Imperial College London
Faculty of Engineering
Department of Aeronautics

A spectral/*hp* element DNS study of flow past low-pressure turbine cascades and the effects of inflow conditions

A thesis submitted for the degree of
Doctor of Philosophy

Andrea Cassinelli

Supervisors:
Prof. Spencer J. Sherwin
Dr. Francesco Montomoli

2020

Abstract

The combined rapid progress of hardware capability and the development of cutting-edge numerical methods have recently provided an opportunity for Computational Fluid Dynamics to be inserted in the design loop, with the role of a virtual wind tunnel. This thesis tackles the development of a validated incompressible Direct Numerical Simulation capability to model complex configurations of interest for the turbomachinery Industry, adopting for the first time the spectral/*hp* element methods implemented in the *Nektar++* software framework.

First, an extensive analysis of the numerical convergence properties is carried out on an open geometry with clean inflow boundary conditions, to establish a set of best practices and relate accuracy and computational cost.

Subsequently, the effect of stochastic and deterministic unsteadiness is analysed in detail, with particular focus on various methodologies to provide physical disturbances, their computational cost and accuracy with respect to reference experimental data. The findings are extended to a range of Reynolds numbers representative of realistic operating conditions, with focus on traditional performance indicators but also unsteady statistics to provide rich insight into the suction surface transition mechanism, which plays a crucial role in the generation of profile losses. As a result, a detailed characterisation of the flow physics is provided in a range of inflow conditions and Reynolds numbers. Excellent agreement with high fidelity experimental data is achieved especially at moderate and high Reynolds numbers, supporting the use of these methodologies in Industry as a preliminary standalone investigation tool.

Declaration of originality

I present this thesis for consideration for the degree of Doctor of Philosophy at Imperial College London exclusively. I certify that it contains the results of my own work and research during my doctoral studies, unless otherwise stated. In such cases, I give proper acknowledgement for methodologies and data that are not my own.

Copyright declaration

This thesis is dedicated to my mother, "mami" Daniela.
Without her endless love I would have never made it this far.

Acknowledgements

First and foremost, my deepest gratitude goes to my supervisors Prof. Spencer J. Sherwin and Dr. Francesco Montomoli. Their guidance, support and encouragement has been invaluable, both from a technical and personal point of view. They have provided many opportunities to be involved in a diverse range of activities, helping me grow as a researcher and an independent thinker.

I would like to thank Dr. Paolo Adami for following my progress through the past four years, with constant and knowledgeable encouragement and for motivating me to work at the best of my abilities. I wish to also thank Dr. Raúl Vázquez Díaz for providing invaluable advice through the project, and always shining light on the next question we were trying to answer. I am also grateful for the guidance of Dr. Luigi Capone during the second year of my PhD, and for long and stimulating conversations with Yuri Frey.

Within the *Nektar++* group, I am immensely grateful to Dr. Chris Cantwell and Dr. David Moxey for their guidance and help in countless occasions, and for always seeing a way through seemingly impossible problems. I would like to thank Dr. Giacomo Castiglioni from whom I have learned a lot, through endless discussions about work, life, future and for always stimulating me to think a bit further. I would also like to acknowledge the help of Dr. Rodrigo Moura and Dr. Douglas Serson at the very beginning of my PhD.
I am grateful to my internal examiner through the various preliminary milestones of the PhD project, Prof. Denis Doorly. His valuable feedback helped me point my research in the right direction. I would also like to thank Dr. Yongyun Hwang, Prof. Michael Leschziner and Dr. Sylvain Laizet for a number of pivotal discussions.

I gratefully acknowledge the financial support by Rolls-Royce plc. through the project and beyond. This has enabled me to attend training courses, travel to many conferences and it has provided all that was necessary for the success of the project.
This research would not have been possible without the access I have been granted to supercomputing facilities at Imperial College London and the UK national Supercomputing Service ARCHER. I would like to acknowledge all the work the Imperial College Research Computing Services have done in the past few years, as well as the UK Turbulence Consortium under grants EP/L000261/1 and EP/R029326/1.

On a more personal note, there are a few people who have made my years in London absolutely unforgettable. I want to thank Michela for dragging me along to countless adventures, and for sharing all the good and bad times, when nothing seemed to go according to plan. Special thanks also go to my other office mates Jan, Julian (the "Nektarines") and Marius for all the long (and often animated) lunch discussions and many evenings together; to Matteo and Pepe, friends and flatmates, for having shared so many fun times in our Elsham Road flats, and having offered each other support through life's best and worst moments.

I am deeply grateful to Asha, for having hosted me through lockdown, and for her constant support and affection in the most difficult part of 2020 and beyond.
I am forever indebted to Anya, without whom I would not have overcome some of the most difficult times in my life, coming out stronger and with a more positive outlook on the future. I am very grateful for her patience and constant support through the (long and intensive) writing process. I truly look forward to a bright future together.

Finally, my family. My parents Daniela and Sergio, and my brother Tommaso have been there for me at every step of the way, supporting every difficult decision and sharing so many of life's best memories during our travels around the world, and at our home in Voghera. Our lives changed forever in 2018 when Daniela left us way too soon, but her love, energy, livelihood and positivity will continue to inspire us and drive us forward every single day.

Contents

List of Figures

List of Tables

Nomenclature

Symbols

α	Flow angle
$\boldsymbol{\omega}$	Vorticity
Δt	Time step
Δx_i^+	Mesh spacing in wall units
δ	BL displacement thickness
ϵ	Turbulence dissipation rate
$\hat{Q}$	SVV kernel
κ	Wave number
κ_η	Kolmogorov length scale
$\mathbf{u}$	Velocity
$\mathcal{T}$	Turbulence time scale
μ	Dynamic viscosity
μ_{SVV}	SVV diffusion coefficient
ν	Kinematic viscosity
Ω	Computational domain
ω	Total pressure loss coefficient
Φ	Flow coefficient
ϕ	SPM concentration field
ρ	Density
τ_w	Wall shear stress
St	Strouhal number
θ	BL momentum thickness
φ	Wake passing phase
ξ	SPM interface thickness
C	Blade chord length
C_d	Drag coefficient
C_f	Skin friction coefficient
C_p	Static pressure coefficient
C_{ax}	Blade axial chord length
d	Cylinder diameter
f	frequency
F_{red}	Reduced frequency
H	BL Shape factor
h	Local element size
I	Momentum forcing intensity
k	Turbulence kinetic energy
L_t	Turbulence length scale
L_z	Spanwise domain size
M	SVV cutoff ratio
n	Blade wall-normal distance
N_{turb}	Number of modes for synthetic turbulence generation
N_z	Number of Fourier planes

P	Polynomial order
p	Pressure
P^k	Turbulence production
P_b	Distance between bars
P_y	Blade pitch
Re	Reynolds number
s	Blade surface distance
S_0	Suction surface perimeter
t, T	Time
U_b	Bar speed
U_∞	Reference speed
Pe	Péclet number

Acronyms

BF	Body Forcing
BL	Boundary Layer
CFD	Computational fluid dynamics
CFL	Courant–Friedrichs–Lewy
CG	Continuous Galerkin
CI	Clean Inflow
DG	Discontinuous Galerkin
DNS	Direct Numerical Simulation
DoF	Degrees of Freedom
FD	Finite Difference
FR	Flux Reconstruction
HIT	Homogeneous Isotropic Turbulence
HOM	High Order Methods
HPC	High Performance Computing
HPT	High Pressure Turbine
IW	Inflow Wakes
KH	Kevin-Helmholtz
LE	Leading Edge
LES	Lage Eddy Simulation
LPT	Low Pressure Turbine
PSD	Power Spectral Density
RANS	Reynolds-Averaged Navier-Stokes
SPM	Smoothed Profile Method
SVV	Spectral Vanishing Viscosity
TE	Trailing Edge
TI	Turbulence intensity
TKE	Turbulence kinetic energy

Subscripts and superscripts

$\overline{q}$, $\langle q \rangle$	Time-averaged quantity
$\tilde{q}$	Phase-locked quantity
q'	Fluctuation around mean value
q^A	Mass flow averaged quantity
q^M	Mixed-out averaged quantity
q_1	Quantity at the inlet measurement plane
q_2	Quantity at the exit measurement plane
q_s	Static quantity
q_t	Stagnation quantity
q_w	Quantity at the wall
q_∞	Inlet reference value
q_{sep}	Quantity at the separation point
q_{TE}	Quantity at the trailing edge

Chapter 1

Introduction

1.1 Motivation

The civil aviation industry is powered by high bypass turbofan engines, where about 80% of the total thrust is generated by the fan. The fan itself is powered by the low pressure turbine (LPT), which accounts for roughly 20-30% of the engine weight, since several stages are needed to achieve the required work output.

The design of these components is an extremely challenging task that needs to reconcile many requirements, and it is the result of decades of research and attempts. Nowadays the process heavily relies on computational methods, which commonly drive the main design decisions of new technological advances. Significant progress has been so far achieved in improving the efficiency of LPTs: a 1% improvement corresponds to a 0.5-1 % reduction of specific fuel consumption [56]. Current design trends are moving towards the reduction of machine dimension and the increase of revolution speed, which in turn results in lower weight and higher power density. However, as recently highlighted by Sandberg & Michelassi [144], the spectral gap existing between deterministic unsteadiness (associated vortex shedding of blade rows and their rotor-stator interaction) and stochastic unsteadiness (associated with the energy-containing eddies of turbulence) is progressively reduced. This trend implies an interaction between the two sources of unsteadiness, which poses a challenge for conventional approaches to turbulence modelling.

To set the context, Figure 1.1 shows the range of frequencies captured by different modelling strategies. Reynolds Averaged Navier-Stokes (RANS) and Unsteady RANS (URANS) approaches capture deterministic phenomena but do not resolve the scales of turbulence beyond the inertial subrange. This modelling approximation holds valid in cases where the deterministic and stochastic unsteadiness are characterised by non-overlapping frequencies, thus in absence of a spectral overlap [170]. This problem is of considerable importance, given that the interaction of these two mechanisms is believed to be linked to the majority of irreversibility generation [144].

While RANS and URANS will continue to be the backbone of computer-driven design, the technological advances of the last decades have created an opportunity for high-fidelity

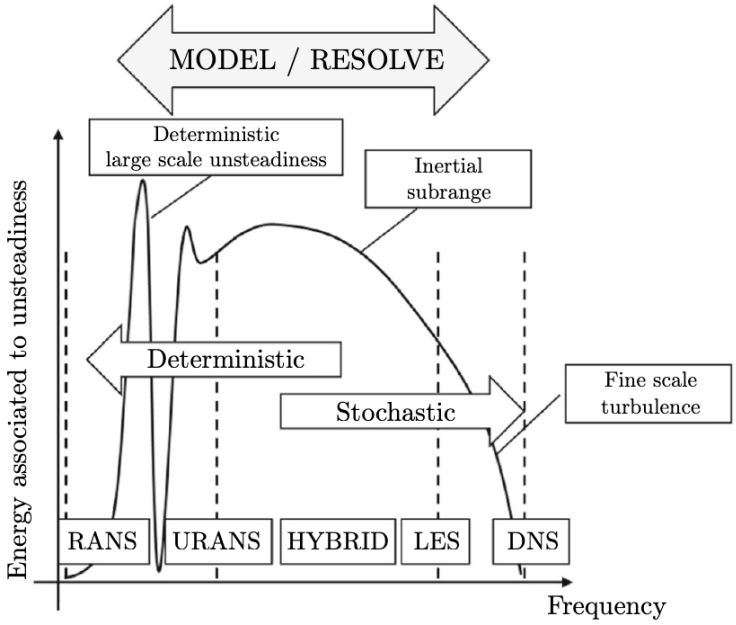

Figure 1.1: Sketch of the turbulent spectrum showing different blade passing frequency ranges. Adapted from [144].

Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS) to be inserted in the design loop as a virtual wind tunnel. The benefit is that of resolving the physics with greater level of accuracy, therefore providing an increased understanding which is necessary to further advance the technological state of the art. A well-known NASA report [157] provides a series of recommendations in a roadmap to enable Computational Fluid Dynamics (CFD) to introduce a radical improvement in the prediction of critical flow phenomena. One of many important remarks is related to the fast-paced advances in hardware capabilities, which will involve heterogeneity in the hardware (hybrid of multiple processors and accelerators), as well as the memory architecture and the network interconnect. Laskowski et al. [87] further detailed how the developing and future interaction of High Performance Computing (HPC), high-fidelity CFD and high-order unstructured grid algorithms has the potential to allow for simulation-based research, analysis and design capability in the context of turbomachinery. In particular, the requirement for effective and trusted validation is highlighted over a broad design space. Another recent review [173] discusses a number of novel numerical approaches suitable for turbomachinery simulations, together with the challenges the community is facing.

The recent advances in hardware technology and the development of novel numerical methods have made LES and DNS affordable, and now widely used to reproduce simple turbomachinery configurations with high degree of accuracy. Since in real gas turbine components the rotor and stator count are usually prime numbers, simulations of a realistic configuration are currently still unfeasible. For this reason, most of the studies employing high fidelity CFD focus on simplified configurations, e.g. linear cascades. Though not fully

representative of realistic geometries, more than 80% of the losses occur in the vicinity of the boundary layer (BL). Therefore, the use of simplified configurations does not hinder the possibility of gaining improved understanding of the flow physics in a more realistic environment.

1.2 Literature review

This review focuses mainly on a description of the fundamental physical challenges to be tackled in the analysis of gas turbine components, more specifically low pressure turbines. Subsequently, after mentioning the main efforts towards the development of high-order solvers, the focus is placed on the most important contributions deriving from the pioneering adoption of novel numerical approaches.

Low pressure turbines have high aspect ratios, which means that the majority of losses are generated in the midspan (profile losses) and they depend particularly on the state of the suction surface boundary layer [38], given the typically thin trailing edges. Especially at low Reynolds numbers, the boundary layer transition mechanism affects the presence of a separation bubble, and whether the separation is open or closed. An open separation bubble results in less flow turning, and therefore less power output. The ability to correctly predict the flow physics associated with the suction surface boundary layer in a realistic flow environment is therefore of great importance and constitutes a challenge for computational methods. RANS methods typically struggle to provide accurate estimation of pressure losses, especially in off-design conditions. This factor, together with the relatively low Reynolds and Mach numbers characteristic of LPTs (compared to high pressure turbines or high pressure compressors), provides an opportunity to leverage the advantages of high fidelity simulations.

1.2.1 Suction surface boundary layer transition mechanisms

The wake passing interactions are the primary form of unsteadiness in turbomachinery, combined with the resulting background turbulence generated by upstream rows of rotating components. A detailed review by Hodson & Howell [60] summarises the wake-induced boundary layer trasition mechanisms in LPTs, briefly reported here. A schematic representation of rotor-stator interaction is shown in Figure 1.2: the wake deficit caused by the stator results in an effect analogous to a negative jet. The convection of the wake segment through the passage is characterised by bowing, reorientation, elongation and stretching [159]. These characteristics of the wake kinematics are visible from the Laser Doppler Anemometry (LDA, also called Laser Doppler Velocimetry) measurements by Stieger & Hodson [163], shown in Figure 1.3.

The wake passing has an important effect on the suction surface, where it modifies the transition mechanism and, if present, the separation bubble dynamics. The ability to

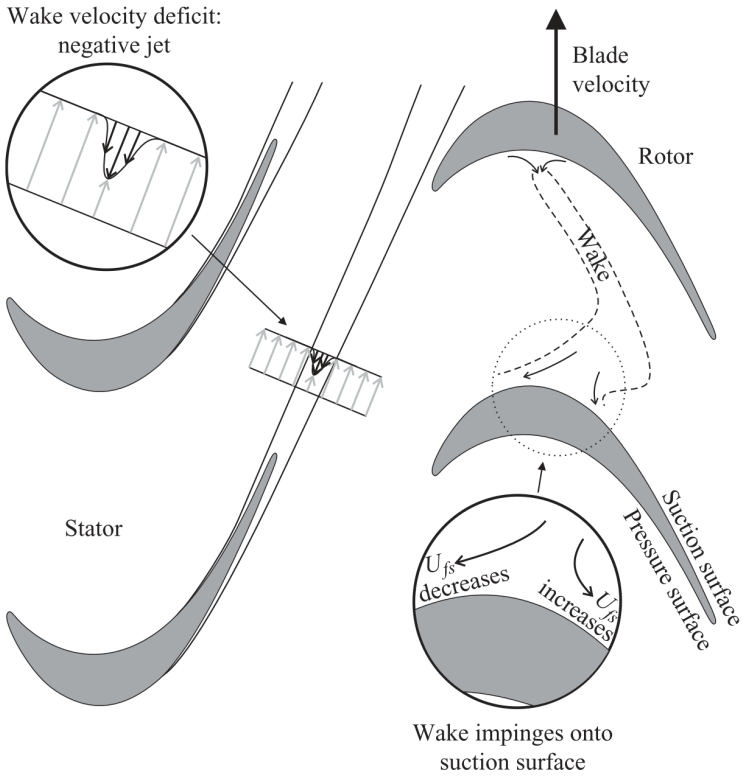

Figure 1.2: Sketch of the wake dynamics due to rotor-stator interaction. From Coull & Hodson [29].

predict boundary layer transition locations accurately on turbomachinery airfoils is critical to accurately evaluate aerodynamic performance. This motivated extensive research efforts aimed at the development of transition models for RANS solvers [137, 138].

According to Hodson & Howell [60], wake-induced transition may occur in three different ways, depending on the Reynolds number and on the blade geometry.

Wake-induced BL transition in attached flow

At high Reynolds number, transition may take place before the occurrence of laminar separation, induced by the high turbulence carried by the wakes; between subsequent wake cycles, the flow remains attached. In this context, the first stages of natural transition are bypassed and turbulent spots are formed within the boundary layer. Having contributed to the analysis of the topology of turbulent spots, Schubauer & Klebanoff [151] also highlighted the existence of a calmed region trailing each spot (shown in the sketch of Figure 1.4). This is a laminar-like region with a full velocity profile [31, 56, 153] and, as such, it can resist transition and separation. Since the trailing edge of the calmed region travels at 20-30% of the free-stream velocity (depending on the local pressure gradient), the calmed region lasts longer than the turbulent region. Schulte & Hodson [153] argued that maximum performance is obtained if the blade design is such that as the new wake

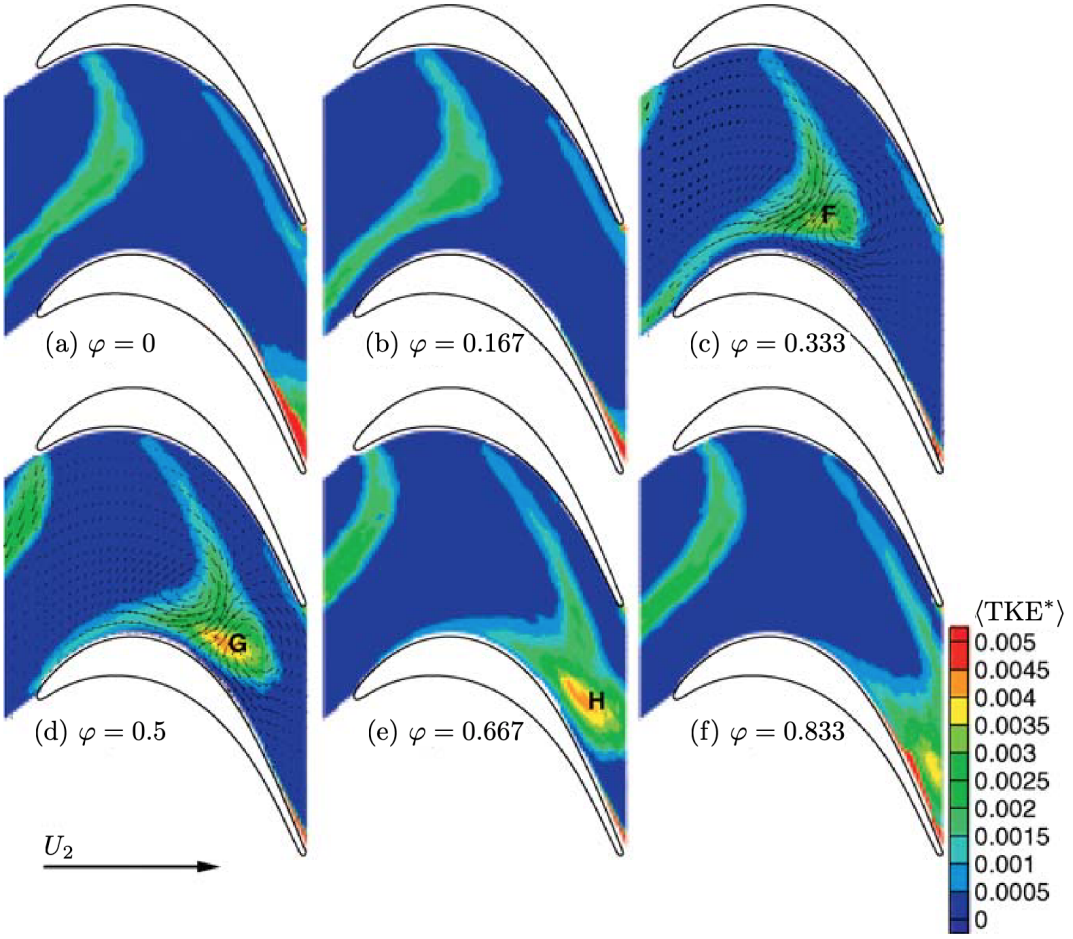

Figure 1.3: LDA measurements of phase-averaged 2D turbulence kinetic energy $k_{2D} = \frac{1}{2}\langle u'^2 + v'^2 \rangle$ nondimensionalised by isoentropic exit velocity. Adapted from Stieger & Hodson [163].

arrives, the calmed region is advected past the trailing edge. Significant research efforts were further dedicated in recent years to in-depth analysis of the wake-induced transition mechanism, discussed in further detail in a later paragraph.

Wake-induced BL transition in attached flow on a separation bubble

At intermediate Reynolds numbers, laminar flow separation is more likely. The effect of periodic disturbances is beneficial in that it reduces the impact of the separation bubble on profile losses. As shown by Schulte & Hodson [152], after the turbulent flow carried by the wakes increases the losses, the calmed region that follows introduces a significant reduction of the momentum thickness at the trailing edge, proportional to profile losses [38]. Before the subsequent wake cycle impinges on the aft portion of the suction surface, the separation bubble tends to re-establish its steady-state behaviour. Furthermore, Stieger & Hodson [162] showed that although the wake-induced turbulent strip is characterised by large levels of dissipation, the calmed flow following the impinging wake and subsequent initial stages

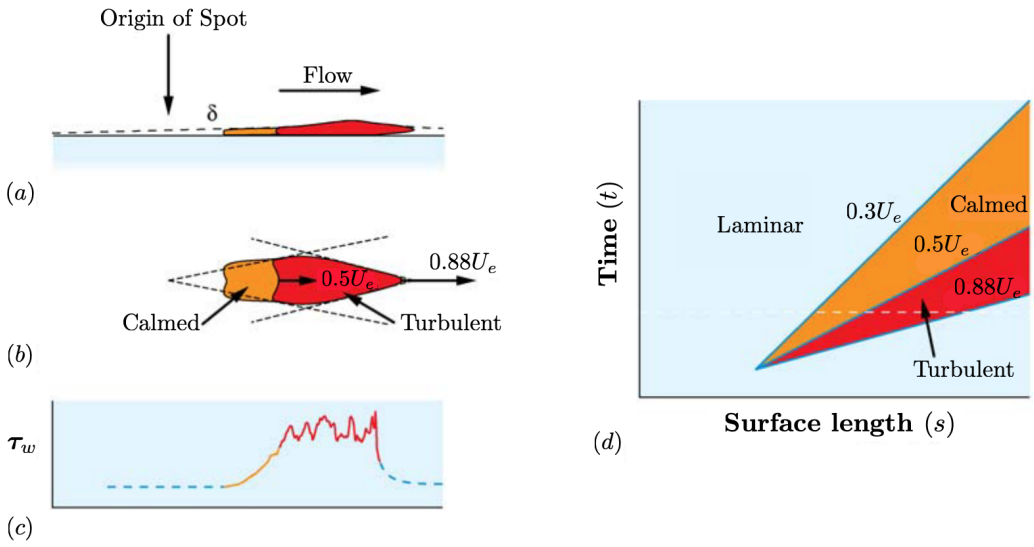

Figure 1.4: Characterisation of a turbulent spot, showing (a) the wall-normal elevation, (b) the shape in the wall-parallel direction, (c) the wall-shear stress along the centreline and (d) a space-diagram. Adapted from Hodson & Howell [60].

of formation of a new separation bubble are characterised by low levels of dissipation. Therefore, overall effect is a reduction in time-averaged losses.

Wake-induced BL transition in separated flow

When the Reynolds number is low, flow transition may not occur prior to separation, and the effect of wake passing leads to transition through a different mechanism compared to the higher flow regimes discussed in the previous paragraphs. This is due to the highly unstable inflexional profiles of the separated flow region. The mechanism of transition in the laminar separation bubble of a T106A low pressure turbine cascade was first described by Stieger & Hodson [163], and schematically represented in Figure 1.5. The negative jet previously shown in Figure 1.2 impinges on the bubble by accelerating the flow downstream of the wake, and slowing down the flow trailing the wake (top of Figure 1.5). The accelerating flow in the outer part of the boundary layer intensifies the shear existing across the separated region. Concurrently, the wall-normal component of the negative jet introduces perturbations in the shear layer (Figure 1.5b), promoting the development of an inviscid Kelvin–Helmholtz roll-up (Figure 1.5c). Since the wake travels faster than the roll-up vortices, it convects over the remaining portion of the separation bubble and triggers the emergence of further roll-up vortices (Figure 1.5d). The trailing region of the wake leaves behind a calmed region which decreases profile losses. Finally, the separation bubble is then progressively formed again (Figure 1.5e) and the cycle repeats itself.

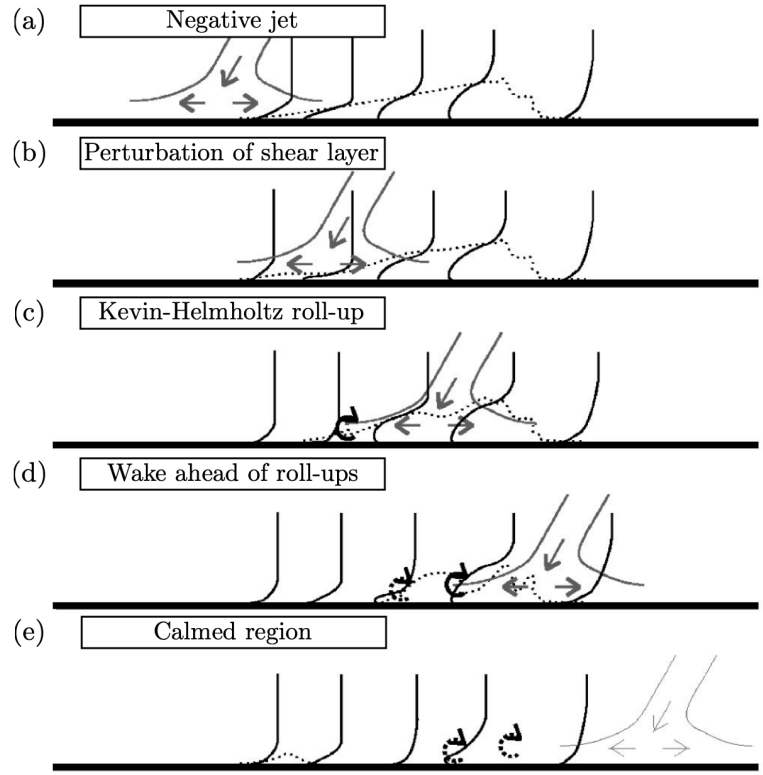

Figure 1.5: Sketch of the wake-induced boundary layer roll-up mechanism. Figure from Stieger & Hodson [163], further adapted from Hodson & Howell [60].

1.2.2 The development of high-order codes and applications to turbines

The tendency in moving toward high-order methods (HOM) for high fidelity simulations owes to their superior dispersion-diffusion properties, offering potential to achieve better accuracy then traditional second order schemes at comparably reduced computational cost. Importantly, comparisons between low-order schemes and high-order methods should be carried out looking at the cost required to achieve a given degree of accuracy, thus not evaluating method efficiency on the same computational mesh [185].

The change in computational technologies and level of fidelity poses new challenges: the new algorithms require specific activities to develop guidance on their best use with verification and validation for consolidated flow problems of industrial relevance, especially for the turbomachinery field; the International High-Order Workshop [185, 26] is an example of these activities.

High-order software frameworks

A number of high-order software frameworks have been developed in the past 20 years, proposing different types of numerical schemes and capabilities to tackle various types of industrial problems, each with specific advantages and limitations. A recent review by Tyacke et al. [173] mentions a number of important examples.

Some of the main high-order finite difference (FD) codes are: *HiPSTAR* [145] (extensively applied to turbomachinery), *OpenSBLI* [67] and *3DNS* [187]. Within the spectral element family, some of the most well-known frameworks are *Nektar++* [118] and *Nek5000* [6]. *Nektar++* (adopted through this work, and further discussed in detail in Chapter 2) supports both continuous (CG) and discontinuous Galerkin (DG) formulations, as well as a variety of solvers. DG solvers are being actively developed at MIT [45], Cenaero (*ARGO*) [22], and NASA AMES [48].

Within the class of finite element methods, the type of lifting operator that propagates information across elements determines the specific family of methods. The Flux Reconstruction (FR) approach was recently introduced by Huynh [64]. Within this formulation, a multi-parameter family of provably stable schemes was identified [177, 178]. FR unifies different numerical schemes: by appropriate choice of correction terms, it was shown to recover DG, spectral difference and spectral volume methods [65], offering the advantage of simplicity over alternative formulations. The connections between FR and DG were also recently extensively explored [37, 105].

Among the FR software frameworks, *PyFR* [194] has been successfully applied to a range of fluid dynamics problems, both of fundamental nature [66] and of aeronautical interest [179, 127]. Among other frameworks currently under active development, *GFR* [160], *Mu2s2T* [15], *hpMusic* [183] and a space-time extension of FR (STEFR) [96] should be mentioned.

High fidelity studies of linear cascades

The adoption of high-fidelity CFD methods in turbomachinery has allowed for significant advances in the understanding of the flow physics, owing to the unprecedented level of detail that can be simulated. However, several seminal contributions were made in the early 2000s through the use of low order codes, and they are mentioned in this section for completeness.

Clean inflow studies Most studies analysing cascades with clean inflow boundary conditions are part of a validation and verification activity linked to software development. In fact, this case bridges industrial geometrical complexity with relative semplicity in setting up the numerical simulations, and the absence of uncertainty on the imposition of unsteady boundary conditions. One of the examples is a study by Raverdy et al. [142], where a T106 LPT linear cascade was analysed with a compressible formulation. Particular attention was placed on the suction surface separation bubble, and Fourier analysis applied to the separation and reattachment point highlighted the existence of coupling between the bubble and the trailing edge vortex shedding. More recently, Marty et al. [102] investigated a high-order MUSCL interpolation implemented in the *elsA* software framework [20], a multi-purpose CFD solver developed by Onera. Among the test cases analysed, a T106C cascade was simulated at third and fifth order of accuracy, highlighting the robustness of the method and demonstrating the capability of the solver to simulate

complex flow geometries. Vincent et al. [179] performed simulations of a linear LPT cascade using *PyFR*, resolving the full span and five repetitions in the pitchwise domain, thus setting an unprecedented milestone in the turbomachinery field by solving for 113 billion Degrees of Freedom (DoF) and showcasing very close agreement in the mean pressure distribution against experimental data. Medic et al. [104] analysed a compressor cascade reporting very good agreement with predicted losses at design conditions. At off-design conditions, laminar separation and transition occur downstream compared to the experiments, thus impacting the accuracy of loss prediction. Bhaskaran et al. [13] compared the accuracy of two different high-order LES methods on a T106A LPT cascade, obtaining very close agreement with experimental measurements of pressure distribution and wake losses. Fernandez et al. [45] discussed the implementation and validation of a high-order ILES methodology in a compressor cascade. The solver was demonstrated to accurately and efficiently capture the transition process up to $Re = 460000$, highlighting one of the advantages of high-order methods over low-order schemes owing to their low numerical dissipation. In particular, this allows to accurately predict transition not because of the small length and time scales, but rather for the ability to capture small instabilities that would otherwise be damped. Cassinelli et al. [23] discussed the resolution requirements of a T106A linear cascade analysed with the spectral/*hp* element methods implemented in the *Nektar++* framework; the main results of this study are presented in detail in Chapter 3. Recently, Bolinches-Gisbert et al. [15] focused on the Reynolds sensitivity of a research LPT cascade as part of validation efforts of the *Mu2s2T* FR solver. Comparison against a comprehensive set of experimental data was presented, showcasing very good agreement in blade wall distributions and wake statistics.

Inflow turbulence In the past 20 years, significant research efforts were dedicated to the development of numerical techniques to introduce stochastic unsteadiness; the main and most successful approaches are reviewed in more detail in Chapter 2, while this section is concerned with the main applications of such techniques to turbomachinery components.

Among the early DNS studies of flow past a turbine blade, Kalitzin et al. [73] compared results obtained with clean inflow, wake passing data from Wu & Durbin [196] and inflow turbulence. In particular, the latter was generated via precursor simulation of homogeneous, decaying turbulence. The natural transition mechanism on the suction surface in the free inlet case is modified by the convecting disturbances, which trigger bypass transition. Zaki et al. [206] performed DNS simulations of flow past a compressor blade, introducing disturbances from a precursor data set of decaying homogeneous isotropic turbulence (HIT), analysing the effect of increasing free-stream turbulence intensity on the transition mechanism. Streamwise-elongated disturbances were amplified upstram of the transition point. On the pressure side, with increasing inflow turbulence level, inner streak instability was replaced with bypass transition via secondary streak instability. On the suction surface, at low turbulence intensity the separation persisted and was modulated by boundary layer streaks. As the intensity was increased, turbulent spots were observed,

and progressively separation was completely suppressed. Medic & Sharma [103] conducted LES simulations of three LPT airfoils in a range of Reynolds numbers, comparing low free-stream and high-free-stream turbulence. Particularly at low Reynolds numbers, significant deviations from measurements were reported.

More recently, the first compressible DNS of a linear LPT cascade was performed by Sandberg et al. [145] with the high-order finite difference code *HiPSTAR*, showcasing excellent agreement with experimentally measured pressure coefficient and wake loss profile. A simple inflow turbulence generation method based on generating harmonic waves at a few discrete wave numbers was described. Within $20\%C_{ax}$ from the inlet, the characteristic broadband spectrum was obtained thus supporting the validity of the proposed methodolody. The same turbulence generation mechanism was employed by Wheeler et al. [186] to perform the first DNS of a HPT vane at transonic conditions. The introduction of disturbances promoted the formation of near-wall streaks thus augmenting the surface heat transfer. Suction surface transition was found to occur via intermittent development of Kelvin–Helmholtz, roll-ups triggered by upstream-propagating pressure waves generated at the trailing edge.
Garai et al. [49] analysed the T106A and T106C linear cascades using a high-order entropy-stable DG spectral element method solver with a space-time formulation. The inflow turbulence was generated by adapting a linear forcing technique for the DG framework. This approach was further applied to a high-pressure turbine cascade [50] with high levels of inflow turbulence. The presence of disturbances modified the suction surface dynamics by promoting transition via bypass mechanism. The numerical results matched experimental measurements at lower levels of inflow turbulence. Comparison with RANS simulations demonstrate the superiority of the proposed DG approach in predicting the separation location and boundary layer thickness in the context of high-Reynolds number flows. As part of their validation efforts, Bhaskaran et al. [13] carried out the first high-order solution of a HPT vane to include stationary turbulence-generating bars in the computational domain. As a result, the blade heat transfer distribution obtained with two different solvers were in excellent agreement. More recently, Cassinelli et al. [24] compared two different approaches to introduce inflow turbulence in a research LPT cascade at various intensity levels; the findings of this study are presented in Chapter 4.

Very often, the turbulence length scale measured experimentally is significantly larger than the computational domain, and therefore reduced in the numerical experiments. The first attempt at estimating the effect of turbulence intensity and length scale on the blade boundary layers and heat flux was carried out by Pichler et al. [134]. A HPT vane at realistic engine conditions ($Re_{is} = 540000$, $M_{is} = 0.92$) was simulated with *HiPSTAR* introducing inflow turbulence via a digital filter method [84, 199, 169]. Increased levels of turbulence intensity and length scale were found to promote greater heat flux levels and earlier boundary layer transition as well as its intermittent behaviour. Recently, Zhao et al. [207] performed a highly resolved LES of a HPT vane at realistic Reynolds and Mach numbers, analogously adopting a digital filter approach and observing the effect of

different turbulence length scale and intensity on the transition mechanism. At relatively low inflow turbulence levels, transition occurred via streak instability. Evidence of both varicose instability (in the adverse pressure gradient region) and sinuous instability (in the weak favourable pressure gradient region) was presented. At the highest levels of free-stream turbulence, no clear streak instability path was observed.

Finally, among the studies focusing on compressor flows with stochastic inflow disturbances, Przytarski & Wheeler [141] discussed the accuracy of the estimation of entropy generation rate, and the effect of resolution on the various loss generation mechanisms; the latter was found to depend strongly on the mechanism and location of transition.

Wake passing effect Rotor-stator interactions are the key feature of turbomachinery flows, and commonly investigated experimentally by means of a set of moving bars situated upstream of the leading edge. The very first computational studies investigating the effect of wake passing were carried out by Wu & Durbin [198, 196], who presented evidence that incoming wakes are responsible for the formation of longitudinal structures on the pressure side. Subsequently, a number of DNS studies were performed on LPTs, providing further contributions to the understanding of the effect of incoming wakes on the pressure and suction side boundary layers [110, 111, 189, 191]. For example, Michelassi et al. [110] performed an LES of T106 airfoil at $Re = 148K$, observing noticeable differences between DNS and LES in the prediction of transition. This analysis confirmed that LES can have difficulties in predicting unsteady transition, also later highlighted by Durbin & Wu [43].

The wake dynamics in the mean passage has been less investigated than the effect of wakes on the suction surface boundary layer. Michelassi & Wissink [112] conducted a study using incompressible LES to identify the turbulence kinetic energy (TKE) generation mechanism outside of the blade boundary layer. The largest TKE production was found to occur at the wake location where the maximum strain tensor eigenvalue approximately coincides with the location where the wake and the direction of compression are aligned. At the higher $Re = 148K$, the turbulent kinetic energy peak is closer to the suction side boundary layer and speculated to improve its stability. These findings provide important information for the design of suction-side-separation-free LPT blades.

From an engineering point of view, improvements in the efficiency of the machine can be largely obtained by optimising profile losses, but a highly resolved and validated DNS is required for this purpose. Michelassi et al. [109] built on top of previous work [145] to analyse the effect of various types of inflow disturbances on losses in a series of compressible DNS simulations. In particular, they observed the effect of three reduced frequencies and two Reynolds numbers ($Re_2 = 60K, 100K$) on various types of loss indicators. The behaviour of high reduced frequencies ($F_{\text{red}} = 1.22$, corresponding to 4 bars per pitch) was found to be most effective in suppressing the suction surface separation bubble by manifesting a very similar behaviour to high levels of inflow turbulence, due to the constant-area mixing prior to the leading edge. The largest difference between mixed-out total pressure losses and profile losses calculated via Denton's approach [38]

was found at low levels of reduced frequency ($F_{\text{red}} = 0.31$, corresponding to 1 bar per pitch): in this case, the wakes remain distinct within the blade passage and experience significant wake distortion losses. Further LES simulations [108] (previously compared to DNS at the same regime, to ensure that the numerical approximations introduced would not significantly affect the accuracy of the flow physics) demonstrated the importance of the reduced frequency and flow coefficient on losses. These important design parameters affect the frequency with which incoming wakes impact on the cascade as well as the wake inclination. The normal distance between incoming wakes was found to correlate very well with mixed-out losses, suggesting that wake merging (which produces a free-stream turbulence-like behaviour) is in general beneficial for the reduction of unsteady losses. This approach can importantly be adopted to inform design optimisation, as the various combinations of reduced frequency and flow coefficient are effectively a consequence of design parameters like revolution speed and flow-through velocity.

Hammer et al. [57] more recently performed LES of a T106A cascade where the effect of bar rotation was analysed in conjunction with their usual translational motion. Clockwise rotation of the bars was found most effective in suppressing the suction surface separation bubble but it presented the highest mixed-out losses, likely due to wake distortion.

When analysing rotor-stator interation experimentally, it is difficult to differentiate between the mean effect of the wake acting as a negative jet and the effect introduced by small-scale fluctuations. This challenge is relatively unexplored in the literature and to the author's knowledge only two studies are available on the subject. Wissink [192] first tackled this research question, focusing on comparing the impact from the wake data of Wu & Durbin [196] and the same wake without velocity fluctuations. As a result, in both simulations a Kelvin–Helmholtz instability of the separated shear layer was triggered by the impinging wakes. The small-scale fluctuations carried by the convective wake influenced the transition to turbulence of the separated shear layer: without external disturbances, the suction side boundary layer retained two-dimensionality and remained laminar almost up to the trailing edge. Thus, the unsteady flow on the suction surface was shown to be dependent on the turbulence intensity and the length scale of passing wakes. Sarkar [146] later analysed a T106A profile under the influence of migrating wakes, analysing the difference of a 2D and 3D precursor simulation of flow past a cylinder to provide the wake deficit mechanism and focused particularly on the kinematics of the migrating wakes as well as the unsteady boundary layer dynamics. Analogously to Wissink [192], the Kevin-Helmholtz instability of the separated shear layer was shown to be triggered by the low-frequency effect of the wake. The small eddies characteristic of fully three-dimensional wakes have significant effects on the levels of turbulent production and dissipation in the aft portion of the suction surface.

The wake passing effect has been analysed mostly in the context of LPT, but some studies also focused on compressor cascades. Among the experimental work carried out in this context, a study by Heigenfeld & Pfitzner [59] is often adopted as a reference. Zaki et al. [202] introduced wakes generated from a separate simulation of flow past a circular

cylinder. The periodic disturbances periodically triggered transition on the suction surface via Kevin-Helmholtz instabilty; on the suction surface, the boundary layer remained attached through the wake passing cycle. Wissink et al. [193] focused on analysing the effect of wake strength on BL transition and separation. Stronger wakes were found to intermittently suppress separation, while weaker wakes marginally affected the size of the separation bubble. More recently, Leggett et al. [90] compared very refined LES and RANS results in a range of off-design configurations, revealing a similar loss estimation between RANS and LES but with different contributions, according to Denton's analysis. The unsteady test case with incoming wakes showed limited sensitivity of the loading and skin friction distribution to the incoming disturbances, but nonetheless was subject to intermittent separation.

Gap-size effect The effect of incoming wakes is commonly studied both numerically and experimentally by introducing upstream moving bars. This approach can produce realistic wakes and thus be adopted to analyse the impact of disturbances on the transition mechanism. However, it cannot reproduce the potential effect that occurs in the axial gap between consecutive stator and rotor rows. Pichler et al. [133] carried out the first simulation of stator-rotor interaction in a realistic LPT vane, analysing the effect of different axial gaps. The turbulence kinetic energy level entering the rotor row was reduced by a factor of three when doubling the gap size. The stronger disturbances of the small axial case entirely suppressed the small time-averaged suction surface separation bubble present in the large axial gap case. However, the overall loss in the former case was higher, due to the TKE amplification in the blade passage. The rise in loss due to smaller axial gap was also reported by Przytarski & Wheeler [140]. The study focused on a multi-stage compressor environment, and the rotor-stator interaction was modelled via the introduction of a wake recycling method, as opposed to explicitly simulating two stages.

The importance of grid resolution

Turbomachinery flows are extremely complex, featuring a range of flow phenomena in a highly turbulent flow environment. As pointed out by Sandberg & Michelassi [144], the claim of DNS should be based on more stringent criteria than simply mean pressure distributions. The Authors subdivided grid requirements in three regions, dominated by different turbulence scales: development region upstream of the blade, boundary layers and blade wake. In-depth details are provided for each of these areas [144]. Supporting the claim that mean blade wall distributions are often not sufficient as convergence indicators, Pichler et al. [135] focused on assessing the impact of various extremely well-resolved numerical setups on turbulence quantities, namely turbulence kinetic energy, turbulence production and dissipation. All setups adopted provided accurate estimation of the mean pressure distribution and wake loss profile, which are conventionally adopted as convergence indicators. For accurate estimation of TKE and turbulence production profiles the

most coarse grid is deemed sufficient, while dissipation (which requires capturing the smallest scales of turbulence) required a grid count increase of a factor of 6. This resulted in an order-of-magnitude increase in computational cost, derived from the combined effect of problem size and timestep limitation due to the use of an explicit solver.

1.3 Objectives, research questions and thesis outline

This thesis presents the fist extensive numerical investigation of the spectral/*hp* element framework *Nektar++* [21, 118] applied to low pressure turbine cascades using unstructured grids. The principal focus is twofold: on one side, to develop a full understanding of the computational requirements to use this class of numerical methods as a virtual cascade for industrial test cases, mapping the computational cost and accuracy together with the necessary technology development to simulate realistic flow physics. On the other side, the project focuses on building on the significant advances of high-fidelity CFD in the last 20 years to contribute to an improved understanding of the physics and the modelling linked to the effect of different inflow conditions that represent deterministic and stochastic sources of unsteadiness.
More specifically, the main research questions addressed through this work are the following:

1. **Development of incompressible DNS capability for LPTs using spectral/*hp* element methods**. What is the best numerical configuration that combines computational efficiency and high accuracy? How can the numerical properties of the method be leveraged in the context of flow past low pressure turbines?

2. **Resolution requirements**. What is the effect of mesh resolution and expansion order, domain size, and time-averaging extent on the accuracy of the results?

3. **Inflow disturbances methodology**. How can inflow disturbances (stochastic disturbances) be modelled physically and with minimal algorithmic overhead?

4. **Wake passing methodology**. How can the wake passing effect (deterministic disturbances) be modelled physically and yield accurate predictions compared to experimental measurements?

5. **Reynolds sensitivity**. What is the Reynolds sensitivity of a LPT cascade with clean inflow and inflow disturbances? How accurate is the numerical prediction of losses across various operating regimes?

6. **Reynolds sensitivity with wake passing**. What is the Reynolds sensitivity of a LPT cascade subject to wake passing? How accurate is the numerical prediction of losses across various operating regimes?

The various research questions are discussed through the methodology and results chapters. The connections between the various topics are highlighted in the roadmap of Figure 1.6, color coded for clarity.

The present chapter has provided a summary of the main recent development efforts towards high-fidelity CFD capability for industrial-scale cases, as well as the most relevant studies of the flow physics in turbomachinery components using high-fidelity CFD. The rest of the thesis is structured as follows: Chapter 2 provides a description of the spectral/*hp* element methods, the numerical techniques adopted to simulate high-Reynolds number flows and the technologies integrated into the *Nektar++* framework to model inflow turbulence and the wake passing effect. Chapter 3 constitutes the core of the first principal focus of this work and the first two research questions, i.e. to provide an understanding of the resolution requirements to accurately and efficiently model low-pressure turbine cascades at moderate Reynolds numbers. The findings of this Chapter provide the indications to inform the choice of numerical configurations adopted in the rest of the thesis. Chapters 4 focuses on the effect of two different methodologies to introduce inflow disturbances (stochastic unsteadiness). Chapter 5 and 6 analyse the Reynolds number sensitivity in a realistic operating range: Chapter 5 focuses on clean inflow boundary conditions and the introduction of stochastic inflow disturbances; Chapter 6 builds on the previous one, combining the analysis of Reynolds sensitivity with the wake passing effect (deterministic unsteadiness). Finally, Chapter 7 draws the main conclusions and suggestions for future research directions.

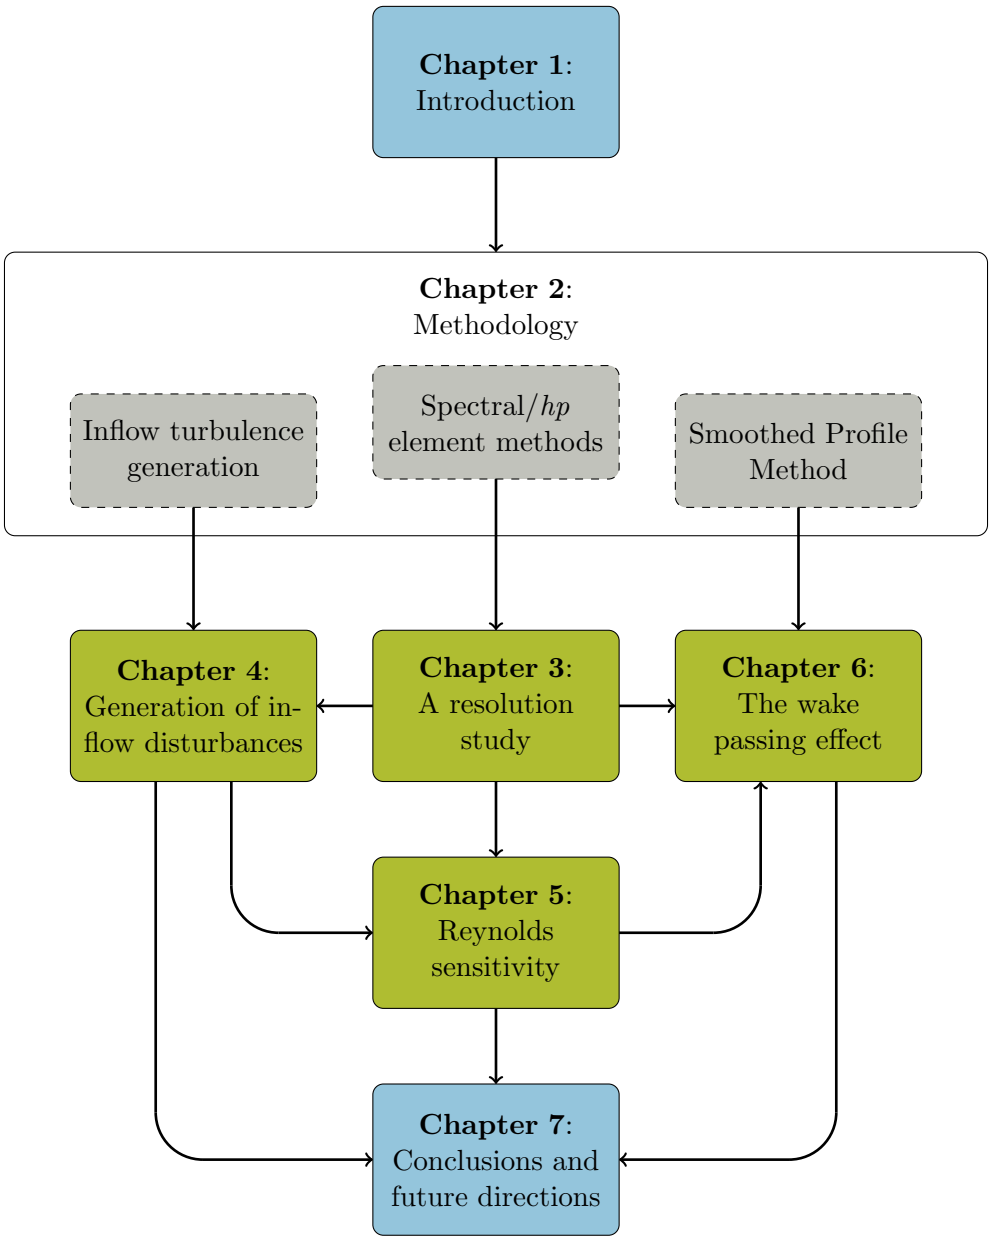

Figure 1.6: Thesis roadmap. In blue: discussion chapters; in grey: methodology and validation chapters; in green: results chapters.

Chapter 2

Numerical methods and computational approach

This chapter discusses the main computational technologies underpinning the rest of the thesis, and it is divided in four main parts. First, a description of the numerical methodology adopted for the discretisation of PDEs using the spectral/*hp* element method is provided, focusing specifically on algorithms for incompressible flows. The convergence properties discussed are further investigated in the context of turbomachinery flows in Chapter 3, which provides the grounding capability to simulate a virtual rig. Special attention is further dedicated to the numerical techniques adopted to tackle high-Reynolds flows and guarantee numerical stabilisation, discussed in Section 2.2. Subsequently, the third part of the Chapter focuses on methodologies to introduce stochastic disturbances (i.e. inflow turbulence), justifying and validating the approaches selected, which are applied and compared in Chapter 4. Finally, the focus is set on the computational approach selected to model deterministic inflow disturbances (i.e. wake passing effect), which is validated in detail and further explored in Chapter 6.

2.1 Spectral/*hp* element methods

In the numerical approximation of a system of PDEs, the mathematical formulation defined in a continuous space is approximated by a discrete representation. The conditions under which the finite representation is defined are determined by the specific choice of numerical method adopted. The modern spectral methods were first presented by Gottlieb and Orszag [55], while the spectral element method was later introduced by Patera [129] and further developed into the spectral/*hp* element method by Karniadakis & Sherwin [77]. In the context of the spectral/*hp* element method, the so-called method of weighted residuals is adopted, leading to the well-known Galerkin formulation which is briefly outlined in the next section.

2.1.1 The method of weighted residuals

The method of weighted residuals shows how by choosing different *test functions* (also called *weight functions*), some of the various methods for solving a system of PDEs can be derived. Complete details can be found in the book by Karniadakis & Sherwin [77], which is taken as a reference throughout this section and the following.
The numerical solution of a PDE of the form $\mathcal{L}(u) = 0$ over a domain Ω is considered, subject to appropriate initial and boundary conditions. The solution is assumed to be accurately represented by the approximation:

$$u^{\delta}(\mathbf{x}, t) = u_0(\mathbf{x}, t) + \sum_{i=1}^{N_{\text{dof}}} \hat{u}_i(t)\phi_i(\mathbf{x})\,, \tag{2.1}$$

where $\phi_i(\mathbf{x})$ are called *trial* (or *expansion*) *functions*, $\hat{u}_i$ are N_{dof} unknown coefficients and $u_0(\mathbf{x}, t)$ satisfies boundary and initial conditions. Substituting this expression into the differential problem produces non-zero residual, $\mathcal{L}(u^{\delta}) = R(u^{\delta})$. In order to uniquely determine the coefficients $\hat{u}_i$ it is necessary to introduce a restriction on the residual, which reduces the problem to a system of ordinary differential equations, to be solved via some time integration scheme. Introducing the Legendre inner product as

$$(f, g) = \int_{\Omega} f(\mathbf{x}), g(\mathbf{x}))\mathrm{d}\mathbf{x}\,, \tag{2.2}$$

the inner product of the residual with respect to a weight (or test) function is required to be zero:

$$(v_j(\mathbf{x}), R) = 0 \qquad j = 1, \cdots, N_{\text{dof}}\,. \tag{2.3}$$

The residual tends to zero as $N_{\text{dof}} \to \infty$, since the discrete approximation tends to the exact solution. The choice of the trial and test functions determines the nature of the numerical scheme and its properties. Most finite element and spectral/*hp* element formulations adopt a Galerkin projection, where the test functions are chosen to be identical to the trial functions: $v_j = \phi_j$.

The convergence of Galerkin methods (as well as pseudo-spectral methods) is exponential. However, the rate of convergence is not fixed and it depends on the regularity of the solution. For applications that lead to sufficiently smooth solutions, this approach has a significant algorithmic advantage.

Spectral/*hp* element methods combine the local nature of the expansion functions typical of finite element methods and the arbitrary expansion type typical of spectral methods. *H*-type refinement is used in conjunction with *p*-type refinement, therefore combining the geometrical flexibility of finite element methods and the superior spatial accuracy properties of spectral methods. Beyond the exponential convergence rate allowed by *p*-refinement [7], some other advantages (common to other high-order methods) deserve a mention: small diffusion/dispersion error, easier implementation of the inf-sup condition

for the incompressible solver as well as better data volume-over-surface ratio for efficient parallel processing.

In general, the spectral/*hp* element method involves a decomposition of the domain into subdomains or elements, each of which is mapped into a reference domain (the standard element) through a parametric mapping. In each standard element, the solution is represented by means of a polynomial function with arbitrary degree (also referred to as *expansion basis function*) and the required elemental operations (multiplications, integrals and differentiations) are performed. The global problem is then finalised by introducing specific connectivity rules at the interfaces between adjacent elements. Connectivity can be enforced in different ways. The classical CG method imposes C^0 continuity across elments by means of a global assembly strategy. On the other hand, if the solution is allowed to be discontinuous across element boundaries (but requiring fluxes to be continuous) the DG or FR approaches are obtained. The different formulations are suitable for different classes of physical problems. The DG method is naturally well-suited for hyperbolic conservation laws typical of compressible flows, and while applicable for solving incompressible flows, the CG method can leverage more advanced maturity in the algorithms for tackling this class of problems and therefore it is adopted in this study.

The computational efficiency of the numerical method and the suitability to CPU-GPU architectures depends on the equations being solved, generally interlinked with the numerical approach. Over the past two decades, significant improvements in arithmetic capability have outpaced the advances in random access memory [194]. This explains how the bottleneck in efficiency of the CG incompressible formulation is currently posed by memory movement, thus limiting the size of the problems that can be tackled in feasible compute times. Modern hardware architectures rely on high-speed caches and shared memory to maintain throughput. For an algorithm to leverage these efficiently, however, the memory access pattern must exhibit a degree of data locality. Compressible solvers relying on a DG or FR formulation with compact stencils demonstrate good scaling properties and can leverage GPU architectures.

Research efforts in the *Nektar++* group have been directed towards the development of implicit solvers [200] to alleviate the significant timestep restrictions of the explicit formulation while emphasizing development patterns to limit the memory footprint, which tends to be high owing to the coupling between elements.

Despite the challenge in providing an accurate and fair comparison across different numerical solvers, the compressible and incompressible *Nektar++* capability was tested on two standard benchmark cases, namely the Taylor-Green vortex and the flow past a circular cylinder at $Re = 3900$ [182]. For the incompressible formulation, both a full 3D discretisation and a Quasi-3D formulation were considered (and discussed in detail in the following section). The comparison was established against the explicit and implicit compressible solvers. In the flow past a circular cylinder, considering the Quasi-3D incompressible approach as benchmark, the full 3D incompressible formulation was found approximately an order of magnitude slower. The same gap separated the 3D incompress-

ible solver and the explicit compressible formulation, while the implicit approach was found twice as fast as the explicit counterpart. It should be highlighted that these estimates are highly dependent on a number of factors, and further developments (e.g. the ability to leverage vector processing units in the explicit compressible solver) may significantly affect the relative efficiency and speed of the various solvers. These considerations highlight how the addition of compressibility effects is not trivial, owing to the fact that it is not simply a small modification of the solver, but it requires an entirely diffent numerical formulation.

The following sections briefly outline the main building blocks of the spatial discretisation. Subsequently, the time-integration scheme adopted in this work for solving the incompressible Navier-Stokes equations is outlined.

2.1.2 Spatial discretisation

The computational domain Ω is subdivided into N non-overlapping elements Ω_n. Among the various shapes that can be used, this work is mainly concerned with quadrilateral and triangular elements. The standard element is denoted by Ω_s. The expansion bases are defined on the standard element, as well as the main operations like integration and differentiation. The standard quadrilateral region is defined as:

$$\Omega_s = \{(\xi_1, \xi_2)| -1 \leq \xi_1, \xi_2 \leq 1\}\,, \tag{2.4}$$

and the standard triangular region is instead:

$$\Omega_s = \{(\xi_1, \xi_2)| -1 \leq \xi_1, \xi_2,\ \xi_1 + \xi_2 \leq 0\}\,. \tag{2.5}$$

A collapsed two-dimensional coordinate system is introduced to describe the standard triangular element with independent coordinate limits. Introducing:

$$\eta_1 = 2\frac{1+\xi_1}{1-\xi_2} - 1 \tag{2.6}$$

$$\eta_2 = \xi_2 \tag{2.7}$$

the definition of the standard element with η_1, η_2 is identical to that of the quadrilateral region with ξ_1, ξ_2.

Expansion bases

The key feature of p-type methods is that of constructing a polynomial expansion of arbitrary order within each elemental region. A suitable expansion is typically an orthogonal or near-orthogonal set of functions. Importantly, the expansion should be amenable to numerical implementation and computational efficiency. The construction of a multi-dimensional expansion basis is in general obtained as an extension of one-dimensional expansion bases, defined on the standard segment $\Omega_s = \{\xi)| -1 \leq \xi \leq 1\}$. For example, well-known examples of choices for the basis functions are the Lagrange polynomials and

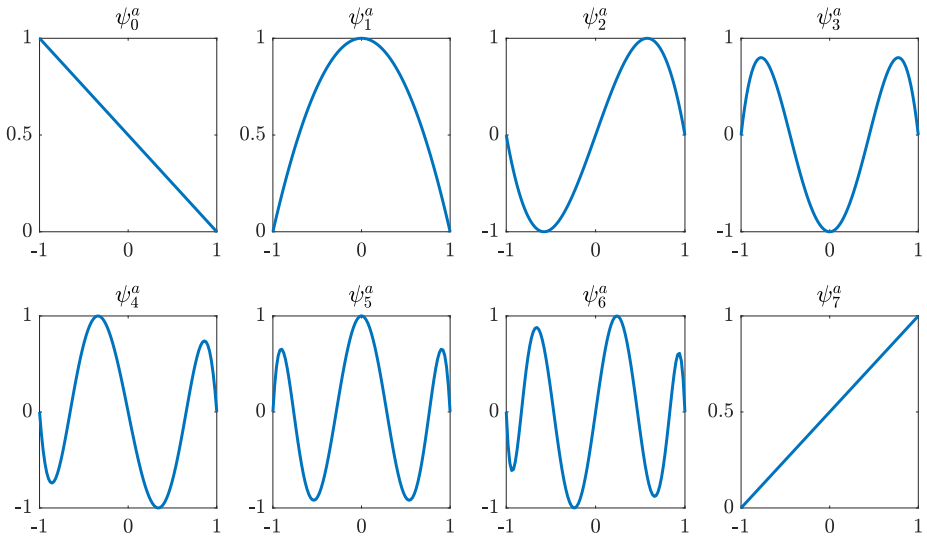

Figure 2.1: One-dimensional modified modal basis functions for $P = 7$.

Legendre polynomials. The former are an example of *nodal* expansion, while the latter is a *modal* expansion. Modal bases are also called *hierarchical* bases, because an expansion of order $P + 1$ contains the expansion of order P. Several factors are taken into account when selecting an expansion set: numerical efficiency, conditioning and linear independence of the basis, as well as its approximation properties [77]. For instance, the orthogonality property of Legendre polynomials is an attractive feature as in general it leads to well-conditioned matrices. However, this is not sufficient to make them the best choice in the context of a CG formulation using spectral/*hp* element methods. In fact, the requirement of C^0 continuity would couple all elemental degrees of freedom to adjacent elements. Therefore, this work adopts a modified modal expansion with boundary-interior decomposition [130, 166, 125], constructed from a family of orthogonal Jacobi polynomials denoted by $\mathcal{P}_P^{\alpha,\beta}$. Indicating a general expansion basis for a one-dimensional standard element with $\phi_p(\xi)$, and the modified modal basis as $\psi_p^a(\xi)$:

$$\psi_p^a(\xi) = \begin{cases} \frac{1-\xi}{2}\,, & p = 0 \\ \left(\frac{1-\xi}{2}\right)\left(\frac{1+\xi}{2}\right)\mathcal{P}_{p-1}^{1,1}(\xi)\,, & 0 < p < P \\ \frac{1+\xi}{2}\,, & p = P \end{cases} \tag{2.8}$$

High-order expansions are leveraged significantly through this work, with polynomial order $P = 7$ being the most frequent choice; the one-dimensional modified basis at this polynomial order is shown in Figure 2.1.

The two-dimensional expansion basis for triangular elements is constructed through

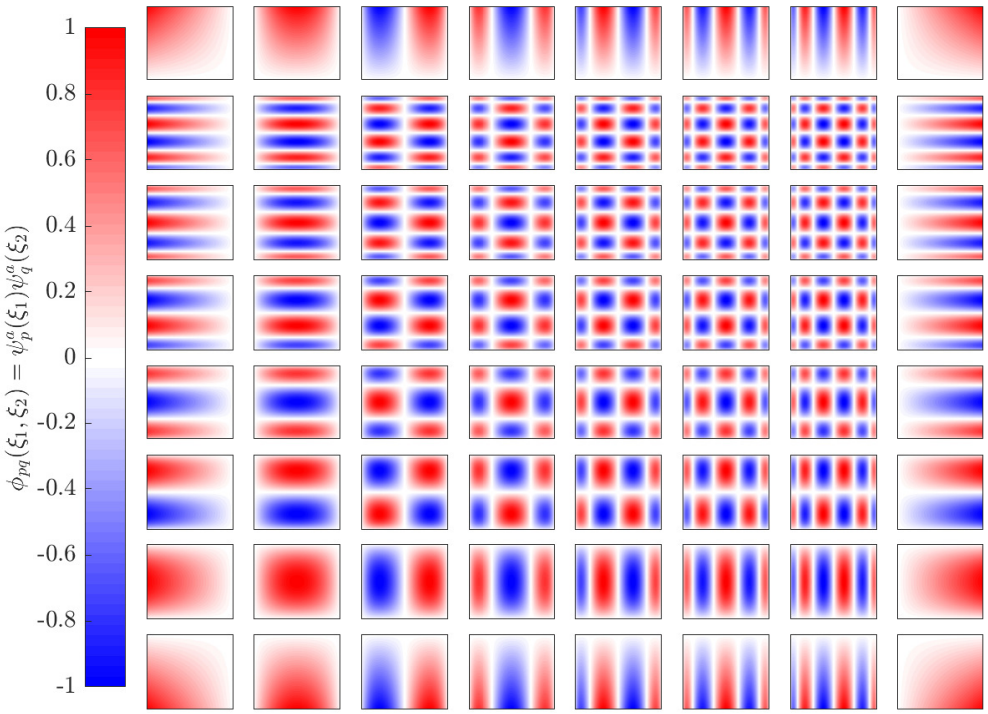

Figure 2.2: Two-dimensional modified modal basis functions for $P = 7$ in a quadrilateral element, with $-1 \leq \xi_1, \xi_2 \leq 1$ The boundary modes are the only ones that have nonzero values at the edge of the standard domain.

the collapsed coordinates in the standard region:

$$\psi_{pq}^b(\eta) = \begin{cases} \psi_q^a\,, & p = 0\,, \quad 0 \leq q \leq Q \\ \left(\frac{1-\eta}{2}\right)^{p+1}\,, & 1 \leq p \leq P\,, \quad q = 0 \\ \left(\frac{1-\eta}{2}\right)^{p+1}\left(\frac{1+\eta}{2}\right)\mathcal{P}_{q-1}^{2p+1,1}(\eta)\,, & 1 < p < P\,, \quad 1 \leq q \leq Q \\ \psi_q^a\,, & p = P\,, \quad 0 \leq q \leq Q \end{cases} \tag{2.9}$$

The two-dimensional expansion bases are obtained via tensor product of the one-dimensional expansion basis. More specifically, for quadrilateral elements:

$$\phi_{pq}(\xi_1, \xi_2) = \psi_p^a(\xi_1)\psi_q^a(\xi_2)\,. \tag{2.10}$$

The two-dimensional basis for a quadrilateral element at $P = 7$ is shown in Figure 2.2, highlighting how only the boundary modes have non-zero values on the boundaries therefore coupling only vertex and edge modes to the neighbouring elements. For triangular elements the basis is obtained via collapsed coordinates:

$$\phi_{pq}(\xi_1, \xi_2) = \psi_p^a(\eta_1)\psi_{pq}^b(\eta_2)\,. \tag{2.11}$$

The type of problems analysed in this work contain a homogeneous direction, associ-

ated with the third coordinate ξ_3. In this case, the three-dimensional expansion can be obtained as a tensor product of the two-dimensional basis with a complete expansion in ξ_3, denoted as $\varphi_r(\xi_3)$:

$$\phi_{pqr}(\xi_1, \xi_2, \xi_3) = \phi_{pq}(\xi_1, \xi_2)\varphi_r(\xi_3)\,. \tag{2.12}$$

The expansion in the third direction is purely spectral: it can be thought of as a single element with (in general) very high polynomial order. If the boundary conditions required for the homogeneous direction are periodic, a Fourier expansion is adopted: $\varphi_r(\xi_3) = e^{ir\beta\xi_3}$, with $\beta = 2\pi/L_{\xi_3}$. L_{ξ_3} is the periodic length, generally referred to as L_z in this work as the third direction coincides with the z-direction. This formulation is particularly attractive for two reasons. First, the use of the Fast Fourier Transform (FFT) allows to quickly transform between Fourier and physical space; second, a linear three-dimensional differential problem can be reduced to $N_z/2$ two-dimensional problems over the Fourier planes, therefore opening ways for more efficient parallelisation strategies, as discussed by Bolis [17] and further investigated in practical applications by Serson [155].

Local to standard element mapping

In order to carry out a domain transformation from the local domain Ω_n to the standard domain Ω_s an appropriate mapping must be defined. For an arbitrary-shaped straight-sided quadrilateral with vertices $\{(x_1^A, x_2^A), (x_1^B, x_2^B), (x_1^C, x_2^C), (x_1^D, x_2^D)\}$, the bilinear mapping is defined as:

$$\begin{aligned} x_i = \chi_1(\xi_1, \xi_2) =& x_i^A \frac{1-\xi_1}{2}\frac{1-\xi_2}{2} + x_i^B \frac{1+\xi_1}{2}\frac{1-\xi_2}{2} && (2.13)\\ &+ x_i^D \frac{1-\xi_1}{2}\frac{1+\xi_2}{2} + x_i^C \frac{1+\xi_1}{2}\frac{1+\xi_2}{2}\,, \qquad i = 1, 2\,. && (2.14)\end{aligned}$$

In the case of triangular element, the mapping is (denoting the collapsed vertex with C):

$$x_i = \chi_1(\eta_1, \eta_2) = x_i^A \frac{1-\eta_1}{2}\frac{1-\eta_2}{2} + x_i^B \frac{1+\eta_1}{2}\frac{1-\eta_2}{2} + x_i^C \frac{1+\eta_2}{2}\,, \qquad i = 1, 2\,. \tag{2.15}$$

The mapping is in practice constructed using the vertex modes of the hierarchical modal expansion. Therefore, the quadrilateral expansion can also be written as:

$$x_i = \chi_1(\xi_1, \xi_2) = \sum_{p=0}^{p=P_1} \sum_{q=0}^{q=P_2} \hat{x}_{pq}^i \psi_p^a(\xi_1)\psi_q^a(\xi_2)\,. \tag{2.16}$$

In particular, $\hat{x}_{pq}^i = 0$ except for the vertex modes where $\hat{x}_{00}^i = x_i^A$, $\hat{x}_{P_1 0}^i = x_i^B$, $\hat{x}_{P_1 P_2}^i = x_i^C$ and $\hat{x}_{0P_2}^i = x_i^D$. This construction allows to account for curvilinear elements through an *iso-parametric* representation, where the expansion is of the same form and order as the unknown variables.

Local operations

Within the standard element, the various differentiation and integration operations are defined by means of appropriate quadrature rules. The interested reader can refer to Karniadakis & Sherwin [77] for an in-depth discussion of the topic. Finally, the operations of forward and backward transformation allow to transform the solution between physical values and expansion coefficients. In particular, the former computes the expansion coefficients from physical values, and vice-versa.

Global assembly

In the CG formulation, elemental contributions of neighbours are summed to enforce C^0-continuity through an assembly operator. Indicating the vector containing the coefficients of the global DoF as $\hat{u}_g$, and the concatenation of local coefficients as $\hat{u}_l$, the correspondence is established through the connectivity map $\mathcal{A}$, such that $\hat{u}_l = \mathcal{A}\hat{u}_g$. The matrix $\mathcal{A}$ is very sparse, and the non-zero entries are either 1 or -1 depending on the orientation of adjacent edges in neighbouring elements. The global system matrix is obtained as:

$$\mathbf{M}_g = \mathcal{A}\underline{\mathbf{M}}_e\mathcal{A}\,, \tag{2.17}$$

where the matrix $\underline{\mathbf{M}}_e$ corresponds to the diagonal concatenation of elemental matrices. Most operations are carried out within elements, but the global system is required to obtain the solution.

2.1.3 Time integration

On top of the spectral/*hp* element spatial discretisation, a strategy for time advancement is needed. In this work, the stiffly stable time discretisation scheme proposed by Karniadakis et al. [76] is adopted and briefly summarised in this section. This approach is also called Velocity-Correction Scheme (VCS).

