005-0879, AIAA, 2005.
[11] W.R. Wolf, J. Azevedo, High-order unstructured essentially nonoscillatory and weighted essentially nonoscillatory schemes for aerodynamic flows, AIAA J. 44 (10) (2006) 2295–2310.
[12] T.J. Barth, Recent developments in high order K-exact reconstruction on unstructured meshes, Paper 93-0668, AIAA, 1993.

[13] T.J. Barth, P.O. Fredrickson, Higher order solution of the Euler equations on unstructured grids using quadratic reconstruction, Paper 90-0013, AIAA, 1990.
[14] C.F. Ollivier-Gooch, High-order ENO schemes for unstructured meshes based on least-squares reconstruction, Paper 97-0540, AIAA, 1997.
[15] C.F. Ollivier-Gooch, Quasi-ENO schemes for unstructured meshes based on unlimited data-dependent least-squares reconstruction, J. Comput. Phys. 133 (1997) 6–17.
[16] 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.
[17] K. Michalak, C.F. Ollivier-Gooch, Accuracy preserving limiter for the high-order accurate solution of the Euler equations, J. Comput. Phys. 228 (23) (2009) 8693–8711, http://dx.doi.org/10.1016/j.jcp.2009.08.021.
[18] G. Capdeville, Towards a compact high-order method for non-linear hyperbolic systems. I: The Hermite Least-Square Monotone (HLSM) reconstruction, J. Comput. Phys. 228 (2009) 3762–3788.
[19] P. Colella, M.R. Dorr, J.A.F. Hittinger, D.F. 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.
[20] P. McCorquodale, P. Colella, A high-order finite-volume method for conservation laws on locally refined grids, Commun. Appl. Math. Comput. Sci. 6 (1) (2011) 1–25.
[21] P. Colella, M.D. Sekora, A limiter for PPM that preserves accuracy at smooth extrema, J. Comput. Phys. 227 (2008) 7069–7076.
[22] B. Cockburn, C.-W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite-element method for conservation laws II: General framework, Math. Comput. 52 (1989) 411.
[23] B. Cockburn, S.-Y. Lin, C.-W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite-element method for conservation laws III: One-dimensional systems, J. Comput. Phys. 84 (1989) 90.
[24] B. Cockburn, S. Hou, C.-W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite-element method for conservation laws IV: The multi-dimensional case, J. Comput. Phys. 54 (1990) 545.
[25] J. Qui, C.-W. Shu, Runge–Kutta discontinuous Galerkin method using WENO limiters, SIAM J. Sci. Comput. 26 (3) (2005) 907–929.
[26] G.E. Barter, D.L. Darmofal, Shock capturing with higher-order, PDE-based artificial viscosity, Paper 2007-3823, AIAA, 2007.
[27] G.E. Barter, D.L. Darmofal, Shock capturing with PDE-based artificial viscosity for DGFEM: Part I. Formulation, J. Comput. Phys. 229 (5) (2010) 1810–1827.
[28] Z. Xu, Y. Liu, C.-W. Shu, Hierarchical reconstruction for discontinuous Galerkin methods on unstructured grids with a WENO-type linear reconstruction and partial neighboring cells, J. Comput. Phys. 228 (2009) 2194–2212.
[29] E.J. Kubatko, C. Dawson, J.J. Westerink, Time step restrictions for Runge–Kutta discontinuous Galerkin methods on triangular grids, J. Comput. Phys. 227 (2008) 9697–9710.
[30] 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 (2008) 8209–8253.
[31] Z.J. Wang, Spectral (finite) volume method for conservation laws on unstructured grids: basic formulation, J. Comput. Phys. 178 (2002) 210–251.
[32] Z.J. Wang, L. Zhang, Y. Liu, Spectral (finite) volume method for conservation laws on unstructured grids IV: Extension to two-dimensional systems, J. Comput. Phys. (ISSN 0021-9991) 194 (2) (2004) 716–741, http://dx.doi.org/10.1016/j.jcp.2003.09.012.
