Enabling Adaptive Mesh Refinement for Spectral-Element Simulations of Turbulence Around Wing Sections

Álvaro Tanarro^{1,2} · Fermín Mallor^{1,2} · Nicolas Offermans^{1,2} · Adam Peplinski^{1,2} · Ricardo Vinuesa^{1,2} · Philipp Schlatter^{1,2}

Received: 6 November 2019 / Accepted: 21 April 2020 / Published online: 19 May 2020
© The Author(s) 2020

Abstract

The implementation of adaptive mesh refinement (AMR) in the spectral-element method code Nek5000 is used for the first time on the well-resolved large-eddy simulation of the turbulent flow over wings. In particular, the flow over a NACA4412 profile with a 5° angle of attack at chord-based Reynolds number $Re_c = 200,000$ is analyzed in the present work and compared with a previous conformal simulation used as baseline. The mesh, starting from a coarse resolution, is progressively and automatically refined by means of AMR, which allows for high resolution near the wall and wake whereas significantly larger elements are used in the far-field. The resulting mesh, which remains unchanged in the production runs (i.e. AMR is used to create the final mesh, which is then fixed), is of higher resolution than those in previous conformal cases, and it allows for the use of larger computational domains, avoiding the use of precursor RANS simulations to determine the boundary conditions. This is achieved with, approximately, 2 times lower total number of grid points if the same spanwise length is used. Turbulence statistics obtained in the AMR simulation show good agreement with the ones obtained with the conformal mesh, and the pressure-coefficient distribution along the wing surface matches pressure-scan experimental data obtained in the MTL wind tunnel at KTH. The use of AMR is therefore expected to enable the simulation of high-Reynolds numbers turbulent flows over complex geometries (such as wings), thus allowing the study of pressure-gradient effects at high Reynolds numbers relevant for practical applications.

Keywords Large-eddy simulation · Airfoil · Turbulent boundary layer · Adverse pressure gradient · Adaptive mesh refinement

✉ Fermín Mallor
mallor@kth.se

¹ SimEx/FLOW, Engineering Mechanics, KTH Royal Institute of Technology, SE-100 44 Stockholm, Sweden

² Swedish e-Science Research Centre (SeRC), Stockholm, Sweden

1 Introduction

Wall-bounded turbulence is of great relevance in a wide range of industrial applications such as turbine blades or aircraft wings among many others, since it is the main source of drag generation which limits the performance of these devices. Drag reduction has become unavoidable in certain fields such as aircraft manufacturing due to the strong restrictions set on fuel consumption and pollution for the upcoming years. One region of improvement is the wings of the aircraft where a considerable amount of drag is generated in the turbulent boundary layers (TBLs) developing around them. These TBLs are very complex due to the strong streamwise pressure gradient (PG) caused by the airfoil curvature, which has a significant effect on the development of the boundary layer. The relevance of PG TBLs is supported by the numerous works related to its analysis, both experimental and numerical. Whereas the first studies of turbulent boundary layers developing over wings were conducted experimentally in wind tunnels (see for instance Coles and Wadcock 1979; Wadcock 1987), recent developments in the available computational resources and methods have allowed accurate numerical simulations to be performed of cases equivalent to those carried out experimentally. Some examples of this are the large-eddy simulations (LESs) performed by Sato et al. (2016), Frère et al. (2018) and Vinuesa et al. (2018), where the flow around airfoils is simulated at Reynolds numbers based on inflow velocity and chord length ($Re_c = cU_\infty/\nu$) of at least 1 million. Furthermore, a number of studies have focused on the complex effects of the adverse pressure gradient (APG) produced along the wing suction side on the characteristics of the developing TBL. The first numerical study of APG TBLs was conducted by Spalart and Watmuff (1993), who reported reduced inner-scaled velocities in the buffer and beginning of the overlap layers as a consequence of the APG. The prominent role of the outer region in APGs, leading to very energetic large-scale motions, was documented in the experiments by Skåre and Krogstad (1994) and Harun et al. (2013). APG TBLs have recently received some renewed attention from the numerical perspective, where the studies by Bobke et al. (2017) and Lee (2017) consider constant and moderate APG magnitudes, and the works by Kitsios et al. (2017) and Maciel et al. (2018) are concerned with much stronger APG conditions. The mechanisms leading to the very energetic large-scale motions in the outer region of APG TBLs developing around wings were recently identified and discussed by Tanarro et al. (2020).

Numerical simulations are an excellent way to study fluid dynamics as they avoid typical wind-tunnel limitations such as the interference of the side walls on the flow, inaccuracies in the measurements or flow conditioning effects that might significantly affect the results, as reported in Ref. Barlow et al. (2015). On the other hand, wind-tunnel experiments are currently the only available option to study high Reynolds numbers due to the extremely high cost of corresponding high-resolution numerical simulations: the number of grid points required for a wall-resolving LES is proportional to $Re^{13/7}$, as reported by Choi and Moin (2012). Despite the fact that it is currently possible to perform high-fidelity simulations of TBLs around wings at moderate Re_c (Sato et al. 2016; Frère et al. 2018; Vinuesa et al. 2018), we are still far from the transportation aircraft range, where $Re_c \sim O(10^7)$, which will only be accomplished through more efficient computational methods and increased computational power. In this project we aim to increase the achievable value of Re_c by using adaptive mesh refinement (AMR), which provides an increased mesh flexibility and improves the current code performance at high Re_c , as discussed below. Berger and Oliger (1984) first proposed AMR in the context of hyperbolic partial differential equations, and later two- and three-dimensional

AMR simulations were carried out by Berger and Colella (1989) and Bell et al. (1994), respectively. In 1990 Mavriplis proposed the use of a-posteriori error estimators for the spectral-element method (Mavriplis 1990). Adaptive finite-element-method (FEM) techniques for general computational mechanics were reviewed by Johnson and Hansbo (1992). In the context of computational fluid dynamics (CFD), the FEniCS project (Hoffman et al. 2013; Alnæs et al. 2015) was one of the first ones to incorporate adjoint error estimators for finite elements (Hoffman 2009). Other examples include three-dimensional aerodynamic applications (Hartmann et al. 2010), shallow-water equations and tsunami modelling (Pranowo et al. 2008; Blaise et al. 2013), as well as shock modelling (Löhner 1987; Devloo et al. 1988), which clearly benefits from localized and moving regions of high resolution.

In this study, which is an extension of the work by Tanarro et al. (2019), we compare the results obtained with a mesh constructed by means of AMR used to simulate the flow around a NACA4412 wing section with 5° angle of attack and $Re_c = 200,000$, with results of the same flow using a conformal mesh based on typical resolution requirements for turbulence simulations. The ultimate aim of this project is to perform well-resolved LES of the turbulent boundary layers developing around a NACA4412 at $Re_c = 1,640,000$, i.e. the same Re_c as that of the experiment by Wadcock (1987), at various angles of attack up to around 12° . This will provide some insight on the effect of the angle of attack and will allow to extend the work by Vinuesa et al. (2018) to analyse the effect of APGs at high Re_c .