The reference equations are the incompressible Navier-Stokes equations, describing the motion of Newtonian fluids with constant properties. Introducing the assumption $\rho = 1$, the problem to be solved is:

$$\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u}\cdot\nabla)\mathbf{u} - \nabla p + \nu\nabla^2\mathbf{u} \qquad \text{in}\,\Omega\,, \tag{2.18a}$$

$$\nabla\cdot\mathbf{u} = 0 \qquad \text{in}\,\Omega\,, \tag{2.18b}$$

$$\mathbf{u} = \mathbf{u}_{\mathcal{D}} \qquad \text{on}\,\Gamma_{\mathcal{D}}\,, \tag{2.18c}$$

$$\frac{\partial \mathbf{u}}{\partial n} = \mathbf{u}_{\mathcal{N}} \qquad \text{on}\,\Gamma_{\mathcal{N}}\,. \tag{2.18d}$$

Equations 2.18a and 2.18b are subject to initial conditions $\mathbf{u}_0$; $\mathbf{u}$ is the velocity vector, p is the pressure and ν the kinematic viscosity. Following the approach of [76], the convective terms are treated explicitly, while pressure and the viscous contribution are treated im-

Table 2.1: Coefficients for stiffly stable time integration

Coeff.	Order 1	Order 2	Order 3
γ_0	1	3/2	11/6
α_0	1	2	3
α_1	0	-1/2	-3/2
α_2	0	0	1/3
β_0	1	2	3
β_1	0	-1	-3
β_2	0	0	1

plicitly. This circumvents stability problems that would significantly reduce the timestep. The time derivative is represented through a backwards differentiation formula, and the convective terms via polynomial extrapolation. Indicating with J_i the integration order of implicit terms, and with J_e the order of explicit terms, the momentum equation is discretised as:

$$\frac{\hat{\mathbf{u}} - \sum_{q=0}^{J_i-1} \alpha_q \mathbf{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J_e-1} \beta_q \left[-(\mathbf{u}\cdot\nabla)\mathbf{u}\right]^{n-q} - \nabla p^{n+1} + \nu\nabla^2\mathbf{u}^{n+1}\,. \tag{2.19}$$

Table 2.1 shows the coefficient for integration up to third order. The solution to the system is obtained by splitting Equation 2.19 into three contributions:

$$\frac{\hat{\mathbf{u}} - \sum_{q=0}^{J_i-1} \alpha_q \mathbf{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J_e-1} \beta_q \left[-(\mathbf{u}\cdot\nabla)\mathbf{u}\right]^{n-q}\,, \tag{2.20a}$$

$$\frac{\hat{\hat{\mathbf{u}}} - \hat{\mathbf{u}}}{\Delta t} = -\nabla p^{n+1}\,, \tag{2.20b}$$

$$\frac{\gamma_0\mathbf{u}^{n+1} - \hat{\mathbf{u}}}{\Delta t} = \nu\nabla^2 u^{n+1}\,. \tag{2.20c}$$

Advection

The space discretisation of Equation 2.20 using a Galerkin projection introduces a mass matrix in the advection equation. For a modal expansion, this would introduce an additional computational cost comparable to the solution of the Poisson equation or a single Helmholtz problem [77]. Therefore, in practice Equation 2.20a is analytically combined with Equation 2.20b.

Pressure

Taking the divergence of Equation 2.20b and imposing $\nabla\cdot\hat{\hat{\mathbf{u}}} = 0$ yields the following Poisson equation:

$$\nabla^2 p^{n+1} = \nabla\cdot\left(\frac{\hat{\mathbf{u}}}{\Delta t}\right)\,. \tag{2.21}$$

At velocity Dirichlet boundaries, the following Neumann boundary conditions are used for the pressure:

$$\frac{\partial p^{n+1}}{\partial \mathbf{n}} = \mathbf{n} \cdot \left[\sum_{q=0}^{J_e-1} \beta_q \left[-(\mathbf{u} \cdot \nabla)\mathbf{u} - \nu \nabla \times (\nabla \times \mathbf{u}) \right]^{n-q} \right] . \tag{2.22}$$

Viscous diffusion

The final step of the algorithm consists of solving the Helmholtz equation with the intermediate velocity $\hat{\mathbf{u}}$ acting as a forcing term:

$$\left(\nabla^2 - \frac{\gamma_0}{\nu \Delta t} \right) \mathbf{u}^{n+1} = \frac{\hat{\hat{\mathbf{u}}}}{\nu \Delta t} . \tag{2.23}$$

2.2 Applications to scale-resolving simulations

This work is concerned with large-scale simulations of turbulent flows at moderate and high Reynolds numbers. The simulations are run on large parallel clusters, and optimally leveraging parallel efficiency is of crucial importance. Furthermore, at these flow regimes specific numerical techniques are required to ensure numerical stability. These two topics are discussed in the following sections.

2.2.1 A note on parallelisation

All simulations presented in this work employ the Fourier-spectral/*hp* element formulation introduced by Karniadakis [75]. In this case, very efficient parallelism can be achieved by distributing each process on different planes: the Fourier modes are not coupled in the linear part of the momentum equation, and communication is only required when performing the FFTs to compute explicitly the nonlinear terms in physical space. However, this approach is feasible for applicative purposes in relatively simple configurations, especially in situations where the mean flow is aligned with the z-direction. In most practical cases (i.e. high Reynolds number flows), this approach is unfeasible, because the maximum number of processes is half the number of Fourier planes. The alternative is to combine partitioning in the Fourier expansion and the $x-y$ planes, referred to as *hybrid* parallelism [16]. The latter approach is adopted in this work. The efficiency of a pure Fourier parallelisation is not trivially retained because the coupling between partitions now occurs in the implicit part of the VCS, which requires efficient parallel solution of linear systems. In this study, the linear systems were solved by means of direct solve via XX^T matrix decomposition [171], as opposed to the more common conjugate gradient method. This solution is only applicable to Fourier-spectral/*hp* element methods. Serson [154] extensively explored hybrid parallelisation efficiency with this formulation. For the cases of interest in this work, the most efficient way to split the process topology was found to be that of maximising the number of partitions in the Fourier direction. This was verified via

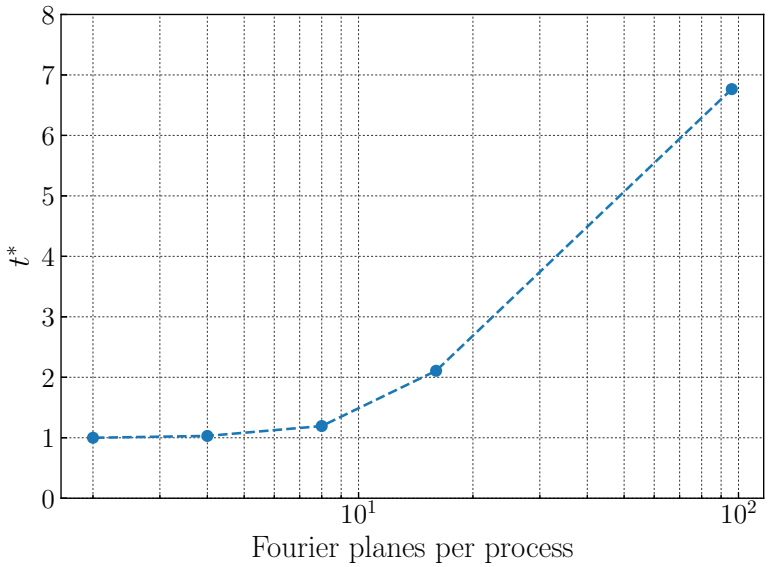

Figure 2.3: Normalised timestep time measured with constant number of total processes and different spanwise partitioning.

scaling test on the UK national supercomputer Archer[1], conducted varying the balancing between $x - y$ and Fourier partitioning with constant total core count (2304 cores). The solution was advanced in time for 200 timesteps, adopting the averaged timestep time as performance metric. The result is shown in Figure 2.3, normalised by the most efficient value of $t^* = 0.1389$. High parallelism in the Fourier direction is desirable in order to achieve maximum computational efficiency. In particular, no major performance degradation is reported when allocating 4 planes per process ($\sim 3\%$ increase in timestep time compared to 2 planes per process), but any partitioning that allocates a higher number of planes per process incurs significant increases in timestep time; this behaviour was found to occur consistently across all test cases analysed in this thesis.

2.2.2 Ensuring numerical stability: spectral vanishing viscosity

Under-resolved simulations of turbulent flows at high Reynolds numbers are prone to numerical instabilities, owing to the low numerical diffusion properties of spectral/*hp* element methods [77]. Therefore, the use of stabilization techniques is required: if the dissipative scales at high wavenumbers are not resolved, a buildup of energy in the small scales will generally cause the solution to diverge. Spectral Vanishing Viscosity (SVV) was first introduced by Tadmor et al. [168], and extended for use with spectral/*hp* element methods by Karamanos & Karniadakis [74]. Kirby & Sherwin [83] further extended the formulation using orthogonal expansions for one-, two- and three-dimensional spectral element discretisations.

[1]ARCHER uses the Cray XC30 architecture consisting of compute nodes connected together by the Aries interconnect, where each compute node contains two 2.7 GHz, 12-core E5-2697 v2 (Ivy Bridge) series processors. Within the node, the two processors are connected by two QuickPath Interconnect (QPI) links.

The oscillations causing divergence of the solution emerge at sub-elemental level; hence SVV is a model-free stabilization technique that acts at the subgrid-scale level. The main idea of SVV consists of introducing an artificial viscosity operator:

$$\frac{\partial \mathbf{u}}{\partial t} = -(\mathbf{u} \cdot \nabla)\mathbf{u} - \nabla p + \nu \nabla^2 \mathbf{u} + S_{vv}(\mathbf{u}) \tag{2.24a}$$

$$\nabla \cdot \mathbf{u} = 0 \tag{2.24b}$$

In the context of spectral/hp element methods, a high-frequency viscosity operator allows for exponential convergence properties to be conserved for sufficiently resolved simulations. The SVV operator is defined as:

$$S_{vv} = \mu_{SVV} \frac{\partial}{\partial x} \left(\hat{Q} \star \frac{\partial u}{\partial x} \right) , \tag{2.25}$$

where μ_{SVV} is a constant coefficient, $\star$ denotes the convolution operator and $\hat{Q}$ is a kernel that regulates which modes receive how much damping. A widely used kernel is of exponential type [99]:

$$\hat{Q}_k = \begin{cases} \exp\left(-\frac{(k-N)^2}{(k-M)^2}\right) , & k > M \\ 0 , & k \leq M \end{cases} \tag{2.26}$$

In this case, artificial viscosity is only applied to modes k above a certain cutoff mode M, with $N = P + 1$ representing the number of modes for a discretisation of polynomials of order P. When applied to the Fourier expansion, k represents the Fourier mode number and N the total number of modes. A more detailed description of the SVV formulation in a single element is found in Kirby & Sherwin [83], and here omitted for brevity.

DG Kernel for SVV

The exponential kernel introduced in the previous paragraph has two main issues, discussed in detail by Moura et al. [117]. To highlight them, the Péclet number is introduced to quantify the ratio of advection and diffusion. This parameter can be thought of as a mesh-spacing based local Reynolds number. Introducing the advection velocity a:

$$\mathrm{Pe} = \frac{ah}{\mu_{\mathrm{SVV}}} . \tag{2.27}$$

Since the eigencurves are uniquely defined by Pe and P, the DoF-based Péclet number $\mathrm{Pe}^* = \mathrm{Pe}/P$ is considered, with $\hbar = h/P$:

$$\mathrm{Pe}^* = \frac{a\hbar}{\mu_{\mathrm{SVV}}} . \tag{2.28}$$

The conventional approach is to use $\mu_{\mathrm{SVV}} = \mu_0/P$, with μ_{SVV} representing the base SVV magnitude and μ_0 a fixed parameter. This in turn makes $\mathrm{Pe}^* = ah/(\mu_{\mathrm{SVV}}P) = ah/\mu_0$ still dependent on the product ah, which in general varies significantly in applications of

practical interest: the two parameters are independent of each other. The dissipation and diffusion eigencurves depend on the Péclet number, and therefore the effect of SVV on a certain physical problem at fixed $\hbar$ is not unique and it depends on the local Péclet number. The second main inconvenience is again due to the fact that the eigencurves depend on Pe^*. For this reason, they provide information that is only relevant for a certain mesh spacing h, and therefore of little practical use.

Furthermore, in general it is difficult to provide physical robust guidelines on how artificial viscosity should be added to provide the sufficient level of stabilisation at the right wavelengths, without impacting the flow physics and regardless of the problem in analysis. When adopting traditional kernels (e.g. exponential kernel), in most cases the user is left to decide on what diffusion and cutoff rate parameters to choose, which typically results in adopting a heuristic approach driven by robustness (the simulation does not "blow up") or agreement of certain flow statistics with reference data in well-known benchmark test cases. However, this approach is arguably not optimal and can be considered a further motivator to adopting a more rigorous approach for numerical stabilisation.

Moura et al. focused on temporal [117] and spatial eigenanalysis of the linear advection-diffusion equation, both with a DG [106] and CG method [116]. Temporal eigenanalysis is traditionally more widely analysed in literature and it assumes periodic boundary conditions. Spatial analysis, on the other hand, concernes inflow/outflow types of problems and is of greater interest for the applications of this work. The outcome of this series of studies is the formulation of a DG-Kernel for SVV, now briefly summarised.

The model problem analysed in [116] is the one-dimensional linear advection-diffusion equation:

$$\frac{\partial u}{\partial t} + a\frac{\partial u}{\partial x} = \mu\frac{\partial^2 u}{\partial x^2}\,, \quad \text{for } \; x > 0\,,\; t > 0\,, \tag{2.29}$$

where μ is the molecular viscosity. This is the main parameter varied in the spatial eigenanalysis, which highlights two important problems that SVV must address. First, a spurious wave mode causing unphysical reflection was found to exist regardless of Pe^*. In realistic test cases, this may arise in regions of variable mesh spacing and interact with incoming turbulent structures, affecting the flow physics. The effect is particularly enhanced in the inviscid limit (i.e. high-Reynolds number flows), where viscous damping is very small. At the same flow regime, another feature observed is the presence of a dissipation bubble (for $P > 2$) at certain temporal frequencies. This affects the streamwise energy spectrum by introducing a "dissipative valley" [116], and it clearly visible from the diffusion curves in the inviscid limit shown in Figure 2.4.

In order to tackle the existing issues in the exponential kernel, Moura et al. [117] proposed to make the base SVV magnitude proportional to both advection speed and mesh spacing, therefore scaling the magnitude of numerical dissipation with the Péclet number: $\mu_{\mathrm{SVV}} = \mu_0 a\hbar$, thus maintaing $\mathrm{Pe}^* = \mu_0^{-1}$ constant regardless of local advection speed and mesh spacing. The issues highlighted through eigenanalysis of the linear advection-diffusion problem (dissipative bubbles and vanishing dissipation of spirious modes) were

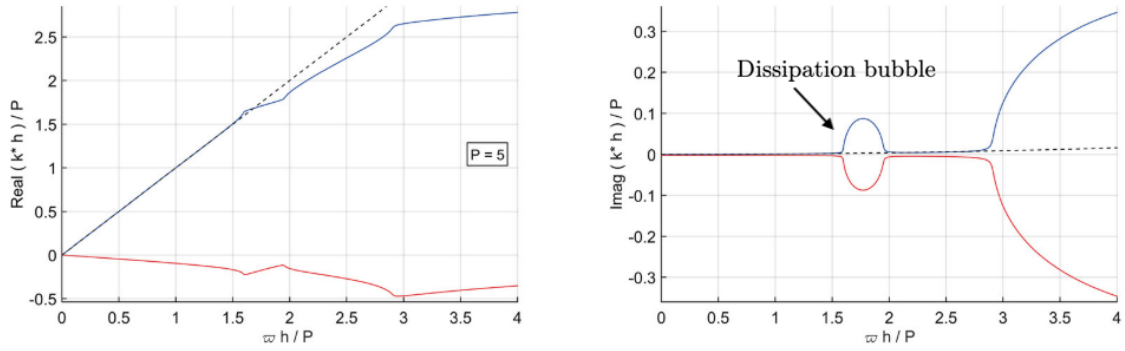

Figure 2.4: Numerical dispersion (left) and diffusion (right) curves for CG-based linear advection-diffusion with $Pe^* = 1000$ and $P = 5$. The blue and red curves indicate respectively the physical and spurious eigenmodes. Adapted from [116].

instead tackled through the consideration that DG schemes are arguably robust and weakly dissipative at higher polynomial orders. Therefore, the novel SVV operator was designed based on DG dissipation curves. Since it is not possible to match same-order curves [116], the dissipation curves of CG were optimised against the DG curves of order $P-2$, using $\mathrm{Pe}^* = \mu_0^{-1}$. This addresses the issue of the dissipative bubbles found in the CG formulation without SVV. Spurious reflections were instead tackled by introducing penalisation in the optimisation process where the dissipation of the spurious mode was too small. The result of this optimisation procedure is shown in Figure 2.5.

It is important to note that lower dissipation levels are achieved at higher polynomial order confirming that, at fixed DoF, better resolution ability is achieved by increasing the polynomial order (both in terms of diffusion and dispersion properties). Therefore, especially when tackling under-resolved high Reynolds number flows, the most sensible approach is to adopt a higher polynomial order in combination with a moderately coarse grid, as opposed to favouring h-refinement. In fact, combining SVV with the latter approach would damp spurious oscillations and unphysical reflections, but it would also suppress physical scales of the problem. In light of these considerations, all significant discussions on the flow physics presented through this work were based on results calculated at polynomial order $P \geq 6$.

Through the thesis, the DG-Kernel is adopted for ensuring numerical stability in the spectral/hp planes. To the authors' knowledge, this is the first study that employs this SVV formulation for scale-resolving simulations of realistic geometries. The traditional exponential kernel is adopted for the Fourier expansion in the spanwise direction. Diffusion coefficient and cutoff ratio are tuned based on robustness and previously consolidated expertise, aiming at a tradeoff between simulation stability and limited damping of high-frequency modes. The topic is further discussed in Chapter 3, where variations of the cutoff ratio within the range of values adopted (which must be reduced at times to ensure numerical stability) are shown to have no first-order effects on the flow physics. All simulations in the thesis also rely on the use of spectral/hp dealiasing for consistent integration

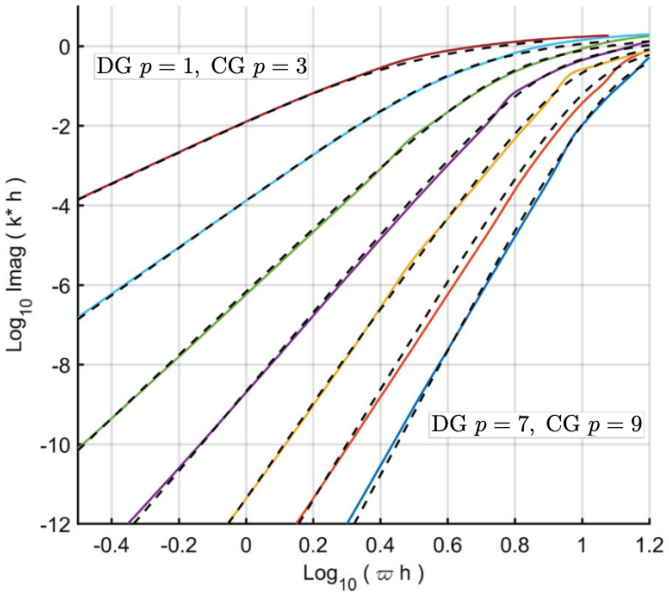

Figure 2.5: Comparison between optimised CG-SVV dissipation (colour) for $P_{CG} = 3, \ldots, 9$ and their reference DG curves (dashed) for $P_{DG} = 1, \ldots, 7$. Adapted from [116].

of the nonlinear terms, based on the work by Mengaldo et al. [107].

2.3 Inflow turbulence generation

In LPTs, the suction surface boundary layer dominates the mechanisms of loss production [33]. Up to 80% of the span may experience boundary layer separation in the adverse pressure gradient region, and in these conditions flow unsteadiness and transition mechanism are of fundamental importance for design decisions. The generation of physical boundary conditions is crucial, and specific validation with assessment is needed. Where possible, this should be informed from the actual conditions encountered in tests, as well as in potential real applications. Inflow turbulence is a key aspects for turbomachinery problems, especially at the Reynolds number conditions involving transitional boundary layers. This section focuses on the approaches here adopted for modelling background turbulence and flow disturbances. First, a brief review of the topic is presented, with enhanced focus on the methodology effectively adopted, which is subsequently described in greater detail alongside with its validation.

2.3.1 The main types of turbulence generation methods

From a computational perspective, inflow turbulence generation methods have been the objectives of significant academic research efforts especially since the 1990s. The reviews by Tabor & Baba-Ahmadi [167] and by Wu [197] (whose classification is adopted for the purposes of this summary) highlight the most significant contributions to the field.

Various methodologies for turbulence generation in spatially developing flows are now available, and they mainly break down into precursor simulation methods and synthetic methods. The former can be further subdivided in strong and weak recycling methods; in general, these approaches are I/O heavy, but intrinsically introduce "true" turbulence in the computational domain. Synthetic methods are less accurate but they are implemented on-the-fly injecting "artificial" turbulence, therefore requiring "some" length to develop the correct spectral characteristics. Finally, although less commonly adopted, a direct approach can be formulated based on the simulation of an upstream turbulence grid [13].

This work is concerned with synthetic methods, owing to their low memory and storage requirements, and simplicity by which their parameters can be tuned. Wu [197] further subdivides synthetic generation methods into various approaches: random Fourier method, digital filtering method, coherent eddy method, and synthetic volume forcing method.

2.3.2 Synthetic inflow turbulence generation

Random Fourier Method (RFM)

In the context of simulating incompressible flows, it is important to adopt a methodology that respects the divergence-free constraint. To this end, before the increasing popularity of DNS studies, Kraichnan [85] developed an approach to represent a velocity field by summation of Fourier modes, providing the basis for research advances in the decades to come. In fact, this technique was later employed to develop the first synthetic random Fourier methods in the early 1990s [78, 89, 12, 9]. In particular, Bechara et al. [12] introduced the core approach adopted in this work, and described below. The following approximation is introduced for the HIT field:

$$u_i'(x_j) = 2 \sum_{n=1}^{N_{\text{turb}}} \hat{u}^n \cos(\kappa_j^n x_j + \psi^n)\sigma_i^n \,, \tag{2.30}$$

where for each mode n, $\hat{u}^n$ is the amplitude, κ_j^n is the wavenumber vector, ψ^n is the random phase and σ_i^n is the direction of the mode associated with κ_j^n. For the generation of an incompressible turbulent field respecting $\partial u_i/\partial x_i = 0$, the unit vector $\boldsymbol{\sigma}^n$ (which determines the direction of the nth wave vector $\boldsymbol{\kappa}^n$) must be constructed orthogonally to $\boldsymbol{\kappa}^n$:

$$\boldsymbol{\kappa}^n \cdot \boldsymbol{\sigma}^n = 0 \,, \qquad n = 1, \ldots, N \,. \tag{2.31}$$

The spatial orientation of these vectors is shown in Figure 2.6. The angles φ^n and θ^n control the orientation of $\boldsymbol{\kappa}^n$, and α^n determines the orientation of the unit vector in the plane orthogonal to $\boldsymbol{\kappa}^n$. To obtain a homogeneous and isotropic random field, for each Fourier mode the orientation angles and mode phase ψ^n must be generated following the probability density functions reported in Table 2.2. In this context, introducing the turbulence kinetic energy $k = \sum_{n=1}^{N_{\text{turb}}} \hat{u}_n^2$ and turbulence dissipation rate $\epsilon \equiv \nu \overline{\frac{\partial u_i'}{\partial x_k} \frac{\partial u_i'}{\partial x_k}}$, the three-dimensional energy spectrum that characterises homogeneous isotropic turbulence

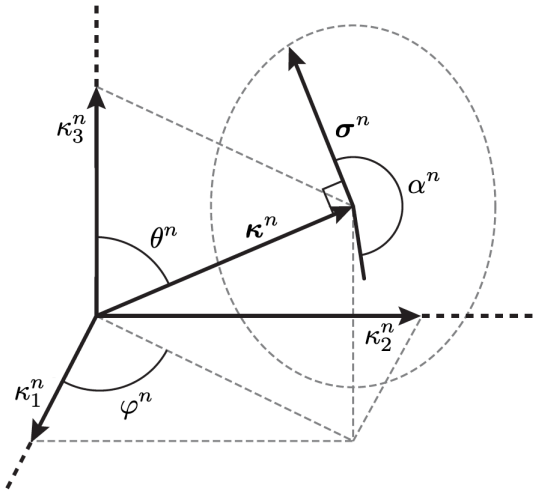

Figure 2.6: Sketch defining the spatial orientation of the wave vector κ^n and unit vector σ^n. Adapted from [197].

Table 2.2: Probability density function for the random variables.

Variable	PDF	Sample space
φ^n	$p(\varphi^n) = 1/(2\pi)$	$0 \leq \varphi^n \leq 2\pi$
α^n	$p(\alpha^n) = 1/(2\pi)$	$0 \leq \alpha^n \leq 2\pi$
ψ^n	$p(\psi^n) = 1/(2\pi)$	$0 \leq \psi^n \leq 2\pi$
θ^n	$p(\theta^n) = \sin(\theta^n)/2)$	$0 \leq \theta^n \leq \pi$

satisfies the following:

$$\int_0^\infty E(\kappa)\mathrm{d}\kappa = k\,, \tag{2.32}$$

$$2\nu \int_0^\infty \kappa^2 E(\kappa)\mathrm{d}\kappa = \epsilon\,. \tag{2.33}$$

From [12], the mode amplitude $\hat{u}^n$ is:

$$\hat{u}^n = \sqrt{E(|\kappa^n|)\Delta\kappa^n}\,, \tag{2.34}$$

with $\Delta\kappa^n$ representing the wavenumber step, which is constant in case of a uniform distribution of modes over the wavenumbers. Bechara et al. [12] suggested employing a logarithmic distribution of wave numbers, which better discretises the spectrum in the low wavenumber range. However, the boundary conditions setup of the current work (i.e. pitchwise and spanwise periodicity on the inflow plane) requires the use of a modified approach, described in detail in Section 2.3.3.

The modified von Kármán spectrum (shown in Figure 2.7) is discretised:

$$E(\kappa) = A\frac{\frac{2}{3}k}{\kappa_e}\frac{(\kappa/\kappa_e)^4}{\left[1+(\kappa/\kappa_e)^2\right]^{17/6}} e^{\left[-2(\kappa/\kappa_\eta)^2\right]}\,, \tag{2.35}$$

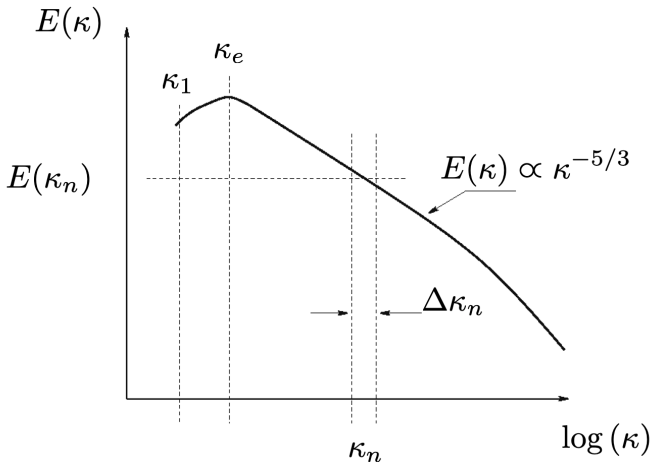

Figure 2.7: Modified von Kármán spectrum, adapted from [35].

with A and κ_e determined from Equations 2.32 and 2.33, and $\kappa_\eta = (\epsilon/\nu^3)^{\frac{1}{4}}$.

Finally, temporal correlation is enforced by means of a digital filter in the frequency domain. This step is further discussed in detail in Section 2.3.3.

Building on the work by Bechara et al. [12], Smirnov et al. [158] presented a procedure for scaling the generated fluctuations, obtaining an inhomogeneous, anistropic field that still satisfies the divergence-free constraint. Other approaches for generating non-isotropic fluctuations were later introduced [10, 14, 36]. As observed by Davidson [35], a disadvantage of the latter approach is that modifying the Reynolds stress tensor to obtain non-homogeneous turbulence also modifies the prescribed spectrum and the two-point correlation.

More recently, Huang et al. [62] and Castro & Paz [25] built on Smirnov's approach introducing a technique that can generate a fluctuating field from an arbitrary given spectrum, appropriately discretised. The main shortcoming of this novel approach is the necessary introduction of additional coefficients that may increase the parameter space to be explored when tackling a new test case. Yu & Bai [201], again building on the work by Smirnov et al. [158] presented a different methodology based on the introduction of a vector potential field, and generating a strictly divergence-free flow field by taking the vector curl of the potential field.

Digital Filtering Method (DFM)

Klein et al. [84] reported on a number of disadvantages of methods based on random Fourier modes, including (but not limited to) programming complexity, and difficulty in selecting the peak wavenumber in the energy spectrum. Extending the temporal coherence approach introduced by Bechara et al. [12] to spatial coherence, they presented a different method based on digital filtering of random data, which can reproduce auto-correlation functions and prescribed second order statistics. The drive for adopting this approach is practicability, i.e. allowing the user to specify statistical quantities that can

be produced with simple experimental procedures or heuristic estimates. This method was further modified to provide practical indications for an efficient implementation [79]. Di Mare et al. [39] showed how in principle the filter coefficients can be determined for any known correlation function, but in some specific cases it is possible to write an explicit approximate solution. When extended to multiple directions, this approach requires significantly higher computational cost [84]; it was modified for this purpose by Xie & Castro [199], and extended to compressible flows by Touber & Sandham [169]. This method was implemented with success by Pichler et al. [134] and used to study the effect of turbulence intensity and length scale on the transition mechanisms close to realistic engine operating conditions in a HPT vane. Among the drawbacks of most digital filtering approaches, the divergence-free constraint is rarely enforced; Kim et al. [82] proposed an approach to account for incompressibility in the synthetically generated signal.

Other synthetic methods

A number of alternative approaches to RFM and DFM were developed over the years. These two methods impose constraints to a random field, thereby constructing a signal with specified properties. A different approach called Synthetic Coherent Eddy Method [70] is based on producing the synthetic signal as a superposition of coherent structures. This method can reproduce prescribed first and second order one-point statistics as well as characteristic length and time scales.

An alternative and possibly less expensive approach consists of adding a body force term to the momentum equations in a control zone. The forcing should be designed to promote rapid transition and concurrently avoid introducing strong distortions in the flow, which may artificially impact its downstream development. Examples of this Volume Forcing Method were applied in the wall-normal direction in a Blasius profile, to promote its development into a zero-pressure-gradient smooth flat-plate boundary layer [148, 149].

Finally, Kempf et al. [80] introduced a methodology that imposes coherence to an initial random field by solving an unsteady diffusion equation for a certain time length, depending on the target length scale. The method is easy to implement but requires the solution of an additional differential equation.

2.3.3 Generation of inflow disturbances in *Nektar++*

As highlighted in the previous section, nowadays a large number of different approaches are available for inflow turbulence generation. Every method is presented with ad hoc test cases, targeting specific aspects of the flow physics and, as such, they are difficult to compare. It is therefore important to establish comparisons between different metodologies adopting simple, canonical test cases (i.e. temporal/spatially decaying generated data). Despite the difficulty of undertaking such study, the work by Dietzel et al. [40] tackled this challenge by looking at a box of decaying turbulence. In particular, the methods by

Klein et al. [84], Kempf et al. [80] and Davidson et al. [36] were analysed with focus on four specific issues: (1) fidelity in representing the characteristic energy spectrum and (2) the velocity correlations, (3) how do these properties decay and (4) what needs to be clear for their practical application. Davidson's approach was found to provide accurate representation of the characteristic energy spectrum as well as the velocity correlations, comparing well against experimental data. The approaches by Klein et al. [84] and Kempf et al. [80] only construct large scale turbulence, and initially the model correlations did not compare accurately against experiments. However, the discrepancies vanished at a later decay stage. Despite the difficulty of transposing the findings related to a simple test case to more complex, spatially evolving flows, this study suggests that the random Fourier approach proposed by Davidson [35] has the best potential to satisfy the requirements of the present study, and it was therefore adopted and incorporated in the *Nektar++* incompressible Navier-Stokes solver.

Concurrently, a volume forcing method was here developed as an alternative approach to promote inflow disturbances. This methodology was first developed by Schlatter et al. [148, 149] and applied to turbulent boundary layer flow. It is here repurposed, simplified and applied to turbomachinery problems for the first time, to explore a novel approach to provide highly unsteady incoming flow at reduced computational cost. Since the approach developed is very problem-specific, it is discussed in detail in Chapter 4 within the context of its application.

Implementation details

The adapted implementation of Davidson's method [36, 35] is briefly described in this section. For the purposes of clarity, user-defined parameters are highlighted in the equations that follow.

First, all parameters that remain constant through the simulation are calculated prior to time advancement:

- In general, care should be taken to ensure that the base mesh is generated with uniform element spacing on the inflow plane. In the algorithm, the grid spacing can be specified to control the smallest scales generated, and within the context of spectral/hp element methods it would generally be defined as $\Delta = h/p$. The maximum wavenumber is then:

$$\kappa_{\max} = \frac{2\pi}{2\Delta}\,. \tag{2.36}$$

  However, if one desires Δ can be differently specified, introducing structures that resolve only the lower portion of the wavenumber range that could potentially be supported by the mesh.

- Define the smallest wavenumber, associated to large-scale spatial features. Introducing the user-defined turbulence length scale L_t, the smallest wavenumber is

$\kappa_1 = \kappa_e / p_e$, where the peak wavenumber κ_e (Figure 2.7) is:

$$\kappa_e = \frac{9\pi A}{55 L_t}, \tag{2.37}$$

with $A = 1.45276$. The factor p_e scales κ_1 to introduce larger scales than those corresponding to κ_e. Davidson [36, 35] suggested adoping $p_e = 2$; care must be taken to ensure that the size of the computational domain is sufficiently large, so as to contain the largest structures generated.

- The wavenumber space is divided in N_{turb} user-defined equally spaced modes, yielding the following uniform wavenumber spacing:

$$\Delta\kappa^n = \Delta\kappa = \frac{\kappa_{\max} - \kappa_1}{N_{\text{turb}} - 1}. \tag{2.38}$$

  Since the problems in this work involve periodicity in both the pitchwise and spanwise direction, an additional option was added to enforce periodicity in the method. This consists of determining the wavenumbers in such a way that the wavelengths selected fit in the computational domain an exact number of times. This means that grid resolution and domain size uniquely determine the wavenumber discretisation. In this case, the specified modes N_{turb} are therefore overwritten, generally yielding a significant reduction of the Fourier modes, which as a consequence increases the loss of energy with respect to the specified target value of inflow turbulence intensity. In case two directions are periodic, in order to adopt this approach the domain size in one direction must be a factor of the other.

- The modified von Kármán spectrum (Equation 2.35) is discretised into N_{turb} modes, with $\kappa_\eta = (\epsilon/\nu^3)^{\frac{1}{4}}$. Indicating with U_∞ the mean Reynolds-averaged velocity (which is typically normalised to 1), the turbulence intensity (or turbulence level) is defined as a function of turbulence kinetic energy as:

$$\text{TI} = \frac{\sqrt{\frac{2}{3}k}}{U_\infty}, \tag{2.39}$$

  so the spectrum sampled can be rewritten as:

$$E(\kappa) = A \frac{U_\infty \, \text{TI}^2}{\kappa_e} \frac{(\kappa/\kappa_e)^4}{[1 + (\kappa/\kappa_e)^2]^{17/6}} e^{\left[-2(\kappa/\kappa_\eta)^2\right]}. \tag{2.40}$$

  In this work, the turbulence dissipation is estimated with the user-input length scale L_t and time scale $\mathcal{T}$: $\epsilon = U_t^3 / L_t$ [49], with $U_t = L_t / \mathcal{T}$.

- The mode amplitude $\hat{u}^n$ can be obtained using Equation 2.34.

- Time correlation is introduced by means of a digital filter in physical space, which

requires the introduction of the following coefficients:

$$a = e^{(-\Delta t/\mathcal{T})}\,, \tag{2.41}$$

$$b = \sqrt{(1-a^2)}\,. \tag{2.42}$$

After a summary of the setup phase of the method, the run-time procedure is now discussed. During the simulation, three main operations are executed at each timestep. First, four random vectors of N_{turb} elements each are generated, containing the randomly generated variables reported in Table 2.2. These are used to compute the components of the wavenumber vector $\boldsymbol{\kappa}^n$ and unit vector $\boldsymbol{\sigma}^n$. Then, for each quadrature point on the inlet plane the fluctuations of the velocity components are computed using Equation 2.30. Finally, time correlation is introduced through a digital filter: the fluctuating component of the inflow velocity at each timestep is a linear combination of the new randomly generated fluctuations and the velocity fluctuations at the previous timestep:

$$(u'_{i,in})^t = a(u'_{i,in})^{t-\Delta t} + bu'_i\,. \tag{2.43}$$

The new fluctuating components are added to the mean flow profile U_∞ at the inflow plane.

This algorithm can be easily parallelised. In fact, once the random vectors have been generated by the process with rank 0 and scattered to all other processes on the inflow, each rank is responsible for generating the new fluctuating component and updating the velocity in a small partition on the inflow plane. When periodicity is enforced in the wavenumber discretisation, the number of modes is typically reduced significantly, lowering the algorithmic cost. The option of optimising the mesh partitioning by increasing the number of partitions including the inflow plane was considered. However, the *SCOTCH* library is employed for this purpose; the task of optimising the partitioning algorithm without affecting the load balancing of the solver would have most likely been extremely complex. Since the performance obtained was deemed satisfactory, this research question was not pursued further.

2.3.4 Validation of Davidson's method

The recognised strenghts of the method are:

- Extremely flexible for fine tuning of the turbulent scales introduced in the computational domain.
- Deemed to be "physical": it generates turbulence in a manner similar to what would be obtained with a traditional experimental grid (square-mesh arrays of round rods or wires), potentially requiring a shorter development region.
- Minor computational cost (especially for massively parallel simulations), provided that the inflow plane discretisation is distributed over a sufficiently high number

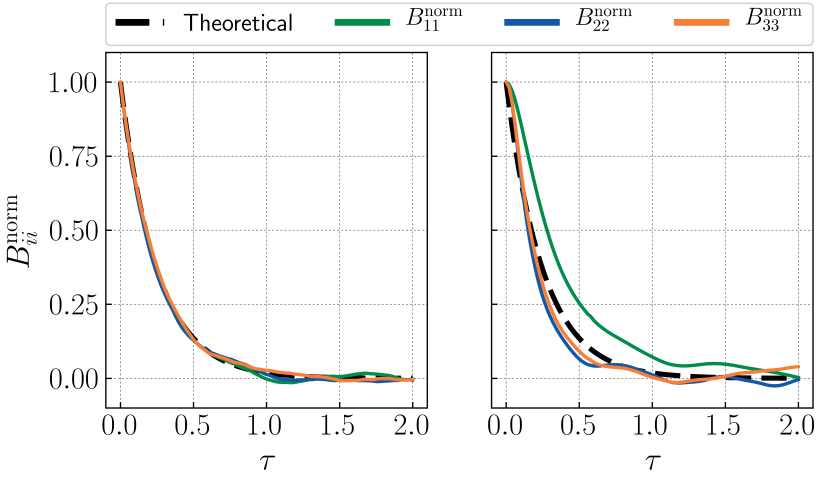

Figure 2.8: Autocorrelation function in the centre of the cross-flow plane, at the inlet (left) and at $x = \pi$ (right). Copyright © 2019 by Rolls-Royce plc.

of processes. For the simulations presented in this thesis, the inflow turbulence generation takes less than 10% of the time-stepping time[2].

In general, a high number of modes should be preferred: not all wavenumbers can be exactly supported by the mesh, so as a consequence the energy allocated to certain modes falls onto a different wavenumber. If little energy is associated to each mode, this causes minor changes to the spectra of synthetically generated turbulence. This conjecture was confirmed by analyzing the TI spatial evolution with various N_{turb} and observing a clustering of the decay curves for $N_{\text{turb}} > 128$.

The synthetic turbulence method theoretically satisfies continuity, but its discrete projection on non-equispaced quadrature points is not strictly divergence free: imposing the three velocity components was found to lead to pressure fluctuations that require long averaging times for convergence of the pressure distribution. This issue is ofter overlooked in the literature, as highlighed by Kim et al. [82]. Mass rescaling on top of the synthetic generation method described above was tested, but without measuring significant improvements, which were instead observed through the introduction of periodicity in the choice of the Fourier modes. In general, removing the fluctuations from the streamwise component and allowing discrete divergence to be naturally enforced by the solver allows to circumvent the issue. Where mode periodicity is not enforced, the pressure statistics in this work are obtained with this approach.

The implementation of the method was extensively tested on a spanwise and vertically periodic development region with domain size $L_x = L_y = L_z = \pi$ at $Re = 5600$, analyzing the effect of all the user-input parameters on auto-correlation, two-point correlation, turbulent intensity decay and velocity spectra. Figure 2.8 shows the imposed auto-correlation function of the method in two different downstream locations, with $L_t = \mathcal{T} = 0.25$. Indic-

[2]Based on tests on the CX2 cluster of Imperial College London.

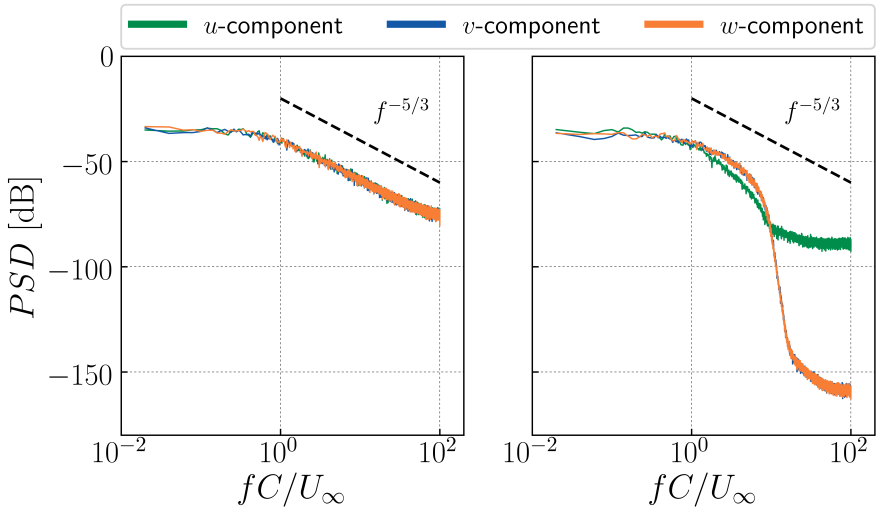

Figure 2.9: Velocity spectra in the centre of the cross-flow plane, at the inlet (left) and at $x = \pi$ (right), with $P = 8$ and $N_z = 128$. Copyright © 2019 by Rolls-Royce plc.

ating with τ the time separation, the normalized auto-correlation is calculated as:

$$B_{ii}^{\text{norm}}(\tau) = \frac{1}{u_{i,\text{rms}}^2} \overline{u_i'(t) u_i'(t+\tau)}\,. \tag{2.44}$$

The Figure shows that the synthetic signal follows the expected correlation profile at the inlet, and coherence is maintained as the turbulent structures evolve spatially. Integration of the autocorrelation function over τ recovers the integral time scale, yielding $\mathcal{T}_u = 0.244$, $\mathcal{T}_v = 0.238$ and $\mathcal{T}_w = 0.254$.

Figure 2.9 shows the velocity spectrum for the three velocity components at the inlet and at $x = \pi$. At the inflow, the broadband spectrum typical of fully developed turbulence is measured, and it agrees well with the $-5/3$ decay law. As the flow evolves spatially, the streamwise velocity spectrum manifests a different behavior from the spanwise and vertical components, which undergo stronger decay at high frequencies. Fourier-spectral/hp element simulations and the corresponding full 3D setup showed a similar behavior, indicating that the use of a hybrid SVV formulation in the Quasi-3D model is not responsible for the discrepancy in the decay of the spectra. Since the same exact decay is observed for the spanwise and vertical velocity spectra, this suggests that the presence of a driving force in the streamwise direction maintains higher energy content at high frequencies. Comparison of results at $P = 8$ with $N_z = 128$ (shown in figures 2.8 and 2.9) with a simulation at $P = 5$ and $N_z = 80$ using structured and unstructured meshes (with both quadrilateral and triangular elements, maintaining the same characteristic dimension) produces exactly the same evolution of the spectra along the streamwise direction. This evidence is used as confirmation that the resolution is sufficient to capture the spectra and that the behavior of spatially evolving turbulence is not appreciably affected by the use of an unstructured mesh.

2.4 Modelling the wake passing effect

In turbomachinery flows, much of the unsteady nature of the suction surface transition process can be attributed to the wakes shed by upstream blade rows, as discussed in Chapter 1. There are virtually no cases of industrial relevance where the inlet flow is not dominated by incoming wakes or pressure waves or, more in general, deterministic unsteadiness. However, it should be highlighted that the flow environment in the back stages of compressors and turbines is characterised by the presence of strong background turbulence, on top of the wakes and potential effects driven by the rows immediately upstream and downstream. In this context, the challenge of generating inlet turbulence (which is traditionally a very academic problem, discussed in detail in the previous section) becomes relevant and should be accounted for. Regardless of the presence of stochastic unsteadiness, the ability to accurately model the wake passing phenomenon is of crucial importance in the context of building the capability for a virtual cascade. A number of studies were dedicated to the development of accurate methodologies to simulate the relative motion typical of rotor-stator interactions. This section first provides a brief summary of the main approaches previously adopted in the context of LPTs, and it further describes the methodology adopted within the *Nektar++* framework.

The first numerical study focusing on the wake passing effect [196] introduced the wake dynamics by means of precursor simulation of a temporally decaying plane wake [198] that was swept across the inflow as a Dirichlet boundary condition. The wake data was subsequently adopted in a number of LPT studies [189, 110, 191, 112].
A viable alternative to the use of precursor data is the adoption of a sliding interface to represent the passing bars in the upstream part of the domain within the same simulation as the linear cascade. Ferrer et al. [46] presented the first implementation of a high-order DG solver with sliding meshes for the incompressible Navier-Stokes equations; more recently, Johnson et al. [72] introduced a sliding grid method in the context of a high-order FD formulation to solve the compressible Navier-Stokes equations, which was later applied to a rotor-stator LPT stage looking at the gap-size effect [133]. More recently, a compressor repeating-stage simulation was employed to study the effect of rotor-stator axial gap [140]. The inflow disturbances were obtained by sampling the velocity profile from a downstream traverse and recycling it by superimposing the approximated fluctuating field onto the inflow boundary conditions.

Lastly, there exists a vast body of literature focusing on immersed boundary methods (IBM), that have been also applied to LPTs. The idea behind an IBM is to model the effect of certain boundary conditions by imposing an external force field as opposed to specifying boundary conditions on the body. Therefore, the computational grid does not need to represent the object in analysis: a conventional Eulerian mesh is adopted to solve for the solvent, and the immersed bodies are represented with Lagrangian markers. A recent example is found in the work by Michelassi et al. [109], where the IB approach first introduced by Goldstein et al. [54] was adopted.

The present thesis adopts an IBM to model the bar passing effect, namely the Smoothed Profile Method (SPM, used intercheangeably with IBM hereinafter) which was first introduced by Nakayama et al. [122]. The method was further developed to adopt a time discretisation based on a high-order semi-implicit splitting scheme [97, 184] and incorporated in the *Nektar++* framework[3]. From a high level perspective, the novelty introduced by the SPM is the idea of representing particle boundaries with a smooth interface layer. The velocity inside each particle is updated through the introduction of a penalty body force term that ensures particle rigidity and imposes the no-slip boundary condition.

Below, a brief summary of the method is presented for the sake of completeness, based on the work by Luo et al. [97]. This section highlights the added computational cost required by the method, which becomes particularly cumbersome in the context of high-Reynolds scale-resolving simulations. The feature is subsequently validated and further analysed in its ability to accurately model the flow past a circylar cylinder at $Re = 3900$, a canonical benchmark flow problem.

2.4.1 The Smoothed Profile Method

Each particle is represented with a smoothed profile that varies between 0 in the fluid domain to 1 inside the rigid particles, as shown in Figure 2.10. Among the several possibilities available, the following is adopted:

$$\phi_i(\mathbf{x}, t) = \frac{1}{2}\left[\tanh\left(\frac{-d_i(\mathbf{x}, t)}{\xi_i}\right) + 1\right]. \tag{2.45}$$

The function $d_i(\mathbf{x}, t)$ represents the signed distance from the boundary, which is positive inside the rigid body and negative in the fluid domain. The overall concentration field is obtained by summation of non-overlapping particles:

$$\phi(\mathbf{x}, t) = \sum_{i=1}^{N_{par}} \phi_i(\mathbf{x}, t)\,. \tag{2.46}$$

The concentration field is used to express the particle velocity field $\mathbf{u}_p(\mathbf{x}, t)$:

$$\phi(\mathbf{x}, t)\mathbf{u}_p(\mathbf{x}, t) = \sum_{i=1}^{N_{par}} \{\mathbf{V}_i(t) + \Omega_i(t) \times [\mathbf{x} - \mathbf{R}_i(t)]\}\phi_i(\mathbf{x}, t)\,, \tag{2.47}$$

with $\mathbf{R}_i$, $\mathbf{V}_i$ and $\mathbf{\Omega}_i$ representing particle position, translational velocity and angular velocity. This velocity can be used to define the total velocity field as follows:

$$\mathbf{u}(\mathbf{x}, t) = \phi(\mathbf{x}, t)\mathbf{u}_p(\mathbf{x}, t) + (1 - \phi(\mathbf{x}, t))\mathbf{u}_f(\mathbf{x}, t)\,. \tag{2.48}$$

The method mathematically enforces both non-slip and non-penetration constraints [97].

[3]The Author acknowledges Mr. Andres Mateo for the implementation of the SPM method.

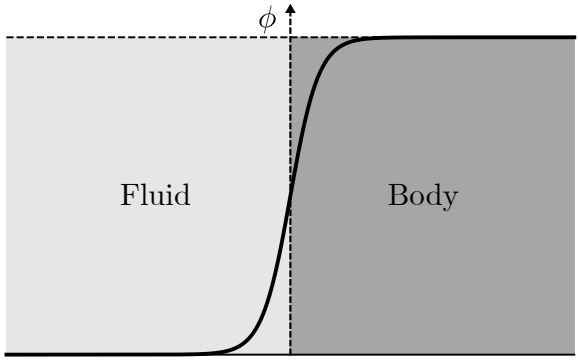

Figure 2.10: Concentration function, smoothly varying between 0 in the fluid to 1 in the rigid particles.

An added force term in the incompressible Navier-Stokes equations is used to solve for the total velocity of Equation 2.48:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla)\mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u} + \mathbf{f} + \mathbf{f}_s \tag{2.49a}$$

$$\nabla \cdot \mathbf{u} = 0 \tag{2.49b}$$

Temporal integration

The method adopts an extension of the VCS scheme introduced in Section 2.1.3. The approach is here included to clarify the modified formulation as well as the sources of increased computational cost.
At each timestep, a first preliminary step consists of calculating the position and orientation of the particles at the new timestep, via an explicit integration scheme. These are then used to recalculate the concentration field ϕ^{n+1}. The time discretisation of Equations 2.49 can be broken down into four macro steps.

Advection and external forces An intermediate velocity $\mathbf{u}_s$ is calculated from explicit integration of the nonlinear terms and body forces:

$$\frac{\hat{\mathbf{u}} - \sum_{q=0}^{J_i-1} \alpha_q \mathbf{u}^{n-q}}{\Delta t} = \sum_{q=0}^{J_e-1} \beta_q \left[-(\mathbf{u} \cdot \nabla)\mathbf{u} + \mathbf{f} \right]^{n-q} . \tag{2.50}$$

Compared to Equation 2.20a, in this case body forces are also included.

Pressure An intermediate pressure field is obtained with an approach analogous to the classical VCS formulation:

$$\nabla^2 p^* = \nabla \cdot \left(\frac{\hat{\mathbf{u}}}{\Delta t} \right) . \tag{2.51}$$

At velocity Dirichlet boundaries, the following Neumann boundary conditions are used:

$$\frac{\partial p^*}{\partial \mathbf{n}} = \sum_{q=0}^{J-1} \beta_q \left[-(\mathbf{u}\cdot\nabla)\mathbf{u} + \mathbf{f} - \nu\nabla\times(\nabla\times\mathbf{u}) \right]^{n-q} \cdot \mathbf{n}\,. \tag{2.52}$$

Solving for p^* allows to update the intermediate velocity:

$$\frac{\hat{\hat{\mathbf{u}}} - \hat{\mathbf{u}}}{\Delta t} = -\nabla p^*\,. \tag{2.53}$$

Viscous effects The viscous contribution is treated implicitly, to obtain the intermediate velocity $\mathbf{u}^*$:

$$\left(\nabla^2 - \frac{\gamma_0}{\nu\Delta t}\right)\mathbf{u}^* = \frac{\hat{\hat{\mathbf{u}}}}{\nu\Delta t}\,. \tag{2.54}$$

The VCS terminates here, but the SPM formulation requires accounting for the presence of rigid particles.

SPM contribution Lastly, the added contribution from the rigid body motion is accounted for. After having updated the concentration ϕ^{n+1}, it is necessary to solve for an extra pressure field p_p:

$$\nabla^2 p_p = \gamma_0 \nabla\cdot\left(\frac{\phi^{n+1}(\mathbf{u}_p^{n+1} - \mathbf{u}^*)}{\Delta t}\right), \tag{2.55}$$

employing the following Neumann BC for pressure:

$$\frac{\partial p_p}{\partial \mathbf{n}} = \frac{\gamma_0\phi^{n+1}(\mathbf{u}_p^{n+1} - \mathbf{u}^*)}{\Delta t}\cdot\mathbf{n}\,. \tag{2.56}$$

The final step consists of updating the total velocity field using the particle velocity field:

$$\frac{\gamma_0\mathbf{u}^{n+1} - \gamma_0\mathbf{u}^*}{\Delta t} = \frac{\gamma_0\phi^{n+1}(\mathbf{u}_p^{n+1} - \mathbf{u}^*)}{\Delta t} - \nabla p_p\,. \tag{2.57}$$

The total pressure field is simply obtained as $p^{n+1} = p^* + p_p$. However, ∇p_p may induce a non-zero velocity inside the particle, which can be suppressed by modifying Equation 2.57 to ensure that the non-penetration constraint is satisfied strictly inside the particle, but introducing some compressibility inside and on the particle boundary:

$$\frac{\gamma_0\mathbf{u}^{n+1} - \gamma_0\mathbf{u}^*}{\Delta t} = \frac{\gamma_0\phi^{n+1}(\mathbf{u}_p^{n+1} - \mathbf{u}^*)}{\Delta t} - (1-\phi)\nabla p_p\,. \tag{2.58}$$

In summary, the SPM formulation solves a similar series of discrete problems as the classical VCS, but with the added contribution of a second Poisson problem for the SPM pressure p_p, in turn producing a significant increase in computational cost.

2.4.2 Flow past a circular cylinder at $Re = 3900$

This work aims at building the underlying knowledge to implement the wake passing effect in a LPT cascade. Extensive validation is first required to assess the validity of the SPM approach. Preliminary validation is based on the analysis of the flow past a stationary cylinder, which is representative of the target geometry. However, a high resolution reference solution is required to discuss the accuracy of the SPM approach.

The flow past a cylinder features a wide variety of flow phenomena, and as such it has traditionally been a reference test case to validate the accuracy and performance of numerical solvers. The wake dynamics changes dramatically depending on the Reynolds number [188]. In the present study the focus is on the popular benchmark value of $Re = 3900$, based on inflow velocity and cylinder diameter; this flow regime is also representative of the relative Reynolds numbers of interest for the wake passing study of Chapter 6. At $Re = 3900$, the flow transitions in the separated shear layers into a fully turbulent wake. A number of experimental studies were carried out [95, 123, 126, 42, 128] and reviewed by Norberg [124]. Several numerical studies focused on this problem adopting LES [11, 19, 86, 47] and, more recently, DNS [98, 42, 190, 91, 195].

Cylinder computational setup

The computational base mesh selected is shown in Figure 2.11, and it is comprised of 1500 quadrilateral elements and 3950 triangular elements. The spanwise domain is extended for $L_z/d = 2\pi$ and a spanwise resolution of $N_z = 64$ is adopted in all the simulations presented. While it is recognised that the spanwise resolution has an important role in capturing the breakdown of structures in the turbulent wake, a full resolution study is out of the scope of the present work, which is aimed at analysing the large parameter space that combines resolution requirements *and* SPM. Uniform inflow boundary conditions are combined with high-order outflow boundary conditions [41]. The choice of boundary conditions for the top and bottom parts of the domain as well as the domain size is important because of their effect on the wake dynamics. Some of the most common approaches include slip-wall [190, 91], farfield [47], periodic [42, 128], Neumann [98] and Riemann invariant for the farfield in a compressible formulation [195]. The end goal of this study is the representation of a moving array of cylinders. Therefore, periodic boundary conditions are adopted for the top and bottom parts of the domain, analogously to the work by Wang et al. [184] which is adopted as reference to compare the SPM results.

Convergence with p-refinement

In the context of spectral/hp element methods, p-refinement is a powerful tool to achieve convergence, therefore employed to ensure DNS-like resolution levels in the near wake. Four polynomial orders are considered: $P = 3, 5, 7, 9$, summarised in Table 2.3 with some of the main quantitative measurements.

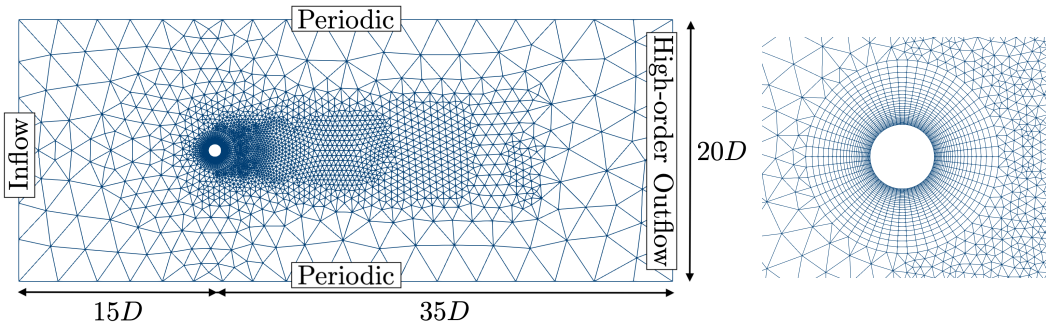

Figure 2.11: Computational base mesh and boundary conditions employed for all cases of flow past a circular cylinder.

Table 2.3: Effect of polynomial refinement on the main flow statistics.

Case	DoF $[\cdot 10^6]$	C_d	$-C_{p_b}$	St	L_{rec}/D	$U_{\min}/U_\infty$	θ_S [°]
$P = 3$	4.05	0.1446	0.8521	0.23	2.0285	-0.312	97.029
$P = 5$	8.72	0.1531	0.8508	0.2075	2.2486	-0.3305	86.007
$P = 7$	15.16	0.1579	0.8822	0.2075	2.1289	-0.3033	86.316
$P = 9$	23.37	0.1556	0.874	0.21	2.1573	-0.314	86.216

In all cases presented, transient simulations were run for at least $T = 100d/U_\infty$. Time independence of the mean wake statistics was verified by running the $P = 7$ case up to $T = 700d/U_\infty$, which corresponds roughly to 146 vortex shedding cycles. However, some statistics are more sensitive than others to the time integration window. Franke & Frank [47] were the first to discuss the importance of the time averaging window, presenting results collected over $T = 200d/U_\infty$ but highlighting that average statistics in the wake would require a longer integration window; Parnaudeau et al. [128] showed that the recirculation length L_r requires over 250 shedding cycles to converge and found strong correlation between recirculation length computed within a certain time window and agreement with experimental results from PIV (Particle Image Velocimetry). Further, they estimated that time averaging over 52 vortex shedding cycles leads to an uncertainty of roughly 10% in the prediction of the peak of fluctuating velocity components. Therefore, caution must be exercised when analysing results sampled with a shorter time window. For further in-depth discussion on the averaging window, the reader may refer to Lehmkuhl et al. [91] where time averages collected over 858 shedding cycles were discussed. In this work, unless otherwise stated, statistics were averaged over $T = 400d/U_\infty$, which corresponds to roughly 83 vortex shedding cycles. The experimental data of Parnaudeau et al. [128] are included in this section; however, it is reminded that, on top of the sources of uncertainties specified above, the experiments analysed the effect of an isolated cylinder as opposed to an array of cylinders. Since the problem is of a slightly different nature, the presence of discrepancies is to be expected, even though the degree of uncertainty is

Table 2.4: Percentage relative error of main flow statistics with respect to case $P = 9$.

Case	C_d	$-C_{p_b}$	St	L_{rec}/D	$U_{\min}/U_\infty$	θ_S [°]
$P = 3$	7.063	2.507	9.524	5.973	0.611	12.542
$P = 5$	1.577	2.659	1.19	4.229	5.284	0.242
$P = 7$	1.492	0.934	1.19	1.319	3.403	0.116

difficult to quantify more precisely.

The main quantitative flow statistics are reported in Table 2.3, and the relative error with respect to the most refined case $P = 9$ is shown in Table 2.4. Under-resolution fails to accurately capture a number of parameters, particularly the drag coefficient, Strouhal number, recirculation length and separation angle. θ_S is defined as the point where the axial wall shear stress component becomes negative, with a reference system starting from the stagnation point where $\theta = 0$. At low order the recirculation length is underpredicted, but the minimum velocity is estimated correctly. An increase in resolution improves all statistics (except for the minimum streamwise velocity, but as highlighted previously the accurate prediction of the recirculation bubble length might require a longer integration time), with most relevant flow parameters at $P = 7$ within 2% of the results at $P = 9$.

Considering the mean wake profiles shown in Figure 2.12, it can be appreciated that the under-resolved case $P = 3$ has a rather different behaviour compared to the more highly refined cases. In terms of streamwise velocity (Figure 2.12a), case $P = 3$ presents a V-shaped profile in the very near wake, while increased resolution produces U-shaped profiles in very close agreement with experiments from Parnaudeau et al [128]. The curves $P = 7$ and $P = 9$ are completely overlapped. In the near wake behind the cylinder, nonlinear effects are evident in the behaviour of vertical velocity (Figure 2.12b): $P = 3$ smears the peaks near the centerline, while $P = 5$ overestimates them. Convergence is achieved at $P = 7$. Analogous behaviours are observed for the evolution of Reynolds shear stress, shown in Figure 2.12e. Concerning the horizontal fluctuating components, the case $P = 9$ clearly shows a double peak feature that is linked to the transitional state of the boundary layer. The emergence of a second peak near the centreline for $x/d > 1.06$ is physically motivated by shear layer flapping that takes place through its inflection by primary vortex formation [128]. Figure 2.12c shows that both $P = 3$ and $P = 5$ underestimate the peak, especially at $x/d = 1.54$, thus failing to accurately capture the transitional state of the shear layer. The agreement with experimental data is very close at $x/d = 1.06$ and $x/d = 2.02$, while the peak value is slightly different at the intermediate location $x/d = 1.54$. In the vertical fluctuating component (Figure 2.12d) case $P = 7$ is fully converged, but the peak is underestimated compared to the experiments. Lastly, the wake profiles of turbulent production are shown in Figure 2.12f, which clearly suggest under-resolution in the near wake for $P = 3$, given its oscillatory behaviour. Analogously to the other wake profiles, case $P = 7$ is overlapped to $P = 9$.

Figure 2.13a shows the horizontal component of velocity along the centreline, and it

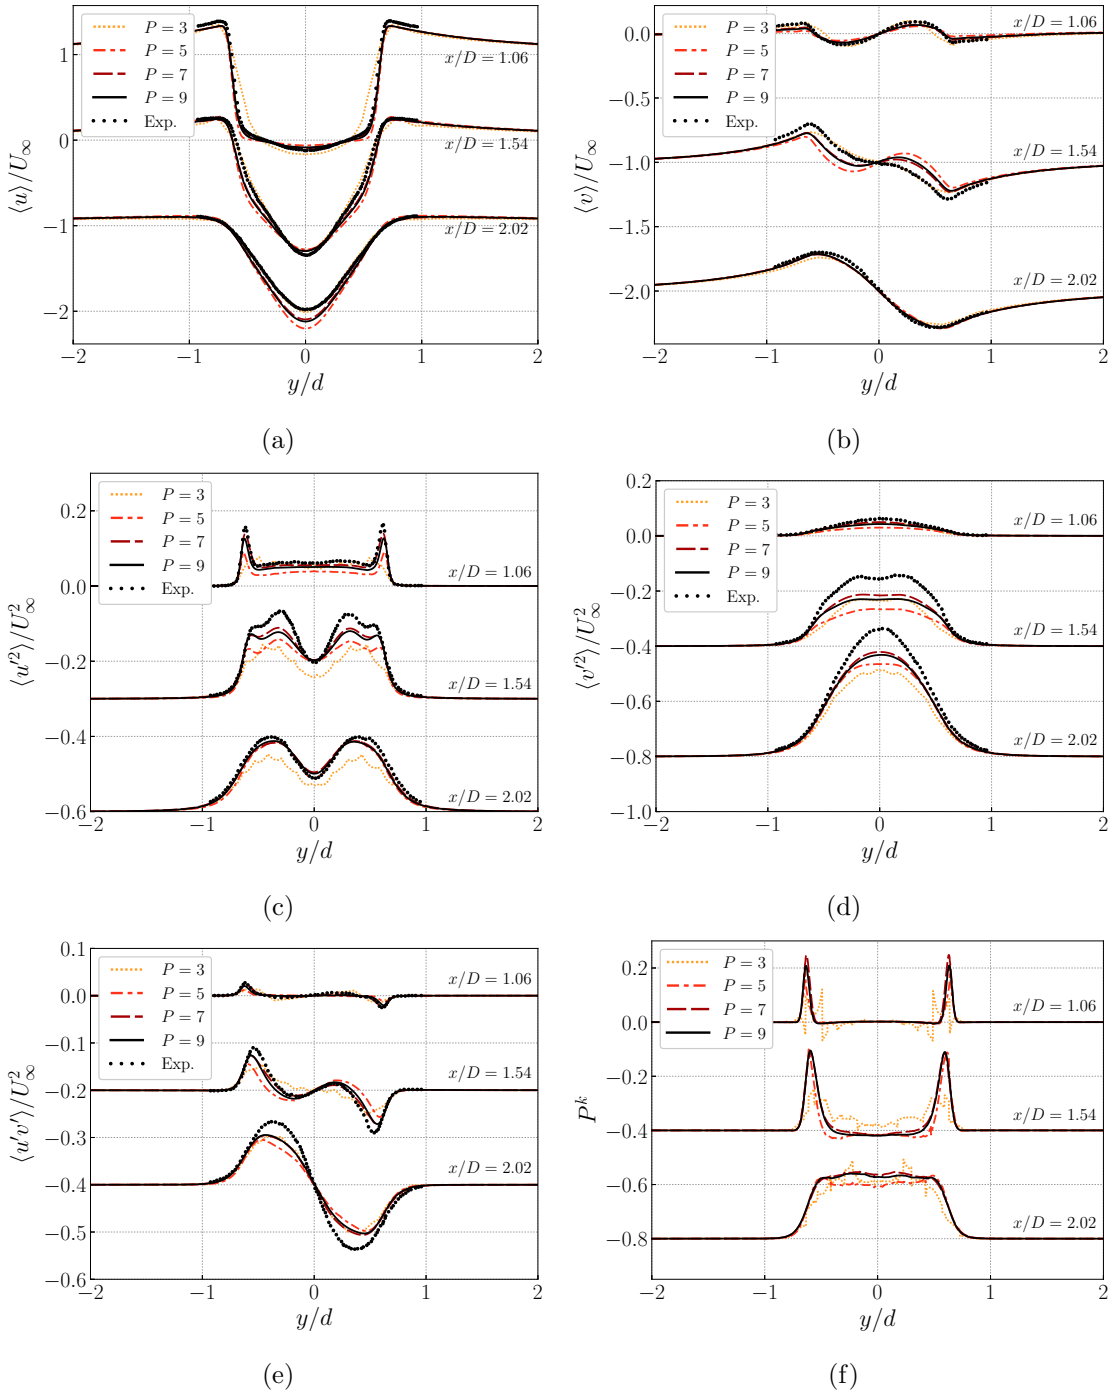

Figure 2.12: Time-averaged wake profiles at three streamwise locations with $P = 3, 5, 7, 9$. Mean velocity profiles: (a) streamwise velocity, and (b) vertical velocity. Reynolds stresses (c) streamwise, (d) vertical, and (e) shear stresses. (f) turbulence production. Experimental data from Parnaudeau et al. [128]. Profiles (a,b) are not scaled, while (c-f) are scaled by a factor of 2.

helps to clarify the reason for the overshooting behaviour of $P = 5$. In fact, this case shows a longer separation bubble. Convergence is again achieved at $P = 7$. The pressure coefficient is finally shown in Figure 2.13b. $P = 3$ is clearly under-resolved along the

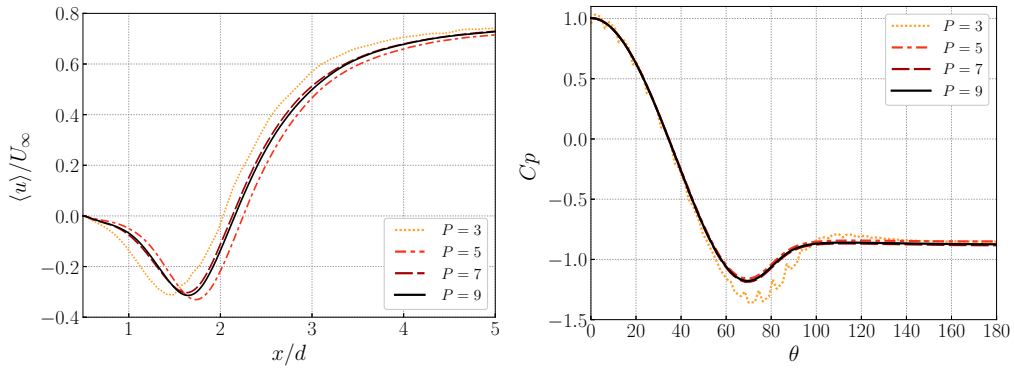

(a) (b)

Figure 2.13: Effect of p-refinement on (a) time-averaged streamwise velocity along the centreline, and (b) pressure coefficient distribution on the cylinder surface.

cylinder walls, as denoted by a strong oscillatory behaviour. Analogously to most of the other statistics, the curves $P = 7$ and $P = 9$ are completely overlapped.

The DNS data at $P = 9$ are taken as reference for validation of the SPM feature, which is discussed in the following section.

Convergence of SPM with stationary circular cylinder at $Re = 3900$

Validation of the SPM feature requires separate convergence phenomena:

1. The SPM method converges to a certain solution with increasing resolution, and at fixed interface thickness ξ. However, this solution may not be correct.

2. As the resolution *and* interface thickness are refined, the method converges to the *same* solution as the reference DNS simulation.

Having previously ensured convergence of DNS results, the attention is now focused on SPM results and their comparison with DNS at $P = 9$. The most significant findings of an exhaustive sensitivity study are included in this section. As part of this preliminary testing phase, both formulations shown in Equations 2.57 and 2.58 were analysed. No significant discrepancy in the main statistics was observed, thus the classical formulation was adopted. The use of regular elements versus Taylor-Hood elements (where pressure is solved at one order lower than velocity) was also investigated, yielding overlapping results in all the main time-averaged statistics, particularly the velocity wakes. To ensure consistent numerical approach with the rest of the thesis, the Taylor-Hood formulation was adopted.

In order to verify points (1) and (2) (i.e. convergence at fixed and variable interface thickness ξ, respectively), four increasingly refined meshes (labelled M1, M2, M3 and M4 in Figure 2.14) were considered. The choice of mesh M1 is justified in greater detail in Chapter 6. The characteristic element size of M2 in the region where ϕ is non-zero was halved compared to M1, significantly increasing the total element count. The process was

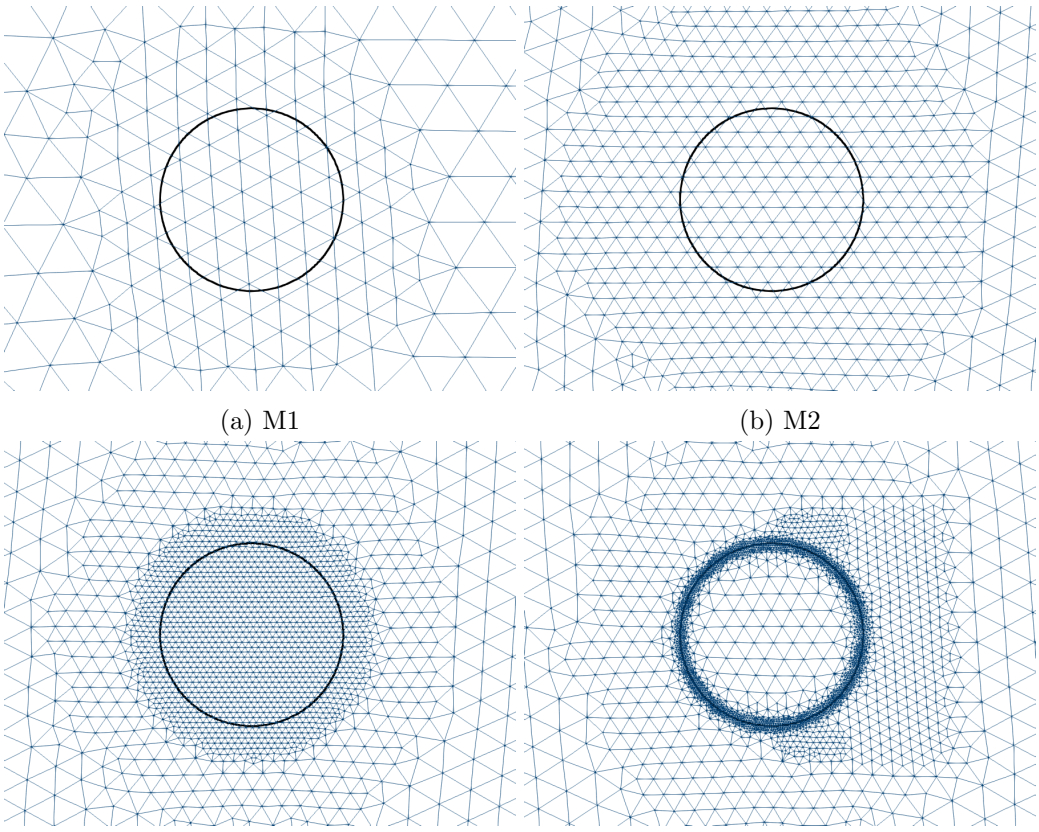

(a) M1 (b) M2

(c) M3 (d) M4

Figure 2.14: Computational base mesh employed to show results convergence at fixed ξ (a, b, c) and variable ξ (c, d).

repeated to generate mesh M3. Mesh M4 was instead obtained by retaining the resolution of M2 in the central region of the cylinder, and considerably increasing the resolution across the interface region. As shown in Table 2.5, M1, M2 and M3 were employed to investigate point (1), while M3 and M4 were employed for point (2). The Table also reports grid specifications, computational cost and details of the SPM setup for the two sets of experiments, labelled "Test 1" and "Test 2".

