

A dynamically load-balanced parallel p -adaptive implicit high-order flux reconstruction method for under-resolved turbulence simulation

Lai Wang^a, Matthias K. Gobbert^b, Meilin Yu^{a,*}

^a Department of Mechanical Engineering, University of Maryland, Baltimore County, United States of America

^b Department of Mathematics and Statistics, University of Maryland, Baltimore County, United States of America



ARTICLE INFO

Article history:

Received 8 October 2019

Received in revised form 29 March 2020

Accepted 18 May 2020

Available online 25 May 2020

Keywords:

Dynamic load balancing

Dynamic p -adaptation

Under-resolved turbulence simulation

Implicit high-order flux reconstruction

Matrix-free GMRES

ABSTRACT

We present a dynamically load-balanced parallel p -adaptive implicit high-order flux reconstruction method for under-resolved turbulence simulation. The high-order explicit first stage, singly diagonal implicit Runge–Kutta (ESDIRK) method is employed to circumvent the restriction on the time step size. The pseudo transient continuation is coupled with the matrix-free restarted generalized minimal residual (GMRES) method to solve the non-linear equations at each stage, except the first one, of ESDIRK. We use the spectral decay smoothness indicator as the refinement/coarsening indicator for p -adaptation. A dynamic load balancing technique is developed with the aid of the open-source library ParMETIS. The trivial cost, compared to implicit time stepping, of mesh repartitioning and data redistribution enables us to conduct p -adaptation and load balancing every time step. An isentropic vortex propagation case is employed to study the impact of element weights used in mesh repartitioning on parallel efficiency. We apply the p -adaptive solver for implicit large eddy simulation (ILES) of the transitional flows over a cylinder when Reynolds number (Re) is 3900 and the SD7003 wing when Re is 60000. Numerical experiments demonstrate that a significant reduction in the run time (up to 70%) and total number of solution points (up to 76%) can be achieved with p -adaptation.

© 2020 Elsevier Inc. All rights reserved.

1. Introduction

Recent decades have witnessed tremendous developments in high-order computational fluid dynamics (CFD) methods, such as discontinuous Galerkin methods (DG) [1–8], spectral difference methods (SD) [9–12], classic compact finite difference methods [13,14], finite difference summation by parts (SBP) operators [15–17], and flux reconstruction/correction procedure via reconstruction methods (FR/CPR) [18–24]. Low-dissipation and low-dispersion properties of high-order methods have made them attractive for ILES of turbulent flows. It has been reported in [25] that DG-based ILES can outperform subgrid-model-based LES for transitional flows and wall bounded flows. The dissipation of high-order methods on low-frequency large-scale flow features is trivial, and it is only significant on high wavenumbers/frequencies. Therefore, the truncation error of high-order methods is considered as an implicit subgrid model for turbulence simulation. We note that high-order methods are prone to suffer from instabilities due to aliasing errors in under-resolved turbulence simulation,

* Corresponding author.

E-mail address: mlyu@umbc.edu (M. Yu).

especially when the spatial polynomial degree exceeds two. It has been shown that with proper de-aliasing techniques, ILES using high-order methods has promising capabilities in under-resolved turbulence simulation [7]. Four popular types of stabilization approaches for under-resolved turbulence simulation can be found in the literature, including (a) the split form [26,27], (b) over integration [28,29], (c) artificial viscosity via spectral vanishing viscosity [30–32], and (d) polynomial filtering [33–35]. In this study, we employ nodal polynomial filtering proposed by Fischer and Mullen [35] when de-aliasing is needed.

A uniformly high-order spatial discretization in the entire flow field for under-resolved turbulence simulation of certain problems, such as wall bounded turbulent flows at high Reynolds numbers, is very expensive. As a matter of fact, the high-order spatial discretization is only needed in the near wall region, such as the turbulent boundary layer, and the wake region where vortex shedding dominates. Collocation schemes favor a straightforward implementation of p -adaptation, which has the potential to significantly decrease computational cost. Instead of refining/coarsening the polynomial degrees, mesh-adaptation either relocates the mesh points or locally refines/coarsens the mesh to adapt the mesh resolution. Comparison of h -, p -, and hp -adaptation for high-order methods has been carried out [36,37]. It is found that p -adaptation is more advantageous for smooth flows, and h - and hp -adaptation perform better for flows with discontinuities. In this work, we only consider p -adaptation for smooth flows.

In the literature, three major groups of adaptation methodologies can be found. The first one is the feature-based adaptation [38–40], the second one is the truncation-error-based or discretization-error-based adaptation [41–43], and the last one is the output-based or adjoint-based adaptation [44–47]. Feature-based adaptation methods are usually *ad hoc* and heavily rely on empirical parameters; however, their ease of implementation and reasonable robustness make them a good choice for adaptation. Truncation-error based approaches usually use the correction from either an additional coarser mesh or a lower-order discretization to estimate the local discretization error, which can serve as the adaptation indicator. A comparison of several feature-based and discretization-error-based adaptation indicators is conducted by Naddei et al. [40]. Adjoint-based adaptation methods are popular for engineering purposes since engineers are more interested in output functionals, such as lift and drag. Their superiority over the former two approaches has been demonstrated for steady flow problems. However, the computational cost of adjoint-based adaptation, especially for unsteady turbulence simulation, can be large. Recently, Bassi et al. [48] employed an efficient entropy-adjoint-based [49] p -adaptive DG solver to conduct scale-resolving turbulence simulation. In this study, we will evaluate the performance of a feature-based adaptation method [40], which employs the spectral decay smoothness indicator [50] as the refinement/coarsening indicator when it is applied to under-resolved turbulence simulation.

Explicit high-order Runge–Kutta (RK) methods [51,52] have been widely applied to unsteady flow simulation. p -adaptation will naturally lead to p -enrichment in near wall regions where the elements are usually clustered, thus worsening the Courant–Friedrichs–Lewy (CFL) condition when explicit methods are employed. Implicit time integrators can essentially circumvent the CFL restriction that explicit methods have. Diagonally implicit RK methods [53] and backward differentiation formula (BDF) methods are among the most popular implicit time integration methods. Recently, linearly implicit Rosenbrock methods have become popular for under-resolved turbulence simulation [54,55]. Matrix-based implicit methods are notorious for the large memory consumption. Therefore, the matrix-free implementation [55,56] is usually employed to reduce memory usage [57] for massive turbulence simulation. In our recent comparative study of BDF, Rosenbrock–Wanner (ROW), and ESDIRK [55], we found that multistage implicit RK is more efficient than multistep BDF, and ESDIRK method is generally more robust than ROW in the context of matrix-free implementation with an element-Jacobi preconditioner. Therefore, we employ ESDIRK in this study. Interested readers are also referred to [58,59] for implicit-explicit (IMEX) Runge–Kutta methods, which use explicit and implicit RK for non-stiff and stiff terms, respectively.

Given that the time step size of an implicit time integrator can be relatively large, dynamic adaptivity is desired to track the rapid change of turbulence features. Consequently, the work loads on all processes in parallel simulation will be imbalanced once p -adaptation takes place. The difference of numbers of degrees of freedom on different processes can be over 500% for a simple isentropic vortex propagation problem [60]. Hence, a dynamic load balancing technique is of crucial importance for the parallel efficiency of p -adaptive methods. Existing publications regarding p -adaptive high-order methods for turbulence simulation use the mean flow field to conduct the adaptation without dynamic adaptivity [39,48]. We utilize the open-source library ParMETIS [61] to achieve dynamic load balancing for p -adaptation. Technical details on using ParMETIS for dynamic adaptivity are presented in Section 3 for interested readers. In our numerical experiments, we find that the cost of mesh repartitioning and data redistribution is trivial compared to that of implicit time stepping. This enables p -adaptation to be conducted every time step to significantly decrease the overall unsteady flow simulation time.

Contributions. We develop a dynamically load-balanced p -adaptation technique for parallel implicit high-order flux reconstruction solution of unsteady Navier–Stokes equations. The implementation regarding the dynamic load balancing technique is presented in detail. We discuss the impact of weight calculation for each element on the parallel efficiency of p -adaptation when implicit time integrators are used. The p -adaptive solver is applied to under-resolved turbulence simulation of the transitional flow over an infinite cylinder at $Re = 3900$ and the transitional flow over the SD7003 wing at $Re = 60000$. Compared to the p -uniform spatial discretization, the p -adaptive FR method can save up to 70% computational cost when $p \leq 3$ in our experiments and provide favorable numerical predictions. We expect more savings when p is higher as demonstrated in a simple experiment of the isentropic vortex propagation problem.

Organization. The remainder of the paper is organized as follows. Section 2 provides the mathematical background of the governing equations, the spatial discretization, and the time integration. Section 3 introduces the p -adaptation algorithm

with the spectral decay smoothness indicator, and then explains the parallel mesh partitioning technique with ParMETIS. A simple example is employed to demonstrate the impact of weight calculation of each element on the parallel efficiency. Applications of the p -adaptive solver to under-resolved turbulence simulation are presented in Section 4. In Section 5, we draw conclusions from the current work.

2. Background

2.1. Governing equations

Using Einstein summation convention, the compressible Navier–Stokes equations can be written as

$$\frac{\partial \rho}{\partial t} + \frac{\partial(\rho u_j)}{\partial x_j} = 0, \quad (1)$$

$$\frac{\partial(\rho u_i)}{\partial t} + \frac{\partial(\rho u_j u_i + \delta_{ji} P)}{\partial x_j} = \frac{\partial \tau_{ji}}{\partial x_j}, \quad (2)$$

$$\frac{\partial(\rho E)}{\partial t} + \frac{\partial(\rho u_j H)}{\partial x_j} = \frac{\partial(u_i \tau_{ij} - K_j)}{\partial x_j}, \quad (3)$$

where $i = 1, \dots, n_d$, and n_d is the dimension number. Herein, ρ is the fluid density, u_i is the velocity component, P is the pressure, $E = \frac{P/\rho}{\gamma-1} + \frac{1}{2}u_k u_k$ is the specific total energy, $H = E + \frac{P}{\rho}$ is the specific total enthalpy, τ_{ij} is the viscous stress and K_j is the heat flux. γ is the specific heat ratio defined as $\gamma = c_p/c_v$, where c_p and c_v are specific heat capacity at constant pressure and volume, respectively. In this study, γ is set as 1.4. The ideal gas law $P = \rho R T$ holds, where R is the ideal gas constant and T is the temperature. The viscous stress tensor and heat flux vector are given by

$$\tau_{ij} = 2\mu \left\{ S_{ij} - \frac{1}{3} \frac{\partial u_k}{\partial x_k} \delta_{ij} \right\}, \quad (4)$$

$$K_j = -\frac{\mu c_p}{Pr} \frac{\partial T}{\partial x_j}, \quad (5)$$

where μ is the fluid dynamic viscosity, Pr is the molecular Prandtl number, and the strain-rate tensor S_{ij} is defined as

$$S_{ij} = \frac{1}{2} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right). \quad (6)$$

In this study, μ and Pr are treated as constants.

