High-Fidelity Aerostructural Design Optimization of Wind Turbine Rotors

by

Marco Mangano

A dissertation submitted in partial fulfillment
of the requirements for the degree of
Doctor of Philosophy
(Aerospace Engineering)
in the University of Michigan
2024

Doctoral Committee:

Professor Joaquim R. R. A. Martins, Chair
Associate Professor Matthew D. Collette
Associate Professor Andrew Ning
Professor Veera Sundararaghavan

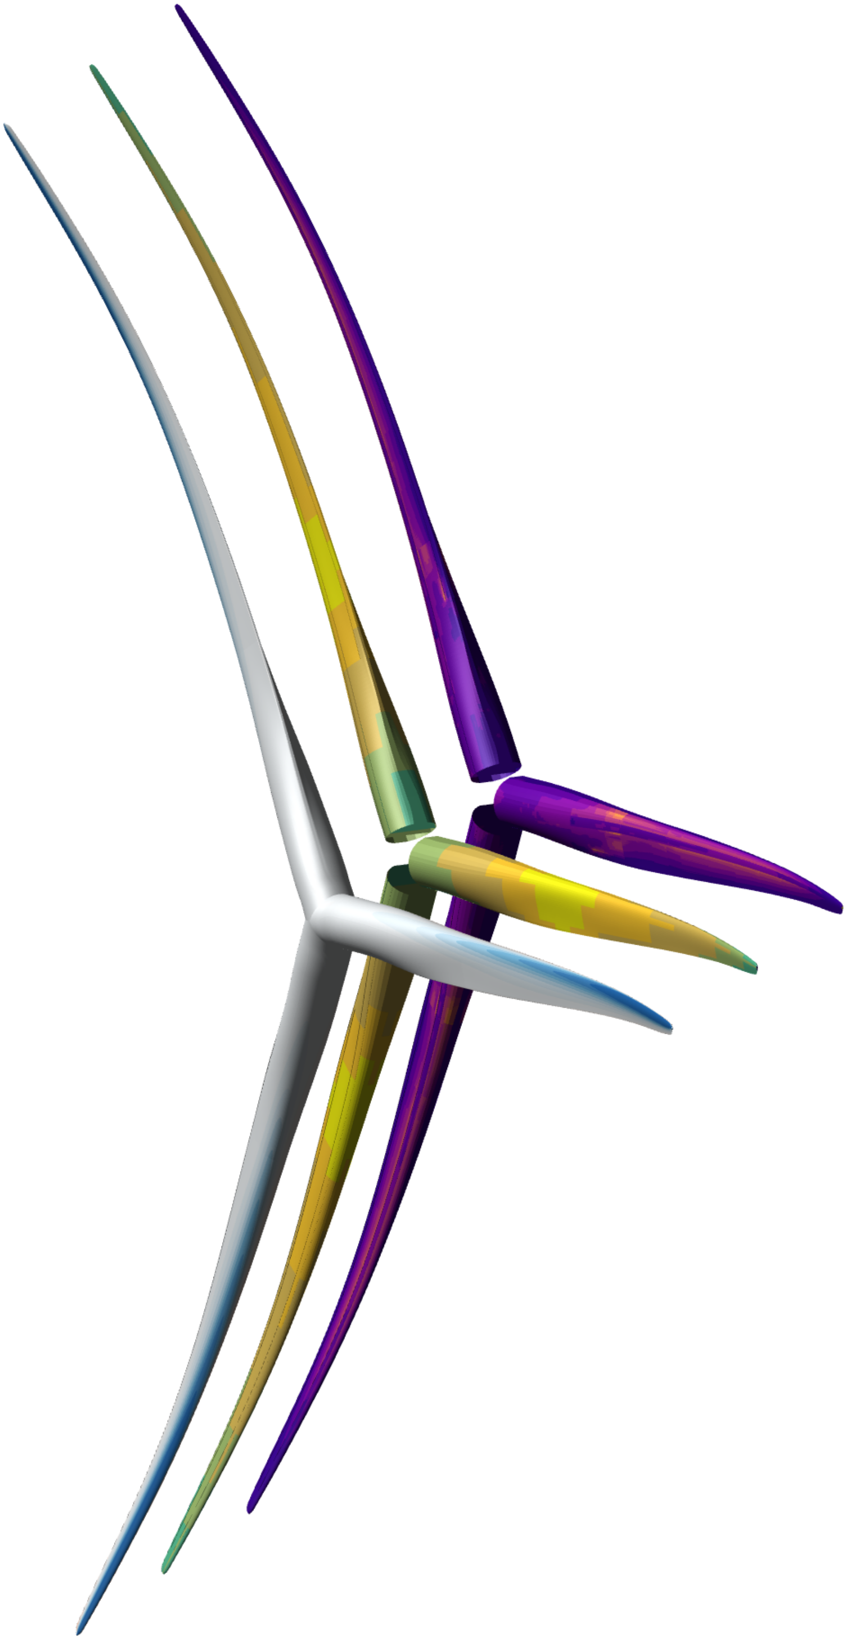

Marco Mangano

mmangano@umich.edu

ORCID iD: 0000-0001-8495-3578

Acknowledgments

This manuscript sums up my research efforts during the last 4.5 years at the University of Michigan. However, it inevitably fails to capture the emotional and personal challenges encountered along the way, and the contributions of the people who shared even a part of this journey with me. Without them, hardly any of this would have been feasible.

In the first place, I want to thank my advisor Prof. Joaquim R. R. A. Martins. By first accepting me as a visiting master's student and then offering me the chance to pursue a PhD, he gave me a life-changing opportunity that I will always be grateful for. His support and trust through the most challenging parts of my PhD boosted my self-confidence and made me grow into a more mature and focused researcher. I am also thankful for the prompt support and insightful feedback of my dissertation committee members Prof. Collette, Prof. Ning, and Prof. Sundararaghavan, who made my manuscript more solid and the path to the defense more enjoyable.

I am endlessly grateful for the warm support and friendship of my "co-authors" over the last few years. They have been a patient sounding board for my often confusing streams of ideas and helped me concretize the most exciting parts of my work. Dr. (now Prof.!) Sicheng He has been instrumental in making the first optimizations possible, helping with derivative debugging and structural meshing. Dr. Yingqian Liao patiently held my hand and shared her expertise as we explored composite structure optimization and untangled the many issues along the way. Dr. Eirikur Jonsson shared his vast experience on best optimization practices, showing me the way forward at times when I was feeling a bit lost. I am grateful for Dr. Denis-Gabriel Caprace's friendship during the ATLANTIS project and beyond, for the countless hours of venting, debugging, and brainstorming, and for the exciting research ideas we hope to keep working on.

The MDO Lab has an incredible pool of knowledge to learn from, but its greatest strength is the deep sense of community developed over the years. Sharing my journey with Sabet and Shugo since the start has been heartwarming and made the road look a bit less scary. John, Sham, Anil, Josh, and Ella acted as elder siblings as I learned my way through lab codes and life as a graduate student. Ali, Andrew, Alex C, Alex K, Bernardo, Eytan, Galen, Hannah, Saja, and all my "younger siblings" and visitors who joined the lab after me brought a wave of energy and fun that gave me a laugh even in the most dull or stressful days. Thanks to Saja and Ella, who have been very supportive office mates, always ready to stop and listen to my technical or personal digressions.

Even if they are not officially part of the lab, thanks to Dr. Mads Madsen for his warm support and the mindful technical advice, and thanks to Prof. Cristina Riso for our long walks and even longer discussions on life in academia far away from home.

There are other folks outside the lab who also made the journey much more meaningful and less stressful than it could have been. Ruthie has been an awesome friend and student manager, helping to deal with the full spectrum of bureaucratic and emotional crises. All the fellow students and faculty who supported GSAC and the DEI committee made me feel part of a broader community of caring people within FXB. I am thankful to have shared the last part of this journey with fellow GEO organizers, who inspired me and opened my eyes on what is worth fighting for, now more than ever. I am grateful for all my roommates, who kept the house lively over the long Michigan winters. Although it lasted less than we hoped, the time at 601 has been great fun, and I am thankful for all the folks who stopped by for dinner, a drink, or to share the warmth of a bonfire.

I feel incredibly lucky to have my partner (and picket-line crush) Evelyn in my life. Her love, sense of humor, determination, and support have made the last few months of my PhD sweet and exciting, despite the tremendous amount of pressure I was under. I am so excited for our journey ahead.

Being this far away from home and so deep into research work has been hard on the folks who love me on the other side of the world. I am very grateful for the loved ones who stayed with me all along, despite not being as present as I could have been - and I still keep in my heart those with whom I grew apart over the years. Matteo, Francesca, Gippo, and Giulia have been supporting me through the roughest spots and I cannot express how important they are to me.

My family's support and love have been unwavering through the years. They have pushed me forward without hesitation, despite what it cost them, and cheered for me all along. Making them proud fills my heart.

To conclude, I want to acknowledge the sponsors of my research. I have been supported by the Department of Energy (DOE) Advanced Research Projects Agency-Energy (ARPA-E) Program award DE-AR0001186 entitled "Computationally Efficient Control Co-Design Optimization Framework with Mixed-Fidelity Fluid and Structure Analysis". Thanks to the DOE ARPA-E Aerodynamic Turbines Lighter and Afloat with Nautical Technologies and Integrated Servo-control (ATLANTIS) team for their support.

This research was partly supported through computational resources and services provided by Advanced Research Computing at the University of Michigan, Ann Arbor.

We heavily relied Texas Advanced Computing Center (TACC) Stampede2 High-Performance Computing system via the Extreme Science and Engineering Discovery Environment (XSEDE), now Advanced Cyberinfrastructure Coordination Ecosystem: Services and Support (ACCESS),

which is supported by the National Science Foundation grants #1548562, #2138259, #2138286, #2138307, #2137603, and #2138296.

Table of Contents

Chapter

List of Figures

List of Tables

List of Appendices

List of Acronyms

AEP annual energy production

BEMT blade element momentum theory

CFD computational fluid dynamics

CSM computational structural mechanics

DEL damage-equivalent loads

DLC design load cases

DTU Danmarks Tekniske Universitet

FFD free-form deformation

IEA International energy agency

IPCC Intergovernmental panel on climate change

IRENA International renewable energy agency

KS Kreisselmeier–Steinhauser

LCOE levelized cost of energy

MACH MDO of aircraft with high-fidelity

MDA multidisciplinary analysis

MDO multidisciplinary design optimization

NASA National Aeronautics and Space Administration

PDF probability density distribution

RANS Reynolds-averaged Navier–Stokes

SA Spalart–Allmaras

TACS toolkit for analysis of composite structures

TSR tip speed ratio

XDSM extended design structure matrix

BR Below-rated

BR-T5 Below-rated, low TSR

BR-T10 Below-rated, high TSR

HW High-wind

MR Min-RPM

Abstract

THE mitigation of climate change effects through the decarbonization of the global production system is one of the most pressing challenges of our time. Wind and solar power will account for tenths of petawatt-hour yearly by the second half of this century. Thus, efficient energy harvesting systems need to be rapidly developed to match the demand for low-cost-per-kilowatt electricity from renewable sources.

Wind energy technology has matured prominently in the last few decades. However, the growing size and complexity of modern wind turbines is pushing the boundaries of conventional design tools and methodologies. Large and highly flexible rotors are characterized by complex multi-physics interactions and unconventional design features that require novel investigation strategies. High-fidelity design tools promise to address these arising challenges. Their application however has been limited by their computational and implementation costs so far.

The outcome of my doctoral work is the first set of large-scale high-fidelity wind turbine rotor aerostructural design optimization studies. I use a framework that couples computational fluid dynamics (CFD) and computational structural mechanics (CSM) analyses with an efficient gradient-based optimization algorithm. The coupled-adjoint sensitivity solver implementation makes the problem numerically tractable and highly scalable. The optimization studies I perform account for multiple representative steady-state below-rated operating conditions and use hundreds of structural and geometric design variables simultaneously.

First, I adapt an existing aircraft design framework to handle a wind turbine rotor and its system sensitivities. Both the aerodynamic solver and geometry manipulation tool are enhanced, and a newly developed aerostructural rotor model is verified.

Following that, I demonstrate the framework optimization capabilities on this benchmark rotor model. I dissect optimal design features and explore tradeoffs between steady-state aerodynamic efficiency and structural weight, quantifying the impact of design freedom and objective formulation on the final design.

A more refined aerostructural model is then used to perform composite fiber angle tailoring. The optimizer leverages material anisotropic properties to increase local stiffness and reduce the sandwich panel thickness. The resulting designs are more than 10% lighter than the benchmark for

the same planform shape.

Next, I optimize the rotor for multiple wind and rotation velocities. Single-point designs are outperformed over a larger operational envelope and flow separation at off-design conditions is mitigated. Enabling modifications to airfoil geometry leads to radical rotor redesign. Using the full set of geometrical and structural variables leads to an aeroelastically tailored rotor with more than 10% higher average torque and a 40% lighter structure than the baseline model.

Finally, I collaborated to the development of a mixed-fidelity framework combining our high-fidelity code with conventional design tools. This prototype software enables high-fidelity optimization including life-cycle sizing constraints. We use state-of-the-art software to predict extreme loads and fatigue damage from time-accurate multiphysics simulations. Two different approaches are proposed to integrate these constraints within our original framework.

Modelling the full range of operating conditions and improving the tradeoffs between accuracy and cost remain open challenges. Nevertheless, this work demonstrates how coupled CFD-CSM approaches can capture unconventional design tradeoffs and provide an unprecedented level of insight on the optimal blade design features. High-fidelity optimization can complement conventional design approaches using more realistic blade-resolved rotor models and leveraging a larger design space. The results I present pave the way for more comprehensive design studies and holistic multifidelity design optimization workflows.

Chapter 1

Introduction

Climate scientists overwhelmingly agree that anthropogenic emissions alter the Earth's atmosphere's average temperature and composition. Such alteration will negatively impact natural resources availability, industrial-scale food production, and ultimately, the well-being of current and future generations [1]. The International energy agency (IEA) [2] has proposed net-zero emissions targets to mitigate global warming and limit the global temperature increase to 2.0°C by the end of the century. To meet these goals, almost 70 % of the global energy needs in 2050 should be harvested from solar and wind resources.

Onshore wind energy is a mature technology on an industrial scale, and there is a growing interest in developing more extensive and efficient offshore installations. Like other carbon-free energy sources, resource intermittency and installation site constraints currently limit the amount of energy captured. Industry trends, driven by the need to mitigate these issues, suggest that wind turbines and their rotors will keep growing in size in the next decades [3, 4]. Large-diameter rotors with low solidity can generate the same power output at a higher tip speed ratio (TSR) and lower torque, ultimately lowering the drivetrain costs [5, Sec. 5.5.5]. Turbines with large rotors can be deployed onshore and offshore. Onshore, these turbines extend the operating range to lower wind speeds. Offshore deployment circumvents transportation constraints and exploits the untapped potential of wind away from the coastline [6].

Numerical simulations are essential to model the multiphysics interactions that affect the performance of wind turbines and accelerate their development [7]. From a design perspective, turbine rotors are characterized by their aerostructural behavior. However, the traditional sequential design of aerodynamic shape and internal structure does not fully account for the interaction between fluid dynamics and structural response [8, 9, 10]. Monolithic multidisciplinary design optimization (MDO) approaches exploit trade-offs in tightly coupled aerostructural models to ultimately drive down the cost of energy of wind turbines [11]. However, such models are challenging to implement within gradient-based MDO because they require accurate and efficient coupled derivative computation to guide the optimizer toward the optimal design.

The core of this dissertation presents the first application of a high-fidelity aerostructural optimization tool for wind turbine design. I build on a previously developed software, MDO of aircraft with high-fidelity (MACH) [12, 13], and previous works on a computational fluid dynamics (CFD)-based approach for aerodynamic shape optimization [14, 15].

Dhert et al. [14] performed the first aerodynamic shape optimization study with MACH. They maximized the torque of the NREL Phase IV turbine using hundreds of geometrical variables and considering up to 25 operational wind speed scenarios. Madsen et al. [15] used a Reynolds-averaged Navier–Stokes (RANS)-based aerodynamic shape optimization framework to optimize a turbine rotor's aerodynamic performance by varying a large set of design variables that parametrized the blade planform and cross-sectional shape.

My work extends the analysis and optimization to a tightly coupled, steady-state aerostructural model. I optimize the rotor blades by considering the aerostructural performance at below-rated operating conditions. The combination of CFD and computational structural mechanics (CSM) expands the optimization design space and improves the fidelity compared to blade element momentum theory (BEMT) and beam-theory-based design tools —occasionally referred to as "low-fidelity", by opposition to the high-fidelity CFD-CSM coupling proposed in the current work. I discuss how my work relates to the broader literature in Section 1.3.

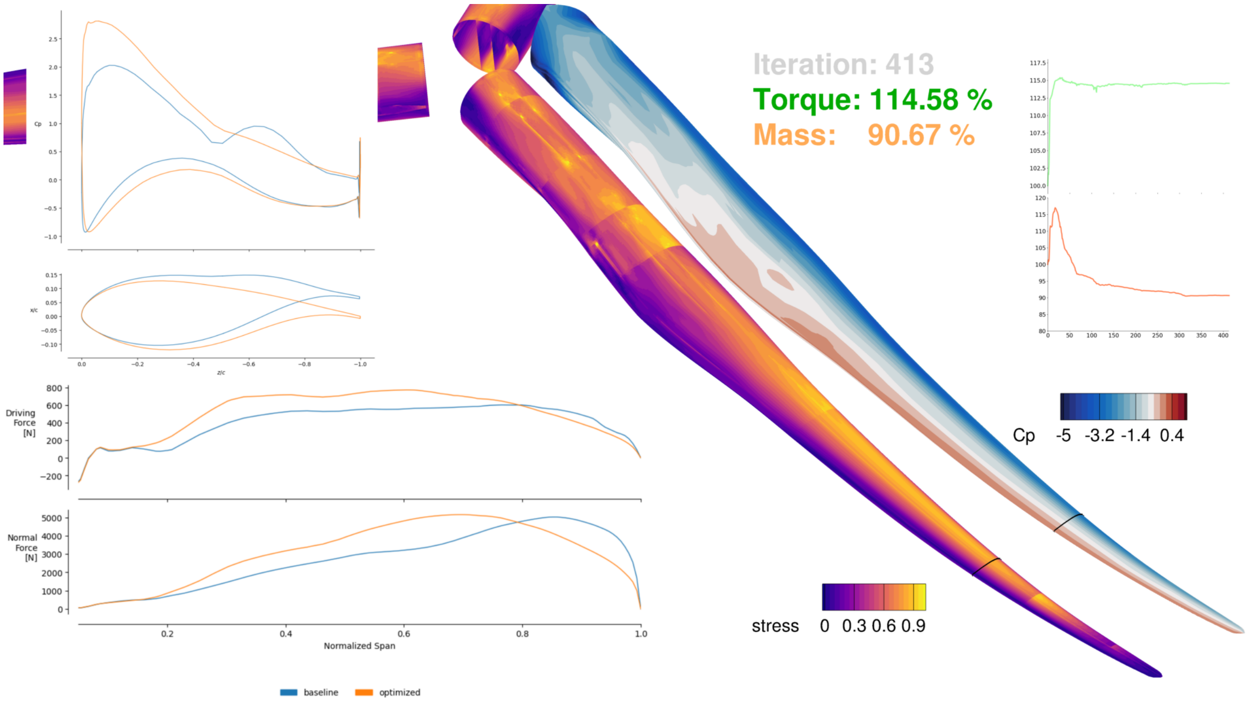

Figure 1.1: The proposed approach simultaneously increases the torque and reduces the mass by optimizing the wind turbine rotor outer mold line and its structural layout. The case in the figure is discussed in Chapter 5.

I first demonstrate the optimization capabilities for a single inflow condition, to limit the computational cost. I simultaneously optimize the blade planform and its structural thickness distribution to design lighter and more efficient wind turbine rotors. I assess the impact of design freedom and problem formulation on the final layout, and explore a Pareto front of optimal designs. A rotor with optimized planform and structural distribution is shown in Figure 1.1.

Following the initial study, I heavily update the structural model to analyze and optimize composite anisotropic properties. I first explore fiber composite tailoring to design lighter and stiffer rotor structures with an unprecedented design freedom. I assess the impact of different structural parametrizations and assumptions on the final design. Optimized structures can withstand higher loads than baseline with thinner structural panels on the same planform.

I then explore more complex multipoint problem formulations. I select additional simulation scenarios, or "design points", with different combinations of wind speed and rotor rotation rates. These design points are included in the optimization to obtain a rotor with better performance over its nominal operational envelope. Local shape sections are manipulated to enable more radical redesigns of the rotor. I again explore different problem formulations, demonstrating that including appropriate design points is essential to mitigate off-design performance drops.

The most complex multipoint optimization problem, including the full set of geometrical and structural design variables, leads to a rotor with an average torque output 10% higher than baseline and a structure almost 40% lighter.

Finally, in collaboration with the FLOW Lab [1], we developed a prototype mixed-fidelity optimization framework to include life-cycle considerations in the design problem. High-fidelity optimization can become excessively expensive when evaluating many design points, regardless of the framework and implementation efficiency. We extend the MACH framework to handle both high-fidelity and conventional BEMT in the optimization, considering fatigue damage and extreme loads in the blade sizing process. This approach combines the efficient implementation and design capabilities of MACH with comparably cheaper time-accurate simulations, capturing the rotor response over a representative interval of its operating life.

In this chapter, I first give an overview of wind energy systems and the necessity of a transition to renewable energy sources (Section 1.1). Following that, I introduce the key concepts of gradient-based MDO in Section 1.2. Section 1.3 covers the most relevant wind turbine optimization works from scientific literature, to the best of our knowledge. Finally, the contributions of my dissertation are listed in Section 1.4.

---

[1] FLOW Lab web page

1.1 Wind energy systems

Wind turbines are complex machinery with a long history. Although their employment for energy production is relatively recent, they have a key role to reduce the global economy dependency on fossil fuels. This section provides a high-level overview of the driving techno-economic factors behind the development and deployment of modern wind turbines.

Section 1.1.1 outlines the motivation for an urgent transition to renewable energy sources, spearheaded by solar and wind systems. An engineering breakdown of a wind turbine is provided in Section 1.1.2. I discuss the design drivers for cost-effective wind energy systems in Section 1.1.3. The fundamental physical principles of wind power extraction are finally presented in Section 1.1.4.

1.1.1 The transition to renewable energy sources

Record-breaking seasonal [2] and average [3] temperatures are prompting appeals to address human-caused climate change through the "*accelerated implementation of adaptation actions in this decade*" [16]. The economy decarbonization and the transition to renewable sources remain primarily a policy-driven effort that must include broader socio-technical considerations [17], ranging from grid integration to social acceptance. Nevertheless, the development of more efficient energy conversion technology has a direct impact on the reduction cost of energy from renewable sources. Despite the current market and supply chain challenges, IRENA [4] recorded a cost-of-energy reduction by more than 13% in 2020-2021, and predicts that more than 10 TW of installed capacity will be needed by 2050 [18].

The Intergovernmental panel on climate change (IPCC) [5] is a UN-sanctioned body aiming at establishing the scientific basis of human-made climate change. The experts who contributed to the most recent report [16] pictured a dire scenario for the upcoming decades. 150 years of intensive fossil fuels use have exponentially increased the concentration of CO_2 and other greenhouse gasses and pollutants in the atmosphere. Unsustainable resource and land use has jeopardized the climate resiliency of large parts of the world, where billions of people live. This climate injustice has disproportionately affected the global south, but the far reaching effects of climate change are affecting the quality of life all over the globe.

Global warming is measured against pre-industrial age average temperatures – the oldest available systemic meteorological observations. The current global surface temperature has increased

---

[2] National Aeronautics and Space Administration (NASA) report, last accessed October 2023

[3] Nature News Explainer, last accessed October 2023

[4] International Renewable Energy Agency website, last accessed November 2023

[5] IPCC website, last accessed November 2023

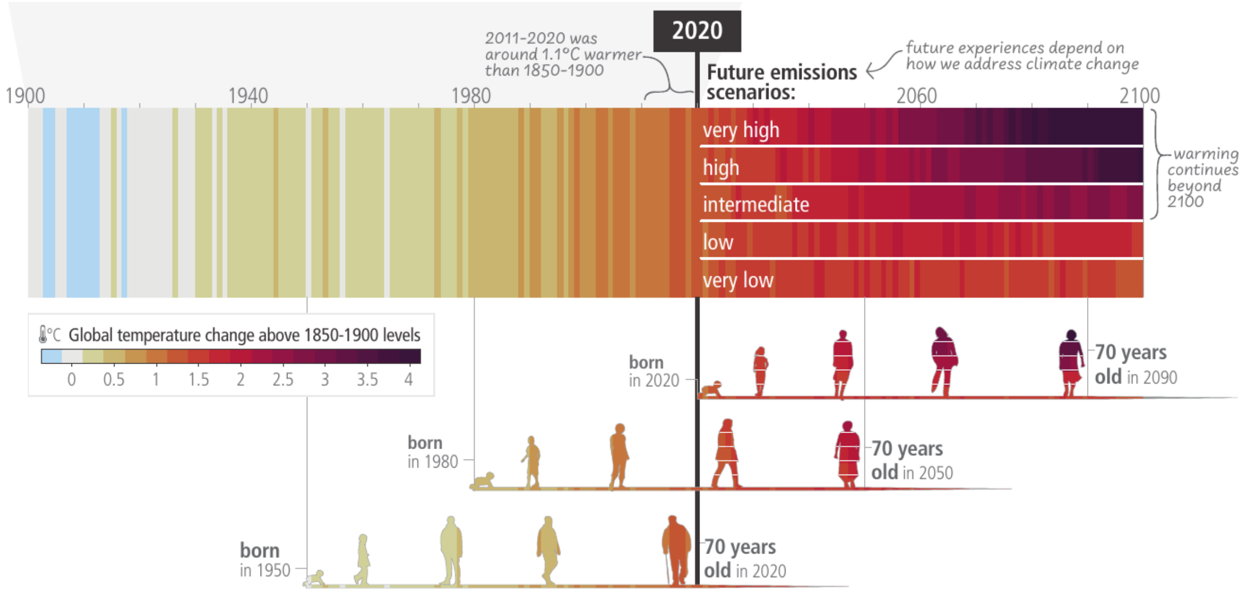

Figure 1.2: Observed and projected global surface temperature changes, from Lee et al. [16, Figure SPM.1].

by more than 1.1 °C from 1850. Even the most optimistic projections estimate the average temperature will increase beyond 1.5 °C in the next decade. Without sufficient mitigation policies, the 2.0 °C barrier will likely be broken well before the end of this century. Figure 1.2 graphically represents the climate change experienced by current and future generations under different emission scenarios.

Human civilization and most of the food sources we depend upon have developed in the last few million years. With temperatures and CO_2 concentrations rapidly increasing to levels unseen on Earth for tenths of millions of years, we cannot accurately predict how and if the ecosystem will adapt to such drastic changes. Mitigating the effects of climate change is, literally, a matter of life and death.

Limiting human-caused global warming requires net zero CO_2 emissions. Large investments and disruptive changes to the productive system are necessary to achieve this goal. Failure to act quickly will tighten our dependency from fossil fuel infrastructure and raise the costs to complete the transition. Despite the international agreements and national policies in place since the early 1990s, we are still falling short of the optimistic goal of not breaking the 1.5 °C barrier. Figure 1.3 highlights how the currently implemented policies will not lead to carbon neutrality by the end of this century.

A recent report from International renewable energy agency (IRENA) [18] breaks down the power generation mix under current plans and more aggressive renewable transition policies. Wind and solar power will share the largest share of energy production, regardless of the total amount

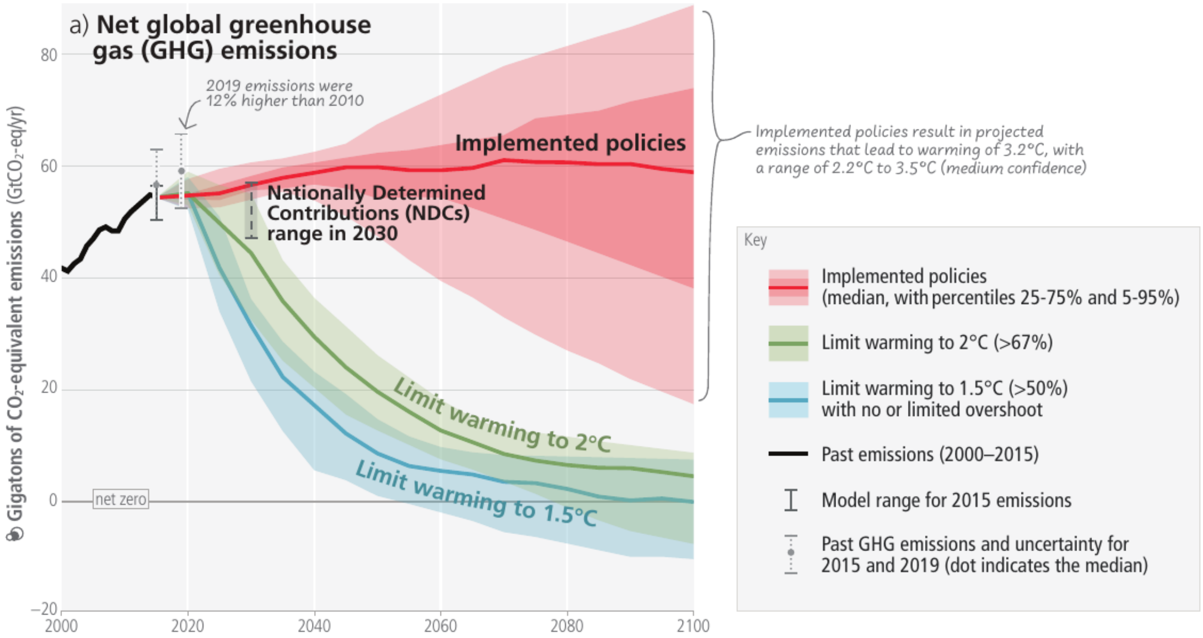

Figure 1.3: Emission pathways under different policy scenarios, from Lee et al. [16, Figure SPM.5].

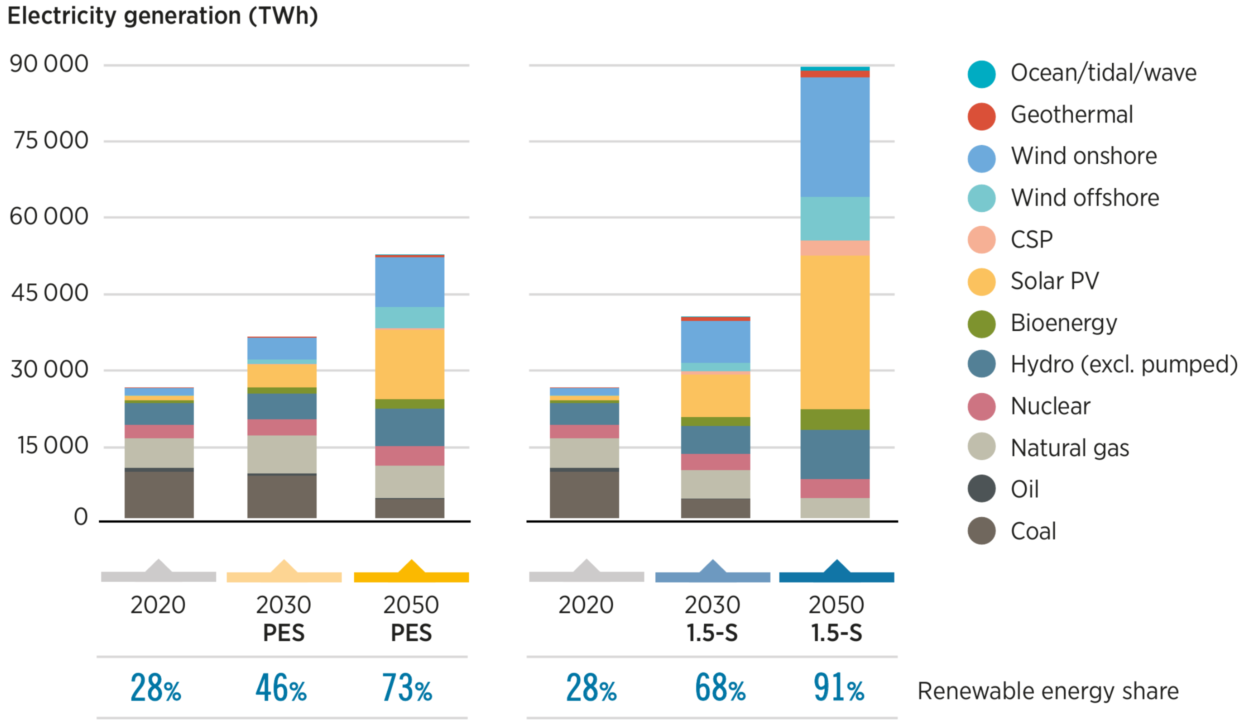

Figure 1.4: Power generation mix for the planned energy scenario (PES) and the 1.5 °C scenario (1.5-S), from IRENA [18, Figure 2.3].

of energy consumption and the specific scenarios. Figure 1.4 highlights the enormous scale of investments that will be required for new, clean renewable sources to be deployed in the next few decades.

Due to the intermittency and geographical availability of renewable energy sources, multiple technologies have to be pursued at the same time. Fossil fuel plants, given an appropriate supply infrastructure and terrain, can be built pretty much anywhere in the globe. Wind and solar power conversely have to "chase" their energy source. Different countries and regions have different resource availability. Northern Europe, for example, has a low solar per-square-meter irradiation but a large availability of wind throughout the year. Moreover, solar panels do not work at night or on cloudy days. Wind turbines conversely could work 24/7 but require a steady wind inflow within a certain speed range. Resilient energy grids thus need to combine multiple sources (and storage systems) to guarantee reliable electricity supply at all times. This paradigm shift is one of the biggest challenges of the transition away from fossil fuels.

The price of solar energy has been reducing at a fast pace in the last decades. While novel technologies can further reduce costs, footprint, and the dependency on hard-to-dispose raw materials, the technology is mature. Wind turbines conversely, despite their widespread installation, are getting closer to a technological bottleneck that limits their increase in size and reduction in cost. They already have some advantages compared to solar power. Their ground footprint is lower than solar energy for the same power and are in general more "energy dense" than solar counterparts. Moreover, they do not rely on potentially toxic or polluting raw materials like some solar panel technologies.

Wind systems however are not inherently carbon-free. The concrete and steel used for the tower and foundations use large amounts of energy and processes that rely on coal. At this scale, power and storage systems require large amounts of rare earth materials, whose extraction often leads to land abuse and social injustice in mining areas. Blade fibers and resins are challenging to recycle and dispose, and circular-economy approaches are still in early development phases [19]. Moreover, wind farms have an impact on wildlife and crops, and their development often encounters strong resistance from local communities.

While most of these challenges go beyond the scope of my dissertation, it is important to keep in mind the big picture scenario where technological and policy advancements concurrently enable systemic changes of the global industrial system. While wind energy has fewer drawbacks than conventional fossil fuels, a holistic socio-technical approach is required to make it effectively sustainable.

1.1.2 History and technical breakdown of wind turbines

The use of mechanical systems to harness power from the wind dates back several centuries. Historical sources indicate the presence of windmills in the area of modern Afghanistan before year 1000 A.D., and their expansion in China and Europe in the following centuries [5]. Far before the spread of fossil fuels with the second industrial revolution in the 1800s, windmills were used to mill grain, pump water, and other power intensive tasks replacing or supporting working animals. In a pre-industrial era with rudimentary machinery, they were one of the key means of production, playing a strategic economical role and being subject to heavy regulation.

Early windmills configurations predominantly had a vertical-axis, drag-based configuration and used simple wooden or cloth sails. Horizontal axis windmills were introduced and perfected in Europe over the Middle Age and Modern Era. Both "post" and "dutch" [6] windmills have aerostructural solutions to align the rotor to the wind and ensure operation continuity regardless of the inflow direction [5, Sec.1.1-1.2]. Compromises between costs, maintainability, and operability were already driving the design and installation of these early machines.

Between the 18th and early 20th centuries, more rigorous mathematical approaches enabled a deeper understanding of the flow physics, propelling the development of new technologies and more powerful systems. By the 1920s, the theory of lift by Prandtl and Betz (mentioned later in Section 1.1.4) was established. Around the same years, the first industrial efforts to generate electrical power from the wind were initiated.

Early attempts to connect dynamos to windmills date back to the 19th century. However, modern wind energy production was established only after WWII and the energy crisis in the 1970s. This effort to reduce dependency from fossil fuels was driven primarily by economical and geopolitical issues, rather than from widespread awareness of human-made climate change.

In the 1980s, the development of large experimental rotors and the first multi-megawatt wind farms established horizontal-axis wind turbines as the dominant configuration in the market. One of the main advantages of this configuration is the possibility to control both rotation speed and pitch for power extraction and protection under extreme scenarios or system faults [5, Sec.1.4]. Moreover, the technological expertise in lifting devices and propeller design can easily be transferred to wind turbine design, boosting technological improvements.

As of the early 2020s, utility-scale wind turbines are rated for powers up to 16 MW [7]. Industry trends [4] and academic studies [20] predict even larger and more powerful systems to be developed in the upcoming years. A breakdown of a modern wind turbine is shown in Figure 1.5. Wind turbine components can be grouped into mechanical, electrical, and control subgroups [21], whose

---

[6] Windmill types from Encyclopedia Britannica, last accessed Nov 2023

[7] Goldwind GWH252-16MW Offshore Wind Turbine Successfully Installed, last accessed Nov 2023

characteristics are dissected in the next paragraphs.

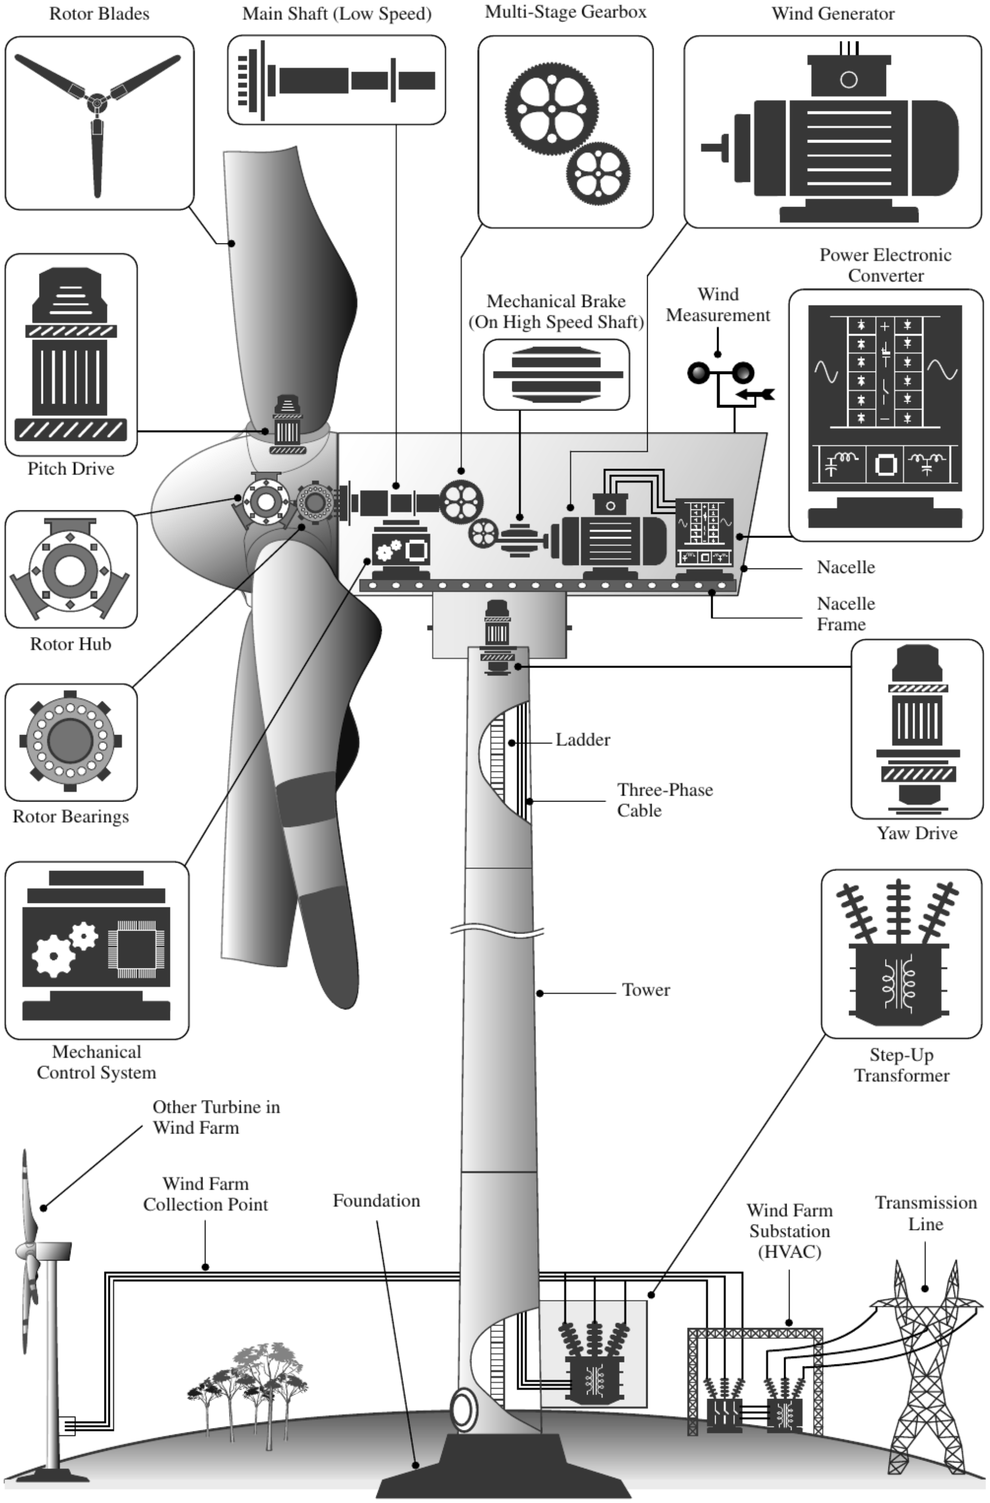

Figure 1.5: Component breakdown of an industrial-scale horizontal wind turbine. From Yaramasu and Wu [21, Figure 1.12].

Mechanical components

Rotors, the main focus of this dissertation, are the primary components that convert the wind kinetic energy into mechanical energy. Considering inertia and force balance, three-bladed solutions are the most common due to the compromise between symmetrical loading and rotor solidity.

Modern rotor blades make large use of composite materials to further reduce the overall mass and inertia of the rotating system. Lighter blades have a beneficial cascade effect on the sizing of the main bearings and shaft. Low rotation speeds reduce the loads on the mechanical components inside the hub. Conventional configurations use gearboxes to increase RPM at the generator by up to a factor 120. However, larger new turbine designs are shifting to direct drive generators to reduce costs, weight, and increase reliability [22].

At the same time, the size of the rotor influences directly the tower and foundation sizing. These components, usually made of steel, concrete, or both, make the largest use of raw materials and drive the capital investment costs. Novel concepts such as floating platforms have additional stability requirements that also increase the mass and upfront costs.

Electrical components

Connected to the rotor through a the main shaft, the generator housed inside the hub at the top of the tower transforms the main shaft rotation into alternate current. Several technologies have been used at industrial scale in the last decades. Synchronous, variable-speed generators are preferred for direct drive transmissions because they can operate at lower rotation regimes [21]. Permanent magnet synchronous generator (PMSG)s have a high power density and efficiency, but are generally heavy and require expensive rare earth materials. High-temperature-superconducting synchronous generators (HTS-SGs)s are smaller and lighter but are a less mature technology, require more complex cooling systems to ensure superconductivity, and have higher overall costs. Converters, transformers, and power cables also need to be accounted for in design packaging and cost estimation.

