



The hybridized Discontinuous Galerkin method for Implicit Large-Eddy Simulation of transitional turbulent flows

P. Fernandez^{a,b,*}, N.C. Nguyen^{a,b}, J. Peraire^{a,b}

^a Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

^b Center for Computational Engineering, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA

ARTICLE INFO

Article history:

Received 2 May 2016

Received in revised form 30 December 2016

Accepted 4 February 2017

Available online 8 February 2017

Keywords:

Discontinuous Galerkin methods

Implicit Large-Eddy Simulation

Navier–Stokes iterative solvers

High-order methods

Laminar separation bubble

Transition to turbulence

ABSTRACT

We present a high-order Implicit Large-Eddy Simulation (ILES) approach for transitional aerodynamic flows. The approach encompasses a hybridized Discontinuous Galerkin (DG) method for the discretization of the Navier–Stokes (NS) equations, and a parallel preconditioned Newton–GMRES solver for the resulting nonlinear system of equations. The combination of hybridized DG methods with an efficient solution procedure leads to a high-order accurate NS solver that is competitive to alternative approaches, such as finite volume and finite difference codes, in terms of computational cost. The proposed approach is applied to transitional flows over the NACA 65–(18)10 compressor cascade and the Eppler 387 wing at Reynolds numbers up to 460,000. Grid convergence studies are presented and the required resolution to capture transition at different Reynolds numbers is investigated. Numerical results show rapid convergence and excellent agreement with experimental data. In short, this work aims to demonstrate the potential of high-order ILES for simulating transitional aerodynamic flows. This is illustrated through numerical results and supported by theoretical considerations.

© 2017 Elsevier Inc. All rights reserved.

1. Introduction

Numerical simulation of transitional turbulent flows is a challenging problem. First, turbulent flows exhibit spatial and temporal scales that differ by several orders of magnitude, thereby making numerical simulation of turbulence computationally demanding. Second, transitional flows are hard to predict due to the complex physical phenomena that take place when the flow undergoes transition from laminar to turbulent regime. Indeed, transitional turbulent flows are not yet well understood theoretically despite being studied for over a century.

In spite of the availability of powerful supercomputers, complete information about turbulent flows through Direct Numerical Simulation (DNS) remains intractable for most practical applications. Large-Eddy Simulation (LES) is a viable alternative to DNS. The central premise of LES is that large-scale eddies dominate the turbulent transport and energy budget, so that a numerical simulation will provide a realistic depiction of the flow if it captures those scales explicitly and somehow accounts for the small scales that cannot be resolved. Also, the small scales tend to be more homogeneous and isotropic, and thus easier to model. Strategies for dealing with the small turbulent scales include explicit subgrid-scale (SGS) modeling and implicit numerical dissipation.

* Corresponding author at: Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 77 Massachusetts Avenue, Cambridge, MA 02139, USA.

E-mail addresses: pablof@mit.edu (P. Fernandez), cuongng@mit.edu (N.C. Nguyen), peraire@mit.edu (J. Peraire).

In the classical (explicit) LES approach, the large-scale eddies of the flow field are resolved and the small scales are modeled using a SGS model. The development of SGS models has been a subject of intense interest for decades, in particular in the 1990s [22,31]. It turns out, however, that the leading-order term of the truncation error introduced by many numerical schemes is similar in form and magnitude to conventional SGS models. As a rule of thumb, a stable numerical scheme often achieves stability by introducing truncation errors that replicate the effect of the subgrid scales into the resolved scales; which corresponds to dissipation in under-resolved turbulent simulations. This has been known to be the case for monotone (and thus stable) numerical schemes for about 20 years [4,9,14,20,27,41] and is now also thought to be the case for stabilized high-order methods [18,32].

A natural alternative to the classical LES approach is therefore to use the numerical dissipation of the discretization scheme to account for the dissipation that takes place in the unresolved scales, leading to the so-called Implicit LES (ILES). The ILES approach was first introduced in 1990 by Boris et al. [3] and has been successfully applied with a number of different schemes, including finite volume methods [13,15,16], standard [19] and compact [17,46] finite difference methods, spectral difference methods [50], spectral/*hp* element methods [24], flux reconstruction methods [36], and discontinuous Galerkin methods [12,33,40,47–49]. ILES benefits from its easy implementation without a SGS model and currently gains considerable attention from researchers in the computational fluid dynamics community. As pointed out by Spalart [45], this increase in popularity may be attributed to the fact that research has failed to show an advantage of sophisticated SGS models over the *same-cost* LES with a simplistic model –or even with no model– and a slightly finer grid.

While second-order finite volume (FV) schemes have been widely used in academia and industry for LES, transitional flows do share many features with wave propagation phenomena for which high-order accuracy is known to be key. Let us consider natural transition to turbulence. The main difficulty to numerically capture transition is the very small magnitude of the perturbations that get exponentially amplified along the unstable portion of the laminar boundary layer. These small perturbations are ultimately responsible for the so-called nonlinear breakdown and transition to turbulence. The amplitude of these instabilities at the location in which the boundary layer becomes unstable is up to ten orders of magnitude below the freestream velocity [44]. As a result, very small amount of numerical dissipation and dispersion is needed to capture them and accurately predict the transition location. Overdissipation of low-order schemes may kill these small perturbations and lead to inaccurate transition prediction. This motivates researchers to consider high-order methods for the simulation of transitional flows.

High-order finite difference (FD), finite volume, and finite element (FE) methods have been developed for numerical simulation of turbulent flows. Among them, we rely on Discontinuous Galerkin (DG) finite element methods for several reasons. First, DG methods are based on a strong mathematical foundation that can be exploited for error estimation and mesh adaptation purposes. Also, they provide local conservation, a stable discretization of the convective operator, and are well-suited for turbulent simulations due to the *ab initio* separation of scales in the variational formulation. Most importantly, DG methods allow for high-order implementations on complex geometries and unstructured meshes. This is crucial for engineering purposes due to the need to simulate complex three-dimensional geometries. On the other hand, DG methods require significantly more operations per computational cell than FD and FV methods. Hence, the need for more computationally efficient DG methods becomes apparent. In the spirit of making DG methods competitive, researchers have recently developed the Hybridizable DG (HDG) method [6,34,37] and the Embedded DG (EDG) method [6,7,38]. These DG methods result in spatial discretizations that have fewer globally coupled unknowns than other DG methods, thereby reducing the computational cost and memory footprint.

At present, ILES of transitional flows using high-order DG methods is limited to Reynolds numbers of 100,000 or less [12,33,40,47–49]. It may be attributed to the fact that higher Reynolds number flows would require significantly more computational effort than standard DG methods could afford in most current computing clusters. This motivates us to employ hybridized DG methods [11,35], which generalize the HDG and the EDG methods, for ILES of transitional turbulent flows. This work contains novel contributions from the points of view of methodology and applications. In terms of methodology, we develop parallel preconditioned Newton-GMRES solvers for the nonlinear system arising from the hybridized DG discretization of the Navier–Stokes equations. In terms of applications, this paper is, together with preliminary results in [10,11], the first attempt at the high-order DG ILES of transitional flows at Reynolds number up to 460,000 that is beyond the existing simulations in the literature. We believe that this work will hold a promising approach to simulate transitional flows and contribute to the ongoing research into ILES of turbulent flows.

The paper is organized as follows. In Section 2, we describe the numerical discretization of the Navier–Stokes equations using hybridized DG methods. In Section 3, we discuss solution methods for the nonlinear system of equations arising from the numerical discretization. In Sections 4 and 5, the proposed approach is applied to the ILES of transitional flows over the NACA 65-(18)10 compressor cascade and the Eppler 387 wing at Reynolds numbers ranging from 100,000 to 460,000. Finally, some concluding remarks and a rationale for the success of high-order ILES for transition prediction are presented in Section 6.

2. Flow discretization

In this section, we present an overview of hybridized discontinuous Galerkin methods for solving the unsteady, compressible Navier–Stokes equations.

2.1. Preliminaries and notation

2.1.1. Finite element mesh

Let $\Omega \subseteq \mathbb{R}^d$ with $d = 3$ be a physical domain with Lipschitz boundary $\partial\Omega$. We denote by \mathcal{T}_h a collection of disjoint, regular, p -th degree curved elements K that partition Ω , and set $\partial\mathcal{T}_h := \{\partial K : K \in \mathcal{T}_h\}$ to be the collection of the boundaries of the elements in \mathcal{T}_h . For an element K of the collection \mathcal{T}_h , $F = \partial K \cap \partial\Omega$ is a boundary face if its $d - 1$ Lebesgue measure is nonzero. For two elements K^+ and K^- of \mathcal{T}_h , $F = \partial K^+ \cap \partial K^-$ is the interior face between K^+ and K^- if its $d - 1$ Lebesgue measure is nonzero. We denote by \mathcal{E}_h^I and \mathcal{E}_h^B the set of interior and boundary faces, respectively, and we define $\mathcal{E}_h := \mathcal{E}_h^I \cup \mathcal{E}_h^B$ as the union of interior and boundary faces. Note that, by definition, $\partial\mathcal{T}_h$ and \mathcal{E}_h are different. More precisely, an interior face is counted twice in $\partial\mathcal{T}_h$ but only once in \mathcal{E}_h , whereas a boundary face is counted once both in $\partial\mathcal{T}_h$ and \mathcal{E}_h .

2.1.2. Finite element spaces

Let $\mathcal{P}_k(D)$ denote the space of complete polynomials of degree k on a domain $D \in \mathbb{R}^n$, let $L^2(D)$ be the space of square-integrable functions on D , and let ψ_K^p denote the p -th degree parametric mapping from the reference element K_{ref} to some element $K \in \mathcal{T}_h$ in the physical domain. We then introduce the following discontinuous finite element spaces:

$$\mathcal{Q}_h^k = \{ \mathbf{r} \in [L^2(\mathcal{T}_h)]^{m \times d} : (\mathbf{r} \circ \psi^p)|_K \in [\mathcal{P}_k(K_{ref})]^{m \times d} \quad \forall K \in \mathcal{T}_h \},$$

$$\mathcal{V}_h^k = \{ \mathbf{w} \in [L^2(\mathcal{T}_h)]^m : (\mathbf{w} \circ \psi^p)|_K \in [\mathcal{P}_k(K_{ref})]^m \quad \forall K \in \mathcal{T}_h \},$$

where m denotes the number of equations of the conservation law, i.e. $m = d + 2$ for the Navier–Stokes system. In addition, let us denote the p -th degree parametric mapping from the reference face F_{ref} to some physical face F by ϕ_F^p . We then introduce the following traced finite element spaces on the mesh skeleton \mathcal{E}_h :

$$\widehat{\mathcal{M}}_h^k = \{ \boldsymbol{\mu} \in [L^2(\mathcal{E}_h)]^m : (\boldsymbol{\mu} \circ \phi_F^p)|_F \in [\mathcal{P}^k(F_{ref})]^m \quad \forall F \in \mathcal{E}_h \},$$

$$\widetilde{\mathcal{M}}_h^k = \{ \boldsymbol{\mu} \in [C^0(\mathcal{E}_h)]^m : (\boldsymbol{\mu} \circ \phi_F^p)|_F \in [\mathcal{P}^k(F_{ref})]^m \quad \forall F \in \mathcal{E}_h \},$$

where $C^0(D)$ is the space of continuous functions on D . Note that $\widehat{\mathcal{M}}_h^k$ consists of functions which are discontinuous at the boundaries of the faces, whereas $\widetilde{\mathcal{M}}_h^k$ consists of functions that are continuous at the boundaries of the faces. We also denote by \mathcal{M}_h^k a traced finite element space that satisfies $\widetilde{\mathcal{M}}_h^k \subseteq \mathcal{M}_h^k \subseteq \widehat{\mathcal{M}}_h^k$. In particular, we define

$$\mathcal{M}_h^k = \{ \boldsymbol{\mu} \in [L^2(\mathcal{E}_h)]^m : (\boldsymbol{\mu} \circ \phi_F^p)|_F \in [\mathcal{P}^k(F_{ref})]^m \quad \forall F \in \mathcal{E}_h, \text{ and } \boldsymbol{\mu}|_{\mathcal{E}_h^E} \in [C^0(\mathcal{E}_h^E)]^m \},$$

where \mathcal{E}_h^E is a subset of \mathcal{E}_h . Note that \mathcal{M}_h^k consists of functions which are continuous on \mathcal{E}_h^E and discontinuous on $\mathcal{E}_h^H := \mathcal{E}_h \setminus \mathcal{E}_h^E$. Furthermore, if $\mathcal{E}_h^E = \emptyset$ then $\mathcal{M}_h^k = \widehat{\mathcal{M}}_h^k$, and if $\mathcal{E}_h^E = \mathcal{E}_h$ then $\mathcal{M}_h^k = \widetilde{\mathcal{M}}_h^k$. Different choices of \mathcal{E}_h^E lead to different discretization methods that have different properties in terms of convergence rates, stability, and number of globally coupled unknowns. We shall discuss various rational choices in Section 2.3.

