Università degli studi di Genova

Scuola Politecnica

DIME

Department of Mechanical, Energy, Management and Transportation Engineering

Thesis submitted for the degree of Doctor of Philosophy
XXXVIII cycle

PhD Program in Marine Sciences and Technologies
Machines and Energy Systems for Marine Applications

Modal Decomposition-Based Post-Processing Techniques of High-Fidelity Data for Detailed Loss Characterization in Gas Turbines for Marine Applications

Supervisors:
Prof. Dario Barsi
Prof. Francesca Satta

Author:
Matteo Russo

Academic Year 2025–2026

A mio fratello Gabriele
M.R.

Abstract

The continuous demand for higher efficiency and reduced emissions in marine propulsion systems is driving renewed interest in gas turbine–based solutions, particularly in comparison with conventional marine internal combustion engines. Achieving competitive performance levels requires a detailed understanding of the complex loss mechanisms that limit the efficiency of turbomachinery components and a clear identification of the flow phenomena on which design improvements should be focused. In this context, high-fidelity numerical simulations represent a key tool, as they provide access to the unsteady, multiscale dynamics governing loss generation, which cannot be fully captured by lower-order modeling approaches. In this work, the low-pressure turbine (LPT) is adopted as a representative case study, due to its central role in determining the overall efficiency and operability of gas turbine systems, particularly in marine propulsion applications. Unsteady Large Eddy Simulations (LES) performed on a dataset of six different LPT stage geometries are exploited to construct a comprehensive and generalizable framework for the characterization, localization, and prediction of total pressure losses. Specifically, an advanced post-processing procedure is developed to analyze loss-related quantities through modal decomposition techniques, including Proper Orthogonal Decomposition (POD), Spectral Proper Orthogonal Decomposition (SPOD), and Phase-Locked Proper Orthogonal Decomposition (PPOD). These methods enable the spatial, spectral, and phase-resolved characterization of loss sources, highlighting the role of unsteady interactions between upstream wakes, large-scale coherent structures, and boundary-layer dynamics. The analysis reveals that differences in aerodynamic performance across the investigated geometries can be primarily attributed to variations in blade loading distribution, which governs the location and intensity of wake–boundary layer interactions. Based on these findings, a reduced-order predictive model is constructed by parameterizing blade loading and loss fields through POD-based representations and linking them via simple regression models. Despite the limited size of the high-fidelity dataset, the model demonstrates good predictive capability and allows the identification of an optimal blade loading distribution associated with a reduction in global losses. Overall, the results demonstrate the potential of combining high-fidelity simulations, modal analysis, and data-driven modeling to support the design of more efficient turbomachinery components. The proposed methodology provides a physically grounded and computationally efficient framework to understand, predict, and mitigate losses in a generic gas turbine component, with direct implications for the improvement of gas turbine performance in marine propulsion applications.

Acknowledgements

I would like to begin this long list of acknowledgments by thanking my supervisors, Professor Dario Barsi and Professor Francesca Satta, as well as Professor Daniele Simoni and Professor Davide Lengani. I am deeply grateful to all of you for your constant support and attentive guidance throughout this wonderful journey, for sharing your passion for research with me, and for giving me the opportunity to learn from you. I am truly happy to have had the chance to work alongside you.

Then, I would like to express my sincere gratitude to Matteo Dellacasagrande, Daniele Petronio, Pawel Przytarski, and Gustavo Lopez, whose experience and guidance have been a constant point of reference for me. I wish you all the very best for your future careers, and I hope to build on what I have learned from you in mine.

I would like to acknowledge all the participants involved in the ACROSS project. This collaborative experience has been highly enriching, providing me with the opportunity to work in a dynamic and interdisciplinary environment. In particular, I would like to thank my closest collaborators from Avio Aero and Morfo Design SRL for their valuable contributions and support. Among them, I would like to extend my sincere thanks to Francesco Bertini, Monica Gily, Andrea Notaristefano, Matteo Giovannini, and Juri Bellucci. I am truly grateful for the opportunity you gave me to work alongside you on this remarkable project. Your expertise, professionalism, and collaboration played an important role in the development of this work and significantly contributed to my growth as a researcher.

I would also like to sincerely thank all my fellow PhD colleagues. A special thank you goes to Andrea Carlucci and Virginia Bologna, who started and completed this journey alongside me, having been by my side since our early years at university. You have always been there, both during the most challenging moments and the happiest ones. Your constant presence, encouragement, and companionship have meant more to me than words can express, and you have been a part of this experience in a way I will never forget.

My final but truly heartfelt thanks go to all the people which I have mentioned so far for the friendships that have grown throughout this journey. I truly hope that there will be many more opportunities in the future to work together again.

I sincerely wish all of you the very best in your future paths.

Thank you,
M.R.

European HPC Joint Undertaking (JU) under grant agreement No. 955648 provided funding within the ACROSS project for this work. We also acknowledge PRACE, which awarded access to the Fenix Infrastructure resources at CINECA, partially funded from the European Union's Horizon 2020 research and innovation program through the ICEI project under the grant agreement No. 800858.

Contents

List of Figures

Symbols and Acronyms

Roman letters

C	Chord [m]
k	Turbulent kinetic energy [J kg^{-1}]
p	Pressure [Pa]
t	Time [s]
T	Temperature [K]
U	Velocity [m s^{-1}]
u	Velocity component [m s^{-1}]
$\mathcal{V}$	Volume [m^3]

Greek letters

μ	Dynamic viscosity [Pa s]
ν	Kinematic viscosity [m^2 s^{-1}]
ρ	Density [kg m^{-3}]
τ	Viscous stress tensor [Pa]
ϕ	Flow coefficient [-]

Dimensionless groups

Re	Reynolds number, based on outlet velocity
Ma	Mach number
St	Strouhal number
AG	Dimensionless rotor–vane axial gap

Acronyms and abbreviations

EEA	European Environment Agency
GHGs	Green House Gases
ICE	Internal Combustion Engine
VOC	Volatile Organic Compound
PM	Particulate Matter
WHO	World Heath Organization
HFO	Heavy Fuel Oil
MGO	Marine Gas Oil
LNG	Liquefied Natural Gas
IMO	International Maritime Organization

ECA	Emission Control Area
MARPOL	International Convention for the Prevention of Pollution from Ships
COGAS	Combined Gas And Steam
COGES	Combined Gas turbine Electric and Steam
CFD	Computational Fluid Dynamics
LES	Large Eddy Simulation
DNS	Direct Numerical Simulation
(U)RANS	(Unsteady) Reynolds Averaged Navier–Stokes
F.A.S.P.	Flux of Average Stagnation Pressure
F.T.K.E.	Flux of Turbulent Kinetic Energy
V.D.	Viscous Diffusion
T.D.	Turbulent Diffusion
P.D.	Pressure Diffusion
WALE	Wall Adapting Local Eddy-viscosity
FV	Finite Volume
POD	Proper Orthogonal Decomposition
SPOD	Spectral Proper Orthogonal Decomposition
PPOD	Phase-locked Proper Orthogonal Decomposition
TKE	Turbulent Kinetic Energy
LPT	Low Pressure Turbine
FFT	Fast Fourier Transform
DFT	Discrete Fourier Transform
DWT	Discrete Wavelet Transform
SVD	Singular Value Decomposition
CSD	Cross Spectral Density
FTT	Flow Through Time
PIV	Particle Image Velocimetry
HPDA	High Performance Data Analytics
HPC	High-Performance Computing
ML	Machine Learning
LOOCV	Leave-One-Out Cross Validation
RMSE	Root Mean Square Error

Superscripts and subscripts

x	Axial component value
cx	Axial chord-defined value
t	Total/stagnation value
1	Inlet section value
2	Outlet section value
i	x-axis component value
j	y-axis component value
$\dot{a}$	Time rate value
$\overline{a}$	Time-averaged value
a'	Time-fluctuating component value

Chapter 1

Introduction

The transportation of goods and people represents a cornerstone of modern society and economy. A well-developed transportation network is essential to support economic growth across all regions of the world, as it facilitates access to key public services such as education and healthcare, and contributes to job creation and overall societal well-being. The past century has witnessed a rapid expansion of the transportation sector driven by an increasing demand for mobility. In Europe alone, the European Environment Agency (EEA) reported for the year 2019 a total exceeding 6 billion passenger-kilometres (pkm) and 3.4 billion tonne-kilometres (tkm) of freight, corresponding to increases of 20.8% and 23%, respectively, compared to the year 2000 [1]. Figures 1.1 and 1.2 show the detailed breakdown of the transport sector between 2000 and 2020.

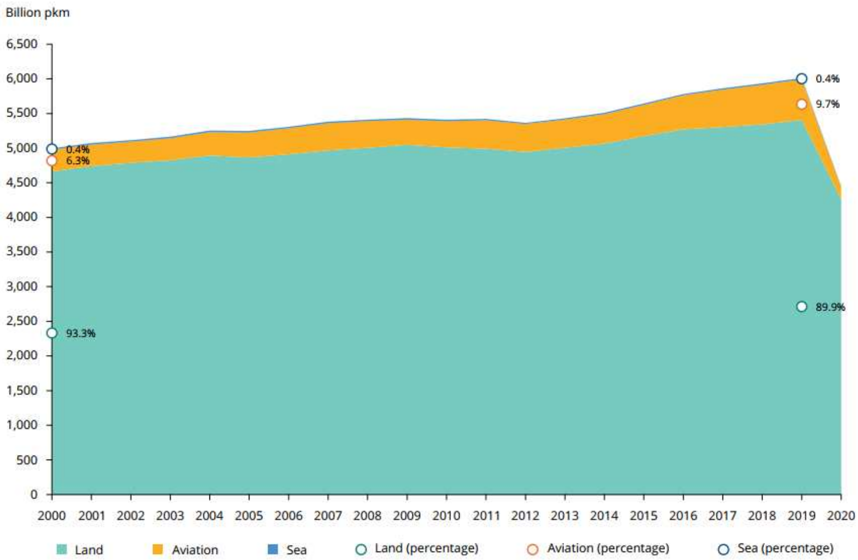

Figure 1.1: Passenger transport in Europe (EU-27) between 2000 and 2020 [1].

Since the SARS-CoV-2 pandemic, projections indicate a continuous increase in both passenger and freight transportation volumes. As an example, Table 1.1 reports the forecasts provided by the EU Reference Scenario 2020 [2] regarding maritime

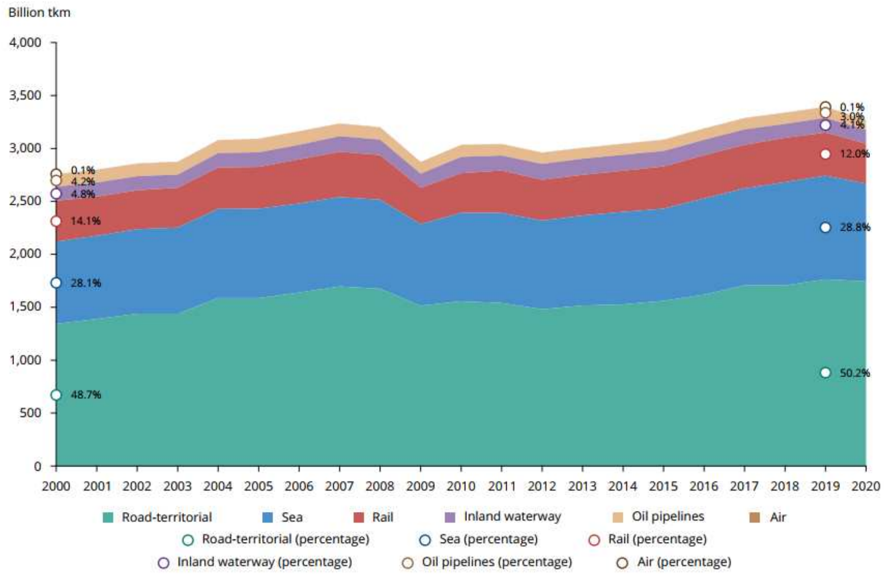

Figure 1.2: Freight transport in Europe (EU-27) between 2000 and 2020 [1].

passenger and freight transport activities up to the year 2050. However, such large-scale development is inevitably accompanied by a parallel rise in energy consumption and, in particular, in the emission of pollutants harmful to both human health and the environment. By the end of 2022, the share of greenhouse gas (*GreenHouse Gases*, GHGs) emissions attributed to the transport sector exceeded 20% [3], as shown in Figure 1.3.

Table 1.1: International maritime passenger and freight transport activity [2].

Activity	**Reference Scenario**			
	2015	2020	2030	2050
Passenger [Gpkm]	1.962	1.186	2.426	3.052
Freight [Gtkm]	14.269	11.327	17.075	21.444

Whether referring to air, land, or maritime transport, the majority of propulsion systems currently in use rely on the combustion of fossil fuels. The burning of these substances leads to the release of hundreds of chemical compounds into the atmosphere, many of which are harmful to the development of all living organisms and have adverse effects on environmental balance and infrastructure integrity [5]. Among the most critical emissions are greenhouse gases (GHGs), whose primary components include carbon dioxide (CO_2), methane (CH_4), chlorofluorocarbons (CFCs), and nitrous oxide (N_2O). These compounds are key contributors to global warming through the greenhouse effect. Specifically, they are capable of trapping infrared radiation reflected from the Earth's surface, thereby maintaining elevated temperatures near the ground. This mechanism underlines the harmful impact of excessive concentrations

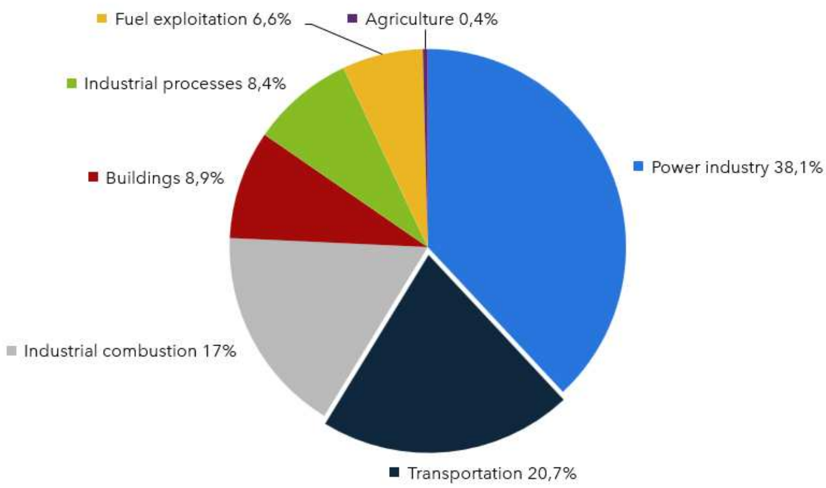

Figure 1.3: Sectoral distribution of global carbon dioxide emissions in 2022 [4].

of these gases in the atmosphere [6, 7]. Among them, CO_2 and N_2O are particularly relevant as they are typically detected in the exhaust emissions of various internal combustion engine (ICE) types used in the transport sector. In addition to these primary emissions, the interaction of hydrocarbon combustion products with atmospheric components can lead to the formation of secondary pollutants, including ozone (O_3). Ozone plays a significant role as a GHG, but with a complex dual effect depending on the altitude at which it is found. In the upper layers of the atmosphere, namely the stratosphere, high concentrations of ozone are beneficial as they absorb ultraviolet radiation, preventing it from reaching the Earth's surface. Conversely, near the ground—in the troposphere—ozone is considered a pollutant due to its toxicity when inhaled and its strong greenhouse effect [7].

Alongside the aforementioned GHGs, fossil-fuel-based engine systems emit several other harmful substances into the atmosphere, such as carbon monoxide (CO), nitrogen oxides (NO_x), sulfur oxides (SO_x), volatile organic compounds (VOCs), and particulate matter (PM). Regarding NO_x and SO_x emissions, their fate after being released is governed by meteorological and atmospheric chemical conditions [8]. Shortly after emission, a fraction of these species is deposited near the source as local deposition. The remaining portion is transported over long distances by wind and undergoes chemical transformations and dry or wet deposition at the surface. It is estimated that sulfur dioxide (SO_2), nitric oxide (NO), and their atmospheric recombination products have an atmospheric lifetime of about 1–3 days, corresponding to an average transport distance of approximately 400–1200 km [9]. Sulfur oxides influence the climate both directly and indirectly. After forming sulfuric acid, their deposition leads to the acidification of water and soil and the degradation of various materials. The fraction that remains suspended in the atmosphere reflects solar radiation and indirectly alters the composition and formation mechanisms of clouds. Nitrogen oxides also contribute to acidification phenomena. They are responsible for eutrophication—an excessive nitrogen enrichment of soils—which has a significant negative impact on biodiversity. In the atmosphere, NO_x, together with VOCs and CO, initiates a chain of photochemical reactions (triggered by sunlight) that lead to the formation of ozone, whose detrimental climate effects have been previously discussed [10, 11, 12]. Finally, particulate matter, which results from the incomplete combustion of carbon-based fuels, poses a serious threat to human health due to its association with premature cardiovascular and respiratory diseases. According to scientific evidence reported by the WHO [13], the most hazardous particles are those with diameters below 10 μm.

The importance of reducing the emission of these harmful substances is currently driving research in the transport sector toward innovative solutions, including the improvement of existing propulsion systems, the adoption of fuels with lower environmental impact, and the implementation of advanced propulsion technologies. Research efforts are focused on the major actors in the transportation landscape, with the involvement of both public and private entities worldwide. This work specifically focuses on the maritime sector, which currently accounts for over 48% of global goods transport [14], as shown in Figure 1.4.

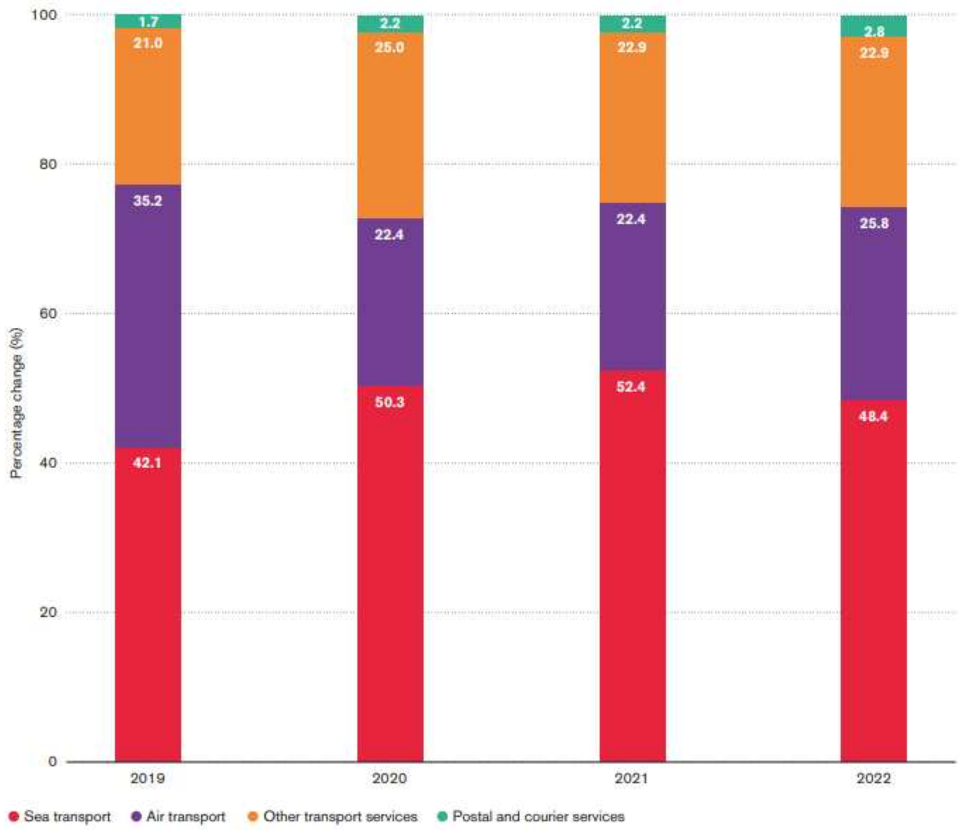

Figure 1.4: Distribution of global trade in goods and transport services between 2019 and 2022 [14].

Although maritime transport plays a fundamental role in the global economy and is among the most energy-efficient modes of transport, it is also a major and growing source of greenhouse gas emissions and other pollutants. In 2018, global shipping emissions amounted to 1,076 million tonnes of CO_2, accounting for approximately 2.9% of total anthropogenic emissions. In Europe, this share rises to 3–4% of total CO_2 emissions, corresponding to over 124 million tonnes in 2021. Projections indicate that these emissions could increase by up to 130% by 2050 compared to 2008 levels. In such a scenario, the objectives of the Paris Agreement would be compromised, and it would not be possible to limit global warming below 2 °C [15, 16]. Therefore, effective global measures are needed to significantly reduce emissions from international maritime transport.

In the maritime sector, reciprocating internal combustion engines are currently the most common choice for both propulsion and onboard power generation [17]. Large marine diesel engines (either two-stroke or four-stroke) have been widely used since World War I. These are highly efficient systems capable of maintaining high performance over a wide range of operating conditions. The primary fossil fuels used in conventional marine engines are HFO (Heavy Fuel Oil) and MGO (Marine Gas Oil). HFO is produced from the residual fractions of crude oil after the higher-quality hydrocarbons have been extracted through thermal distillation or catalytic cracking processes [18]. Although economically advantageous (which explains its widespread use), HFO contains high levels of sulfur and other impurities.

On the other hand, MGO refers to marine fuels composed solely of distillates—i.e., crude oil components that evaporate during fractional distillation and are subsequently condensed [19]. The increased use of these lower-sulfur fuels follows the restrictions imposed by the International Maritime Organization (IMO) on emissions from marine transport. The current IMO regulations, outlined in MARPOL Annex VI and in effect since May 2005 [20], set limits on global emissions of sulfur oxides, nitrogen oxides, and substances that may deplete the ozone layer. Some areas of the world, designated as ECAs (Emission Control Areas), enforce even stricter standards for SO_x, NO_x, and particulate matter.

To comply with increasingly stringent environmental regulations, one of the approaches pursued by the scientific community involves replacing HFO and MGO with alternative fuels capable of reducing pollutant emissions. An immediately feasible approach involves retrofitting existing systems to allow for the use of lower-impact fuels, with higher hydrogen-to-carbon ratios and no sulfur content. This strategy enables the reduction of carbon and sulfur oxide emissions, but does not improve NO_x or particulate emissions. In this respect, diesel engines are at a disadvantage compared to other technologies available on the market. Specifically, the operational and design characteristics that enable high efficiency in these engines also make them prone to high NO_x and particulate emissions [12], which are closely linked to the peak temperatures reached within the combustion chamber.

The growing need to adopt cleaner fuels in order to comply with increasingly stringent maritime emission regulations is guiding the scientific community to explore, at the same time, alternative propulsion systems to the more common diesel internal combustion reciprocating engines (ICEs). One promising solution is the use of renewable-based propulsion systems: ideally, basing maritime transport on renewable energy sources would solve most of the environmental issues under consideration. However, the implementation of new energy sources such as solar and wind still represents a significant challenge. Among the various difficulties to overcome, the main ones are both technical and economic: corrosion due to saltwater exposure, the need for high-efficiency solar panels or wind turbines, limited deck space for optimal installation of such components, the need to redesign the vessel to ensure stability and resilience to sudden weather changes, and high capital costs [21, 22, 23, 24]. Moreover, the intermittent nature of energy generation from these sources arguably represents the main reason why they have not yet become the dominant energy system onboard, as they must be constantly supplemented by dispatchable power systems that can meet the instantaneous demand of onboard consumers.

At present, the most promising dispatchable systems for playing a key role in the energy transition include nuclear plants (which have zero CO_2 emissions), combined cycle power plants, and fuel cells.

Fuel cell propulsion systems exhibit efficiencies comparable to those of large marine diesel engines. Fuel cells convert various types of fuels and oxygen into electricity through an electrochemical reaction, offering several advantages such as higher electrical efficiency (40%–65%) [25], greater system reliability, and lower maintenance costs [26]. However, their adoption is currently limited by the complexities associated with the onboard storage and transportation of the fuel typically used, i.e., hydrogen (H_2). Additionally, the application of fuel cells in maritime contexts is hindered by their low power density, which necessitates the installation of large modules onboard.

A possible interim solution to reduce the environmental impact of maritime transport, particularly in the case of large vessels, is the use of aero-derivative gas turbines for propulsion. Interest in gas turbine systems lies especially in their rapid start-up capability. This characteristic, uncommon among power plants, allows them to cover sudden peaks in energy demand that cannot be addressed by systems operating under steady output conditions, such as nuclear power plants. Many other characteristics make gas turbines ideal candidates for serving as a transitional energy system in maritime transport. In particular, their high power-to-weight ratio and compact design result in reduced onboard space requirements, which represent a critical advantage in naval applications where volume and mass constraints are stringent. Moreover, one of the main advantages of gas turbines lies in their ability to operate on a wide variety of fuels without requiring significant modifications. This fuel flexibility not only offers a broad margin to achieve the cleanest possible combustion [27], but also positions gas turbines as a valuable testing ground for the development of new technologies. In this regard, they could play a key role in enabling the safe and widespread adoption of hydrogen as an energy carrier, paving the way for its future integration into other propulsion systems as well.

While their performance can be significantly enhanced through integration into combined-cycle systems, the main drawback of using gas turbines for marine propulsion lies in their inability to achieve the high efficiencies typical of large diesel ICEs. Modern large-scale terrestrial combined cycle power plants, with capacities in the hundreds of megawatts, can reach thermal efficiencies of approximately 60%. However, when adapted for marine propulsion, these systems are scaled down to sizes on the order of 20–30 MW, which entails smaller turbomachinery and, generally, lower efficiency.

Building upon the considerations discussed so far, the growing interest in gas turbines for marine propulsion, especially in commercial applications, is accompanied by the urgent need to improve their overall efficiency in order to meet increasingly stringent environmental and economic targets. This challenge drives extensive efforts within the scientific community to better understand and mitigate performance losses occurring throughout the machine. In this context, each component—ranging from compressors to turbines—must be analyzed in detail to identify the underlying mechanisms of aerodynamic inefficiencies, unsteady phenomena, and thermal or acoustic losses.

The operation of each gas turbine component is governed by a wide range of interacting variables, which makes a fully analytical description of their behavior unfeasible. To overcome this limitation, high-fidelity simulations such as Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS) provide powerful tools for capturing the complex, multi-scale dynamics of the flow within turbomachinery. Traditionally, loss production mechanisms have been assessed as a collective entity; however, one of the main engineering goals today (especially given the availability of large databases from LES or DNS) is to decompose and identify the distinct sources of losses, thereby offering new perspectives for optimizing the efficiency of modern turbomachines. Indeed, the development of optimization procedures that selectively target and mitigate specific loss sources, rather than focusing solely on overall loss, may lead the designers to more effectively improve the performance of new-generation turbines. This approach requires the introduction of advanced techniques for the analysis of extensive datasets, among which the Proper Orthogonal Decomposition (POD) technique, introduced by Lumley in 1967 [28], is now established as a mature tool for identifying various dynamics driving complex turbulent flows and related loss production mechanisms [29, 30, 31, 32, 33, 34, 35, 36]. By leveraging the mathematical properties of POD, it is possible to reformulate classical loss metrics (such as stagnation pressure drop or turbulent kinetic energy production) in terms of mode contributions, allowing for the targeted analysis of distinct phenomena acting within the computational domain.

The present thesis is focused on the development and application of different modal decomposition techniques for the analysis of unsteady flow fields in turbomachinery, with the aim of providing an advanced tool to support design. In particular, such techniques are here applied to high-fidelity LES datasets of highly loaded low-pressure turbine (LPT) blades, to identify entropy generation sources from a spatial, temporal, and dynamical perspective, as well as in relation to the frequency content of the associated vortical structures, ultimately providing a physically grounded basis for loss reduction strategies. A dataset constructed from the results of such detailed analyses can also be employed to train a predictive model for the loss distribution starting from a given blade geometry. The intrinsic features of the post-processing routine adopted, common to all the loss decomposition techniques presented, make it easily extendable to the analysis of any component of a gas turbine. The ultimate goal is to provide a powerful and versatile tool for the design of next-generation aeroderivative gas turbines for naval propulsion, with improved performance and efficiency, through a more selective and physically informed approach to the identification and mitigation of specific loss sources.

The manuscript is organized as follows: Chapter 2 provides an overview of gas turbine systems for marine propulsion, with particular emphasis on aeroderivative gas turbines and combined-cycle configurations. Their main advantages and limitations are discussed, together with representative applications drawn from the open literature. Chapter 3 introduces low-pressure turbines, focusing on the principal sources of efficiency losses, which represent the physical background of the present investigation. In Chapter 4, the problem of loss evaluation is addressed. The role of turbulence in loss generation is discussed, together with an overview of numerical approaches for flow-field modeling and a comparison between global and local loss assessment methodologies. Chapter 5 presents the modal decomposition techniques employed throughout the work. The high-fidelity database used in this study is described in Chapter 6, while Chapter 7 is devoted to the post-processing routine developed for loss analysis. The overall workflow is presented, together with implementation details of the MATLAB and Python procedures and scalability considerations. The results of the loss analysis are presented and discussed in Chapter 8. A comparative assessment of the different modal decomposition techniques is provided, highlighting the physical insights gained from the proposed methodology. Finally, Chapter 9 introduces the predictive model for loss characterization and design optimization, trained from the results of the presented detailed loss analysis procedure.

Chapter 2

Gas Turbine Systems in Marine Propulsion

2.1 Aeroderivative gas turbines

As already mentioned in the Introduction, a possible temporary solution to mitigate the environmental impact of maritime transport, particularly in the case of large vessels, is the use of gas turbines derived from the aeronautical sector. The main interest in gas turbine plants lies in their fast start-up capability. This feature, uncommon among other types of power plants, ensures the ability to handle sudden peaks in demand, which cannot be managed by systems operating predominantly under steady production regimes, such as nuclear plants. Moreover, these machines offer a promising solution to the transitional needs of maritime transport, bridging the gap between conventional marine engines and future low-emission technologies. Among gas turbines, aeroderivative models (originally developed for aviation) stand out due to their high power-to-weight ratio, modularity, and adaptability to various fuels. One notable example is the GE *LM2500* (Figure 2.1).

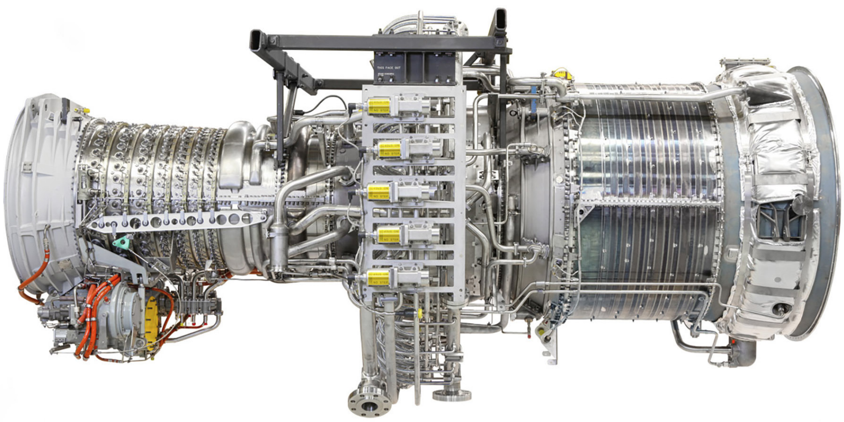

Figure 2.1: Aeroderivative gas turbine *LM2500* by General Electric [37].

Gas turbines are continuous-flow machines whose operation is based on the open

Brayton-Joule cycle. Referring to Figure 2.2, the cycle involves the adiabatic compression of a gaseous fluid, typically atmospheric air, which serves as the oxidizer in a combustion chamber. Within the chamber, the chemical energy of a fuel is utilized to increase the enthalpy of the incoming fluid. This enthalpy rise is achieved through combustion in an ideally isobaric process, and the resulting high-enthalpy combustion gases are expanded through a turbine to produce mechanical work. The work extracted from the gas expansion is used to drive the compressor and to produce the desired useful effect.

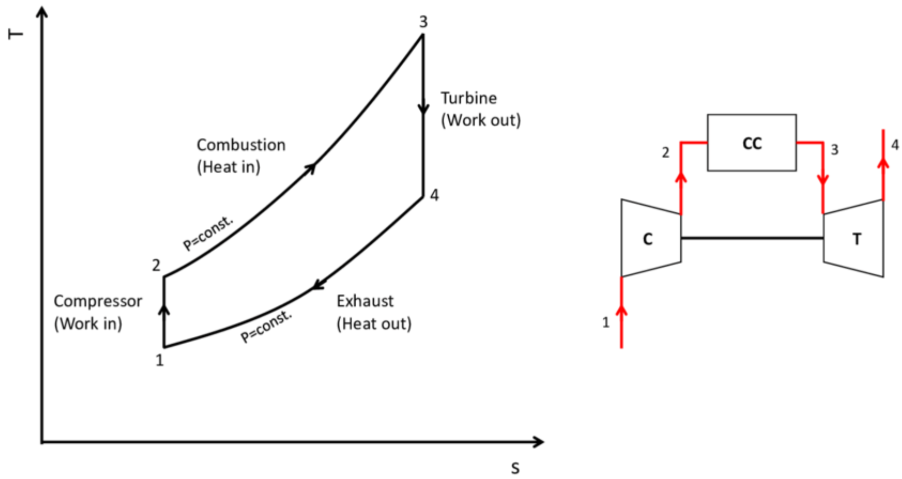

Figure 2.2: T-s diagram and process scheme of the ideal Brayton-Joule cycle [38].

Depending on the application, downstream of the expansion stage that provides power to the compressor, another power turbine and/or a propulsive nozzle may be found (the latter only in aeronautical applications). Large gas turbines, due to the high mass flow rates involved, are generally of the axial type: in these machines, the main direction of the fluid is parallel to the rotation axis. In such cases, both the compression and expansion stages (including any power turbine) are axial machines, as schematically shown in Figure 2.3. The engine depicted is a simple turbojet, in which the residual enthalpy of the gases leaving the expansion turbine is utilized in a nozzle to produce thrust. In other types of aircraft engines and in heavy-duty turbines, this component is either preceded or entirely replaced (in terrestrial and marine applications) by a power turbine, which is generally mounted on a shaft coaxial with the main turbine. This power turbine is used to provide mechanical power to a generic user, such as a propeller or an alternator. Being mechanically decoupled from the high-pressure compressor and turbine, the power turbine can rotate at a different speed than the rest of the machine. This allows a design tailored specifically to the requirements of the connected user.

The main difference between heavy-duty and aeroderivative turbines lies in their technological level. Aeroderivative turbines, in fact, benefit from advanced technologies that make them particularly suitable for naval applications. They are considerably lighter and more compact, as previously mentioned, while also offering higher thermal efficiency. Furthermore, aeroderivative turbines can respond rapidly

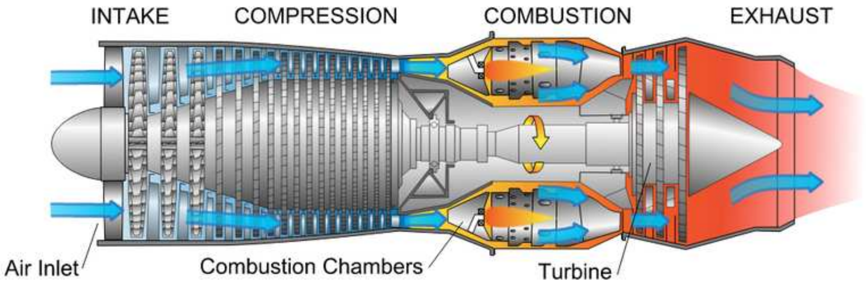

Figure 2.3: Simplified meridional section of a simple turbojet engine for aeronautical applications [39].

to load variations, providing high operational flexibility, and their modularity facilitates maintenance and repair, reducing downtime and improving the operational availability of the ship.

The use of aeroderivative turbines in marine applications, however, requires some adaptations to ensure that engines originally designed for aviation can operate effectively under maritime conditions. The marine environment is highly corrosive due to the presence of salt water, making it essential to apply specialized coatings and corrosion-resistant materials on turbine components to prevent damage and extend service life. In this regard, advanced air filtration systems are necessary to remove salt, dust, and other impurities from the intake air that could otherwise damage internal turbine components.

2.2 Combined cycles for marine propulsion

Although aeroderivative gas turbines do not reach the high efficiencies of large reciprocating diesel ICEs, their performance can be significantly improved when integrated into combined cycles.

A combined cycle power plant generally consists of two thermodynamic cycles arranged in series: a "top" cycle generates power, and a portion of the heat rejected to the lower-temperature sink is recovered and used to supply heat to a "bottom" cycle operating beneath it [40]. The goal of combined cycles is to increase the total useful work output, either by extracting more work from the same amount of fuel, as in *unfired* configurations, or by achieving a larger overall power output through additional fuel input in *fired* configurations. In both cases, the thermodynamic benefit arises from a more favorable distribution of heat addition and discharge over temperature, which can be achieved either by increasing the average temperature at which heat is provided or by decreasing the average temperature at which heat is discharged, thereby approaching the ideal Carnot cycle. Combined cycles using one or more gas turbines for the top cycle and a steam cycle for the bottom cycle are a widely adopted solution for power generation due to their high efficiency and intrinsically low pollutant emissions. The basic configurations for applying combined cycle systems to marine propulsion are known by the acronyms COGAS (*Combined*

Gas And Steam) and COGES (*Combined Gas turbine Electric and Steam*), whose simplified layouts are shown in Figure 2.4.

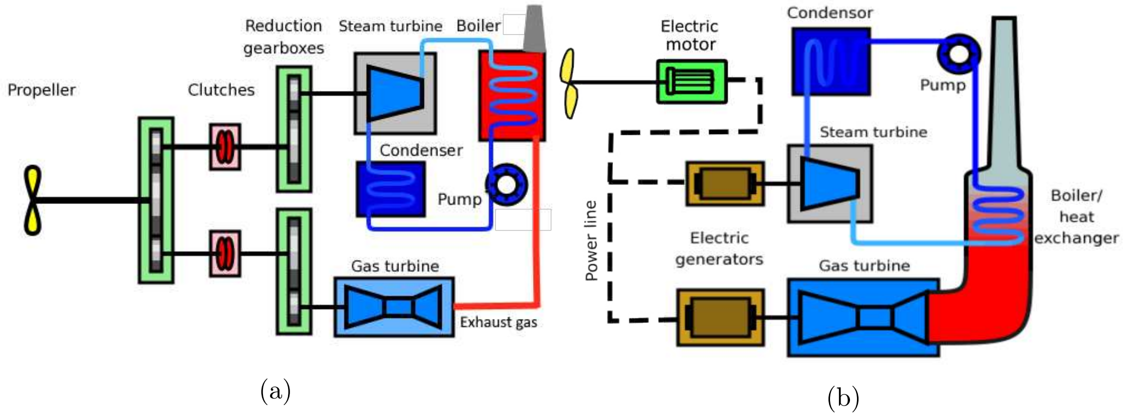

Figure 2.4: Simplified layouts of COGAS (a) and COGES (b) marine propulsion systems [41].

In COGAS systems (Figure 2.4a), a mechanical coupling via a reduction gearbox connects the propeller to the steam and gas turbines. In contrast, the COGES configuration (Figure 2.4b) converts the mechanical work produced by the gas and steam turbines into electrical energy via generators. This electrical energy can be stored or used directly to drive the propeller through an electric motor [42, 43, 44].

COGES systems are often regarded as a promising solution for marine propulsion, as they allow the turbines to operate at their optimal rotational speed over a wide range of vessel operating conditions, provided that a sufficiently high energy conversion efficiency is achieved. Furthermore, since the system is decoupled from the actual propeller shaft, the installation layout can be selected independently of the propeller position, thereby alleviating some of the plant design constraints associated with the use of reduction gearboxes, as required in COGAS systems. However, the additional energy conversion stages from mechanical to electrical power and back to mechanical power may lead to efficiency penalties and increased investment costs, which currently prevent COGES configurations from being considered a universally optimal solution.

2.3 Advantages/limitations of turbogas applications

The potential of gas turbines in maritime transport makes them an ideal candidate for serving as a transitional energy system. While the use of hydrogen as an energy carrier holds great promise in terms of emission reduction, gas turbines could provide a testing ground for the development of new technologies enabling its safe and widespread future use in other propulsion systems as well.