Controls

Finally, control systems include pitch and yaw actuators, mechanical brakes, and sensors, including (on advanced models) LIDAR systems for gust and turbulence detection and mitigation [23]. Wind turbine control is a vast research field that primarily focuses on robust strategies to maximize power extraction and enhance reliability. We neglect control-related considerations in this work since we focus on steady state analyses.

Benchmark models

Benchmark models such as the DTU 10 MW [24] (used in this dissertation) and the more recent IEA 15 MW [22] embed the main characteristics of modern turbines. They both have an upwind, three-bladed, horizontal axis configuration, with rotor diameters in the order of 200 m installed at more than 120 m from the ground or sea level. Both designs are classified as category I-B according to the IEC standards [25]. This categorization refers to turbines installed in sites with annual average wind speed of at least 10 $\mathrm{m\,s^{-1}}$, with turbulence levels up to 16% and wind gusts up to 70 $\mathrm{m\,s^{-1}}$. They use a mix of RPM-based control below rated power, and active pitch control at their rated generator power. The more recent IEA 15 MW configuration is intended to use a direct drive generator rather than a more expensive and complex geared drivetrain.

As mentioned above, power extraction and manufacturing costs depend on the design and integration of several electromechanical components. Design features aimed at increasing power harvesting often are in contrast with cost-saving solutions. In the following Section 1.1.3, we briefly discuss the factors contributing to the cost-of-energy from wind systems.

The present dissertation work focuses on rotor design. Other turbine components are assumed fixed over the optimization and are neglected for simplification purposes. The tools we developed need to be integrated in a more holistic design cycle in order to effectively deliver innovative designs and make the next generations of wind turbines more affordable and reliable.

1.1.3 The cost of wind energy

As summarized in Section 1.1.2, utility scale wind turbines are complex electromechanical machinery. They require a large amount of raw materials, lengthy design processes, and expensive installation and maintenance operations. Despite the "fuel" being freely available, wind energy systems have considerable upfront investment costs. Industrial power plants must be carefully planned to ensure the energy yield is maximized over their operating life and the initial investment is fully recovered.

The cost of energy depends on both the energy yield and the capital and operating costs [26]. Lower costs and higher yield lead to lower energy market price, making renewable energy sources competitive compared to the fossil fuel alternatives. Profitability drives more investments, which are needed to ultimately decarbonize the economy. Aside the obvious social benefits of reducing emissions (discussed in Section 1.1.1), lower energy costs also have direct social justice implications. Profits hardly ever trickle down, but costs almost always do. Affordable clean energy leads to cheaper utility bills and decrease the burden on lower income families.

Costs

Manufacturing cost estimations are usually grounded in the preliminary mass sizing of the components listed in Section 1.1.2. Capital expenditures also depend on the cost of raw materials used in production systems. The rotor blades are usually made of fiberglass composites, while the generators require more expensive rare earth metals. However, the large majority of components are made of steel, whose price has a strong impact on capital costs [11].

Taxes and interests on these capital intensive investments have a sizable impact on the cost breakdown [26, Ch.8, Fig.6]. Variable operating costs depend on design, environmental conditions, infrastructure,and site accessibility. For example, harsh environments increase the failure rate and blade performance degradation. One of the main challenges of new concepts like floating wind farms is to incorporate maintenance and deployment considerations early in the conceptual design of platform configurations [27].

Market and interest rate fluctuations, together with supply chain issues, affect mid- and long-term cost predictions and affect the market prices. Site characteristics affect the cost of power distribution lines and grid integration. Regardless of these fluctuating parameters, the component and installation costs per kW generally decrease with the turbine size [26]. This has fostered the development of increasingly large wind energy systems [4].

Yield

The energy production depends both on the system efficiency, reliability, and the resource intermittency. Conversely to conventional power plants, the generator rotor power is not the main indicator for for energy yield. Rotor diameter and its installation site meteorological conditions are a more useful metric [5]. Designers aim at maximizing the turbine output over the broad operational range of the system. The annual energy production (AEP) is a common metric used to quantify the energy output of a single turbine or a power plant.

Combining the turbine design-specific power curve and the site-specific wind characteristics, developers can estimate the energy yield over a given time frame. The role of site-specific considerations in design is discussed more in detail in Section 1.3.1. Wind farm wake interactions affect the layout of power plants and the design of control strategies, but their effect can be usually quantified during the design stage.

Design uncertainty and other stochastic parameters, such as blade soiling, impact the turbine energy production in a less predictable way. Resource intermittency brings grid integration and short-term demand challenges, but it is usually included in long-term yield estimation through meteorological assessments. Maintenance and other unplanned down times reduce the theoretical capacity factor of a wind turbine over its lifetime.

Levelized cost of energy

The levelized cost of energy (LCOE) is a practical metric that aggregates AEP and costs. It has been used in the last few decades to select the best combination of sites and technologies for plant development. LCOE depends on market variability and site characteristics [11]. It is also a useful parameter to compare different technologies. Myhr et al. [28] for example performed a life cycle assessment of different floating turbine concepts. They identified a range of site and operational conditions where floating systems can be competitive with bottom fixed counterparts. This analytical metric also facilitates sensitivity analysis during the conceptual design stage.

The technological advantages and market development have lead to a substantial decline of wind LCOE. Taller, larger, and more powerful turbines have concurrently increased the energy yield and reduced the cost-per-kW [3]. Veers et al. [4] predict that these trends will continue in the near future, highlighting the need for novel analysis and design tools to prevent LCOE stagnation for the next generations of wind energy systems.

Recent works are exploring novel cost of energy metrics beyond LCOE. Dykes [29] claims that lifetime profitability estimation must include market-based metrics. With larger shares of the market dominated by intermittent renewable resources, its dynamics have a feedback effect on the value of energy. This in turn affects wholesale prices, curtailment policies, and the overall market capacity. Related research by Simpson et al. [30] has proposed Levelized Avoided Cost of Energy simplified (LACEs) and Cost of Valued Energy (COVE) as alternative metrics to LCOE. Their work shows how these metrics can capture energy value trends better than conventional approaches.

Alternative production and storage approaches are also active research fields. Hybrid power plants can mitigate resource intermittency and losses due to curtailment by combining wind and solar systems with battery and hydrogen-based storage systems [31]. The growth of this relatively novel energy production concept will go hand in hand with the rise of hydrogen-based economy [8]

Despite the wind cost of energy kept decreasing through the early 2020s, as shown in Figure 1.6, reliability and supply chain issues have been plaguing the industry during and after the Covid-19 pandemic [9]. Recent auctions for wind energy projects in the US and Europe have seen lack of bidders due to minimum price supports that hindered profitability [10] These road bumps are hopefully temporary. However, it is evident that strong public policies [32] and resilient productive systems are necessary to support the growth of renewable energy sources. Without them, the goals outlined in Section 1.1.1 are not reachable.

---

[8] Gov. Whitmer Announces Michigan Wins Funding for Clean Hydrogen Hub, State of Michigan website, October 13 2023.

[9] Wind turbine troubles have sent one stock tumbling, CNBC, July 3 2023

[10] No Bidders in British Offshore Wind Auction, New York Times, Sept 8 2023

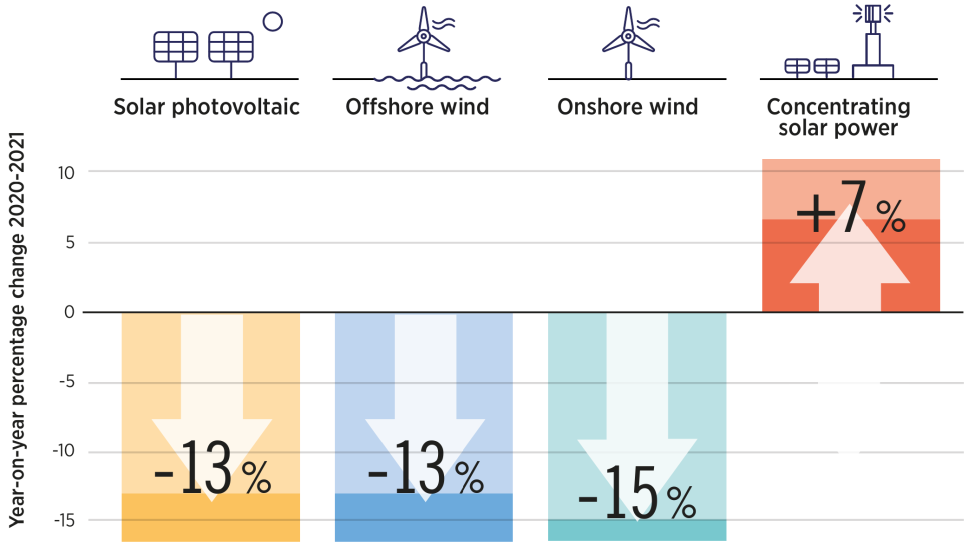

Figure 1.6: Global changes in LCOE for solar and wind technologies in 2020-2021, from IRENA [18, Figure 2.2].

Many of the factors contributing to LCOE are not considered in this work. Design optimization efforts hereby presented are focused on a single component and on a subset of operating conditions. However, we consider high-level yield and cost parameters in our optimization studies. Rotor torque acts as a proxy of rotor efficiency. The mass contribution to the objective function and the thrust constraint account for (or limit) turbine costs. More details on the performance metrics we use are outlined in Section 4.3. The next section outlines some key physical concepts behind wind power harvesting.

1.1.4 The physics of wind energy

A wind turbine is an electromechanical system that converts kinetic energy from the wind into mechanical power and, subsequently, alternate electrical current. Throughout history (see Section 1.1.2), a wide range of concepts have been explored, not always resulting in practical designs. Regardless of their configuration, efficient wind turbines leverage aerodynamic lift to extract energy from the wind. An overview of different concepts goes beyond the scope of this chapter, since the optimization studies we perform focus on the horizontal-axis configuration that dominates the utility-scale market. In this section we briefly describe the physical principles behind “propeller-like” wind turbines, although the same concept can be extended to vertical-axis or more exotic configurations.

Momentum theory

Elementary momentum theory provides a first-order understanding of the energy conversion process. Albert Betz is credited to first applying these principles to wind turbines, identifying a theoretical maximum amount of energy that can be harvested from a moving mass of fluid [33].

The rotor can be considered as an ideal permeable disk in a 1D model of the velocity and pressure along the flow streamlines. The rotor acts essentially as a drag device [34], slowing down the flow of the streamtube and introducing a pressure drop at its rotation plane, as shown in Figure 1.7. Assuming the flow is in the incompressible regime, the velocity decreases without discontinuities.

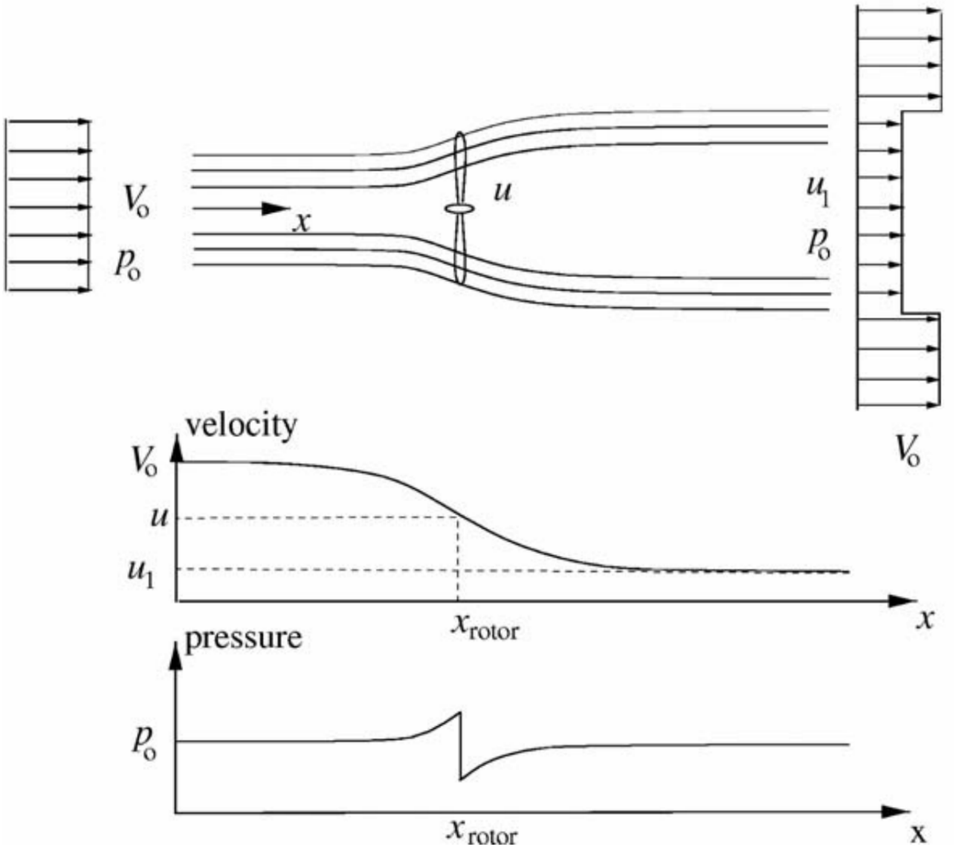

Figure 1.7: Streamtube, velocity, and pressure trends of an air mass flowing through a wind turbine rotor, from Hansen [34, Fig.4.1].

Considering the kinetic power as the amount of energy that flows at velocity v over section A per unit time, and mass flow rate as $\dot{m} = \rho v A$

$$P_{\text{kin}} = \frac{1}{2}\dot{m}v^2 = \frac{1}{2}\rho A v^3, \tag{1.1}$$

the balance between kinetic and extracted mechanical power is given by:

$$P_{\text{mech}} = P_{\text{kin}_u} - P_{\text{kin}_d} = \frac{1}{2}\rho \left(A_u v_u^3 - A_d v_d^3\right) \tag{1.2}$$

with A the streamtube area, v the local flow velocity, and u and d subscripts indicating upstream and downstream sections respectively. From the mass flow continuity in the streamtube:

$$v_u A_u = v_d A_d \tag{1.3}$$

Substituting Eq. 1.3 into Eq. 1.2, and using the mass flow rate definition, we obtain that the power depends on the flow (squared) velocity difference between sections upstream and downstream of the turbine:

$$P_{\text{mech}} = \frac{1}{2}\dot{m}\left(v_u^2 - v_d^2\right) \tag{1.4}$$

Since $v_d = 0$ would not be physically possible (wind would be completely stopped!), Betz equated this energy-based power expression with a force-based power expression:

$$P_{\text{mech}} = F v_r = \dot{m}\left(v_u - v_d\right) v_r \tag{1.5}$$

with F the rotor thrust given by the momentum change (see Hansen [34, Eqs.4.5–4.10]) and v_r the flow velocity at the rotor plane. Equating Eq. 1.2 and Eq. 1.5, the velocity at the rotor plane becomes the average of the upstream and downstream velocities $v_r = (v_u + v_d)/2$.

Defining the power coefficient as the ratio between mechanical power and total kinetic wind power at the rotor plane, and substituting the expression for v_r to get:

$$P_{\text{mech}} = (1/4)\rho A_r \left(v_u^2 - v_d^2\right)\left(v_u + v_d\right), \tag{1.6}$$

we finally end up with:

$$c_P = \frac{P}{P_0} = \frac{\frac{1}{4}\rho A_r \left(v_u^2 - v_d^2\right)\left(v_u + v_d\right)}{\frac{1}{2}\rho A_r v_u^3} = \frac{1}{2}\left[1 - \left(\frac{v_d}{v_u}\right)^2\right]\left[1 + \frac{v_d}{v_u}\right] \tag{1.7}$$

The maximum achievable power coefficient is $c_P = 0.593$, obtained with $v_d/v_u = 1/3$, is called Betz's limit.

Forces on the rotor blades

Momentum theory gives a high-level understanding of the energy extraction process. However, we need to look at the force distribution over the blades to understand how to design more efficient turbine rotors. Lift-based wind turbines extract mechanical power from the in-plane component of aerodynamic lift over the blades. A breakdown of the wind velocity and resulting forces components is shown in Figure 1.8.

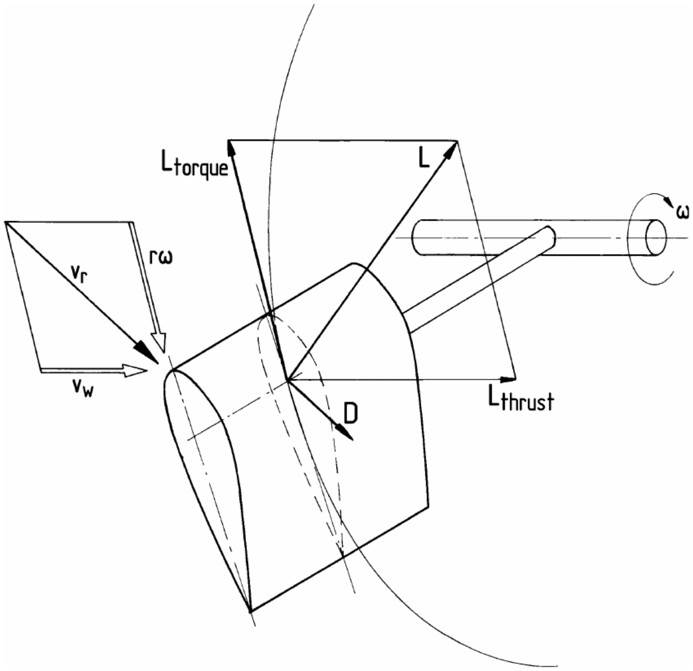

Figure 1.8: Schematics of the velocity triangle over a representative blade section, from Hau [5, Fig.4.6].

We assume horizontal wind direction for sake of simplicity in this explanation. At each spanwise location, the inflow angle of attack is given by the combination of the horizontal wind velocity and the tangential rotation speed $V = \omega r$, with ω the blade rotation rate and r the local radius. The resulting lift has an in-plane component, contributing to the driving torque, and a out-of-plane component, building up the rotor horizontal thrust.

BEMT analysis tools combine momentum theory with blade section properties at discretized radial sections over the rotor span. Balancing the two models leads to radial tangential and axial force distributions, used to calculate the rotor performance and structural loads on the blades. These analysis software are included in state-of-the-art holistic turbine analysis frameworks due to their ease of implementation and low computational cost. Existing tools include corrections for tip loss, wake vorticity, and unsteady flows. We redirect the reader to resources such as Moriarty and Hansen [35] for a comprehensive description of BEMT theory and implementation

From Figure 1.8 it is clear that the shape, size, and pitch angle of the blade determine the load distribution over the rotor. Shape optimization thus has a key role in increasing the efficiency of a wind turbine. Higher local L_{torque} increases the power extracted by the turbine at the same rotation rate. Reducing L_{thrust} leads to lower structural loads, and thus lower component mass and cost.

Moreover, aerostructural tradeoffs can be exploited to improve the design of the blade. The aerodynamic and structural response are tightly coupled. The aerodynamic loads determine not just the energy yield, but the structural sizing of the rotor. As discussed in Section 1.1.2, a heavier rotor increases capital and operating costs. Conversely, the structural response of the blade has a large impact on the aerodynamic performance. Blade deflection under loading can be designed to obtain more favorable local angles of attack for the same sectional airfoil shape, or increase the aeroelastic stability of the system.

It is not possible to design the turbine aerodynamics without structural considerations, and vice-versa. The interaction between aerodynamics and structures is sometimes very subtle to capture. Traditional, sequential design strategies might not fully take advantage of such a tight coupling. The goal of this dissertation work is to tackle this design problem using high-fidelity design tools, fitting in a larger research thrust discussed in Section 1.3. In the next section, we introduce the concept of multidisciplinary design optimization, which will be leveraged for wind turbine rotor design in this dissertation.

1.2 System design through multidisciplinary design optimization

Modern, complex engineering products such as aircraft and wind turbines exploit multi-component and multiphysics tradeoffs to reduce costs and increase system efficiency. Advanced numerical analysis tools have been developed in the last few decades, supported by hardware developments that cut computational wall time. These tools are used by engineers to gain insight on the system behavior while varying a range of design parameters.

Despite these increasingly refined simulation tools, product development processes have evolved at a slower pace. Conventional design approaches still rely on lengthy manual iterations between different discipline-specific teams to converge to the final system design. Multidisciplinary optimization strategies can facilitate the design of increasingly complex systems with multiple components. Conversely to conventional approaches that sequentially iterate between sub-systems, MDO considers all the components simultaneously - as illustrated later in this section. MDO formalizes the coupling between these different components and couples the multidisciplinary system analysis with mathematical optimization algorithms.

Optimization is often an abused term. In common language, it describes a general performance “improvement” or the specific “tailoring” of a product of procedure for a specific task. In this work, we refer to mathematical definition of optimization discussed by Martins and Ning [36].

Conventional and optimization-driven design workflows are shown in Figure 1.9. Design optimization proposes to replace human-centered low-level decision-making with a deterministic mathematical process. Given a formal optimization problem, numerical algorithms update the design parameters until an optimum configuration is identified. Correctly implemented MDO strategies return the best achievable configuration according to physical analysis and system parametrization.

MDO must not be intended as a "push-button" solution. Human intervention is still required to define and update the problem formulation and assess that the final system configuration effectively meets the intended design goals.

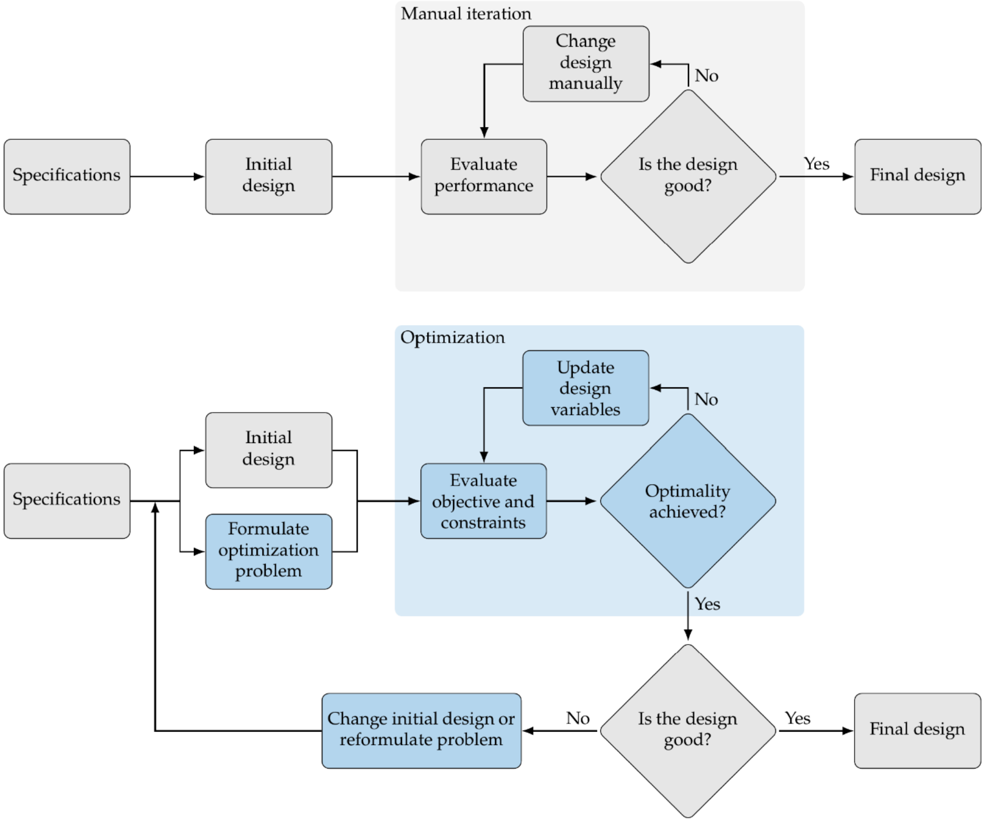

Figure 1.9: Traditional and optimization-driven design approaches, from Martins and Ning [36, Fig.1.2].

Pironneau [37], Haftka [38], and Jameson [39] pioneered the application of MDO and numerical optimization for aircraft design. In the last decades, MDO has slowly found application in an increasing number of engineering fields, including wind turbine engineering (Section 1.3) Transitioning from conventional to optimization-based workflows, despite growing adoption of general frameworks such as OpenMDAO [40] and GEMSEO [41], comes however with a heavy overhead implementation cost.

Nevertheless, MDO has several advantages, some of them hinted above. Leveraging simulation-based design metrics and sensitivities enables the identification of unconventional designs based on tradeoffs that conventional approaches fail to capture. By removing the human intervention on the iterative design updates, the entire workflow is accelerated. Both conceptual and detailed design studies can benefit from reduced turnaround times and fully automated decision making. Finally, the engineers energies and expertise can be completely focused on the more creative aspects of system engineering design, relieving them from the tedious and repetitive steps of the design process.

Stemming from system specifications, the optimization workflow requires the formal definition of an objective, constraints, and design parameters (or "variables").

The design variables are the parameters that describe the system and are updated during the optimization loop. The choice of variables can be arbitrary, although an educated choice must be made to avoid convergence issues. For wind turbine aerostructural design in this dissertation, we use structural thicknesses and composite ply angles, and a range of planform and local section geometry variables.

The objective function is a metric used to identify the best design. It is usually a scalar function of the design variables and the system state. The function selection must be consistent with the practical engineering goals to return meaningful designs. Cost-of-energy metrics, discussed in Section 1.1.3, are a popular choice in energy system engineering. In this work, we use torque and mass as a proxy for the value of energy.

Finally, constraints are used to bound the design space to ensure the optimization process returns practical designs. For rotor design, we constrain structural stresses and deflections, aerodynamic loads and, in Chapter 8, life-cycle equivalent damage and extrapolated extreme loads. More details on the formal problem definition for this work are presented in Section 4.3 and Chapters 5 to 7 for each individual study.

Depending on the system properties, problem formulation, physical models, and design variables, the design space explored by the optimizer might have different properties. Optimizations can turn into a challenging nonlinear problem, sometimes with multiple possible optimal solutions.

The considerations drawn so far apply to general design optimization problems. Multidisciplinary optimization has the additional challenge of modeling and coupling different components, increasing the implementation effort and problem complexity. Except for trivially coupled systems, practical engineering systems are composed of multiple sub-systems affecting each other during operations and, consequently, during the design stage.

Figure 1.10 shows a notional multi-component, tightly coupled system. Each block represents different analysis tools or physical systems, coupled to each other though a flow of input and output data. For example, in a simplified aircraft design problem, r_1 could be an aerodynamic

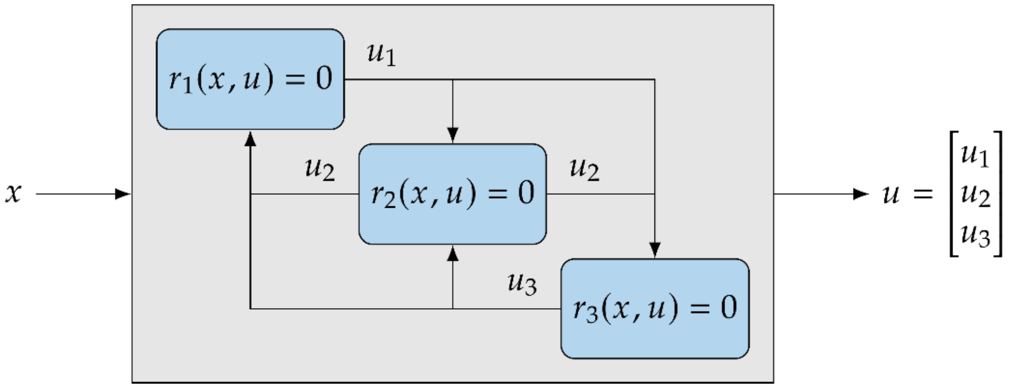

Figure 1.10: An example of tightly coupled multidisciplinary system, from Martins and Ning [36, Fig.13.2]. Each block is a discipline (e.g. an analysis tool) connected to other blocks of the system through a flow of input and output data.

tool, r_2 a structural model, and r_3 a fuel consumption or weight model. In that case, aerodynamic loads and drag affect both the structural design and fuel consumption. The structural analysis affects both the aerodynamic shape and performance at operating conditions, and the weight of the aircraft. Changing the weight and fuel volume affect the internal structure and aerodynamic lift requirements. These blocks are tightly interconnected and must be analyzed together to obtain a feasible and well designed system.

A holistic multidisciplinary model has the benefit of capturing mutual interaction between components and including broader design consideration, such as economical and human factors. Moreover, the design variables affecting different components are optimized simultaneously – conversely to conventional sequential design approaches that might converge to sub-optimal solutions [36, Fig.13.2]. Correctly coupling multidisciplinary models requires a careful implementation effort. However, from a mathematical standpoint, a coupled model is just a larger system of equations to be solved.

Several MDO architectures have been used in academic and industrial scenarios. Distributed approaches allow for more flexible work compartmentalization and intellectual property management. Monolithic strategies however guarantee faster and more robust convergence [10]. Among a range of monolithic architectures, MACH implements the multidisciplinary feasible (MDF) approach. MDF solves the tightly coupled multiphysics model at every optimization iteration, before moving to the following design step.

1.2.1 The benefits of gradient-based optimization

Gradient-based algorithms fall under the “local” optimization categorization [36, Sec.1.4.2]. They use sensitivity information from the analysis tools to search the design space and advance towards the optimal design. Compared to global approaches that explore the full available design space, these algorithms provide a potentially optimization-path-dependent solution. Moreover, the design space must be smooth, i.e. without discontinuities that would prevent the optimizer to converge. Their implementation cost is also higher compared to simpler gradient-free strategies.

Nevertheless, by using higher-order information, gradient-based schemes can traverse the design space quicker and converge with fewer design evaluations. Their computational efficiency makes them the only practical solution for large-scale design optimization problems [36]. As shown in Figure 1.11, gradient-based algorithms scale better with increasing design variables compared to gradient-free approaches. However, the way gradients are computed has a large impact on the optimization robustness and speed.

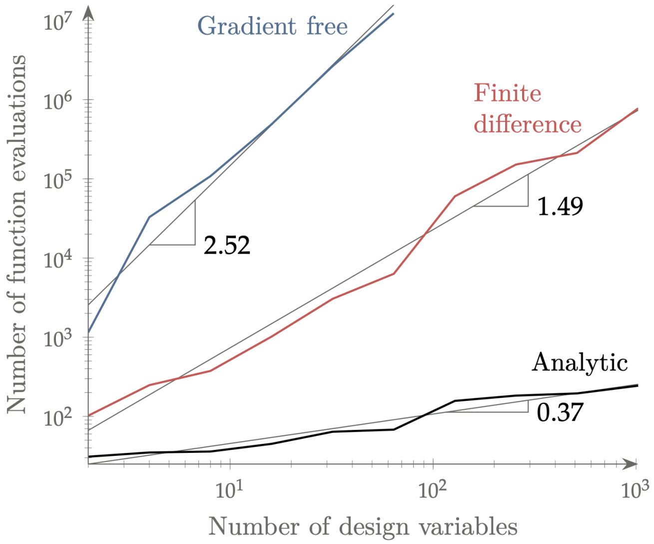

Figure 1.11: Gradient-based optimization strategies with analytic sensitivities are highly scalable and outperform gradient-free and finite-difference-based approaches, from Martins and Ning [36, Figure 7.1].

Finite-difference derivatives are easier to implement as they only require additional system analyses using a perturbed design vector. For this reason however, their computational cost scales linearly with the number of design variables. Moreover, they are subject to numerical truncation errors that might slow or prevent optimization convergence. Analytical approaches are more involved but have been proven to have excellent accuracy and scalability properties.

The adjoint method

MACH leverages an efficient adjoint solver coupled with automatically-differentiated code base to obtain the design sensitivities of the coupled aerostructural system. Kenway et al. [12] present the tool implementation in more detail. In this section, we briefly describe the general adjoint approach and its benefits.

In general, the derivatives of a function f with respect to the full set of design variables x can be expressed via the chain rule as:

$$\frac{\mathrm{d}f}{\mathrm{d}x} = \frac{\partial f}{\partial x} + \frac{\partial f}{\partial u}\frac{\mathrm{d}u}{\mathrm{d}x} \tag{1.8}$$

where u is a vector of system states. The Jacobian $\mathrm{d}f/\mathrm{d}x$ has size $n_f \times n_x$. Using a residual form to express the system components [36, Sec.6.7.1], we assume that the system residual is $r(x, u) = 0$ even after a small perturbation of the design variables. Expanding and rearranging the residual gradient we obtain:

$$\mathrm{d}r = \frac{\partial r}{\partial x}\mathrm{d}x + \frac{\partial r}{\partial u}\mathrm{d}u = 0 \quad \rightarrow \quad \frac{\partial r}{\partial u}\frac{\mathrm{d}u}{\mathrm{d}x} = -\frac{\partial r}{\partial x} \tag{1.9}$$

By manipulating Eq. 1.9 and substituting $\mathrm{d}u/\mathrm{d}x$ in Eq. 1.8, we obtain the system in Figure 1.12, which conveniently shows the size of the matrix components of the system.

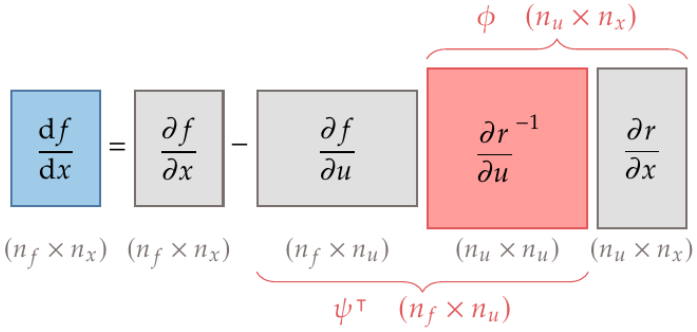

Figure 1.12: Matrix size for components of the adjoint equation, from Martins and Ning [36, Figure 6.27].

Using the adjoint approach, we first solve the linear system for ψ:

$$\frac{\partial r}{\partial u}^T \psi = \frac{\partial f}{\partial u}^T \tag{1.10}$$

one column at a time. Assuming that the partial derivative terms $\partial f/\partial x$, $\partial f/\partial u$, $\partial r/\partial x$, and $\partial r/\partial u$ are available from the analysis tools, the remaining implementation challenge is to use fast and memory-efficient linear solvers.

From a computational standpoint, each adjoint solution consists of n_f linear solutions. Conversely, the direct approach (solving for Φ in Figure 1.12) requires the solution of n_x linear systems. The adjoint approach thus is particularly convenient when $n_f \ll n_x$, as it is the case for large-scale high-fidelity MDO problems. Once the adjoint matrix ψ is obtained, the total function sensitivities $\mathrm{d}f/\mathrm{d}x$ are obtained from Eq. 1.8.

As mentioned earlier in this section, from a mathematical standpoint coupled systems are essentially just larger systems of equations. Numbering the components of a general multidisciplinary system from 1 to n, the adjoint system in Eq. 1.10 becomes:

$$\begin{bmatrix} \frac{\partial r_1}{\partial u_1}^T & \cdots & \frac{\partial r_n}{\partial u_1}^T \\ \vdots & \ddots & \vdots \\ \frac{\partial r_1}{\partial u_n}^T & \cdots & \frac{\partial r_n}{\partial u_n}^T \end{bmatrix} \begin{bmatrix} \psi_1 \\ \vdots \\ \psi_n \end{bmatrix} = \begin{bmatrix} \frac{\partial f}{\partial u_1}^T \\ \vdots \\ \frac{\partial f}{\partial u_n}^T \end{bmatrix}. \tag{1.11}$$

Finally, the Jacobian of the coupled system is given by:

$$\frac{\mathrm{d}f}{\mathrm{d}x} = \frac{\partial f}{\partial x} - \left[\psi_1^T \dots \psi_n^T\right] \begin{bmatrix} \frac{\partial r_1}{\partial x} \\ \vdots \\ \frac{\partial r_n}{\partial x} \end{bmatrix} \tag{1.12}$$

The components of MACH are presented later in Chapter 2. To better frame the research niche explored in this work, an overview of the state-of-the-art for wind turbine design optimization is presented in the next Section 1.3. The high-fidelity gradient based optimization capabilities of our tool enable optimization studies with a smaller scope but a higher level of detail than what can be obtained with conventional approaches. Large scale studies like the ones presented in this dissertation are unprecedented in the field.

1.3 MDO and wind energy

This chapter so far has discussed the characteristics of wind energy systems and the fundamentals of MDO In this section we give an overview of the design studies and MDO applications to wind turbine systems available in the scientific literature.

Despite ongoing efforts to shift the focus to demand-based value metrics, mentioned in Section 1.1.3, the LCOE remains the primary design driver for wind energy systems [3]. Veers et al. [4] outlined the challenges and unknowns for future wind turbine systems design with competitive LCOE. From a design perspective, growing rotor sizes are particularly pushing conventional tools to their application limit. Consequently, new high-fidelity solvers and design strategies are necessary to develop innovative designs. Numerical optimization can help engineers to exploit design tradeoffs and shorten the design cycle.

Several tools have been developed to design wind turbines over the last decade. State-of-the-art software, such as OpenFAST [42], Cp-Max [8], HAWTopt2 [43], ATOM [44], SHARPy [45], Qblade [46], and MoWit [47], simulate the full turbine system, including aerostructural rotor and tower models, power generators, and fixed-bottom foundations or floating platforms. Linear and nonlinear controllers are also included for generator and pitch controller development, adding more practical considerations to the modeling of the turbine operating behavior. Such a holistic approach is necessary to capture the system's multibody dynamics and the unsteady aerostructural phenomena, which affect the performance and drive the sizing of a turbine over its life cycle. We refer to these tools as *conventional* to distinguish them from the framework presented in this dissertation.

These conventional tools combine models of variable fidelity to analyze different system components. OpenFAST [42] has set the standard for turbine analyses, providing frequency- and time-domain aero-hydro-servo-elastic analysis capabilities, including all the main turbine components. OpenFAST models the structure using the BeamDyn and ElastoDyn modules, which use beam models to capture the turbine dynamic response. However, these modules are limited in terms of degrees of freedom (ElastoDyn) and accuracy of local stress values (ElastoDyn and BeamDyn). OpenFAST's aerodynamic predictions through BEMT-based AeroDyn agree well with experiments for conventional designs. However, BEMT assumptions become increasingly poor for larger blades with large deflections and more significant 3D effects [48, 49, 50].

The blade-resolved rotor model used by Madsen et al. [15] and in the present work neglects the hub geometry. However, the CFD analysis captures the spanwise and tip flow phenomena with a more significant impact on in-plane and out-of-plane loads. Horcas et al. [48] investigated how tools with different modeling fidelity predict the impact of curved tips on turbine performance. They showed that HAWC2, a software based on the same physical models of OpenFAST, overpredicts the loads on a 10 MW turbine at below-rated wind speeds and underestimates the benefits of curved wing tips, even for a steady power curve. Validation campaigns highlighted how different BEMT model implementations affect the simulation results [49]. Moreover, tools like OpenFAST are limited when used with numerical optimization because they cannot compute accurate gradients. Furthermore, these tools use precomputed wing section drag polars and thus cannot optimize the airfoil shapes unless they are coupled to an airfoil solver.

Several efforts combined high- and low-fidelity tools to balance accuracy and computational cost. Ramos-García et al. [51] coupled finite-element structural models with a hybrid lifting-line-vortex method in a multibody aero-hydro-servo-elastic framework, showing good agreement with Heinz et al. [52]. This effort addressed some limitations of BEMT at a lower cost than blade-resolved solvers. However, applying such tools is primarily limited to analysis because computing derivatives remains challenging. Design frameworks based on surrogate models [53] and artificial neural networks [54] are promising approaches to reduce optimization computational time. However, the metamodel generation and design space sampling usually has high computational overhead costs [55]. Such strategies are substantially different from the simulation-based approach we present in this work, and we omit more detailed discussion on the topic in this manuscript.

Other recent efforts focused on a narrower but higher-fidelity approach to capture the complex fluid-structure interaction of the turbine blades [56]. Wainwright et al. [57] combined a commercial CFD software and a finite-element modal solver with graphics processing unit acceleration to study the aerostructural behavior of a rotor in the wake of an upstream turbine. They used a strong coupling strategy to resolve time-accurate analyses, using a radial basis function to pass load-transfer information between the solvers. Cheng et al. [58] validated a model that couples an actuator line model, an aerodynamic RANS solver, and a hydrodynamic CFD tool in OpenFOAM to study the behavior of a floating wind turbine. While BEMT and beam models have limited accuracy for complex flow conditions, high-fidelity codes are limited by their computational cost and implementation complexity. Hence, despite the promising results, none of these high-fidelity models have been used for optimization.

Some work has applied numerical analysis tools to wind turbine optimization problems. Ning et al. [59] identified how sequential design optimization strategies and a focus on the sole energy production lead to sub-optimal designs. They used a ratio of mass over AEP as objective, where the latter metric is obtained based on a Rayleigh wind distribution. Including mass in the objective optimization can be used as a proxy for costs, which are driven by the price and quantity of raw materials for each component of the turbine.

Gray et al. [60] and Ning and Petch [61] showed how analytical gradients can be efficiently computed to handle optimization problems beyond rotor design and include a more extensive set of constraints. The problem formulation has been extended to include fatigue considerations in the structural sizing [62]. The design space exploration strategy developed by Bortolotti et al. [63] efficiently sized the rotor at a conceptual design stage. However, they could not fully exploit the tradeoffs between aerodynamic and structural design, yielding suboptimal designs from an aerostructural standpoint. Bottasso et al. [64] proposed a bi-level design approach in a loosely coupled nested optimization using a comprehensive turbine model with a broad range of design constraints. They found the design maximizing AEP over the mass by interpolating a family of designs generated via a low-fidelity optimization.

A more compact sequential approach using the same tool was presented by Bortolotti et al. [65] while another work from the same research group extended the model to a monolithic approach that includes a more accurate finite element method analysis to evaluate the structural properties of selected blade cross sections [8]. The latter method eliminated any simplifying load-displacement transfer assumption between the aerodynamic and structural models. However, the algorithm was described as “more complicated and in general less robust” than the sequential counterparts.

To avoid the implementation effort of a tightly coupled aerostructural approach, the sequential approaches use a fixed or partially updated set of optimized aerodynamic loads between structural optimization subiterations. An iterative algorithm drives the aerodynamic and structural models to numerical convergence, but the resulting configuration does not achieve the performance possible with a tightly coupled approach. Scott et al. [66] used a similar approach for the sequential optimization of a large turbine using a linearized aerostructural model to speed up aero-servo-elastic simulations.

Zahle et al. [43] and McWilliam et al. [67] showed the potential of aeroelastic tailoring and passive load alleviation in blade design when utilizing composite material anisotropy. Although they used conventional low-fidelity models and finite-difference gradients, which ultimately limit the optimization accuracy and robustness, their results were promising. Heinz et al. [52] investigated the coupling of a CFD and an aeroelastic solver, showing good agreement with conventional approaches despite a nonconservative force transfer scheme and a loosely coupled aerostructural solution.

Commercial CFD codes and reduced-order structural models have also been used to analyze the blade-tower interactions and quantify acceptable tip clearances and efficiency losses [68]. Even these latter efforts did not perform CFD-based multidisciplinary optimization using a monolithic architecture, which is the main contribution of this paper.

1.3.1 Site-specific design

In Section 1.1.3 we stressed on how identifying appropriate design scenarios is essential to reduce LCOE for industrial-scale wind energy plants. A seminal work by Fuglsang et al. [69] in the early 2000s obtained large performance and component-level differences by including site-specific characteristics in the turbine design process. Both terrain features and Weibull-based wind distributions were used to characterize the design conditions at several onshore and offshore locations. Cost of energy reductions by up to 15% were obtained by simultaneously increasing AEP and reducing manufacturing costs.

Kenway and Martins [70] performed wind turbine rotor aerostructural MDO considering wind conditions at two different sites and 12 wind speeds for each optimization. The blade planform and main spars were tailored to minimize the cost of energy while satisfying load and power constraints. Despite an overall yield increase, the turbines designed for a specific site showed a power output reduction around 3-4% when analyzed at a different location.