Next, we define several inner products associated with these finite element spaces. In particular, given $\mathbf{w}, \mathbf{v} \in \mathcal{V}_h^k$, $\mathbf{W}, \mathbf{V} \in \mathcal{Q}_h^k$ and $\boldsymbol{\eta}, \boldsymbol{\zeta} \in \mathcal{M}_h^k$, we write

$$\langle \mathbf{w}, \mathbf{v} \rangle_{\mathcal{T}_h} = \sum_{K \in \mathcal{T}_h} \langle \mathbf{w}, \mathbf{v} \rangle_K = \sum_{K \in \mathcal{T}_h} \int_K \mathbf{w} \cdot \mathbf{v}_K, \quad (1a)$$

$$\langle \mathbf{W}, \mathbf{V} \rangle_{\mathcal{T}_h} = \sum_{K \in \mathcal{T}_h} \langle \mathbf{W}, \mathbf{V} \rangle_K = \sum_{K \in \mathcal{T}_h} \int_K \mathbf{W} : \mathbf{V}, \quad (1b)$$

$$\langle \boldsymbol{\eta}, \boldsymbol{\zeta} \rangle_{\partial\mathcal{T}_h} = \sum_{K \in \mathcal{T}_h} \langle \boldsymbol{\eta}, \boldsymbol{\zeta} \rangle_{\partial K} = \sum_{K \in \mathcal{T}_h} \int_{\partial K} \boldsymbol{\eta} \cdot \boldsymbol{\zeta}, \quad (1c)$$

where $:$ denotes the Frobenius inner product.

2.2. Governing equations

We consider the unsteady, compressible Navier–Stokes equations written in non-dimensional conservation form as

$$\begin{aligned} \mathbf{q} - \nabla \mathbf{u} &= \mathbf{0}, \quad \text{in } \Omega \times (0, T), \\ \frac{\partial \mathbf{u}}{\partial t} + \nabla \cdot \mathbf{F}(\mathbf{u}, \mathbf{q}) &= \mathbf{0}, \quad \text{in } \Omega \times (0, T). \end{aligned} \quad (2)$$

Here, $\mathbf{u} = (\rho, \rho v_j, \rho E)$, $j = 1, \dots, d$ is the m -dimensional vector of non-dimensional conserved quantities, and $\mathbf{F}(\mathbf{u}, \mathbf{q})$ are the Navier–Stokes fluxes of dimension $m \times d$

$$\mathbf{F}(\mathbf{u}, \mathbf{q}) = \begin{pmatrix} \rho v_j \\ \rho v_i v_j + \delta_{ij} p \\ v_j (\rho E + p) \end{pmatrix} - \begin{pmatrix} 0 \\ \tau_{ij} \\ v_i \tau_{ij} + f_j \end{pmatrix}. \tag{3}$$

For a Newtonian, calorically perfect gas in thermodynamic equilibrium, the non-dimensional viscous stress tensor, heat flux, and pressure are given by

$$\tau_{ij} = \frac{1}{Re} \left[\left(\frac{\partial v_i}{\partial x_j} + \frac{\partial v_j}{\partial x_i} \right) - \frac{2}{3} \frac{\partial v_k}{\partial x_k} \delta_{ij} \right], \quad f_j = -\frac{\gamma}{Re Pr} \frac{\partial T}{\partial x_j}, \quad p = (\gamma - 1) \rho \left(E - \frac{1}{2} v_k v_k \right), \tag{4}$$

respectively. Here, Re denotes the freestream Reynolds number, Pr the Prandtl number, and γ the specific heat ratio. In particular, $Pr = 0.72$ and $\gamma = 1.4$ for air. The freestream Mach number M_∞ enters into the Navier–Stokes system through the non-dimensional freestream pressure $p_\infty = 1/(\gamma M_\infty^2)$.

2.3. The hybridized DG discretization

The hybridized DG discretization [35] of the governing equations (2) reads as follows: Find $(\mathbf{q}_h(t), \mathbf{u}_h(t), \hat{\mathbf{u}}_h(t)) \in \mathcal{Q}_h^k \times \mathcal{V}_h^k \times \mathcal{M}_h^k$ such that

$$(\mathbf{q}_h, \mathbf{r})_{\mathcal{T}_h} + (\mathbf{u}_h, \nabla \cdot \mathbf{r})_{\mathcal{T}_h} - \langle \hat{\mathbf{u}}_h, \mathbf{r} \cdot \mathbf{n} \rangle_{\partial \mathcal{T}_h} = 0, \tag{5a}$$

$$\left(\frac{\partial \mathbf{u}_h}{\partial t}, \mathbf{w} \right)_{\mathcal{T}_h} - \left(\mathbf{F}(\mathbf{u}_h, \mathbf{q}_h), \nabla \mathbf{w} \right)_{\mathcal{T}_h} + \langle \hat{\mathbf{f}}_h(\hat{\mathbf{u}}_h, \mathbf{u}_h, \mathbf{q}_h), \mathbf{w} \rangle_{\partial \mathcal{T}_h} = 0, \tag{5b}$$

$$\langle \hat{\mathbf{f}}_h(\hat{\mathbf{u}}_h, \mathbf{u}_h, \mathbf{q}_h), \boldsymbol{\mu} \rangle_{\partial \mathcal{T}_h \setminus \partial \Omega} + \langle \hat{\mathbf{b}}_h(\hat{\mathbf{u}}_h, \mathbf{u}_h, \mathbf{q}_h), \boldsymbol{\mu} \rangle_{\partial \Omega} = 0, \tag{5c}$$

for all $(\mathbf{r}, \mathbf{w}, \boldsymbol{\mu}) \in \mathcal{Q}_h^k \times \mathcal{V}_h^k \times \mathcal{M}_h^k$ and all $t \in (0, T)$. Equations (5a)–(5b) weakly impose the Navier–Stokes equations, whereas Eq. (5c) weakly enforces the flux conservation across elements and the boundary conditions. Also, \mathbf{n} is the unit normal vector pointing outwards from the elements and $\hat{\mathbf{b}}_h$ is the boundary flux, whose precise definition depends on the type of boundary condition. In particular, the following approximate non-reflecting boundary condition is imposed at the inlet and outlet sections of the flow field:

$$\hat{\mathbf{b}}_h = \frac{1}{2} \left(\mathbf{A}_n(\hat{\mathbf{u}}_h) + |\mathbf{A}_n(\hat{\mathbf{u}}_h)| \right) \cdot (\hat{\mathbf{u}}_h - \mathbf{u}_h) + \frac{1}{2} \left(\mathbf{A}_n(\hat{\mathbf{u}}_h) - |\mathbf{A}_n(\hat{\mathbf{u}}_h)| \right) \cdot (\hat{\mathbf{u}}_h - \mathbf{u}_b), \tag{6}$$

where \mathbf{u}_b is a given boundary state and $\mathbf{A}_n = \partial(\mathbf{F}_{inv} \cdot \mathbf{n})/\partial \mathbf{u} = \mathbf{X} \mathbf{A}_n \mathbf{X}^{-1}$ denotes the Jacobian of the inviscid flux normal to the boundary. Also, $|\mathbf{A}_n| = \mathbf{X} |\mathbf{A}_n| \mathbf{X}^{-1}$. This corresponds, up to first order, to choosing the incoming Riemann invariants from the boundary state \mathbf{u}_b and extrapolating the outgoing Riemann invariants from the interior of the domain. Also, a non-slip, adiabatic wall boundary condition is used on the airfoil surface, and periodicity is imposed on the remaining boundaries of the domain. Finally, $\hat{\mathbf{f}}_h$ in Eq. (5) denotes the numerical flux defined as

$$\hat{\mathbf{f}}_h(\hat{\mathbf{u}}_h, \mathbf{u}_h, \mathbf{q}_h) = \mathbf{F}(\hat{\mathbf{u}}_h, \mathbf{q}_h) \cdot \mathbf{n} + \mathbf{S}(\hat{\mathbf{u}}_h, \mathbf{u}_h) \cdot (\mathbf{u}_h - \hat{\mathbf{u}}_h). \tag{7}$$

$\mathbf{S} \in \mathbb{R}^{m \times m}$ is the so-called stabilization tensor and plays an important role in the accuracy and stability of the scheme. In particular, the following global Lax–Friedrichs-like approach is used in this work:

$$\mathbf{S}_{ij}(\hat{\mathbf{u}}_h, \mathbf{u}_h) = \mathbf{S}_{ij} = \lambda_{\max}(\mathbf{u}_{in}) \delta_{ij}, \tag{8}$$

where $\lambda_{\max}(\mathbf{u}_{in})$ denotes the maximum eigenvalue of $|\mathbf{A}_n(\mathbf{u}_{in})|$ and \mathbf{u}_{in} is the inflow state.

The semi-discrete system (5) is further discretized in time using L -stable, diagonally implicit Runge–Kutta (DIRK) schemes [1]. The use of L -stable schemes allows for choosing the time-step size based on physical considerations instead of on numerical stability issues.

It remains to analyze the choice of \mathcal{E}_h^E in the definition of the space \mathcal{M}_h^k . Let us consider two simple choices. The first one is $\mathcal{E}_h^E = \emptyset$ and yields $\mathcal{M}_h^k = \widehat{\mathcal{M}}_h^k$. This corresponds to the Hybridizable Discontinuous Galerkin (HDG) method [6,34, 37]. The second choice is $\mathcal{E}_h^E = \mathcal{E}_h$ and implies $\mathcal{M}_h^k = \widetilde{\mathcal{M}}_h^k$. This corresponds to the Embedded Discontinuous Galerkin (EDG) method [6,7,38] and makes the approximation space \mathcal{M}_h^k continuous over \mathcal{E}_h . Hence, the HDG method and the EDG method can be considered as two particular instances of a larger family of methods. In particular, the only difference between HDG and EDG lies in the definition of the approximation space \mathcal{M}_h^k for the trace of the solution. This subtlety is responsible for important differences between both schemes in terms of accuracy, stability, and computational efficiency. We shall refer to this family of methods as *hybridized Discontinuous Galerkin* methods.




Fig. 1. Illustration of the degrees of freedom for the HDG method, the IEDG method, and the EDG method. The blue nodes represent the degrees of freedom of the approximate solution $(\mathbf{u}_h, \mathbf{q}_h)$, while the red nodes represent the degrees of freedom of the approximate trace $\hat{\mathbf{u}}_h$. (For interpretation of the references to color in this figure legend, the reader is referred to the web version of this article.)

Another interesting choice of the approximation space \mathcal{M}_h^k is obtained by setting $\mathcal{E}_h^E = \mathcal{E}_h^I$. This implies $\tilde{\mathcal{M}}_h^k \subset \mathcal{M}_h^k \subset \hat{\mathcal{M}}_h^k$, and the resulting approximation space consists of functions that are continuous everywhere but at the borders of the boundary faces. The resulting method has an HDG flavor on the boundary faces and an EDG flavor on the interior faces, and is therefore referred to as the Interior Embedded DG (IEDG) method [35]. Fig. 1 illustrates the degrees of freedom for the HDG method, the IEDG method, and the EDG method in a four-element mesh. Again, the three schemes differ from each other only in the degrees of freedom of the approximate trace. We note that the IEDG method enjoys advantages of both HDG and EDG. First, IEDG inherits the reduced number of global degrees of freedom, and thus the computational efficiency, of EDG. In fact, the degrees of freedom of the approximate trace on \mathcal{E}_h^B can be locally eliminated without affecting the sparsity pattern of the Jacobian matrix; which yields an even smaller number of global degrees of freedom than in the EDG method. Second, the IEDG scheme enforces the boundary conditions as strongly as the HDG method, hence retaining most of the stability and robustness of HDG. These features make the IEDG method an excellent alternative to the HDG and EDG schemes.

3. Solution method

The temporal discretization of the semi-discrete system (5) using a DIRK scheme yields a nonlinear system of equations at every stage of every time step. In this section, we propose a parallel Newton-GMRES method for the solution of this nonlinear system of equations.

3.1. Nonlinear solver

3.1.1. Newton's method

Let us denote $\mathbf{z}_h = (\mathbf{q}_h, \mathbf{u}_h) \in \mathcal{Q}_h^k \times \mathcal{V}_h^k$. At any given step $n = q(\ell - 1) + s$ of a DIRK(q, r) scheme (q is the number of stages, r is the order of accuracy, ℓ is the current time step, and $s = 1, \dots, q$ is the current stage within the current time step), the nonlinear system of equations that arises from the temporal discretization of the semi-discrete form (5) can be written as

$$\mathbf{R}_{\text{VL}}(\mathbf{z}_h^n, \hat{\mathbf{u}}_h^n) = \mathbf{0}, \quad (9a)$$

$$\mathbf{R}_{\text{FC}}(\mathbf{z}_h^n, \hat{\mathbf{u}}_h^n) = \mathbf{0}, \quad (9b)$$

where \mathbf{R}_{VL} and \mathbf{R}_{FC} are the non-linear residuals associated to (5a)–(5b) and (5c), respectively. The Newton's method is used to solve the nonlinear system (9). Linearizing (9) around the solution $(\mathbf{z}_h^{n,m}, \hat{\mathbf{u}}_h^{n,m})$ at the current Newton iteration $m = 0, 1, \dots$, we arrive at the following linear system

$$\begin{bmatrix} \mathbf{A}^{n,m} & \mathbf{B}^{n,m} \\ \mathbf{C}^{n,m} & \mathbf{D}^{n,m} \end{bmatrix} \begin{pmatrix} \delta \mathbf{z}_h^{n,m} \\ \delta \hat{\mathbf{u}}_h^{n,m} \end{pmatrix} = \begin{pmatrix} \mathbf{F}^{n,m} \\ \mathbf{G}^{n,m} \end{pmatrix}, \quad (10)$$

where $\delta \mathbf{z}_h^{n,m}$ is the vector of degrees of freedom for $\delta \mathbf{z}_h^{n,m} = (\delta \mathbf{q}_h^{n,m}, \delta \mathbf{u}_h^{n,m}) \in \mathcal{Q}_h^k \times \mathcal{V}_h^k$, and $\delta \hat{\mathbf{u}}_h^{n,m}$ is the vector of degrees of freedom for $\delta \hat{\mathbf{u}}_h^{n,m} \in \mathcal{M}_h^k$. The next Newton iterate is defined as $(\mathbf{z}_h^{n,m+1}, \hat{\mathbf{u}}_h^{n,m+1}) := (\mathbf{z}_h^{n,m} + \delta \mathbf{z}_h^{n,m}, \hat{\mathbf{u}}_h^{n,m} + \delta \hat{\mathbf{u}}_h^{n,m})$, and this process is repeated until the norm of the full residual vector $\mathbf{R}_{\text{NS}} := (\mathbf{R}_{\text{VL}}, \mathbf{R}_{\text{FC}})$ is less than a specified tolerance.