[33] Z.J. Wang, Y. Liu, G. May, A. Jameson, Spectral difference method for unstructured grids II: Extension to the Euler equations, J. Sci. Comput. (ISSN 0885-7474) 32 (1) (2007) 45–71, http://dx.doi.org/10.1007/s10915-006-9113-9.
[34] Y. Liu, M. Vinokur, Z. Wang, Discontinuous spectral difference method for conservation laws on unstructured grids, in: Proceedings of the 3rd International Conference on Computational Fluid Dynamics, Toronto, Canada, July 2004, pp. 12–16.
[35] H. Huynh, A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods, Paper 2007-4079, AIAA, 2007.
[36] M. Yang, Z.J. Wang, A parameter-free generalized moment limiter for high-order methods on unstructured grids, Tech. Rep. 2009-605, AIAA, 2009.
[37] M. Yang, Z.J. Wang, A parameter-free generalized moment limiter for high-order methods on unstructured grids, Adv. Appl. Math. Mech. 1 (4) (2009) 451–480.
[38] Z.J. Wang, H. Gao, A unifying lifting collocation penalty formulation including the discontinuous Galerkin, spectral volume/difference methods for conservation laws on mixed grids, J. Comput. Phys. (ISSN 0021-9991) 228 (21) (2009) 8161–8186, http://dx.doi.org/10.1016/j.jcp.2009.07.036.
[39] R. Abgrall, P.L. Roe, High order fluctuation schemes on triangular mesh, J. Comput. Phys. 19 (1–3) (2003) 3–36.
[40] R. Abgrall, Essentially non-oscillatory residual distribution schemes for hyperbolic problems, J. Comput. Phys. 214 (2006) 773–808.
[41] S. Guzik, C.P.T. Groth, Comparison of solution accuracy of multidimensional residual distribution and Godunov-type finite-volume methods, Int. J. Comput. Fluid Dyn. 22 (1) (2008) 61–68.
[42] R. Abgrall, A. Larat, M. Ricchiuto, Construction of very high order residual distribution schemes for steady inviscid flow problems on hybrid unstructured meshes, J. Comput. Phys. 230 (11) (2011) 4103–4136.
[43] S.C. Dias, D.W. Zingg, A high-order parallel Newton–Krylov flow solver for the Euler equations, Paper 2009-3657, AIAA, 2009.
[44] W.J. Coirier, An adaptively-refined, Cartesian, cell-based scheme for the Euler and Navier–Stokes equations, Ph.D. thesis, University of Michigan, 1994.
[45] M. Delanaye, M.J. Aftomis, M.J. Berger, Y. Liu, T.H. Pulliam, Automatic hybrid-Cartesian grid generation for high-Reynolds number flows around complex geometries, Paper 99-0777, AIAA, 1999.
[46] M. Barad, P. Colella, A fourth-order accurate local refinement method for Poisson's equation, J. Comput. Phys. 209 (2005) 1–18.
[47] C.F. Ollivier-Gooch, M. Van Altena, A high-order accurate unstructured mesh finite-volume scheme for the advection–diffusion equation, J. Comput. Phys. 181 (2) (2002) 729–752.
[48] F. Bassi, S. Rebay, A high-order accurate discontinuous finite element method for the numerical solution of the compressible Navier–Stokes equations, J. Comput. Phys. 131 (1997) 267–279.
[49] B. Cockburn, C.-W. Shu, The local discontinuous Galerkin method for time-dependent convection diffusion system, SIAM J. Numer. Anal. 35 (6) (1998) 2440–2463.
[50] B. van Leer, M. Lo, M. van Raalte, A discontinuous Galerkin method for diffusion based on recovery, Paper 2007-4083, AIAA, 2007.
[51] M. van Raalte, B. van Leer, Bilinear forms for the recovery-based discontinuous Galerkin method for diffusion, Commun. Comput. Phys. 5 (2–4) (2009) 683–693.