The article is organized as follows: first, the methodology used in the simulations is explained in Sect. 2; the numerical setup is discussed, as well as the refinement process and the error estimators used for the adaptive mesh refinement. Then, as a means of validation of the AMR setup, results for a test-case are compared to those obtained using a conformal mesh in Sect. 3. Lastly, the main results of this work are summarized in Sect. 4, and the steps to continue this work are discussed.

2 Simulations with Adaptive Mesh Refinement (AMR)

Solution-aware simulation methods such as AMR are a promising way to achieve high- Re simulations of complex cases with sufficient well-controlled accuracy. The idea behind this method is to automatically refine the mesh in the regions where higher resolution is needed, whereas the areas with less flow variations become coarsened. Although this idea is not new, its application to the spectral-element code Nek5000 (Fischer et al. 2008) employed in the present work is very recent. Nek5000 has shown very high accuracy and efficiency for simulating turbulent flows with excellent parallel scalability [see the work by Offermans et al. (2016)]. Nevertheless, one of the main limitations of Nek5000 when simulating complex geometries is the lack of mesh flexibility and thus the difficulty of creating well-resolved boundary-layer meshes, since Nek5000 requires conforming hexahedral meshes (although it also supports unstructured meshes). Thus, the AMR implementation provides a significant flexibility improvement to the code. In addition, previous conformal meshes of the flow over wing sections, such as the ones used by Vinuesa et al. (2018), exhibit two main issues: first, the conforming nature of the mesh leads to an over-refinement of the far field, which considerably increases the number of grid points; and second, the large aspect ratio of the elements in the far field which increases the number of iterations required to solve the pressure, therefore decreasing significantly the code performance.

2.1 Description of the Flow Case

The flow under study consists of the turbulent boundary layers developing around a NACA4412 wing section at $Re_c = 200,000$ with an angle of attack of 5° . Here we consider a well-resolved LES, which allows to accurately compute the largest scales of the turbulent flow while the smallest, and most isotropic scales are modelled using a comparably simple relaxation filter, described in Schlatter et al. (2004). The well-resolved LES approach was validated against fully resolved direct numerical simulations (DNS) of the TBLs around a wing section by Negi et al. (2018). The velocity components in each spectral element are represented in terms of Lagrange interpolants of polynomial order $N = 7$. The flow around the airfoil is simulated inside a rectangular domain with streamwise and vertical lengths $L_x = L_y = 40c$, as illustrated in Fig. 1 (left), and spanwise length $L_z = 0.6c$. The airfoil is located at the centre of the domain and rotated in order to introduce the angle of attack of 5° . The boundary conditions are inflow with constant velocity U_∞ on the left, normal outflow condition with tangential U_∞ on top and bottom ends of the domain, stabilised outflow boundary condition on the right side of the domain (see Ref. Dong et al. (2014) for the details) and periodic boundary conditions in the spanwise direction. The latter sets a limitation in the number of elements in the spanwise direction for the initial mesh. Since three elements are needed in order to impose periodicity, a compromise between the size of the elements in the far-field and the size of the homogeneous direction must be found in order to have a low aspect ratio. In the present case, this leads to the use of a wider domain in the AMR case than in the conformal case used as baseline, in which $L_z = 0.2c$. In order to set the transition location of the boundary layers, both of them (on the suction and pressure sides of the airfoil) are tripped at $x/c = 0.1$. The tripping is modeled by introducing a weak random time-dependent volume forcing as discussed in Ref. Schlatter and Örlü (2012). The forcing also varies in the spanwise direction, replicating the tripping strips commonly used in experiments, and its shape is represented by a two-dimensional Gaussian distribution. For the non-conformal AMR implementation, the forcing is cut-off at a

Fig. 1 Comparison of computational domain using AMR for the simulation of a NACA4412 wing section at $Re_c = 200,000$ (marked in black) and the reference conformal mesh computational domain (marked in red). (Left) Zero-level AMR mesh before refinement and (right) final AMR mesh. The conformal mesh has been shifted and rotated in order to match the coordinates of the AMR case

fixed distance from the source (i.e. the Gaussian is truncated), whereas in the previous conformal implementation in Nek5000, the forcing was smoothly attenuated by the aforementioned Gaussian distribution.

2.2 Refinement Method

The various strategies for mesh refinement with high-order methods are *r*-refinement, where the grid points are displaced without changing the topology of the mesh, *h*-refinement, where the number of grid points is increased locally, and *p*-refinement, where the order of the polynomial expansion is increased locally. The latter two methods are commonly combined in the so-called *hp*-refinement method. In the present work, *h*-refinement is used for mesh adaptation. There are two main reasons behind favouring this method over *p*-refinement: the use of *h*-refinement (together with the implementation of non-conformal meshing) allows a higher degree of mesh flexibility and eases the initial mesh design process. Secondly, the current implementation of data structures in Nek5000 does not support elements with unequal polynomial order.

Refinement (of *h*-type) is performed by splitting selected elements with an oct-tree (3D) or quad-tree (2D) structure isotropically, which means that children elements will maintain the aspect ratio of the parent element. This type of refinement introduces non-conformal meshing when two adjacent elements have different levels of refinement: As a consequence of the splitting, elements might share a fraction of an edge only and hanging nodes appear at the interface between fine and coarse elements. The technique is chosen for its flexibility but requires significant modifications in Nek5000, which does not support such a feature by default. First, the one-to-one correspondence between grid points along the element faces is lost and interpolation operators are introduced to ensure continuity at the interface between coarse and fine elements as in the work by Kruse (1997). The continuity of the solution can be clearly observed in both Figs. 6 and 15, in which instantaneous fields of the streamwise velocity and coherent vortical λ_2 structures [introduced by Jeong and Hussein (1995)], respectively, are shown. From a solver performance standpoint, we verified in a previous work that the efficiency of the pressure solver is maintained despite the introduction of those operators [the details can be found in Peplinski et al. (2018)]. Then, we rely on the library p4est, developed by Burstedde et al. (2011), which represents the mesh as a recursive tree structure and keeps track of the changing grid hierarchy and connectivity. Moreover, a well-balanced grid partitioning among the processes is ensured by the library ParMETIS (Karypis et al. 2009) to maintain good parallel scaling. This load balancing is performed after each refinement step, in order to maintain an optimal parallel performance throughout the whole simulation. A two-level partitioning strategy is used, where the distribution of spectral elements is done in two steps: an inter-node partitioning occurs first, followed by an intra-node partitioning. The resulting grid management is done fully in parallel and a good scalability of the code is maintained. One last change to the Nek5000 code is the modification of the preconditioner for the pressure equation. Nek5000 uses the combination of an overlapping Schwarz decomposition, resulting in small local problems, and a global coarse-grid solver. The local problems are modified to include the interpolation operator at refined interfaces. The coarse-grid problem must be adapted to exclude contributions from hanging nodes, which are not true degrees of freedom, as discussed by Peplinski et al. (2018).