In addition to the aforementioned fast start-up capability, one of the main advantages of applying gas turbines to marine propulsion lies in their ability to operate on a variety of fuels without requiring significant modifications. Gas turbines, in particular, are continuous-flow machines characterized by a stable flame dynamic under steady-state conditions. These features reduce constraints on the required fuel properties for efficient, safe, and reliable combustion (for instance, there are no strict requirements related to octane number or cetane index), thus offering a broad margin to achieve the cleanest possible combustion [27].

The use of pure gaseous hydrogen is still difficult to implement due to the molecule's intrinsic physicochemical properties, which introduce several technical challenges. However, several studies [45, 46, 47, 48, 49] have explored the feasibility of blending increasing amounts of H_2 into other conventional gaseous fuels, such as Liquefied Natural Gas (LNG). This approach could temporarily mitigate the issues related to storing large volumes of hydrogen and burning a particularly unstable gas, while awaiting the development and consolidation of innovative solutions.

When considering conventional fuels only, gas turbines offer the advantage of reaching lower maximum temperatures in the combustion chamber compared to typical reciprocating ICEs. This mainly stems from the differences in how these two types of machines perform the thermodynamic transformations that characterize their operation. In both cases, a compression phase of the oxidizer is followed by fuel injection and combustion, resulting in a sharp temperature rise of the working fluid. In reciprocating internal combustion engines, this process occurs cyclically inside the cylinder, whose walls are thus exposed to the peak combustion temperatures for only a limited portion of each cycle. This is advantageous because it allows for higher maximum source temperatures, thereby improving thermodynamic efficiency.

From an emissions perspective, however, high combustion temperatures promote the formation of so-called thermal NO_x. In contrast, in a continuous-flow machine like a gas turbine, the flame is continuously present and anchored at a fixed location within the combustion chamber, requiring lower peak temperatures to be maintained. This ultimately leads to lower NO_x concentrations in the exhaust gases [50].

From a plant engineering perspective, gas turbines also require less space in the engine room compared to typical reciprocating diesel engines. This is primarily due to their higher specific power. Figure 2.5 shows a comparison between the footprints of a diesel ICE and a combined gas–steam cycle within the engine room of the same cargo vessel. The availability of additional onboard space can be exploited either to increase payload capacity or to store extra fuel, thus enabling longer voyages for the same external hull volume. This advantage is accompanied by a significant weight reduction of the entire system (especially when the gas turbine is derived from aeronautical applications) which in turn leads to lower fuel consumption and reduced pollutant emissions. The long-term economic benefits that arise from these improvements can offset the higher fuel costs and the substantial initial capital required for design and installation [12, 51, 52].

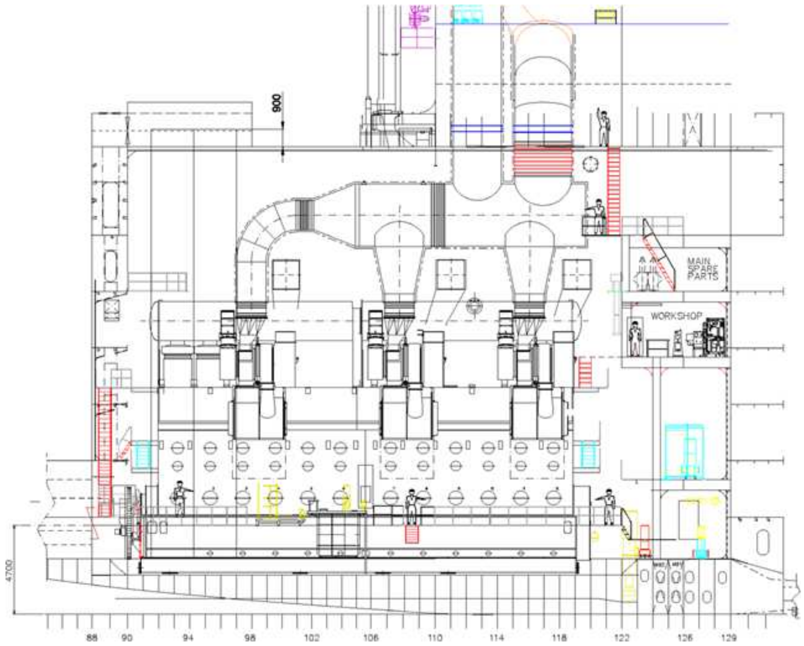

(a)

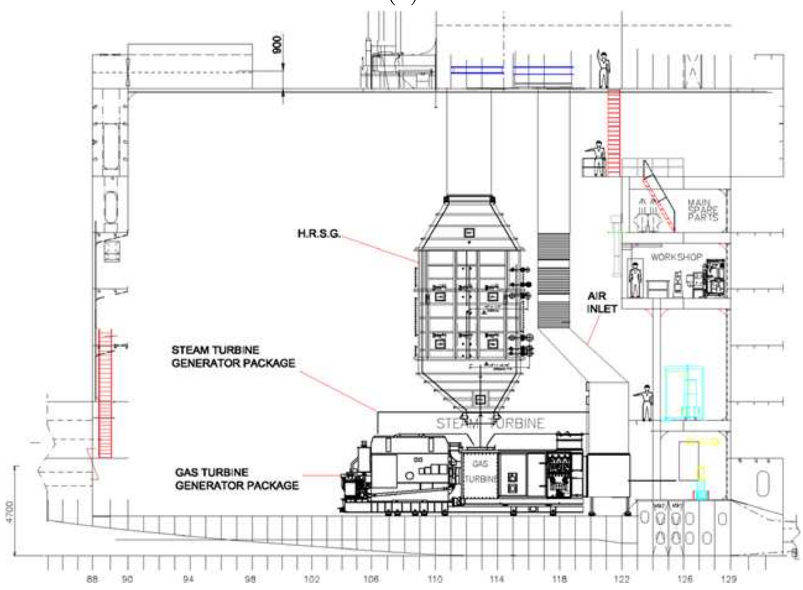

(b)

Figure 2.5: Engine room of a cargo vessel equipped with a reciprocating ICE (a) and a combined gas–steam cycle (b) [51].

2.4 Examples of applications from the literature

In the following, several examples of ships propelled by gas turbines are presented, some of which are integrated into combined gas-steam cycle plants. These vessels, built over the past twenty years, illustrate the practical applications and operational versatility of gas turbines in modern naval propulsion.

Gas turbines have become the propulsion system of choice for military vessels due to

their high power-to-weight ratio, excellent reliability, fast start-up times, and rapid response to load changes. In this context, their relatively low efficiency compared to reciprocating diesel engines becomes a secondary issue, outweighed by all the aforementioned advantages.
The *LM2500* by General Electric (GE) Marine Solutions, for instance, is the main gas turbine model powering the surface fleet of the United States Navy. More than 700 units of this system are installed aboard a wide range of platforms, including the *Ticonderoga*-class cruisers, *Arleigh Burke*-class destroyers, the *USS Makin Island*, the *America*-class amphibious assault ships, and the *Legend*-class National Security Cutters of the U.S. Coast Guard.
Globally, the LM2500 (in its various configurations, such as the *LM2500+* and *LM2500+G4*) is also installed on several other naval vessels: the Canadian Halifax-class frigates, the Italian and French FREMM-class frigates, the Italian PPA-class light frigates, the Turkish MILGEM-class corvettes, the German F124 and F125 frigates, the South Korean KDX-class destroyers, and the Indian P17-class frigates, with a total of 1250 units delivered worldwide [53].
The *MT30* gas turbine developed by Rolls-Royce is currently the most powerful marine gas turbine available in terms of power-to-weight ratio. It is employed on some of the most advanced naval platforms in the world, including the *Queen Elizabeth*-class aircraft carriers, the Royal Navy's *Type 26 Global Combat Ships* and their international variants, the U.S. Navy's *Freedom*-class Littoral Combat Ships and the advanced *DDG-1000* destroyer, the *Daegu*- and *Ulsan*-class frigates of the Republic of Korea Navy, the Italian Navy's landing helicopter dock *Trieste*, and the *Mogami*-class 30FFM frigates of the Japan Maritime Self-Defense Force [54].

Table 2.1: Technical specifications of the aeroderivative gas turbines LM2500 and MT30.

Model	Dry weight [mt]	Dimensions LxWxH [m]	Net power output [MW]	Output shaft speed [rpm]	Thermal efficiency
GE LM2500	18.4	8.0 x 2.6 x 2.4	22.7	3600	Up to 38.5%
Rolls-Royce MT30	30	8.1 x 3.5 x 3.1	38	3600	Up to 40%

In the commercial and cruise sectors, the first gas turbine-powered cruise ship was the *GTS Millennium*, operated by Celebrity Cruises (Figure 2.6), which entered service in 2000. The vessel was powered by a COGES (Combined Gas turbine Electric and Steam) system with turbo-electric transmission, comprising two LM2500+ aeroderivative gas turbines rated at 25 MW each and one 9 MW steam turbine. Each gas turbine was coupled to an electrical generator and supplied all onboard power, including that required by the electric propulsion motors. Including service steam extracted from the bottoming cycle, the fuel consumption was up to 7% lower than an equivalent diesel installation, achieving an overall efficiency of around 40% [55, 56]. As expected, exhaust emissions were significantly reduced: nitrogen oxides were lowered by 80% and sulfur dioxide (SO_2) by 98% [57]. According to Sannemann [55], sea trials of the GTS Millennium measured total NO_x emissions of 5 g/kWh for the COGES system, already approaching the IMO Tier III regulatory limits. No visible smoke emissions were recorded at any power output level.

Figure 2.6: The cruise ship *GTS Millennium* operated by Celebrity Cruises [58].

Despite its promise, the Millennium underwent a retrofit in 2007 to install a heavy fuel oil (HFO)-fired diesel engine, aimed at reducing the ship operational costs. However, under current regulations that call for a paradigm shift in marine fuels, a propulsion system of this kind would likely be favored again.

Another application in which gas turbine systems have been considered as the primary propulsion solution is for LNG (Liquefied Natural Gas) carriers. LNG is transported in non-pressurized, insulated tanks, in which the temperature is maintained constant via passive evaporation. A fundamental requirement for propulsion systems in this field is the ability to utilize boil-off gas as fuel. Both reciprocating diesel engines and combined gas–steam cycles could replace the traditional steam turbine plants, hence reducing the amount of LNG that must be intentionally evaporated to supplement the boil-off gas. However, in order to comply with current emission regulations, diesel engines must be equipped with expensive and bulky emission reduction systems (especially for NO_x). The economic benefits, lower emissions, and intrinsic redundancy of the COGES configuration make it the most suitable alternative to steam power plants on this type of vessel [44, 59, 60, 61].

Chapter 3

Low-Pressure Turbines

As anticipated above, this work focuses on the development and application of modal decomposition techniques to improve the analysis of unsteady flow fields in aeroderivative gas turbine components. For the purposes of this study, a specific component of an aeroderivative gas turbine was chosen as the case study: the low-pressure turbine (LPT).

Among the different components contributing to engine efficiency, turbine blades play a crucial role in converting the internal energy of high-velocity gas into mechanical work. In particular, the low-pressure turbine is one of the components with the greatest impact on the overall performance of a gas turbine. In modern aeroderivative applications, the LPT can account for up to 30% of the total engine weight and is fully responsible for converting the residual kinetic energy of the hot gases coming from the combustion chamber and the high-pressure turbine into mechanical work to drive the connected components. Specifically, the low-pressure turbine is responsible for extracting useful energy from the flow to generate the power required to drive, depending on the system architecture, electric generators in COGES applications or directly the propeller shafts in COGAS architectures. Figure 3.1 highlights the low-pressure turbine of the General Electric *GE9X* turbofan engine, emphasizing the size of the selected component relative to the rest of the engine.

In pursuing this useful output, it is necessary to consider that, in a real machine, a portion of the energy introduced into the system through fuel is not converted into mechanical work but is instead dissipated. This dissipation is associated with various loss mechanisms that cause the system's behavior to deviate from the ideal one.

3.1 Loss sources in LPTs

In turbines, irreversibilities result in a reduction of the actual work extracted from the fluid compared to the ideal value. The enthalpy–entropy diagram in Figure 3.2 illustrates this phenomenon by showing the difference in enthalpy drop between the

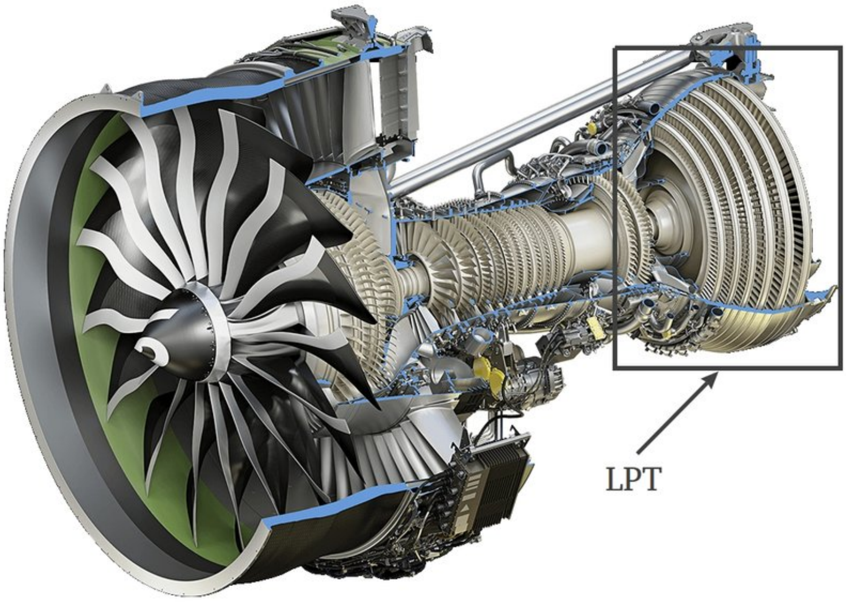

Figure 3.1: Cutaway of the GE9X aeronautical turbofan engine, with the low-pressure turbine highlighted as the focus of this study.

real expansion 1–3 and the ideal isentropic expansion 1–3ss: this difference results in a reduction of recoverable energy.

By adopting a reference frame fixed with the blade rows (absolute for stators and relative for rotors), the total enthalpy (absolute or relative, respectively) remains constant. According to the Gibbs relation for total quantities:

$$dh_t = T_t dS + \frac{dp_t}{\rho_t} \tag{3.1}$$

it is evident that, under adiabatic conditions ($dh_t = 0$), a variation in total pressure is an indicator of entropy increase, and hence of irreversibility. For this reason, in turbomachinery, losses are usually quantified in terms of total pressure drop, using appropriate coefficients.

In low-pressure turbines operating in the marine environment, viscous mechanisms are the predominant source of irreversibility. Unlike high-pressure stages, where shock waves play a significant role, the subsonic flow regime of LPTs makes shock-related losses negligible, while the limited residence times and negligible heat exchange justify the adiabatic assumption. Viscous effects convert kinetic and pressure energy into internal energy, which cannot be effectively recovered as useful work in the rotor.

3.1.1 Profile losses

The unsteady aerodynamics developing in a low-pressure passage arise from the superposition of various phenomena taking place in different regions of the passage, each characterized by distinct temporal and spatial scales. Each of them differently

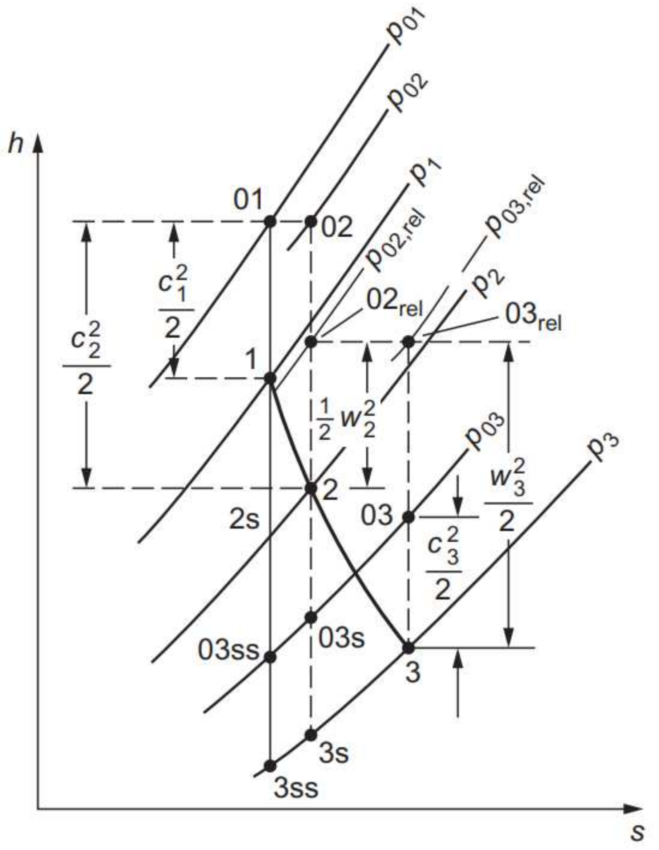

Figure 3.2: Representation of real and ideal expansion on the $h - s$ diagram.

affects the generation of losses, i.e. entropy, throughout the entire passage, within and outside the blade boundary layers. The onset of these phenomena, as shown in Figure 3.2, causes the real expansion curve to end with a higher enthalpy than the ideal one, meaning that the subsequent nozzle (an isenthalpic transformation) releases a more energetic fluid to the environment or to the next stage, resulting in a net efficiency loss. All these mechanisms can be traced back to the so-called profile losses, which typically represent the dominant contribution in LPT blades with a high aspect ratio. In such configurations, the radial flow variation is limited, and a large portion of the blade operates under quasi-two-dimensional flow conditions. Profile losses arise from the formation and growth of the boundary layer along the suction and pressure sides of the blade, which may undergo separation near the trailing edge under adverse pressure gradients. The wake generated at the blade exit then propagates downstream and periodically interacts with subsequent rows. The energy contained in the unsteady turbulent structures ultimately resulting from this interaction is eventually dissipated into internal energy and is no longer available for work extraction.

Denton's seminal work [62] clarified how different mechanisms, strongly dependent on blade geometry, Reynolds number, Mach number, and pitch-to-chord ratio, govern the onset of viscous losses in LPT cascades. Going into further detail, as shown in Figure 3.3, the boundary layer grows on both the pressure and the suction sides, creating a wake (i.e., a velocity defect marked in cyan in the figure) downstream of the blade. The amount of loss is therefore determined by the evolution of the boundary layer and depends on its state (laminar or turbulent); thus, it also relies on its transition process. The flow conditions and the shape of the blades fundamentally influence these phenomena, especially on the suction side where the flow is subjected to an adverse pressure gradient, as can be clearly seen by the blade loading distribution reported as an example in Fig. 3.4. If the adverse pressure gradient is particularly strong, flow separation may also occur.

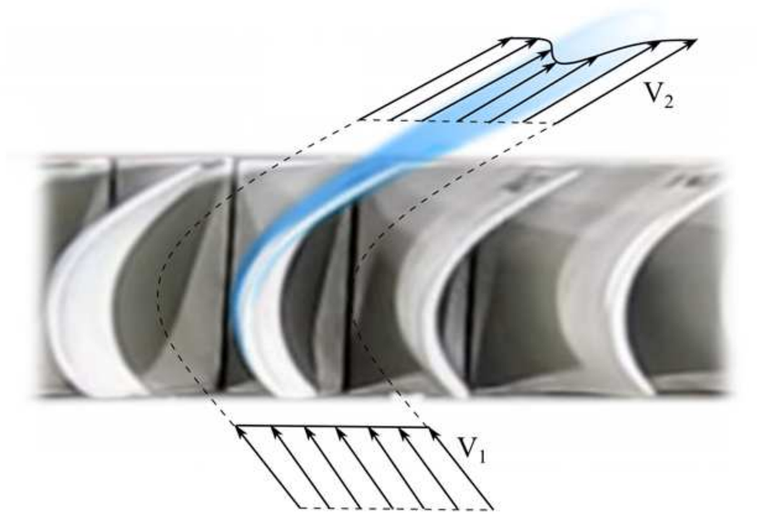

Figure 3.3: Sketch of the velocity defect created by the boundary layer development.

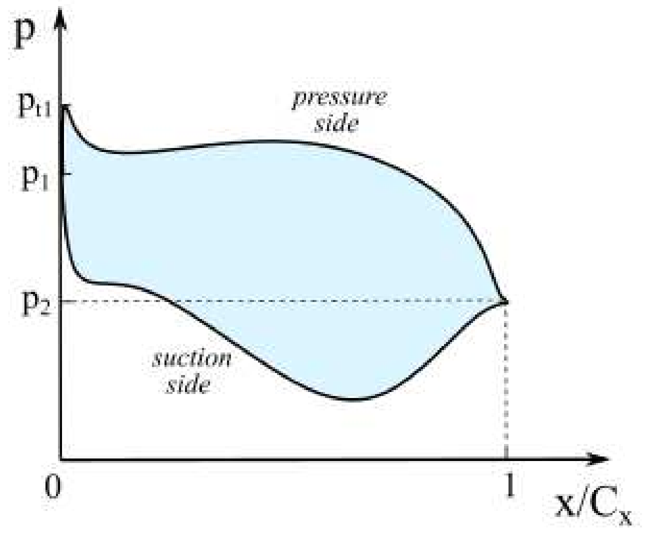

Figure 3.4: Example of pressure distribution over an LPT blade.

Figure 3.5 shows the evolution of the boundary layer on a flat plate. A first region of laminar boundary layer is evidenced on the left. Here, the flow motion is parallel to the wall and the velocity smoothly decreases from the external velocity at the edge of the boundary layer to zero at the wall. Then, the appearance of flow disturbances and instabilities defines a transition region, that leads the flow toward the turbulent regime. In the end of the plate, the boundary layer is fully turbulent, i.e., the flow motion is chaotic and characterized by oscillating velocity components. A laminar sub-layer, however, persists near the wall. The figure also points out the correlations for the laminar and the turbulent boundary layer growth.

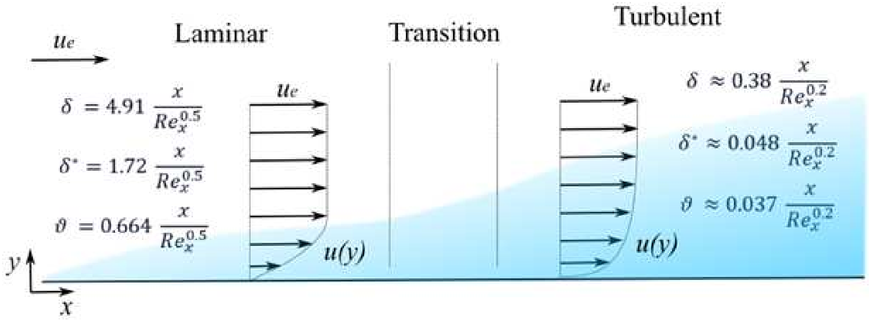

Figure 3.5: Boundary layer evolution on the flat plate.

In the figure, δ is the boundary layer thickness, δ^* the displacement thickness, i.e., the mass flow rate loss due to the presence of the boundary layer (Eq. (3.2)), and θ is the momentum defect (Eq. (3.3)). The correlations valid for flat plate show that the turbulent boundary layer is subjected to a higher growth rate, and higher mass flow and momentum defects, with respect to the laminar boundary layer.

$$\delta^* = \int_0^\delta \left(1 - \frac{u(y)}{u_e}\right) dy \tag{3.2}$$

$$\theta = \int_0^\delta \frac{u(y)}{u_e} \left(1 - \frac{u(y)}{u_e}\right) dy \tag{3.3}$$

The flow behavior is more complex in the case of the presence of an external pressure (or velocity) gradient, as along turbomachinery blade surfaces, especially when the pressure gradient is unfavorable (i.e., $d_x p > 0$, $d_x u_e < 0$). In this case, the velocity profile in the boundary layer is impoverished, and flow separation may occur. The flow separates when the deceleration is strong enough to induce a reverse flow region near the wall (as depicted in Fig. 3.6), i.e., the flow inside the separated region moves in the opposite direction with respect to the main flow. Separation may occur for both the laminar and the turbulent boundary layers, however the turbulent boundary layer is generally more resistant to separation with respect to the laminar one.

Since the boundary layer growth, the generated losses, and also the resistance to separation are very different between the laminar and the turbulent states, it is fundamental to investigate where and how transition takes place. As anticipated, transition is the process by which the boundary layer evolves from the laminar to the turbulent regime. Different mechanisms may occur:

- Natural transition

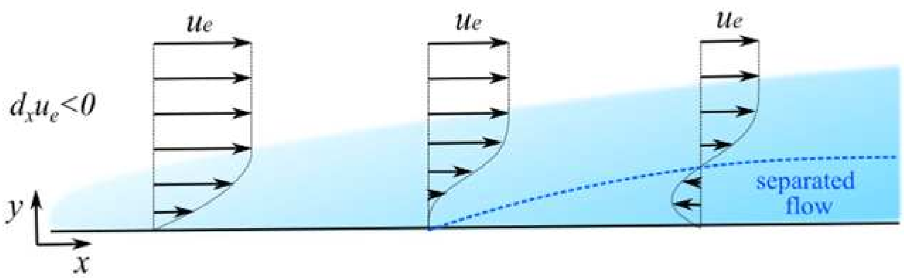

Figure 3.6: Boundary layer separation.

- By-pass transition
- Separation-induced transition
- Unsteady wake-induced transition

The natural transition is a relatively slow process in which linear oscillations, referred to as Tollmien–Schlichting waves, are amplified to the point of generating turbulent spots that propagate and, ultimately, define a fully turbulent layer. It occurs for low values of the free-stream turbulence intensity ($Tu < 0.1\%$, where Tu is the ratio between the RMS and the mean of the velocity), and for this reason it is not particularly interesting in the field of turbomachinery. Indeed, the internal aerodynamics of turbomachines, and especially of turbines, is characterized by much higher Tu levels. Thus, transition in turbomachines is usually a faster process where the amplification of the Tollmien-Schlichting waves is by-passed and the turbulent spots are directly induced by the free-stream disturbances (by-pass transition). The mechanisms by which the free-stream disturbances induce the formation of boundary layer streaks [63] and the instabilities that cause their breakdown into turbulent spots are complex and, in recent years, they became the subject of many detailed investigations [64].

The separation-induced transition may instead occur when the laminar boundary layer separates. This is the case when the boundary layer is subjected to a strong adverse pressure gradient, as in high-loaded turbine profiles. The separation causes the presence of a shear layer between the flow inside and outside of the separation bubble, that is subjected to Kelvin-Helmholtz instabilities. The amplification of such oscillations induces the formation of vortices, whose breakdown causes transition in the flow over the separation region. Downstream, the mixing effect induced by the turbulence may also promote flow reattachment.

The unsteady wake-induced transition is proper of the unsteady turbomachinery environment. As anticipated, the real operation of a turbine blade row is characterized by the aerodynamic interaction between stator and rotor, which makes the flow through the blade passage unsteady. The relative motion between stationary and rotating blades indeed causes the presence of upstream wakes, that periodically evolve in the downstream blade passage. Figure 3.7 schematically presents the upstream wake dynamics. As shown in many recent studies [65, 66, 67, 68, 69, 70, 71, 72], the upstream wakes are conveyed by the pressure gradient toward the blade suction side, where they interact with the evolving suction side boundary layer. Since the wake is a flow region characterized by high turbulent fluctuations, the impinging wakes can alter the boundary layer transition process by directly inducing the formtaion of local turbulent spots. This spots propagates similarly to those generated during the by-pass transition, periodically anticipating the transition.

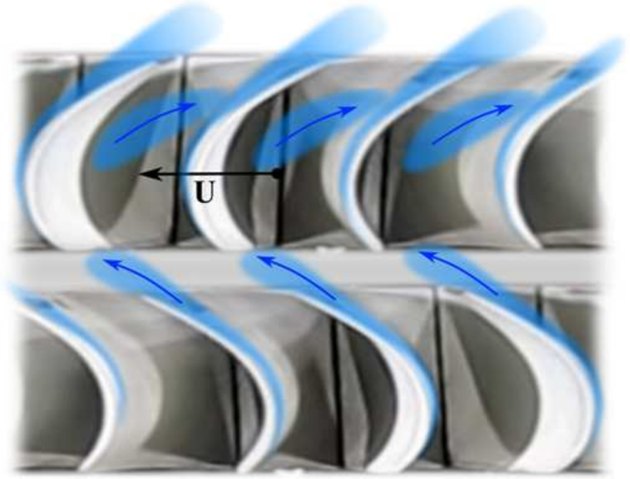

Figure 3.7: Sketch of the unsteady operation of LPT rows.

The periodic nature of the wake-boundary layer interaction is highlighted in the time-space plot of Figure 3.8, where it is shown that the passage of the wake induce a local turbulent state, which is followed by a calmed region, repeated in a periodic cycle. This anticipates the mean onset location of the transition process, and, in the case of separated flow, it may also promote flow reattachment, with overall positive effects (see, for instance, the work of Gompterz and Bons [73]). In all other scenarios, these mechanisms have a large negative impact on the profile losses. The experimental test carried out by Simoni et al. [74] revealed an increment up to the 50% of the losses in the unsteady operation with respect to the steady one. The authors stated that the observed increment is due to the upstream wakes promoting an earlier transition of the suction side boundary layer, as also pointed out by the investigations carried out by Stieger and Hodson [67]. Moreover, for highly loaded turbine blades, the boundary layer growing between successive wakes may undergo an intermittent separation, depending on the loading and diffusion level (Satta et al. [75]), that evidently imposes an inviscid-like instability on the separated shear layer, as shown in the works of Sarkar and Voke [71] and Simoni et al. [76]. In this case, large scale coherent structures are shed as a consequence of the separated shear layer roll-up during interaction between wake and boundary layer.

In recent years, detailed experimental campaigns and high fidelity simulations provided deep insight into the complex wake-related flow dynamics responsible for loss production. The 2-D like low momentum region representing the wake approaching the downstream cascade undergoes strong bowing, tilting, dilatation and stretching processes, and it finally becomes fragmented into segments as a consequence of the

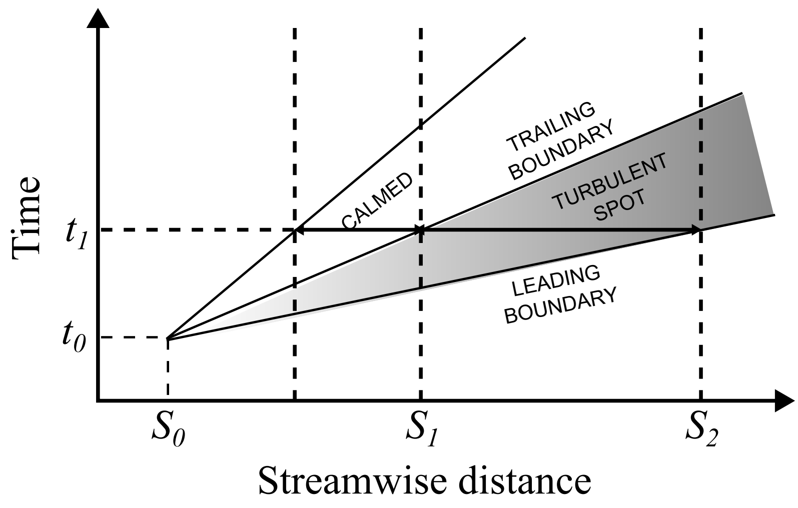

Figure 3.8: Time-space evolution of the boundary layer unsteady wake-induced transition.

relative motion between the stator and rotor rows, as discussed by Binder et al. [77, 78]. Stieger and Hodson [67] provided a detailed experimental characterization of the kinematics governing wake convection through an LPT cascade, clarifying the physical mechanisms by which upstream wakes influence loss generation beyond their direct interaction with the blade boundary layers. Their measurements showed that the wake segment convected through the blade passage behaves as a perturbation jet directed toward its source, commonly referred to as a *negative jet* [79]. When defined with respect to the time-averaged flow field, this perturbation exhibits a significant cross-passage velocity component, which drives the wake fluid toward the suction side of the downstream blade. The negative jet effect causes the wake fluid to impinge on the suction surface, where it splits into two perturbation streams: one convecting downstream along the blade surface and the other directed upstream. As a consequence, the perturbation induced by the wake accelerates the flow downstream of the wake center while decelerating it upstream, locally strengthening the adverse pressure gradient and causing the boundary layer velocity profiles to be more inflectional. Beyond its role in boundary-layer transition, the negative jet dynamics were shown to directly contribute to turbulence production within the blade passage. The wake fluid, characterized by elevated and anisotropic turbulent stresses, experiences strong production of turbulent kinetic energy when convected into regions of high mean strain. This production mechanism occurs predominantly outside the boundary layer and is maximized where the principal turbulent stresses within the wake align with the local strain-rate field [80]. The accumulation of wake fluid near the suction surface, driven by the combined effects of negative jet convection and wake stretching, therefore results in localized regions of enhanced turbulence intensity upstream of the trailing edge. These regions subsequently interact with the suction-side boundary layer, further influencing transition, separation behavior, and loss generation.

3.1.2 Secondary flows

In addition to profile losses, other relevant loss mechanisms in LPTs include secondary flow losses and tip leakage losses, particularly in three-dimensional flow fields.

Secondary flows originate from the interaction between the flow and the endwalls, namely hub and casing, which impart strongly three-dimensional features to the main flow in the regions close to these surfaces. In these areas, complex vortical structures arise, with velocity components transverse to the main flow direction, causing unwanted deflections and total pressure losses. Among the most characteristic structures are the passage vortex, generated by the pressure gradient between the suction and pressure sides, the horseshoe vortex, formed near the leading edge due to the pressure non-uniformity along the blade height, and the corner and trailing edge vortices, related to the interaction of the main vortices with viscous effects at the trailing edge. The magnitude of secondary flow losses strongly depends on the blade aspect ratio: in low-pressure turbines, where the height-to-chord ratio is high, their contribution is reduced compared to other turbomachinery components, although still non-negligible in the overall loss balance.

Another important source of inefficiency is represented by tip leakage losses, which arise from the clearance between the rotor blade tips and the casing wall. The pressure difference between the suction and pressure sides drives a leakage flow through the gap, leading to flow separation, vortex formation, and interactions with other secondary flows. The extent of these losses depends on the ratio between clearance size and blade geometry, being particularly significant in blades with low height. Several design strategies have been developed to mitigate tip leakage effects, such as blade shrouding, the injection of cooling air near the tip to form a fluid-dynamic barrier, or the adoption of squealer tip configurations, which introduce geometric obstacles to reduce the leakage flow across the gap.

Chapter 4

Loss Evaluation Metrics

4.1 The problem of turbulence

In pursuit of ever-increasing performance, fluid dynamic research has progressively moved away from the simplifying assumptions underlying traditional loss coefficients. Nowadays, thanks to the growth in computational power and the availability of advanced measurement systems, even highly complex flow problems are modeled with increasing fidelity to physical reality.

A complete description of fluid motion in space requires the solution of the Navier–Stokes equations, a system of partial differential equations valid for continuous fluids. These include the conservation of mass, momentum, and energy. To date, analytical solutions of the Navier–Stokes equations are only available for a limited number of idealized cases. In all other situations, characterized by complex geometries and boundary conditions, numerical simulations or experimental investigations must be employed to accurately reproduce the relevant flow dynamics. The inherent mathematical complexity of the Navier–Stokes equations has earned them a place among the "Millennium Prize Problems," representing a major challenge for the scientific community.

Turbulent flows, in particular, remain one of the most difficult classes of problems to model due to their highly irregular and chaotic behavior. As of now, reliable predictions of turbulent flows rely on numerical approaches requiring substantial computational resources.

Turbulence is characterized by the presence of vortical structures with seemingly stochastic behavior, which arise when inertial forces dominate over viscous forces—that is, at high Reynolds numbers. This regime is inherently unsteady, irregular, and dissipative, as illustrated in Figure 4.1.

In many practical applications, especially those involving strong flow instabilities, the onset of turbulence is unavoidable. Such instabilities can either trigger or amplify irregularities in the velocity field. Turbulence theory is therefore concerned with the analysis of unsteady and non-uniform flow phenomena. A commonly used tool for

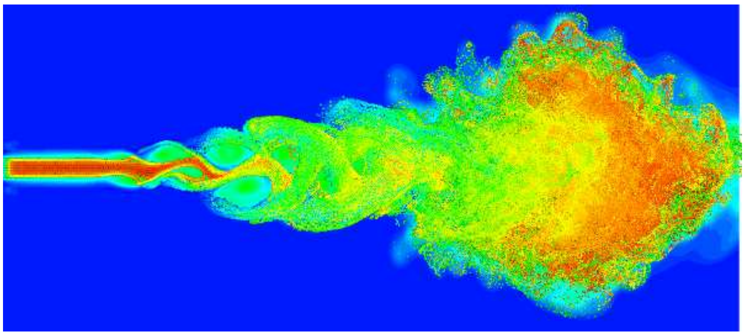

Figure 4.1: Example of a flow undergoing transition from laminar to turbulent regime.

their study is the Reynolds decomposition, whereby any instantaneous flow variable is split into a time-averaged (mean) component and a fluctuating one.

The kinetic energy associated with the fluctuating components is known as turbulent kinetic energy (TKE), typically denoted by k. For a three-dimensional flow with velocity components u, v, and w, the turbulent kinetic energy is defined as:

$$k = \frac{1}{2}\left(\overline{u'^2 + v'^2 + w'^2}\right) \tag{4.1}$$

From the Navier–Stokes equations, a simplified expression for the temporal variation of the turbulent kinetic energy can be derived in the absence of spatial gradients. This can be written as:

$$\frac{\partial k}{\partial t} = \int_V \rho\, \mathbf{u} \cdot \mathbf{f}\, dV - \int_V \rho\, \varepsilon\, dV \tag{4.2}$$

Since the dissipation rate ε is always positive, energy is transferred into vortices of progressively smaller characteristic size, a mechanism known as *vortex stretching*. The energy is cascaded to ever-smaller scales until it is ultimately dissipated by viscous effects. This process is known as the *energy cascade*, first conceptualized by Richardson and later formalized by Kolmogorov [81]. In two-dimensional flows, a dual-cascade process is instead observed.

Kolmogorov's work provides a clearer representation of the range of turbulent scales and the kinetic energy associated with each vortical structure. He introduces the wavenumber κ as:

$$\kappa = \frac{2\pi}{\lambda} \tag{4.3}$$

where λ denotes the wavelength of a turbulent eddy. According to this definition, progressively smaller vortices correspond to shorter wavelengths and thus higher wavenumbers. This relationship allows for the construction of the turbulent energy spectrum, which describes how kinetic energy is distributed among the various spatial scales of turbulence. A qualitative representation of this spectrum is shown in Fig. 4.2.

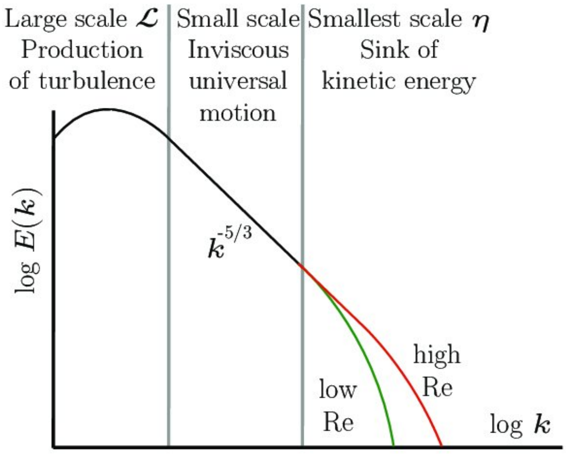

Figure 4.2: Turbulent kinetic energy spectrum as a function of the wavenumber [82]

The characteristic scales of turbulence are defined through key length scales such as the integral scale l_0 and the Kolmogorov microscale l_k. The integral scale l_0 corresponds to the scales at which turbulent kinetic energy is produced, i.e., where energy is injected into the cascade. It is defined in a statistical sense and has the advantage of being experimentally determinable through the integration of a velocity fluctuation correlation coefficient R_x, as follows:

$$l_0 = \int_0^\infty R_x(r)\, dr \qquad \text{with} \qquad R_x(r) = \frac{\overline{v'_x(0)\, v'_x(r)}}{v'_{x,\mathrm{rms}}(0)\, v'_{x,\mathrm{rms}}(r)} \tag{4.4}$$

The Kolmogorov microscale l_k is associated with the smallest eddies responsible for dissipating kinetic energy and cannot be measured directly. However, a theoretical relationship was derived by Kolmogorov using dimensional analysis and several simplifying assumptions, allowing an estimate of this scale:

$$l_k = \frac{l_0}{\mathrm{Re}_{l_0}^{3/4}} \qquad \text{with} \qquad \mathrm{Re}_{l_0} = \frac{v'_{\mathrm{rms}}\, l_0}{\nu} \tag{4.5}$$

These two scales define the boundaries of a range of scales known as the inertial subrange. Within this subrange, each scale is characterized by its own length, where transport processes occur and eddies are formed from larger structures. The turbulent kinetic energy is transferred from large to smaller scales via the nonlinear terms of the Navier–Stokes equations. However, energy transfer is not strictly unidirectional: the reverse process, known as *backscatter*, has also been observed, in which energy is transferred from smaller to larger scales.