2.2. The FR/CPR method

For completeness, a brief review of the FR/CPR method [20] is presented in this section. A symbolic form of the compressible Navier–Stokes equations (1), (2), (3) is written as

$$\frac{\partial \mathbf{q}}{\partial t} + \nabla \cdot \mathbf{f} = 0, \quad (7)$$

which is defined in domain Ω . Ω is partitioned into N non-overlapping and conforming elements Ω_e , where $e = 0, 1, \dots, N-1$. After multiplying each side by the test function ϑ and integrate over Ω_e , one obtains

$$\int_{\Omega_e} \frac{\partial \mathbf{q}_e}{\partial t} \vartheta dV + \int_{\Omega_e} \vartheta \nabla \cdot \mathbf{f}_e dV = 0. \quad (8)$$

On applying integration by parts and Gauss divergence theorem, Eq. (8) reads

$$\int_{\Omega_e} \frac{\partial \mathbf{q}_e}{\partial t} \vartheta dV + \int_{\partial \Omega_e} \vartheta \mathbf{f}_e \cdot \mathbf{n} dS - \int_{\Omega_e} \mathbf{f}_e \cdot \nabla \vartheta dV = 0, \quad (9)$$

where \mathbf{n} is the outward-facing unit normal vector of the faces of the element Ω_e . In the discrete form, we assume \mathbf{q}_e^h is the approximate solution in element Ω_e . The solution and the test function belong to the polynomial space of degree k , i.e., $\mathbf{q}_e^h \in p^k$ and $\vartheta^h \in p^k$. To ensure conservation, $\mathbf{f}_e \cdot \mathbf{n}$ in Eq. (9) is replaced with \mathbf{f}_n^{com} , the common flux in the normal direction of the element surfaces. Eq. (9) then reads

$$\int_{\Omega_e} \frac{\partial \mathbf{q}_e^h}{\partial t} \vartheta^h dV + \int_{\partial \Omega_e} \vartheta^h \mathbf{f}_n^{com} dS - \int_{\Omega_k} \mathbf{f}_e^h \cdot \nabla \vartheta^h dV = 0. \quad (10)$$

After applying integration by parts and Gauss divergence theorem again to the last term of Eq. (10), one obtains

$$\int_{\Omega_e} \frac{\partial \mathbf{q}_e^h}{\partial t} \boldsymbol{\vartheta}^h dV + \int_{\Omega_e} \boldsymbol{\vartheta}^h \nabla \cdot \mathbf{f}_e^h dV + \int_{\partial\Omega_e} \boldsymbol{\vartheta}^h [\mathbf{f}] dS = 0, \quad (11)$$

where $[\mathbf{f}] = \mathbf{f}_n^{com} - \mathbf{f}_n^{loc}$ with $\mathbf{f}_n^{loc} = \mathbf{f}_e^h \cdot \mathbf{n}$. In the FR/CPR method, the correction field $\delta_e \in p^k$ is defined as [20]

$$\int_{\partial\Omega_e} \boldsymbol{\vartheta}^h [\mathbf{f}] dS = \int_{\Omega_e} \boldsymbol{\vartheta}^h \delta_e dV. \quad (12)$$

Therefore, Eq. (11) can be expressed as

$$\int_{\Omega_e} \left(\frac{\partial \mathbf{q}_e^h}{\partial t} + \nabla \cdot \mathbf{f}_e^h + \delta_e \right) \boldsymbol{\vartheta}^h dV = 0. \quad (13)$$

The differential form can then be employed as

$$\frac{\partial \mathbf{q}_e^h}{\partial t} + \mathbb{P} \left(\nabla \cdot \mathbf{f}_e^h \right) + \delta_e = 0. \quad (14)$$

Herein, $\mathbb{P} \left(\nabla \cdot \mathbf{f}_e^h \right)$ is the projection of the flux divergence $\left(\nabla \cdot \mathbf{f}_e^h \right)$, which may not be a polynomial, onto an appropriate polynomial space. We note that Eq. (14) can be directly derived from the differential form; their equivalence has been established in [62]. Specifically, for quadrilateral and hexahedral elements, the correction field can be obtained by means of the tensor product of the one dimensional correction polynomials; for triangular and tetrahedral elements, the readers are referred to [20,63]. Only hexahedral elements are considered in this study.

The Roe approximate Riemann solver [64] is used to calculate the common inviscid fluxes at the cell interfaces in their normal directions as

$$\mathbf{f}_{n,inv}^{com} = \frac{\mathbf{f}_{n,inv}^+ + \mathbf{f}_{n,inv}^-}{2} - \mathbf{R} |\Lambda| \mathbf{R}^{-1} \frac{\mathbf{q}^+ - \mathbf{q}^-}{2}, \quad (15)$$

where superscripts ‘-’ and ‘+’ denote the left of right side of the current interface, the subscript \mathbf{n} is the unit normal direction from left to right, Λ is a diagonal matrix consisting of the eigenvalues of the Jacobian $\partial \mathbf{f}_n / \partial \mathbf{q}$, and \mathbf{R} consists of the corresponding right eigenvectors evaluated with the Roe averaged values. The common viscous fluxes at the cell interfaces are $\mathbf{f}_{n,vis}^{com} = \mathbf{f}_{vis}(\mathbf{q}^+, \nabla \mathbf{q}^+, \mathbf{q}^-, \nabla \mathbf{q}^-)$. Here, we need to define the common solution \mathbf{q}^{com} and common gradient $\nabla \mathbf{q}^{com}$ at the cell interface. On simply taking average of the primitive variables, we get

$$\mathbf{q}^{com} = \frac{\mathbf{q}^+ + \mathbf{q}^-}{2}. \quad (16)$$

The common gradient is computed as

$$\nabla \mathbf{q}^{com} = \frac{\nabla \mathbf{q}^+ + \mathbf{r}^+ + \nabla \mathbf{q}^- + \mathbf{r}^-}{2}, \quad (17)$$

where \mathbf{r}^+ and \mathbf{r}^- are the corrections to the gradients on the interface. The second approach of Bassi and Rebay (BR2) [5] is used to calculate the corrections.

2.3. ESDIRK methods with pseudo transient continuation

The ESDIRK methods for the compressible Navier-Stokes equation (7) can be written as

$$\begin{cases} \mathbf{q}^{n+1} = \mathbf{q}^n + \Delta t \sum_{i=1}^s b_i \mathbf{R}(\mathbf{q}^i), \\ \mathbf{q}^i = \mathbf{q}^n, i = 1, \\ \mathbf{q}^i = \Delta t \omega \mathbf{R}(\mathbf{q}^i) + \mathbf{q}^n + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathbf{R}(\mathbf{q}^j), i = 2, \dots, s, \end{cases} \quad (18)$$

where i is the stage number, s is number of total stages, n denotes the physical time step, and $\mathbf{R} = -\nabla \cdot \mathbf{f}$. The second-order, three-stage ESDIRK2 [53], and fourth-order, six-stage ESDIRK4 [65] methods are employed in this paper. A comparative study of different implicit time integration methods can be found in [55]. In every stage except the first one, a nonlinear system, which can be expressed as

$$\mathbf{F}(\mathbf{q}^i) = \left(-\frac{1}{\omega \Delta t} \mathbf{q}^i + \mathbf{R}(\mathbf{q}^i) \right) + \frac{1}{\omega \Delta t} \left(\mathbf{q}^n + \Delta t \sum_{j=1}^{i-1} a_{ij} \mathbf{R}(\mathbf{q}^j) \right), \quad i = 2, \dots, s, \tag{19}$$

needs to be solved for \mathbf{q}^i .

The pseudo transient continuation for the i -th stage reads

$$\frac{\mathbf{q}^{m+1,i} - \mathbf{q}^{m,i}}{\Delta \tau} = \mathbf{F}(\mathbf{q}^{m+1,i}), \tag{20}$$

where m is the iteration step. Eq. (20) can be linearized as

$$\left(\frac{1}{\omega \Delta t} + \frac{1}{\Delta \tau} - \frac{\partial \mathbf{R}}{\partial \mathbf{q}} \right)^m \Delta \mathbf{q}^{m,i} = \mathbf{F}(\mathbf{q}^{m,i}), \tag{21}$$

where $\Delta \mathbf{q}^{m,i} = \mathbf{q}^{m+1,i} - \mathbf{q}^{m,i}$. We employ the successive evolution relaxation (SER) algorithm [66] to update the pseudo time step size as

$$\Delta \tau^0 = \Delta \tau_{init}, \text{ and } \Delta \tau^{m+1} = \min \left(\Delta \tau^m \frac{\|\mathbf{F}\|_{L_2}^{m-1}}{\|\mathbf{F}\|_{L_2}^m}, \Delta \tau_{max} \right). \tag{22}$$

In all the numerical experiments conducted in this study, we set the convergence tolerance of relative residual of the pseudo transient continuation as $tol_{rel}^{pseudo} = 10^{-4}$.

We use the restarted GMRES framework in the portable, extensible toolkit for scientific computation (PETSc) package [67] with user-defined functions to conduct the matrix-vector product and preconditioning. In Krylov subspace methods, the Jacobian matrix only appears in the matrix-vector product. A finite difference approximation of the matrix-vector product reads

$$\left(\frac{\partial \mathbf{R}}{\partial \mathbf{q}} \right) \mathbf{X} = \frac{\mathbf{R}(\mathbf{q} + \varepsilon \mathbf{X}) - \mathbf{R}(\mathbf{q})}{\varepsilon} + O(\varepsilon), \tag{23}$$

where $\varepsilon = 10^{-6}$ in this study. The element-Jacobi preconditioner, i.e., the inverse of the diagonal blocks of $\left(\frac{1}{\omega \Delta t} + \frac{1}{\Delta \tau_{max}} - \frac{\partial \mathbf{R}}{\partial \mathbf{q}} \right)$, is used for left preconditioning in this study. The preconditioner is only evaluated once at the starting stage of each physical time stepping.

The pseudo transient continuation is an inexact Newton's method. Therefore, we assign a relatively large tolerance to the GMRES solver, i.e., $tol_{rel}^{gmres} = 10^{-1}$, to save computational cost [55]. However, for stiff problems, $tol_{rel}^{gmres} = 10^{-1}$ may lead to divergence of the pseudo transient continuation [68]. In this case, we will decrease tol_{rel}^{gmres} to 10^{-2} . If not specifically mentioned, the restart number is 60 for all numerical experiments. We note that the performance of the element-Jacobi preconditioner will quickly deteriorate as $\Delta \tau$ increases to large values $\Delta \tau \gg \Delta t$. Therefore, in the pseudo transient continuation, we do not increase the pseudo time step $\Delta \tau$ to large values to ensure that the relative tolerance of GMRES tol_{rel}^{gmres} can always be met within 100 iterations. Otherwise, we will decrease the current pseudo time step size by half and redo the current pseudo-time iteration. In this study, we set $\Delta \tau_{max}$ as $\Delta \tau_{max} = O(\Delta t)$ and $\Delta \tau_{min}$ is usually one magnitude smaller than Δt . We have developed a p -multigrid solver for coarsely-resolving simulation of low-Mach-number turbulent flows in [69]. Applying the p -multigrid solver as a preconditioner for Newton-Krylov methods will be our future work.

3. Dynamically load-balanced p -adaptation for high performance computing

3.1. p -adaptation using spectral decay smoothness indicator

The spectral decay smoothness indicator has been successful used to detect trouble cells for shock-capturing [50]. It is defined as

$$\eta_k = \frac{\|s_p - s_{p-1}\|_{L^2}}{\|s_p\|_{L^2}} \tag{24}$$

for one element. For a hexahedral element, $\|\cdot\|_{L^2}$ is defined as

$$\|\cdot\|_{L^2} = \frac{\sum_{\xi=1}^{p+1} \sum_{\eta=1}^{p+1} \sum_{\zeta=1}^{p+1} [(\cdot)^2 \omega_\xi \omega_\eta \omega_\zeta |J|_{\xi,\eta,\zeta}]}{\sum_{\xi=1}^{p+1} \sum_{\eta=1}^{p+1} \sum_{\zeta=1}^{p+1} [\omega_\xi \omega_\eta \omega_\zeta |J|_{\xi,\eta,\zeta}]}, \tag{25}$$

where \mathbf{J} is the Jacobian matrix of the coordinate transformation from a physical element to the standard element, $|J|$ is the determinant of \mathbf{J} , and $\omega_\xi/\eta/\zeta$ are the quadrature weights in the $\xi/\eta/\zeta$ directions, respectively. s_{p-1} is obtained by projecting the solution from the degree p polynomial space to the degree $p - 1$ space. With the spectral decay smoothness indicator, the adaptation procedure can be achieved as follows:

- calculate the smoothness indicator of every element;
- adjust the polynomial degree of every element according to the adaptation criteria;
- limit the difference of polynomial degrees at non-conforming interfaces to one; and
- project or prolong the solutions when the polynomial is decreased or increased, respectively.

The adaptation criteria we employ in the present study are organized as follows:

- increasing the polynomial degree by one when $\eta_k > \nu_{max}\eta_{k,max}$; and
- decreasing the polynomial degree by one when $\eta_k < \nu_{min}\eta_{k,max}$.