Wang et al. [184] introduced the following criteria to ensure resolution requirements

Table 2.5: Summary of mesh specifications adopted for validation of SPM.

Mesh	# Elements	P	N_z	DoF $[\cdot 10^6]$	Δt	T	Test 1	Test 2
M_1	1550	7	64	3.57	1e-3	400	$\xi = 0.01$	
M_2	4906	7	64	11.30	1e-3	400	$\xi = 0.01$	
M_3	7612	7	64	17.54	2.5e-4	200	$\xi = 0.01$	$\xi = 0.01$
M_4	10036	7	64	23.12	1.25e-4	200		$\xi = 0.0048$
		8	64	28.90	1e-4			$\xi = 0.0015$

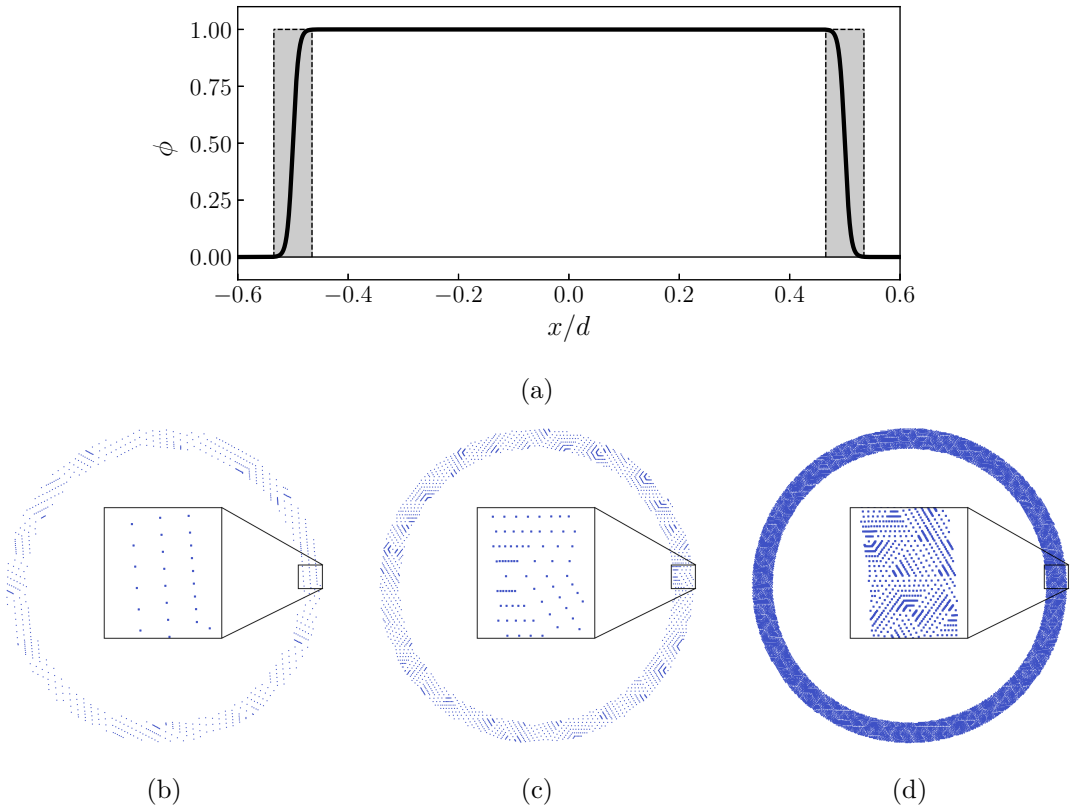

Figure 2.15: (a) Sketch of the smoothed profile ϕ in the case of a stationary cylinder with $\xi = 0.01$. The grey areas are bounded with a tolerance of $10^{-}3$. (b-d) Quadrature points of M1, M2 and M3 within the grey regions with $P = 7$.

are met:

$$\frac{h}{P} \leq \xi\,, \tag{2.59}$$

$$\frac{L_z}{N_z} \leq 6\xi\,. \tag{2.60}$$

where h here refers to the length of the element edge across the boundary of the SPM. In order to test the behaviour of the method in the context of a relatively large interface thickness, $\xi = 0.01$ is selected for Test 1. With this choice of the interface thickness, case M1 at $P = 7$ does not satisfy 2.59, while M2 and M3 do. To provide a visual representation of the resolution level in this case, Figure 2.15a shows the particle function ϕ, highlighting in grey the transition area defined as the region where ϕ is within 10^{-3} of its value outside the interface region. Below, the effective high-order resolution of meshes M1-M3 within the grey region is shown. It should be noted that the equispaced output of the quadrature points is shown for visual clarity, and therefore they do not exactly correspond to the physical solution points. Nonetheless, this provides an indication of the level of resolution available in the three cases.

The time-averaged wakes of streamwise velocity are shown in Figure 2.16, where the

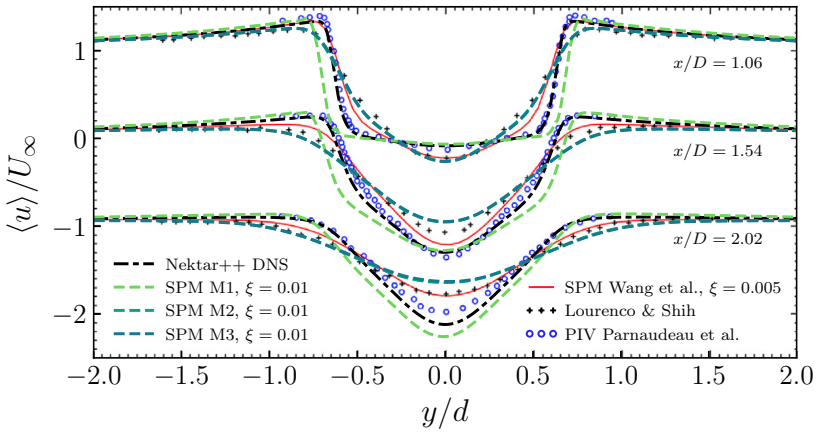

Figure 2.16: Wake profiles of streamwise velocity with increasing resolution in the $x - y$ plane and constant interface thickness $\xi = 0.01$. Figure edited from [184].

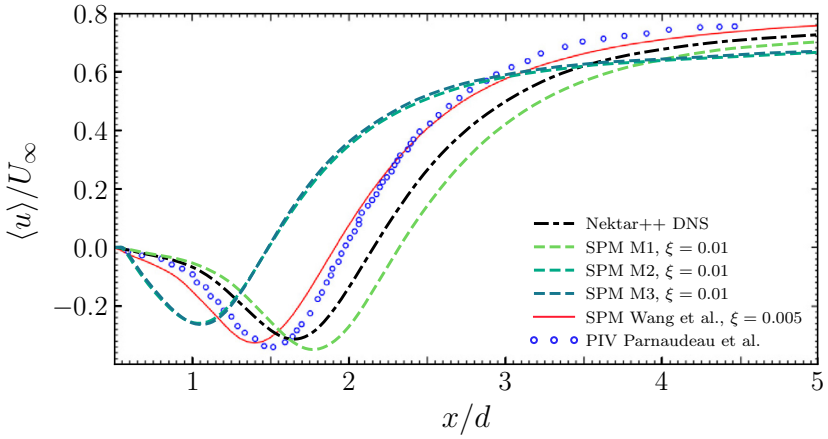

Figure 2.17: Streamwise velocity along the centreline with increasing resolution in the $x - y$ plane and constant interface thickness $\xi = 0.01$. Figure edited from [184].

Nektar++ results are superimposed with previous SPM results [184] and experimental measurements from Parnaudeau et al. [128] and Lourenco & Shih [95]. The DNS data agree very closely with experiments [128] both in terms of wake width and depth. The SPM results obtained in the coarse mesh (M1) retain the same shape as the DNS results, while overestimating both wake width and depth. The results saturate in M2 and M3, where an increase in resolution does not show any further changes in the wake profiles. This suggests that the resolution of M1 is insufficient to accurately represent an interface thickness of $\xi = 0.01$, but resolution levels superior to M2 introduce no further improvement in the agreement with DNS results. The same considerations apply to the streamwise velocity along the centreline, shown in Figure 2.17. In this case, the M1 profile is actually closer to the DNS values compared to M2 and M3. More specifically, the coarse case

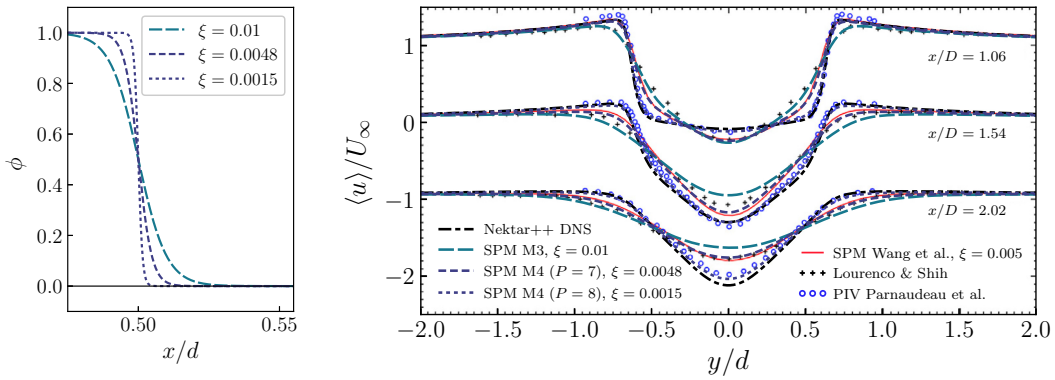

(a) (b)

Figure 2.18: (a) Analytical SPM profiles across the cylinder boundary with decreasing thickness ξ, and (b) resulting wake profiles of streamwise velocity with increasing resolution in the $x - y$ plane and reduced interface thickness ξ. Figure edited from [184].

provides a closer agreement in terms of recirculation length, although it overestimates the maximum recirculation speed compared to DNS results. The offset between DNS data and the experimental measurements may be linked to the presence of background noise in the experiments, causing early transition [190]. Higher resolution in the SPM cases underestimates the bubble length as well as the far-wake centreline velocity. This figure highlights that the use of SPM with large ξ converges to a solution as the resolution is increased, but different from the reference, and therefore supports the verification of point (1).

In order to verify point (2), it is necessary to introduce further mesh refinement the spectral/*hp* planes while retaining the same number of DoF in the spanwise direction, and concurrently reducing the interface thickness. It should be highlighted again that SPM simulation results are very sensitive to both ξ and mesh resolution. Since SPM was first proposed [122], a number of studies made attempts to find an optimal way to tune ξ (Luo et al. [97], Mohaghegh & Udaykumar [113] to name a few), but most of the estimates depend on the simulation timestep Δt, which is not a desirable property.
Wang et al. [184] proposed an approach based on an estimate of the momentum thickness in the context of the flow past a bluff body at moderate Reynolds number ($80 \leq Re \leq 10^4$), which is within the range of interest of this work. Schlichting & Gersten [150] provided an estimate of the smallest value of the momentum thickness at $x = \frac{\pi d}{4}$:

$$\theta = \frac{0.664}{\sqrt{0.25 Re \cdot \pi}} \,. \tag{2.61}$$

The interface thickness is then calculated as $\xi = \epsilon\,\theta$, by introducing a tuning parameter with value $\epsilon = 0.2$ for two-dimensional simulations, and $\epsilon = 0.4$ for three-dimensional

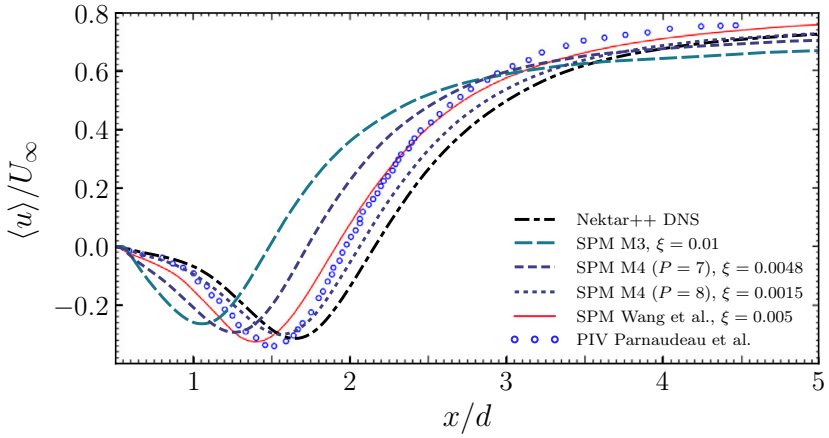

Figure 2.19: Streamwise velocity along the centreline with increasing resolution in the $x - y$ plane and reduced interface thickness ξ. Figure edited from [184].

simulations. According to this estimate, at $Re = 3900$ interface thickness should be $\xi = 0.0048$ for accurate results. Therefore, this setup is implemented in M4 with $P = 7$, where the $x - y$ resolution is largely sufficient to capture the interface. This case provides insight into the capability of a *physics-based* approach to select ξ, upon the condition that the resolution is sufficient. Lastly, the condition previously shown in Equation 2.59 is employed to find the smallest ξ that can be theoretically be represented with the resolution of M4 at $P = 8$, to assess the result of adopting a *resolution-based* approach to the selection of ξ. Equation 2.59 at the resolution level of M4 with $P = 8$ yields $\xi = 0.0015$. The three cases compared are summarised in Table 2.5, and their analytical ϕ function is shown in Figure 2.18a. For M3 and M4, given the high computational cost of this simulation, time averages were collected over a reduced time interval of $T = 200d/U_\infty$, which was deemed sufficient to estimate a number of mean flow features including velocity wakes at the level of accuracy required in this section.

Time-averaged streamwise velocity wakes at three streamwise stations are shown in Figure 2.18b. At $x/d = 1.06$, the SPM wake is V-shaped with large $\xi = 0.01$, and U-shaped with small $\xi = 0.0015$. With $\xi = 0.0048$ the wake profile at all streamwise stations reproduces the results obtained by Wang et al. [184] where a similar value of $\xi = 0.005$ is adopted. Case $\xi = 0.0015$ is overlapped with DNS wakes at $x/d = 1.06$ and $x/d = 1.54$, with a small discrepancy in the centreline of the wake at $x/d = 2.02$. Therefore, progressive convergence towards DNS results is demonstrated in the near-wake streamwise velocity at all locations analysed. The streamwise evolution along the centreline is shown in Figure 2.19. Case M3 with $\xi = 0.01$ underestimates the recirculation length and the maximum recirculation velocity, as well as the far wake velocity at $x/d = 5$. In case M4 with $\xi = 0.0048$, the centreline velocity is not as close to the results by Wang et al. [184] as apparent from the near wake profiles. However, the decrease in interface thickness produces a notable improvement compared to $\xi = 0.01$. The most refined case (M4

with $\xi = 0.0015$) is extremely close to DNS results until the negative peak in centreline recirculation velocity at $x/d \approx 1.7$, and further downstream for $x/d > 3.5$.

This section has provided a validation of the SPM feature aimed at supporting the adoption of this method to model the wake passing effect. An important limitation can be highlighted: in the flow past a circular cylinder, the *physics-based* approach to the tuning of ξ provides less accurate time-averaged velocity results compared to the *resolution-based* approach. In other words, to obtain comparable accuracy to a wall-resolved DNS solution where conventional boundary conditions are imposed, a combination of extremely high resolution for the interface region and small interface thickness ξ are required. This in turn yields a steep increase in computational cost, which becomes practically unfeasible for all realistic applications where the SPM body is not stationary. This topic is explored in further detail in Chapter 6, where a novel practical approach is developed to accurately model the wake passing effect despite under-resolution in the SPM region.

Chapter 3

A resolution study

A relatively small number of previous studies have devoted attention to extensively exploring the convergence properties of HOM. The effect of a larger spanwise computational domain and various polynomial order setups were investigated by Uranga et al. [174] on a wing profile at low Reynolds number using a DG method. Different mesh configurations were also tested by Raverdy et al. [142]. A more recent paper by Pichler et al. [135] confirms that the grid resolution used in several previous studies is qualitatively sufficient for accurately capturing the main flow features dominated by large scale structures (e.g. pressure distribution and mean wake profiles). Second order statistics in the wake are recommended for assessing grid convergence: TKE and TKE production are large scale turbulence quantities, while turbulent dissipation is the most stringent measure to determine grid convergence to DNS level.

The T106A LPT test case [165] is recognized and adopted as an ideal benchmarking case. The potential of the numerical method is leveraged through the built-in refinement capability of element refinement (h) and polynomial order (P). The aim is to assess a framework of best practices for scale-resolving simulations of turbomachinery components using HOM and verifying recommendations that provide accuracy with optimized computational cost and convergence. The number of computational parameters that can be explored is substantial, and independently analysing all relevant combinations is unfeasible. Therefore, several key computational setups are selected for a parametric study, and the results are compared and discussed to verify statistical convergence. The broad range of length and time scales that dominate turbomachinery flows is such that each flow property has different requirements in terms of grid resolution and time averaging window, which complicates most attempts of generalization.

The next section discusses the computational approach in detail, as well as the main physical features of the linear cascade at a moderate Reynolds number. The techniques for data reduction and calculation of performance parameters employed across all following Chapters are introduced together with quantitative results, subdivided into various parts: (1) $x - y$ plane p-refinement, also discussing the effect of different choices for the SVV parameters in the spanwise direction, and the importance of an appropriate averaging

window; (2) effect of different choices of the spanwise domain L_z, and (3) impact of the number of Fourier planes N_z in the spanwise direction. Then, a comparison with an experimental campaign from Cambridge University [27] is established, and a brief overview of the main findings is presented.

3.1 Problem setup and numerical approach

The test case considered in this Chapter is a T106A LPT, a well known openly available geometry. Previous works focusing on the wake passing effects adopted a modified inflow angle, following the experiments by Stadtmüller et al. [161].

In this Chapter, clean inflow boundary conditions are employed: they allow for high sensitivity of the flow statistics to the effect of the various setups, without any influence introduced by background fluctuations or wake passing effects. However, this setup arguably is not closely representative of a realistic industrial flow. The design inflow angle of $\alpha = 37.7°$ was employed in combination with a moderate inflow Reynolds number of $Re_\infty = 50000$ (based on inflow velocity $U_\infty = 1$ and chord $C = 1$), which corresponds approximately to $Re_2 = 88450$, based on mixed-out velocity magnitude at the exit measurement plane. Further details on various averaging techniques are discussed in Section 3.1.1. The design exit angle is 63.2°. The inflow plane is situated $1\,C_{ax}$ upstream of the leading edge ($C_{ax} = 0.859\,C$). High-order outflow conditions [41] were applied to the outflow plane, located $1.5\,C_{ax}$ downstream of the trailing edge, while periodic conditions were enforced in the pitchwise direction, with pitch-to-chord ratio $P_y/C = 0.799$. The upstream and downstream measurement planes are respectively $0.5\,C_{ax}$ upstream of the leading edge (LE) and downstream of trailing edge (TE).

One of the bottlenecks in the development of high-order solvers and their incorporation in an industrial setting is efficient and robust high-order mesh curving capability [13]; the *NekMesh* platform [172] allows to tackle this challenge. The computational mesh of the T106A blade is shown in Figure 3.1. It was generated following a series of criteria, resulting from extensive trial and error testing. First, the geometrical points distribution was slightly smoothed near the trailing edge and leading edge, since a few of the original data points [165] presented an unusual distribution. The resulting 2D loading of the original geometry and the modified one showed no differences. To generate curved high-order elements, four splines joined in low curvature regions were fitted through the blade surface points. The O-mesh around the blade was generated with 8 base mesh layers with a geometric progression of 1.5 and total thickness corresponding approximately to the boundary layer thickness at the separation point ($n_{e,\mathrm{sep}}/C \approx 0.012$). The wake mesh is highly resolved in the whole pitchwise domain up to a streamwise extent of $0.75\,C_{ax}$ downstream of the trailing edge. The resulting base mesh has 2632 quadrilateral elements in the O-mesh and 6031 triangular elements in the unstructured mesh. The coarsening outflow region also promotes the dissipation of strong vortical features, enhancing the stability of the simulations especially during transient flow evolution.

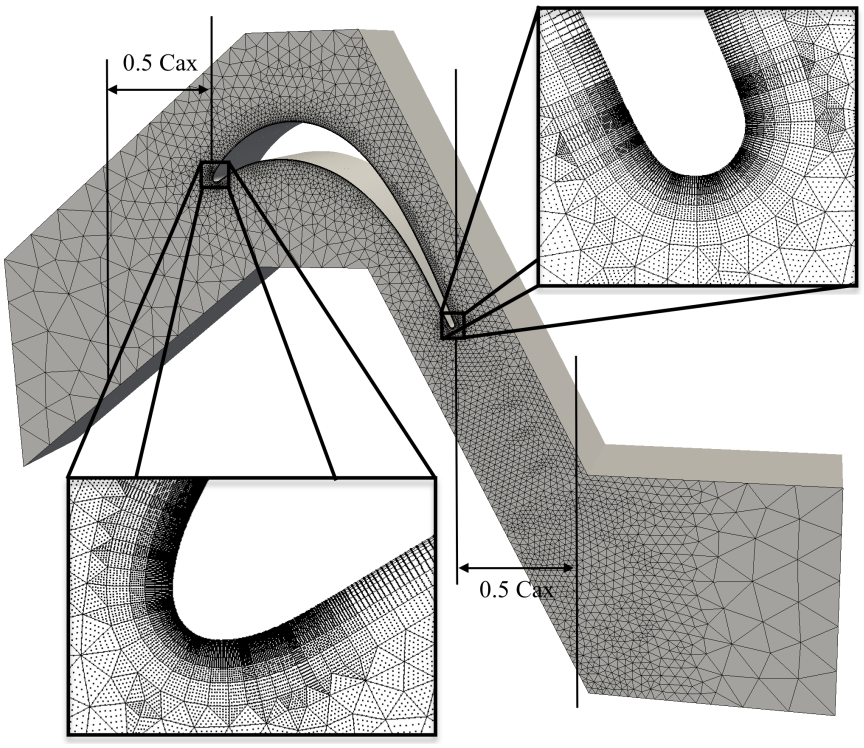

Figure 3.1: Computational base mesh and (zoomed) high-order LE and TE mesh with $P = 7$. The boundary layer O-mesh has 8 layers with geometric progression 1.5.

In the framework of high-order spectral/*hp* element methods, the local mesh spacing h needs to be selected at a scale comparable to that of the flow features, on top of which p-refinement at fixed h allows for exponential convergence with a sufficiently smooth solution. To ensure that the base mesh spacing selected is sufficiently refined in the near-wall region, another mesh with approximately half cell size along the blade and twice as many layers in the O-mesh was tested with the same setup, yielding identical results in the mean 2D surface distributions.

The near wall resolution of the high-order mesh is shown in Figure 3.2. For the calculation of the resolution in wall units, the wall-shear stress in the blade-normal direction $\hat{\mathbf{n}}$ is computed:

$$\tau_w = \mu \frac{\partial u}{\partial n}\bigg|_{\text{wall}}, \tag{3.1}$$

and used to estimate the friction velocity u_τ:

$$u_\tau = \sqrt{\frac{\tau_w}{\rho}}. \tag{3.2}$$

The mesh spacing in wall units can be written along the curvilinear coordinate s as:

$$\Delta x_i^+(s) = \frac{(\Delta x_i(s)/P)u_\tau}{\nu}. \tag{3.3}$$

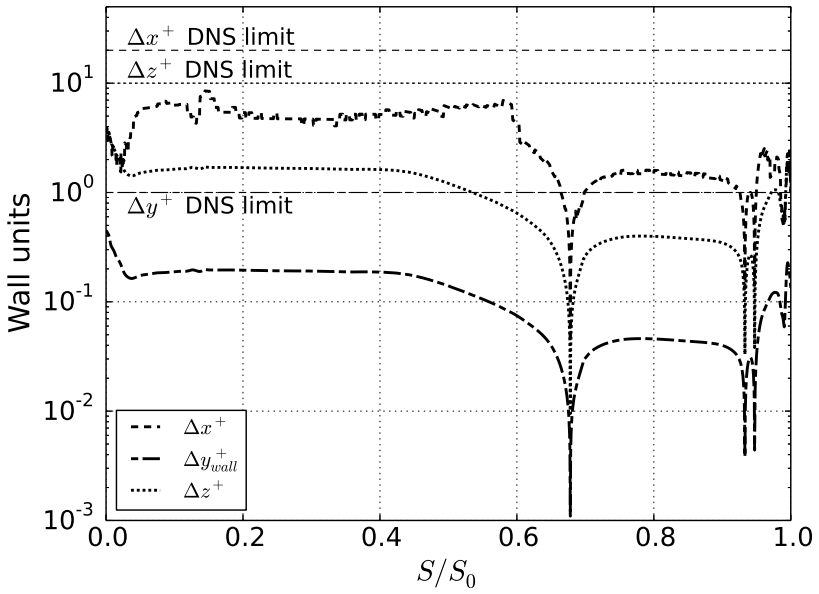

Figure 3.2: Near wall mesh resolution for the case P7L02N96. Copyright © 2018 by Rolls-Royce plc.

The ratio h/P is obtained by dividing the base mesh spacing $\Delta x_i(s)$ by the polynomial order P. The resolution obtained is within typical limits recommended for DNS [51], suggesting that the current setup achieves high resolution standards on the blade surface. Nonetheless, the viscous wall coordinates are derived under the assumption of fully turbulent boundary layer; it is not yet clear how generally accepted resolution requirements for turbulent BL would translate to the case of transitional and separated boundary layers, so this result is mainly indicative. However, it suggests that the resolution employed is sufficient to capture the turbulent structures evolving in the separated portion of the suction surface.

When simulating a new test case or geometry, the simulations are impulsively started in 2D from uniform initial conditions at very low Reynolds number, progressively increasing P and Re concurrently, until the flow is fully established at the resolution level and Reynolds numbers of interest. Subsequently, the computational domain is extended in the third dimension, perturbing the flow in the spanwise direction with white noise to trigger turbulent mixing and breaking down coherent two-dimensional structures. Due to the high sensitivity of the flow statistics to inflow conditions and parameters variation, it is important to ensure that transient phases are extinguished and the flow is statistically converged when sampling time averages. Unless otherwise specified, in every case presented in this work the time evolution of three quantities was monitored to determine when time-averaged statistics could be sampled. The quantities considered are: (1) the volume integral of the kinetic energy, (2) the evolution of aerodynamic forces and (3) the modal energy distribution across the Fourier modes. In general, the highly unsteady nature of these measures limits possibility to identify the end of the transient phase rig-

Table 3.1: List of cases analysed, grouping the setups compared in the Chapter, excluding cases for experimental comparison. The column DoF indicates the degrees of freedom *per variable*. The computational time is calculated based on average timestepping only, and it indicates the approximate computing time required for $T = 1C/U_\infty$ on 100 Intel® Xeon® E5-2680 v4 processors, assuming linear scaling. Case P7L04N192 was ran on a different system and therefore omitted for consistency.

Case	P	L_z/C	N_z	M/N_z	DoF $[\cdot 10^6]$	Comp. time [h]	CFL
P3L02N96	3	0.2	96	0.75	9.83	2 h 18'	0.17393
P5L02N96	5	0.2	96	0.75	21.2	5 h 33'	0.48492
P7L02N96	7	0.2	96	0.75	37.0	20 h 48'	0.47516
P9L02N96	9	0.2	96	0.5	57.1	56 h 16'	0.62859
P7L02N32	7	0.2	32	0.5	12.3	13 h 2'	0.47543
P7L02N48	7	0.2	48	0.5	18.5	17 h 48'	0.47514
P7L02N72	7	0.2	72	0.5	27.8	19 h 47'	0.47505
P7L02N128	7	0.2	128	0.75	49.3	31 h 51'	0.47508
P7L0025N12	7	0.025	12	0.5	4.63	2 h 36'	0.47437
P7L005N24	7	0.05	24	0.5	9.25	4 h 51'	0.47531
P7L01N48	7	0.1	48	0.75	18.5	8 h 42'	0.47499
P7L015N72	7	0.15	72	0.75	27.8	16 h 18'	0.47576
P7L04N192	7	0.4	192	0.75	74.0	-	0.47540
P7SVV	7	0.2	96	0.5	37.0	23 h 13'	0.47527
VARP	Variable	0.2	96	0.5	-	7 h 42'	0.33626

orously, but such an approach was found to be robust and suggested that $T = 20C/U_\infty$ chord-based flow through times were sufficient in most cases analysed. Time-averages were subsequently sampled for further 24 chord-based flow through times in all cases, having ensured convergence of all properties within such averaging frame. Further details on time convergence are provided in Section 3.3.6.

As mentioned in the introduction of the Chapter, the analysis focuses on the effect of resolution on the main blade statistics. Table 3.1 presents a summary of the test cases analysed, together with their respective performance measurements. As a side note, the CFL number indicated in the column on the right is relative to the timestep employed for each specific case. In particular, in cases P3L02N96 and P5L02N96 it was possible to double the timestep Δt (while retaining the same sampling frequency for all statistics).

3.1.1 A note on various averaging techniques

Traditionally, the type of averaging scheme for various flow quantities has predominantly relied on matters of opinion or convenience. This is particularly true for the inherent limitations present in experimental work, often limited to a crude area average. However, the increasing role of CFD provides an opportunity to select the correct averaging technique, given the non-intrusive availability of large datasets. A paper by Cumpsty & Horlock [32] focuses on this topic in detail. Each of the various averaging techniques available is appropriate depending on the physical problem being analysed. The topic is of considerable

importance: internal flows are typically strongly nonuniform, and component performance is assessed based on some type of data reduction to a single value. Therefore, the way averaging is carried out has important repercussions in the estimation of performance.

This work is mainly concerned with mixed-out averaging and, in some cases, mass averaging. However, very often in practical applications the quantitative difference between the various averaging approaches is very limited [32]. This was checked in all problems analysed in this thesis. Despite the fact that values of total pressure and velocity were observed to differ no more than 0.2% by comparing the various averaging methods, the uncertainty associated to these variables is enhanced when derived variables are calculated. For example, a relative difference of $\approx 5\%$ was found in the calculation of total pressure losses. It is therefore important to clarify the distinction and purpose of various averaging methodologies.

Area average The appropriate average for the static pressure is derived from the net force exterted on the control domain [32]:

$$F = \overline{p} P_y = \int_0^{P_y} \langle p(y) \rangle_{z,t} \mathrm{d}y \,. \tag{3.4}$$

Mass average Cumpsty & Horlock [32] argued that for compressor and turbine flow, where the focus is on the work input or output, pressure ratio and efficiency, the correct total pressure averaging approach is called work-average. However, this averaging technique involves the use of thermodynamic quantities that assume flow compressibility. In practice, mass-averaging is very commonly adopted to estimate stagnation pressure: in the limit as the pressure variation tends to zero, the work-average tends to the mass-average:

$$p_t^A = \frac{\int_0^{P_y} \langle p_t(y) \rangle_{z,t} \langle u(y) \rangle_{z,t} \mathrm{d}y}{\int_0^{P_y} \langle u(y) \rangle_{z,t} \mathrm{d}y} \,. \tag{3.5}$$

Mixed-out average Several averaging methods produce mean values that are not consistent with all conservation laws [132]. The mixed-out state is instead determined from the initial nonuniform state by enforcing conservation of mass, momentum and energy between the two measurement stations, yielding consistent results for the loss of an airfoil row. The mixed-out approach accounts for all sources of loss downstream of the blade, and it is invariant to the axial location [139]. Since this work focuses on incompressible flows, a number of simplifications can be introduced. The assumption $\rho = 1$ is made. Applying mass conservation between any station and a point further downstream where the flow is assumed to be completely mixed out yields:

$$P_y \overline{u} = \int_0^{P_y} \langle u \rangle_{z,t} \mathrm{d}y \,. \tag{3.6}$$

This is effectively an area average applied to the streamwise component of velocity. Applying momentum conservation and isolating the x- and y- components provides two further equations for the vertical component of velocity and for pressure:

$$P_y(\overline{p} + \overline{u}^2) = \int_0^{P_y} (\langle p \rangle_{z,t} + \langle u \rangle_{z,t}^2) \mathrm{d}y \,, \tag{3.7}$$

$$P_y \overline{u}\,\overline{v} = \int_0^{P_y} (\langle u \rangle_{z,t} \langle v \rangle_{z,t}) \mathrm{d}y \,. \tag{3.8}$$

From these three equations one finds the mixed-out velocity $U^M = \sqrt{\overline{u}^2 + \overline{v}^2}$, static and total pressure:

$$p_t^M = \overline{p} + \frac{1}{2}\rho(\overline{u}^2 + \overline{v}^2) \,, \tag{3.9}$$

as well as the mixed-out exit angle:

$$\alpha^M = \arctan\left(\frac{\overline{v}}{\overline{u}}\right) \,. \tag{3.10}$$

From a practical point of view, at times it is difficult to conduct experiments without including a fluctuating component. In that instance, Equations 3.7 and 3.8 become:

$$P_y(\overline{p} + \overline{u}^2) = \int_0^{P_y} (\langle p \rangle_{z,t} + \langle u \rangle_{z,t}^2 + \langle u'^2 \rangle_{z,t}) \mathrm{d}y \,, \tag{3.11}$$

$$P_y \overline{u}\,\overline{v} = \int_0^{P_y} (\langle u \rangle_{z,t} \langle v \rangle_{z,t} + \langle u'v' \rangle_{z,t}) \mathrm{d}y \,. \tag{3.12}$$

In this work, the mixed-out values are predominantly used. However, the type of normalisation is pointed out where a different approach is adopted.

3.2 Cascade flow features

Given the inflow and Reynolds conditions considered, the suction side boundary layer remains attached and laminar up to a rear separation point in the aft portion of the suction surface, as visible in the instantaneous contours of Figure 3.3. The Figure represents the mean mode (i.e. spanwise average) and thus shows higher coherence in the turbulent structures compared to a single $x-y$ slice; however, the large-scale vortex shedding remains coherent as the vortical structures propagate downstream, which raises the question of whether the numerical dissipation introduced is sufficient. In general, spectral methods have no numerical dissipation. In under-resolved simulations, energy accumulates in the high wavenumber region, therefore requiring the introduction of a subgrid-scale (SGS) model. In *Nektar++*, the use of SVV (a model-free stabilisation technique) introduces numerical dissipation, which needs to be sufficiently representative of physical dissipation. The combination of insufficient numerical dissipation and coarse resolution would result in an energy peak in the high wavenumber region of the spectrum. Velocity spectra are

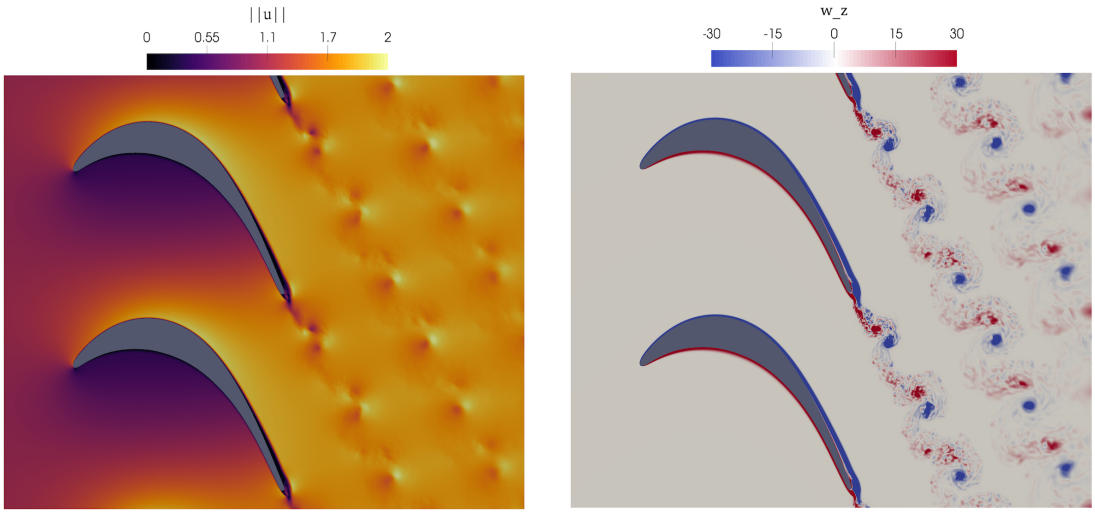

(a) (b)

Figure 3.3: Contours of the mean mode (mode zero in Fourier space) of instantaneous velocity magnitude (a) and spanwise vorticity in case P7L02N96.

discussed later in the Chapter, and such behaviour is not observed. However, in order to conclusively determine the potential benefit of a SGS model it would be necessary to estimate the numerical dissipation and compare it with physical dissipation. One such recent study by Dairay et al. [34] compared numerical dissipation via SVV and SGS modelling for LES. The authors argued that the relevant scale selectivity for the artificial dissipation introduced via SVV is even more crucial than the ability to represent inertial SGS contribution. In fact, these terms are comparable to aliasing and differentiation errors when the mesh size and filter size are comparable.

The underlying reason for large-scale structures in the LPT wake is deemed to be of physical nature, owing to the large open separation merging into the wake. The elongated separation bubble in the aft portion of the suction side is revealed in the time-averaged statistics, shown in Figure 3.4. This highlights how at the moderate reference Reynolds number of $Re_2 = 88450$, the separation indeed remains open merging into the trailing edge wake, therefore resulting in a highly oscillatory wake of alternating spanwise vortices. Conversely, at the same Reynolds the flow is attached on the whole pressure side and remains laminar until the trailing edge.

The skin friction coefficient is defined as:

$$C_f = \frac{\tau_w}{\frac{1}{2}\rho U_\infty^2}, \tag{3.13}$$

where τ_w indicates wall-shear stress, introduced in Equation 3.1. The time-averaged 3D C_f map of Figure 3.5 shows the location of flow separation. In the proximity of the trailing edge, there also exists a small region of mean positive shear, denoting the presence of a minor bubble circulating in the direction opposite to the main separation bubble, in a time-averaged sense. In the region very close to the trailing edge, the turbulence kinetic

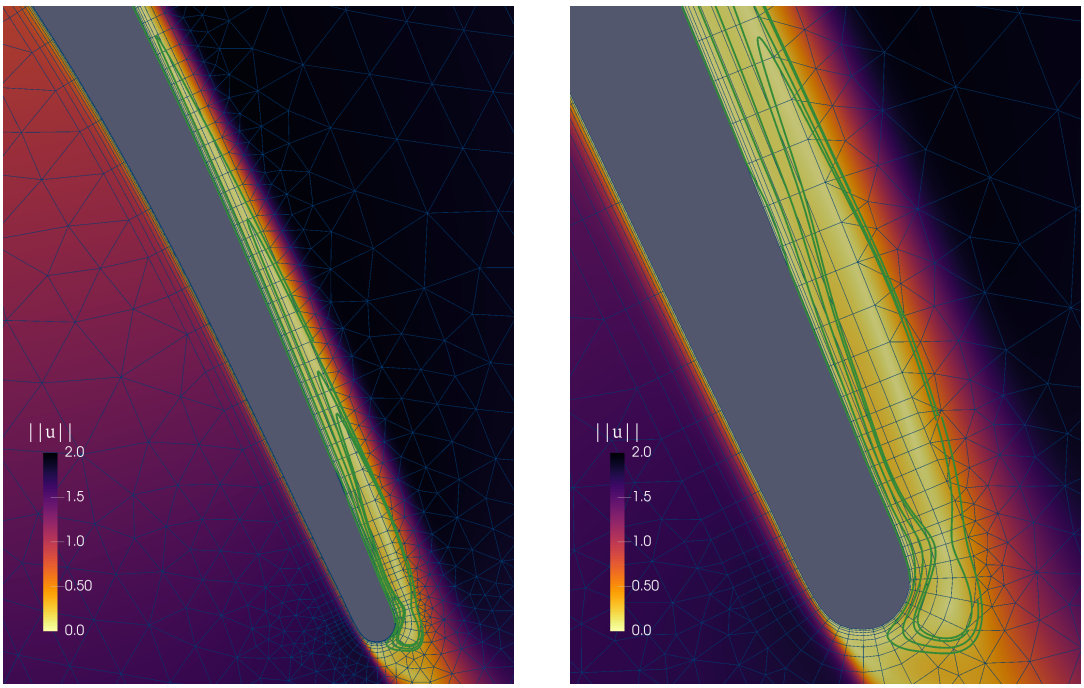

(a) (b)

Figure 3.4: Contours of the mean mode of instantaneous velocity magnitude superimposed with the base mesh (in dark blue) and flow streamlines (dark green), showing the extent and shape of the separation bubble in case P7L02N96 (a); the zoomed region near the TE region is shown in (b). The velocity magnitude colourmap is reversed with respect to Figure 3.3 for graphical clarity.

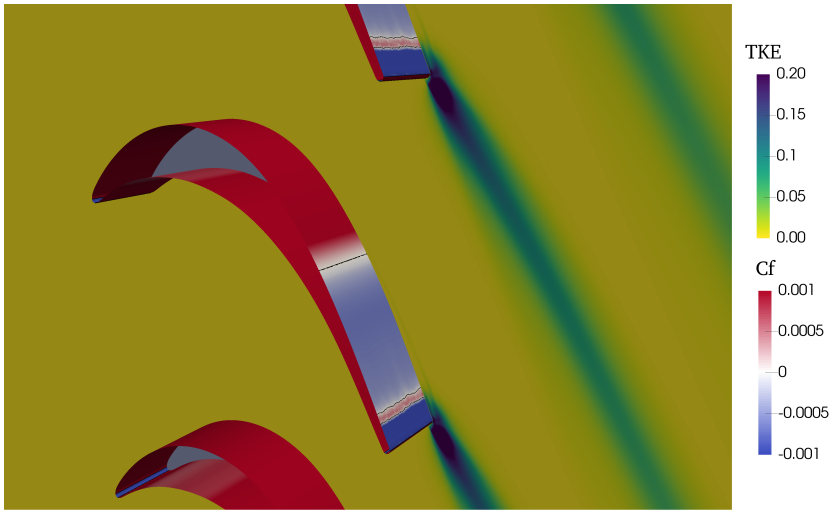

Figure 3.5: Time-averaged skin friction coefficient map on top of turbulent kinetic energy contours ($T = 24C/U_\infty$, case P7L02N128). The black lines indicate isolines of zero wall shear stress.

energy increases abruptly and the flow undergoes turbulent transition, consistently with the darker blue contours in the skin friction map showing the beginning of a reattachment process. Increasing the Reynolds number moves the separation forward affecting its length,

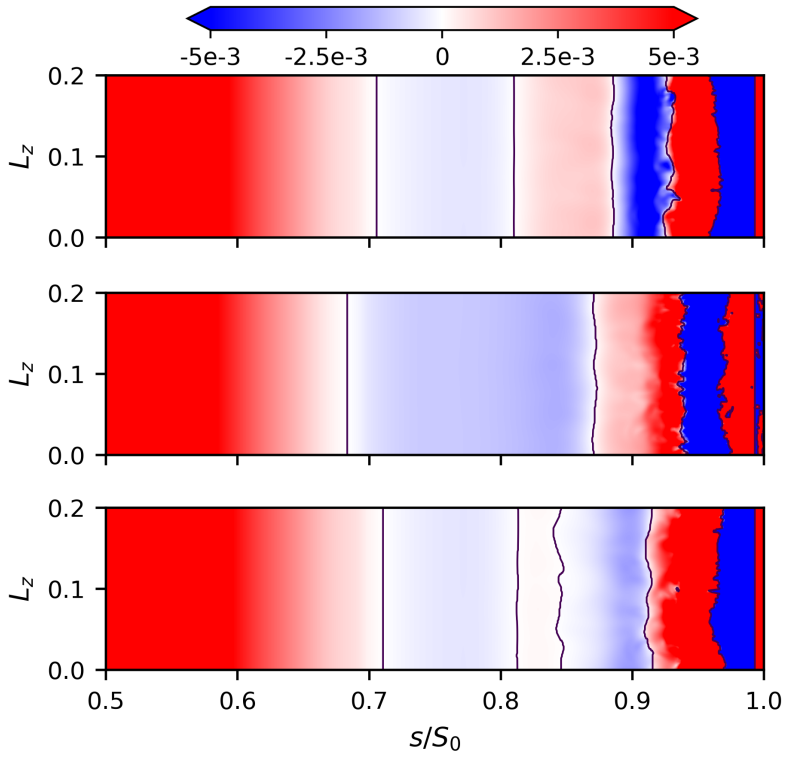

Figure 3.6: Instantaneous skin friction coefficient map on the aft portion of the suction surface in consecutive instants spaced by $T = 0.08C/U_\infty$, from top to bottom. The black lines are isolines of zero wall shear stress. Copyright © 2018 by Rolls-Royce plc.

the losses and the flow dynamics of the boundary layer approaching the trailing edge.

The laminar inflow promotes large unsteadiness of the flow structures, which appear to be very sensitive to the numerical setup in the range of Reynolds numbers considered. The instantaneous spanwise-independent separation line shown in Figure 3.6 has strong streamwise oscillatory behavior, and it is subject to movements in the range of over 4.5% of its average location. The saturated colour map allows to better isolate the reversed flow region and understand its inherent unsteadiness. Downstream of the separation point, emerging spanwise waviness accompanies highly unsteady streamwise-alternating regions of attached and separated flow convecting towards the trailing edge. The high sensitivity of such flow environment requires long and careful averaging procedures. This is unlikely to be an issue in a more realistic turbomachinery environment, where the presence of background turbulence modifies the transition mechanism, fixing the flow properties of interest.

The Q criterion [71] visualisation of Figure 3.7 provides further qualitative confirmation of the instantaneous coherent structures developed in the flow. Transition takes place only in the final part of the suction surface, where the separated shear layer rolls up due to Kevin-Helmoltz instability.

In this work, all profile results presented refer to the curvilinear system s/S_0, which is

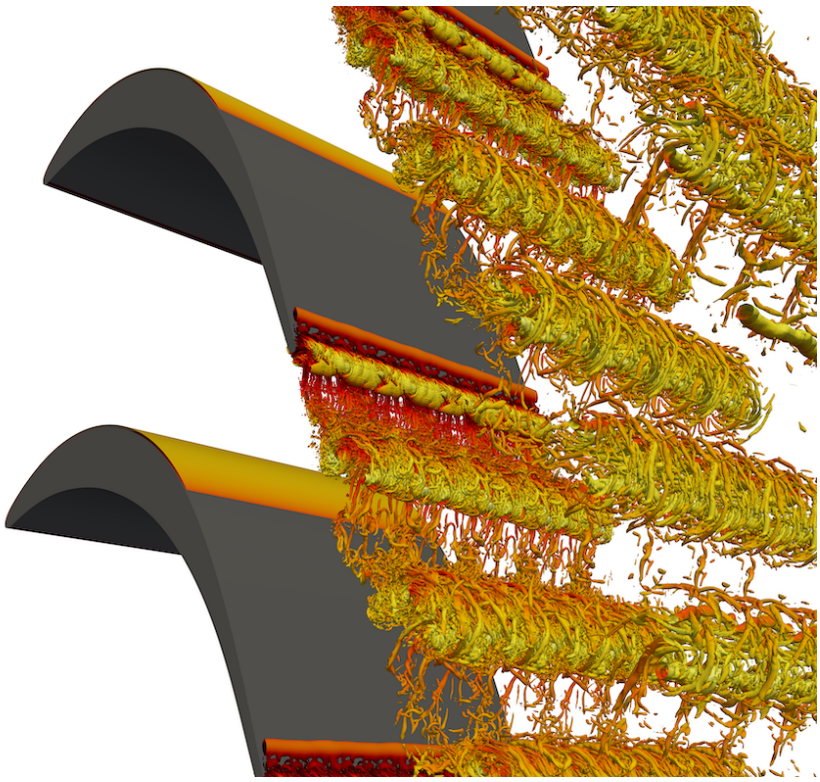

Figure 3.7: Instantaneous iso-surfaces of second invariant of the velocity-gradient tensor ($Q = 500$) coloured by velocity magnitude, case P7L02N96. The computational domain is replicated in the spanwise and pitchwise directions for graphical purposes.

preferred to the streamwise coordinate x/C_{ax} often used in the literature. The choice aims to better show variations on the aft portion of the suction surface. The leading edge and trailing edge are defined as the first point of contact with the blade of a line perpendicular to the design inflow and outflow angles, respectively.

3.3 P-refinement

This section provides a reference for the rest of the thesis, helping ensure mesh independence to obtain the most accurate flow statistics at reduced computational cost. The effect of polynomial order refinement is assessed for a given spanwise number of Fourier planes, progressively increasing the resolution in the blade-to-blade plane. This operation yields a significant increase in computational cost (as demonstrated in Table 3.1), which needs to be accounted for in analysing its tradeoff against the increasing accuracy gained owing to higher order expansions.

3.3.1 Blade wall distributions

A first quantitative estimation of the improvement of the performance prediction is visible by the analysis of the average pressure coefficient and skin friction coefficients. The

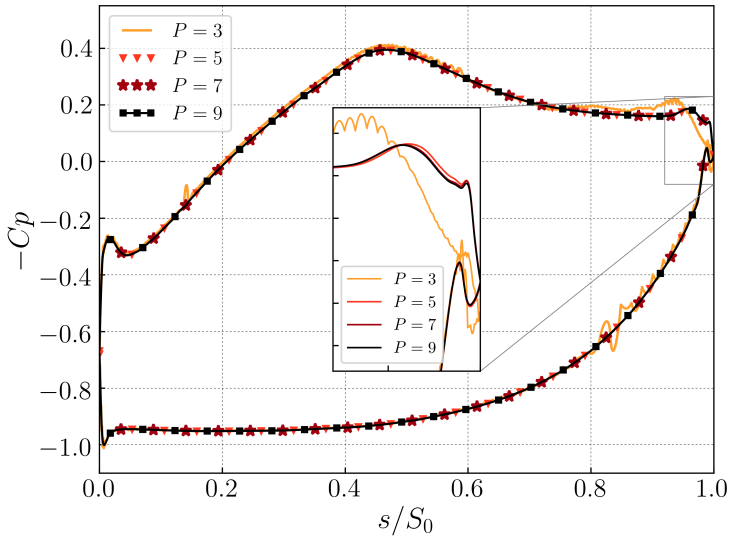

Figure 3.8: pressure coefficient with increasing polynomial order P. Copyright © 2018 by Rolls-Royce plc.

pressure distribution is calculated as [109]:

$$C_p = \frac{p_s - p^A_{s,2}}{p^A_{t,1} - p^A_{s,2}} \,. \tag{3.14}$$

This definition is adopted through this work, unless specified otherwise. The line integration employed to extract reference quantities at the inflow and exit measurement planes guarantees accuracy by using a very high number of equispaced interpolation points in the pitchwise direction. Figure 3.8 suggests that convergence is achieved at a relatively low order, and the distributions predicted with $P = 5$, $P = 7$ and $P = 9$ match closely along the whole blade. As expected, the most unstable region is the rear suction surface close to the TE, where discrepancies between cases $P = 5$ and $P = 9$ prove to be negligible as highlighted by the inset plot.

Similar observations apply to the skin-friction coefficient, shown in Figure 3.9. Analogously to the pressure distribution, the case $P = 3$ oscillates due to the under-resolution in the near-wall region and this is particularly noticeable in the separated region as well as in the aft portion of the pressure side. An order-of-magnitude improvement in numerical accuracy (attained by refining to $P = 5$) is sufficient to remove the under-resolution instabilities and a much better agreement is achieved with the highest order case.

The C_f distribution is flat in the region between $(s/S_0) = 065$ and $(s/S_0) = 0.9$, so the separation point ($(s/S_0)_{\text{sep}} = 0.6777$ with $P = 9$) is expected to be affected by a large uncertainty in the prediction. A robust location of the separation point $(s/S_0)_{\text{sep}}$ is computed through a spline passing through 4 points on each side of abscissa where the shear stress reaches zero and after having interpolated the distributions to $P = 15$.

The errors in Table 3.2 demonstrate the p-refinement rate of convergence to the asymp-

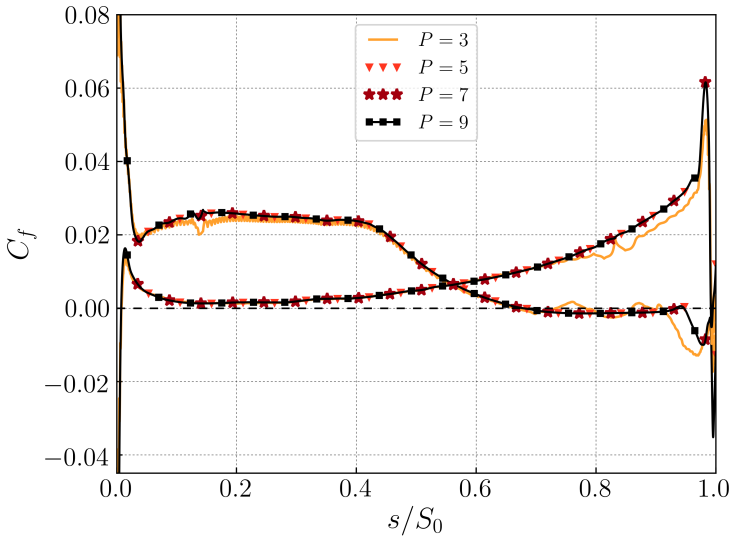

Figure 3.9: Skin friction coefficient distribution with increasing polynomial order P. Copyright © 2018 by Rolls-Royce plc.

totic solution. The error is calculated by interpolating all cases to a high-order distribution ($P = 15$), so that a point-wise correspondence is achieved allowing the correct calculation of the RMS relative error. The comparison of p-refinement results confirms the theoretical expectation that the separation location is captured with an order-of-magnitude accuracy improvement as P is increased by 2.

In this case, the solution with $P = 9$ is taken as reference as no exact solution is known. It is very likely that the error of solution $P = 9$ with respect to $P = 11$ would be smaller than the error of $P = 7$ with respect to $P = 9$. Therefore, the error of case $P = 7$ can be interpreted as a conservative uncertainty bound for the coarser cases.

Table 3.2: RMS of the relative error for various statistics compared to the reference case $P = 9$ (P9L02N96). The separation point error is the relative error.

Property	$P = 3$	$P = 5$	$P = 7$
C_p	0.03668	0.002618	0.000939
C_f	0.1956	0.00797	0.00221
$(s/S_0)_{\text{sep}}$	0.0221	0.00400	0.000512
θ	0.2158	0.01314	0.003608
H	0.1527	0.01185	0.003049

3.3.2 Boundary layer parameters

The evolution of the boundary layer parameters on the suction side of the turbine blade is an important parameter linked to profile loss estimation [38]. However, a robust and generally accurate estimation of the BL thickness [180] is not an obvious task especially for high fidelity time-varying velocity profiles in thin laminar, transitional or separated

flow regions, where no reference farfield velocity can be identified.

In this work, the method proposed by Uranga et al. [174] is adopted to provide an estimation of BL thickness. This approach is based on a pseudo-velocity which always asymptotes to a constant outside the boundary layer, hence reducing the arbitrariness in the definition of the boundary layer edge. Even though the tuning parameters proposed in the original paper were found to be slightly conservative when applied to a flat-plate configuration, they were adopted in this work. More specifically, the pseudo-velocity is defined as:

$$\overline{\mathbf{u}}^*(s,n) \equiv \int_0^n \overline{\boldsymbol{\omega}} \times \hat{\mathbf{n}}\, \mathrm{d}n\,, \tag{3.15}$$

where $\hat{\mathbf{n}}$ is the local vector perpendicular to the surface, and $\overline{\boldsymbol{\omega}}$ is the flow vorticity. The boundary layer edge n_e is defined as the first location where the magnitude of vorticity and vorticity variation in the direction normal to the surface simultaneously satisfy the conditions:

$$|\overline{\boldsymbol{\omega}}|\, n < \epsilon_1 \left|\overline{\mathbf{u}}^*\right|\,, \qquad \left|\frac{\partial \overline{\boldsymbol{\omega}}}{\partial n}\right| n^2 < \epsilon_2 \left|\overline{\mathbf{u}}^*\right|\,, \tag{3.16}$$

with $\epsilon_1 = 0.01$ and $\epsilon_2 = 0.1$. The edge velocity magnitude u_e at a distance of n_e from the blade is used for the calculation of boundary layer parameters. Defining $u_{\parallel}(n)$ as the velocity in the direction parallel to the blade at a certain BL location, the boundary layer displacement thickness and momentum thickness are calculated as:

$$\delta = \int_0^{n_e} \left(1 - \frac{u_{\parallel}(n)}{u_e}\right) \mathrm{d}n\,, \tag{3.17}$$

$$\theta = \int_0^{n_e} \left(1 - \frac{u_{\parallel}(n)}{u_e}\right) \frac{u_{\parallel}(n)}{u_e} \mathrm{d}n\,. \tag{3.18}$$

The shape factor is finally obtained as:

$$H = \frac{\delta}{\theta}\,. \tag{3.19}$$

The BL analysis was carried out by isolating the suction surface and distributing a large number of sampling stations ($\mathcal{O}(10^3)$) with increasing clustering in the aft portion of the blade, near the trailing edge. At each station, the boundary layer profile was interpolated to 500 points distributed along the normal to the surface following Robert's distribution [3] (specifying a coefficient $\delta = 0.02$), and up to a wall-normal extent of $n = 0.05C$. The staggered wall-normal interpolated grid allows to capture the viscous sublayer with greater accuracy than a uniform distribution. The number of points employed for the profile extraction and the stagger coefficient were decided by exploring a range of different combinations. The most refined case (P9L02N96) was taken as reference for estimating the boundary layer thickness in each station, which is then employed for calculating the integral parameters in all other cases. To further verify independence of the integration height from the parametric setup, the BL thickness was estimated independently in the case $P = 7$ and cross checked against $P = 9$. The negligible difference confirmed the

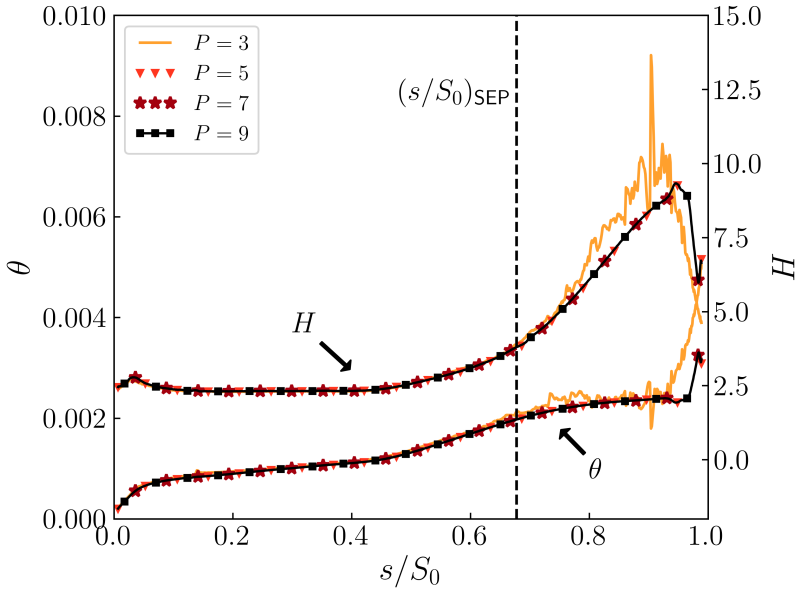

Figure 3.10: Momentum thickness θ and shape factor H along the suction surface with increasing P. The vertical dashed line indicates the separation point. Copyright © 2018 by Rolls-Royce plc.

robustness of the method.

Figure 3.10 shows the evolution of the momentum thickness and shape factor along the suction surface. In agreement with the blade wall distributions, the case $P = 3$ is strongly under-resolved and introduces large errors in the integral parameters. Solutions associated with higher order are almost entirely overlapped, with both momentum thickness and shape factor showing, again, the target order-of-magnitude improvement in the RMS relative error as P is increased by 2 (Table 3.2).

Further understanding of the accuracy in representing the boundary layer evolution along the suction surface can be gained by analysis of the relative error distribution of each case against the reference case P9L02N96. The momentum thickness is shown in Figure 3.11, and the shape factor is shown in Figure 3.12. The visualization in log scale helps to single out the increasing accuracy attained by increasing P. An order-of-magnitude decrease in error is found in the laminar portion of the suction surface, and in the trailing edge region the error tends to increase relatively to the front portion of the suction side. The maximum relative error in the TE region drops from around 60% with $P = 3$ to 1% with $P = 7$, which can also be intepreted as a conservative delta of uncertainty in the whole error estimation for cases $P = 3$ and $P = 5$. The presence of highly oscillatory behaviour of the error distribution is likely due to the boundary layer interpolation procedure. This is especially necessary due to the unstructured nature of the mesh outside the boundary layer mesh, and it introduces a component of numerical error.

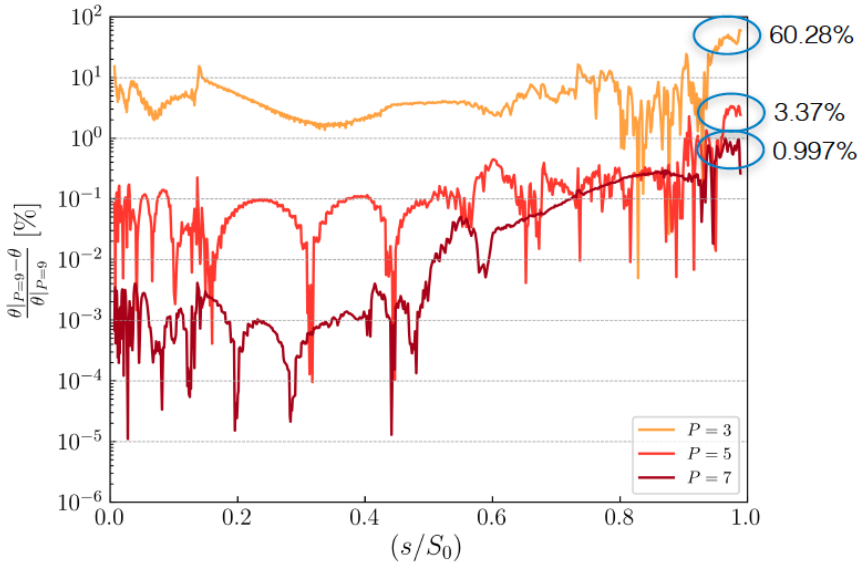

Figure 3.11: Relative error of the momentum thickness θ against $P = 9$ with increasing polynomial order. The circled areas represent the regions of maximum error in each case.

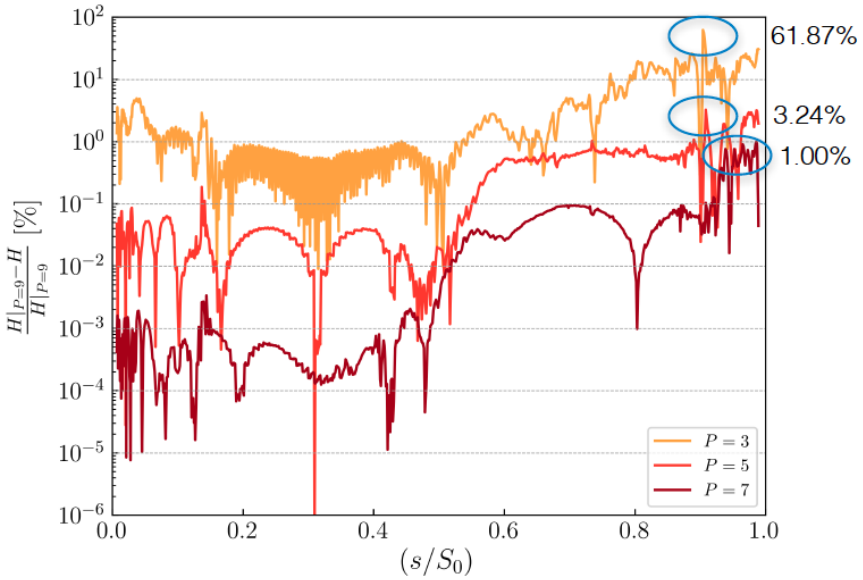

Figure 3.12: Relative error of the shape factor H against $P = 9$ with increasing polynomial order. The circled areas represent the regions of maximum error in each case.

3.3.3 Velocity spectra

Power spectra of streamwise, vertical velocity and velocity magnitude are analysed by sampling the flow variables in 11 points scattered in the trailing edge and wake region (Figure 3.13a), with a sampling rate of $fC/U_\infty = 10000$. For simplicity, one spectrum

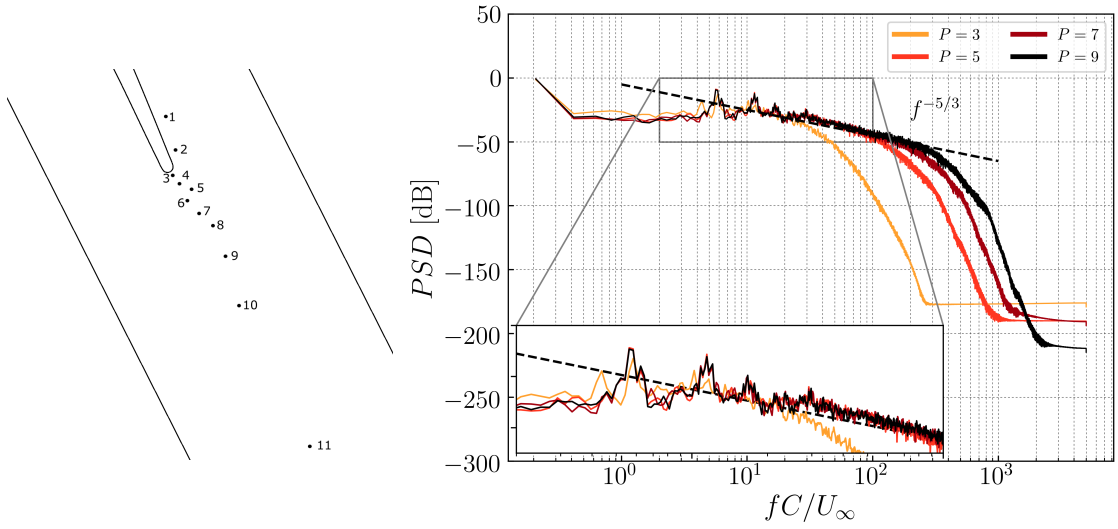

(a) (b)

Figure 3.13: (a): Probe points spatial distribution in the trailing edge and wake regions. (b): Streamwise velocity spectra in the turbulent wake with increasing P, in point 9. Copyright © 2018 by Rolls-Royce plc.

representing the dominant trends is presented in each case. To obtain a smoother spectrum, the PSD of the velocity signals are processed by partitioning the complete time history in 5 sections, on top of which Hann windows are applied with 50% overlap. At each (x, y) location, multiple samples are recorded in the z-direction, and the resulting spectra are arithmetically averaged.

Figure 3.13b shows progressive matching with the most refined case as P is increased, particularly at high frequency. The main peaks at low frequency are captured accurately with $P = 5$, but the spectrum differs significantly for fC/U_∞, where instead $P = 7$ is closer to the most refined case, up to roughly $fC/U_\infty \approx 300 - 400$.

When the $x - y$ plane resolution is very high compared to the spanwise resolution, the solution may diverge due to instability of the high frequency modes. This issue is further addressed in Section 3.3.5, showing that changes in the SVV cutoff can overcome this issue, and only affect the spectra in the high frequency end of the spectrum. Case $P = 9$ requires such further stabilization. It should be noted that, due to the finite sampling period, the phenomenon known as "leakage" in the spectra cannot be avoided because the signal is not exactly periodic in the time frame considered for sampling. The effect of leakage is that of smearing energy across the frequency spectrum. For instance, in this case it occurs at $fC/U_\infty > 3000$ in the case $P = 9$ and $fC/U_\infty > 2000$ for $P = 7$. Two measures can be undertaken to reduce this phenomenon: (1) extending the sampling period, or (2) selecting a good window function, whose Fourier transform is closest to a delta function. Since the first measure is impractical, several functions were tested; the Hann window, commonly adopted owing to its good spectral properties, was found to produce the best attenuation of the phenomenon.

3.3.4 Spatially adaptive polynomial order

In spectral/hp element methods, h-type refinement is used in conjunction with p-type refinement, thereby leveraging the geometrical flexibility of finite element methods and the superior spatial accuracy properties of spectral methods. P-refinement is a powerful tool to increase accuracy in representing turbulent phenomena with lower total degrees of freedom (DoF) compared to low-order FD and FV methods (thus reducing diffusion and dispersion error). For applications that lead to sufficiently smooth solutions (i.e. in the absence of shocks), the rate of convergence of Galerkin methods is exponential, highlighting the algorithmic advantage of p-refinement [93].

A typical approach to hp-adaptive techniques is based on iterating a local refinement procedure based on *a posteriori* error estimation. The main refinement indicators are generally grouped in three categories: feature based indicators, Discretisation-Error based indicators and goal oriented indicators. More details on the differences, advantages and drawbacks can be found in the recent paper by Naddei et al. [121].

Spatially variable (and adaptive) polynomial order was implemented in the *Nektar++* framework [155, 119]. The compressible solver features a goal-based error estimator based on the adjoint solution [44], while in the incompressible Navier-Stokes solver a DE-based local sensor as error estimator [131] is available:

$$S_e = \frac{||u_e^p - u_e^{p-1}||_{L_2}}{||u_e^p||_{L_2}}\,, \tag{3.20}$$

where the sensor variable can be specified; the error estimator is evaluated separately within each element. Based on specified upper and lower error tolerances, the local polynomial order is increased or decreased by one order at each adaptation step. The refinement operation is expensive (typically 200 times a single time step) and therefore the ratio between number of computing iterations and steps of polynomial adaptation must be balanced. Very stringent tolerances on the error may also introduce severe CFL limitations for local stability. If optimally tuned, the adaptation procedure allows for a significant reduction in computational cost for a given accuracy (See Table 3.1 for evidence of the reduced cost against the rest of the computations performed). In the current work, upper and lower tolerances of 10^{-8} and 10^{-10} were adopted, respectively. The expansion order was allowed to range between $P = 3$ and $P = 11$, and it was restarted every $T = 0.2C/U_\infty$ for at least 50 times. Both streamwise velocity and the pressure field were separately tested as sensor variables, and the resulting elemental distribution of P from the former is shown in Figure 3.14.

As a result of the refinement algorithm, the polynomial order at the inflow is the lowest allowed; this introduces significant computational savings in a relatively uniform region of the flow domain. High-order refinement is triggered in the region adjacent to the TE, and progressively diminishes with the downstream distance. In the outflow region, past $0.75C_{ax}$ downstream of the TE, the base mesh becomes coarse, and the algorithm attempts

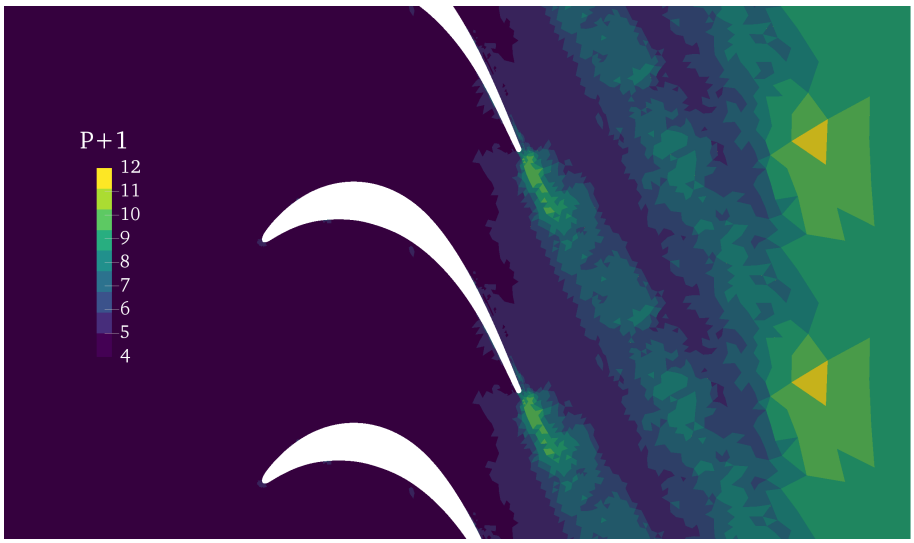

Figure 3.14: P distribution with error estimation based on streamwise velocity. Copyright © 2018 by Rolls-Royce plc.

to retain high resolution of the turbulent structures by increasing the order again, which justifies the high P distribution near the outflow plane. As a result, the simulation time in this case (indicated with VARP in Table 3.1) is roughly a third of case P7L02N96. At constant tolerance bounds, the use of pressure as sensor variable is more constraining: the outflow regions where the pressure is close to zero require very low errors, therefore suggesting a less stringent tolerance would be more appropriate in this case.