4.2 Numerical methods for flow field modeling

In particularly complex fluid dynamic scenarios, such as those typical of gas turbine operation, where turbulence plays a fundamental role in governing the flow dynamics, experimental analyses would require highly expensive and complex setups to provide comprehensive insight into the actual phenomena occurring inside the machine. In such cases, it is now common practice to rely on Computational Fluid Dynamics (CFD), that is, a numerical analysis approach for simulating fluid flow fields and associated phenomena, aimed at iteratively solving the Navier–Stokes equations using various mathematical models and approaches. The development of CFD has been driven by the evolution of computational power, and thanks to its widespread accessibility, it has become prevalent across many industrial sectors.

As a numerical method, each CFD approach requires the preliminary generation of a computational grid that discretizes the continuous fluid domain to be simulated, along with the application of appropriate boundary conditions and the definition of the flow properties. The level of mesh resolution is particularly important: an excessively coarse discretization can lead to significant errors due to severe approximations, whereas an overly refined mesh demands substantial computational resources, often without fully justifying the improvements obtained compared to coarser meshes or even experimental results.

The limitations of computational fluid dynamics lie in the ability of the employed mathematical models to effectively represent the phenomena being analyzed. Therefore, validation of the numerical results through experimental analyses remains necessary. Over the years, various approaches have been developed, including the RANS method, which is less accurate but computationally cheaper; the LES method, which offers an intermediate level of accuracy and computational cost; and the DNS method, which is the most accurate but also the most demanding in terms of computational resources.

The Reynolds Averaged Navier-Stokes (RANS) approach has enjoyed and continues to enjoy widespread popularity in turbomachinery applications. Its foundation lies in Reynolds' decomposition theory [83], which separates the flow into two distinct components: a mean component and a fluctuating component. To better understand this principle, one can consider a generic instantaneous flow quantity:

$$\phi(x,t) = \overline{\phi}(x) + \phi'(x,t) \tag{4.6}$$

where $\phi(x,t)$ is the instantaneous quantity, $\overline{\phi}(x)$ is the time-averaged component (independent of time), and $\phi'(x,t)$ is the fluctuating component. Considering a time interval T sufficiently long compared to the typical scales of the fluctuations, the mean value can be written as:

$$\overline{\phi}(x) = \lim_{T\to\infty} \frac{1}{T} \int_t^{t+T} \phi(x,t)\, dt \tag{4.7}$$

Time averaging over T represents a key strategy to reduce computation time and resources. However, this approach cannot resolve the small-scale turbulent motions that occur at scales much smaller than the mean flow. Consequently, turbulence modeling becomes necessary. Applying Reynolds decomposition to the velocity and pressure vectors in the continuity and momentum equations yields the Reynolds-averaged Navier–Stokes equations:

$$\nabla \cdot \overline{\mathbf{u}} = 0 \tag{4.8}$$

$$\frac{\partial \overline{\mathbf{u}}}{\partial t} + \nabla \cdot (\overline{\mathbf{u}}\,\overline{\mathbf{u}}) = -\frac{1}{\rho}\nabla \overline{p} + \nu \nabla^2 \overline{\mathbf{u}} + \frac{1}{\rho}\nabla \cdot \boldsymbol{\tau}^R \tag{4.9}$$

where $\boldsymbol{\tau}^R$ is the Reynolds stress tensor:

$$\boldsymbol{\tau}^R = -\begin{pmatrix} \rho\,\overline{u'^2} & \rho\,\overline{u'v'} & \rho\,\overline{u'w'} \\ \rho\,\overline{v'u'} & \rho\,\overline{v'^2} & \rho\,\overline{v'w'} \\ \rho\,\overline{w'u'} & \rho\,\overline{w'v'} & \rho\,\overline{w'^2} \end{pmatrix} \tag{4.10}$$

To account for turbulence effects, this tensor is not computed directly but modeled through closure models. The most widespread approach is based on the Boussinesq hypothesis, which assumes linearity between the Reynolds stress tensor and the strain rate tensor, introducing a turbulence viscosity that depends on the flow rather than the fluid. In general, closure models rely on turbulence quantities such as turbulent kinetic energy k, its dissipation rate ε, or the vortex frequency ω.

RANS simulations are suited for steady or slowly varying flows, providing a statistical solution. Extensions such as URANS (Unsteady RANS) have been developed to better represent unsteady flows, in which the time derivative in the Navier–Stokes equations is retained.

Direct Numerical Simulation (DNS) is the most accurate approach, as it resolves all spatial and temporal scales directly, solving the Navier–Stokes equations without turbulence models or approximations. Conceptually, this is the simplest method, but also the most computationally expensive. The computational cost of DNS depends strongly on the Reynolds number [84], often requiring computing power beyond that of current supercomputers. As such, DNS is primarily used as a research tool to understand turbulence and develop improved models.

Large Eddy Simulation (LES) offers a compromise between DNS and RANS. In LES, the Navier–Stokes equations are filtered to resolve only the large energy-containing eddies, while the smaller, sub-grid scale motions are modeled. According to Kolmogorov's theory, the small scales follow universal laws, while the large scales depend on geometry, making them worth resolving.

Filtering can be performed through various techniques. A common one is a temporal moving average via convolution of the signal ϕ with a window function $h(\tau)$:

$$\check{\phi}(t) = \int_0^T \phi(t - \tau)\, h(\tau)\, d\tau \tag{4.11}$$

To preserve the energy content of the signal, $h(\tau)$ must be a probability density function. Filtering intensity can be tuned by changing T: increasing T approximates a RANS solution, while decreasing T moves toward DNS. Spectral filtering can also be used, e.g., via Fourier transforms to isolate the frequencies corresponding to scales to be modeled.

In analogy with Reynolds decomposition but retaining time dependence, a generic quantity is written as:

$$\phi(x, t) = \check{\phi}(x, t) + \phi'_R(x, t) \tag{4.12}$$

where $\check{\phi}(x, t)$ is the filtered component and $\phi'_R(x, t)$ the sub-grid scale residual. From Eq. (4.12), along with Eqs. (4.8)–(4.9), one can derive the filtered Navier–Stokes equations, where the filtered terms are solved numerically and the residuals are modeled.

In the field of turbomachinery, the low fidelity of RANS simulations (which have traditionally formed the basis for turbine design), combined with the increasing availability of computational resources compared to the past, is progressively encouraging the adoption of LES approaches. These latter methods have been shown to provide higher accuracy and enable a more in-depth investigation of unsteady phenomena that have so far been neglected or approximated using empirical correlations. Moreover, the computational cost is now sufficiently manageable to allow for parametric studies as well.

4.3 Global and local approaches

In order to assess the efficiency of an open system through which a fluid evolves, it is necessary to evaluate the mechanical energy losses experienced by the fluid as it flows through the system. In the context of turbomachinery, two main approaches are generally adopted for the quantification of energy losses:

- **Local approach:** This method relies on the evaluation of local quantities such as viscous dissipation and turbulent dissipation. When integrated over the computational domain, these quantities can provide an estimate of the total mechanical energy losses. This approach is particularly suited for numerical simulations based on high-fidelity methods, such as Large Eddy Simulations or Direct Numerical Simulations, where full spatial fields of velocity and pressure are available.
- **Global approach:** This method evaluates the total pressure drop between two measurement planes located upstream and downstream of the system. The mechanical energy loss is assumed to be proportional to the difference

in total pressure across these planes. This approach is typically employed in experimental settings, where local flow field data is not accessible.

The local approach allows for a detailed analysis of the loss mechanisms within the domain, offering insights that are inaccessible to global methods. In particular, losses can be evaluated by integrating and time-averaging the instantaneous total pressure derivative throughout the domain:

$$\frac{\int_D \overline{\frac{Dp_t}{Dt}}\, d\mathcal{V}}{\dot{\mathcal{V}} \cdot p_{dyn,ref}} = \text{Loss coefficient} \tag{4.13}$$

where $\mathcal{V}$ denotes the domain volume, $\dot{\mathcal{V}}$ is the volumetric flow rate across the domain, and $p_{dyn,ref}$ is a reference dynamic pressure. This formulation corresponds to what is typically measured in experimental investigations through total pressure loss coefficients, under certain assumptions. The mathematical steps that demonstrate the equivalence between the local and global approaches are reported in Appendix A, where the derivation is carried out in detail to ensure full consistency between the two formulations. From this point onward, the analysis is carried out within an incompressible-flow framework. While a compressible formulation is required when addressing the thermodynamic cycle and turbine expansion, the focus here is restricted to the flow within a single low-pressure turbine stage, where density variations are negligible. Under the low-Mach-number assumption, the flow can therefore be considered incompressible, allowing losses to be evaluated based solely on mass and momentum conservation, without explicitly invoking the energy equation or thermodynamic relations.

By applying Reynolds decomposition to the velocity and pressure fields and performing a temporal average of the material derivative of the total pressure p_t, one obtains:

$$\overline{\frac{Dp_t}{Dt}} = \underbrace{\overline{u}_i \frac{\partial \overline{p}}{\partial x_i} + \frac{D}{Dt}\left(\frac{1}{2}\overline{\rho}\,\overline{u}_i\overline{u}_i\right) + \frac{D}{Dt}\left(\frac{1}{2}\overline{\rho}\,\overline{u_i'u_i'}\right)}_{\frac{D\overline{p_t}}{Dt}} + \frac{\partial}{\partial x_j}\left[\overline{u}_i\left(\overline{\rho u_i'u_j'}\right)\right] + \frac{\partial \overline{\rho u_i' k}}{\partial x_i} + \frac{\partial \overline{u_i' p'}}{\partial x_i} \tag{4.14}$$

The first group of terms, highlighted by the brace, represents the time-averaged total pressure change that can be captured experimentally through total pressure probes. The remaining terms, which account for turbulent and viscous diffusion, cannot be directly measured and are typically neglected in experimental assessments. Therefore, experimental methods provide an accurate estimation of losses only when these diffusive contributions are negligible [85, 86].

The material derivative of the mean total pressure can be rewritten as the sum of two transport equations: one governing a pseudo-total pressure (mean pressure and mean kinetic energy), and another for the turbulent kinetic energy k. Specifically:

$$\underbrace{\overline{u}_i \cdot \nabla \left(\overline{p} + \frac{1}{2}\rho \overline{V}^2\right)}_{F.A.S.P.} = -\underbrace{\overline{\tau}_{ij}\frac{\partial \overline{u}_i}{\partial x_j}}_{\Phi_m} - \underbrace{\left(-\rho\overline{u'_i u'_j}\frac{\partial \overline{u}_i}{\partial x_j}\right)}_{P_k} - \underbrace{\frac{\partial \overline{\tau}_{ij}\overline{u}_i}{\partial x_j}}_{V.D.} - \underbrace{\frac{\partial}{\partial x_j}\left[\overline{u}_i\left(\rho\overline{u'_i u'_j}\right)\right]}_{T.D.} \tag{4.15}$$

$$\underbrace{\overline{u}_i \cdot \nabla \overline{k}}_{F.T.K.E.} = \underbrace{\nu\frac{\partial^2 k}{\partial x_j^2}}_{V.D._k} + \underbrace{\left(-\overline{u'_i u'_j}\frac{\partial \overline{u}_i}{\partial x_j}\right)}_{P_k} - \underbrace{\frac{\partial \overline{u'_i k}}{\partial x_i}}_{T.D._k} - \underbrace{\frac{1}{\rho}\frac{\partial \overline{u'_i p'}}{\partial x_i}}_{P.D._k} - \underbrace{\nu\overline{\frac{\partial u'_i}{\partial x_j}\frac{\partial u'_i}{\partial x_j}}}_{\overline{\Phi}_k} \tag{4.16}$$

By summing Eqs. (4.15) and (4.16), and assuming an adiabatic flow with no external work exchange, the simplified expression for the time-averaged material derivative of total pressure becomes:

$$\overline{\frac{Dp_t}{Dt}} = -\overline{\tau}_{ij}\frac{\partial \overline{u}_i}{\partial x_j} - \mu\overline{\frac{\partial u'_i}{\partial x_j}\frac{\partial u'_i}{\partial x_j}} \tag{4.17}$$

This final expression highlights that the dominant loss mechanisms are:

- Viscous dissipation due to the mean flow;
- Viscous dissipation due to turbulent fluctuations.

The problem of loss estimation reduces to the accurate evaluation of these two contributions. Mean viscous dissipation depends on the mean flow, so the numerical simulation must be able to accurately capture the mean velocity profile, especially within the boundary layer, where steep gradients make this term particularly relevant. On the other hand, turbulent dissipation depends on the fluctuating field and, to date, only a DNS simulation can provide a complete representation of it. In a LES simulation, the dissipation contribution of the smallest scales is lost, as these scales are not resolved. In a RANS simulation, since only the mean velocity field is available, this term cannot be directly computed. Experimentally, unless strong assumptions are made, it cannot be measured either.

The standard approach when relying on a RANS simulation is to estimate turbulent dissipation through the turbulent kinetic energy production term P_k. By analyzing the transport equation for turbulent kinetic energy (Eq. (4.16)), if turbulence is in equilibrium (the left-hand side of the equation vanishes) and diffusive terms are neglected, the two terms become equal and the turbulent dissipation term can be approximated by the turbulent kinetic energy production term P_k. This provides a convenient way to estimate losses from turbulent effects without explicitly computing all terms of the k-equation.

However, when turbulence is not in equilibrium and DNS data are not available (i.e., a direct evaluation of the turbulent dissipation term is not possible), the correct practical estimation of the substantial derivative of total pressure must rely on an extended formulation that accounts for the imbalance between the production and dissipation of turbulent kinetic energy. As discussed in the following sections, this formulation provides a more accurate representation of the mean and turbulent dissipation mechanisms under non-equilibrium conditions:

$$\frac{\overline{Dp_t}}{Dt} = -\overline{\tau}_{ij}\frac{\partial \overline{u}_i}{\partial x_j} - \left(-\rho\overline{u_i'u_j'}\frac{\partial \overline{u}_i}{\partial x_j}\right) + \rho\overline{u}_i \cdot \nabla \overline{k} \tag{4.18}$$

Equation (4.18) is the result of neglecting the diffusion terms $V.D._k$, $T.D._k$, and $P.D._k$ in Eq. (4.16).

Chapter 5

Modal Decomposition Techniques

5.1 Introduction

In physics, mathematics, and engineering, a topic of great relevance concerns the transformation of physical equations into coordinate systems that simplify their resolution. One of the most important contributions in this field is attributed to J.B. Joseph Fourier, who introduced the concept that sinusoidal functions of increasing frequency form an orthogonal basis for the solution space of the problem at hand. Fourier's work is today a cornerstone of computational mathematics, large-scale numerical physics and engineering, real-time image and audio compression, and advanced data analysis.

The potential of the Fourier Transform forms the foundation of modern modal decomposition techniques. Specifically, from a data analysis perspective, the Fourier Transform in the form of the *Discrete Fourier Transform* (DFT) can be used to extract information about the characteristic frequencies of a flow field. Typically, when performing DNS or LES simulations on a computational domain discretized into M spatial points over N time steps, the resulting dataset consists of a matrix (hereinafter referred to as the *snapshot matrix*) of dimensions $M \times N$ for each extracted variable (velocity components, pressure, temperature, etc.). By multiplying the *snapshot matrix* by the Fourier coefficient matrix, which has dimensions $N \times N$, one obtains a projection of the full flow field onto a basis of sinusoidal functions. This projection yields a matrix of Fourier modes, with the same dimensions as the *snapshot matrix*, but containing frequency-domain information that is physically insightful.

The operation is carried out as schematically shown in Figure 5.1, where U denotes the data matrix, while Ψ indicates the DFT matrix. The latter is essentially a Vandermonde matrix with complex entries, where:

$$w_N^{(j,k)} = e^{-2\frac{\pi i}{N}jk}$$

The matrix F, i.e., the output of the operation in Figure 5.1, is the aforementioned

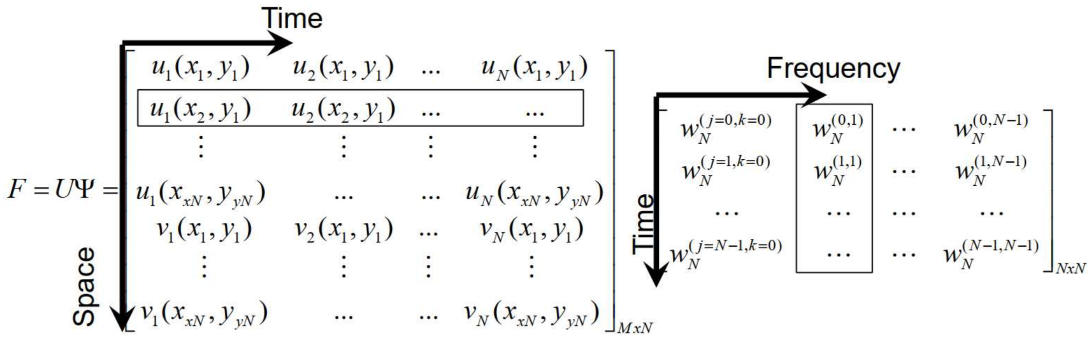

Figure 5.1: Discrete Fourier Transform.

matrix of Fourier modes, which, when considered as individual snapshots, can be interpreted as specific flow structures recurring over time with a characteristic frequency.

Although the Discrete Fourier Transform offers a powerful framework for identifying the dominant frequencies in a flow field, its basis functions are fixed a priori and independent of the actual data. This means that the sinusoidal modes may not optimally represent the energetic structures of the specific flow under investigation. In particular, for flows characterized by broadband or localized features, such a basis can lead to inefficient or physically unrepresentative decompositions. To overcome this limitation, data-driven techniques such as Proper Orthogonal Decomposition (POD) have been developed, aiming to extract an optimal set of modes directly from the dataset based on energy content.

5.2 Proper Orthogonal Decomposition

The study of turbulence poses two major challenges. The first is mathematical in nature, due to the intrinsic non-linearity of the governing differential equations. The second, more significant from a physical standpoint, lies in the fact that turbulence involves the interaction of a vast number of degrees of freedom across a wide range of spatial and temporal scales. Progress in understanding turbulent flows, therefore, requires some form of complexity reduction, whereby non-essential degrees of freedom are discarded.

The first substantial evidence that such a reduction could be achieved in fully developed turbulent flows came from the experimental observation of coherent structures, as documented by J. T. C. Liu [87]. Coherent structures are spatially organized features that appear repeatedly and undergo a characteristic temporal evolution. Proper Orthogonal Decomposition has emerged as a suitable method for extracting such features. In particular, POD is a linear technique that identifies an optimal modal basis from a set of signals or flow realizations. Its appeal lies in its mathematical simplicity and the fact that it does not require any assumption on the linearity of the underlying physical problem.

The POD technique features the following key characteristics:

- It has a statistical foundation, as it extracts information directly from experimental or numerical datasets.
- Its analytical basis provides clear insights into its capabilities and limitations.
- It enables the identification of spatial and temporal structures that are deemed essential in turbulent flows.

POD was introduced into the context of turbulence by Lumley in 1967 [28], although it has since found applications in many other fields. According to Lumley, POD adopts an energetic perspective to identify coherent structures, which are defined as those associated with the largest fraction of the flow's kinetic energy. By selecting an appropriate basis, the projections of the velocity field onto the modes are maximized for these dominant structures. The resulting basis consists of orthonormal modes, obtained by solving an eigenvalue problem associated with the spatial correlation matrix between two points in the domain.

This classical formulation of POD, based on spatial cross-correlation at a fixed time, is often referred to as Lumley's POD. An alternative formulation, proposed by Sirovich [88] and commonly known as *Snapshot POD*, instead relies on a temporal cross-correlation at a fixed spatial location. Both approaches are mathematically equivalent and lead to the same modal basis. However, the Snapshot POD is typically preferred in the context of numerical simulations, due to its computational efficiency when dealing with large datasets.

When applied to the fluctuating velocity components, the Snapshot POD generates two complete orthogonal bases, namely the POD modes (φ^k) and the related eigenvectors (χ^k), which retain the mode dynamics. The starting point is the snapshot matrix $\mathbf{U} \in \mathbb{R}^{M\times N}$, where each column represents a flattened three-dimensional velocity field (e.g., from a DNS or LES simulation), and M is the number of spatial degrees of freedom, while N is the number of time snapshots. The method seeks to extract the most energetic spatial modes by solving an eigenvalue problem derived from the temporal correlation matrix $\mathbf{C}$, defined as:

$$\mathbf{C} = \mathbf{U}^\top \mathbf{U} \in \mathbb{R}^{N\times N}$$

The eigenvalue problem to solve is:

$$\mathbf{C}\chi^k = \lambda_k \chi^k$$

where λ_k and χ^k are the eigenvalues and corresponding eigenvectors, respectively. The eigenvectors χ^k represent the temporal coefficients of the modes, while the spatial modes φ^k are obtained as linear combinations of the original snapshots:

$$\boldsymbol{\varphi}^k = \mathbf{U}\chi^k$$

The set $\{\boldsymbol{\varphi}^k\}$ forms an orthonormal basis that optimally captures the kinetic energy content of the original dataset, ordered by decreasing λ_k. In other words, the POD modes are ranked based on the turbulent kinetic energy associated with each of them.

POD (and other modal decomposition techniques) is traditionally employed to identify and rank coherent structures based on their energetic content, and its use in the literature remains largely limited to flow characterization within the chosen modal basis. In the present work, the novelty does not lie in the implementation of the decomposition itself, but in its rigorous projection onto the governing balance equations, enabling a direct quantification of the contribution of each mode to loss generation mechanisms. To the best of the author's knowledge, such a direct linkage between modal decomposition and loss production has not been systematically established in previous studies.

Specifically, the bi-orthogonality of the POD modes and eigenvectors allows the Reynolds stresses to be expressed in terms of modal contributions. The fluctuating velocity field is decomposed as:

$$\mathbf{u'}(x, y, z, t_i) = \sum_k \chi^k(t_i)\varphi^k(x, y, z) \tag{5.1}$$

The Reynolds stress tensor is obtained by time-averaging the product of velocity fluctuations:

$$\overline{u_i'u_j'} = \overline{\sum_k \chi^k\varphi_i^k \ \sum_l \chi^l\varphi_j^l} \tag{5.2}$$

Expanding the product:

$$\overline{u_i'u_j'} = \sum_k \sum_l \varphi_i^k\varphi_j^l, \overline{\chi^k\chi^l} \tag{5.3}$$

Exploiting the orthogonality of the temporal coefficients:

$$\overline{u_i'u_j'} = \sum_k \varphi_i^k(x, y, z)\varphi_j^k(x, y, z) \tag{5.4}$$

Finally, the TKE production term can be written as:

$$P_k = -\sum_k \varphi_i^k(x, y, z)\varphi_j^k(x, y, z)\frac{\partial \overline{u}_i}{\partial x_j} \tag{5.5}$$

Observing equation (5.4), it is evident that for $i = j$, we obtain the trace of the tensor $\overline{u_i'u_j'}$, which is twice the total turbulent kinetic energy corresponding to the sum of the TKE associated with each mode φ^k. Equation (5.5) shows how each spatial POD mode $\varphi^k(x, y, z)$ contributes to the turbulent kinetic energy production in the domain of interest. A similar decomposition can be applied to the turbulent kinetic energy substantial derivative term of equation (4.16), i.e. the divergence operator can be applied to the partial TKE associated with a specific POD mode family, as for contraction of eq. (5.4) for $i = j$. The ability to identify and separate the various coherent structures within the analyzed flow makes the POD a powerful tool for detecting loss-generating mechanisms in a LES or DNS computational domain.

While POD is highly effective in identifying the most energetic coherent structures within a turbulent flow, it does not provide direct information about their frequency content or the precise timing of their occurrence. This limitation arises from the purely energetic nature of the decomposition, which ranks modes by their average contribution to the total kinetic energy, but discards any spectral or phase-related information. To address these aspects, the POD framework can be extended or reformulated in various ways, depending on the specific objectives of the analysis. Among these, the Spectral POD and the Phase-locked POD techniques, detailed in the following sections, represent two notable developments aimed at enriching the physical interpretability of the modes by incorporating frequency and temporal resolution, respectively.

5.3 Spectral POD

While the classical POD decomposes a dataset into spatial modes and corresponding temporal coefficients based on energy content, it does not explicitly provide information on the frequency content of the flow structures. In fact, each POD mode generally contains a mixture of frequencies, rather than being associated with a single, well-defined spectral component. This limitation can be particularly restrictive when analyzing unsteady or multi-scale phenomena, such as those occurring in turbomachinery flows, where coherent structures may evolve on distinct temporal scales.

The *Spectral Proper Orthogonal Decomposition* (SPOD) addresses this limitation by extending the POD framework to the frequency domain, combining spatial optimality with a rigorous spectral characterization of the flow structures [89]. In SPOD, the flow is interpreted as a realization of a stationary stochastic process, and the decomposition is formulated in terms of the cross-spectral density (CSD) of the fluctuating field.

From a practical standpoint, the velocity fluctuation field $\mathbf{u}'(\mathbf{x}, t)$ is first divided into a set of temporal blocks. For each block, a temporal Fourier transform is applied, yielding the frequency-domain representation $\widetilde{\mathbf{u}'}(\mathbf{x}, \omega)$. The cross-spectral density tensor is then defined as:

$$\mathbf{S}(\mathbf{x}, \mathbf{x}', \omega) = \mathbb{E}\left[\widetilde{\mathbf{u}'}(\mathbf{x}, \omega)\widetilde{\mathbf{u}'}^{*}(\mathbf{x}', \omega)\right], \tag{5.6}$$

where $\mathbb{E}[\cdot]$ denotes ensemble averaging and $(\cdot)^*$ the complex conjugate. In practice, the ensemble average is approximated by averaging over the available temporal blocks.

The SPOD modes are obtained by solving, at each angular frequency ω, the eigenvalue problem associated with the CSD operator:

$$\int \mathbf{S}(\mathbf{x}, \mathbf{x}', \omega)\, \boldsymbol{\varphi}_\omega^k(\mathbf{x}')\, \mathrm{d}\mathbf{x}' = \lambda_\omega^k\, \boldsymbol{\varphi}_\omega^k(\mathbf{x}), \tag{5.7}$$

where $\boldsymbol{\varphi}_\omega^k$ are the frequency-dependent SPOD modes and λ_ω^k are the corresponding eigenvalues, representing the modal energy content at frequency ω. For each frequency, the modes are optimally ranked in terms of energy, in direct analogy with classical POD.

The fluctuating velocity field in the frequency domain can therefore be expressed as:

$$\widetilde{\mathbf{u}'}(x, y, z, \omega) = \sum_k \chi_\omega^k(\omega) \varphi_\omega^k(x, y, z), \tag{5.8}$$

where χ_ω^k are the SPOD expansion coefficients, analogous to the temporal coefficients of POD but defined in the spectral space.

In this work, following the approach adopted in [90], the velocity fluctuation field is partitioned into non-overlapping temporal blocks. Unlike the more common practice of using overlapping segments, this choice ensures a complete and strictly orthonormal decomposition in both the frequency domain and the associated SPOD modal space.

By projecting the TKE production onto SPOD modes, one can isolate the contribution of structures associated with specific spectral content. In analogy with the POD-based analysis, the SPOD modal decomposition allows for a quantitative assessment of the turbulent dissipation and thus of the total loss associated with each frequency-resolved structure. This frequency-based perspective is particularly useful for flows exhibiting tonal or periodic features, where distinct spectral signatures play a key role in loss generation. SPOD has proven particularly valuable in the analysis of high-fidelity simulations of turbomachinery components, such as LPT cascades, where unsteady phenomena and multi-scale interactions strongly influence overall losses. In these cases, the frequency-resolved nature of SPOD makes it capable of selectively targeting turbulent structures sharing the same frequency content and, therefore, possibly triggered by the same flow mechanism.

5.4 Phased-locked POD

As previously mentioned, POD allows for a statistical representation of the turbulence-related loss sources and their spatial locations, but does not provide information on the temporal sequence of turbulent events. The *Phase-Locked POD* (PPOD) method extends the capabilities of classical POD by enabling the decomposition of the flow field with respect to specific phases of a periodic or quasi-periodic event, such as the wake-passing cycle in a low-pressure turbine passage.

The procedure starts by applying the snapshots POD method [88] to the fluctuating velocity components. This produces a set of spatial POD modes $\boldsymbol{\varphi}^k(\mathbf{x})$ and associated temporal coefficients $\chi^k(t)$, which retain the dynamic information of each mode. For flow fields dominated by a strong periodic component, such as those in the analyzed LPT cascades, the first two POD modes typically capture the evolution of the periodic portion of the flow. The corresponding temporal eigenvectors exhibit a periodic behavior and, therefore, they can be used to sort the snapshots based on their relative position within one cycle. Further details are provided by Lengani et al. [91], who applied a similar procedure to PIV data. In this work, the first two modes obtained from the application of POD to every LPT cascade dataset are representative of the wake evolution in the inter-blade channel. The CFD-based results allowed the exploitation of the $\pi/2$ phase shift between the first two eigenvectors, enabling a robust temporal assignment based on corresponding angular phases. In particular, the angular phase θ_i of a snapshot at time t_i can be obtained as

$$\theta_i = \arctan\left(\frac{\chi^1(t_i)}{\chi^2(t_i)}\right) \tag{5.9}$$

Snapshots are then grouped into N_{bin} angular bins, each corresponding to a sector of the wake cycle. By doing so, phase-locked ensembles of snapshots are defined, resulting in N_{bin} subsets per full wake cycle for a single dataset. This phase-conditioned formulation is consistent with the concept of conditional modal decomposition for periodically forced flows, as discussed by Legrand et al. [92], who demonstrated that phase-resolved statistical analyses provide improved insight into the coupling between coherent unsteadiness and turbulence dynamics in turbomachinery configurations. A simplified diagram of the sorting procedure is provided in Figure 5.2.

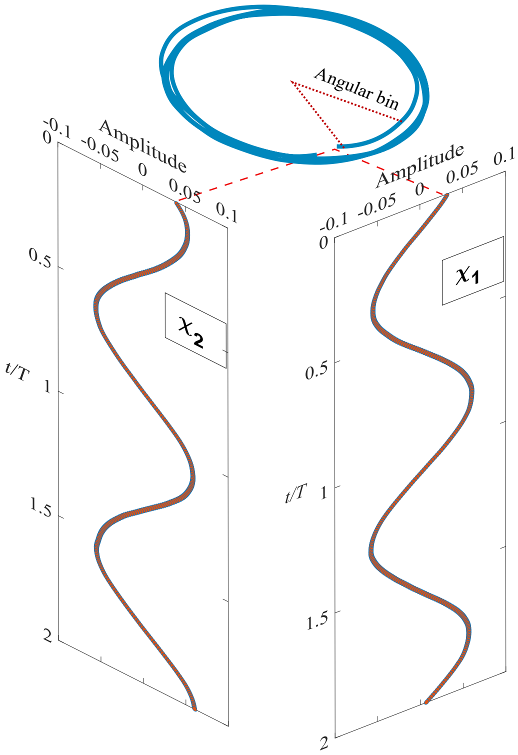

Figure 5.2: Simplified diagram representative of the snapshot sorting procedure, based on the $\pi/2$ phase shift between the first two POD eigenvectors

Once it is verified that the number of snapshots per period satisfies the convergence criteria reported in [93], the procedure concludes with the application of POD to each phase-locked ensemble of data, allowing for the identification, for each phase of the wake cycle, of the flow structures responsible for the higher quota of entropy generation and loss production:

$$\mathbf{u}'_\theta(x, y, z, \theta_i) = \sum_k \chi^k_\theta(\theta_i)\varphi^k_\theta(x, y, z) \tag{5.10}$$

5.5 Discrete Wavelet Transform

In more general cases, where it is crucial to gain insights into both the spectral and temporal organization of the loss mechanisms, it becomes advantageous to combine the properties of SPOD and PPOD within a unified framework capable of separating frequency content from temporal evolution. Wavelet transforms represent a powerful mathematical tool for decomposing signals into both time and frequency domains simultaneously. Unlike the Fourier transform, which provides a purely spectral description, wavelet-based approaches allow for a localized representation of the signal's frequency content, capturing both transient events and scale-dependent phenomena. This property makes wavelets particularly suitable for the analysis of turbulent flows, where coherent structures coexist across a wide range of scales and their temporal evolution plays a crucial role in determining entropy production.

The Discrete Wavelet Transform (DWT) performs a hierarchical decomposition of the signal through cascaded high-pass and low-pass filtering. This yields a set of detail and approximation coefficients that progressively isolate flow features associated with different characteristic time scales:

$$\mathbf{u}'(x, y, z, t_i) = \sum_j \sum_n c_{j,n}(x, y, z)\, \psi_{j,n}(t_i). \tag{5.11}$$

In Eq.(5.11), $c_{j,n}(x, y, z)$ denotes the wavelet coefficient associated with scale j and translation n at the spatial location (x, y, z), while $\psi_{j,n}(t)$ represents the corresponding wavelet basis function.

A comparative overview of the time–frequency resolution associated with different analysis techniques is provided in Fig. 5.3. The left panel of the top row of the figure illustrates a raw time signal, which lacks any explicit spectral information. The right panel shows the Fourier transform, where frequency content is fully resolved but all temporal localization is lost. As shown in the bottom panel, the multi-resolution map associated with the DWT highlights the adaptive, scale-dependent tiling that characterizes wavelet-based decompositions, enabling the identification of transient or intermittent features that are not adequately captured by purely Fourier-based methods.

When an orthonormal mother wavelet is adopted, the DWT representation allows the production of turbulent kinetic energy P_k to be expressed in terms of wavelet modes exactly in the same way as shown for the POD decomposition in Eq. (5.5). In practice, it is possible to construct a set of DWT modes that are spatially localized and selectively capture the distinct spectral and temporal scales of loss mechanisms.

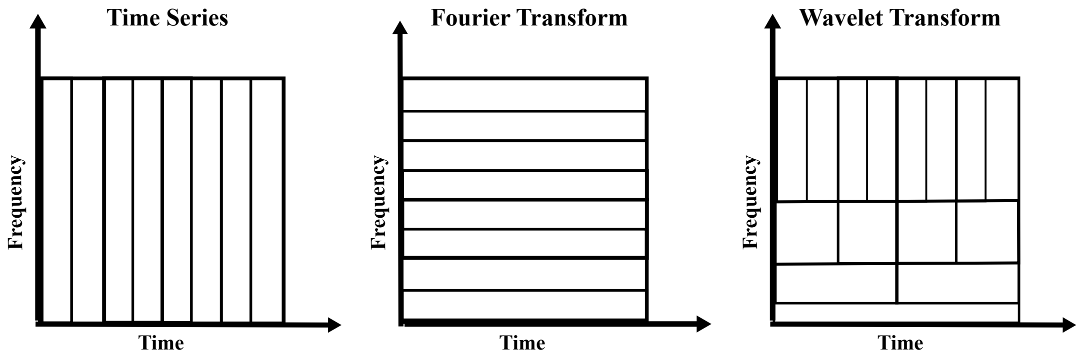

Figure 5.3: Comparison of time–frequency localization properties across different signal representations.

Chapter 6

LPT High-Fidelity Database

6.1 Large Eddy Simulations

The datasets used in this study were obtained from a series of Large Eddy Simulations carried out on multiple Low-Pressure Turbine profiles investigated throughout the PhD activity. Simulations were performed by the University of Florence's spin-off *Morfo Design SRL* using the commercial FV software STAR-CCM+ and employed a segregated flow model with bounded-central differencing. A 2^{nd} order implicit backward differentiation scheme was adopted for the time derivative, while the WALE model [94] was selected as the subgrid-scale model, in combination with the STAR-CCM+ Synthetic Eddy Model (SEM) [95] for turbulence generation at the domain inlet. Given the low Mach numbers of the investigated operating conditions, simulations were conducted assuming an incompressible flow regime.

To evaluate the aerodynamic performance of each profile under conditions representative of real LPT environments, stage-like configurations were developed. As shown in Fig. 6.1, each computational domain consists of a 2D repeating stage with a degree of reaction equal to 0.5. The profiles, mirrored between the two adjacent blade rows, operate under equivalent relative conditions, creating what is referred to as "natural" inflow conditions. Specifically, each original profile is duplicated and mirrored to form the stage configuration, enabling the investigation of key unsteady phenomena in LPTs, such as wake-blade interaction, which are critical for assessing aerodynamic performance.

The computational grids employed are 2D-extruded polyhedral meshes extending for about 25% of the axial chord in the span-wise direction, consisting of approximately 25 million cells. Meshes were designed to satisfy $y^+ < 1$ and $x^+ = z^+ \approx 20$ across the blade surface. Preliminary URANS simulations were used to optimize grid resolution, and the resulting URANS mean-flow fields were applied as initial conditions for the LES. The computational time step was selected to maintain a CFL number close to 1 in the region of interest, corresponding to approximately 500 samples per blade-passing period. Once the LES was initialized, a transient period equivalent to 4 flow-through times (FTTs) was simulated to ensure the development

Figure 6.1: Schematic of an exemplary stage-like simulation environment for unsteady inflow conditions

of the turbulent field, followed by snapshot collection and time-averaging of the flow solution. This allowed for independent and efficient post-processing of both instantaneous and mean flow quantities. It is worth noting that the numerical setup (mesh resolution, time step, etc.) adopted for all profiles was selected based on a sensitivity study conducted on comparable test cases [30].

For each configuration, instantaneous data were collected over approximately 8 blade-passing periods, saving one snapshot every ten time steps, resulting in a total of about 400 snapshots per simulation. Two exemplary time snapshots of the instantaneous flow field around the vane blade are reported in Fig. 6.2. The picture shows the contour plots of the fluctuating axial velocity u' (according to the Reynolds decomposition $u' = u(t_i) - \overline{u}$). The wakes generated by the upstream blade can be observed in the two snapshots at different phases (θ_1, θ_2) as blue patterns at the inlet of the domain (as marked up over the picture).

Time-averaged quantities were monitored to verify statistical convergence, and post-processing operations were focused on evaluating aerodynamic performance metrics, which were subsequently compared with the results obtained through the POD-based modal analysis. The flow conditions of the tested profiles vary within narrow ranges typical of realistic LPT cascades under normal operating conditions, with the Reynolds number ranging between 7.5×10^4 and 1.1×10^5. All blade profiles share identical inlet and outlet flow angles, while differing in pitch-to-chord ratio, camberline geometry, and thickness distribution, thus enabling a systematic assessment of geometric effects under comparable loading conditions.

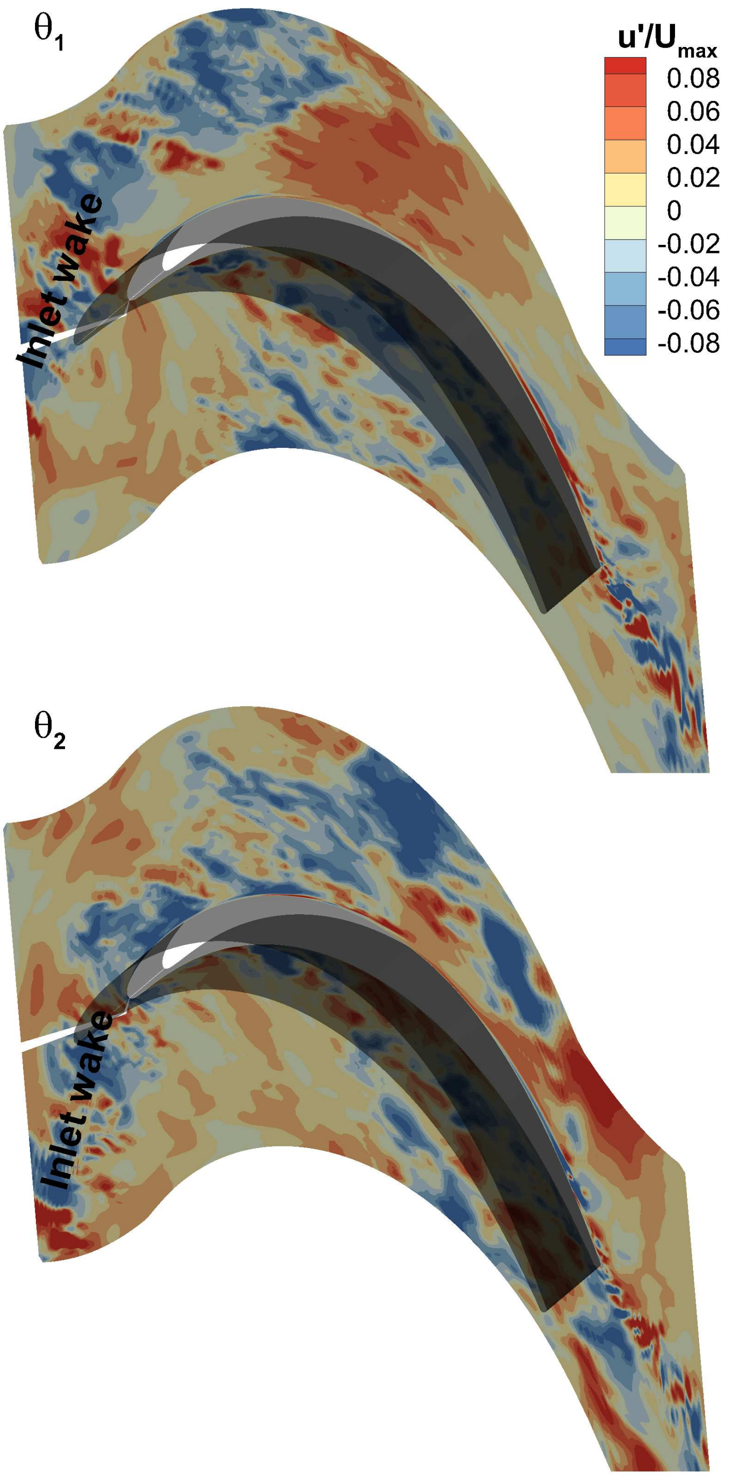

Figure 6.2: Instantaneous flow field realizations at two different positions of the inlet wake. Instantaneous fluctuating axial velocity.