2.3 Spectral Error Indicators

To estimate the error used to decide how to modify the mesh, we employ the a-posteriori error indicators developed by Mavriplis (1990). Considering the exact solution u to a system of partial differential equations and an approximate spectral-element solution at polynomial order N , i.e. u_N , the estimators provide an estimate of the L^2 -norm between u and u_N . For the sake of clarity, we present the method for a 1D problem. We can expand a solution $u(x)$ on a reference element in terms of the Legendre polynomials and express the associated spectral coefficients $\hat{u}_k$ as:

$$u(x) = \sum_{k=0}^{\infty} \hat{u}_k L_k(x) \quad \text{and} \quad \hat{u}_k = \frac{1}{\gamma_k} \int_{-1}^1 u(x) L_k(x) dx, \quad (1)$$

where L_k is the Legendre polynomial of order k and $\gamma_k = \|L_k\|_{L^2}^2$. The error on $\|u - u_N\|_{L^2}$ is attributed to two contributions: a truncation error due to the finite number of coefficients in the spectral expansion and a quadrature error. The estimated error indicator $\varepsilon \approx \|u - u_N\|_{L^2}$ is given by:

$$\varepsilon = \left(\int_N^{\infty} \frac{\hat{u}(k)^2}{\frac{2k+1}{2}} dk + \frac{\hat{u}_N^2}{\frac{2N+1}{2}} \right)^{\frac{1}{2}}, \quad (2)$$

where it is assumed that the decay of the spectral coefficients follows the relation $\hat{u}_k \sim \hat{u}(k) = c \exp(-\sigma k)$. The parameters c and σ are obtained via a least-square best fit on the last four spectral coefficients. The computations are readily extended to the 2D and 3D cases. The main advantage of these estimators is their low computational cost, which is negligible in comparison to the time spent in the solver. From these indicators, we obtain an element-wise measure of the representation error at each time step. The error is then averaged in time and mesh adaptation is performed periodically, at a given frequency. The estimators were used for steady adaptive simulations and compared to dual-weighted error estimators by Offermans et al. (2020) and they were shown to successfully identify poorly resolved regions.

2.4 Design of the Non-conformal Mesh

The main objective with AMR is to use the tools for mesh refinement and error estimators to design an optimal mesh, where the error in the solution is as low and uniform as possible at a reasonable computational cost. Starting from the initial really coarse mesh, a minimum refinement level of 6 is imposed along the walls of the airfoil to ensure that the initial run is sufficiently resolved and stable. An additional limitation in the refinement level around the tripping lines is imposed, so as to avoid the over-refinement in these areas which arises from the high-frequency oscillations introduced by the tripping body force. Then, the simulation proceeds and error indicators are collected in time. With a fixed frequency the refinement process is performed: a defined percentage (which is decreased as the number of elements increases) of elements are refined or coarsened based on the value of the spectral error indicators: the elements with the lowest error are coarsened, and those with the highest one are refined. The process is repeated until the refinement at the wall reaches at least the well-resolved LES resolution described by Vinuesa et al. (2018). In the

present simulation, the resulting maximum level of refinement along the wall is 7. Once the refinement process is over, the mesh is “frozen”, i.e. mesh refinement is disabled, and the production runs are carried out. Let us note that this is a deliberate choice and not a limitation of our tools for AMR, which support dynamic mesh adaptation during the course of the simulation.

The resulting final mesh is shown in Fig. 1, with an illustration of the conformal mesh used in one of our previous studies (Atzori et al. 2018) superposed over it for comparison. Note that AMR enables the use of a significantly larger computational domain, where boundary conditions do not have to be extracted from a precursor Reynolds-averaged Navier–Stokes (RANS) simulation (Vinuesa et al. 2017, 2018) and where the far-field elements are not over-refined. At the same time, high resolution is achieved near the wall in the boundary-layer region, as well as downstream of the airfoil in the wake. In order to limit the computational cost, we apply refinement up to two chords downstream of the airfoil, and consider a coarser mesh downstream of this point. Doing so, we avoid over-refining the wake, which is not relevant to the scope of this project. Figure 2 illustrates the different refinement levels for the final mesh at $Re_c = 200,000$, in which the elements with the highest refinement levels are located close to the wall or in certain regions of the wake due to the large velocity gradients present in those regions. Moreover, the refinement constraint set in the wake of the airfoil can be clearly seen. A validation of the results for the conformal and non-conformal meshes is performed in terms of the LES of the NACA4412 at $Re_c = 200,000$ and it is discussed next.

The wing setup as discussed in this work was also used as one of the ExaFLOW European Union (EU) project flag-ship cases used to perform detailed studies on the parallel performance of the AMR version of Nek5000 [see the discussion in Ref. Offermans (2019)]. The code performance has two components: the algorithmic performance (such as the number of pressure iterations and the overhead induced by the non-conformal meshes) and the parallel scaling. As discussed in Ref. Offermans (2019), the non-conformal solver does not lead to any increase of pressure iterations, and the interpolation for hanging nodes is an entirely local operation executed at very

Fig. 2 Regions of the domain covered by refinement levels higher than or equal to four for the final AMR mesh. The colours indicate the refinement level of the elements going from green (refinement level 4) to red (refinement level 7)

high speed. Therefore, we could not detect any algorithmic slowdown due to the non-conforming solver. Regarding the parallel performance, we focus mostly on the strong scaling properties of the non-conforming solver, and consider the time evolution loop only, excluding code initialisation, finalisation and any input/output (I/O) operations. The target architectures are petascale Cray XC40 systems, such as the Beskow system at PDC (KTH Stockholm). We paid attention to the different components of the non-conforming solver, e.g. the coarse grid solver setup times, and the work balance as a result of the two-level partitioning. The scaling tests were run for 100 steps starting from the same initial conditions, with a single MPI process per core and the number of compute nodes increasing as a power of two. Results of the scaling tests executed on Beskow are presented in Fig. 3. In the left panel, the timings for the coarse-grid setup, executed once after each mesh refinement, is shown. The relative cost of this setup with increasing number of nodes is apparent, even though the actual execution time decreases with increasing number of nodes. For a large number of elements and MPI ranks it becomes the most expensive operation of the whole refinement process. At the same time the solver itself scales very well, as evidenced by the time per time step presented in the right panel of Fig. 3, showing even slight super-linear scaling which may be caused by network interference while performing the runs. Furthermore, we compared these results with the scaling data of the conforming Nek5000 presented in Ref. Offermans et al. (2016). The most interesting case presented in this work is the turbulent pipe flow at $Re_\tau = 360$ (where the friction Reynolds number Re_τ is defined and discussed below), especially the upper-right plot in their Figure 5, as it was similar in size to the case discussed here. Note that the mentioned plot is almost identical to our strong scaling plot, a fact that shows that the parallel performance of non-conforming and conforming solvers is almost the same and proves the efficiency of our implementation.