It should be mentioned that this approach has limitations concerning its adoption in cases with incoming disturbances. In fact, with clean inflow the relatively frequent error estimation and consequent change of order settles quickly on the spatial distribution of P shown in Figure 3.14. In a case with incoming wakes, instead, this type of adaptation procedure would attempt to place more DoF at the passing bar's location at the time of mesh refinement, without accounting for its vertical motion. For this purpose, a more advanced approach would be required, to account of the inherent deterministic flow unsteadiness. Therefore, the results presented in the rest of the thesis do not rely on adaptive polynomial order, which is left to future research efforts.

3.3.5 SVV cutoff for the exponential kernel

Tuning of the SVV parameters is required for the exponential kernel adopted in the spanwise direction. The DG-kernel in the $x-y$ plane does not entail manual parameters setting, as discussed in Section 2.2.2. From the time evolution of the modal energy, it was observed that when the in-plane resolution is high relatively to the spanwise resolution, energy tends to build up in the high-frequency modes and cause the solution to diverge. This issue can be overcome in two ways. One approach is to rebalance the resolution accuracy by increasing the number of Fourier planes. The alternative solution (adopted here for the sake of comparison and understanding) is that of applying numerical diffusion

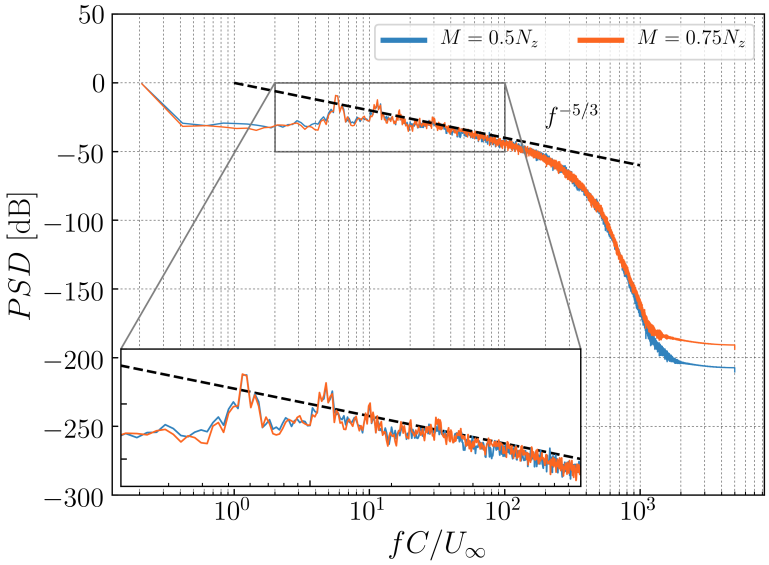

Figure 3.15: Streamwise velocity spectra in the turbulent wake with different SVV cutoff coefficients in the spanwise direction, in point 9 (Figure 3.13a). Copyright © 2018 by Rolls-Royce plc.

on a wider range of high frequency modes, by lowering the SVV cutoff in the homogeneous direction. This expedient was implemented in a series of cases in this study. To understand to what extent the statistics of interest are affected by this necessary change, two cases are compared: $M/N_z = 0.75$ (P7L02N96) and $M/N_z = 0.5$ (P7SVV). Pressure and skin friction coefficient distributions are identical, and therefore omitted.

The velocity spectra (Figure 3.15 for an example) manifest a very similar behavior up to about $fC/U_\infty \approx 1000$ in all the 11 probe points. Due to mixed contributions of leakage and SVV effects at high frequency, it is difficult to identify a trend caused by changes in the cutoff coefficient alone; clearly, it contributes to a variation of the energy contained in the small scales. However, the statistics depending mainly on low frequency mechanisms (pressure distribution, skin friction coefficient) are not affected by slightly decreasing the SVV cutoff parameter for the homogeneous direction. Therefore, when analysing statistics dominated by large scale structures, it was found acceptable to establish comparisons between computational setups where M/N_z was slightly modified to ensure robust numerical stabilization. Such comparisons are to be considered with caution only in the range of frequencies higher than approximately $fC/U_\infty \approx 1000$ for the current $x - y$ plane resolution; it is worth reminding that the higher the in-plane resolution, the higher the frequencies where differences are detected.

3.3.6 Time averaging window

When the simulation has gone through the initial transient period, convergence is achieved in a subsequent number of axial through-flows that depends on the individual property of interest. In general, mean flow quantities require shorter time averaging than turbu-

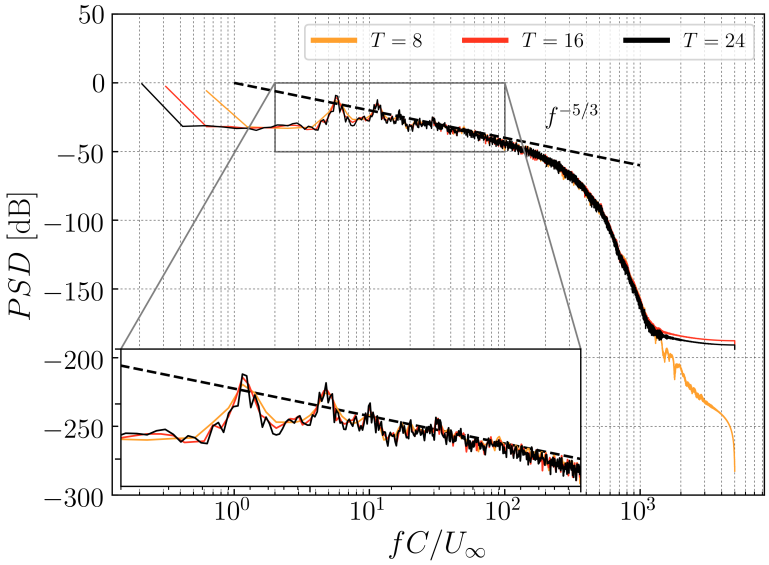

Figure 3.16: Streamwise velocity spectra in the turbulent wake with various sampling intervals T, in point 9 (Figure 3.13a). Copyright © 2018 by Rolls-Royce plc.

lence quantities, with the most stringent requirements associated to turbulence dissipation [135]. In the present study, blade distributions were found to converge quickly (as consistently observed across the literature), so the velocity spectra were analysed to identify the minimum averaging frame to be employed to ensure convergence of all types of statistics within the averaging window. Figure 3.16 shows that the mean peaks at low frequencies are captured accurately with relatively short time averaging (8 chord-based lengths), while even the high frequency end of the spectrum becomes almost independent of the averaging time for $T > 16C/U_\infty$. In the current study, $T = 24C/U_\infty$ was employed to guarantee a conservative averaging for comparisons. It should be noted that this choice is not necessarily the most cost-effective.

Time convergence measured by velocity spectra provides some level of guidance, but is not directly linked to flow quantities of practical interest for designers. In performance assessment, the estimation of losses plays a dominant role. The mixed-out loss coefficient is defined as [109]:

$$\omega^M = \frac{p_{t,1}^M - p_{t,2}^M}{p_{t,1}^M - p_2^M} \,. \tag{3.21}$$

The time convergence of ω^M with increasing polynomial order is reported in Figure 3.17a. Achieving convergence requires a much longer time frame than blade wall distributions. Comparing the mixed-out loss values measured for $T \geq 16C/U_\infty$ and the values measured at $T = 24C/U_\infty$, the maximum relative errors for cases $P = 3, 5, 7, 9$ are respectively:

$$\frac{|\omega^M - \omega^M(T = 24C/U_\infty)|}{\omega^M(T = 24C/U_\infty)} \cdot 100 = 1.594, 1.057, 2.093, 1.196 \,. \tag{3.22}$$

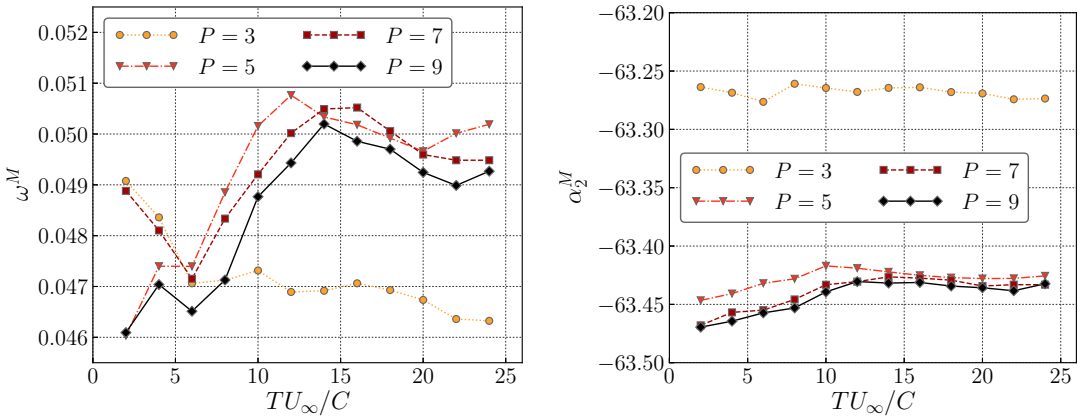

(a) (b)

Figure 3.17: Mixed-out total pressure loss coefficient (a) and exit angle(b) calculated over accumulating time averages, comparing the effect of p-refinement.

Table 3.3: Mixed-out total pressure loss and exit angle with increasing order P, and relative error with respect to $P = 9$. Results averaged over $T = 24C/U_\infty$.

Property	$P = 3$	$P = 5$	$P = 7$	$P = 9$
ω^M	0.04632	0.05019	0.04948	0.04927
ω^M_{err} [%]	5.976	1.873	0.4386	-
α_2^M [°]	-63.2736	-63.4255	-63.4331	-63.4324
$\alpha^M_{2,err}$ [%]	0.2503	0.01068	0.001209	-

This result provides an approximate uncertainty bound for the estimation of total pressure losses and highlights their sensitivity. The large variability still measured after a large number of convective time scales is attributed to the specific flow configuration analysed, where the open separation merging into the wake gives rise to large-scale coherent structures which are convected towards the outflow. Therefore, in terms of time convergence (in the absence of discrete periodic disturbances, i.e. wake passing) the case analysed provides the most challenging flow conditions. Analogous results for the mixed-out exit angle are shown in Figure 3.17b. Differently from total pressure, the exit angle converges more quickly, visibly settling for $T \geq 12C/U_\infty$. in both cases, a sharp improvement in the agreement with $P = 9$ is achieved with orders higher than $P = 3$.

Considering the time-converged results at $T = 24C/U_\infty$, Table 3.3 shows the values and relative errors of the loss coefficient and exit angle. The relative error of ω^M with respect to $P = 9$ is reduced by $\sim 3 - 4$ times every time the order is increased by 2, while the exit angle achieves an order-of-magnitude improvement. The results obtained at $P = 7$ and $P = 9$ are well within experimental tolerances, suggesting that case $P = 7$ is fully converged in space when its sensitivity is analysed based on mixed-out quantities.

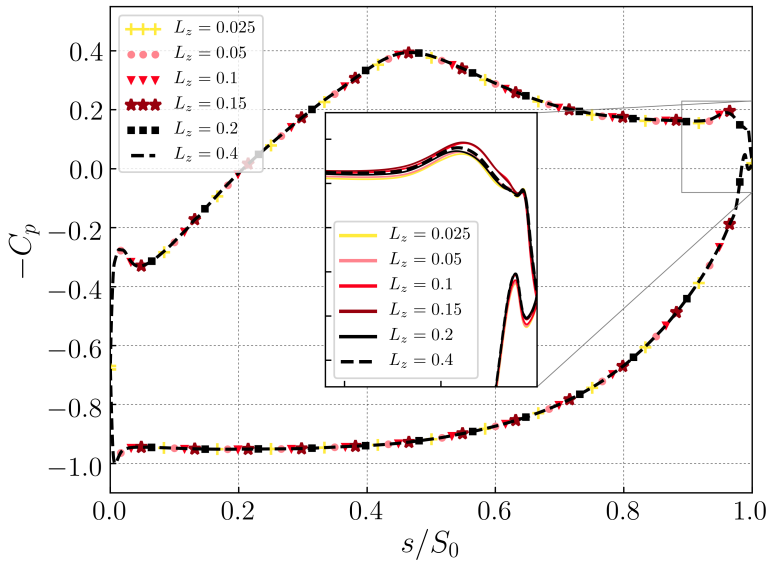

Figure 3.18: Pressure coefficient distribution with various spanwise domains L_z/C. Copyright © 2018 by Rolls-Royce plc.

Table 3.4: RMS of the relative error for various statistics compared to the case $L_z = 0.4C$.

Property	$L_z = 0.025C$	$L_z = 0.05C$	$L_z = 0.1C$	$L_z = 0.15C$	$L_z = 0.2C$
C_p	0.008151	0.005651	0.004925	0.003360	0.001578
C_f	0.01711	0.01406	0.02772	0.01618	0.008564
$(s/S_0)_{\text{sep}}$	0.01073	0.007986	0.01073	0.01106	0.003378

3.4 Spanwise domain effect

This section analyses the effect of statistical changes due to variations of the spanwise domain L_z, by maintaining the polynomial order at $P = 7$ and the ratio $N_z/(L_z/C) = 480$ constant. The pressure coefficient (Figure 3.18) shows marginal changes between the cases considered, as visible in the inset box. In the TE region, two cases are not on top of the reference case: $L_z = 0.1C$ and $L_z = 0.15C$, against expectations on convergence properties. Analogously, the skin friction coefficient distribution (Figure 3.19) is well captured in all the cases considered.

Table 3 shows that the RMS relative error with respect to the case $L_z = 0.4C$ converges slowly in the pressure distribution as the spanwise domain is extended, while no clear trend can be inferred from the skin-friction coefficient estimates. Raverdy et al. [142] argued that the BL thickness at the TE may indicate the characteristic length of turbulent structures, therefore suggesting if spanwise filtering of the structures is taking place due to the constraining size of the spanwise domain. In case P9L02N96, $n_{e,\text{TE}} \approx 0.04$; the mean flow properties are accurately captured even with a smaller spanwise domain extent, suggesting a way to preliminarily assess such properties at low computational cost.

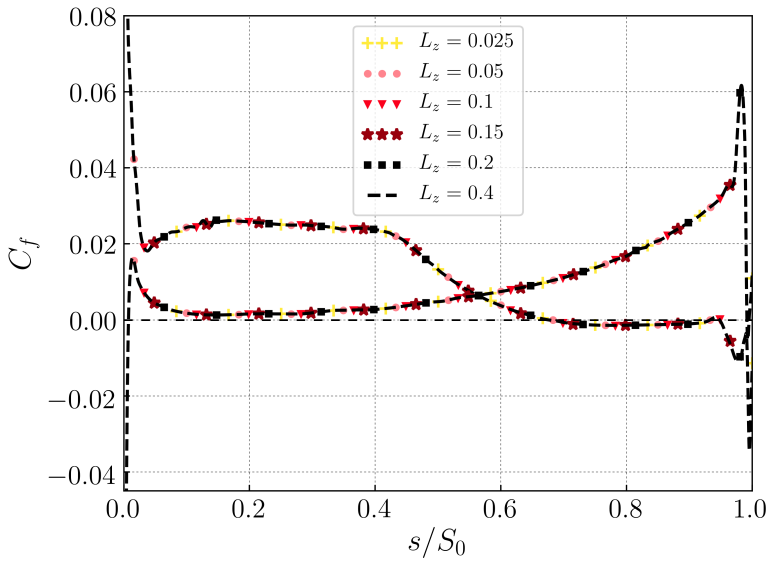

Figure 3.19: Skin friction coefficient distribution with various spanwise domains L_z/C.

The velocity spectra are reported in Figure 3.20. The main trend is consistent across the cases analysed, again suggesting that a very constraining domain is sufficient for an initial estimation of the flow physics, and good agreement with the $-5/3$ power law is verified. However, the main peaks at low frequency of cases $L_z = 0.1C$ and $L_z = 0.15C$ are slightly shifted by a mutually consistent value towards the low frequency end of the spectrum. This behavior is aligned with the discrepancies observed in the pressure distribution, and it is currently not fully understood.

At this Re, the separation bubble is open, therefore giving way to complex and highly unsteady dynamics in the TE region. Comparison of the two spanwise extents $L_z = 0.1C$ and $L_z = 0.2C$ at higher $Re_2 = 160000$ (presented in a later section) does not manifest this behavior. At such higher Re, however, the separation bubble is reattached near the TE, where the boundary layer is fully turbulent, and these TE features might be more robust to changes in the computational setup. Therefore, a possibility is that the complex dynamics of the open separation bubble is very sensitive to a constraining domain size, producing more involved behaviors in the series of setups analysed. It is also reminded that these conclusions are only valid for laminar inflow: the presence of background turbulence would introduce additional transition modes at various lengthscales, potentially requiring an even larger domain extent to contain the very large turbulent scales.

3.4.1 Two-point correlation

The analysis of two-dimensional, time-averaged performance indicators like the blade wall distributions are not sufficient in this case to provide conclusive evidence of whether the flow structures are fully contained in the computational domain, i.e. that a constraining

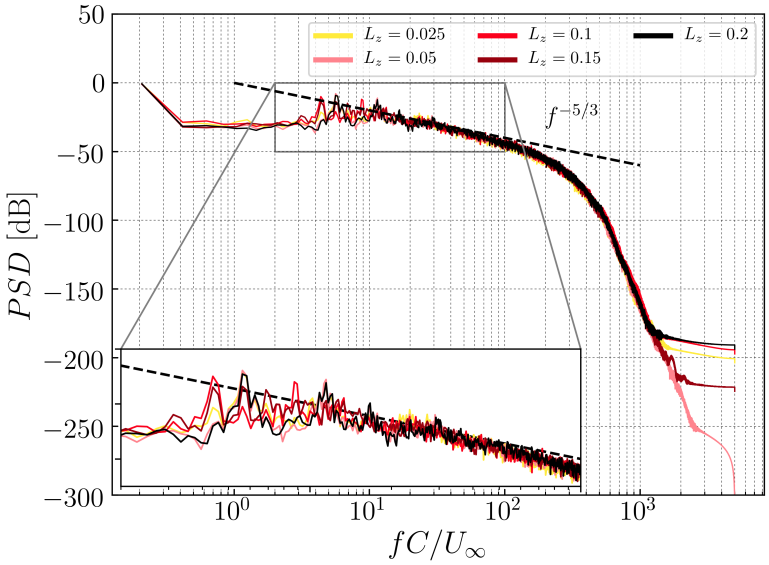

Figure 3.20: Streamwise velocity spectra in the turbulent wake with increasing L_z/C, in point 9 (Figure 3.13a). Copyright © 2018 by Rolls-Royce plc.

domain size is not producing artificial flow structures thus modifying the flow dynamics along the suction surface. The analysis of two-point correlations in the spanwise direction can provide further insight into this question by quantifying the level of correlation of turbulent structures in a certain spatial direction. Since the flow is homogeneous in the z-direction, the two-point correlation function only depends on the distance between two points Δz. Therefore, it is calculated by averaging the standard two-point correlation function obtained by rolling over all points in the spanwise direction:

$$R_{ij}^{\text{norm}}(\Delta z) = \frac{1}{N_z} \sum_{k=1}^{N_z} \frac{\langle u_i'(z_j) u_i'(z_j + \Delta z) \rangle_t}{\langle u_i'^2 \rangle_{z,t}}, \tag{3.23}$$

and further mirroring the resulting distribution in the spanwise direction, exploiting its periodicity.

The test case representative of the most widely adopted setup through the thesis is P7L02N96. An additional case was considered (P7L04N192), with doubled spanwise domain $L_z = 0.4C$, and retaining a constant number of DoF per spanwise length, therefore resulting in $N_z = 192$ and DoF $= 74M$. This is the largest case analysed in this work, in terms of total number of degrees of freedom. Samples are collected at three $x - y$ locations at every Fourier plane, at the same frequency as the probes employed to analyse energy spectra (every $t = 10^{-4} C/U_\infty$). In case P7L02N96, line samples are collected for $T = 20C/U_\infty$, while for P7L04N192 they are sampled for $T = 16C/U_\infty$ due to the higher computational cost. Time independence in this range of windows was verified in Section 3.3.6. The three probes analysed correspond to probes $2, 4, 9$ of Figure 3.13a. In particular, the first point is in the transition region of the separation bubble; the

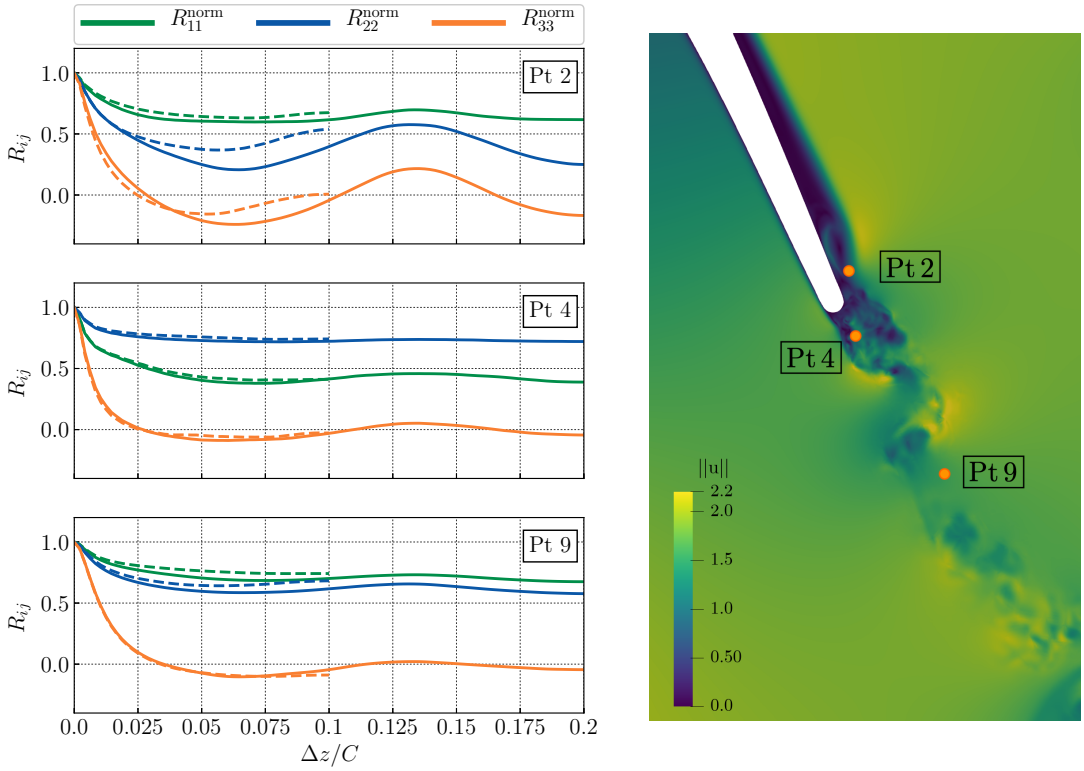

Figure 3.21: Left: Two-point correlation function for cases $L_z = 0.2C$ (dashed lines) and $L_z = 0.4C$ (solid lines). Right: contour plot of velocity magnitude showing the location of probe lines.

second probe location is very close to the trailing edge, while the third probe is further downstream in the turbulent wake, and it corresponds to the location where all velocity spectra are shown. In all three correlations, shown in Figure 3.21, the w component of velocity reaches zero within $\Delta z = 0.025C$, while the in-plane components u, v do not decay to zero even with a larger spanwise domain, for both domain sizes. This suggest the presence of coherent structures extending for the whole spanwise domain. Its motivation is mainly justified by the clean inflow boundary conditions: the lack of a physical meachanism providing noise to excite three-dimensional breakdown of the coherent structures results in Keving-Helmholtz roll-ups characterised by weak waviness in the spanwise direction, as previouly highlighted in Figure 3.6.

Probe line 2. The decay rate of the correlation curves with $L_z = 0.4C$ and $L_z = 0.2C$ is similar up to $\Delta z \approx 0.03C$, after which the smaller domain shows a higher value of R_{22}^{norm} and R_{33}^{norm}. This probe is a region of the flow where strong recirculation occurs, and the separation bubble is in the process of reattaching. The flow direction is predominantly aligned with the vertical and spanwise direction, which may motivate this correlation trend. The curves for $L_z = 0.4C$ show a peak at $\Delta z \approx 0.13C$, which may suggest the

presence of coherent structures more strongly correlated at this larger length scale. Since this spanwise distance is larger than half of the smaller domain size, this suggests that $L_z = 0.2C$ is not sufficient to fully contain all coherent structures in the TE region of the separation bubble. In order to fully verify this, it would be necessary to simulate a larger domain (i.e. $L_z = 0.8C$, with $N_z = 384$), but the high computational cost of this numerical experiment proved unfeasible for this study.

Probe lines 4 and 9. The decay rate is very accurately matched, especially at Pt 4. At this location, the curves for $L_z = 0.2C$ and $L_z = 0.4C$ are completely overlapped, while at Pt 9 there are small discrepancies for the u- and v- components. This suggests that the smaller domain $L_z = 0.2C$ is here sufficient to fully represent the dynamics of the wake.

In summary, it seems that the dynamics of the separation bubble contains low wavenumber structures with a spanwise extent that is larger than $L_z = 0.2C$. However, the introduction of background disturbances is likely to be sufficient to break down this dynamics; however, this statement remains a speculation and further analysis would be required to verify it in a disturbed environment. In order to confirm that $L_z = 0.4C$ is itself large enough, a wider domain should also be investigated. However, this would require significant computational resources that were not available for this purpose. In the wake, both in the TE region and further downstream, the two-point correlation curves are very well overlapped, suggesting that the smaller domain size is not limiting the three-dimensional development of turbulent structures.

3.5 Number of spanwise modes

Taking as reference parameters $P = 7$ and $L_z = 0.2C$, the effect of changes in the number of spanwise planes on spanwise- and time-averaged statistics is addressed. The pressure and skin friction coefficients appear relatively insensitive to changes in this parameter (Figures 3.22 and 3.23). The resolution per unit span employed is higher than that of several studies at this Reynolds number. From visual inspection of the blade distributions and the RMS error norms in Table 3.5 slow convergence is observed, with very low tolerances being already achieved on 2D statistics at low resolution.

The velocity spectra shown in Figure 3.24 present very similar trends in all the cases analysed. The low resolution cases accurately capture the main peaks at low frequencies. This reinforces what already outlined in the previous section: constraining spanwise domain and spanwise resolution do not have a major impact on the 2D flow statistics driven by low frequency dynamics. This means that a low resolution computational setup can be employed for testing purposes, considerably reducing the cost of performing the preliminary simulations.

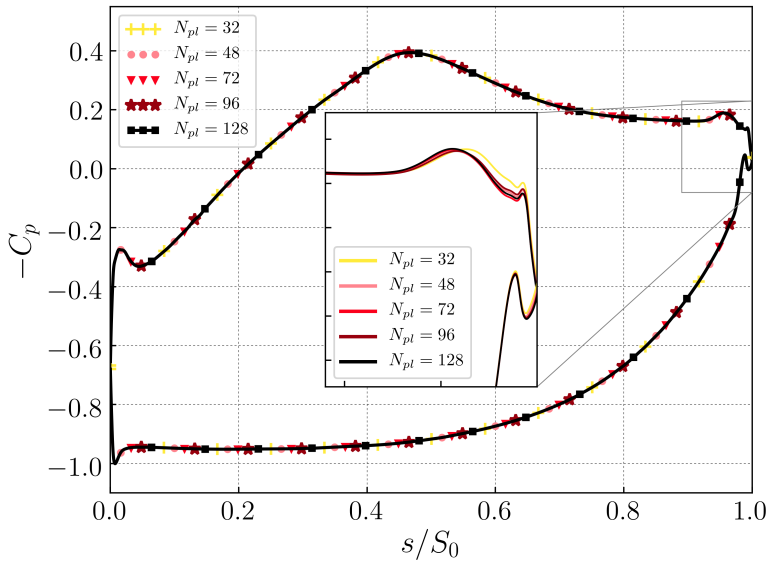

Figure 3.22: Pressure coefficient distribution with various spanwise Fourier planes N_z. Copyright © 2018 by Rolls-Royce plc.

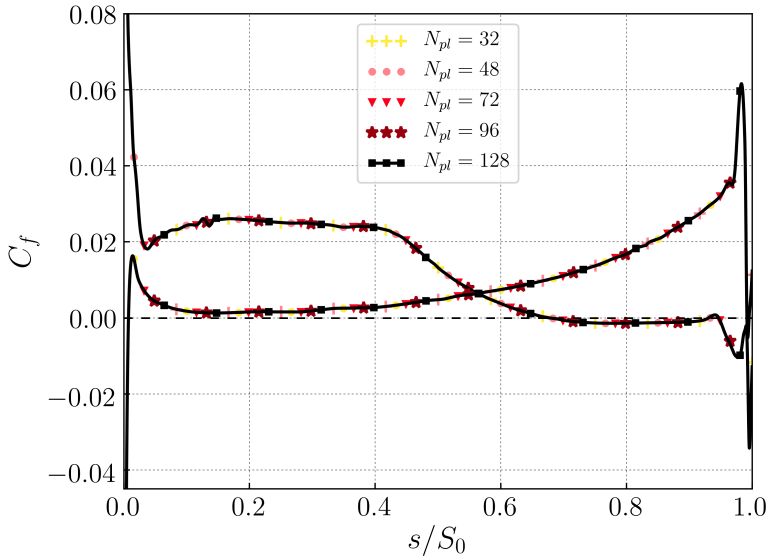

Figure 3.23: Skin friction coefficient distribution with various spanwise Fourier planes N_z. Copyright © 2018 by Rolls-Royce plc.

Table 3.5: RMS of the relative error for various statistics compared to the case P7L02N128 ($N_z = 128$).

Property	$N_z = 32$	$N_z = 48$	$N_z = 72$	$N_z = 96$
C_p	0.00629	0.00236	0.000990	0.00210
C_f	0.0173	0.00974	0.00694	0.00553
$(s/S_0)_{\text{sep}}$	0.00592	0.00219	0.00117	0.000952

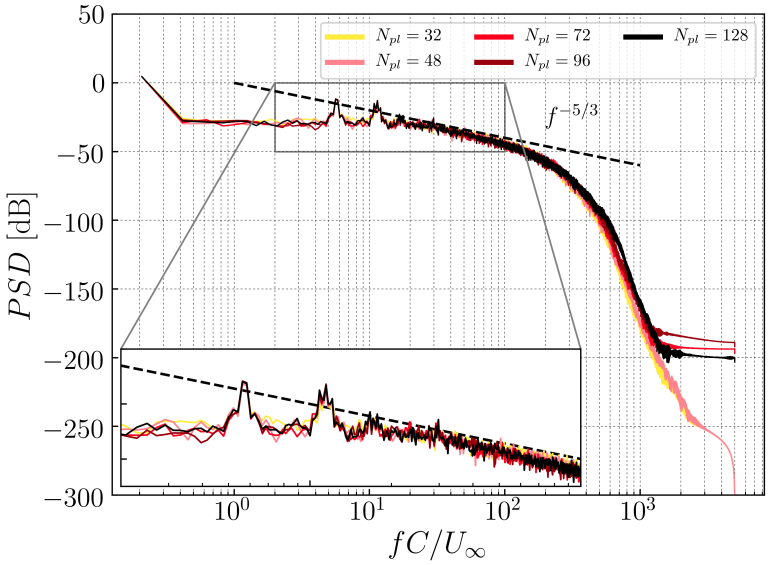

Figure 3.24: Vertical velocity spectra in the turbulent wake with various spanwise Fourier planes N_z, in point 9 (Figure 3.13a). Copyright © 2018 by Rolls-Royce plc.

3.6 Comparison with experiments

This section presents a comparison against an experimental campaign performed at Cambridge University [27]. The experiments were carried out at incompressible conditions. The pressure distribution was computed following the definition adopted in the report, and the downstream plane at $0.25C_{ax}$ was used to compute the mixed-out velocity reference to estimate the exit Reynolds number. Tests confirmed that very close results are obtained employing the area- and mass-averaged velocity. It is worth pointing out that, despite the care placed in these comparisons, some uncertainties still exist on how exactly the numerical data reduction is reproducing the experimental data regression and its uncertainties.

The set of boundary conditions declared in the experimental data was found to deliver an inaccurate matching of the blade loading distribution. A range of calculations were systematically used to identify a set of modified numerical conditions to better match the experimental loading, by varying the inflow angle and Reynolds number. It should be highlighted that the purpose of this section is that of verifying agreement with the experiments rather than an extensive and detailed investigation of the inflow angle and the Reynolds number (discussed in Chapter 5).

Table 3.6 summarizes the three computational setups that are compared with experiments. It is assumed that the in-plane resolution requirements discussed in the previous section can be extended to a moderately higher Reynolds number; the spanwise resolution was increased to maintain numerical stability. As seen from the summary, the matching of the incidence and front portion of the suction surface was obtained with an inflow angle approximately three degrees higher than the nominal value. Despite that, when COMP1

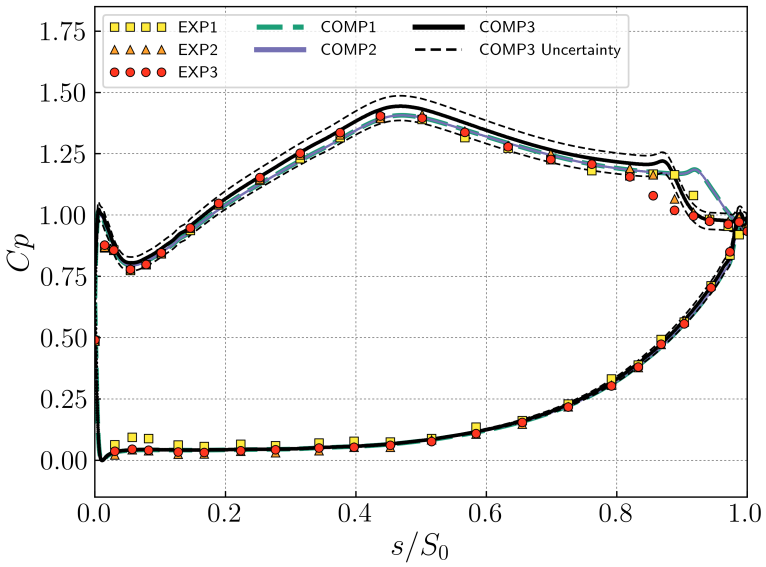

Figure 3.25: Comparison between experimental and numerical pressure distributions, with reference to $Re_2 = 160000$. Copyright © 2018 by Rolls-Royce plc.

Table 3.6: Summary of computational and experimental test cases compared.

Case	Re_1	Re_2	α_∞	P	L_z/C	N_z	M/N_z
COMP1	93000	158889	41	7	0.1	64	0.5
COMP2	93000	158882	41	7	0.2	128	0.5
COMP3	135000	228896	40.7	7	0.1	64	0.5
EXP1	-	159813	37.7	-	-	-	-
EXP2	-	213943	37.7	-	-	-	-
EXP3	-	258795	37.7	-	-	-	-

achieves the nominal exit Reynolds number of case EXP1, the rear separation is delayed and remains open until the TE (Figure 3.25). The remaining part of the profile compares with remarkable accuracy. The comparison between cases COMP1 and COMP2 is reported to uphold that (in view of the observations made above on spanwise accuracy) there are no appreciable effects due to the 3D confinement. Another qualitative arguments supports this hypothesis: the Q isosurfaces represented for the three flow regimes analysed (Figure 3.26) show that as Re is increased the spanwise waviness of the KH roll-ups disappears, and transition is anticipated yielding smaller scale turbulent structures, confirming that at the regime considered for comparison with experiments $L_z = 0.1C$ is most likely adequate to capture the blade flow properties.

Comparison of numerical and experimental pressure coefficient profiles is shown in Figure 3.25. Increasing Re_2 moves the separation onset predicted in the numerical simulations, following the same trend as in the experiments. At higher Re_2 the separation bubble reattaches and the simulation COMP3 reproduces the experimental distribution of EXP1. As for the inflow angle, the simulations seem to match experiments at a different

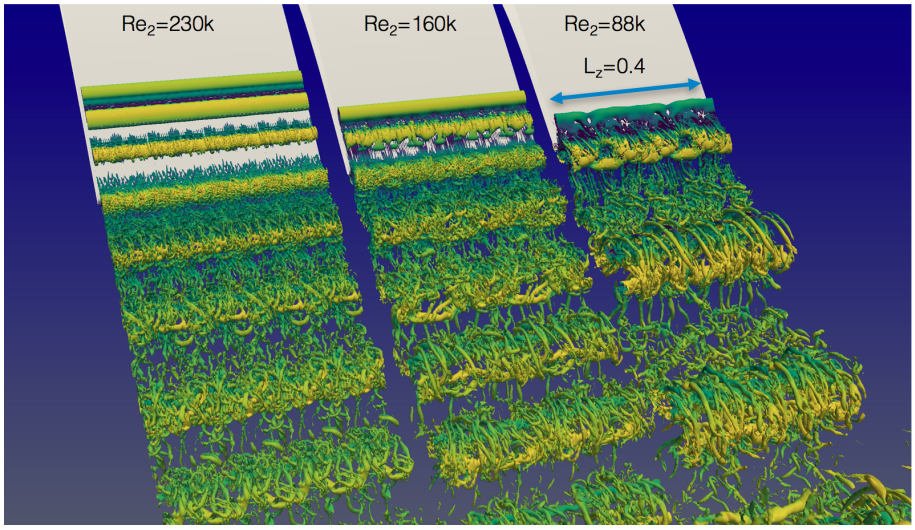

Figure 3.26: Instantaneous iso-surfaces of Q (Q=500) coloured by velocity magnitude. Increasing Re_2 is shown from left to right. Original computational domain (replicated for graphical purposes): (right) $Re_2 = 88k$: $L_z = 0.2C$; (middle) $Re_2 = 160k$: $L_z = 0.2C$; (left) $Re_2 = 230k$: $L_z = 0.1C$.

Re_2 compared to the experimentally declared one. Extensive testing allowed to rule out h- and p-convergence issues, as well as the stabilisation technique employed. The clean inflow setup behaves like an extremely silent wind tunnel, and the absence of noise allows for natural transition to occur. Likely, even a low level of background turbulence ($TI = 0.5\%$, for instance) would be sufficient to trigger anticipated transition and reattachment at lower Reynolds number, yielding more accurate matching of the pressure distribution at the correct Reynolds number. Another important factor introducing discrepancies is related to the uncertainties observed in the definition of the reference values used for numerical data reduction, in respect to position and type of averaging for the reference pressure and velocity values. For this reason, uncertainties were estimated and included to the comparison in figure by adding an envelope to the C_p range of the COMP3 curve (dashed black lines in Figure 3.25) around the mean value derived from area-mean integrals at section $0.25C_{ax}$. The ranges are not based on formal uncertainty quantification methods, but they are derived from evaluation of maximum and minimum local static pressure values in the extraction plane $0.25C_{ax}$ downstream of TE. On the pressure surface, the C_p profile is not sensitive to these effects as for the suction side, and good agreement is found with the experimental data. The small "kink" in the LE region existing in data set EXP1 around $s/S_0 = 0.05$ disappears in both cases EXP2 and EXP3. None of the preliminary simulations performed reproduced such feature within the Re spanned, even considering further local mesh refinements targeting that area. Another element of uncertainty is the streamtube contraction of the experimental rig, which is not accounted for in the numerical simulations. Chapters 5 and 6 discuss this topic in further detail.

The time- and spanwise-averaged boundary layer profiles from case COMP3 and data set EXP1 are compared in Figure 3.27. Good agreement is found in the evolution of

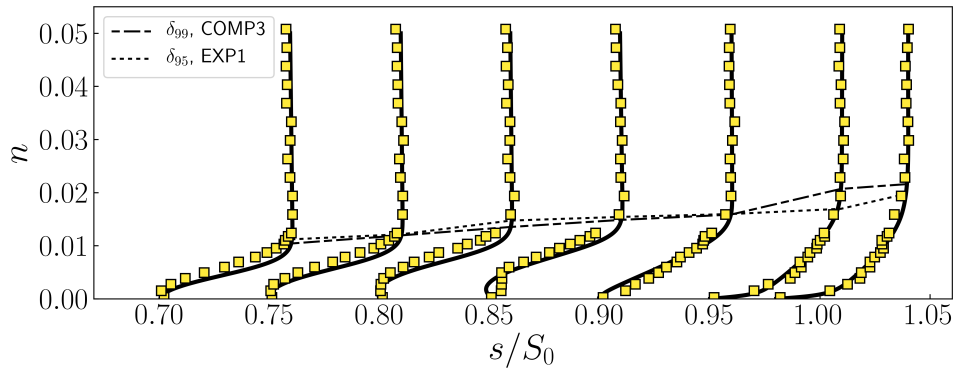

Figure 3.27: Boundary layer profiles on the suction surface. Dashed lines show experimental and computational BL edge. Copyright © 2018 by Rolls-Royce plc.

the profiles along the suction surface, and the BL thickness predicted in close agreement with the experiments. The remarkable matching of these profiles confirms the accuracy quality achieved in the numerical simulation COMP3 and justifies its comparison to EXP1, regardless of the offset required to the Re_2 and incidence nominal values.

3.7 Discussion

This chapter presented a systematic investigation of the T106A blade at fixed Reynolds number. The effect of the numerical setup and scheme features was discussed, looking at mean flow statistics and turbulent spectral properties. A total of 18 cases were presented, with averages sampled over 24 chord based flow through time units in all the cases concerning the mesh independence study (15 cases).

The concept of convergence of the results depends on the type of statistics of interest and the required level of accuracy. In line with theoretical expectations, the impact of p-refinement on various statistics is demonstrated to be a powerful tool enabling exponential error reduction, as shown by the RMS of blade wall distributions and BL parameters which gain an order-of-magnitude increase in accuracy every time P is increased by 2. The high efficiency of the numerical implementation allows to run cases at very high order in feasible computation times.

Variations of the computational setup in the spanwise direction (spanwise domain L_z and number of Fourier planes N_z) yield minor discrepancies in the mean flow properties and suggest that even extremely constraining setups may provide a cost-effective way of assessing the main blade performance indicators, without substantial losses in accuracy. The analysis of two-point correlations highlighted that in absence of inflow disturbances, large-scale coherent structures arise in the aft portion of the suction surface, requiring a large spanwise domain to avoid introducing constraints on the maximum wavelenth allowed to develop.

A comparison against an experimental data set is presented, building on the knowledge

developed in the preceding sections of the Chapter. Modified boundary conditions were employed due to uncertainties in the experimental definition of reference parameters for numerical data reduction. The computational results demonstrate good agreement against experiments at $Re_2 = 160000$, which were found within an uncertainty range appropriately defined.

All results discussed in this chapter were obtained with clean inflow boundary conditions to analyse convergence properties independently from time-varying boundary conditions, which would introduce additional complications into the problem. However, clean inflow does not accurately reflect realistic flow physics. The following Chapter is therefore dedicated to a discussion of the introduction of inflow turbulence in a linear LPT cascade.

Chapter 4

Generation of inflow disturbances

This Chapter focuses on the topic of inflow turbulence generation, with application to a representative industrial LPT cascade, building on the assessment reported in Chapter 3. The cascade design and regime selected are intentionally challenging both for the accuracy and flow stability. High fidelity experimental data are available for the case, that also offers the possibility for exploring the Reynolds number sensitivity and wake passing interactions; these last two aspects are discussed in Chapters 5 and 6. Two methodologies to introduce inflow disturbances are explored. The first approach belongs to the synthetic volume forcing category, based on the idea of adding momentum forcing terms to the Navier-Stokes equation system in a control zone. This should promote the development of turbulent structures and ideally have a limited impact on the downstream turbulent region. In this work, the approach by Schlatter et al. [148, 149] was repurposed and simplified to be explored for the first time in turbomachinery flows. Leveraging its reduced computational cost, it can be used as a fast way to simulate the cleaning of turbulence from the LE to peak suction, and therefore determine the best suction side design to prevent the emergence of an open separation bubble in the adverse pressure gradient region. Subsequently, as part of the effort to develop a robust and systematic inflow turbulence generation methodology via a random Fourier method, Davidson's approach [35] (previously introduced and validated in Section 2.3) is explored in the context of turbomachinery applications.

The Chapter is organized in the following sections: first, the computational approach and the numerical setup of the inflow disturbances is discussed. Results are then presented and subdivided in three parts, analyzing the effect of inflow turbulence on (1) blade distributions, (2) wake profiles and (3) boundary layer parameters. Comparisons against experimental data are also established. The final section concludes the Chapter with an overview of the main findings.

4.1 Problem setup and numerical approach

The test case considered is a representative industrial (LPT) cascade [88]. A moderate inflow Reynolds number of $Re_\infty = 58581$ is employed, based on inflow velocity $U_\infty = 1$

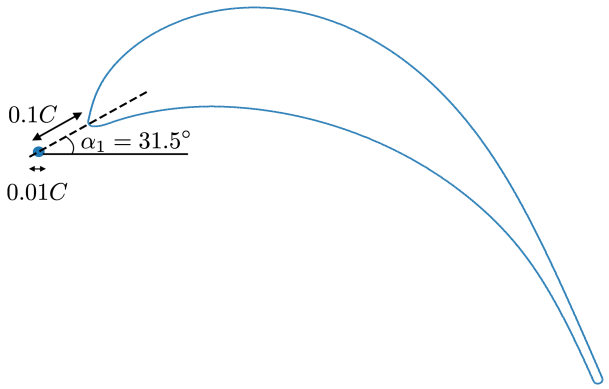

Figure 4.1: Momentum forcing location in the LE region. Figure not to scale.

and chord $C = 1$. This corresponds approximately to $Re_2 = 111200$, based on mixed-out velocity magnitude at the exit measurement plane. The setup of the simulation model, inflow-outflow average conditions and extraction planes were informed using RANS simulation matched to test-data for the actual 3D test rig configuration (Hydra R-R solver and recommended quality procedures).

The inflow plane is situated $1C_{ax}$ upstream of the leading edge. The numerical setup is analogous to that of Chapter 3: high-order outflow conditions [41] are applied to the outflow plane, located $1.5C_{ax}$ downstream of the trailing edge, while periodic conditions are enforced in the pitchwise and spanwise directions, with pitch to chord ratio of $P_y/C = 0.84412$. The inflow angle is $\alpha_\infty = 31.5°$. The upstream and downstream measurement planes are located respectively $0.5C_{ax}$ upstream of the LE and downstream of the TE. The $x - y$ plane is discretised with an expansion of order $P = 7$, the spanwise domain is extended for $L_z = 0.2C$ and discretised with $N_z = 96$ Fourier planes. For the exponential kernel of SVV in the spanwise direction, diffusion coefficient $\mu_{\text{SVV}} = 1$ and cutoff ratio of $M/N_z = 0.5$ are employed.

NekMesh [172] was employed to generate the high-order mesh, following the same criteria previously highlighted in Chapter 3, resulting in a mesh with 2704 quadrilateral elements in the O-mesh and 7188 triangular elements in the unstructured mesh.

4.1.1 Momentum forcing setup

A localized synthetic forcing with spanwise-varying u and v velocity components is added to the right-hand side of the Navier-Stokes equations, with the aim of disturbing the flow

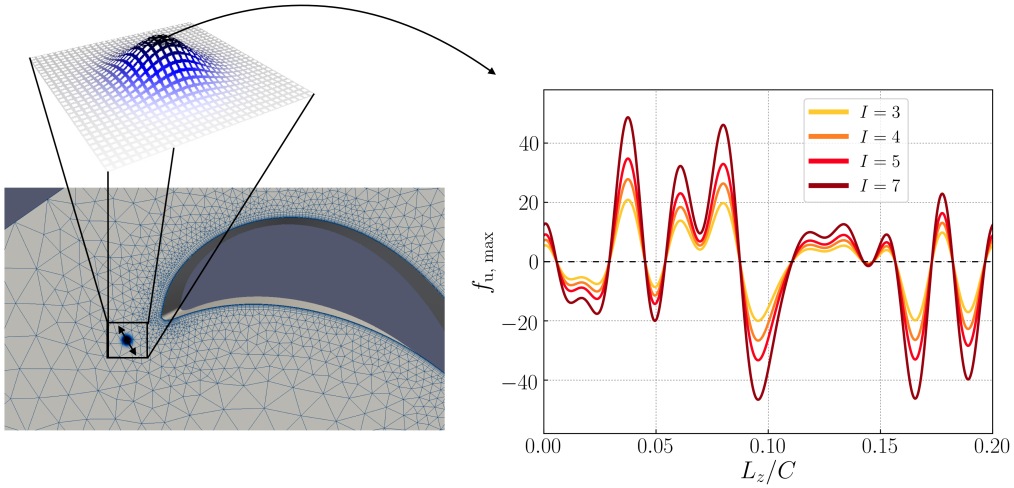

Figure 4.2: Momentum forcing in the LE region. On the left, the blue dot indicates the location of the forcing. on the right, the various intensities of the forcing function are shown along the spanwise direction. Copyright © 2019 by Rolls-Royce plc.

and triggering symmetry breakdown along the blade:

$$\mathbf{f_u}(x,y,z,t) = \begin{cases} I(t) \cdot \frac{g(z)}{\int_0^{L_z} \sqrt{g(z)^2}} \cdot \alpha \cdot e^{\frac{-\left[(x-x_c)^2+(y-y_c)^2\right]}{\delta^2}} \\ I(t) \cdot \frac{g(z)}{\int_0^{L_z} \sqrt{g(z)^2}} \cdot \beta \cdot e^{\frac{-\left[(x-x_c)^2+(y-y_c)^2\right]}{\delta^2}} \\ 0 \end{cases} \tag{4.1}$$

In this formulation, $g(z) = \sum_{i=1}^{N_{\text{body}}} A_i \sin\left(\frac{2\pi}{L_z} iz + \phi_i\right)$ and (x_c, y_c) represents the location of the forcing. $N_{\text{body}} = 10$, $A_i = 1$ and random phase ϕ_i were employed.

This approach is similar to the one adopted in [149]. The tripping effect is addressed in both steady (mimicking an experimental-like trip-wire configuration) and unsteady configurations. Overall, the implementation consists of spanwise-varying modes (the red component of Equation 4.1) with a Gaussian bell in the $x-y$ plane, restricting the action of the forcing to a region near the leading edge (the blue component of Equation 4.1), which proved to be very effective in perturbing the symmetry of the clean laminar incoming flow. By appropriate choice of the coefficients α, β, the disturbance was made perpendicular to the flow direction. After preliminary sensitivity testing for the location and size of the forcing, Figure 4.1 shows it was situated $0.1C$ from the leading edge (measured in the direction of the incoming flow), with a characteristic dimension of $\delta = 0.01C$: the pressure distribution in the front portion of the blade showed no appreciable effect with this choice.

Figure 4.2 shows the location of the forcing and its qualitative structure in the $x-y$ plane on the left, while the function $g(z)$ employed for the spanwise direction is shown on the right.

The effect of increasing forcing amplitude is analysed in four steady cases ($I = 3, 4, 5, 7$) and qualitatively shown in Figure 4.3. The baseline case ($I = 0$, on the left) sheds coherent

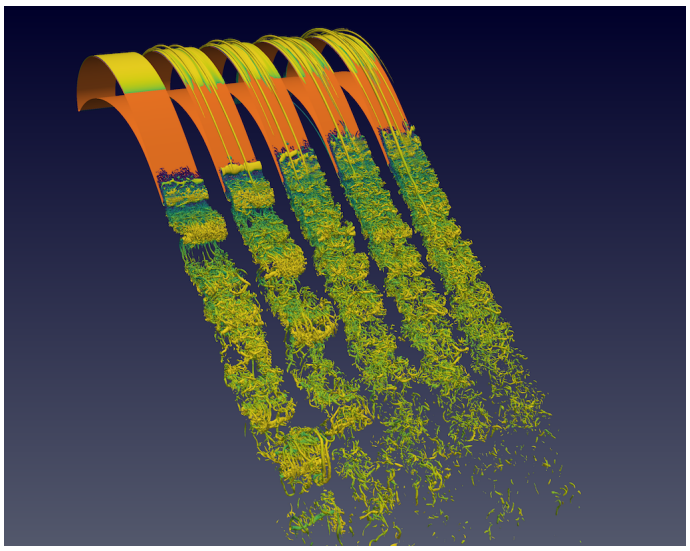

Figure 4.3: Instantaneous isosurfaces of Q contoured by velocity magnitude. left to right: clean inflow case, and increasing forcing intensity. Copyright © 2019 by Rolls-Royce plc.

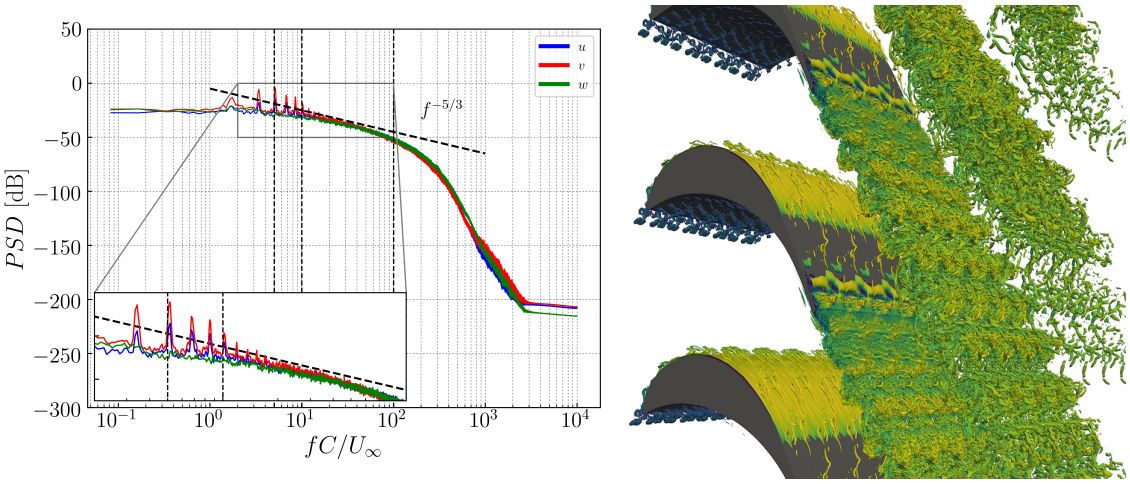

(a) (b)

Figure 4.4: (a) velocity spectra in the turbulent wake of the LPT with clean inflow. The frequencies $fC/U_\infty = 5, 10, 100$ are highlighted by the vertical dashed lines. (b) Isosurfaces of Q ($Q = 200$) contoured by velocity magnitude in case $I^* = 4$. The computational domain is replicated in the spanwise and pitchwise directions for graphical purposes.

large-scale vortical structures from the TE, while higher forcing intensity modifies the transition mechanism and increases the content of small coherent structures in the wake of the blade, by closing the separation bubble that is otherwise open in the unforced case.

The amplitude $I = 4$ yields the best overall agreement with experimental data. Further experiments were carried out by introducing unsteadiness in the forcing function: $I(t) = I \cdot \sin(2\pi fC/U_\infty t)$, with t representing nondimensional time. The dimensionless frequencies $fC/U_\infty = 5, 10, 100$ were selected; the first two values correspond to low-frequency peaks

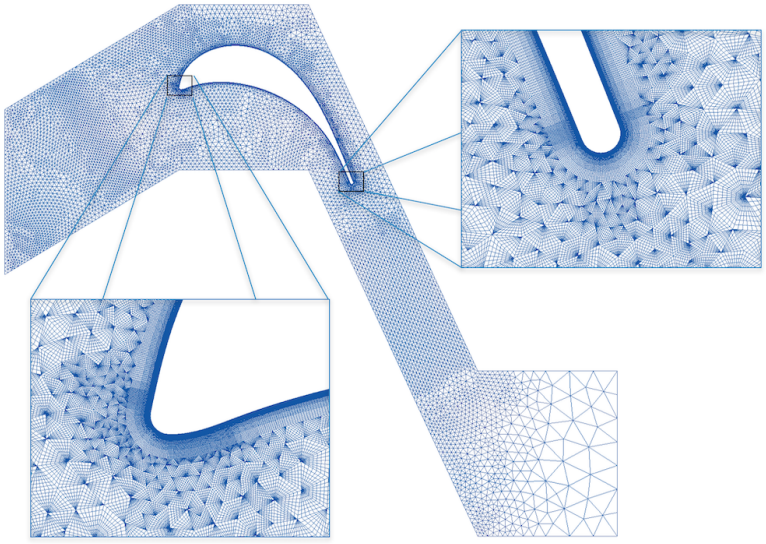

Figure 4.5: Base mesh for the inflow turbulence cases and (zoomed) high order TE and LE mesh with $P = 7$. Figure not to scale. Copyright © 2019 by Rolls-Royce plc.

in the velocity spectra, measured in various locations of the separated portion of the suction surface (shown in Figure 4.4a). The effect of $fC/U_\infty = 10$ with $I = 4$ provided the best improvement in the agreement with experimental data, and is referenced through the Chapter as $I^* = 4$.

Figure 4.4b provides a qualitative representation of the time-varying body forcing effect. This method provides a fast way to test the response of a profile (and in particular the suction side) to inflow disturbances: the forcing is confined to the near-wall region and therefore enhanced mesh resolution in the free-stream is not required, but it can nontheless provide indications of the suction side response to incoming turbulence.

4.1.2 Synthetic inflow turbulence setup

The algorithm and implementation of Davidson's method [35] was described in Section 2.3.3. To accurately resolve the development of turbulent structures in the inflow region of the computational domain, a modified mesh was adopted, as shown in Figure 4.5. The base mesh has 3024 quadrilaterals elements and 15663 triangular elements, yielding a total of 72.7 million DoF per variable at $P = 7$. Table 4.1 reports the choice of parameters adopted for synthetic turbulence generation. In general, length scales found in experiments or real engines are difficult to reproduce in scale-resolving simulations because of the demanding computational cost that they require in a suitable large domain. Sensitivity of the statistics to the domain size is therefore left to further analysis. The chosen numerical setup was retained throughout the study to provide a consistent comparison of the simulated flow fields.

The evolution of the turbulent fluctuations was monitored from the inflow plane along the streamline leading to the LE. It is reminded that the turbulence intensity is commonly

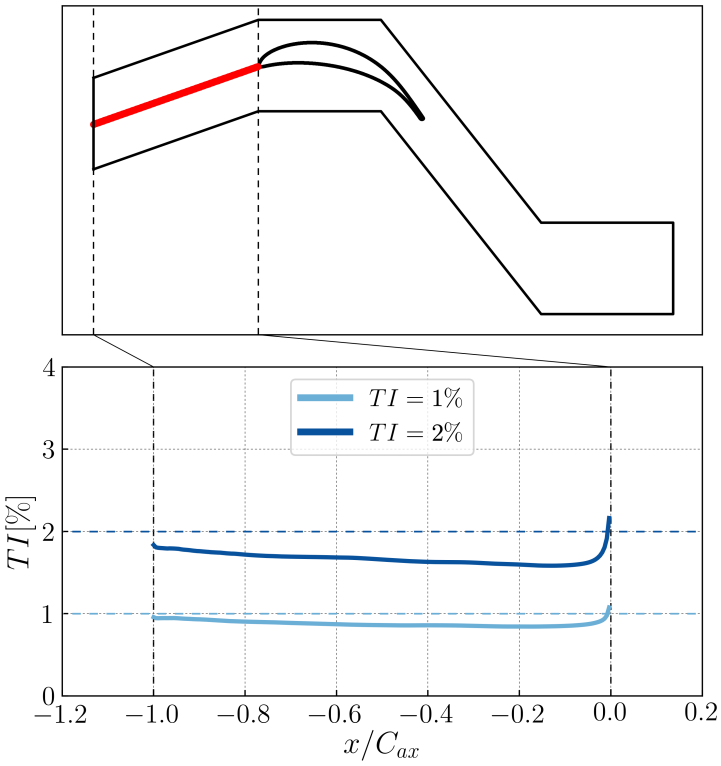

Figure 4.6: TI decay in the leading edge path, averaged in the spanwise direction. Top figure not to scale. Copyright © 2019 by Rolls-Royce plc.

Table 4.1: Parameters of inflow turbulence test cases analysed.

N_{turb}	L_t	$\mathcal{T}$	TI[%]
1024	$0.1C$	$0.1C$	1
1024	$0.1C$	$0.1C$	2

defined as $\text{TI} = (\sqrt{2/3k})/U_\infty$, where $k = 1/2(\overline{u'^2} + \overline{'v^2} + \overline{w'^2})$ is the TKE. The TI is subject to very little decay in the development region with both turbulent intensities analysed, as shown in Figure 4.6. The injected turbulence intensity levels are close to the user specified values, with a discrepancy of 4.18% and 9.07% for the cases TI = 1% and TI = 2%, respectively. The TI recovers in the region immediately preceding the leading edge due to the presence of the blade and reaches a level that slightly exceeds the prescribed value at the inflow. In under-resolved simulations during sensitivity tests, the evolution of the decaying TI curve was found to present irregular peaks. Therefore, a smooth decay curve is regarded as a necessary (although not sufficient) indication of an adequately refined mesh.

The streamwise velocity spectra are shown in Figure 4.7 at various locations along the development region, along the red line of Figure 4.6. The spectrum quickly adjusts itself and the profile at $x/C_{ax} = -0.75$ is qualitatively the same as the profile at $x/C_{ax} = -0.1$. The streamwise velocity spectra show minor decay at high frequency; this effect is

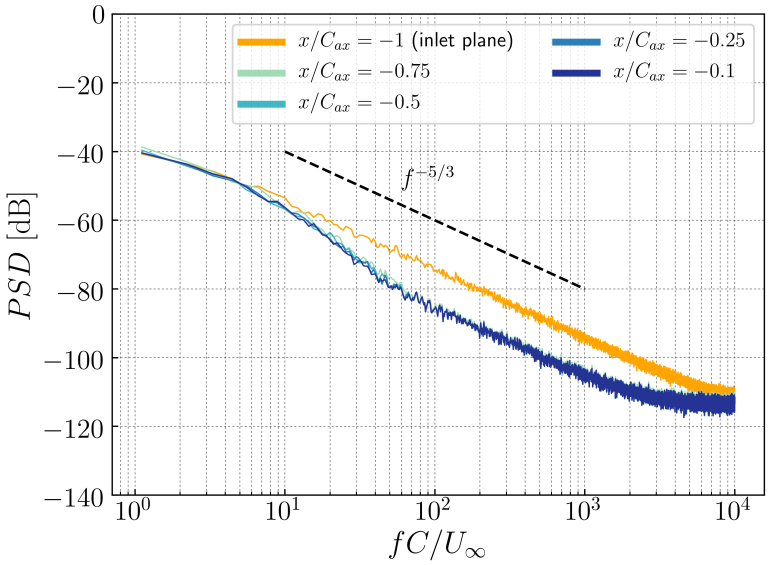

Figure 4.7: Streamwise velocity spectra in the centre of the development region, in different streamwise location. The orange line is the imposed velocity spectrum at the inlet. Copyright © 2019 by Rolls-Royce plc.

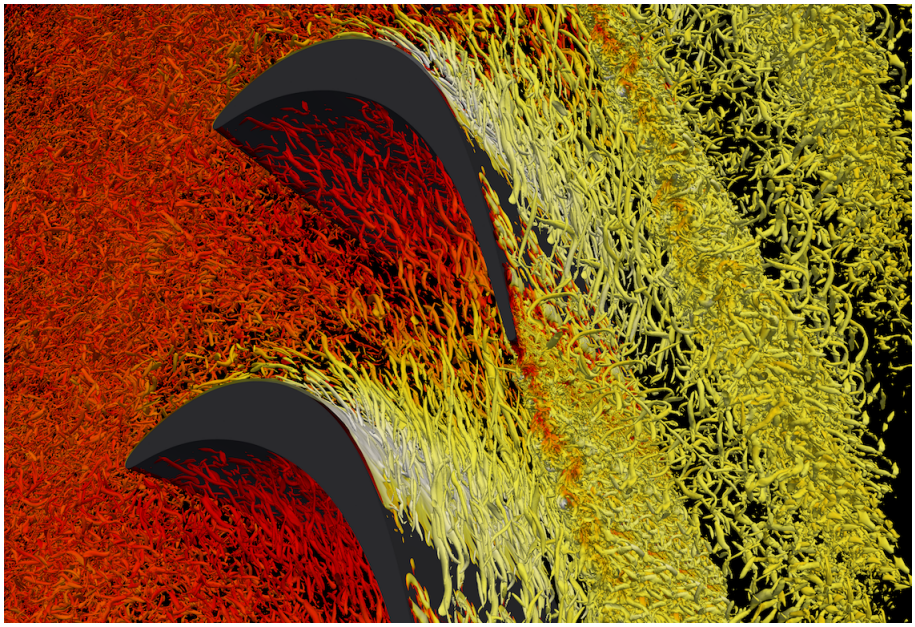

Figure 4.8: Instantaneous isosurfaces of Q contoured by velocity magnitude. T106A cascade with inflow TI = 10% for visual representation of the synthetically generated turbulent structures at high inflow turbulence levels. The computational domain is replicated in the spanwise and pitchwise directions for graphical purposes.

more enhanced in the pitchwise and spanwise velocity spectra, as previously shown in the periodic channel test case described in the validation of Section 2.3.4.

Figure 4.8 shows the qualitative effect of the introduction of the synthetic turbulence boundary conditions on the T106A cascade of Chapter 3: Q isocontours highlight the

presence of turbulent structures of different scales being advected through the passage and over the blade.

Given the additional cost of generating synthetic turbulence and resolving the flow scales in the devleopment region, there is particular interest in minimising the transient times. In addition to the quantities enumerated in Section 3.1, the indication for statistical convergence was also based on the time evolution of enstrophy:

$$\varepsilon(\omega) = \frac{1}{2\mu(\Omega)} \int_{\Omega} ||\omega^2|| \mathrm{d}x \,, \tag{4.2}$$

where $\mu(\Omega)$ is the volume of the computational domain Ω. This approach suggested that $T = 8C/U_\infty$ is a sufficient transient time for the inflow turbulence simulations. Time averages were subsequently sampled for further 9 chord-based flow through times, while the simulations using the momentum forcing model were advanced for 18 chord-based flow through times. Progressively averaged statistics were compared to confirm time convergence.

The rest of the chapter focuses on the variation of the main performance indicators subject to the two inflow methodologies, progressively increasing the disturbance intensity. The numerical data processing is performed analogously to Chapter 3.

4.2 Time-averaged results

4.2.1 Blade wall distributions

The pressure distribution subject to increasing forcing intensity is shown in Figure 4.9. The flow distortion near the leading edge does not impact the distribution in the front portion of the blade. However, C_p is modified in the aft portion of the suction surface, especially in the separated region. A shift to a noisier environment causes a change in the transition mechanism, which is associated to an extended transition length, as suggested by the decrease in C_p diffusion gradient downstream of the turning point (see inset of Figure 4.9).

The experimental data points are shown in Figure 4.10 and compared to a low-magnitude forcing intensity ($I = 4$, both stationary and time-varying) and low level of inflow turbulence (TI = 1%). The latter case introduces minor changes in the C_p distribution, visible only in the reattachment region of the suction surface. Instead, in case $I = 4$ and $I^* = 4$ the separation bubble (plateau near the TE) shrinks owing to earlier reattachment, providing an improved agreement with experimental data in the TE region.

Having previously assessed the convergence properties of the numerical setup employed in this range of Re, the remarkable agreement of the disturbed setup with the experimental data set suggests that low levels of physical disturbances are necessary to capture the physics correctly in the separated region, even when the experiments are conducted in absence of turbulence grids.

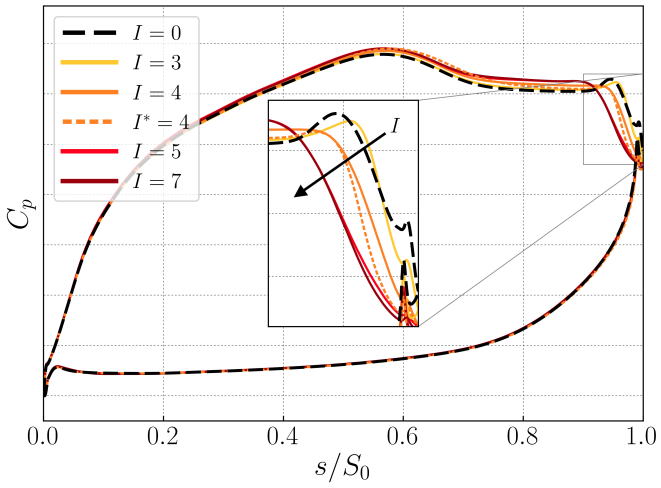

Figure 4.9: Pressure coefficient with increasing body forcing amplitude. Y-axis tick labels are omitted due to data sensitivity. Copyright © 2019 by Rolls-Royce plc.

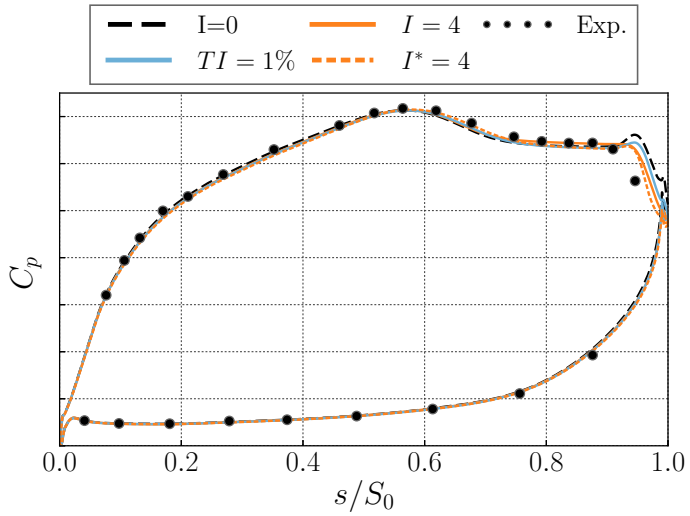

Figure 4.10: Pressure coeffcient: comparison between the baseline clean case, $I = 4$ and $I^* = 4$, TI = 1% and experimental results. Y-axis tick labels are omitted due to data sensitivity. Copyright © 2019 by Rolls-Royce plc.

Unlike the pressure distribution, the skin-friction coefficient C_f is distorted along the whole blade by the steady forcing (Figure 4.11a) on both the suction and pressure surface. The small secondary recirculation bubble found at $s/S_0 \approx 0.95$ in the clean inflow case is suppressed (in a time-averaged sense) as I is increased. The synthetic disturbances energize the flow, resulting in delayed transition and anticipated reattachment of the separated boundary layer before the TE. The effect of the disturbance at $I = 3$ is that of enhancing the negative C_f near the TE (yellow curve), but the negative peak is reduced as I is further increased. The curves $I = 5$ and $I = 7$ are overlapped, suggesting that the effects from this type of forcing may saturate. The introduction of an unsteady component

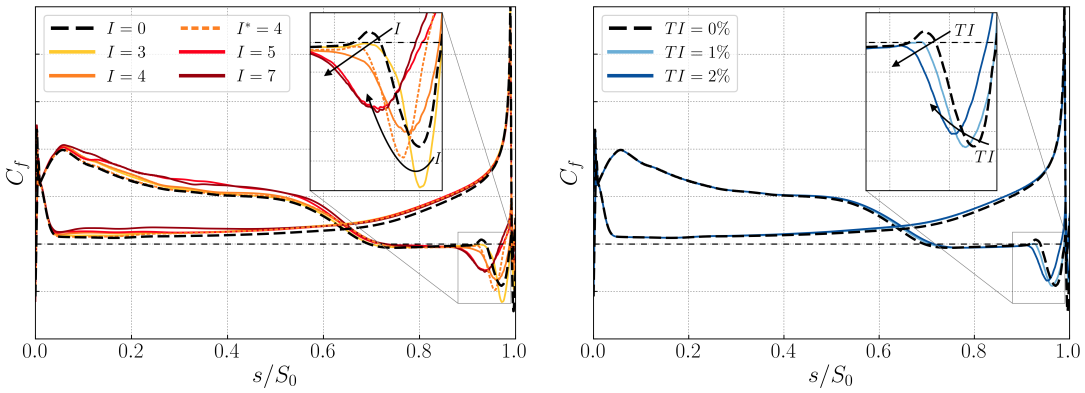

(a) (b)

Figure 4.11: Skin friction coefficient with increasing body forcing amplitude (a), and inflow turbulence intensity (b). Y-axis tick labels are omitted due to data sensitivity. Copyright © 2019 by Rolls-Royce plc.

Table 4.2: Time-averaged separation point with increasing inflow disturbance intensity. All values are normalised by the baseline case $I = 0$.