6.2 Experimental validation

To assess the accuracy of the LES database, three out of the six profiles considered in this study were characterized in a dedicated experimental campaign. Specifically, for each geometry, the inlet-to-outlet loss coefficient predicted by the LES simulation was compared with the experimental loss coefficient obtained from a dedicated wind-tunnel campaign carried out at the Aerodynamics and Turbomachinery Laboratory of the University of Genoa.

The experimental facility consists of a blow-down wind tunnel which can be equipped with a large-scale linear cascade specifically designed for profile loss measurements. The planar cascade features seven constant-section blades mounted between two transparent side plates, with the geometry of the supports tailored to ensure both optical accessibility and accurate alignment of the instrumentation. Since a linear cascade cannot reproduce the circumferential periodicity of an actual turbomachine, side bleed openings and adjustable tailboards are installed at the inlet and outlet of the cascade. Their position is tuned to replicate the periodic flow conditions over at least the three central blades, thus ensuring that the central passage provides representative, quasi-periodic flow conditions.

Upstream of the cascade, a wake generator, schematized in Figure 6.3, reproduces the inflow unsteadiness associated with stator–rotor interaction. The system consists of cylindrical rods mounted on two moving belts driven by an electric motor. This configuration allows the operator to control both the rod pitch, its diameter, and the belt rotational speed, enabling accurate reproduction of the wake–blade passing frequency and the inflow turbulence characteristics for a wide range of reduced frequencies, Reynolds numbers, and flow coefficients. Although the unsteady inflow in the experiments is generated by cylindrical bars rather than by an actual upstream blade row (as in the LES environment), bar-generated wakes are widely recognized to provide a representative approximation of rotor wakes in terms of velocity deficit, turbulence intensity, and boundary layer interaction, provided that the relevant non-dimensional parameters are properly matched. In the present campaign, the wake generator is therefore tuned to reproduce these key flow characteristics, ensuring a meaningful comparison with the LES configuration.

As depicted in Figure 6.4, profile losses were evaluated by acquiring traverses of total pressure upstream and downstream of the cascade using a Kiel probe mounted on a two-axis automated traversing system. At each pitch-wise measurement location, the probe acquires ten thousand samples at 1 kHz, which are subsequently averaged to obtain a converged estimate of the local total pressure. This procedure is repeated over a sufficiently dense discretization of the pitch to accurately capture the wake deficit and the freestream regions. In addition to the dynamic and total-pressure measurements obtained from the downstream traversing probe, the acquisition system also records a downstream static-pressure signal and two upstream total-pressure signals. The last ones are measured by two fixed Kiel probes located at the top and bottom of the inlet section, respectively. These fixed Kiel probes serve a crucial diagnostic purpose: they allow continuous monitoring of the quality of the inlet flow, ensuring that no significant spatial non-uniformity (differences between top and bottom) or temporal drift of the overall facility affects the reference condition used for the loss evaluation.

Once the upstream and downstream traverses are acquired, the local total pressure

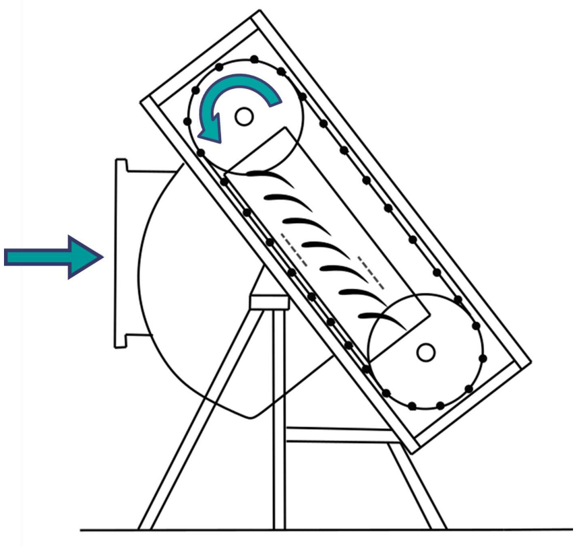

Figure 6.3: Simplified schematic of the upstream rotor-wake generation system for the linear cascade.

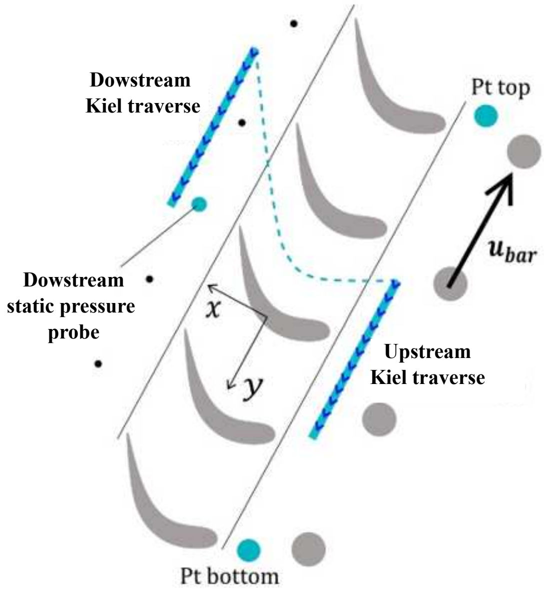

Figure 6.4: Schematic of the pressure signal acquisition system used for the experimental assessment of LPT profile performance.

loss coefficient is computed as:

$$C_{pt} = \frac{p_{t1}(y) - p_{t2}(y)}{p_{\text{dyn.ref.}}} \tag{6.1}$$

The integration of the loss coefficient along the pitch then provides the global inlet-to-outlet loss coefficient ω for the Re, f^+ and ϕ conditions under which the test was performed.

The experimental setup of the test section ensures an accuracy in evaluating the total pressure loss coefficient of $\pm 2\%$. Table 6.1 shows the comparison between

LES predictions and the reference measurements for the three experimentally tested profiles and reports, for all six considered blade geometries, the main qualitative characteristics of the blade loading, namely the location of the loading peak and the corresponding Zweifel number classification (low, medium, or high). The agreement between numerical and experimental loss coefficients confirms the reliability of the LES dataset for the subsequent analysis presented in this work.

Table 6.1: LES–experiment loss coefficient comparison and qualitative blade loading classification of the investigated profiles.

Profile	Loading type	Zweifel	LES-to-experiment l.c. ratio
1	Aft-loaded	High	0.952
2	Aft-loaded	Low	0.940
3	Mid-loaded	High	0.961
4	Mid-loaded	Medium	/
5	Mid-loaded	Low	/
6	Aft-loaded	Medium	/

6.3 Data interpolation on structured grid

Each dataset, as obtained from LES calculations, consists of snapshots of velocity and pressure data at every point of the LES unstructured mesh, saved within CSV files. These datasets were post-processed following the local loss evaluation approach described in Section 4.3. This technique is usually exploited in high-fidelity numerical simulations, where local evaluation is made possible by having access to the instantaneous spatial distribution of both the pressure and velocity fields. When integrated over the computational domain, as already discussed, loss-related quantities - particularly, viscous and turbulent dissipation - can also provide a global loss assessment. For all geometries, the post-processing domain was defined as the region surrounding the stator blade, extending from a $y-z$ plane located approximately 25% of the axial chord upstream of the leading edge to another positioned about 20% downstream of the trailing edge. An example of computational domain boundaries is depicted in Fig. 6.5, superimposed on the full LES unstructured grid.

To compute the derivatives of the velocity and pressure fields imported from the CSV files, in order to evaluate the terms of Eq. (4.15) and (4.16) depending on these quantities, data were interpolated onto a structured grid where the nodes are indexed following known criteria. The interpolation was necessary as the CSV files originated from a commercial, non-open-source software, meaning that the information regarding the indexing criterion of the cells in the unstructured grid was not directly available. Knowing the node indexing logic, it was possible to compute the derivatives using built-in functions available in MATLAB and/or Python environments that employ the finite difference method. To this end, a tool developed in the

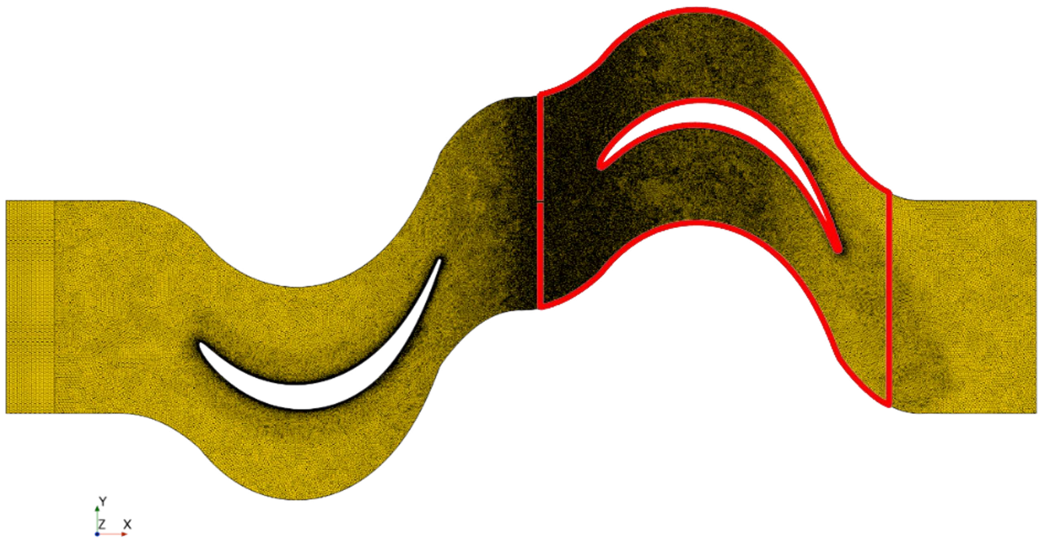

Figure 6.5: Post-processing domain boundaries (red line), superimposed to the corresponding LES computational mesh.

MATLAB environment, as described in [96], was employed to automatically generate structured meshes based on the blade geometry and domain boundaries provided as inputs. For all profiles, the blade geometry was assigned point-wise, as was the boundary of the computational domain. First, the domain boundary is divided by the code into four regions: *inlet*, *outlet*, *up*, and *down*. Figure 6.6 illustrates the external domain partitioning for an exemplary blade from the dataset. It's important to note that the *outlet* region extends into the final portion of the *down* region. This extension was implemented to gain better control over the structured mesh near the blade's trailing edge, thus reducing mesh distortion in this critical area.

In each region, the distribution of points is non-uniform, with a higher clustering near the boundaries and a lower clustering in the central part. This approach allows a greater concentration of nodes in the most critical areas, namely the leading and trailing edges, while nodes are more sparsely distributed in the central parts of the pressure side and suction side. These latter regions are less critical due to their relatively large curvature radius.

The structured mesh generation code requires the following input parameters:

- $nptX_i$: total number of points used to discretize the blade geometry
- $nptY_i$: number of points used to discretize the lines connecting the blade surface to the external domain
- $nptZ_i$: number of discretization points along the Z direction
- $nptTE$: number of points within $nptX_i$ used to discretize the trailing edge
- $nptLE$: number of points within $nptX_i$ used to discretize the leading edge

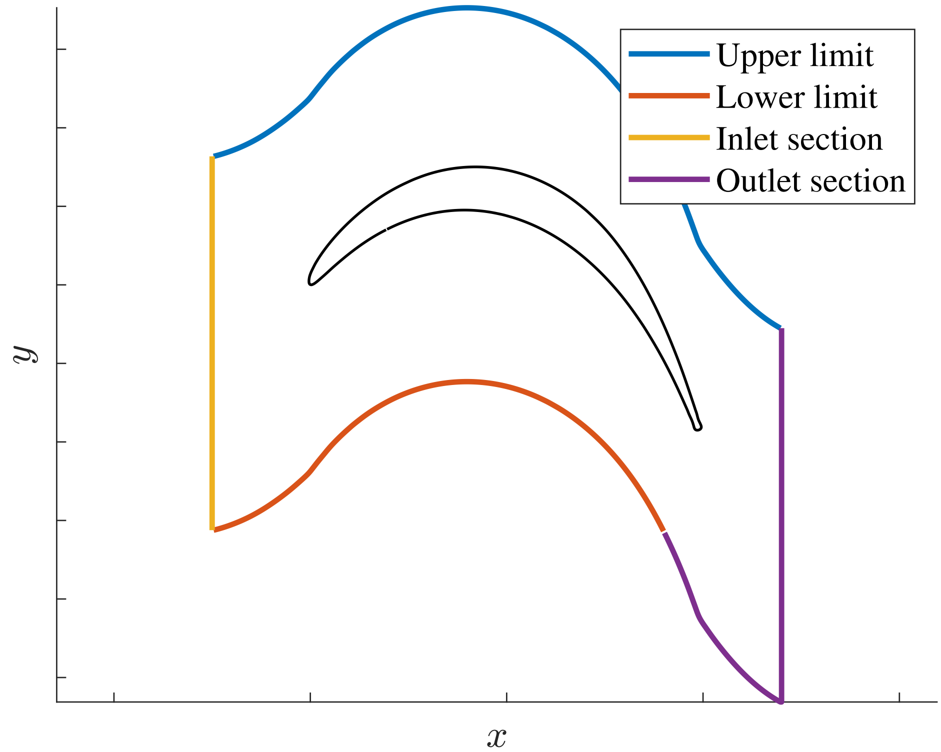

Figure 6.6: Regions of the post-processing domain's external boundary.

- $npti_{BL}$: number of points within $nptY_i$ used to refine the region near the blade wall, i.e., the boundary layer region
- k_{in}: coefficient controlling point clustering in the inlet region
- k_{out}: coefficient controlling point clustering in the outlet region

The generated mesh consists of two O-type blocks: an inner block adjacent to the blade wall and an outer block adjacent to the external domain. For instance, Fig. 6.7 shows the structured mesh for a blade geometry from the dataset. As can be observed, the interface line between the two O-type mesh blocks is defined as the locus of points equidistant from the blade surface. Specifically, it is constructed by projecting each point on the blade geometry along the surface normal vector at that point, at a distance corresponding to a specified fraction of the axial chord length. Consequently, in the inner block, the mesh lines connecting the blade surface to the interface between the blocks are all normal to the blade surface. Conversely, in the outer block, this constraint is not imposed, resulting in the mesh lines connecting the interface line to the outer domain being neither normal to the interface nor to the outer domain boundaries. In order to ensure a comparable spatial resolution, the structured mesh parameters were set individually for each geometry so that the final grid contained a number of points comparable to that of the unstructured mesh used in the LES simulations.

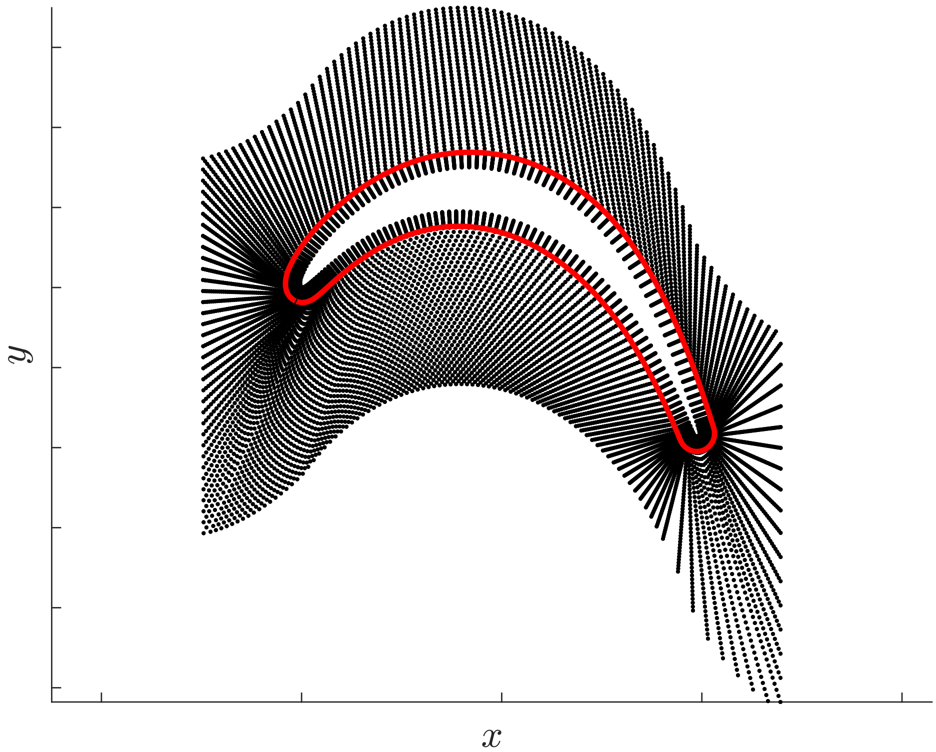

Figure 6.7: O-type structured mesh used as interpolation grid for the original LES data. Red line highlights the interface between the two grid blocks.

In the inner O-block, the discretization is carried out as follows: the blade is discretized into $nptX_i$ points (discretization with index i), from which normals are traced up to the separation line (discretization with index j). These normals are then uniformly divided into $npti_{BL}$ points. This parameter allows increasing the number of cells in the inner O-block, particularly in the proximity of the blade wall.

In the outer O-block, the discretization with index i is analogous to that of the inner block, whereas for the discretization with index j, the lines are uniformly divided into $(nptY_i - npti_{BL})$ points, in a different manner than in the inner region. Consequently, the discretization with index j of the entire mesh consists of $nptY_i$ points.

The discretization along the Z direction is uniform and extends from a value of $Z_{min}/C_x = 0$ to $Z_{max}/C_x = 0.25$, where C_x denotes the axial chord of the blade. Within this interval, $nptZ_i$ points are assigned.

Chapter 7

Post-Processing Routine

7.1 Workflow implementation

After generating the structured mesh, for each blade the following procedure was carried out: the unstructured mesh used for the LES simulations, the corresponding flow field data, and the structured mesh were provided as inputs to a dedicated MATLAB code. This code interpolates the flow field data from the unstructured mesh onto the structured one. For each point of the structured mesh, the five nearest neighbors of the unstructured mesh are considered, and the data at these points are combined using radial basis functions based on the inverse of the distances between the unstructured mesh points and the structured mesh point under consideration. In this way, the flow field data for the structured meshes are obtained, namely a matrix where the rows contain the velocity components at all the points of the domain and the columns correspond to the time instants.

The next step consists of performing the POD of the flow field in order to obtain the eigenvectors and the spatial modes, which are subsequently employed for the loss calculation, as described in Section 5.2. In the following, two implementations of the POD technique, developed in MATLAB and Python, are presented in detail.

The post-processing through POD was carried out using parallel computing on state-of-the-art high-performance computing (HPC) machines. In particular, resources were provided by CINECA in their *GALILEO100* and *LEONARDO* supercomputer infrastructures, and by the IT4Innovations National Supercomputing Center with their *Karolina* system.

GALILEO100 is a Tier-1 HPC and cloud system co-funded by the European ICEI project and engineered by Dell, operational since September 2021 at CINECA. The infrastructure comprises:

- Thin nodes: 422 nodes, each with 2 × Intel CascadeLake 8260 CPUs (24 cores each @ 2.4 GHz, 48 cores total), 384 GB RAM, and 480 GB SSD.
- Fat/data-processing nodes: 180 nodes, same CPUs and RAM plus 2 TB SSD and 3 TB Intel Optane persistent memory.

- GPU/visualization nodes: 34–36 nodes, with 2 × CascadeLake 8260 CPUs, 384 GB RAM, 2 × NVIDIA V100 GPUs, and 2 TB SSD; interconnected via 100 Gb/s Infiniband.
- OpenStack cloud servers: 77 machines, each with 2 × CascadeLake 8260 CPUs (48 cores), 768 GB RAM, connected by 100 Gb Ethernet.

The storage infrastructure includes 20 PB of active shared storage, 5 PB of fast HPC storage, 720 TB of fast DDN storage, and 1 PB of NVMe/SSD Ceph storage for the cloud partition. The system is interconnected via 100 Gb/s Infiniband and achieves a theoretical peak performance of 2 PFlop/s (3.53 TFlop/s per node) [97].

LEONARDO is a pre-exascale EuroHPC supercomputer based on a modular architecture combining general-purpose and accelerated nodes, and is currently one of the most powerful systems in Europe, ranking among the top positions of the TOP500 list. It comprises two main compute partitions:

- Booster Module: built on BullSequana X2135 “Da Vinci” single-node GPU blades, with 3 456 nodes, each equipped with one Intel Xeon Platinum 8358 CPU (32 cores @ 2.6 GHz), 512 GB DDR4-3200 MHz RAM, and four custom NVIDIA Ampere A100 GPUs (64 GB HBM2e, NVLink 3.0 @ 200 GB/s); interconnection is via dual-port NVIDIA HDR InfiniBand (aggregated 400 Gb/s).
- Data-Centric (DCGP) Module: composed of 1 536 nodes, each a three-node BullSequana X2140 blade containing two Intel Sapphire Rapids (Xeon Platinum 8480+) CPUs (total 112 cores @ 2.0 GHz), with 512 GB DDR5-4800 MHz RAM per node, plus 3 TB local NVMe SSD storage, and connected via single-port HDR100 InfiniBand (100 Gb/s).

The storage subsystem features a fast tier of approximately 5.4 PB NVMe/SSD (bandwidth 1.4 TB/s) and a capacity tier of about 106 PB HDD (bandwidth 620-700 GB/s). These tiers, along with login and visualization nodes, collectively support high-throughput I/O workloads [98, 99].

Karolina, operated by IT4Innovations, is a petascale EuroHPC system with 15.7 PFlop/s peak performance. Its compute resources include:

- 756 CPU-only nodes, each with 2× AMD EPYC 7H12 (64 cores each @ 2.6 GHz), and 256 GB RAM.
- 72 GPU-accelerated nodes, each with 2× AMD EPYC 7763 (64 cores each @ 2.45 GHz), 1 TB RAM, and 8× NVIDIA A100 GPUs (40 GB HBM2).
- One data-analytics node based on Intel Xeon-SC 8628 (24 cores @ 2.9 GHz) with 24 TB RAM.
- Two visualization nodes with AMD EPYC 7452 (32 cores @ 2.35 GHz), 256 GB RAM, and NVIDIA RTX 6000 GPU.

The system features an Infiniband HDR interconnect (200 Gb/s) and storage comprising 30 TB home and 1275 TB scratch (NVMe) with sequential throughput of

731 GB/s (write) and 1198 GB/s (read) [100].

The use of supercomputers was necessary since the meshes involved contain a very large number of grid points and the temporal information is stored in a high number of snapshots (the size of each dataset is approximately 10 TB). Consequently, performing such computations on a standard workstation would have been unfeasible due to the prohibitive computational effort.

7.1.1 MATLAB procedure

One possible implementation of the POD-based loss transport equation decomposition relies on a procedure developed in the MATLAB environment (e.g., [80, 101]). The procedure takes advantage of parallel computing (similarly to what is done by Sayadi and Schmid in [102]) by dividing the domain into subdomains, with the user specifying the number of subdomains and processors operating on each of them. The code exploits the `spmd` construct to distribute the computation among workers across multiple computational nodes. Although this procedure can also be executed on a local workstation, it requires long computation times, as intermediate results must be written to disk. A schematic representation of the steps performed by the MATLAB code is provided in Figure 7.1. These are the main steps:

(I) The database obtained from high-fidelity simulations consists of a large number of snapshots N_t. Each snapshot contains the solution of the flow field (velocity components and gas properties) for a different time step.

(II) In this phase, the files are read by different workers in an `spmd` loop to extract the velocity component and coordinate information of the flow field. Then the velocity components are stored separately.

(III) The velocity components are stacked horizontally to create the submatrices U, V and W for the velocity components u, v and w, respectively. These submatrices are tall and skinny and contain the fluctuating velocity components in all points of the domain on the rows and all snapshots on the columns. Then, these submatrices are stacked vertically to create a single tall and skinny matrix MM (a $m \times N_t$ matrix, with $m >> N_t$) which includes the entire dataset. It should be noted that keeping in memory the full matrix MM is often not feasible in standard computational nodes.

(IV) Then, the submatrices U, V and W are divided into k blocks, where k is the number of subdomains for parallel computing. Then a data reshuffling is performed, and the k-th blocks of the 3 matrices are stacked vertically to obtain the matrix M_k.

(V) Each M_k matrix contains all the snapshots of the MM matrix and all three velocity components related to the points in the k-th subdomain. In other words, each matrix M_k contains the entire dataset of the k-th subdomain. Each M_k matrix is a tall and skinny matrix (a $K \times N_t$ matrix, with $N_t << K < m$).

(VI) For each matrix M_k the cross-correlation matrix C_k is calculated by performing the matrix product between M_k^T and M_k. The C_k cross-correlation matrices related to all k subdomains are thus obtained.

(VII) The C_k matrices are summed up to obtain the CC matrix, which is the cross-correlation matrix related to the entire computational domain.

(VIII) From the CC matrix, the eigenvalues and eigenvector matrix X are calculated.

(IX) The matrix product between the M_k matrices, which contain the velocity components related to the different subdomains, and the matrix of eigenvectors X is performed.

(X) In this way we obtain the matrices φ_k, which contain the POD modes related to each subdomain.

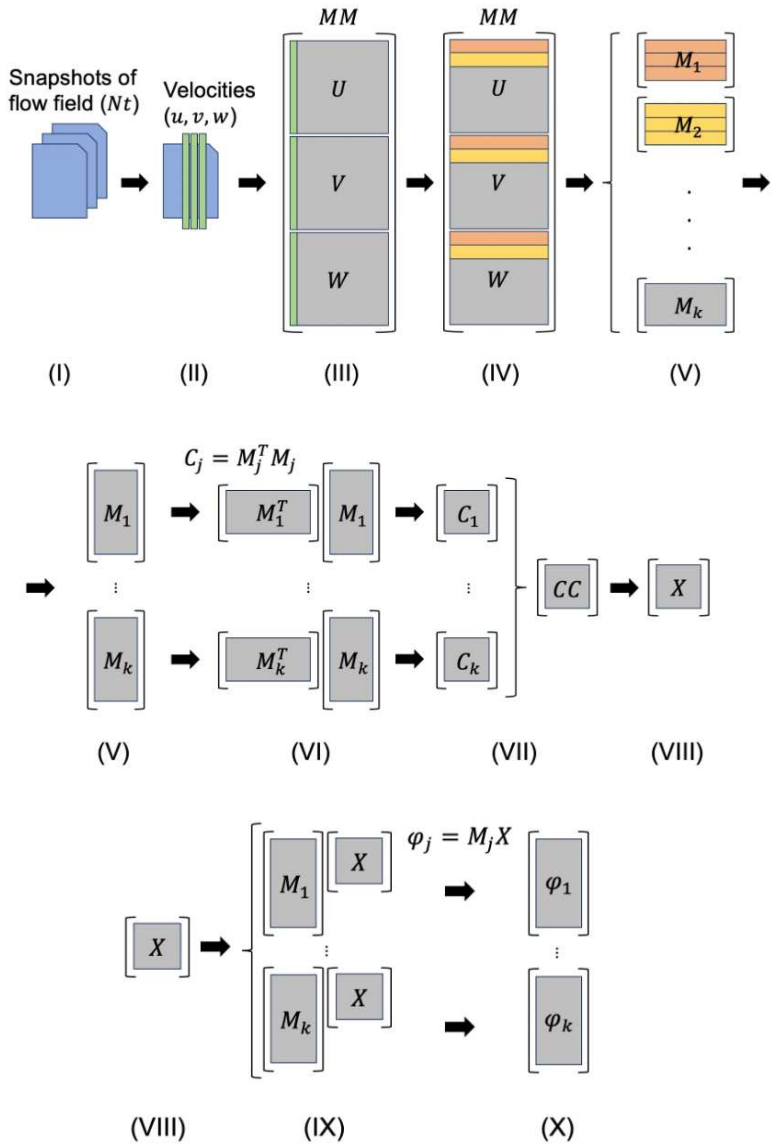

Figure 7.1: Representation of the MATLAB POD procedure considered in this work.

After calculating the POD modes, it is possible to use the equations in the previous sections (Eqs. (4.15), (4.16), (4.17) and (5.5)) to calculate the viscous dissipation associated with the mean flow and the quota of dissipation due to turbulence associated with each POD mode.

For all the blade profiles analyzed during the doctoral activity, the computational domain was systematically divided into eight subdomains. Specifically, four of them were defined in direct proximity to the blade surface, while the remaining four were arranged to surround the inner region. On each of these eight subdomains, a single worker was assigned through the `spmd` construct, ensuring parallel execution of the procedure. This subdivision scheme, consistently adopted throughout the study, is illustrated in Fig. 7.2 for an exemplary domain.

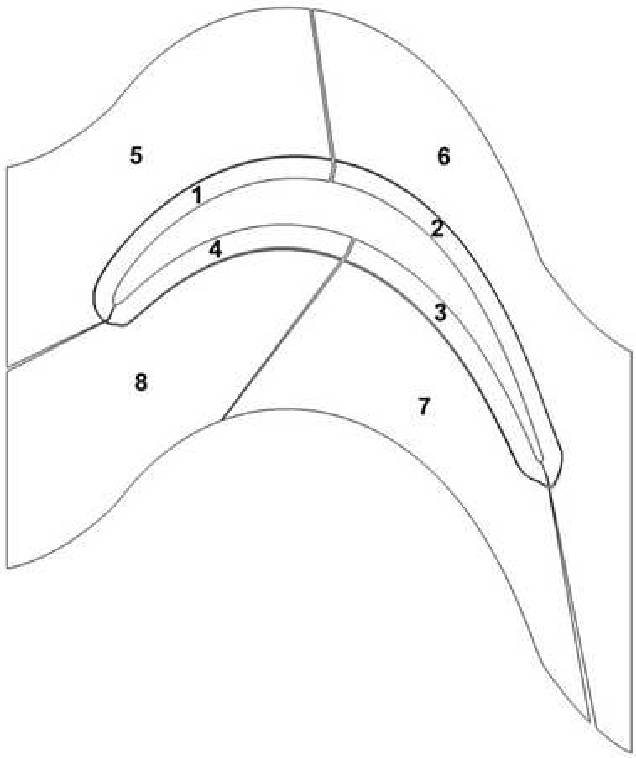

Figure 7.2: Subdivision of the computational domain into eight subdomains, each assigned to a single worker via the `spmd` construct.

In the following section, an alternative implementation of the same procedure in the Python environment will be presented.

7.1.2 Python procedure

The procedure developed in the Python environment is based on Singular Value Decomposition (SVD) and represents an improvement over the MATLAB implementation. The code has been developed within the EU-funded ACROSS project and is openly available in an online repository at the present link.

The Python-Dask library is used to parallelize the computation. The advantage of Dask is that it is no longer necessary for the user to manually partition domain data structures into subdomains since this is done automatically by Dask. Another advantage of Dask is that one can start from a serial prototype based on NumPy and Pandas libraries and convert it very easily to a fairly large-scale distributed code.

The procedure is visualized in Figure 7.3 and described in Algorithm 1. It is worth noting that the formulation with the SVD is equivalent to the one described in the previous paragraph, using the cross-correlation matrix CC.

HPDA algorithm for POD

Inputs: Output files from CFD
% Read input files with Dask lazy read
$lazy\_read = extract.files$
$for\ item\ in\ lazy\_read$
$da\_temp = da.from\_delayed(item)$
$Us_i = append(da\_temp[:, i])$
end
$u_i = concatenate(Us_i).rechunk$
% Compute the time-mean and
%fluctuating velocity components
$u_i\_m = mean(U_i)$
$U_i = U_i - u_i\_m$
% Concatenate the velocity matrices with dask
$MM = da.concatenate([u, v, w])$
% SVD
$svd\_u, svd\_s, svd\_v = svd(MM)$
% Transpose svd_v
$svd\_vt = transpose(svd\_v)$
% Compute POD modes
$\varphi_{Ui} = dot(U_i, svd\_vt)$
% Normalize POD modes
$\varphi_{Ui} = \varphi_{Ui}/N_t\hat{}0.5$
% Compute derivative of mean flow
$dU_i\_dx_j = derivative(U_i, x_j)$
% Compute intergral of Φ_M and P_k
$\Phi_M = \sum((du_i/dx_j)^2 * vol)$
$P_k = \sum(phi\_u_i * phi\_u_j * du_i\_dx_j * vol)$
Output: u_i, u_i_mean, svd_v, Φ_M, P_k

Algorithm 1: Open HPDA Python algorithm

(I) As in the MATLAB procedure we start from N_t files, where each file refers to a different snapshot

(II) There is a phase of reading data single snapshot at a time to extract the velocity components. After this phase, as reported in Algorithm 1, the time-averaged

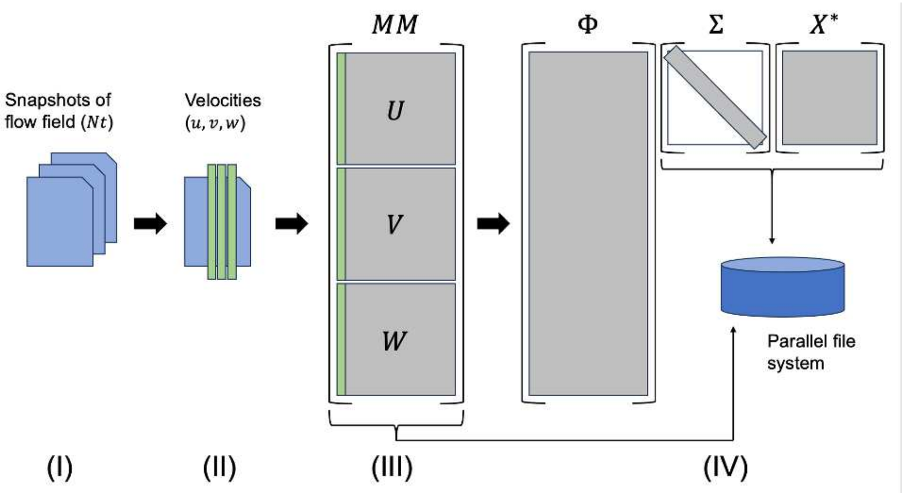

Figure 7.3: Representation of the Python SVD procedure considered in this work.

components of the velocity field are calculated. Then the fluctuating velocity components are calculated as the difference between the flow field components and the time-averaged components. In this way, we obtain the submatrices U, V and W, containing the fluctuating velocity components. It is worth noting that snapshots are read in parallel from disk and all the computation up to the SVD is formalized by Dask as a task graph[1], where each node of the graph is executed as soon as the necessary data becomes available (i.e., the construction of the MM matrix can start as early as some snapshots are available in memory, it does not have to wait for the whole dataset to be loaded).

(III) Then the submatrices U, V and W are stacked vertically to create a single tall and skinny matrix MM (a $m \times N_t$ matrix, with $m >> N_t$) that contains the entire dataset. This matrix exists as a distributed Dask array, arranged in chunks that live distributed among the the memory of the compute nodes.

(IV) In this phase we see the main difference between the procedures described in this article. In fact, with the Python procedure, the SVD of the MM matrix is done, obtaining the Φ, Σ and and X^* matrices (respectively called *svd_u*, *svd_s*, *svd_v* in Algorithm 1). The SVD is also computed in parallel way over the Dask distributed data structure.

In this way, SVD modes can be calculated. These are normalized by the square root of the number of snapshots N_t, as given at the end of Algorithm 1.

7.2 Scalability tests

In this section, the results of several scalability tests performed using data from the typical LES cases described here are presented. Strong and weak scalability have

[1] https://docs.dask.org/en/latest/graphs.html

been considered. In the first case, the LES database is analyzed with an increasing number of computational nodes to observe the speed-up provided by the increase of CPUs. In the second case, the number of snapshots is progressively increased while increasing the number of computational nodes in order to evaluate the efficiency of the code. All tests here described were carried out on *GALILEO100* supercomputer.

The computation times for the MATLAB and Python procedures are reported in the (a) panel of Fig. 7.4. The MATLAB implementation can be executed on a single cluster node (384 GB of RAM), since in this case all intermediate steps are written to disk. Conversely, the Python implementation requires at least three nodes, as data are kept in memory. The reduction in computation time achieved with the HPDA Python procedure compared to the MATLAB one is substantial: the Python version is about five times faster when using six nodes, and up to seven times faster with twelve nodes. As already described in the previous sections, the G100 system is also equipped with so-called fat nodes (3 TB RAM each), which make it possible to run the HPDA Python code on a single node. In this configuration, the procedure is still 1.7 times faster than the MATLAB implementation. However, since only a limited number of clusters provide such large memory nodes, the following results refer to standard nodes only.

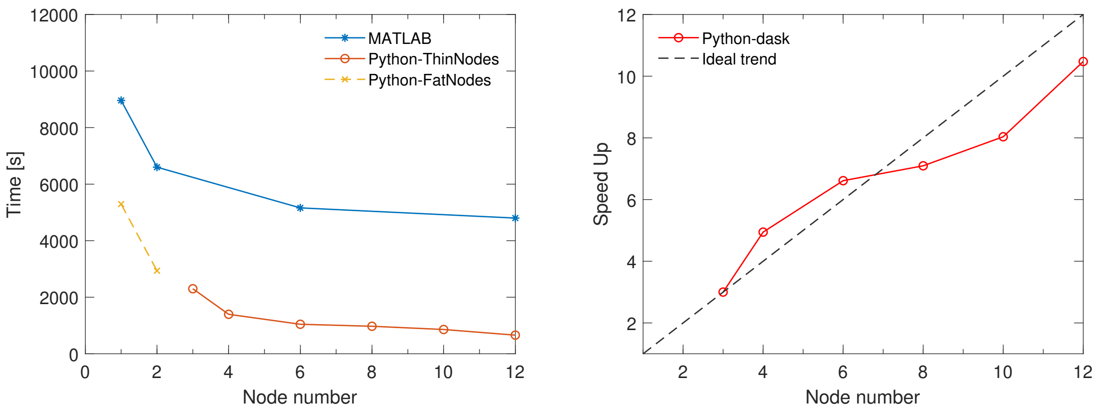

(a) (b)

Figure 7.4: Strong scaling results: (a) time employed by MATLAB and Python procedures; (b) speed up compared to the ideal trend, reference speed is normalized to 3 for three compute nodes.

The speed-up of the HPDA procedure, compared against the ideal trend, is reported in the (b) panel of Fig. 7.4. The ideal scaling assumes a linear speed-up with a slope of one, with a reference value of 3 for three nodes (the minimum required due to memory constraints). The Python code exhibits a super-linear speed-up when moving from three to four nodes, achieving a factor greater than 4. A similar behavior was already observed in the preliminary implementation of the procedure [103], which did not yet include the computation of Φ_M and P_k. This effect is explained by the fact that, with three nodes, the available memory is barely sufficient, forcing the Dask library to spill part of the data to disk. With additional nodes, this partial writing is avoided, resulting in the observed super-linear scalability. The speed-up remains close to the ideal trend up to twelve nodes, for which the total number of CPUs is nearly equal to the number of snapshots. Beyond this point, performance begins to degrade.