Herein, ν_{max} and ν_{min} are problem dependent. The polynomial degree p of an element in the flow field is $p \in [p_{min}, p_{max}]$, where p_{min} and p_{max} are the minimum and maximum polynomial degree, respectively. If not specifically mentioned, we choose momentum in the x direction, i.e., ρu , as the variable for smoothness indicator calculation; $\nu_{max} = 0.1$ and $\nu_{min} = 0.001$. In this study, all adaptive solvers will have $p_{min} = 1$. When the description adaptive p^k FR or p^k FR with p -adaptation is used, we are referring to the adaptive FR method with $p_{min} = 1$ and $p_{max} = k$. Some preliminary results using the current p -adaptation method to solve 2D unsteady Navier–Stokes equations have been presented in [60]. Since no dynamic load balancing was employed there, the differences of the numbers of degrees of freedom on different processors in parallel simulation can be over 500% for a simple isentropic vortex propagation problem. In this study, we propose to develop a dynamic load balancing strategy for parallel simulation with p -adaptation.

3.2. Implementation of parallel mesh partitioning

To achieve dynamic load balancing, `ParMetis_V3_AdaptiveRepart()` in the open source library ParMETIS [61] is employed for efficient parallel mesh partitioning. This application programming interface (API) is particularly developed to repartition locally adapted mesh in parallel computing. It allows one to use n_{proc} processes to partition the mesh into n_{part} parts. In this work, $n_{proc} = n_{part}$ is used to assure the load of repartitioning is balanced among all processes. A distributed mesh is required as one of the inputs of `ParMetis_V3_AdaptiveRepart()`. We employ `METIS_PartMeshDual()` in serial METIS [61] to partition the mesh to obtain the initial distributed mesh and no weights are assigned to any elements. On using `ParMetis_V3_AdaptiveRepart()`, each element of the unstructured mesh is regarded as a vertex in the graph. An illustration of the parallel mesh partitioning is presented in Fig. 1 to explain the technical details. Following C++ convention, all indices start from 0. There are n_{proc} processes and each process possesses one mesh partition. Assume the i -th process has n_i^e elements. The global index of the j -th local element in the i -th process must be $\text{Index}(j, i) = \sum_{m=0}^{i-1} n_m^e + j$ to ensure that the distributed mesh is a legal input of `ParMetis_V3_AdaptiveRepart()`. One output of `ParMetis_V3_AdaptiveRepart()` is an array of size n_i^e which stores the process indices of the local elements after parallel mesh partitioning. As shown in the second row in Fig. 1, before the parallel mesh partitioning, each process has four elements. Process 0 has Elements 0 to 3, Process 1 has Elements 4 to 7, etc. After parallel mesh partitioning, the index of the process that an element belongs to is stored locally. Due to change of element weights resulting from p -adaptation, elements could appear to be ‘randomly’ distributed to all processes. In other words, some processes will possess a part of the elements that they have before partitioning and some will obtain all elements from other processes. As shown in the fourth row in Fig. 1, Process 0 has four elements and two of them, Elements 4 and 5, are obtained from Process 1. Process 3 needs to fetch Elements 2, 3 from Process 0, Elements 6, 7 from Process 1, and Element 11 from Process 2. To make the new distributed mesh as a legal input for `ParMetis_V3_AdaptiveRepart()` in the next parallel mesh partitioning, one needs to reorganize the global element indices as illustrated in the last row of Fig. 1. Corresponding CFD data should also be reorganized following the mapping between the old and new global element indices.

For data redistribution, we use a collect-and-distribute strategy. We utilize `MPI_Allgather()` to gather all the conservative variables on all the processes and each process will fetch the corresponding working variables from the collected data pool. The aforementioned randomness of redistributing elements to all processes leads to the fact that when it comes to data redistribution, the elements on many of the processes could be totally different from those before the parallel mesh partitioning. This implies that the cost of a process-to-process communication strategy to exchange CFD data could possibly be close to that of the collect-and-distribute strategy. We refer interested readers to several advanced scalable distribution approaches developed in [70,71]. In our numerical experiments, we have found that the total amount of run time needed for mesh partitioning and data redistribution is trivial when compared to that of implicit time stepping. We would like to clarify that in the spatial and temporal solvers, communication among all processes is done in a process-to-process manner to maximize efficiency.

An important input of `ParMetis_V3_AdaptiveRepart()` is the weight of each element in the distributed mesh. For the FR/CPR discretization, the number of solution points n_{sp} within a element is $(p + 1)^3$ for a hexahedral element. And there are five equations at each solution point to be solved in three dimensional problems. Hence, the number of degrees of freedom in one element is $n_{dof} = 5n_{sp}$. For a hexahedral element, all operations of the FR/CPR methods are indeed conducted dimension by dimension. Therefore, we roughly estimate the computational complexity of one-time residual evaluation as $O(n_{dof} \cdot (p + 1))$. When implicit time integrators are employed, the cost of one-time residual evaluation is trivial compared to three major parts, (a) evaluating the element Jacobian matrix, (b) calculating the element-Jacobi preconditioner

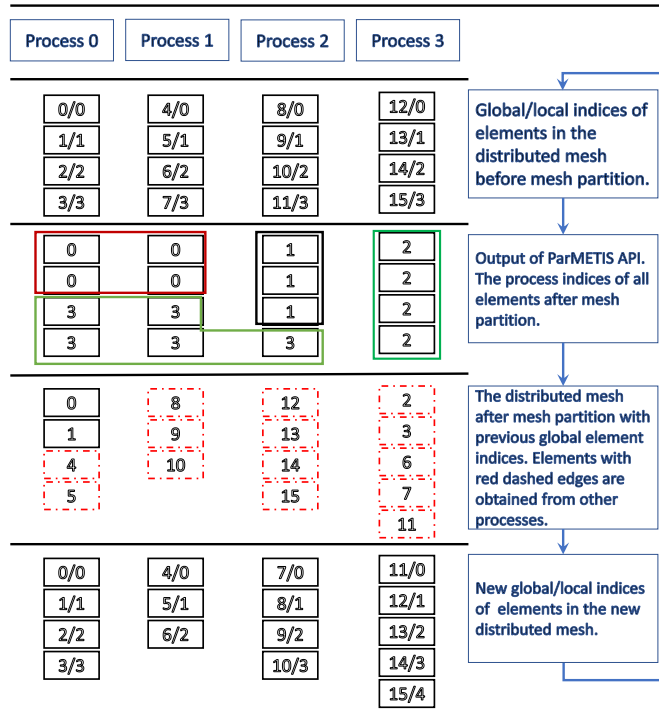


Fig. 1. Illustration of the parallel mesh partitioning using the ParMETIS API.

and (c) solving the nonlinear/linear equations using Newton/Krylov methods. When evaluating the element Jacobian matrix, the finite difference approach is used. For each element, there will be n_{dof} times the residual evaluation and the computational complexity is $O(n_{dof}^2 \cdot (p + 1))$. We use lower-upper (LU) decomposition to invert the element Jacobian matrix to obtain the element-Jacobi preconditioner. The computational complexity of is $O(n_{dof}^3)$. In the matrix-free implementation of the GMRES solver, the approximation of matrix-vector product and preconditioning will contribute to the computational cost dominantly. For the matrix-vector product approximation, the computational cost will be that of one-time residual evaluation. The left preconditioning is used in our approach and the complexity is $O(n_{dof}^2)$. Overall, the complexity of the Newton-Krylov solver is $G(O(n_{dof}^2) + O(n_{dof} \cdot (p + 1)))$, where G is the total number of GMRES iterations in the pseudo transient continuation. G is highly problem dependent. Thus, three candidates to calculate the weight of each element, namely, $\omega_e = n_{dof}/5 = n_{sp}$, $\omega_e = n_{sp}^2$, and $\omega_e = n_{sp}^3$, will be investigated.

Though we can sketch the computational complexity within each element to pursue an optimal candidate for weight calculation, the parallel performance is only directly related to the output distributed mesh of ParMETIS. In our numerical experiments, we observe that when the disparity of element weights is excessively large, e.g., $w_e = n_{sp}^3$, $p_{min} = 2$, and $p_{max} = 5$, the output mesh will lead to degraded parallel efficiency and the large disparity will occasionally lead to failure of ParMETIS, even when p is smaller than 4 and $w_e = n_{sp}$ is used. Specifically, there will be processes which have no elements after parallel mesh partitioning. Therefore, when ParMETIS failure is encountered, we will decrease the weight of each element to $w_e = p + 1$ and redo the mesh partitioning for the current time step. Note that the computational cost of parallel mesh partitioning is trivial compared to implicit time stepping. In the following subsection, a simple example is employed to demonstrate the proposed dynamic load balancing strategy.

3.3. A simple example of dynamic load balancing

We simulate the 2D isentropic vortex propagation on a 3D mesh (obtained by extruding a 2D mesh in the z direction for two layers) to demonstrate the dynamic load balancing strategy. The free stream condition is $(\rho, u, v, w, Ma)^T = (1, 1, 1, 0, 0.5)^T$. The fluctuation is defined as [54]

$$\begin{cases} \delta u = -\frac{\alpha}{2\pi} (y - y_0) e^{\phi(1-r^2)}, \\ \delta v = \frac{\alpha}{2\pi} (x - x_0) e^{\phi(1-r^2)}, \\ \delta w = 0, \\ \delta T = -\frac{\alpha^2(\gamma-1)}{16\phi\gamma\pi^2} e^{2\phi(1-r^2)}, \end{cases} \quad (26)$$

where $\phi = \frac{1}{2}$ and $\alpha = 5$ are parameters that define the vortex strength. $r = (x - x_0)^2 + (y - y_0)^2$ is the distance from any point (x, y, z) to the center of line the vortex $(x_0, y_0, z) = (0, 0, z)$ at $t = 0$. The domain is within $[-10, 10] \times [-10, 10] \times [0, 0.8]$. A uniform mesh of $50 \times 50 \times 2$ elements is used for the numerical experiments. Periodic boundary conditions are imposed on all boundaries. We only simulate this problem for 60 steps with a time step size $\Delta t = 0.05$. In the pseudo transient continuation, $\Delta \tau_{init} = 0.05$ and $\Delta \tau_{max} = 10$ are used for SER. For the adaptive solver, the flow field is initialized using uniform p_{max} discretization.

72 processes are employed in this section. Herein, we take p^5 FR with p -adaptation as an example to show how the elements will be distributed to all processes. Elements in the first seven processes and the corresponding order-of-accuracy distribution, i.e., $(p + 1)$ distribution, of four consecutive time steps are presented in Fig. 2. From $t = 0.15$ to $t = 0.20$, the local polynomial degree at non-critical region will be coarsened to $p_{min} = 1$ while a circular region surrounding the vortex will maintain high polynomial degrees. Due to this coarsening, the redistributed mesh change drastically. From $t = 0.20$ to $t = 0.25$, all the first seven processes except the third one do not exchange any elements with other processes. From $t = 0.25$ to $t = 0.3$, even though the order-of-accuracy distribution only changes slightly, almost all the seven processes will send a large portion of local elements to other processes and obtain a significant amount of elements from other processes. For turbulence simulation, we anticipate that the change of the distributed mesh will be more dramatic than this simple problem.

In Fig. 3, numerical results from p/p_{max} -refinement studies for the FR solver with/without adaptation are presented. From Fig. 3(a), it is observed that both the p -uniform solver and p -adaptive solver have spectral convergence. The errors of the p -adaptive and p -uniform solvers are of the same magnitude when the highest polynomial degrees are the same. Run time of different solvers and the run time reduction of the p -adaptive solver using different weight algorithms (with respect to the p -uniform solver) are illustrated in Fig. 3(b) and Fig. 3(c), respectively. Overall, with proper weights assigned to all the elements, around 80% run time reduction can be achieved via p -adaptation when $p_{max} \geq 3$ (highest order of accuracy is no smaller than 4). For $\omega_e = n_{sp}$, when $p_{max} > 3$, the parallel efficiency will keep decreasing as p_{max} increases. This is due to the fact the computational cost of the Jacobian matrix and preconditioner evaluation will grow at much larger rates than that of the one-time residual evaluation. When $p_{max} = 5$, $\omega_e = n_{sp}^2$ and $\omega_e = n_{sp}^2(p + 1)$ have better performance than $\omega_e = n_{sp}$. However, it is shown that $\omega_e = n_{sp}^3$ generally degrades the efficiency than other candidates. Especially, when $\omega_e = n_{sp}^3$ and $p_{max} = 5$, the p -adaptive solver fails to finish the simulation within the time that the p -uniform solver needs. In order to achieve optimal performance, both the weight calculation and the parallel mesh partitioning algorithm should be taken into account. In the following section, we use $\omega_e = n_{sp}$ for the p -adaptive solver when it is applied to under-resolved turbulence simulation since we only consider $p_{max} \leq 3$ there.