Rezaei Mirghaed and Roshandel [71], building on other works from the early 2000s, performed single-turbine and farm optimization studies to minimize the cost of energy at three different sites in Iran. They coupled a BEMT aerodynamic model and a cost model within a nested optimization approach, taking into account rotor and tower size, and generator characteristics. Optimized hub height and rotor diameter vary by 20% and 7% respectively for different sites.

Holistic optimization frameworks include the whole operational envelope to size the system and predict its AEP. Ashuri et al. [72] minimized the LCOE for an offshore wind turbine using a multilevel approach to tailor the size, planform shape, and structure of a full turbine model. They modeled the wind conditions based on a Dutch North Sea location, splitting the operating envelope in bin intervals of 2 $\mathrm{m/s}$. Their approach also included fatigue and extreme loads, but the use of simplified disciplines and finite-difference-based gradients limited the design space exploration. Design studies using Cp-Max [64, 8] include both AEP in the objective and the full set of design loads used for certification purposes. Their sophisticated energy yield model includes regulation policy calculations accounting for the multibody response of the system. This tool includes a larger set of turbine components and design considerations than what used in the present work but, similarly to other BEMT-based codes, has limited geometry and structural detailed tailoring capabilities.

The effects of bend-twist coupling and site-specific wind distributions on blade design have been investigated in more recent works [73, 74]. Cognet et al. [73] demonstrated a two-parameter, scale-independent design approach for conceptual blade design, identifying optimal material properties to increase the power output at different offshore sites while reducing the rotor mass. Zhu et al. [74] investigated the blade and tower interaction impact on AEP and mass. They used a coupled BEM and a FEM solver and considering wind speeds at different rotor heights. A genetic algorithm was used to optimize 38 geometrical and structural design variables accounting for strain, deflection, and stability constraints. The resulting rotor designs display a marked shift in the C_P curve, with maximum values up to 7.7% higher and increased low-speed performance.

Incorporating a broader range of inflow conditions is also crucial for reliability and life-cycle considerations. Meng et al. [75] performed a Reliability-Based MDO (RBMDO) study of the tower of a 5 MW wind turbine, decoupling the problem into a sequence of deterministic optimizations that satisfy reliability constraints using Kriging-based function models. They considered a stochastic distribution of metocean conditions, buckling and frequency constraints, and a subset of 5 structural design variables to reduce the tower mass by 6%.

Few works in the literature tackled wind rotor design using a CFD solver in the loop. Dhert et al. [14] optimized the NREL VI two-bladed rotor using ADflow, a CFD tool with a discrete adjoint solver also used in the present work. Their optimizations used more than 250 design variables to modify the twist distribution and airfoil section shape, increasing the torque by more than 20%. However, the rotor thrust also increased by almost 70% due to the lack of an adequate constraint. They compared airfoil sections obtained through single point and multipoint optimizations, showing minor design differences that might be stemming from early optimization termination.

1.3.2 Structural tailoring

Slender, flexible turbine blades have inherent aeroelastic coupling that affect local angles of attack, and thus torque output, during operation. The idea of steering composite fibers to improve wind turbine rotor response and obtain lighter blades dates back to the 1980's [76]. Veers et al. [77] described aeroelastic tailoring as "*a cost effective, passive means to shape the power curve and reduce the loads*". Their review focused on a number of design solutions from wind energy researchers and the helicopter industry, including asymmetric fiber layups to improve power production, stall control, and blade vibration response. Solutions like non-orthotropic ply layups were intended to induce favorable twist-bending response without major blade structural modifications. Such concepts were considered promising but challenging to be reliably manufactured at industrial scale. Multidisciplinary design optimization was also considered instrumental in capturing complex aeroelastic tradeoffs.

Manufacturing techniques have improved in the last decades. Modern turbines heavily rely on composite materials based on thermoset resins with glass or carbon fiber reinforcements [4]. The latter material has higher passive tailoring potential but still remains more expensive than its fiberglass counterpart. It is expected that future blades will exploit aeroelastic coupling with passive and active approaches implemented earlier in the design process [78].

Several works in the last decade have investigated structural and aerostructural design optimization with a focus on composite layups. Bottasso et al. [79] investigated passive load alleviation on rotor blades by rotating composite fiber angles. They performed multilevel structural thickness optimizations for a fixed blade outer mold line and prescribed fiber angles, accounting for stress, displacement, buckling, fatigue, and frequency constraints [80]. They observed how fiber angle rotation has a different impact on bend-twist coupling depending on the component and the spanwise location over the blade. Large fiber angle rotations over the entire blade lead to heavier designs due their lower flapwise bending stiffness. This suggests that only limited fiber steering should be applied on spar caps, while larger rotations on blade skins could be beneficial for mass reduction. Under gust conditions, local pitch can vary by up to 2° depending on the fiber angle distribution. Applying fiber rotations to only the outer sections of the blade was also demonstrated to be beneficial for mass reduction.

Albanesi et al. [81] used a FEM code to optimize ply order, number, and drop on a small wind turbine blade. They reverse-designed the structural layout to match a prescribed deflected shape, while enforcing stress, and frequency constraints. A genetic optimization algorithm was selected due to the presence of integer design variables. Optimizing up to 9 plies simultaneously, the final design was more than 15% lighter than their baseline.

Anderson et al. [82] coupled a CFD solver and a 3D FEM model to minimize structural stress through composite fiber steering, including dynamic scenarios. A scalar off-axis stress criterion is used to account for fatigue-critical loads. They compared two parametrizations with different design freedoms under different loads. The target stresses were reduced by between 20% and 60% for different load and optimization scenarios The use of a steepest descent optimization algorithm run for 10 iterations most likely limited further design improvements.

HAWTOpt2, developed from Danmarks Tekniske Universitet (DTU), uses FEM cross sectional tools and unsteady BEMT analysis code combined into a monolithic MDO strategy. Zahle et al. [43] demonstrated its capabilities increasing the energy yield of a reference 10 MW rotor without incurring in mass penalties. The blade outer shape, thickness distribution, and internal topology were optimized simultaneously, using finite-difference-based design sensitivities. The optimized design displayed a more aggressive pitch distribution up to the tip region, which is unloaded. Combined with the forward shift of the shear web and higher spar cap thickness at the root, the resulting design generate higher aerodynamic loads while satisfying stress, strain, and root moment constraints. McWilliam et al. [67] used the same tool to compare the effect of fiber-based and geometry-based passive bend-twist coupling. The annual energy production was increased by up to 17% and the mass was reduced by 6% without exceeding baseline platform loads. Fiber angle tailoring showed marginal advantages compared to conventional aerostructural optimization approaches or a novel swept blade concept. The tip displacement constraint was once again one of the main design drivers, limiting the overall mass reduction.

1.3.3 Fatigue and extreme loads

Due to implementation complexity and computational cost, only a few optimization studies in literature effectively consider ultimate load and fatigue constraints. Running the full set of IEC [25] design load cases (DLC)s would be not practical in large-scale optimization context [43]. A subset of DLCs including DLC1.x, DLC2.x, and DLC6.x is usually selected to consider relevant wind conditions without including turbulent scenarios. Pavese et al. [83] demonstrated that a deterministic reduced design load base can be extracted from the full DLC within a wind turbine MDO framework. At every optimization step, the load changes on the reduced load base are propagated to the full load envelope through a set of correction factors. They used a prescribed, spatially varying shear load to replicate peak loads in the flapwise and edgewise directions at operating and fault conditions. An extreme operating gust is used to evaluate the loads on a parked rotor. Removing stochasticity from design loads improves the optimization robustness and greatly reduces the computational cost of ultimate load evaluations.

Steady state simulations can also be used to enforce tip displacement or tip clearance constraints, as done by Zahle et al. [43] for an aerostructural optimization study on a 10 MW turbine comparable to the model used in this work.

The above mentioned integrated aerostructural approaches by Bottasso et al. [8] include a fatigue constraint in their optimization problems. They combined 2D and 3D FEM analysis tools with lifting line and BEMT codes to analyze the turbine performance at certification-level design load cases. System derivatives are obtained through finite differences, thus introducing accuracy and design variables scalability issues. They observed how "mode-switching" for critical ultimate design loads can make the design space non-smooth, so they opted to freeze the DLC selection at every iteration to avoid convergence issues with their SQP-based optimizer. The necessity of fatigue constraints was also discussed by Bottasso et al. [79], who showed how both fatigue and tip displacement constraints drive the optimal structural design.

Finally, Ingersoll and Ning [62] proposed a sequential optimization strategy where the equivalent fatigue damage is evaluated outside the optimization loop. They use FAST for dynamic rotor simulations and RotorSE for the steady-state, gradient-based optimization loop. They demonstrated that the sequential optimization converges in 4 to 5 steps, obtaining similar rotor LCOE of dynamic-solver-in-the-loop optimizations at a fraction of the computational cost. A surrogate-model-based approach also showed promising results, but involved high upfront costs for the training phase. The staggered approach they propose is further discussed and expanded in Chapter 8. We swap the BEMT-based optimization loop with MACH to perform fatigue-compliant high-fidelity gradient-based optimization.

1.4 Contributions

The previous sections outlined the need for more more efficient wind turbines and identified a research gap for achieving this through high-fidelity optimization strategies. My dissertation aims to answer a few overarching research questions:

Can high-fidelity MDAO be effectively used for the design of wind turbine rotors? What are its pros and cons compared to conventional design tools? What is the impact of the problem formulation and design freedom on the optimized configurations? How can composite structural properties be leveraged in aerostructural optimization? Can high-fidelity MDO be integrated with other tools in the design process?

This dissertation presents the first extensive aerostructural optimization studies for a wind turbine rotor using a coupled CFD-CSM solver. In this work, I leverage MACH high-fidelity aerostructural analysis capabilities to capture steady-state aeroelastic interactions over the rotor of the DTU 10 MW configuration. The optimization concurrently varies hundreds of geometrical and structural design variables. Such a large problem is computationally tractable thanks to the efficient computation of gradients via a coupled-adjoint method with a gradient-based optimizer.

This work has several original contributions. As discussed in Section 1.3, only a handful of previous efforts focus on comparable optimization problems with a gradient-based approach and a monolithic or semi-monolithic architecture. None of these efforts use a 3D coupled high-fidelity model, which has fewer assumptions than low-order models and provides greater design freedom. Our CFD-based aerodynamic analysis does not require precomputed lift-drag polars, enabling the parametrization of the blade airfoil sections. Similarly, the 3D structural model enables a component-based parametrization that is more detailed than an equivalent beam model and avoids sectional stiffness property condensation. Thus, this work is the first to perform concurrent aerodynamic and structural design optimization of a blade-resolved turbine rotor model.

The multidisciplinary coupling encompasses multiple aspects of the optimization framework. The aerodynamic and structural mesh geometries are updated consistently at every optimization iteration. The force-displacement transfer between the two solvers is conservative, ensuring physical consistency between the analyses. The tight coupling of these two models ensures the physical and numerical consistency of steady-state rotor aeroelastic response. The adjoint method for computing derivatives ensures scalability with the number of design variables [36, Sec. 6.7]. Thus, we can explore rotor design in a way that conventional optimization strategies cannot.

I do not investigate the rotor dynamic response and focus on its steady-state analysis, although aerodynamic shape optimization with flutter suppression has been previously performed using MACH [84, 85, 86]. Including stability constraints is a promising future extension of my work. Since I only model steady-state performance, there is no controller in the loop. However, the load case selection is informed by the original benchmark stability limits and control laws.

The most complex optimizations presented in this dissertation concurrently tailor the blade planform, sectional shape, structural thickness, and local fiber angle. Compared to a baseline aerostructural layout obtained through loosely coupled structural optimization, I obtain almost 40% mass reduction and an average load increase of more than 10%.

Despite the modeling capabilities of our tool, researchers still need to bridge the gap between high-fidelity MDO and design workflow that include practical, real-world sizing constraints. I have been involved in a side-effort to integrate MACH with OpenFAST and include life-cycle considerations in the loop.

Several key contributions are presented in this dissertation. To summarize, I:

1. Developed and verified a CFD-CSM wind turbine model to perform MDO studies with the MACH framework.

2. Demonstrated the advantages of tightly coupled high-fidelity optimization by performing the first aerostructural wind turbine rotor design study of its kind.

3. Tailored the blade composite fiber orientation and obtained lighter and stiffer turbine rotors using high-fidelity MDO.
4. Investigated the impact of geometrical parametrization, design objective, and site-specific conditions on the optimized rotor layout.
5. Contributed to the development of a mixed-fidelity framework to include life-cycle sizing constraints (for fatigue damage and extreme loads) in rotor design optimization.

1.5 Dissertation outline

The research contributions are outlined as following in the following chapters:

- Chapter 2 discusses the details of the MACH framework components.
- Chapter 3 covers the aerostructural model development and verification, including the code enhancements implemented in MACH
- Chapter 4 characterizes the benchmark model features, design load selection, and optimization metrics.
- Chapter 5 covers the first capability demonstration and Pareto front exploration.
- Chapter 6 presents composite fiber tailoring studies leveraging an updated aerostructural model.
- Chapter 7 extends the previous studies to a multipoint formulation including airfoil shape design variables.
- Chapter 8 details the implementation of the mixed-fidelity framework from Caprace et al. [87] with a new formulation to include life-cycle sizing constraints within high-fidelity optimization.
- Chapter 9 summarizes the concluding remarks and novel contributions, and outlines future research opportunities.

Chapter 2

The Optimization Framework

Gradient-based optimization is the only viable strategy to solve high-dimensional and high-fidelity optimization problems [36, Sec. 1.5]. The cost of the adjoint method for coupled system derivative computation is independent of the number of geometric and structural design variables, making the approach advantageous for optimization problems where the number of design variables is greater than the number of functions of interest [36, Sec. 6.7]. A brief overview of the mathematical formulation is discussed in Section 1.2.1.

We use our MDO framework, MACH, to perform aerostructural optimization studies using high-fidelity CFD and CSM analysis tools [12]. MACH has been developed by the MDO Lab for more than 15 years. The publicly available part of the code, MACH-Aero[1], has been applied to a broad range of aerodynamic shape optimization studies, including unconventional aircraft [88, 89, 90], hydrofoils [91], and wind turbines [14, 15].

MACH uses a coupled-adjoint approach for gradient computation and has been used extensively for wing [13, 92, 93, 94, 95, 96, 97, 98] and hydrofoil design [99, 100, 101, 102]. This dissertation work applies the same methodology to a high-fidelity wind turbine aerostructural model to demonstrate the potential performance and mass reduction benefits of integrating this approach in rotor blade design.

As shown in the extended design structure matrix (XDSM) diagram [103] in Figure 2.1, MACH is composed of a set of tightly integrated sub-modules that enable geometry parametrization and deformation, coupled aerostructural analysis, and efficient derivative computation for gradient-based optimization.

---

[1] `https://github.com/mdolab/MACH-Aero`

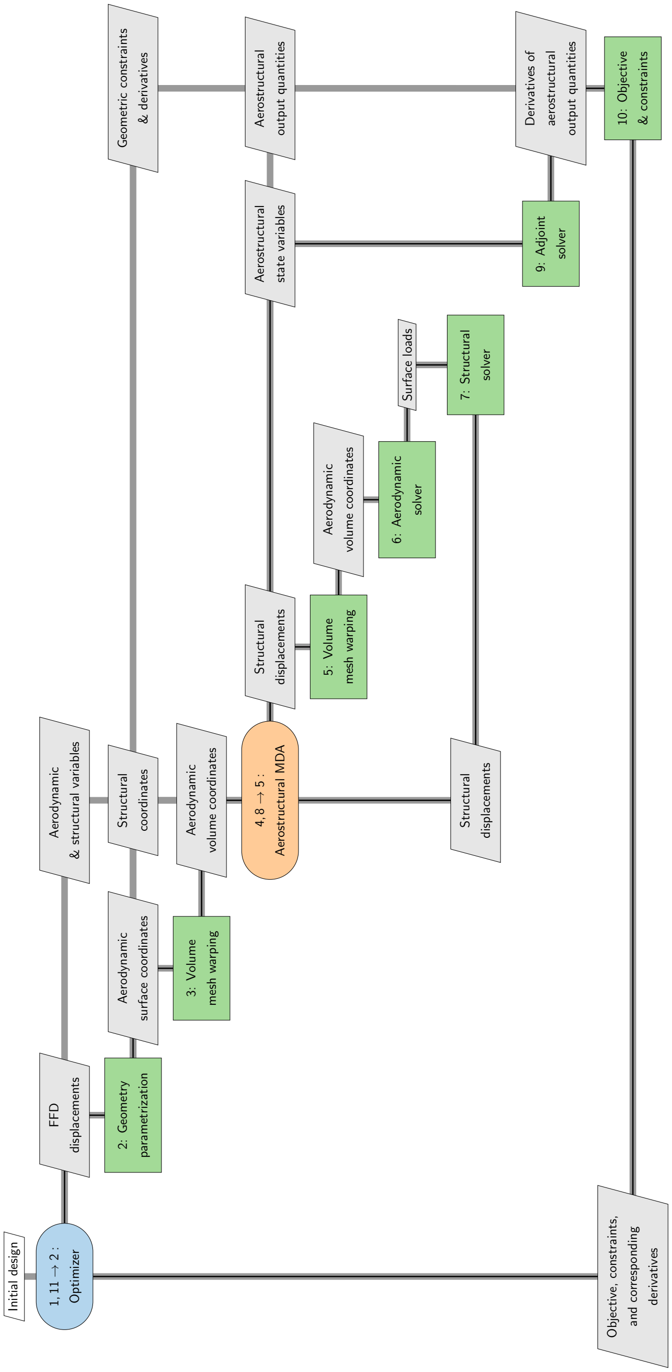

Figure 2.1: The MACH optimization framework, represented schematically in an XDSM diagram above, combines several modules for model manipulation and analysis.

2.1 Geometry parametrization

The geometry manipulation module, pyGeo[2] [104] (block 2 in Figure 2.1), is based on the free-form deformation (FFD) approach [105] as implemented by Kenway et al. [106]. The aerodynamic surface mesh and structural mesh points are embedded in a control volume defined by a set of FFD nodes (or "control points") on its surface. Having both meshes embedded in the same object is convenient for aerostructural optimizations, because it ensures consistent deformations between the two models.

The cloud of mesh points is mapped to a parametric tri-variate spline space when the geometry object is first initialized. The displacement of the FFD control points deforms the volume as if it was a flexible transparent "jelly-like" box. The volume deformation is propagated to the embedded cloud of mesh points using:

$$V(u,v,w) = \sum_{i=0}^{\mathrm{N}_u-1} \sum_{j=0}^{\mathrm{N}_v-1} \sum_{k=0}^{\mathrm{N}_w-1} \mathcal{N}_{i,p_u}(u)\mathcal{N}_{j,p_v}(v)\mathcal{N}_{k,p_w}(w)P_{i,j,k} \tag{2.1}$$

with V the deformed mesh points coordinates, P the control points coordinates, $\mathcal{N}$ the polynomial basis functions of degree p_um p_v, and p_w, and u, v, w the parametric coordinates of the embedded points.

Geometric design variables are defined as prescribed displacements of these control points. The optimizer can displace control points individually or collectively through user-defined sequences of scaling and rotations. More details about the geometric parametrization used in this work are discussed in Section 3.3.2 and in Chapters 5 and 7.

2.2 Mesh deformation

The aerodynamic surface node displacements are propagated to the volume mesh through IDwarp[3](blocks 3 and 5 in Figure 2.1), which uses an interpolation algorithm based on the inverse distance weighting method [107]. This avoids regenerating the CFD mesh at each iteration, increasing the optimization robustness and reducing the computational cost, or using FFD control volumes to manipulate the volume mesh [108]. IDWarp also leverages reverse-mode automatic differentiation to calculate machine-precision-accurate derivatives.

---

[2] `https://github.com/mdolab/pygeo`

[3] `https://github.com/mdolab/idwarp`

2.3 Aerodynamic solver

For aerodynamic analyses, we use ADflow[4] [109] (block 6 in Figure 2.1), a second-order, finite-volume compressible solver developed at the MDO Lab. We use the Spalart–Allmaras (SA) [110, 111] turbulence model and leverage on a low-speed preconditioner and an approximate Newton–Krylov method [112] to improve solver performance at low inflow speed [15, 113]. Our analyses do not model the laminar-turbulent transition, although the capability is under development [114, 115]. ADflow includes an efficient formulation to compute flow derivatives via the adjoint method [116], which makes it ideal for gradient-based optimization. More considerations on the use of ADflow at incompressible regime are outlined in Section 3.1 and in Chapters 5 and 7.

2.4 Structural solver

The structural solver used in this work is the toolkit for analysis of composite structures (TACS) [117] (block 7 in Figure 2.1). This solver is scalable, robust, and computes accurate machine-precision derivatives using the adjoint approach. It has been used in several previous studies and includes nonlinear [118] and stability [119] analysis capabilities. In this work, we perform linear structural analyses with a preconditioner based on the direct factorization of the global Schur complement matrix and a modified linear solver based on the generalized conjugate residual method with inner orthogonalization and outer truncation (GCROT) [120]. The structural stress constraints we use in this work are based on either the von Mises criterion or the Tsai-Wu criterion for the isotropic and anisotropic constitutive elements respectively [121]. The assumptions behind our structural model are discussed in Section 4.1.1.

We use the Kreisselmeier–Steinhauser (KS) aggregation function [122] to reduce the cost of the adjoint-based gradient computations [36, Sec. 5.7]. The KS function returns a smooth approximation of the maximum value of constraints over the complete set of element-wise local function values [123, 124]. We use an overflow-safe formulation:

$$\mathrm{KS}(\mathbf{g}(\mathbf{x})) = g_{\max}(\mathbf{x}) + \frac{1}{\rho_{\mathrm{KS}}} \ln\left(\sum_{j=1}^{m} e^{\rho_{\mathrm{KS}}[g_j(\mathbf{x}) - g_{\max}(\mathbf{x})]}\right) \tag{2.2}$$

with $\mathbf{x}$ the current set of design variables, $\mathbf{g}(\mathbf{x})$ the selected constraint, $g_{\max}(\mathbf{x})$ the highest constraint value from the structural analysis, and ρ_{KS} the KS aggregation parameter – equivalent to a “penalty” parameter. For $\rho_{\mathrm{KS}} \to \infty$, the aggregated constraint approaches the “true” maximum but makes the optimization more challenging due to the sharper curvature introduced in the design space. We use $\rho_{\mathrm{KS}} = 100$ throughout this work.

---

[4] `https://github.com/mdolab/adflow`

2.5 Aerostructural solver

A nonlinear block Gauss–Seidel solver with Aitken acceleration converges the coupled aerostructural analysis [36, Sec.13.2.4] (block 8 in Figure 2.1). The linear coupled-adjoint system is solved via the coupled Krylov subspace method, as implemented by Kenway et al. [12] (block 9 in Figure 2.1). The coupled system leverages matrix-free automatically differentiated routines and PETSc's GMRES for memory efficiency. The aerodynamic and structural models are coupled through a rigid link scheme [125] that is demonstrably consistent and conservative, as shown later in Section 3.3.1.

2.6 Optimization framework

Finally, pyOptSparse[5] [126] (block 1 in Figure 2.1) passes the function values and sparse Jacobian from MACH to the optimizer of choice, SNOPT [127].

pyOptSparse is a popular software package that wraps a set of gradient-free and gradient-based optimizers. It is leveraged by a number of independently developed design optimization frameworks, including NASA's OpenMDAO [40]. The use of object-oriented language enables a uniform user interface for different optimizers, including consistent optimization history recording. MPI wrapping and an extensive use of sparse matrices throughout the code make pyOptSparse flexible and efficient for CFD-based optimizations like the ones enabled by MACH.

SNOPT terminology

SNOPT implements a sequential quadratic programming (SQP) algorithm with an augmented Lagrangian merit function and an approximated quasi-Newton Hessian obtained with BFGS updates [36, Algorithm 4.22]. SNOPT automatically satisfies linear constraints at each iteration and is well suited for problems with sparse nonlinear constraints, which are handled through the active-set method.

Two metrics are relevant to monitor for optimization convergence and termination. Using the notation and definitions from Wu [128] and Gill et al. [127], the *primal feasibility* is given by:

$$\tau_{fea} = \max_i \frac{v_i(\boldsymbol{x})}{\|\boldsymbol{x}\|_2} \tag{2.3}$$

and the dual feasibility, or *optimality*:

$$\tau_{opt} = \max_j \frac{\mathrm{Comp}_j}{\|\boldsymbol{\lambda}\|_2} \tag{2.4}$$

---

[5] https://github.com/mdolab/pyoptsparse

with $\|\boldsymbol{x}\|_2$ and $\|\boldsymbol{\lambda}\|_2$ the L_2 norm of the design vector and Lagrange multipliers respectively. The constraint violation v_i for the i^{th} constraint is given by:

$$v_i(\boldsymbol{x}) = \begin{cases} |c_i(\boldsymbol{x})|\,, & i \in \mathcal{E} \\ \max\{0, c_i(\boldsymbol{x})\}\,, & i \in \mathcal{I} \end{cases}, \tag{2.5}$$

with $\mathcal{E}$ and $\mathcal{I}$ the set of equality and inequality constraints, while the complementary slackness Comp_j is:

$$\mathrm{Comp}_j = \begin{cases} d_j \min\{x_j - l_j, 1\} & \text{if } d_j \geq 0 \\ -d_j \min\{u_j - x_j, 1\} & \text{if } < 0 \end{cases}, \tag{2.6}$$

with the reduced gradient:

$$d_j = \frac{\mathrm{d}f(\boldsymbol{x})}{\mathrm{d}x_j} + \lambda_j \frac{\mathrm{d}\boldsymbol{c}(\boldsymbol{x})}{\mathrm{d}x_j}. \tag{2.7}$$

At the optimum the KKT conditions [36, Eq.5.26] are satisfied and $\tau_{fea} = \tau_{opt} = 0$. Otherwise these metrics indicate how “far” the optimization is from the optimal solution. Design variable and function scaling affect these values [128, Sec.2.5.7.1], so the optimizations in this work consider relative (compared to the initial optimality τ_{opt}) convergence as a metric for how tightly the optimization has converged. Feasibility can be easily satisfied in general, but for stiff aerostructural problems it is hard to obtain relative convergence higher than two or three orders of magnitude. As shown in Section 5.4, from an engineering standpoint the design is practically unchanged between optimization iterations at that point.

Chapter 3

Model Development and Verification

A sizable part of my research consisted in adapting MACH to handle a wind turbine optimization problem. The framework core capabilities have been outlined in the previous Chapter 2. Aircraft wings and turbine blades are lifting devices that leverage the same physical phenomena and share several design features. Nevertheless transitioning to a different model presented a number of practical and numerical challenges.

On the aerodynamics side, we observed the same mesh sensitivity described by Madsen et al. [15]. We performed additional verification studies and verified that, despite their known load overestimation, we could use more practical coarse meshes during the optimization process. Moreover, the fixes on the moving reference frame derivatives initially implemented for Madsen et al. [15] have been fully implemented, verified, and released to the public.

From the structural standpoint, the main effort consisted in developing the blade-resolved structural model for the reference turbine. We developed a workflow to mesh and verify the structural model from scratch, leveraging on existing in-house-developed meshing tools. We outline practical considerations for mesh convergence study and aerostructural coupling.

The details of the reference turbine are presented later in Chapter 4. The present chapter summarizes the key development steps for the wind turbine model used in this dissertation. Section 3.1 presents the aerodynamic mesh features, the analysis validation efforts, and the sensitivity bug-fix and verification process. The details of the structural model development are summarized in Section 3.2, including the constitutive elements and material properties used in different studies. Finally, Section 3.3 briefly details the coupled aerostructural model, including its geometrical parametrization.

3.1 Aerodynamic model

This section summarizes the features the aerodynamic meshes used in this work. We also briefly describe the derivatives verification effort completed early in the framework development phase.

ADflow, described in Section 2.3, is used to perform steady state RANS flow analysis assuming uniform inflow conditions. Unsteady or time spectral simulations still remain numerically impractical for large scale optimization studies. Non-uniform inflow conditions are also challenging to simulate with this solver. The viscous dissipation intrinsic to RANS models would smear the turbulent inflow distribution usually set for BEMT-based simulations. At the same time using laminar non-uniform velocity distributions in the far field for steady-state simulations would require multiple analyses to evaluate the rotor response and performance over a complete revolution. Such studies go beyond the scope of this dissertation.

Pros and cons of using a CFD solver for wind turbine design

CFD-based tools such as MACH cannot replace the BEMT-based tools mentioned in Section 1.3. These analysis software can complement each other, as demonstrated later in Chapter 8. Nevertheless, due to the intrinsic differences between the two methods, a one-to-one result comparison using the same load cases is only useful to identify common trends rather than detailed design features [15]. To investigate how the different modeling approaches influence the turbine performance predictions, we re-analyze the optimized rotors using both a finer CFD mesh and OpenFAST later in Section 5.6. Identifying the pros and cons of different models remains a challenging task, but we observe comparable performance improvements between different solvers.

Moreover, ADflow, and RANS solvers in general, have known accuracy and robustness limitations at low-speed inflow conditions. Seraj et al. [129] recently presented developments on ADflow low-speed preconditioner to address these issues. They proposed to modify the scalar dissipation term of the JST scheme, reducing or eliminating its dependency on the Mach number. This approach significantly improved accuracy at Mach ≤ 0.1. Moreover, they implemented a characteristic time stepping approach to accelerate the convergence of the ANK solver [112]. Future extensions of the wind turbine optimization studies presented in this work could leverage on these new capabilities.

As discussed later in Chapter 5, the SA turbulence model tends to postpone the stall occurrence. We verified that the optimized layouts presented in this work do not operate at the maximum lift of the airfoils (Section 5.4). Nevertheless, the absence of flow separation indicates generally overestimated lift values.

Keeping these limitations in mind, CFD tools, combined with FEM structural solvers, move design optimization a step forward. As mentioned in Section 1.4, using a blade-resolved 3D aerostructural model in the optimization has two main benefits. On the one hand, the analysis captures spanwise flow characteristics and coupled aerostructural interactions without fundamental assumptions or empirical corrections on the flow behavior and local stiffness properties. The result discussion in this dissertation highlight how designers can obtain deeper insight when using CFD-CSM, such as the local pressure distribution and stress concentrations.

On the other hand, the parametrizations described later in Sections 5.1.2, 6.2.3 and 7.1.3 enable the optimizer to explore a broader, more realistic design space maintaining consistency between the aerodynamic and structural models. Lower fidelity tools rely on lower-order models that, while capturing high-level trends, do not include the airfoil shape and detailed structural sizing variables used in the present work.

The modeling assumptions of this work are necessary to limit the complexity of the study at the current development stage but do not prevent the extension of the same approach to more complex and practical optimization scenarios.

3.1.1 Mesh family

The set of DTU 10 MW turbine meshes we use have been developed by Madsen et al. [15] in their previous work on aerodynamic shape optimization, in collaboration with the MDO Lab. The original DTU 10 MW chord and twist distribution are slightly altered to allow for more room for improvement during the optimization [15, Figure 2]. The structured mesh is compatible with both ADflow and EllipSys3D.

The multiblock meshes use an O-O topology and extend radially up to approximately 20 rotor radii, where the farthest nodes are 1740 m away from the rotor center. Rotation is modeled through a non-inertial reference frame attached to the blades, with related rotational terms added to the standard RANS equations. The rotation rate of the rotor is a fixed input, together with the freestream thermodynamic properties.

Once the most refined mesh (L0) is extruded, the rest of the family is generated according to the grid refinement approach by Roache [130]. At each coarsening step, every other node in the three dimensions is removed. This reduces the number of cells by a factor of 8. A family of four meshes was defined, starting from the most refined L0 ($14,155,776$ cells) to L3 ($27,648$ cells) in increasing coarsening levels.

A detail of the “hub” section is highlighted in Figure 3.1 for the most refined blades, L1 and L0. Note that we do not model an actual hub and nacelle, as the three blades “blend” together at the hub location instead. This simplifies the geometry generation with negligible impact on the load output [24, 15].

Due to their size, any analysis on meshes more refined than L2 needs HPC resources. We used both UM’s Great Lakes and TACC’s Stampede2, with the first preferred for cheaper or urgent runs, and the latest providing enough computational power to run the most expensive cases. The difference in computational cost for different grids significant, as reported in Tab. 3.1.

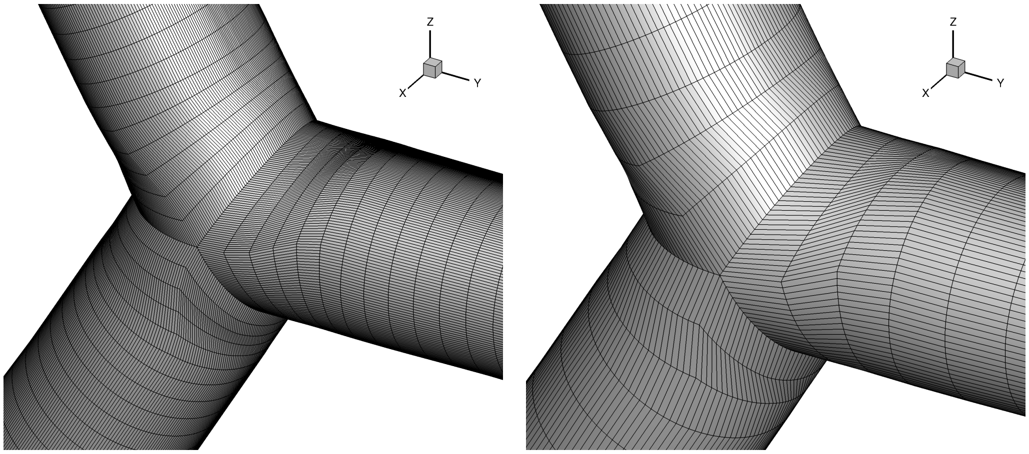

(a) L0 (b) L1

Figure 3.1: Detail of the root section of the L0 and L1 aerodynamic meshes used in this work.

Mesh Details		
Level	N° cells	Cost (CPUh)
L0	14155776	~ 350
L1	1769472	~ 30
L2	221184	~ 0.5

Table 3.1: Size and approximate cost of the mesh family for DTU 10 MW reference wind turbine.

The computational cost increases non linearly with the number of elements due to the increasingly stiff non-linear system to be solved. Moreover, ADflow parallel preconditioner has some known weaknesses that will be addressed in future works.

3.1.2 Previous verification and validation

As mentioned earlier in this chapter, low-speed inflow conditions are challenging for ADflow due to its density-based solver formulation. While its robustness [112] and accuracy at transonic [131] and supersonic regimes [132] have been thoroughly investigated, fewer works have used ADflow in the incompressible regime. Dhert et al. [14] verified the tool accuracy on their NREL phase VI model against experimental data [133], showing good agreement for inflow conditions without separation.

Nevertheless, they also observed a large solution sensitivity to mesh resolution. However, as further discussed later in Chapters 5 and 7, the performance improvements obtained through optimization were consistent between the finest and a reasonably coarse mesh. Studies such as Lyu et al. [134] have used different meshes for optimization and verification, showing consistency between mesh levels. The discretization error has been quantified and leveraged within a multifidelity aerostructural optimization framework developed by Wu et al. [135] for efficient transport aircraft design optimization.

Dhert et al. [14] also addressed early ADflow convergence issues by removing the cylindrical root section, which was triggering large flow separation. Flow solver robustness issues related to blade root flow separation have been observed in incompressible flow solvers as well [136]. The blending at the root for our model, shown in Figure 3.1, is a compromise between modeling complexity and flow solver robustness.

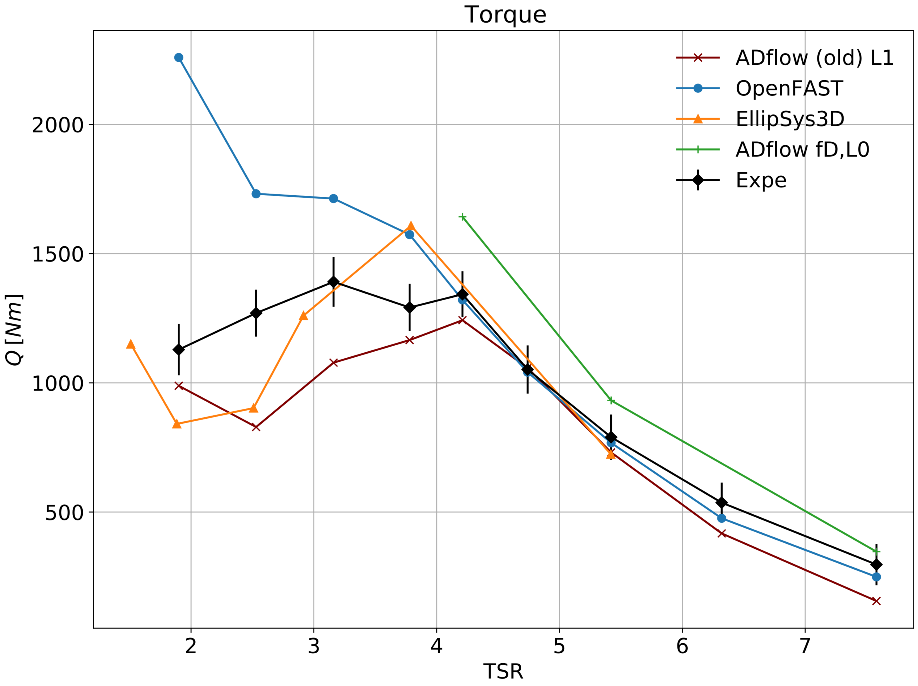

Figure 3.2: Torque-vs-TSR comparison between experiment, AeroDyn and ADflow for the NREL Phase VI UAE case H configuration. Credits: Denis-Gabriel Caprace

The ATLANTIS project [1] that funded the early phase of my research also involved tool verification tasks. In collaboration with BYU's FLOW Lab [2], we developed and tested our own NREL phase VI model and compared it with experimental data [137] and numerical results from Ellip-

---

[1] Project website, last accessed Nov 2023

[2] FLOW Lab Website, last accessed Nov 2023

Sys3D [138] and OpenFAST. The output torque for a range of operational TSR is illustrated in Figure 3.2.

The experimental and numerical data match between TSR=4 and TSR=8, with ADflow under-predicting the torque output at higher TSR. For lower rotation rates, all numerical tools fail to capture the torque values and trends. OpenFAST largely overpredicts the rotor output at TSR<4, while the CFD solvers mostly under-predict the available power.

Shifting the focus on the DTU 10 MW model used in this work, Madsen et al. [15] verified the same structured mesh family presented in Section 3.1.1 on ADflow and EllipSys3D. They highlighted how even the L0 solution displayed some mesh dependency, and they ran an additional analysis on a more refined mesh to identify the asymptotic region and extrapolate a discretization-agnostic thrust and torque prediction using Richardson's extrapolation [130].

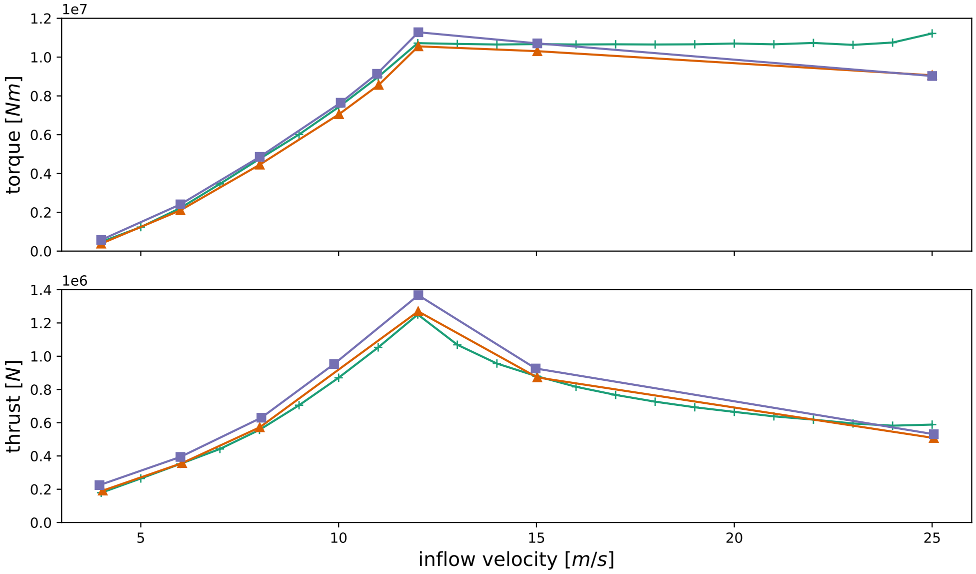

Figure 3.3: Aerodynamic torque and thrust of the baseline turbine, obtained with AeroDyn (+), ADflow (■) and EllipSys3D (▲) (L0). From Caprace et al. [87].

While they identify L2 meshes as too coarse to be reliable, they also show that aerodynamic shape optimizations using the L1 mesh, despite the discretization error, essentially capture the same design trends obtained with the higher-resolution L0 grid. Moreover, in their comparison with EllipSys3D, they acknowledge the limitations of the latter in predicting compressibility effects on large turbine applications. The large mesh sensitivity has been observed in other compressible solvers [139].

Additional verification with EllipSys3D (using more data from Madsen et al. [15]) and AeroDyn [42] has been performed by Caprace et al. [87]. The results in Figure 3.3 show torque and thrust values for the three solvers over a range of wind speeds and rotation rates[15, Table 10].

ADflow again predicts higher loads than EllipSys3D due to the mix of discretization errors, different turbulence models, and lack of compressible correction on EllipSys3D. The agreement between EllipSys3D and Aerodyn, especially for the thrust, comes from the fact that the latter uses EllipSys2D-calculated polars in its BEMT analysis. Nevertheless, the three solvers show consistent agreement over the full range of the DTU 10 MW model.

Improving ADflow accuracy and mesh-independence at low speed regimes remains an open challenge for future work. Nevertheless, the studies mentioned in this section highlight consistent behavior and the solver capability to capture the rotor performance trends over its operating envelope. In this work, the computational complexity and cost is largely increased by the use of the coupled aerostructural model. In particular, the coupled adjoint solutions are significantly more expensive than the aerodynamic counterparts. For practical purposes, we use the L1 aerodynamic mesh for the optimization studies presented in this dissertation, and leave the L0 mesh for post-optimality result verification studies.

3.1.3 Derivative verification

The main advantage of using ADflow for wind turbine optimization studies, despite the limitations outlined in Section 3.1.2, is its efficient sensitivity solver based on the adjoint method. Dhert et al. [14], in the first application of ADflow to wind turbine design, used a rather inefficient matrix-based approach. Flow sensitivities were calculated using automatic differentiation in forward mode, with significant memory usage computational cost.

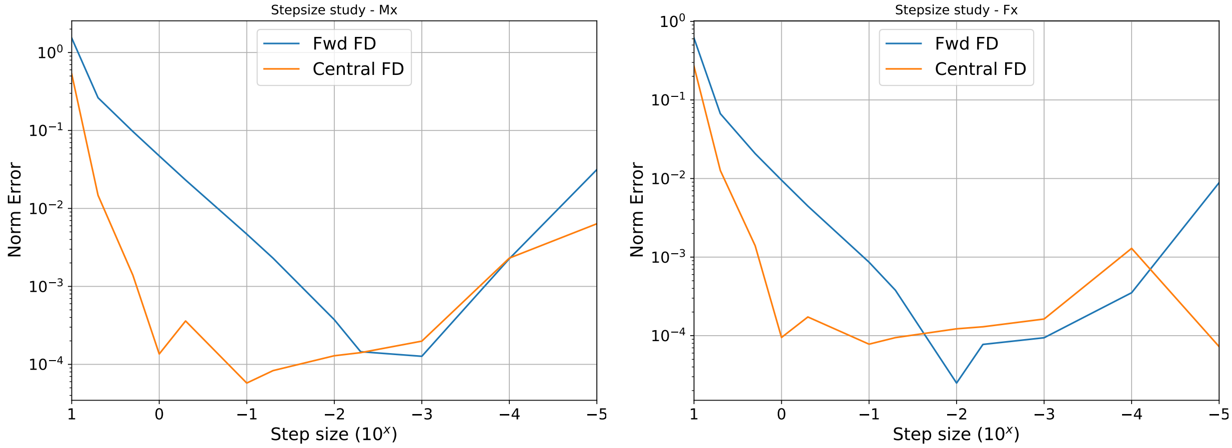

(a) (b)

Figure 3.4: Stepsize study for the sensitivity of torque (a) and thrust (b) with respect to pitch for L1 mesh, using forward and central finite differences.