3 Results

In this study we present the first wing simulations using AMR in Nek5000, therefore the first step in this work is to compare the results obtained using AMR with simulations with a conformal mesh. The conformal case is the one reported by Atzori et al. (2018) at $Re_c = 200,000$.

Fig. 3 Parallel performance of non-conforming solver for the wing case in its final configuration executed on a Cray XC40 computer at KTH, Stockholm (Beskow); note that each node consists of 32 compute cores. (Left) Strong scaling of the coarse-grid setup phase and (right) time per time step

3.1 Mesh Characteristics

The main differences between the AMR and the conformal cases, in addition to the employed spectral-element mesh, are the polynomial order ($N = 7$ in this work and $N = 11$ in the conformal case), the tripping function (which was slightly modified in the AMR version by removing the forcing in regions far from the smoothing length) and the resulting resolution. The resolution requirements are extracted from the work by Negi et al. (2018), in which the optimal resolution for the LES in order to match DNS results was found to be $\Delta y_w^+ = 0.64$, $\Delta x^+ = 18$ and $\Delta z^+ = 9$, in accordance with the resolution used in the well-resolved LES of TBLs in Refs. Schlatter et al. (2010) and Eitel-Amor et al. (2014). The superscript $+$ denotes inner scaling, which is obtained in terms of the friction velocity $u_\tau = \sqrt{\tau_w/\rho}$ (where τ_w is the wall-shear stress and ρ the fluid density) and the viscous length $\ell^* = \nu/u_\tau$. The grid resolution in the three spatial directions is reported in Fig. 4, where the aforementioned resolution goals are marked as a black dashed line. The resolution in the wall-normal direction is measured as the distance between the first grid point (not element center) and the wall. Even though the conformal case employs a considerably higher polynomial order, the significant refinement of the AMR case (needed to properly resolve the velocity gradients in the near-wall region) around the airfoil surface results in a similar resolution on the suction side, in which the acceleration and later deceleration of the flow are higher. This high concentration of elements around the suction side and wake region can be clearly seen in Fig. 5, in which the two meshes around the airfoil are compared, as well as in Fig. 6, where an instantaneous streamwise velocity field is shown. In the latter, the working principle and scope of the spectral error indicators used as a basis for refinement become clear: regions in which (locally) high velocity fluctuations exist need a higher spatial resolution in order to reduce the spectral representation error caused by the larger amplitude of the higher modes, driving the refinement of such regions. In the present case, these regions correspond to the wing's boundary layers and wake. On the other hand, the resolution on the pressure side is lower in the AMR mesh compared to the conformal one. On the pressure side, the mild favorable pressure gradient results in lower turbulent fluctuations compared to the suction side, relaxing the local resolution requirements. In fact, the refinement, as seen in the right panel of Fig. 5, is one level lower on the pressure side than on the suction side. The effect of refinement level on the resolution can be clearly seen in the area close to the trailing edge of the suction side. The sudden increases in grid spacing are caused by changes in the refinement level. In order to assess the resolution in the streamwise and spanwise directions (shown in Fig. 4), the maximum spacing between grid points is used, as this constitutes the most strict condition. Due to the low-aspect-ratio elements of the non-conformal mesh, and the isotropic nature of the refinement process (i.e. elements are split into eight equal elements when refined), the spanwise and tangential resolutions of the AMR case are considerably finer than in the conformal case, particularly on the suction side, in which the highest refinement levels are achieved by the AMR.

Even if the resolution of the non-conformal (AMR) mesh around the wing is higher than that of the conformal mesh, the total number of grid points in the mesh is lower if one takes into account the domain size in the spanwise direction. The non-conformal mesh has a total of 323 million grid points (with a spanwise width of $L_z/c = 0.6$) whereas the conformal mesh, with $L_z/c = 0.2$, has a total of around 211 million grid points. This implies that the resolution of the conformal case, extended to the full

Fig. 4 Inner-scaled grid resolution at the wing surface in the streamwise (top) and spanwise (middle) directions, together with spacing of the first grid point from the wall (bottom), for the AMR (red color line) and conformal (blue color line) cases. The sudden spacing increases observed for the AMR case relate to changes in refinement levels

width of $L_z/c = 0.6$, would require 634 million grid points (a number almost 2 times larger than the one of the AMR mesh). Therefore, the use of AMR in Nek5000 provides numerous advantages in comparison to the conformal mesh for the same wall-normal resolution: lower number of grid points, over 60 times larger computational domain (and, in turn, no need for the RANS-extracted boundary condition) and low-aspect-ratio elements (which leads to an increased efficiency of the pressure solver).

Fig. 5 Spectral-element distribution around the wing and wake regions for the conformal (left) and non-conformal AMR (right) meshes. The differences in the coordinates are due to the different positioning of the wing in the domain: with the leading edge (LE) at $x/c = -0.5$ and rotated 5° in the AMR case, and with the LE at $x/c = 0$ and not rotated in the conformal case. The 99% boundary-layer thicknesses close to the trailing edge are 0.055 and 0.023 in both simulations on the suction and pressure sides, respectively

Fig. 6 Instantaneous field of the streamwise velocity component normalized with U_∞ around the wing and wake, with the spectral elements superimposed in black

3.2 Integral Quantities

The analysis of the AMR results is performed by computing several integral quantities and turbulence statistics of the TBLs around the wing, and comparing them with those obtained from the simulation using the conformal mesh. In order to obtain the boundary-layer thickness in a robust way, the method proposed by Vinuesa et al. (2016a) is used. This method is based on the diagnostic-plot scaling, reported in Ref. Alfredsson et al. (2011), and allows the determination of a boundary-layer thickness similar to the usual 99% boundary-layer thickness δ_{99} even in flows subjected to a wide range of pressure-gradient conditions. The turbulence statistics are averaged in the spanwise direction due to the homogeneity of the flow in this direction, and in time after initial transients related to startup after each refinement step. In order to quantify the temporal convergence of the results, the length of the time average is expressed in terms of the normalized eddy-turnover time scale $ETT^* = \delta_{99}/u_\tau L_{z,ref}/L_z$ as proposed by Vinuesa et al. (2016b). This definition of the eddy-turnover time normalises the length of the homogeneous direction (in this case L_z) with respect to a reference length $L_{z,ref} = 3\delta_{99}$, as defined by Flores and Jiménez (2010) as the minimum box size required to fully capture the largest structures in the logarithmic region. In terms of convergence as a function of this normalized eddy-turnover time, we compare our results to the ones of the temporally well-converged high- Re DNS performed by Silero et al. (2014). At higher Reynolds number of about $Re_\theta = 4500$ they reach an averaging time of about $39ETT^*$, whereas our averaging time at $x/c = 0.6$ on the suction side is about

43 normalized eddy-turnover time units. In the more obvious units of flow-over times T_{FO} , which are defined terms of the inflow velocity and the chord length, this averaging corresponds to $2.7 T_{FO}$. Note that the simulation was run for a total of 18 flow-over times, of which the first 15 were discarded as they were used in the refinement process as described above, as well as in an initial unsuccessful AMR run, which is detailed in the “Appendix” added to the end of the document.