Case	I					TI [%]	
	3	4	*4	5	7	1	2
$(s/S_0)_{\text{sep}}$	1.00645	1.03624	1.03409	1.04208	1.05068	1.01815	1.02983

in the forcing ($I^* = 4$) shows a delayed but shorter reattachment region compared to the steady forcing at the same intensity level, and it does not modify the C_f distribution along the front portion of the blade with respect to the clean inflow case. No significant changes are observed along the whole blade in the synthetic turbulence cases (Figure 4.11b), except for the aft portion of the suction surface. Analogously to the body forcing case, the secondary recirculation bubble quickly disappears. An increase in TI causes a self-similar shift of C_f in the TE region towards anticipated reattachment, sufficient to close the separation bubble before the TE. The time- and spanwise-averaged separation point is reported in Table 4.2 for all the cases considered, and it shows that stronger disturbances promote increasingly delayed separation.

4.2.2 Wake profiles

The normalized distance from the trailing edge is defined as $\hat{x} = (x - x_{\text{TE}})/C_{ax}$. The wake profiles are extracted by interpolating the time-averaged velocity fields from the unstructured mesh to traverses of equispaced points in the pitchwise direction; four extraction stations with streamwise locations $\hat{x} = 0.1, 0.3, 0.5, 0.7$ are analysed. Subsequently, the clean case (taken as reference) is vertically shifted so that the profile peak is shown at $y/P_y = 0.5$; the other curves are rearranged based on the reference. This ensures that the plot is independent of the location of the wake relative to the domain.

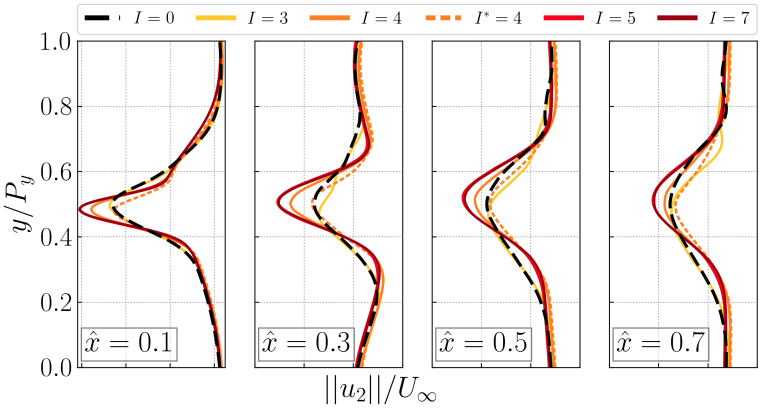

(a)

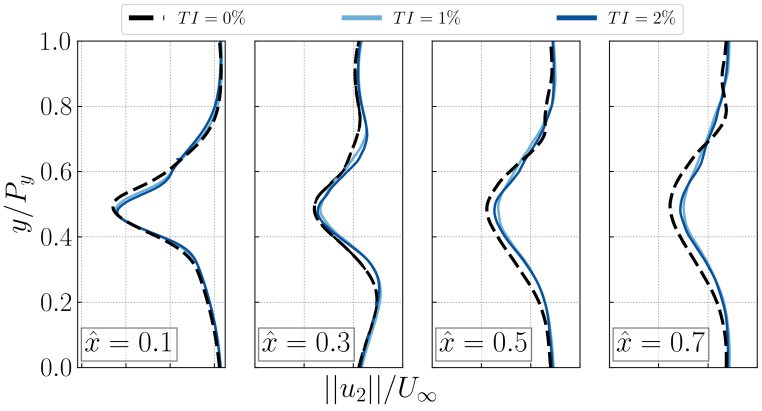

(b)

Figure 4.12: Velocity profiles in four streamwise locations downstream of the trailing edge. (a): body forcing. (b): inflow turbulence. X-axis tick labels are omitted due to data sensitivity. Copyright © 2019 by Rolls-Royce plc.

The velocity profiles are shown in Figure 4.12. In the body forcing simulations, as the forcing amplitude increases the peak is initially damped and shifted ($I = 3$, visible for $\hat{x} > 0.3$), and then significantly enhanced; the flanks of the profile concurrently become thinner. The intensities $I = 5$ and $I = 7$ are almost completely overlapped. $I^* = 4$ has the same wake width as $I = 4$, but reduced peak. In the synthetic turbulence simulations, the difference in velocity wake between the clean case and the ones with inflow disturbances is minor, with the peak becoming less pronounced and the width slightly reduced. The trend followed by the peak is different from the one observed in the body forcing case.

The TKE wakes are shown in Figure 4.13. Differently from the velocity profiles, both the body forcing cases and the synthetic turbulence cases show the same trend. As the disturbance intensity is increased, the bypass transition mechanism progressively moves the reattachment point upstream, closing the separation bubble. Therefore, the width of the TKE wake is reduced and the peak is progressively less pronounced as the intensity of the disturbances increases.

The profiles of cases $I = 4$, $I^* = 4$ and TI = 1% at $\hat{x} = 0.5$ are compared to LDA

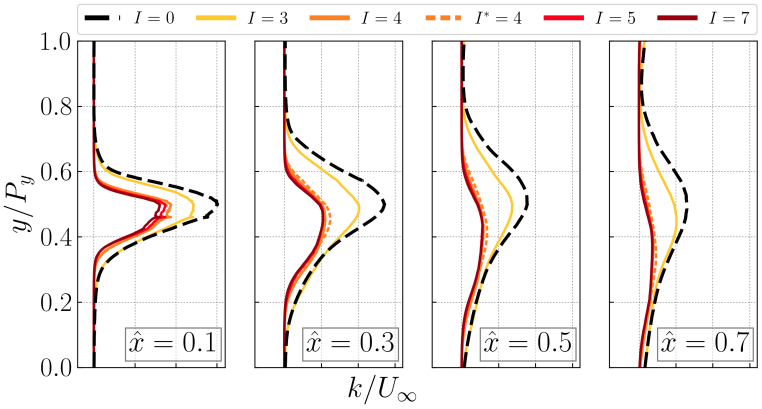

(a)

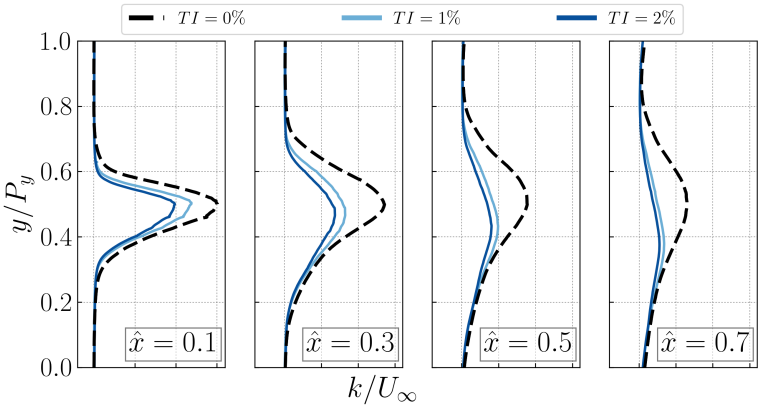

(b)

Figure 4.13: Turbulence kinetic energy profiles in four streamwise locations downstream of the trailing edge. (a): body forcing. (b): inflow turbulence. X-axis tick labels are omitted due to data sensitivity. Copyright © 2019 by Rolls-Royce plc.

experimental data in Figure 4.14. Three instrumented blades are reported to provide a qualitative indication of the measurement uncertainty, and labelled S1, S2 and S3. The numerical data reduction followed in this Figure is modified to match the experiments, with U_2 indicating mixed-out exit velocity. The wake profiles are shown in the Figure on the left: while $I = 4$ overestimates the wake peak and TI = 1% slightly underestimates it, $I^* = 4$ follows the experimental profile extremely closely. The clean inflow case also accurately captures the wake peak, but predicts wider wake flanks. As previously highlighted from Figure 4.11a, the introduction of momentum forcing closes the separation bubble by moving the reattachment point upstream. In particular, the introduction of unsteady forcing promotes earlier reattachment compared to steady forcing, and thus the dynamics of the wake is more strongly affected. The middle Figure compares the TKE profiles. In this case, the pointwise normalisation with the exit velocity was required to match the experiments. The clean inflow case accurately captures the wake peak but it is shifted towards the suction surface, and the wake width is overpredicted. The introduction of disturbances promotes earlier reattachment, improving the agreement with experiments:

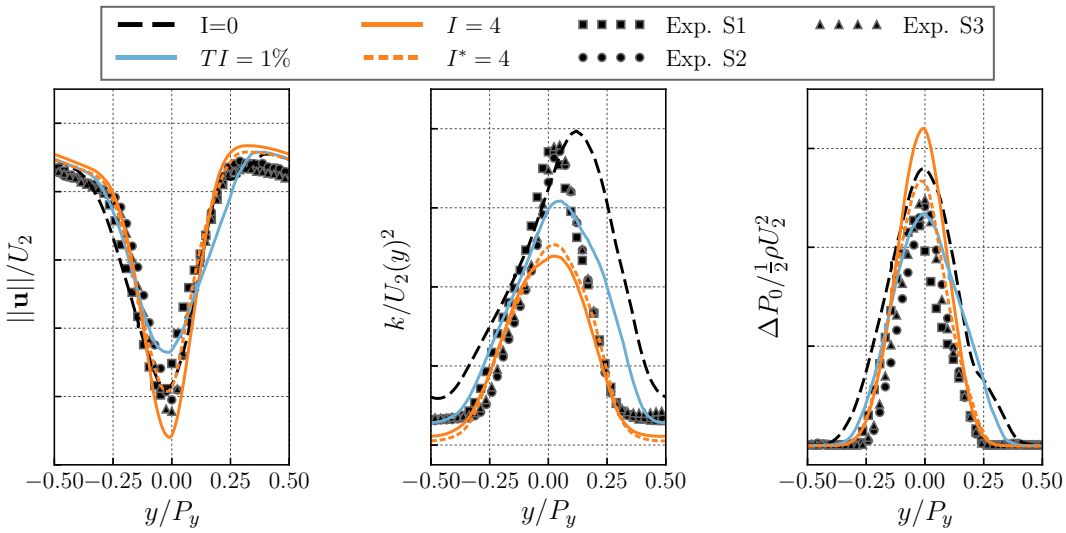

Figure 4.14: Comparison of selected cases with experimental data at $\hat{x} = 0.5$. Left: velocity magnitude; middle: turbulence kinetic energy, and right: total pressure loss coefficient. Y-axis tick labels are omitted due to data sensitivity. Copyright © 2019 by Rolls-Royce plc.

the peak is captured well by the inflow turbulence simulations, while the momentum forcing allows for higher accuracy in the representation of the profile width. Differently from the velocity wakes, the introduction of an unsteady disturbance $I^* = 4$ does not cause significant changes in the TKE wake.

Finally, on the right is the total pressure loss coefficient, sometimes referred to as KSI:

$$\text{KSI} = \frac{\Delta P_0}{\frac{1}{2}\rho U_2^2} = \frac{p_{01}^M - (p_2(y) + \frac{1}{2}\rho||\mathbf{u}_2(y)||^2)}{\frac{1}{2}\rho U_2^2}\,. \tag{4.3}$$

Consistently with velocity and TKE wakes, cases $I = 4$ and $I^* = 4$ accurately follow the width of the profile, with the latter providing a very close match with the experimental loss profiles. As in the other statistics, the synthetic turbulence approach well represents the peak but fails to accurately captures the wake width. The results of three adjacent blade rows are reported to provide approximate indication of the uncertainty associated to the experimental measurements.

Extensive mesh independence studies confirmed that the resolution in the wake is sufficient to capture first and second order statistics. The physical effect underlying the mismatch between experimental data and synthetic turbulence simulations in predicting the wake would require futher investigation: the length scale prescribed might not be large enough, which would require a wider spanwise domain to accurately model the physical effects of turbulence. It should also be noted that any uncertainties in the exact streamwise location of the extraction plane would cause a discrepancy between the virtual cascade and experimental data.

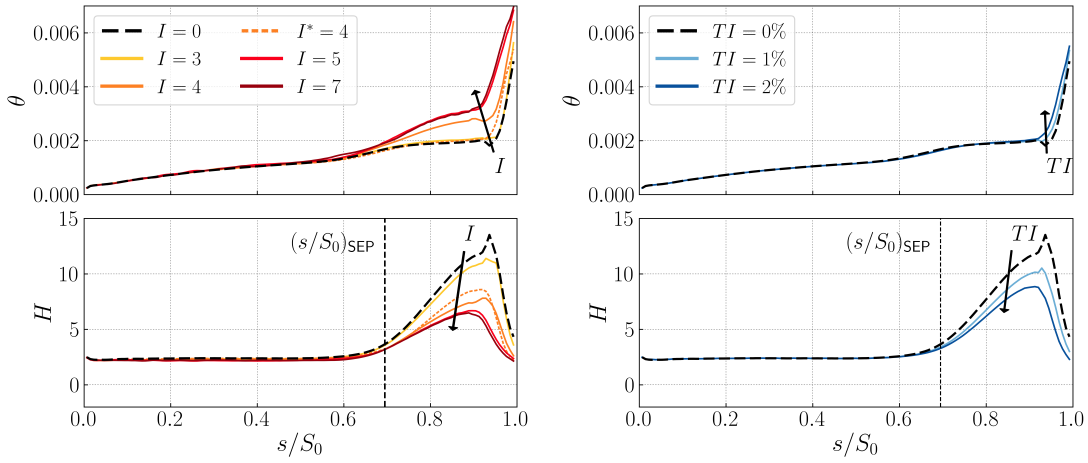

(a) (b)

Figure 4.15: Momentum thickness θ and shape factor H along the suction surface with increasing p. The vertical dashed line indicates the separation point in the baseline case. Copyright © 2019 by Rolls-Royce plc.

4.2.3 Boundary layer parameters

The BL parameters on the suction side of the turbine blade are linked to profile loss estimation [38] and flow separation. With an approach analogous to that reported in Chapter 3, the clean case is employed as reference for the BL integration limit.

Figure 4.15 shows the evolution of the momentum thickness and shape factor along the suction surface. In the body forcing case (Figure 4.15a), the momentum thickness is modified in the whole separated region, with $I = 3$ almost coinciding with the clean case and $I = 5$ overlapped with $I = 7$. This reinforces the idea that nonlinearities are responsible for the evolution of the boundary layer: in this case, a minimum level of forcing intensity is necessary to introduce a visible modification, and no further change is appreciated after reaching a certain value. The same behavior is observed in the evolution of the shape factor, with the peak being anticipated and reduced respectively by roughly 15.9%, 42.0%, 50.0% and 52.0% in cases $I = 3, 4, 5, 7$ and 36.5% in case $I^* = 4$, with respect to the reference clean case. Introducing unsteadiness in the forcing function ($I^* = 4$) substantially modifies the evolution of the momentum thickness in the separated region: θ remains almost unaltered until the very aft portion of the blade, which corresponds to the location where the secondary recirculation bubble disappears as the forcing is increased in amplitude (more clearly shown in the C_f distribution of Figure 4.11).

In the synthetic turbulence cases (Figure 4.15b), the momentum thickness behaves very similarly to $I^* = 4$, subject to a modification in the reattachment region only. The trend of the shape factor is in general analogous to the body forcing case, with peak reductions of 21.4% and 33.9% for TI = 1% and TI = 2%, respectively.

Table 4.3 reports the averaged BL parameters over four profiles extracted at the trailing edge. As previously mentioned, the momentum thickness at the trailing edge is linked to

Table 4.3: Boundary layer parameters at the trailing edge. All values are normalised by the baseline case $I = 0$.

Case	I					TI [%]	
	3	4	*4	5	7	1	2
θ_{TE}	1.147	1.347	1.159	1.474	1.512	1.113	1.170
H_{TE}	0.810	0.575	0.476	0.505	0.455	0.665	0.504

profile losses, and increasing I and TI corresponds to progressively larger θ_{TE}. On the contrary, the shape factor is subject to a decrease, as high levels of disturbances allow for earlier reattachment and therefore the boundary layer at the trailing edge is more skewed. It should be noted that the change in momentum thickness for the body forcing case is comparable to that of the turbulent simulations only for low disturbance intensities, suggesting that stronger forcing is unphysical. The introduction of an unsteady forcing maintains θ_{TE} in the same range of values as the synthetic turbulence cases. This supports the use of the unsteady artificial forcing an alternative strategy to trip the boundary layer transition and obtain the same BL evolution as in a more expensive setup, that requires a highly refined mesh in the front part of the domain to support the evolution and decay of synthetic turbulent structures.

4.3 Discussion

This chapter presented an investigation of a representative industrial LPT cascade at moderate Reynolds number, exploring the effect of two different methodologies for the introduction of physical disturbances. The first consists of a localized momentum forcing located near the leading edge and perpendicular to the incoming flow, that breaks the flow symmetry along the blade and promotes earlier reattachment of the boundary layer on the suction surface. The unsteady formulation of the forcing provided excellent agreement with experimental results, supporting the use of this methodology as an alternative to more expensive methods, owing to its ease of implementation and less stringent resolution requirements. The second approach explored is a random Fourier method for the generation of synthetic inflow turbulence. This specific formulation is applied for the first time to the industrial LPT test case at low levels of turbulence intensity.

Not surprisingly, in this case of low inflow turbulence (typical of low speed cascades) the two methodologies produce similar trends in the modification of pressure and skin friction coefficient, turbulence kinetic energy wakes and boundary layer parameters. However, the evolution is different when looking in more detail at the flow structures and in particular the velocity wakes. The physical mechanism driving the two approaches is fundamentally different: the momentum forcing cases introduce spanwise distortions of the flow, that cannot be fully appreciated in the conventional spanwise-averaged statistics. However, it provides the necessary physical disturbance to initiate the breakdown of the turbulent

structures as expected, and therefore improves the agreement with experimental data, with little cost increase with respect to a clean inflow case. On the other hand, the synthetic inflow turbulence provides a more physical inflow condition: it requires higher resolution and computing cost to support turbulence development in the region upstream of the blade. This is necessary to allow the vortical structures to evolve and organize according to the Navier-Stokes equations from the assumed inflow synthetic generation model. The cost of generating turbulent structures on-the-fly proved to be low owing to the current parallel implementation.

Chapter 5

Reynolds sensitivity

This chapter focuses on the analysis of the Reynolds sensitivity of a LPT vane, with the aim of validating the accuracy of the setup developed in Chapter 3 over a range of flow regimes. Particular focus is placed on wake profiles and loss estimation, given the availability of experimental data. The representative LPT introduced in Chapter 4 is considered at three flow regimes, which correspond respectively to subcritical, critical and supercritical state of the suction surface separation bubble. The different behaviour of the boundary layer near the trailing edge translates into important changes of profile losses and wake characteristics, which motivates the importance of assessing the capability of the numerical solver to accurately capture the flow physics across the entire operating range.

5.1 Problem setup and numerical approach

5.1.1 Low-speed experimental testing of LPTs

The present generation of LPTs typically operates at subsonic conditions, characterised by exit Mach numbers of $Ma_2 \sim 0.6$. However, the cost of experimental programs can be significantly reduced by performing tests in low-speed wind tunnels, at essentially incompressible conditions. However, the pressure distribution is in general strongly dependent on the Mach number [175]: the peak suction Mach number ultimately controls the adverse pressure gradient flow region, where most losses are generated. Therefore, experimental testing cannot be carried out on the same blade shape and cascade setup. A number of scaling techniques as well as more advanced redesign strategies were developed to derive a modified profile shape and flow conditions, which allow to compensate for the effects of compressibility.

Among the most relevant redesign approaches, Vera & Hodson [176] matched the isentropic Mach ratio distribution in a linear cascade with design exit Mach number $Ma_2 = 0.64$, also modifying the inflow angle and pitch-to-chord ratio. More recently, Marconcini et al. [101] presented an artificial neural network approach to obtain a low-speed airfoil which matches the blade loading and the loss coefficient versus Reynolds number, and is characterised by similar boundary layer behaviour. This was demonstrated by developing

modified profiles matching specific target properties, i.e. pressure distribution, isentropic velocity ratio and isentropic Mach number ratio. The open separation bubble characteristic of low Reynolds was shown to pose additional challenges. These techniques are targeted at linear cascades. Giovannini et al. [53] introduced a methodology to scale a 3D high-speed rotor to low-speed conditions, also aiming to achieve comparable boundary layer behaviour (and therefore profile losses) by matching the pressure distribution. A response-surface approach based on artificial neural-networks was used to exploit the three-dimensional shaping of the blade, affecting the radial equilibrium of the flow field and adopting a shape refinement near the endwalls. Very good agreement was reported over the span at fixed Reynolds $Re_2 = 80K$.

The cascades in analysis in this thesis were redesigned to match the design high Mach number distribution when tested at incompressible conditions, modifying profile shape and increasing the inlet flow angle. This has critical implications for the applicability of the capability developed in this thesis. Since one of the main research purposes is the validation of a virtual wind tunnel capability, the use of an incompressible flow solver is not a road blocker: leveraging the scaling techniques mentioned above, realistic LPT geometries with a peak suction Mach number up to $0.6 - 0.7$ can be accurately simulated in their low-speed testing conditions.

5.1.2 Description of the experimental setup

The test facility at UPM (Polytechnic University of Madrid) is an open return, impulse wind tunnel that operates at very low Mach numbers (in a range between $Ma = 0.02-0.08$, depending on the Reynolds number). The approaching flow has a 495x240 mm rectangular cross-section. Prior to the test section, conditioning elements ensure a measured inlet turbulence intensity of $TI = 0.1\%$ and length scale $L_t = 0.5C$.

The rig was operated in a closed test-cell that can be isotropically seeded with oil steam, enabling the use of laser velocimetry techniques through the perspex test section. In order to study the effect of incoming wakes, a mechanism was included to introduce upstream generated wakes shed from uniformly spaced cylindrical bars, controlled by an electric motor. The design of the facility is analogous to other existing experimental rigs [152]. Flow diagnostics available in the experimental facility include measurements of wall pressure, inlet to outlet total pressure losses, and temporally resolved anemometry through hot-wire and laser-Doppler velocimetry; the latter was used to obtain the wake traverses presented in this work. For optimal flow stability, even when the cascade is nominally not subject to incoming disurbances, an isolated cylinder is swept through the inlet, resulting in a low reduced frequency of $F_{\text{red}} = U_b C/(P_b U_2) = 0.04$, where U_b and P_b indicate the bar speed and distance. In this case, the bar is estimated to impact on the blade every approximately $T = 16C/U_\infty$, therefore introducing minimal impact on the time-averaged statistics. This effect is conventionally not accounted for in computational studies focusing on clean inflow problems: it would require simulating an unfeasibly large number of blades

Table 5.1: Summary of setups analysed in this Chapter. The first column indicates the nominal values adopted for reference.

Re_2 (nominal)	Experimental	Re_2 - Clean inflow (CI)	Re_2 - body forcing (BF)
$83K$	833386.8	83126.45	83647.92
$155K$	155715.8	155691.33	156680.58
$290K$	289996.5	290005.61	284550.54

in the pitchwise direction. The wake passing phenomenon, its effect and its dynamics are explored in greater detail in the next chapter.

5.1.3 Computational approach

The three flow regimes considered are $Re_2 = 83K, 155K, 290K$. In this chapter as well as the following, the definition of the Reynolds number is modified for consistency with respect to the reference experimental data, and it is based on the suction surface perimeter:

$$Re_2 = \frac{U_2 S_0}{\nu}\,, \tag{5.1}$$

as opposed to the true chord adopted in Chapters 3 and 4. In the simulations, the inflow Reynolds number is specified. In the cases analysed, this corresponds respectively to $Re_1 = U_\infty S_0/\nu = 46K, 84K, 155K$. Six sets of numerical experiments were conducted: for each reference Re_2, the simulations targeted both clean inflow and the effect of incoming disturbances via momentum forcing. The latter was implemented in analogous manner to case $I^* = 4$ from Chapter 4. Complete details of the reference parameters of the simulations performed are shown in Table 5.1.

Based on prior experience, an improved mesh with 3066 quadrilateral elements in the O-mesh and 8276 triangular elements was produced for this test case. The resolution is particularly enhanced in the TE region and the wake. Based on the considerations of Chapter 3, the simulations are set up with $P = 7$, extending the domain for $L_z = 0.2C$ in the spanwise direction and discretising it with $N_z = 96$ Fourier planes. The numerical setup results in $47.4M$ DoF per variable. Each flow condition was run for at least 16 convective time scales to allow transients out of the computational domain, and further time-averaged for $T = 16C/U_\infty$, previously shown to be sufficient to achieve time convergence for the range of statistics of interest (see Section 3.3.6 for further details).

Figure 5.1 shows the effect of increasing Reynolds numbers on the near wall resolution. In the separated flow region, the resolution in wall units is within the limits identified for DNS [51], suggesting that even at the higher flow regime analysed $Re_2 = 290K$ the blade resolution is sufficient to capture the relevant scales of turbulence.

Figure 5.2 shows instantaneous contours of spanwise vorticity, which provide qualitative indication of the Reynolds effect on the cascade flow features. The shear layer in the aft portion of the suction surface becomes thinner and moves closer to the wall. Concurrently,

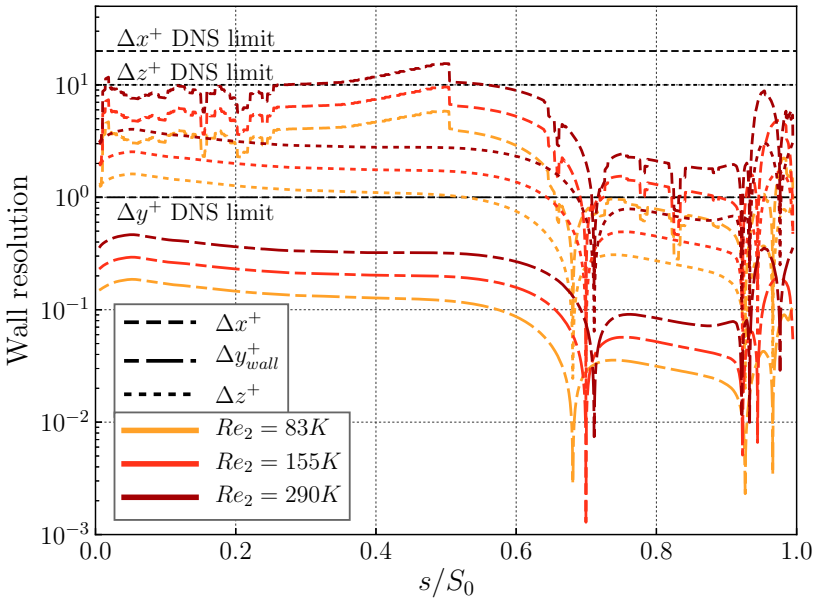

Figure 5.1: Near wall mesh resolution with clean inflow and increasing Re_2.

the wake dynamics is modified: the vortex shedding quickly breaks down to small-scale turbulence, while it remains significantly more coherent in the low Reynolds number case.

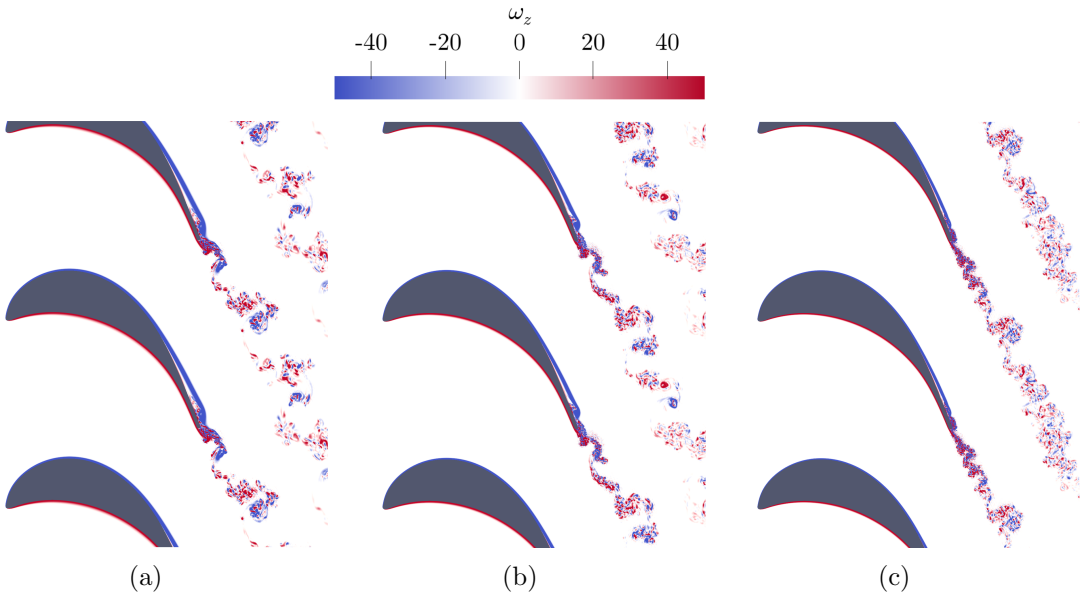

Figure 5.2: Contours of spanwise vorticity with increasing Reynolds numbers. (a) $Re_2 = 83K$, (b) $Re_2 = 155K$, (c) $Re_2 = 290K$. Figures not to scale.

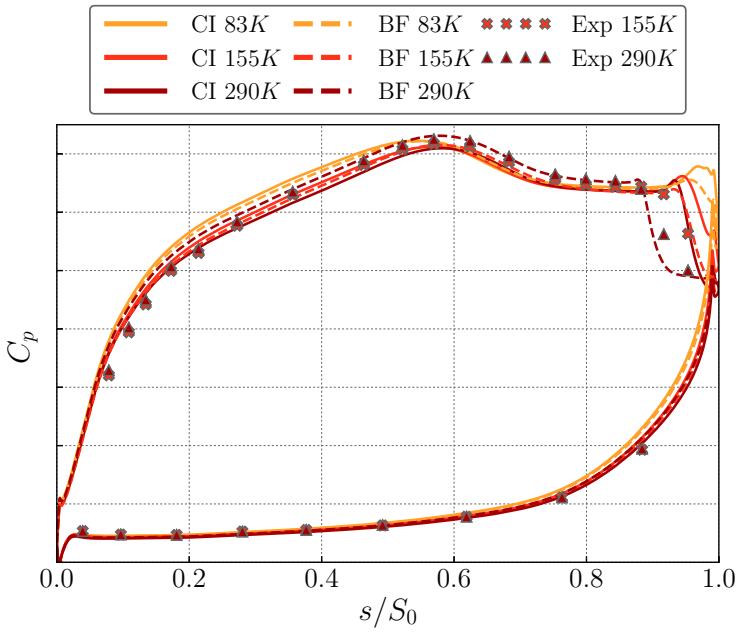

Figure 5.3: Pressure coefficient with increasing Re_2, comparing clean inflow cases (CI), body forcing near the leading edge (BF) and experimental results. Y-axis tick labels are omitted due to data sensitivity.

5.2 Blade wall distributions

5.2.1 Pressure coefficient

In this chapter, the pressure coefficient is calculated adopting the following definition, in line with experimental measurements:

$$C_p = \frac{p^A_{t,1} - p_s}{p^A_{t,2} - p^A_{s,2}} \,. \tag{5.2}$$

The Reynolds sensitivity of the pressure distribution is shown in Figure 5.3. Experimental data is available for $Re_2 = 155K, 290K$ only.

The agreement between experiments and numerical results is very close along the whole pressure surface. The suction surface peak is located at $s/S_0 \approx 0.6$ for all flow regimes analysed. Along the suction surface, good agreement with experiments is observed for $s/S_0 \leq 0.9$. Among the different cases shown, discrepancies can be observed in the front part of the suction surface owing in part to small differences in the inflow angle specified from the experiments. In the aft part, all cases accurately capture the separation location, where the distribution reaches a plateau. Further downstream, the clean inflow setup does not accurately capture the location of reattachment of the recirculation bubble. As previously demonstrated in Chapter 4, the introduction of inflow disturbances provides a sufficient level of noise.

Streamtube effect and definition

The absence of physical disturbances is not the only contributor to the discrepancy in the reattachment location, due to a phenomenon induced by the experimental setup here referred to as *streamtube effect*. Although the cascade has a relatively high aspect ratio, a boundary layer develops on the sidewalls of the passage, giving rise to secondary flows that induce blockage on the flow in the mean passage where flow statistics are measured. This is shown qualitatively in Figure 5.4. The result is a spanwise contraction of the streamtube, which is not strong enough to affect large-scale flow structures in the mean passage but should nonetheless be accounted for. In the context of a compressible formulation, this effect can be analytically incorporated in the governing equations [52]. This formulation was recently introduced in the software Mu^2s^2t [15] and applied to a LPT case looking at the Reynolds sensitivity of performance indicators. In particular, the contracted streamtube induces an enhanced streamwise pressure gradient, and therefore the onset of the laminar separation and the associated transition-reattachment points are moved upstream with respect to the case without streamtube. Despite the different geometries and choice of coordinate system adopted to visualise the pressure distributions (axial vs. suction surface perimeter), this observation is consistent with the findings of Bolinches-Gisbert et al. [15]. There is no trivial approach to incorporate the streamtube formulation into the incompressible Navier-Stokes equations without violating the incompressibility constraint. Tackling this issue is outside of the scope of this work, but two avenues of further investigation can be proposed. First of all, a coordinate mapping [154] could be employed to transform a conventional geometry to mimick the streamtube contraction illustrated in Figure 5.4. However, coordinate transformations of the type $z = z(x, y)$ are currently not supported in the *Nektar++* framework with Quasi-3D simulations. Alternatively, a full 3D formulation (without mapping) could be adopted in combination with a computational domain that accounts for the predicted contraction at a given Reynolds number. A number of numerical experiments would be required to analyse various approaches to the spanwise contraction (e.g. linearly varying from LE to TE). However, the adoption of slip-wall boundary conditions that would be required in the z-direction to model the effect of the streamtube is not currently available in the *Nektar++* v 5.0 code release [118]. Both modelling strategies require further development efforts, and the former is identified as a better choice for two reasons: the computational cost of Quasi-3D setup is lower by roughly an order of magnitude compared to a 3D setup with the same number of DoF, and the use of a Fourier expansion allows for flexibility in the imposition of the spanwise contraction $z = z(x, y)$, while a full 3D approach would require a new computational mesh for every different choice.

5.2.2 Skin-friction coefficient

Several physical mechanisms can be identified as Re is increased, and the skin friction coefficient distribution shown in Figure 5.5 provides some insight. Lower Re correspond

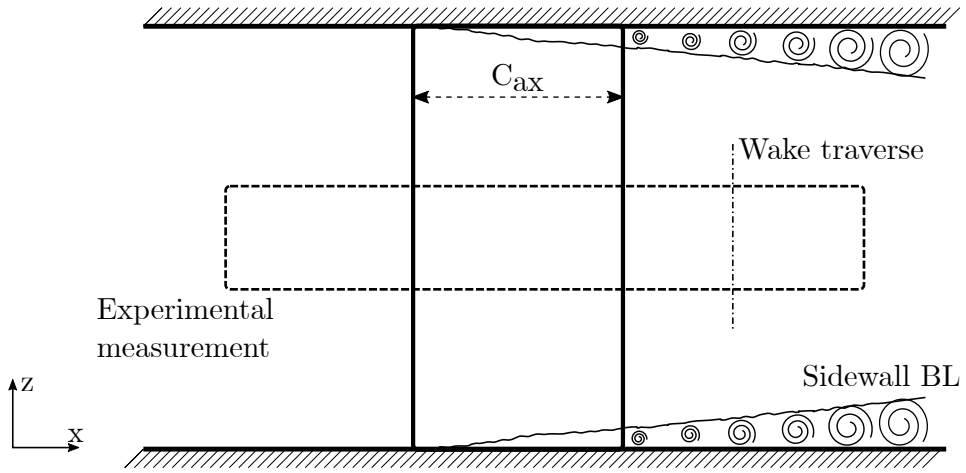

Figure 5.4: Sketch of the streamtube contraction occurring due to the presence of sidewalls in the experimental rig, with first-order repercussions on the blade wall and traverse measurements.

to increased shear stress in the front portion of the suction surface as well as the aft portion of the pressure surface. Increasing Re shortens the separation bubble by delaying the separation onset and anticipating reattachment. As visible in the zoomed inset in Figure 5.5a, the regime $Re_2 = 83K$ corresponds to an open bubble configuration regardless of the presence of momentum forcing (Figure 5.5b), which in this case delays separation by 3.9%. The transition location (identified by the negative peak in skin friction coefficient) also moves upstream as the Reynolds number is increased, and the secondary separation bubble visible at $Re_2 = 83K$ is suppressed. The case $Re_2 = 155K$ represents the critical regime where the separation bubble closes at the TE, as highlighted quantitatively in Table 5.2; the presence of the forcing promotes slightly earlier reattachment and shortening of the separation bubble from $(s/S_0)_{\text{bub}} \approx 0.29$ to $(s/S_0)_{\text{bub}} \approx 0.25$. In turn, the turbulent boundary layer at the trailing edge is not fully developed. A further increase in Reynolds number to $Re_2 = 290K$ moves the reattachment location upstream of the TE at $(s/S_0)_{\text{reat}} \approx 0.97$) (supercritical regime). The introduction of body forcing significantly shifts the reattachment location to $(s/S_0)_{\text{reat}} \approx 0.9$, shortening the separation bubble by 37.7%. Before reaching the TE, the C_f profile flattens, suggesting that a fully developed turbulent BL is established. As in the pressure distribution, the effects of body forcing are significantly evident in the final portion of the suction surface (Figure 5.5b), and higher sensitivity is captured at high Reynolds number, where the separation bubble is shrunk and moved further upstream.

Figure 5.6 shows an instantaneous carpet plot of the skin-friction coefficient on the suction surface, focusing in particular on the aft portion. The introduction of momentum forcing has an equivalent effect to introducing low levels of inflow turbulence (discussed in Chapter 4). The boundary layer distortion occurs as a consequence of low-frequency

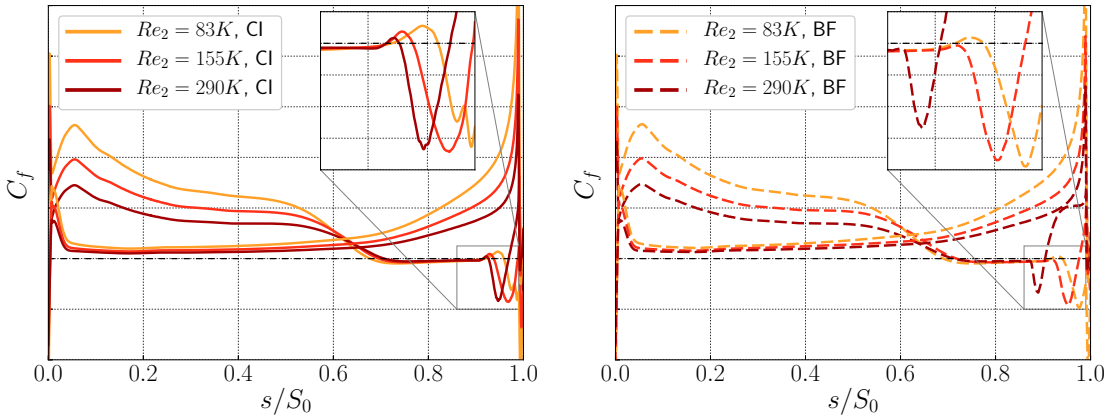

Figure 5.5: Skin friction coefficient with increasing Re_2 with (a) Clean inflow and (b) body forcing near the LE. Y-axis tick labels are omitted due to data sensitivity.

Table 5.2: Time-averaged separation bubble statistics: separation point and reattachment point. Due to data sensitivity, all values were intentionally nondimensionalised by the respective values of case BF at $Re_2 = 290K$.

Re_2	$83K$		$155K$		$290K$	
Inflow	CI	BF	CI	BF	CI	BF
$(s/S_0)^*_{\text{sep}}$	0.93166	0.96801	0.97737	1.0090	0.99412	1.0
$(s/S_0)^*_{\text{reat}}$	-	-	1.09283	1.07885	1.07155	1.0

perturbations, owing to the shear sheltering role of the boundary layer shear [63, 205], which acts as a filter. The impinging disturbances therefore result in elongated regions of spanwise-alternating high- and low-speed fluid, called streaks. Through rapid distortion theory it was demonstrated that the low-frequency disturbances that penetrate the boundary layer are not subject to the restoring action of pressure and thus futher amplify [43]. As a consequence, linear theory is not sufficient to accurately capture the evolution of streaks. If the streak amplitude grows beyond a certain treshold, secondary instability may develop and cause early breakdown, despite theoretical prediction of modal decay [5]. Accurate simulation of the secondary instability of streaks requires nonlinear analysis and has been the topic of extensive investigations in the past two decades. Providing an in-depth analysis of the transition mechanism is outside the scope of this work, which is instead concerned with motivating the effect of inflow disturbances at various regimes on the accuracy of the prediction of the main design statistics. For further details on the physics of bypass transition to turbulence the reader can refer to the review by Zaki [203].

Between $s/S_0 = 0.4$ and $s/S_0 = 0.6$, boundary layer streaks are visible in the cases with momentum forcing. Although no clear evidence of one physical mechanism in particular can be identified from a single instantaneous field, the introduction of physical disturbances causes the separation line to become non-uniform in the spanwise direction, and

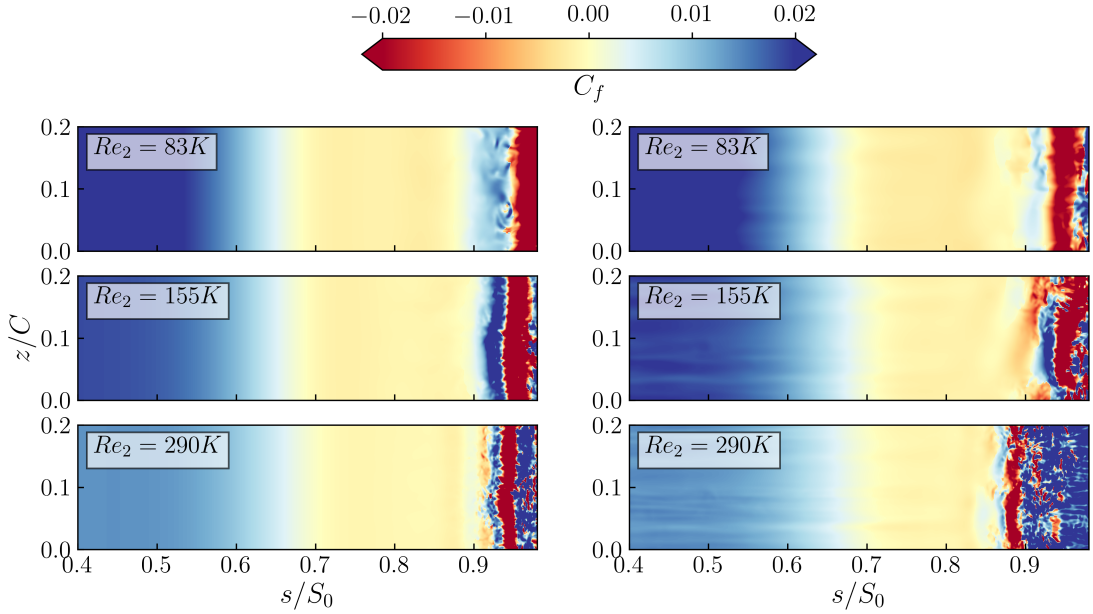

Figure 5.6: Carpet plot of instantaneous skin friction coefficient. Left: clean inflow; right: momentum forcing.

the transition and reattachment process is initiated further upstream: comparing at each Reynolds number the cases with and without inflow turbulence, the red region denoting strong negative recirculation near the TE highlights that the boundary layer distortions are sufficient to trigger anticipated transition and reattachment.

5.3 Boundary layer parameters

The integral parameters of the BL provide further indications on the effect of increasing Reynolds number and the presence of body forcing. Figure 5.7 reports momentum thickness and shape factor on the suction surface. The front portion is considered first (before peak suction, i.e. $s/S_0 < 0.6$). The introduction of body forcing does not yield appreciable differences in the evolution of θ and H in this region. This is expected, as the time-varying forcing is calibrated to be low enough to avoid introducing large-scale distortions to the incoming flow. As the Reynolds number is increased, the momentum thickness is reduced. In this region, the shape factor remains constant in all cases ($H \approx 2.5$), denoting self-similar growth of the BL profiles. At peak suction, $\theta(0.6) \approx 0.001$ at the highest Reynolds number is about half the value for $Re_2 = 86K$, the lowest Reynolds case.

Downstream of peak suction, the presence of an adverse pressure gradient is captured in both statistics. Between peak suction and the separation point (which is between $s/S_0 = 0.65$ and $s/S_0 = 0.72$ in the flow regimes analysed), the slope of momentum thickness increases with inverse proportionality to the Reynolds number, and it decreases again past the separation point.

In all cases, the shape factor highlights how the presence of inflow disturbances pro-

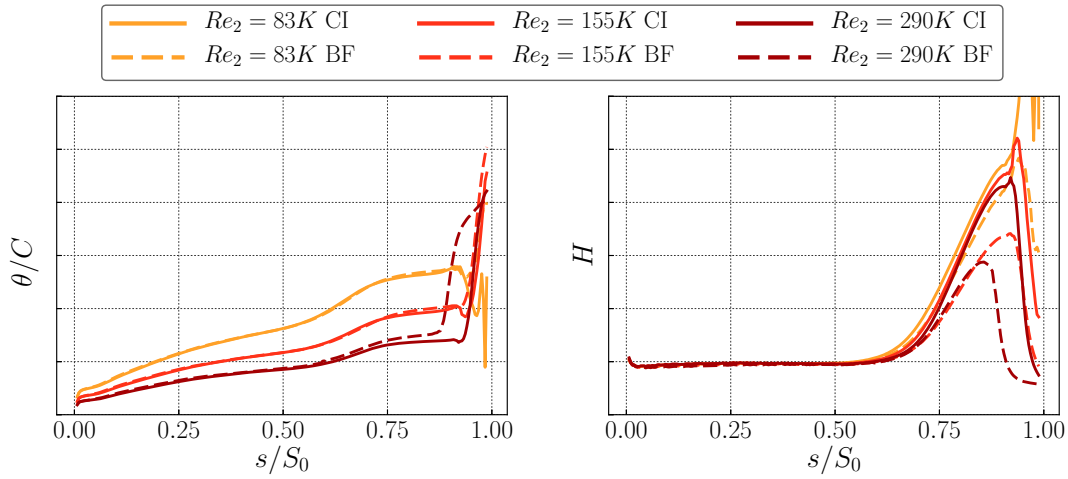

Figure 5.7: Boundary layer parameters with increasing Re_2, comparing clean inflow (CI) and momentum forcing near the LE (BF). Left: momentum thickness, right: shape factor. Y-axis tick labels are omitted due to data sensitivity.

motes a delay in the separation point, and anticipates reattachment in those cases where the separation bubble is closed. Especially at $Re_2 = 297K$ this is reflected in the anticipated rapid growth of θ in the region $s/S_0 \approx 0.9$.

5.4 Wake profiles

This section focuses on a detailed discussion of developed wake profiles, measured at $\hat{x} = 0.513$. The experimental measures were sampled on two consecutive blades and they are both reported. In order to provide a first-order quantification of the scatter of experimental data, a new parameter is introduced to estimate the RMS of the difference between the instrumented experimental traverses in a given quantity u_i. Introducing the difference between the traverses as $u_i^{\text{diff}} = ||u_i^{S1} - u_i^{S2}||$:

$$\epsilon_{u_i} = \sqrt{\frac{1}{N_y^{\text{exp}}} \sum_{i=0}^{N_y^{\text{exp}}} \left(u_i^{\text{diff}} - \langle u_i^{\text{diff}} \rangle_y\right)^2}. \tag{5.3}$$

The accuracy of the experimental measurement chain is not always trivially available. Therefore, for each parameter u_i presented, a grey shaded region is introduced showing the area within the curves $u_i(y) \pm \epsilon_{u_i}/2$, with the purpose of providing an indication of the *precision* of the experimental data. Where available, the accuracy is instead indicated by a light orange shaded region to distinguish it from the approach discussed above.

5.4.1 Velocity wakes

The velocity wake profiles are shown in Figure 5.8, also distinguishing the axial and pitchwise components. The use of pitchwise periodic boundary conditions allows to shift the

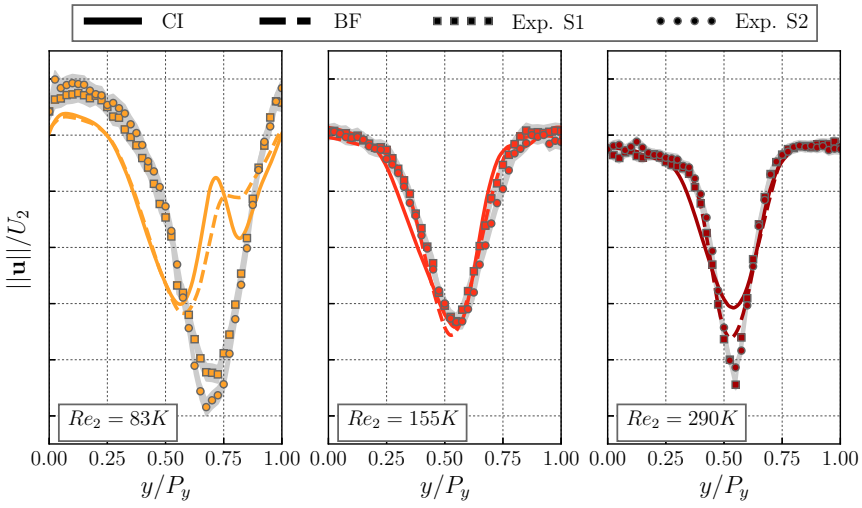

(a) Velocity magnitude

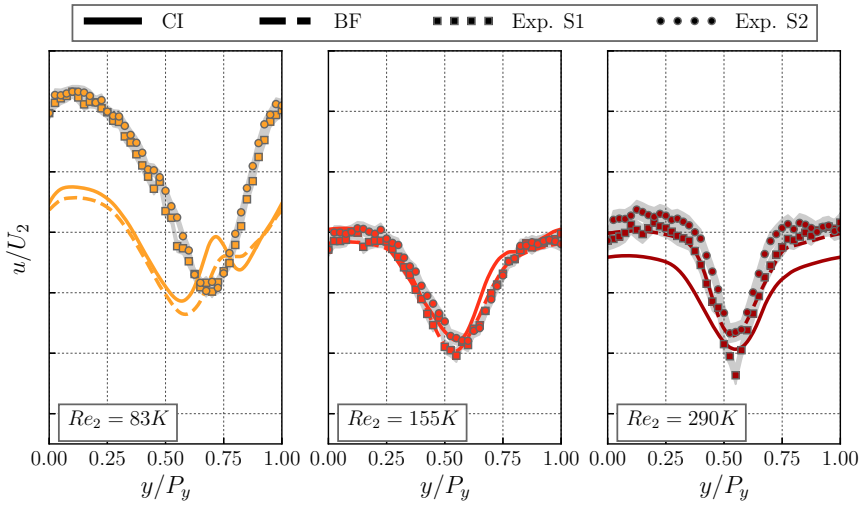

(b) Axial velocity

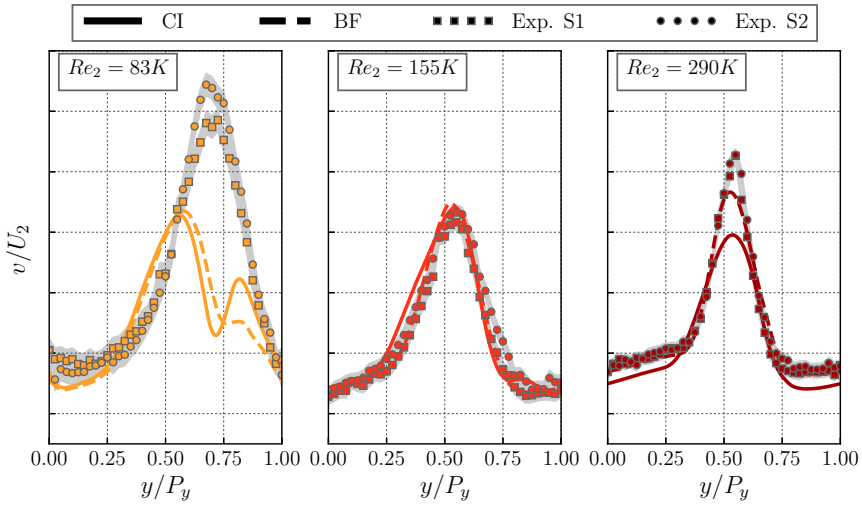

(c) Pitchwise velocity

Figure 5.8: Time-averaged velocity profiles at $\hat{x} = 0.513$. Y-axis tick labels are omitted due to data sensitivity.

computational profiles in the pitchwise direction within the figure to match the experimental result, since the location of the peak depends on the choice of the computational domain.

Clearly highlighted in Figure 5.8 is a discrepancy between numerical results and experiments at low Reynolds number. The main reason behind the discrepancy is the state of the boundary layer at the suction surface trailing edge. Blade wall measurements would provide some indication on this, but no pressure tappings are available at $Re_2 = 83K$. Therefore, at this regime it is difficult to discuss the accuracy of the numerical solver in capturing whether the experimental BL reattachment occurs prior to the TE. Both numerical results at $Re_2 = 83K$ (clean inflow and momentum forcing) show an open suction surface separation bubble, which results in a double-peak feature of the velocity profiles owing to the suction side (above $y/P_y > 0/5$), particularly pronounced in the vertical component. This effect is less intense when inflow disturbances promote anticipated transition. The large-scale mixing that occurs in the wake spreads in the pitchwise direction, therefore decreasing the axial velocity component on the pressure side as well.

This low accuracy in matching experimental measurements is further discussed in the following section, and it is only verified at subcritical Reynolds number.

In fact, at the critical regime $Re_2 = 155K$ the closest agreement with experimental data is achieved. At this Reynolds number, both clean inflow and body forcing well capture the wake flank width and negative peak, both in the axial and pitchwise directions. The effect of the inflow disturbances at this regime is limited, as previously seen in the skin friction coefficient behaviour near the TE.

At high Reynolds number, the agreement is very close especially in terms of velocity magnitude. The axial velocity component is underestimated across the pitch in the clean inflow case, which suggests that not accounting for the streamtube effect plays a role in causing this discrepancy. This topic is discussed in greater detail in Chapter 6. As highlighted in Figure 5.5, the introduction of momentum forcing promotes more significant changes in the boundary layer behaviour at this regime, compared with the lower case $Re_2 = 155K$. The increased level of turbulence intensity highlights high sensitivity of the wake profiles to anticipated reattachment of the separation bubble and consequent developmet of a turbulent boundary layer before the TE.

5.4.2 TKE wakes

The profiles of TKE are shown in Figure 5.9, rolled in the pitchwise direction in the same way as the velocity profiles. The region of most intesense turbulence activity at low Reynolds number is near $y/P_y \approx 0.75$, which corresponds to the positive peak in the velocity wakes where the discrepancy with experiments is maximum. As the Reynolds number increases, the turbulence intensity decreases significantly, but the clean inflow case consistently overestimates the TKE peak and, for $Re_2 = 155K, 290K$, the width as well. At these Reynolds numbers, the introduction of the forcing improves the agreement with

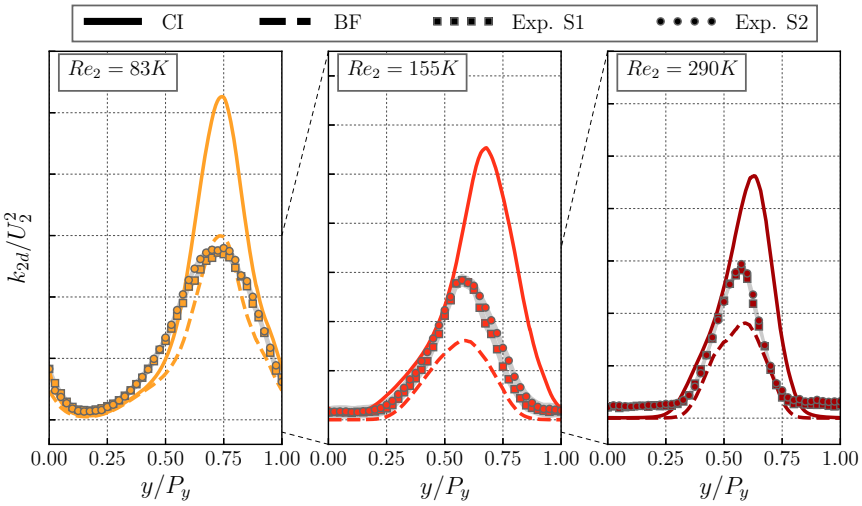

Figure 5.9: Turbulence kinetic energy profile at $\hat{x} = 0.513$. The dashed lines rescale the y-axis of the cases $Re_2 = 155K$ and $Re_2 = 290K$. Y-axis tick labels are omitted due to data sensitivity.

the wake width but underestimates the peak. Instead, at $Re_2 = 83K$ the TKE level in the peak region is close the experimental one. An even lower forcing level would likely yield more limited effects on the suction surface boundary layer, potentially providing closer agreement in the TKE traverses at moderate to high Reynolds numbers.

In order to further investigate the large discrepancy reported in the velocity wakes at low Reynolds number, Figure 5.10 focuses in detail on the qualitative topology of the wake with and without forcing, as well as the various fluctuating components. The top two subfigures (5.10a, 5.10b) show the time-averaged velocity magnitude in the near-wake region. In the absence of forcing, the two dashed lines track the development of local negative peaks in the velocity profiles, which are formed in the near TE region and captured downstream at the measurement traverse shown in Figure 5.8. When the forcing is introduced, this mechanism is strongly reduced and a single peak can be tracked.

Below, Figures 5.10c-5.10e show the three normal components of the Reynolds stress tensor with clean inflow, while the momentum forcing is reported in Figures 5.10f-5.10h. The introduction of forcing promotes three-dimensional mixing. This is reflected in the concurrently lower levels of in-plane fluctuations and higher levels of spanwise fluctuating activity. High spanwise resolution alone is not sufficient to promote mixing, which is verified by comparing an additional clean inflow simulation with only 4 planes in the spanwise direction, yielding very similar distributions of the fluctuating components to the case with $N_z = 96$. The white dashed line shown in the velocity magnitude contours is included in the $\langle w'^2 \rangle$-plots, highlighting a connection between the diffent behaviour with and without inflow turbulence.

This evidence suggests that the introduction of LE disturbances has two-fold effects on the wake statistics. On one hand, it promotes three-dimensional breakdown of turbulent structures, which results in significantly lower levels of TKE in the near- and far-wake re-

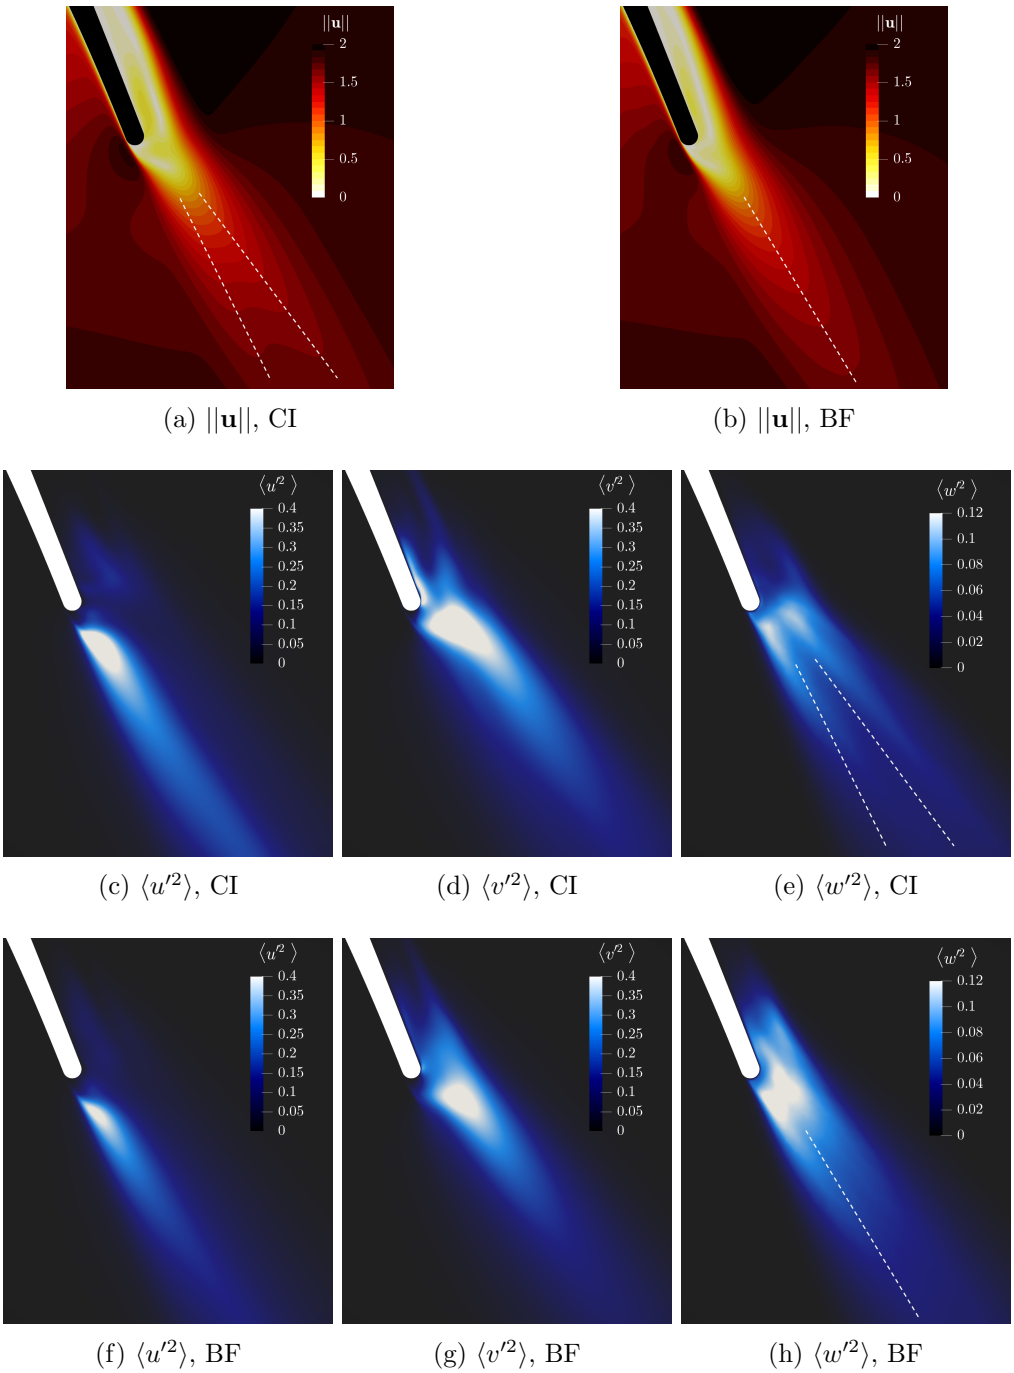

(a) $||\mathbf{u}||$, CI (b) $||\mathbf{u}||$, BF

(c) $\langle u'^2 \rangle$, CI (d) $\langle v'^2 \rangle$, CI (e) $\langle w'^2 \rangle$, CI

(f) $\langle u'^2 \rangle$, BF (g) $\langle v'^2 \rangle$, BF (h) $\langle w'^2 \rangle$, BF

Figure 5.10: Time-averaged velocity magnitude and Reynolds stresses comparing clean inflow (a, c, d, e) and body forcing (b, f, g, h) at $Re_2 = 83K$. Figures not to scale.

gions, improving the agreement with experimental measurements. However, the colourbar scale for $\langle u'^2 \rangle$ and $\langle v'^2 \rangle$ is larger than for $\langle w'^2 \rangle$: fluctuations occur mostly in the flow direction, even in the presence of inflow forcing. On the other hand, the enhanced mixing due to the forcing helps provide an explanation for the physical mechanism underlying the emergence of double negative peaks in the velocity traverses with clean inflow. In fact, in-plane fluctuations exhibit similar spatial patterns (with different intensity) regardless of the presence of the forcing. However, in the case of spanwise fluctuations, the introduction of the forcing both enhances the peak intensity and promotes a different spatial behaviour, as seen by comparing Figures 5.10e and 5.10h. In particular, CI shows a double peak in the spanwise fluctuations that is stretched in the flow direction and well aligned with the negative velocity peaks. The physical underlying mechanism is therefore driven by the presence of an open separation bubble and the absence of external disturbances: the presence of background turbulence and higher Reynolds number are both sufficient to suppress this flow feature. Another likely contributor to the discrepancy is the presence of a constraining spanwise domain in combination with low Reynolds number and clean inflow boundary conditions.

In order to verify that no effect of the SVV stabilisation strategy is responsible for an erroneous prediction of the physical behaviour in the TE region, an additional simulation was carried out with modified SVV parameters in the exponential kernel adopted for the Fourier expansion. The diffusion coefficient was reduced from $\mu_{SVV} = 1$ to $\mu_{SVV} = 0.1$, while the cutoff ratio was increased from $M = 0.5$ to $M = 0.75$. The pressure distribution and wake profiles are nearly identical to the results presented in this chapter and therefore omitted. However, this test demonstrates that the discrepancy here discussed is of physical nature rather than numerical.

5.4.3 Loss profiles and exit angle

The loss profiles are shown in Figure 5.11. Analogously to the velocity wakes, at $Re_2 = 83K$ the loss profile capture a double peak, with lower magnitude than the experimental measurements. The introduction of body forcing eliminates the double peak but still does not achieve the correct level of loss in the mid-wake region. At higher Reynolds numbers, the wake is captured well, but with an overestimation of the width. The introduction of momentum forcing allows to improve the agreement showing remarkable accuracy at $Re_2 = 155K$ and $Re_2 = 290K$ both in terms of peak and flank width.

The dependence of the level of accuracy is demonstrated to depend strongly on the Reynolds number. This behaviour remains consistent in the exit angle (Figure 5.12). At low Reynolds, both CI and BF fail to capture the correct trend and magnitude of the exit angle profile. At the higher regimes $Re_2 = 155K, 290K$ the agreement is improved, particularly with the introduction of momentum forcing at $Re_2 = 290K$. The scatter of experimental data in the measurements of α_2 is relatively high, and for $Re_2 = 155K, 290K$ BF falls within the shaded region suggesting that they are within measurement uncertainty.

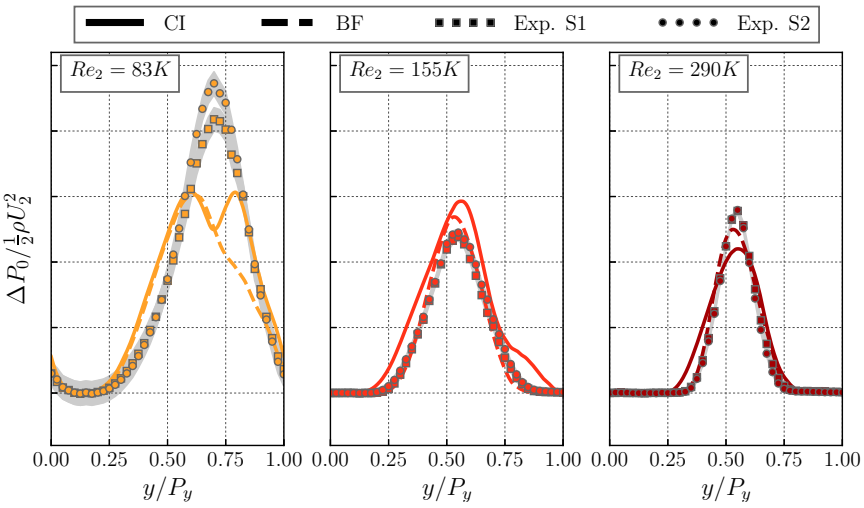

Figure 5.11: Total pressure loss profile at $\hat{x} = 0.513$. Y-axis tick labels are omitted due to data sensitivity.

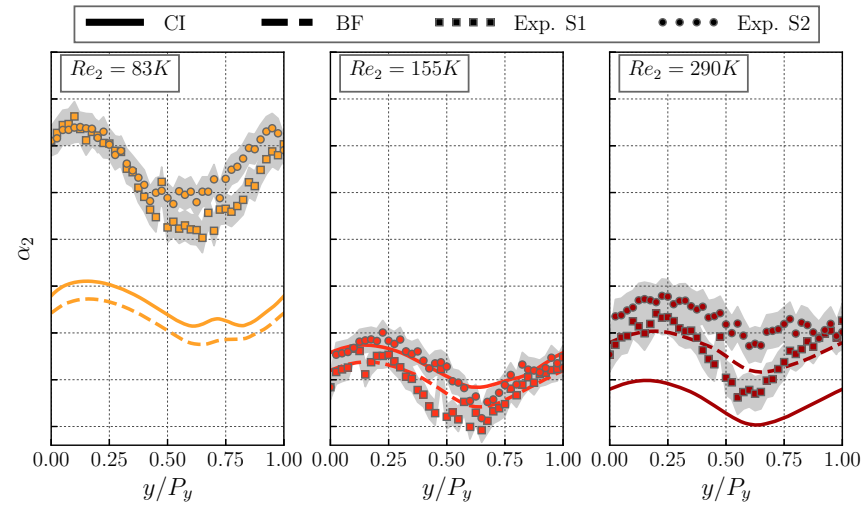

Figure 5.12: Exit angle profiles at $\hat{x} = 0.513$, Y-axis tick labels are omitted due to data sensitivity.

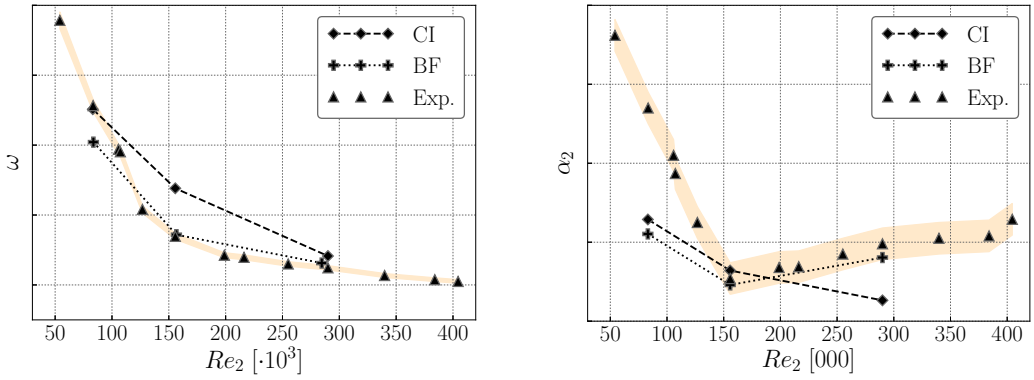

(a) (b)

Figure 5.13: (a) Mixed-out total pressure loss coefficient. The orange area indicates an uncertainty of 2.5% associated with the measurement chain. (b) Mixed-out exit angle, and highlighted uncertainty of ±0.2°. Y-axis tick labels are omitted due to data sensitivity.

5.4.4 Mixed-out measurements

Finally, the mixed-out values are reported in Figure 5.13, which shows mixed-out total pressure loss coefficient and exit angle. These parameters are of fundamental importance in the design process to ensure aerodynamic performance. At the critical regime $Re_2 \approx 150000$, a sharp change occurs in the loss trend as well as the exit angle. The state of the suction surface separation bubble is linked to different physical loss generation mechanisms. In the subcritical regime, the open separation bubble dominates the loss production mechanism by yielding a thick trailing edge boundary layer. Accurate loss estimation at this regime is challenging because it depends on the prediction of the separation onset and the thickness of the separation bubble. In this regime, the loss scaling with the Reynolds number is $\omega \propto Re_2^{-1}$ [15]. The CI case at $Re_2 = 83K$ misleadingly provides close agreement in the estimation of ω compared to experimental measurements. The wake traverses previously discussed show how the agreement of the loss profile is not close, but its mixed-out average yields the same value. At higher Reynolds, the cases with inflow disturbances agree very closely with experimental measurements both in terms of loss coefficient and exit angle, and well within experimental uncertainty. In this regime, the most relevant loss generation mechanism is laminar viscous friction and the losses are proportional to $\omega \propto Re_2^{-1/2}$ [15].

5.5 Discussion

The three flow regimes analysed ($Re_2 = 83K, 155K, 290K$) are strategically placed to test the capability and accuracy of the numerical techniques discussed in Chapter 3 over a range of Reynolds numbers, within which the numerical setup is deemed to be sufficiently resolved. The three cases analysed correspond to subcritical, critical and supercritical states of the suction surface boundary layer, with each case being analysed with clean

inflow and with momentum forcing near the LE. The flow regimes $Re_2 = 155K, 290K$ with inflow disturbances are found to provide very close agreement with experiments in terms of blade loading as well as a range of traverse measurements. Importantly, high accuracy is achieved in the prediction of mixed-out loss coefficient, within measurement uncertainty.

Momentum forcing is thus confirmed to be a useful technique to introduce controlled disturbances and trigger anticipated boundary layer transition, that has an important role in achieving the close agreement with experimental results at supercritical regime.

At subcritical regime, the presence of an open separation bubble poses a challenge. The measured wake profiles are not in close agreement with experiments and a number of factors are identified as contributors to this uncertainty. Among these, the experimental streamtube contraction is not modelled within the numerical solver. Accounting for this first-order effect is expected to modify the separation bubble length, promoting flow transition and reattachment prior to the trailing edge [15], with significant repercussions on the pressure distribution, and even more dramatic effect on wake profiles and therefore the accuracy in the prediction of total pressure losses.

The discussion on the Reynolds effect introduced in this chapter provides the grounding for Chapter 6, where the effect of wake passing is explored at the same flow regimes analysed in this Chapter.