3.1.2. Minimal residual algorithm for computing the initial guess

The convergence of Newton's method depends on the initial guess $(\mathbf{z}_h^{n,0}, \hat{\mathbf{u}}_h^{n,0})$. One common initialization strategy is to take the initial guess to be the solution at the previous DIRK step, namely, $(\mathbf{z}_h^{n,0}, \hat{\mathbf{u}}_h^{n,0}) := (\mathbf{z}_h^{n-1}, \hat{\mathbf{u}}_h^{n-1})$. In order to reduce the number of Newton iterations relative to this simple initialization strategy, we propose a minimal residual algorithm to compute a better initial guess. In particular, we express the initial guess as a linear combination of the solutions at the

J previous DIRK steps, namely, $(\mathbf{z}_h^{n,0}, \hat{\mathbf{u}}_h^{n,0}) := \sum_{j=1}^J \alpha_j (\mathbf{z}_h^{n-j}, \hat{\mathbf{u}}_h^{n-j})$. The coefficients α_j are found as the minimizer of the following nonlinear least squares problem

$$(\alpha_1, \dots, \alpha_J) = \arg \min_{(\beta_1, \dots, \beta_J) \in \mathbb{R}^J} \left\| \mathbf{R}_{\text{NS}} \left(\sum_{j=1}^J \beta_j (\mathbf{z}_h^{n-j}, \hat{\mathbf{u}}_h^{n-j}) \right) \right\|^2. \tag{11}$$

This optimization problem is solved by using the Levenberg–Marquardt (LM) algorithm [23,28]. The LM algorithm requires us to compute the gradient vectors $\partial \|\mathbf{R}_{\text{NS}}\| / \partial \beta_j$; which are approximated by finite differences. Since this requires a small number of residual evaluations and each residual evaluation is about two orders of magnitude faster than the Jacobian computation, this minimal residual procedure adds very little to the overall computational cost of the solution method.

3.2. Linear solver

We describe next our approach for solving the linear system (10). Hereinafter, we shall drop the superscripts n, m to simplify the notation.

3.2.1. Static condensation

Due to the discontinuous nature of the finite element spaces \mathcal{Q}_h^k and \mathcal{V}_h^k , the matrix \mathbf{A} has block-diagonal structure. This allows us to inexpensively compute its inverse and eliminate $\delta \mathbf{Z}$ to obtain a smaller linear system in terms of $\delta \hat{\mathbf{U}}$ only as follows:

$$\mathbf{K} \delta \hat{\mathbf{U}} = \mathbf{R}, \tag{12}$$

where $\mathbf{K} = \mathbf{D} - \mathbf{C} \mathbf{A}^{-1} \mathbf{B}$ and $\mathbf{R} = \mathbf{G} - \mathbf{C} \mathbf{A}^{-1} \mathbf{F}$.

This is the global system to be solved at every Newton iteration. One of the advantages of the hybridized DG methods lies in the fact that the size and number of nonzeros of \mathbf{K} are much smaller than those of the global matrix in other DG methods. This is in turn due to the introduction of the approximate trace $\hat{\mathbf{u}}_h$ as an explicit unknown. As a result, the computational cost and memory footprint of the hybridized DG methods can be several times smaller than in other DG methods [35].

3.2.2. Parallel GMRES solver with restricted additive Schwarz preconditioning

The linear system (12) is solved in parallel using the restarted Generalized Minimal Residual GMRES(β) method [43] with iterative classical Gram–Schmidt (ICGS) orthogonalization. A left preconditioner $\mathbf{M}^{-1} \approx \mathbf{K}^{-1}$ is used in order to accelerate GMRES convergence, so that the linear system (12) is replaced by

$$\mathbf{M}^{-1} \mathbf{K} \delta \hat{\mathbf{U}} = \mathbf{M}^{-1} \mathbf{R}. \tag{13}$$

In particular, we employ the restricted additive Schwarz (RAS) method [5] in order to achieve excellent performance of the linear solver. This approach relies on a decomposition of the unknowns in $\delta \hat{\mathbf{U}}$ among parallel workers; which is performed as described below.

Let \hat{P} denote the set of nodes $\{\hat{p}_1, \hat{p}_2, \dots\}$ used to approximate the solution trace $\hat{\mathbf{u}}_h$, i.e., the red dots in Fig. 1. These nodes will be referred to as *traced nodes*. We first partition \hat{P} into N nonempty, nonoverlapping subdomains $\hat{P}_i^{(0)}$, that is,

$$\hat{P} = \bigcup_{i=1}^N \hat{P}_i^{(0)}, \quad \hat{P}_i^{(0)} \cap \hat{P}_j^{(0)} = \emptyset \text{ for } i \neq j, \quad \hat{P}_i^{(0)} \neq \emptyset \text{ for } i = 1, \dots, N. \tag{14}$$

This is the “zero-overlap” decomposition and is computed by minimizing the edge-cut of the \hat{p} -to- \hat{p} connectivity graph through a multilevel k -way partitioning algorithm [25] and the METIS software [26]. The “one-overlap” decomposition is then defined as the collection of supersets of $\hat{P}_i^{(0)}$ obtained by including all the traced nodes immediately neighboring the nodes in $\hat{P}_i^{(0)}$. Using this idea recursively, we define a δ -overlap decomposition $\hat{P}_i^{(\delta)}$ for some nonnegative integer δ . Note that the δ -overlap decomposition depends on which hybridized DG method is used to discretize the Navier–Stokes equations. Fig. 2 depicts an illustration of the one-overlap decomposition for the HDG and IEDG schemes.

The RAS preconditioner is then defined as

$$\mathbf{M}^{-1} := \sum_{i=1}^N \mathbf{R}_i^0 \mathbf{K}_i^{-1} \mathbf{R}_i^\delta, \tag{15}$$

where $\mathbf{K}_i = \mathbf{R}_i^\delta \mathbf{K} \mathbf{R}_i^\delta$ is the so-called subdomain problem and \mathbf{R}_i^δ is the restriction operator onto the subspace associated to the traced nodes in the δ -overlap subdomain number i . We note that the choice $\delta = 0$ leads to a block Jacobi (BJ) preconditioner at the subdomain level. From our experience, $\delta = 1$ provides the best balance between computational cost




Fig. 2. Illustration of the one-overlap decomposition for the HDG and IEDG schemes. The traced nodes are first decomposed into two nonoverlapping subdomains. The one-overlap decomposition is obtained by adding all the traced nodes immediately neighboring nodes in the nonoverlapping subdomains.

per GMRES iteration and number of iterations for the flow regimes considered in this work, and is therefore employed here. This is consistent with the observation in [42].

In practice, we replace \mathbf{K}_i^{-1} by the inverse of the block incomplete LU factorization with zero fill-in, BILU(0), of \mathbf{K}_i , that is, $\mathbf{K}_i^{-1} \approx \tilde{\mathbf{U}}_i^{-1} \tilde{\mathbf{L}}_i^{-1}$. The BILU(0) factorization in each subdomain is performed in conjunction with the Minimum Discarded Fill (MDF) ordering algorithm [39]. Our parallel preconditioner is thus given by

$$\mathbf{M}^{-1} := \sum_{i=1}^N \mathbf{R}_i^0 \tilde{\mathbf{U}}_i^{-1} \tilde{\mathbf{L}}_i^{-1} \mathbf{R}_i^\delta. \quad (16)$$

Each GMRES iteration requires us to compute $\mathbf{y} = \mathbf{M}^{-1} \mathbf{K} \mathbf{r}$, where \mathbf{r} is the current Krylov vector. This is computed in two steps as

$$\mathbf{x} = \mathbf{K} \mathbf{r}, \quad \mathbf{y} = \mathbf{M}^{-1} \mathbf{x}. \quad (17)$$

It remains to describe the implementation of these two operations: Matrix–vector product and preconditioner solve.

Matrix–vector product and preconditioner solve

Fast matrix–vector products and preconditioner solves are required for an efficient solution of the linear system (13). To that end, we exploit computer memory hierarchy and cache policies. First, a block compressed row format is used to store \mathbf{K}_i , as well as its BILU(0) factors $\tilde{\mathbf{L}}_i$ and $\tilde{\mathbf{U}}_i$. Second, in order to take advantage of spatial locality, the blocks are stored in memory in the same order as they will be accessed for the matrix–vector product and the preconditioner solve. For the BILU(0) factors, this is in turn determined by the MDF ordering. This storage strategy increases the hit rate and thus improves the performance of these memory bound operations.

The block matrix–vector products that comprise the full matrix–vector product and preconditioner solve are performed first for the nodes on the interface between subdomains and then for the interior nodes in order to overlap communication and computation as much as possible. These block matrix–vector products are computed through the Basic Linear Algebra Subprograms (BLAS). This way, we exploit the structure of \mathbf{K} and \mathbf{M} while performing the actual computations through highly optimized dense linear algebra libraries.

Mixed-precision algorithm

A mixed-precision algorithm is used for the GMRES iteration. In particular, the matrix–vector products and preconditioner solves are performed in single precision, while the orthogonalization is performed in double precision. Since the former are memory bound operations, a factor of two speedup is achieved. In our experience, the number of GMRES and Newton iterations are largely unaffected by the type of arithmetic thanks to the following factors. First, we solve a suitable non-dimensional form of the Navier–Stokes equations, so that all the non-dimensionalized variables are of order one. Second, the time-step size is chosen small enough to guarantee accuracy and ensure the smooth convergence of the GMRES and Newton iterations even when single precision is used.

Table 1

Details of the computational meshes considered for the NACA 65-(18)10 compressor cascade. Global unknowns indicates the number of unknowns in Eq. (13) and the $\times 5$ factor accounts for the five components in the Navier–Stokes system. The computational cost corresponds to an Intel Xeon E5-2650 architecture.

Mesh No.	k	p	No. elements	Element type	Global unknowns	Cores \times hours
1	2	2	95,400	Hexahedra	$810,600 \times 5$	64×20.3
2	2	2	201,180	Hexahedra	$1,632,828 \times 5$	128×21.6
3	2	2	427,040	Hexahedra	$3,365,784 \times 5$	256×23.0
4	2	2	909,360	Hexahedra	$6,918,290 \times 5$	512×24.5

Stopping criterion and adaptive tolerance

The stopping criterion for the GMRES iteration is based on the relative magnitude of the norm of the residual and the right-hand side in Eq. (13). In particular, the GMRES iteration is performed until a quadratically convergent Newton iterate is obtained, without the need for solving the linear system to machine precision. This is typically achieved by reducing the initial residual by five or six orders of magnitude. Also, the GMRES tolerance is adaptively increased to avoid unnecessary iterations when the nonlinear residual is close to the nonlinear tolerance.

4. NACA 65-(18)10 compressor cascade

The proposed approach is first applied to transitional flows over a NACA 65-(18)10 linear compressor cascade at Reynolds number $Re = 250,000$ and Mach number $M_\infty = 0.081$. The blade solidity is $\sigma = 1.0$, the stagger angle $\xi = 28.3$ deg., and the freestream turbulence intensity $I_\infty = 0.0$. Two operating conditions are considered. First, the relative angle between the freestream and the blade chord is set to 16.7 deg. This corresponds to the *design* condition of the NACA 65-(18)10 compressor cascade at this Reynolds number. Next, the angle of attack is increased to 25.7 deg. for an analysis in *off-design* condition. This is 9.0 deg. larger than in design condition, and will lead to early separation on the suction side and a large laminar separation bubble on the pressure side. Our simulation results will be compared to those of a subgrid-scale LES model discretized using a second-order finite volume code and 31,000,000 elements [30].

4.1. Details of the numerical discretization

The IEDG method and the third-order, three-stage DIRK(3,3) scheme are used for the spatial and temporal discretization, respectively. The polynomial degree of the IEDG approximation is set to $k = 2$ and the non-dimensional time-step to $\Delta t^* = \Delta t \cdot U_\infty / c = 0.005$, where c denotes the blade chord and U_∞ is the freestream velocity. The resulting discretization is therefore third-order accurate in space and time.

The computational domain is partitioned using isoparametric hexahedral elements. In particular, the three-dimensional meshes are generated through extrusion of two-dimensional quadrilateral meshes with extrusion length equal to $0.1c$. Due to the chaotic nature of LES, the integration time is sufficiently large to ensure statistical output convergence. The simulation is considered statistically converged when the L^2 -norm of the covariance field of spanwise- and time-averaged pressure coefficient is below $2c \cdot 0.01$. The variance of an ILES estimator is computed from the instantaneous fields and the Generalized Central Limit Theorem.