In Fig. 4, we present the distribution of order, smoothness indicator in the log scale, and point-wise error of ρu defined as $Err(\rho u) = (\rho u)_{exact} - (\rho u)$. Large η resides at the vortex region and the transition region of high-degree polynomials to low-degree polynomials. Note that this is consistent with the error distribution. It indicates that the smoothness indicator can be used to effectively adjust the polynomial degrees in the flow field to maintain simulation accuracy, thus enhancing the efficiency of p -adaptive solvers.

4. Applications to under-resolved turbulence simulation

The numerical studies in this section use seven computational nodes in a distributed-memory cluster. Each node has two 18-core Intel Xeon Gold 6140 Skylake CPUs (2.3 GHz clock speed, 24.75 MB L3 cache) and 384 GB memory (12 \times 32 GB DDR4). The nodes are connected by a network of four 36-port EDR (Enhanced Data Rate) InfiniBand switches (100 Gb/s bandwidth, 90 ns latency). 252 processes are used for every simulation. The CFD codes and third party libraries, such as ParMETIS and PETSc, are compiled using MPICH 3.2.1 and GCC 7.3.0 compilers. The C++11 standard is used in the compilation of the CFD codes.

4.1. Under-resolved simulation of the flow over an infinite cylinder

Long time simulations of the transitional flow over an infinite cylinder are conducted to validate the reliability of the p -adaptive solver in this section. The diameter of the cylinder is $d = 1$. The inflow conditions are set as the vector $(\rho_\infty, u_\infty, v_\infty, w_\infty, Ma_\infty)^T = (1, 1, 0, 0, 0.1)^T$. The Reynolds number of the inflow with respect to the diameter of the cylinder is $Re_d = \rho_\infty u_\infty d / \mu = 3900$. The Prandtl number is $Pr = 0.71$. A 2D view of the mesh is illustrated in Fig. 5. The center of the cylinder sits at the origin. The 3D mesh is obtained by extruding the 2D mesh along the z direction, i.e., $(0, 0, 1)^T$, for eight layers and the thickness of each layer is $0.25d$. There are 17694 hexahedral elements and the curved wall boundary is represented by p^3 elements. p^2 and p^3 FR are employed with/without p -adaptation. The cylinder surface is treated as a no-slip adiabatic wall. Farfield boundary conditions are applied to outer boundaries. Periodic boundary conditions are imposed at the front and back sides. We employ ESDIRK4 for time integration and $\Delta t = 0.025$. In the pseudo transient continuation, $\Delta \tau_{init} = 0.001$ and $\Delta \tau_{max} = 0.01$ are used for SER. The tolerance for the pseudo transient continuation is $tol_{rel}^{pseudo} = 10^{-4}$ and that of the GMRES solver is $tol_{rel}^{gmres} = 10^{-1}$. We run all simulations until $t_{end} = 800$. The instantaneous solutions in $t \in (100, 800]$ are used for time averaging. For p^2 FR with adaptation, $(v_{max}, v_{min}) = (0.1, 0.001)$ and $(v_{max}, v_{min}) = (0.1, 0.01)$ are tested. For p^3 FR, only $(v_{max}, v_{min}) = (0.1, 0.01)$ is used to carry out the p -adaptation. The flow field is initialized uniformly with the inflow conditions and the p -adaptive solver starts from a uniform p^1 discretization.

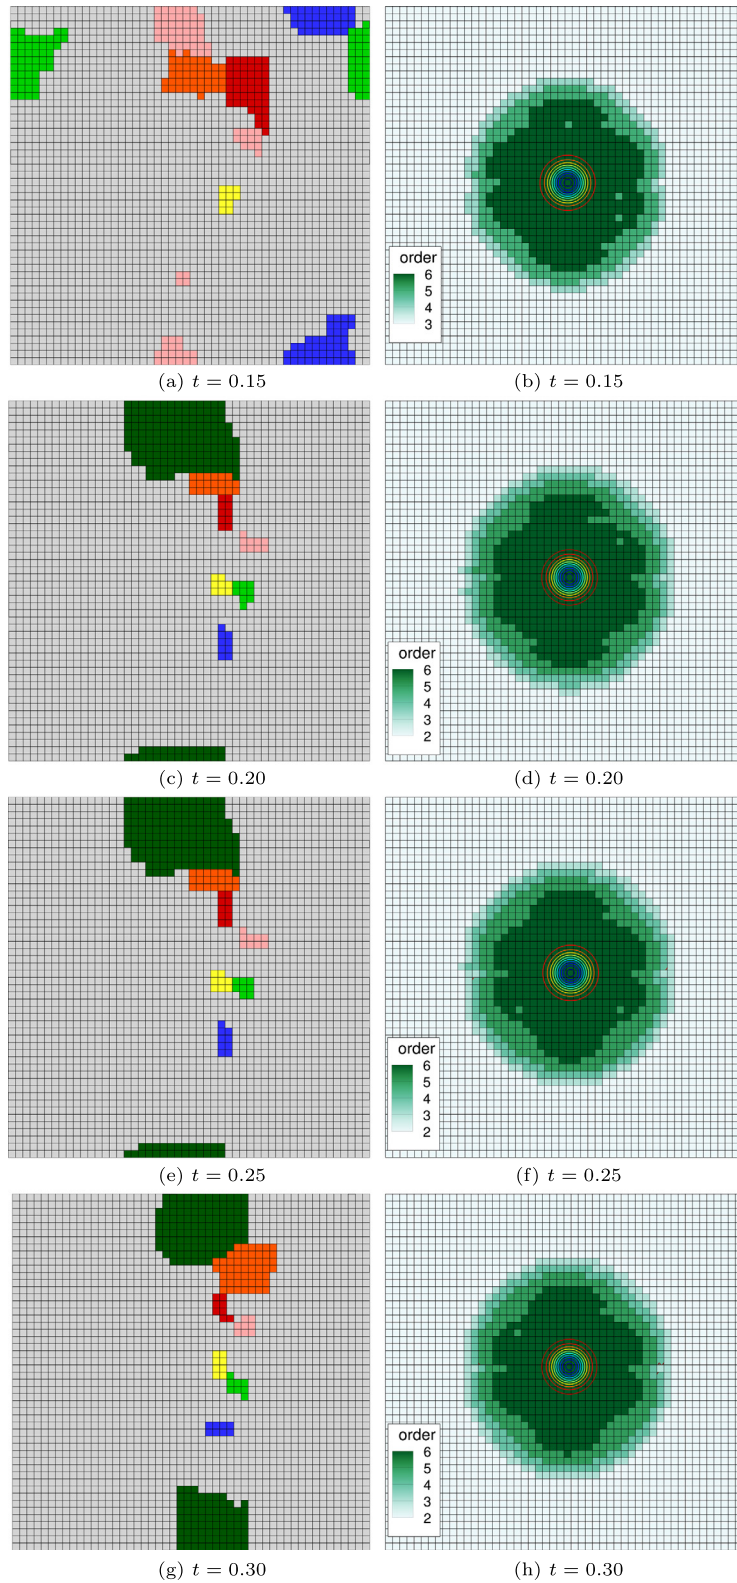


Fig. 2. (a), (c), (e), and (g) are elements in Processes 0–6 at four consecutive time steps. Processes 0–6 are colored by red, yellow, green, blue, orange, pink, and dark green, respectively. (b), (d), (f), and (h) are corresponding instantaneous order-of-accuracy distributions. (For interpretation of the colors in the figure(s), the reader is referred to the web version of this article.)

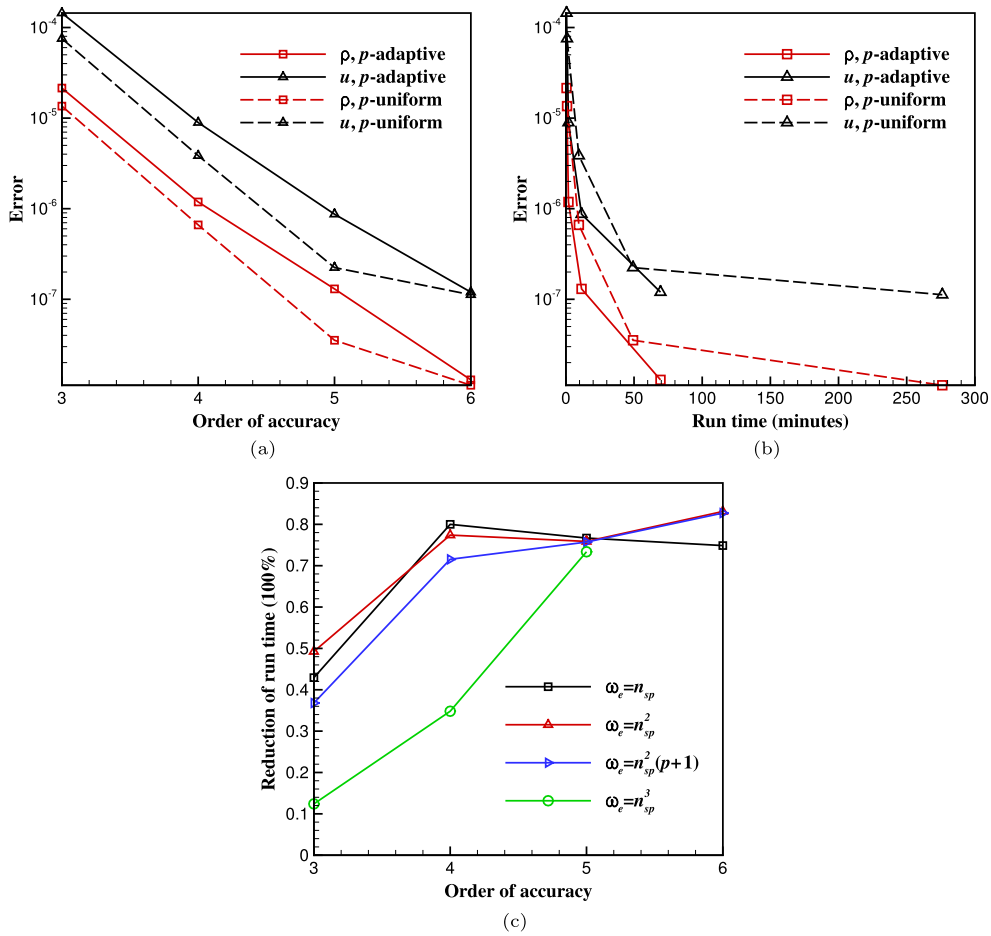


Fig. 3. p/p_{max} -refinement of the p -uniform and p -adaptive FR solvers for the isentropic vortex propagation. (a) Error vs. order of accuracy, (b) error vs. run time, and (c) reduction of run time vs. order of accuracy. For the p -adaptive solver, the order of accuracy indicates the highest order of accuracy, i.e., $p_{max} + 1$ in the flow field.

Table 1
Run time of all simulations of the transitional flow over the infinite cylinder.