Weak scalability results are shown in Fig.7.5. Two databases were considered: the "typical" database, identical to the one used in the strong scalability test, and a "large" database representative of DNS-level simulations, consisting of 80 million spatial points and 960 snapshots. In the latter case, the total size of the input CSV files exceeds 11 TB. For the "typical" database, the reference time t_1 was obtained by running the code on one node with 48 snapshots. The number of snapshots was then progressively increased (96 snapshots on two nodes, and so on) up to the full case on twelve nodes. For the "large" database, the code was executed with only ten CPUs per node, starting from a reference case of 30 snapshots on one node, up to 32 nodes for the full 960-snapshot database. This restriction on multi-threading was necessary to optimize memory requirements per Dask worker. In both cases, efficiency remained around 80% for the largest configurations tested. Notably, the HPDA procedure enabled the post-processing of the "large" database (over 11 TB of CSV input files) in less than 30 minutes. This represents a remarkable improvement in performance, especially considering the simplicity and flexibility of the Python implementation.

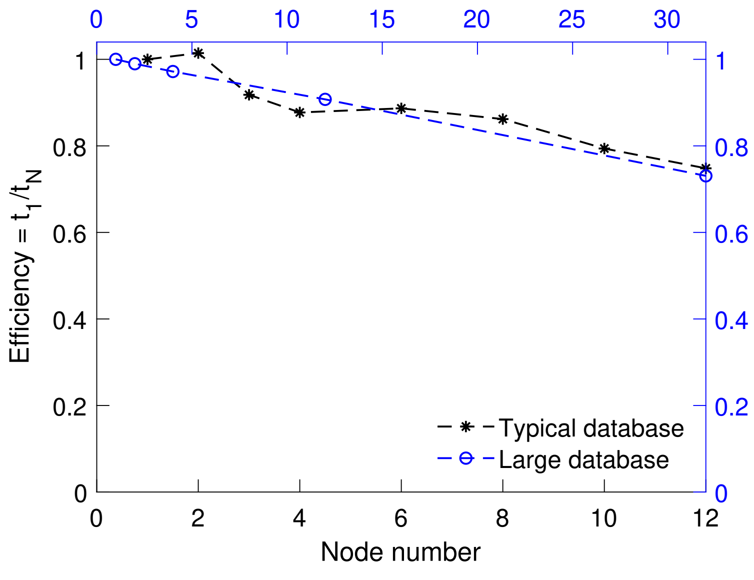

Figure 7.5: Weak scaling. Bottom x-axis (black) is related to the typical database, while the top one (blue) concerns a larger one. Efficiency is computed against the time employed by 1 node with the smallest database that fits its memory.

7.3 Implementation aspects of POD variants

To extend the classical POD procedure to the spectral (SPOD) and phase-locked (PPOD) frameworks, several modifications were implemented in the original code. In this section, these aspects are briefly discussed.

7.3.1 SPOD

Starting from the SPOD, the implemented procedure follows the sequence of steps summarized in Figure 8.1. On this basis, two algorithms for the computation of SPOD modes were developed, with the second representing an improved version of the first.

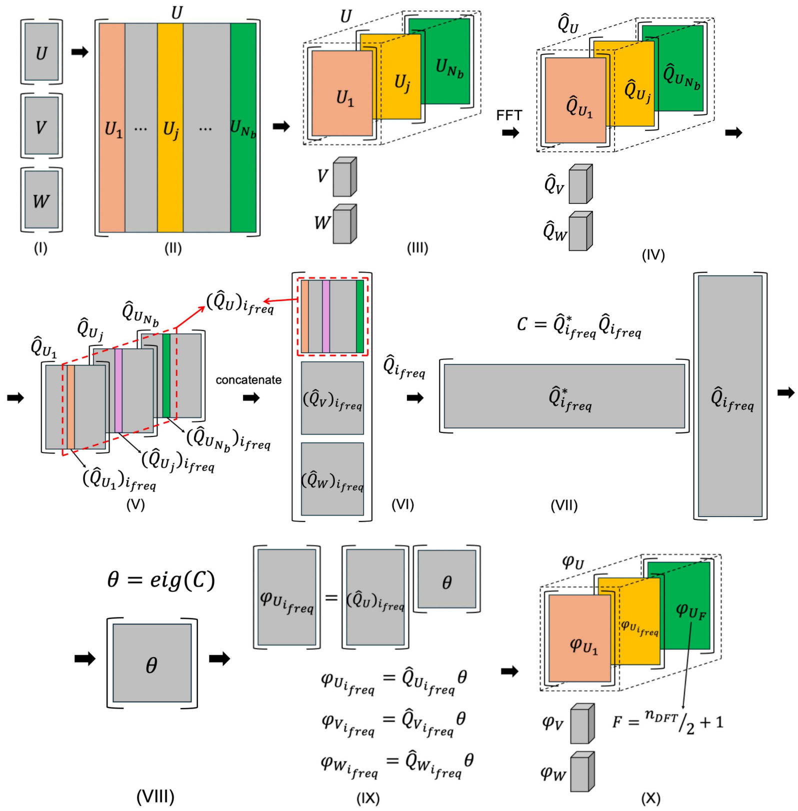

Figure 7.6: Schematic overview of the Spectral-POD procedure.

The baseline SPOD implementation follows the theoretical formulation of [89]. For the practical *serial* implementation, the works of [104], [105], and [106] can be considered as references. The common practice of overlapping data in time, often used in such implementations, was not adopted for the analyzed cases. This choice ensured a fully orthonormal decomposition in both frequency and SPOD modes. Distributed computation is handled directly via Dask, while the variables in the SPOD kernel are kept separate whenever possible. Specifically, the operations managed by Dask are highlighted in Algorithm 2 with the suffix “da”, whereas operations explicitly using the NumPy library are indicated with “np”.

The input files are read by Dask and stored in three data matrices. The three fluctuating velocity components (U_i in the algorithm) are stored in the three data matrices U, V, W (step I of Figure 7.6). Each column vector of the matrix retains the spatial information. The rows got dimension N_t (i.e., the number of instantaneous realizations or snapshots), while the columns got dimension M_s, the total number of computational mesh points. The first step of SPOD is fragmenting the time sequence in a discrete number of blocks N_b (step II), which is obtained by means

HPDA algorithm for SPOD

Inputs: Output files from CFD
% Read input files with Dask to obtain data matrices
$U_i.shape = [M_s, N_t]$
% Compute and remove the mean
$u_i\_m = da.mean(u_i)$
$U_i = U_i - u_i\_m$
% Compute SPOD modes
$U_i = U_i.reshape(M_s, N_b, n_{DFT})$
$\hat{Q}_{Ui}$ =da.fft(U_i,axis= 2)
For i_{freq} in range$(0, n_{DFT}/2 + 1)$
  $(\hat{Q}_{Ui})_{ifreq}$ =da.squeeze($\hat{Q}_{Ui}(:, :, i_{freq})$)
  $\hat{Q}_{ifreq}$=da.concatenate($(\hat{Q}_{Ui})_{ifreq}$)
  C_{ifreq} =da.dot($\hat{Q}^H_{ifreq}, \hat{Q}_{ifreq}$)
  θ =np.eigh(C_{ifreq})
  φ_{Ui} =da.dot($(\hat{Q}_{Ui})_{ifreq}, \theta$)
End
% SPOD modes are normalized
% End of SPOD function
% Compute derivative of mean flow
$du_i\_dx_j = derivative(u_i, x_j)$
% Compute P_k per mode
$P_k = conj(\varphi_{Ui}) * \varphi_{Uj} * du_i\_m\_dx_j$
% Integral of Φ_M and P_k
$\phi_M = \sum((du_i\_m/dx_j)^2 * vol)$
$P_k = \sum(P_k * vol)$
Output: Φ_M, P_k

Algorithm 2: HPDA algorithm for SPOD computation.

of an array reshape operation (step III). Dask then manages the discrete Fourier transform (DFT) using the FFT algorithm along the temporal direction of size n_{DFT} (with $n_{DFT} * N_b = N_t$) to obtain the matrices $\hat{Q}_{Ui}$ for each velocity component (step IV). Each 3D array $\hat{Q}_{Ui}$ contains the frequency information along its last dimension (shape 2 according to Python indexing, which starts from 0).

In the baseline implementation of SPOD, the data for each frequency are then collected within a *for* loop in order to store the matrices $(\hat{Q}_{Ui})_{ifreq}$ at the frequency loop index *ifreq* (step V). This operation is done for all variables at hand (step VI). The number of frequencies of interest is $n_{DFT}/2+1$, as specified in the algorithm, and this is because of the symmetry properties of the DFT. The matrices $(\hat{Q}_{Ui})_{ifreq}$ are then concatenated to form the $\hat{Q}_{ifreq}$ matrix (from step VI to VII) in order to compute the cross-correlation matrix C_{ifreq} within the *for* loop (step VII) and its eigenvectors θ_{ifreq} (step VIII). Note that, at present, Dask does not provide the equivalent of the eigenvector calculation of the NumPy library, and therefore the computation of step VIII is managed by NumPy. Nevertheless, the C_{ifreq} matrix is very small with dimension $[N_b, N_b]$. The SPOD modes are computed separately for each variable as $(\varphi_{Ui})_{ifreq} = (\hat{Q}_{Ui})_{ifreq}\theta$ (step IX) and stored in a 3D array (step X) with the same shape as matrices in step III. Finally, to ease the computation of the P_k, the SPOD modes are reshaped to a matrix with dimension $[M_s, N_t]$.

The SPOD computation, detailed in steps V to IX of Fig. 7.6, has been optimized by eliminating the use of *for* loops. The improved algorithm, outlined in Alg. 3, takes advantage of the execution order in the eigenvector solver from the NumPy library. The input data are rearranged using stacking and transposing operations, resulting in a three-dimensional array, C, with dimensions $[n_{DFT}/2+1, N_b, N_b]$. The first dimension corresponds to the frequency index, and the other two dimensions represent the cross-correlation matrix for each frequency, C_{ifreq}. The numpy function *eigh* is then applied, solving the eigenvector problem for each C_{ifreq}. This optimization ensures that the output SPOD modes are equivalent to those produced by the baseline implementation.

Optimized HPDA algorithm for SPOD

Inputs: Output files from CFD
% Read input files with Dask to obtain
% data matrices
$U_i.shape = [M_s, N_t]$
% Compute and remove the mean
$u_i\_m = da.mean(U_i)$
$U_i = U_i - u_i\_m$
% Compute mean viscous dissipation
% Φ_M and save to file
% Compute SPOD modes
$U_i = U_i.reshape(M_s, N_b, n_{DFT})$
$\hat{Q}_{Ui}$ =da.fft(U_i,axis= 2)
$\hat{Q}_{Ui} = \hat{Q}_{Ui}[:,:,0 : n_{DFT}/2+1]$
$\hat{Q}$ =da.stack($\hat{Q}_{Ui}$,axis= 1)
$\hat{Q} = \hat{Q}.reshape(M_s * 3, N_b, n_{DFT}/2+1)$
$\hat{Q} = \hat{Q}.transpose$(2,0,1)
C =da.matmul($\hat{Q}^H, \hat{Q}$)
θ =np.eigh(C)
φ_{Ui} =da.matmul($(\hat{Q}_{Ui}).T, \theta$)
% with $\hat{Q}_{Ui}.T = \hat{Q}_{Ui}.transpose$(2,0,1)
$\varphi_{Ui} = \varphi_{Ui}.transpose$(1,0,2))
  % Normalize $\varphi_u, \varphi_v, \varphi_w$
  % Compute P_k using φ_{Ui}
Output: Φ_M and P_k

Algorithm 3: Optimized HPDA algorithm for SPOD computation.

Strong and weak scalability tests for baseline and optimized SPOD implementations are shown in Figure 7.7. In the former, the LES database is analyzed with an increasing number of computational nodes to observe the speed-up obtained by increasing the CPUs. Conversely, in the latter, the number of snapshots is progressively increased alongside the number of computational nodes to assess the code efficiency. All tests described here were executed on the CINECA cluster LEONARDO, utilizing the general-purpose CPU partition (DCGP). For further details, the reader is referred to [90].

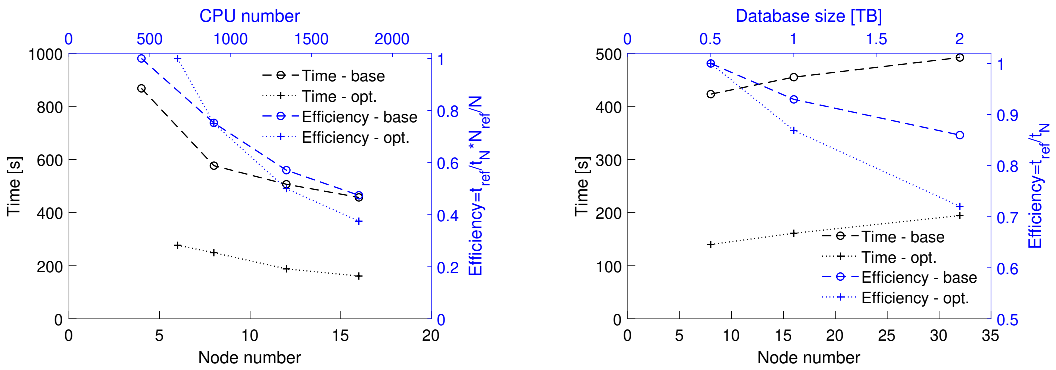

(a) Strong scalability (b) Weak scalability

Figure 7.7: Running time and efficiency of scalability tests: strong scalability (a); weak scalability (b).

7.3.2 PPOD

Similarly to what was done with SPOD, the classical POD framework was further extended to the phase-locked formulation (PPOD). In this case, the snapshots are first sorted according to their phase within the periodic cycle, and a separate POD decomposition is applied to each phase bin. The sequence of steps implemented in the PPOD procedure is schematically represented in Figure 7.8.

The PPOD implementation builds on the existing Dask-based workflow. Initially, the three velocity components and their mean fields are read as Dask dataframes and converted to Dask arrays, which are reshaped according to the structured mesh and rechunked to optimize blockwise parallelism (step I of Fig. 7.8). Key arrays are persisted in memory to accelerate repeated access during computations.

The phase-locked procedure starts by computing the instantaneous phase of each snapshot from the first two temporal coefficients of the precomputed right-singular vectors. The resulting phase angles are binned into a fixed number of sectors, and a phase-sorted sequence of snapshot indices is built, ensuring that each bin contains an equal number of snapshots (step II).

For each phase bin, the corresponding subset of velocity snapshots is selected and concatenated to form a phase-specific data matrix (step III). A distributed singular value decomposition is then performed on this matrix (step IV), yielding right-singular vectors that are used to project the snapshots onto spatial modes for each velocity component. The modes are normalized by the square root of the number of snapshots in the phase to maintain energy consistency. After reshaping back to the 3D mesh layout (step V), the modes are saved to disk for subsequent analysis.

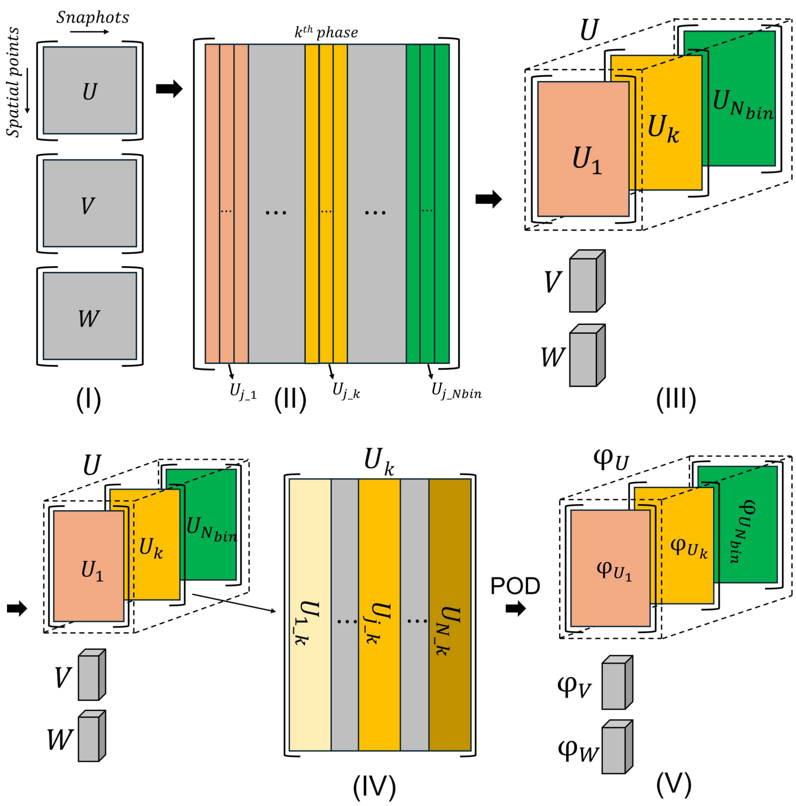

Figure 7.8: Schematic overview of the Phase-locked POD procedure.

Finally, the phase-resolved turbulent kinetic energy production is evaluated by contracting the modal velocity fields with the mean-flow gradients, accounting for all cross-terms and weighting by local cell volumes. This produces an array of phase-resolved TKE contributions, which is also stored for further post-processing. The code structure, as outlined above, is summarized in Algorithm 4.

7.4 Implementation aspects of DWT

The DWT procedure was implemented within the HPDA framework developed in Python for the use of POD and its variants. The Discrete Wavelet Transform was implemented using the Python package `PyWavelets` (`pywt`), leveraging the same HPDA framework used for POD, SPOD, and PPOD. To ensure a fair comparison with the other modal decomposition techniques, only orthogonal mother wavelets were considered. Several mother functions were tested, including Daubechies (`db1`, `db2`, `db4`), Coiflets (`coif2`, `coif3`), and Symlets (`sym4`, `sym8`). No appreciable differences were observed in the scale-separation capability or in the identification of dominant flow structures. For this reason, the `db1` wavelet was selected as the reference for the comparative analysis. The workflow is articulated into the steps reported in Algorithm 5.

HPDA algorithm for PPOD

Inputs: Output files from CFD, SVD modes
% Read input files with Dask to obtain data matrices
% Load right-singular vectors svd_V from previous POD
% Compute mean viscous dissipation
% Φ_M and save to file
% — Phase sorting —
% Extract two leading temporal coefficients from svd_V
$eig1 = svd_v[0, :]$
$eig2 = svd_v[1, :]$
% Compute instantaneous phase angle
$\theta = np.degrees(np.arctan2(eig2, eig1))$
% Sort snapshots into N_{bin} phase bins
$edges = np.linspace(min(\theta), max(\theta), N_{bin} + 1)$
$N, bin_{edges} = np.histogram(\theta, bins = edges)$
$bin_{indices} = np.digitize(teta, bin_{edges}) - 1$
% — Phase-locked POD —
For $phase$ in range$(0, N_{bin})$:
  $U_{i\_phase} \leftarrow$ select snapshots in bin $phase$
  $U_{i\_phase}.shape = [Ms, N]$
  % Perform SVD: $U_{i\_phase} \approx \Phi\Sigma V^T$
  % Normalize φ_{Ui_phase}
  % Compute P_k using φ_{Ui_phase}
End
Output: Φ_M and P_k for each phase bin

Algorithm 4: HPDA algorithm for PPOD computation.

HPDA algorithm for DWT

Inputs: Output files from CFD
% Read input files with Dask to obtain data matrices
$U_i.shape = [M_s, N_t]$
% Compute and remove the mean
$u_i\_m = da.mean(U_i)$
$U_i = U_i - u_i\_m$
% Compute mean viscous dissipation % Φ_M and save to file
% Define DWT parameters
$n_{levels} = \log_2(N_t)$
$\Delta t \rightarrow f_s$
% Loop over spatial chunks (lazy)
$for\ k\ in\ n_{blocks}$
$U_i^{(k)} = delayed(block_k)$
$\varphi_{U_i}^{(k)} = \mathrm{DWT}(U_i^{(k)})$
end
% Assemble DWT features
$\varphi_{U_i} = concatenate(\varphi_{U_i}^{(k)})$
% Reshape to physical domain
% Normalize $\varphi_u, \varphi_v, \varphi_w$
% Compute DWT mode energy
$E_l = \sum_{x,y,z}(\varphi_u^2 + \varphi_v^2 + \varphi_w^2)$
% Compute TKE production per scale
$P_k = \sum(\varphi_{Ui} * \varphi_{Uj} * du_i\_dx_j * vol)$
Output: Φ_M and P_k

Algorithm 5: HPDA algorithm for DWT computation.

7.5 Spatial localization of loss sources

To enable the spatial localization of losses, a filtering procedure was implemented in MATLAB through logical masks, allowing the contribution of the coherent structures identified by POD and its variants to be associated with distinct physical regions of the flow. The computational domain was divided into five regions of interest, schematically illustrated in Figure 7.9. The first step consisted in separating the boundary-layer zones from the undisturbed flow by means of a logical condition applied to the local vorticity magnitude. Grid points where the vorticity exceeded a prescribed threshold—calibrated on the basis of several reference cases from the literature—were classified as belonging to the boundary layer in the vicinity of the blade. This region was then further divided into the pressure-side boundary layer and the suction-side boundary layer, the latter being split into an accelerating and a decelerating section. The subdivision was performed automatically by identifying the streamwise grid index corresponding to the peak suction location on the blade, determined from the static pressure distribution obtained from the LES database.

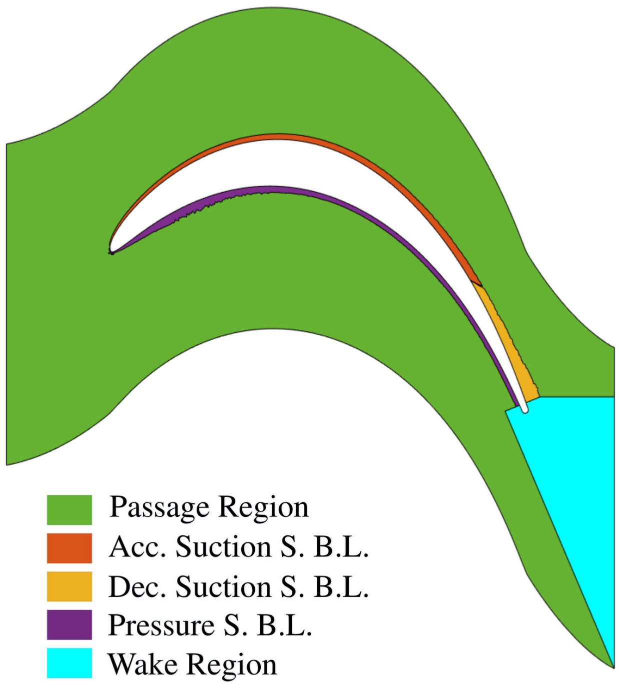

Figure 7.9: Sub-regions for integration of viscous and turbulent losses.

Outside the boundary layer, an additional trapezoidal region was defined to isolate the wake downstream of the trailing edge. This area was delimited by four vertices: two located at the streamwise coordinate of the suction and pressure sides in the immediate vicinity of the trailing edge curvature, with a wall-normal position equal to the maximum boundary-layer thickness on the blade; and two placed at the upper outlet boundary, the first at the minimum y coordinate and the second at a distance of approximately 20% of the blade pitch from the maximum y position. The remaining portion of the computational domain was finally identified as the passage region.

The subregions thus obtained were employed as integration subdomains for the evaluation of the loss contribution associated with each POD mode. This approach enabled a concurrent partitioning of the losses, both in terms of their spatial localization and with respect to the dynamical features representative of the underlying loss-generation mechanisms, providing a consistent framework for a more detailed physical interpretation of the modal energy distribution. This approach joins the methodology employed by Lengani et al. [33], Dellacasagrande et al. [30], and Leggett et al. [107].

Chapter 8

Loss Analysis Results

The developed HPDA code aims to achieve an almost real-time decomposition of the data produced by high-fidelity simulations. As explained, data in the time domain (e.g., the snapshots of Fig. 6.2) cannot be directly used to separate the dynamical features. Namely, they may easily provide an inlet-to-outlet loss coefficient or an integral value when transport equations are applied. However, as it may be observed from the instantaneous P_k of Fig. 8.1, it is not possible to decouple the physical features. Specifically, Fig. 8.1 shows that high (red) values of turbulent kinetic energy production may be observed in the middle of the passage, above the front part of the suction side. This local production is related to the position of the inlet wake (as identified from the instantaneous flow visualization of Fig. 6.2) but highlights both small and large scale structures. Similarly, in the rear part of the boundary layers, a multitude of structures of different scales may be observed. To overcome this limit, the HPDA adopts the POD framework for decoupling dynamical features. Furthermore, as explained in Chapter 7, different decomposition frameworks can be easily implemented in the HPDA code by replacing the SVD operation with others, leveraging the ease of implementation offered by the Python-Dask library.

The present Chapter is divided into two sections. First, in Section 8.1, the Discrete Fourier Transform is considered and compared to the POD and its variants to further investigate the capabilities of the presented loss analysis tool. In Section 8.2, the results of applying these techniques to the available dataset are presented.

8.1 Comparison between modal decompositions

Different decomposition frameworks provide different rankings of flow dynamics. Figure 8.2 shows the frequency-mode relationship for both POD and DFT. In the latter case, the relationship is linear (as highlighted by the red dots), since each Fourier mode represents a specific frequency, which is a multiple of the frequency resolution. The relationship between POD mode and frequency is obtained from the DFT analysis of the temporal coefficients.

As POD modes are ranked by energy, it is possible to find a maximum frequency for

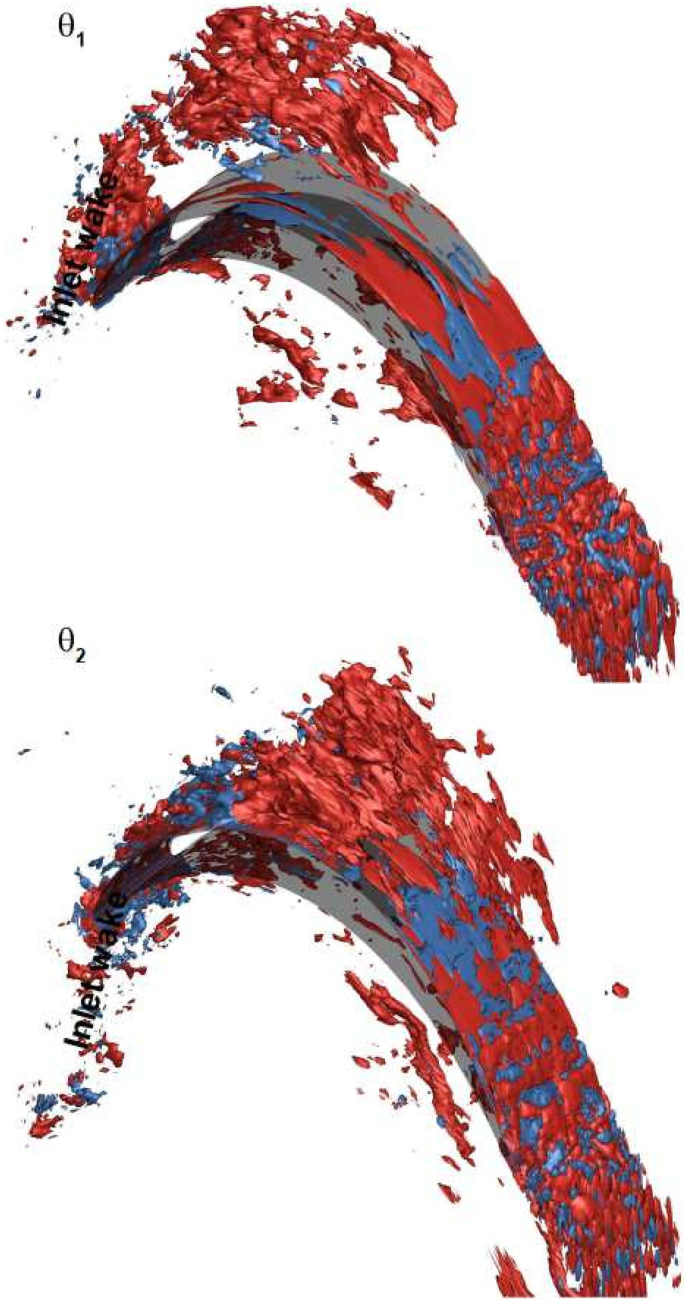

Figure 8.1: Iso-surfaces of instantaneous turbulent kinetic energy production.

each mode (represented by the black dots), while more often multiple frequencies coexist in the POD mode (gray shade in Fig. 8.2). The first mode is not related to the smallest observable frequency, but rather to the principal fluctuation of the system. POD modes are real valued and often paired [108], namely two modes with similar energy have the same main frequency. Therefore, the inclination of the POD mode frequency fit (blue dotted line of Fig. 8.2) is half of that of DFT since the Fourier modes are complex valued and symmetric above the Nyquist frequency (half of the sampling frequency).

The spatial distribution of turbulent kinetic energy (TKE) production captured by

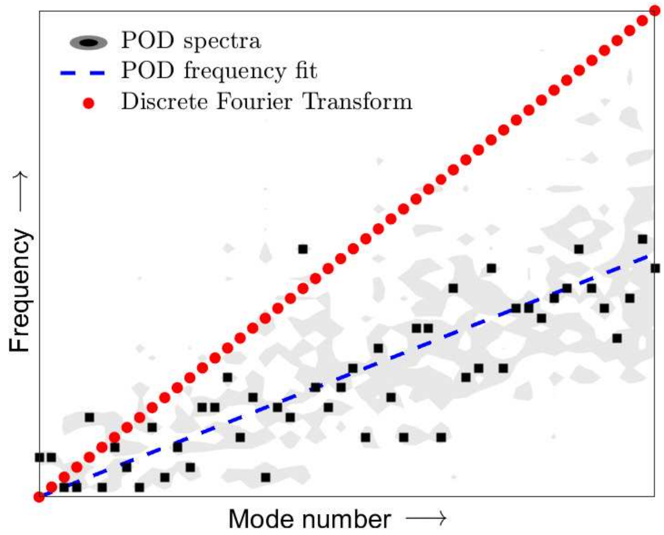

Figure 8.2: Relation between POD mode and Fourier transform frequency.

the POD and DFT modes is reported in Figs. 8.3a and 8.3b, respectively. Two modes are shown as representative examples of the procedure. The first POD mode captures the periodic disturbances of the incoming wakes, as also discussed in Lengani et al. [101] and Dellacasagrande et al. [30]. This decomposition clearly highlights the TKE production in the blade passage due to large-scale disturbances induced by the inlet wakes, thus enabling a decoupling of loss mechanisms that cannot be directly inferred from the instantaneous flow field.

The frequency associated with this POD mode corresponds to Fourier mode index 5, as shown by the DFT mode in Fig. 8.3b. The DFT representation of P_k at that frequency is in good agreement with the POD mode. Minor differences arise because two POD modes exist at the same frequency, whereas the DFT requires only one mode. Additionally, each Fourier mode contains real and imaginary components, while P_k computed from the DFT is a real value due to the orthogonality of the basis. In essence, one Fourier mode encodes the same information as two POD modes. A similarly good agreement is observed for higher-frequency phenomena shown in the lower panels of Figs. 8.3a and 8.3b, where both decompositions clearly highlight elongated streaky structures undergoing breakdown, representative of the bypass transition scenario [109].

As a validation of the approach, the cumulative sum of losses for POD and DFT is compared with the inlet-to-outlet coefficient ω of Eq. (4.13) in Fig. 8.4. According to Eq. (4.17), the volume integral approach is adequately performed by computing the viscous dissipation due to the mean flow Φ_M and the dissipation due to turbulence $\overline{\Phi}_k$. As explained in Chapter 3, the term $\overline{\Phi}_k$ can be directly computed only under the condition that a DNS is performed, as it is mainly associated with the small-scale turbulent structures. In the absence of highly resolved flow field data, the dissipation due to turbulence can be evaluated with a sufficient degree of approximation through

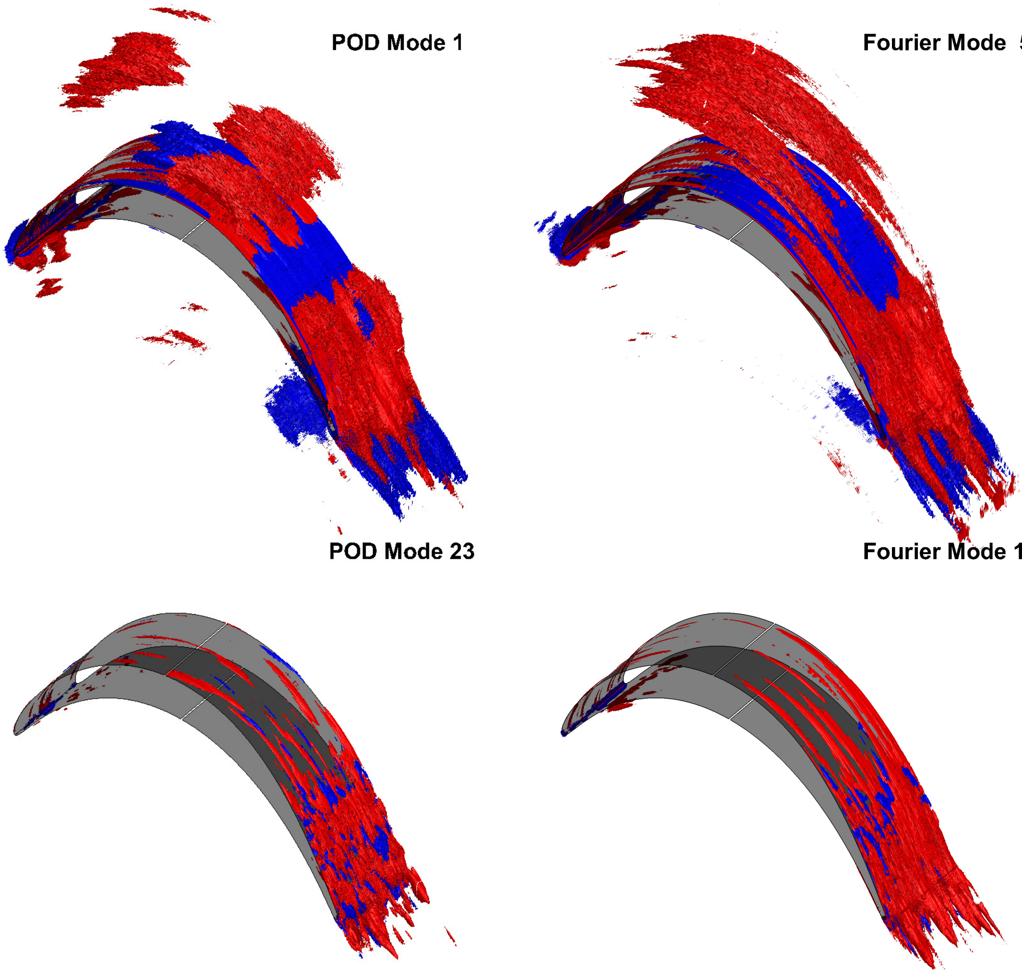

(a) POD

(b) DFT

Figure 8.3: Iso-surfaces of turbulent kinetic energy production from modal decompositions: POD (a) and DFT (b).

the production of TKE P_k when turbulence is in equilibrium over the computational domain. If this is not the case, $\overline{\Phi}_k$ can be retrieved from the summation of P_k and the substantial derivative of TKE (as for contraption of Eq. (4.16)), providing consistency with the definition of the ω coefficient.

To further clarify this aspect, tables 8.1 and 8.2 report the integral budgets of the pseudo-stagnation pressure and turbulent kinetic energy transport equations, respectively. The presented values refer to a representative case from the dataset. The integration domain corresponds to the stator passage alone and extends 25% of the axial chord upstream of the leading edge and 20% downstream of the trailing edge.

Both diffusion terms, $V.D.$ and $T.D.$, appearing in Eq. (4.15), are negligible, and the same applies to the diffusion terms $V.D._k$, $T.D._k$, and $P.D._k$ in Eq. (4.16). Their values are consistently below 5% of the left-hand side term of the corresponding equation. However, the flow is not in perfect turbulent equilibrium, as the substantial derivative of turbulent kinetic energy $F.T.K.E.$ is non-zero. This behavior reflects the inherent characteristics of a Large Eddy Simulation (LES), where the dissipation effect of small-scale vorticity is only partially modeled, leading to a residual imbalance between the production and dissipation of turbulent kinetic energy.

Consequently, the correct evaluation of the time-averaged material derivative of total pressure, as introduced above, is obtained from the modified formulation of Eq. (4.18).

Table 8.1: Pseudo-stagnation pressure transport equation budgets.

$$F.A.S.P. = -\Phi_m - \rho P_k - V.D. - T.D. \quad [W/m^3]$$

$F.A.S.P.$	Φ_m	ρP_k	$V.D.$	$T.D.$
-0.086	0.041	0.055	0.003	-0.004

Table 8.2: Turbulent kinetic energy transport equation budgets.

$$F.T.K.E. = V.D._k + P_k - T.D._k - P.D._k - \overline{\Phi}_k \quad [m^2/s^3]$$

$F.T.K.E.$	$V.D._k$	P_k	$T.D._k$	$P.D._k$	$\overline{\Phi}_k$
-0.0054	-0.0002	0.0449	0.0007	-0.0009	0.0149

The blue and red curves in Fig. 8.4, corresponding to the POD and DFT transforms of the transport equation (Eq. (4.15)), are obtained by computing the dissipation due to the mean flow (where the curves intercept the y-axis) and the cumulative summation over the mode index of the production of TKE P_k associated with each one of the modes identified through POD and DFT. As can be observed, the curves converge to the same final value. As both are orthogonal bases, the reconstruction with all modes yields the time-averaged loss coefficient. For the DFT, only half of the modes were retained (with double amplitude) to exploit Fourier symmetry and reduce memory usage, since each mode carries both real and imaginary parts, whereas all N_t POD modes are needed for full reconstruction. Differences appear only in the first few modes, where the POD ranks the mode related to incoming wakes at frequency index 5 as dominant (Fig. 8.2).

While computed through the summation of Φ_M and P_k (instead of $\overline{\Phi}_k$), the total POD and DFT cumulative losses closely match the inlet-to-outlet surface integral coefficient ω (orange square in Fig. 8.4), with discrepancies below 3%, mainly due to interpolation and differentiation accuracy. When including the substantial derivative of TKE in the computation of cumulative losses for POD and DFT (purple dot), the error reduces below 2%. Notably, this derivative can also be decomposed onto an orthogonal basis and efficiently implemented in the HPDA code, as done in Russo et al. [85], to limit errors in the volume integral approach.

Overall, the proposed procedure allows for a rapid and accurate computation of the decomposed transport equation of total pressure loss.

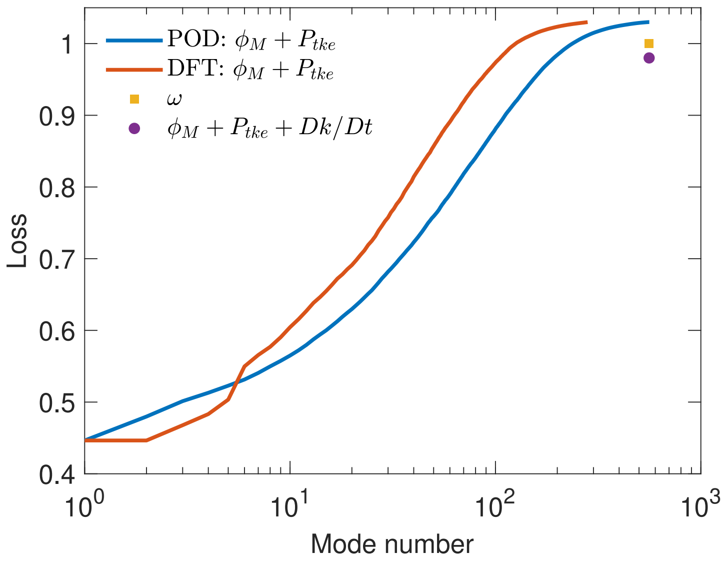

Figure 8.4: Cumulative curves obtained from different integration methods.

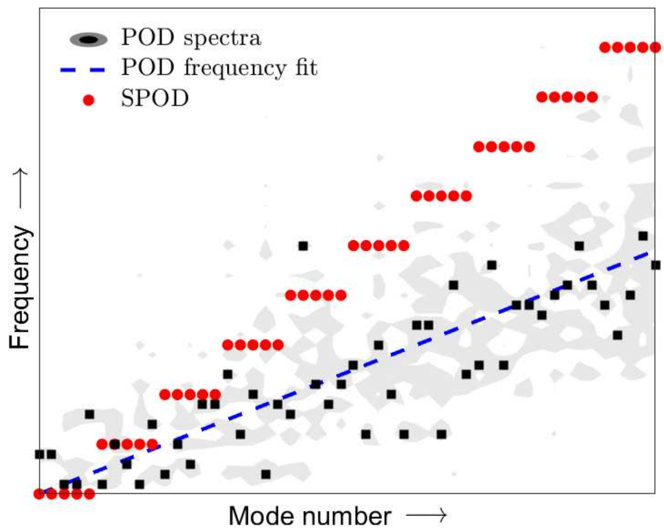

Figure 8.5: Relation between POD and SPOD mode frequency.