The design and off-design conditions have been both simulated on a sequence of finer and finer meshes until the transition location¹ and the pressure coefficient are grid independent and thus *grid converged*. The simulation is considered grid converged when both the L^2 -norm of the difference of spanwise- and time-averaged pressure coefficient between consecutive meshes is below $2c \cdot 0.005$, and the difference in the transition location is below $0.0025c$. This leads to the four computational meshes whose details are summarized in Table 1. An estimate of the computational cost on an Intel Xeon E5-2650 architecture is included as well. Also, Fig. 3 shows a 2D slice of mesh No. 2.

4.2. Grid convergence study and comparison to second-order SGS-LES

Fig. 4 (left) shows the time and spanwise average of the negative pressure coefficient in design condition computed with our high-order ILES and the second-order FV SGS-LES in [30]. The ILES and the SGS-LES results show very good agreement, except on the transition location on the pressure side. This in turn induces a discrepancy on the reattachment location and the subsequent pressure rise, as illustrated in Fig. 4. Similarly, Fig. 5 shows the negative of the spanwise- and time-averaged pressure coefficient in off-design condition. The disagreement on the transition location between the high-order ILES and the second-order SGS-LES is more profound than in design condition. Since the transition location and the pressure coefficient seem grid converged in the high-order ILES results for both operating conditions, these discrepancies can be attributed to (1) unphysical dissipation introduced by the SGS model in the laminar portion of the boundary layer, (2) under-resolution

¹ Since the transition location is not a well-defined location, we use here the end of the separation bubble as indicator of transition.




Fig. 3. 2D slice of mesh No. 2 for the NACA 65-(18)10 compressor cascade.

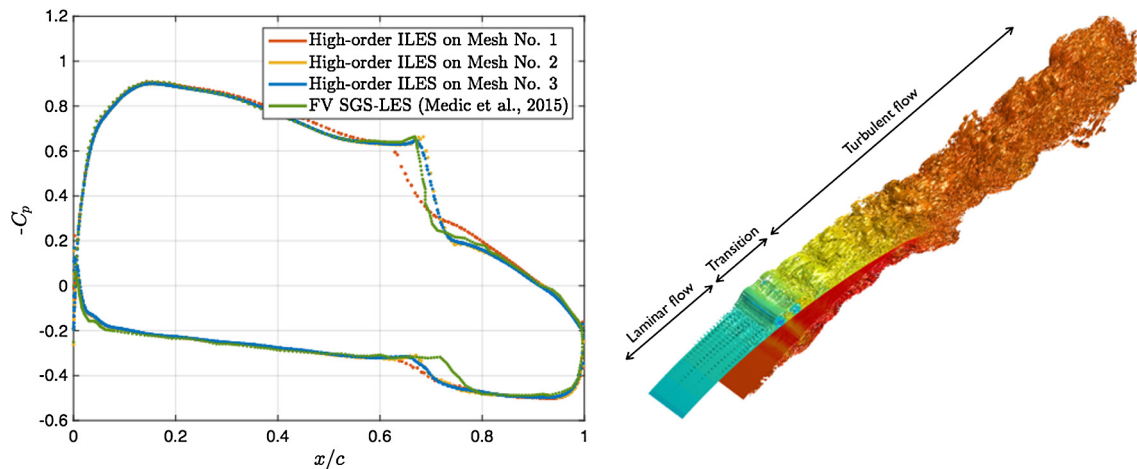


Fig. 4. Negative spanwise- and time-averaged pressure coefficient computed with high-order ILES and second-order FV SGS-LES for NACA 65-(18)10 in design condition. Iso-surface of the Q-criterion [21] colored by pressure is shown on the right.

in the FV SGS-LES results, or (3) other numerical artifacts in either simulation, such as spurious reflections on the inflow and outflow boundaries of the computational domain.

We emphasize that high-order ILES achieves grid convergence on the transition location and the pressure coefficient with two orders of magnitude fewer elements and one order of magnitude fewer degrees of freedom² (DOFs) than those used with the second-order FV SGS-LES. While few grid convergence studies are available in the literature of LES for turbomachinery flows, the spatial resolution required by the high-order ILES is well below the ones typically used with low-order LES approaches. An explanation for high-order ILES being able to accurately capture transition with such few elements and degrees of freedom will be presented in Section 6.

4.3. Boundary layer analysis

We investigate the structure of the boundary layer (BL) and the transition mechanism using the high-order ILES results. The BL analysis is presented first for the design condition and then for the off-design condition. The nomenclature and details of the post-processing strategy are described in Appendix A.

The flow separates on the suction side at $x_s/c = 0.50$ in design condition due to the adverse pressure gradient (see Fig. 4). This produces a laminar separation bubble (LSB) and strongly destabilizes the boundary layer; which eventually transitions to turbulence. After transition, the turbulent mixing leads to rapid reattachment at $x_r/c = 0.71$ and the separation bubble ends. The turbulent BL then remains attached all the way until the trailing edge thanks to the resistance to separation provided by the turbulent mixing. A similar behavior is observed on the pressure side, with BL separation also at $x_s/c = 0.50$ and reattachment at $x_r/c = 0.71$. The laminar separation and turbulent reattachment are illustrated in Fig. 6 through the

² We use the term “degrees of freedom” here as an indicator of the resolution provided by our high-order DG discretization in comparison with a second-order FV method. For example, a “5M DOFs” high-order solution has the same number of effective grid points as a cell-centered finite volume discretization with 5M elements.

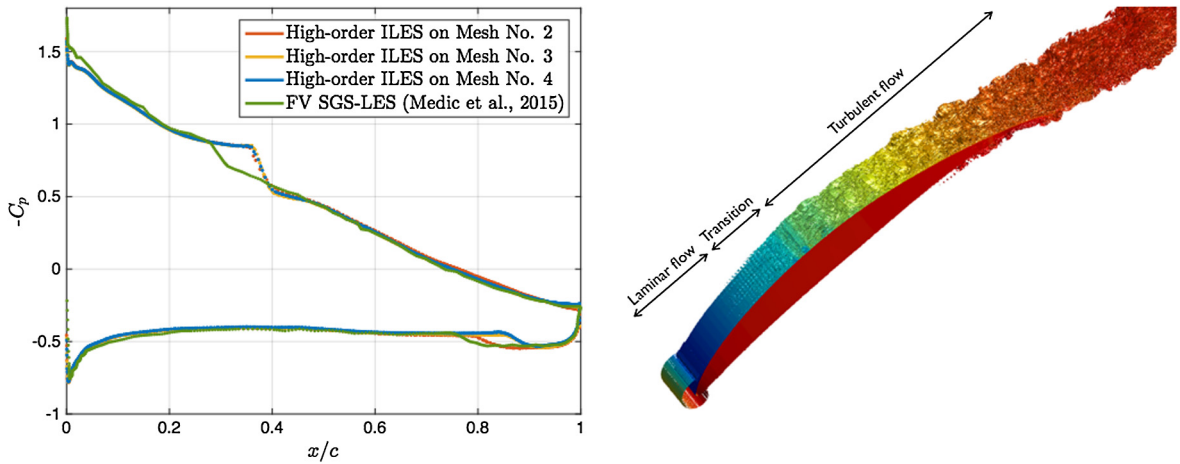


Fig. 5. Negative spanwise- and time-averaged pressure coefficient computed with high-order ILES and second-order FV SGS-LES for NACA 65-(18)10 in off-design condition. Iso-surface of the Q-criterion [21] colored by pressure is shown on the right.




Fig. 6. Instantaneous (left) and time-averaged (right) velocity magnitude fields for NACA 65-(18)10 in design condition.

instantaneous and time-averaged velocity magnitude fields. Similarly, the LSBs translate into the pressure plateaux in Fig. 4, followed by a rapid pressure rise after transition. These plateaux are a consequence of the nearly-still fluid inside the laminar part of the bubble being unable to sustain any significant pressure gradients. The turbulent mixing in the turbulent portion of the bubble, however, can support the strong adverse pressure gradient in Fig. 4 [8].

Also, BL separation leads to an increase in the displacement thickness and thus the shape parameter, as illustrated in Fig. 7 for the suction side. The increase in the shape parameter will be (indirectly) responsible for the rapid growth of instabilities along the bubble, as discussed below. The large value of the displacement thickness along the adverse pressure gradient in the turbulent portion of the bubble is responsible for the so-called *bubble drag* and produces additional losses in the cascade. Once the LSB ends, the displacement thickness and the shape parameter reduce back to attached regime values. Then, the strong pressure gradient from the reattachment location to the trailing edge leads to the final increase in the displacement and momentum thickness in Fig. 7.

Due to the lack of bypass and forced transition mechanisms and the quasi-2D nature of the flow, natural transition through two-dimensional unstable modes is expected. This is numerically confirmed by the high-order ILES results. The two-dimensional nature of transition is illustrated in Fig. 8 through the much larger amplitude of the streamwise instabilities compared to the cross-flow instabilities. In particular, Tollmien–Schlichting (TS) waves form before the boundary layer separates, and Kelvin–Helmholtz (KH) instabilities are ultimately responsible for transition after separation.³ The former are shown in Fig. 9 (left) through the Gaussian-like shape of the root mean square (RMS) in time of the fluctuating velocity magnitude at different BL locations prior to separation. More specifically, the left plot in Fig. 9 shows the superposition of (1) TS waves and (2) pressure waves generated in the turbulent boundary layer of the blade at hand and the neighboring blades. The latter effect is responsible for the nonzero fluctuating velocity outside the boundary layer.

The growth rate of TS waves along the BL is exponential, as shown on the right of Fig. 9 and predicted by linear stability theory. It is worth noting the small magnitude of the instabilities compared to the freestream velocity. This shows why very small amount of numerical dissipation is required for transition prediction. Similarly, very low numerical dispersion

³ We consider TS and KH waves to be different phenomena. In particular, we refer to the unstable modes of the Orr–Sommerfeld equation as TS modes if the boundary layer is attached, and as KH modes if the boundary layer is separated.

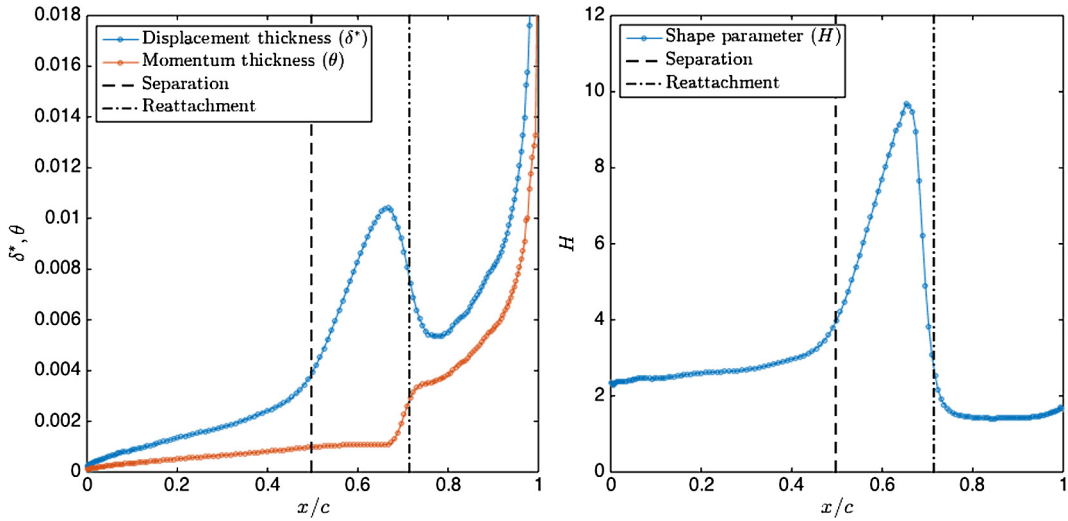


Fig. 7. Streamwise displacement and momentum thickness (left) and shape parameter (right) along the suction side for NACA 65-(18)10 in design condition.

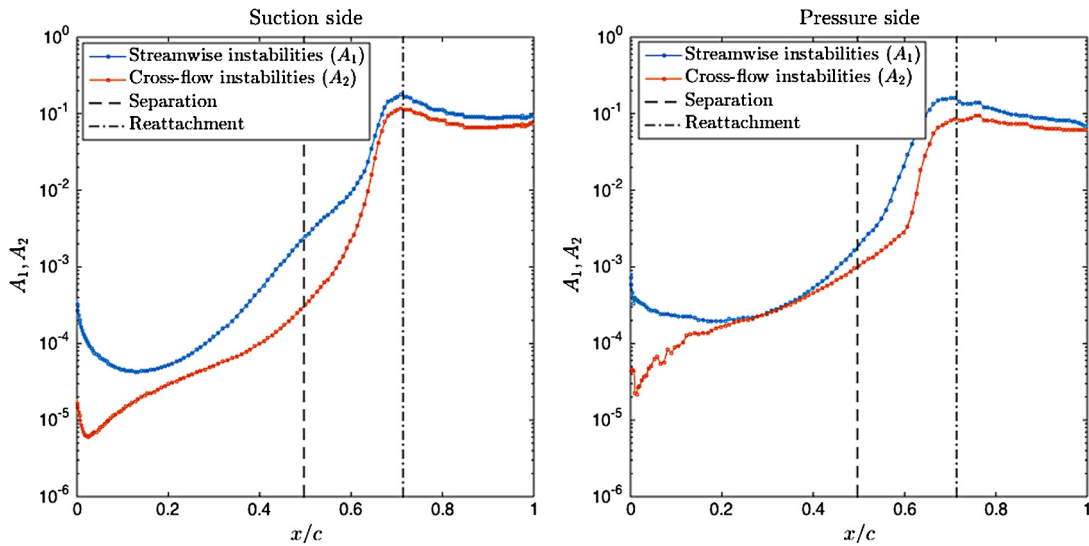


Fig. 8. Amplitude of streamwise and cross-flow instabilities on the suction (left) and pressure (right) sides for NACA 65-(10) in design condition.

is needed to properly resolve all the frequencies present in the transition process. After separation, TS waves turn into KH instabilities as illustrated in Fig. 10. KH modes produce very rapid vortex growth and are ultimately responsible for natural transition in the separated shear layer. Turbulence constitutes a new, “stable” flow regime that prevents instabilities to keep growing exponentially in space and time.