All the turbulence statistics are expressed in terms of the surface-local tangential (t) and normal directions (n), and the corresponding tensors are rotated as discussed by Vinuesa et al. (2017). In the following, the subscripts ss and ps are employed to denote suction and pressure sides of the wing, respectively. The first quantity under consideration is the Clauser pressure-gradient parameter $\beta = \delta^*/\tau_w dP_e/dx_t$, which represents the magnitude of the adverse pressure gradient along the chord. The variables that define the Clauser pressure-gradient parameter are the displacement thickness δ^* , the pressure at the boundary-layer edge P_e and the distance from the leading edge x_t measured along the direction tangential to the wing surface. Figure 7 shows β for both cases and, although there are small differences mainly at the beginning of the TBL and as it approaches the trailing edge, it can be stated that the results show a generally good agreement. Note that in this case the inflection point at the location of maximum profile camber (i.e. $x_{ss}/c = 0.4$) observed at higher Reynolds numbers (see Ref. Vinuesa et al. (2018)) is only visible in the AMR case, which has higher resolution close to the wing.

Next, the skin-friction coefficient normalised with the free-stream velocity U_∞ , i.e. $C_f = 2(u_\tau/U_\infty)^2$, is provided in Fig. 8. Our results show that C_f is slightly larger in the simulation with AMR than in the conformal case in the majority of the chord on both sides of the wing. This discrepancy can be attributed to the higher near-wall resolution of the former which allows to resolve higher velocity gradients increasing the friction. The rapidly decreasing skin friction towards the trailing edge on the suction side, due to the progressively stronger APG, is very well reproduced in both simulations. The difference in the skin-friction coefficient leads to a similar discrepancy in the friction Reynolds number $Re_\tau = \delta_{99}u_\tau/\nu$, as can be observed in Fig. 9. Furthermore, this figure reveals some differences between both simulations in the range $x/c \simeq 0.2 - 0.3$, which are due to the differences in the tripping of the boundary layer of both simulations. The strength of the tripping in the conformal case is higher than that of the AMR case, having a more noticeable effect both upstream and downstream of the tripping location

Fig. 7 Clauser pressure-gradient parameter β on the suction side as a function of the distance from the leading edge at $Re_c = 200,000$

Fig. 8 Skin-friction coefficient C_f distribution scaled with the free-stream dynamic pressure ($1/2\rho U_\infty$) as a function of the distance from the leading edge

Fig. 9 (Left) Streamwise evolution of the friction (Re_τ) and (right) momentum-thickness (Re_θ) Reynolds numbers

at $x/c = 0.1$. The differences in the tripping implementation lead to different transition lengths (as seen in Fig. 8): The more refined tripping in the AMR case leads to a shorter transition region, and it has an overall smaller impact on the the flow-field around it, both upstream and downstream. Note that the differences in the Re_τ curves farther downstream are mainly due to the friction velocity u_τ , since the streamwise evolution of the boundary-layer thickness is in very good agreement in both simulations. This can be further observed in the right plot of Fig. 9, which shows the streamwise evolution of the Reynolds number based on momentum thickness $Re_\theta = U_e \theta / \nu$ in both cases. This figure shows an excellent agreement of both wing simulations on the suction and pressure sides. Therefore, we can conclude from these integral boundary-layer parameters that the global evolution of the boundary layer is captured well with both meshes, in particular the global pressure gradient and the various thicknesses are in very good agreement. The advantages of the AMR, allowing refined resolution close to the wall, is highlighted by the differences in the wall-shear stress, for which the AMR data is assumed to be more accurate.

3.3 Turbulence Statistics

Regarding turbulence statistics, inner-scaled mean tangential velocity and non-zero Reynolds stresses are discussed next at $x/c = 0.4$ and $x/c = 0.6$, on both sides of the wing. Starting with the inner-scaled tangential mean velocity profiles shown in Fig. 10, the profile on the suction side at $x/c = 0.4$ shows a fair agreement between both cases. The pressure side shows a minor discrepancy in the mean velocity. When scaled in outer units (not shown here) this difference is not observed, therefore, the discrepancy arises from the difference in wall-shear stress discussed previously. Figure 8 shows smaller differences in C_f at $x/c = 0.6$, which result in a better agreement of the inner-scaled mean velocity profile on the pressure side. Furthermore, the mean velocity profile at $x_{ss}/c = 0.6$ is in reasonably good agreement in both simulations, given the small differences in wall-shear stress. This is a remarkable result, keeping in mind the strong local pressure gradient at this location and the general complexity of the flow case.

Figure 11 shows the inner-scaled Reynolds stresses of both cases on the suction side. The results at $x_{ss}/c = 0.4$ exhibit an excellent agreement in all the stresses, except for a slightly larger value of the near-wall peak of u_t^2 predicted by the AMR simulation. This is also connected to the higher resolution in the spanwise and streamwise directions in the AMR case which, as mentioned before, allows to resolve more scales of the turbulent flow. The Reynolds stresses are also in reasonably good agreement at $x_{ss}/c = 0.6$, in particular when it comes to the accumulation of energy in the outer region as a consequence of the streamwise APG. Note, however, the small deviations in the spanwise velocity fluctuations in the outer region. The discontinuity observed in the tangential and spanwise stresses at $y_n^+ \approx 30$ for $x/c = 0.6$ coincides with a change in refinement level in the AMR mesh. Strategies to mitigate such an imprint of the refinement onto the fluctuating quantities is part of our present work. Nonetheless, both simulations are able to capture well the main aspects of APG boundary-layer flow.