What has been highlighted for POD in comparison with DFT, in terms of the relationship between mode index and captured frequency, can be extended to the other techniques discussed in this thesis, namely SPOD and PPOD.

As can be deduced by the plot of Figure 8.5, SPOD provides a direct association between each mode and a specific spectral interval, thus enabling the isolation of coherent structures characterized by well-defined frequencies.

PPOD, on the other hand, exploits a snapshot synchronization logic based on a periodic reference event (e.g., rotor blade passing). As a consequence, the chronological sequence of snapshots is completely lost and any direct frequency resolution is inherently absent. This decomposition can only provide a statistical description of the flow structures conditioned on the selected phase of the period. In other words, the PPOD is the only technique among those considered in this work that does not retain a frequency–mode relationship, but instead resolves the flow dynamics exclusively in terms of their phase-dependent statistics.

8.2 Insights gained from the loss analysis procedure

To better illustrate the potential of the proposed loss analysis framework, this section presents the results obtained from the application of the different modal decomposition techniques to the available dataset.

For confidentiality reasons related to the proprietary nature of the geometries and high-fidelity simulations provided by Avio Aero, the detailed results are shown only for a single representative profile from the examined LPT cascade set. The blade loading distribution of this reference profile, shown in Figure 8.6, provides an overview of the main inviscid features of the flow and serves as a basis for interpreting the loss structures revealed by the subsequent modal analyses. As can be observed, the profile's blade loading is characterized by an aft-loaded distribution, with the peak suction occurring at approximately 74% of the axial chord. Compared to the other profiles in the dataset, this profile exhibits a relatively high overall loading magnitude.

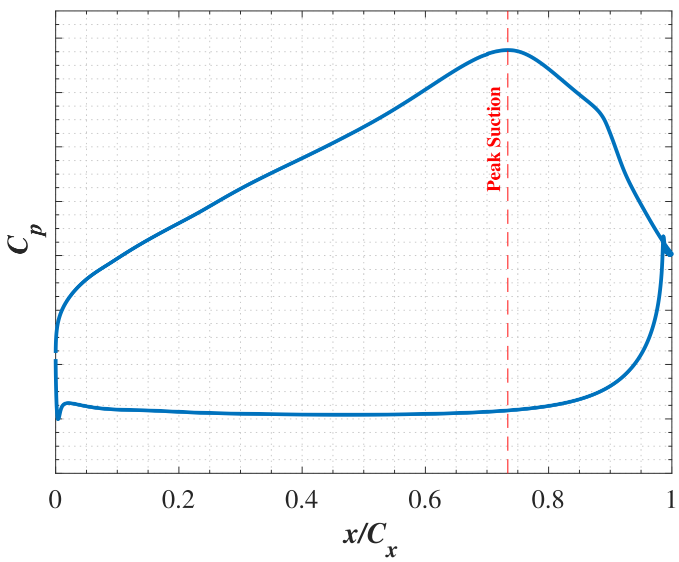

Figure 8.6: Blade loading distribution of the reference LPT profile selected as exemplary case for the modal analysis.

The application of Proper Orthogonal Decomposition to the dataset allows for a detailed understanding of the mechanisms of loss production in the considered profile. Following the mathematical framework detailed in Chapter 5, by projecting the high-fidelity data onto the orthogonal POD basis and integrating the dissipation term quota associated with every POD mode (i.e, coherent flow structure) over the physically relevant sub-regions of the computational domain described in Section 7.5, it is possible to obtain the loss distributions per region, as shown in Figure 8.7. This visualization highlights which regions of the domain contribute most to the overall dissipation and how the contributions of individual modes evolve across the mode hierarchy. Such an approach enables the decomposition of the total losses in terms of both the spatial locations where they are generated and the coherent flow structures responsible for them.

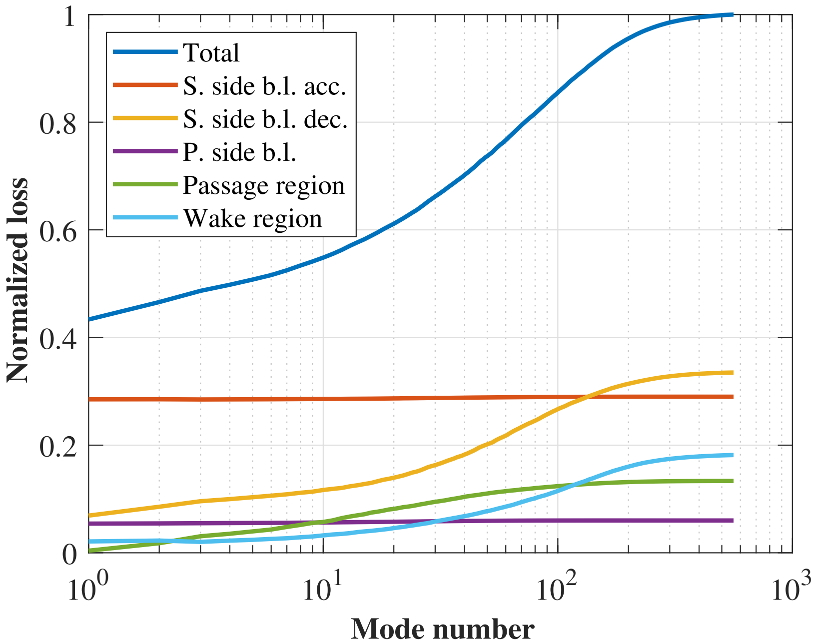

Figure 8.7: Cumulative contribution of POD modes to total losses, split by spatial regions of interest: suction side boundary layer, pressure side boundary layer, passage, and wake.

From a purely spatial perspective (i.e. by inspecting the final level reached by each curve), it is immediately evident that the largest share of total losses originates from the suction-side boundary layer (both accelerated and decelerated portions) and from the wake region downstream of the trailing edge.

From a modal distribution point of view, by construction of the loss distribution curves, the influence of a single mode or a set of modes on the losses within a given region is directly associated with the local slope of the corresponding cumulative curve: the steeper the slope, the larger the contribution of the associated dynamics to entropy generation. As expected, the accelerated suction-side boundary layer (red curve) and the pressure-side boundary layer (purple curve) are nearly independent of the mode index, as their curves remain almost flat across the modal range. This behavior indicates that losses in these two regions are almost exclusively related to the mean-flow contribution of viscous dissipation, with limited influence from turbulent fluctuations captured by higher-order modes. In all remaining regions, the cumulative curves exhibit a much stronger dependence on the POD mode index. In particular, both the inter-blade passage (green curve) and the decelerated portion of the suction-side boundary layer (yellow curve) are significantly affected by the first POD modes, which capture the highest share of turbulent kinetic energy. This trend is in contrast with the wake region (cyan curve), where the contribution of the first modes remains marginal. In fact, the first POD mode, illustrated through the three-dimensional iso-surfaces of the streamwise velocity component in Fig. 8.8a, captures large-scale coherent structures occupying almost the entire passage and generating a strong impingement effect on the decelerating portion of the suction-side boundary layer. Conversely, the wake region shares with the decelerated suction-side layer the fact that it is mostly influenced by higher-order modes, representative of smaller turbulent scales, which have limited impact within the passage region itself.

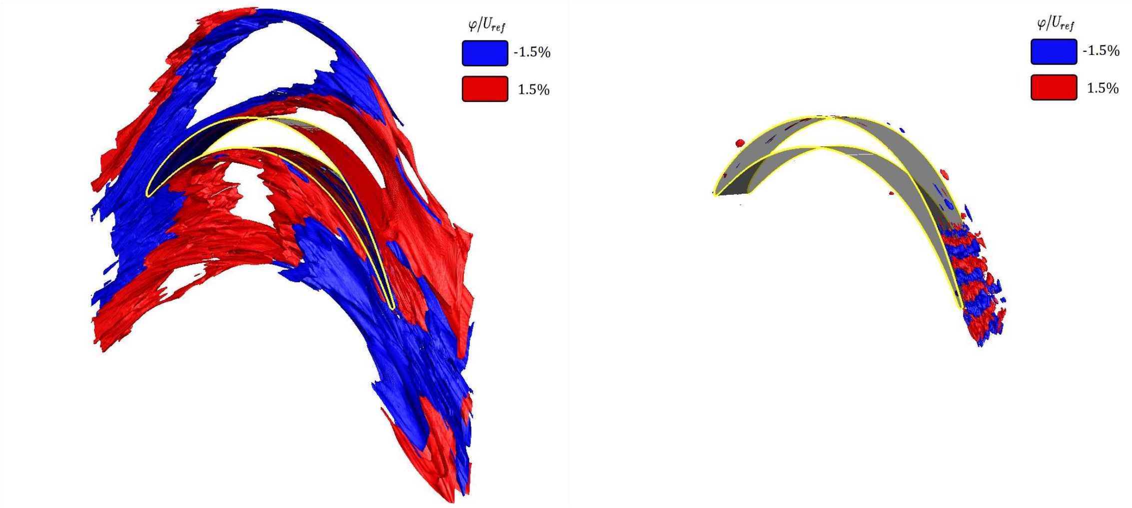

(a) Mode 1 (b) Mode 60

Figure 8.8: Isosurfaces of POD modes of the stream-parallel velocity for the analyzed profile.

The large loss levels observed in the decelerated suction-side boundary layer and in the wake region are therefore indicative of highly dissipative dynamics associated with the smallest turbulent scales. By inspecting the spatial structure of the POD modes at the index corresponding to the rapid change in slope in the suction-side boundary-layer curve—an example of which is reported in Fig. 8.8b—it is evident that the captured dynamics correspond to the evolution of the suction-side boundary layer, which ultimately breaks down into the mixing region downstream of the trailing edge.

These observations suggest that the overall losses are primarily driven by the interaction of the large-scale structures convected through the passage with the boundary layer. The resulting mechanisms presumably involve a cross-scale interaction at different frequencies and, due to the strong flow unsteadiness, take place at different time instants. However, as already mentioned, the classical POD does not provide temporal or spectral localization of the captured dynamics. A high-level assessment of the dominant frequencies and their temporal evolution can instead be obtained through a multi-resolution decomposition based on the wavelet transform. The resulting DWT loss map reported in Figure 8.9 highlights the contribution to P_k of the different DWT modes for all frequency bands and temporal locations. The plot displays the normalized frequency f/BPF (where BPF denotes the blade-passing frequency) on the vertical axis, and the normalized time t/T (with T being the blade-passing period) on the horizontal axis. Two key features emerge from this representation: (i) a distributed energy content across multiple frequency levels, with pronounced peaks located around $f/BPF \approx 1$, confirming the strong influence of dynamics directly tied to the blade-passing event; and (ii) a clear temporal and periodic modulation of these energetic contributions as a function of t/T, consistent with the four blade-passing periods characterizing the LES dataset.

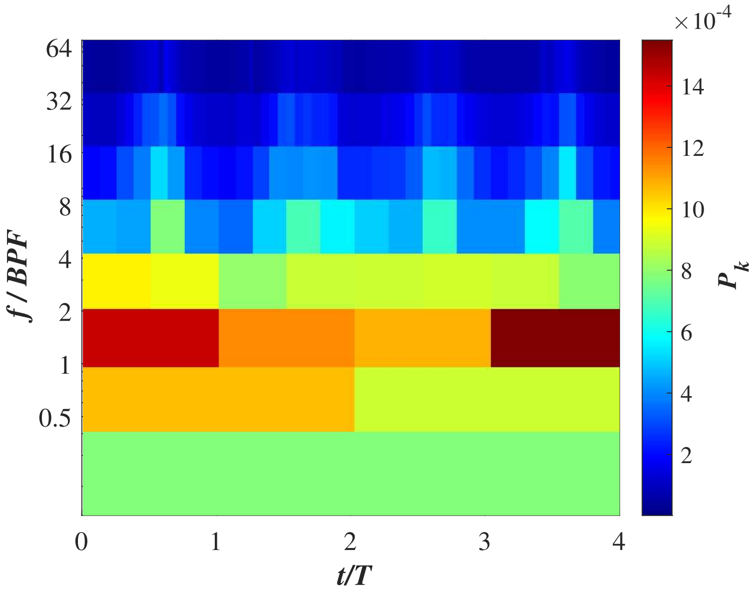

Figure 8.9: Spectral and temporal distributions of P_k obtained from the wavelet-based multi-resolution analysis

These observations, common to all six profiles tested, highlight the need to rigorously isolate coherent spectral components and focus on phase-locked statistics. To this end, SPOD and PPOD are employed, respectively. The insights gained from the application of these techniques to the selected exemplary profile are reported in the following sections.

8.2.1 SPOD results

The results presented in this section were obtained using $N_b = 128$ temporal blocks in the SPOD analysis. To provide a first qualitative comparison between the multi-scale analysis based on the Discrete Wavelet Transform and the modal decomposition performed through SPOD, Figure 8.10 reports the distribution of the turbulent kinetic energy production P_k as a function of the normalized frequency f/BPF, as obtained from both techniques. In the case of the DWT, the representation is obtained by integrating the two-dimensional map of P_k over the temporal dimension, whereas the SPOD spectrum results from the summation of the contribution to total P_k from the resulting $n_{DFT} = N_t/N_b$ modes for each frequency level. Besides small differences in terms of distribution of P_k over the lowest frequency bands, which can attributed to the different spectral resolution of the two approaches, both DWT and SPOD assign the largest share of P_k to the frequency corresponding to the blade-passing event. The SPOD spectrum exhibits a sharp and well-localized peak at $f/BPF = 1$, whereas the DWT distribution shows a broader spectral localization of energy around the same frequency level, with part of the content shifted toward slightly lower frequencies.

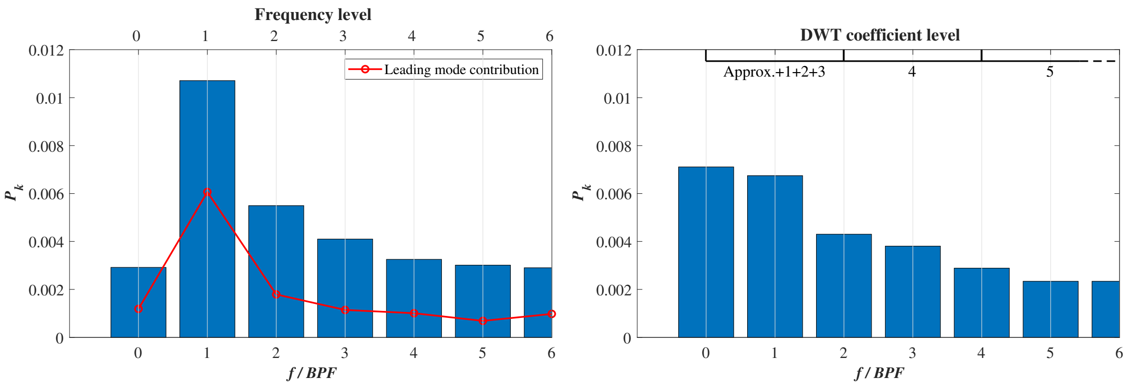

(a) SPOD

(b) DWT

Figure 8.10: Comparison of P_k modal contribution as a function of the normalized frequency f/BPF for (a) SPOD and (b) DWT.

To further localize the energetic contribution within the computational domain, Figure 8.11 decomposes the SPOD bar-plot into subregions. At the blade-passing frequency, the largest share of P_k originates from the passage (green sector), followed (with marginal difference) by the contribution of the decelerated suction-side boundary layer (yellow sector). The spatial structure of the leading SPOD mode at this frequency, reported in Figure 8.12a, confirms that the dynamics is governed by a large-scale coherent structure convected along the passage. From a statistical viewpoint, this mode is equivalent to the first POD mode (see Fig. 8.8a), whose dominant contribution to P_k was also concentrated in the same two regions. This indicates that both the POD and SPOD representations reflect the same underlying mechanism, namely the convection of the incoming wake within the inter-blade channel.

Although the energetic content progressively decreases at frequencies higher than the blade-passing peak, the distribution of P_k within the subregions reveals an inversion of the trend starting from approximately $f/BPF = 6$. At these frequencies, the relative contribution associated with the passage decreases, while the fraction attributed to the decelerated suction-side boundary layer increases. This behavior is consistent with the trend previously identified through classical POD, where

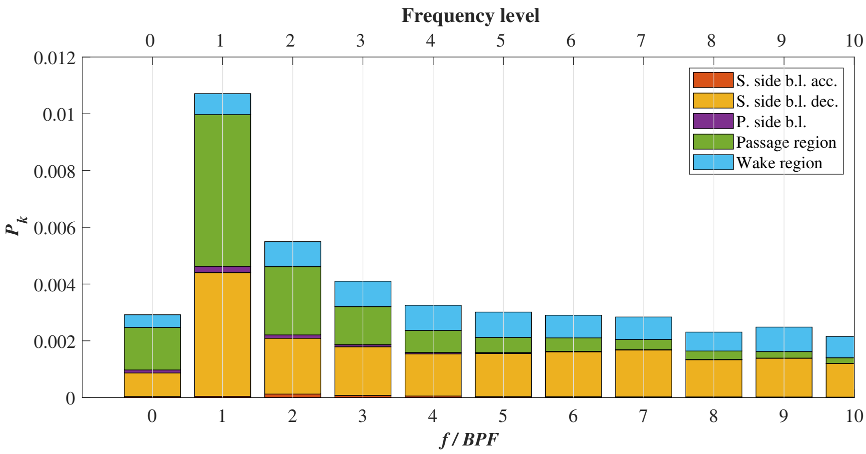

Figure 8.11: SPOD spectral distribution of P_k for the different subregions of the computational domain.

the large-scale wake structures are progressively broken down into smaller coherent features within the boundary layer.

Focusing on the dominant mode at this frequency level, depicted in Figure 8.12b, the SPOD representation highlights coherent patterns associated with the evolution of the suction-side boundary layer, analogous to those observed in high-index POD modes (see Fig. 8.8b). However, in this case, the SPOD mode also reveals the coexistence of coherent structures at the same frequency within the passage, located immediately upstream of the onset of boundary-layer development.

Furthermore, a complementary insight can be gained by examining the two-dimensional DWT map. At the frequencies associated with the inversion of the trend observed in the SPOD bar plot (starting from $f/BPF = 6$), the modulation of P_k becomes clearly periodic with respect to time, reaching a local maximum once per blade-passing cycle. This behavior confirms that the redistribution of P_k between the passage region and the decelerated suction-side boundary layer is closely linked to a phase-locked interaction between the incoming wake dynamics and the boundary-layer evolution. To further clarify the mechanisms driving this periodic coupling, a phase-conditioned analysis is required. In this respect, the PPOD framework provides a natural extension of the previous results, enabling a statistical characterization of how coherent structures evolve across the phase of an individual blade-passing event.

8.2.2 PPOD results

The results presented in this section were obtained using $N_{bin} = 10$ phase bins for the PPOD analysis, ensuring a statistically converged modal decomposition. To establish a direct comparison between PPOD and the wavelet-based analysis, the

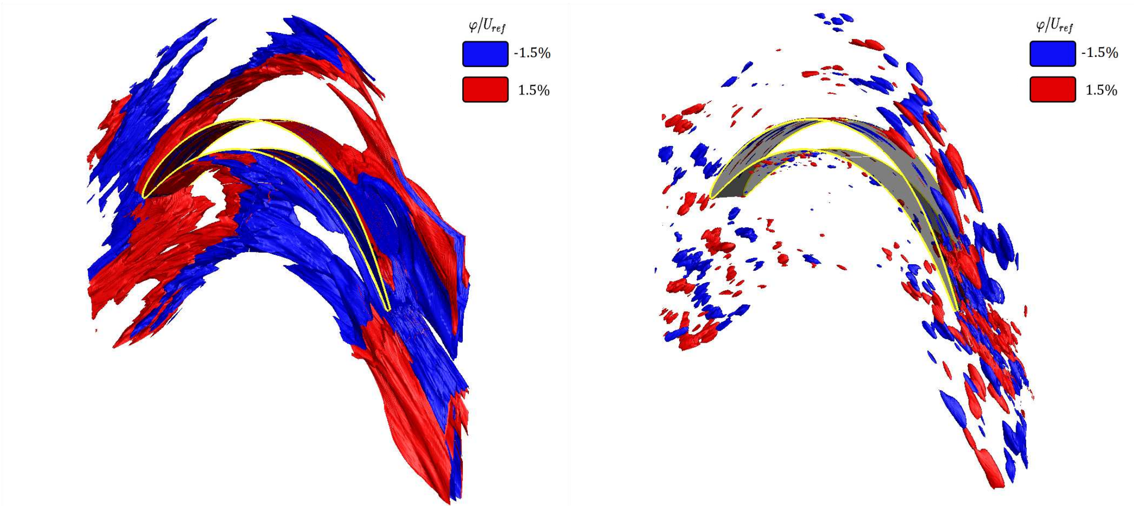

(a) Mode 1, frequency 1 (b) Mode 1, frequency 6

Figure 8.12: Isosurfaces of SPOD modes of the stream-parallel velocity for the analyzed profile.

two-dimensional DWT map of P_k was integrated along the frequency axis and evaluated at fixed phase positions. The resulting phase-resolved distribution was obtained by summing the contributions associated with the four blade-passing periods contained in the dataset, yielding an equivalent representation of a single, phase-normalized cycle. This allows for a one-to-one comparison with the PPOD results, which inherently provide phase-locked statistics over a single blade-passing event. The phase-resolved bar plots reported in Fig. 8.13 show good agreement between the two approaches. Both the DWT-based reconstruction and the PPOD analysis reveal a strong variation of P_k over the blade-passing phase, with pronounced peaks and minima occurring within a single cycle.

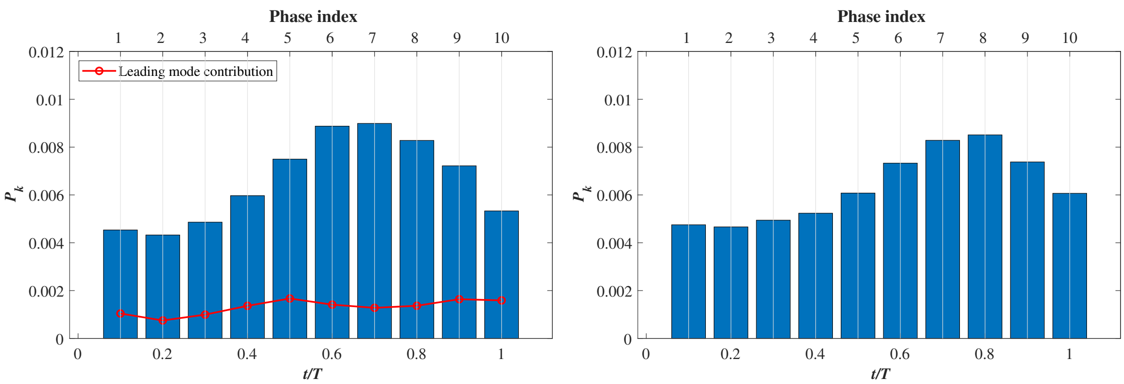

(a) PPOD (b) DWT

Figure 8.13: Comparison of P_k modal contribution as a function of t/T for (a) PPOD and (b) DWT.

As described in Chapter 5, the application of the PPOD procedure results in a modal decomposition of the flow dynamics for a given position of the wake within the inter-blade channel. For instance, Figure 8.14a-d shows the three-dimensional isosurfaces of the profile's first PPOD mode for the stream-parallel velocity component for phase indexes from 5 to 8, capturing the migration of the upstream wake along the blade suction side. As can be observed, the first mode of each phase provides a filtered (or phase-averaged) view of the flow field for each relative position of the wake (the wake centerline is schematically represented by a cyan dashed line). The phases shown in Fig. 8.14 correspond to those exhibiting the highest turbulent kinetic energy production, as indicated by the phase-resolved distribution reported in Fig. 8.13a. In these phases, the wake centerline crosses the aft region of the suction side, where the local adverse pressure gradient favors the amplification of unsteady shear-layer dynamics.

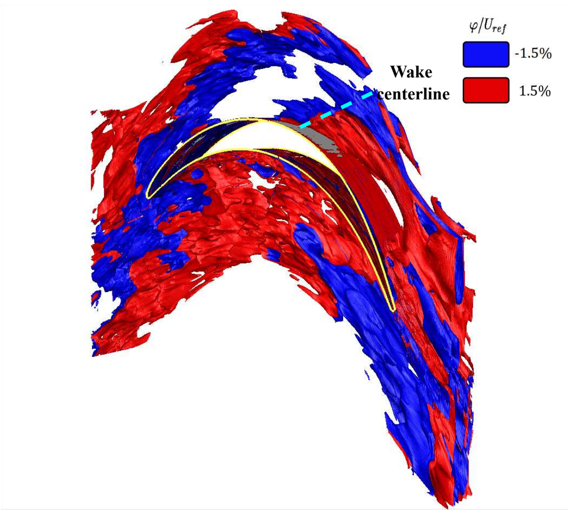

(a) Mode 1, phase 5

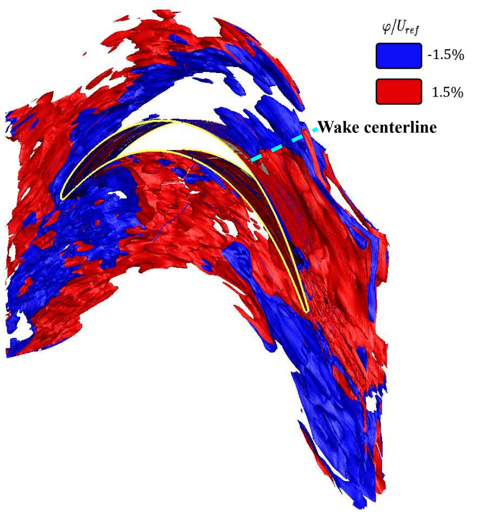

(b) Mode 1, phase 6

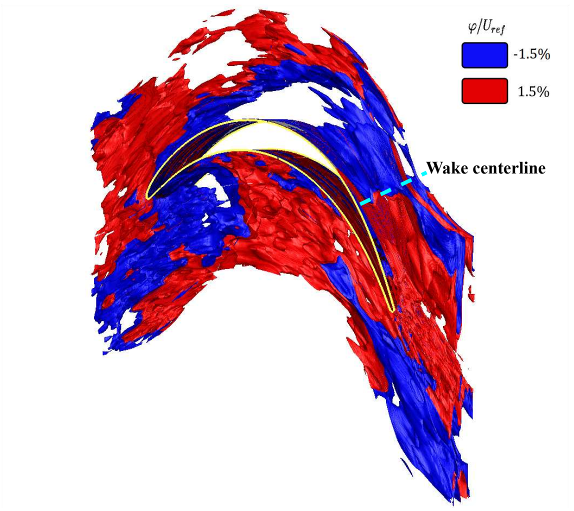

(c) Mode 1, phase 7

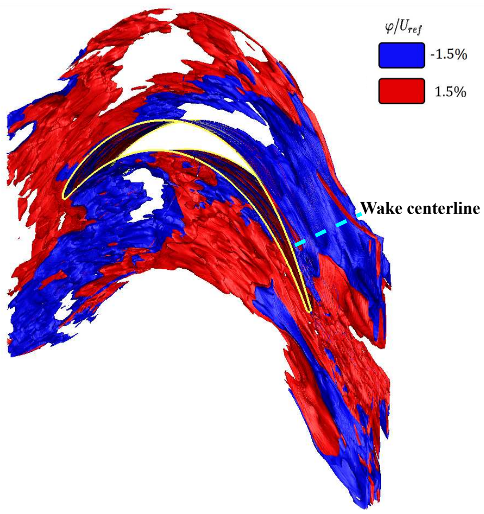

(d) Mode 1, phase 8

Figure 8.14: Isosurfaces of PPOD modes of the stream-parallel velocity for the analyzed profile.

By applying the spatial filtering strategy over the selected subregions of the computational domain (see Fig. 8.15), a clear redistribution of P_k among the different flow regions can be observed. Specifically, the phase characterized by the minimum turbulence-related loss is phase 2, which corresponds to the wake located close to the leading edge of the blade (not shown here for brevity). Here, turbulent production is primarily concentrated in the passage region (green bar). As anticipated, with the wake centerline moving towards the trailing edge of the profile, losses increase and become dominated by events in the decelerating region of the blade suction side and within the profile's wake, with the passage region no longer being the area characterized by the maximum turbulent production. During the phases of maximum wake–blade interaction, the contribution associated with the decelerated suction-side boundary layer (yellow bar) increases sharply and rapidly becomes dominant when compared to that of all other regions. The phase index 7 represents the condition of maximum total P_k, characterized by a simultaneous minimum contribution from the passage region and a peak contribution from the decelerated suction-side boundary layer.

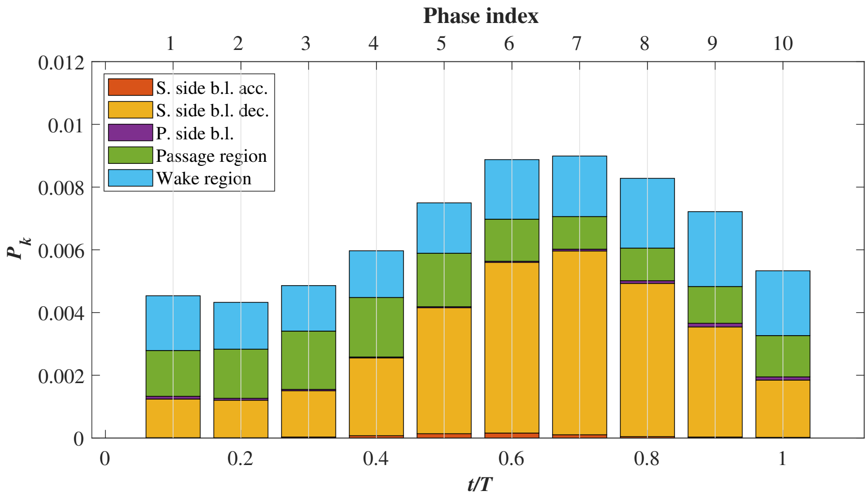

Figure 8.15: PPOD temporal distribution of P_k for the different subregions of the computational domain.

These observations are reflected in the plots of P_k as a function of the PPOD mode index at a fixed phase. Figure 8.16a shows this trend for the phase characterized by the lowest turbulent production (phase 2) for the analyzed profiles. As can be observed, the most significant contribution to turbulent kinetic energy production is primarily associated with the low-order modes (i.e., for mode indices smaller than 10), which predominantly act in the passage region (green curve), with a minor additional contribution in the decelerating part of the boundary layer (yellow curve). As previously noted, at this phase, the upstream wake is located near the leading edge of the cascade, thus the momentum impingement effects on the rear suction side are limited. The boundary layer is stabilized by the acceleration of the flow imposed by the blade loading, which keeps its laminar characteristics as the wake passes. This results in the boundary layer losses being predominantly of a viscous nature, as highlighted by POD results (red curve of Fig. 8.7), while turbulence-related events are observed essentially outside of the boundary layer, within the bulk flow region.

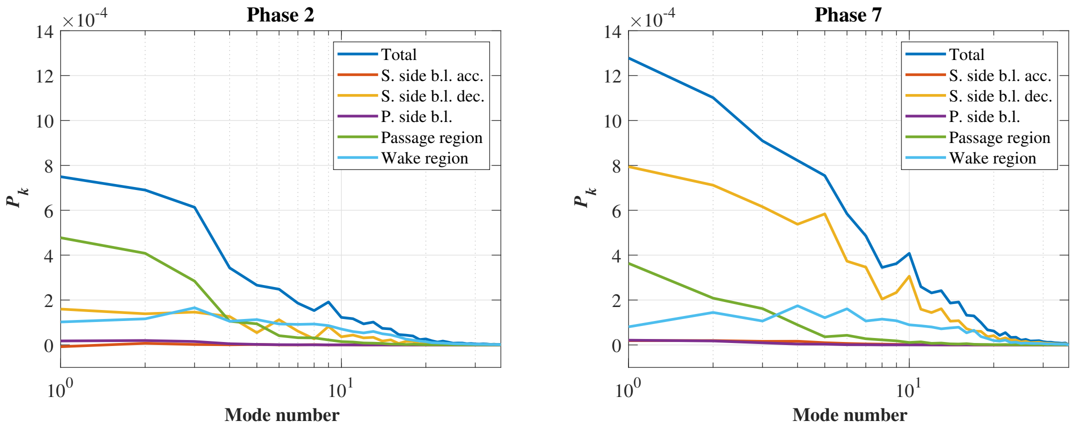

(a) Phase 2 (b) Phase 7

Figure 8.16: Plots of P_k as a function of (a) phase 2 and (b) phase 7 mode index. Blue curves represent the total mode contribution from the five sub-regions.

From the P_k plot obtained for phase 7, reported in Figure 8.16b, a new scenario emerges. Maximum turbulence production occurs on the rear suction side (yellow curve), where the low-order modes associated with the upstream wake continue to dominate the loss generation mechanisms. However, during this phase, higher-order modes contribute substantially to the loss generation process. To better understand the different loss mechanisms involved, the contour plots in Figure 8.17 show the $\hat{P}_k$ levels for two PPOD modes observed at phase 7, in a meridional section (located at midspan) of the rear part of the computational domain. This quantity represents a normalization of the production of TKE defined as:

$$\hat{P}_k = \frac{\rho P_k}{0.5 U_{ref}^2 \dot{m} / \mathcal{V}} \tag{8.1}$$

The contours are overlaid with streamlines computed from the low-order reconstruction associated with each mode and centered around the peak suction position. As illustrated in Fig. 8.17a, streamlines of mode 1 highlight the formation of the so-called "negative-jet" structure. The momentum transfer associated with the wake velocity defect generates a jet which, around the wake centerline, directs toward the blade suction side. This jet splits into two branches, forming two distinct large-scale counter-rotating vortices at the leading and trailing edges of the wake, as known from previous literature results. In the figure, these structures are referred to as "Q1" and "Q2" vortices. The negative jet is known to play a key role in the mechanism of wake-induced boundary layer transition. As the wake convects over the boundary layer separation point, the wall-normal component of the negative jet distorts the shear layer, which is inherently unstable. This interaction triggers an inviscid Kelvin-Helmholtz roll-up, where the resulting vortex moves at half the freestream

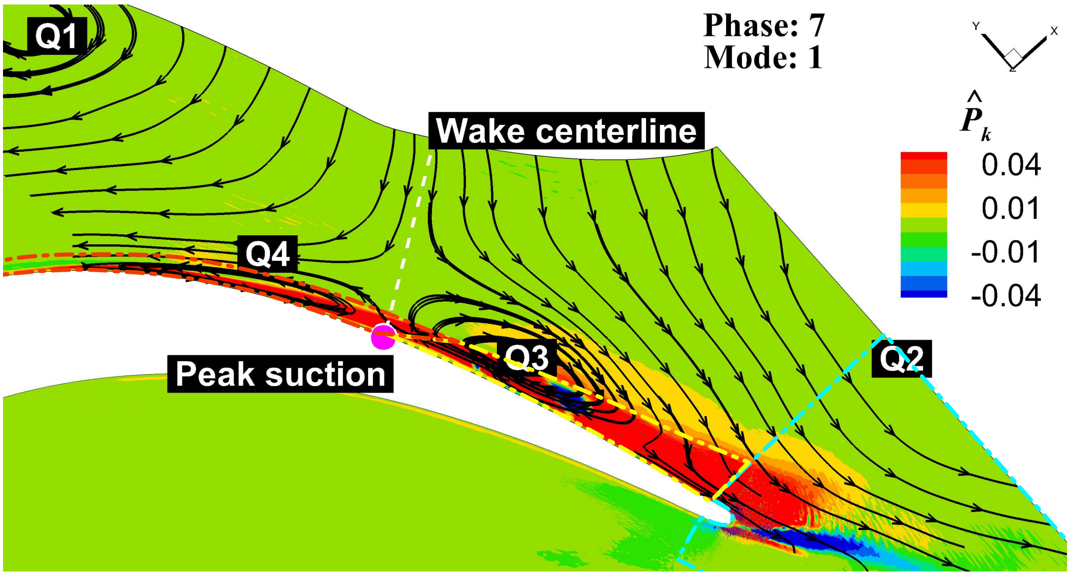

(a)

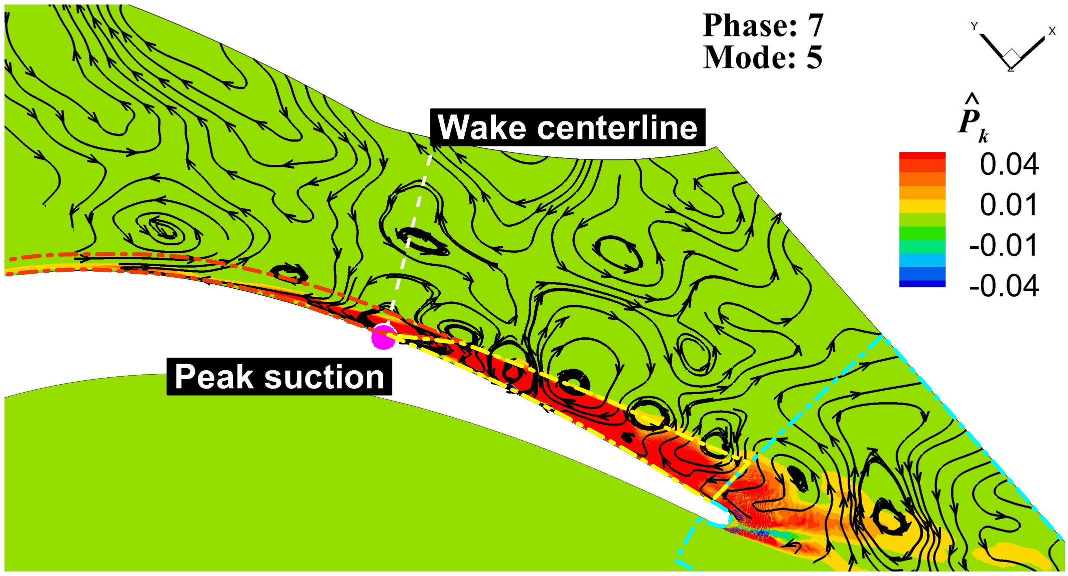

(b)

Figure 8.17: Contour plots of the analyzed profile's mode 1 (a) and 5 (b) of the dimensionless TKE production at phase 7, overlaid to the corresponding mode streamlines. The extent of the integration sub-regions is marked by dashed lines adopting the color scheme as in Fig. 7.9.

velocity. The wake, traveling at the freestream speed, advances ahead of the roll-up and further disturbs the separated shear layer downstream, leading to the formation of additional roll-up vortices that quickly break down into turbulence. In addition to these vortices, the figures highlight two other counter-rotating vortical structures, generated upstream and downstream of the wake centerline within the boundary layer. These are labeled as "Q3" and "Q4" vortices, and their formation is linked to the negative jet of the incoming wake. The modal representation of the flow obtained from the PPOD procedure enables the direct identification of the phase of maximum loss as the one characterized by the greatest amplification of these large-scale vortical structures.

The analysis of PPOD mode 1 for the phase with the highest turbulent production (phase 7) allows us to pinpoint the position of the wake centerline, which is located almost exactly over the peak suction position. This occurs not only for the profile considered in this analysis, but also for all the other geometries in the dataset. In each case, the position of the identified vortical structures (in particular that of the Q3 vortex) is determined by the specific blade-loading distribution: at the phase of maximum turbulence-related losses, profiles with a more front-loaded distribution (i.e., characterized by a peak suction located closer to the leading edge) exhibit an earlier interaction between the incoming wake and the boundary layer, and vice versa. This results in differences in the flow dynamics associated with the specific interaction mechanism, primarily due to the different spatial location of the Q3 vortex.

For the profile considered here, the blade-loading configuration causes this vortex (which is strongly amplified at this phase) to experience a pronounced deceleration, but at a curvilinear coordinate sufficiently upstream of the blade trailing edge. As shown in Figure 8.17b, which reports the contour plot of $\hat{P}_k$ for a higher-order PPOD mode at this phase (mode 5), this leads to the onset of characteristic structures originating into the rear suction side. They are medium-scale roll-up vortices with axes aligned with the spanwise direction, which can be typically observed in the case of bypass-like transition mechanisms. This is due to the fact that the maximum spatial extent of the Q3 vortex is limited, above by the Q2 vortex and below by the blade wall. Under these conditions, the evolution of this vortex is largely confined within the boundary layer region.