Chapter 6

The wake passing effect

6.1 Introduction

In a real engine, the pressure expansion through the high- and low- pressure turbines is achieved in a number of subsequent stages. The interaction of multiple stages of rotors and stators is a crucial source of unsteadiness which has important repercussions on the loss production mechanisms, and it is thus of great importance to designers. Leveraging the increase in available compute resources, recent studies focused for the first time on developing an approach to support the LPT design process via high-fidelity simulations [109, 108] focusing in particular on the profile loss sensitivity to characteristic types of inflow disturbances. In the present work, building on consolidated expertise within the *Nektar++* framework and existing work in the literature, the wake passing effect is analysed in the context of the representative LPT profile discussed in Chapters 4 and 5, with particular focus on the Reynolds sensitivity of the transition mechanism on the suction surface and, mainly, the loss prediction. The optimal and most realistic combination of load and vane/blade count (i.e. degree of reaction) corresponds to the reduced frequency range where unsteady losses are highest [109]. Such realistic range of reduced frequencies is targeted in this Chapter. The availability of experimental data in the wake traverses at the flow regimes analysed allows to quantify the accuracy of the proposed numerical approach.

The Chapter is organised as follows: first, a detailed introduction on the wake modelling strategy is presented, discussing the necessary compromises required to find a tradoff between accuracy and computational cost. Subsequently, the flow dynamics on the suction surface is characterised in the various phases of the wake passing cycle, discussing the Reynolds sensitivity of key features of the separation bubble. The following sections focus on the time-averaged effect of the wake passing, analysing blade wall distributions, boundary layer statistics and finally wake profiles and mixed-out variables, establishing a comparison with highly accurate experimental data. Furthermore, an estimation of the streamtube effect is produced and a first-order correction is discussed. The final part of the loss discussion focuses on the phase-locked estimations which allow some insight

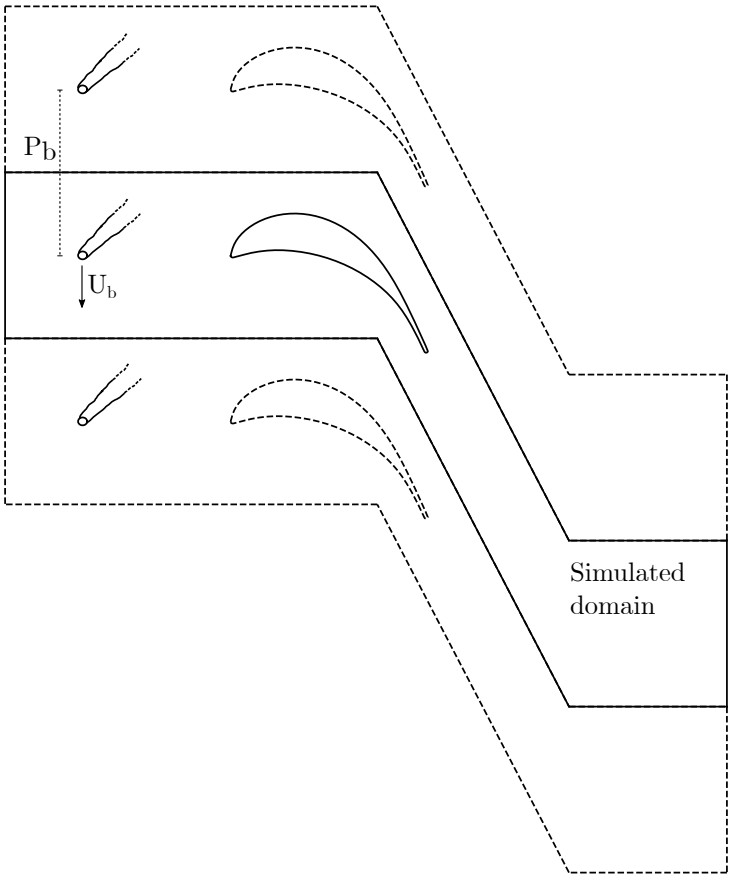

Figure 6.1: Sketch of the computational setup. Cylinders and blade outline not to scale.

into the strong modification to the loss-generating mechanisms introduced by the periodic disturbances. The Chapter finally concludes by providing some remarks and highlighting potential future research avenues to build on the contributions here presented.

6.2 Problem setup and numerical approach

6.2.1 LPT setup with wake passing

The configuration in analysis is shown schematically in Figure 6.1. The availability of experimental data motivated the choice of the flow regimes simulated. Three Reynolds numbers are considered, corresponding closely to the flow regimes discussed in Chapter 5: $Re_2 = 86K, 157K, 297K$, based on mixed-out exit velocity U_2 and suction surface perimeter S_0. The LPT cascade is scaled by the true chord and inflow velocity, so that $U_\infty = 1$. The resolution requirements dictated by the need to capture the wake passing mechanism introduced a steep increase in computational cost, thus limiting the computational domain to a single pitchwise realisation of the cascade.

When simulating the wake passing phenomenon, two nondimensional parameters are of fundamental importance: the reduced frequency F_{red}, and flow coefficient Φ, defined as:

$$F_{\text{red}} = \frac{U_b C}{P_b U_2}\,, \tag{6.1}$$

$$\Phi = \frac{U_{\infty,x}}{U_b}\,, \tag{6.2}$$

where U_b indicates the bar velocity, P_b the vertical distance between bars, U_2 the mixed-out exit velocity and $U_{\infty,x}$ the axial inlet velocity. These two parameters effectively control the bar speed (which affects the inclination of the wake when it impinges on the blade) and their relative distance. Together, these affect the frequency at which the wakes impact on the blade. Therefore, in general there are two degrees of freedom in the specifications of the bar parameters. Experimentally, this is simply obtained by adjusting the cylinders' distance and speed, allowing to analyse more realistic situations in which rotor and stator count are different. However, from a computational perspective there are further limitations due to the nature of the experiment, which requires pitchwise periodic boundary conditions. In order to simulate experiments with unequal rotor and stator count, it would be necessary to increase the number of simulated vanes, but this is typically difficult due to the resulting high computational cost. Therefore, the cascade pitch P_y must be a multiple of the distance between bars P_b.

From the available experimental data, the reduced frequency can be calculated from various quantities that are either prescribed or measured:

$$F_{\text{red}} = \frac{U_b C}{P_b U_2} = \frac{C\cos(\alpha_2)}{P_b \Phi}\,. \tag{6.3}$$

In the available data, the difference is within 0.02%, confirming the accuracy of the measurements. From a computational perspective, the control parameter in the simulations is the bar speed alone:

$$U_b^{\text{sim}} = \frac{F_{\text{red}} U_{\infty,x} P_b}{C\cos(\alpha_2)}\,. \tag{6.4}$$

In this study, matching the exact value of the reduced frequency was deemed to be necessary, while the flow parameter was enforced as a consequence of single-pitch periodicity and therefore introducing a small discrepancy. The bar speed was calculated by substituting the bar distance with cascade pitch ($P_b = P_y$):

$$U_b^{\text{sim}} = \frac{F_{\text{red}} U_{\infty,x} P_y}{C\cos(\alpha_2)}\,. \tag{6.5}$$

The numerical values imposed in the simulations are reported in Table 6.1. In the three cases analysed, the relative error between nominal flow coefficient and effective flow coefficient is 6.748%. The lower value of the computational flow coefficient means the wake inclination is slightly higher than the experimental one, introducing an element of uncertainty in the comparison.

Table 6.1: LPT bar passing setup, with cylinder parameters in the upper portion of the table. The lowest Reynolds number is simulated both with inflow wakes (IW) and inflow turbulence (IT), while other regimes analyse IW alone. The compute time is estimated on 1000 cores on the UK national supercomputer Archer, assuming linear scaling and including problem setup, checkpoints and filters, thus providing a conservative estimate.

Re_2	$86K$ (IW, IW+IT)	$157K$ (IW)	$297K$ (IW)
d_b/C	0.022973416	0.022973416	0.022973416
x_b/C	-0.656383328	-0.656383328	-0.656383328
F_{red} (C-based)	0.462034	0.464656	0.468738
F_{red} (S_0-based)	0.624132	0.627675	0.633188
$U_b^{\text{sim}}/U_\infty$	0.705339	0.706414	0.708116
Φ^{sim}	1.17731	1.17414	1.16966
$(\Phi - \Phi^{\text{sim}})/\Phi$ [%]	6.7478	6.7478	6.7478
α_1 [°]	33.86	33.96	34.08
Δt	$2.5 \cdot 10^5$	$2.5 \cdot 10^5$	$2 \cdot 10^5$
Compute time for $T = 1C/U_\infty$	8 h 40 min	8 h 40 min	10 h 45 min

The simulations were run for $12C/U_\infty$ convective timescales to allow transients out of the computational domain, and time-averaged samples were subsequently collected over $T = 24C/U_\infty$. The wake passing period is defined as $\mathcal{T}_b = P_b/U_b$. The simulation time can be rewritten as a function of wake passing periods as $t = m\mathcal{T}_b + \varphi\mathcal{T}_b$, where m is an integer representing the number of periods $\mathcal{T}_b$ elapsed, and $0 \leq \varphi \leq 1$ is the phase. Phase-locked averaging capability was incorporated in the *Nektar++* framework, and averages were sampled in 8 equispaced phases ($\Delta\varphi = 0.125$) over $m = 20$ bar passing periods, which correspond to just under $T = 24C/U_\infty$ for the specific parameters of the problem. At phase $\varphi = 0$, the bar is situated at the same pitchwise location as the leading edge, that is $y = 0$.

The inflow wakes are modelled by means of the Smoothed Profile Method introduced in Section 2.4. The formulation of the time integration scheme is a modified version of the VCS (described in Section 2.4.1) which requires solving the Poisson problem for pressure twice. This approach introduces a significant increase in the computational cost, for algorithmic reasons as well as due to the increased resolution requirements to accurately model the cylinders. The limited computational resources available constrained the numerical setup for the LPT cascade, compared to the standard maintained through previous Chapters: the simulations were run at $P = 6$ and with a spanwise resolution of $N_z = 72$, combined with a particularly refined base mesh shown in Figure 6.2. This setup results in 52.67 M DoF. Despite the reduced resolution, Figure 6.3 shows that the resolution requirements adopted through this work are still respected along the suction surface, where the wake passing impact is most important.

The spanwise domain was modified from the value of $L_z = 0.2C$ adopted in Chapter 5 to $L_z = P_y/4 = 0.21164C$. Introducing a geometrical configuration where the pitchwise

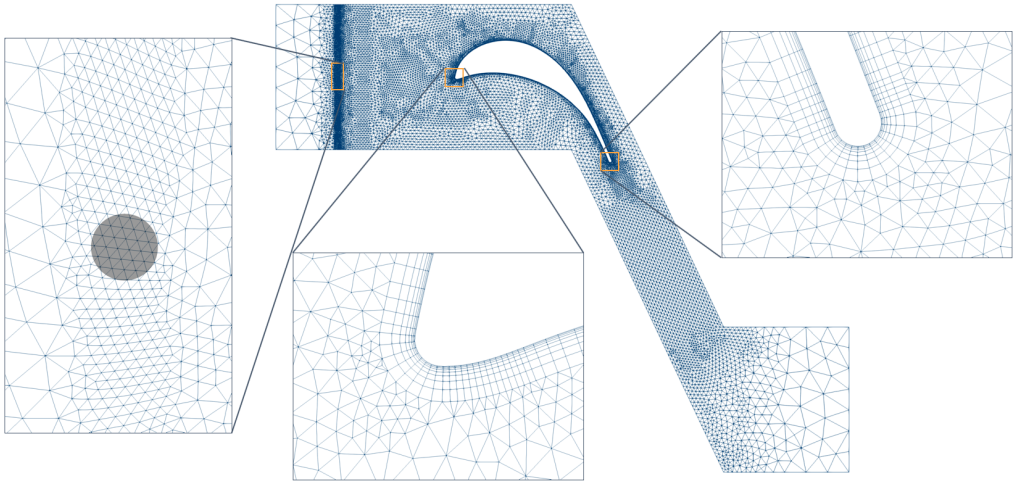

Figure 6.2: Base computational mesh, with 3191 quadrilateral elements in the O-mesh and 20541 triangular elements. The zoomed insets show details on the wake passing region, as well as the LE and TE. Figure not to scale.

domain is a multiple of the spanwise domain allowed to adopt the modified version of Davidson's inflow turbulence generation algorithm [35] introduced in Chapter 2, where only the modes that can be exactly contained in the computational domain were selected, thereby enforcing spanwise and pitchwise periodicity in the synthetic velocity signal generated. The low Reynolds number case $Re_2 = 86K$ is anaysed with and without inflow turbulence (on top of the discrete disturbances shed by passing bars). This simulates the flow environment in the back rows of the turbine, characterised by strong background turbulence built up by the upstream airfoils. The inflow turbulence case required the generation of an additional computational mesh to resolve the turbulent structures introduced at the inlet, and is denoted with IT. This mesh is omitted and it differs from Figure 6.2 only in the region comprised between the inflow and the vertically refined line where the bar passing occurs. A nominal turbulence intensity TI = 3.5% is prescribed in combination with a length scale $L_t = 0.05C$ to guarantee domain independence [134]. The algorithm selected 30 modes to discretise the modified von Kármán spectrum, compared to 1024 specified in Chapter 4. As a consequence of a coarser discretisation of the wavenumber space, the level of turbulence intensity effectively introduced at the inlet of the computational domain is measurably lower than prescribed: TI = 2.578%.

6.2.2 Numerical setup of SPM to model the bar passing effect

Adopting the Smoothed Profile Method as a way to generate inflow disturbances in the form of periodic passing wakes, the computational mesh cannot ensure a local resolution level sufficient to fully capture the cylinder's dynamics at realistic flow regimes. Multiple reasons can be identified. First of all, since the cylinders move across the whole pitchwise domain, the entire stripe extending across the domain would require extremely high

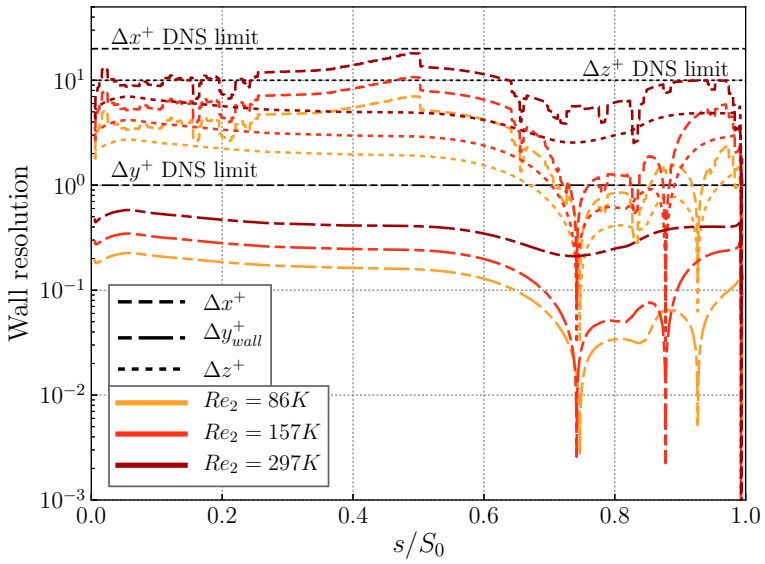

Figure 6.3: Wall resolution along the suction surface with increasing Re_2 and incoming wakes.

resolution. For example, the adoption of a high-resolution configuration similar to the one shown in mesh M3 of Chapter 2 (Figure 2.14) would require $\approx$ 90000 elements just to capture the motion of the cylinder, equivalent to over 180.746M DoF per variable. Beyond the increase in computational cost, the element size (with element side h approximately 6 times smaller than the smallest elements in the TE region) would introduce a significant CFL barrier, therefore effectively resulting in a further increase of the computational cost to achieve the target time integration extent. However, an extremely high-resolution is likely not necessary in this region, as the cylinder wake impacts on the turbine blade in the very far wake. Given the parameters of the problem and accounting for the correct wake inclination, the cylinder wake reaches the leading edge of the turbine around 52 diameters downstream. Therefore, any computational effort to capture the correct dynamics around the moving cylinders would require being supported for a considerable distance (in the cylinder's frame of reference). On top of this, due to the shear sheltering mechanism [69, 63, 202] the high frequency perturbations do not penetrate the boundary layer, and therefore the added computational cost would actually not produce justifiable improvement in the results.

Therefore, the approach adopted is driven by practicality requirements, and it is aimed at ensuring that the wake width and spectral content at low frequencies (i.e. vortex shedding) is preserved. The element size in the SPM region should *not* be smaller than the smallest elements in the trailing edge region, sufficiently far from the blade (where the CFL number is measured to be highest). This guarantees that no further timestepping restrictions are imposed by the introduction of the immersed boundaries. The computational mesh produced by imposing this constraint is such that in every pitchwise location the region within the cylinder is represented with around 300 quadrature points, and the

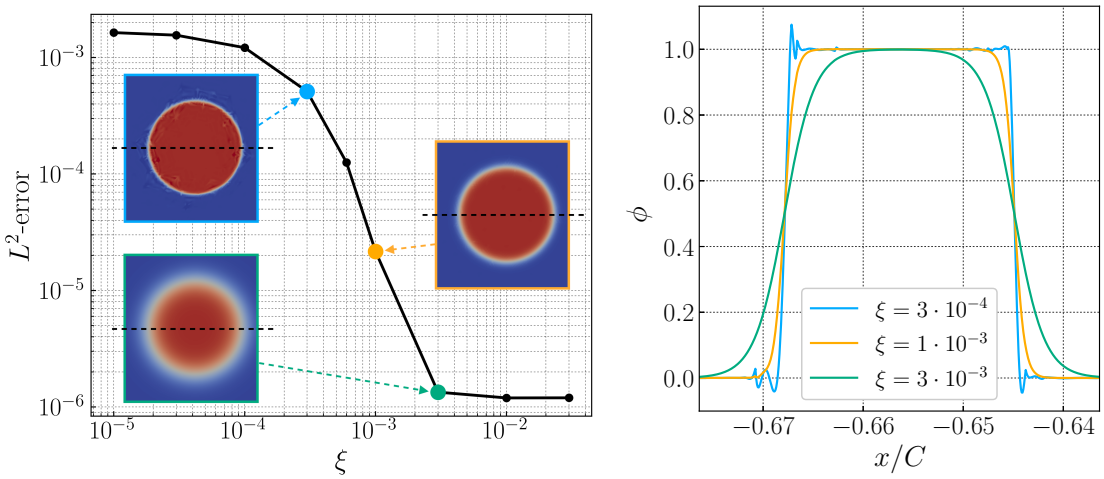

(a) (b)

Figure 6.4: (a) L^2-error of the projection of the ϕ function onto the LPT mesh, with varying interface thickness parameter ξ. The contour of the ϕ profile shown for three examples of ξ is extracted in correspondence to the dashed black lines and reported in (b).

relative resolution in the cylinder's frame of reference is analogous to that of mesh M1 in Figure 2.14.

Since the choice of resolution is driven by the setup of the problem, this poses the challenge of how to best select the interface thickness parameter ξ and the diameter of the immersed boundary. Figure 6.4 shows the evolution of the L2 error of the projection of the ϕ function describing the immersed boundary onto the discrete computational mesh employed for the LPT simulations, at order $P = 6$. Lower interface thickness corresponds to sharper definition of the body, but also higher error. The error itself tends to saturation for $\xi < 1 \cdot 10^{-4}$ and $\xi > 3 \cdot 10^{-3}$. The effect of ξ can be visually appreciated in the three contours of Figure 6.4a. The ϕ profiles extracted along the black dashed lines are shown in Figure 6.4b, which highlights that the effect of under-resolution translates into wiggles in the SPM profile, in case $\xi = 3 \cdot 10^{-4}$. Case $\xi = 1 \cdot 10^{-3}$ is the smallest choice of interface thickness that accurately represents the shape of the smoothed profile, and therefore $\xi = 1 \cdot 10^{-3}$ is selected for the LPT test case. However, the effect of employing a relatively large interface thickness is that of increasing the effective diameter of the cylinder. Therefore, a series of preliminary numerical experiments were carried out to ensure the generation of realistic wake dynamics.

Auxiliary cylinder simulations

For each of the Reynolds numbers considered in the LPT cascade, a series of auxiliary simulations was targeted at observing the effect of diameter variations of the SPM profiles at *fixed* interface thickness and resolution, analysing the mean wake profiles and the velocity spectra and comparing them with results obtained in a corresponding DNS simulation.

Table 6.2: Parameters employed for preliminary cylinder simulations.

Re_2	Re_1	$\nu = \frac{S}{Re_1}$	U_c	Re_c
86000	48000	$2.8142 \cdot 10^{-5}$	1.51112003	1233.56734
157000	84000	$1.6081 \cdot 10^{-5}$	1.51269516	2160.99304
297000	166000	$8.1376 \cdot 10^{-6}$	1.51493002	4276.84316

The auxiliary simulations (both DNS and SPM) were conducted on a stationary cylinder test case with unitary diameter, with domain analogous to that adopted for the SPM validation (Section 2.4.2). The DNS was carried out with the same computational setup described in Section 2.4.2, at polynomial order $P = 7$ and $N_z = 64$, resulting in 15.16M DoF per variable. Instead, the SPM setup was tested with a mesh size analogous to that of the LPT (scaled appropriately by the bars' diameter and relative velocity in the bars' moving frame of reference) in order to mimic the effect of different SPM setups in the cascade simulations at reduced computational cost, necessary to perform a parameter sweep. Importantly, in order to fully mimic the resolution of the LPT case, the spanwise domain of the SPM cases was scaled so that $L_z^c = L_z/d = 9.191667$, while maintaing the same number of spanwise planes as in the LPT case: $N_z = 72$. The resulting setup has 3.12M DoF per variable. The interface thickness $\xi = 1 \cdot 10^{-3}$ was scaled appropriately, yielding $\xi = 0.0435286$.

Table 6.2 reports (in the LPT setup) the parameter scaling required to calculate Re_c. U_c is the true speed in the cylinder's frame of reference, calculated from the triangle constructed from the inflow velocity $U_\infty = 1$ and the translational velocity of the bars. These parameters were dictated by the experiments, reported in Table 6.1. The cylinder test case is simulated at the flow regimes indicated by Re_c.

For each Re_c, five diameter sizes were tested ranging from $d_\mathrm{SPM} = 1$ to $d_\mathrm{SPM} = 0.6$, for a total of 3 DNS cases and 15 SPM simulations. Analogously to the cylinder simulations of Section 2.4.2, all cases were averaged over $T = 400d/U_\infty$, which corresponds to $\simeq 83$ vortex shedding cycles (assuming a Strouhal number St = 0.208). It should be highlighted that the resolution of the SPM cases is extremely coarse, and therefore differences are expected in the comparison with DNS results.

Figure 6.5 shows the mean streamwise velocity wake in the three Re_c cases, extracted at three streamwise locations. At $Re_\mathrm{c} = 1233, 2161$, in the near wake region the nominal diameter $d_\mathrm{SPM} = 1$ produces a wider profile compared to DNS, while the wake is accurately reproduced with a diameter of $d_\mathrm{SPM} = 0.6$. This highlights the effect of under-resolution in the SPM region which yields a higher effective diameter. In the $Re_\mathrm{c} = 4277$ case the negative peak in the near-wake is captured well by the $d_\mathrm{SPM} = 1$ case, which overestimates the wake width significantly. Smaller diameters improve on the accuracy of the wake flanks. For the application of interest, the behaviour in the under-resolved far wake is particularly important. Introducing the adimensional streamwise coordinate $x^* = x/d$, at $x^* = 10$ (shown at the bottom of Figure 6.5) a decreasing SPM size corresponds to

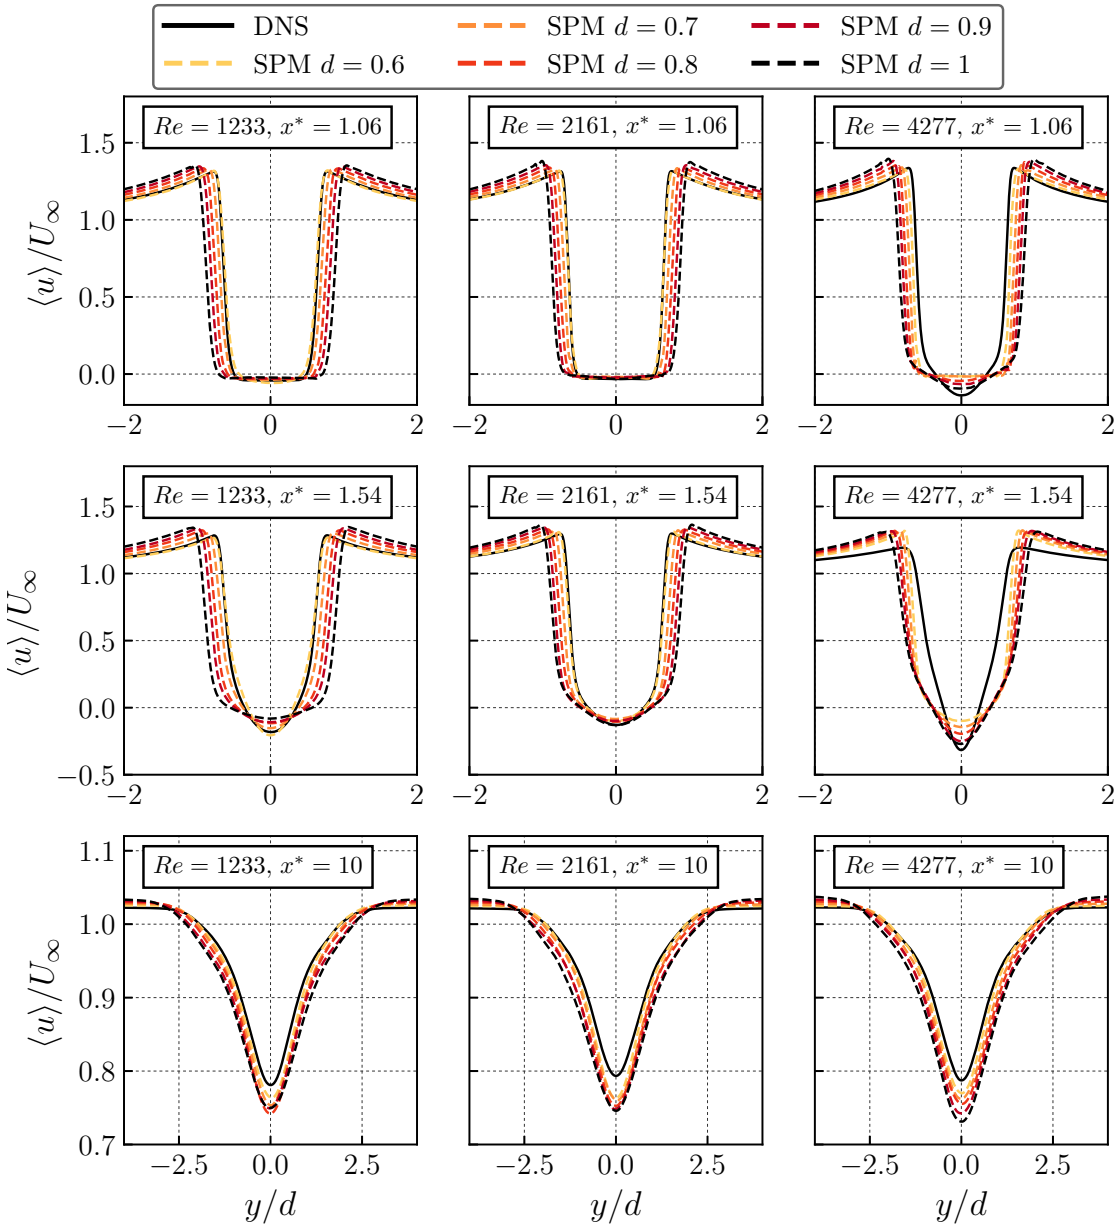

Figure 6.5: Time-averaged streamwise velocity, comparing (left to right) three increasing $Re_c = 1233, 2161, 4277$ and (top to bottom) three streamwise stations $x^* = 1.06, 1.54, 10$, with the centre of the cylinder at $x^* = 0$. The streamwise stations are sampled assuming $d = 1$ in all cases. The black solid line indicate DNS results, while dashed lines are SPM simulations with decreasing diameter corresponding to lighter colour.

improved agreement with the negative wake peak (which is overpredicted with $d_{\mathrm{SPM}} = 1$), as well as the wake flanks and the progressively undisturbed velocity away from the body ($|y/d| > 3$). Overall, the case $d_{\mathrm{SPM}} = 0.6$ produced the best agreement with the DNS in all the streamwise velocity far wake comparisons.

To obtain a clearer picture of the wake behaviour, the intermediate location $x^* = 5$ is considered to further discuss other types of wake statistics, shown in Figure 6.6. The top row shows the cross-flow velocity wake. In all Re_{c} cases, the nominal value $d_{\mathrm{SPM}} = 1$ fails to capture the correct peak intensity, which is overpredicted by a factor larger than two. Decreasing values of d_{SPM} progressively improve the agreement with DNS data, which is extremely close in case $d_{\mathrm{SPM}} = 0.6$ at Reynolds numbers $Re_{\mathrm{c}} = 1233, 2161$. The highest Reynolds case $Re_{\mathrm{c}} = 4277$ is less sensitive to changes in the SPM diameter. The relevant components of the Reynolds stress tensor are reported below the vertical wake statistics in Figure 6.6. At the lower $Re_{\mathrm{c}} = 1233, 2161$ cases the same trend was found to occur for $\langle u'^2 \rangle$, $\langle v'^2 \rangle$ and $\langle u'v' \rangle$, with $d_{\mathrm{SPM}} = 0.6$ producing extremely accurate agreement with DNS data, both in terms of peak intensity and flank width. This is consistent with previous findings on the mean streamwise velocity wakes. At the higher regime $Re_{\mathrm{c}} = 4277$, the normal stresses achieve an improvement similar to the lower Reynolds cases, while in the Reynolds stress wake $\langle u'v' \rangle$, the improvement derived from decreasing the SPM diameter is less drastic, analogously to mean velocity statistics.

The streamwise velocity measured along the centreline is shown in the left column of Figure 6.7. In all cases the nominal diameter $d_{\mathrm{SPM}} = 1$ results in an overprediction of the separation bubble length, as well as an underprediction of the streamwise velocity in the far wake. At $Re_{\mathrm{c}} = 1233$, decreasing the SPM diameter allows for significant improvement with DNS results, but no diameter exactly captures the correct maximum recirculation velocity (negative peak) and separation length. Overall, the configuration $d_{\mathrm{SPM}} = 0.6$ produces the best agreement in the near and far wake. The same observations apply to $Re_{\mathrm{c}} = 2161$, but in this case the improvement is not as significant, with all SPM diameters producing close agreement with DNS data. At higher $Re_{\mathrm{c}} = 4277$, the trend in the recirculation region is inverted with respect to $Re_{\mathrm{c}} = 1233$: a reduction in SPM diameter results in marginal extension of the recirculation length. However, the far wake is still better predicted by smaller SPM diameters, with the shift occurring at $x^* \approx 5$. All in all, despite opposite trends observed in the recirculation region, the far wake effect across the range of Reynolds numbers in analysis is that of marginal improvements in the streamwise velocity with smaller diameter, supporting the findings previously discussed on mean velocity wakes and turbulence quantities.

Lastly, the velocity spectra in the near wake are shown in the right column of Figure 6.7. The behaviour is extremely similar in all flow regimes. At low frequency, the vortex shedding frequency corresponds to the dominant peak and is shown in the inset plot. At all Re_{c}, the diameter $d_{\mathrm{SPM}} = 0.6$ captures the vortex shedding frequency very accurately compared to the DNS case: the curves are nearly overlapping. As highlighted in the case of streamwise velocity wakes, larger SPM diameters produce a larger effective diameter,

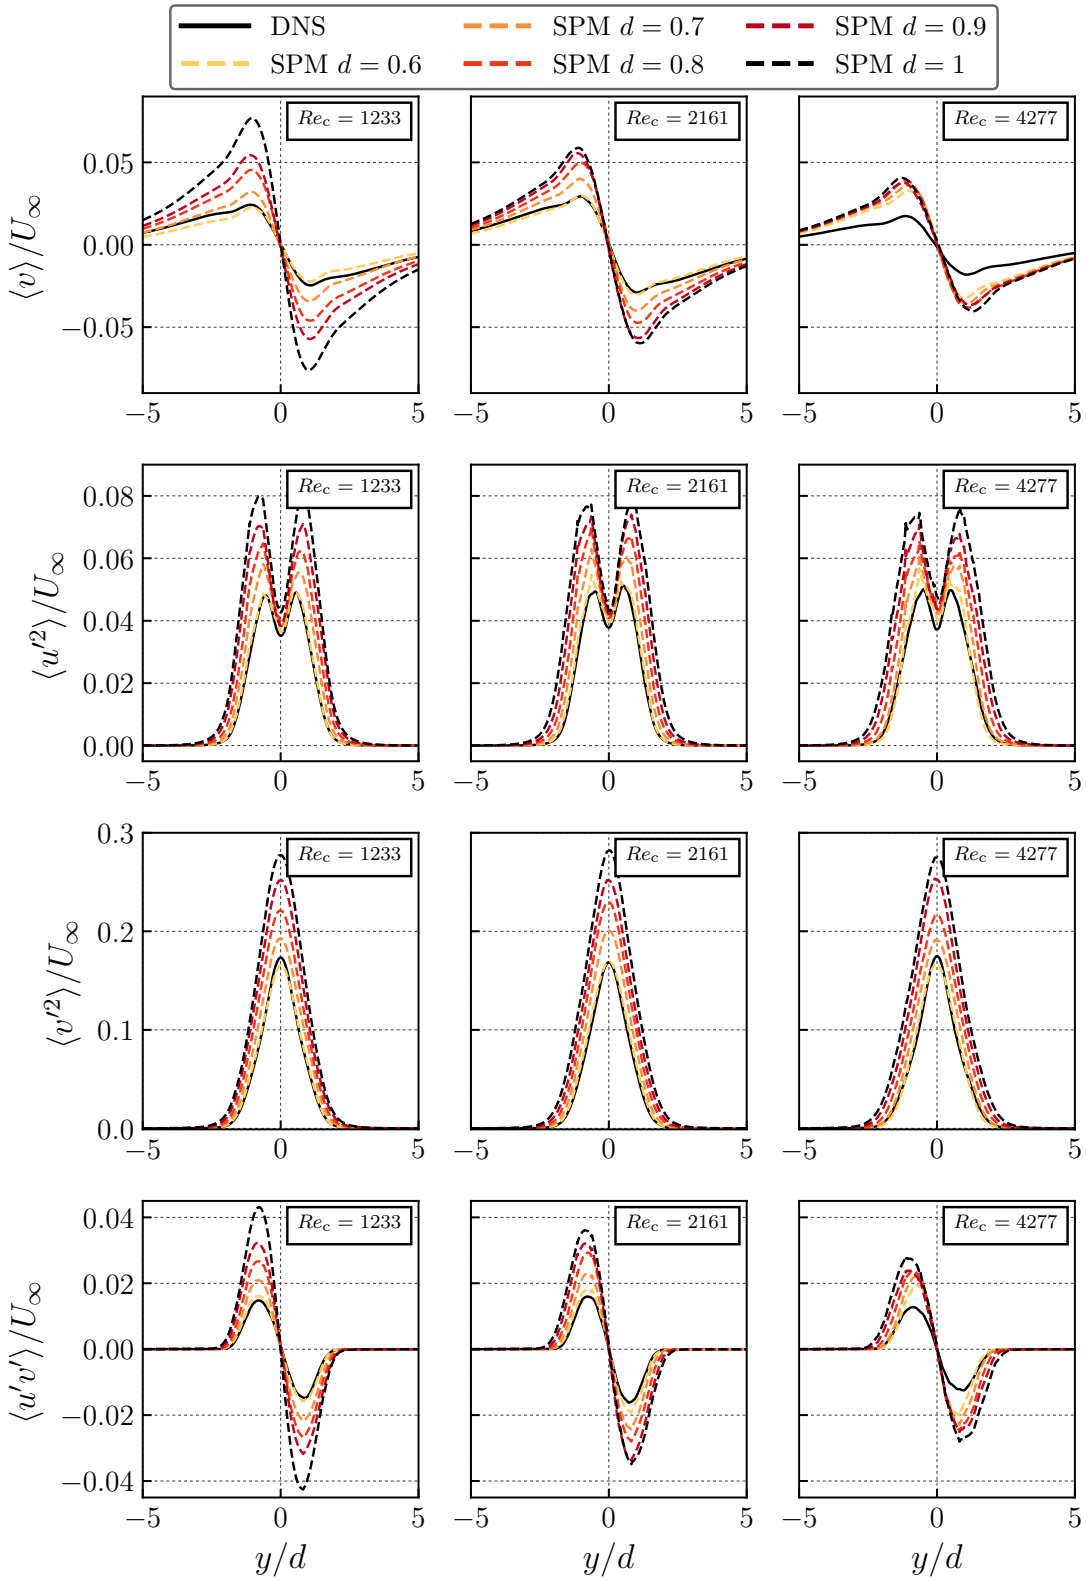

Figure 6.6: Time-averaged wake profiles sampled at $x^* = 5$ comparing (left to right) three increasing $Re_c = 1233, 2161, 4277$. From top to bottom: cross-stream velocity, streamwise fluctuating velocity, cross-stream fluctuating velocity and Reynolds shear stress. The black solid line indicate DNS results, while dashed lines are SPM simulations with decreasing diameter corresponding to lighter colour.

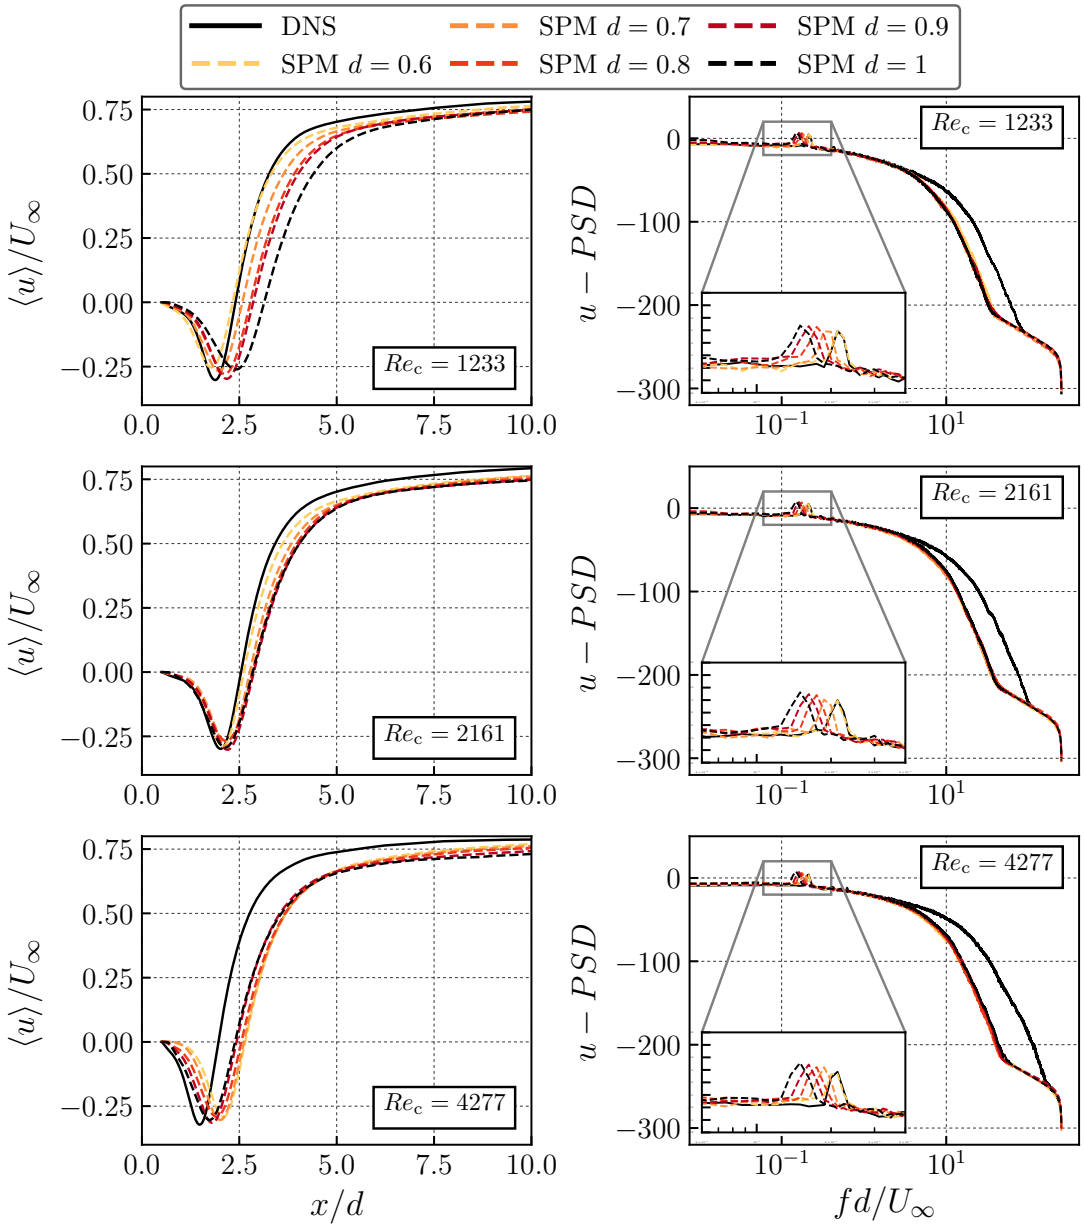

Figure 6.7: Left: time-averaged streamwise velocity along the wake centreline. Right: streamwise velocity spectra at $x/d = 2.57$, $y/d = 0.52$. From top to bottom, three increasing $Re_c = 1233, 2161, 4277$.

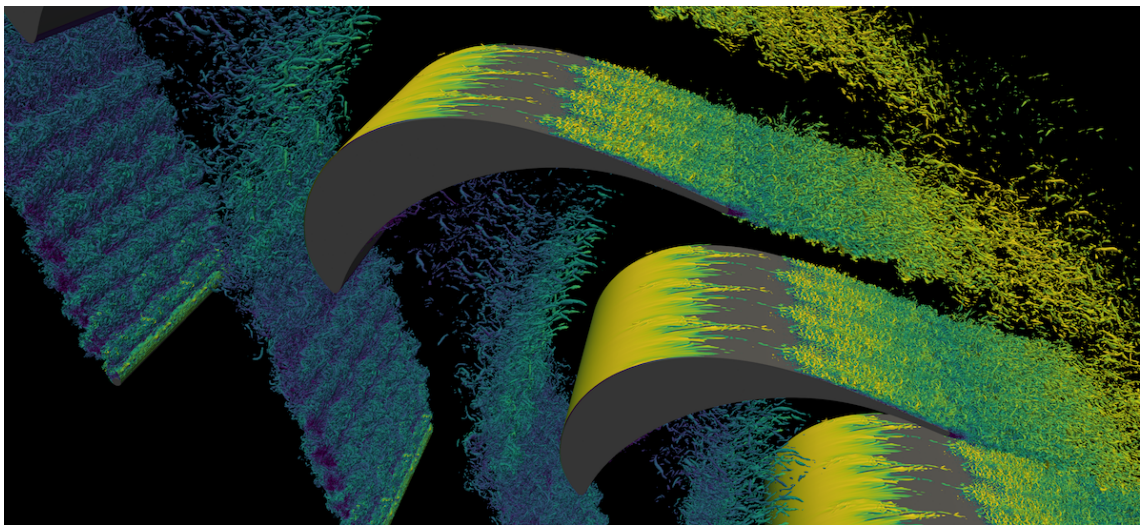

Figure 6.8: Isosurfaces of $Q = 200$ contoured by velocity magntitude in case $Re_2 = 297K$. The computational domain is replicated in the spanwise and pitchwise directions for graphical purposes.

thus reducing the frequency sampled. At higher frequency, the velocity spectra in the SPM cases are overlapped, and differ from DNS spectra for $f \geq 5U_\infty/C$ because of the significant difference in mesh resolution: the SPM mesh is extremely coarse and therefore high frequencies are strongly damped.

Considering the various combinations analysed, both in terms of various Re_c and SPM diameters, it appears that with the fixed interface thickness parameter $\xi = 1 \cdot 10^{-3}$ the SPM diameter $d_{\text{SPM}} = 0.6$ produces the best agreement with DNS results. Therefore, the nominal bar diameter in the LPT simulations was scaled by 60%. The improvement arising in the near wake from smaller diameters is more pronounced at low Reynolds number, while less sensitivity was measured at $Re_c = 4277$. However, consistently improved results are observed in the far wake region.

Overall, this section demonstrated the validity of the approach adopted for simulating the wake passing effect in the LPT test case, and in particular how the SPM parameters were tuned to guarantee the generation of a realistic cylinder wake, without additional compromises in the computational cost of the LPT simulations owing to further timestep restrictions and additional mesh refinement.

Practical implementation of SPM with pitchwise-periodic simulations

The setup of the SPM ϕ function in the case of vertically translating bodies in a pitchwise-periodic simulation requires to be prescribed in a fictitious domain three times as large as the actual domain. The simple reason is highlighted in Figure 6.1. Since the analytical function that describes the cylinders refers to the centre of the cylinder, by prescribing a single domain when the shape shifts out of the bottom and is moved to the top of the domain, the top portion of the semicircle is removed and inserted instantaneously in the

top domain. This requires the introduction of a fictitious domain of height P_y above and below, to ensure continuity in the translational motion of the SPM profiles.

The qualitative result of the implementation discussed above is shown in Figure 6.8 at high Reynolds number, focusing on the suction surface of the LPT in analysis. The initial stage of wake bending is visible in the Figure, while the suction surface transition leads to fully turbulent boundary layer at the trailing edge. More details on the flow features and the transition mechanism are discussed in the following section.

6.3 The wake passing effect on the suction surface

The time averaging operation masks the wealth of flow phenomena occurring in the boundary layer of LPTs subject to incoming disturbances, and it particularly fails to highlight the presence of high-amplitude events that precede the onset of turbulence. For this reason, a discussion on some instantaneous flow statistics is presented before an in-depth analysis of time-averaged statistics, to highlight some of the flow phenomena occurring on the suction surface of the cascade at the various stages of the wake passing cycle.

6.3.1 Evidence of the transition mechanism

Figures 6.9, 6.10 and 6.11 provide qualitative insight into the suction surface dynamics at consecutive phases φ for $Re_2 = 86K, 157K, 297K$ respectively. Case $Re_2 = 86K$ is considered first (Figure 6.9a). The incoming wakes impinge on the suction surface separation bubble between $\varphi = 0.25$ and $\varphi = 0.375$. The extended region of weakly negative recirculation (the pale yellow region) visible at $\varphi = 0, 0.125$ between $s/S_0 = 0.7 - 0.95$ is largely suppressed by the incoming disturbance. Concurrently, at $\varphi = 0.25$ the shear layer rolls up into two separate regions of recirculating flow, as previously highlighted in a number of studies [111, 189, 108]. The upstream recirculation region is found at $s/S_0 \approx 0.85$, while the second merges into the trailing edge. This mechanism can also be identified at $Re_2 = 157K$ at $\varphi = 0.375$ but not at the higher regime $Re_2 = 297K$.

As the wake impacts on the separation bubble, the region of high shear in the upstream portion of the suction surface (the dark blue region) moves towards the trailing edge, following the impinging disturbance. At $\varphi \approx 0.375$, the attached flow region develops a spanwise pattern of alternating high- and low-speed flow, called streaks, which play an important role in the bypass transition mechanism [68, 203]. Streaks are found at all the flow regimes analysed, and their spanwise length scale is inversely proportional to the Reynolds number. The generation mechanism of elongated streaks in the suction surface in presence of inflow turbulence was recently discussed in detail by Zhao et al. [207]. Their inception is linked to large magnitude of the spanwise fluctuating velocity component measured at the leading edge of the blade. The tangential vorticity ω_t (aligned with the tangent to the blade wall) can be related to the spanwise fluctuating velocity as $\omega_t' \sim \partial w'/\partial n$. As the streamwise-aligned vortices are stretched around the high-curvature

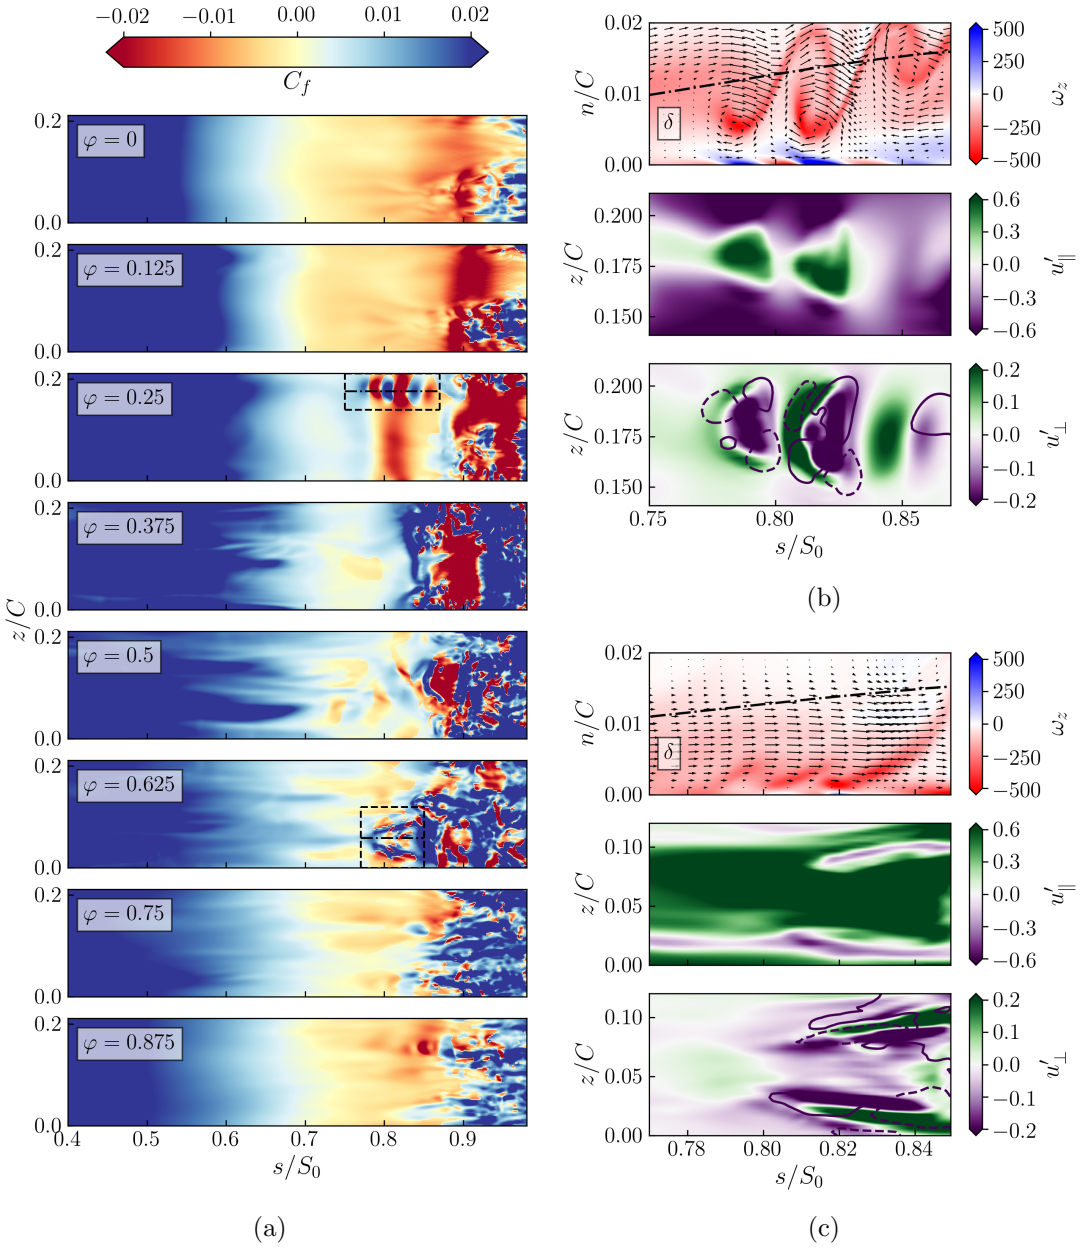

Figure 6.9: Instantaneous suction surface statistics for $Re_2 = 86K$. (a) Instantaneous skin-friction coefficient carpet plot on the aft portion of the suction surface, at 8 different phases. The dashed boxes in phases $\varphi = 0.25$ and $\varphi = 0.625$ are shown in detail in (b) and (c), respectively. For (b) and (c) the top figure shows spanwise vorticity in the blade-normal plane denoted with dash-dotted lines in (a), superimposed with fluctuating velocity vectors. The middle and bottom figures show respectively wall-parallel and wall-normal fluctuating velocity $n/C = 0.01$ away from the wall. The solid and dashed lines in the bottom figures in (b) and (c) are isolines of $w' = \pm 0.15$.

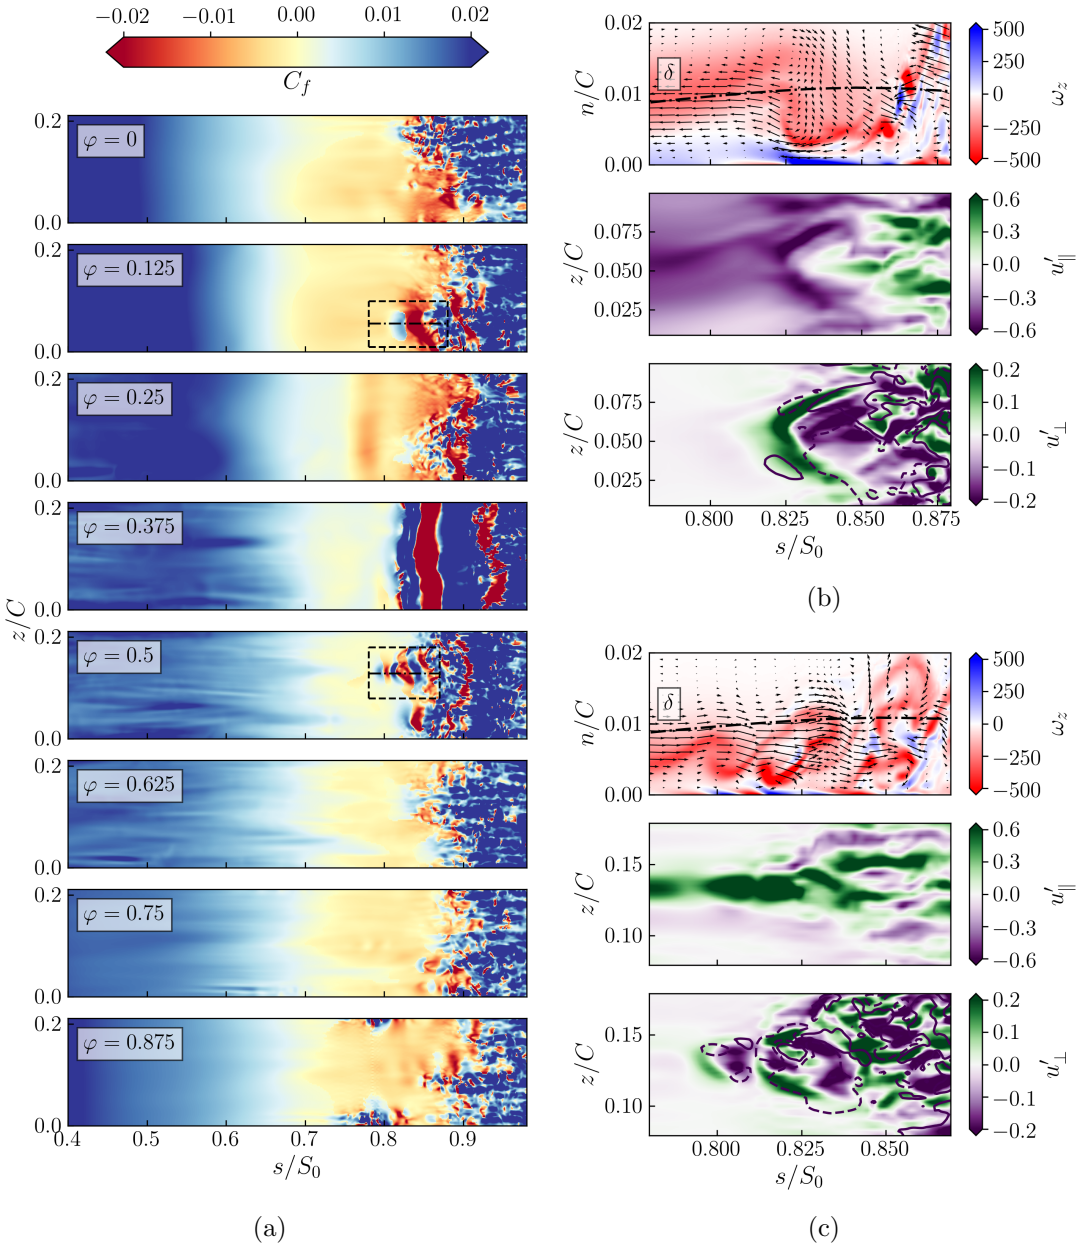

Figure 6.10: Instantaneous suction surface statistics for $Re_2 = 157K$. (a) Instantaneous skin-friction coefficient carpet plot on the aft portion of the suction surface, at 8 different phases. The dashed boxes in phases $\varphi = 0.125$ and $\varphi = 0.5$ are shown in detail in (b) and (c), respectively. For (b) and (c) the top figure shows spanwise vorticity in the blade-normal plane denoted with dash-dotted lines in (a), superimposed with fluctuating velocity vectors. The middle and bottom figures show respectively wall-parallel and wall-normal fluctuating velocity $n/C = 0.005$ away from the wall. The solid and dashed lines in the bottom figures in (b) and (c) are isolines of $w' = \pm 0.15$.

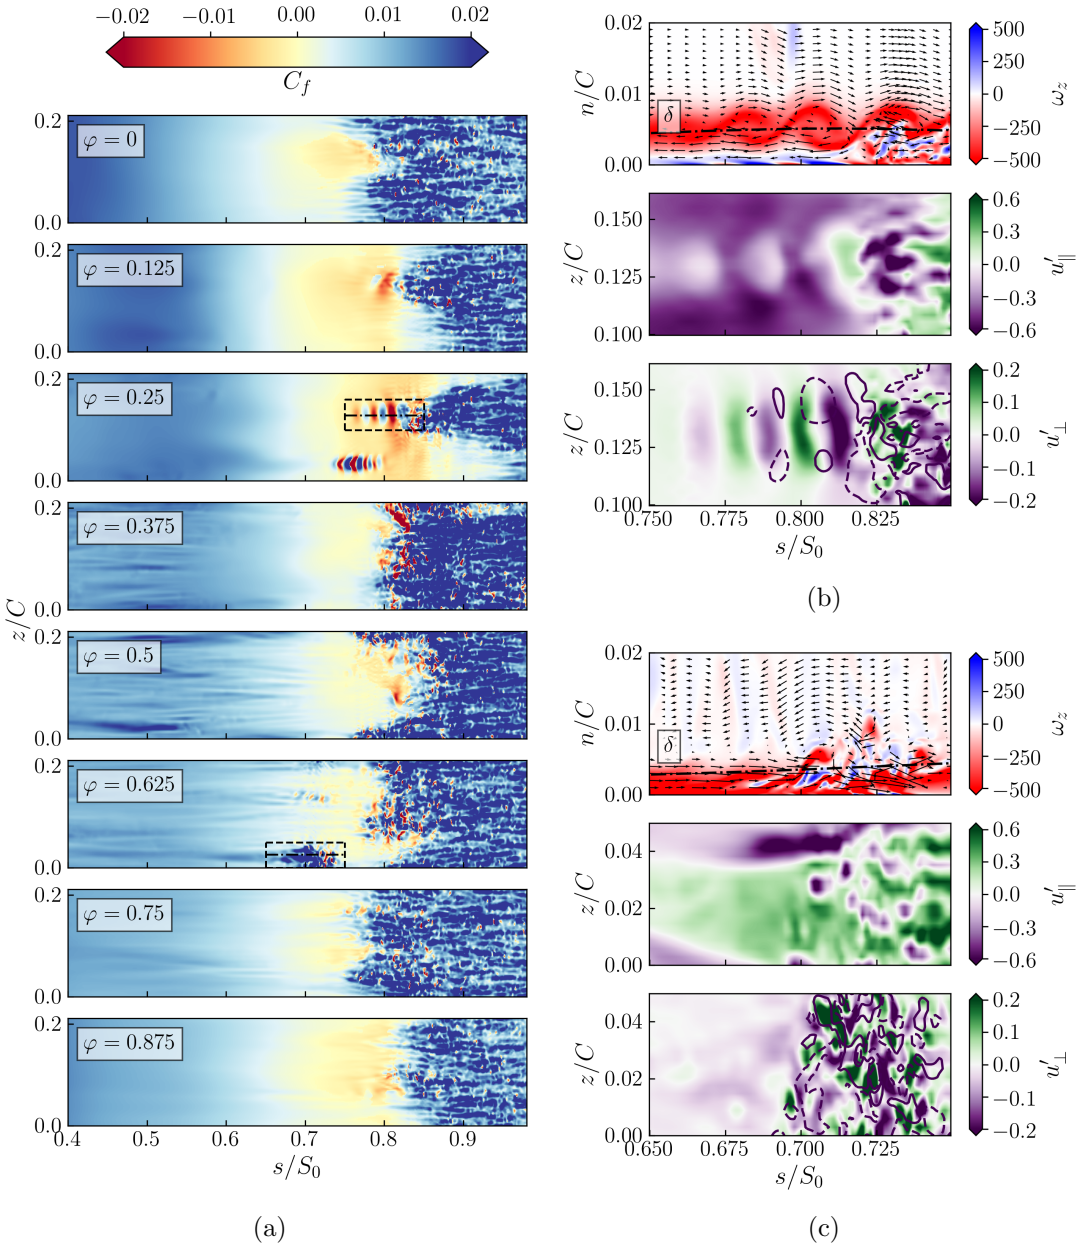

Figure 6.11: Instantaneous suction surface statistics for $Re_2 = 297K$. (a) Instantaneous skin-friction coefficient carpet plot on the aft portion of the suction surface, at 8 different phases. The dashed boxes in phases $\varphi = 0.25$ and $\varphi = 0.625$ are shown in detail in (b) and (c), respectively. For (b) and (c) the top figure shows spanwise vorticity in the blade-normal plane denoted with dash-dotted lines in (a), superimposed with fluctuating velocity vectors. The middle and bottom figures show respectively wall-parallel and wall-normal fluctuating velocity $n/C = 0.005$ away from the wall. The solid and dashed lines in the bottom figures in (b) and (c) are isolines of $w' = \pm 0.15$.

region of the front portion of the suction surface, they generate streamwise streaks via linear lift-up effect [18]. In the context of wake passing, the generation of streaks occurs periodically, and its effect on the transition mechanism in the aft portion of the suction surface is evident at $0.375 \leq \varphi \leq 0.75$.

The nature of the disturbances that penetrate the boundary layer shear is determined by the shear sheltering mechanism [63, 43, 205], further discussed in the context of transition in a compressor cascade by Zaki et al. [206, 203]: the boundary layer shear acts as a filter for the high-frequency vortical disturbances. As previously highlighted by Wissink [193], transition originates behind the wakes and it is due to local forcing. The low-frequency perturbations that penetrate into the boundary layer promote strong shear associated with the spanwise distortions (termed Klebanoff streaks) which may develop into turbulent spots through localized secondary streak instability. These steps characterise the early stages of the bypass transition mechanism, described as a secondary instability of lifted shear layers when they reach the top of the boundary layer, and they are subject to high-frequency free-stream disturbances [198, 68, 204].

The physical mechanisms governing streak instability can be revealed through Floquet analysis of secondary instability [94]. In the presence of Klebanoff streaks, the two most unstable modes of the boundary layer are the inner mode (varicose instability) and the outer mode (sinuous instability) [203], so called due to the wall-normal location of their critical layer. The inner mode is most amplified in the overlap region between low- and high-speed streaks, with highest growth rate at low streak amplitude. The outer mode has high growth rate for high streak amplitude, and it is hosted by low-speed streaks, as also highlighted previously by Schlatter et al. [147].

In the early stages of the wake passing cycle, the penetration of streaks within the boundary layer has not reached the TE region yet. However, the formation of instabilities is detected at $\varphi \approx 0.25$ and highlighted in Figure 6.9b. The vorticity contour suggests the presence of part-span K-H instability, which lifts up from the wall and develops across the boundary layer edge. The alternating parallel fluctuating velocity patterns shown in Figures 6.9b, 6.11b are consistent with the findings of Zhao et al. [207], where they are argued to be evidence of varicose streak instability. However, in this case the instability is most likely due to part-span K-H type instability which occurs naturally in the separation bubble.

At the later phases of the wake passing cycle, the presence of developed turbulent spots is detected (e.g. Figure 6.11c). However, the evidence here presented is not sufficient to uniquely identify the physical streak instability mechanism (i.e. sinuous or varicose) preceeding the inception of the spots. This would require the use of further advanced analytical tools to compute the phase speed and track the time-space development of the localised instability.

Common to the turbulent spot instances shown is their tendency to rapidly lift up in the wall-normal direction and surpass the boundary layer edge, as shown in Figure 6.9c, and at higher Reynolds in Figures 6.10c and 6.11c. Moreover, they are found to be within

regions of positive fluctuating parallel velocity. The opposite behavious is highlighted when the passing wake first impacts on the separated region (Figures 6.9b, 6.10b and 6.11b), and the instability arises within roll-ups of the separation bubble, which explains the presence of negative tangential fluctuating velocity.

6.3.2 Space-time boundary layer behaviour

The time evolution of the spanwise-averaged effect of the wake interactions on the suction surface boundary layer is shown in the space-time plots of momentum thickness and shape factor of Figure 6.12. The BL evolution is captured over 4 wake passing periods $\mathcal{T}_b$ with a resolution of 80 flow snapshots per period. Instantaneous statistics are here preferred over phase-averaged realisations because of their ability to showcase the highly unsteady nature of the wake interactions.

The momentum thickness pattern (shown in Figures 6.12a, 6.12c, 6.12e) is similar across the Reynolds number ranges analysed. The presence of the separation bubble is evident in Figures 6.12b, 6.12d, 6.12f, where solid isolines of $C_f = 0$ are superimposed with the contours of shape factor, identifying the spanwise-averaged extent of separation. When the wake interacts with the separation bubble, transition occurs and the flow reattaches for the whole extent of the suction surface, as shown by the dashed arrow on Figures 6.12b, 6.12d and 6.12f. Concurrently, the momentum thickness is increased starting from a location of $s/S_0 \approx 0.7$ to the TE, which is linked to a periodic increase of profile losses. Following the wake-driven reattachment and the subsequent calmed region, the separation bubble recovers as the separation line moves upstream. The topology of these phenomena is similar across all Reynolds numbers, but at increased Re_2 the momentum thickness is lower over the entire extent of the suction surface shown. The separation bubble captured in Figure 6.12f at $Re_2 = 297K$ is significantly smaller than at lower Reynolds numbers, due to increased BL shear sheltering, but it is also subject to larger relative variability in terms of both streamwise and temporal extent: the second instance of separation visualised in in Figure 6.12f is roughly twice as persistent compared to the following instance. This suggests that the dynamics of the boundary layer between $s/S_0 \approx 0.7 - 0.9$ is largely affected by the presence or absence of streak instabilities trailing the passing of the wake in every cicle (e.g. as previously shown in Figure 6.11a at $\varphi = 0.625$).

The solid box in the bottom-right of Figure 6.12b also highlights another flow feature that occurs predominantly at low Reynolds number. The dashed lines mark the presence of space-time regions of strong negative C_f. In particular, two thin, elongated areas can be identified within the separated flow region. The first (in a temporal sense) begins at $s/S_0 \approx 0.9$, and it is followed by another occurring upstream at $s/S_0 \approx 0.85$. The former represents the inception of the effect of external disturbances on the state of the boundary layer. Nested between the two dashed regions is a thin area of low shape factor and thus reattached flow. This corresponds to the inception of two large-scale roll-ups with a region of attached flow in between. The rolls appear consecutively and they are short-

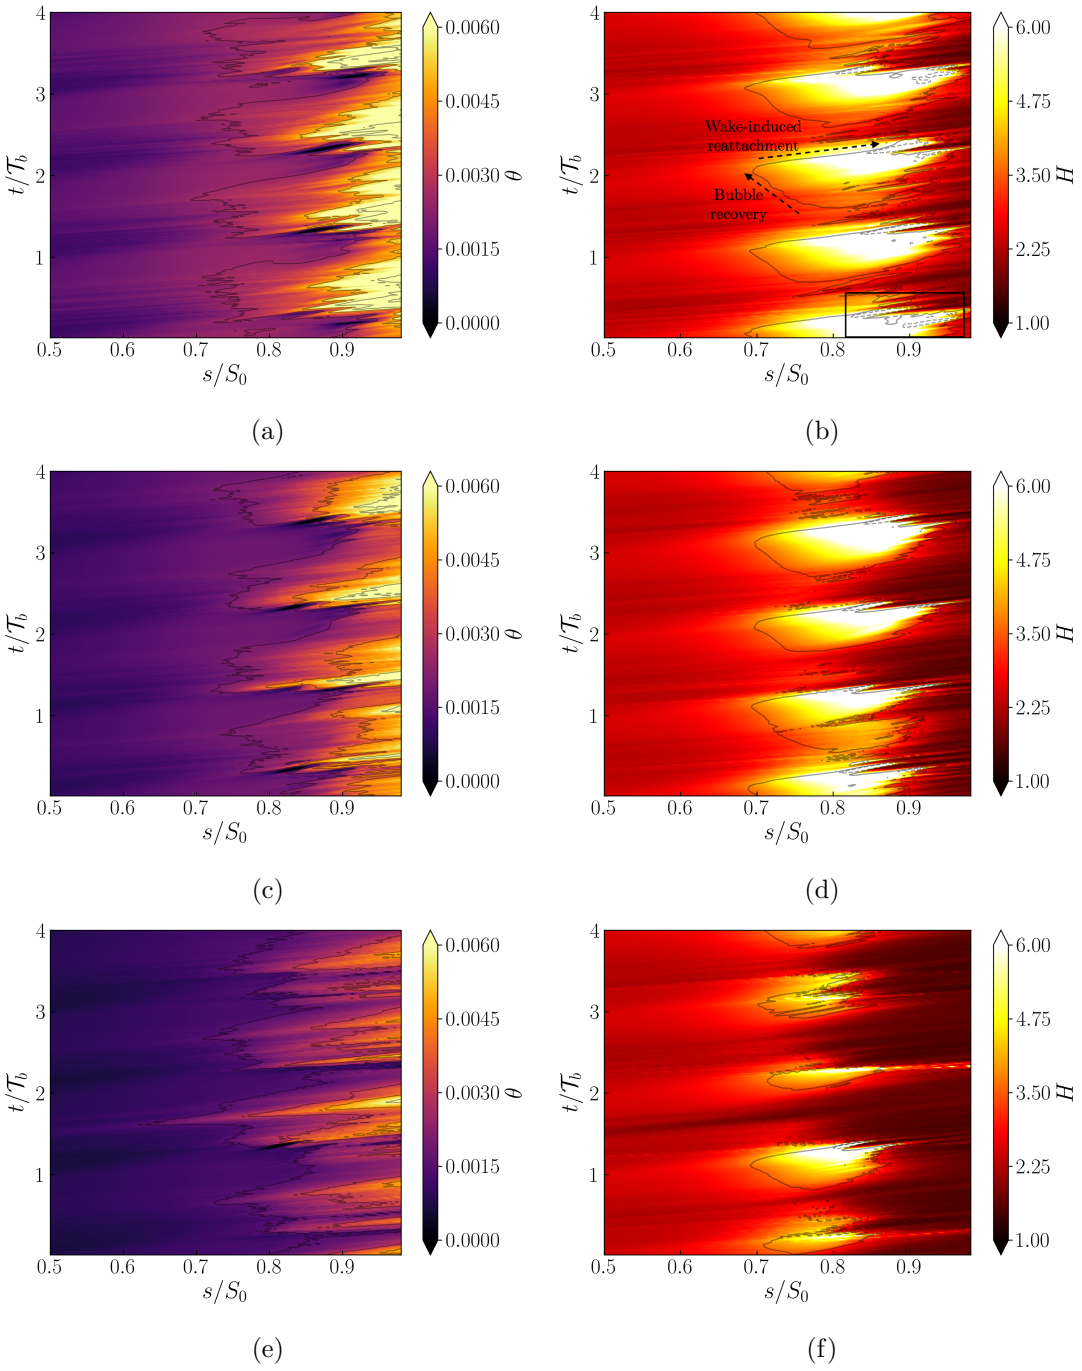

Figure 6.12: Instantaneous spanwise-averaged space-time contour of suction surface boundary layer parameters. Left: θ. Right: H superimposed with isolines of wall-shear stress at two levels: $C_f = 0$ (continuous line), $C_f = -0.024$ (dashed line). From top to bottom: (a,b) $Re_2 = 86K$, (c,d) $Re_2 = 157K$, (e,f) $Re_2 = 297K$.

lived. Following the second roll-up, the boundary layer undergoes transition to turbulence and remains attached up to the TE. This pattern cannot be identified at higher Reynolds number ($Re_2 = 297K$), where the smaller separation bubble is periodically suppressed due to bypass transition and without the formation of large-scale K-H roll-ups.