Fig. 10 Inner-scaled tangential mean velocity profiles at (left) $x/c = 0.4$ and (right) $x/c = 0.6$, for both suction and pressure sides

Fig. 11 Inner-scaled non-zero Reynolds stresses corresponding to (left) $x_{ss}/c = 0.4$ and (right) $x_{ss}/c = 0.6$ on the suction side. The colors indicate: tangential (blue color line), wall-normal (red color line), spanwise (green color line) velocity fluctuations and Reynolds-shear stress (cyan color line)

3.4 Mean Flow Fields

The mean flow around and downstream of the airfoil is analyzed in Figs. 12, 13 and 14. In Fig. 12 the streamwise velocity defect $(U - U_\infty)/U_\infty$ is plotted at different horizontal and vertical profiles, whereas in Fig. 13, the difference in the velocity and pressure-coefficient (c_p) fields between the conformal and AMR cases are shown. Lastly, the pressure-coefficient distributions around the wing surfaces are reported in Fig. 14, and they are compared with experimental pressure-scan data obtained in an on-going campaign at KTH (see Mallor, 2019 for additional details Mallor 2019). Note that the pressure coefficient $c_p = 2(P - P_{ref})/(\rho U_\infty^2)$ is defined in terms of a reference pressure P_{ref} such that $c_p = 1$ at the stagnation point. Small deviations in the velocity defect close to the leading-edge region of the airfoil can be seen in the left panel of Fig. 12, which can be attributed to the fact that the excessive strength of the tripping in the conformal case interfered with the laminar-flow region upstream of it, affecting notably the pressure distribution around the leading edge as seen in Fig. 14. The higher strength of the body-force tripping can more clearly be observed in the right panel of Fig. 13, in which

Fig. 12 Outer-scaled horizontal velocity defect profiles at different horizontal (left) and vertical (right) profiles around and downstream of the airfoil

Fig. 13 Contour plots of the difference in streamwise velocity U (left) and pressure coefficient c_p (right) between the AMR and conformal cases ($U_{AMR} - U_{Conf}$ and $p_{AMR} - p_{Conf}$, respectively) scaled with the inflow velocity (U_∞) and inflow dynamic pressure ($1/2\rho U_\infty^2$), respectively. The domain and coordinate system correspond to that of the conformal case

Fig. 14 Pressure-coefficient c_p distribution along the NACA4412 profile for the AMR and conformal cases, compared to recently obtained experimental data (Mallor 2019)

the difference in the pressure coefficient is shown. The largest difference between both cases can be found on both sides of the wing around the tripping position ($x/c = 0.1$), confirming that the employed tripping is one of the main sources of differences between both cases. The excellent agreement between the AMR and experimental c_p data observed in Fig. 14, together with the disturbance of the laminar region in the conformal case, allows us to state that the new implementation of the tripping forcing provides a smoother and less intrusive transition to turbulence. On the other hand, the evolution of the turbulent wake downstream of the airfoil can be observed in the right panel of Fig. 12, showing the same downwash magnitude for both the conformal and AMR simulations. However, the vertical location of the wake as it is advected downstream varies slightly between both cases, as it can be seen in the left panel of Fig. 13 in which the differences in streamwise velocity between the conformal and AMR cases are plotted. The wake in the AMR case exhibits a slightly higher displacement downwards, a fact that can be attributed to the confinement effect present in the conformal simulation (with a narrower domain in z). Furthermore, the conformal case exhibits small pressure differences near the corners of the outflow (Fig. 13, right), which are an indication of the shortcomings of the use of a small computational domain (with the corresponding

Fig. 15 Instantaneous visualization of the flow around a NACA4412 wing section at $Re_c = 850,000$, showing coherent vortical structures identified with the λ_2 criterion (Jeong and Hussein 1995). The spectral-element mesh is also shown (although not the individual grid points within elements), illustrating the various required refinement levels. The color indicates streamwise velocity, where blue and red denote low- and high-velocity regions, respectively

Dirichlet boundary condition). This can be avoided through the use of non-conformal meshing. Nevertheless, note that the differences between both cases are always below 5% in both velocity and pressure coefficient, and they are almost zero in most parts of the flow field.

4 Conclusions and Outlook

In this study we have shown that, through the use of adaptive mesh refinement, it is possible to obtain high-quality meshes to simulate the turbulent boundary layers around wing sections using well-resolved LES. The employed code is Nek5000, a high-performance spectral-element code, suitable for moderately complex geometries, such as wing profiles. We have compared turbulence statistics at $Re_c = 200,000$ for a conformal case and an AMR-based configuration, which allows to reduce the required number of grid points by approximately a factor of 2. The agreement between both simulations is good in all the analyzed quantities, and the small observed discrepancy in the predicted skin-friction coefficient is connected to resolution differences. The aim is that of exploiting the computational savings of AMR in order to enable well-resolved simulations at higher Reynolds numbers. In Fig. 15 we show an instantaneous visualization of the coherent vortical structures on a NACA4412 wing section, also with 5° angle of attack, at $Re_c = 850,000$. This on-going simulation is also obtained with AMR, and requires a total of 3.3 million spectral elements with $N = 8$, i.e. a total of 2.4 billion grid points. Note the high level of detail of the turbulent structures around the wing obtained in this simulation. On-going work is aimed at reproducing the higher Reynolds number $Re_c = 1,640,000$ put forward by the experiment by Wadcock (1987). The same base mesh, with additional refinements, will be used for that case.

Acknowledgements Open access funding provided by Royal Institute of Technology.

Funding The simulations were performed on resources provided by the Swedish National Infrastructure for Computing (SNIC) at PDC Center for High Performance Computing, KTH, Stockholm. This research is funded by the Knut and Alice Wallenberg Foundation (Grant Number 2018.0151), and was partially supported by the ExaFLOW H2020 Project (Project Number 671571).

Code Availability The spectral-element code Nek5000 is open-source and is publicly available at <https://nek5000.mcs.anl.gov/>. The authors are currently working with the Nek5000 development team to add the AMR code into the main repository.

Compliance with Ethical Standards

Conflict of interest The authors declare that they have no conflict of interest.

Availability of Data and Materials Not applicable.

Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

Appendix: Cautionary Note on Adaptive Refinement Strategy

As discussed in Sect. 2.4, the refinement process followed to obtain the final mesh used for the production runs was completely automatic: A defined percentage of elements were refined or coarsened in each refinement step (carried out after averaging the spectral error estimators during a fixed number of simulation steps) based on their local error estimator. The only constraints applied to the process were the minimum refinement level at the wall, as well as a limitation in the maximum refinement in a small area around the tripping position. Following this strategy, the mesh generated allowed to obtain a very good agreement with the reference data from a previous conformal simulation, as evidenced in Sect. 3. Previously, a different refinement strategy had been used to generate a production mesh: Elements were coarsened or refined based on maximum and minimum error estimator thresholds. Hence, in our previous approach, an extra step was added to the AMR process, namely defining ad-hoc relevant thresholds which were case-dependent. Moreover, as seen in Fig. 16, additional limitations were imposed on the mesh: Maximum and minimum refinement levels were defined in two ellipses surrounding the wing, with the aim of obtaining a fine enough mesh in the areas in which velocity fluctuations were small (e.g. near the leading edge, or at a small distance from the wing but outside of the boundary layer). An undesirable effect of using such restrictions on the minimum refinement level in broad areas is that as elements are concentrated in non-essential parts (which could be the case when considering a lower refinement level), the critical areas (like the TBL or wake) may be under-refined. Contrary to this, the completely automatic refinement avoids this over-refinement and concentrates the elements in the areas in which they are most needed (i.e. where small turbulent structures are present, namely, in the TBL and wake regions), albeit at potentially requiring a larger grid count. In terms of the quality of the simulation,