The non-dimensional velocity profile at different locations along the turbulent portion of the suction side BL are displayed in Fig. 11 for the inner (left) and outer (right) layers. High-order ILES properly captures the viscous sublayer $u^+ = n^+$; which extends from $n^+ = 0$ to $n^+ \approx 8$. Also, a log-layer $u^+ = (1/\kappa) \log(n^+) + C^+$ is observed from $n^+ \approx 20$ to $n^+ \approx 200$, where the exact extremes depend on the local Reynolds number and the pressure gradient. For moderate pressure gradients, the numerical results fit well the experimentally measured value for the von Kármán constant $\kappa = 0.40 \pm 0.02$ [2], while smaller values of κ are predicted for strong adverse pressure gradients.

Next, we discuss the structure of the BL and the transition mechanism in off-design condition. Some details will be omitted in this case due to the similarities with the design condition. Again, the boundary layer undergoes laminar separation and turbulent reattachment on both the suction and pressure sides, hence leading to a LSB. The separation and reattachment locations are $x_s/c = 0.20$ and $x_r/c = 0.38$ on the suction side, and $x_s/c = 0.57$ and $x_r/c = 0.89$ on the pressure side. This is illustrated in Fig. 12 through the instantaneous (left) and time-averaged (right) velocity magnitude fields. The existence of separation bubbles is also clear from the pressure plateaux in Fig. 5, as well as from the streamwise displacement thickness and shape parameter in Fig. 13. It is worth noting the large size of the separation bubble on the pressure side; which is typical of off-design operating conditions at low Reynolds numbers. Also, the larger angle of attack with respect

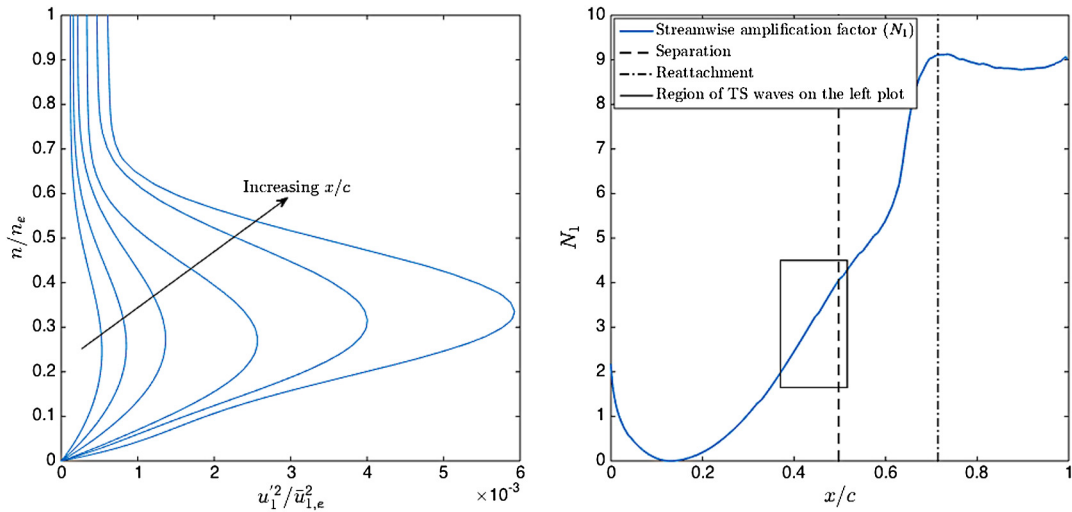


Fig. 9. TS waves (left) and streamwise amplification factor (right) on the suction side for NACA 65-(18)10 in design condition. The box on the right figure indicates the region of the BL in which the TS waves on the left are located.

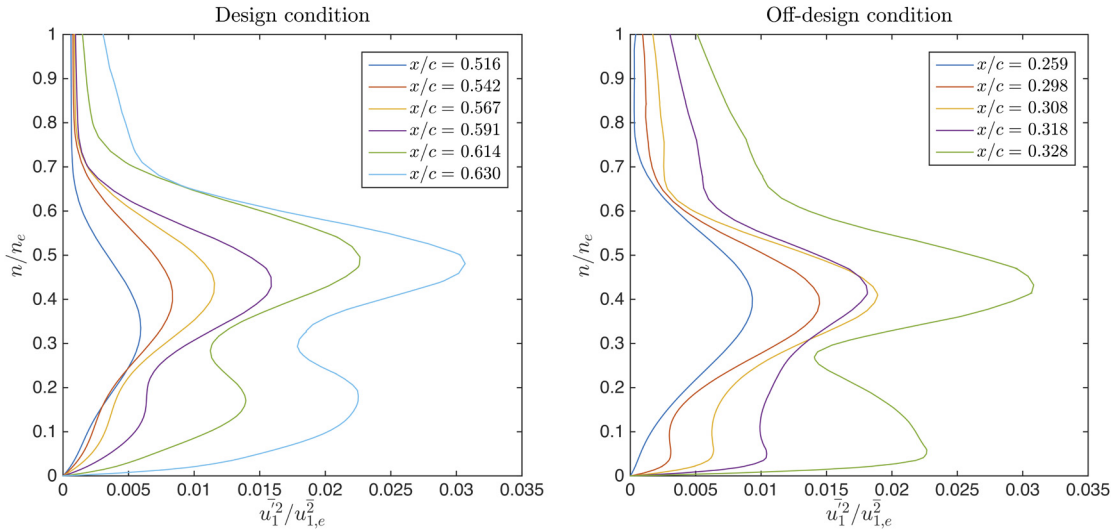


Fig. 10. Transition from TS to KH modes along the separated, suction side boundary layer for NACA 65-(18)10 in design (left) and off-design (right) conditions.

to the design condition moves the LSB downstream on the pressure side and upstream on the suction side. In particular, early separation and transition allow for a more gradual pressure recovery along the suction side, and this translates into a suboptimal tangential force on the blade. Again, TS waves and KH instabilities are responsible for natural transition to turbulence, as illustrated in Fig. 14 for the suction side. Finally, Fig. 15 shows the non-dimensional velocity profiles at different locations along the suction side turbulent BL for the inner (left) and outer (right) layers.

5. Eppler 387 wing

The results for the NACA 65-(18)10 compressor cascade demonstrate the rapid grid convergence of high-order ILES to capture transition. In this section, we examine the accuracy of the high-order ILES transition prediction capability by comparing it to experimental data. In this spirit, we consider transitional flows over the Eppler 387 wing at Reynolds numbers of 100,000, 300,000, and 460,000. The angle of attack is set to $\alpha = 4.0$ deg. and the freestream turbulence intensity to $I_\infty = 0.0$. The Mach number M_∞ is 0.08, 0.09, and 0.13 for the three Reynolds numbers considered, respectively. As it happens, the experimental data for the Eppler wing [29] are arguably the most accurate and reliable in the literature of transitional flows. This justifies the choice of this test case for validation purposes.

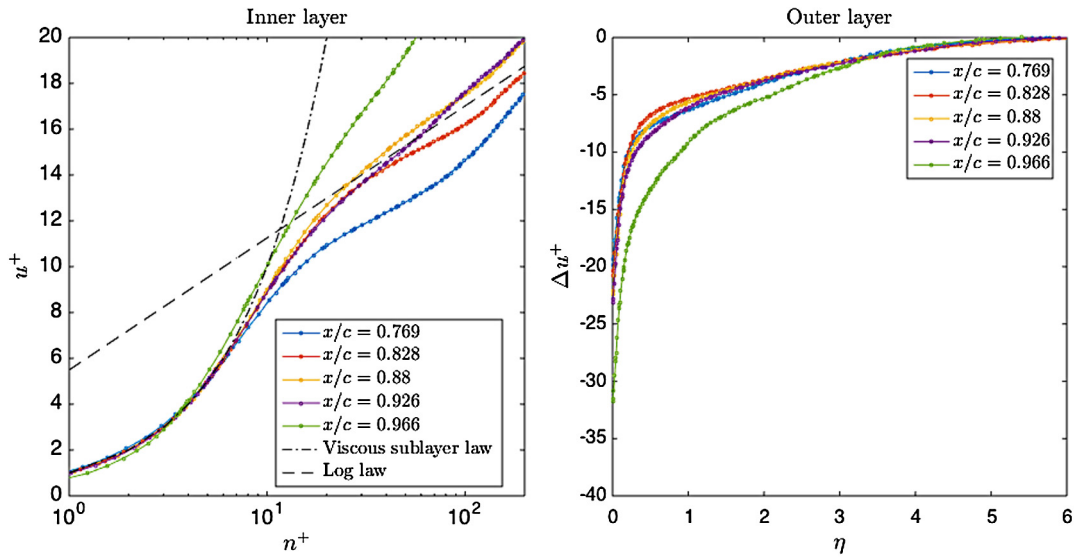


Fig. 11. Non-dimensional turbulent velocity profiles in the inner (left) and outer (right) layers at different locations along the suction side turbulent BL for NACA 65-(18)10 in design condition.




Fig. 12. Instantaneous (left) and time-averaged (right) velocity magnitude fields for NACA 65-(18)10 in off-design condition.

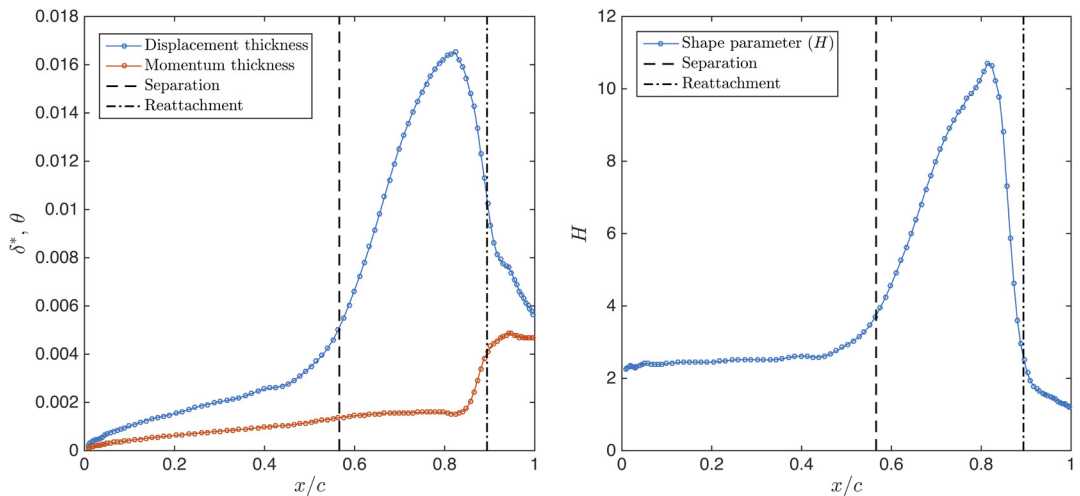


Fig. 13. Streamwise displacement and momentum thickness (left) and shape parameter (right) along the pressure side for NACA 65-(18)10 in off-design condition.

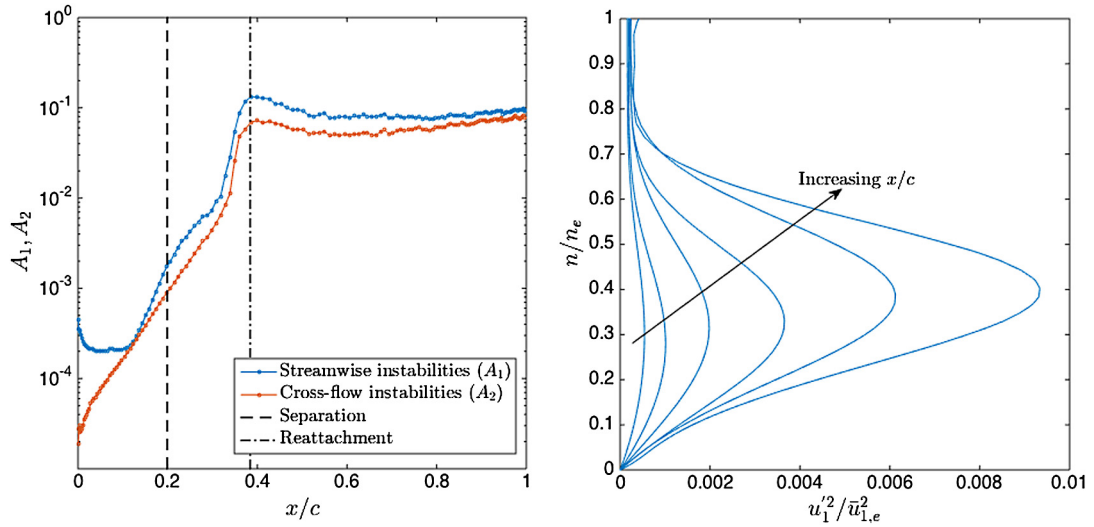


Fig. 14. Amplitude of streamwise and cross-flow instabilities (left) and TS waves (right) on the suction side for NACA 65-(10)18 in off-design condition.