Madsen et al. [15] re-implemented the reverse mode Jacobian-free approach for rotating reference frames. This feature was however not included in the publicly available code base. The original code updates were also not directly compatible with the current ADflow version. One of the early development thrusts of my research work focused on addressing these issues.

We initially verified that the original code was effectively providing consistent and accurate derivatives. The adjoint-based sensitivities from this early code version are compared to finite-difference sensitivities in Figure 3.4. This sanity check shows the expected trends, with the error initially reducing to 10^{-4} with step size decrease, then increasing due to finite-difference truncation errors.

Table 3.2: Aerodynamic derivatives verification. We report torque and thrust sensitivities with respect to a set of geometrical design variables. Sensitivities are obtained through finite-difference (FD), complex-step (CS) [140], and adjoint method (Adj).

DV	Method	$\frac{\mathrm{dTorque}}{\mathrm{dDV}}$	$\frac{\mathrm{dThrust}}{\mathrm{dDV}}$
Pitch (before fix)	FD	+116607.51705952	+30163.95325537
	Adj	-894829.77473474	-10360.33613246
Pitch (after fix)	FD	116607.51705952	30163.95325537
	CS	116932.38914284	30172.60830132
	Adj	116932.37831244	30172.60825961
Twist	FD	883.22199881	687.86973716
	CS	599.02723727	661.19523656
	Adj	599.02731769	661.19523822
Chord	FD	477391.6509002	38638.5008460
	CS	477638.9593960	38638.4694959
	Adj	477638.9474204	38638.4692759
Shape	FD	-314760.434	-5776.91779
	CS	-314646.769	-5765.26134
	Adj	-314646.774	-5765.26135

The core of the bugfix [3] consisted in correctly differentiating missing subroutines with the mixed analytical and automatic-differentiation-based methodology described by Kenway et al. [116]. Fortran subroutines that compute the grid velocities for the moving reference frame approach were first differentiated through Tapenade [141]. Parts of the original Fortran code had to be reformatted to ensure Tapenade would modify the source code consistently. Finally, the differentiated routines have been manually added to the master file to make sure the reverse-mode seeds are propagated correctly at each call of the Jacobian-free GMRES algorithm for the solution of the

---

[3] Bugfix PR on ADflow GitHub repository

adjoint system.

We verified the sensitivities for different sets of design variables (DV), as shown in Table 3.2. We obtain at least a 7-digit match between the adjoint and complex step derivatives, despite the reference solution with complex variables could not converge to machine precision. Note that the issues with the bugged version of ADflow ("Pitch (before fix)") were severe to the point that the derivative sign was wrong. This would lead to immediate optimization failures due to the incorrect gradient information. With the bugfix finally completed and released to the public in ADflow v2.7.3 [4], we could finally move on to aerostructural optimization studies, initially presented by Mangano et al. [142] and then further expanded in Chapters 5 to 7.

3.2 Structural model development

Conversely to the aerodynamic meshes, the structural model for this work has been developed from scratch. This section summarizes the workflow we put in place to generate and test the model. We also discuss the key modeling features of the blade structure and the areas of improvement for future studies.

Section 3.2.1 outlines the meshing process and presents some topological details blade structure, including a grid convergence study to select the mesh refinement level. In Section 3.2.2 we discuss the two types of TACS structural elements we use in our optimizations. Finally, the material properties of the two models we used in our studies are listed in Section 3.2.3.

The structural mesh parametrization and layup are presented later in Section 4.1.1. The problem-specific details are further discussed in Chapters 5 to 7.

3.2.1 Meshing

We based the mesh topology on the benchmark description by Bak et al. [24] and the publicly available data on the DTU GitLab repository [5].

The blade has a single shear web, made of the two main spars, extending from the root up to the tip. Conversely to the simpler straight spars of the IEA 15 MW [22] model, the DTU 10 MW spars have a curved section between approximately the span 15% and 25%. A reinforcement spar in the trailing edge area starts from approximately the span 23% and also extends to the tip. We manually sampled the 3D geometry model to extract the leading and trailing edge curve of the blade. We use the information in Bak et al. [24, Chapter 4] to parametrically generate the curves defining the

---

[4] ADflow v2.7.3

[5] DTU 10 MW GitLab repository, last accessed November 2023

internal topology, i.e. the chordwise location of the spars. The resulting blade planform is shown in Figure 3.5.

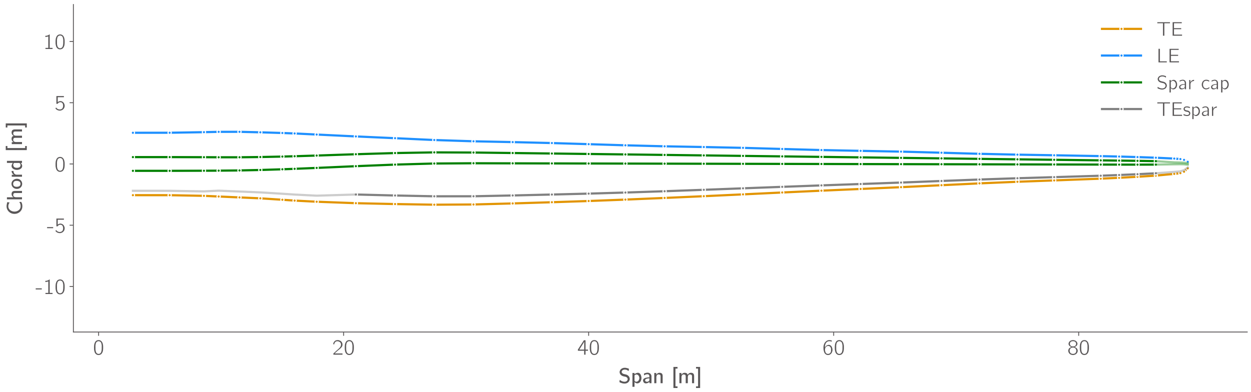

Figure 3.5: Planform topology of the DTU 10 MW structural model.

We use an in-house developed tool to project point from the 2D planform onto the available 3D geometry to generate the blade-resolved structural mesh. The blade is duplicated and rotated to obtain the final 3-bladed layout consistent with the aerodynamic mesh discussed in Section 3.1.1.

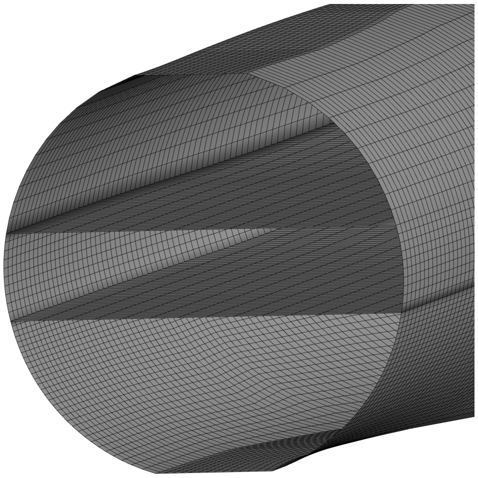

Figure 3.6: Detail of the structural mesh root section.

This procedure is manually intensive but ensures a tight fit between the tow models. Note that to avoid skewed elements at the root leading and trailing edge section, we leave a gap between the structural mesh and the circular section on the aerodynamic model, as shown in detail in Figure 3.6.

The blade degrees of freedom (DOFs) are fixed at the root section. To be consistent with real world blade installation, we do not limit the DOFs of the spars at the root, leaving the outer skins as the only constrained elements.

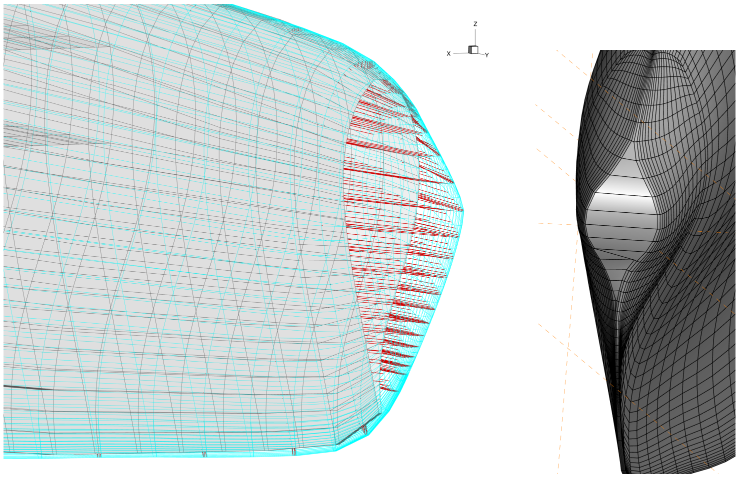

(a) Rigid links connecting to the strucutural mesh edge (b) Mesh failure on the deflected tip

Figure 3.7: Mesh deformation failures due to the missing rib. The aerodynamic mesh is shown in light blue, the structural mesh in gray, and the aerostructural rigid links are highlighted in red.

The modeling considerations discussed above have been informed by analyses with fixed aerodynamic loads. However, our goal is to use the structural model within an aerostructural multidisciplinary analysis (MDA). The first coupled analyses on the structural model revealed a numerical weakness in the tip region. The rigid links connecting the aerodynamic and structural model (discussed in Section 2.5) were defined to connect the tip of the aerodynamic surface mesh with the edge of the upper and lower surfaces on the structural model, shown in Figure 3.7(a). However, even minor deflections lead to inconsistent aerodynamic mesh deformation and ultimately mesh failures at the tip, highlighted in Figure 3.7(b).

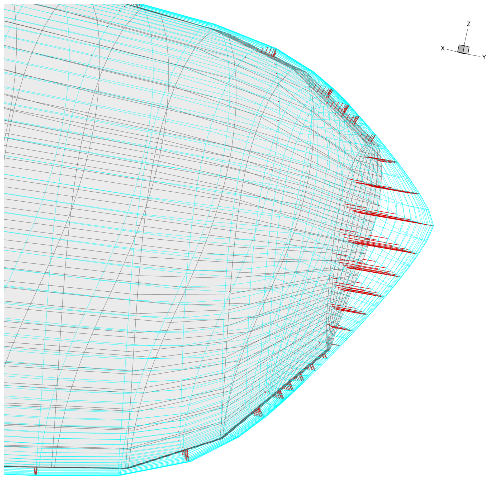

Figure 3.8: Final structural mesh with a rib at the tip and more robust rigid links.

For this reason, we opted to include a rib to "close" the tip of the structural model. The rigid links, shown in Figure 3.8 are evenly distributed over the rib, preventing the aerodynamic nodes from being stretched inconsistently. The impact on the structural behavior of the model is negligible since this area has low aerodynamic loads. The adjacency structural constraints are not enforced on this element of the blade, which is then allowed to hit the lower thickness bound.

Grid convergence study

Starting from the most refined mesh, L0, we halve the number of elements in both chordwise and spanwise direction to obtain a coarser L1 mesh. The same procedure is performed on the L1 mesh to obtain an even coarser mesh L2. A 3D view of the mesh family is illustrated in Figure 3.9 and Table 3.3 reports the number of elements for each grid.

Figure 3.9: The structural mesh family, showing L0, L1, and L2 from top to bottom.

The most refined L0 mesh has almost one million elements and is unnecessarily refined for the kind of analysis we are performing in this work. Moreover, excessive refinement leads to artificial stress concentrations on the structural computational model. We include some considerations on this issue in Appendix A.

Table 3.3: Structural mesh family grid size.

Level	Element number
L0	943 344
L1	235 836
L2	58 959

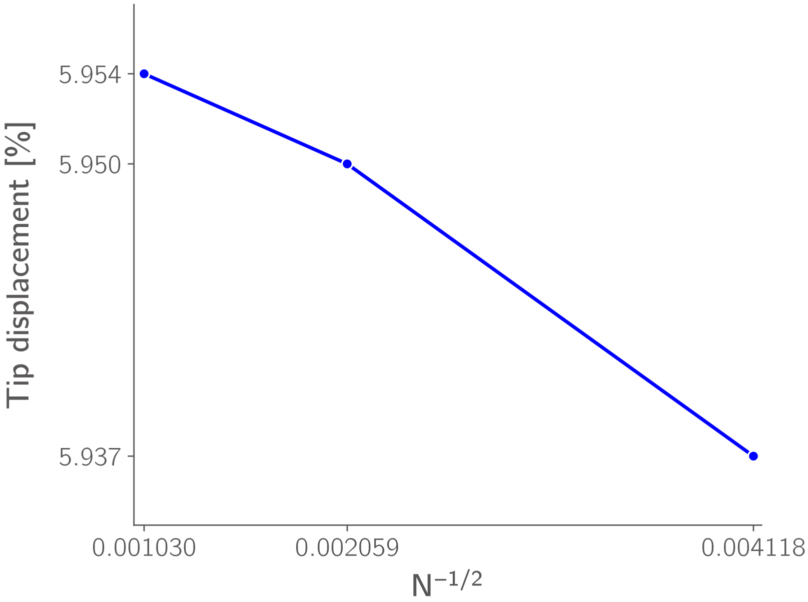

Figure 3.10: Grid convergence study for the tip displacement, with increasing mesh refinement from right to left.

For grid convergence purposes, we use fixed aerodynamic loads from an aerodynamic simulation on the L2 aerodynamic mesh from Section 3.1.1. This mesh overpredicts the loads compared to the L1 grid used for production cases, resulting in higher stress value than those experienced by the structure during optimization.

The tip displacement is a good metric to assess grid convergence due to the above mentioned issue with local stress concentrations. Despite the lower refinement, the structural L2 mesh predicts a tip displacement less than 0.02% different from the most refined L0 grid, as shown in Figure 3.10. Considering this, and the absence of spurious stress concentrations at the skin spar junction (see Appendix A), we select the structural L2 mesh for the studies in Chapters 6 and 7.

3.2.2 Finite element constitutive models

TACS implements a set of different element types that enable different structural parametrizations and material definitions. Solid elements, modeled as 8-node brick elements (CHEXA8), have been implemented by Garg et al. [99] to model hydrofoil solid construction. Beam elements using the classical Euler–Bernoulli theory are also available. However, most of the applications of MACH relied on shell element formulation, more appropriate for thin walled structures such as aircraft wingbox and wind turbine blades.

Shell elements use first-order shear deformation theory (FSDT) from Uflyand–Mindlin theory. This approach is the plate equivalent of the Timoshenko–Ehrenfest beam theory, assuming linear in-plane strain variation over the shell. When modeling composite structures, we define "panel" the whole element modeled in TACS, "laminate" the composite plies full layup, and "ply" layers of individual material.

The optimizations in this dissertation use two types of constitutive elements. The prototype model, detailed later in Section 4.1.1.1, uses an isotropic "isoFSDT" element. This model assumes isotropic material properties to model the panel stiffness - hence the self explanatory name.

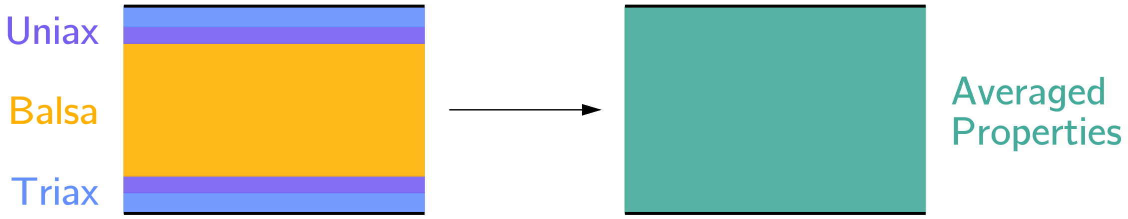

Figure 3.11: The smeared stiffness approach averages the ply material properties into an element with thickness-wise uniform properties.

The DTU 10 MW blades are made of balsa wood and fiberglass/epoxy composite sandwich panels. To model composite properties, we switched to an element type that enables individual ply material definition and fiber direction. This "smearedFSDT" element, used for the model described in Section 4.1.1.2, is based on the same FSDT theory and uses classical laminate theory to average material properties. We use a smeared stiffness approach to model the panel properties in our finite element structural solver. Figure 3.11 describes the ply properties averaging graphically.

Kennedy et al. [143] and Gray et al. [144] describe the mathematical modeling of smeared laminates implemented in TACS. The generalized stiffness matrix for a panel can be split into 4 sub matrices:

$$\mathbf{C} = \begin{bmatrix} \mathbf{A} & \mathbf{B} \\ \mathbf{B} & \mathbf{D} \end{bmatrix} \tag{3.1}$$

where

$$\mathbf{Q} = \sum_{i=1}^{3} f_i \hat{\mathbf{Q}}_i, \quad \mathbf{A} = t_p \mathbf{Q}, \quad \mathbf{B} = 0, \quad \mathbf{D} = \frac{t_p^3}{12}\mathbf{Q}.$$

$\hat{\mathbf{Q}}_i$ is the individual ply stiffness, t_p is the panel thickness, and the f_i is the ply fraction of the i^{th} ply. The ply fraction values are normalized so that:

$$\sum_{i=1}^{3} f_i = 1. \tag{3.2}$$

Conversely to Kennedy et al. [143], the constitutive element we use in this work does not include panel stiffeners but allows the user to arbitrarily define the principal direction of each laminate ply. The stiffness properties of the laminate are defined with respect to this reference axis. The blade span axis is fixed during the optimization. In practice, changing the axis direction without altering the material properties is equivalent to steering the ply fiber angle. The stress for this constitutive model is based on the Tsai-Wu criterion [143, 145]:

$$\sigma_{\mathrm{TW}} = \sqrt{\frac{\sigma_{11}^2 + (X_c - X_t)\,\sigma_{11}}{X_t X_c} + \frac{\sigma_{22}^2 + (Y_c - Y_t)\,\sigma_{22}}{Y_t Y_c} + \left(\frac{\sigma_{66}}{S_{12}}\right)^2} \tag{3.3}$$

with σ_{11}, σ_{22}, and σ_{66} the tangential, transverse, and shear stresses in the local frame of reference. X_i and Y_i are the ply strengths parallel and transversal to the fibers in tension (t) and compression (c), and S_{12} the ply shear strength. The properties of different materials are reported in the next section.

3.2.3 Material properties

There are 4 different material types defined in the DTU 10 MW benchmark documentation [24]. The core of the sandwich panel is made out of balsa wood. The fiberglass skins can have fibers aligned in a single direction (0°, Uniax), two directions (±45°, Biax), or three directions (0°, ±45°, Triax). For the isotropic prototype model with "isoFSDT" elements, we use aluminum characteristics. The material properties of different ply materials are in Table 3.4.

Due manufacturing assumptions on the original model [24], in practice there are never more than two different fiberglass skins in the same panel component – hence why the $i_{\text{max}} = 3$ in Eq. 3.2.

Table 3.4: Properties of the materials used in the optimization studies.

Property	Symbol	unit	Aluminum	Balsa	Uniax	Biax	Triax
Density	ρ	$\left[\frac{\text{kg}}{\text{m}^3}\right]$	2780	110	1915.5	1845.0	1845.0
Longitudinal Young's modulus	E_x	[GPa]	73.1	2.73	41.63	13.92	21.79
Transversal Young's modulus	E_y	[GPa]		0.05	14.93	13.92	14.67
Poisson's ratio	ν_{12}		0.33	0.5	0.241	0.533	0.478
In-plane shear	G_{xy}	[GPa]		0.15	5.047	11.5	9.413
Out-of-plane Shear	G_{yz}	[GPa]		0.0167	5.046	4.538	4.538

3.3 Aerostructural model

Once the aerodynamic and structural models have been independently set up and generated, we combine them together using pyAerostructure (Section 2.5). A set of rigid links connecting the two meshes is generated when the model is first initialized. These links are used to map the load-displacement transfer consistently throughout the optimization [12].

We use a Gauss–Seidel (GS) algorithm to converge the MDA, looping the loads and displacement between the aerodynamic and the structural modules respectively until the residuals in the two solvers are below a prescribed tolerance. Generally a relative convergence of 10^{-8} from the initial residual corresponds to individual discipline absolute residuals below 10^{-3}–10^{-4}.

Section 3.3.1 describes a verification study on the numerical consistency of the virtual work on the two models, to further quantify the numerical convergence of the analysis.

3.3.1 Load-displacement transfer verification

The load-displacement module used by MACH [12] is provably conservative and consistent. Nevertheless, we assess the numerical error affecting the turbine model used for our optimizations by comparing the work on the structural and aerodynamic models.

The MDA displayed in Figure 3.12 uses the most refined aerodynamic mesh (L0) and the isotropic model used for coupled optimizations in Chapter 5. The surface pressure distribution and structural stress are juxtaposed to highlight the consistent load and deformation trends.

For this analysis, we define an arbitrary input thickness distribution, which results in an oversized structure for the given inflow condition. The "stress con" variable shown in Figure 3.12 indicates the local margin from structural failure based on the von Mises stress criterion. Without

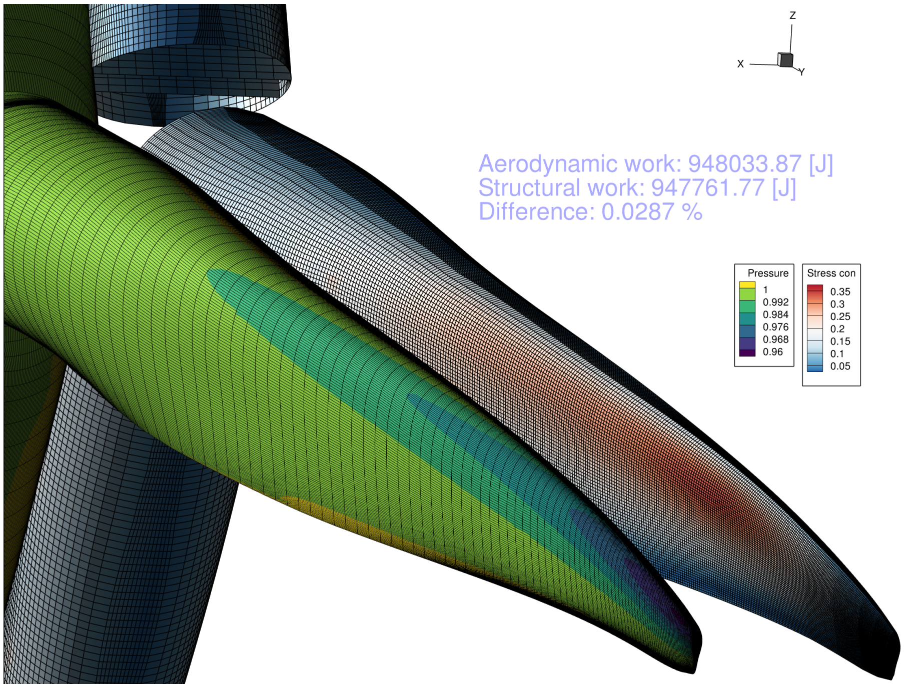

Figure 3.12: Converged MDA using the L0 aerodynamic mesh and the prototype structural mesh, including the numerical error estimation of the load-displacement transfer scheme.

considering a safety factor, if this "stress con" value exceeds one, it indicates that material failure occurs. We consider a 1.5 safety factor for the failure indicator evaluation in this example analysis, so the maximum allowable indicator is $\sigma_{\text{stress con}} = 1/1.5 = 0.667$. As shown in Figure 3.12, the example MDA analysis results in a maximum value $\sigma_{\text{stress con}} \leqslant 0.38$. The measured tip displacement is 4.1 m and the loss of torque with respect to the rigid aerodynamic model of approximately 1.9%.

The virtual work for the two models is given by:

$$W_{\text{aero}} = \mathbf{F}_{\text{aero}} \cdot \mathbf{u}_{\text{aero}}$$
$$W_{\text{struct}} = \mathbf{F}_{\text{struct}} \cdot \mathbf{u}_{\text{struct}}$$

The results of the dot product of the nodal displacement and force vectors (i.e. work) as obtained by the two solvers on the final MDA iteration are reported in Figure 3.12. We consider the relative error of 0.029% between W_{aero} and W_{struct} satisfying for the current set of models and optimization studies.

The MDAs for the load conditions and designs in Chapter 5, as the one shown in Figure 3.12, converge between 6 and 25 iterations. The iteration limit is increased to more than 40 to enable tighter convergence for more complex designs and inflow conditions investigated in Chapters 6 and 7.

Close to convergence, the computational cost of each iteration typically drops by one or two orders of magnitude. Overall, we typically observe an increase in the computational cost of approximately 80% compared to a stand-alone CFD solution. The cost of a full analysis is not directly proportional to the number of GS iterations. After the initial solution, the following iterations restart from the final states of the previous iteration.

Nevertheless, the increased analysis cost and the higher number of output functions with respect to the aerodynamic shape optimization studies in [15] makes these optimization studies numerically tractable only on High-Performance Computing clusters. Early optimizations from Mangano et al. [142] have been run on the University of Michigan's Great Lakes cluster, using 1 (L3 mesh) or 2 (L2 mesh) Skylake nodes with 36 CPUs for each node, with a wall time between 4 and 16 hours depending on the optimization cases. Larger cases using L1 meshes reported in this dissertation relied on the Texas Advanced Computing Center's Stampede2 cluster. The single point optimizations required up to 38000 CPUh using IceLake (Intel Xeon Platinum 8380) nodes with 80 CPU for each node, running for up to 5 consecutive 48-hour jobs. Multipoint optimizations have essentially the same wall time but a computational cost increased by a factor 2-3 depending on the number of design points evaluated in parallel and other load balancing considerations.

3.3.2 Free-form deformation grid setup

The rotor geometry is handled through the FFD approach [105, 106] introduced in Section 2.1. We leverage the hierarchical approach described by Hajdik et al. [104] to manipulate each blade consistently.

In general, deformations on the "parent" FFD block are extended to the "children". The displacement of "children" control points allows for localized manipulation on a subset of the aerodynamic and structural mesh points. Each FFD block has its own "reference axis", a user-defined curve that is used to define the displacement of multiple control points at the same time. Children FFDs are particularly suited to handle multiple rotor blades because of their independent reference axis.

We can define both local and global design variables. In the first case, individual control points are displaced along a selected axis. The shape variables used in Chapter 7 use this approach. In the second case, a group of control points is simultaneously scaled or rotated around the block reference axis. Planform and twist variables used throughout this work are defined as global variables. A random combination of global design variables is applied to the baseline rotor in Figure 3.13

The parametrization used in this work, an extension of Madsen et al. [15] original setup, embeds the whole rotor in a 8-point FFD parallelepiped block (not shown in Figure 3.13). This block

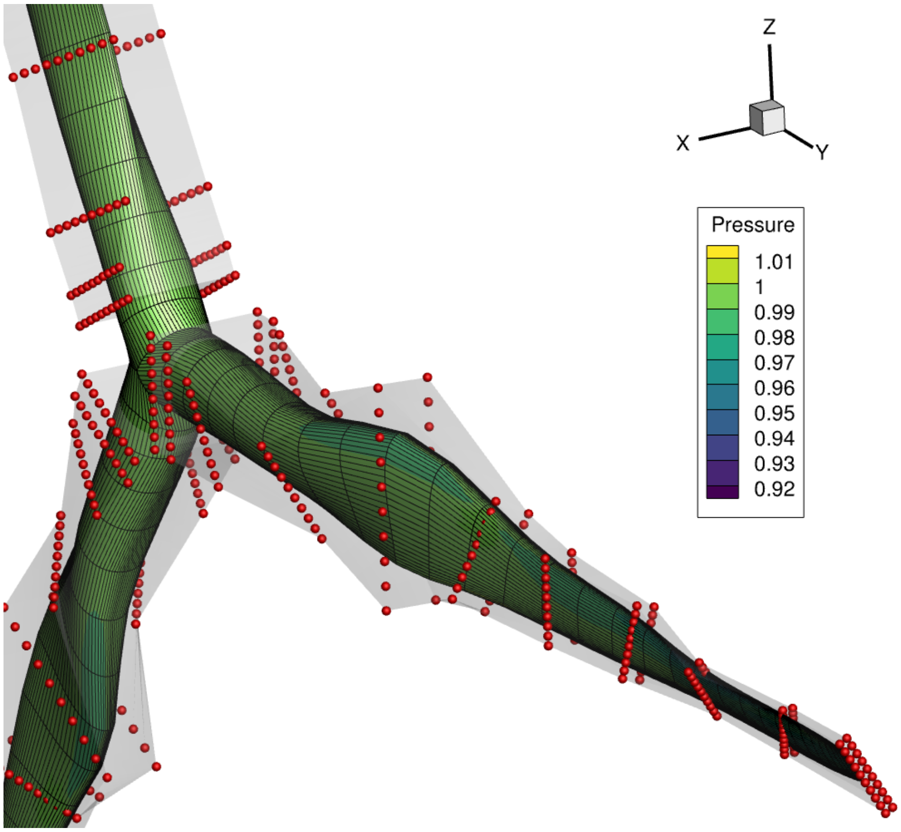

Figure 3.13: Example random geometrical deformation of the L2 aerodynamic mesh.

is only used to group the individual blade FFDs. Its control points are fixed during the optimization.

Each rotor blade is embedded in a single child FFD, whose reference axis is aligned with the blade spanwise axis at each section quarter-chord. The FFD deformations are applied consistently on each blade. The twist, planform, and shape deformation on blade #1 are exactly mirrored on blade #2 and blade #3.

We implemented a new feature in pyGeo [6] to handle global variables on FFD blocks whose reference axis is not aligned with the main system of reference. The user can now provide the axis angle offset when the FFD object is set up.

pyGeo will first rotate the FFD block to be aligned with the main reference axes, apply global and local deformations, and then rotate the FFD back in the original position. The control points deformation are propagated to the embedded mesh points only at this stage. This approach guarantees consistency between different children FFDs, without the need to define different design variables depending on the original reference axis angle. The bugfix also enables a more flexible reference axis definition within the FFD box along the x-y-z axes.

---

[6] pyGeo new feature PR on GitHub

Chapter 4

Turbine Performance and Operating Conditions

The previous chapters described the computational framework (Chapter 2) and the development and verification (Chapter 3) of the model used in this work. In this chapter, we outline some practical design, formulation, and setup considerations before we delve into the optimization results discussed in the rest of the dissertation.

The design and parametrization details of the rotor configuration are outlined in Section 4.1. The operating conditions, in terms of wind and RPM combination, used in our optimization studies are presented in Section 4.2. Finally, a derivation of more practical optimization metrics from Garcia-Sanz [11]'s LCOE definition is summarized in Section 4.3.

4.1 Baseline turbine rotor configuration

This section connects some of the modeling details from Chapter 3 with practical considerations for our optimization studies.

In this study, we use the DTU 10 MW configuration modified by Madsen et al. [15, Figure 2] starting from the original outer mold line and structural layup defined by Bak et al. [24] This is a three-blade rotor with 178.3-meter diameter and tailored chord, airfoil, and twist distributions.

As previously shown in Figure 3.1, the blades are blended at the root to limit the meshing effort and prevent CFD solver convergence issues [14]. This neglects the hub effects on the aerodynamic performance. Although the aerodynamic blade-nacelle interactions are a relevant design factor, they have a limited impact on the extracted torque for large turbines. The influence of these interactions on the blade structural sizing is negligible. Tower-blade interactions are also ignored.

Our simulations use the same aerodynamic meshes as those of Madsen et al. [15, Sec. 4.1], as described in Section 3.1.1. That previous work compares ADflow and DTU's Ellipsys3D analysis on the same meshes and flow conditions. We verified that we can replicate their mesh convergence study. As mentioned in Section 2.3, we use the SA turbulence model for our adjoint-based optimization. Our analyses do not consider laminar-turbulent transition, and the flow over the blades is assumed to be fully turbulent. The implications of these model assumptions have been outlined in Section 1.4 and are further addressed in Chapter 5.

As mentioned in the previous chapter, ADflow is particularly sensitive to mesh resolution in the incompressible flow regime [15]. We use the L1 mesh for optimizations to balance accuracy and computational cost. This choice enables a larger set of optimization studies than what could be obtained with the L0 mesh at a comparable computational expense. Even considering the other assumptions on the structural and aerodynamic models and solvers, the results presented in this paper are sufficiently accurate to discuss design trends with different problem formulations and highlight the benefits of our proposed methodology. This reflects the approach used by Madsen et al. [15] and is consistent with previous MACH-based work [88].

Using the L1 mesh ($1,769,472$ cells), a single aerodynamic analysis converges in less than 30 CPU-hours. This coarser mesh tends to overpredict the loads over the span by a consistent offset compared to the L0 mesh, but preserves the flow characteristics over the rotor. The computational cost of an optimization iteration for the L0 mesh can be up to 10 times higher (Table 3.1) than for the L1 mesh used for this study, which would extend the optimization time to over a month (see Section 3.3). ADflow is a density-based flow solver originally developed for transonic flow; this poor cost scaling is partly due to the solver behavior at incompressible regime (see more in Section 3.1).

Improving the efficiency of ADflow for large meshes at low-speed conditions using an improved preconditioner is an active area of development [146]. Although the L0 mesh is not used in optimizations, we re-analyze the optimized designs with the L0 mesh for verification. In the future, we might consider using pressure-based solvers such as DAFoam [147] or EllipSys3D [148]. DAFoam has been coupled to structural and acoustic analysis tools [149] through MPhys [1] and OpenMDAO [40].

The discussion by Madsen et al. [15] shows that the resulting high-level aerodynamic design is consistent with the more expensive L0 mesh result. A convergence study with an even finer mesh (47.8 M cells) showed that the L0 mesh is within the asymptotic region, verifying that this mesh is adequate for capturing practical design trends.

4.1.1 Structural model

The development and verification of a MACH-compatible structural model for the DTU 10 MW is one of the early contribution of my dissertation work. The definition of the structural topology, grid refinement, and material properties and layup definition presented in Section 3.2 has been an iterative process resulting in two configurations.

---

[1] https://github.com/OpenMDAO/mphys

A first mesh prototype has been used for the optimization studies in Chapter 5. This mesh uses ‘isoFSDT’ elements and isotropic material properties reflecting aluminum-2024. Instead of performing a full grid convergence study, we opted to build an overly-refined mesh and move ahead with the load-displacement verification (Figure 3.12) and a first set of optimization studies. This mesh is presented in Section 4.1.1.1.

For the following multipoint and fiber tailoring studies, presented in Chapter 6-Chapter 7, we developed an updated grid. We switched to ‘smearedFSDT’ elements to enable composite layup stack modeling. Both constitutive models are discussed in Section 3.2.2.

The structural topology had minor adjustments to reflect more accurately the spar location from the original benchmark [24]. The same document has been used as a reference to define the composite layup fraction and material properties. The grid update included a convergence study, which has been discussed in Section 3.2.1. The updated structural model is presented in Section 4.1.1.2.

4.1.1.1 Prototype isotropic model

The model presented in this section has been used by Mangano et al. [142] to perform the first set of high-fidelity optimization studies. The mesh follows the outer mold line contour and includes an internal main shear web and a reinforcement spar in the trailing edge area. CQUAD hexahedral linear shell elements are used to model the structural components. The grid has $208,464$ elements. Three identical blades are fixed at the root, located 3 m from the rotation axis. All the degrees of freedom of the nodes are fixed at the root. Similar to the aerodynamic analysis, the structural analysis is static; dynamic stability is not considered in this work. The sizing process for the baseline rotor through sequential structural optimization is outlined later in Section 5.2

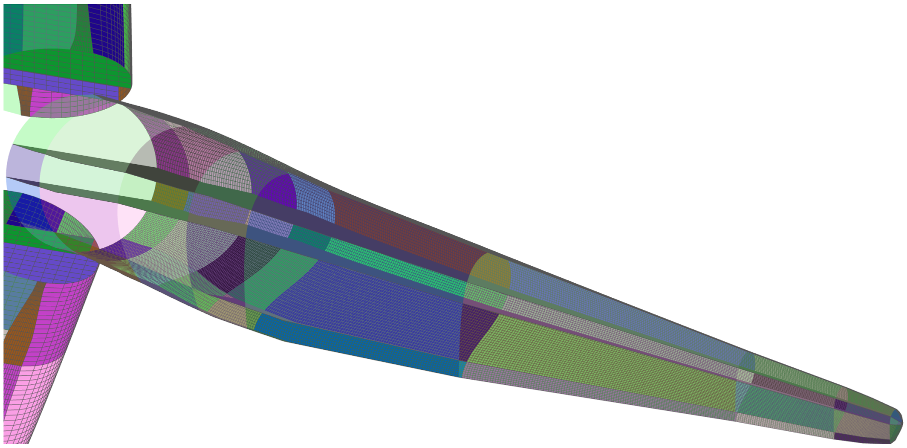

Figure 4.1: Prototype structural mesh and parametrization. Color patches show sections controlled by different independent DVs.

We use isotropic shell elements made with 2024 aluminum alloy properties. This structural model is relatively simple because Chapter 5 primary objective is to demonstrate our optimization framework's capabilities. Although this isotropic material assumption might lead to less practical designs, the design trends and insights are valid and valuable.

Figure 4.1 shows this prototype structural model. The parametrization of the different structural panels is also simplified at this stage. The layout is a compromise between design freedom and complexity. We selected a panel discretization that follows the main geometry breaks of the outer mold line, leading to a total of 9 spanwise sections. For a given spanwise section, the optimizer sizes every spar, spar cap, and leading and trailing edge panels independently. As discussed in detail later in Section 5.1.2, this amounts to a total of 117 structural thickness design variables.

4.1.1.2 Composite structural model

Composite structural models including fiber angle or ply fraction design variables have been used in previous works with MACH [150, 95, 101, 102], which leveraged structural anisotropy and tailored bend-twist coupling properties.

We combined the existing numerical tools with a revised workflow to finally include anisotropic material properties in our wind turbine optimization studies. The results illustrated in Chapters 6 and 7 use the mesh and panel discretization shown in Figure 4.2

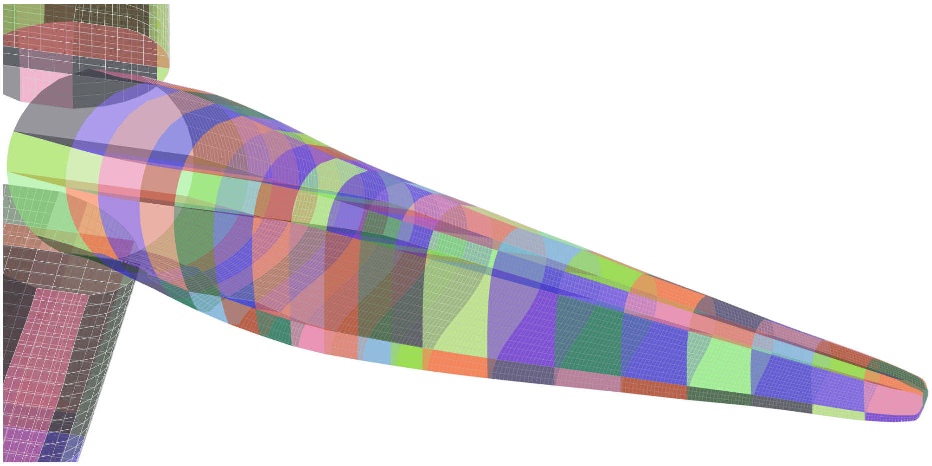

Figure 4.2: Updated structural model parametrization, discussed in Appendix B.

The meshing workflow (Section 3.2.1) initially developed for the prototype mesh in Section 4.1.1.1 has been extended and refined to generate a coarser grid without losing accuracy. The computational cost of TACS simulations is already orders of magnitude lower than CFD simulations.

Nevertheless, a coarser mesh is advantageous for aerostructural MDAO. As discussed in Appendix A, an overly refined mesh can lead to spurious numerical stress peaks in a FEM simulation. Larger structural elements also reduce the number of rigid links of our load-displacement transfer (Section 2.5), reducing initialization time and complexity, with possible benefits for the aerostructural analysis convergence. Details on the parametrization selection process are discussed in Appendix B.

4.2 Design load cases

Wind turbines experience a range of wind conditions over their operating life. As discussed in Section 1.3.1, designing the rotor using site-specific meteorological assessments is crucial to maximize the energy yield of the system. At the same time, structural sizing must account for nominal, extreme, and cyclical loads with a reasonable margin of safety.

Load paths on the turbine structure cannot be identified a priori. Realistic scenarios need to be simulated on coupled aerostructural systems to assess the static and dynamic response of the system. The metocean conditions experienced by the turbine are defined as “load cases” [5, Sec. 6.4]

International standards such as IEC61400-3 [25] outline the certification requirements for industrial wind systems. They define ranges of wind average velocity, turbulence, and operating conditions to be tested. These load cases include nominal power production (DLC 1.1-1.5), electrical faults (DLC 2.1-2.4), start up (DLC 3.1-3.3), shut down (DLC 4.1-4.2 and 5.1 for emergency), and parked conditions (DLC 6.1-6.4).

We use a subset of these DLCs for life-cycle damage estimation in Chapter 8. Due to the steady-state nature of our CFD-based aerodynamic analyses, we limit our high-fidelity design points to nominal production cases without accounting for turbulent flow variations.

Wind systems can operate within a range of wind velocities. Below a certain wind intensity, called *cut-in* speed, there is not enough energy available to start-up and operate the turbine. Conversely, the turbine is shut down when the wind goes beyond a certain *cut-off* threshold to preserve the electromechanical integrity of the system. The DTU 10 MW has a cut-in speed of $4\mathrm{m\,s^{-1}}$ and a cut-off speed of $25\mathrm{m\,s^{-1}}$ [24, Table 2.1].

Within this operating range, three regions are defined. An example from the IEA 15 MW benchmark is reported in Figure 4.3. The control policy for the DTU 10 MW is comparable with this newer benchmark. In the low speed region 1.5, the controller adjusts the pitch at constant rotation speed, but the TSR remains sub-optimal. The minimum RPM is set to avoid structural resonance issues. In region 3, the turbine operates at rated generator power and fixed RPM. The pitch decreases with wind speed to not exceed the prescribed torque.

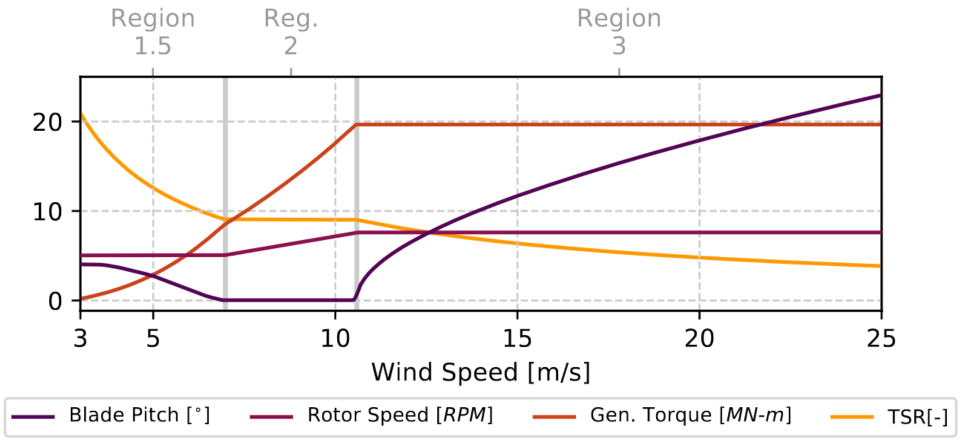

Figure 4.3: Operating regions for the IEA 15 MW turbine, from Gaertner et al. [22, Figure 3.1].