[52] G. Gassner, F. Lörcher, C.-D. Munz, A contribution to the construction of diffusion fluxes for finite volume and discontinuous Galerkin schemes, J. Comput. Phys. 224 (2007) 1049–1063.
[53] F. Lörcher, G. Gassner, C.-D. Munz, An explicit discontinuous Galerkin scheme with local time-stepping for general unsteady diffusion equations, J. Comput. Phys. 227 (2008) 5649–5670.
[54] H. Liu, J. Yan, The direct discontinuous Galerkin (DDG) methods for diffusion problems, SIAM J. Numer. Anal. 41 (1) (2009) 675–698.
[55] T.A. Oliver, D.L. Darmofal, An unsteady adaptation algorithm for discontinuous Galerkin discretizations of the RANS equations, Paper 2007-3940, AIAA, 2007.
[56] Y. Sun, Z.J. Wang, Y. Liu, Spectral (finite) volume method for conservation laws on unstructured grids VI: Extension to viscous flow, J. Comput. Phys. (ISSN 0021-9991) 215 (1) (2006) 41–58, http://dx.doi.org/10.1016/j.jcp.2005.10.019.

[57] G. May, A. Jameson, High-order accurate methods for high-speed flow, Paper 2005-5251, AIAA, 2005.
[58] G. May, A. Jameson, A spectral difference method for the Euler and Navier–Stokes equations on unstructured meshes, Paper 2006-304, AIAA, 2006.
[59] R. Kannan, Z.J. Wang, A study of viscous flux formulations for a p-multigrid spectral volume Navier–Stokes solver, J. Sci. Comput. 41 (2) (2009) 165–199.
[60] R. Kannan, Z.J. Wang, The direct discontinuous Galerkin (DDG) viscous flux scheme for the high order spectral volume method, 39 (10) (2010) 2007–2021.
[61] R. Kannan, Z.J. Wang, A variant of the LDG flux formulation for the spectral volume method, J. Sci. Comput. 46 (2) (2011) 314–328.
[62] R. Kannan, Z.J. Wang, A high order spectral volume solution to the Burgers' equation using the Hopf–Cole transformation, Int. J. Numer. Methods Fluids 69 (4) (2012) 781–801.
[63] T. Haga, H. Gao, Z.J. Wang, A high-order unifying discontinuous formulation for the Navier–Stokes equations on 3D mixed grids, Math. Model. Nat. Phenom. 6 (3) (2011) 28–56.
[64] S. De Rango, D. Zingg, Higher-order spatial discretization for turbulent aerodynamic computations, AIAA J. 39 (2001) 1296–1304.
[65] L. Ivan, C.P.T. Groth, High-order central ENO finite-volume scheme with adaptive mesh refinement, Paper 2007-4323, AIAA, 2007.
[66] 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, 2011.
[67] 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.
[68] 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.
[69] X. Gao, S. 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.
[70] J. Shi, C. Hu, C. Shu, A technique of treating negative weights in WENO schemes, J. Comput. Phys. 175 (2002) 108–127.
[71] A. Harten, B. Enquist, S. Osher, S.R. Chakravarthy, Uniformly high order accurate essentially non-oscillatory schemes, III, Report 86-22, ICASE, 1986.
[72] M. Dumbser, M. Käser, Arbitrary high order non-oscillatory finite volume schemes on unstructured meshes for linear hyperbolic systems, J. Comput. Phys. 221 (2007) 693–723.
[73] G. Capdeville, A central WENO scheme for solving hyperbolic conservation laws on non-uniform meshes, J. Comput. Phys. 227 (2008) 2977–3014.
[74] I.G. Petrovsky, Lectures on Partial Differential Equations, Interscience, New York, 1954.
[75] L. Bers, F. John, M. Schechter, Partial Differential Equations, Interscience, New York, 1966.
[76] J. McDonald, C.P.T. Groth, Numerical modeling of micron-scale flows using the Gaussian moment closure, Paper 2005-5035, AIAA, 2005.