6.3.3 Phase-averaged flow visualisations

In order to provide further evidence of the change in the dynamics with increasing Re_2, phase-averaged detailed views of the spanwise vorticity, TKE and turbulence production P^k in the TE region are shown in Figure 6.13.

Two distinct regions of large vorticity associated with the interaction of the wake and the separation bubble are visible in the vorticity contours for $Re_2 = 83K, 157K$. In these regions, strong turbulent activity is detected. As the Reynolds number is increased, the transitional region is shifted upstream and away from the TE. At high Reynolds number ($Re_2 = 297K$), only the first K-H roll-up structure is visible but of significantly reduced size, comparable to that of the thin shear layer. At the two lower Reynolds numbers, higher levels of turbulence production are associated with the presence of roll-ups; two thin regions of negative turbulence production are nested in the near-wall region, preceding and following the first K-H roll-up. This phenomenon was documented in an adverse pressure gradient turbulent boundary layer in the separated flow region upstream of reattachment [156, 2] and more recently investigated in depth by Cimarelli et al. [28].

The kinematics of the wake convection through the blade passage was extensively explored experimentally in its various phases [163, 164]. A comparison of the differences in the kinematics of wake convection through the mean passage at different Reynolds numbers is shown in Figure 6.14, where the phase-locked perturbation velocity vectors are superimposed to phase-averaged TKE. The negative jet perturbation decelerates the flow upstream of the wake center, and accelerates is downstream. Wake elongation and convective transport away from the pressure surface promote a reduction of the turbulence intensity on the pressure side, while wake stretching causes the same effect on the front portion of the suction surface [60]. High levels of aligned turbulent stress and spatial velocity gradients (which yield higher turbulent production) within the mean passage promote higher levels of TKE [111] in the downstream portion of the bowing wake, which has an impact on the subsequent blade row. From the Figure, the shape of the fluctuating ensemble-averaged velocity vectors (i.e. negative jet effect on the suction surface, and large contra-rotating vortical structures) is retained across the Reynolds numbers analysed. The peak turbulence intensity at $Re_2 = 86K$ is larger than at higher Reynolds numbers, but no significant changes can be visibly appreciated between $Re_2 = 157K$ and $Re_2 = 297K$.

6.3.4 Fluctuating blade wall distributions

The results so far analysed provide insight into the effect of the periodic disturbances on the suction surface separation bubble, importantly linked to the state of the boundary layer

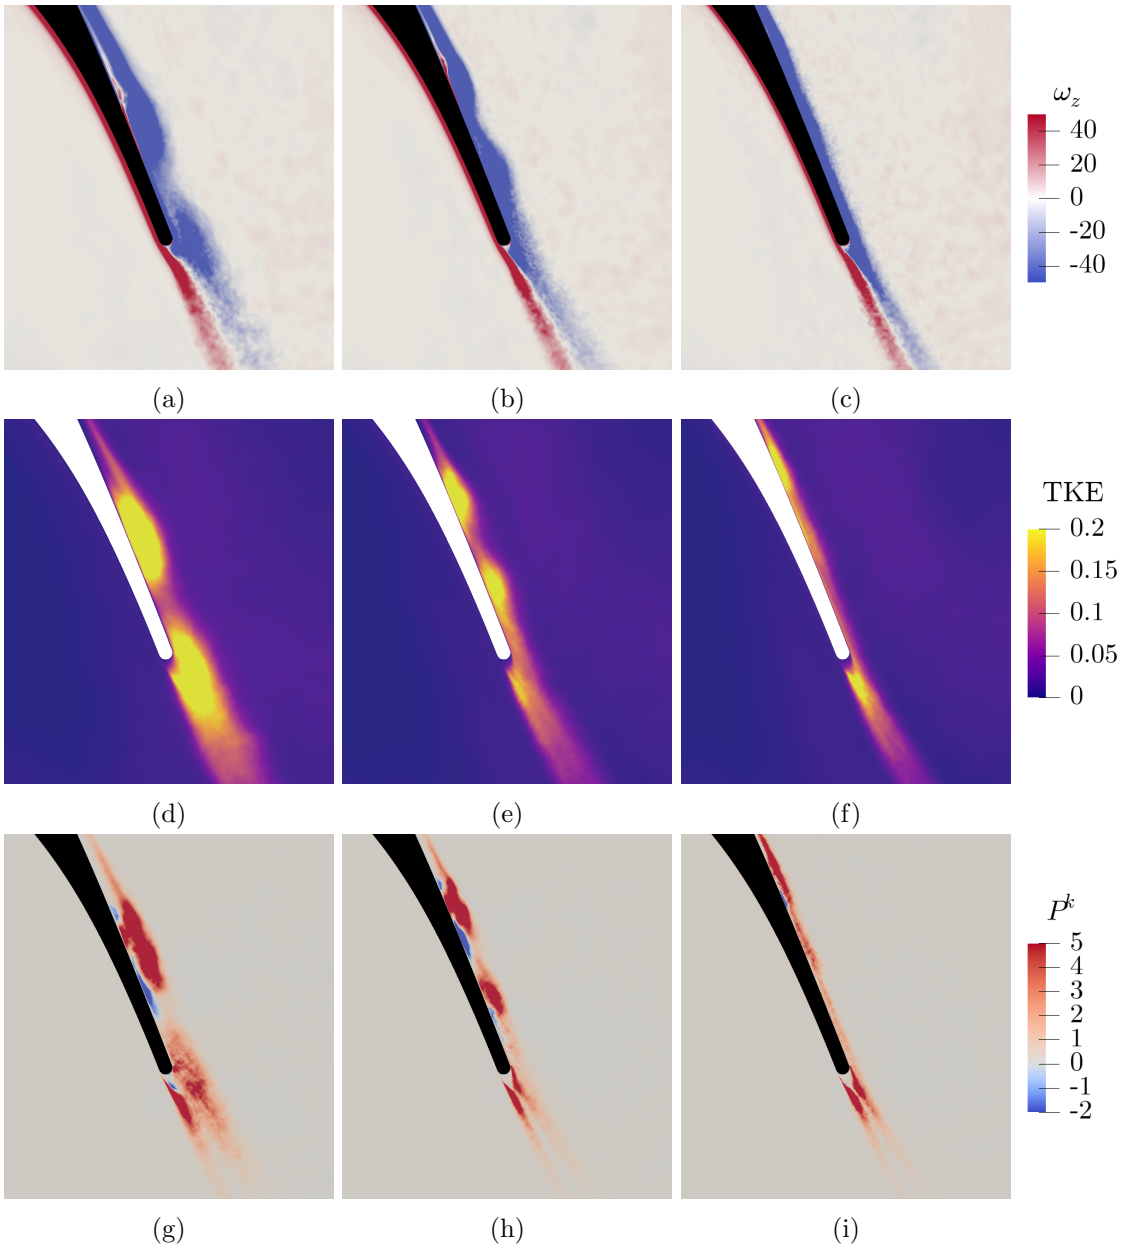

Figure 6.13: Phase-locked fields at $\varphi = 0.375$ in the TE region. From top to bottom: (a, b, c) spanwise vorticity, (d, e, f) TKE and (g, h, i) turbulence production P^k. From left to right: (a, d, g) $Re_2 = 86K$, (b, e, h) $Re_2 = 157K$ and (c, f, i) $Re_2 = 297K$. Figures not to scale.

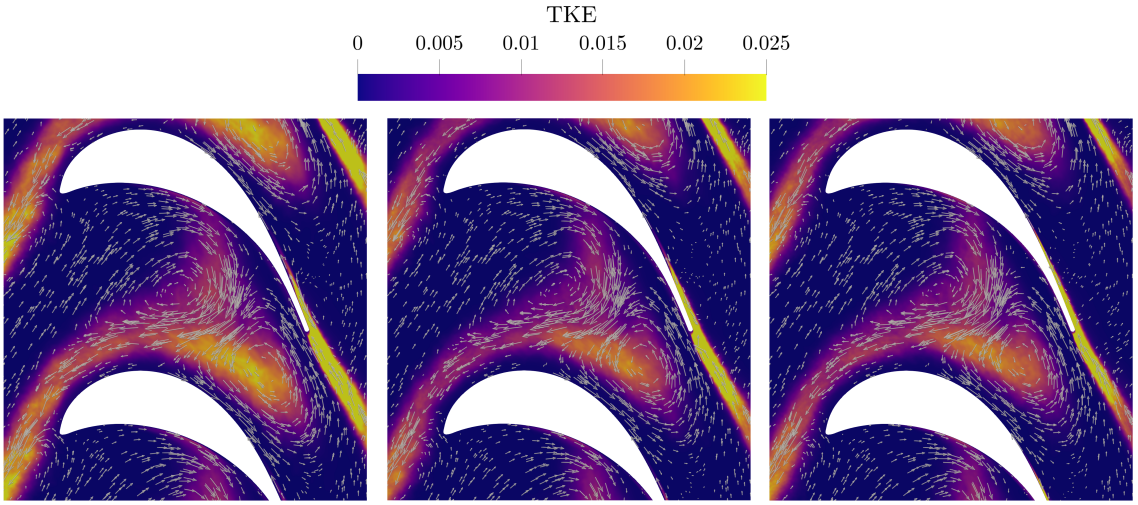

Figure 6.14: Phase-locked TKE at $\varphi = 0$ superimposed with phase-locked fluctuating velocity vector $\widetilde{\mathbf{u}}_{\varphi=0} - \overline{\mathbf{u}}$. The length of the arrows is proportional to the velocity magnitude. From left to right: $Re_2 = 86K, 157K, 297K$. Figures not to scale.

at the trailing edge and therefore profile losses. However, it is necessary to investigate other quantitative results to gain insight into the potential effect of the wake passing on the blade loading. Figure 6.15 shows the pressure and skin-friction coefficient on both suction and pressure surface, with three types of information overlayed on every subfigure. The black solid line indicates the time-averaged distribution, which is discussed in detail in the next section. In order to highlight unsteady effects, the standard deviation of the distributions (low transparency) and the envelope of instantaneous flow fields (high transparency) are superimposed. Each Reynolds number is shown with the same shade of colour adopted through the rest of the chapter. The pressure distribution alone is not sufficient to isolate the potential effect on the blade loading, and the skin-friction coefficient provides additional information.

The standard deviation of C_p remais extremely similar across all flow regimes: the distribution is rather constant along the pressure surface and it becomes thinner in the aft portion for $s/S_0 > 0.9$. Analogously, the impact of the wake on the pressure distribution in the front portion of the suction surface ($s/S_0 < 0.1$) remains limited. The instantaneous envelope trend is similar to the standard deviation in the entirety of the pressure surface and the suction surface up to $s/S_0 \approx 0.8$, which corresponds approximately to the time-averaged separation location. The impact of the wake increases significantly in the adverse pressure gradient region of the suction surface. Comparing the pressure distributions and skin-friction coefficient highlights the areas where flow separation occurs, which is the only portion of the blade where significant differences are induced by an increase in the Reynolds number. The overall behaviour of the fluctuating distributions in cases $Re_2 = 86K$ and $Re_2 = 157K$ is extremely similar, and it differs with $Re_2 = 297K$ in the final part of the suction surface, for $s/S_0 > 0.8$. In this region, the shear layer periodically rolls up, which is highlighted by the strong negative values reached by the C_f evelope, which are about

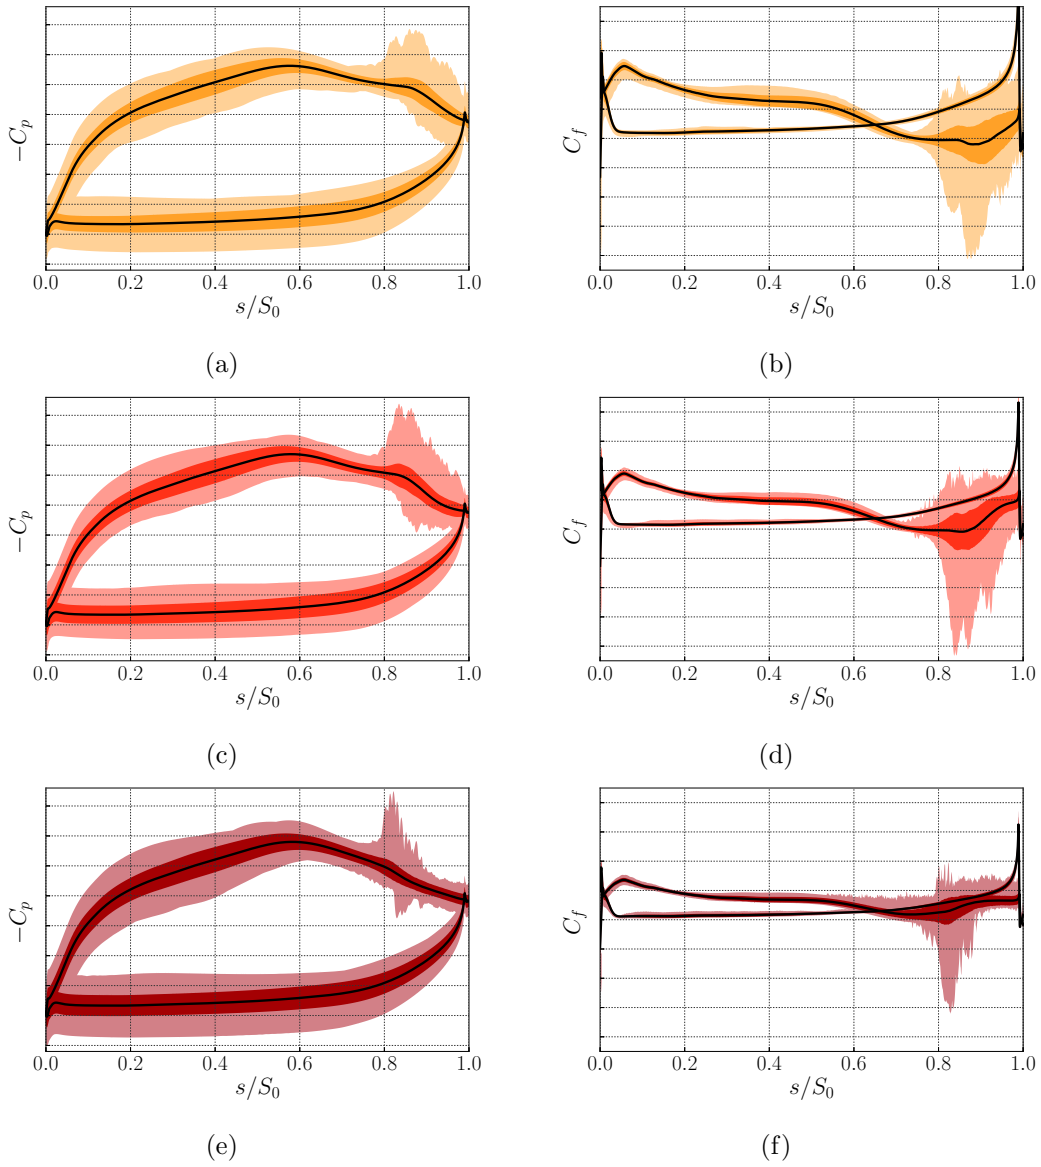

Figure 6.15: Black line: time-averaged pressure distribution (a,c,e) and skin-friction coefficient (b,d,f) for increasing Reynolds numbers: (a,b) $Re_2 = 86K$, (c,d) $Re_2 = 157K$, (e,f) $Re_2 = 297K$. Dark area: $C_p \pm \sigma$ and $C_f \pm \sigma$; light area: minimum and maximum value envelope. Y-axis tick labels are omitted due to data sensitivity.

three times as large in magnitude as the maximum positive fluctuating difference with the time-averaged distribution. The region of large negative fluctuating wall-shear stress reaches the trailing edge at lower Reynolds, while it is significantly shorter at $Re_2 = 297K$, reaching up to $s/S_0 \approx 0.9$ only. The increased positive wall-shear stress region extends to the TE, thus denoting transition to a fully turbulent boundary layer. Having discussed the impact on the blade wall distributions with focus on the unsteady effects on the suction surface, the analysis of time-averaged statistics is the focus of the following section.

6.4 Blade wall distributions

The time-averaged and spanwise-averaged pressure distribution and skin-friction coefficient are shown in Figure 6.16. The data with LE forcing from Chapter 5 are included to highlight the impact of the periodic disturbances on time-mean flow performance indicators. Due to the presence of passing bars, the effective inflow angle is decreased, from $\alpha_1 = 33.86°$ to $\alpha_1^{\text{eff}} = 31.73°$ in $Re_2 = 86K$, $\alpha_1 = 33.96°$ to $\alpha_1^{\text{eff}} = 32.17°$ in $Re_2 = 157K$ and $\alpha_1 = 34.08°$ to $\alpha_1^{\text{eff}} = 32.18°$ in $Re_2 = 297K$. The difference translates into discrepancies in the front portion of the suction surface between the cases with and without passing wakes, evident in the suction surface skin friction peak at $s/S_0 \approx 0.05$, where the IW profiles are slightly lower than the cases without incoming wakes. In this same region of the suction surface, lower Reynolds number corresponds to higher wall-shear stress.

The major differences due to the introduction of periodic disturbances are in the adverse pressure gradient part of the suction surface. Especially at low Reynolds number, in the C_p distribution the BF cases presents a short plateau region at $s/S_0 \approx 0.85$, indicative of a weak separation bubble. This region is delayed and shortened with the introduction of wake passing. The skin-friction coefficient distribution further clarifies the differences. Due to the periodically modified transition mechanism on the suction surface, the wake passing causes an upstream shift of the time-averaged separation bubble at $Re_2 = 86K$ and $Re_2 = 157K$. In both cases flow separation occurs at $s/S_0 \approx 0.74$, and the extent of separation is almost halved compared to the cases with body forcing, as highlighted by the quantitative values reported in Table 6.3. The negative C_f peak denoting the presence of roll-ups and anticipating reattachment is moved upstream, flattened and only visible at low Reynolds number. This is a consequence of the wake motion along the suction surface and it does not correspond to stationary flow features. In both cases $Re_2 = 86K$ and $Re_2 = 157K$ the flow is attached at the trailing edge, but not fully developed. Further increasing the Reynolds number ($Re_2 = 297K$) removes separation in a time-averaged sense and yields a fully turbulent boundary layer at the trailing edge, as highlighted by the flat C_f profile for $s/S_0 > 0.9$.

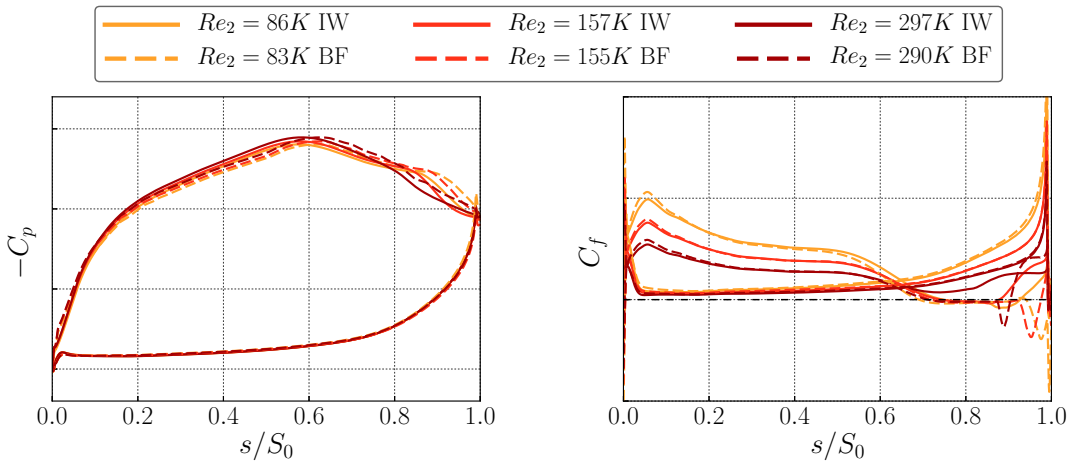

Figure 6.16: Blade wall distributions with increasing Re_2, compared with momentum forcing cases of Chapter 5. Left: pressure distribution, and right: skin-friction coefficient. Y-axis tick labels are omitted due to data sensitivity.

Table 6.3: Time-averaged separation bubble statistics with inflow wakes and momentum forcing (presented in Chapter 5): separation point and reattachment point. Due to data sensitivity, all values were intentionally nondimensionalised by the respective values of case BF at $Re_2 = 297K$, consistently with Chapter 5.

Re_2	$86K$		$157K$		$297K$	
Inflow	IW	BF	IW	BF	IW	BF
$(s/S_0)^*_{\text{sep}}$	1.04969	0.96801	1.04287	1.00905	-	1.0
$(s/S_0)^*_{\text{reat}}$	1.02477	-	0.97038	1.07885	-	1.0

6.5 Boundary layers

The impact of wake passing on the boundary layer behaviour is compared against the body forcing data of Chapter 5. Figure 6.17 shows momentum thickness and shape factor at the three Reynols numbers analysed.

The front portion of the suction surface before peak suction (for $s/S_0 \lessapprox 0.6$) is considered first. As the Reynolds number is increased, the momentum thickness decreases indicating a thinner shear layer. No appreciable difference exists between the cases with and without incoming wakes: in the favourable pressure gradient region of the profile, the high shear sheltering mechanism prevents large boundary layer distortions. The shape factor is similar in all cases, indicating self-similar boundary layer growth.

The aft portion of the blade ($s/S_0 \gtrapprox 0.6$) is considered next. The time-averaged momentum thickness grows gradually at all Reynolds numbers. In the cases with no wakes, downstream of $s/S_0 \gtrapprox 0.75$ its growth decreases up to the TE region, where BL transition occurs. Downstream of transition, θ rises sharply, and only at the highest Reynolds number $Re = 297K$ the body forcing case presents a higher momentum thickness at the trailing

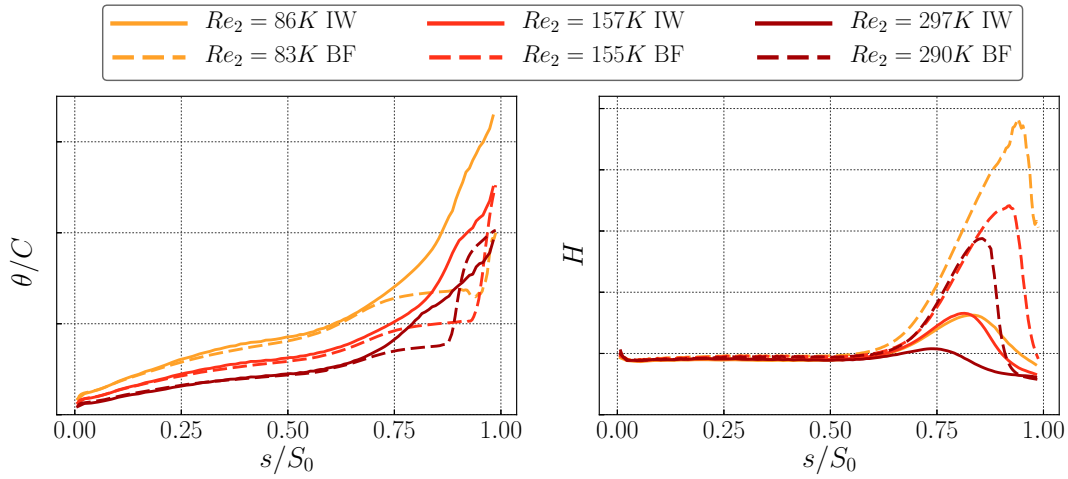

Figure 6.17: Boundary layer parameters with increasing Re_2, compared with clean inflow cases of Chapter 5. Left: momentum thickness, and right: shape factor. Y-axis tick labels are omitted due to data sensitivity.

edge. In the same region, the absence of incoming wakes yields a sharp increase in the shape factor, which indicates a region of separated flow already discussed in Chapter 5. Instead, the presence of periodic disturbances and their effect on the separation bubble contribute to maintaining the shape factor at significantly lower values in the adverse pressure gradient region. As previously highlighted by the skin-friction distribution in Figure 6.16, in case $Re = 297K$ the BL reaches a fully turbulent state before reaching the TE, which is also reflected in the shape factor that settles for $s/S_0 > 0.9$.

Further details on the boundary layer profiles are provided in Figure 6.18. The top row shows velocity profiles; the last two BL profiles shown prior to the TE are identical at $Re = 297K$, indicating that its state is fully developed. The same Figure shows that the reattachment process at low Reynolds number occurs between $s/S_0 \approx 0.85 - 0.9$ and the BL at the trailing edge is not fully turbulent.

The TKE and P^k boundary layer profiles are also shown in Figure 6.18. In the three stations between $0.7 \leq s/S_0 \leq 0.8$, the case $Re = 297K$ initiates transition, as highlighted by a pronounced peak in the near-wall region. At $s/S_0 \approx 0.8$ the same mechanism also develops at lower Reynolds numbers, but the peak is increasingly displaced away from the blade surface. At $s/S_0 \approx 0.85$, a second peak develops in the high Reynolds number case. This occurs at $n/C \approx 0.01$ for the TKE profiles, and in the very near wall region for the P^k profiles. Therefore, the emergence of the second peak for TKE corresponds to the damping of the first peak of P^k, and vice versa. As the turbulent profile develops, the fluctuations are enhanced away from the wall but the production mechanism is confined in the region of large pressure gradients induced by the wall. This behaviour occurs in the presence of a fully developed turbulent BL under adverse pressure gradient, and it is not observed at lower Reynolds numbers, where the main peak of turbulence production is driven by the presence of migrating roll-up vortices.

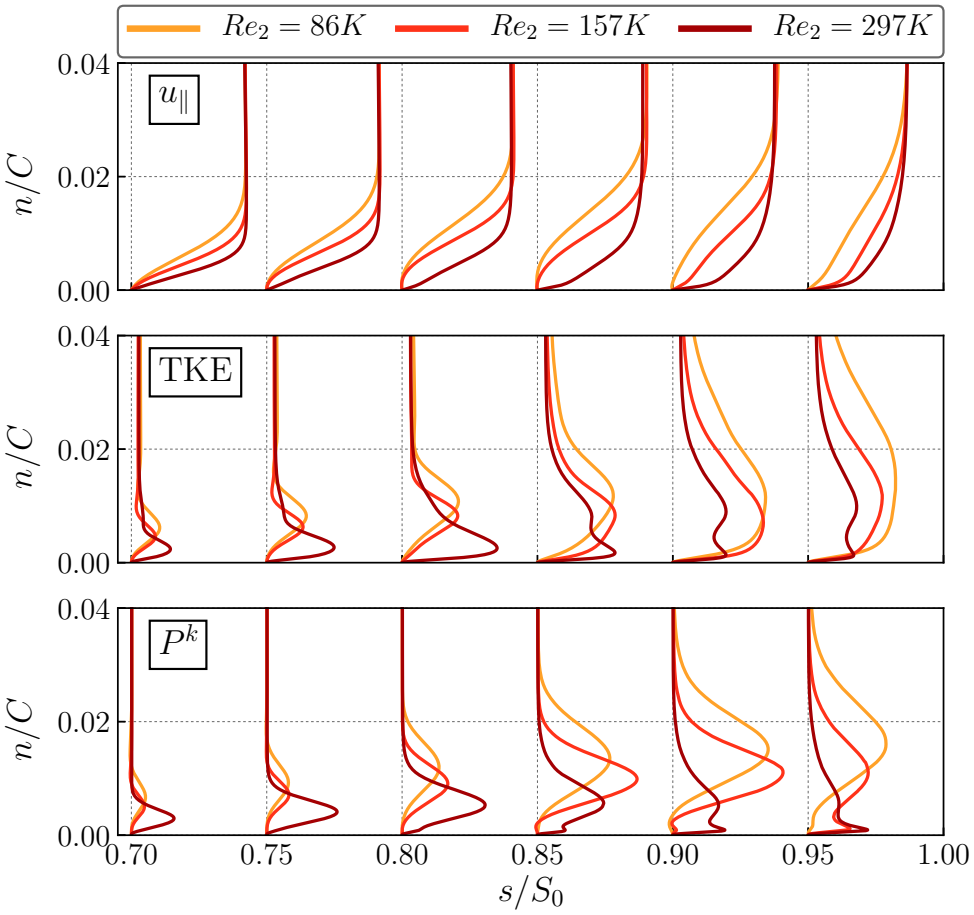

Figure 6.18: Time-averaged boundary layer profiles in the aft portion of the suction surface subject to incoming wakes, for increasing Re_2. Each quantity is scaled by an appropriate factor to improve visual clarity. Top: Parallel velocity component; middle: turbulence kinetic energy; bottom: turbulence production.

In particular, at $Re_2 = 157K$, a second peak of moderate intensity in the P^k profile is identified in the TE region, but the TKE profiles do not reflect the same topology as $Re_2 = 297K$ and rather show a single peak at roughly half the height of the boundary layer edge. This behaviour was analogously documented by Sarkar [146].

6.6 Wake traverses and experimental comparison

Extensive experimental data is available from LDA measurements and it is presented in this section. As a reference for clarity of interpretation of the figures, the suction side is located towards the right side of the profiles.

6.6.1 Velocity wakes

The velocity wakes are compared with experiments in Figure 6.19a. The agreement is very close at all flow regimes, accurately capturing both flank width and negative peak. The same level of accuracy is achieved in the vertical velocity wakes reported in Figure 6.19c. The agreement observed at $Re_2 = 86K$ represents a substantial improvement compared to the clean inflow analysis of Chapter 5. The introduction of periodic disturbances and the modified BL state at the trailing edge allow for a more faithful representation of the flow physics measured in the experiments. The presence of background turbulence was simulated at $Re = 86K$ only and represented with dashed lines. The wake profiles overlap closely to the curves obtained without incoming HIT, with marginal discrepancy observed in the mean passage on the pressure side. This highlights that discrete disturbances are the dominant mechanism in the dynamics of the suction surface separation bubble.

The axial flow component (Figure 6.19b) is underestimated both in the wake region and especially in the mean passage. Particularly at $Re_2 = 86K$, the experimental scatter (defined in Equation 5.3 and represented as a grey shaded area) is higher compared to the velocity magnitude and vertical component profiles. However, experimental precision is not resposible for the full-pitch discrepancy. This effect is due to the experimental streamtube contraction introduced in Section 5.2.1; the discussion is expanded in greater detail later in the Chapter.

6.6.2 Loss profiles

The loss profiles are reported in Figure 6.20, showcasing excellent agreement with the experiments. The trend captured is very similar to that of the velocity wakes, with agreement improving at higher Reynolds numbers. The entity of losses in the mean passage remains the same across the three flow regimes explored. The computed results capture a slight decrease of the loss profile peak by $\sim 6\%$ and 10% in cases $Re_2 = 157K$ and $Re_2 = 297K$, respectively. The experimental values remain however approximately constant across the three flow regimes. At higher Re_2, the wake width decreases significantly. It is therefore evident that the smaller momentum thickness at the trailing edge typical of high-Reynolds configurations, combined with a fully suppressed separation bubble due to bypass transition results in a thinner loss profile.

6.6.3 TKE wakes

The turbulence length scale leaving the turbine measured from a spanwise line of probe points is smaller in the high-Reynolds case, where the two-point correlation decays faster. This in turn promotes lower levels of fluctuations in the wake. At low Re the separation bubble is also periodically open, giving rise to large-scale vortex shedding which in turn is visible in higher levels of fluctuating turbulence kinetic energy (Figure 6.21). The higher Reynolds cases $Re_2 = 157K$ and $Re_2 = 297K$ have very similar profiles. At all flow regimes, the TKE profile in the mean passage on the pressure side retains marginally

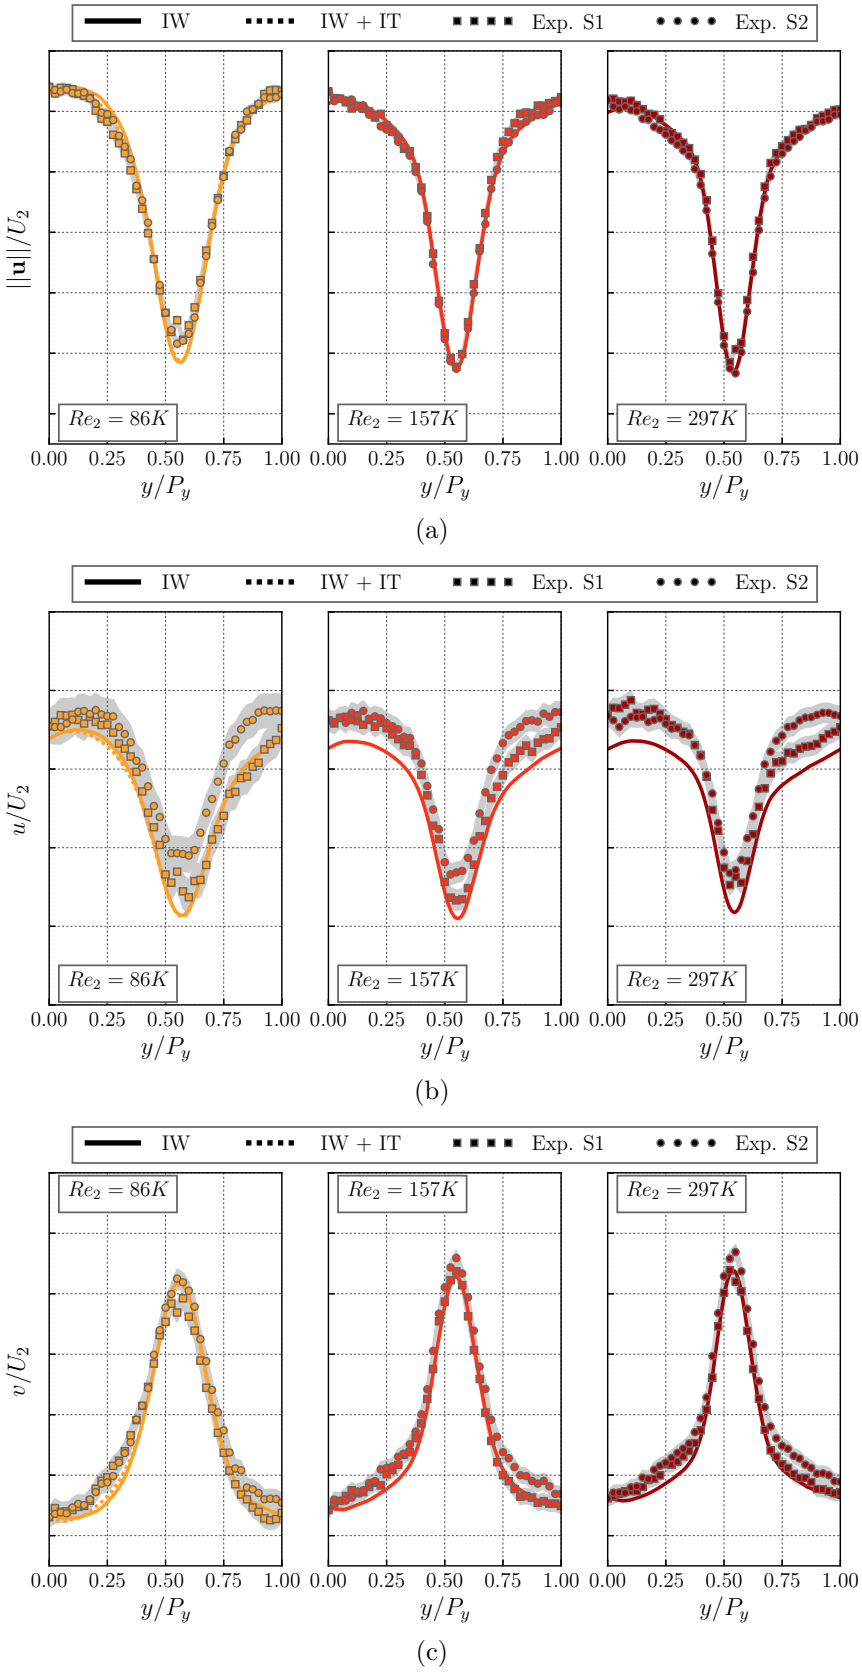

Figure 6.19: Velocity profiles at $\hat{x} = 0.513$. Y-axis tick labels are omitted due to data sensitivity.

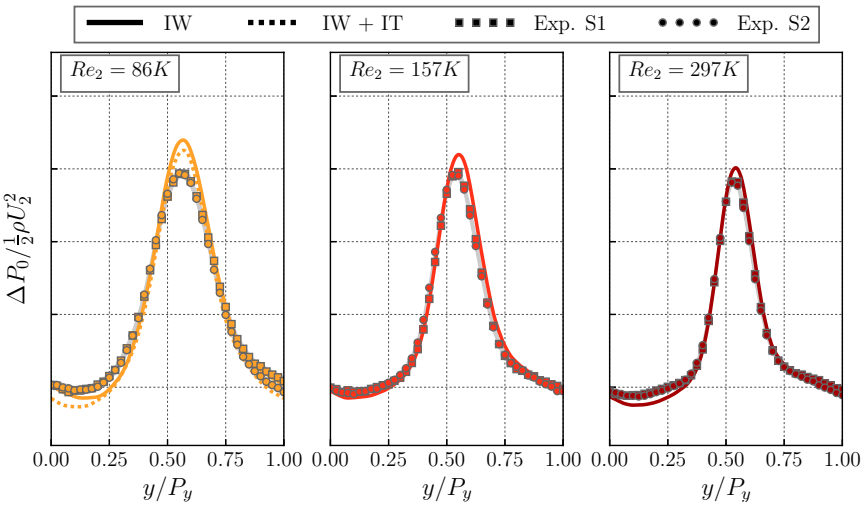

Figure 6.20: Total pressure loss profiles at $\hat{x} = 0.513$. Y-axis tick labels are omitted due to data sensitivity.

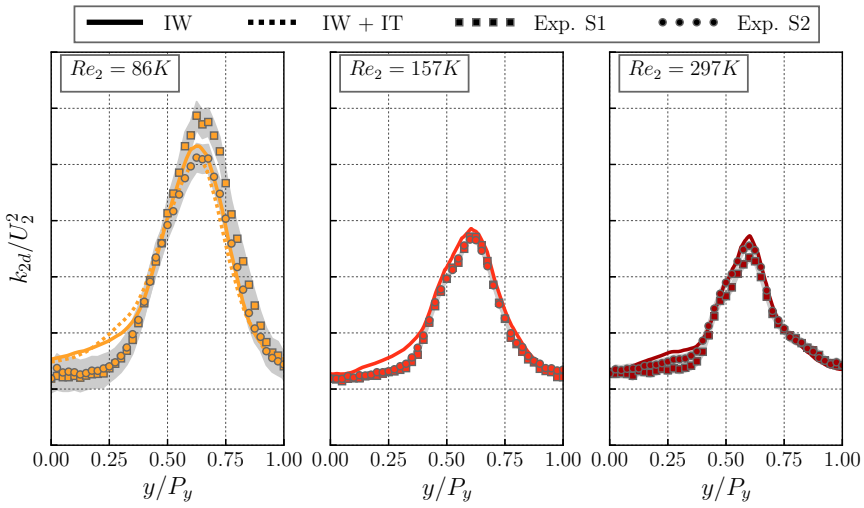

Figure 6.21: Turbulence kinetic energy profiles at $\hat{x} = 0.513$. Y-axis tick labels are omitted due to data sensitivity.

higher values compared to the experiments; the presence of background fluctuations does not modify the TKE profile in this region.

6.6.4 Mixed-out measurements

The mixed-out loss coefficient is shown in Figure 6.22a. The numerical results are very close to the narrow band of uncertainty of the experimental measurements, overall showcasing excellent agreement, particularly in the supercritical regime. At low Reynolds numbers, where the separation periodically extends to the trailing edge, it is traditionally very challenging to capture losses accurately. Case IW + IT at $Re_2 = 86K$ does not improve the agreement with experiments, which were conducted in absence of background fluctuations.

The mixed-out loss coefficient is an important design parameter, relevant in the context of isolated blade rows. However, the realistic gap between rotor and stator rows is between $20 - 35\%C_{ax}$. Therefore, the losses when the wake impacts on the following blade row do not exactly correspond to what is predicted by the mixed-out measurements (here carried out at $\hat{x} = 0.513$). For this reason, the mass-averaged total pressure loss coefficient was also calculated at $\hat{x} = 0.2$ for the IW cases and included in Figure 6.22a. The maturity of the wake in a location representative of the leading edge of the following airfoil row is shown by comparison of the mass-averaged and the mixed-out loss coefficient, the latter calculated at the experimental measurement location. As shown in the Figure, at all flow regimes the two values are extremely similar, suggesting that the wake quickly develops downstream of the TE. This type of analysis is of relevance in the design of the axial gap, and its impact on machine performance.

The relative error of the mixed-out total pressure loss coefficient with respect to the experimental measurements is reported in Table 6.4, which highlights that the two extreme regimes $Re_2 = 86K$ and $Re_2 = 297K$ provide the closest agreement. The strongest impact deriving from the introduction of the wake passing effect is in the subcritical regime. In fact, at $Re_2 < 150000$, the mixed-out loss is drastically reduced with the introduction of periodic disturbances, while in the supercritical regime the difference is smaller. This can be appreciated in Figure 6.22b, which shows the relative difference of the body forcing results of Chapter 5 (which provided the closest agreement with clean inflow experiments) and the findings of the present Chapter, with inflow wakes. In the subcritical regime (the light blue area) the loss coefficient measured in the experiments reaches values $\approx 60\%$ higher in the clean inflow case than the setup with incoming wakes. The low Reynolds numerical results without inflow wakes did not achieve the same level of agreement as the higher Reynolds cases, which explains the lower predicted difference of $\approx 35\%$ at $Re_2 = 86K$. In the critical regime, this difference drops to nearly zero both in experimental and numerical results. At supercritical Reynolds numbers, the experimental loss in the absence of incoming wakes is $\approx 10\%$ lower than the value at $Re = 157K$, while no significant difference is shown by the computational results, highlighting the lower

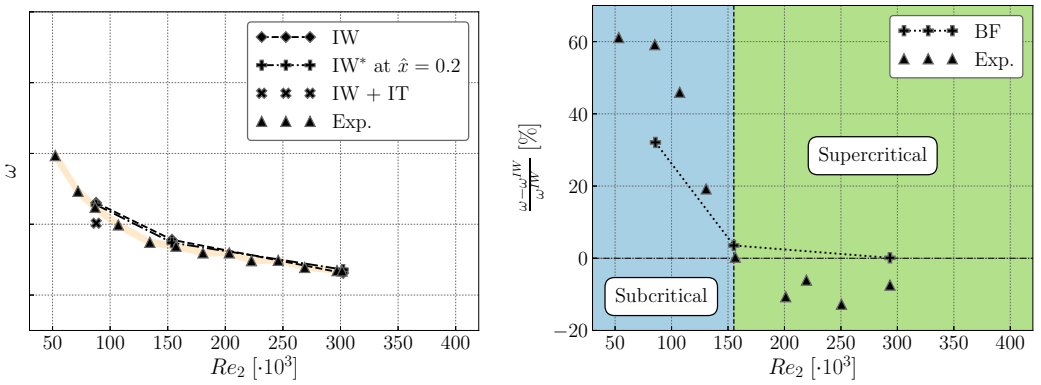

(a) (b)

Figure 6.22: (a) Mixed-out total pressure loss coefficient. The orange area indicates an uncertainty of 2.5% associated with the measurement chain. The nomenclature IW* indicates mass-averaged loss coefficient. (b) Relative difference of loss coefficient, comparing the inflow wakes and experimental results with the body forcing and experiments of Chapter 5. Y-axis tick labels are omitted due to data sensitivity.

sensitivity of the loss coefficient at higher Reynolds numbers.

6.6.5 Exit angle and streamtube correction

The exit angle profiles are shown in Figure 6.23. The experimental data shows larger uncertainty than in all other types of wake statistics, with significant differences between the two experimental traverses S1 and S2. Within each traverse, the scatter (and therefore the measurement precision) is also significant. The experimental set S1 shows a more strongly varying profile in the pitchwise direction and is adopted as main reference. In the experiments, as Re_2 is increased the absolute variation of the exit angle is roughly halved. In the computational cascade, the traverses do not capture such evolution but remain similar in all flow regimes analysed. The superposition of inflow wakes and inflow turbulence only introduces marginal differences in the flow angle profile at $Re_2 = 86K$, analogously to all previously shown statistics.

The streamwise component of the velocity wakes 6.19b is in deficit compared to the experimental data, especially in the mean passage and at higher Re_2. Therefore, the exit angle is overestimated, predicting excessive flow turning. The various sources of uncertainty in the wake profiles are dominated mainly by two factors. The first contributor is the uncertainty in the experimental data, which is small for the velocity wakes: $\epsilon_{u/U_2} = 0.00426, 0.00256, 0.00243$ for $Re_2 = 86K, 157K, 297K$ respectively. The exit angle presents higher scatter in the experimental measurements, of $\epsilon_{\alpha_2} = 0.3°, 0.195°, 0.188°$ for the three Reynolds numbers considered. The second contributor to measurement discrepancies is the streamtube effect, introduced conceptually in Section 5.2.1. A first-order correction is now introduced to show that the absence of a modelling feature accounting for the streamtube effect is mostly responsible for the lack of agreement in the streamwise velocity component

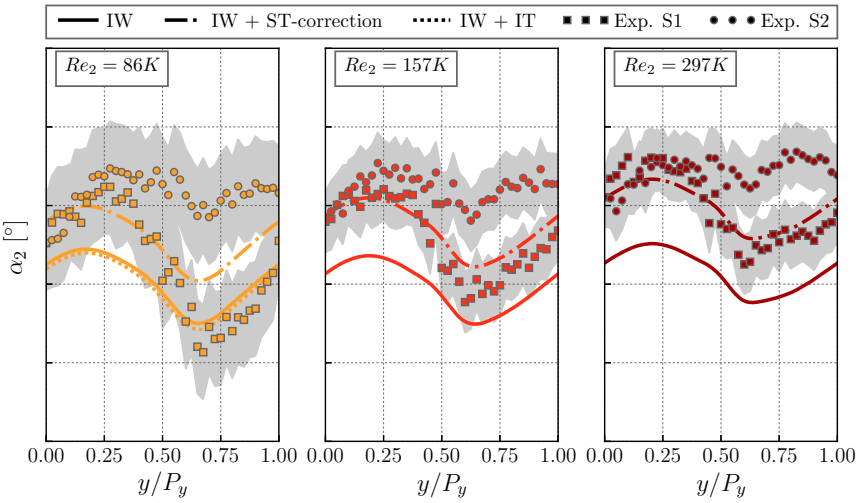

Figure 6.23: Exit angle profiles at $\hat{x} = 0.513$. Y-axis tick labels are omitted due to data sensitivity.

and the exit angle.

The ratio between the experimental and computational mass-flow rate can be estimated as:

$$k_s = \frac{\int_0^{p_y} \overline{\mathbf{u}}_{\text{Exp}} \cdot \hat{\mathbf{n}} \, \mathrm{d}y}{\int_0^{p_y} \overline{\mathbf{u}}_{\text{IW}} \cdot \hat{\mathbf{n}} \, \mathrm{d}y} = \frac{\frac{1}{2}(\int_0^{p_y} \overline{u}^{S1}_{\text{Exp}} \, \mathrm{d}y + \int_0^{p_y} \overline{u}^{S2}_{\text{Exp}} \, \mathrm{d}y)}{\int_0^{p_y} \overline{u}_{\text{IW}} \, \mathrm{d}y} \,, \tag{6.6}$$

considering the average of the mass flow rate calculated in the experimental passages S1 and S2. The streamwise velocity extracted from the numerical simulations is stretched:

$$\overline{u}^{\text{mod}}_{IW} = k_s \overline{u}_{IW} \,, \tag{6.7}$$

and the exit angle profile as well as the mixed-out values are recalculated. The contracted spanwise domain can also be estimated by imposing mass conservation:

$$L'_z = L_z / k_s \,. \tag{6.8}$$

The corrected exit angle profile is represented by the dash-dotted lines in Figure 6.23. The agreement with S1 is significantly improved, especially at high Reynolds numbers, with the computational traverses well within the scatter bounds. The corrected streamwise velocity profile (not shown) is completely overlapped to the experimental measurements. However, this was expected since the mass flow rate is the metric adopted to adjust the wakes extracted from the numerical experiments.

The mixed-out exit angle is shown in Figure 6.24. The values computed without streamtube correction are offset by almost 1° but they capture the correct trend, indicating the presence of a bias. Introducing the streamtube correction shifts the predicted exit angle to a region within the error bounds of the experimental measurements, therefore suggesting that a first-order estimation is sufficient to account for the correct physical

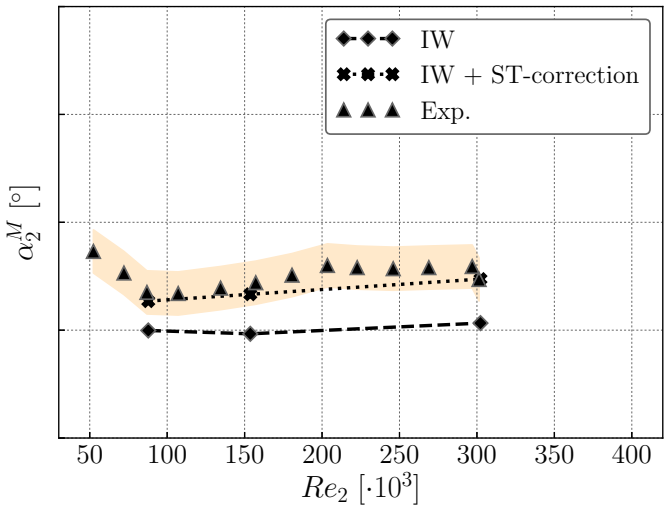

Figure 6.24: Mixed-out exit angle as function of the Reynolds number. The orange area indicates an uncertainty of ±0.2° associated with the measurement chain. Y-axis tick labels are omitted due to data sensitivity.

Table 6.4: Summary of relative error between experimental and computational mixed-out quantity, as well as summary of the streamtube contraction factors.

Parameter	$Re_2 = 86K$	$Re_2 = 157K$	$Re_2 = 297K$
$\frac{\|\omega^M_{\mathrm{IW}} - \omega^M_{\mathrm{Exp}}\|}{\omega^M_{\mathrm{Exp}}}$ [%]	2.91906711	5.38763974	2.51585411
$\frac{\|\alpha^M_{2,\mathrm{IW}} - \alpha^M_{2,\mathrm{Exp}}\|}{\alpha^M_{2,\mathrm{Exp}}}$ [%]	0.56329861	0.75750294	0.84060225
k_s	1.0117741	1.01588824	1.01764736
$\frac{\|\alpha^{M,\mathrm{mod}}_{2,\mathrm{IW}} - \alpha^M_{2,\mathrm{Exp}}\|}{\alpha^M_{2,\mathrm{Exp}}}$ [%]	0.12881935	0.17132023	0.18656431
L'_z/C	0.20870667	0.20786145	0.20750214
$\frac{L'_z - L_z}{L_z}$ [%]	-1.16370795	-1.56397498	-1.73413268

mechanism. Quantitatively, the mixed-out exit angle error is reduced by $\approx 80\%$ in all cases, as reported in Table 6.4. The streamtube contraction estimated to yield the correct mass flow rate is between $1 - 2\%$ and it is larger at higher Reynolds number.

6.6.6 Phase-locked losses

The previous section presented a time-averaged estimation of wake traverses and losses. However, in the context of the wake passing effect, the unsteady effects are dominant and hidden in the time-averaging process. There exists extensive literature on the effect of discrete incoming disturbances on losses in LPTs [109, 108, 92, 57] and compressors [90]. The classical Denton loss analysis [38] was shown to be accurate only in cases without wakes [109]: when the incoming wakes remain discrete, the discrepancy between mixed-out total pressure losses and total Denton loss was found to be maximum, and the difference is

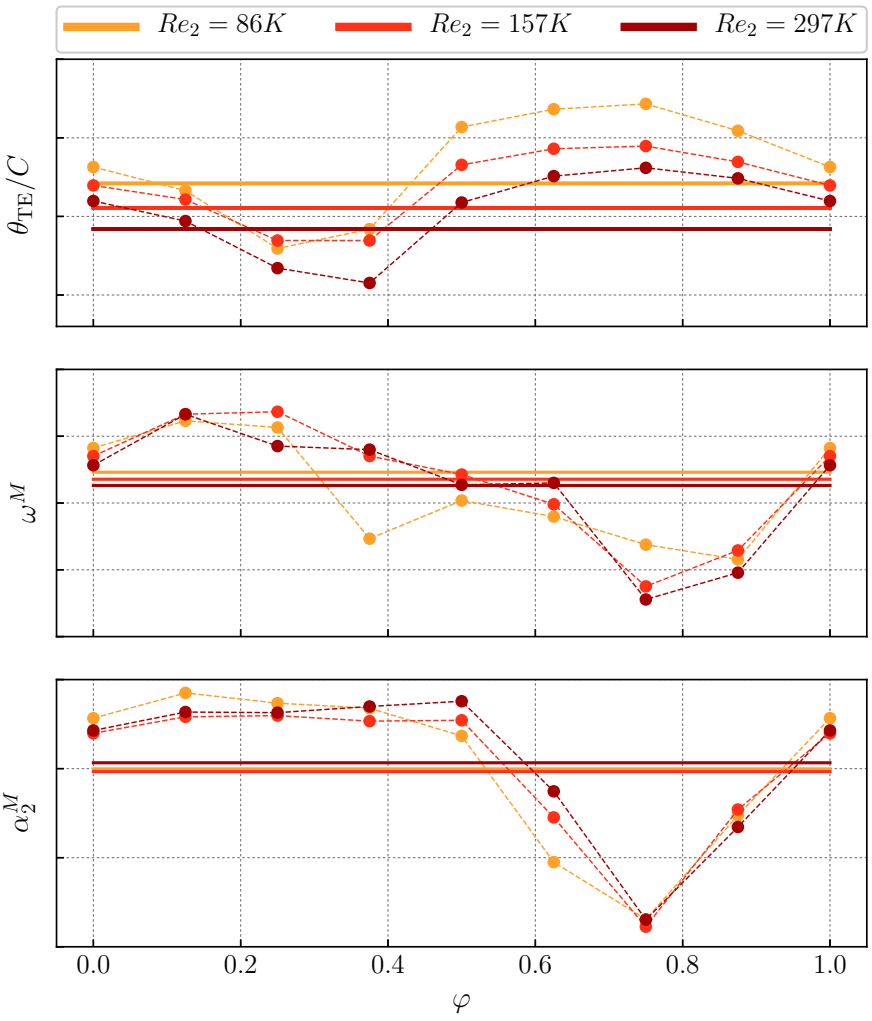

Figure 6.25: Phase-average (top) trailing edge momentum thickness, (middle) mixed-out loss and (bottom) exit angle. Horizontal lines indicate time-averaged values. Y-axis tick labels are omitted due to data sensitivity.

mainly attributable to distortion losses due to variable area mixing within the blade passage. In other words, loss analysis that depends on the state of the boundary layer fails to capture the multiple loss generation mechanisms in play. Figure 6.25 shows the behaviour of phase-locked quantities related to losses (top and mid-figure), reporting the time-averaged values as a comparison metric. The impact of incoming wakes is extremely significant: both boundary layer thickness at the trailing edge and mixed-out loss coefficient incur significant variations throughout the phase cycle. The lag that exists between the boundary layer state and mixed-out loss is due to the distance between the TE and the measurement traverse where losses are sampled: the wake impacting the aft portion of the suction surface at $\varphi \approx 0.3$ reaches the traverse at $\varphi \approx 0.75$. For most of the phase cycle, there exists a vertical shift in the trailing edge momentum thickness, between the low Reynolds and high Reynolds numbers. The only exception is found in correspondence

of the wake passing on the trailing edge at $\varphi \approx 0.3$, where similar θ_{TE} are measured in cases $Re_2 = 86K$ and $Re_2 = 157K$, followed by larger recovery in the low Reynolds number case. At this stage, profile losses associated with the boundary layer trailing edge thickness reach the lowest point. When the impact of bubble suppression is measured at the downstream wake traverse ($\varphi \approx 0.75$) the loss coefficient is reduced to negative values, and the flow turning is increased by almost 3%.

6.7 Discussion

This chapter discussed the introduction of the wake passing effect on a LPT cascade, focusing in particular on the Reynolds sensitivity of the flow in the suction surface trailing edge region, and introducing an extensive comparison of wake measurements with experiments. The three Reynolds numbers analysed correspond to the flow regimes discussed in Chapter 5, representative of subcritical, critical and supercritical states of the suction surface boundary layer.

In order to assess the accuracy of the representation of the bars (implemented by leveraging a high-order formulation of the Smoothed Profile Method, SPM), a set of auxiliary stationary cylinder simulations were performed on a computational mesh that mimics the resolution of the LPT setup. The interface thickness was selected to ensure accurate representation of the SPM ϕ function (and it was thus driven by resolution requirements), but not necessarily producing physical results in the wake. Therefore, the auxiliary study focused on varying the diameter of the cylinders at fixed interface thickness. A smaller bar diameter of $d_{\mathrm{SPM}} = 0.6d$ was shown to produce wake profiles and spectral characteristics that very closely match those of a DNS simulation over the range of Reynolds numbers analysed. Therefore, the smaller modified diameter was adopted to accurately represent the bar passing.

The analysis of the LPT flow features first focused on instantaneous statistics. Despite the difficulty in tracking the time- and space-evolution of boundary layer instabilities without the use of more advanced post-processing analytical tools, the aft portion of the suction surface was shown to feature a range of different transition mechanisms, driven by the periodic impingment of the wake passing. At low Reynolds number, two roll-up regions are periodically identified where the wake suppresses the suction surface separation bubble, but at higher Reynolds number this mechanism is not retained. This highlights the high sensitivity of the suction surface separation to external disturbances in the subcritical flow regime, which in turn has important implications on the cascade losses.

The comparison of wake traverses from the numerical simulations demonstrates extremely accurate agreement with experimental measurements, with significant improvements with respect to the clean inflow results of Chapter 5. The introduction of background turbulence alongside with wake passing does not produce appreciable effects in the wake statistics, highlighting the dominant effect of discrete disturbances on the boundary layer and wake dynamics. The discrepancies in the axial velocity traverses are owed to the

experimental streamtube contraction, which is not modelled in the simulations. A first-order estimation suggests that a spanwise contraction of the spanwise domain between 1% and 2% (inversely proportional to the Reynolds number) would adjust the mass-flow rate and achieve the same level of accuracy in the axial velocity component u/U_2 as in the other wake profiles discussed. Minor discrepancies are observed on the pressure side of the fluctuating TKE statistics, which are damped at higher Reynolds number.

Extremely accurate total pressure loss prediction is achieved relative to experiments, with the numerical results within measurement uncertainty. Comparison of the loss coefficient in the presence of wake passing and body forcing (i.e. the clean inflow formulation with leading edge perturbation, introduced in Chapter 4 and applied to three Reynolds numbers in Chapter 5) confirms that the loss regime strongly depends on the state of the suction surface separation bubble: in the subcritical regime, the periodic suppression of the open separation bubble results in significant loss reduction. In the supercritical regime ($Re_2 \approx 300K$), where the boundary layer is attached and fully turbulent at the trailing edge, the experiments highlight a small increase in losses due to the presence of periodic wakes, while numerical experiments predict very similar values to the case $Re_2 = 157K$.

Although the approach presented in the Chapter yields very accurate representation of the flow physics demonstrated by the close agreement with experimental measurements, a number of limitations can be outlined, highlighting potential future research avenues to enhance the robustness of the methodology here introduced:

- The diameter scaling adopted to accurately model the bar passing is specifically tuned for the resolution levels adopted in this study. This approach should be generalised to simply relate local resolution levels and Reynolds number to a modified bar diameter with equivalent characteristics to a well-resolved DNS simulation of a cylinder at that flow regime.

- The use of periodic boundary conditions in the pitchwise direction introduces an inherent limitation in the bar distances that can be simulated, which always need to be an exact dividend of the pitch. A bar distance larger than P_y could be simulated by replicating the blade multiple times, thus multiplying the computational cost.

- As previously discussed in Chapter 5, the accuracy in the exit mass flux prediction can be improved by accounting for the streamtube contraction. This is an important effect because in turn the boundary layer behaviour in the trailing edge region may be significantly affected; future research efforts should address this issue.

- In the longer term, the introduction of sliding planes [72, 46] is a necessary technological requirement to enable simulations of multiple rotor-stator stages. However, for conformal discretisations the presence of a geometric discontinuity (e.g. hanging nodes) in turn creates a geometric incompatibility because C^0 continuity cannot be enforced across multiple elements [77]. The natural solution would be the adoption

of a compressible formulation via Discontinuous Galerkin projection, as opposed to Continuous Galerkin. The choice of numerical methodology adopted to represent the sliding planes poses further challenges, but it is an important stepping stone in building the capability towards large-scale high-fidelity simulations of rotating engine components.

Chapter 7

Conclusions and future directions

This thesis presented an extensive study of the first application of the spectral/*hp* element methods implemented in the *Nektar++* framework to low pressure turbine cascades, in a range of flow conditions and Reynolds numbers. The aim of the project was that of establishing the grounding knowledge for the use of high-fidelity CFD in industrial setting, aiding the design process and taking the role of a virtual wind tunnel. Throughout the project, a large valuable database was generated, leveraging both the Imperial College Research Computing Service [1] and the UK National Supercomputing Centre (ARCHER), the latter through cumulative allocations of about 5 million core hours.

7.1 General conclusions

7.1.1 Mesh resolution

Approaching the task of simulating a low pressure turbine vane with the use of spectral/*hp* element methods for the first time, the initial important requirement was to establish the necessary level of resolution to accurately capture the flow physics of interest. This was the topic of Chaper 3, which presented an extensive resolution study focusing on the effect of different numerical setups on the main performance indicators. The polynomial order in the spectral/*hp* element planes was shown to be a powerful tool to achieve order-of-magnitude convergence, especially in the context of clean inflow simulations where the flow is predominantly two-dimensional. With simple boundary conditions, the use of spatially adaptive polynomial order was shown to drastically reduce the computational cost while retaining high accuracy levels. However, it was not further investigated in the rest of the thesis due to the difficulty of extending this approach to retain its algorithmic advantage in the context of time-varying inflow boundary conditions. The numerical setup in the spanwise direction (i.e. resolution and domain extent) was shown to have little effect on commonly analysed blade wall distributions, useful notion to enable first order estimations of the flow physics at low computational cost.

7.1.2 Stochastic and deterministic unsteadiness

The lack of realistic inflow boundary conditions in Chapter 3 naturally suggested the avenues of investigation presented in Chapters 4 and 6, respectively focusing on the development of a validated capability to model stochastic and deterministic unsteadiness.

For the generation of inflow turbulence, no general consensus exists on which numerical approach is best for the specific target application. As such, two different methodologies were explored: the introduction of localised momentum forcing near the leading edge, and synthetic inflow turbulence via Random Fourier Method. For each approach, advantages and disadvantages were discussed, concurrently highlighting the difficulty of developing a robust, physical and computationally efficient formulation.

The approach adopted to model the wake passing effect was an immersed boundary method called Smoothed Profile Method, validated in length in the context of the flow past a circular cylinder. The exceptional agreement of the LPT results with experimental measurements was well within uncertainty levels, supporting the methodology here developed for use in Industry with realistic geometries. The wake passing effect was shown to dominate the flow physics when analysed in conjunction with inflow turbulence, highlighting the relative impact of the two sources of unsteadiness on the loss production mechanism.

7.1.3 The flow physics across the operating range

In order to validate the accuracy of the numerical methodology developed in Chapter 3 across the entire operating range, a set of experiments was designed to analyse the blade performance in the subcritical, critical and supercritical regimes with clean inflow and momentum forcing near the leading edge (in Chapter 5), and wake passing (in Chapter 6). In the latter case, excellent agreement with experiments was verified at every flow regime, while in absence of inflow disturbances the dynamics of the separated shear layer was not accurately captured at subcritical Reynolds number, with repercussions on the measured wake traverses. The main discrepancies were attributed to the lack of modelling of the experimental streamtube contraction, indicated below as one of the main future avenues of investigation.

7.2 Limitations and recommendations for future work

The complexity of accurately modelling the flow past a turbine blade under realistic flow conditions poses a number of significant challenges, highlighted in detail through the thesis. Not all avenues of investigation could be explored through the course of this PhD project, thus a number of limitations can be identified, and used to provide important indications that may inform future research directions.

Three main limitations of the approach adopted through this work are identified, not directly related to the content of any specific chapter, but rather arising from the modelling approach itself:

1. Through the thesis the incompressible Navier-Stokes solver was adopted, as opposed to a compressible formulation. At the time of writing, the incompressible solver is the most mature and most highly opimised solver within the *Nektar++* software framework: the execution time of various compressible and incompressible Navier-Stokes solvers implemented in *Nektar++* was measured in the context of the flow past a circular cylinder [182]. The simulation time of the explicit compressible solver was found to be over two orders of magnitude larger than the incompressible Quasi-3D solver. The selection of the incompressible formulation was made on the basis that the more efficient formulation would enable a first thorough investigation of a wide range of phenomena without introducing an excessively constraining hypothesis. In fact, LPTs generally experience significantly lower Mach and Reynolds numbers compared to HPTs and HP compressors. The experimental data presented in this work for validation purposes was collected under very low Mach numbers, supporting the validity of this approximation. However, despite the low speed of these geometries, at design condition the flow speed at peak suction may reach high subsonic values. This was investigated by Michelassi et al. [111], who compared compressible URANS with incompressible DNS and LES. The incompressible simulations were shown to overpredict the strength and the extent of the adverse pressure gradient acting on the suction surface after peak suction, causing an overestimation of the losses. The low Mach number (but not quasi-incompressible) nature of the flow is peculiar of direct-drive fan architectures, while geared fan architectures require adopting LPTs designed with a reduced number of stages at lower work coefficient and higher Mach number, thus highlighting the importance of a validated high-fidelity compressible capability. It would be instructive to perform a comparative DNS study of compressible and incompressible formulations in a benchmark test case for a range of Mach numbers.

   However, as discussed in Chapter 5, the use of redesign techniques allows to transform blade shape and flow conditions so that experimental testing can be carried out in low-speed wind tunnels. Therefore, the use of an incompressible approach is not inherently a limitation: the validated incompressible DNS capability presented in this thesis remains applicable in a wider range of cases, as a preliminary design and analysis tool for low-speed turbines and the back stages of compressors, as well as their interaction with cavities (which are typically characterised by very low Mach number).

2. The limit in computational resources available often introduced constraints in the spanwise extent that could be simulated. This consideration is particularly valid at low Reynolds numbers and in two scenarios. First, when considering clean inflow boundary conditions, the structures formed on the suction surface are not perturbed by external disturbances, and therefore present weaker waviness in the spanwise direction when compared to simulations that include inflow disturbances. This is

supported by the two-point correlations presented in Chapter 3, showing that a spanwise domain of at least $L_z = 0.4C$ might be required to fully contain the turbulent structures developing in the transitioning separated shear layer. However, the clean inflow scenario is more important for validation and verification activities rather than realistic simulations looking to provide insight into the flow physics. The second case is that of inflow turbulence simulations as discussed in Chapter 4, where the computational length scale must be chosen so that the largest structures are contained in the domain. However, in order to simulate a realistic turbulent length scale of 20%, Pichler et al. [134] showed that a spanwise domain of $L_z = 0.8C_{ax}$ is required. When a more constraining spanwise domain is adopted, the turbulence length scale must be consequently adjusted, and this affects the rate of decay in the development region, which is increased with inverse proportionality to the length scale and direct proportionality to the turbulence intensity. This argument was found to be less problematic in the context of the wake passing effect, where the lengthscales generated in the frame of reference of the moving bars are generally well-contained in the spanwise domain.

3. The simulations presented in this thesis did not introduce a systematic approach to model the streamtube effect. The topic was discussed in greater detail in Sections 5.2.1 and 6.6.5. This effect is important when comparing with rig data, and it is deemed to be one of the most likely causes of discrepancy in the wake profiles and exit angle estimation. Introducing a methodology to model this effect from rig data would allow for further improvements in accuracy when comparing numerical and experimental results, especially at low Reynolds numbers.

A number of further research avenues and opportunities for methodology development can be identified, at times arising from the limitations of this work but also as a continuation of the efforts presented in the thesis. The are presented below as a list of topics, in no particular order of importance:

- The resolution study of Chapter 3 was concerned with only one flow condition. The outcome of the study was used to determine the numerical setup adopted through the following Chapters, wholly dedicated to the analysis of the flow physics. As such, an implicit assumption was made to extend the setup discussed with clean inflow at low to moderate Reynolds number to higher flow regimes. Therefore, the resolution study presented in Chapter 3 (or part of it) should be repeated in a number of additional flow configurations. In particular, a more realistic physical setup should be adopted (i.e. including wake passing disturbances), and the extremes of the operating envelope should be explored to allow the construction of a more complete map of the resolution requirements. Additional metrics based on turbulence measurements (i.e. turbulence production and dissipation, as previously highlighted by Pichler et al. [135]) should be adopted, enabled by the introduction of the combined filter and post-processing tool described in Appendix A into *Nektar++*. Furthermore, in the current

study the resolution in the $x-y$ planes was only explored by means of P-refinement, building on extensive prior theoretical findings supporting the algorithmic advantage of higher order spatial discretisations. The analysis of combinations of base mesh resolution and polynomial orders at fixed DoF could potentially confirm the approach presented as the most efficient one, or rather indicate any better alternatives.

- Additional research efforts could be invested in further developing the inflow turbulence capability discussed in Chapter 4. Concerning the body forcing approach, future efforts could establish a direct connection between the intensity levels of the forcing, and equivalent turbulence intensity levels introduced via synthetic inflow turbulence. Regarding the latter approach, this thesis only investigated the effect of homogeneous isotropic fluctuations on the blade statistics. The capability for generating anisotropic fluctuations could be introduced and validated in the context of realistic industrial geometries.

- All simulations in this work were concerned with the midspan section, because in blades with high aspect ratio the majority of stage losses are generated in this section. Although the role of endwalls in LPTs is relatively minor compared to HPTs [144], it is worthy of investigation. Due to the three-dimensional nature of the problem and the added cost of simulating the whole span, only few studies are currently available on the topic. Cui and Tucker [30] carried out the first endwall study using LES, and observing its interaction with purge flows as well as its impact on loss generation. More recently, in the first comparative endwall LES and RANS study, the whole span was simulated analysing the dynamics of secondary flows [136] and the associated loss generation mechanisms [100]. The combination of full-span and sliding planes is an important avenue of investigation that will enable to advance the understanding of rotor-stator interaction on endwall flows. Within the *Nektar++* framework, the Quasi-3D incompressible solver was estimated to run approximately an order of magnitude faster than the full 3D solver. Therefore, increased availability of compute power is required to tackle the analysis of three-dimensional flows while retaining the same level of accuracy as discussed in this thesis.

- On the side of methodology development, all simulations in this work adopted a direct parallel system solve in conjunction with hybrid parallelisation, which was shown to very efficient for a reasonably low number of partitions in the spectral/hp planes. However, scaling up to larger simulations, in order to retain parallel efficiency it may be necessary to adopt an alternative iterative approach, which will need improvements: for the test cases of interest in this thesis, the performance of the direct solver is significantly higher than the alternative conjugate gradient approach.

- Most of the discussion on technology development in this thesis was concerned with the solver, and relatively less attention was devoted to pre- and post-processing. A number of Python scripts were developed through the project to assist in mesh

generation and data processing and analysis. However, the level of direct assistance required in the various phases of the CFD cycle is still not negligible, and this is an important consideration to be made when adopting a novel framework in industrial setting. As such, some effort could be made to further develop the existing capability into a *Nektar++* toolchain that is specific for turbomachinery research and analysis, to automate, facilitate and simplify meshing and post-processing activities, possibly exploiting the ongoing development of a Python API for *NekMesh* and *FieldConvert*. This could be extended to also address the current level of interaction required to solve a new test case, having previously built a fundamental understanding on how to robustly initialise new geometries and flow conditions. This would also help to minimise the failure modes that an inexperienced user might encounter, despite the fact that the SVV formulation adopted to stabilise the incompressible Navier-Stokes solver was found to work very robustly.

- In recent years the use of machine learning algorithms gained widespread popularity across several disciplines. A widely investigated application in CFD is the use of learning agorithms to inform low order RANS and URANS models, fully enabled by the presence of large data sets generated by high-fidelity DNS simulations. However, there is added value in high-fidelity CFD, which can be extracted through the development of simple 1D surrogate models for preliminary design capability. These correlation models are easy to derive compared to tuning turbulence models, and they can gather the result of virtual tunnel testing (i.e. the capability developed through the thesis) as an addedum to experimental testing, with promising potential to deliver very valuable design indications in Industry.

Bibliography

[1] Imperial College Research Computing Service, DOI: 10.14469/hpc/2232.

[2] B. H. Abe, Y. Mizobuchi, and Y. Matsuo. DNS and modeling of a turbulent boundary layer with separation and reattachment over a range of Reynolds numbers. *CTR Annual Research Briefs*, pages 311–322, 2012.

[3] A. K. Agrawal and R. S. Peckover. Nonuniform grid generation for boundary-layer problems. *Computer Physics Communications*, 19(2):171–178, 1980.

[4] J. Ahrens, B. Geveci, C. Law, J. Ahrens, B. Geveci, and C. Law. ParaView: An End-User Tool for Large Data Visualization, 2005.