Simulating operating conditions in these regions requires time-accurate analyses and a controller model. Conversely, the controller at the below-rated region 2 tracks the set point TSR. For performance analysis, we can select a combination of wind speed and rotation rate for our coupled aerostructural analyses. Madsen et al. [15, Figure 16] select two points in region 2 and one at the border with region 1.5 for their CFD-based multipoint optimization studies. Section 4.2.1 details how we adapt and extend their design point selection to our multipoint optimization studies.

4.2.1 Inflow conditions for high-fidelity MDO

In Chapter 5, we limit our optimization studies to a single design point to demonstrate the framework capabilities. The "aerodynamic" objective for our optimizer is solely torque maximization at a prescribed inflow condition. We use the same reference below-rated-power design condition from the previous work on aerodynamic shape optimization [15] – namely a uniform inflow velocity V=8 $\mathrm{m/s}$ and a TSR = 7.8.

The approach is extended to a multipoint strategy [151, 132] in Chapters 6 and 7, where we average the performance of the turbine over a few selected design points. The point selection and weighting in the objective function is discussed later in this section. Optimizations in Chapter 6 just use two design points at the same time, to assess the impact of fiber tailoring on bend-twist coupling under different loads. The studies in Chapter 7 conversely investigate the impact of the design point selection.

In the optimizations presented in Chapter 8, we include fatigue and extreme load conditions to size the internal structure of the rotor. The ongoing work presented by Caprace et al. [87] overcomes the shortcomings of the high-fidelity approach by including low-fidelity-informed load distributions in the optimization loop.

We optimize the baseline DTU 10 MW rotor in a constant uniform flow for a set of wind speeds at a prescribed rotation rate, mirroring the RPM-based control schedule of the benchmark rotor. Since we perform steady-state analyses and do not model controls in-the-loop, we ignore the dynamic pitch response at rated power conditions.

In Chapter 7 we compare single point optimizations, 3-point optimizations with different weights on torque, mass, and design points, and a 5-point optimization. The wind speeds and TSR for each design point are listed in Table 4.1.

Table 4.1: Wind speeds and rotation rates for the design points included in the optimizations.

Name	Focus	V [$\mathrm{m/s}$]	TSR	RPM
High-wind (HW)	Structural Sizing	11.0	7.8	9.19
Below-rated (BR)	Performance	8.0	7.8	6.68
Min-RPM (MR)	Low-speed performance	5.0	11.2	6.00
Below-rated, low TSR (BR-T5)	Off-design performance	10.2	5.5	6.00
Below-rated, high TSR (BR-T10)	Off-design performance	8.9	10.1	9.60

The BR case (V=8 $\mathrm{m/s}$, TSR=7.8) has been used in previous DTU 10 MW optimization studies [15, 142] and for the studies in Chapter 5. We consider two other combinations of wind speed and TSR used by Madsen et al. [15] for the DTU 10MW rotor: a ‘HW’ condition at the same optimum TSR and V=11 $\mathrm{m/s}$, and a MR condition at V=5 $\mathrm{m/s}$ and TSR adjusted to match the minimum design rotation rate of 6 RPM. These additional design points correspond to different angle of attack distributions over the blade. Considering them simultaneously in the optimizations ensures better performance over the below-rated operational envelope [15]. Moreover, the HW steady-state loads, just below peak thrust, are used for structural sizing.

The optimizations in Section 7.5 consider two additional wind conditions, labeled BR-T5 and BR-T10 respectively. Our analyses do not model gusts or inflow variations over a rotor revolution. In these scenarios, the local angle of attack over the blade can vary by up to 3-4° [15]. Since the TSR is linked to the local inflow angle distribution over the blade, we perturb the design TSR=7.8 by ± 2.3 to account for local inflow angle variations of roughly ± 3° over the blade. This value has been selected taking into account both the “distance” from other design points, and the design RPM upper (9.6 RPM) and lower (6 RPM) limits. The BR-T5 case is particularly relevant due to the lower TSR corresponding to higher angles of attack over the blade. This design point is considered a proxy for a stall margin constraint.

The relative weight of individual flow conditions, and thus their impact on the final design, is defined by a set of γ_k coefficients in the weighted-average objective function (Eq. 7.1 discussed later in Section 7.1.2). For the simplest approach (case 3pt-Eq in Table 7.3), we assign the same relative weight for every design point. However, wind engineering best practices (discussed in

Section 1.3.1) recommend using site-specific wind distributions to tailor turbines and farms layouts to available resources.

For the reference optimization 3pt-Wb (see Table 7.3), the γ_k coefficients are adjusted to mimic a notional wind probability distribution. We use a Weibull function - Eq. 4.1 - with average speed $A = 8\mathrm{m/s}$ and form factor $k = 2$, consistently with Madsen et al. [15]:

$$f(x; k, A) = \frac{k}{A}\left(\frac{x}{A}\right)^{k-1} e^{-(k/A)^k} \tag{4.1}$$

Due to the computational cost of the CFD-CSM solver, the present work considers a limited set of wind speeds compared to BEMT-based studies.

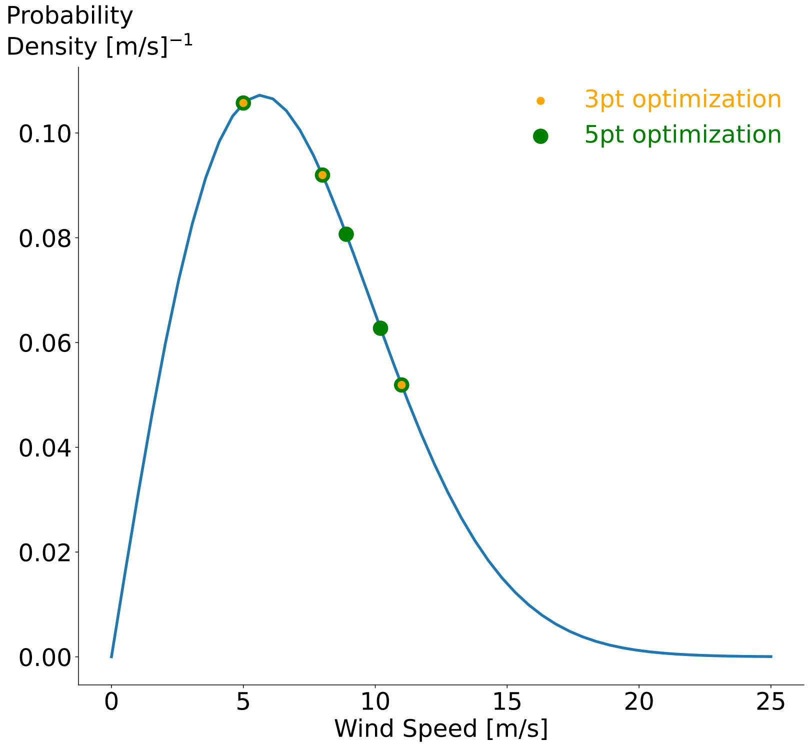

Figure 4.4: The design points from Table 4.1 are highlighted on the selected Weibull curve describing wind probability distribution at a nominal site.

The selected set of wind speeds is too small for a conventional binning strategy to provide meaningful probability values. Instead, we directly use the probability at a given wind speed to weigh design points in the objective. The values are then normalized so that $\sum_{k \in K} \gamma_k = 1$ and used as weights in the objective function. The curve and selected design points are highlighted in Figure 4.4.

4.3 Optimization metrics

The general constrained optimization problem formulation outlined in Section 1.2 is defined as follows:

$$\begin{aligned} &\text{minimize} \quad \mathbf{f}(\mathbf{x}) \\ &\text{by varying} \quad \mathbf{x}_{\min} \leqslant \mathbf{x} \leqslant \mathbf{x}_{\max} \\ &\text{subject to} \quad \mathbf{g}(\mathbf{x}) \leq 0 \end{aligned}$$

where $\mathbf{f}(\mathbf{x})$ is a designer-defined objective function that quantifies the merit of the design, $\mathbf{x}$ the complete set of design variables, and $\mathbf{g}(\mathbf{x})$ a set of linear and nonlinear constraints. The upper and lower bounds of the design variables $\mathbf{x}_{\max}$ and $\mathbf{x}_{\min}$ are handled separately from other analysis-dependent constraints $\mathbf{g}(\mathbf{x})$. Design variables and constraints have been outlined in Chapter 3 and are further detailed of each optimization study in the following chapters.

Concerning the overall design objectives, there is a general consensus in academia and industry about using cost metrics (Section 1.1.3) as drivers for wind turbines and farms, often defined starting from a system-engineering approach [152, 72, 59, 28]. In practical design terms, the cost of energy is driven by the power extraction of the turbine and the sum of investment and operating costs. Neglecting the market fluctuations, the cost-of-energy corresponding to a given design is affected by both the pure aerodynamic efficiency under given conditions and the overall mass of the system. In this work, we derive our optimization metrics from Garcia-Sanz [11] to measure the performance of our designs.

This new approach proposes to decouple the primary design drivers determining the levelized cost of energy (LCOE) from other environmental and economic factors. In its conventional definition, the LCOE is given by:

$$\text{LCOE} = \frac{\text{FCR} \times \text{CapEx} + \text{OpEx}}{\text{AEP}} \tag{4.2}$$

where OpEx stands for operation costs, AEP stands for the annual energy production, FCR stands for fixed charge rate, and CapEx stands for the capital expenditure. Assuming fixed OpEx, AEP is influenced by site-specific factors such as the wind probabilistic distribution, while CapEx is strongly dependent on the cost of steel, which in turn follows the market fluctuations.

The cost and performance components that determine the LCOE are split into two separate metrics in this work. Our focus is restricted to the technological primary design factors without considering these fluctuations. M_1 is introduced to assess the “air-to-electron” generation efficiency

$$M_1 = \frac{\sum_{k=1}^{N} P_e(k)}{\sum_{k=1}^{N} P_w(k)} \tag{4.3}$$

where the extracted electrical power, $P_e(k)$, and the total wind power, $P_w(k)$ are defined as

$$\begin{aligned} P_e(k) &= \frac{1}{2}\rho A_r(k) C_{p_{\max}}(k)\mu(k)V^3 \\ P_w(k) &= \frac{1}{2}\rho A_r(k) V^3 \end{aligned} \tag{4.4}$$

Here k is the turbine index, N is the total number of turbines under evaluation, V is the inflow velocity, A_r is the swept area, $C_{p_{\max}}$ is the maximum turbine power coefficient, and μ is the electromechanical loss coefficient.

The generalized cost of the farm is defined as the ratio of its swept area over its equivalent mass of steel (M_2):

$$M_2 = \frac{\sum_{k=1}^{N} A_r(k)}{\sum_{k=1}^{N} M_{eq}(k)}$$

where the equivalent mass of steel M_{eq}, and the component mass m_j are defined as:

$$\begin{aligned} M_{eq} &= \sum_{j=1}^{Z} m_j \\ m_j &= m_{cj} f_j \end{aligned}$$

with Z the total number of components, f_j the aggregated material, manufacturing, and installation factors and j the component index.

As we aim to analyze a single, isolated turbine rotor, these metrics are simplified to give:

$$M_1 = C_{p_{\max}} \tag{4.5}$$

$$M_2 = \frac{A_r}{M_{eq}} \tag{4.6}$$

where M_{eq} refers to the sole rotor mass for our optimization problem. A_r can be neglected since we do not change the rotor size. The loss coefficient μ is neglected as the design optimization does not affect the gearbox and the generator. The M_1 function is not as comprehensive as the AEP metric to assess the “lifetime” power generation of a turbine, which must also take the probabilistic distribution of the wind resource into account. However, it is representative of the power production in region II (Section 4.2). As we are not considering the peak torque for every configuration, we ultimately use $M_1 = C_P \propto Q_x$ evaluated at a chosen inflow velocity and TSR to compare the different designs in this paper.

Regardless of the number of inflow conditions considered, having two independent system metrics gives us the opportunity to compare different optimized designs with different weighted combinations of the metrics to be minimized, for example:

$$f(\mathbf{x}) = \omega\, M_1 + (1 - \omega)\frac{1}{M_2} \tag{4.7}$$

with ω the weighting parameter. Shifting ω between 0 and 1, the designer can obtain turbine rotor designs that emphasize mass reduction or increased energy yield respectively.

The optimization studies presented in the following chapters use this high-level formulation. In Chapters 5 and 8, M_1 and $\frac{1}{M_2}$ correspond to the torque and mass of the rotor respectively. The optimizations in Chapters 6 and 7 replace M_1 with a weighted average of the rotor torque at the different inflow conditions considered.

Chapter 5

High-Fidelity Rotor Aerostructural Optimization

In this chapter, we present the first set of optimization studies using our high-fidelity model.

On the one hand, we showcase the capabilities of the MACH optimization framework, whose features are summarized in Chapter 2 and detailed in the literature [12]. Its geometry manipulation and aerodynamic modules have been updated to handle wind turbine optimization studies. The structural solver has also been updated to include centrifugal forces on the rotating blades. We further implemented changes originally made by Madsen et al. [15] in the sensitivity analysis code and verified that the adjoint-based derivatives for rotating frames are accurate (see Section 3.1.3). These updates are publicly available on the ADflow open-source repository.

On the other hand, we discuss the results of a set of optimization problems to highlight the benefits of this high-fidelity approach for wind turbine design.

As discussed in Section 4.3, the objective function of these optimizations combines mass and rotor torque at a single inflow condition. The resulting design features are thus specific to the selected inflow regime, and we do not consider the performance at other operating points in this chapter. The objective weight ω from Eq. 4.7 is adjusted to prioritize aerodynamic efficiency or mass reduction.

The problem solved in this chapter is the first step for this type of CFD-CSM wind turbine optimization and is thus simplified. The structural model uses isotropic material properties based on Aluminum-2024. This is done to reduce the problem's complexity and simplify model pre-processing. The studies in Chapter 6 and Chapter 7 relax this assumption by including material anisotropy in the structural model. Preliminary work on composite materials has been presented by Mangano et al. [142].

Including additional constraints or improving the aerodynamic predictions with a transition model would likely lead to a more conservative blade design than the ones presented in the current work. Including stall margins would also have comparable effect. Results discussion later in Section 7.5 investigate this issue more in detail.

To address these concerns for the study in this chapter, we inspect the flow over the blade, estimate local margins from the local $C_{l_{\max}}$ angle of attack, and compare integrated loads to verify consistent load increase patterns. We also perform post-optimality analyses with higher and lower fidelity simulations to rule out aerodynamic code weaknesses exploited by the optimizer.

5.1 Problem Formulation

The turbine efficiency in harvesting power from the wind and the sum of investment and operating costs drive the cost of energy, conventionally condensed in the LCOE presented in Section 1.1.3. Since we are limited to a single turbine rotor at a single inflow condition, LCOE and comparable metrics cannot be directly computed for our model. The derivation presented in Section 4.3 illustrates how we simplify the metric proposed by Garcia-Sanz [11] for an individual turbine with fixed inflow velocity and swept area. The cost estimate assumes that the manufacturing process does not substantially change for the optimized layouts.

Thus, this multidisciplinary design optimization study explores tradeoffs between the rotor mass and its power output as a proxy for costs and energy yield respectively. We use structural thickness and different sets of geometrical variables to tailor the blade design and its steady-state aeroelastic response while satisfying nonlinear structural and performance constraints.

5.1.1 Optimization objective

We use two approaches to explore the design space. The first approach combines the torque and mass metrics into a single objective function as a weighted sum,

$$f(\boldsymbol{x}) = -\omega \frac{Q_x}{Q_{x_i}} + (1-\omega)\frac{M}{M_i}, \tag{5.1}$$

where Q_x and M are the torque and mass, respectively, and $0 \leq w \leq 1$ is an arbitrary coefficient. Both metrics are normalized with respect to their baseline values Q_{x_i} and M_i. The negative sign in the torque component is necessary because the optimizer minimizes the function $f(\boldsymbol{x})$. Comparing Eq. 5.1 with Eq. 4.7, we see how $M_1 = -Q_x/Q_{x_i}$ and $\frac{1}{M_2} = M/M_i$ Solving the optimization problem defined by Eq. 5.1 for values of ω ranging between 0.7 (emphasis on torque) and 0.5 (same emphasis on torque and mass) results in a Pareto front defined by the torque and mass metrics.

The second approach uses the epsilon-constraint method to explore the Pareto front [36, Sec.9.3.2]. Using this method, we perform mass minimizations ($f(\boldsymbol{x}) \equiv M/M_i$) with a constraint on the torque. We vary the torque constraint value to more accurately add points in the Pareto front. We compare the results of these approaches in Section 5.5.

5.1.2 Design variables

The FFD approach (Chapter 2 and Section 3.3.2) applied to the blade-resolved layout of the aerodynamic and structural meshes enables a flexible and effective parametrization of geometry components whose shape would otherwise be approximated or kept fixed through the optimization. Conventional approaches for turbine design are either limited in the number of high-level parameters directly used within an optimization or require additional external iterations to update blade section drag polars and stiffness properties between optimization loops. The high-fidelity approach enables the optimizer to alter detailed layout features, such as airfoil shape and individual structural panel thicknesses, within a monolithic optimization run. This reduces human intervention in the loop and avoids the discontinuities in the design space introduced by conventional multi-level strategies. In addition to the increased control on local variables for design and manufacturing considerations, using a representative 3D model gives a more accurate system mass estimation using an element-by-element calculation based on the element area and the thickness design variables.

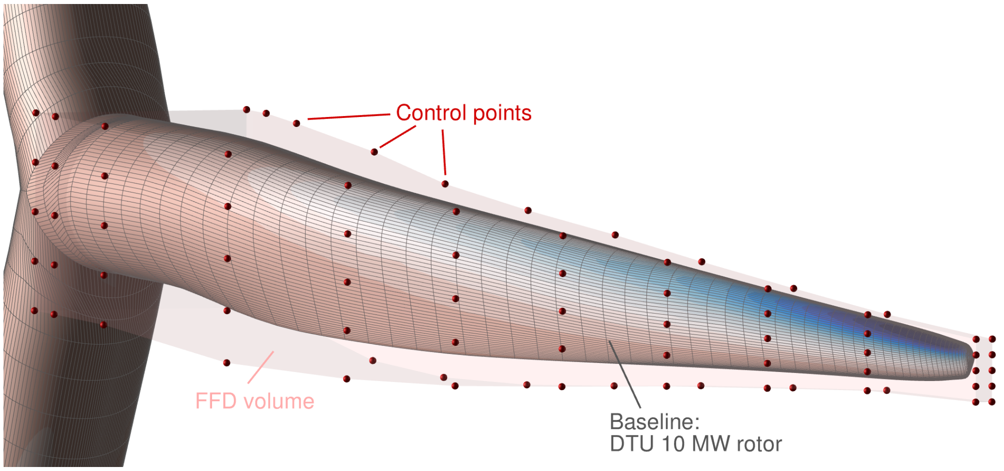

Figure 5.1: The rotor blades (with highlighted pressure contour here) are embedded in the FFD volume, and the shape is manipulated through control points displacement.

A snapshot of the aerodynamic model embedded in the FFD control box is illustrated in Figure 5.1. We group the control points of the pre-processed FFD grid and parametrize their displacement through user-defined functions passed to pyGeo - as discussed in Section 3.3.2. In this way, we model planform design variables and local airfoil shape deformations using the same geometry object. The pitch and twist variables are defined as rotations of the FFD grid control sections around a user-defined axis - located at the airfoil quarter chord in this case. Chord and thickness changes are obtained by scaling the control sections along the z-axis and x-axis respectively. Results in this chapter use either the same factor to scale chord and thickness (preserving the airfoil shape) or decoupled coefficients to give the optimizer more design freedom. The former approach mimics the parametrization capabilities of conventional design tools. With the latter approach, the independent scaling of chord and thickness changes the airfoil section shape in a way that goes beyond the capabilities of these lower-fidelity tools.

The parametrization scheme can model additional geometric design variables not currently included in our optimizations. Pre-cone and pre-bend can be obtained by displacing airfoil sections off-plane. Airfoil shape can be changed by displacing single control points [15]. The latter set of variables is used later in Chapter 7.

We split the mesh into nine spanwise sections to parametrize the structural model. We further split the upper and lower skins into four chordwise sections. Each of these structural panels is assigned an independent design variable. This results the 117 structural thickness design variables, shown as different color patches in Figure 4.1.

5.1.3 Constraints

The structural model addresses structural feasibility considerations and avoids the need for geometrical thickness and root bending moment constraints that are usually enforced in aerodynamic-only shape optimization studies [15, 14]. TACS is used for structural sizing, taking into account local element-wise stresses.

To implement stress constraints efficiently and effectively for gradient-based optimization, we use the KS constraint aggregation technique mentioned in Chapter 2. This is useful for two reasons. Firstly, constraining the mechanical stress shell-by-shell would increase the number of functions of interest for the optimization by up to six orders of magnitude, negating the benefits of the adjoint method for derivative computation. Secondly, a function that only considers the overall maximum stress over the blade would be non-smooth because the most critically stressed element can vary between successive optimization iterations, hindering optimization convergence.

We define three aggregated KS stress constraints for this problem: the blade upper skin, lower skin, and spars, respectively. The inflow condition we are considering is not critical for structural sizing, so the optimized rotors are characterized by thin structural panels.

As a consideration for tower structure safety, we enforce a thrust constraint in our formulation to ensure that the loads from the rotor to the tower remain within a threshold relative to the baseline. We limit the increase in the net rotor thrust at the given operating conditions to 14%, as done by Madsen et al. [15] for the same rotor configuration. Given the lower loads generated by the flexible baseline rotor, this constraint is more conservative than in the previous aerodynamic-only shape optimization study [15].

Finally, to ensure that there are no abrupt thickness changes between adjacent panels, we enforce linear adjacency constraints in the structural model so that two structural panels next to each other have less than a 5 mm difference in thickness for the current problems.

As mentioned in Section 5.1.1, the mass minimization optimizations enforce a torque constraint. Similar to the objective function, the constraint is normalized with respect to the baseline configuration performance.

5.1.4 Design load case

In this work, we limit our optimization studies to a single design point to demonstrate the framework's capabilities. As mentioned in Section 4.2.1, we use $V = 8\ \mathrm{m/s}$ and TSR$= 7.8$ as done by Madsen et al. [15]. The same approach could be extended to a multipoint strategy [151, 132, 15], averaging the turbine performance over 3 to 10 selected inflow cases.

Later Chapter 7 presents 3-point and 5-point optimizations results, while the fiber angle tailoring studies in Chapter 6 consider only 2 design points simultaneously.

5.1.5 Problem summary

Table 5.1 lists the problem objective, variables, and constraints.

Table 5.1: Aerostructural optimization problem statement.

	Name	Symbol	Qty
Objectives	Torque	Q	1
	Mass	M	1
Design Variables	Panel thickness	$\mathbf{t_p}$	117
	Twist	$\mathbf{x_{tw}}$	7
	Airfoil scale	$\mathbf{x_{sc}}$	7
	Chord	$\mathbf{x_{ch}}$	7
	Airfoil thickness	$\mathbf{x_{tk}}$	7
Constraints	Stress	$\mathrm{KS}_\sigma \leqslant 1$	3
	Adjacency		318
	Torque	$Q_x \geqslant Q_{x_{\mathrm{ref}}}$	1
	Thrust	$F_x/F_{x_\mathrm{i}} \leqslant 1.14$	1

We allow the displacement of the outermost seven FFD sections shown in Figure 5.1, while the other sections are fixed. The twist derivatives on these fixed sections are flatter than their outboard counterparts. Enabling twist in this area harms the optimization convergence without noticeable performance benefits. The rest of the DVs are applied to the same outboard sections to ensure a fair comparison.

Such a large nonlinear problem requires some precautions to ensure good numerical behavior. The planform variables derivatives are scaled by a factor of 10 in pyOptSparse to speed up numerical convergence. We set the SNOPT convergence tolerances for both the feasibility and optimality conditions to 10^{-6}. While feasibility is quickly satisfied, optimality is slower to converge because of the high number of design variables and the design space non-linearity. To limit the number of subsequent optimization restarts on HPC systems, we generally accept optimized layouts showing a plateau in the objective function and an optimality close to 10^{-4}, yielding a decrease of more than two orders of magnitude with respect to the optimality in the baseline design. The computational cost and convergence are discussed in more detail in Section 5.4.

Because of the assumptions on material properties for the structural model and the single below-rated inflow used in our studies, we do not start monolithic aerostructural optimizations from the benchmark turbine described by Bak et al. [24]. Instead, we use a loosely-coupled approach based on a sequence of structural optimizations to define a baseline structural layout. The resulting structural thickness distribution constitutes a more reasonably sized layout for the inflow conditions reported above. This preliminary step, discussed in Section 5.2, reduces the overall computational cost. Moreover, it improves the robustness of the subsequent coupled optimization, which has to traverse a smaller portion of the design space. We show how this latter approach is necessary to fully exploit the aerostructural trade-offs in the rotor design by comparing the loosely-coupled and monolithic, tightly-coupled optimizations.

This chapter presents multiple optimizations. We first compare the loosely coupled sizing approach to a tightly coupled optimization strategy that optimizes only structural variables. We discuss the trade-off between computational cost and design performance for these two approaches. Then, we present several optimization studies to quantify the trade-off between rotor mass and torque for different design variables. The three different sets of design variables used for optimization — Tw, TwSc, and TwChTk — are listed in Table 5.2.

Table 5.2: Design variable combinations for the optimization studies.

	Panel thickness	Twist	Airfoil scale	Chord	Airfoil thickness
Tw	✓	✓			
TwSc	✓	✓	✓		
TwChTk	✓	✓		✓	✓

Design trends over the span and blade sections are analyzed in detail, leveraging the blade-resolved model used by the CFD-CSM solver. Finally, we explore the effect of varying the optimization problem formulation on the rotor mass and torque.

As mentioned in Section 5.1.4, we optimize for a single inflow condition. Therefore, we focus our discussion on the tool capabilities and design trends rather than actual design parameters, which require broader life-cycle operating conditions to be practical.

5.2 Preliminary structural sizing through loosely-coupled approach

This first section explains the initial structural sizing obtained through a sequential, loosely coupled aerostructural design strategy and compares it to the tightly coupled MACH approach. The steps of this sizing process are illustrated in Algorithm 1.

We run a sequence of structural optimizations and aerostructural MDAs instead of a single structural optimization. In this way, we consider the change in aerodynamic loads as the single-discipline optimization alters the spanwise structural displacement. With this approach, we use only a subset of MACH packages. The aerostructural analysis is performed outside the optimization loop. The structural model is connected directly to the optimizer through pyOptSparse. Only structural design variables are included in this optimization. The blade geometry is not altered at this stage, and the aerostructural coupling is not accounted for during each optimization. This is why we refer to this approach as "loosely coupled" optimization.

Algorithm 1 Loosely coupled structural optimization

```
Define initial layout
for $i \leftarrow 0$ to $N$ do
    1: Run MDA
    2: Update aerodynamic input forces
    3: Run structural optimization
    4: Update structural layout
end for
```

A simplified XDSM matrix representing this loosely coupled approach is shown in Figure 5.2. The MDA in the upper left corner is identical to the tightly coupled MDA shown in Figure 2.1.

The first aerodynamic load distribution is obtained by running a coupled aerostructural analysis on a rotor with an arbitrary uniform thickness distribution. The optimizer minimizes the blade mass using this initial load while enforcing constraints on the aggregated stress. Once the structural optimization is completed, a new aerostructural MDA runs using the optimized thickness distribution just obtained. The loads from the last MDA are passed to the structural solver for new structural optimization. We run a sequence of 13 MDA-optimization iterations until convergence.

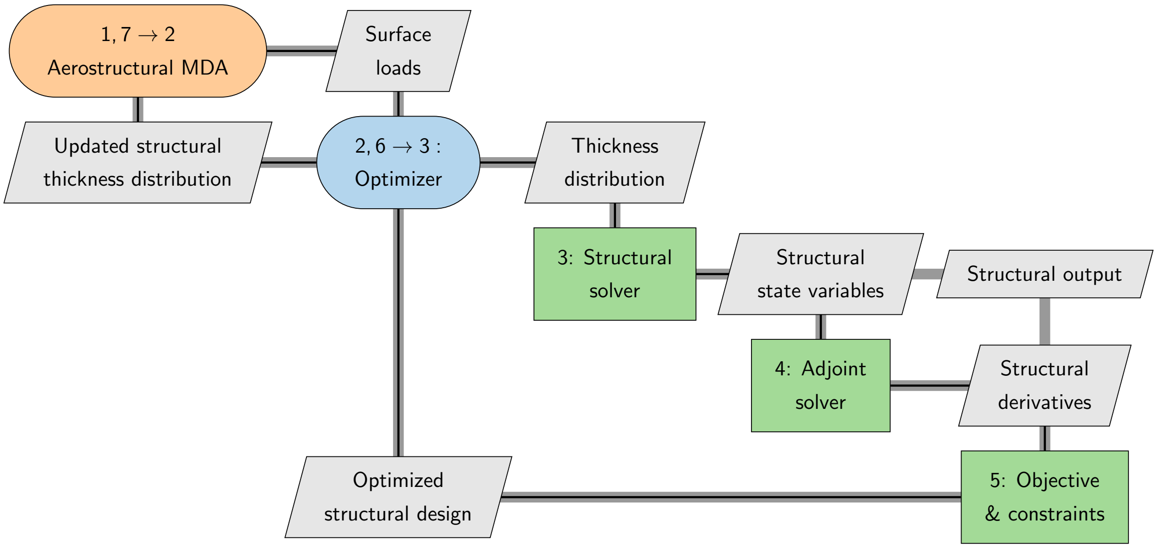

Figure 5.2: XDSM for the loosely coupled optimization approach.

For this loosely coupled approach, convergence means that any new structural optimization using an updated set of loads converges at the first (optimizer) iteration.

Because the optimization is purely structural, we cannot include geometric design variables and torque or thrust constraints in the optimization problem. The structural thickness distribution obtained through the loosely coupled approach is suboptimal from an aerostructural standpoint, but it is a cheap and reasonable initial guess for the more expensive tightly coupled problem. This preliminary step saves hundreds of fully coupled optimization iterations or time-consuming manual structural sizing. The mass and torque values for the configuration we obtained serve the normalization values (Q_{x_i} and M_i in Section 5.1.1, and thrust F_{x_i}) for tightly coupled optimizations in the rest of this chapter.

Table 5.3: The tightly coupled optimization is more expensive but yields additional mass reduction.

	Uniform arbitrary	Loosely coupled baseline	Tightly coupled restart
Mass [kg]	89 399	35 774	35 060 (−2 %)
Time [CPU-h]	–	1 332	7 791

We now compare the results of the loosely coupled sizing from algorithm 1 to a subsequent tightly coupled optimization. Both optimizations minimize mass. The latter case uses structural thickness distribution from the loosely coupled approach as the initial design. This problem includes the structural thickness variables and a torque constraint to ensure that the aerodynamic performance does not degrade. The results of this study are presented in Table 5.3 and Figure 5.3

The initial structure with an arbitrary uniform thickness distribution is oversized for the inflow condition for which we are sizing the turbine. The loosely coupled optimization reduces the rotor mass from more than 89 t to less than 36 t. The computation cost of this optimization is 1,332 core hours on 160 Intel Xeon Platinum 8380 (Ice Lake) cores. This corresponds to less than 9 hours of computational wall time. Unless otherwise specified, the loosely coupled optimization result is the baseline case for all subsequent optimizations.

The tightly coupled optimization uses the full aerostructural model at each design iteration. The MDA and coupled derivatives capture more subtle aerostructural trade-offs and further reduce the mass by 2 % — roughly 700 kg for this layout. This additional mass reduction is only possible through coupled optimization because the loosely coupled approach could not further reduce the rotor mass.

Figure 5.3 shows where the optimizer reduces the blade thickness and mass. The view is perpendicular to the rotor plane. The loosely coupled optimum has the highest thickness over the spar caps away from the root and tip regions and on a large section aft of the front spar on the lower skin (high-pressure side). The tightly coupled optimization reduces this section's structural thickness while reinforcing the upper skin counterpart. The area closer to the root shows thickness reduction on both sides. Conversely, more material is added in front of the spar caps on both skins in the 50–80 % span range. This added material is due to the torque constraint. The optimizer has no direct control over the geometry in this design problem, so it modifies the torque output by altering the blade elastic displacement and the bend-twist coupling through changes in the local structural stiffness.

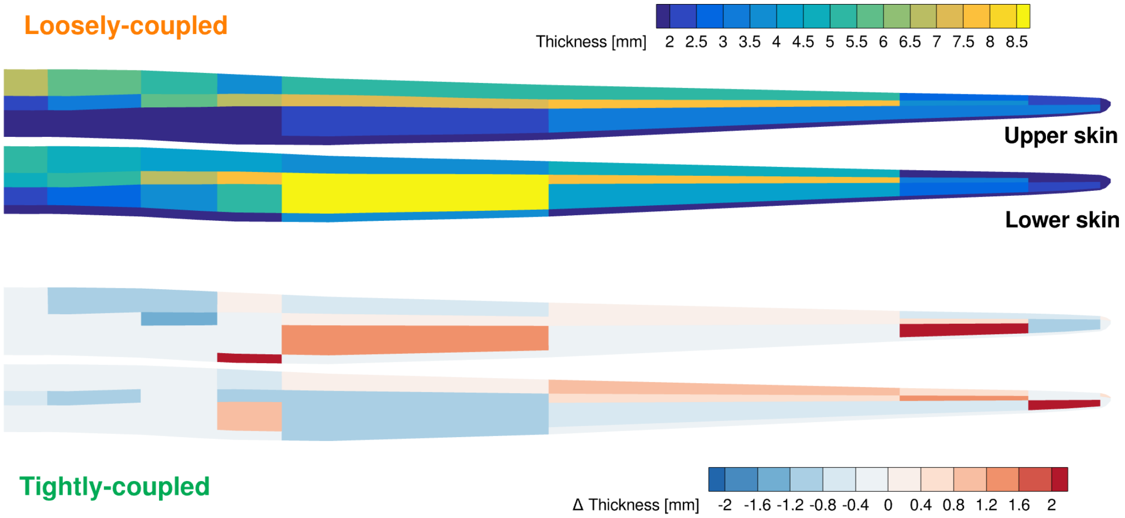

Figure 5.3: The optimizer adjusts the thickness to save 2 % more mass in the tightly coupled model compared to the loosely coupled model.

The specific layout modifications depend on load case, parametrization, and constraints. Mangano et al. [142] showed that the additional mass reductions achieved by the tightly coupled model are consistent over different torque and displacement constraint values and structural parametrizations. Recall that we are sizing the turbine for a below-rated inflow condition. In addition, we use a material with higher stiffness and comparable density properties compared to the fiberglass-based composites commonly used for industrial-scale wind turbines. Structural buckling is also not considered at this stage. These factors explain unrealistically low rotor structural thicknesses (1.6–8.5 mm). However, we expect the tightly coupled optimization benefits over the loosely coupled approach to remain similar regardless of the material used.

5.3 Impact of constraints on structural sizing

In Mangano et al. [142] we performed a range of optimizations including a tip displacement constraint. This constraint is used in Chapters 6 and 7, but the remaining sections of the current chapter do not enforce it on the aluminum-based model. In this early study we also revert to the original objective formulation from Garcia-Sanz [11] shown in Eq. 5.1.

The advantage of the coupled-system optimization approach and its sensitivity to nonlinear torque and displacement constraints is outlined in Figure 5.4 over a broader set of optimization problems. Δd_{tip} is defined as the increase (in percentage) of the tip displacement with respect to the baseline value (5.46 m for this case).

We explore the design space by varying the prescribed constraint value, ranging from 100% to 85% of initial torque values, and from 0% to 15% of tip displacement increase with respect to the initial layout.

There are a few key insights to unpack in Figure 5.4. In the first place, the results marked as $\Delta d_{\text{tip}} = 0\%$ highlight the advantage of using the coupled model rather than just the single-discipline approach used in Section 5.2. The rightmost point of this set refers to an optimization where we maintain the same displacement and torque as the initial configuration, and are nevertheless able to reduce the total mass by 4.1%. Note that the mass is expressed by the normalized M_2 metric from Garcia-Sanz [11] as the inverse of the mass. Thus, an increase in M_2 indicates a mass reduction. As hinted earlier in this section, when the coupled model is used the optimizer can take full advantage of the reciprocal effects of the aerodynamic loads on the stress distribution and, vice-versa, of the structural displacements on the aerodynamic loads.

Secondly, the availability of the torque constraint allows us to perform structural sizing while practically enforcing a prescribed torque value (normalized M_1 metric) and obtain the lightest blade possible with the given parametrization. In the range we are investigating, we observe a linear relationship between the rotor torque and the available mass decrease. The mass decreases

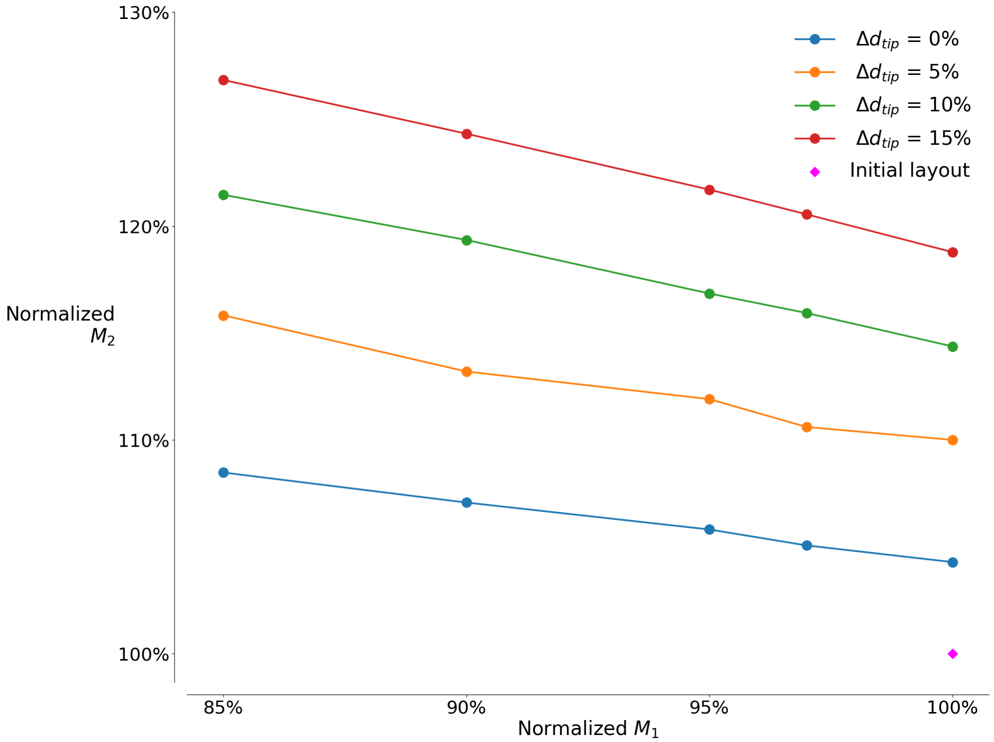

Figure 5.4: Design space exploration for optimally-sized rotors using M_1 (torque) and M_2 (inverse of the mass) metrics. Different Pareto fronts refer to different optimization sets with different maximum tip displacement constraint.

amounts up to 7.8% for a 15% reduction of the initial torque. This trend appears to be loosely related to the prescribed displacement constraint. The final mass difference between the 100% and the 85% torque cases amounts to 3.7% for $\Delta d_{\text{tip}} = 0\%$ and 5.3% for $\Delta d_{\text{tip}} = 15\%$, showing that the optimizer can also take advantage of the coupled effects of simultaneously relaxing the two constraints.

Looking at the general trends, the optimal designs obtained with this formulation are more sensitive to the displacement constraint than to the overall torque. This can be explained with the same considerations reported in Section 5.2. For the prescribed flapwise tip displacement, the optimizer tends to reinforce the outer section of the blades by thickening the panels by up to 14 mm.

5.4 Tightly coupled aerostructural optimization

We now perform tightly coupled aerostructural optimization, including geometric design variables. We refer to this optimization approach as ”aeroelastic tailoring”. The optimizer concurrently alters the blade shape and its structural thickness distribution. This way, the relative position of the center of pressure and elastic axis is adjusted so that the static aerostructural response of the blade at the given operating condition is tuned to maximize the performance of the rotor. The objective function is the weighted average of torque and rotor mass defined by Eq. (5.1). We use a weight of $\omega = 0.7$, emphasizing torque over mass. The effects of varying ω are discussed in Section 5.5.

Adding geometric variables yields higher performance improvements but increases the problem complexity and slows the optimization. As mentioned in Section 5.1, we set tight convergence tolerances for SNOPT. The optimizer consistently designs configurations that satisfy the linear and nonlinear constraints within thousandths of one percent.

Figure 5.5: The Tw optimization converges to a feasible design in fewer than 300 iterations.

However, as shown in Figure 5.5 for the case using structural and twist variables (Tw), it takes almost 300 design iterations to reduce the optimality by more than two orders of magnitude. Although the optimizer could further improve the design from a numerical standpoint, the mass and torque are practically constant beyond 200 iterations. The variations beyond this point are likely smaller than the modeling error. The overall cost of this specific optimization is $25,000$ core hours. Future work on parametrization and tighter analysis convergence could address these robustness and efficiency issues, but the effect on the rotor design would be negligible from a practical design standpoint.

We now discuss the optimized layouts. Figure 5.6 compares the baseline rotor and the Tw optimum (Table 5.2). The semi-transparent plots refer to the undeflected blade shape on which the geometry deformations are applied. The other lines refer to the same blade deflected under the design conditions. The undeflected and deflected shapes are shown for both the baseline and the optimized blades. The pitch sign convention is such that the zero-pitch position is at an angle of $90°$ relative to the incoming airflow, with the leading edge pointing upwards. The twist angle increases as the sections align with the freestream direction. We extract these distributions directly from the aerodynamic meshes of the undeflected and deflected solutions.

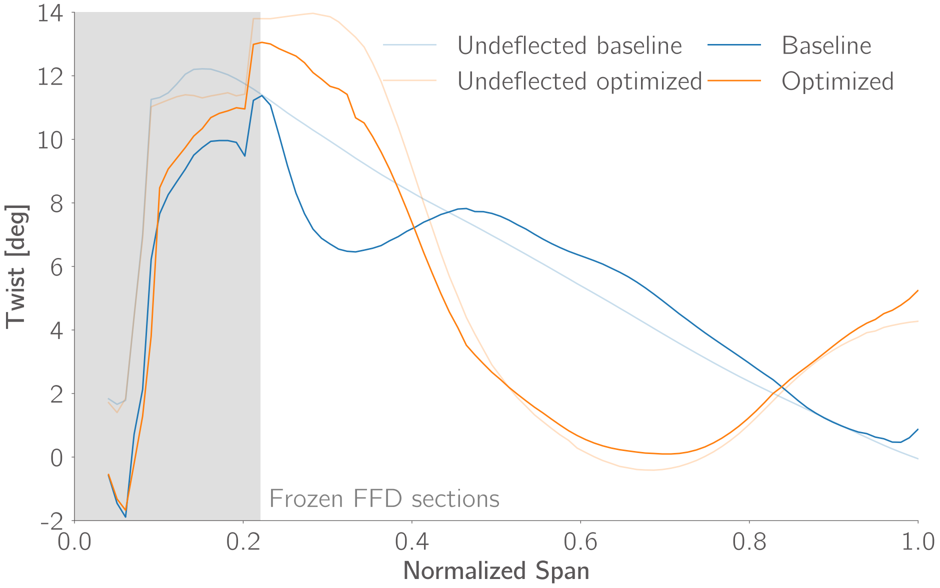

Figure 5.6: For the Tw optimum, the optimizer unloads the tip and reduces the twist between 45 % and 85 % of the span.

The baseline undeflected DTU 10 MW rotor outer mold line has a linear twist distribution [15]. The blade flexibility exacerbates the load-induced twist in the baseline design. On the one hand, the twist in the blade outboard is higher than the rigid counterpart, contributing to local load alleviation. On the other hand, the twist is reduced over the inboard half of the blade. These trends might be explained by the relative position of the elastic center and center of pressure at different spanwise locations as the blade tapers from a circular section at the root to an FFA-W3 airfoil.