Method	(v_{max}, v_{min})	Run time (hours)	Reduction of run time	Reduction of n_{sp}^{tot}
p^2 FR	no adaptation	27.43	0	0
	(0.1, 0.001)	22.43	18.23%	49.96% at $t = 800$
	(0.1, 0.01)	16.55	39.66%	63.43% at $t = 800$
p^3 FR	no adaptation	149.78	0	0
	(0.1, 0.01)	45.56	69.58%	75.98% at $t = 800$

The run time of all simulations is documented in Table 1. Overall, the p -adaptive solver can reduce a significant amount of run time. Particularly, for p^3 FR, the adaptive solver can reduce the run time by 69.58% when $(v_{max}, v_{min}) = (0.1, 0.01)$. At $t = t_{end}$, the adaptive p^2 FR solver has 239,073 solution points when $(v_{max}, v_{min}) = (0.1, 0.001)$ and 17,4701 solution points when $(v_{max}, v_{min}) = (0.1, 0.01)$; the adaptive p^3 FR solver has 271,958 solution points. When turbulence is fully developed, the total number of p -refined elements will be similar at different time steps. In general, the reduction of run time and that of n_{sp}^{tot} are consistent with each other and p -adaptation with larger p_{max} is encouraged as shown in Table 1. From the isosurfaces of Q -criterion, where $Q = 0.5$, illustrated in Fig. 6, it is intuitive that the p -adaptive solver is more dissipative than p -uniform solver in the wake region away from the cylinder. The order-of-accuracy distributions of the adaptive solver with different adaptation parameters at slice $z = 0$ are also presented. When v_{min} is decreased from 0.01 to 0.001, the p^2 region will substantially extend into the wake region away from the cylinder. Thus, the reduction in run time will decrease. The order-of-accuracy distributions of the adaptive p^3 FR solver at different slices are shown in Fig. 7 to give a better presentation of the local p -adaptation.

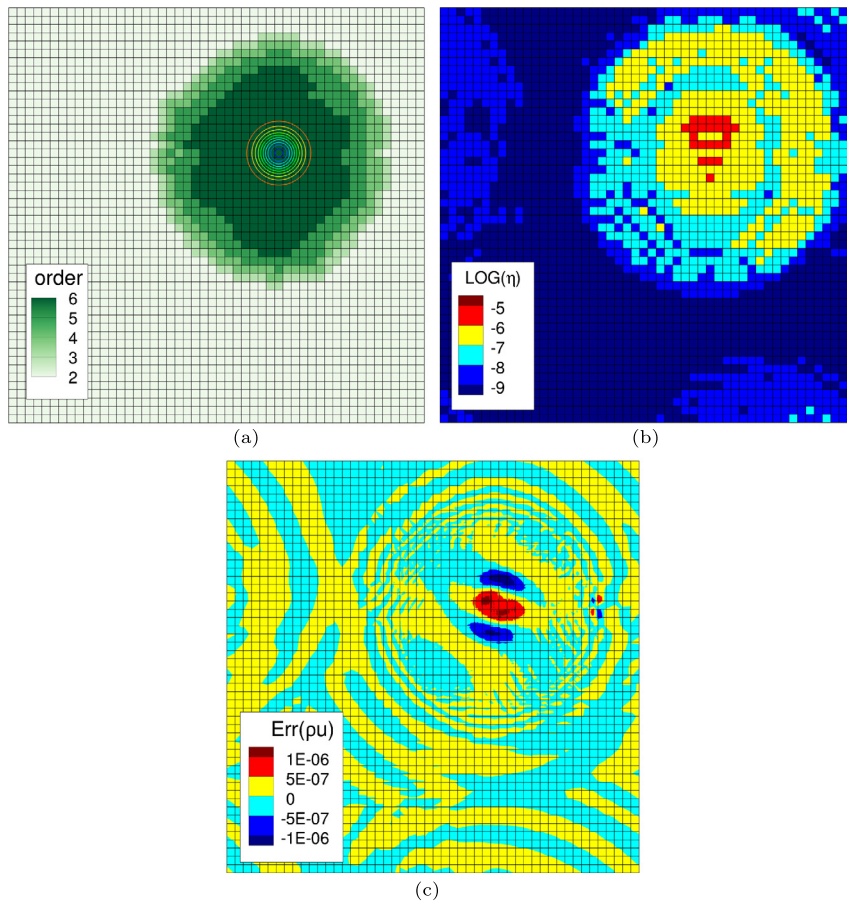


Fig. 4. The distribution of (a) order of accuracy, (b) smoothness indicator η , and (c) point-wise error of ρ_u at $t = 3$ using the adaptive p^5 FR method. Note that the smaller the η is, the smoother the flow is.




Fig. 5. 2D views of the unstructured mesh around a circular cylinder.

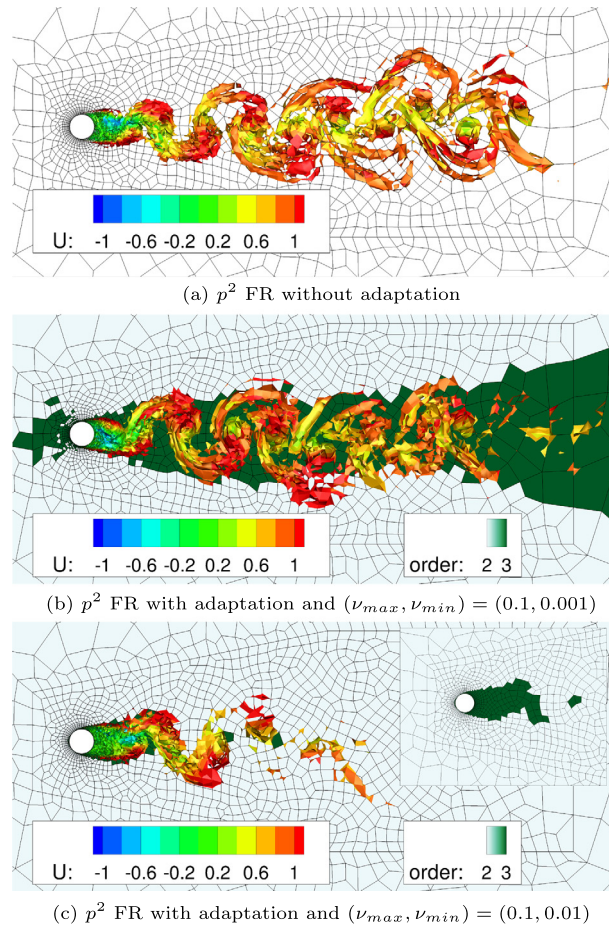


Fig. 6. Transitional flow over the infinite cylinder at $Re = 3900$. Instantaneous isosurfaces of $Q = 0.5$ colored by velocity component in the x direction at $t = 800$. Order-of-accuracy distribution at slice $z = 0$ is turned on in (b) and (c). A close-up view of the near wall region is also presented in (c).

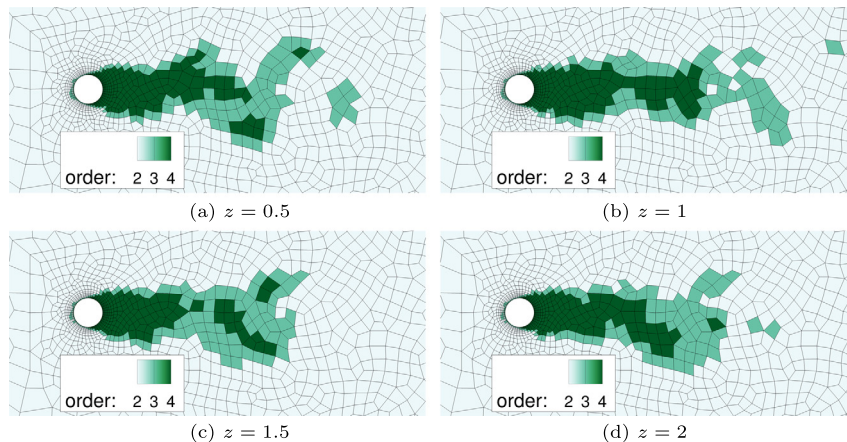


Fig. 7. Instantaneous order-of-accuracy distribution of p^3 FR with adaptation at different slices when $t = 800$.

We further examine the power spectral density (PSD) of the total velocity at four locations in the wake region, namely $(0.58, 0, 1)^T$, $(1.54, 0, 1)^T$, $(6, 0, 1)^T$, and $(10, 0, 1)^T$ as presented in Fig. 8. Compared to the DNS results in [72], a large portion of the inertial range can be resolved at the first two points. At the last two points, $(6, 0, 1)^T$, and $(10, 0, 1)^T$, the adaptive solver gets more dissipative as the parameter ν_{min} increases. The velocity profiles at different positions on the x -axis in Fig. 9, where $y/d \in [-3, 3]$, further demonstrates this observation. At $x = 0.58$ and $x = 1.54$, the velocity profiles




Fig. 8. Power spectral density of the total velocity at different locations in the wake region. (a)–(d) adaptive p^2 FR without adaptation, (e)–(h) adaptive p^2 FR with $(v_{max}, v_{min}) = (0.1, 0.001)$, (i)–(l) adaptive p^2 FR with $(v_{max}, v_{min}) = (0.1, 0.01)$, (m)–(p) p^3 FR without adaptation, and (q)–(t) adaptive p^3 FR with $(v_{max}, v_{min}) = (0.1, 0.01)$. A line of slope $-5/3$ is added to every graph as a reference.

of the p -adaptive FR methods are close to those of the p -uniform FR methods. At $x = 6$ and $x = 10$, the local extrema of the p -adaptive FR are largely dissipated, even when $(v_{min}, v_{max}) = (0.1, 0.001)$.

The surface pressure coefficient C_p and surface friction coefficient C_f on the $y > 0$ side of the cylinder are presented in Fig. 10. When $(v_{max}, v_{min}) = (0.1, 0.001)$, the results of the adaptive p^2 FR method is close to that of the p -uniform p^2 FR method. When $(v_{max}, v_{min}) = (0.1, 0.01)$, the results of adaptive solver are still comparable to those of the p -uniform solver even though a large portion of the wake region uses p^1 polynomials only.




Fig. 9. Profiles of velocity component u in x -direction at different locations. Legend 1, p -uniform p^2 FR; Legend 2, p -adaptive p^2 FR with $(v_{max}, v_{min}) = (0.1, 0.001)$; Legend 3, p -adaptive p^2 FR with $(v_{max}, v_{min}) = (0.1, 0.01)$; Legend 4, p -uniform p^3 FR; Legend 5, p -adaptive p^3 FR with $(v_{max}, v_{min}) = (0.1, 0.01)$.

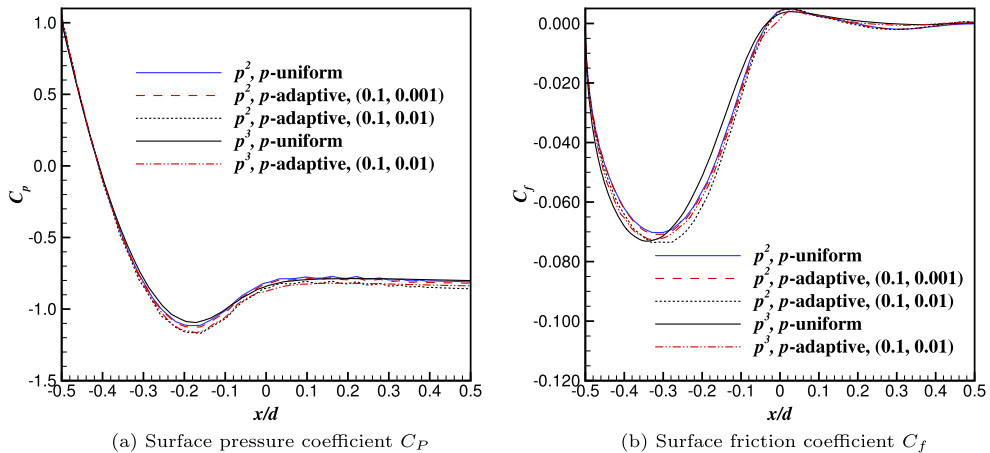


Fig. 10. Surface C_p and C_f in the averaged field of the transitional flow over the infinite cylinder at $Re = 3900$.




Fig. 11. 2D views of the unstructured mesh around the SD7003 wing.