[5] P. Andersson, L. Brandt, A. Bottaro, and D. S. Henningson. On the breakdown of boundary layer streaks. *Journal of Fluid Mechanics*, 428:29–60, 2001.

[6] I. Argonne National Laboratory. Nek5000. https://nek5000.mcs.anl.gov/.

[7] I. Babuska and M. Suri. The P and H-P Versions of the Finite Element Method, Basic Principles and Properties. *SIAM Review*, 36(4):578–632, 1994.

[8] C. Bailly and G. Comte-Bellot. Statistical Description of Turbulent Flows. In *Turbulence*, pages 33–50. 2015.

[9] C. Bailly and D. Juvé. A Stochastic Approach To Compute Subsonic Noise Using Linarised Euler's Equations. In *AIAA paper 99-1872*, 1999.

[10] P. Batten, U. Goldberg, and S. Chakravarthy. Interfacing Statistical Turbulence Closures with Large-Eddy Simulation. *AIAA Journal*, 42(3):485–492, 2004.

[11] P. Beaudan and P. Moin. Numerical experiments on the flow past a circular cylinder at sub-critical Reynolds number. Technical Report Report No. TF-62, Stanford University, 1994.

[12] W. Béchara, C. Bailly, P. Lafon, and S. M. Candel. Stochastic approach to noise modeling for free turbulent flows. *AIAA Journal*, 32(3):455–463, 1994.

[13] R. Bhaskaran, F. Jia, G. M. Laskowski, Z. J. Wang, and U. Paliath. Towards High-Order Large Eddy Simulation of Aero-Thermal Flows for Turbomachinery Applications. In *Proceedings of the ASME Turbo Expo*, pages GT2017–63358, Charlotte, USA, 2017.

[14] M. Billson, L. E. Eriksson, and L. Davidson. Jet noise modeling using synthetic anisotropic turbulence. In *10th AIAA/CEAS Aeroacoustics Conference*, volume 3, page 15, 2004.

[15] M. Bolinches-Gisbert, D. C. Robles, R. Corral, and F. Gisbert. Prediction of Reynolds Number Effects on Low-Pressure Turbines Using a High-Order ILES Method. *J. Turbomach.*, 142(3), 2020.

[16] A. Bolis, C. D. Cantwell, D. Moxey, D. Serson, and S. J. Sherwin. An adaptable parallel algorithm for the direct numerical simulation of incompressible turbulent flows using a Fourier spectral/hp element method and MPI virtual topologies. *Computer Physics Communications*, 206:17–25, 2016.

[17] A. Bolis. *Fourier Spectral - hp Element Method : Investigation of Time-Stepping and Parallelisation Strategies*. PhD thesis, 2013.

[18] L. Brandt. The lift-up effect: The linear mechanism behind transition and turbulence in shear flows. *European Journal of Mechanics, B/Fluids*, 47:80–96, 2014, 1403.4657.

[19] M. Breuer. Large eddy simulation of the subcritical flow past a circular cylinder: Numerical and modeling aspects. *International Journal for Numerical Methods in Fluids*, 28(9):1281–1302, 1998.

[20] L. Cambier, S. Heib, and S. Plot. The Onera elsA CFD software: Input from research and feedback from industry. *Mechanics and Industry*, 14(3):159–174, 2013.

[21] C. D. Cantwell, D. Moxey, A. Comerford, A. Bolis, G. Rocco, G. Mengaldo, D. De Grazia, S. Yakovlev, J.-E. Lombard, D. Ekelschot, B. Jordi, H. Xu, Y. Mohamied, C. Eskilsson, B. Nelson, P. Vos, C. Biotto, R. M. Kirby, and S. J. Sherwin. Nektar++: An open-source spectral/hp element framework. *Computer Physics Communications*, 192:205–219, 2015.

[22] C. Carton de Wiart, K. Hillewaert, E. Lorriaux, and G. Verheylewegen. Development of a Discontinuous Galerkin Solver for High Quality Wall-Resolved/Modelled DNS and LES of Practical Turbomachinery Flows on Fully Unstructured Meshes. In *Proceedings of the ASME Turbo Expo*, pages GT2015–43428, Montreal, Canada, 2015.

[23] A. Cassinelli, F. Montomoli, P. Adami, and S. J. Sherwin. High Fidelity Spectral/hp Element Methods for Turbomachinery. In *Proceedings of the ASME Turbo Expo*, pages GT2018–75733, Oslo, Norway, 2018.

[24] A. Cassinelli, H. Xu, F. Montomoli, P. Adami, R. Vazquez Diaz, and S. J. Sherwin. On the Effect of Inflow Disturbances on the Flow Past a Linear LPT Vane Using Spectral/hp Element Methods. In *Proceedings of the ASME Turbo Expo*, pages GT2019–91622, Phoenix, Arizona, USA, 2019.

[25] H. G. Castro and R. R. Paz. A time and space correlated turbulence synthesis method for Large Eddy Simulations. *Journal of Computational Physics*, 235:742–763, 2013.

[26] Cenaero. 5th International Workshop on High Order CFD Methods, 2018.

[27] G. Cicatelli, R. D. Stieger, R. J. Howell, H. P. Hodson, and W. N. Dawes. Turbulence Modelling for Unsteady Flows in Axial Turbines - Wake-blade interference. Technical report.

[28] A. Cimarelli, A. Leonforte, E. De Angelis, A. Crivellini, and D. Angeli. On negative turbulence production phenomena in the shear layer of separating and reattaching flows. *Physics Letters, Section A: General, Atomic and Solid State Physics*, 383(10):1019–1026, 2019.

[29] J. D. Coull and H. P. Hodson. Unsteady boundary-layer transition in low-pressure turbines. *Journal of Fluid Mechanics*, 681(1):370–410, 2011.

[30] J. Cui and P. Tucker. Numerical study of purge and secondary flows in a low-pressure turbine. *J. Turbomach.*, 139(2):1–10, 2017.

[31] N. A. Cumpsty, Y. Dong, and Y. S. Li. in the Presence of. In *Proceedings of the ASME Turbo Expo*, pages 95–GT–443, 1995.

[32] N. A. Cumpsty and J. H. Horlock. Averaging Nonuniform Flow for a Purpose. *J. Turbomach.*, 128(1):120, 2006.

[33] E. M. Curtis, H. P. Hodson, M. R. Banieghbal, J. D. Denton, R. J. Howell, and N. W. Harvey. Development of Blade Profiles for Low-Pressure Turbine Applications. *J. Turbomach.*, 119:531–538, 1997.

[34] T. Dairay, E. Lamballais, S. Laizet, and J. C. Vassilicos. Numerical dissipation vs. subgrid-scale modelling for large eddy simulation. *Journal of Computational Physics*, 337:252–274, 2017.

[35] L. Davidson. Using isotropic synthetic fluctuations as inlet boundary conditions for unsteady simulations. *Advances and Applications in Fluid Mechanics*, 1(1):1–35, 2007.

[36] L. Davidson and M. Billson. Hybrid LES-RANS using synthesized turbulent fluctuations for forcing in the interface region. *International Journal of Heat and Fluid Flow*, 27(6):1028–1042, 2006.

[37] D. De Grazia, G. Mengaldo, D. Moxey, P. E. Vincent, and S. J. Sherwin. Connections between the discontinuous Galerkin method and high-order flux reconstruction schemes. *International Journal for Numerical Methods in Fluids*, 75(12):860–877, aug 2014.

[38] J. D. Denton. The 1993 IGTI Scholar Lecture: Loss Mechanisms in Turbomachines. *J. Turbomach.*, 115(4):621–656, 1993.

[39] L. di Mare, M. Klein, W. P. Jones, and J. Janicka. Synthetic turbulence inflow conditions for large-eddy simulation. *Physics of Fluids*, 18(2), 2006.

[40] D. Dietzel, D. Messig, F. Piscaglia, A. Montorfano, G. Olenik, O. T. Stein, A. Kronenburg, A. Onorati, and C. Hasse. Evaluation of scale resolving turbulence generation methods for Large Eddy Simulation of turbulent flows. *Computers and Fluids*, 96:116–128, 2014.

[41] S. Dong, G. E. Karniadakis, and C. Chryssostomidis. A robust and accurate outflow boundary condition for incompressible flow simulations on severely-truncated unbounded domains. *Journal of Computational Physics*, 261:83–105, 2014.

[42] S. Dong, G. E. Karniadakis, A. Ekmekci, and D. Rockwell. A combined direct numerical simulation–particle image velocimetry study of the turbulent near wake. *Journal of Fluid Mechanics*, 569:185, 2006.

[43] P. Durbin and X. Wu. Transition Beneath Vortical Disturbances. *Annual Review of Fluid Mechanics*, 39(1):107–128, 2007.

[44] D. Ekelschot, D. Moxey, S. J. Sherwin, and J. Peiró. A p-adaptation method for compressible flow problems using a goal-based error indicator. *Computers and Structures*, 181:55–69, 2017.

[45] P. Fernandez, N. C. Nguyen, and J. Peraire. The hybridized Discontinuous Galerkin method for Implicit Large-Eddy Simulation of transitional turbulent flows. *Journal of Computational Physics*, 336:308–329, 2017.

[46] E. Ferrer and R. H. Willden. A high order Discontinuous Galerkin - Fourier incompressible 3D Navier-Stokes solver with rotating sliding meshes. *Journal of Computational Physics*, 231(21):7037–7056, 2012.

[47] J. Franke and W. Frank. Large eddy simulation of the flow past a circular cylinder at Re=3900. *Journal of Wind Engineering and Industrial Aerodynamics*, 90:1191–1206, 2002.

[48] A. Garai, L. Diosady, S. Murman, and N. Madavan. DNS of Flow In a Low-Pressure Turbine Cascade Using a Discontinuous-Galerkin Spectral-Element Method. In *Proceedings of the ASME Turbo Expo*, pages GT2015–42773, Montreal, Canada, 2015.

[49] A. Garai, L. T. Diosady, S. M. Murman, and N. K. Madavan. DNS of Low-Pressure Turbine Cascade Flows With Elevated Inflow Turbulence Using a Discontinuous-Galerkin Spectral-Element Method. In *Proceedings of the ASME Turbo Expo*, pages GT2016–56700, Seoul, South Korea, 2016.

[50] A. Garai, L. T. Diosady, S. M. Murman, and N. K. Madavan. Scale-Resolving Simulations of Bypass Transition in a High-Pressure Turbine Cascade Using a Spectral Element Discontinuous Galerkin Method. *J. Turbomach.*, 140(3):031004, 2018.

[51] N. J. Georgiadis, D. P. Rizzetta, and C. Fureby. Large-Eddy Simulation: Current Capabilities, Recommended Practices, and Future Research. *AIAA Journal*, 48(8):1772–1784, 2010.

[52] M. Giles. UNSFLO: A Numerical Method For The Calculation Of Unsteady Flow In Turbomachinery. Technical report, 1991.

[53] M. Giovannini, M. Marconcini, F. Rubechini, A. Arnone, and F. Bertini. Scaling three-dimensional low-pressure turbine blades for low-speed testing. *Journal of Turbomachinery*, 138(11):1–9, 2016.

[54] D. Goldstein, R. Handler, and L. Sirovich. Modeling a No-Slip Flow Boundary with an External Force Field. *Journal of Computational Physics*, 105(2):354–366, 1993.

[55] D. Gottlieb and S. A. Orszag. *Numerical analysis of spectral methods: theory and applications*. SIAM-CMBS, 1977.

[56] D. E. Halstead, D. C. Wisler, T. H. Okiishi, G. J. Walker, H. P. Hodson, and H. W. Shin. Boundary Layer Development in Axial Compressors and Turbines: Part 1 of 4 - Composite Picture. *J. Turbomach.*, 119(January):114–127, 1997.

[57] F. Hammer, N. Sandham, and R. Sandberg. The Influence of Different Wake Profiles on Losses in a Low Pressure Turbine Cascade. *International Journal of Turbomachinery, Propulsion and Power*, 3(2):10, apr 2018.

[58] M. Harrower and C. A. Brewer. ColorBrewer.org: An online tool for selecting colour schemes for maps. *Cartographic Journal*, 40(1):27–37, 2003.

[59] L. Hilgenfeld and M. Pfitzner. Unsteady Boundary Layer Development Due to Wake Passing Effects on a Highly Loaded Linear Compressor Cascade. *J. Turbomach.*, 126(4):493–500, 2004.

[60] H. P. Hodson and R. J. Howell. Bladerow Interactions, Transition, and High-Lift Aerofoils in Low-Pressure Turbines. *Annual Review of Fluid Mechanics*, 37(1):71–98, 2005.

[61] S. Hoyas and J. Jiménez. Reynolds number effects on the Reynolds-stress budgets in turbulent channels. *Physics of Fluids*, 20(10), 2008.

[62] S. H. Huang, Q. S. Li, and J. R. Wu. A general inflow turbulence generator for large eddy simulation. *Journal of Wind Engineering and Industrial Aerodynamics*, 98(10-11):600–617, 2010.

[63] J. C. Hunt and P. A. Durbin. Perturbed vortical layers and shear sheltering. *Fluid Dynamics Research*, 24(6):375–404, 1999.

[64] H. T. Huynh. A flux reconstruction approach to high-order schemes including discontinuous Galerkin methods. *Collection of Technical Papers - 18th AIAA Computational Fluid Dynamics Conference*, 1(June):698–739, 2007.

[65] H. T. Huynh, Z. J. Wang, and P. E. Vincent. High-order methods for computational fluid dynamics: A brief review of compact differential formulations on unstructured grids. *Computers and Fluids*, 98:209–220, 2014.

[66] A. S. Iyer, F. D. Witherden, S. I. Chernyshenko, and P. E. Vincent. Identifying eigenmodes of averaged small-amplitude perturbations to turbulent channel flow. *Journal of Fluid Mechanics*, 875:758–780, 2019.

[67] C. T. Jacobs, S. P. Jammy, and N. D. Sandham. OpenSBLI: A framework for the automated derivation and parallel execution of finite difference solvers on a range of computer architectures. *Journal of Computational Science*, 18:12–23, 2017, 1609.01277.

[68] R. G. Jacobs and P. A. Durbin. Simulations of bypass transition. *Journal of Fluid Mechanics*, 428(2001):185–212, 2001.

[69] R. G. Jacobs and P. A. Durbin. Shear sheltering and the continuous spectrum of the Orr-Sommerfeld equation. *Physics of Fluids*, 10(8):2006–2011, 1998.

[70] N. Jarrin, S. Benhamadouche, D. Laurence, and R. Prosser. A synthetic-eddy-method for generating inflow conditions for large-eddy simulations. *International Journal of Heat and Fluid Flow*, 27(4):585–593, 2006.

[71] J. Jeong and F. Hussain. On the identification of a vortex. *Journal of Fluid Mechanics*, 285:69–94, 1995.

[72] R. Johnstone, L. Chen, and R. D. Sandberg. A sliding characteristic interface condition for direct numerical simulations. *Computers and Fluids*, 107:165–177, 2015.

[73] G. Kalitzin, X. Wu, and P. A. Durbin. DNS of fully turbulent flow in a LPT passage. *International Journal of Heat and Fluid Flow*, 24(4):636–644, 2003.

[74] G.-S. Karamanos and G. E. Karniadakis. A Spectral Vanishing Viscosity Method for Large-Eddy Simulations. *Journal of Computational Physics*, 163(1):22–50, 2000.

[75] G. E. Karniadakis. Spectral element-Fourier methods for incompressible turbulent flows. *Computer Methods in Applied Mechanics and Engineering*, 80(1-3):367–380, 1990.

[76] G. E. Karniadakis, M. Israeli, and S. A. Orszag. High-order splitting methods for the incompressible Navier-Stokes equations. *Journal of Computational Physics*, 97(2):414–443, 1991.

[77] G. E. Karniadakis and S. J. Sherwin. *Spectral/hp Element Methods for Computational Fluid Dynamics*. Oxford University Press, Oxford, 2nd edition, 2005.

[78] M. Karweit, P. Blanc-benon, D. Juvé, and G. Comte-Bellot. Simulation of the propagation of an acoustic wave through a turbulent velocity field : A study of phase variance. *The Journal of the Acoustical Society of America*, 89(1):52–62, 1991.

[79] A. M. Kempf, S. Wysocki, and M. Pettit. An efficient, parallel low-storage implementation of Klein's turbulence generator for LES and DNS. *Computers and Fluids*, 60:58–60, 2012.

[80] A. Kempf, M. Klein, and J. Janicka. Efficient generation of initial- and inflow-conditions for transient turbulent flows in arbitrary geometries. *Flow, Turbulence and Combustion*, 74(1):67–84, 2005.

[81] J. Kim, P. Moin, and R. Moser. Turbulence statistics in fully developed channel flow at low reynolds number. *Journal of Fluid Mechanics*, 177:133–166, 1987.

[82] Y. Kim, I. P. Castro, and Z. T. Xie. Divergence-free turbulence inflow conditions for large-eddy simulations with incompressible flow solvers. *Computers and Fluids*, 84:56–68, 2013.

[83] R. M. Kirby and S. J. Sherwin. Stabilisation of spectral/hp element methods through spectral vanishing viscosity: Application to fluid mechanics modelling. *Computer Methods in Applied Mechanics and Engineering*, 195(23-24):3128–3144, 2006.

[84] M. Klein, A. Sadiki, and J. Janicka. A digital filter based generation of inflow data for spatially developing direct numerical or large eddy simulations. *Journal of Computational Physics*, 186(2):652–665, 2003.

[85] R. H. Kraichnan. Diffusion by a Random Velocity Field. *The Physics of Fluids*, 13(1):22–31, 1970.

[86] A. G. Kravchenko and P. Moin. Numerical studies of flow over a circular cylinder at ReD=3900. *Physics of Fluids*, 12(2):403–417, 2000.

[87] G. M. Laskowski, J. Kopriva, V. Michelassi, S. Shankaran, U. Paliath, R. Bhaskaran, Q. Wang, C. Talnikar, Z. J. Wang, and F. Jia. Future Directions of High Fidelity

CFD for Aerothermal Turbomachinery Analysis and Design. In *46th AIAA Fluid Dynamics Conference*, pages 1–30, Washington, D.C., USA, 2016.

[88] B. J. Lazaro, E. Gonzalez, and R. Vazquez. Unsteady Loss Production Mechanisms in Low Reynonlds Number, High Lift, Low Pressure Turbine Profiles. In *Proceedings of the ASME Turbo Expo*, pages GT2007–28142, 2007.

[89] S. Lee, S. K. Lele, and P. Moin. Simulation of spatially evolving turbulence and the applicability of Taylor's hypothesis in compressible flow. *Physics of Fluids A: Fluid Dynamics*, 4(7):1521–1530, 1992.

[90] J. Leggett, S. Priebe, A. Shabbir, V. Michelassi, R. Sandberg, and E. Richardson. Loss prediction in an axial compressor cascade at off- design incidences with free stream disturbances using large eddy simulation. *J. Turbomach.*, 140(7), 2018.

[91] O. Lehmkuhl, I. Rodríguez, R. Borrell, and A. Oliva. Low-frequency unsteadiness in the vortex formation region of a circular cylinder. *Physics of Fluids*, 25(8), 2013.

[92] D. Lengani, D. Simoni, M. Ubaldi, P. Zunino, F. Bertini, and V. Michelassi. Accurate Estimation of Profile Losses and Analysis of Loss Generation Mechanisms in a Turbine Cascade. *J. Turbomach.*, 139(12):121007, 2017.

[93] Y. Li, S. Premasuthan, and A. Jameson. Comparison of h- and p- Adaptations for Spectral Difference Methods. (July):1–15, 2010.

[94] Y. Liu, T. A. Zaki, and P. A. Durbin. Floquet analysis of secondary instability of boundary layers distorted by Klebanoff streaks and Tollmien-Schlichting waves. *Physics of Fluids*, 20(12), 2008.

[95] L. M. Lourenco and C. Shih. Characteristics of the plane turbulent near wake of a circular cylinder, a particle image velocimetry study. 1993.

[96] Y. Lu and W. N. Dawes. High Order Large Eddy Simulations for a Transonic Turbine Blade using Hybrid Unstructured Meshes. In *Proceedings of the ASME Turbo Expo*, pages GT2015–42283, Montreal, Canada, 2015.

[97] X. Luo, M. R. Maxey, and G. E. Karniadakis. Smoothed profile method for particulate flows: Error analysis and simulations. *Journal of Computational Physics*, 228(5):1750–1769, 2009.

[98] X. Ma, G.-S. Karamanos, and G. E. Karniadakis. Dynamics and low-dimensionality of a turbulent near wake. *Journal of Fluid Mechanics*, 410(May):29–65, 2000.

[99] Y. Maday, S. M. O. Kaber, and E. Tadmor. Legendre pseudospectral viscosity method for nonlinear conservation laws. *SIAM Journal on Numerical Analysis*, 30:321–342, 1993.

[100] M. Marconcini, R. Pacciani, A. Arnone, V. Michelassi, R. Pichler, Y. Zhao, and R. D. Sandberg. LES and RANS Analysis of the End-Wall Flow in a Linear LPT Cascade: Part II — Loss Generation. In *Proceedings of the ASME Turbo Expo*, pages GT2018–76450, Oslo, Norway, jun 2018. American Society of Mechanical Engineers.

[101] M. Marconcini, F. Rubechini, R. Pacciani, A. Arnone, and F. Bertini. Redesign of High-Lift Low Pressure Turbine Airfoils for Low Speed Testing. *Journal of Turbomachinery*, 134(5):1–8, 2012.

[102] J. Marty, N. Lantos, B. Michel, and V. Bonneau. LES and Hybrid RANS/LES Simulations of Turbomachinery Flows Using High Order Methods. In *Proceedings of the ASME Turbo Expo*, pages GT2015–42134, Montreal, Canada, 2015.

[103] G. Medic and O. P. Sharma. Large-eddy simulation of flow in a low-pressure turbine cascade. In *Proceedings of the ASME Turbo Expo*, pages GT2012–68878, Copenhagen, 2012.

[104] G. Medic, V. Zhang, G. Wang, J. Joo, and O. P. Sharma. Prediction of transition and losses in compressor cascades using large-eddy simulation. *J. Turbomach.*, 138(12):1–9, 2016.

[105] G. Mengaldo, D. De Grazia, P. E. Vincent, and S. J. Sherwin. On the Connections Between Discontinuous Galerkin and Flux Reconstruction Schemes: Extension to Curvilinear Meshes. *Journal of Scientific Computing*, 67(3):1272–1292, 2016.

[106] G. Mengaldo, R. C. Moura, B. Giralda, J. Peiró, and S. J. Sherwin. Spatial eigensolution analysis of discontinuous Galerkin schemes with practical insights for under-resolved computations and implicit LES. *Computers and Fluids*, 169:349–364, 2018.

[107] G. Mengaldo, D. De Grazia, D. Moxey, P. E. Vincent, and S. J. Sherwin. Dealiasing techniques for high-order spectral element methods on regular and irregular grids. *Journal of Computational Physics*, 299:56–81, 2015.

[108] V. Michelassi, L. Chen, R. Pichler, R. Sandberg, and R. Bhaskaran. High-Fidelity Simulations of Low-Pressure Turbines: Effect of Flow Coefficient and Reduced Frequency on Losses. *J. Turbomach.*, 138(11):111006, 2016.

[109] V. Michelassi, L. Chen, R. Pichler, and R. D. Sandberg. Compressible Direct Numerical Simulation of Low-Pressure Turbines—Part II: Effect of Inflow Disturbances. *J. Turbomach.*, 137(7):071005–041005–12, 2015.

[110] V. Michelassi, J. G. Wissink, J. Fröhlich, and W. Rodi. Large-Eddy Simulation of Flow Around Low-Pressure Turbine Blade with Incoming Wakes. *AIAA Journal*, 41(11):2143–2156, 2003.

[111] V. Michelassi, J. G. Wissink, and W. Rodi. Direct numerical simulation, large eddy simulation and unsteady Reynolds-averaged Navier-Stokes simulations of periodic unsteady flow in a low-pressure turbine cascade: A comparison. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 217(4):403–412, 2003.

[112] V. Michelassi and J. G. Wissink. Turbulent kinetic energy production in the vane of a low-pressure linear turbine cascade with incoming wakes. *International Journal of Rotating Machinery*, 2015(case C), 2015.

[113] F. Mohaghegh and H. S. Udaykumar. Comparison of sharp and smoothed interface methods for simulation of particulate flows I: Fluid structure interaction for moderate reynolds numbers. *Computers and Fluids*, 140:39–58, 2016.

[114] K. Moreland. Diverging Color Maps for Scientific Visualization. In *Proceedings of the 5th International Symposium on Visual Computing*, 2009.

[115] K. Moreland. Why we use bad color maps and what you can do about it. *Human Vision and Electronic Imaging*, pages 262–267, 2016.

[116] R. C. Moura, M. Aman, J. Peiró, and S. J. Sherwin. Spatial eigenanalysis of spectral/hp continuous Galerkin schemes and their stabilisation via DG-mimicking spectral vanishing viscosity for high Reynolds number flows. *Journal of Computational Physics*, 406:109112, 2020.

[117] R. C. Moura, S. J. Sherwin, and J. Peiró. Eigensolution analysis of spectral/hp continuous Galerkin approximations to advection-diffusion problems: Insights into spectral vanishing viscosity. *Journal of Computational Physics*, 307:401–422, 2016, fld.1.

[118] D. Moxey, C. D. Cantwell, Y. Bao, A. Cassinelli, G. Castiglioni, S. Chun, E. Juda, E. Kazemi, K. Lackhove, J. Marcon, G. Mengaldo, D. Serson, M. Turner, H. Xu, J. Peiró, R. M. Kirby, and S. J. Sherwin. Nektar++: Enhancing the capability and application of high-fidelity spectral/hp element methods. *Computer Physics Communications*, 249:107110, 2019.

[119] D. Moxey, C. D. Cantwell, G. Mengaldo, D. Serson, D. Ekelschot, J. Peiró, S. J. Sherwin, and R. M. Kirby. Towards p-Adaptive Spectral/hp Element Methods for Modelling Industrial Flows. In *Spectral and High Order Methods for Partial Differential Equations ICOSAHOM 2016*, volume 119, pages 63–79, 2017.

[120] D. C. Moxey. *Spatio-Temporal Dynamics in Pipe Flow*. PhD thesis, 2011.

[121] F. Naddei, M. de la Llave Plata, V. Couaillier, and F. Coquel. A comparison of refinement indicators for p-adaptive simulations of steady and unsteady flows using

discontinuous Galerkin methods. *Journal of Computational Physics*, 376:508–533, 2019.

[122] Y. Nakayama and R. Yamamoto. Simulation method to resolve hydrodynamic interactions in colloidal dispersions. *Physical Review E*, 71(3):036707, 2005, 0403014.

[123] C. Norberg. An experimental investigation of the flow around a circular cylinder: Influence of aspect ratio. *Journal of Fluid Mechanics*, 258:287–316, 1994.

[124] C. Norberg. Fluctuating lift on a circular cylinder: review and new measurements. *J. Fluids and Structures.*, \bf 17:57–96, 2002.

[125] J. T. Oden. Optimal h-p finite element methods. *Computer methods in applied mech- anics and engineering*, 112(1):309–331, 1994.

[126] L. Ong and J. Wallace. The velocity field of the turbulent very near wake of a circular cylinder. *Experiments in Fluids*, 20(6):441–453, 1996.

[127] J. S. Park, F. D. Witherden, and P. E. Vincent. High-order implicit large-eddy simulations of flow over a NACA0021 aerofoil. *AIAA Journal*, 55(7):2186–2197, 2017.

[128] P. Parnaudeau, J. Carlier, D. Heitz, and E. Lamballais. Experimental and numerical studies of the flow over a circular cylinder at Reynolds number 3900. *Physics of Fluids*, 20(8), 2008.

[129] A. T. Patera. A spectral element method for fluid dynamics: Laminar flow in a channel expansion. *Journal of Computational Physics*, 54(3):468–488, 1984.

[130] A. Peano. Hierarchies of conforming finite elements for plane elasticity and plate bending. *Computers and Mathematics with Applications*, 2(3-4):211–224, 1976.

[131] P. O. Persson and J. Peraire. Sub-cell shock capturing for discontinuous Galerkin methods. In *44th AIAA Aerospace Sciences Meeting*, page 112, 2006.

[132] M. Pianko. Modern Methods of Testing Rotating Components of Turbomachines (Instrumentation). Technical report, 1975.

[133] R. Pichler, V. Michelassi, R. Sandberg, and J. Ong. Highly resolved large eddy simulation study of gap size effect on low-pressure turbine stage. *J. Turbomach.*, 140(2), 2018.

[134] R. Pichler, R. D. Sandberg, G. Laskowski, and V. Michelassi. High-Fidelity Simulations of a Linear HPT Vane Cascade Subject to Varying Inlet Turbulence. In *Proceedings of the ASME Turbo Expo*, pages GT2017–63079, Charlotte, USA, 2017.

[135] R. Pichler, R. D. Sandberg, and V. Michelassi. Assessment of Grid Resolution Requirements for Accurate Simulation of Disparate Scales of Turbulent Flow in Low-Pressure Turbines. In *Proceedings of the ASME Turbo Expo*, pages GT2016–56858, Seoul, South Korea, 2016.

[136] R. Pichler, Y. Zhao, R. Sandberg, V. Michelassi, R. Pacciani, M. Marconcini, and A. Arnone. Large-eddy simulation and RANS analysis of the end-wall flow in a linear low-pressure turbine cascade, Part I: Flow and secondary vorticity fields under varying inlet condition. *J. Turbomach.*, 141(12):1–10, 2019.

[137] T. J. Praisner and J. P. Clark. Predicting transition in turbomachinery - Part I: A review and new model development. *Journal of Turbomachinery*, 129(1):1–13, 2007.

[138] T. J. Praisner, E. A. Grover, M. J. Rice, and J. P. Clark. Predicting transition in turbomachinery - Part II: Model validation and benchmarking. *Journal of Turbomachinery*, 129(1):14–22, 2007.

[139] A. Prasad. Calculation of the mixed-out state in turbomachine flows. *J. Turbomach.*, 127(3):564–572, 2005.

[140] P. J. Przytarski and A. P. S. Wheeler. The Effect of Rotor-Stator Gap on Repeating-Stage Compressor. In *Proceedings of the ASME Turbo Expo*, pages GT2019–91007, 2019.

[141] P. J. Przytarski and A. P. Wheeler. Accurate prediction of loss using high fidelity methods. In *Proceedings of the ASME Turbo Expo*, pages GT2018–77125, 2018.

[142] B. Raverdy, I. Mary, P. Sagaut, and N. Liamis. High-Resolution Large-Eddy Simulation of Flow Around Low-Pressure Turbine Blade. *AIAA Journal*, 41(3):390–397, 2003.

[143] S. Robinson. Coherent Motions in the Turbulent Boundary Layer. *Annual Review of Fluid Mechanics*, 23(1):601–639, 1991.

[144] R. D. Sandberg and V. Michelassi. The Current State of High-Fidelity Simulations for Main Gas Path Turbomachinery Components and Their Industrial Impact. *Flow, Turbulence and Combustion*, 102(4):797–848, 2019.

[145] R. D. Sandberg, V. Michelassi, R. Pichler, L. Chen, and R. Johnstone. Compressible Direct Numerical Simulation of Low-Pressure Turbines—Part I: Methodology. *J. Turbomach.*, 137(5):051011–0501011–10, 2015.

[146] S. Sarkar. Influence of wake structure on unsteady flow in a low pressure turbine blade passage. *J. Turbomach.*, 131(4):1–14, 2009.

[147] P. Schlatter, L. Brandt, H. C. de Lange, and D. S. Henningson. On streak breakdown in bypass transition. *Physics of Fluids*, 20(10), 2008.

[148] P. Schlatter, E. Deusebio, R. De Lange, and L. Brandt. Numerical study of the stabilisation of boundary-layer disturbances by finite amplitude streaks. *International Journal of Flow Control*, 2(4):259–288, 2010.

[149] P. Schlatter and R. Örlü. Turbulent boundary layers at moderate Reynolds numbers: inflow length and tripping effects. *Journal of Fluid Mechanics*, 710:5–34, 2012.

[150] H. Schlichting and K. Gersten. *Boundary Layer Theory*. Springer, 9th edition, 2017.

[151] G. B. Schubauer and P. S. Klebanoff. Contributions on the Mechanics of Boundary-Layer Transition. NACA-TR-12, 1955.

[152] V. Schulte and H. P. Hodson. Unsteady Wake-Induced Boundary Layer Transition in High Lift LP Turbines. *J. Turbomach.*, 120(1):28–35, 1998.

[153] V. Schulte and H. P. Hodson. Prediction of the becalmed region for lp turbine profile design. *J. Turbomach.*, 120, 1998.

[154] D. Serson, J. R. Meneghini, and S. J. Sherwin. Velocity-correction schemes for the incompressible Navier-Stokes equations in general coordinate systems. *Journal of Computational Physics*, 316:243–254, 2016.

[155] D. Serson. *Numerical Study of Wings with Wavy Leading and Trailing Edges*. PhD thesis, 2016.

[156] M. Skote and D. S. Henningson. Direct numerical simulation of a separated turbulent boundary layer. *Journal of Fluid Mechanics*, 471:107–136, 2002.

[157] J. Slotnick, A. Khodadoust, J. Alonso, D. Darmofal, W. Gropp, E. Lurie, and D. Mavriplis. CFD Vision 2030 Study: A Path to Revolutionary Computational Aerosciences. Technical Report NASA/Cr-2014-21878, 2014.

[158] A. Smirnov, S. Shi, and I. Celik. Random Flow Generation Technique for Large Eddy Simulations and Particle-Dynamics Modeling. *Journal of Fluids Engineering*, 123(2):359, 2001.

[159] L. H. Smith Jr. Wake Dispersion in Turbomachines. *ASME J. Basic Eng.*, 88(3):688–690, 1966.

[160] S. C. Spiegel, J. R. DeBonis, and H. T. Huynh. Overview of the NASA Glenn flux reconstruction based high-order unstructured grid code. In *54th AIAA Aerospace Sciences Meeting*, 2016.

[161] P. Stadtmüller and L. Fottner. A Test Case for the Numerical Investigation of Wake Passing Effects on a Highly Loaded LP Turbine Cascade Blade. In *Proceedings of the ASME Turbo Expo*, pages 2001–GT–0311, 2001.

[162] R. D. Stieger and H. P. Hodson. Unsteady dissipation measurements on a flat plate subject to wake passing. *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy*, 217(4):413–420, 2003.

[163] R. D. Stieger and H. P. Hodson. The Transition Mechanism of Highly-Loaded LP Turbine Blades. *J. Turbomach.*, 126(4):536–543, 2004.

[164] R. D. Stieger and H. P. Hodson. The Unsteady Development of a Turbulent Wake Through a Downstream Low-Pressure Turbine Blade Passage. *J. Turbomach.*, 127(2):388–394, 2005.

[165] R. D. Stieger. *The Effects of Wakes on Separating Boundary Layers in Low Pressure Turbines*. PhD thesis, 2002.

[166] B. A. Szabó and I. Babuška. *An Introduction to Finite Element Analysis*. John Wiley & Sons, 1991.

[167] G. R. Tabor and M. H. Baba-Ahmadi. Inlet conditions for large eddy simulation: A review. *Computers and Fluids*, 39(4):553–567, 2010.

[168] E. Tadmor. Convergence of Spectral Methods for Nonlinear Conservation Laws. *SIAM J. Numer. Anal.*, 26(1):30–44, 1989.

[169] E. Touber and N. D. Sandham. Large-eddy simulation of low-frequency unsteadiness in a turbulent shock-induced separation bubble. *Theoretical and Computational Fluid Dynamics*, 23(2):79–107, 2009.

[170] P. G. Tucker. Computation of unsteady turbomachinery flows: Part 1 - Progress and challenges. *Progress in Aerospace Sciences*, 47(7):522–545, 2011.

[171] H. M. Tufo and P. F. Fischer. Fast Parallel Direct Solvers for Coarse Grid Problems. *Journal of Parallel and Distributed Computing*, 61(2):151–177, 2001.

[172] M. Turner, D. Moxey, S. J. Sherwin, and J. Peiró. Automatic Generation of 3D Unstructured High-Order Curvilinear Meshes. *VII European Congress on Computational Methods in Applied Sciences and Engineering*, pages 428–443, 2016.

[173] J. Tyacke, N. Vadlamani, W. Trojak, R. Watson, Y. Ma, and P. Tucker. Turbomachinery simulation challenges and the future. *Progress in Aerospace Sciences*, (July):100554, 2019.

[174] A. Uranga, P.-O. Persson, M. Drela, and J. Peraire. Implicit Large Eddy Simulation of transition to turbulence at low Reynolds numbers using a Discontinuous Galerkin method. *Int. J. Numer. Meth. Engng.*, 87:232–261, 2010, 1201.4903.

[175] R. Vázquez Diaz and D. Torre. The Effect of Mach Number on the Loss Generation of LP Turbines. In *Proceedings of the ASME Turbo Expo*, pages GT2012–68555, 2012.

[176] M. Vera and H. P. Hodson. Low Speed Vs High Speed Testing of LP Turbine Blade-Wake Interaction. In *Symposium on Measuring Techniques in Transonic and Supersonic Flow in Cascades and Turbomachines*, number September, pages 1–10, 2002.

[177] P. E. Vincent, P. Castonguay, and A. Jameson. A new class of high-order energy stable flux reconstruction schemes. *Journal of Scientific Computing*, 47(1):50–72, 2011.

[178] P. E. Vincent, A. M. Farrington, F. D. Witherden, and A. Jameson. An extended range of stable-symmetric-conservative Flux Reconstruction correction functions. *Computer Methods in Applied Mechanics and Engineering*, 296:248–272, 2015.

[179] P. E. Vincent, F. Witherden, B. Vermeire, J. S. Park, and A. Iyer. Towards Green Aviation with Python at Petascale. In *International Conference for High Performance Computing, Networking, Storage and Analysis*, pages 1–11, Salt Lake City, UT, USA, 2016.

[180] R. Vinuesa, A. Bobke, R. Örlü, and P. Schlatter. On determining characteristic length scales in pressure-gradient turbulent boundary layers. *Physics of Fluids*, 28(5):055101–055101–13, 2016.

[181] R. Vinuesa, L. Fick, P. Negi, O. Marin, E. Merzari, and P. Schlatter. Turbulence statistics in a spectral element code: a toolbox for High-Fidelity Simulations. Technical report, 2017.

[182] G. Wang, A. Cassinelli, and S. J. Sherwin. Speed Comparison among Nektar++ Solvers. https://www.nektar.info/speed-comparison-among-nektar-solvers/, 2020.

[183] Z. J. Wang and H. Gao. A unifying lifting collocation penalty formulation including the discontinuous Galerkin, spectral volume/difference methods for conservation laws on mixed grids. *Journal of Computational Physics*, 228(21):8161–8186, 2009.

[184] Z. Wang, M. S. Triantafyllou, Y. Constantinides, and G. E. Karniadakis. A spectral-element / Fourier smoothed profile method for large-eddy simulations of complex VIV problems. *Computers and Fluids*, 172:84–96, 2018.

[185] Z. Wang, K. Fidkowski, R. Abgrall, F. Bassi, D. Caraeni, A. Cary, H. Deconinck, R. Hartmann, K. Hillewaert, H. Huynh, N. Kroll, G. May, P.-O. Persson, B. van Leer, and M. Visbal. High-order CFD methods: current status and perspective. *International Journal for Numerical Methods in Fluids*, 72(8):811–845, jul 2013.

[186] A. P. S. Wheeler, R. D. Sandberg, N. D. Sandham, R. Pichler, V. Michelassi, and G. Laskowski. Direct Numerical Simulations of a High-Pressure Turbine Vane. *J. Turbomach.*, 138(7):071003, 2016.

[187] A. P. Wheeler, A. M. Dickens, and R. J. Miller. The effect of nonequilibrium boundary layers on compressor performance. *J. Turbomach.*, 140(10):1–10, 2018.

[188] C. H. K. Williamson. Vortex Dynamics in the Cylinder Wake. *Annu. Rev. Fluid. Mech*, 28:477–539, 1996.

[189] J. G. Wissink. DNS of separating, low Reynolds number flow in a turbine cascade with incoming wakes. *International Journal of Heat and Fluid Flow*, 24(4):626–635, 2003.

[190] J. G. Wissink and W. Rodi. Numerical study of the near wake of a circular cylinder. *International Journal of Heat and Fluid Flow*, 29(4):1060–1070, 2008.

[191] J. G. Wissink and W. Rodi. Direct numerical simulation of flow and heat transfer in a turbine cascade with incoming wakes. *Journal of Fluid Mechanics*, 569:209–247, 2006.

[192] J. G. Wissink, W. Rodi, and H. P. Hodson. The influence of disturbances carried by periodically incoming wakes on the separating flow around a turbine blade. *International Journal of Heat and Fluid Flow*, 27(4):721–729, 2006.

[193] J. G. Wissink, T. A. Zaki, W. Rodi, and P. A. Durbin. The effect of wake turbulence intensity on transition in a compressor cascade. *Flow, Turbulence and Combustion*, 93(4):555–576, 2014.

[194] F. D. Witherden, A. M. Farrington, and P. E. Vincent. PyFR: An open source framework for solving advection-diffusion type problems on streaming architectures using the flux reconstruction approach. *Computer Physics Communications*, 185(11):3028–3040, 2014, 1312.1638.

[195] F. D. Witherden, B. C. Vermeire, and P. E. Vincent. Heterogeneous computing on mixed unstructured grids with PyFR. *Computers and Fluids*, 120:173–186, 2015, 1409.0405.

[196] X. Wu and P. A. Durbin. Evidence of longitudinal vortices evolved from distorted wakes in a turbine passage. *Journal of Fluid Mechanics*, 446(2001):199–228, 2001.

[197] X. Wu. Inflow Turbulence Generation Methods. *Annual Review of Fluid Mechanics*, 49(1):23–49, 2017.

[198] X. Wu, R. G. Jacobs, J. C. R. Hunt, and P. A. Durbin. Simulation of boundary layer transition induced by periodically passing wakes. *Journal of Fluid Mechanics*, 398(1999):109–153, 1999.

[199] Z. T. Xie and I. P. Castro. Efficient generation of inflow conditions for large eddy simulation of street-scale flows. *Flow, Turbulence and Combustion*, 81(3):449–470, 2008.

[200] Z. G. Yan, Y. Pan, G. Castiglioni, K. Hillewaert, J. Peiró, D. Moxey, and S. J. Sherwin. Nektar++: Design and implementation of an implicit, spectral/hp element, compressible flow solver using a Jacobian-free Newton Krylov approach. *Computers and Mathematics with Applications*, 81:351–372, 2021, 2002.04222.

[201] R. Yu and X. S. Bai. A fully divergence-free method for generation of inhomogeneous and anisotropic turbulence with large spatial variation. *Journal of Computational Physics*, 256:234–253, 2014.

[202] T. A. Zaki, J. G. Wissink, P. A. Durbin, and W. Rodi. Direct computations of boundary layers distorted by migrating wakes in a linear compressor cascade. *Flow, Turbulence and Combustion*, 83(3):307–322, 2009.

[203] T. A. Zaki. From streaks to spots and on to turbulence: Exploring the dynamics of boundary layer transition. *Flow, Turbulence and Combustion*, 91(3):451–473, 2013.

[204] T. A. Zaki and P. A. Durbin. Continuous mode transition and the effects of pressure gradient. *Journal of Fluid Mechanics*, 563:357–388, 2006.

[205] T. A. Zaki and S. Saha. On shear sheltering and the structure of vortical modes in single- and two-fluid boundary layers. *Journal of Fluid Mechanics*, 626:111–147, 2009.

[206] T. A. Zaki, J. G. Wissink, W. Rodi, and P. A. Durbin. Direct numerical simulations of transition in a compressor cascade: the influence of free-stream turbulence. *Journal of Fluid Mechanics*, 665(2010):57–98, 2010.

[207] Y. Zhao and R. D. Sandberg. Bypass transition in boundary layers subject to strong pressure gradient and curvature effects. *Journal of Fluid Mechanics*, 888, 2020.

Appendix A

Turbulence kinetic energy budget

The calculation of turbulence statistics is of vital importance to scale-resolving simulations of turbulent flows. This appendix shows the approach followed for the calculation of Reynolds stresses and the components of the turbulence kinetic energy transport equation, how it was incorporated into the *Nektar++* framework and how it was validated.

Turbulence statistics are based on fluctuating velocity components. Adopting the conventional Reynolds decomposition, $u = U + u'$, the mean velocity field $U_i(x, y, z)$ is only available at the end of the computation. Therefore, in order to compute these quantities accurately, they must be obtained as a post-processing step.

The two-step process was developed and integrated in the incompressible Navier-Stokes solver implemented in the *Nektar++*platform. The run-time sampling of the necessary quantities is carried out by a filter (*TKEBalance*). The turbulence statistics are then obtained in a post-processing phase, by adding a new module to *Nektar++* post-processing utility, *FieldConvert*. By default, this feature calculates the Reynolds Stresses, turbulence production and dissipation. Additionally, the other components of the turbulence kinetic energy transport equation can optionally be calculated: viscous diffusion, transport, velocity-pressure gradient, convection.

The process is based in part on the report by Vinuesa et al. [181] and illustrated in more detail in the following section. Validation is performed on a well-known test case: turbulent channel at $Re_\tau = 180$ [81]. The important properties of the averaging operators employed for the following derivations are listed in [8].

A.1 TKEBalance filter

In runtime, the following quantities are accumulated and then time-averaged:

- $\overline{u}$, $\overline{v}$, $\overline{w}$
- $\overline{uu}$, $\overline{uv}$, $\overline{uw}$, $\overline{vv}$, $\overline{vw}$, $\overline{ww}$
- $e_{11} = \overline{\left(\frac{\partial u}{\partial x}\right)^2 + \left(\frac{\partial u}{\partial y}\right)^2 + \left(\frac{\partial u}{\partial z}\right)^2}$

- $e_{22} = \overline{\left(\frac{\partial v}{\partial x}\right)^2 + \left(\frac{\partial v}{\partial y}\right)^2 + \left(\frac{\partial v}{\partial z}\right)^2}$
- $e_{33} = \overline{\left(\frac{\partial w}{\partial x}\right)^2 + \left(\frac{\partial w}{\partial y}\right)^2 + \left(\frac{\partial w}{\partial z}\right)^2}$

If all statistics are to be calculated (i.e. including viscous diffusion, transport, velocity-pressure gradient and convection), the third-order moments are additionally sampled, as well as the products of velocity and pressure gradients:

- $\overline{uuu}$, $\overline{uuv}$, $\overline{uuw}$, $\overline{uvv}$, $\overline{vvv}$, $\overline{vww}$, $\overline{uww}$, $\overline{vww}$, $\overline{www}$
- $\overline{u\frac{\partial p}{\partial x}}$, $\overline{v\frac{\partial p}{\partial y}}$, $\overline{w\frac{\partial p}{\partial z}}$

A.2 TKEBalance post-processing module

The transport equation of the Reynolds Stress tensor can be written as:

$$\frac{\partial}{\partial t}\overline{u_i' u_j'} = P_{ij} + \epsilon_{ij} + D_{ij} + T_{ij} + \Pi_{ij} - C_{ij}. \tag{A.1}$$

The equation is governed by the following tensors: production P_{ij}, dissipation ϵ_{ij}, viscous diffusion D_{ij}, transport T_{ij}, velocity-pressure gradient $\prod_{ij}$ and convection C_{ij}. Taking the trace of Eq. A.1 leads to the transport equation for the turbulence kinetic energy k:

$$\frac{\partial k}{\partial t} = P^k + \epsilon^k + D^k + T^k + \Pi^k + C^k \tag{A.2}$$

where each term of the respective tensor is obtained as $A^k = \frac{1}{2}(A_{11} + A_{22} + A_{33})$.

Reynolds Stresses

The Reynolds stresses are computed simply employing the relationship:

$$\overline{u_i u_j} = U_i U_j + \overline{u_i' u_j'} \tag{A.3}$$

Turbulence production

The turbulence production tensor can be written as:

$$P_{ij} = -\overline{u_i' u_k'}\frac{\partial U_j}{\partial x_k} - \overline{u_j' u_k'}\frac{\partial U_i}{\partial x_k} \tag{A.4}$$

Turbulence kinetic energy production is therefore: $P^k = \frac{1}{2}(P_{11} + P_{22} + P_{33})$.

$$\begin{aligned} P^k = -\Bigg[& \overline{u'^2}\frac{\partial U}{\partial x} + \overline{u'v'}\frac{\partial U}{\partial y} + \overline{u'w'}\frac{\partial U}{\partial z} \\ & + \overline{u'v'}\frac{\partial V}{\partial x} + \overline{v'^2}\frac{\partial V}{\partial y} + \overline{v'w'}\frac{\partial V}{\partial z} \\ & + \overline{u'w'}\frac{\partial W}{\partial x} + \overline{v'w'}\frac{\partial W}{\partial y} + \overline{w'^2}\frac{\partial W}{\partial z}\Bigg] \end{aligned} \tag{A.5}$$

The mean flow gradients need to be computed. Since the Reynolds Stresses were calculated at the previous step, turbulence kinetic energy production can then be obtained.

Turbulence dissipation

The turbulence dissipation tensor can be written as:

$$\epsilon_{ij} = -2\nu \overline{\frac{\partial u_i'}{\partial x_k} \frac{\partial u_j'}{\partial x_k}} \tag{A.6}$$

Turbulence kinetic energy dissipation is: $\epsilon^k = \frac{1}{2}(\epsilon_{11} + \epsilon_{22} + \epsilon_{33})$. Expanding this equation:

$$\begin{aligned} \epsilon^k = -\nu \Bigg\{ & \left[\overline{\left(\frac{\partial u'}{\partial x}\right)^2} + \overline{\left(\frac{\partial u'}{\partial y}\right)^2} + \overline{\left(\frac{\partial u'}{\partial z}\right)^2} \right] \\ & + \left[\overline{\left(\frac{\partial v'}{\partial x}\right)^2} + \overline{\left(\frac{\partial v'}{\partial y}\right)^2} + \overline{\left(\frac{\partial v'}{\partial z}\right)^2} \right] \\ & + \left[\overline{\left(\frac{\partial w'}{\partial x}\right)^2} + \overline{\left(\frac{\partial w'}{\partial y}\right)^2} + \overline{\left(\frac{\partial w'}{\partial z}\right)^2} \right] \Bigg\} \end{aligned} \tag{A.7}$$

Consider for the sake of the argument the quantity $\overline{\frac{\partial u}{\partial x}\frac{\partial u}{\partial x}}$, sampled during simulation time. It is can be expanded as follows:

$$\begin{aligned} \overline{\frac{\partial u}{\partial x}\frac{\partial u}{\partial x}} &= \overline{\frac{\partial (U+u')}{\partial x}\frac{\partial (U+u')}{\partial x}} = \overline{\left(\frac{\partial U}{\partial x} + \frac{\partial u'}{\partial x}\right)\left(\frac{\partial U}{\partial x} + \frac{\partial u'}{\partial x}\right)} \\ &= \overline{\left(\frac{\partial U}{\partial x}\right)^2 + \left(\frac{\partial u'}{\partial x}\right)^2 - 2\frac{\partial U}{\partial x}\frac{\partial u'}{\partial x}} = \overline{\left(\frac{\partial U}{\partial x}\right)^2} + \overline{\left(\frac{\partial u'}{\partial x}\right)^2} - 2\overline{\frac{\partial U}{\partial x}\frac{\partial u'}{\partial x}} \\ &= \left(\frac{\partial U}{\partial x}\right)^2 + \overline{\left(\frac{\partial u'}{\partial x}\right)^2} - 2\frac{\partial U}{\partial x}\overline{\frac{\partial u'}{\partial x}} = \left(\frac{\partial U}{\partial x}\right)^2 + \overline{\left(\frac{\partial u'}{\partial x}\right)^2} - 2\frac{\partial U}{\partial x}\cancelto{0}{\overline{\frac{\partial u'}{\partial x}}} \\ &= \left(\frac{\partial U}{\partial x}\right)^2 + \overline{\left(\frac{\partial u'}{\partial x}\right)^2} \end{aligned}$$

Therefore:

$$\overline{\left(\frac{\partial u'}{\partial x}\right)^2} = \overline{\left(\frac{\partial u}{\partial x}\right)^2} - \left(\frac{\partial U}{\partial x}\right)^2 \tag{A.8}$$

Substituting this in all the components of ϵ^k yields:

$$\begin{aligned} \epsilon^k = -\nu \Bigg\{ & \left[\overline{\left(\frac{\partial u}{\partial x}\right)^2} - \left(\frac{\partial U}{\partial x}\right)^2 + \overline{\left(\frac{\partial u}{\partial y}\right)^2} - \left(\frac{\partial U}{\partial y}\right)^2 + \overline{\left(\frac{\partial u}{\partial z}\right)^2} - \left(\frac{\partial U}{\partial z}\right)^2 \right] \\ & + \left[\overline{\left(\frac{\partial v}{\partial x}\right)^2} - \left(\frac{\partial V}{\partial x}\right)^2 + \overline{\left(\frac{\partial v}{\partial y}\right)^2} - \left(\frac{\partial V}{\partial y}\right)^2 + \overline{\left(\frac{\partial v}{\partial z}\right)^2} - \left(\frac{\partial V}{\partial z}\right)^2 \right] \\ & + \left[\overline{\left(\frac{\partial w}{\partial x}\right)^2} - \left(\frac{\partial W}{\partial x}\right)^2 + \overline{\left(\frac{\partial w}{\partial y}\right)^2} - \left(\frac{\partial W}{\partial y}\right)^2 + \overline{\left(\frac{\partial w}{\partial z}\right)^2} - \left(\frac{\partial W}{\partial z}\right)^2 \right] \Bigg\} \end{aligned} \tag{A.9}$$

$$\begin{aligned}\epsilon^k = -\nu\Bigg\{\Bigg[&\overbrace{\overline{\left(\frac{\partial u}{\partial x}\right)^2+\left(\frac{\partial u}{\partial y}\right)^2+\left(\frac{\partial u}{\partial z}\right)^2}}^{e_{11}}-\left(\frac{\partial U}{\partial x}\right)^2-\left(\frac{\partial U}{\partial y}\right)^2-\left(\frac{\partial U}{\partial z}\right)^2\Bigg]\\ +\Bigg[&\overbrace{\overline{\left(\frac{\partial v}{\partial x}\right)^2+\left(\frac{\partial v}{\partial y}\right)^2+\left(\frac{\partial v}{\partial z}\right)^2}}^{e_{22}}-\left(\frac{\partial V}{\partial x}\right)^2-\left(\frac{\partial V}{\partial y}\right)^2-\left(\frac{\partial V}{\partial z}\right)^2\Bigg]\\ +\Bigg[&\overbrace{\overline{\left(\frac{\partial w}{\partial x}\right)^2+\left(\frac{\partial w}{\partial y}\right)^2+\left(\frac{\partial w}{\partial z}\right)^2}}^{e_{33}}-\left(\frac{\partial W}{\partial x}\right)^2-\left(\frac{\partial W}{\partial y}\right)^2-\left(\frac{\partial W}{\partial z}\right)^2\Bigg]\Bigg\}\end{aligned} \tag{A.10}$$

Using the definitions of e_{ii} introduced in Section A.1, turbulent dissipation can be finally written as:

$$\begin{aligned}\epsilon^k =\nu\Bigg\{\Bigg[&\left(\frac{\partial U}{\partial x}\right)^2+\left(\frac{\partial U}{\partial y}\right)^2+\left(\frac{\partial U}{\partial z}\right)^2-e_{11}\Bigg]\\ +\Bigg[&\left(\frac{\partial V}{\partial x}\right)^2+\left(\frac{\partial V}{\partial y}\right)^2+\left(\frac{\partial V}{\partial z}\right)^2-e_{22}\Bigg]\\ +&\left(\frac{\partial W}{\partial x}\right)^2+\left(\frac{\partial W}{\partial y}\right)^2+\left(\frac{\partial W}{\partial z}\right)^2-e_{33}\Bigg]\Bigg\}\end{aligned} \tag{A.11}$$

Velocity-pressure gradient

The velocity-pressure gradient tensor is:

$$\Pi_{i,j} = -\frac{1}{\rho}\left(\overline{u_i'\frac{\partial p'}{\partial x_j}}+\overline{u_j'\frac{\partial p'}{\partial x_i}}\right) \tag{A.12}$$

Taking its trace $\Pi^k = \frac{1}{2}(\Pi_{11}+\Pi_{22}+\Pi_{33})$ yields:

$$\begin{aligned}\Pi^k =&\frac{1}{2}\Bigg[-\frac{1}{\rho}\left(\overline{u'\frac{\partial p'}{\partial x}}+\overline{u'\frac{\partial p'}{\partial x}}\right)\\ &-\frac{1}{\rho}\left(\overline{v'\frac{\partial p'}{\partial y}}+\overline{v'\frac{\partial p'}{\partial y}}\right)\\ &-\frac{1}{\rho}\left(\overline{w'\frac{\partial p'}{\partial z}}+\overline{w'\frac{\partial p'}{\partial z}}\right)\Bigg]\\ =&-\frac{1}{\rho}\left(\overline{u'\frac{\partial p'}{\partial x}}+\overline{v'\frac{\partial p'}{\partial y}}+\overline{w'\frac{\partial p'}{\partial z}}\right)\end{aligned} \tag{A.13}$$

To calculate this quantity, the relationship below is employed:

$$\overline{u_i\frac{\partial p}{\partial x_j}} = U_i\frac{\partial P}{\partial x_j}+\overline{u_i'\frac{\partial p'}{\partial x_j}} \tag{A.14}$$

The mean pressure gradient is calculated in the post-processing step, while the last three

components of the equation are accumulated during the simulation, obtaining:

$$\begin{aligned}\Pi^k = &-\frac{1}{\rho}\Bigg(U\frac{\partial P}{\partial x}+V\frac{\partial P}{\partial y}+W\frac{\partial P}{\partial z}\\ &+\overline{u\frac{\partial p}{\partial x}}+\overline{v\frac{\partial p}{\partial y}}+\overline{w\frac{\partial p}{\partial z}}\Bigg)\end{aligned} \tag{A.15}$$

Convection

The convection tensor is the following:

$$C_{ij} = U_k\frac{\partial}{\partial x_k}\overline{u_i'u_j'} \tag{A.16}$$

Convection in the TKE equation is: $C^k = \frac{1}{2}(C_{11}+C_{22}+C_{33})$. Writing out the components:

$$\begin{aligned}C^k =&\frac{1}{2}\Bigg[\left(U\frac{\partial}{\partial x}\overline{u'^2}+V\frac{\partial}{\partial y}\overline{u'^2}+W\frac{\partial}{\partial z}\overline{u'^2}\right)\\ &+\left(U\frac{\partial}{\partial x}\overline{v'^2}+V\frac{\partial}{\partial y}\overline{v'^2}+W\frac{\partial}{\partial z}\overline{v'^2}\right)\\ &+\left(U\frac{\partial}{\partial x}\overline{w'^2}+V\frac{\partial}{\partial y}\overline{w'^2}+W\frac{\partial}{\partial z}\overline{w'^2}\right)\Bigg]\end{aligned} \tag{A.17}$$

Having already calculated the Reynolds stresses, the spatial gradients of the normal stresses are sufficient to compute C^k .

Viscous diffusion

The viscous diffusion tensor is the following:

$$D_{i,j} = \nu\frac{\partial^2}{\partial x_k^2}\overline{u_i'u_j'} \tag{A.18}$$

Viscous diffusion in the TKE equation is: $D^k = \frac{1}{2}(D_{11}+D_{22}+D_{33})$. Expanding the components:

$$\begin{aligned}D^k =&\frac{1}{2}\Bigg[\nu\left(\frac{\partial^2}{\partial x^2}\overline{u'^2}+\frac{\partial^2}{\partial y^2}\overline{u'^2}+\frac{\partial^2}{\partial z^2}\overline{u'^2}\right)\\ &+\nu\left(\frac{\partial^2}{\partial x^2}\overline{v'^2}+\frac{\partial^2}{\partial y^2}\overline{v'^2}+\frac{\partial^2}{\partial z^2}\overline{v'^2}\right)\\ &+\nu\left(\frac{\partial^2}{\partial x^2}\overline{w'^2}+\frac{\partial^2}{\partial y^2}\overline{w'^2}+\frac{\partial^2}{\partial z^2}\overline{w'^2}\right)\Bigg]\end{aligned} \tag{A.19}$$

The Reynolds stresses were made available from previous processing, as well as the first derivative of the normal stresses. Therefore, the second derivative of $\overline{u_i'^2}$ must be calculated.

Transport

The TKE transport tensor is:

$$T_{ij} = -\frac{\partial}{\partial x_k}\overline{u_i' u_j' u_k'} \tag{A.20}$$

Turbulence transport in the TKE equation is: $T^k = \frac{1}{2}(T_{11} + T_{22} + T_{33})$. The trace of the tensor is:

$$\begin{aligned} T^k =& \frac{1}{2}\Bigg[-\left(\frac{\partial}{\partial x}\overline{u'^3} + \frac{\partial}{\partial y}\overline{u'^2 v'} + \frac{\partial}{\partial z}\overline{u'^2 w'}\right) \\ & -\left(\frac{\partial}{\partial x}\overline{u' v'^2} + \frac{\partial}{\partial y}\overline{v'^3} + \frac{\partial}{\partial z}\overline{v'^2 w'}\right) \\ & -\left(\frac{\partial}{\partial x}\overline{u' w'^2} + \frac{\partial}{\partial y}\overline{v' w'^2} + \frac{\partial}{\partial z}\overline{w'^3}\right) \end{aligned} \tag{A.21}$$

In order to calculate this quantity, the third order moments are required. Considering for example $\overline{u'u'v'}$, it can be written as:

$$\begin{aligned} \overline{u'u'v'} =& \overline{(u-U)^2(v-V)} = \overline{(u^2 - 2uU + u^2)(v-V)} \\ =& \overline{u^2 v - 2uvU + vU^2 - u^2 V + 2uUV - U^2 V} \\ =& \overline{u^2 v} - 2\overline{uv}U + \overline{v}U^2 - \overline{u^2}V + 2\overline{u}UV - U^2 V \\ =& \overline{u^2 v} - 2\overline{uv}U - \overline{u^2}V + 2U^2 V \end{aligned} \tag{A.22}$$

Therefore, in general:

$$\overline{u_i' u_i' u_j'} = \overline{u_i^2 u_j} - 2\overline{u_i u_j}U_i - \overline{u_i^2}U_j + 2U_i^2 U_j \tag{A.23}$$

All these quantities are known, so T^k can be calculated by computing $\overline{u_i'^2 u_j'}$, taking its derivative in direction j and combining it all together. The full expression very long, therefore a more compact version is reported:

$$\begin{aligned} T^k =& -\frac{1}{2}\Bigg[\left(\frac{\partial}{\partial x}\overline{u'^3} + \frac{\partial}{\partial y}\overline{u'^2 v'} + \frac{\partial}{\partial z}\overline{u'^2 w'}\right) \\ & +\left(\frac{\partial}{\partial x}\overline{u' v'^2} + \frac{\partial}{\partial y}\overline{v'^3} + \frac{\partial}{\partial z}\overline{v'^2 w'}\right) \\ & +\left(\frac{\partial}{\partial x}\overline{u' w'^2} + \frac{\partial}{\partial y}\overline{v' w'^2} + \frac{\partial}{\partial z}\overline{w'^3}\right) \end{aligned} \tag{A.24}$$

A.3 Validation on turbulent channel flow at $Re_\tau = 180$

The validation of the *TKEBalance* filter and *FieldConvert* module is carried out by analysing a canonical turbulent channel flow problem. DNS of turbulent channel flow is one of the most widely analysed geomemtries of the past few decades, widely adopted to support the development of the state-of-art of turbulence research [81]. The main results are here compared with turbulent channel flow data at $Re_\tau = \frac{u_\tau h}{\nu} = 180$ [81, 61].

The problem setup generally allows to discretise the computational domain using a

Fourier Spectral method, leveraging the two periodic dimensions (streamwise and spanwise directions). However, in this study the numerical setup is analogous to that employed with LPT cases, combining a spectral/hp element discretisation in the $x-y$ planes with a Fourier expansion in the spanwise direction. The computational mesh employed to discretise each $x-y$ plane has 28 quadrilateral elements in the streamwise direction and 25 in the wall-normal direction (Figure A.1), stretched analogously to Kim et al. [81]. The computational domain is $L_z \times L_y \times L_z = 4\pi h \times 2h \times 2\pi h$, with $h = 1$. An expansion order of $P = 8$ is adopted in conjunction with $N_z = 168$ spanwise planes yielding $N_x \times N_y \times N_z = 252 \times 225 \times 168$ and thus ensuring DNS resolution ($\Delta y^+ \approx 0.18$). The flow is driven in the streamwise direction by imposing constang mass flux, adopting the approach described in detail in [120].

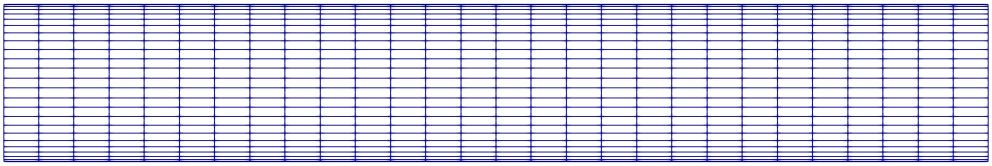

Figure A.1: Computational base mesh employed for the $x-y$ planes.

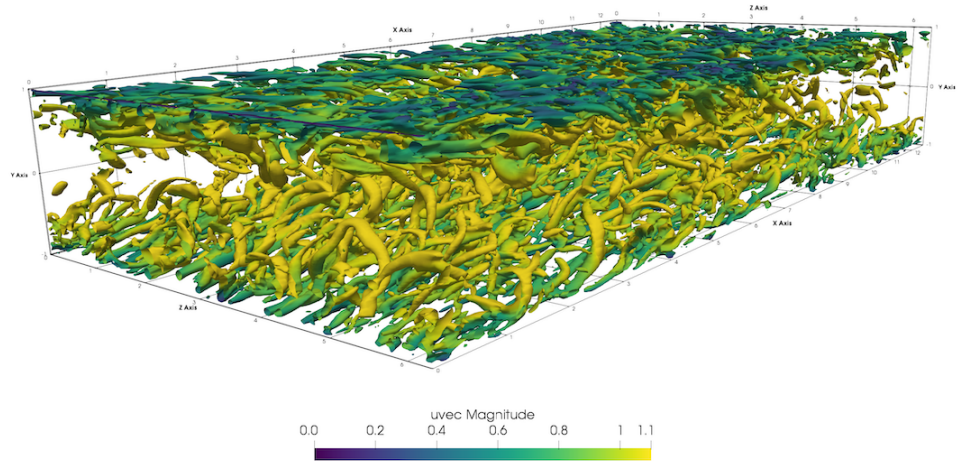

Figure A.2: Isosurface of Q-criterion contoured by velocity magnitude. The increasing size of turbulent scales away from the wall is clearly visible, alongside with the varied topology of structures that populate the various regions of the flow [143].

Figure A.2 qualitatively shows the turbulent flow field after transition. Small-scale structures are visible in the near-wall region, while larger scales dominate the outer region.

The mean streamwise velocity profiles in inner units is shown in Figure A.3a, showing close agreement with the reference data. Analogous levels of accuracy are achieved in the spanwise and vertical stresses $\langle v'^2 \rangle$ and $\langle w'^2 \rangle$ in Figure A.3b. The slight discrepancy

found in the streamwise component $\langle u'^2 \rangle$ may be due to a different computational approach adopted to imposing constant mass flow rate compaed to the reference data. This figure also validates the two-step approach introduced in this Appendix for accurate calculation of Reynolds stresses against the existing approach based on accumulating the time-averaged results in run-time. The two results are overlapped, validating the new implementation.

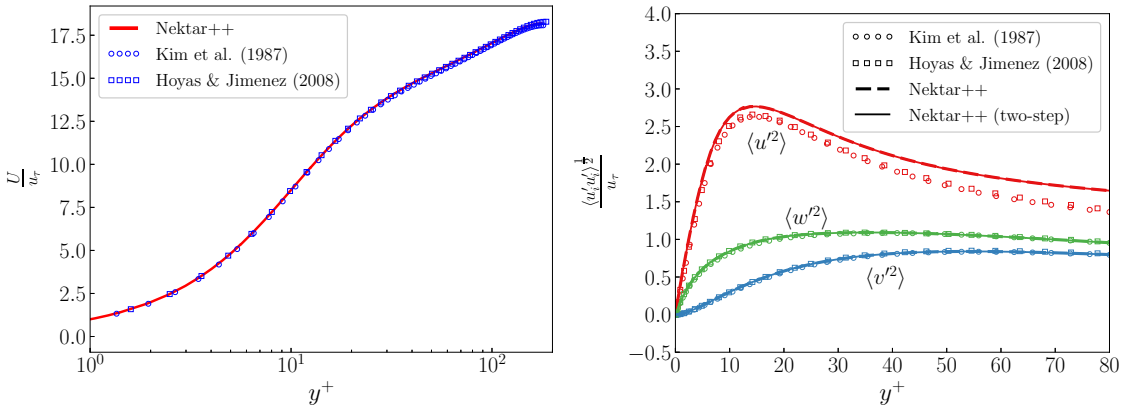

(a) (b)

Figure A.3: (a) Mean streamwise velocity profile, and (b) mean rms profile of the three velocity components.

The components of the kinetic energy transport equation are shown in Figure A.4, where very close agreement is demonstrated against existing data.

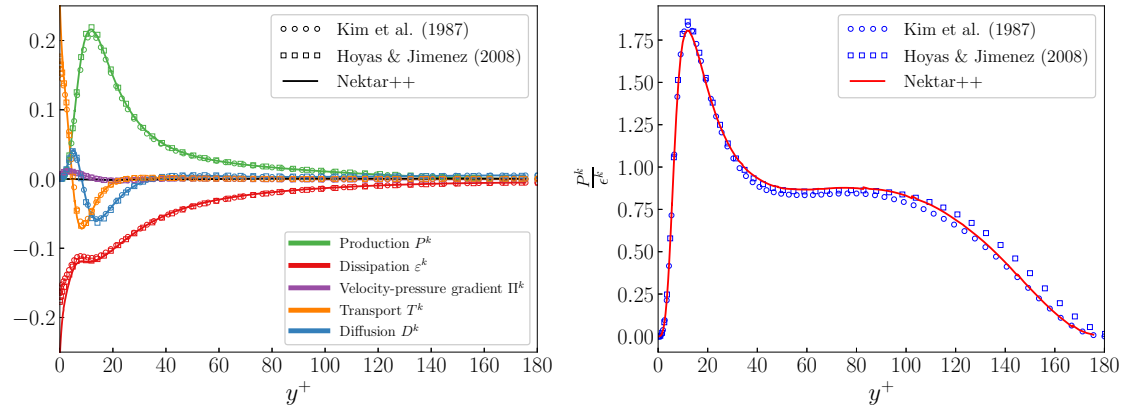

(a) (b)

Figure A.4: (a) Turbulence kinetic energy budget normalised by viscous scales, and (b) ratio of production to dissipation P^k/ϵ^k.

Appendix B

A brief note on the use of colourmaps

This thesis includes a large number of figures representing 2D and 3D contour maps of scalar quantities. The visualisation process consists of mapping scalar values to colours. Extensive research exists to support the choice of a suitable colourmap, to better allow the viewer to mentally perform reverse mapping to the scalar values. This section aims to briefly outline the principles adopted to inform the choice of colourmaps in this work, but without delving into thorough analysis. In the scientific community, the most popular has traditionally always been the rainbow colourmap, which is however a poor colourmap for a number of reasons, summarised by Moreland [115]: lack of natural perceived ordering, nonuniformity of perceptual changes in the colours [1] and its sensitivity to deficiencies in the human vision.

However, there is no clear "best" alternative colourmap, since every candidate has its features and flaws. Therefore, a possible approach is to adopt a number of general guidelines, and adapt them to the specific data set being analysed. The desirable characteristic of a colour palette are [114]:

- The map yields images that are aesthetically pleasing.
- The map has a maximal perceptual resolution.
- Interference with the shading of 3D surfaces is minimal.
- The map is not sensitive to vision deficiencies.
- The order of the colours should be intuitively the same for all people: in other words, there should be uniform luminance gradient.
- The perceptual interpolation matches the underlying scalars of the map without creating artifacts.

[1] The human eye perceives stronger contrast in the yellow-orange-red part of the spectrum compared to green-blue hues, creating artificial gradients in the yellow region of the spectrum.

These general indications are translated into practice by following a few guidelines. The type of data being visualised introduces a first discrimination in the type of acceptable colourmaps, which ideally should have an intuitive relationship with the field being visualised. Therefore, depending on the kind of data in analysis, a few suggestion can be made for suitable colourmaps (as available in Paraview [4]):

- Data in a range $[-a_1, a_2]$; diverging palette: *Cool to Warm, RdYlBu, Blue Orange (divergent)*.
- Data in a range $[0, a_1]$ (or vice versa); sequential palette: *2hot, Black-Body Radiation*, and *Inferno, Viridis, Magma, Plasma* from matplotlib.
- Discrete data sets without inherent ordering: qualitative palette. The online tool *ColorBrewer.org* [58] was found extremely useful for this purpose, (well as the generation of discrete sequential and diverging palettes).
- Cyclic data set: cyclic palette (not used in this work).

It should be noted that the above is not aiming to provide a comprehensive list of the best colourmaps, but simply some that were found suitable for the applications here presented.