The optimizer decreases the twist by up to $5.7°$ between 45 % and 85 % of the span, increasing the aerodynamic loads in this section. This increase in the local angle of attack is consistent with the pitch optimization studies performed by Madsen et al. [15, Sec. 6.1]. Conversely, the tip unloads, and the local twist increases by up to $4°$. The optimizer adds a smaller $2°$ twist in the 20–40 % span range, where the aerodynamic loads tend to realign the local sections to the freestream. Because the optimized layout has minor structural changes (discussed later in Figure 5.11), the twist deformation trend on the deflected blade is consistent with that of the baseline rotor.

ADflow uses the SA turbulence model and most likely overpredicts the stall angle of attack, as is the case with many RANS solvers [146]. The optimizer could exploit this weakness by making the airfoils operate at their maximum lift. However, the pitch-only aerodynamic optimization by Madsen et al. [15] identified an optimal pitch for the same configuration at more than 7° for the L1 mesh, higher than the local values we observe in this section. That set of results suggests that local sections on our optimized blade have an even larger margin from $\alpha_{C_{l_{\max}}}$. To further dissipate concerns about excessive pitch angles of the optimized designs, Figure 5.7 reports the estimated angle of attack of local airfoil sections over the blade.

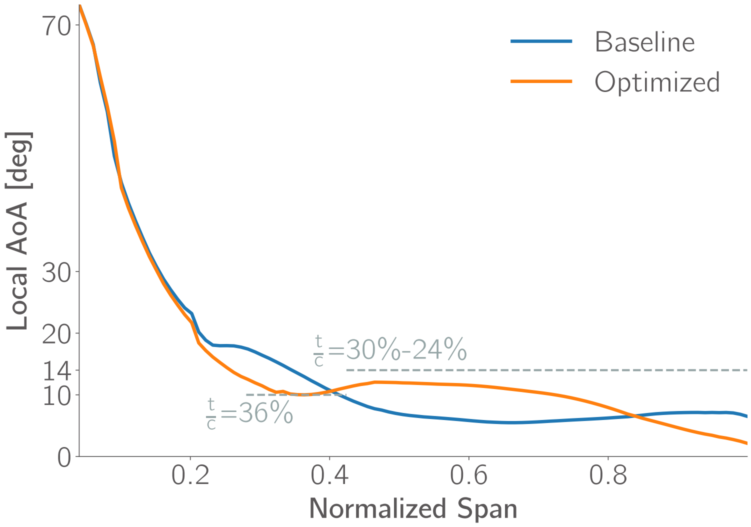

Figure 5.7: The local angle of attack for the layouts in Figure 5.6 does not exceed the airfoil $C_{l_{\max}}$ angles reported by Bak et al. [24] (gray dashed lines) in the reference DTU 10 MW report.

We sum together the pitch distribution from Figure 5.6 with the local inflow angle over the blade, obtained as the arctangent of the ratio of the local rotation speed and the inflow velocity of 8 $\mathrm{m/s}$ (see Hau [5, Fig. 4.6] for reference). This semi-analytical estimate is conservative, since we do not consider induction: in both numerical simulations and practical applications, the horizontal wind velocity at the rotor plane is lower than the freestream, reducing the local angles of attack on the blades. Even with this angle overestimation, the optimized design remains at least 2° below the $Cl_{\max}$ angles of the airfoil family (shown as gray dashed lines in Figure 5.7) used on the DTU 10 MW benchmark [24]. Figure 5.7 further highlights how the optimizer reduces the local pitch in the inboard section that is typically more prone to stall, resulting in a larger stall margin [15, Fig. 24] (see Figure 5.6 and Figure 5.8).

Fully turbulent RANS simulations tend to delay flow separation onset prediction, which might lead to higher lift values for a given angle of attack. This suggests that ADflow overpredicts the lift values in this region of the design space. Multipoint optimization and the inclusion of transition model in the aerodynamic solver would help address these limitations. Nevertheless, as further shown later in this section, we can confidently state that no flow separation occurs over the optimized blade at the design pitch angle.

The optimizer identifies the best trade-off according to the function weights selected for the objective function. A lower moment arm for the aerodynamic loads on the blade helps reduce the mass because of the lower root bending moment. However, it also potentially reduces the in-plane moment that drives the torque generation. The blade twist is adjusted so that the overall loads and torque are increased. The resulting design's torque is 14.3 % higher at the selected inflow condition. The structure thickens because of the higher loads, resulting in a mass increase of less than 1 %.

Table 5.4: More design freedom yields similar torque improvements with increasing mass reduction compared to the baseline.

Case	Tw	TwSc	TwChTk
Torque	$+14.3$ %	$+14.3$ %	$+14.6$ %
Mass	$+0.7$ %	-6.3 %	-9.3 %

Next, we investigate how increasing design freedom benefits the optimized rotor performance. In Table 5.4, we compare the objective metrics for three optimizations with increasing levels of design freedom (see acronyms in Table 5.2). With the selected objective function and thrust constraint, the increase in torque is comparable over different optimizations. However, the optimizer utilizes the increased design freedom to reduce the mass by more than 9 % for the TwChTk optimum.

The twist, chord, and relative airfoil thickness distribution are compared in Figure 5.8. The twist distribution on the deflected blades matches within 1–2° for the two optimizations that include planform design variables. The trend is comparable to the Tw case in the mid-span section, but the twist increase is lower by up to 5° in the inboard section and up to 2 °the outer 20 % of the blade. More specifically, the maximum decrease in pitch is 5.7° for Tw, 5.5° for TwSc, and 4.5° for TwChTk. Even though we verified that the airfoils do not operate at their $C_{l_{\max}}$, the optimizer is likely taking advantage of the absence of stall margin to extract more torque by approaching the sectional $C_{l_{\max}}$ closer than stall margins would allow. A multipoint optimization formulation with a higher-load design point would likely lead to a more practical and less aggressive pitch decrease over the blade, something observed later in Chapter 7.

When the optimizer controls the blade planform, chord variations over the span compensate for these differences in the twist. Both TwSc and TwChTk cases reduce the chord (and relative

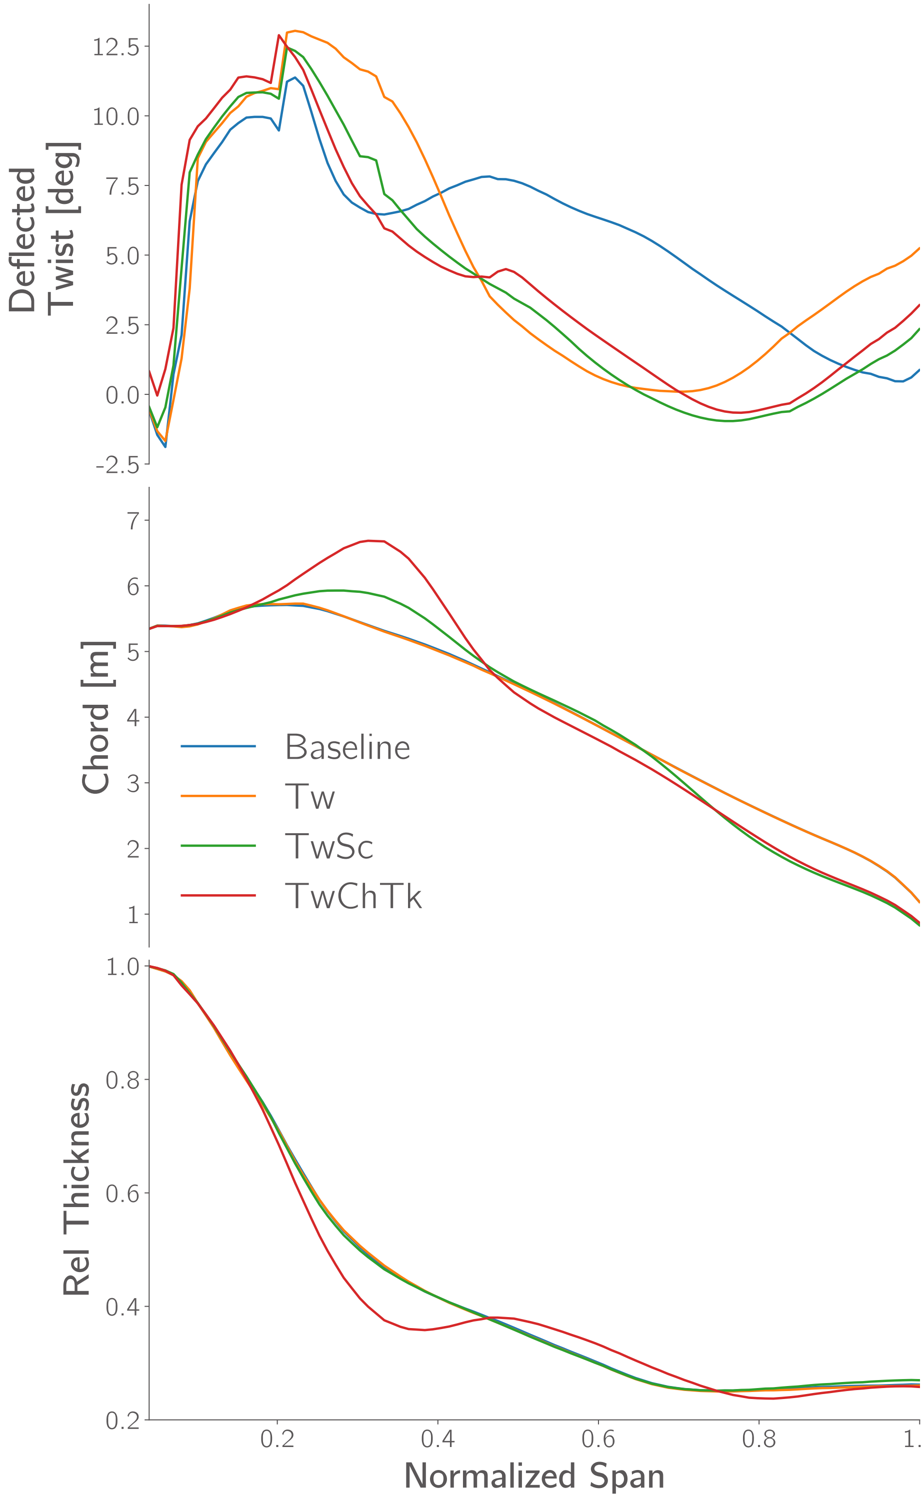

Figure 5.8: The additional design freedom from planform variables decreases the twist and chord at the tip compared to the Tw counterpart.

thickness for TwSc) by 30 % close to the tip. Between 22 % and 45 % of the span, the TwSc case increases the chord by 15 % while TwChTk case increases by 40 %, reaching the design variable upper bound. The most complex case (TwChTk) displays the more aggressive chord increase with an up-to-10 % decrease in relative thickness, stretching the airfoil in the chordwise direction. In contrast, the airfoil section gets shorter and thicker than the Tw and TwSc cases between 45 % and 65 % span, where the twist decrease is the highest.

The two optimizations with planform design variables are subject to the same bounds discussed in Section 5.1.3. The TwSc case does not increase the maximum chord to the upper bound because the corresponding thickness increase incurs penalties. MACH-based analysis shows that a large chord with a smaller relative thickness, as obtained in the TwChTk case, is more beneficial for this design point and problem formulation. Our blade-resolved tightly coupled model captures these trade-offs, demonstrating the advantages of high-fidelity MDO in the rotor design process.

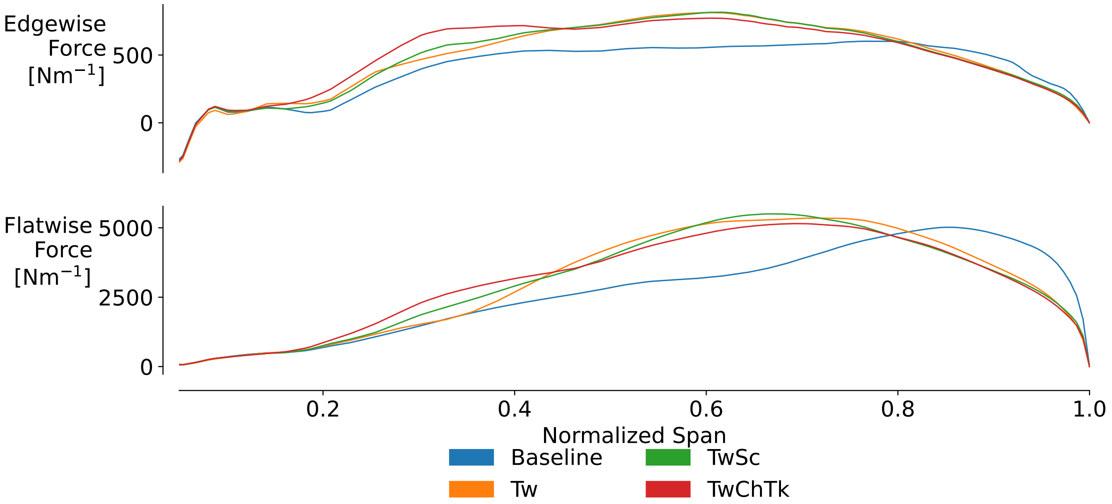

Figure 5.9: The optimization reshapes the spanwise load to move its peak inboard.

The effects of the structural and geometrical changes on the load distribution over the blade are shown in Figure 5.9. Both edgewise and flatwise force peaks shift inboard, reducing their moment arm. Forces in the span outermost 20 % are lower than the baseline, but higher loads in the 20–80 % span region compensate for the reduction of edgewise torque. The flatwise force peak is higher than the baseline but shifted inboard by 20 % of the span. The corresponding integrated load does not exceed a 14 % increase compared to the baseline because of the thrust constraint.

More design freedom enables more aggressive load shifts. The TwChTk case shows a marked load increase between 20 % and 45 % span due to the planform changes shown in Figure 5.8. The load shift is more significant for flatwise forces, ultimately driving the blade sizing. This explains how the optimizer achieves the highest mass reduction in this case. These trends could be specific to the selected inflow condition and rely on ADflow correctly capturing the aerodynamic loads at low speed near the root. However, these trends are consistent with engineering intuition.

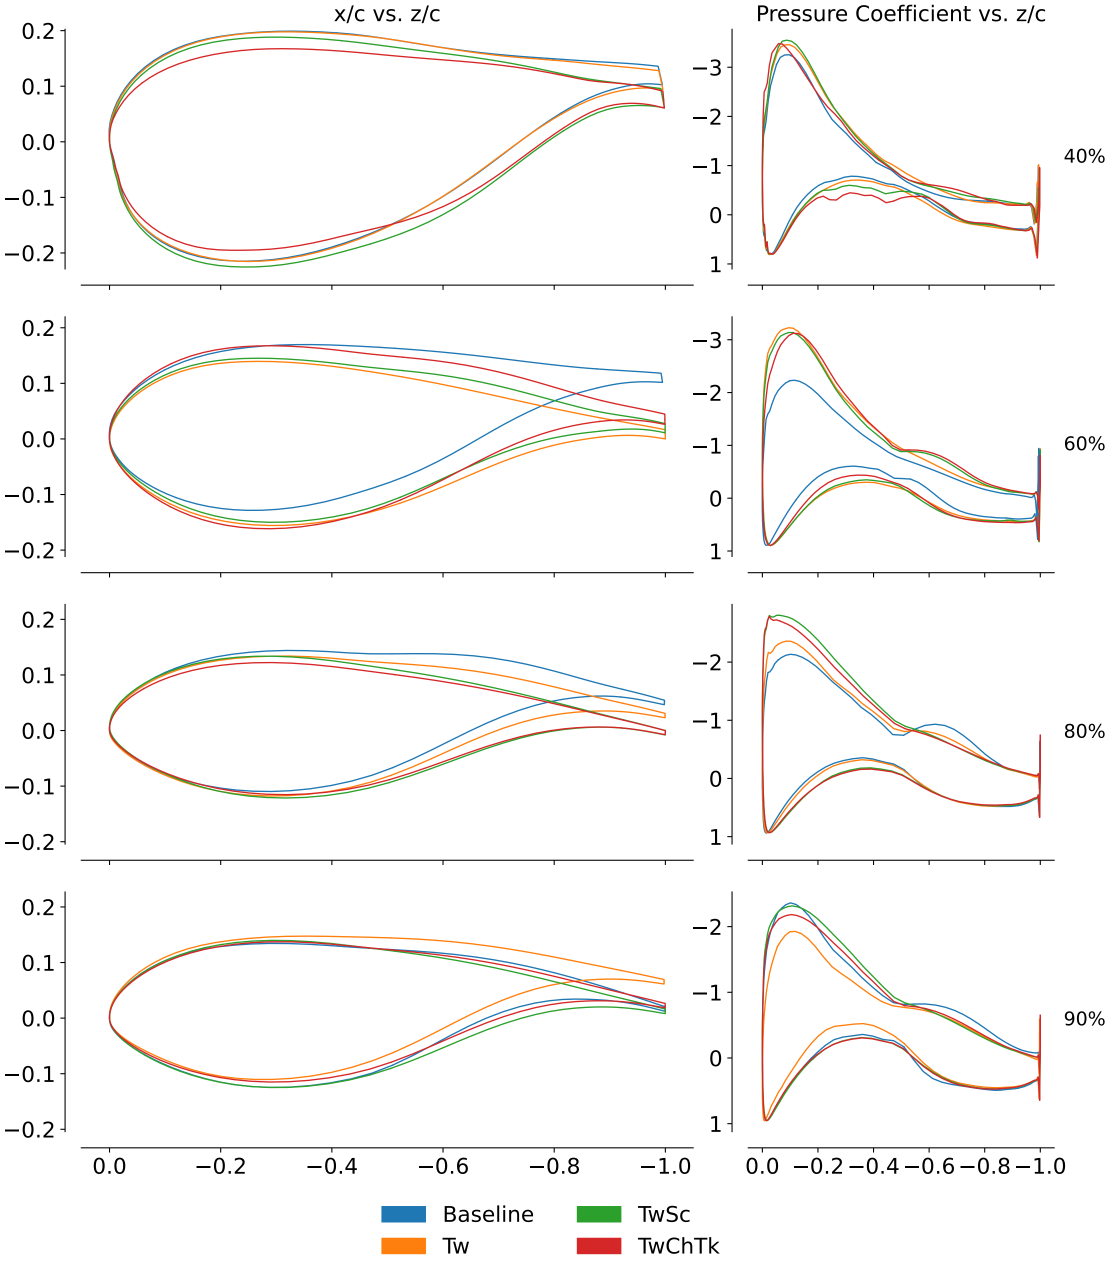

Figure 5.10: Cross-section pressure distributions show consistent optimization trends; when planform variables are optimized, the pressure peak at the leading edge increases. Section locations on the platform are highlighted in Figure 5.11.

Thanks to MACH's high-fidelity analysis capabilities, we can investigate how the optimizer alters the local airfoil shape and pressure distribution, reported in Figure 5.10. The spanwise location of these sections is also highlighted in Figure 5.11. At 40 % span, the C_p over the normalized chord is comparable in suction peak location and magnitude between different optimized designs. The TwSc and TwChTk optimizations use planform variables to obtain this distribution with a lower twist than the Tw case, as shown previously in Figures 5.8 and 5.9. The TwChTk airfoil has a lower thickness-to-chord ratio and generates a higher pressure difference around the mid-chord.

At 60 % of the span, the optimized layouts show a 50 % increase in the suction peak compared to the baseline. Optimal airfoils show similar C_p trends despite the differences in local twist on the deflected blades. Unlike the previous case, the thickness-to-chord ratio for the TwChTk is higher, as highlighted earlier in Figure 5.8. At 80 % of the span, the planform-optimized cases still show a higher suction peak than the baseline and the Tw case. However, closer to the tip, the C_p distributions for TwSc and TwChTk more closely resemble the baseline pressure distribution, and the force reduction shown in Figure 5.9 comes from the chord and thickness reduction. The Tw case has the same baseline planform and must reduce local lift through a more aggressive twist increase, which lowers the C_p peak.

The two outermost baseline sections display a bump in the shape in the aft part of the airfoil, which results in a corresponding change in C_p. This bump appears due to the airfoil elastic deformation under the current loading condition. This is because of the thin baseline structural layout (obtained through the loosely coupled approach), which reduces the chordwise stiffness. This issue can be addressed by enforcing stricter structural constraints.

Nevertheless, this phenomenon is interesting for two reasons. First, this level of detail is captured only through high-fidelity coupled analysis. Combined BEMT and beam codes cannot model this aerostructural interaction because the airfoil shape is fixed during analysis, and the airfoil drag polars do not account for chordwise deflections — only empirical corrections for 3D effects are available.

Second, the optimizer modifies this bump in the coupled aeroelastic solution. The optimizer removes the bump because it is detrimental to the aerodynamics (see the 80 % section in Figure 5.10). Conversely, the optimizer adds a small bump at the 60 % section, where the local load increases are the highest. However, the higher local torque contribution compensates for the effects of this local deformation. This detailed tailoring is possible because our high-fidelity optimization approach accurately captures the two-way-coupled aerostructural elastic deformation.

To understand how the optimizer reduces the rotor weight in the different optimization cases, we compare the structural thickness distribution between the baseline layout and the three optimized configurations discussed in this section. In Figure 5.11, we observe that the Tw optimization only made minor adjustments over the blade, for an increase of 242 kg over the whole rotor.

The planform variables enable a structural thinning over the thick lower skin panel, with reductions up to 3 mm for the TwChTk optimization. The thickness distribution trends are consistent between the TwSc and TwChTk over the rest of the blade. The thinning of this large section drives

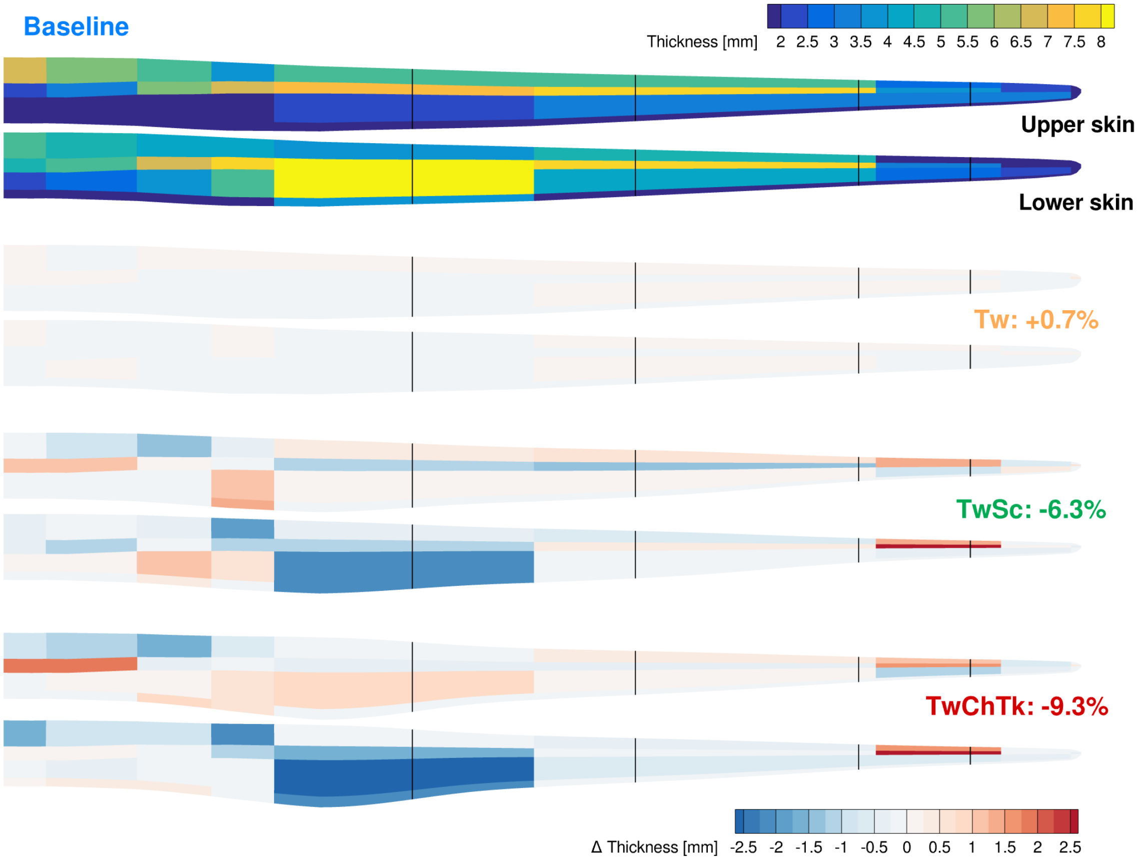

Figure 5.11: Increased geometric design freedom decreases the mass, particularly on the thick lower surface panel.

the mass reduction of the overall blade. This thinning might raise a concern about buckling, which is not considered in our analysis. However, this specific section would very likely not be subject to buckling because it is loaded in tension at the load condition we are using for analysis and optimization.

The main structural reinforcements occur at the leading edge section between 80 % and 90 % of the span, most likely to improve the bend-twist coupling, and on the upper skin down to 30 % of the span. The upper spar cap is also reinforced in both cases. The use of isotropic aluminum properties and the absence of life-cycle sizing considerations, as mentioned in Section 4.1.1, result in designs that are optimal for the given problem and model but might not be practical. However, the optimizer identifies the design trends and the structural components most critically stressed at the current design condition. This is partially addressed in the following Chapters 6 and 7.

The changes on the cross-section made by the optimizer for the TwChTk case, in particular, affect both the local airfoil aerodynamics and the sectional bending stiffness. Because of the coupling of the aerodynamics and the structural response, it is hard to dissect which changes benefit the torque output and the mass reduction. The cases presented in this section exemplify the potential of our tightly coupled high-fidelity approach to capture geometric and structural trade-offs. This approach can be extended to more elaborate design scenarios and rotor models.

Finally, in Figure 5.12 we illustrate the C_p and structural stress distribution on the upper skin of the TwChTk optimized rotor detailed earlier in this section. From an aerodynamic standpoint, the figure at the top extends the insights highlighted in Figure 5.10 to the full blade-resolved model. The surface streamlines dispel concerns of flow separation occurring over the blade. However, different flow solvers could more accurately capture the flow characteristics and the impact of laminar-turbulent transition in the inboard section, as shown by Madsen et al. [15, Figs. 23–24]. At the same time, the absence of cross-flow at the given design point on the optimized design suggests that the advantages of this high-fidelity approach would be better exploited through a multipoint formulation, including more challenging flow conditions.

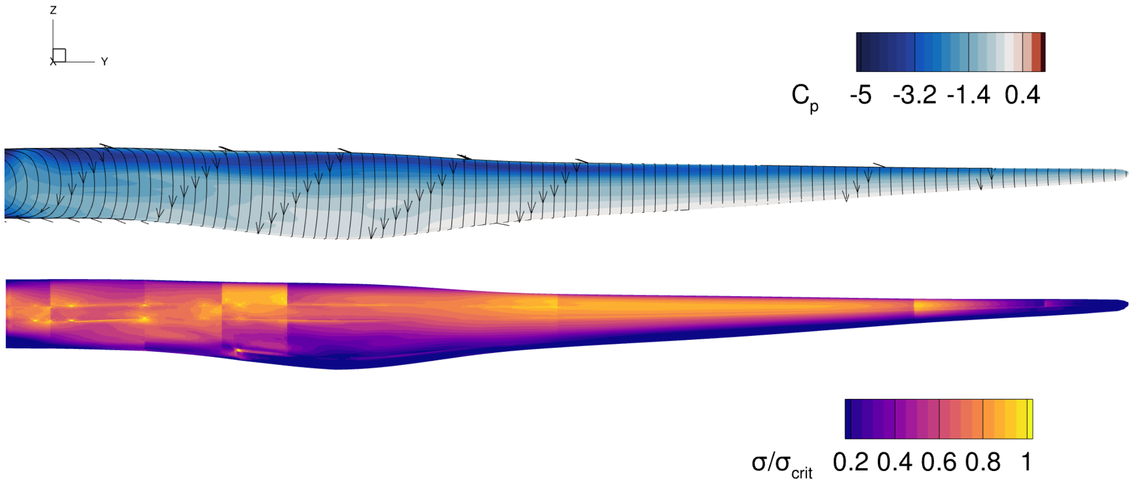

Figure 5.12: Pressure coefficient, velocity streamlines, and stress distribution over the upper surface of the TwChTk optimized design in Figure 5.11

On the structural side, we observe stress concentrations at the panel junctions over the skin and the spar-skin connection. This suggests that a more refined parametrization could smooth out these differences in local stress through a more gradual thickness variation over the span. Larger mass reductions could be obtained, as the optimized results satisfy aggregated stress constraints (Section 5.1.3) over the entire blade. In particular, the large panel corresponding to the maximum chord location appears oversized in the trailing edge section, as the local stress in that area is less than 50% of the critical load.

5.5 Design space exploration

This section explores how the objective formulation affects the optimized design's torque output and total rotor mass. The results are summarized as a Pareto front over the mass-torque plane in Figure 5.13, where every point corresponds to a different optimization case, color-coded to match the design variable sets discussed in Section 5.4. The designs closer to the lower right corner of the figure are lighter and more aerodynamically efficient, outperforming the other optimized layouts.

The two cases in light blue (closest to the baseline in the top left corner) use a smaller subset of design variables and consider either only torque ($\omega = 1$) or mass ($\omega = 0$) in the objective function. The "Mass min" case is the tightly coupled mass minimization problem discussed in Section 5.2. "Torque max" refers to a pure torque maximization problem ($\omega = 1$ in Eq. 5.1) using twist design variables without changing the internal structural layout. The baseline rotor structure can tolerate only small load increases; thus, the optimizer increases the torque by no more than 4 %.

The case in the upper right corner of the figure (part of the orange Tw set) refers to another torque maximization case where the optimizer can change the twist and structural thickness distribution without mass penalty in the objective. The optimizer increases the torque by 18 % without violating stress or thrust constraints. However, the total mass almost doubles because the structural thickness increases to a point where the aerodynamic efficiency no longer improves with higher blade stiffness. Because this is not a relevant design case, we did not run TwSc and TwChTk design variable sets with the same torque-only objective. These results reinforce the need to explore mass and torque trade-offs through objective function formulation.

The remaining optimizations use both structural and geometric design variables. They also include the rotor mass in the objective. The cases ranging from 10 to 15 % torque increase (see Figure 5.13) minimize Eq. 5.1 using ω = 0.5, ω = 0.6, and ω = 0.7 respectively.

The rightmost cases in this set with ω = 0.7, thus using a 0.7–0.3 weighting for torque and mass, were discussed in Section 5.4. The mass reduction discussed in the previous section is consistent between different design variable sets, and the torque outputs are within 1–2 % for the same objective weighting. Including planform variables reduces the mass by 6–8 % compared to their Tw counterparts. Decoupling chord and thickness yields an additional 2–3 % reduction. An objective with equally weighted torque and mass yields designs with a torque output between 9.4 % and 10.5 % higher than the baseline rotor and mass decreases of 5.4 % to 16.4 %, depending on the parametrization.

The relation between objective weights and final layout metrics becomes highly nonlinear when the objective emphasizes mass. With this formulation, the optimizer tends to aggressively reduce the aerodynamic loads to achieve a lower mass. A case with $\omega = 0.4$ (not shown in Figure 5.13) had a torque output that is 10 % lower than the baseline. For this reason, optimizations exploring

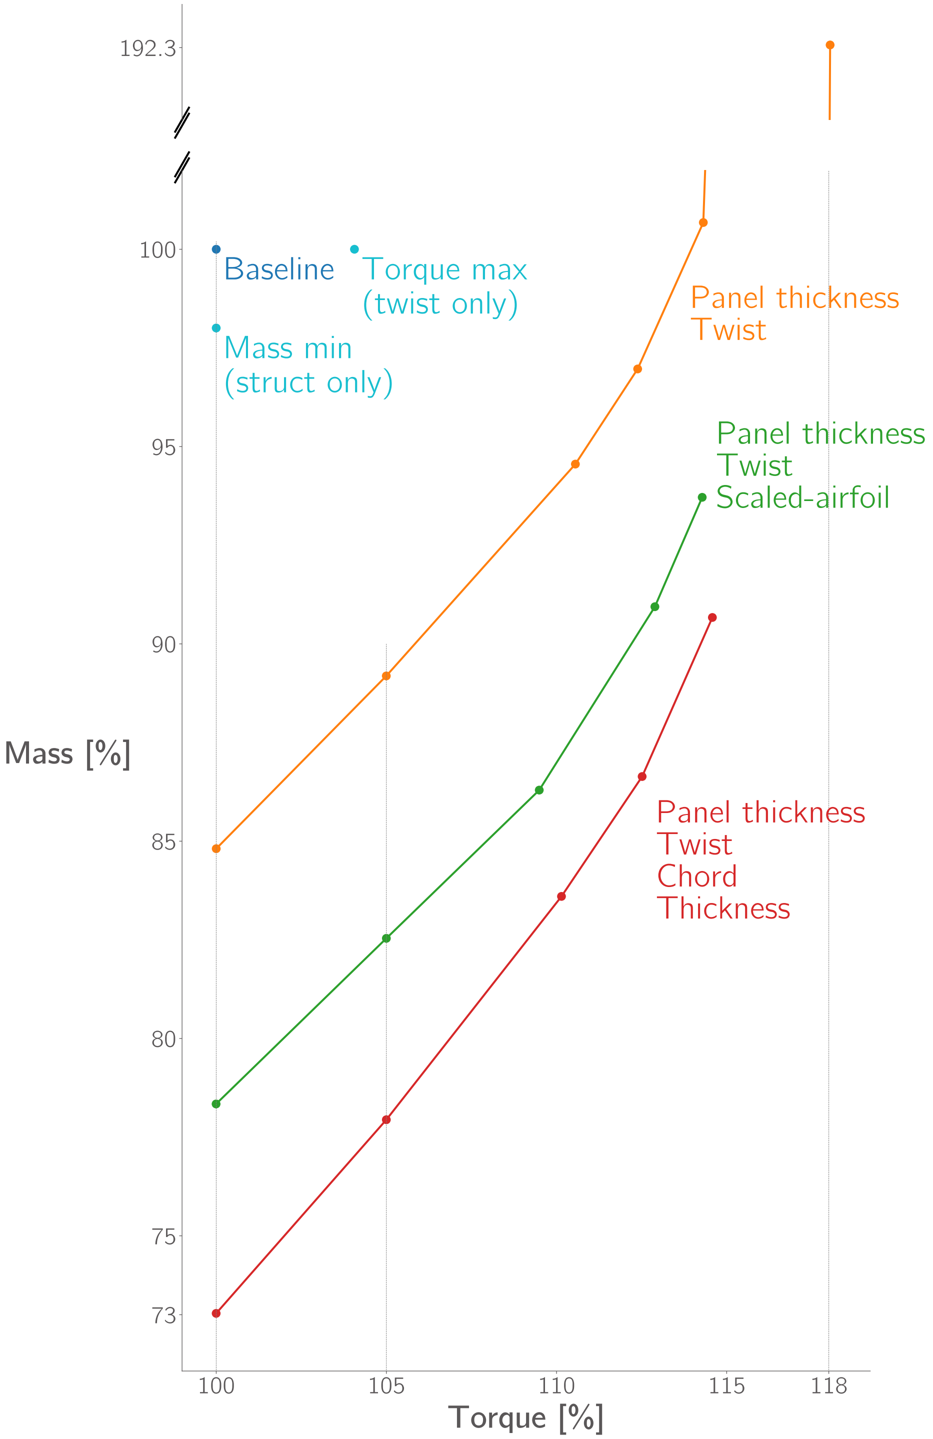

Figure 5.13: Mass and torque trends for different design variable sets are consistent over the design space.

this design space area should prescribe the torque output rather than include it in the objective function.

The two leftmost sets of optimizations in Figure 5.13 consist of mass minimization problems ($\omega = 0$ in Eq. (5.1)) with a torque constraint that is either the same as the baseline or 5 % higher. When the torque is increased by 5 %, rotor mass increases by about 3 % relative to the corresponding cases with baseline torque output. The load alleviation induced by the geometric design variables reduced the mass by 15 % when only changing the twist distribution - a 13 % additional reduction relative to the design of Section 5.2.

When the optimizer can scale the airfoil sections (TwSc), the mass decreases by an additional 7 %. In the highest design freedom scenario with decoupled chord and relative thickness variables (TwChTk), the optimizer achieves a 27 % lighter structure than the baseline while maintaining the same torque output.

As mentioned in Section 5.1, several factors justify the large changes in mass. The suboptimal baseline blade layout from Madsen et al. [15] and the loosely coupled optimization approach for initial structural sizing discussed in Section 5.2 allowed large optimization improvements. The inboard load shift illustrated in Figure 5.9 is consistent over the optimal layouts. The different spanwise load distribution and the consequent reduction in bending moment drives the mass minimization.

Figure 5.14 shows how the problem formulation affects the optimal twist distribution more in detail. Three optimizations with the Tw formulation are compared

1. M3T7: the case with $\omega = 0.7$ discussed in Figure 5.6;
2. T1: the case with $\omega = 1$ (no mass in the objective) the top right case in Figure 5.13 discussed earlier in this section;
3. T1-U: a new optimization with with $\omega = 1$ but a fixed, uniform thickness distribution - i.e. the baseline structural layout is not defined through the loosely coupled approach from Section 5.2.

These three cases have an increasing mass and blade stiffness, which ultimately enables the optimizer to achieve a higher torque increase on T1 and T1-U compared to the M3T7 design. When mass is not in the objective (cases T1 and T1-U), the optimized layouts do not show the previously observed inboard load shift. The highest local pitch (which corresponds to the lowest twist value in the figure) is located around 85% of the span This twist distribution more closely resembles the optimization results from Madsen et al. [15, Fig. 10],where the highest pitch is below zero degree and located between 90% and 95% of the span.

As discussed for Figure 5.13, the different designs shown in Figure 5.14 are tailored for the specific combination of operating conditions and objective formulation. Higher design freedom

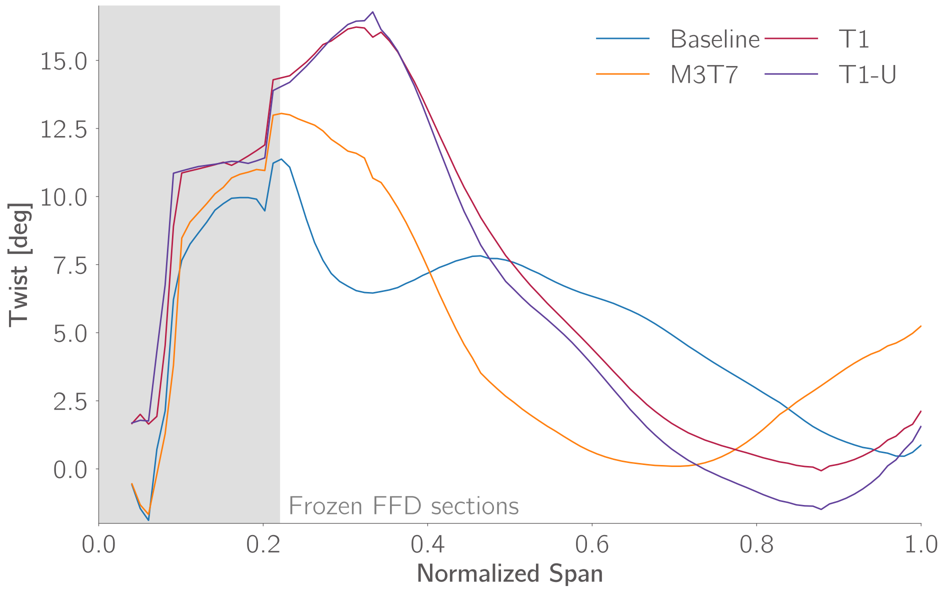

Figure 5.14: The optimal twist distribution on heavier, stiffer blades is comparable to aerodynamic shape optimization results from Madsen et al. [15].

enables more effective tailoring for the specific optimization problem. Our optimization framework can explore more objective function combinations and rotor parametrizations. Such exploration would be impossible with gradient-free approaches because the computational cost would be intractable.

5.6 Optimal designs verification

We re-analyze selected optimized designs presented in Sections 5.4 and 5.5 to verify that the gains obtained with the coarse CFD mesh are observed on a finer CFD mesh and a BEMT solver. On the one hand, we use the more refined L0 aerodynamic mesh discussed in Chapter 2 to confirm the aerodynamic performance improvements for a more accurate model. The geometry and structural layout used for these aerostructural analyses are identical to the output of the optimizations discussed above.

On the other hand, we use AeroDyn to assess whether the performance improvements for the optimized design are consistent with a conventional analysis code. Because we use a fixed set of airfoil polars taken from the baseline model, we choose to only compare Tw and TwSc optimization studies that do not directly alter the airfoil relative thickness. As shown in Figure 5.10, the CFD-CSM aerostructural analysis captures chordwise structural deflections, while BEMT codes assume rigid chordwise sections. The twist and planform scaling distributions shown in Figure 5.8 are provided as inputs to AeroDyn, together with the original, publicly-available airfoil polars computed with EllipSys2D. Although the BEMT model is run in rigid mode, the reference axis of the blade is set to be consistent with the blade deflection resulting from the CFD-CSM analysis. Therefore, the twist distribution accounts for the bend-twist coupling over the deflected blade.

Ensuring the consistency of the CSM and beam structural models is beyond the scope of this work and has been partially addressed by Caprace et al. [87]. The main goal of the comparison with BEMT is to ensure that the optimized blade layouts show consistent aerodynamic behavior when analyzed with a conventional design tool. Table 5.5 and Table 5.6 list the values of torque and thrust increases. All increases are relative to the respective baseline model and mesh in MACH or AeroDyn.

Table 5.5: Optimum for the Tw case ($\omega = 0.7$) analyzed with L0 mesh and AeroDyn.

	Torque	Thrust
L1	$+14.3$ %	$+14.0$ %
L0	$+10.1$ %	$+14.1$ %
AeroDyn	$+6.5$ %	$+12.5$ %

For the Tw case in Table 5.5, the L0 mesh predicts a thrust increase within 0.1% of the L1 analysis. The difference in torque is larger, but there is a significant relative increase when evaluating with the L0 mesh (10.1%). The loads predicted by AeroDyn are lower, with a torque increase of 6.5% and a thrust increase above 12%. Different fidelities show that the optimized design has consistently higher loads than the baseline, despite the unconventional twist distribution. The greatest improvements are obtained with the fidelity used in the optimization loop.

Different models identify different optimized designs or, conversely, predict different improvements compared to the model used for the optimization itself (L1 in this case). An optimization using the L0 mesh could lead to higher relative torque increase while maintaining the thrust within the prescribed constraints. Madsen et al. [15] showed how L1 and L0 aerodynamic optimizations lead to similar trends, but the optimizer can further adapt an L1-optimum when a more accurate flow resolution is available.

The twist distribution shown in Figure 5.6 (Tw), also reported in Figure 5.14 (M3T7), leads to a lower torque increase for the BEMT model. This can be related to the high amplitude changes in the twist distribution (compared to the baseline) that the CFD code deems beneficial. The large variations in the twist angle result in the airfoils operating far away from the baseline case. This might exacerbate potential errors stemming from fixed airfoil polars (on the BEMT side) and overestimated lift coefficients (on the CFD-CSM side).

Similar conclusions can be drawn for the TwSc case in Table 5.6. When planform design variables are also considered, both the finer mesh and the BEMT code show closer agreement to

Table 5.6: Optimum for the TwSc case ($\omega = 0.7$) analyzed with L0 mesh and AeroDyn.

	Torque	Thrust
L1	+14.3 %	+14.0 %
L0	+11.8 %	+15.5 %
AeroDyn	+8.6 %	+16.4 %

the L1 optimum torque, suggesting that the larger discrepancy observed in the Tw optimization is related to the larger inboard twist variation. The L0 mesh torque is 1.7% higher than for the Tw case despite the L1 mesh predicting an almost identical torque increase in both cases. The thrust is more than 1% higher for L0 and 2% higher for AeroDyn, indicating that further optimization is needed in this case to guarantee that the constraint is not exceeded.