[77] S.A. Northrup, C.P.T. Groth, Solution of laminar diffusion flames using a parallel adaptive mesh refinement algorithm, Paper 2005-0547, AIAA, 2005.
[78] S.K. Godunov, Finite-difference method for numerical computations of discontinuous solutions of the equations of fluid dynamics, Mat. Sb. 47 (1959) 271–306.
[79] J.J. Gottlieb, C.P.T. Groth, Assessment of Riemann solvers for unsteady one-dimensional inviscid flows of perfect gases, J. Comput. Phys. 78 (1988) 437–458.
[80] E.F. Toro, Riemann Solvers and Numerical Methods for Fluid Dynamics, Springer-Verlag, New York, 1999.
[81] W.H. Press, S.A. Teukolsky, W.T. Vetterling, B.P. Flannery, Numerical Recipes. The Art of Scientific Computing, 3rd edition, Cambridge University Press, New York, 2007.
[82] A. Harten, S.R. Chakravarthy, Multi-dimensional ENO schemes for general geometries, Report 91-76, ICASE, 1991.
[83] A. Harten, High resolution schemes for hyperbolic conservation laws, J. Comput. Phys. 49 (1983) 357–393.
[84] X. Zhang, C.-W. Shu, On maximum-principle-satisfying high order schemes for scalar conservation laws, J. Comput. Phys. 229 (2010) 3091–3120.
[85] X. Zhang, C.-W. Shu, On positivity preserving high order discontinuous Galerkin schemes for compressible Euler equations on rectangular meshes, J. Comput. Phys. 229 (2010) 8918–8934.
[86] D.S. Balsara, Self-adjusting, positivity preserving high order schemes for hydrodynamics and magnetohydrodynamics, J. Comput. Phys. 231 (22) (2012) 7504–7517.
[87] D.J. Mavriplis, Revisiting the least-squares procedure for gradient reconstruction on unstructured meshes, Paper 2003-3986, AIAA, 2003.
[88] C. Lawson, R. Hanson, Solving Least Squares Problems, Prentice Hall, Inc, 1974.
[89] R. Penrose, A generalized inverse for matrices, Proc. Camb. Philos. Soc. 51 (1955) 406–413.
[90] A. Jalali, C. Ollivier-Gooch, Higher-order finite volume solution reconstruction on highly anisotropic meshes, Paper 2013-2565, AIAA, 2013.
[91] B. van Leer, Towards the ultimate conservative difference scheme. V. A second-order sequel to Godunov's method, J. Comput. Phys. 32 (1979) 101–136.
[92] V. Venkatakrishnan, On the accuracy of limiters and convergence to steady state solutions, Paper 93-0880, AIAA, 1993.
[93] P.L. Roe, Approximate Riemann solvers, parameter vectors, and difference schemes, J. Comput. Phys. 43 (1981) 357–372.
[94] B. Einfeldt, On Godunov-type methods for gas dynamics, SIAM J. Numer. Anal. 25 (1988) 294–318.
[95] T.J. Linde, A three-dimensional adaptive multifluid MHD model of the heliosphere, Ph.D. thesis, University of Michigan, 1998.
[96] T. Linde, A practical, general-purpose, two-state HLL Riemann solver for hyperbolic conservation laws, Int. J. Numer. Methods Fluids 40 (2002) 391–402.
[97] D.J. Mavriplis, Unstructured mesh discretizations and solvers for computational aerodynamics, Paper 2007-3955, AIAA, 2007.
[98] H. Lomax, T. Pulliam, D. Zingg, Fundamentals of Computational Fluid Dynamics, Springer-Verlag, New York, 2001.
[99] D.S. Balsara, C. Meyer, M. Dumbser, H. Du, Z. Xu, Efficient implementation of ADER schemes for Euler and magnetohydrodynamical flows on structured meshes — Speed comparisons with Runge–Kutta methods, J. Comput. Phys. 235 (2013) 934–969.