Fig. 16 Two-dimensional cut of the AMR meshes, when (left) introducing ad-hoc thresholds and (right) with a fully-automatic refinement process. The latter is the mesh considered throughout the manuscript. Note that the spectral-element boundaries are represented as black lines, and the bottom line provides a detailed view around the wing. Overall, the mesh to the left has approximately 50% more spectral elements

Fig. 17 Inner-scaled non-zero Reynolds stresses corresponding to (left) $x/c = 0.4$ and (right) $x/c = 0.6$ on the suction side. The colors indicate: tangential (blue color line), wall-normal (red color line), spanwise (green color line) velocity fluctuations and Reynolds-shear stress (cyan color line). Numbers indicate refinement level, and vertical lines the non-conformal interface between different refinement levels for (red) the case with ad-hoc thresholds and (blue) the fully-automatic configuration

the use of ad-hoc thresholds may lead to the presence of “kinks” in the turbulence statistics and deviations from the reference data, which can be observed in Fig. 17. The deviations from the reference data (and from the automatically-refined AMR mesh) coincide with a reduction in the refinement level caused by the presence of this ad-hoc refinement limitations. By removing the constraints on the refinement, and allowing the refinement process to be driven by the relative value of the error indicator, these problems disappear.

Nonetheless, it is important to note that minimum refinement levels should be set at the wing boundaries so that the turbulent boundary layer can be properly resolved, according to conventional resolution criteria in wall-bounded turbulence [see e.g. Vinuesa et al. (2018)].

References

- Alfredsson, P.H., Segalini, A., Örlü, R.: A new scaling for the streamwise turbulence intensity in wall-bounded turbulent flows and what it tells us about the “outer” peak. *Phys. Fluids* **23**(4), 041–702 (2011)
- Alnæs, M.S., Blechta, J., Hake, J., Johansson, A., Kehlet, B., Logg, A., Richardson, C., Ring, J., Rognes, M.E., Wells, G.N.: The FEniCS project version 1.5. *Arch. Numer. Softw.* **3**, 9–23 (2015)
- Atzori, M., Vinuesa, R., Stroh, A., Frohnapfel, B., Schlatter, P.: Assessment of skin-friction-reduction techniques on a turbulent wing section. In: Proceedings of the 12th ERCOFTAC Symposium on Engineering Turbulence Modelling and Measurements (ETMM12), Montpellier, France. arXiv preprint [arXiv:181203762](https://arxiv.org/abs/181203762) (2018)
- Barlow, J.B., Rae Jr., W.H., Pope, A.: Low speed wind tunnel testing. *INCAS Bull.* **7**(1), 133 (2015)
- Bell, J., Berger, M.J., Saltzman, J., Welcome, M.: Three-dimensional adaptive mesh refinement for hyperbolic conservation laws. *SIAM J. Sci. Comput.* **15**, 127–138 (1994)
- Berger, M.J., Colella, P.: Local adaptive mesh refinement for shock hydrodynamics. *J. Comput. Phys.* **82**, 68–84 (1989)
- Berger, M.J., Oliger, J.: Adaptive mesh refinement for hyperbolic partial differential equations. *J. Comput. Phys.* **53**, 484–512 (1984)
- Blaise, S., St-Cyr, A., Mavriplis, D., Lockwood, B.: Discontinuous Galerkin unsteady discrete adjoint method for real-time efficient tsunami simulations. *J. Comput. Phys.* **232**, 416–430 (2013)
- Bobke, A., Vinuesa, R., Örlü, R., Schlatter, P.: History effects and near equilibrium in adverse-pressure-gradient turbulent boundary layers. *J. Fluid Mech.* **829**, 667–692 (2017)
- Burstedde, C., Wilcox, L.C., Ghattas, O.: *p4est*: Scalable algorithms for parallel adaptive mesh refinement on forests of octrees. *SIAM J. Sci. Comput.* **33**(3), 1103–1133 (2011)
- Choi, H., Moin, P.: Grid-point requirements for large eddy simulation: Chapman’s estimates revisited. *Phys. Fluids* **24**(1), 011–702 (2012)
- Coles, D., Wadcock, A.J.: Flying-hot-wire study of flow past an NACA 4412 airfoil at maximum lift. *AIAA J.* **17**(4), 321–329 (1979)
- Devloo, P., Oden, J.T., Pattani, P.: An $h - p$ adaptive finite element method for the numerical simulation of compressible flow. *Comput. Methods Appl. Mech. Eng.* **70**, 203–235 (1988)
- Dong, S., Karniadakis, G.E., Chrysosostomidis, C.: A robust and accurate outflow boundary condition for incompressible flow simulations on severely-truncated unbounded domains. *J. Comput. Phys.* **261**, 83–105 (2014)
- Eitel-Amor, G., Örlü, R., Schlatter, P.: Simulation and validation of a spatially evolving turbulent boundary layer up to $Re_\theta = 8300$. *Int. J. Heat Fluid Flow* **47**, 57–69 (2014)
- Fischer, P., Kruse, J., Mullen, J., Tufo, H., Lottes, J., Kerkemeier, S.: Nek5000: open source spectral element CFD solver. <https://nek5000.mcs.anl.gov/> (2008)
- Flores, O., Jiménez, J.: Hierarchy of minimal flow units in the logarithmic layer. *Phys. Fluids* **22**(7), 071–704 (2010)
- Frère, A., Hillewaert, K., Chatelain, P., Winckelmans, G.: High Reynolds number airfoil: from wall-resolved to wall-modeled LES. *Flow Turbul. Combust.* **101**, 457–476 (2018)
- Hartmann, R., Held, J., Leicht, T., Prill, F.: Error estimation and adaptive mesh refinement for aerodynamic flows. In: Kroll, N., Bieler, H., Deconinck, H., Couaillier, V., van der Ven, H., Sørensen, K. (eds.) ADIGMA—A European Initiative on the Development of Adaptive Higher-Order Variational Methods for Aerospace Applications: Results of a collaborative research project funded by the European Union, 2006–2009, pp. 339–353. Springer, Berlin (2010)
- Harun, Z., Monty, J.P., Mathis, R., Marusic, I.: Pressure gradient effects on the large-scale structure of turbulent boundary layers. *J. Fluid Mech.* **715**, 477–498 (2013)
- Hoffman, J.: Efficient computation of mean drag for the subcritical flow past a circular cylinder using general Galerkin G2. *Int. J. Numer. Methods Fluids* **59**, 1241–1258 (2009)