The BEMT code also shows a closer match, with a torque improvement of more than half of the L1 value and 3.2% less than the L0 prediction. As shown in Figure 5.8, the highest local pitch for this layout is located more outboard than for the Tw case. This reinforces the idea that AeroDyn and ADflow exhibit different sensitivities to twist variations in the outboard span section. Hence, the higher the twist excursions, the larger the errors.

The lower performance improvements obtained through the BEMT code are consistent with the comparison performed by Madsen et al. [15]. The authors compared a BEMT optimization with ADflow-based results. Despite the more aggressive twist and planform distribution of the BEMT optima, these optimized designs resulted in approximately 8% torque increase compared to the 12% predicted by ADflow on the high-fidelity optimum. This suggests that performing a BEMT optimization could lead to similar torque improvements while matching the thrust constraint.

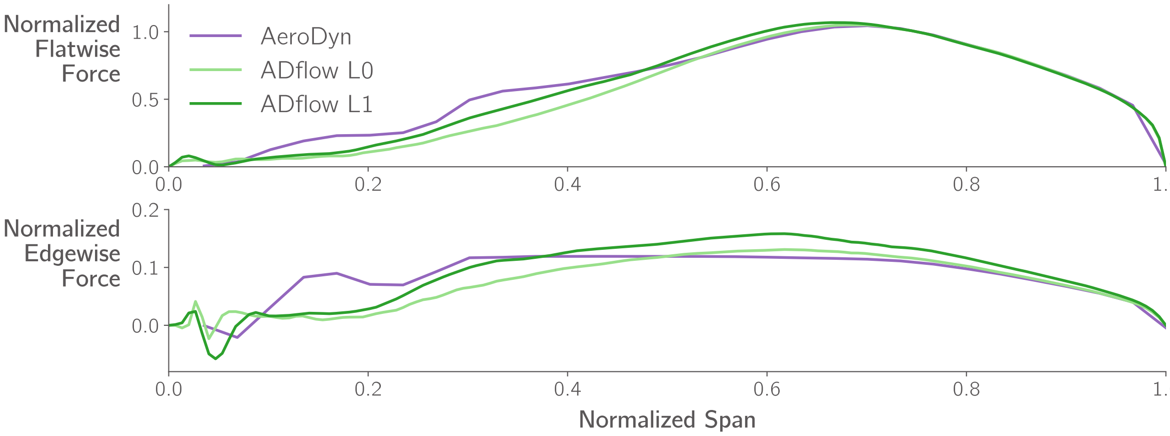

Figure 5.15: Spanwise loads distributions for the TwSc, predicted by the BEMT and the CFD solver. Values are normalized with respect to the flatwise force at 75% of the span.

Finally, in Figure 5.15 we compare the spanwise loads distribution for the AeroDyn and ADflow results on the optimized TwSc configuration summarized in Table 5.6. Since we are focusing on trends and not on the absolute values, the force values are normalized with respect to the flatwise force at 75% of the span. Along the inboard 40% of the blade, AeroDyn predicts higher loads than ADflow. This area is characterized by circular or thick airfoil sections and bluff bodies. The observed discrepancy might arise from different aerodynamic estimations between Ellypsys2D-based polars and our solver at high angles of attack.

In the outboard half, the flatwise force trends match closely between the BEMT solver and both ADflow mesh levels. ADflow estimates higher edgewise loads when using the L1 mesh. Conversely, the refined L0 mesh agrees more closely with the AeroDyn load distributions in the outboard 50% of the span. Between the span 20% and 50%, the L0 mesh has a lower normalized load than both other models. These trends are reflected by the integrated forces reported in Table 5.6.

A thorough tool cross-verification is beyond the scope of the optimization studies presented in this chapter. However, the consistency of distributed and integrated forces is promising for more complicated CFD-CSM optimization studies, and their integration with conventional design approaches.

5.7 Conclusions

In this chapter we presented the first set of optimization studies using the high-fidelity optimization framework and rotor model outlined in Chapters 2 and 3 respectively.

We use a gradient-based optimization strategy with a tightly coupled CFD-CSM solver in the loop to simultaneously minimize the mass and maximize the torque of a benchmark rotor at a set of operating conditions.

A preliminary sizing approach based on a conventional loosely coupled approach leverages the blade-resolved 3D structural model without using the more accurate but expensive tightly coupled model. The tightly coupled model reduces the mass by 2 % more than the loosely coupled approach for a fixed blade geometry. This improvement can only be achieved through the aerostructural solver and its coupled derivatives.

Changes in the twist, chord, and relative thickness distribution reduce the mass by up to 9 % and increase the torque by up to 14 %. The flatwise and edgewise loads are moved away from the tip to find the best compromise between aerodynamic efficiency and structural sizing. The airfoil layout and pressure distribution over the blade are also discussed in detail, contributing insights that cannot be obtained with conventional design methods.

Finally, the Pareto front analysis highlights optimal rotor mass and torque output trade-offs.

The benefits of more complex geometrical parametrization are consistent over the design space. Decoupling chord and relative thickness reduces the mass by 27 % without reducing the initial torque output. This further 5 % reduction relative to the scaled airfoil case demonstrates the benefit of modifying the airfoil shapes. The single-point problem formulation, the absence of life-cycle sizing considerations, and an overly conservative baseline design give the optimizer additional room for these significant mass reductions. Nevertheless, the re-analyses of the optimized design with a finer mesh and a lower-fidelity tool show consistent improvements and further support the case for high-fidelity optimization despite the limitations mentioned above.

Chapter 6

Tailoring Fiber Orientation for Composite Rotors

This chapter investigates wind turbine blade composite fiber tailoring using a tightly coupled CFD-CSM solver. We tailor the structural thickness, fiber angle, and planform twist distribution. We quantify the impact of structural parametrization and objective formulation on the final design. Our results show that fiber tailoring can reduce the mass of the rotor for the same or higher aerodynamic loads.

The optimizer exploits material anisotropic properties during the design process to exploit aerostructural tradeoffs over the blade-resolved rotor model. This addresses one of the main limitations of Chapter 5, where the blade planform is optimized with more geometrical design variables but relied on a simplified isotropic structural model.

The fiber angles in the blade structural panels are now optimized together with the panel thickness. The rotor mass and torque are again combined in the objective function to design lighter and more efficient wind turbines. We consider two design points in our optimizations. In this way we capture bend-twist coupling under different loads without increasing the problem complexity and computational cost.

As shown in Chapter 5, using a blade-resolved high-fidelity solver enables large-scale optimization studies while enabling more detailed component design than BEMT-based conventional analysis tools. We first quantify the impact of increasing structural design freedom on the optimized designs. The thickness and fiber angle patterns are dissected to understand the aerostructural tradeoffs exploited by the optimizer. Different objective formulations and design point selection are compared to understand the main design drivers and how they affect the blade structure. Finally, we investigate how different ply fraction layups affect the optimal panel thickness distribution and blade weight.

6.1 A preliminary loosely coupled optimization comparison

A a preliminary composite structure made of balsa core and fiberglass plies based on Bak et al. [24] with simplifications has been presented by Mangano et al. [142]. The example is based on the structural model used in Chapter 5 using the smeared stiffness element formulation discussed in Section 3.2.2. Although presented as a proof of concept rather than the more thorough composite-focus studies in this chapter and Chapter 7, we first demonstrate the tool capabilities and highlighted some key differences in the model behavior. This early study does not include ply angles as design variable, keeping the fiber aligned with the spanwise axis similarly to what done in Chapter 7.

We compare three different structural models that differ for design variables definition and material properties. Two of these models use isotropic shell elements with aluminum density and stiffness properties from Chapter 5. They differ in the way the design variables are defined, with one model having independent design variables for upper and lower skins, as summarized in Section 5.1.5, while the second has a different spanwise variables distribution and mirrored thicknesses between suction side and pressure side, a set up similar to what has been used in Chapter 6 and Chapter 7. The third model is a preliminary composite structure made of balsa core and fiberglass plies based on Bak et al. [24] with simplifications. This composite model uses the same "mirrored" variables distribution as the second isotropic structural model.

	Mass $[kg]$	
Model	$d_{\text{tip}} = \ 2\ m$	$d_{\text{tip}} = \ 3\ m$
Isotropic	$157\,013.2$	$97\,879.1$
Isotropic (mirrored DVs)	$160\,026.3$	$101\,244.5$
Composite	$196\,326.7$	$114\,865.2$

Table 6.1: Mass of the optimized rotor models, as obtained via the loosely-coupled strategy in Algorithm 1.

Using the loosely coupled optimization from Section 5.2, we compare the optimized mass of our three structural models in Tab. 6.1. To quantify the impact of the arbitrarily-defined displacement constraints, we run two sets of optimization with $d_{\text{tip}}2\ \mathrm{m}$ and d_{tip} = 3 m respectively. Given the selected inflow condition, these values represent rotor deflection limit at below-rated operating conditions and do not directly correspond to a rotor-tower clearance constraint, which needs to be evaluated under extreme load conditions. The impact of the stricter displacement constraint amounts to an average 63% increase in mass for the three models considered, with the composite model showing the highest sensitivity to the displacement constraint with a 70% weight increase between the two optimized layouts. The composite model is also the heaviest of these optimal configurations in both cases. This is likely because of the higher flexibility of the composite structure and the need for additional material to stiffen the blade and satisfy the displacement constraints. We also observe that enforcing the same thickness on upper and lower skins adds a weight penalty, which amounts to more than 3 t when comparing the two “aluminum” configurations.

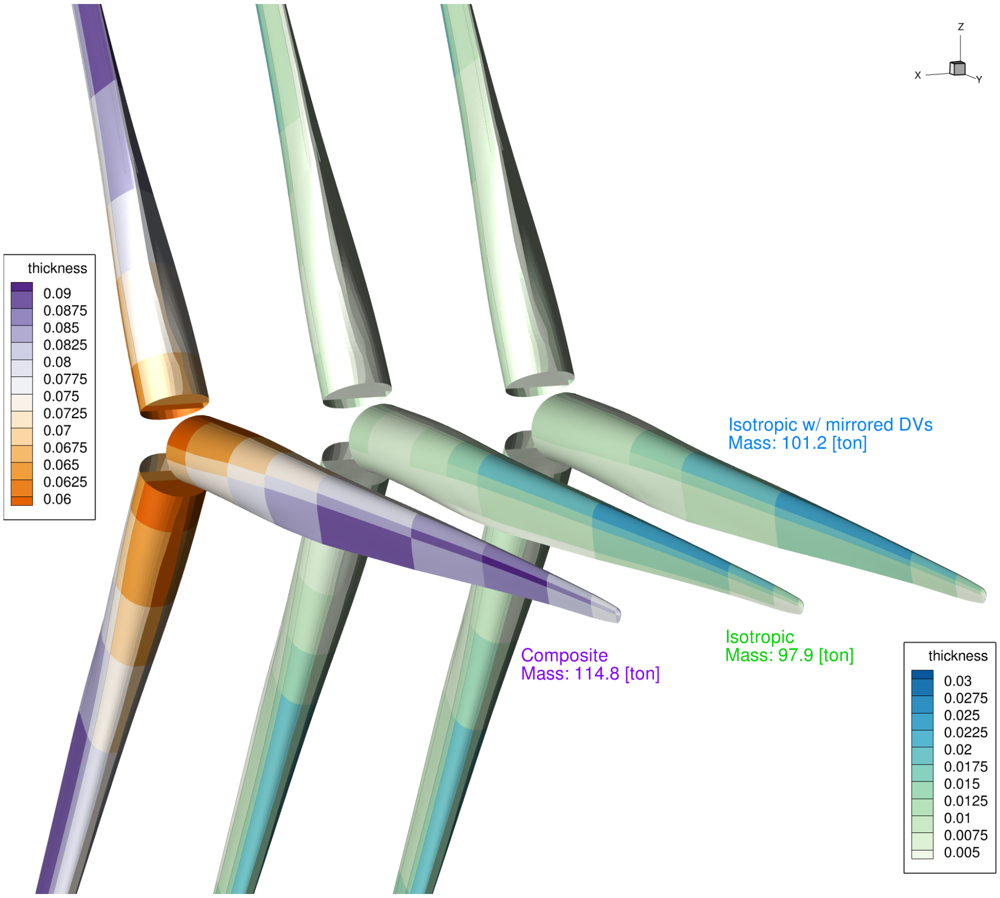

Figure 6.1: Thickness distribution for the composite and isotropic models, optimized using the loosely coupled strategy from Figure 5.2.

In Fig. 6.1, we observe the impact of material properties, design variables, and constraints more closely. We report the thickness distribution for the three different models and $d_{\text{tip}} = 3\ \mathrm{m}$. Because of the lower yield strength of the balsa core and the overly conservative failure criterion, the composite panel thicknesses are more than 3 mm higher over the entire blade with respect to the metallic isotropic configurations, requiring a different color map for clarity. All models present lower thickness close to the root, where the second moment of area of the blade sections with respect to edgewise bending is higher, increasing the local blade stiffness. Panel thickness is highest at mid span with the current parametrization. The optimizer reinforces the leading edge panels on the isotropic models, while focusing on the spar caps and a trailing edge section for the composite model.

6.2 Problem formulation

The goal of this chapter is to assess the impact of fiber angle tailoring on the final rotor mass without compromising the torque output. We only modify the twist distribution of the reference turbine blade. Substantial planform or sectional shape changes would impact the blade stiffness for the same panel thickness and fiber angle distribution. This would make it harder to identify the benefits of fiber tailoring. Changing the twist still gives the optimizer some freedom to adjust the load distribution and exploit spanwise aerostructural tradeoffs, while the structural design variables drive changes in local blade stiffness and total mass.

Two design points are evaluated concurrently to ensure the optimized blade is not maximizing torque output at a certain wind speed at the detriment of the rest of the operational envelope [15]. Moreover, the optimizer can leverage passive aeroelastic tailoring to optimize the blade response under different load conditions.

In this section we discuss the optimization design points (Section 6.2.1), its objective (Section 6.2.2), design variables (Section 6.2.3), and constraints (Section 6.2.4).

6.2.1 Wind conditions

The aerodynamic shape optimization study by Madsen et al. [15] used three representative wind speeds and rotation rates at below-rated conditions, where the controller follows an RPM schedule and does not adjust blade pitch. The wind speed, TSR, and RPM for each design point have been discussed in Section 4.2.1 and are listed in Table 4.1.

In this chapter, we only use two design points in the objective to consider multiple load scenarios without significantly increasing the optimization computational cost. More complex multipoint cases are discussed in Chapter 7. The majority of studies discussed in this chapter use the HW and BR design points in Table 4.1. BR has been used for single-point optimizations in Chapter 5 and is representative of the rotor performance at below-rated conditions. HW is a high load scenario used primarily as a proxy for structural sizing, without considering life-cycle and extreme loads. The MR condition is a low-speed, high TSR case used to investigate the impact of design point selection on the final design. A case with BR and MR is discussed in Section 6.2.2.

Similarly to what done by Madsen et al. [15], we use a Weibull distribution with average velocity $A = 8$ and form factor $k = 2$ to weigh the two design points in the objective, as further explained in Section 6.2.2.

6.2.2 Objective

The objective function from Section 5.1.1 is extended to account for multiple inflow conditions. Using a multipoint optimization formulation, every load condition is evaluated in parallel at each design iteration. We use a weighted average to combine the rotor mass and its torque at different wind speeds into a single metric, as shown in Eq. 6.1:

$$f(\mathbf{x}) = (1-\omega)\frac{M}{M_0} - \omega \sum_{k \in K} \left(\gamma_k \frac{Q_k}{Q_{k_0}} \right) \tag{6.1}$$

with K =[MR,BR, HW] the set of design points considered in each optimization. The mass M and torque Q_k are normalized with respect to their baseline values, M_0 and Q_{k_0}. The parameter ω emphasizes the torque output over the mass in the final objective. $\omega = 0$ corresponds to a pure mass minimization problem, while $\omega = 1$ lets the optimizer increase the torque without considering the rotor mass. We use $\omega = 0.7$ unless explicitly specified, consistent with Section 5.4. The coefficients γ are extracted from the Weibull curve mentioned in Section 6.2.1, normalized so that $\sum \gamma_k = 1$. The majority of the cases use $\gamma_{\text{HW}} = 0.36$ and $\gamma_{\text{BR}} = 0.64$. When MR is part of the optimization, the coefficient become $\gamma_{\text{BR}} = 0.535$ and $\gamma_{\text{MR}} = 0.465$.

6.2.3 Design variables

The structural patches used as independent design variables during the optimization are highlighted in Figure 6.2.

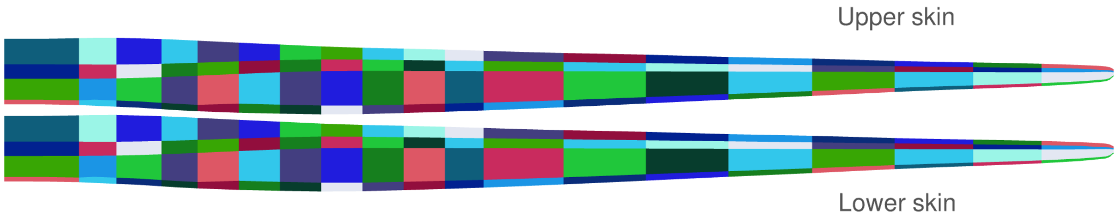

Figure 6.2: Skin structural design variables highlighted in different colors.

This is the result of the iterative process discussed in Appendix B. Spanwise, we opted to aggregate panels closer to the root and over the outer 50% of the span, consistently with the original DTU layup scheme. Chordwise, upper and lower skins are mirrored, and we combine together the trailing panels, leading panels, and spar caps respectively. This is also consistent with the approach by Bak et al. [24].

The structural layout is discretized into 148 independent patches, highlighted in different colors in Figures 6.2, 6.3 and 6.5. The upper and lower surface and the blade individual spars are split into 19 spanwise sections. The upper and lower surfaces are further split into 4 chordwise sections.

The section schematics in Figure 6.3 highlight individual design variables over an example blade section. The spar caps have an independent design variable. The elements in front of the main spars are grouped together, except for the nose reinforcement that remains independent. We split the area downstream of the main spars in two panels, one down to the reinforcement spar, and another covering the whole trailing edge area, including the TE itself. The panels on the suction and pressure side are handled by the same design variable. This simplifies the problem and obtain a more practical and manufacturable structural layout.

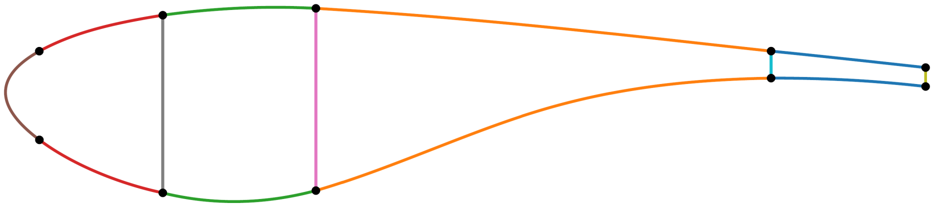

Figure 6.3: Structural breakdown of the blade section. Each independent design variable is highlighted with a unique color.

Each of the 148 panels described above has an independent thickness variable. We selected four different parametrizations for the ply angles in each panel – shown in Table 6.2. The fiber angle is defined on the shell element plane, with positive sign counterclockwise. Only the fiberglass ply angles are modified during the optimization. The principal axis of the balsa core is fixed and aligned with the blade spanwise axis.

Table 6.2: Fiber angle parametrization nomenclature.

Case name	# t_p	# θ	# x_{tw}
No	148	0	9
Uniform	148	1	9
Single	148	148	9
All	148	186	9

The simplest approach, "**N**o", does not include fiber angles as design variables. The principal axis of the sandwich panels is aligned with the blade spanwise axis and fixed during the optimizations. Parametrization "**U**niform" uses a unique fiber angle over the entire blade. Increasing the problem complexity, the "**S**ingle" parametrization assigns a single fiber angle for each of the 148 structural panels. In this scenario, the fiber angles of multiple fiberglass plies in the same panel skin are rotated by the same amount. Finally, the "**A**ll" approach has the largest design freedom. We decouple the angles of different material plies in the same patch. Since only a subset of panels has a multi-material fiberglass skin, the total number of structural design variables in this case is 186.

We do not make a distinction between plies on different sides of the balsa core. Due to the smeared stiffness approach, the element used in our structural analyses has averaged, uniform properties - as described more in detail in Section 3.2.2. Defining two separate angles for the same material would double the number of design variables without any meaningful benefit or insight on the final design. When the skin is composed of multiple fiberglass types, we only distinguish between outer and inner ply to dissect the optimized layout of parametrization "**A**ll". The structural layout does not have more than two different skin materials on the same patch. The spar caps skins are made of Uniax fibers, with additional Triax plies in the root area. There is no balsa core for the spars, consistently with the benchmark design. Conversely, the trailing and leading edge sections have a large balsa core and Triax skins, with Uniax reinforcements at the root. Finally, the main spars are made of a large balsa core and Biax fiberglass skins. The reader is redirected to Bak et al. [24, Fig.4.6] for more details.

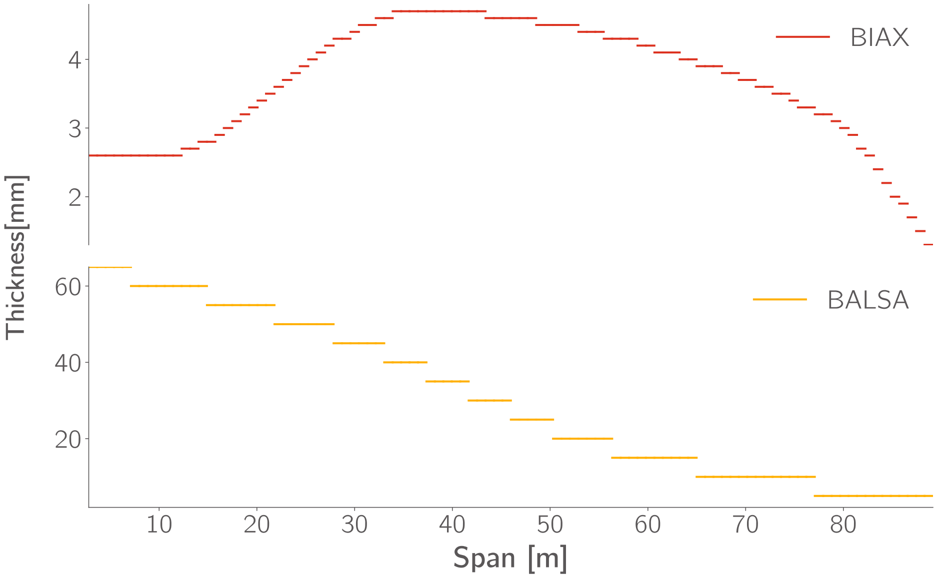

Figure 6.4: Thickness distribution for the Balsa and Biax layers of the rotor main spars.

The ply fractions are defined in the model initialization and are not modified during the optimization. We investigate the impact of this modeling limitation in Section 6.5. An example of spanwise ply thickness distribution for the main spars is shown in Figure 6.4.

Using a coupled aerostructural model allows to simultaneously modify the outer mold line and

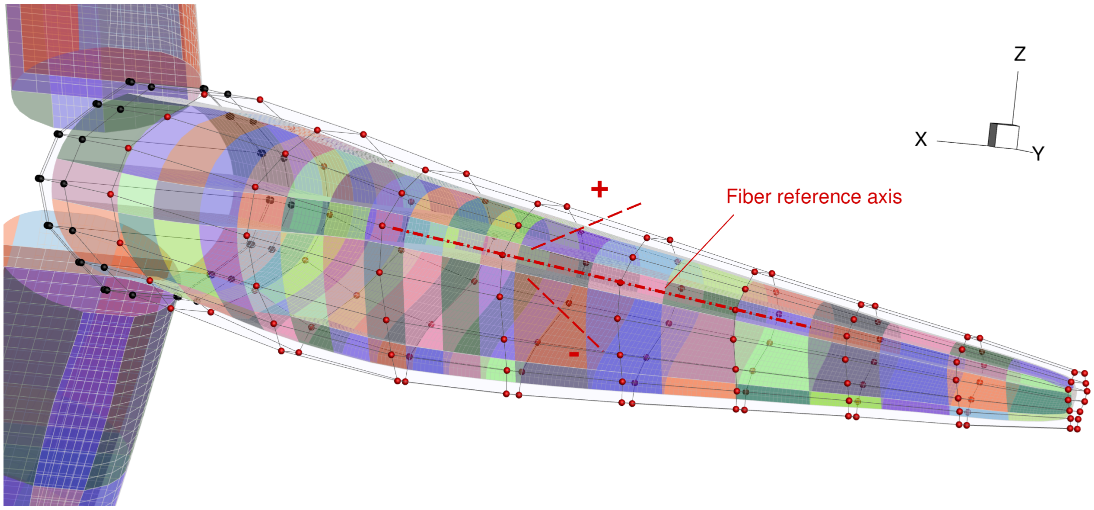

Figure 6.5: Active (red spheres) and inactive (black spheres) FFD control points and structural design variable patches over the baseline DTU 10 MW rotor configuration. The fiber angles are defined positive in the counterclockwise direction around the reference axis highlighted in the figure.

the structural layout of the rotor. Since the focus of this chapter is structural fiber tailoring, we only use a set of twist design variables to alter the geometry of the blade. FFD control sections highlighted in Figure 6.5 rotate around the section quarter-chord to modify the spanwise twist distribution. The 3 sections closer to the root are fixed to prevent geometry deformations to propagate to the root region. There are 9 active FFD sections between the span 10% and the blade tip.

6.2.4 Constraints

Linear constraints are enforced on both the thickness and fiber angle design variables. This way, the final design does not present large changes in thickness and fiber angle between adjacent panels. We limit the thickness change between adjacent panels to $\pm 2\ \mathrm{mm}$ and fiber angle changes to $\pm 12^\circ$. The latter only apply to the outer skin ply, with the inner skin intended as a “reinforcement” layer on certain blade sections.

We enforce a set of aerodynamic and structural nonlinear constraints to ensure the feasibility of optimized rotor layouts. The thrust at the BR condition cannot increase more than 14% than the baseline value, consistently with the original optimization studies by Madsen et al. [15] and Chapter 5.

We include a tip displacement constraint in the problem formulation, limiting the deflection to $5\ \mathrm{m}$ at BR. This constraint, not used in Chapter 5 due to structural model simplifications, has a dual role. Firstly, it acts as a proxy for a rotor-tower clearance constraint at extreme load conditions, not modeled in the present work. Secondly, it ensures that the optimized layouts remain within the linear structural analysis capabilities of TACS. The optimized layouts we present never exceed a tip displacement over span ratio of 10% even at HW condition. We set the displacement constraint at BR for computational load balancing purposes.

Finally, we enforce structural failure constraints based on the Tsai-Wu criterion from Eq.3.3. Ply strengths are extracted from the material properties in Table 3.4 and yield strains using the linear Young's relation. The KS aggregation approach mentioned in Chapter 2 is used to define three separate constraints similarly to Chapter 5- one each for the blade upper and lower surfaces, and one for the spars. We add a load safety factor based on which condition from Table 4.1 is used as the structural sizing load. We estimate these coefficient by running an analysis on the reference model and tuning the safety factor to get $\mathrm{KS}_\sigma = 1$. Using the benchmark original thickness distribution, we obtain a safety factor of 2.45 when HW is part of the optimization problem. In Section 6.4 we discuss a case using MR and BR load cases, so the latter is used to evaluate the stress constraints. A safety factor of 4.6 is used for that specific optimization.

6.2.5 Summary and results outline

In this chapter we present a set of optimizations with different parametrizations, setups, and objective function coefficients. Table 6.3 summarizes the problem formulation.

Table 6.3: Aerostructural optimization problem formulation.

	Name	Symbol	Qty
Objective	Equation (6.1)	Obj	1
Design Variables	Panel thickness	$\mathbf{t_p}$	148
	Fiber angles	θ	0/1/148/186
	Twist	$\mathbf{x_{tw}}$	9
			$\mathbf{157/158/305/343}$
Aerostructural Constraints	Stress	$\mathrm{KS}_\sigma \leq 1$	3
	Displacement	$d_{tip} \leq 5\ \mathrm{m}$	1
	Thrust	$F_x/F_{x_0} \leq 1.14$	1
			$\mathbf{5}$
Adjacency Constraints	Panel Thickness	$-2mm \leq \Delta\ t_p \leq +2mm$	319
	Fiber Angle	$-12^\circ \leq \Delta\theta \leq 12^\circ$	319
			$\mathbf{638}$

We gain insight on how the optimizer leverages our coupled CFD-CSM model and its structural design freedom to design lighter blades without compromising their torque output. Optimized layouts are compared looking at both integrated quantities, such as weight and torque, and the optimized design features.

Different fiber angle parametrizations are compared in Section 6.3. In Section 6.4 we quantify the effect of increasing fiber angle design freedom on the optimized rotor. The effect of design point selection and objective design weights is discussed in Section 6.4, and the impact of the ply fraction assumption is investigated in Section 6.5.

Table 6.4: Optimization cases summary. We use different sets of ply fractions, objective weight ω, and design points from Table 4.1.

Case	DVs (from Table 6.2)	Ply fractions	ω	Design points	Sections
2pt	N,U,S,A	Bak et al. [24]	0.7	BRHW	Sections 6.3 to 6.5
2ptω_9	S	Bak et al. [24]	0.9	BRHW	Section 6.4
2ptMR	S	Bak et al. [24]	0.7	BRMR	Section 6.4
2ptEQ	S	Equal	0.7	BRHW	Section 6.5

We introduce a consistent nomenclature in Table 6.4, which will be used in the remainder of this chapter. The initial design is sized through the staggered structural optimization approach discussed in Section 5.2. Optimizations using different parametrization listed in Table 6.2 start from different pre-optimized thickness and fiber angle distribution. This is to provide a better initial guess for the optimizer and speed up the tightly coupled aerostructural optimization. Nevertheless, to facilitate comparisons, all the mass and torque variations discussed later in this chapter are reported against the same baseline values for the "**N**o" parametrization. The baseline reference rotor weight for this configuration is 92 256 kg.

6.3 Increasing the number of independent fiber angles

In this section we investigate the impact of fiber angle design freedom on the optimized layouts. We compare four optimizations, one for each parametrization listed in Table 6.2. All cases use the same design points, γ_k, and ω objective coefficients defined in Section 6.2.2. The mass reduction and torque increase for each optimized rotor are listed in Table 6.5. Increasing the optimizer design freedom leads to increasingly lighter rotors, at the cost of lower torque outputs.

Without modifying the fiber angle (2pt-N), the optimizer adjusts the structural thickness and twist distribution to obtain a 2.9% mass reduction. Nevertheless, this lighter structure is able to withstand loads 9% higher than baseline. This concurrent mass reduction and torque increase is

Table 6.5: Higher design freedom leads to lighter blades and lower torque increases. * Objective values are based on each design initial layout and not the reference baseline for 2pt-N as done for torque and mass.

	Mass [%]	Torque [%]		Obj*
		HW	BR	
2pt-N	−2.9	+9.1	+10.5	−0.479
2pt-U	−3.5	+8.7	+10.2	−0.479
2pt-S	−9.7	+7.6	+9.5	−0.491
2pt-A	−19.3	+4.1	+7.1	−0.480

consistent with what observed for optimization cases in Section 5.4. The torque output for the 2pt-N case is the highest of this set of results. Given the smaller design freedom to further reduce mass, the optimizer prioritizes torque to achieve a final objective function value comparable to other cases.

Adding one angle design variable for case 2pt-U does not lead to significant mass reduction. The rotor is 0.6% lighter than the 2pt-N case while its torque is just 0.3-0.4% lower. The optimized fiber angle value is $\theta = -1.2°$. Since different components withstand different loads, a large uniform fiber angle change would be beneficial in some sections of the blade while weakening other areas. The small optimized angle and the negligible (compared to other cases) mass reduction suggest that spanwise-aligned fibers remain the best choice when local fiber tailoring is not considered in the design process.

The largest mass reductions in Table 6.5 are obtained by leveraging local fiber angle steering. Having an individual fiber angle design variable for each structural panel (2pt-S) leads to a mass decrease of almost 10% compared to baseline. At the same time, torque increases remain within 1.5% of the 2pt-N case. When the optimizer is given control of individual material fiber angles in each panel (2pt-A), the mass is reduced by 19.3% compared to the baseline layout. The torque output for the HW and BR design points is increased by 4.1% and 7.1% respectively, compared to baseline. In this case the optimizer exploits the structural design freedom to further reduce the rotor mass rather than increasing torque, despite the higher emphasis of the latter in the objective. The final objective function value for the 2pt-A case is consistent with 2pt-N.

The objective function we describe in Section 6.2 can be used to explore an optimal design envelope, as done in Section 5.5. The problem could be reformulated as a mass minimization problem with a prescribed torque output. Results in Table 6.5 suggest that fixing the torque output to match the initial value, fiber tailoring would enable substantial mass gains.

The thickness and fiber angle distribution for the cases in Table 6.5 are illustrated in Figure 6.6. The thickness distribution for the 2pt-N and 2pt-U is barely distinguishable, so we only report the

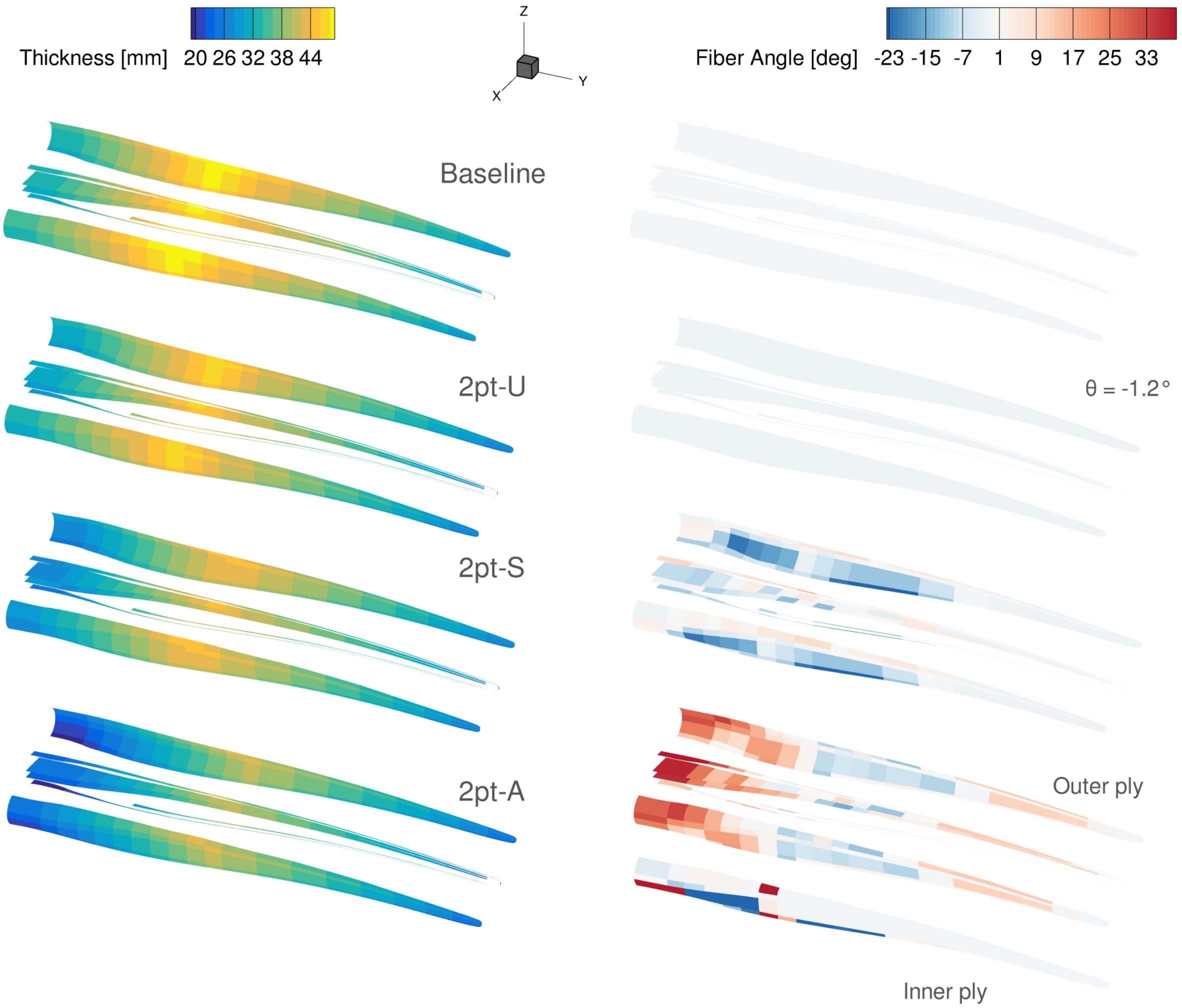

Figure 6.6: Structural thickness and fiber angle distribution for the optimizations in Table 6.5.

latter. Warm colors indicate counterclockwise fiber rotation, colder colors a clockwise rotation (shown in Figure 6.5).

The baseline panel thickness ranges between 30 mm and 50 mm. The thickest panels are located around the span 35%. The adjacency constraint mentioned in Section 6.2.4 ensures a smooth spanwise thickness variation. The thickness trends observed on the sequentially-optimized baseline rotor are consistent with the original benchmark [24]. This highlights both our modeling consistency with benchmark data, and the design effort that went into reference model definition. The absolute thickness values are lower than the publicly available DTU 10 MW data because our simulations do not include extreme and fatigue loads, which drive the structural sizing in practical applications.

The optimized designs maintain comparable thickness trends but display lower local panel thickness. The 2pt-U design (and similarly the 2pt-N case) has a few panels thicker than 48 mm. This area is shifted approximately 1-2m more outboard than on the baseline.

The lighter 2pt-S configuration maximum panel thickness is just below 47 mm at 40% of the span. The inboard panels are also thinner, with a minimum thickness of 26 mm in the aft root section of the blade. Looking at the fiber angle distribution, we observe that the spar caps remain mostly aligned with the spanwise axis – except for a clockwise rotation of less than 10° at the span 15%. The leading edge panels are rotated counter clockwise over the entire span, never exceeding 10° rotation. Conversely, the inboard trailing edge panels and the trailing edge section at mid-span are rotated clockwise by more than 20°. Panels with the larger fiber rotation, between the span 12% and 25%, show the largest thickness reduction – up to more than 5 mm compared to baseline.

The most complex 2pt-A case shows different fiber angle trends. The additional 38 fiber angle design variables, present on panels with multi-material skins, are primarily located near the root and along the trailing edge. We distinguish between an outer, main ply, and an inner, "reinforcement" ply where multiple plies are present. In this model, the inner ply is made out of Uniax fiberglass, reinforcing Triax skins along a specific direction.

Looking at the outer ply, the fibers in both skins and spars in the root region are rotated counterclockwise between 20° and 40°. This applies to the main spars as well, conversely to what observed for the 2pt-S case. Moving away from the root, fibers are rotated clockwise by up to 9° between the span 30% and 60%. Leading and trailing edge panels are then rotated counterclockwise by more than 10° up to the span 85%.

The inner plies, not subject to the adjacency constraints, show more abrupt orientation changes. At the root trailing edge, the Uniax layer is rotated 40° clockwise, while the panels right outboard are rotated clockwise by almost 30°. The latter pattern is consistent with the 2pt-S layout. The thickness in this area is more than 10 mm lower than the baseline layout, with a minimum thickness of 21 mm at the root aft section.

Decoupling ply layers in 2pt-A allows the optimizer to apply non-monotonic fiber angle patterns on the outer layer, while maintaining patterns similar to 2pt-S on the Uniax inner plies. This is favored by the absence of adjacency constraints for this set of angles. The optimizer tries to align the principal axis of individual structural panels to the main load path over the blade. As a result, the areas with the highest fiber steering show the largest thickness reduction. The structural panel discretization and linear constraints over the blade have a major impact on the design space of these optimization problems. Future investigation could further explore different parametrizations and relax certain thickness or fiber angle constraints.

Figure 6.7 provides insight on how the optimizer changes the twist distribution to redistribute the aerodynamic loads and increase the torque output. By convention, as mentioned in Section 5.4,local sections are perpendicular to the inflow for Twist=0° and are aligned with the hor-

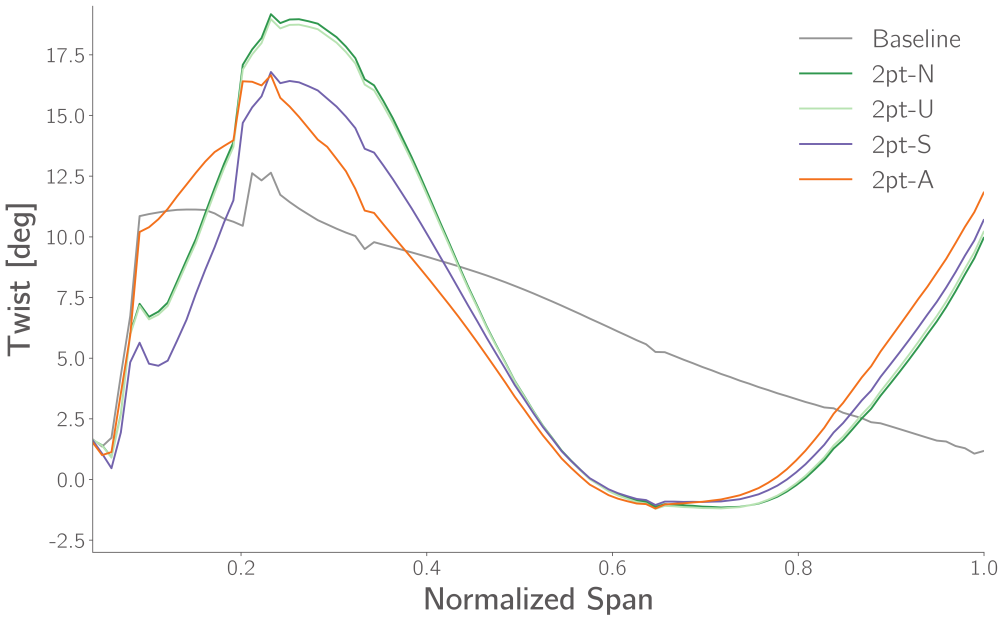

Figure 6.7: Twist distribution for optimizations with different number of fiber angle design variables

izontal axis at Twist=+90°. This means that a higher twist in Figure 6.7 corresponds to lower local angles of attack. All the twist distributions reported in this work are extracted from the blade deflected at operating condition (HW, if not otherwise specified) and thus account for twist-bend coupling and passive load alleviation effects.

All the designs considered in this section have comparable minimum twist between the span 60% and 75%. Without sectional changes, the airfoil polars used for the benchmark design by Bak et al. [24] still apply. The optimizer increases the local angle of attack in this part of the blade while avoiding efficiency losses closer to $C_{l_{max}}$ and stall. This twist distribution is consistent with what discussed in our previous investigation (Chapter 5). The optimizer shifts the aerodynamic loads inboard to mitigate out-of-plane bending moment increases without compromising the torque output. The $\approx 2^\circ$ twist offset between the 2pt-N and 2pt-A cases in the outermost 25% of the blade explains the lower torque output of the latter configuration. As expected from Table 6.5 and Figure 6.6, the twist differences between 2pt-N and 2pt-U are negligible.