For different blade loading distributions, distinct flow scenarios arise. For instance, in the case of more aft-loaded profiles, where at the phase of maximum loss the Q3 vortex is located at the very limit of the blade trailing edge, the lack of a physical constraint to its spatial extent results in its free evolution within the wake region. Higher-order modes, in fact, capture Von Kármán vortex street-like structures, likely originating from the interaction between the Q3 vortex and the low-momentum region downstream of the trailing edge. For these geometries, the POD analysis shows that the wake region is characterized by the main sources of loss not directly related to the large-scale structures that define the upstream wake, unlike what happens for more front-loaded profiles. This, ultimately, is the reason behind different geometries having different performances in terms of global inlet-to-outlet loss coefficient.

Chapter 9

Predictive Model for Loss Characterization

The development of advanced tools for the prediction and characterization of total pressure losses is a natural progression of the analysis presented up to this point. After investigating the physical mechanisms responsible for loss generation and exploring their localization through modal decomposition techniques, the next step is to build a framework capable of translating such knowledge into a predictive model. The objective is to move from a purely diagnostic approach—where flow structures and their contributions to losses are identified a posteriori—to a more prescriptive one, where the essential relationships between blade geometry and loss mechanisms can be exploited to anticipate performance trends. In this context, machine learning has emerged as a promising approach to the analysis and optimization of turbomachinery components, enabling the construction of surrogate models capable of capturing complex correlations between geometry and performance with reduced computational effort. Among the available methods, autoencoders—and, in particular, their convolutional (CNN) variants—have been widely adopted for compressing high-dimensional flow fields and for feature extraction [110, 111, 112, 113]. CNN-based encoders have been successfully employed to identify coherent structures in turbulent flows, reconstruct pressure and velocity fields, or perform reduced-order modeling in complex aerodynamic configurations. In this environment, however, the model learns nonlinear and implicit basis functions derived from the training data, which makes the encoded space difficult to interpret from a physical standpoint. The resulting latent variables represent hidden statistical correlations rather than clearly defined flow or geometric features, thereby hindering the physical interpretability of the relationships established by the model.

From a mathematical perspective, CNNs can be viewed as a local projection of the available dataset onto convolutional kernels that act as predefined filters, implicitly defining a latent space through nonlinear transformations. Alternatively, introducing data-driven basis functions, obtained directly from the statistics of the dataset, can shift the problem toward a more interpretable and physically consistent latent representation. A linear and physically consistent alternative to nonlinear autoencoders is provided by POD. Specifically, POD can be used not only for the decomposition of turbulent flow fields into spatial and temporal modes, identifying the dominant coherent structures that contribute to loss production, but it has also been successfully applied as a multi-purpose dimensionality-reduction tool to extract the main features of a dataset. As reported in [114, 115, 116], an orthogonal basis, such as the one spanned by POD modes, can be effectively used to efficiently parameterize the geometry of blade profiles. The use of a mathematically well-defined basis allows for a direct and interpretable mapping between the geometric variations of a blade and the corresponding flow responses in terms of spatial loss distributions and dynamical behavior.

Building on this, the present section introduces a predictive framework based on the combination of reduced-order modeling and machine learning strategies. By leveraging the properties of POD, the idea is to extract compact, low-dimensional representations of both the blade geometry and the associated loss distribution in space and dynamics, and then establish a functional link between them. In this way, the complexity of high-fidelity simulations can be distilled into a small set of parameters, enabling a faster and more efficient exploration of the design space.

The approach is particularly suited to scenarios where only a limited number of high-fidelity datasets are available, a condition often encountered in practical design campaigns due to the cost of advanced numerical simulations. Once trained, the predictive model can then be applied to interpolate or extrapolate the behavior of new geometries, providing valuable insights into loss localization and intensity across different regions of the computational domain.

9.1 Geometry parametrization

A turbine cascade can be described by a large number of parameters, such as inlet and outlet flow angles, thickness distribution, camber line curvature, and pitch-to-chord ratio. Handling such a wide set of variables makes the search for functional relationships between geometry and losses a highly complex, multidimensional problem, especially when the available dataset is limited. From the results obtained through the diagnostic techniques presented in the previous sections, it clearly emerges that the performance differences observed across the geometries in the dataset can be primarily attributed to variations in their blade loading distributions. The shape and streamwise position of the loading peak were shown to govern the location, intensity, and interaction of the dominant vortical structures, ultimately controlling the mechanisms responsible for loss generation.

To further substantiate this point, a detailed comparison is carried out between two representative profiles of the dataset, both classified as aft-loaded in the blade loading overview presented in Table 6.1, but characterized by markedly different Zweifel number levels. As can be observed in Figure 9.1, the first profile (profile "1", used as reference for the analysis of Chapter 8.2) belongs to the high-Zweifel group and exhibits a loading peak located at approximately 74% of the axial chord, as previously discussed. The second profile (profile "2"), despite being characterized by a low Zweifel number, features a loading peak shifted further downstream, at about 83% of the axial chord, and can therefore be regarded as strongly aft-loaded. The upstream shift of the point of maximum flow velocity for profile "1" allows for a more gradual deceleration in the rear portion of the suction side: despite a higher overall diffusion, a 10% reduction in local diffusion rate with respect to profile "2" is observed, with the red curve having a greater maximum negative slope than the blue curve near the trailing edge. Furthermore, the higher loading of profile "1" allows for a greater flow acceleration in the front part of the suction side.

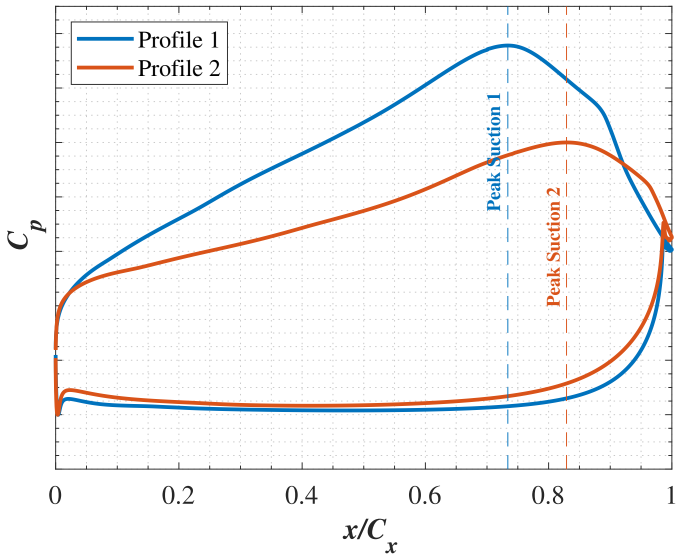

Figure 9.1: Comparison of the blade loading distributions for two LPT profiles from the dataset.

The impact of these differences in loading distribution can be highlighted by comparing the POD-based loss decompositions reported in Fig. 9.2, where the modal and spatial distribution of the losses is shown for the two profiles, partitioned according to the previously defined sub-regions of the computational domain.

The cumulative curves of total dissipation for profile 1 highlight a clear improvement in overall performance compared to profile 2. The total loss, represented by the blue curve accounting for the entire set of modes, amounts to approximately 87% of the value observed for profile 2. This behavior can be directly related to the blade loading characteristics of the two profiles, which results in two significantly different distributions of losses in space and dynamics.

For profile 1, the stronger acceleration imposed in the front portion of the suction side effectively suppresses turbulence-related losses associated with incoming wakes and large-scale structures, resulting in a reduced contribution from wake migration into the passage region (green curve). Viscous dissipation in the accelerating suction side boundary layer (red curve) is also lower, owing to the upstream shift of the peak suction position, which limits the extent of the accelerated front section. Conversely, the more moderate but spatially extended acceleration characterizing profile 2 in the

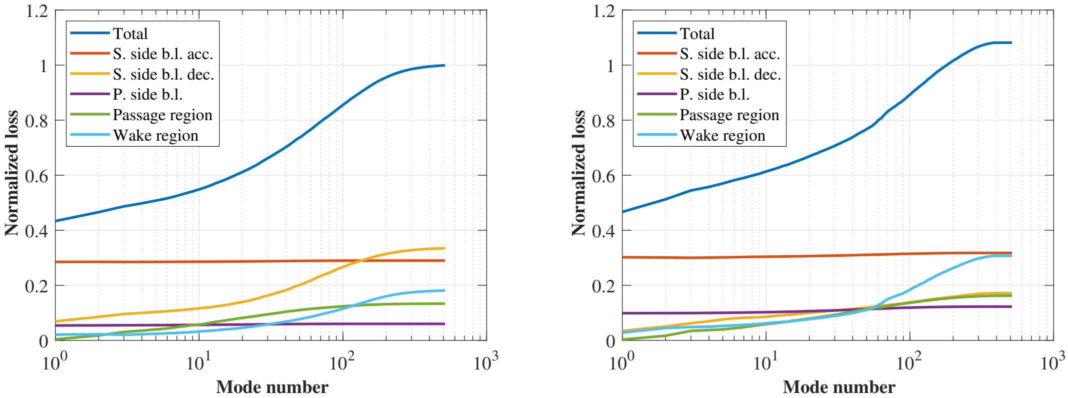

(a) (b)

Figure 9.2: Cumulative contribution of POD modes to total losses, split by spatial regions of interest, for (a) profile "1" and (b) profile "2".

early part of the vane passage allows passage losses to play a significant role in the overall entropy production. As shown by Simoni et al. [117], insufficient flow acceleration in the initial portion of the blade does not suppress turbulence generation associated with wake bowing and tilting processes, thereby increasing total losses. The dominant contribution to total losses for profile 1 arises from the decelerating part of the suction side (yellow curve), where a bypass-like boundary layer transition occur, as described in Chapter 8.2. In contrast, for profile 2, considerable losses are observed in the wake region (cyan curve) downstream of the trailing edge. This is due to the the downstream shift of the maximum velocity location of profile 2, which, despite the limited overall diffusion associated with the lower blade loading, results in a higher local diffusion rate near the trailing edge and promotes intense vortex shedding and a related break-up process in the wake. This ultimately leads to an increase in total pressure losses.

Building on this, blade loading (C_p) was selected as the key geometric indicator for the construction of the predictive model. Beyond its proven relevance in distinguishing the loss-generation mechanisms of the analyzed profiles, blade loading represents a potential-flow characteristic and is therefore largely independent of flow conditions such as the Reynolds number at the high-Re regime of the present LES data. Under these conditions, where large-scale boundary-layer separations on the decelerated suction side can be neglected, the C_p effectively captures the intrinsic geometric characteristics of each blade. This makes it an ideal descriptive parameter for establishing a functional relationship between geometry and loss behavior within the proposed predictive framework.

A blade loading distribution can be described by a vector containing pairs of coordinate points $[x/C_x, C_p]$. In this study, a common x-coordinate array was defined, and each blade loading distribution was represented by $M = 1600$ C_p-coordinate values. The blade loading coordinates $C_p(x/C_x, p_i)$ of each blade (p_i being the representative index of the i-th profile in the dataset) were stored in the column space of the matrix $\boldsymbol{B}$. Figure 9.3 graphically illustrates the results of the discretization of the profile's blade loading for three exemplary geometries from the dataset.

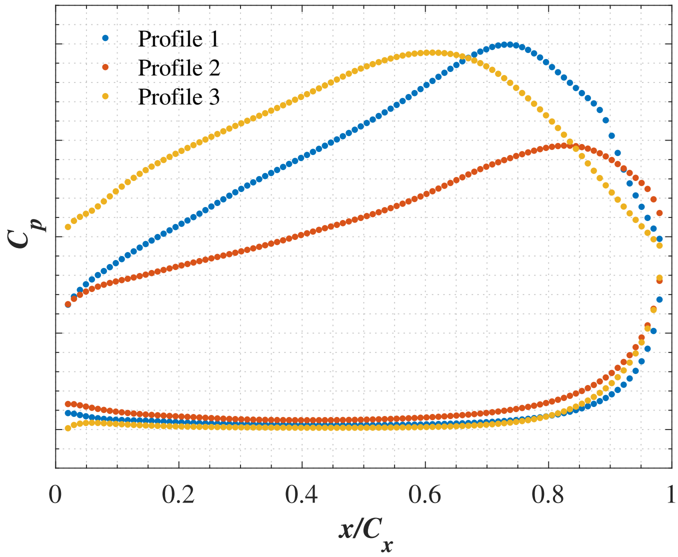

Figure 9.3: Blade loading distributions of three example profiles after discretization.

From a practical standpoint, the parametrization of the blade loading distributions was not carried out through classical physics-based shape descriptors (such as the Zweifel number, the diffusion factor, or the streamwise location and magnitude of the loading peak), but rather through a purely data-driven procedure that exploits the capabilities of the POD. The snapshot POD method was applied to the centered blade loading matrix $\boldsymbol{B'}$:

$$\boldsymbol{B'} = \boldsymbol{B} - \overline{\boldsymbol{B}} \tag{9.1}$$

where $\overline{\boldsymbol{B}}$ is the mean of the matrix $\boldsymbol{B}$. The POD generates two complete orthogonal bases, namely the spatial modes (φ_B^k) and the corresponding eigenvectors ($\chi_{B,k}$), such that:

$$\boldsymbol{B'}(x, p_i) = \sum_k \chi_{B,k}(p_i)\boldsymbol{\varphi}_B^k(x) \tag{9.2}$$

By truncating the blade loading eigenvector matrix $\boldsymbol{\chi}_B$ at an index smaller than k, an approximated reconstruction of the matrix $\boldsymbol{B'}$ can be obtained. The accuracy of this reduced representation was assessed by computing the root mean squared error (RMSE) between the original interpolated C_p and those reconstructed using a limited number of modes. As shown in Figure 9.4, truncating to the first three POD modes was sufficient to keep the reconstruction RMSE below 0.5% of the maximum dataset variation for all profiles. On this basis, $k = n_B = 3$ coefficients $\chi_{B,(1,..,3)}(p_i)$ were selected for each blade loading distribution as the descriptive parameters of the corresponding cascade. The concept behind this kind of parametrization is schematically shown in Figure 9.5.

The reduced-order representation of the blade loading distributions through the first three POD coefficients provides a compact yet meaningful description of the design space explored in this study.

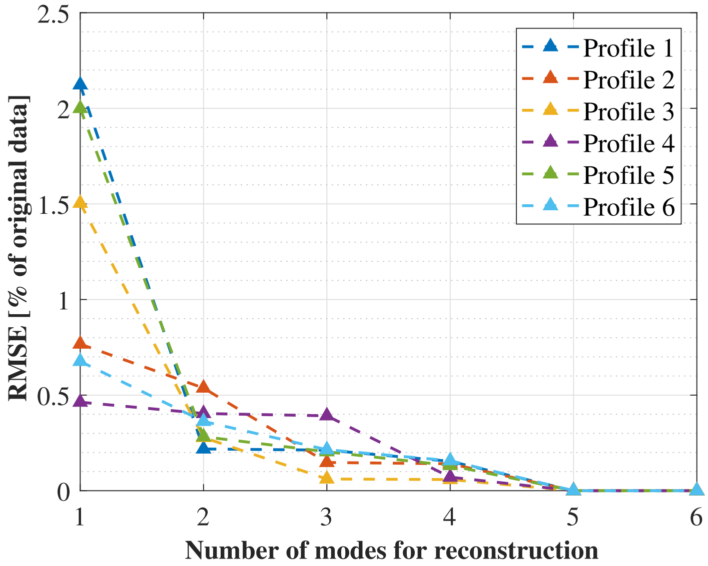

Figure 9.4: Reconstruction error (RMSE) of the blade loading as a function of the number of retained POD modes. The 0.5% threshold is consistently satisfied with $n_B = 3$, ensuring an accurate reduced-order parametrization.

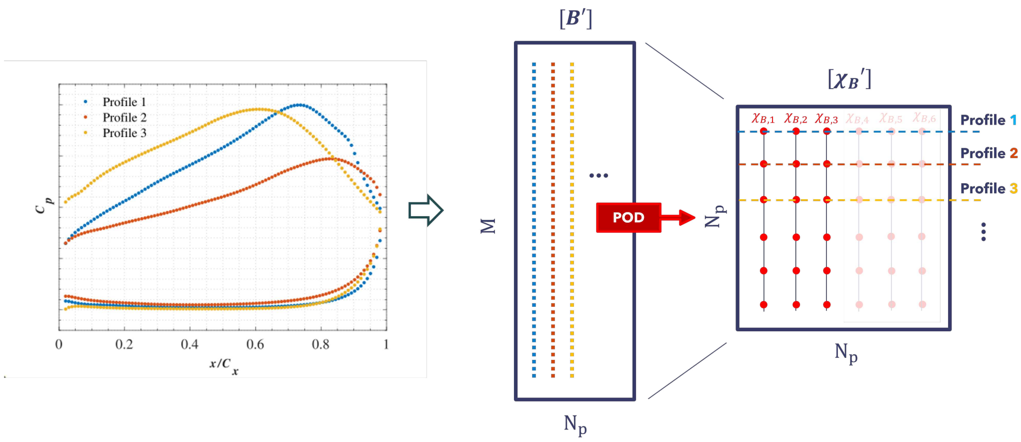

Figure 9.5: Conceptual illustration of the compression process of the high-dimensional blade loading matrix.

To provide a more intuitive interpretation of the physical meaning of each coefficient, Figure 9.6 shows the effect of varying $\chi_{B,1}$, $\chi_{B,2}$, and $\chi_{B,3}$ individually on a representative profile from the design space. For each coefficient, three blade loading distributions are reconstructed by setting the selected coefficient to its minimum, mean, and maximum value, while keeping the remaining two fixed at their mean.

As shown in Figures 9.6a, the first mode $\chi_{B,1}$ primarily affects the overall blade loading and the location of the peak suction point. The second coefficient $\chi_{B,2}$ (Fig. 9.6b) mainly affects the C_p maximum value at the peak suction point, influencing the magnitude of the deceleration in the rear section of the profile's suction side. The third coefficient $\chi_{B,3}$ (Fig. 9.6c) introduces higher-order and less relevant variations, mainly influencing the suction-side local C_p.

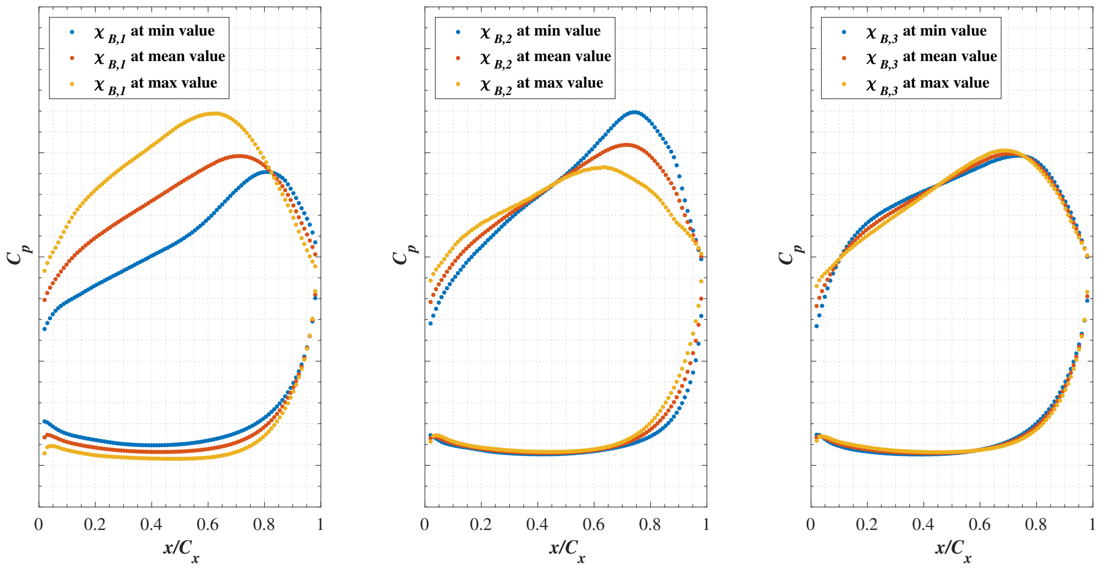

(a) Effect of $\chi_{B,1}$ variation. (b) Effect of $\chi_{B,2}$ variation. (c) Effect of $\chi_{B,3}$ variation.

Figure 9.6: Effect of the variation of each of the first three blade loading coefficients on the reconstructed profile. In each plot, the selected coefficient varies between its minimum, mean, and maximum value, while the other two are fixed at their mean values.

9.2 Loss distribution parametrization

The loss distributions obtained from the LES dataset using the advanced post-processing tool discussed in this work can be further reduced to a compact set of descriptive parameters through a procedure entirely analogous to that applied to the blade loading distributions. In this case, instead of handling loading-related quantities, the focus is placed on the spatially resolved contributions to the total dissipation, previously attributed to distinct coherent flow structures and cascade sub-regions as a result of the application of the HPDA procedure. For the sake of result compactness and interpretability, the results obtained within the classical POD framework are considered here. As will be clarified in the following discussion, results derived from other modal decomposition–based analyses could also be adopted as objective functions of the model.

The normalized total dissipation values associated with each coherent flow structure (i.e., each POD spatial mode into which the turbulent flow has been decomposed) and each integration sub-region were collected in the columns of a matrix $\boldsymbol{L}$, from which the mean was subsequently subtracted. A schematic representation of this procedure is reported in Figure 9.7, where the organization of modal and spatial contributions into the loss matrix is illustrated.

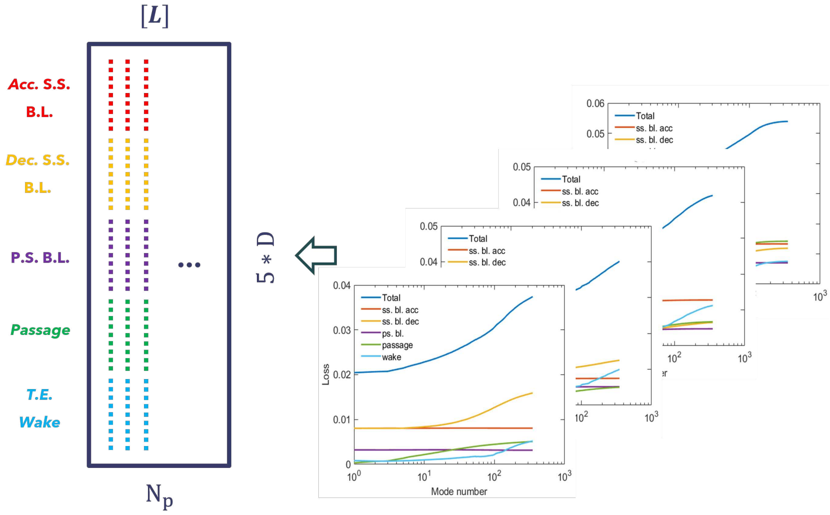

Figure 9.7: Schematic representation of the construction of the loss matrix $\boldsymbol{L}$, obtained by collecting the normalized dissipation contributions associated with each POD mode and each integration sub-region.

The resulting matrix $\boldsymbol{L}$*'* was then decomposed into orthogonal modes and associated eigenvectors, exploiting again POD properties, in the form:

$$\boldsymbol{L}'(j, p_i) = \sum_k \chi_{L,k}(p_i) \boldsymbol{\varphi}_L^k(j) \tag{9.3}$$

with j denoting the index spanning both the modal decomposition and the integration regions.

This operation effectively compresses the information contained in the original high-dimensional matrix into a reduced number of coefficients. The concept is schematically shown in Figure 9.8, where the projection of the full dataset onto a limited set of POD modes is represented.

An error metric based on the reconstruction RMSE of the loss distribution was considered for selecting the truncation order of the $\boldsymbol{\chi}_L$ matrix. The minimum number of modes required for a reconstruction with an MSE consistently below 0.5% was determined to be equal to $k = n_L = 3$. The trend of the reconstruction error with respect to the number of retained modes is reported in Figure 9.9, confirming that three parameters are sufficient to ensure an accurate representation. Thus, three coefficients $\chi_{L,(1,..,3)}(p_i)$ were selected for each profile's loss distribution as descriptive parameters. This allowed each cascade to be characterized by a limited set of descriptors, preserving the essential features of the loss production mechanisms.

The reduced-order representation of the loss distributions through the first three POD coefficients provides a compact yet physically interpretable parametrization of the dissipation mechanisms occurring within the cascade. To gain insight into the role played by each coefficient, a parametric reconstruction of the loss distribution curves was carried out by individually varying $\chi_{L,1}$, $\chi_{L,2}$, and $\chi_{L,3}$ between their

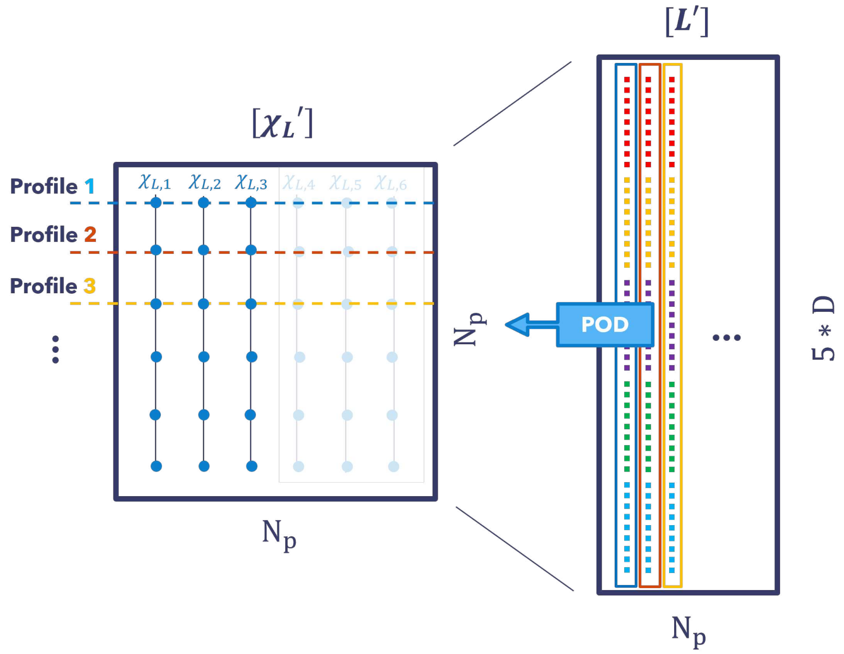

Figure 9.8: Conceptual illustration of the compression process of the high-dimensional loss matrix.

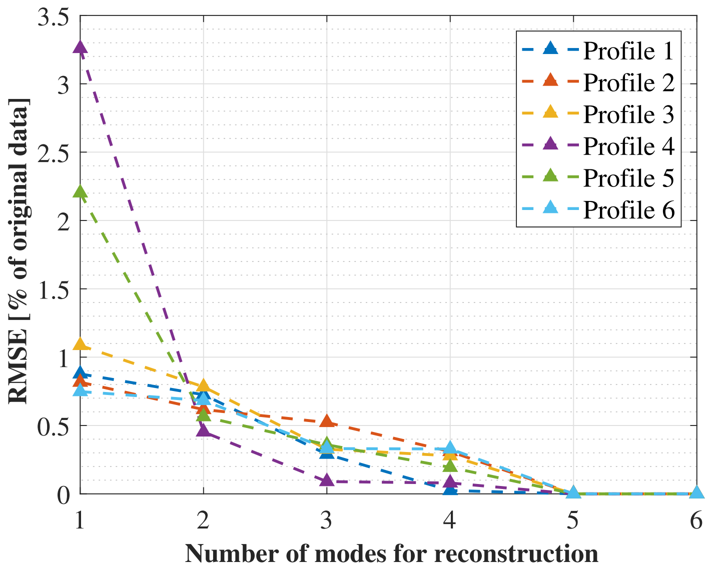

Figure 9.9: Reconstruction error (MSE) of the loss distribution curves as a function of the number of retained POD modes. The threshold of 0.5% is consistently satisfied with $n_L = 3$, ensuring an accurate reduced-order parametrization.

respective minimum, mean, and maximum values, while keeping the remaining two fixed at their mean values.

Figures 9.10–9.12 illustrate the reconstructed loss distributions obtained by varying each of the three coefficients in turn. The first coefficient $\chi_{L,1}$ (Figure 9.10) mainly acts on the viscous losses along the suction side boundary layer. An increase in $\chi_{L,1}$ leads to a marked rise in dissipation within the decelerated portion of the suction-side boundary layer, accompanied by a reduction in losses over the accelerated portion and in the passage region. The second coefficient $\chi_{L,2}$ (Figure 9.11) primarily influences the passage and wake losses. Increasing $\chi_{L,2}$ results in higher dissipation levels in both these regions, although the magnitude of the variation is smaller than that observed for $\chi_{L,1}$. Overall, this coefficient introduces a global increase in total losses across the domain. Finally, the third coefficient $\chi_{L,3}$ (Figure 9.12) affects losses within the passage, wake region, and the dynamically driven losses in the decelerated boundary layer. An increase in $\chi_{L,3}$ produces a strong rise in wake dissipation, together with a significant reduction in the boundary layer losses and a moderate decrease in the passage losses. The impact of this coefficient is comparable in magnitude to that of $\chi_{L,1}$, although it acts through different physical mechanisms.

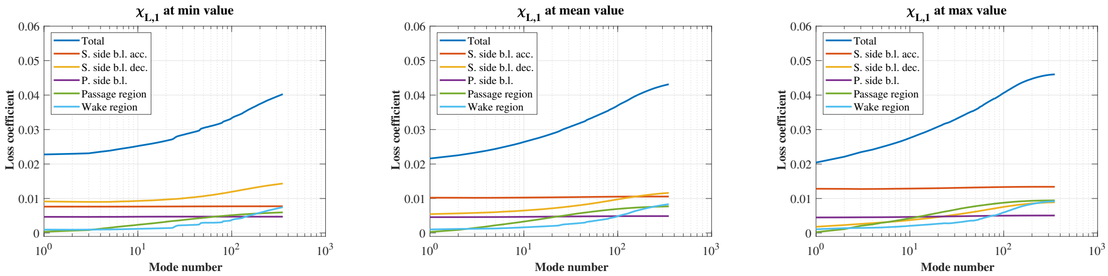

Figure 9.10: Effect of $\chi_{L,1}$ variation on loss distribution curves. The other coefficients are kept fixed at their mean values.

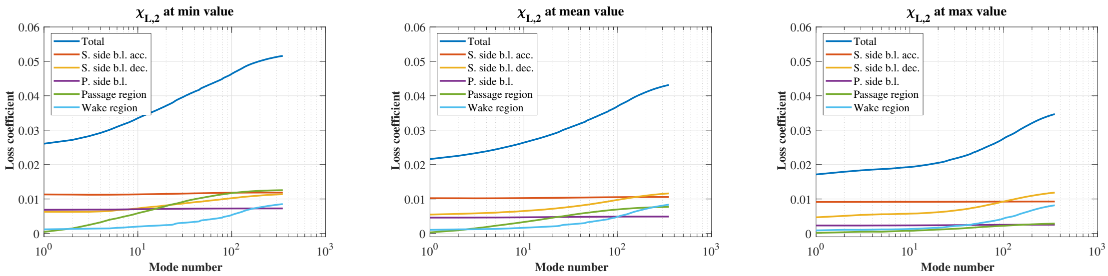

Figure 9.11: Effect of $\chi_{L,2}$ variation on loss distribution curves. The other coefficients are kept fixed at their mean values.

As anticipated, this parametrization strategy is not necessarily limited to classical POD. The same framework can be extended to alternative decomposition techniques discussed in this work, such as SPOD or PPOD, by simply adapting the construction of the loss matrix $\boldsymbol{L}$ to the specific definition of modes provided by the method. In all cases, the outcome is a compact representation of the loss dynamics, suitable for establishing predictive relationships with the corresponding blade loading parameters.

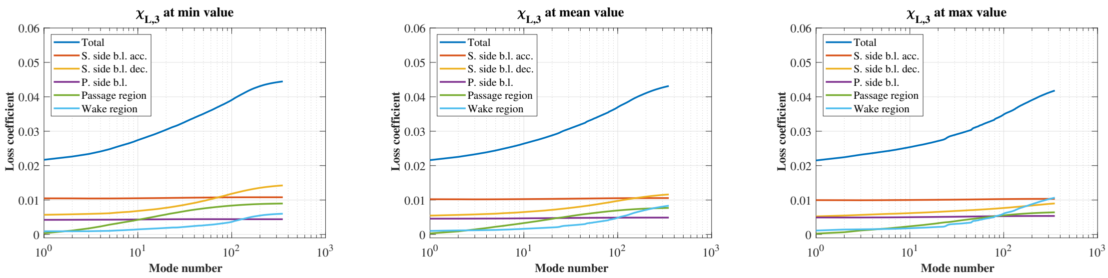

Figure 9.12: Effect of $\chi_{L,3}$ variation on loss distribution curves. The other coefficients are kept fixed at their mean values.

9.3 Tuning of the model

The reduced sets of parameters obtained through the encoding of both blade loadings and loss fields were employed as input and output datasets for the construction of a low-dimensional predictive model linking the blade C_p distribution to the results of the advanced post-processing routine. Figure 9.13 schematically summarizes the overall structure of the POD-encoder model, highlighting the connection between geometry and loss distribution parameterization.

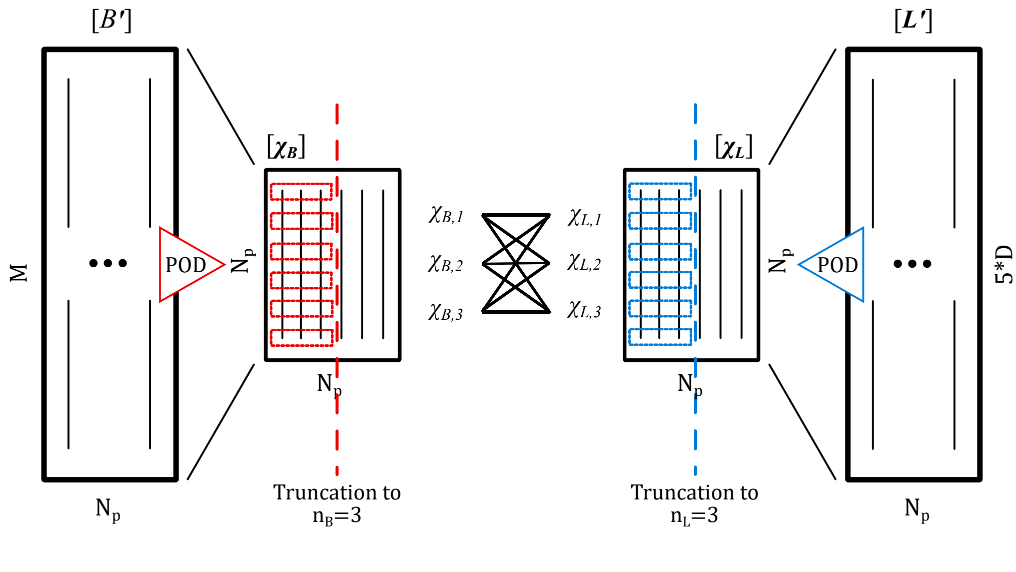

Figure 9.13: Schematic representation of the POD-encoder model.

The main objective was to define a model that remains as simple and physically interpretable as possible, consistent with the philosophy of the overall methodology here discussed. Instead of relying on complex nonlinear regressors or "black-box" neural networks, various linear models were tested by systematically evaluating all possible combinations of preset predictors. Specifically, for each loss coefficient $\chi_{L,(1,..,3)}$ (representing the objective function), models were constructed using different subsets of two to four predictors, chosen among the three first-order blade loading coefficients $\chi_{B,(1,..,3)}$, their squared values, and cross-terms up to the third order. The model selection criterion was based on a Leave-One-Out Cross-Validation (LOOCV) procedure: for each candidate combination, the mean Root Mean Square Error (RMSE) and coefficient of determination (R^2) on the validation sample were computed over all folds, and the configuration that minimized the cross-validated error was selected. This ensured that the optimal configuration of predictors was identified without compromising generalization capability. The resulting models provide a compact functional mapping between the reduced geometrical and loss spaces, ensuring that the key underlying relationships between profile geometry and different loss mechanisms (observed in the 5 different spatial sub-regions) are effectively captured while maintaining computational efficiency and avoiding unnecessary complexity.

Once the optimal configuration was identified, three separate models—one for each loss-related coefficient $\chi_{L,(1,..,3)}$—were obtained. Their performance metrics, reported in Table 9.1, show consistently high R^2 values and low RMSE.

Table 9.1: Performance metrics of the three best regression models obtained through LOOCV.

Model	**RMSE**	R^2
1 ($\chi_{L,1}$)	0.093	0.95
2 ($\chi_{L,2}$)	0.211	0.73
3 ($\chi_{L,3}$)	0.111	0.93

Figure 9.14 reports the comparison between the predicted and actual values for the three loss-related coefficients, showing a satisfactory level of agreement and confirming the model's capability to capture the main physical correlations between geometry and loss behavior across the analyzed cases.

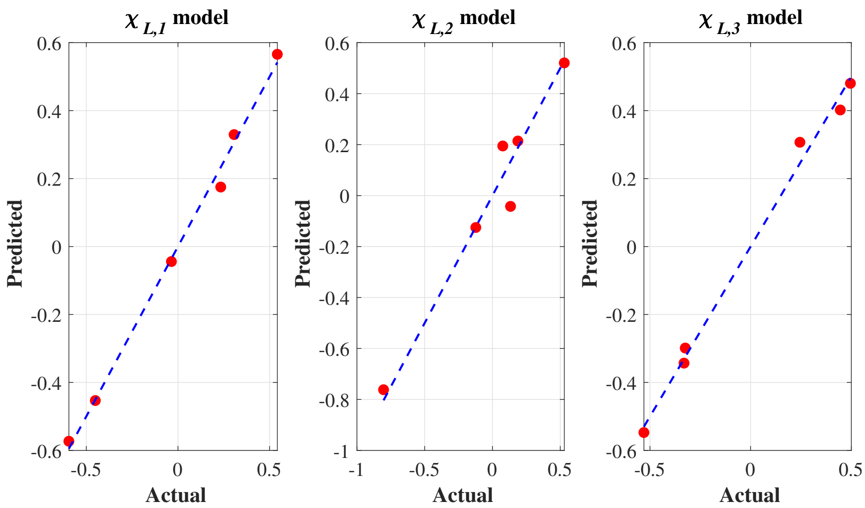

Figure 9.14: Comparison between actual and predicted loss-related parameter values for the three different regression models.

9.4 Results

To further interpret the predictive capability of the proposed model, the behavior of the global loss coefficient was analyzed within the reduced blade-loading parameter space. A response surface was constructed to visualize the dependence of the reconstructed global loss coefficient on the dominant blade-loading modes, providing a compact and physically interpretable representation of the model output.

The response surface was obtained by evaluating the predictive model over a three-dimensional grid defined in the space of the blade-loading coefficients $\chi_{B,1}$, $\chi_{B,2}$, and $\chi_{B,3}$. For each parameter, the grid spans the full range of values observed within the training dataset. At each point of the grid in the blade-loading parameter space, the three loss-related coefficients $\chi_{L,1}$, $\chi_{L,2}$, and $\chi_{L,3}$ were obtained by independently querying the corresponding regression models, each trained to predict one loss coefficient. The resulting set of coefficients was then used as weighting factors for the associated POD loss modes $\boldsymbol{\varphi}_L^k$, which represent waveform-like patterns of the characteristic spatial distributions of loss extracted through the advanced post-processing procedure. Through this POD-based reconstruction, the original loss field was recovered, from which the global inlet-to-outlet loss coefficient ω_{tot} was finally retrieved.

The resulting surface is shown in Figure 9.15, where the normalized global loss coefficient is represented as a function of the first two blade-loading coefficients, while the third coefficient is fixed at its mean value across the dataset.