4.2. Under-resolved simulation of the transitional flow over the SD7003 wing

In this section, we apply the p -adaptive solver to the simulation of the transitional flow over the SD7003 wing. The geometry of the wing is obtained from the 1st International Workshop on High-Order CFD Methods. The chord length of the wing is $c = 1$ with the sharp trailing edge rounded by an arc of radius $r \approx 0.0004$. The inflow conditions are $(\rho_\infty, u_\infty, v_\infty, w_\infty, \text{Ma}_\infty)^T = (1, 1, 0, 0, 0.1)^T$. The angle of attack of the inflow is 8° . The Reynolds number of the inflow with respect to the chord length of the wing is $\text{Re}_c = \rho_\infty u_\infty c / \mu = 60000$. The Prandtl number is $\text{Pr} = 0.72$. 2D views of the unstructured mesh used for under-resolved simulation are illustrated in Fig. 11. The height of the first layer of elements in the normal direction of the wing is $0.0003c$. We extrude the 2D mesh along the z direction to obtain the 3D mesh. The first 3D mesh has 20 layers in the z direction and 109,540 hexahedral elements in total. The thickness of each layer is $0.01c$. The second one has 10 layers in the spanwise direction and each layer has a thickness of $0.02c$. The p -adaptation parameters are set as $(v_{\max}, v_{\min}) = (0.1, 0.001)$.

The time step $\Delta t = 0.002$ and tolerance for the pseudo transient continuation $\text{tol}_{\text{rel}}^{\text{pseudo}} = 10^{-4}$ are employed for all numerical experiments (except when computational cost profiling was conducted). For the simulation on the 20-layer mesh, ESDIRK2 is employed for time integration with $\text{tol}_{\text{rel}}^{\text{gmres}} = 10^{-1}$ for the GMRES solver. In the pseudo transient continuation, $\Delta\tau_{\min} = 0.0002$ and $\Delta\tau_{\max} = 0.004$ are used for SER. We first run the simulation with adaptive p^2 FR until $t_1 = 26$. The instantaneous values of conservative variables in $t \in (20, 26]$ are used for averaging. Then we increase the p_{\max} to $p_{\max} = 3$ and resume the simulation until $t_2 = 42$. Averaging is done for $t \in (36, 42]$. When p^3 FR with p -adaptation is used to simulate this problem on the 10-layer mesh, aliasing errors will lead to failure. Therefore, we employ a simple nodal polynomial filtering method proposed by Fischer and Mullen [35] for every element whose polynomial degree exceeds two. The p^2 polynomial is employed as a basis to perform a cut-off as $\tilde{\mathbf{q}}_{p^3} = (1 - \alpha)\mathbf{q}_{p^3} + \alpha\mathbb{P}_{p^2}^{\mathbb{P}_{p^3}}\mathbf{q}_{p^3}$, where $\mathbb{P}_{p^2}^{\mathbb{P}_{p^3}}$ is the projection operation from p^3 to p^2 and $\alpha = 0.2$. With this nodal polynomial filtering, a small amount of dissipation is introduced to stabilize the numerical methods. ESDIRK4 serves as the time integrator for the simulation on the 10-layer mesh with adaptive p^3 FR. To increase the robustness of the pseudo transient continuation, $\text{tol}_{\text{rel}}^{\text{gmres}} = 10^{-2}$, $\Delta\tau_{\min} = 0.0002$, and $\Delta\tau_{\max} = 0.002$ are used. We run the simulation until $t_1 = 26$ only. Solutions in $t \in (20, 26]$ are averaged for statistics.

Two snapshots of the instantaneous isosurfaces of the Q -criterion, where $Q = 500$, are presented in Fig. 12 for the simulation conducted on the 20-layer mesh. One visible difference is that more finer structures are resolved using adaptive p^3 FR. From Fig. 13 and Fig. 14, we observe that polynomials of degree $p > 1$ are clustered in regions near the stagnation point, turbulent boundary layers as well as the wake region. The order-of-accuracy distributions at different slices in the spanwise direction are not exactly the same since the p -adaptation is conducted locally. Overall, the feature-based adaptation method can give an feature-tracking p -distribution. Unlike the transitional flow over the cylinder, the choice of $(v_{\max}, v_{\min}) = (0.1, 0.001)$ actually clusters all the polynomials of degree $p > 1$ in a small domain and only a few high-order elements can be found in the wake region far away from the wing. The power spectral density of the total velocity at four locations close to the suction side of the wing, i.e., $(0.3, 0.057, 0.1)^T$, $(0.5, 0.048, 0.1)^T$, $(0.7, 0.032, 0.1)^T$, and $(0.9, 0.012, 0.1)^T$ are illustrated in Fig. 15. The first point is near the end of the separation bubble, where the transition from laminar flow to turbulent flow takes place. The slopes of PSDs are generally steeper than $-5/3$ at high frequencies at the first point in all simulation. At the other three points, PSDs align with the reference lines in a certain range of high frequencies, which agrees with the results in [73].

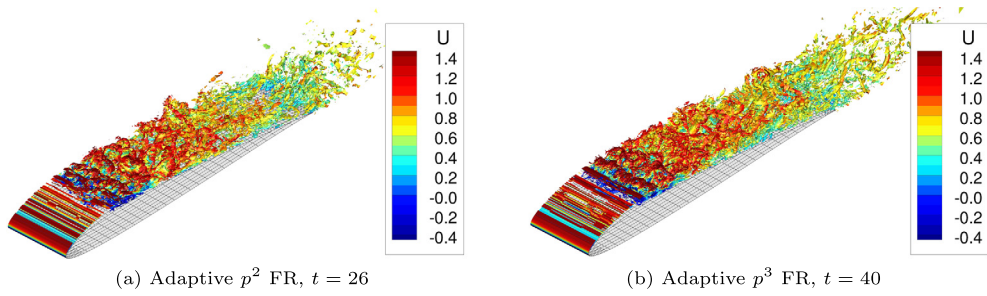


Fig. 12. Instantaneous Q -isosurfaces colored by velocity component u in the x direction when simulating the transitional flow on the 20-layer mesh.




Fig. 13. Instantaneous order-of-accuracy distributions of adaptive p^2 FR at different slices in the spanwise direction when simulating the transitional flow over the SD7003 wing. $t = 26$.




Fig. 14. Instantaneous order-of-accuracy distributions of adaptive p^3 FR at different slices in the spanwise direction when simulating the transitional flow over the SD7003 wing. $t = 40$.

The mean field of the averaged velocity component u in the x direction are presented in Fig. 16. Predictions of the time-averaged flow features, namely, lift coefficient C_l , drag coefficient C_d , separation point x_s , and reattachment point x_{re} , are documented in Table 2. The reduction of the total number of solution points n_{sp}^{tot} at specific time instances are also provided in Table 2. The time-averaged surface pressure coefficient C_p and surface friction coefficient C_f on the SD7003 wing are illustrated in Fig. 17. All numerical experiments over-predict the drag when compared to the experimental result [74]. Simulation of the incompressible Navier–Stokes equations done by Bassi et al. [54] gave larger C_d than those of the compressible Navier–Stokes equations when $Ma = 0.1$. The numerical results of our current work have a decent agreement with those in [75]. For simulation conducted on the 20-layer mesh, the adaptive p^2 FR has 1,644,414 solution points at $t_1 = 26$; the adaptive p^3 FR has 2,570,297 solution points at $t_2 = 42$. The reductions compared to p -uniform p^2 and p^3 FR are 44.40% and 63.34%, respectively. The reduction of solution points on the 10-layer mesh using p -adaptive p^3 FR is 60.57% at




Fig. 15. Power spectral density of the total velocity at different locations in the wake region. (a)–(d) adaptive p^2 FR and (e)–(h) adaptive p^3 FR. A line of slope $-5/3$ is added to every graph as a reference. Simulations are conducted on the 20-layer mesh.

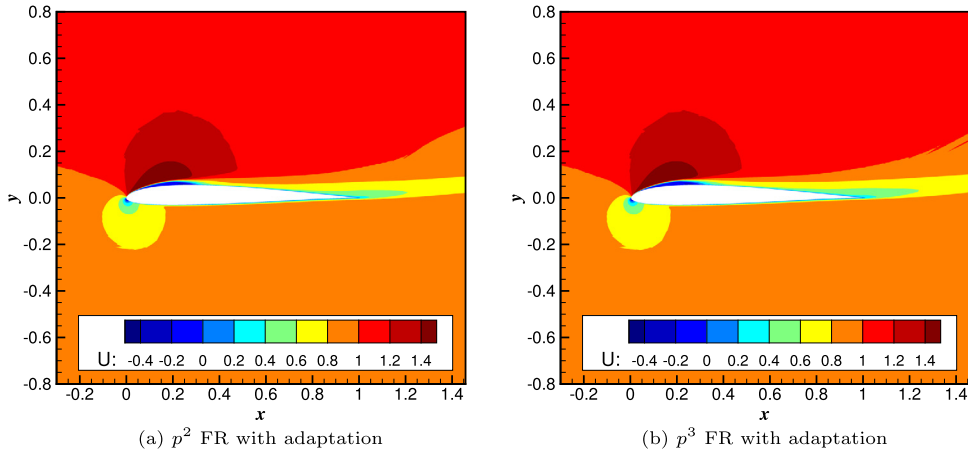


Fig. 16. Averaged velocity component in the x direction on the 20-layer mesh.

$t_1 = 26$. From the reduction in the total number of solution points, we speculate a similar reduction in the computational cost or run time. We note that insufficient resolution of the physical scales will lead to failure of the under-resolved turbulence simulation due to insufficient grid resolution in the 10-layer mesh. However, with a small dissipation introduced by nodal polynomial filtering, the force prediction can be accurate to 0.01 and the length of separation bubble is only slightly shorter. It is hard to know whether one under-resolved turbulence simulation will fail due to aliasing errors. Therefore, in the practice of performing under-resolved turbulence simulation using high-order methods, we would recommend to employ proper de-aliasing techniques.

We profile our code for simulating this problem on the 20-layer mesh using adaptive p^2 , p^3 , and p^4 solvers. We restart our simulations from previous results and average the run time of each components for 20 time steps. $(v_{max}, v_{min}) = (0.1, 0.001)$ is employed for this test. Specifically, we document the cost of mesh repartitioning \hat{t}_1 , the cost of data redistribution \hat{t}_2 , the total cost of adaptation $\hat{t}_{1,2} = \hat{t}_1 + \hat{t}_2$, the cost of element Jacobi evaluation \hat{t}_3 , the cost of preconditioner evaluation \hat{t}_4 , and the cost of the linear solver \hat{t}_5 , as well as the total cost of one implicit time stepping \hat{t}_0 , in Table 3. ESDIRK2 and ESDIRK4 are employed as the time integrators and $\Delta t = 0.002$ and $\Delta t = 0.004$ are tested. Note that nonlinear

Table 2

Predictions of the transitional flow using the p -adaptive solver. Rows 1–3 are from current work using p -adaptive FR. Rows 4–9 are previous numerical results using p -uniform high-order methods. Row 10 presents the results from experiment. The abbreviation Inc. stands for “incompressible”.