Fig. 15. Non-dimensional turbulent velocity profiles in the inner (left) and outer (right) layers at different locations along the suction side turbulent BL for NACA 65-(18)10 in off-design condition.

5.1. Details of the numerical discretization

The HDG method and the DIRK(3,3) scheme are used for the spatial and temporal discretization of the Navier–Stokes equations, respectively. The polynomial degree of the numerical approximation is $k = 4$, hence yielding a scheme that is fifth-order accurate in space and third-order accurate in time. The computational domain is partitioned using isoparametric tetrahedral elements. The extrusion length in the spanwise direction is $0.1c$ and the computational domain extends about 10 chords away from the wing in the radial direction. Three meshes and non-dimensional time-steps are considered; which correspond to uniform refinement in space and time. The details of these meshes are summarized in Table 2, and Fig. 16 shows a 2D slice of mesh No. 2. The integration time is sufficiently large to ensure statistical convergence of the transition location and the pressure coefficient. Grid convergence and statistical convergence are evaluated based on the same criteria as described in Section 4.1 for the NACA 65-(18)10 compressor cascade.

5.2. Grid convergence study and comparison with experimental data

The negative spanwise- and time-averaged pressure coefficient at Reynolds numbers 100,000, 300,000, and 460,000 are shown in Figs. 17, 18 and 19, respectively. The ILES simulations agree very well with the experimental data reported in [29]. In particular, the simulation results converge to the experimental data as the mesh is refined. More importantly, the error




Fig. 16. 2D slice of mesh No. 2 for the Eppler 387 wing. Tetrahedral meshes are created by subdividing each hexahedron into 6 tetrahedra.

Table 2

Details of the computational meshes considered for the Eppler 387 wing. Global unknowns indicates the number of unknowns in Eq. (13) and the $\times 5$ factor accounts for the five components in the Navier–Stokes system.

Mesh No.	k	p	No. elements	Element type	Global unknowns	$\Delta t \cdot U_\infty / c$
1	4	4	64,800	Tetrahedra	$1,959,600 \times 5$	$7.937\text{E}-3$
2	4	4	126,360	Tetrahedra	$3,814,380 \times 5$	$6.300\text{E}-3$
3	4	4	254,976	Tetrahedra	$7,687,680 \times 5$	$5.000\text{E}-3$

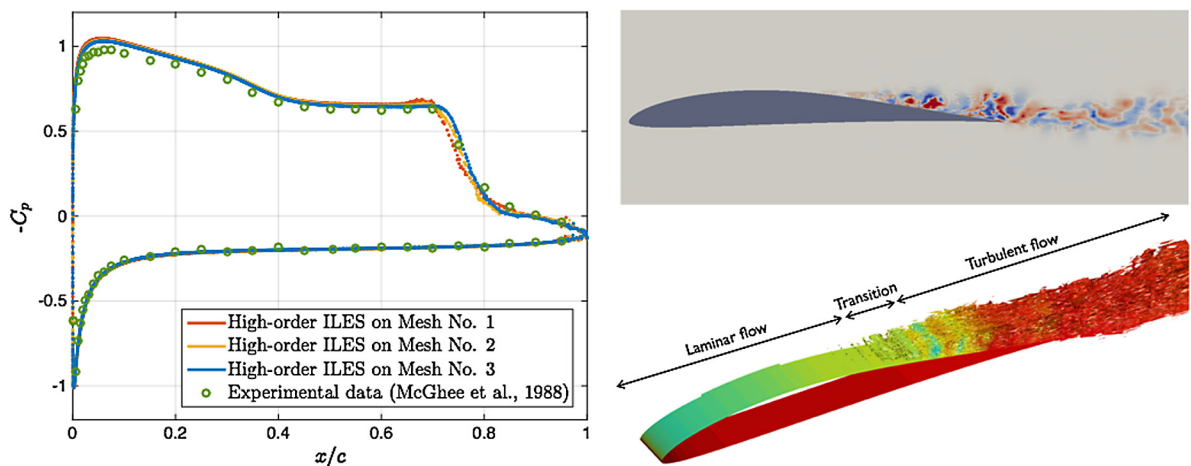


Fig. 17. ILES prediction of the transitional flow over the Eppler 387 wing at $Re = 100,000$: Pressure coefficient (left), instantaneous spanwise velocity (top right), and iso-surface of the Q-criterion colored by pressure (bottom right).

in the transition location is below $0.01c$, $0.005c$, and $0.01c$ at Reynolds number 100,000, 300,000, and 460,000, respectively, even with mesh No. 1. The effective resolution of this mesh is equivalent to a cell-centered finite volume discretization with 691,200 elements. Again, these numbers are much below those typically needed with a low-order SGS-LES approach. Also, the transition from laminar to turbulent regime can be more accurately predicted at Reynolds number of 300,000 than 100,000. As discussed in the next section, this is due to the differences in the transition process between both Reynolds numbers.

The missing vortex upwash due to the finite extent of the computational domain leads to a change in the effective angle of attack. This affects the pressure field around the wing and is responsible for the mismatch near the leading edge with respect to the experimental data. This motivates us to extend the computational domain 50 chords away from the wing, and compare the results to those with a 10-chord domain and the experimental data. In particular, the 50-chord mesh is an extension of the 10-chord mesh No. 3 using the same geometric growth rate for the element size along the radial direction as in the first 10 chords. The results at Reynolds number 100,000 are shown in Fig. 20. The change in the effective angle of attack is inversely proportional to the radius of the computational domain and, in particular, a 50-chord domain leads to $\Delta\alpha = 0.0709$ deg. Also, the wind tunnel walls in [29] were located 13 chords away from the wing; which makes the experimental data susceptible to small inaccuracies as well. Overall, the agreement between the experimental data and the

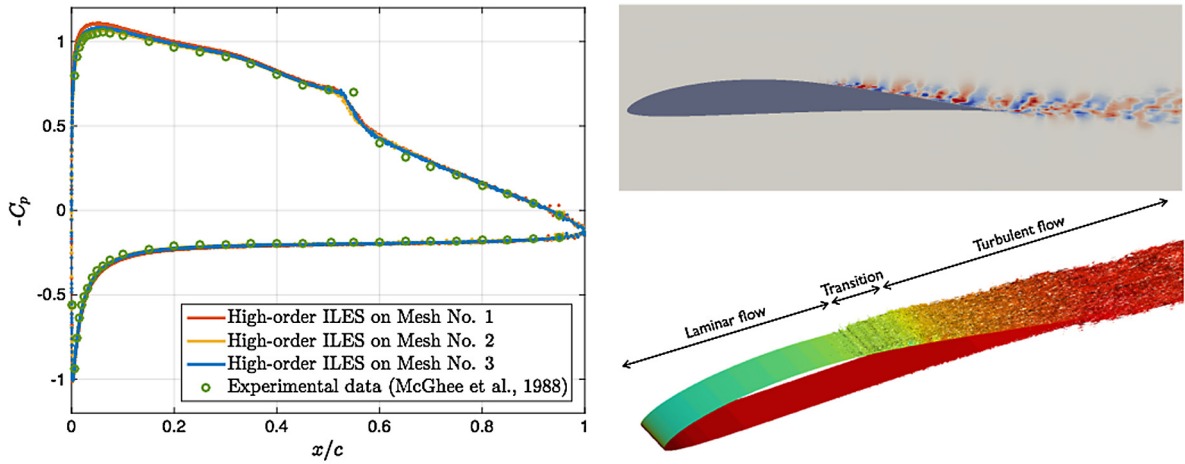


Fig. 18. ILES prediction of the transitional flow over the Eppler 387 wing at $Re = 300,000$: Pressure coefficient (left), instantaneous spanwise velocity (top right), and iso-surface of the Q-criterion colored by pressure (bottom right).

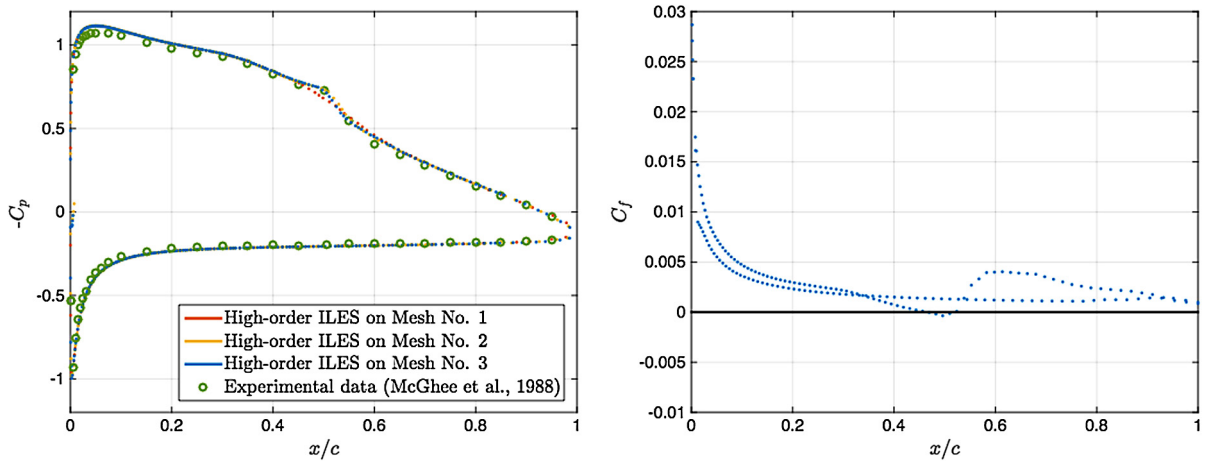


Fig. 19. ILES prediction of the pressure coefficient (left) and skin friction coefficient (right) for the Eppler 387 wing at $Re = 460,000$.

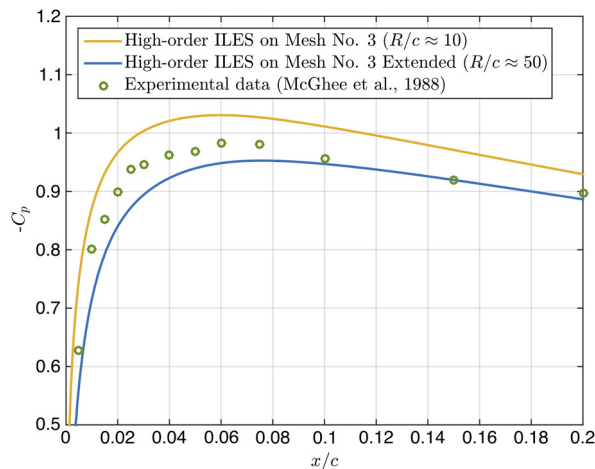


Fig. 20. Pressure coefficient near the leading edge for different sizes of the computational domain for the Eppler 387 wing at $Re = 100,000$.




Fig. 21. Amplitude of streamwise and cross-flow instabilities along the suction (left) and pressure (right) sides for the Eppler 387 wing at Reynolds number 100,000.

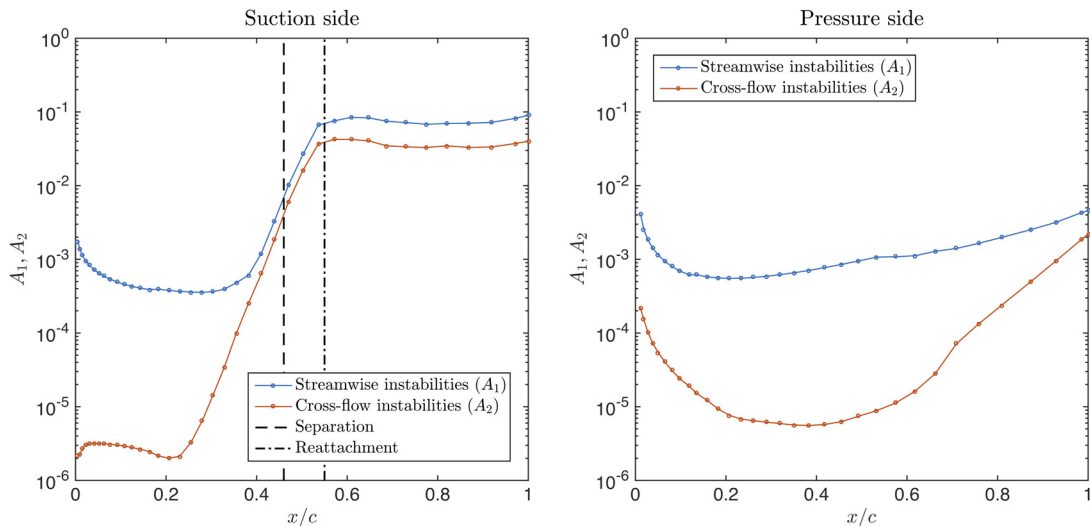


Fig. 22. Amplitude of streamwise and cross-flow instabilities along the suction (left) and pressure (right) sides for the Eppler 387 wing at Reynolds number 300,000.

ILES results is very good, and the minor discrepancies are within expected due to experimental errors, the finite extent of the physical and computational domains, etc.

5.3. Boundary layer analysis

The boundary layer remains attached along the pressure side for the three Reynolds numbers considered. In particular, the design of the Eppler 387 wing is such that the adverse pressure gradient on the lower side takes place only in a very small region near the leading edge, where the momentum thickness is small and the boundary layer therefore very resistant to separation. The greater stability of the attached BL translates into a fully laminar flow along the pressure side despite the relatively large value of the Reynolds number. In fact, the amplitude of the TS waves remains at least two orders of magnitude below the freestream velocity, as shown on the right of Figs. 21 and 22.