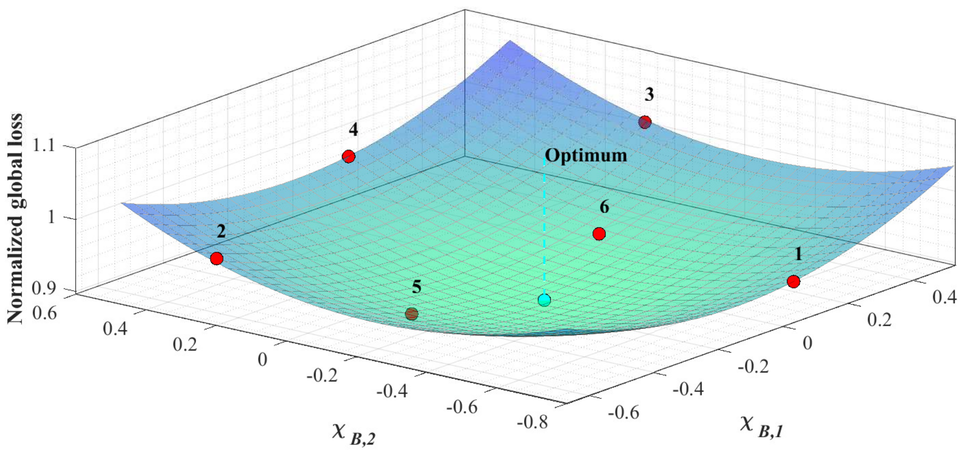

Figure 9.15: Overall response surface of the model in terms of normalized global loss coefficient. Red dots highlight the relative location of the six training points.

By applying a gradient-based minimization algorithm to the reconstructed response surface, a minimum of the global loss coefficient was identified. The corresponding optimal point is highlighted in Fig. 9.15 by a cyan marker. From the set of blade loading parameters associated with this optimum, the corresponding blade loading distribution can be explicitly reconstructed, providing the loading shape that minimizes the global inlet-to-outlet losses within the explored parameter space. The resulting C_p distribution is shown in Figure 9.16, where it is compared against the blade loading adopted as reference for the analyses discussed in the previous sections and the blade loading of the best-performing profile from the training dataset.

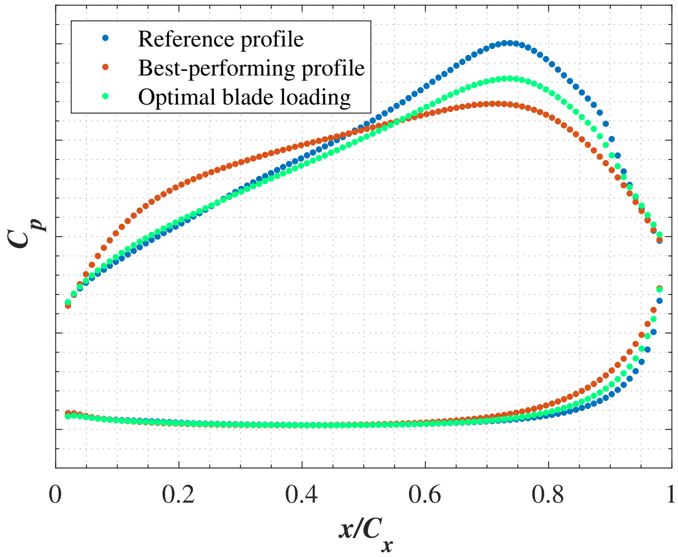

Figure 9.16: Blade loading distributions of the optimized profile (green curve), compared to the blade distribution of the reference profile (blue curve) and the best-performing profile from the dataset (red curve).

Compared to the reference blade loading, the optimal configuration exhibits a marginally reduced loading in the vicinity of the peak suction point, which, however, remains located at the same axial position. This results in a less intense flow acceleration along the fore portion of the suction side. Similarly, the decelerating region of the suction side is characterized by a slightly reduced overall diffusion factor, while a more pronounced reduction is observed in terms of the local maximum deceleration rate.

Despite the relatively small reduction in the overall blade loading magnitude, the localized modifications of the C_p distribution around the peak suction point lead to a markedly different partitioning of losses from both a modal and a spatial perspective. By comparing the cumulative loss curves for the new design point, reported in Figure 9.17, with those of the reference profile (see Fig. 8.7), a clear reduction in the losses associated with the accelerated suction-side boundary layer is first observed, together with a more moderate (but still evident) reduction in the losses occurring in the decelerating suction-side region.

Overall, the modification of the blade loading distribution results in an approximately 10% reduction of the total loss coefficient with respect to the reference profile, and about a 5% reduction relative to the dataset geometry characterized by the minimum loss coefficient. This improvement is achieved with a reduction of the overall blade loading of approximately 8% compared to the reference profile. Compared to the best-performing geometry in the dataset (see Fig. 9.16), differences in

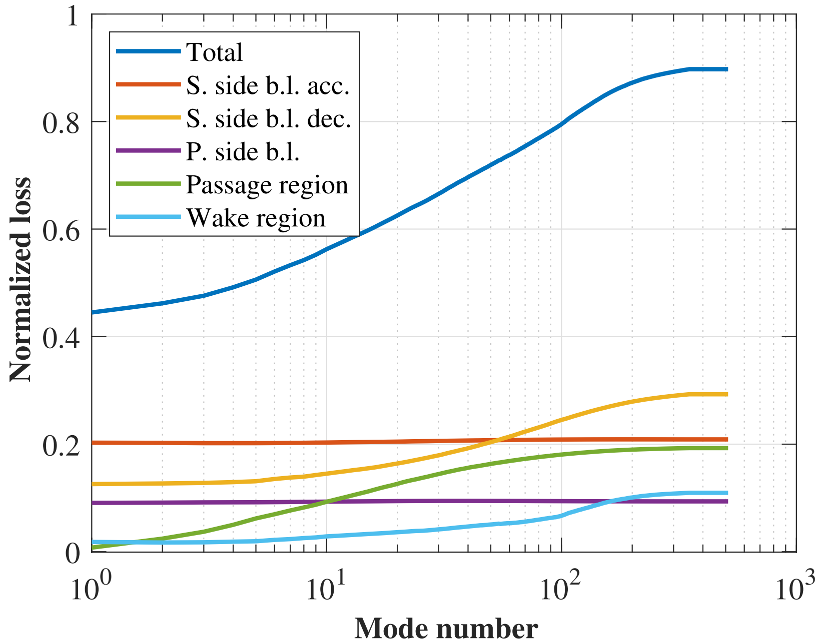

Figure 9.17: Cumulative contribution of POD modes to total losses, split by spatial regions of interest, for the optimal blade loading distribution.

the overall blade loading are observed, with the optimal distribution exhibiting lower C_p values over the front portion of the profile. This effect is, however, compensated by higher suction peaks toward the rear part of the blade.

As previously mentioned, since blade loading is primarily a potential-flow quantity, the optimal distribution identified through the predictive framework can be directly translated into a corresponding blade geometry. Starting from the optimal C_p distribution, the associated blade shape was therefore reconstructed through an inverse-design procedure based on consolidated machine-learning tools, which are described in [110]. These models, trained on a CFD-generated database, enable the real-time prediction of blade loading for a given geometry, effectively acting as high-fidelity surrogate models. Building upon these tools, the inverse-design procedure is formulated as an optimization problem in which standard geometric parameters (such as inlet/outlet flow angles, TE thickness, peak suction location, etc.) are iteratively varied. At each iteration, the corresponding loading distribution is instantaneously estimated by the ML model, and a loss function defined as the discrepancy between the target C_p and the predicted distributions is evaluated. The optimal geometry is then obtained by minimizing this loss function through a lightweight optimization algorithm. The comparison with the geometries available in the original dataset confirmed that the optimal profile does not introduce radically new geometric features, but rather represents a refined redistribution of curvature and camber with respect to existing designs. In this sense, the proposed methodology does not yield an abstract optimum in the reduced space, but leads to a physically realizable blade geometry, fully consistent with classical aerodynamic design principles.

Chapter 10

Conclusions

In the context of marine propulsion, the continuous demand for higher efficiency, reduced fuel consumption, and lower emissions has renewed interest in advanced gas turbine technologies, particularly aeroderivative-based architectures and combined-cycle solutions. Within this framework, the aerodynamic performance of low-pressure turbines plays a key role, as even moderate improvements can translate into significant gains at the system level. The present work was motivated by this perspective, aiming to improve the understanding and prediction of loss mechanisms in low-pressure turbine cascades operating under conditions representative of marine propulsion applications.

The research addressed the problem of loss generation and characterization by combining high-fidelity numerical simulations and advanced modal decomposition techniques. The primary objective was twofold: on the one hand, to gain physical insight into the mechanisms responsible for loss production and their spatial and temporal localization; on the other hand, to exploit this knowledge to construct a reduced-order predictive framework capable of linking blade geometry to loss behavior. Particular emphasis was placed on methodologies that remain compatible with the constraints of industrial design environments, where the availability of high-fidelity data is often limited and interpretability of results is essential.

A comprehensive loss analysis methodology was developed, starting from Large Eddy Simulation data and relying on an advanced post-processing routine designed to estimate local loss contributions through viscous dissipation and turbulent kinetic energy production. Proper Orthogonal Decomposition was first employed as a diagnostic tool to identify the dominant flow structures contributing to loss generation. The analysis of cumulative loss distributions across different spatial regions provided a consistent physical interpretation of loss mechanisms in terms of interactions between wake-induced unsteadiness and boundary-layer evolution.

To overcome the intrinsic limitations of classical POD in terms of spectral and temporal localization, a multi-resolution analysis based on the Discrete Wavelet Transform was introduced within the same post-processing framework. The DWT enabled a time–frequency characterization of turbulent kinetic energy production, revealing a strong modulation of losses over the blade-passing cycle and a redistribution of energy across multiple frequency bands. This analysis provided a high-level overview of loss dynamics, motivating the subsequent use of frequency- and phase-conditioned modal decompositions. Spectral Proper Orthogonal Decomposition and Phase-Locked Proper Orthogonal Decomposition were then applied to further isolate coherent (in frequency and time) loss-generating mechanisms. SPOD results showed a consistent redistribution of energy from passage-dominated dynamics toward boundary-layer-dominated mechanisms at increasing frequencies. PPOD analysis confirmed the phase-dependent nature of these interactions, demonstrating that peak loss production occurs when the upstream wake interacts most intensely with the decelerating region of the suction-side boundary layer. These considerations led to attributing differences in loss behavior across the entire dataset to variations in blade loading distribution, which control the relative position and intensity of wake–boundary-layer interaction mechanisms.

On the basis of these findings, blade loading was identified as the key geometric indicator governing loss generation. Being primarily a potential-flow effect and largely independent of Reynolds number at the operating conditions considered, blade loading provided a compact and physically meaningful parametrization of blade geometry. This insight motivated the development of a reduced-order predictive model linking blade loading distributions to loss characteristics.

A low-dimensional predictive framework was constructed by combining POD-based encodings of blade loading and loss distributions with simple, interpretable regression models. Despite the limited size of the available dataset, the proposed approach proved capable of accurately reproducing the dominant loss-related coefficients, achieving good predictive performance under Leave-One-Out Cross-Validation. The resulting response surface of the global loss coefficient enabled the identification of an optimal blade loading configuration associated with reduced losses. Analysis of the corresponding loading distribution showed that relatively small modifications in the vicinity of the peak suction region can lead to a significant redistribution of losses among different flow regions and modal contributions, ultimately yielding a measurable reduction in global loss. Finally, exploiting the fact that blade loading is a potential-flow quantity, the optimal loading distribution was used to reconstruct the corresponding new and optimized blade geometry.

Overall, this work demonstrates that the combined use of multiscale modal decomposition techniques and reduced-order modeling offers a powerful framework for both understanding and predicting loss mechanisms in low-pressure turbine cascades. Beyond the specific application considered, the methodology is general and can be extended to larger datasets, different operating conditions, or more complex turbine components. Future developments may include the integration of additional training points for the regression model, the extension of the procedure to the analysis of fully three-dimensional cases, and the coupling with optimization strategies for automated design exploration.

Appendix A

Equivalence Between Local and Global Approaches

The starting point for validating both loss evaluation approaches is the first law of thermodynamics, expressed as:

$$\frac{DE_t}{Dt} = \dot{Q} + \dot{W} \tag{A.1}$$

where the heat and work exchange terms are:

$$\dot{Q} = \left(-\frac{\partial q_x}{\partial x} - \frac{\partial q_y}{\partial y} - \frac{\partial q_z}{\partial z} \right) dxdydz = -\nabla \cdot \vec{q} \cdot dV \tag{A.2}$$

$$\dot{W} = \nabla \cdot (\boldsymbol{\sigma} \vec{V}) \cdot dV \tag{A.3}$$

with the stress tensor defined as:

$$\boldsymbol{\sigma} = -p\boldsymbol{I} + \boldsymbol{\tau} \tag{A.4}$$

Substituting the expressions for $\dot{Q}$ and $\dot{W}$ into Eq. A.1, rewriting in specific terms and rearranging yields the transport equation for total internal energy:

$$\rho \frac{De_t}{Dt} + \frac{Dp}{Dt} = \frac{\partial p}{\partial t} - \nabla \cdot \vec{q} + \nabla \cdot (\boldsymbol{\tau} \vec{V}) \tag{A.5}$$

Recalling the fundamental thermodynamic relation:

$$e_t + \frac{p}{\rho} = h_t \tag{A.6}$$

Equation A.5 can be rewritten as:

$$\rho \frac{Dh_t}{Dt} = \frac{\partial p}{\partial t} - \nabla \cdot \vec{q} + \nabla \cdot (\boldsymbol{\tau} \vec{V}) \tag{A.7}$$

The instantaneous kinetic energy of the flow is defined as:

$$K = \frac{1}{2}\rho ||\vec{V}||^2 = \frac{1}{2}\rho(u^2 + v^2 + w^2) \tag{A.8}$$

The transport equation for K is obtained by multiplying the momentum equations by the corresponding velocity components:

$$u \cdot NSE(u) + v \cdot NSE(v) + w \cdot NSE(w) = 0 \tag{A.9}$$

After algebraic manipulations, the transport equation for kinetic energy becomes:

$$\frac{DK}{Dt} = -\vec{v} \cdot \nabla p + \mu \left(u\nabla^2 u + v\nabla^2 v + w\nabla^2 w \right) \tag{A.10}$$

This can also be expressed in terms of the total pressure:

$$\frac{Dp_t}{Dt} = \frac{\partial p}{\partial t} + \mu \left(u\nabla^2 u + v\nabla^2 v + w\nabla^2 w \right) \tag{A.11}$$

Introducing a dissipation function ϕ [118]:

$$\phi = \nabla \cdot (\boldsymbol{\tau} \vec{V}) - \mu \left(u\nabla^2 u + v\nabla^2 v + w\nabla^2 w \right) \tag{A.12}$$

Equation A.11 becomes:

$$\frac{Dp_t}{Dt} = \frac{\partial p}{\partial t} + \nabla \cdot (\boldsymbol{\tau} \vec{V}) - \phi \tag{A.13}$$

Combining this with the total enthalpy transport equation gives:

$$\frac{Dh_t}{Dt} = \frac{Dp_t}{Dt} - \nabla \cdot \vec{q} + \phi \tag{A.14}$$

From the Gibbs relation, if the flow is adiabatic:

$$T_t \frac{Ds}{Dt} = \frac{1}{\rho_t}\phi \tag{A.15}$$

In the absence of work exchange, as in stator blades or in the relative frame for rotor blades:

$$T_t \frac{Ds}{Dt} = \frac{1}{\rho_t}\phi = -\frac{1}{\rho_t}\frac{Dp_t}{Dt} \tag{A.16}$$

Averaging over time:

$$\overline{T_t \frac{Ds}{Dt}} = -\frac{1}{\rho_t}\overline{\frac{Dp_t}{Dt}} \tag{A.17}$$

The integration of the right-hand side over the domain corresponds directly to loss generation, thus validating the local loss estimation based on total pressure variation.

However, for the global loss evaluation, additional considerations are required. The global loss coefficient ω is not strictly proportional to the time-average of the material derivative of p_t, but rather to the material derivative of the time-averaged p_t:

$$\int_D \frac{D\overline{p_t}}{Dt} dV \propto \omega \tag{A.18}$$

In general, time-averaging and material derivative do not commute:

$$\overline{\frac{Dp_t}{Dt}} \neq \frac{D\overline{p_t}}{Dt} \tag{A.19}$$

Applying Reynolds decomposition to velocity and pressure fields, and averaging the material derivative of p_t, we obtain:

$$\overline{\frac{Dp_t}{Dt}} = \underbrace{\overline{u}_i \frac{\partial \overline{p}}{\partial x_i} + \frac{D}{Dt}\left(\frac{1}{2}\overline{\rho}\, \overline{u}_i \overline{u}_i\right) + \frac{D}{Dt}\left(\frac{1}{2}\overline{\rho}\, \overline{u'_i u'_i}\right)}_{\frac{D\overline{p_t}}{Dt}} + \frac{\partial}{\partial x_j}\left[\overline{u}_i \left(\overline{\rho u'_i u'_j}\right)\right] + \frac{\partial \overline{\rho u'_i k}}{\partial x_i} + \frac{\partial \overline{u'_i p'}}{\partial x_i} \tag{A.20}$$

For the local and global formulations to be equivalent, the diffusive terms on the right-hand side must vanish or be negligible. If this condition is satisfied, then also the global approach based on total pressure loss is validated.

Bibliography

[1] European Environmental Agency. *Transport and environment report 2022 - Digitalisation in the mobility system: challenges and opportunities.* Tech. rep. Publications Office of the European Union, 2022.

[2] European Commission et al. *EU reference scenario 2020 – Energy, transport and GHG emissions – Trends to 2050.* Publications Office, 2021. DOI: 10.2833/35750.

[3] M. Crippa et al. *GHG emissions of all world countries.* Tech. rep. Publications Office of the European Union, 2023. DOI: 10.2760/953322.

[4] EDGAR/JRC. *Distribution of carbon dioxide emissions worldwide in 2022, by sector.* Accessed: July 09, 2024. 2023. URL: https://www.statista.com/statistics/1129656/global-share-of-co2-emissions-from-fossil-fuel-and-cement/.

[5] P. Baltrenas, P. Vaitiekūnas, and I. Mincevič. "Investigation on the impact of transport exhaust emissions on the air". In: *Journal of Environmental Engineering and Landscape Management* 12.1 (2004), pp. 3–11. DOI: 10.1080/16486897.2004.9636809.

[6] R.D. Morgenstern and D. Tirpak. "The greenhouse gases". In: *EPA J.* 16 (1990), p. 8.

[7] EPA Climate Change Indicators. "Atmospheric Concentrations of Greenhouse Gases". In: *EPA-United States Environmental Protection Agency* (2022).

[8] S.G. Tsyro and E. Berge. "Estimation of acidifying deposition in Europe due to international shipping emissions in the North Sea and the North East Atlantic Ocean". In: *WIT Transactions on Ecology and the Environment* 25 (1970), pp. 175–184.

[9] S.E. Schwartz. "Acid deposition: unraveling a regional phenomenon". In: *Science* 243.4892 (1989), pp. 753–763.

[10] Ø. Endresen et al. "Emission from international sea transportation and environmental impact". In: *Journal of Geophysical Research: Atmospheres* 108.D17 (2003).

[11] P. Elvingson and C. Ågren. *Air and the environment.* The Swedish NGO Secretariat on acid rain, 2004.

[12] F. Haglind. "A review on the use of gas and steam turbine combined cycles as prime movers for large ships. Part I: Background and design". In: *Energy Conversion and Management* 49.12 (2008), pp. 3458–3467.

[13] World Health Organization et al. *WHO global air quality guidelines: particulate matter (PM2. 5 and PM10), ozone, nitrogen dioxide, sulfur dioxide and carbon monoxide*. World Health Organization, 2021.

[14] World Trade Organization (WTO). *World Trade Statistical Review 2023*. Accessed: July 10, 2024. 2023. URL: https://www.wto.org/english/res_e/publications_e/wtsr_2023_e.htm.

[15] UN Trade and Development (UNCTAD). *REVIEW OF MARITIME TRANSPORT 2023: Towards a green and just transition*. Accessed: July 10, 2024. 2023. URL: https://unctad.org/publication/review-maritime-transport-2023.

[16] European Commission. *Reducing emissions from the shipping sector*. Accessed: July 10, 2024. 2024. URL: https://climate.ec.europa.eu/eu-action/transport/reducing-emissions-shipping-sector_en.

[17] V. Eyring et al. "Emissions from international shipping: 2. Impact of future technologies on scenarios until 2050". In: *Journal of Geophysical Research: Atmospheres* 110.D17 (2005).

[18] Wikipedia contributors. *Heavy fuel oil — Wikipedia, The Free Encyclopedia*. [Accessed: July 10, 2024]. 2024. URL: https://en.wikipedia.org/w/index.php?title=Heavy_fuel_oil&oldid=1199758804.

[19] A. Wankhede. *A Guide To Marine Gas Oil and LSFO Used On Ships*. Accessed: July 11, 2024. 2020. URL: https://www.marineinsight.com/guidelines/a-guide-to-marine-gas-oil-and-lsfo-used-on-ships/.

[20] IMO MARPOL. "MARPOL Annex VI: Regulations for the Prevention of Air Pollution from Ships". In: *Consolidated Edition* (2006).

[21] P. Pan et al. "Research progress on ship power systems integrated with new energy sources: A review". In: *Renewable and Sustainable Energy Reviews* 144 (2021), p. 111048.

[22] W.P. Symington et al. "Emerging technologies in marine electric propulsion". In: *Proceedings of the Institution of Mechanical Engineers, Part M: Journal of Engineering for the Maritime Environment* 230.1 (2016), pp. 187–198. DOI: 10.1177/1475090214558470.

[23] Z.Y. Tay and D. Konovessis. "Sustainable energy propulsion system for sea transport to achieve United Nations sustainable development goals: a review". In: *Discover Sustainability* 4.1 (2023), p. 20.

[24] Z. Wang et al. "A review of marine renewable energy storage". In: *International Journal of Energy Research* 43.12 (2019), pp. 6108–6150. DOI: https://doi.org/10.1002/er.4444.

[25] N.C. Shih et al. "Development of a 20 kW generic hybrid fuel cell power system for small ships and underwater vehicles". In: *International Journal of Hydrogen Energy* 39.25 (2014), pp. 13894–13901.

[26] S.Y. Gómez and D. Hotza. "Current developments in reversible solid oxide fuel cells". In: *Renewable and Sustainable Energy Reviews* 61 (2016), pp. 155–174.

[27] M. Moliere. "Expanding fuel flexibility of gas turbines". In: *Proceedings of the Institution of Mechanical Engineers, Part A: Journal of Power and Energy* 219.2 (2005), pp. 109–119.

[28] J. L. Lumley. "The structure of inhomogeneous turbulent flows". In: *Atmospheric turbulence and radio wave propagation* (1967), pp. 166–178.

[29] E. Canepa et al. "Flow coefficient and reduced frequency effects on low pressure turbine unsteady losses". In: *Journal of Propulsion and Power* 38.1 (2022), pp. 18–29.

[30] M. Dellacasagrande et al. "Analysis of Unsteady Loss Sensitivity to Incidence Angle Variation in Low Pressure Turbine". In: *Turbo Expo: Power for Land, Sea, and Air*. American Society of Mechanical Engineers, paper No. GT2023-103770. 2023.

[31] *Unsteady Flow Interactions Between a High- and Low-Pressure Turbine: Part 2 — Rotor-Synchronic Averaging and Proper Orthogonal Decomposition of the Unsteady Flow Fields*. Vol. Volume 2A: Turbomachinery. Turbo Expo: Power for Land, Sea, and Air. 2019, V02AT45A011.

[32] J.W. Kurelek, A.R. Lambert, and S. Yarusevych. "Coherent structures in the transition process of a laminar separation bubble". In: *AIAA Journal* 54.8 (2016), pp. 2295–2309.

[33] D. Lengani et al. "Identification and quantification of losses in a LPT cascade by POD applied to LES data". In: *International Journal of Heat and Fluid Flow* 70 (2018), pp. 28–40.

[34] Z. Liu, R.J. Adrian, and T.J. Hanratty. "Large-scale modes of turbulent channel flow: transport and structure". In: *Journal of Fluid Mechanics* 448 (2001), pp. 53–80.

[35] B.W. van Oudheusden et al. "Phase-resolved characterization of vortex shedding in the near wake of a square-section cylinder at incidence". In: *Experiments in Fluids* 39 (2005), pp. 86–98.

[36] R. Perrin et al. "Phase-averaged measurements of the turbulence properties in the near wake of a circular cylinder at high Reynolds number by 2C-PIV and 3C-PIV". In: *Experiments in fluids* 42 (2007), pp. 93–109.

[37] GE Vernova. *LM2500 aeroderivative gas turbine*. Accessed: July 22, 2024. 2024. URL: https://www.gevernova.com/gas-power/products/gas-turbines/lm2500.

[38] K. Ranasinghe et al. "Review of Advanced Low-Emission Technologies for Sustainable Aviation". In: *Energy* 118 (2019). DOI: 10.1016/j.energy.2019.115945.

[39] Kumar A. *Turbojet Engine: Definition, Parts or Construction, Working Principle, Advantages*. 2022. URL: https://themechanicalengineering.com/turbojet-engine/.

[40] J. Horlock. *Combined power plants: including combined cycle gas turbined CCGT plants*. Elsevier, 2013.

[41] Wikipedia contributors. *Combined steam and gas — Wikipedia, The Free Encyclopedia*. Accessed: July 16, 2024. 2024. URL: https://en.wikipedia.org/w/index.php?title=Combined_steam_and_gas&oldid=1212184577.

[42] K. Kosowski. *Ship Turbine Power Plants – Fundamentals of Thermodynamical Cycles*. Foundation for the Promotion of Maritime Industry, 2005.

[43] R. Cwilewicz and Z. Górski. "Coges as a future of marine propulsion plants". In: *Zeszyty Naukowe Akademii Morskiej w Gdyni* 84 (2014), pp. 5–13.

[44] W. Nirbito, M.A. Budiyanto, and R. Muliadi. “Performance analysis of combined cycle with air breathing derivative gas turbine, heat recovery steam generator, and steam turbine as LNG tanker main engine propulsion system”. In: *Journal of Marine Science and Engineering* 8.9 (2020), p. 726.

[45] A. Dotto, F. Satta, and U. Campora. “Energy, environmental and economic investigations of cruise ships powered by alternative fuels”. In: *Energy Conversion and Management* 285 (2023), p. 117011.

[46] K. Lee et al. “Hydrogen enriched LNG fuel for maritime applications–A life cycle study”. In: *International Journal of Hydrogen Energy* 78 (2024), pp. 333–343.

[47] Min Chul Lee et al. “Effect of hydrogen content on the gas turbine combustion performance of synthetic natural gas”. In: *Turbo Expo: Power for Land, Sea, and Air*. Vol. 55119. American Society of Mechanical Engineers. 2013, V01BT04A067.

[48] Y. Park, M. Choi, and G. Choi. “Thermodynamic performance study of large-scale industrial gas turbine with methane/ammonia/hydrogen blended fuels”. In: *Energy* 282 (2023), p. 128731.

[49] Y. Park et al. “Performance analysis of large-scale industrial gas turbine considering stable combustor operation using novel blended fuel”. In: *Energy* 236 (2021), p. 121408.

[50] D. Steffens. “The diesel engine and the environment”. In: *Wartsila North America,“Using Multi-Zone Simulations". Lawrence Livermore National Laboratory* (2003).

[51] J. Dahlkvist. “Examination Project, ship design, combined cycle vs. diesel engine”. B.S. thesis. Technical University of Denmark, DTU, DK-2800 Kgs. Lyngby, Denmark, 2007.

[52] A. Rivera-Alvarez, M. J Coleman, and J.C. Ordonez. “Ship weight reduction and efficiency enhancement through combined power cycles”. In: *Energy* 93 (2015), pp. 521–533.

[53] *Gas Turbines Keep The Surface Navy Moving*. Vol. 144/10/1,388. U.S. Naval Institute - The Independent Forum of the Sea Services. 2018.

[54] Rolls-Royce. *MT30 Marine Gas Turbine - Vessel references*. Accessed: July 21, 2024. 2024. URL: `https://www.rolls-royce.com/products-and-services/defence/naval/gas-turbines/mt30-marine-gas-turbine.aspx`.

[55] B.N. Sanneman. “Pioneering gas turbine-electric system in cruise ships: a performance update”. In: *Marine technology and SNAME news* 41.4 (2004), p. 161.

[56] F. Haglind. “A review on the use of gas and steam turbine combined cycles as prime movers for large ships. Part II: Previous work and implications”. In: *Energy Conversion and Management* 49.12 (2008), pp. 3468–3475.

[57] D. Tinsley. “Seminal gas turbine application arrives in Millennium cruise class COGES plant”. In: *Marine Propulsion Int* (2000), pp. 14–5.

[58] Ship Technology. *Celebrity Millennium*. Accessed: July 21, 2024. 2001. URL: `https://www.ship-technology.com/projects/millennium/?cf-view`.

[59] D. Barsi et al. "Gas turbine prime movers fuelled by LNG as a future alternative for sustainable power in marine propulsion: current emission policy assessment and exhaust quality evaluation". In: *E3S Web of Conferences*. Vol. 113. EDP Sciences. 2019, p. 02018.

[60] M.A. Budiyanto et al. "Techno-Economic Analysis of Combined Gas and Steam Propulsion System of Liquefied Natural Gas Carrier". In: *Energies* 17.6 (2024), p. 1415.

[61] L. Ellington et al. "Gas turbine propulsion for LNG transports". In: *Turbo Expo: Power for Land, Sea, and Air*. Vol. 42401. 2006, pp. 65–76.

[62] J.D. Denton. *Loss mechanisms in turbomachines.* American Society of Mechanical Engineers, paper No. 93-GT-435, 1993.

[63] T.A. Zaki. "From streaks to spots and on to turbulence: exploring the dynamics of boundary layer transition". In: *Flow, turbulence and combustion* 91.3 (2013), pp. 451–473.

[64] J. Verdoya. "EXPERIMENTAL ANALYSIS OF TRANSITIONAL FLOWS UNDER TURBINE-LIKE CONDITIONS VIA APPLICATION AND DEVELOPMENT OF ADVANCED POST-PROCESSING TECHNIQUES". PhD thesis. Università degli studi di Genova, 2022.

[65] F. Kameier and W. Neise. "Experimental Study of Tip Clearance Losses and Noise in Axial Turbomachines and Their Reduction". In: *J. Turbomach.* 119.3 (1997), pp. 460–471.

[66] M.G. Rose and N.W. Harvey. "Turbomachinery wakes: Differential work and mixing losses". In: *J. Turbomach.* 122.1 (2000), pp. 68–77.

[67] R.D. Stieger and H.P. Hodson. "The unsteady development of a turbulent wake through a downstream low-pressure turbine blade passage". In: *J. Turbomach.* 127.2 (2005), pp. 388–394.

[68] X. Wu and P.A. Durbin. "Evidence of longitudinal vortices evolved from distorted wakes in a turbine passage". In: *Journal of Fluid Mechanics* 446 (2001), pp. 199–228.

[69] H.P. Hodson and R.J. Howell. "Bladerow interactions, transition, and high-lift aerofoils in low-pressure turbines". In: *Annu. Rev. Fluid Mech.* 37 (2005), pp. 71–98.

[70] E. Canepa et al. "Experimental study of the effect of the rotor-stator gap variation on the tonal noise generated by low-speed axial fans". In: *19th AIAA/CEAS Aeroacoustics Conference*. 2013, p. 2046.

[71] S. Sarkar and P.R. Voke. "Large-eddy simulation of unsteady surface pressure over a low-pressure turbine blade due to interactions of passing wakes and inflexional boundary layer". In: *J. Turbomach.* 128.2 (2006), pp. 221–231.

[72] P. Z. Sterzinger et al. "Impact of Varying High- and Low-Pressure Turbine Purge Flows on a Turbine Center Frame and Low-Pressure Turbine System". In: *Journal of Turbomachinery* 142.10 (2020), p. 101011.

[73] K.A. Gompertz and J.P. Bons. "Combined unsteady wakes and active flow control on a low-pressure turbine airfoil". In: *Journal of Propulsion and Power* 27.5 (2011), pp. 990–1000.

[74] D. Simoni et al. "Turbine cascade profile loss sensitivity to incoming wake parameters: effects of reduced frequency, wake momentum defect and axial

gap". In: *Turbo Expo: Power for Land, Sea, and Air*. Vol. 2A: Turbomachinery. American Society of Mechanical Engineers. 2019, V02AT45A018.

[75] F. Satta et al. "Loading distribution effects on separated flow transition of ultra-high-lift turbine blades". In: *Journal of Propulsion and Power* 30.3 (2014), pp. 845–856.

[76] D. Simoni, M. Ubaldi, and P. Zunino. "Experimental investigation of the interaction between incoming wakes and instability mechanisms in a laminar separation bubble". In: *Experimental thermal and fluid science* 50 (2013), pp. 54–60.

[77] A. Binder et al. "An experimental investigation into the effect of wakes on the unsteady turbine rotor flow". In: *Journal of Engineering for Gas Turbines and Power* 107.2 (1985), pp. 458–465.

[78] A. Binder, T. Schröeder, and J. Hourmouziadis. "Turbulence measurements in a multistage low-pressure turbine". In: *J. Turbomach.* 111.2 (1989), pp. 153–161.

[79] R.X. Meyer. "The Effect of Wakes on the Transient Pressure and Velocity Distributions in Turbomachines". In: *Transactions of the American Society of Mechanical Engineers* 80.7 (2022), pp. 1544–1551. DOI: 10.1115/1.4012797.

[80] D. Lengani et al. "Investigation on Strain and Stress Principal Axes in Unsteady DNS Turbine Data". In: *Turbo Expo: Power for Land, Sea, and Air*. Vol. 86120. American Society of Mechanical Engineers. 2022, V10DT37A031.

[81] A.N. Kolmogorov. "On degeneration (decay) of isotropic turbulence in an incompressible viscous liquid". In: *Dokl. Akad. Nauk SSSR*. Vol. 31. 1941, pp. 538–540.

[82] L. Stein. "A skew-symmetric, conservative Finite Difference Scheme for the simulation of reactive Flow". PhD thesis. Master's thesis, TU-Berlin, 2013.

[83] O. Reynolds. "IV. On the dynamical theory of incompressible viscous fluids and the determination of the criterion". In: *Philosophical transactions of the royal society of london (A)* 186 (1895), pp. 123–164.

[84] P. Moin and K. Mahesh. "Direct numerical simulation: a tool in turbulence research". In: *Annual review of fluid mechanics* 30.1 (1998), pp. 539–578.

[85] M. Russo et al. "Optimization of Low-Pressure Turbine Blade by Means of Fine Inspection of Loss Production Mechanisms". In: *Journal of Turbomachinery* 147.4 (Oct. 2024), p. 041005. DOI: 10.1115/1.4066818.

[86] P.J. Przytarski et al. "The role of turbulence transport in mechanical energy budgets". In: *Journal of Turbomachinery* 146.10 (2024), p. 101008.

[87] J.T.C. Liu. "Contributions to the understanding of large-scale coherent structures in developing free turbulent shear flows". In: *Advances in applied mechanics* 26 (1988), pp. 183–309.

[88] L. Sirovich. "Turbulence and the dynamics of coherent structures, Parts I, II and III". In: *Quart. Appl. Math.* (1987), pp. 561–590.

[89] A. Towne, O.T. Schmidt, and T. Colonius. "Spectral proper orthogonal decomposition and its relationship to dynamic mode decomposition and resolvent analysis". In: *Journal of Fluid Mechanics* 847 (2018), pp. 821–867. DOI: 10.1017/jfm.2018.283.

[90] D. Biassoni, M. Russo, and D. Lengani. "A high-performance code for the use of spectral pod on the analysis of turbomachinery high fidelity simulations". In: *Journal of the Global Power and Propulsion Society* 9 (2026), pp. 330–340.

[91] D. Lengani et al. "Synchronization of multi-plane measurement data by means of POD: application to unsteady boundary layer transition". In: *Experiments in Fluids* 59 (2018), pp. 1–18.

[92] Mathieu Legrand, José Nogueira, and Antonio Lecuona. "Flow temporal reconstruction from non-time-resolved data part I: mathematic fundamentals". In: *Experiments in fluids* 51.4 (2011), pp. 1047–1055.

[93] A. Hekmati, D. Ricot, and P. Druault. "About the convergence of POD and EPOD modes computed from CFD simulation". In: *Computers & Fluids* 50.1 (2011), pp. 60–71. ISSN: 0045-7930. DOI: `10.1016/j.compfluid.2011.06.018`.

[94] F. Nicoud and F. Ducros. "Subgrid-scale stress modelling based on the square of the velocity gradient tensor". In: *Flow, turbulence and Combustion* 62.3 (1999), pp. 183–200.

[95] S. Benhamadouche et al. "Synthetic turbulent inflow conditions based on a vortex method for large-eddy simulation". In: *Progress in Computational Fluid Dynamics, an International Journal* 6.1-3 (2006), pp. 50–57.

[96] D. Biassoni. "Development of an Advanced Data Analysis Procedure for High-Fidelity Simulations of Aeronautical Turbine Blades". Supervisors: Daniele Simoni and Dario Barsi. Master's Thesis. Università di Genova, June 2022.

[97] CINECA. *GALILEO100 Supercomputer - Hardware Overview*. `https://www.hpc.cineca.it/systems/hardware/galileo100`. Accessed: 2025-09-03.

[98] CINECA. *LEONARDO Supercomputer - Hardware Overview*. `https://www.hpc.cineca.it/systems/hardware/leonardo`. Accessed: 2025-09-03.

[99] CINECA. *LEONARDO Supercomputer Official Portal*. `https://leonardo-supercomputer.cineca.eu`. Accessed: 2025-09-03.

[100] IT4Innovations. *Karolina Supercomputer - System Overview*. `https://www.it4i.cz/en/infrastructure/karolina`. Accessed: 2025-09-03.

[101] D. Lengani et al. "On the identification and decomposition of the unsteady losses in a turbine cascade". In: *Journal of Turbomachinery* 141.3 (2019).

[102] T. Sayadi and P. J Schmid. "Parallel data-driven decomposition algorithm for large-scale datasets: with application to transitional boundary layers". In: *Theoretical and Computational Fluid Dynamics* 30 (2016), pp. 415–428.

[103] P. Viviani et al. "Taming Multi-node Accelerated Analytics: An Experience in Porting MATLAB to Scale with Python". In: *Computational Intelligence in Security for Information Systems Conference*. Springer. 2022, pp. 200–210.

[104] M. Fiore et al. "LES of the T106 low-pressure turbine: Spectral proper orthogonal decomposition of the flow based on a fluctuating energy norm". In: *Computers & Fluids* 252 (2023), p. 105761.

[105] X. He et al. "Spectral proper orthogonal decomposition of compressor tip leakage flow". In: *Physics of Fluids* 33.10 (2021).

[106] O.T. Schmidt and T. Colonius. "Guide to spectral proper orthogonal decomposition". In: *AIAA journal* 58.3 (2020), pp. 1023–1033.

[107] J. Leggett et al. "Loss prediction in an axial compressor cascade at off-design incidences with free stream disturbances using large eddy simulation". In: *J. Turbomach.* 140.7 (2018).

[108] D. Lengani et al. "Coherent structures formation during wake-boundary layer interaction on a LP turbine blade". In: *Flow, Turbulence and Combustion* 98 (2017), pp. 57–81.

[109] D. Lengani et al. "On the receptivity of low-pressure turbine blades to external disturbances". In: *Journal of Fluid Mechanics* 937 (2022), A36.

[110] J. Bellucci et al. "A Real Time AI-Based Strategy for the Design of a Low-Pressure Turbine Profile". In: *Turbo Expo: Power for Land, Sea, and Air*. Vol. 88087. American Society of Mechanical Engineers. 2024, V12DT34A026.

[111] S. Lee et al. "CNN-based image recognition for topology optimization". In: *Knowledge-Based Systems* 198 (2020), p. 105887.

[112] M. Morimoto et al. "Convolutional neural networks for fluid flow analysis: toward effective metamodeling and low dimensionalization". In: *Theoretical and Computational Fluid Dynamics* 35.5 (2021), pp. 633–658.

[113] H. Eivazi et al. "Deep neural networks for nonlinear model order reduction of unsteady flows". In: *Physics of Fluids* 32.10 (2020).

[114] D. Petronio et al. "A ML strategy for the identification of optimal LPT design region and related blade shape". In: *Aerospace Science and Technology* 148 (2024), p. 109118. DOI: 10.1016/j.ast.2024.109118.

[115] G.M. Robinson and A.J. Keane. "Concise orthogonal representation of supercritical airfoils". In: *Journal of Aircraft* 38.3 (2001), pp. 580–583.

[116] D.J.J. Toal et al. "Geometric Filtration Using Proper Orthogonal Decomposition for Aerodynamic Design Optimization". In: *AIAA Journal* 48.5 (2010), pp. 916–928. DOI: 10.2514/1.41420.

[117] D. Simoni et al. "Off-design performance of a highly loaded low pressure turbine cascade under steady and unsteady incoming flow conditions". In: *Journal of Turbomachinery* 137.7 (2015).

[118] H. Schlichting and J. Kestin. *Boundary layer theory*. Vol. 121. Springer, 1961.