	Spatial discretization	Ma	C_l	C_d	x_s	x_{re}	Reduction of n_{sp}^{tot}
1	p -adaptive, p^2 FR (20-layer)	0.1	0.9289	0.0459	0.0321	0.3075	44.40% at $t_1 = 26$
2	p -adaptive, p^3 FR (20-layer)	0.1	0.9270	0.0470	0.0301	0.3123	63.34% at $t_2 = 42$
3	p -adaptive, p^3 FR (10-layer)	0.1	0.9316	0.0419	0.0336	0.2735	60.57% at $t_1 = 26$
4	p^4 FR (Vermeire et al. [76])	0.2	0.941	0.049	0.045	0.315	
5	p^3 DG (Beck et al. [75])	0.1	0.923	0.045	0.027	0.310	
6	p^7 DG (Beck et al. [75])	0.1	0.932	0.050	0.030	0.336	
7	$O(h^6)$ FD (Galbriath & Visbal [13])	0.1	0.91	0.043	0.04	0.28	
8	p^3 DG (Bassi et al. [54])	Inc.	0.962	0.042	0.027	0.268	
9	p^4 DG (Bassi et al. [54])	Inc.	0.953	0.045	0.027	0.294	
10	Experiment (Selig et al. [74])		0.92	0.029			

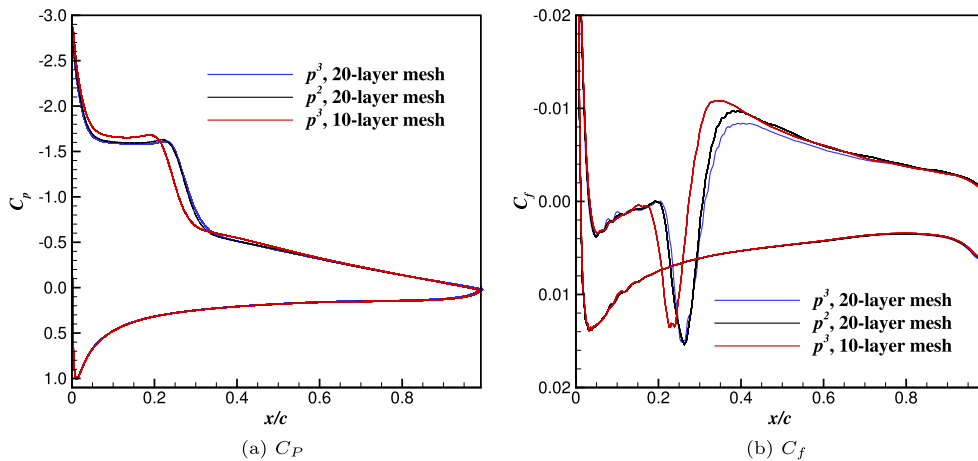


Fig. 17. Time-averaged surface pressure coefficient C_p and surface friction coefficient C_f on the SD7003 wing.

tolerance $tol_{rel}^{pseudo} = 10^{-4}$ is employed for all cases. $tol_{rel}^{gmres} = 10^{-1}$ is used for ESDIRK2 and $tol_{rel}^{gmres} = 10^{-2}$ is employed for ESDIRK4. The averaged number of solution points for adaptive p^2 , p^3 , and p^4 solvers are around 1.6×10^6 , 2.6×10^6 , and 4.0×10^6 , respectively. We find that the cost of adaptation is trivial compared to other operations in one implicit time stepping. We observe that either increasing the polynomial order p or the number of stages of ESDIRK will decrease the ratio $\hat{t}_{1,2}/\hat{t}_0$, which is only 0.004 when ESDIRK4 with $\Delta t = 0.004$ and adaptive p^4 FR are used. This justifies performing p -adaptation and load balancing every time step. As the polynomial degree increases, the cost (i.e., the ratio \hat{t}_4/\hat{t}_0) of evaluating the preconditioner becomes more and more dominant in one implicit time stepping, and the relative cost (i.e., the ratio \hat{t}_5/\hat{t}_0) of the linear solver decreases. This indicates that research efforts are still needed to further improve the performance of preconditioners. As mentioned previously, applying the p -multigrid solver [69] as a preconditioner for Newton-Krylov methods will be our future work.

5. Conclusion

In this work, a dynamically load balanced parallel p -adaptive implicit high-order flux reconstruction method is developed and applied to under-resolved turbulence simulation. The parallel mesh partitioning API in ParMETIS, i.e., `ParMetis_V3_AdaptiveRepart()`, is utilized for efficient parallel mesh partitioning. A collect-and-distribute strategy is used to redistribute the working variables to different processes. We have discussed the impact of weight calculation for each element on the parallel efficiency in the context of matrix-free implementation of the ESDIRK method. We investigate different weights related to the cost of residual evaluation, Jacobian matrix and preconditioner evaluation, and GMRES iterations. For $p \leq 3$, we recommend $\omega_e = n_{sp}$, and as p grows larger, $\omega_e = n_{sp}^k$, where $k > 1$, is more preferable. Overall, a significant reduction in the run time and total number of solution points can be achieved via p -adaptation for turbulence simulation and favorable results can be obtained.

When the adaptive solver is applied to solving the transitional flow over an infinite cylinder, due to the presence of large flow separation, the featured-based solver can result in a large domain where the polynomial degrees are refined. One can adjust the adaptation criteria towards wall-resolving to save computational cost; however, the accuracy would be compromised. When the flow separation is small, e.g., transitional flow over the SD7003 wing at a moderate angle of attack,

Table 3Profiling results of the adaptive p^2 , p^3 , and p^4 FR methods on simulating the transitional flow over the SD7003 wing. Time unit is second.

Adaptive p^2 FR								
ESDIRK2, $\Delta t = 0.002$		ESDIRK2, $\Delta t = 0.004$		ESDIRK4, $\Delta t = 0.002$		ESDIRK4, $\Delta t = 0.004$		
Run time	Ratio	Run time	Ratio	Run time	Ratio	Run time	Ratio	
\hat{t}_1	0.392	0.035	0.393	0.025	0.375	0.021	0.386	0.014
\hat{t}_2	0.178	0.016	0.178	0.011	0.143	0.008	0.184	0.007
$\hat{t}_{1,2}$	0.570	0.050	0.571	0.036	0.518	0.030	0.570	0.021
\hat{t}_3	1.507	0.133	1.549	0.099	1.375	0.079	1.372	0.051
\hat{t}_4	1.126	0.099	1.174	0.075	1.015	0.058	1.016	0.038
\hat{t}_5	8.132	0.717	12.376	0.790	14.562	0.834	23.919	0.890
\hat{t}_0	11.335	1.000	15.670	1.000	17.470	1.000	26.876	1.000
Adaptive p^3 FR								
ESDIRK2, $\Delta t = 0.002$		ESDIRK2, $\Delta t = 0.004$		ESDIRK4, $\Delta t = 0.002$		ESDIRK4, $\Delta t = 0.004$		
Run time	Ratio	Run time	Ratio	Run time	Ratio	Run time	Ratio	
\hat{t}_1	0.433	0.013	0.436	0.010	0.436	0.007	0.417	0.006
\hat{t}_2	0.270	0.008	0.271	0.006	0.271	0.004	0.258	0.004
$\hat{t}_{1,2}$	0.703	0.021	0.707	0.016	0.707	0.011	0.675	0.009
\hat{t}_3	4.588	0.136	4.569	0.106	4.620	0.072	4.403	0.061
\hat{t}_4	9.761	0.289	9.626	0.223	9.692	0.152	9.506	0.132
\hat{t}_5	18.722	0.554	28.267	0.655	48.885	0.765	57.653	0.798
\hat{t}_0	33.775	1.000	43.169	1.000	63.903	1.000	72.237	1.000
Adaptive p^4 FR								
ESDIRK2, $\Delta t = 0.002$		ESDIRK2, $\Delta t = 0.004$		ESDIRK4, $\Delta t = 0.002$		ESDIRK4, $\Delta t = 0.004$		
Run time	Ratio	Run time	Ratio	Run time	Ratio	Run time	Ratio	
\hat{t}_1	0.523	0.004	0.538	0.004	0.536	0.003	0.679	0.002
\hat{t}_2	0.381	0.003	0.389	0.003	0.408	0.002	0.512	0.002
$\hat{t}_{1,2}$	0.904	0.007	0.927	0.007	0.945	0.005	1.191	0.004
\hat{t}_3	12.943	0.104	12.890	0.091	12.983	0.065	17.701	0.059
\hat{t}_4	67.902	0.548	67.652	0.478	68.083	0.339	77.380	0.256
\hat{t}_5	42.141	0.340	60.072	0.424	118.913	0.592	206.231	0.682
\hat{t}_0	123.890	1.000	141.540	1.000	200.924	1.000	302.502	1.000

the feature-based p -adaptation method is able to confine the p -refined region close to the wing, thus significantly reducing the cost while providing good predictions. We also show, with the SD7003 case, that insufficient mesh resolution can lead to instabilities triggered by aliasing errors of high-order methods in under-resolved turbulence simulation. A proper de-aliasing technique can overcome this issue and provide acceptable predictions. According to the computational cost profiling results from transitional flow simulations, we find that the cost of adaptation is trivial compared to that of one implicit time stepping. This justifies our approach to conduct p -adaptation and load balancing every time step. We also find that the cost of preconditioner evaluation increases when the spatial polynomial degree increases. This indicates that research efforts are still needed to further improve preconditioners' performance for high-order numerical simulation with Newton-Krylov methods.

The framework of dynamically load balanced p -adaptive implicit high-order methods developed in this study paves the way towards robust and efficient ILES of turbulent flows at higher Reynolds numbers with the high-order FR/CPR method. The dynamic load balancing technique presented here can be easily extended to other types of high-order collocation methods.

Declaration of competing interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Acknowledgements

Wang and Yu gratefully acknowledge the support of the Office of Naval Research through the award N00014-16-1-2735, and the faculty startup support from the department of mechanical engineering at the University of Maryland, Baltimore County (UMBC). The hardware used in the computational studies is part of the UMBC High Performance Computing Facility (HPCF). The facility is supported by the U.S. National Science Foundation through the MRI program (grant nos. CNS-0821258, CNS-1228778, and OAC-1726023) and the SCREMS program (grant no. DMS-0821311), with additional substantial support from UMBC.