The suction side boundary layer, however, undergoes laminar separation and turbulent reattachment for the three Reynolds numbers. In particular, the separation bubble extends from $x_s/c = 0.38$ to $x_r/c = 0.80$ at $Re = 100,000$, from $x_s/c = 0.46$ to $x_r/c = 0.55$ at $Re = 300,000$, and from $x_s/c = 0.46$ to $x_r/c = 0.52$ at $Re = 460,000$. This is reflected in the pressure plateaux in Figs. 17, 18 and 19, the average velocity magnitude fields in Fig. 23, and the increase in the displacement thickness and the shape parameter in Figs. 24 and 25. Also, it results in rapid amplification of KH instabilities in the




Fig. 23. Spanwise- and time-averaged velocity magnitude field for the Eppler 387 wing at Reynolds number 100,000 (top), 300,000 (middle) and 460,000 (bottom).

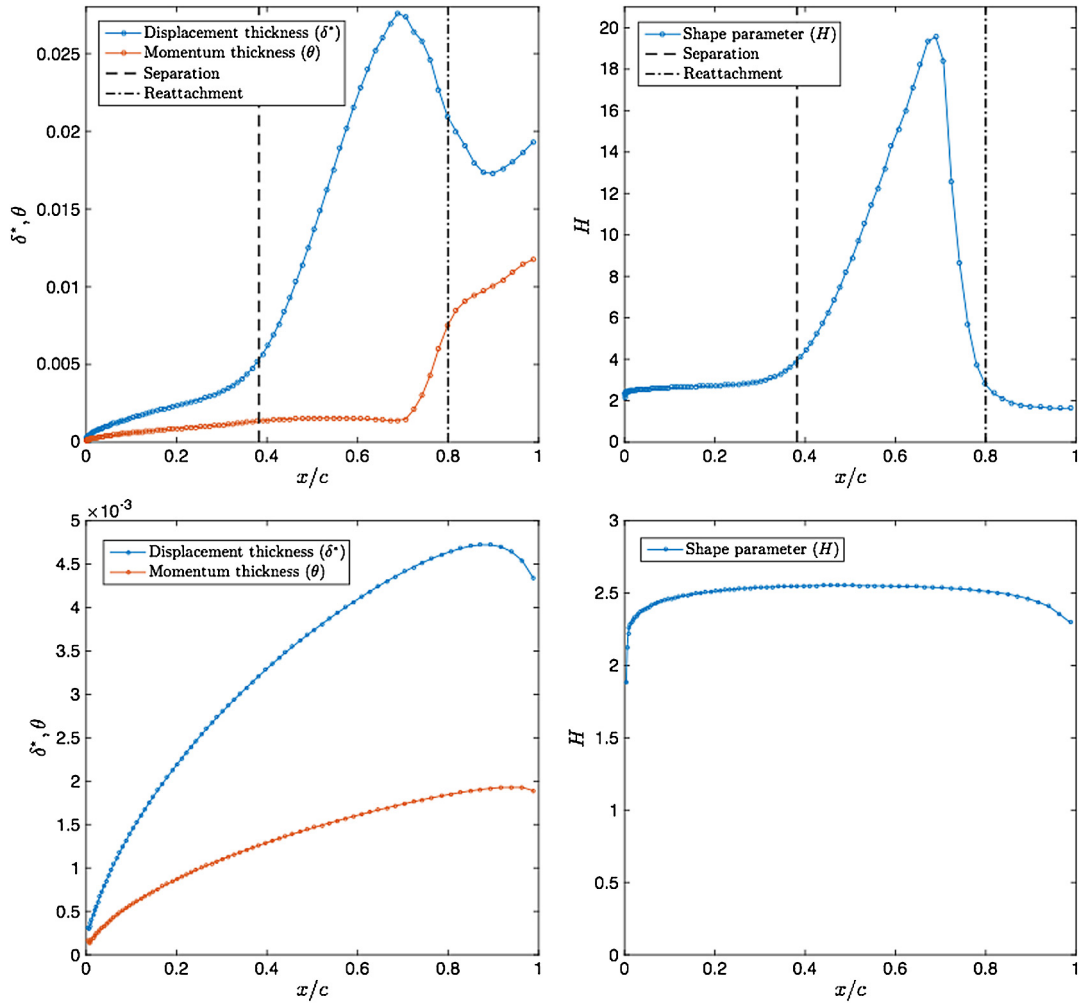


Fig. 24. Streamwise displacement and momentum thickness (left) and shape parameter (right) along the suction (top) and pressure (bottom) sides for the Eppler 387 wing at Reynolds number 100,000.

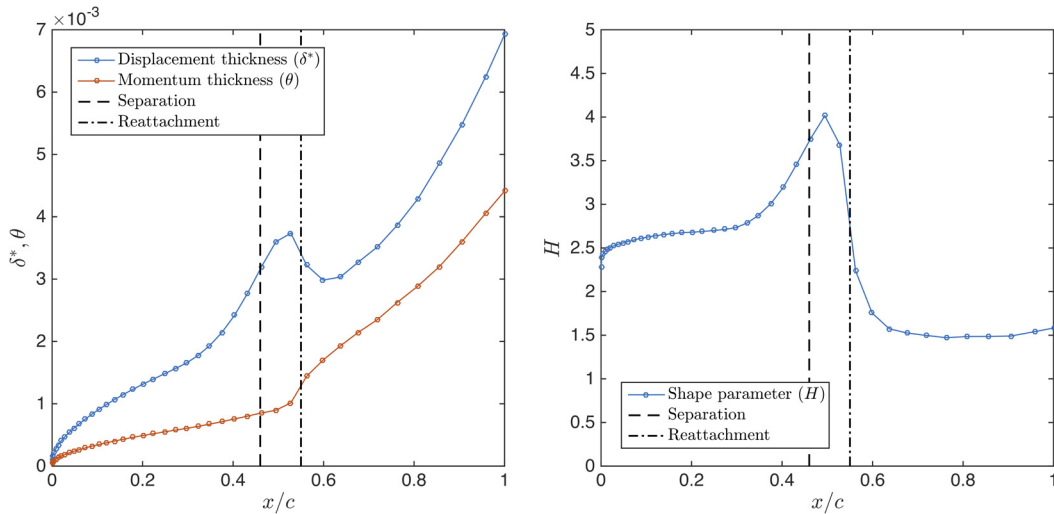


Fig. 25. Streamwise displacement and momentum thickness (left) and shape parameter (right) along the suction side for the Eppler 387 wing at Reynolds number 300,000.

separated shear layer and eventual transition to turbulence, as shown on the left of Figs. 21 and 22. Since the process is analogous to that described for the NACA 65-(18)10 compressor blade, the details are omitted here.

At Reynolds numbers of 300,000 and 460,000, the amplitude of the BL instabilities at the separation location is only one order of magnitude below the transition threshold, hence leading to rapid transition and reattachment. The small size of the LSB has important consequences from a physical and numerical point of view. From a physical perspective, it reduces the bubble drag and thus the overall drag coefficient. From a numerical perspective, it makes transition *easier* to predict. This is in turn due to the transition process taking place under smaller values of the shape parameter H . Indeed, although the length and time scales of the BL instabilities scale as $Re^{-1/2}$ for a fixed value of H , this may be leveraged by a smaller value of the shape parameter itself and translate into the most unstable BL modes having larger length and time scales. This is the case at $Re = 300,000$ vs. 100,000, and demonstrates that the transition mechanism plays a key role in the resolution needed to capture transition and that resolution estimates solely based on dimensional analysis are likely to fail for the transition problem.

We finally note that the separation location coincides at Reynolds numbers of 300,000 and 460,000. This is consistent with the small separation bubble having a negligible effect on the potential flow field around the wing, as illustrated in Fig. 23, and the Reynolds number Independence of laminar boundary layer flows [8]. As such, the displacement and momentum thickness (upon scaling by $Re^{1/2}$) and the shape parameter at $Re = 460,000$ look indistinguishable from those at $Re = 300,000$, and for this reason have been omitted here.

6. Discussion and conclusions

We conclude the paper with a discussion on the rationale for the success of the high-order ILES approach. Regarding high- vs. low-order, two main advantages could be argued in favor of high-order methods for transition prediction. First, they may be able to capture transition in situations in which low-order schemes fail to give accurate results. Second, for some intended accuracy, they could provide a numerical solution in less computational time. These two points are actually related. Low-order schemes can always be accurate enough if the spatial and temporal resolutions are *sufficiently fine*. However, this may render them so computationally expensive that they become less competitive than their high-order counterparts. As discussed before, transition is a very subtle process and high-order accuracy is key to capture the delicate physical phenomena that take place during transition from laminar to turbulent regime.

Indeed, our numerical results showed that high-order methods require significantly fewer elements and degrees of freedom than low-order methods by one or two orders of magnitude. This is justified by the following observation: *Simulating transition is challenging mostly due to the small magnitude of the instabilities involved, rather than due to their length and time scales.* A low-order scheme may kill the small instabilities because of high numerical dissipation even when the mesh size and time-step size are sufficiently small to represent the length and time scales of the instabilities. The number of degrees of freedom to capture transition can therefore be significantly reduced by increasing the accuracy order of the numerical scheme.

It is known that high-order methods become less and less efficient as the order of accuracy increases beyond a certain value. Hence, there will be an optimal range for the order of accuracy to yield the best trade-off between accuracy and efficiency. Through our own experience, the hybridized DG methods reach their best performance for transitional flows when the polynomial degree k is in the range between 2 and 4; that is, when the order of accuracy is between 3 and 5.

Numerical results presented in this paper indicate that this range is adequate to accurately and efficiently resolve the transition process at Reynolds number up to 460,000.

The ILES approach for transition prediction, as opposed to explicit SGS modeling, is justified by the following observation: *While some dissipation is required to account for the unresolved scales in a turbulent flow, this is not the case along the laminar portion of the boundary layer in which the transition process takes place.* As such, directly removing spurious diffusion coming from a SGS model, as ILES does, turns out to be a better approach. Whether ILES is also preferred for the turbulent portion of the boundary layer is the subject of ongoing research.

In summary, this work demonstrates the potential of high-order ILES as a promising approach to simulate transitional turbulent flows. This was illustrated through numerical results and supported by theoretical considerations. Future work will focus on the extension of the approach to separated and fully turbulent flows at high Reynolds numbers.

Acknowledgements

The first author would like to thank “la Caixa” Foundation for the Graduate Studies Fellowship that support his work. We also gratefully acknowledge Pratt & Whitney and the Air Force Office of Scientific Research (FA9550-16-1-0214) for partially supporting this effort. Finally, we would like to thank Prof. M. Drela, Dr. M. Sadeghi and H.-M. Shang for their useful comments and suggestions, as well as Dr. C. Hill for providing with the computational resources to perform some of the simulations presented in this paper.

Appendix A. Boundary layer post-processing

This appendix introduces the nomenclature and methodology employed for the boundary layer (BL) post-processing in Sections 4 and 5. The BL analysis is based on a pseudo-velocity defined as

$$\mathbf{u}^*(\mathbf{s}, n) := \int_0^n (\boldsymbol{\omega} \times \hat{\mathbf{n}}) dn', \tag{A.1}$$

where $\boldsymbol{\omega}$ denotes vorticity, and (\mathbf{s}, n) is the set of curvilinear coordinates associated to the airfoil surface. In particular, $\mathbf{s} = (s_1, s_2)$ and n are the coordinates along the streamwise, cross-flow, and outward normal to the airfoil directions, respectively. Also, the unit vectors associated to these coordinates are denoted by $\hat{\mathbf{s}}_1$, $\hat{\mathbf{s}}_2$ and $\hat{\mathbf{n}}$, respectively. Unlike the actual velocity, this pseudo-velocity asymptotes to a constant outside the boundary layer, even with strong curvature, thus making the edge of the boundary layer a well-defined location [48].

In particular, the BL edge n_e is computed as the first location along n -direction simultaneously satisfying

$$\|\bar{\boldsymbol{\omega}}\| n < \epsilon_1 \|\bar{\mathbf{u}}^*\|, \quad \left\| \frac{\partial \bar{\boldsymbol{\omega}}}{\partial n} \right\| n^2 < \epsilon_2 \|\bar{\mathbf{u}}^*\|, \tag{A.2}$$

where $\epsilon_1 = 0.01$ and $\epsilon_2 = 0.1$ are some properly tuned constants for a systematic and robust detection of the BL edge, and the overbar denotes temporal and cross-flow averaging, i.e.,

$$\bar{\boldsymbol{\omega}}(s_1, n) := \frac{1}{T \cdot \Delta s_2} \int_0^T \int_0^{\Delta s_2} \boldsymbol{\omega}(\mathbf{s}, n, t) ds_2 dt. \tag{A.3}$$

Cross-flow averaging corresponds to ensemble averaging due to the quasi-2D nature of the geometry and boundary conditions in this work, and is therefore used to accelerate the convergence of the statistics of the turbulent flow.