- Hoffman, J., Jansson, J., de Abreu, R.V., Degirmenci, N.C., Jansson, N., Müller, K., Nazarov, M., Spühler, J.H.: Unicorn: Parallel adaptive finite element simulation of turbulent flow and fluid-structure interaction for deforming domains and complex geometry. *Comput. Fluids* **80**, 310–319 (2013)
- Jeong, J., Hussein, F.: On the identification of a vortex. *J. Fluid Mech.* **285**, 69–94 (1995)
- Johnson, C., Hansbo, P.: Adaptive finite element methods in computational mechanics. *Comput. Methods Appl. Mech. Eng.* **101**, 143–181 (1992)
- Karypis, G., Kumar, V.: MeTis: Unstructured Graph Partitioning and Sparse Matrix Ordering System, Version 4.0. <http://www.cs.umn.edu/~metis> (2009)
- Kitsios, V., Sekimoto, A., Atkinson, C., Sillero, J.A., Borrell, G., Gungor, A.G., Jiménez, J., Soria, J.: Direct numerical simulation of a self-similar adverse pressure gradient turbulent boundary layer at the verge of separation. *J. Fluid Mech.* **829**, 392–419 (2017)
- Kruse, G.W.: Parallel nonconforming spectral element solution of the incompressible Navier–Stokes equations in three dimensions. Ph.D. thesis, Brown University (1997)
- Lee, J.H.: Large-scale motions in turbulent boundary layers subjected to adverse pressure gradients. *J. Fluid Mech.* **810**, 323–361 (2017)
- Löhner, R.: An adaptive finite element scheme for transient problems in CFD. *Comput. Methods Appl. Mech. Eng.* **61**, 323–338 (1987)
- Maciel, Y., Wei, T., Gungor, A.G., Simens, M.P.: Outer scales and parameters of adverse-pressure-gradient turbulent boundary layers. *J. Fluid Mech.* **844**, 5–35 (2018)
- Mallor, F.: Enabling high-fidelity measurements of turbulent boundary layer flow over wing sections in the MTL wind tunnel. Master's thesis, KTH Royal Institute of Technology, Stockholm (2019)
- Mavriplis, C.: A posteriori error estimators for adaptive spectral element techniques. In: Wesseling, P. (ed.) *Notes on Numerical Fluid Mechanics*, pp. 333–342. Vieweg, Philadelphia (1990)
- Negi, P.S., Vinuesa, R., Hanifi, A., Schlatter, P., Henningson, D.S.: Unsteady aerodynamic effects in small-amplitude pitch oscillations of an airfoil. *Int. J. Heat Fluid Flow* **71**, 378–391 (2018)
- Offermans, N., Marin, O., Schanen, M., Gong, J., Fischer, P., Schlatter, P., Obabko, A., Peplinski, A., Hutchinson, M., Merzari, E.: On the strong scaling of the spectral element solver Nek5000 on petascale systems. In: *Proceedings of the Exascale Applications and Software Conference 2016*. ACM, p. 5 (2016)
- Offermans, N.: Aspects of adaptive mesh refinement in the spectral element method. Ph.D. thesis, KTH Royal Institute of Technology (2019)
- Offermans, N., Peplinski, A., Marin, O., Schlatter, P.: Adaptive mesh refinement for steady flows in Nek5000. *Comput. Fluids* **197**, 104352 (2020)
- Peplinski, A., Offermans, N., Fischer, P.F., Schlatter, P.: Non-conforming elements in Nek5000: pressure preconditioning and parallel performance. In: *Proceedings of the International Conference On Spectral And High Order Methods*, London, UK, ICOSAHOM '18 (2018)
- Pranowo, W.S., Behrens, J., Schlicht, J., Ziemer, C.: Adaptive mesh refinement applied to tsunami modeling: Tsunafish. In: *Proceedings of the International Conference on Tsunami Warning (ICTW)* (2008)
- Sato, M., Asada, K., Nonomura, T., Kawai, S., Fujii, K.: Large-eddy simulation of NACA 0015 airfoil flow at Reynolds number of 1.6×10^6 . *AIAA J.* **55**(2), 673–679 (2016)
- Schlatter, P., Örlü, R.: Turbulent boundary layers at moderate Reynolds numbers: inflow length and tripping effects. *J. Fluid Mech.* **710**, 5–34 (2012)
- Schlatter, P., Stolz, S., Kleiser, L.: LES of transitional flows using the approximate deconvolution model. *Int. J. Heat Fluid Flow* **25**(3), 549–558 (2004)
- Schlatter, P., Li, Q., Brethouwer, G., Johansson, A.V., Henningson, D.S.: Simulations of spatially evolving turbulent boundary layers up to $Re_\theta = 4300$. *Int. J. Heat Fluid Flow* **31**(3), 251–261 (2010)
- Sillero, J.A., Jiménez, J., Moser, R.D.: Two-point statistics for turbulent boundary layers and channels at Reynolds numbers up to $\delta^+ \simeq 2000$. *Phys. Fluids* **26**(10), 105–109 (2014)
- Skåre, P.E., Krogstad, P.A.: A turbulent equilibrium boundary layer near separation. *J. Fluid Mech.* **272**, 319–348 (1994)
- Spalart, P.R., Watmuff, J.H.: Experimental and numerical study of a turbulent boundary layer with pressure gradients. *J. Fluid Mech.* **249**, 337–371 (1993)
- Tanarro, Á., Mallor, F., Offermans, N., Peplinski, A., Vinuesa, R., Schlatter, P.: Using adaptive mesh refinement to simulate turbulent wings at high Reynolds numbers. In: *TSFP11 conference*, July 30–August 2, 2019, Southampton, UK (2019)
- Tanarro, A., Vinuesa, R., Schlatter, P.: Effect of adverse pressure gradients on turbulent wing boundary layers. *J. Fluid Mech.* **883**, A8 (2020)
- Vinuesa, R., Bobke, A., Örlü, R., Schlatter, P.: On determining characteristic length scales in pressure-gradient turbulent boundary layers. *Phys. Fluids* **28**(055), 101 (2016a)

- Vinuesa, R., Prus, C., Schlatter, P., Nagib, H.M.: Convergence of numerical simulations of turbulent wall-bounded flows and mean cross-flow structure of rectangular ducts. *Meccanica* **51**(12), 3025–3042 (2016b)
- Vinuesa, R., Hosseini, S.M., Hanifi, A., Henningson, D.S., Schlatter, P.: Pressure-gradient turbulent boundary layers developing around a wing section. *Flow Turbul. Combust.* **99**, 613–641 (2017)
- Vinuesa, R., Negi, P., Atzori, M., Hanifi, A., Henningson, D., Schlatter, P.: Turbulent boundary layers around wing sections up to $Re_c = 1,000,000$. *Int. J. Heat Fluid Flow* **72**, 86–99 (2018)
- Wadcock, A.J.: Investigation of low-speed turbulent separated flow around airfoils. NASA technical report (NASA-CF- 177450) (1987)