[100] H. Bijl, M.H. Carpenter, V.N. Vatsa, C.A. Kennedy, Implicit time integration schemes for the unsteady compressible Navier–Stokes equations: laminar flow, J. Comput. Phys. 179 (2002) 313–329.
[101] B. van Leer, C.H. Tai, K.G. Powell, Design of optimally-smoothing multi-stage schemes for the Euler equations, Paper 89-1933-CP, AIAA, 1989.
[102] B. van Leer, W.-T. Lee, P.L. Roe, K.G. Powell, C.-H. Tai, Design of optimally smoothing multistage schemes for the Euler equations, Commun. Appl. Numer. Methods 8 (10) (1992) 761–769.
[103] C. Hirsch, Numerical Computation of Internal and External Flows, vol. 1, Fundamentals of Numerical Discretization, John Wiley & Sons, Toronto, 1989.
[104] C. Hirsch, Numerical Computation of Internal and External Flows, vol. 2, Computational Methods for Inviscid and Viscous Flows, John Wiley & Sons, Toronto, 1990.
[105] 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.
[106] M.J. Berger, J. Oliger, Adaptive mesh refinement for hyperbolic partial differential equations, J. Comput. Phys. 53 (1984) 484–512.
[107] M.J. Berger, P. Colella, Local adaptive mesh refinement for shock hydrodynamics, J. Comput. Phys. 82 (1989) 67–84.
[108] D. De Zeeuw, K.G. Powell, An adaptively refined Cartesian mesh solver for the Euler equations, J. Comput. Phys. 104 (1993) 56–68.
[109] W.J. Coirier, K.G. Powell, An accuracy assessment of Cartesian-mesh approaches for the Euler equations, J. Comput. Phys. 117 (1995) 121–131.
[110] M.J. Aftomis, M.J. Berger, J.E. Melton, Robust and efficient Cartesian mesh generation for component-base geometry, AIAA J. 36 (6) (1998) 952–960.

[111] M.J. Aftomis, M.J. Berger, G. Adomavicius, A parallel multilevel method for adaptively refined Cartesian grids with embedded boundaries, Paper 2000-0808, AIAA, 2000.
[112] AGARD Subcommittee C, Test cases for inviscid flow field methods, AGARD Advisory Report 211, 1986.
[113] D.S. Balsara, C.-H. Shu, Monotonicity preserving weighted essentially non-oscillatory schemes with increasingly high order of accuracy, J. Comput. Phys. 160 (2) (2000) 405–452.
[114] D.J. Mavriplis, A. Jameson, L. Martinelli, Multigrid solution of the Navier–Stokes equations on triangular meshes, Paper 89-0120, AIAA, 1989.
[115] Y. Sun, Z.J. Wang, Y. Liu, Efficient implicit non-linear LU-SGS approach for compressible flow computation using high-order spectral difference method, Commun. Comput. Phys. 5 (2–4) (2009) 760–778.
[116] C. Williamson, Vortex dynamics in the cylinder wake, Annu. Rev. Fluid Mech. 28 (1996) 477–539.
[117] R.D. Henderson, Details of the drag curve near the onset of vortex shedding, Phys. Fluids 7 (9) (1995) 2102–2104.
[118] H. Schlichting, K. Gersten, Boundary-Layer Theory, 8th ed., Springer, New York, 2000.
[119] C. Wieselsberger, New data on the laws of fluid resistance, Technical Note 84, NACA, 1922.
[120] A. Roshko, On the development of turbulent wakes from vortex streets, Technical Note 1191, NACA, 1954.
[121] G.J. Sheard, K. Hourigan, M.C. Thompson, Computational of the drag coefficients for low-Reynolds-number flow past rings, J. Fluid Mech. 526 (2005) 257–275.
[122] 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.
[123] T. Barth, D.C. Jespersen, The design and application of upwind schemes on unstructured meshes, AIAA Paper, 86-0366, 1989.