The local streamwise and cross-flow unit vectors are defined as

$$\hat{\mathbf{s}}_1(s_1) := \bar{\mathbf{u}}_e / \bar{u}_e, \quad \hat{\mathbf{s}}_2(s_1) := \hat{\mathbf{s}}_1 \times \hat{\mathbf{n}}, \tag{A.4}$$

where $\mathbf{u}_e = \mathbf{u}^*(n_e)$ is the pseudo-velocity at the edge of the boundary layer and $u_e = \|\mathbf{u}_e\|_2$ is its magnitude. Also, the average \bar{u}_1 and fluctuating u'_1 streamwise velocities are given by

$$\bar{u}_1(s_1, n) = \bar{\mathbf{u}}^*(s_1, n) \cdot \hat{\mathbf{s}}_1(s_1, n), \quad u'_1(\mathbf{s}, n, t) := u_1(\mathbf{s}, n, t) - \bar{u}_1(s_1, n), \tag{A.5}$$

whereas the streamwise displacement thickness, momentum thickness, and shape parameter read as

$$\begin{aligned}\delta^*(s_1) &:= \int_0^{n_e} \left(1 - \frac{\bar{u}_1}{u_e}\right) dn, \\ \theta(s_1) &:= \int_0^{n_e} \left(1 - \frac{\bar{u}_1}{u_e}\right) \frac{\bar{u}_1}{u_e} dn, \\ H(s_1) &:= \frac{\delta^*}{\theta}.\end{aligned}\tag{A.6}$$

The amplitude A_1 and amplification factor N_1 of streamwise perturbations at the boundary layer location s_1 are given by

$$A_1(s_1) = \frac{1}{u_e(s_1)\sqrt{n_e(s_1)}} \sqrt{\int_0^{n_e} \frac{1}{u_1'^2} dn}, \quad N_1(s_1) := \ln\left(\frac{A_1(s_1)}{A_{1,0}}\right),\tag{A.7}$$

where $A_{1,0}$ is some reference amplitude. We note that $A_{1,0}$ shifts $N_1(s_1)$ by a constant factor but it does not affect its growth rate. The cross-flow version of the previous quantities are defined in an analogous manner.

Finally, the non-dimensional velocity u^+ and distance to the wall n^+ in the inner layer of the turbulent BL read

$$u^+ := \frac{\bar{u}_1}{u_\tau}, \quad n^+ := \frac{n}{l_\tau},\tag{A.8}$$

where $u_\tau = \sqrt{\tau_w/\rho}$ denotes the shear velocity and $l_\tau = \nu/u_\tau$ is the wall unit length. Similarly, the non-dimensional velocity defect Δu^+ and distance to the wall η in the outer layer are given by

$$\Delta u^+ := \frac{\bar{u}_1 - u_e}{u_\tau}, \quad \eta := \frac{n}{\delta^*}.\tag{A.9}$$

These BL quantities are necessary to analyze the structure of the boundary layer and the transition mechanism, as discussed in Sections 4 and 5.

References

- [1] R. Alexander, Diagonally implicit Runge–Kutta methods for stiff ODEs, *SIAM J. Numer. Anal.* 14 (6) (1977) 1006–1021.
- [2] S.C.C. Bailey, M. Vallikivi, M. Hultmark, A.J. Smits, Estimating the value of von Kármán's constant in turbulent pipe flow, *J. Fluid Mech.* 749 (2014) 79–98.
- [3] J.P. Boris, On Large Eddy Simulation using subgrid turbulence models, in: J.L. Lumley (Ed.), *Whither Turbulence: Turbulence at the Crossroads*, Springer, New York, 1990, pp. 344–353.
- [4] J.P. Boris, F.F. Grinstein, E.S. Oran, R.L. Kolbe, New insights into large eddy simulation, *Fluid Dyn. Res.* 10 (4) (1992) 199–228.
- [5] X.C. Cai, M. Sarkis, A restricted additive Schwarz preconditioner for general sparse linear systems, *SIAM J. Sci. Comput.* 21 (2) (1999) 792–797.
- [6] B. Cockburn, J. Gopalakrishnan, R. Lazarov, Unified hybridization of discontinuous Galerkin, mixed and continuous Galerkin methods for second order elliptic problems, *SIAM J. Numer. Anal.* 47 (2) (2009) 1319–1365.
- [7] B. Cockburn, J. Guzman, S.C. Soon, H.K. Stolarski, An analysis of the embedded discontinuous Galerkin method for second-order elliptic problems, *SIAM J. Numer. Anal.* 47 (4) (2009) 2686–2707.
- [8] M. Drela, *Flight Vehicle Aerodynamics*, The MIT Press, 2014.
- [9] D. Drikakis, H. Marco, F.F. Grinstein, C.R. DeVore, C. Fureby, M. Liefvendahl, D.L. Youngs, Numerics for ILES: limiting algorithms, in: F.F. Grinstein, L.G. Margolin, W.J. Rider (Eds.), *Implicit Large-Eddy Simulation: Computing Turbulent Flow Dynamics*, Cambridge University Press, New York, 2007, pp. 94–129.
- [10] P. Fernandez, N.C. Nguyen, X. Roca, J. Peraire, Implicit large-eddy simulation of compressible flows using the Interior Embedded Discontinuous Galerkin method, in: 54th AIAA Aerospace Sciences Meeting, San Diego, USA, 2016, <http://dx.doi.org/10.2514/6.2016-1332>.
- [11] P. Fernandez, High-Order Implicit-Large Eddy Simulation for Transitional Aerodynamic Flows, Master Thesis, Department of Aeronautics and Astronautics, Massachusetts Institute of Technology, 2016, <http://hdl.handle.net/1721.1/105613>.
- [12] A. Frere, K. Hillewaert, H. Sarlak, R.F. Mikkelsen, Cross-validation of numerical and experimental studies of transitional airfoil performance, in: 33rd ASME Wind Energy Symposium, Kissimmee, USA, 2015.
- [13] C. Fureby, G. Taylor, H.G. Weller, A.D. Gosman, A comparative study of subgrid scale models in homogeneous isotropic turbulence, *Phys. Fluids* 9 (5) (1997) 1416–1429.
- [14] C. Fureby, F.F. Grinstein, Monotonically integrated Large Eddy Simulation of free shear flows, *AIAA J.* 37 (5) (1999) 544–556.
- [15] C. Fureby, F.F. Grinstein, Large Eddy Simulation of high-Reynolds-number free and wall-bounded flows, *J. Comput. Phys.* 181 (1) (2002) 68–97.
- [16] C. Fureby, N. Alin, N. Wikstrom, S. Menon, N. Svanstedt, L. Persson, Large-Eddy simulation of high-Reynolds-number wall-bounded flows, *AIAA J.* 42 (3) (2004) 457–468.
- [17] M.C. Galbraith, M.R. Visbal, Implicit Large-Eddy Simulation of low Reynolds number flow past the SD7003 airfoil, in: 46th AIAA Aerospace Sciences Meeting and Exhibit, Reno, USA, 2008.
- [18] G.J. Gassner, A.D. Beck, On the accuracy of high-order discretizations for underresolved turbulence simulations, *Theor. Comput. Fluid Dyn.* 27 (3) (2013) 221–237.
- [19] F.F. Grinstein, C.R. DeVore, Dynamics of coherent structures and transition to turbulence in free square jets, *Phys. Fluids* 8 (5) (1996) 1237–1251.

- [20] F.F. Grinstein, L.G. Margolin, W.J. Rider, A rationale for implicit LES, in: F.F. Grinstein, L.G. Margolin, W.J. Rider (Eds.), *Implicit Large-Eddy Simulation: Computing Turbulent Flow Dynamics*, Cambridge University Press, New York, 2007.
- [21] J. Jeong, F. Hussain, On the identification of a vortex, *J. Fluid Mech.* 285 (1995) 69–94.
- [22] M. Lesieur, O. Metais, New trends in Large-Eddy Simulations of turbulence, *Annu. Rev. Fluid Mech.* 28 (1996) 45–82.
- [23] K. Levenberg, A method for the solution of certain non-linear problems in least squares, *Q. Appl. Math.* 2 (1944) 164–168.
- [24] J.E.W. Lombard, D. Moxey, S.J. Sherwin, J.F.A. Hoessler, S. Dhandapani, M.J. Taylor, Implicit Large-Eddy Simulation of a wingtip vortex, *AIAA J.* 54 (2) (2016) 506–518.
- [25] G. Karypis, V. Kumar, Multilevel k -way partitioning scheme for irregular graphs, *J. Parallel Distrib. Comput.* 48 (1) (1998) 96–129.
- [26] G. Karypis METIS, A Software Package for Partitioning Unstructured Graphs, Partitioning Meshes, and Computing Fill-Reducing Orderings of Sparse Matrices, Version 5.1.0, Department of Computer Science & Engineering, University of Minnesota, 2013.
- [27] L.G. Margolin, W.J. Rider, Numerical regularization: the numerical analysis of implicit subgrid models, in: F.F. Grinstein, L.G. Margolin, W.J. Rider (Eds.), *Implicit Large-Eddy Simulation: Computing Turbulent Flow Dynamics*, Cambridge University Press, New York, 2007, pp. 195–221.
- [28] D.W. Marquardt, An algorithm for least-squares estimation of nonlinear parameters, *J. Soc. Ind. Appl. Math.* 11 (2) (1963) 431–441.
- [29] R.J. McGhee, B.S. Walker, B.F. Millard, Experimental Results for the Eppler 387 Airfoil at Low Reynolds Numbers in the Langley Low-Turbulence Pressure Tunnel, Langley Research Center, NASA Technical Memorandum 4062, 1988.
- [30] G. Medic, V. Zhang, G. Wang, G. Tang, J. Joo, O. Sharma, Prediction of transition and losses in compressor cascades using Large-Eddy-Simulation, *J. Turbomach.* 138 (12) (2016) 121001.
- [31] C. Meneveau, J. Katz, Scale-invariance and turbulence models for Large-Eddy Simulation, *Annu. Rev. Fluid Mech.* 32 (2000) 1–32.
- [32] R.C. Moura, S.J. Sherwin, J. Peiró, Linear dispersion–diffusion analysis and its application to under-resolved turbulence simulations using discontinuous Galerkin spectral/hp methods, *J. Comput. Phys.* 298 (2015) 695–710.
- [33] S.M. Murman, L.T. Diosady, A. Garai, M. Ceze, A space–time discontinuous–Galerkin approach for separated flows, in: 54th AIAA Aerospace Sciences Meeting, San Diego, USA, 2016.
- [34] N.C. Nguyen, J. Peraire, Hybridizable discontinuous Galerkin methods for partial differential equations in continuum mechanics, *J. Comput. Phys.* 231 (18) (2012) 5955–5988.
- [35] N.C. Nguyen, J. Peraire, B. Cockburn, A class of embedded discontinuous Galerkin methods for computational fluid dynamics, *J. Comput. Phys.* 302 (1) (2015) 674–692.
- [36] J.S. Park, F.D. Witherden, P.E. Vincent, High-order accurate Implicit Large Eddy Simulations of flow over a NACA0021 aerofoil in deep stall, *AIAA J.* (2017), submitted for publication.
- [37] J. Peraire, N.C. Nguyen, B. Cockburn, A hybridizable discontinuous Galerkin method for the compressible Euler and Navier–Stokes equations, in: *AIAA Conference*, Orlando, USA, 2010.
- [38] J. Peraire, N.C. Nguyen, B. Cockburn, An embedded discontinuous Galerkin method for the compressible Euler and Navier–Stokes equations, in: 20th AIAA Computational Fluid Dynamics Conference, Honolulu, USA, 2011.
- [39] P.-O. Persson, J. Peraire, Newton-GMRES preconditioning for discontinuous Galerkin discretizations of the Navier–Stokes equations, *SIAM J. Sci. Comput.* 30 (6) (2008) 2709–2733.
- [40] F. Renac, M. de la Llave Plata, E. Martin, J.-B. Chapelier, V. Couaillier, Aghora: a high-order DG solver for turbulent flow simulations, in: *IDIHOM: Industrialization of High-Order Methods – A Top-Down Approach*, in: Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 128, 2015, pp. 315–335.
- [41] W.J. Rider, L.G. Margolin, From numerical analysis to implicit subgrid turbulence modeling, in: 16th AIAA Computational Fluid Dynamics Conference, Orlando, USA, 2003.
- [42] X. Roca, N.C. Nguyen, J. Peraire, Scalable parallelization of the hybridized discontinuous Galerkin method for compressible flow, in: 43rd AIAA Fluid Dynamics Conference and Exhibit, Orlando, USA, 2013.
- [43] Y. Saad, M.H. Schultz, GMRES: a generalized minimal residual algorithm for solving nonsymmetric linear systems, *SIAM J. Sci. Stat. Comput.* 7 (3) (1986) 856–869.
- [44] H. Schlichting, K. Gersten, *Boundary-Layer Theory*, 8th edn., Springer-Verlag, 1999.
- [45] P.R. Spalart, Strategies for turbulence modelling and simulations, *Int. J. Heat Fluid Flow* 21 (2000) 252–263.
- [46] M.R. Visbal, R.E. Gordnier, M.C. Galbraith, High-fidelity simulations of moving and flexible airfoils at low Reynolds numbers, *Exp. Fluids* 46 (5) (2009) 903–922.
- [47] A. Uranga, P.-O. Persson, M. Drela, J. Peraire, Implicit Large Eddy Simulation of transitional flows over airfoils and wings, in: 19th Computational Fluid Dynamics Conference, San Antonio, USA, 2009.
- [48] A. Uranga, P.-O. Persson, M. Drela, 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 (2011) 232–261.
- [49] C.C. de Wiart, K. Hillewaert, Development and validation of a massively parallel high-order solver for DNS and LES of industrial flows, in: *IDIHOM: Industrialization of High-Order Methods – A Top-Down Approach*, in: Notes on Numerical Fluid Mechanics and Multidisciplinary Design, vol. 128, 2015, pp. 251–292.
- [50] Y. Zhou, Z.J. Wang, Implicit Large Eddy Simulation of transitional flow over a SD7003 wing using high-order spectral difference method, in: 40th Fluid Dynamics Conference and Exhibit, Chicago, USA, 2010.