References

- [1] 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) 411–435.
- [2] B. Cockburn, S. Hou, C.-W. Shu, The Runge-Kutta local projection discontinuous Galerkin finite element method for conservation laws. IV. The multidimensional case, *Math. Comput.* 54 (190) (1990) 545–581.
- [3] 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 (2) (1997) 267–279.
- [4] B. Cockburn, C.-W. Shu, Runge–Kutta discontinuous Galerkin methods for convection-dominated problems, *J. Sci. Comput.* 16 (3) (2001) 173–261.
- [5] F. Bassi, A. Crivellini, S. Rebay, M. Savini, Discontinuous Galerkin solution of the Reynolds-averaged Navier–Stokes and $k-\omega$ turbulence model equations, *Comput. Fluids* 34 (4–5) (2005) 507–540.
- [6] J.S. Hesthaven, T. Warburton, *Nodal Discontinuous Galerkin Methods: Algorithms, Analysis, and Applications*, Springer-Verlag, New York, 2008.
- [7] G.J. Gassner, A.D. Beck, On the accuracy of high-order discretizations for underresolved turbulence simulations, *Theor. Comput. Fluid Dyn.* 27 (3–4) (2013) 221–237.
- [8] A. Uranga, P.-O. Persson, M. Dreha, J. Peraire, Implicit large eddy simulation of transition to turbulence at low Reynolds numbers using a discontinuous Galerkin method, *Int. J. Numer. Methods Eng.* 87 (1–5) (2011) 232–261.
- [9] Y. Liu, M. Vinokur, Z.J. Wang, Spectral difference method for unstructured grids I: basic formulation, *J. Comput. Phys.* 216 (2) (2006) 780–801.
- [10] Y. Zhou, Z.J. Wang, Implicit large Eddy simulation of low Reynolds number transitional flow over a wing using high-order spectral difference method, in: 40th Fluid Dynamics Conference and Exhibit, 2010, p. 4442.
- [11] P. Castonguay, C. Liang, A. Jameson, Simulation of transitional flow over airfoils using the spectral difference method, in: 40th Fluid Dynamics Conference and Exhibit, 2010, p. 4626.
- [12] M.L. Yu, Z.J. Wang, H. Hu, A high-order spectral difference method for unstructured dynamic grids, *Comput. Fluids* 48 (2011) 84–97.
- [13] M. Galbraith, M. Visbal, Implicit large eddy simulation of low Reynolds number flow past the SD7003 airfoil, in: 46th AIAA Aerospace Sciences Meeting and Exhibit, 2008, p. 225.
- [14] D.J. Garmann, M.R. Visbal, P.D. Orkwis, Comparative study of implicit and subgrid-scale model large-eddy simulation techniques for low-Reynolds number airfoil applications, *Int. J. Numer. Methods Fluids* 71 (12) (2013) 1546–1565.
- [15] P. Boom, D. Zingg, Time-accurate flow simulations using an efficient newton-krylov-schur approach with high-order temporal and spatial discretization, in: 51st AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 2013, p. 383.
- [16] M.H. Carpenter, T.C. Fisher, E.J. Nielsen, S.H. Frankel, Entropy stable spectral collocation schemes for the Navier–Stokes equations: discontinuous interfaces, *SIAM J. Sci. Comput.* 36 (5) (2014) B835–B867.
- [17] M. Svärd, J. Nordström, Review of summation-by-parts schemes for initial-boundary-value problems, *J. Comput. Phys.* 268 (2014) 17–38.
- [18] H.T. Huynh, A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods, in: 18th AIAA Computational Fluid Dynamics Conference, 2007, p. 4079.
- [19] H.T. Huynh, A reconstruction approach to high-order schemes including discontinuous Galerkin for diffusion, in: 47th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 2009, p. 403.
- [20] 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.* 228 (21) (2009) 8161–8186.
- [21] P.E. Vincent, P. Castonguay, A. Jameson, A new class of high-order energy stable flux reconstruction schemes, *J. Sci. Comput.* 47 (1) (2011) 50–72.
- [22] J. Romero, K. Asthana, A. Jameson, A simplified formulation of the flux reconstruction method, *J. Sci. Comput.* 67 (1) (2016) 351–374.
- [23] L. Wang, M. Yu, Compact direct flux reconstruction for conservation laws, *J. Sci. Comput.* (2017) 1–23.
- [24] F.D. Witherden, A.M. Farrington, P.E. Vincent, PyFR: an open source framework for solving advection–diffusion type problems on streaming architectures using the flux reconstruction approach, *Comput. Phys. Commun.* 185 (11) (2014) 3028–3040.
- [25] P. Fernandez, N.-C. Nguyen, J. Peraire, Subgrid-scale modeling and implicit numerical dissipation in DG-based large-Eddy simulation, in: 23rd AIAA Computational Fluid Dynamics Conference, 2017, p. 3951.
- [26] G.J. Gassner, A skew-symmetric discontinuous Galerkin spectral element discretization and its relation to SBP-SAT finite difference methods, *SIAM J. Sci. Comput.* 35 (3) (2013) A1233–A1253.
- [27] G.J. Gassner, A.R. Winters, D.A. Kopriva, Split form nodal discontinuous Galerkin schemes with summation-by-parts property for the compressible Euler equations, *J. Comput. Phys.* 327 (2016) 39–66.
- [28] R.M. Kirby, G.E. Karniadakis, De-aliasing on non-uniform grids: algorithms and applications, *J. Comput. Phys.* 191 (1) (2003) 249–264.
- [29] G. Mengaldo, D. De Grazia, D. Moxey, P.E. Vincent, S.J. Sherwin, Dealiasing techniques for high-order spectral element methods on regular and irregular grids, *J. Comput. Phys.* 299 (2015) 56–81.
- [30] E. Tadmor, Convergence of spectral methods for nonlinear conservation laws, *SIAM J. Numer. Anal.* 26 (1) (1989) 30–44.
- [31] G. Karamanos, G.E. Karniadakis, A spectral vanishing viscosity method for large-eddy simulations, *J. Comput. Phys.* 163 (1) (2000) 22–50.
- [32] R. Pasquetti, Spectral vanishing viscosity method for large-eddy simulation of turbulent flows, *J. Sci. Comput.* 27 (1–3) (2006) 365–375.
- [33] D. Gottlieb, J.S. Hesthaven, Spectral methods for hyperbolic problems, *J. Comput. Appl. Math.* 128 (1–2) (2001) 83–131.
- [34] J. Hesthaven, R. Kirby, Filtering in Legendre spectral methods, *Math. Comput.* 77 (263) (2008) 1425–1452.
- [35] P. Fischer, J. Mullen, Filter-based stabilization of spectral element methods, *C. R. Acad. Sci., Ser. 1 Math.* 332 (3) (2001) 265–270.
- [36] Y. Li, S. Premasathan, A. Jameson, Comparison of h - and p -adaptations for spectral difference methods, *AIAA Pap.* 4435 (2010).
- [37] M. Woopen, A. Balan, G. May, J. Schütz, A comparison of hybridized and standard DG methods for target-based hp -adaptive simulation of compressible flow, *Comput. Fluids* 98 (2014) 3–16.
- [38] G. Gassner, M. Staudenmaier, F. Hindenlang, M. Atak, C.-D. Munz, A space–time adaptive discontinuous Galerkin scheme, *Comput. Fluids* 117 (2015) 247–261.
- [39] M. Tognoli, A. Abbà, L. Bonaventura, M. Restelli, A locally p -adaptive approach for large Eddy simulation of compressible flows in a DG framework, *J. Comput. Phys.* 349 (2017) 33–58.
- [40] F. Naddei, M. de la Llave Plata, V. Couaillier, A comparison of refinement indicators for p -adaptive discontinuous Galerkin methods for the Euler and Navier–Stokes equations, in: 2018 AIAA Aerospace Sciences Meeting, 2018, p. 0368.
- [41] R. Hartmann, P. Houston, Adaptive discontinuous Galerkin finite element methods for the compressible Euler equations, *J. Comput. Phys.* 183 (2) (2002) 508–532.
- [42] H. Gao, Z. Wang, A residual-based procedure for hp -adaptation on 2-D hybrid meshes, in: 49th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, 2011, p. 492.
- [43] M. Kompenhans, G. Rubio, E. Ferrer, E. Valero, Comparisons of p -adaptation strategies based on truncation- and discretisation-errors for high order discontinuous Galerkin methods, *Comput. Fluids* 139 (2016) 36–46.
- [44] D.A. Venditti, D.L. Darmofal, Anisotropic grid adaptation for functional outputs: application to two-dimensional viscous flows, *J. Comput. Phys.* 187 (1) (2003) 22–46.

- [45] L. Wang, D.J. Mavriplis, Adjoint-based h–p adaptive discontinuous Galerkin methods for the 2D compressible Euler equations, *J. Comput. Phys.* 228 (20) (2009) 7643–7661.
- [46] K.J. Fidkowski, D.L. Darmofal, Review of output-based error estimation and mesh adaptation in computational fluid dynamics, *AIAA J.* 49 (4) (2011) 673–694.
- [47] K.J. Fidkowski, Y. Luo, Output-based space–time mesh adaptation for the compressible Navier–Stokes equations, *J. Comput. Phys.* 230 (14) (2011) 5753–5773.
- [48] F. Bassi, A. Colombo, A. Crivellini, K. Fidkowski, M. Franciolini, A. Ghidoni, G. Noventa, An entropy-adjoint p-adaptive discontinuous Galerkin method for the under-resolved simulation of turbulent flows, in: *AIAA Aviation 2019 Forum*, 2019, p. 3418.
- [49] K.J. Fidkowski, P.L. Roe, An entropy adjoint approach to mesh refinement, *SIAM J. Sci. Comput.* 32 (3) (2010) 1261–1287.
- [50] P.-O. Persson, J. Peraire, Sub-cell shock capturing for discontinuous Galerkin methods, in: *44th AIAA Aerospace Sciences Meeting and Exhibit*, 2006, p. 112.
- [51] B. Cockburn, C.-W. Shu, TVB Runge–Kutta local projection discontinuous Galerkin finite element method for conservation laws. II. General framework, *Math. Comput.* 52 (186) (1989) 411–435.
- [52] S. Gottlieb, C.-W. Shu, E. Tadmor, Strong stability-preserving high-order time discretization methods, *SIAM Rev.* 43 (1) (2001) 89–112.
- [53] C.A. Kennedy, M.H. Carpenter, Diagonally implicit Runge–Kutta methods for ordinary differential equations. A review, *NASA/TM–2016–219173*, 2016.
- [54] F. Bassi, L. Botti, A. Colombo, A. Ghidoni, F. Massa, Linearly implicit Rosenbrock-type Runge–Kutta schemes applied to the discontinuous Galerkin solution of compressible and incompressible unsteady flows, *Comput. Fluids* 118 (2015) 305–320.
- [55] L. Wang, M. Yu, A comparative study of implicit Jacobian-free Rosenbrock–Wanner, ESDIRK and BDF methods for unsteady flow simulation with high-order flux reconstruction formulations, preprint, arXiv:1904.04825, 2019.
- [56] M. Franciolini, A. Crivellini, A. Nigro, On the efficiency of a matrix-free linearly implicit time integration strategy for high-order discontinuous Galerkin solutions of incompressible turbulent flows, *Comput. Fluids* 159 (2017) 276–294.
- [57] D.A. Knoll, D.E. Keyes, Jacobian-free Newton–Krylov methods: a survey of approaches and applications, *J. Comput. Phys.* 193 (2) (2004) 357–397.
- [58] C. Kennedy, M. Carpenter, Additive Runge–Kutta schemes for convection–diffusion–reaction equations, *Appl. Numer. Math.* 44 (2003) 139–181.
- [59] L. Pareschi, G. Russo, Implicit–explicit Runge–Kutta schemes and applications to hyperbolic systems with relaxation, *J. Sci. Comput.* 25 (1) (2005) 129–155.
- [60] L. Wang, M. Yu, Jacobian-free implicit p-adaptive high-order compact direct flux reconstruction methods for unsteady flow simulation, in: *AIAA Aviation 2019 Forum*, 2019, p. 3062.
- [61] G. Karypis, METIS and ParMETIS, in: *Encyclopedia of Parallel Computing*, 2011, pp. 1117–1124.
- [62] M. Yu, Z.J. Wang, On the connection between the correction and weighting functions in the correction procedure via reconstruction method, *J. Sci. Comput.* 54 (1) (2013) 227–244.
- [63] D.M. Williams, P. Castonguay, P.E. Vincent, A. Jameson, Energy stable flux reconstruction schemes for advection–diffusion problems on triangles, *J. Comput. Phys.* 250 (2013) 53–76.
- [64] P.L. Roe, Approximate Riemann solvers, parameter vectors, and difference schemes, *J. Comput. Phys.* 43 (2) (1981) 357–372.
- [65] 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 (1) (2002) 313–329.
- [66] W.A. Mulder, B. Van Leer, Experiments with implicit upwind methods for the Euler equations, *J. Comput. Phys.* 59 (2) (1985) 232–246.
- [67] S. Balay, S. Abhyankar, M.F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, V. Eijkhout, W.D. Gropp, D. Kaushik, M.G. Knepley, L.C. McInnes, K. Rupp, B.F. Smith, S. Zampini, H. Zhang, PETSc users manual, Tech. Rep. ANL-95/11 - Revision 3.6, Argonne National Laboratory, 2015.
- [68] L. Wang, M. Yu, An implicit high-order preconditioned flux reconstruction method for low-Mach-number flow simulation with dynamic meshes, *Int. J. Numer. Methods Fluids* (2019).
- [69] L. Wang, M. Yu, An implicit P -multigrid flux reconstruction method for simulation of locally preconditioned unsteady Navier–Stokes equations at low Mach numbers, preprint, arXiv:1908.03972, 2019.
- [70] C. Burstedde, L.C. Wilcox, O. Ghattas, p4est: scalable algorithms for parallel adaptive mesh refinement on forests of octrees, *SIAM J. Sci. Comput.* 33 (3) (2011) 1103–1133.
- [71] M. Lange, L. Mitchell, M.G. Knepley, G.J. Gorman, Efficient mesh management in firedrake using petsc dmpex, *SIAM J. Sci. Comput.* 38 (5) (2016) S143–S155.
- [72] O. Lehmkuhl, I. Rodríguez, R. Borrell, A. Oliva, Low-frequency unsteadiness in the vortex formation region of a circular cylinder, *Phys. Fluids* 25 (8) (2013) 085109.
- [73] F. Bassi, L. Botti, A. Colombo, A. Crivellini, A. Ghidoni, F. Massa, On the development of an implicit high-order discontinuous Galerkin method for DNS and implicit LES of turbulent flows, *Eur. J. Mech. B, Fluids* 55 (2016) 367–379.
- [74] M.S. Selig, Summary of Low Speed Airfoil Data, Vol. 1, SoarTech Publications, 1995.
- [75] A.D. Beck, T. Bolemann, D. Flad, H. Frank, G.J. Gassner, F. Hindenlang, C.-D. Munz, High-order discontinuous Galerkin spectral element methods for transitional and turbulent flow simulations, *Int. J. Numer. Methods Fluids* 76 (8) (2014) 522–548.
- [76] B.C. Vermeire, F.D. Witherden, P.E. Vincent, On the utility of GPU accelerated high-order methods for unsteady flow simulations: a comparison with industry-standard tools, *J. Comput. Phys.* 334 (2017) 497–521.