Figure 6.8 shows that, for the same configuration, the effective twist on the deflected blade varies by approximately 1° between the HW and BR design points. Bend-twist coupling has a minor impact compared to fiber angle parametrization or, as shown later in Section 6.4, the objective function coefficients. Shell elements are however not the most suited model to capture large bend-twist coupling effects from the structure. Future studies could investigate different modeling approaches to further exploit passive load alleviation to substantially modify twist distribution

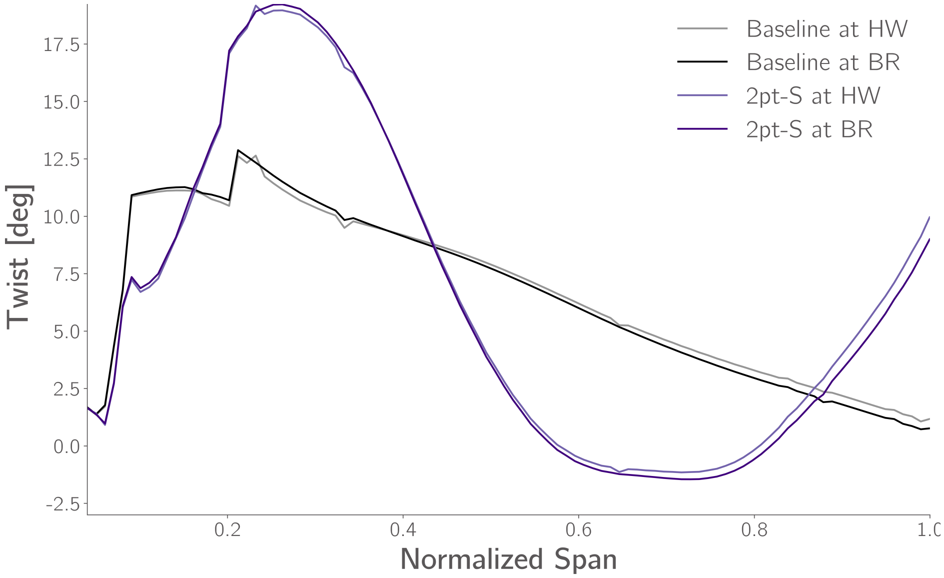

Figure 6.8: Twist distribution for the Baseline and 2pt-S rotors at HW and BR.

under different load conditions.

6.4 Tailoring fibers for different objective functions

In this section we investigate the impact of changing objective function coefficients and design points on the optimized layouts. The 2ptω_9-S case considered below differs from 2pt-S only for the coefficient $\omega = 0.9$ in Eq. 6.1, as shown in Table 6.4. We compare these designs with another case, named 2ptMR-S, identical to 2pt-S except for replacing the HW design point with MR. The coefficients γ_k and structural safety factor for this case have been discussed in Section 6.2.2 and Section 6.2.4 respectively. The mass and torque variations with respect to baseline are reported in Table 6.6.

Table 6.6: The objective weight ω has a higher impact on the final design than changing the set of optimization design points. * Objective values are based on each design initial layout and not the reference baseline for 2pt-N as done for torque and mass.

	Mass [%]	Torque [%]			Obj*
		HW	BR	MR	
2pt-S	-9.7	+7.6	+9.5	-	-0.491
2ptω_9-S	+9.3	+14.2	+13.6	-	-0.915
2ptMR-S	-4.4	-	+10.2	+13.3	-0.496

Increasing the objective emphasis further away from mass in case 2ptω_9-S leads to a torque output 4.1% and 6.8% higher than 2pt-S at BR and HW respectively. The rotor mass however is more than 9.3% higher than baseline, and 19% more than 2pt-S. Since the mass accounts for only 10% of the objective function value in this formulation, the optimizer has a weaker incentive to design a thinner structure.

Case 2ptMR-S mass is just 4.4% lower than baseline. The averaged torque increase from Table 6.6 is higher than for 2pt-S. This latter configuration shows lower increases at the HW condition, so likely the case selection biases the optimizer towards torque gains for the 2ptMR-S layout. The final design is 5% heavier but 0.7% more efficient than 2pt-S at BR.

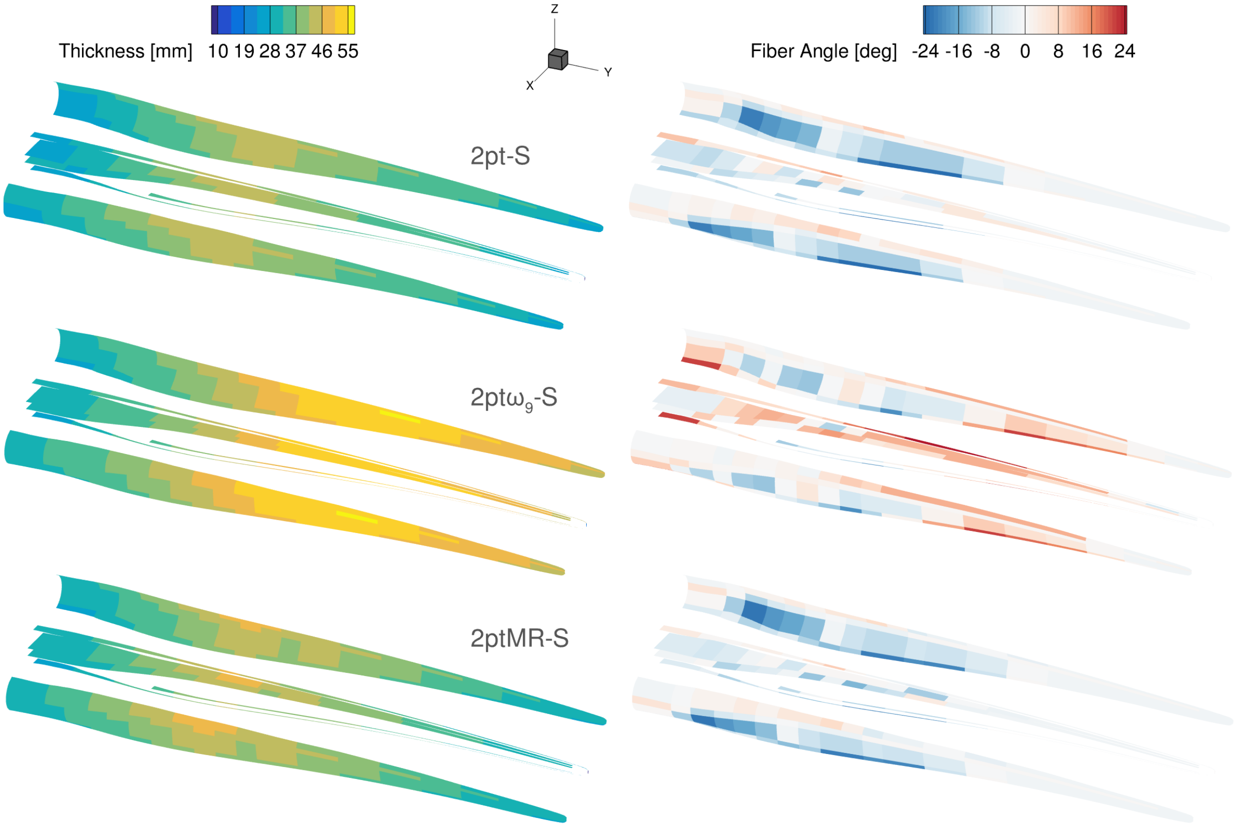

Figure 6.9: Comparing thickness and fiber angles of optimizations with different objective weight ω and design points

The structural layouts shown in Figure 6.9 reflect the mass variations discussed above. The panel thickness for 2ptω_9-S is more than 55 mm thick at the span 60%. The spar fibers are predominantly rotated counterclockwise. Similarly, the fiber angles on the leading an trailing edge panels are rotated counterclockwise by up to 12° in this section. Conversely, trailing edge panels in the span inboard half are rotated clockwise by approximately the same amount. The adjacency constraint is driven by the spar caps (which are within $\pm 1^\circ$ of the spanwise axis), hence the similar opposing values around these elements.

The fiber pattern of 2ptω_9-S resembles a tradeoff between 2pt-S and 2pt-A, without the latter large counterclockwise rotation in the root area. Conversely, the fiber patterns of 2pt-S and 2ptMR-S show little differences. The higher mass of the latter layout mostly comes from leading edge panels more than 3 mm thicker than the reference 2pt-S case.

From a structural standpoint, the design point selection has negligible impact on the final layout with the proposed problem formulation. Due to the limited geometrical design freedom, the blade optimized for the lower speed scenarios (MR,BR) does not show large differences from the reference optimizations. Conversely, the objective weight ω has a large impact on both the structure thickness distribution and maximum values.

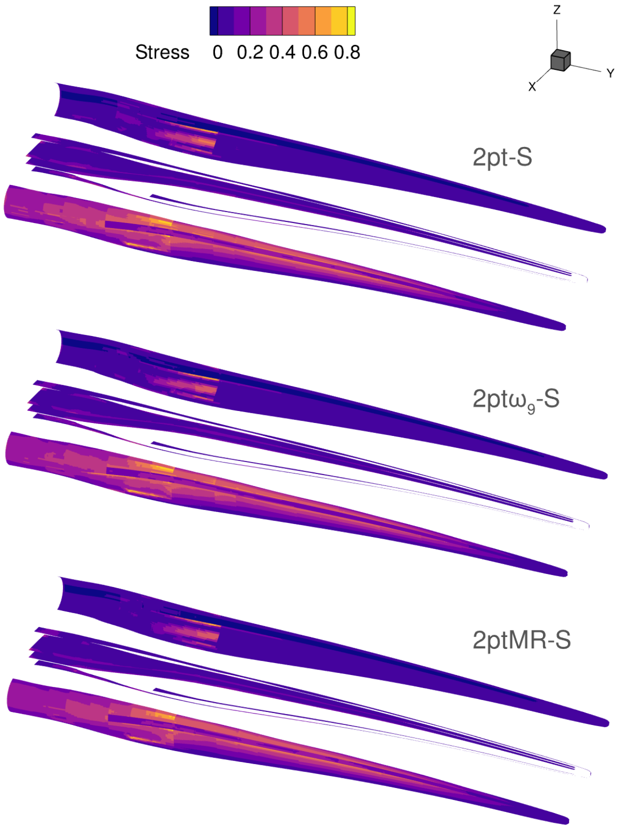

Figure 6.10: Comparing stress patterns of optimizations with different objective weight ω and design points.

The stress distribution for the three designs discussed in this section are reported in Figure 6.10.

Despite the different load and thickness distributions, the blade components show comparable stress peaks and patterns. The blade upper skin is the element bearing the highest loads. Both skins have stress peaks around 30% of the span.

All stress constraints are active at the optimum. Thus, the main differences in thickness distribution discussed in Figure 6.9 must come from the displacement constraint. Different aerodynamic load distributions induce larger local blade deflections at different spanwise locations. The optimizer rearranges the maximum thickness location to maintain the tip deflection below 5 m at BR.

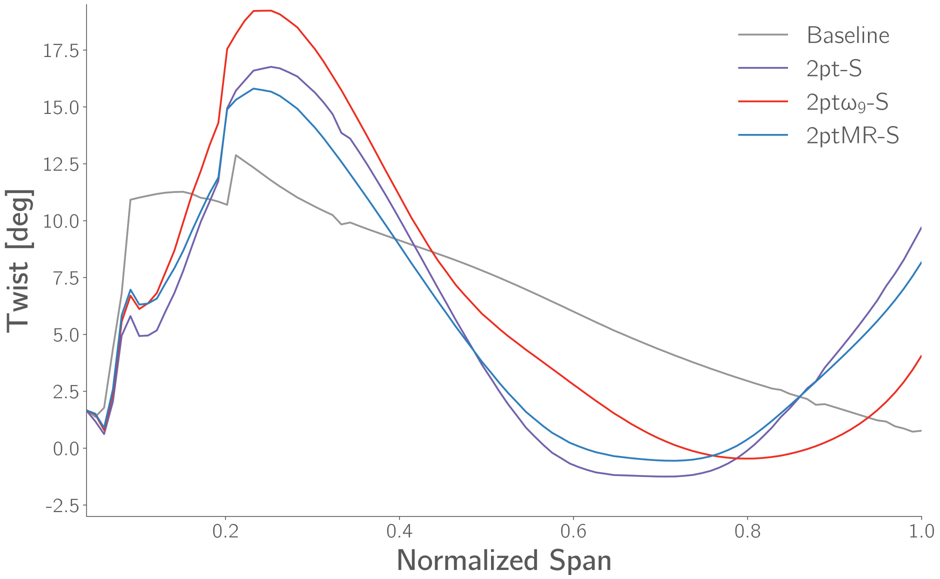

Figure 6.11: Twist distributions of optimizations with different objective weight ω and design points

Similarly to what observed for the structural layout, the twist distributions shown in Figure 6.11 is mostly influenced by the objective weight ω. The 2ptMR-S twist is close to the 2pt-S case, with an approximately 1.5° lower twist between the span 60% and 80%, and a higher tip twist.

The 2ptω_9-S case lowest twist value is consistent with the 2pt-S reference, but its location is shifted further outboard by more than the span 20% compared to the other cases in the figure. This trend is consistent with what observed in Chapter 5. When the torque output has a larger weight in the objective, the optimizer increases the local angle of attack further outboard, accepting a mass increase penalty. The consequent aerodynamic load increase is highlighted in Figure 6.12 for both in-plane and out-of-plane loads.

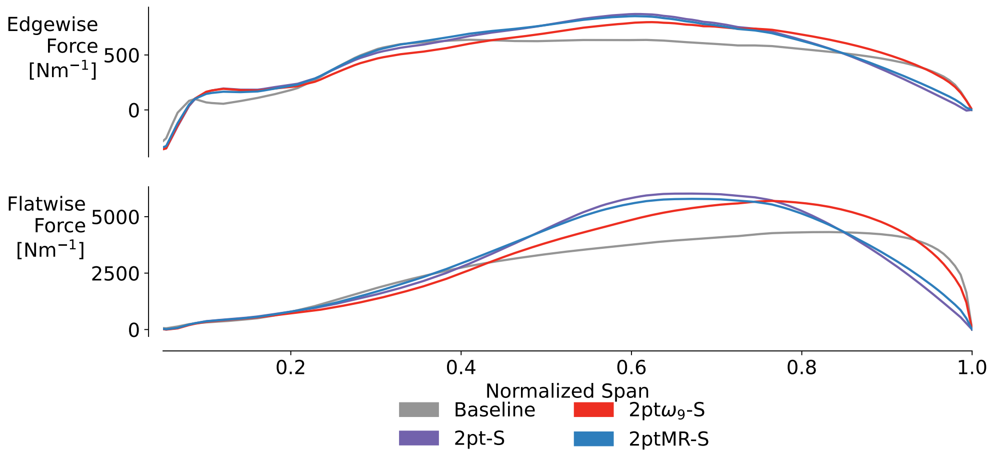

Figure 6.12: Load distribution of optimizations with different objective weight ω and design points

6.5 The effect of ply fractions

Finally, in this section we investigate the impact of fixed ply fractions during the optimization. We compare the 2pt-S layout with another case using a different fraction of skin and core in the sandwich panel. For all the optimization studies presented so far, we based the ply fractions on the original DTU 10 MW benchmark documentation – see Section 6.2.3 and Figure 6.4. For case 2ptEQ-S, conversely, we assign equal fractions to each ply. This case is a simplification and likely does not translate into practical applications. Nevertheless, we use it to emphasize the impact of non-parametrized components of our structural model, paving the way for future investigations aimed at ply fraction optimization.

Table 6.7: The reference 2pt-S case outperforms a design with "equal" skin and core fractions. * Objective values are based on each design initial layout and not the reference baseline for 2pt-N as done for torque and mass.

	Mass [%]	Torque [%]		Obj*
		HW	BR	
2pt-S	−9.7	+7.6	+9.5	−0.491
2ptEQ-S	−4.4	+5.4	+6.7	−0.492

The 2ptEQ-S case is outperformed by the 2pt-S reference both in terms of rotor mass and torque output, as shown in Table 6.7. Torque increases are below 7% at both design points and the mass is, coincidentally, within 100 kg of the 2ptMR-S case described in Section 6.4.

The thickness distribution of the two layouts, reported in Figure 6.13, shows marked differences. The higher fraction of fiberglass plies in 2ptEQ-S leads to a thinner structure, with values ranging between 10 mm and 30 mm. The larger balsa core in 2pt-S has poorer mechanical properties than the thick fiberglass/epoxy skins in 2ptEQ-S. However, as reported in Table 3.4, balsa wood density is an order of magnitude lower than the fiberglass/epoxy composite. Unsurprisingly, the more conventional sandwich panel we use for 2pt-S leads to a lighter rotor design.

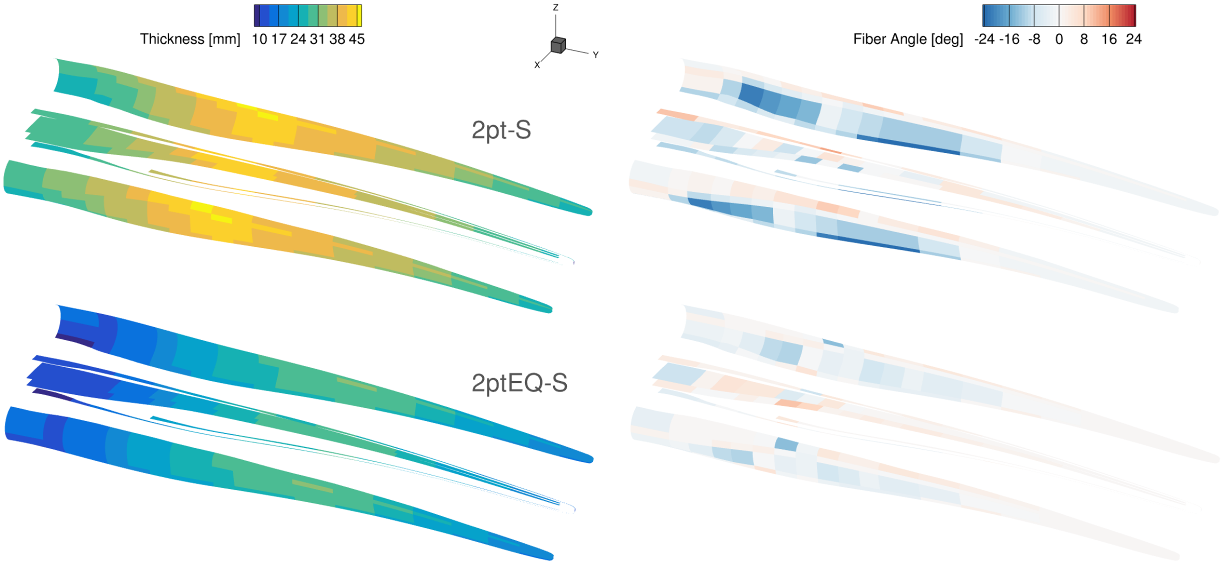

Figure 6.13: Thickness distribution and fiber angles of optimizations with different ply fractions.

Consistently, the lower range of fiber angle values on the 2ptEQ-S layout is an effect of the larger use of composite skins. All fiber angles are within $\pm 10^\circ$ of the spanwise axis in this configuration. Due to the higher strength of the fiberglass material compared to the balsa core, optimizing the fiber angles has a lower impact on the final layout.

The twist distributions in Figure 6.14 show that the optimizer drives the loads further inboard with a more aggressive twist reduction at the span 65%. The tip on the 2ptEQ-S case is unloaded through a twist increase more than 2° higher than 2pt-S. The angle differences are higher between these two cases than for the comparison between 2pt-S and 2ptMR-S in Section 6.4. These results highlight the mutual interaction between optimized outer mold line and ply fraction distribution. Future investigations could leverage the combination of ply fraction, angles, and planform design variables to further exploit aerostructural design tradeoffs.

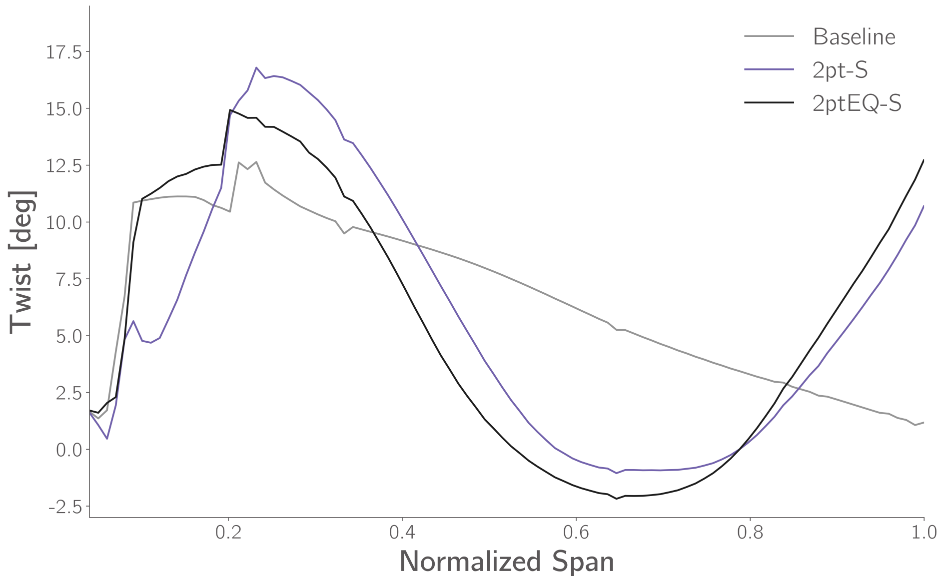

Figure 6.14: Twist distributions of optimizations with different ply fractions.

6.6 Conclusions

This chapter presents a range of high-fidelity aerostructural optimization studies of wind turbine rotors, with an emphasis on structural composite fiber tailoring. The tightly coupled CFD-CSM approach and a gradient-based optimization strategy enable us to simultaneously optimize the blade structural thickness, fiber angle, and twist distribution. Final designs generate more than 10% higher torque at the prescribed inflow conditions. Depending on the problem formulation, the optimizer can reduce the rotor mass by almost 20% without modifying the blade planform.

We compare several optimization cases throughout the chapter. First, we quantify the impact of additional structural design freedom on the final design. When given control of individual panel and ply fiber angles, the optimizer obtains large mass reductions while increasing the torque output by more than 4%. While structural thickness patterns are consistent for different cases, different parametrizations lead to large changes in fiber angle patterns. Future investigations could explore more complex fiber parametrizations including more practical manufacturing constraints.

As shown in previous investigations, the problem formulation has a large impact on the final design. A larger weight of the torque output in the objective function leads to a more aggressive twist distribution and a thicker structure with non-monotonic fiber angle distribution over the span. The resulting design is heavier, but has a torque output more than 13% higher than baseline. Given consistent structural safety factors, changing design points in the objective does not have a large impact on the structural layout and twist distribution. We expect larger differences to emerge when planform and local section shape variables are included in the optimization.

Finally, we explore how changing different ply fractions affect the final design. We select a simplified ply stack where each ply has the same ply fraction. The larger role of fiberglass/epoxy skins leads to a structure that is thinner but heavier than other reference optimizations that use the benchmark sandwich panel setup. The fiberglass composite has better mechanical properties than the balsa core, but a much higher density, driving the mass up. This also affects the optimal ply angle distribution, whose range is lower than other cases discussed in this chapter.

The present study highlights the potential of blade-resolved aerostructural optimizations and paves the way for more complex studies leveraging more refined and practical structural parametrizations. The hereby proposed formulation can be enhanced with additional geometrical design variables to further exploit aeroelastic tradeoff and design lighter and more efficient wind turbine rotors.

Chapter 7

Robust Rotor Design through Multipoint Aerostructural Optimization

In this chapter we present the first aerostructural optimization studies which simultaneously enable large structural, blade planform, and airfoil shape modifications. The benchmark rotor is optimized for a set of wind speeds, demonstrating an effective design approach that can be extended to site-specific studies. We investigate how emphasizing the rotor mass and torque at different wind speeds impacts the final designs. Moreover, we discuss how the design point selection affects the off-design performance of the rotor. Our high-fidelity model enables the comparison of blade-resolved geometrical, structural, and aerodynamic load features. The characteristics of the optimized blade sections are further discussed using conventional analysis tools.

The coupled CFD-CSM model enables granular control of the structural thickness over different components of the blade and, as a first in this study, the local shape of blade airfoil sections. However, optimizing the blades for a single wind speed would likely lead to impractical design features such as sharp surface curvature Madsen et al. [15]. Single-point optimization of lifting devices often leads to poor or sub-optimal off-design performance. This is particularly an issue when the airfoil shape is modified by the optimizer [153, 15].

Multipoint optimizations address this by evaluating the performance of the rotor at several wind speeds. Madsen et al. [15] showed the benefits of this strategy for a CFD-based aerodynamic shape optimization study. We extend that approach to a coupled aerostructural model of the same turbine rotor, increasing torque more consistently over the below-rated operational envelope compared to the studies in Chapters 5 and 6.

The studies in this chapter use the same structural model and parametrization from the previous Chapter 6. Unless otherwise specified, we only use thickness design variables and keep the ply angles fixed and aligned with the blade spanwise axis. The composite layup fraction, based on Bak et al. [24], is not altered during the optimization process. We primarily use this simplified parametrization because the focus of this work is the impact of the multipoint problem formulation and airfoil design variables. Only in a final study we include fiber angle design variables to investigate a more complex aeroelastic tailoring problem.

7.1 Problem formulation

In this section we dissect our approach to the design optimization problem. The objective function from Section 5.1.1 is extended to consider multiple wind speeds, as already shown in Chapter 6. The torque outputs are aggregated into a weighted average, together with the rotor mass, to form a single objective function. We refer to this as a multipoint optimization problem.

The selection of relevant design points for high-fidelity multipoint optimization has been already discussed in Section 4.2.1. Section 7.1.1 summarizes the design point weights for the optimizations presented in this chapter. The objective, design variable, and constraint set up are discussed in Section 7.1.2, Section 7.1.3, and Section 7.1.4 respectively.

7.1.1 Design points

Table 7.1: Design point weights for the objective function. Cases are further discussed in Table 7.3.

γ_k	HW	BR	MR	BR-T10	BR-T5
1pt	-	-	1.000	-	-
1pt-PS	0.000	0.000	1.000	-	-
3pt-Eq	0.333	0.333	0.333	-	-
3pt-Wb(T)	0.208	0.368	0.424	-	-
5pt	0.132	0.233	0.269	0.206	0.16

The coefficients used to weigh different inflow conditions in our 1-point, 3-point, and 5-point optimizations are reported in Table 7.1. This table also introduces the nomenclature for the cases presented in this chapter.

Single point optimizations weights are either 1 or 0. The 3pt-Eq case uses equal weights for all the design points. The other 3-point cases and the 5-point cases use weights based on the Weibull distribution presented in Section 4.2.1 and Figure 4.4. The effect of different weights in the objective function on the final design is discussed in Section 7.3. The impact of additional design points in the optimization is illustrated in Section 7.5.

7.1.2 Objective

The objective for our multipoint optimization problem, same as what used for Chapter 6, is formulated as following:

$$f(\mathbf{x}) = (1-\omega)\frac{M}{M_0} - \omega \sum_{k\in K}\left(\gamma_k \frac{Q_k}{Q_{k_0}}\right) \tag{7.1}$$

with Q_k the torque at a given inflow condition k (with K =[HW,BR, MR], except for Section 7.5 where BR-T10 and BR-T5 are also included) and M the rotor mass. Both metrics are normalized with respect to the baseline values, Q_{k_0} and M_0 respectively. The γ_k coefficients are the multipoint weights related to the inflow probability distribution discussed in Section 7.1.1.

7.1.3 Design variables

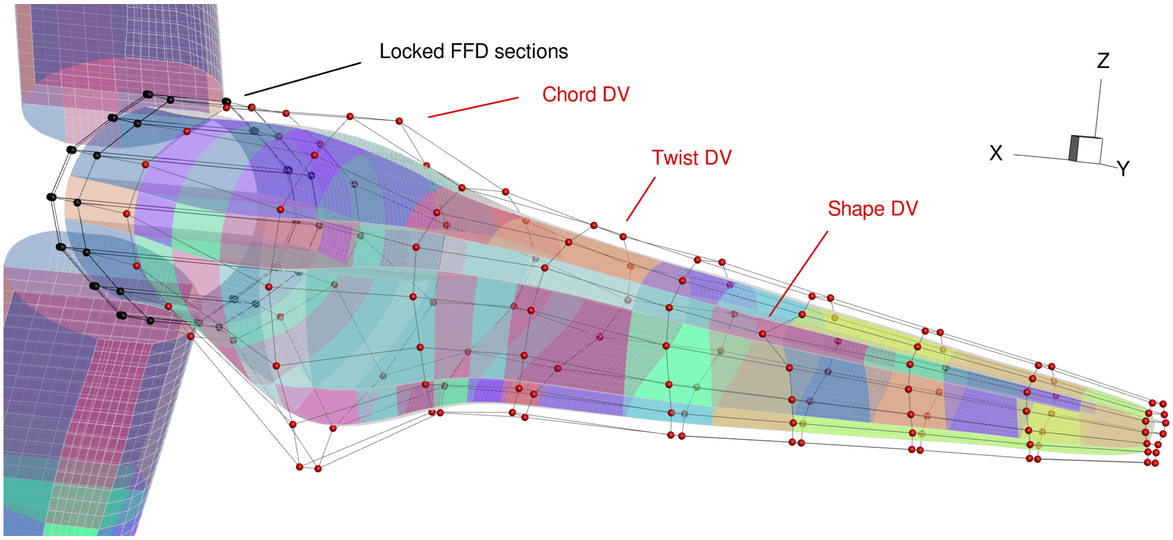

Figure 7.1: Geometrical and structural parametrization. Each colored patch has an independent structural thickness variable. Example geometrical design variables are included in the figure.

The optimizer simultaneously handles the blade shape variables and structural thickness distribution. FFD control points are highlighted as red spheres in Figure 7.1. Of the total 12 sections, the 3 closest the root are kept fixed to prevent FFD deformations to extend to the blades intersection. The twist distribution is modified by rigidly rotating the 9 spanwise active FFD sections around the quarter chord of the local airfoil. In addition to what used for Chapter 6, the chord distribution is modified by stretching the same FFD sections along the z-axis.

The optimizer also controls the individual displacement of FFD points to modify the local airfoil section. The node pairs along the trailing edge cannot move relative to each other to prevent mesh deformation failures. The leading edge node pairs are constrained to move the same amount in opposite directions to prevent spurious twist modifications overlapping with the twist design variable. There are 12 active control points for each section (6 each for the upper and lower surfaces) for a total of 108 local shape variables. We add a set of constraints on the chordwise airfoil thickness to prevent mesh failures and impractically thin sections.

The structural mesh is the same as the one used in Chapter 6. The colored patches in Figure 7.1 highlight the skin and spars partition for parametrization purposes. The optimizer directly controls 148 structural thickness variables. We constrain the thickness to vary by a maximum of ± 2 mm between adjacent panels. This is done to prevent unrealistic discontinuities in the structural layout, which would be unpractical to manufacture. Future investigations could relax this constraint between different components to ensure smooth variations over the span, but larger thickness differences between spar caps, spars, and skin panels.

7.1.4 Constraints

Similarly to what done in previous chapters, we enforce a set of nonlinear constraints on the optimized layouts, in addition to the linear geometrical and structural constraints on mentioned in Section 7.1.3. We constrain structural stress, tip displacement, and aerodynamic thrust.

The simulations at different wind speeds (Section 7.1.1) are run in parallel since the aerostructural solutions are independent from each other. We make an educated choice on which constraint is evaluated at a given inflow condition, to avoid costly coupled adjoint solutions for constraints that would inevitably be inactive at the optimum.

We enforce the stress constraints at HW, the case with the highest aerodynamic loads. The structural stress is evaluated element using the Tsai-Wu criterion and then aggregated using the KS formulation [124]. We define 3 stress constraints: one each for upper and lower skins, and a third one for the set of spars. A safety factor of 2.45 is included in the constraint evaluation - obtained using the reference thickness distribution as explained in Section 6.2.4.

The thrust constraint is evaluated at the BR inflow conditions, mirroring what done for the same configuration in previous optimization studies [15]. This constraint prevents an excessive increase of the axial loads that would require a structural re-design of the tower - not modeled in our simulations.

The blade tip displacement is also evaluated at BR. We enforce a maximum displacement of 5 m at this flow regime. This corresponds to a deflection equal or lower than 9 m at the highest HW load case on the baseline design. While there is in general no guarantee that this would still apply on the optimized layouts, we verified that the displacement at HW remains consistently below 9 m for every optimized design presented in this paper. This ensures that structural analysis remains within the linear domain defined by $d_{\text{tip}}/\text{span} \leq 10\%$.

7.1.5 Summary and results outline

The full set of objective, design variables, and constraints are reported in Table 7.2.

Table 7.2: Aerostructural optimization problem formulation.

	Name	Symbol	Qty
Objective	Equation (7.1)	$f(\mathbf{x})$	1
Design Variables	Panel thickness	$\mathbf{t_p}$	148
	Twist	$\mathbf{x_{tw}}$	9
	Chord	$\mathbf{x_{ch}}$	9
	Shape	$\mathbf{x_{sh}}$	108
	Fiber angles	θ	0/148
			$\mathbf{274/422}$
Aerostructural Constraints	Stress at HW	$\mathrm{KS}_\sigma \leq 1$	3
	Displacement at BR	$d_{tip} \leq 5$ m	1
	Thrust at BR	$F_x/F_{x_\mathrm{i}} \leq 1.14$	1
			$\mathbf{5}$
Geometric Constraints	Airfoil Thickness	$0.6 \leq \mathrm{T_f/T_i} \leq 3.0$	24
	Airfoil Thickness (Tip)	$0.95 \leq \mathrm{T_f/T_i} \leq 3.0$	72
	LE constraints		9
			$\mathbf{105}$
Adjacency Constraints	Panel Thickness	$-2mm \leq \Delta\, t_p \leq +2mm$	319
	Fiber Angle	$-12° \leq \Delta\theta \leq 12°$	319
			$\mathbf{319/638}$

We perform a set of optimizations to understand the impact of the different objective and formulation parameters described in this section. We first discuss single and multipoint optimizations in Section 7.2. Three different designs are compared, highlighting the advantages of the multipoint formulation. In Section 7.3 we compare two multipoint optimized rotors with different design point weights γ_k. We investigate how the optimizer can tailor the rotor to potentially different site-specific scenarios. Similarly to what done in Chapter 5, we compare two optimizations with different objective weight ω in Section 7.4. We discuss how the optimized layout changes the blade planform and airfoils when torque maximization is prioritized over mass reduction. Finally, two additional inflow conditions, BR-T10 and BR-T5, are included to perform a 5-point optimization in Section 7.5. This design performs better at off-design conditions, preventing flow separation to occur at BR-T5. The optimized layouts presented in this sections are summarized in Table 7.3.

All results in this section are reported as percentage difference with respect to the baseline mass

Table 7.3: Summary of optimized configurations.

Name	Design points	ω	γ_k	Sections
1pt	MR	0.7	-	Section 7.3
1pt-PS	HW, BR, MR	0.7	$\gamma_{MR} = 1$	Section 7.2
3pt-Wb	HW, BR, MR	0.7	Weibull	Sections 7.3 to 7.5
3pt-Eq	HW, BR, MR	0.7	Equal	Section 7.3
3pt-WbT	HW, BR, MR	0.9	Weibull	Section 7.4
5pt	HW, BR, MR BR-T5, BR-T10	0.7	Weibull	Section 7.5

and torque. We compare sectional and spanwise aerodynamic loads, and the structural thickness and stress distribution on the final designs. The baseline geometry is consistent with Madsen et al. [15]. The initial structural thickness layout is obtained by performing a loosely coupled mass minimization study (see Section 5.2) at HW regime enforcing stress constraints and a maximum tip displacement of 9 m.

7.2 Single and multipoint optimization

The reference 3pt-Wb design is compared to two single-point optimizations at the MR condition. We select this inflow condition for comparison since it has the highest γ_k value for the 3pt-Wb optimization. The 1pt case only evaluates the design at MR. We allow a maximum thrust increase of 17$\%$ at this condition, consistent with the 3pt-Wb rotor. The 1pt-PS (for "pseudo 1-point") is effectively a multipoint optimization identical to 3pt-Wb, except for the design point weights set as $\gamma_{MR} = 1$, and $\gamma_{BR} = \gamma_{HW} = 0$. Additionally, we enforce the torque at BR and HW to be equal or higher than the baseline values, to prevent unrealistic configurations.

Table 7.4: Torque and mass variation for 1-point and 3pt-Wb optimizations. Gray values are not evaluated during the optimization.

	Torque [%]			Thrust [%]			Mass [%]
	HW	BR	MR	HW	BR	MR	
1pt	+11.8	+14.4	+16.7	+13.4	+16.2	+17.0	−39.4
1pt-PS	+2.3	+0.0	+29.4	+14.8	+13.5	+26.3	−1.8
3pt-Wb	+10.5	+13.2	+17.1	+10.9	+14.0	+17.2	−36.2

The integrated loads and mass in Table 7.4 summarize the characteristics of the three different designs. The optimizer reduces the mass by more than one third on both the 3pt-Wb and 1pt case and by less than 2% for the 1pt-PS case.

We observe a torque increase of almost 30% at MR on the 1pt-PS case, while the off-design points show little to no improvement. The torque constraint at BR is active for the 1pt-PS case.

As for the torque values of the 1pt and 3pt-Wb rotors, we see consistent load increases, with the highest improvement at the MR condition. Surprisingly, the 1pt case outperforms the 3pt-Wb case at the HW and BR conditions, despite them not being part of the optimization formulation. However, we also observe that thrust at BR is higher than the 14% maximum allowed value on the 3pt-Wb case.

Inconsistencies like this one highlight the limitations of the single point approach to deliver a robust design that is also comparable with multipoint optima.

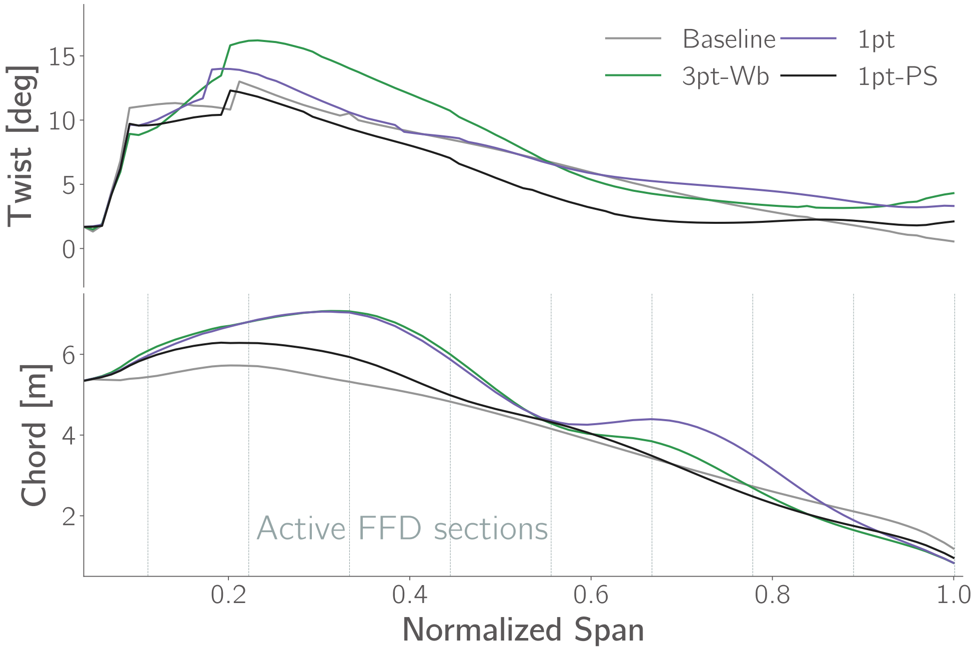

Figure 7.2: Optimized twist, chord, and relative thickness for multipoint and single-point optimized designs.

The blade planform characteristics of the three blades reported in Figure 7.2 are also radically different. The 3pt-Wb and 1pt cases have an almost-identical chord distribution in the inboard half, with maximum chord at 40% of the span. The 1pt planform has another chord increase between the span 60% and 70%, which is present to a smaller extend in the 3pt-Wb case. A similar chord distribution was obtained by Bottasso et al. [64, Figures 6 and 9], who showed how enforcing a taper constraint lead to heavier blades. The 1pt-PS optimized chord is closer to baseline, but the twist reduction (corresponding to a higher local angle of attack) is more aggressive than for the other designs.

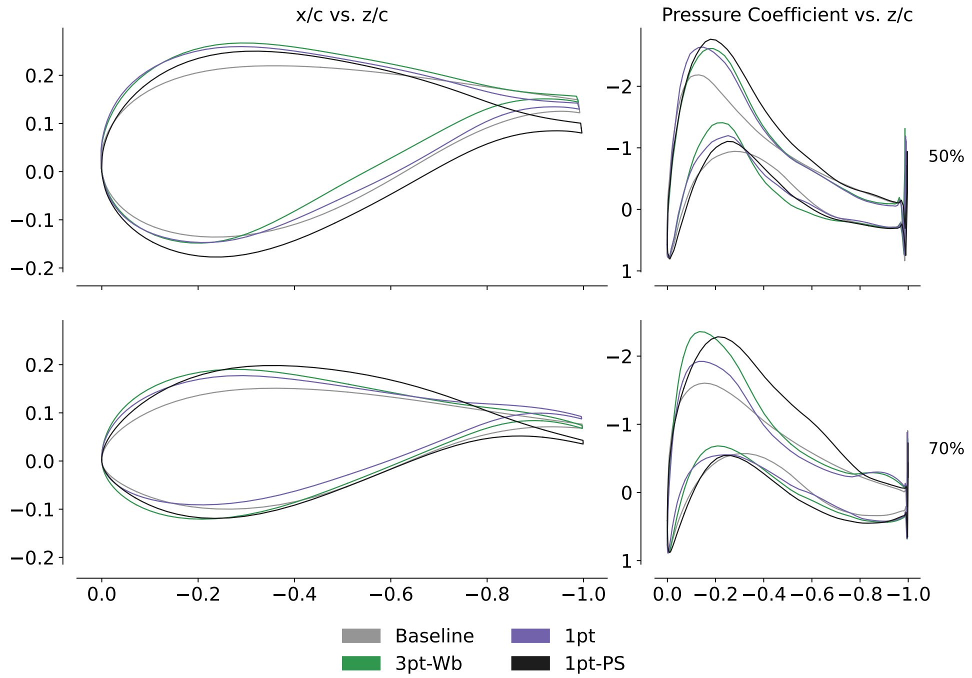

Figure 7.3: Airfoils and pressure distributions at MR for 1-point and 3pt-Wb optimizations.

From the section comparison in Figure 7.3, we observe that at mid span the 1pt-PS case has lower pressure on the suction side due to higher local pitch. At the span 70%, the higher loads are obtained through a higher camber and thickness in the aft part of the airfoil.

The design tradeoffs on the 3pt-Wb case lead to a rotor that has a robust torque increase at different design points (compared to 1pt-PS) while satisfying the constraints of the original formulation in Table 7.2 (conversely to 1pt). We use this design as comparison reference in the following sections.

7.2.1 Results verification

Before dissecting additional scenarios, we re-analyze the 3pt-Wb layout using the L0 mesh. Loads reported in Table 7.5 confirm the trends for integrated loads. The torque increase at MR is almost 5% lower than what predicted by the L1 mesh, but still 2.2% higher than the BR case and 3.9% higher than the MR case. The thrust value below the 14% constraint threshold suggests that the optimizer would have an even larger margin of improvement tailoring the design for this specific mesh.

Table 7.5: The aerodynamic loads increase consistently for different DLCs on the more refined L0 mesh

	Mesh	Torque [%]			Thrust [%]
		HW	BR	MR	BR
Weibull	L1	+10.5	+13.2	+17.1	+14.0
	L0	+8.5	+10.2	+12.4	+13.5

7.3 Varying design point weights

In this section we discuss the impact of the parameter γ_k in Eq. 7.1, which defines the relevance of a specific inflow condition within an optimization. We compare the 3pt-Wb case discussed in Section 7.2 with the 3pt-Eq case with equal design point weights, as shown in Table 7.1. The latter formulation, comparable to Dhert et al. [14], is the simplest multipoint choice and does not account for any site-specific information. The integrated loads of these two cases are reported in Table 7.6

Table 7.6: Torque and mass variation for 3-point optimizations with different sets of design point weight γ_k.

	Torque [%]			$\sum \gamma_k \Delta Q_k$ [%]	Mass [%]
	HW	BR	MR		
3pt-Wb	+10.5	+13.2	+17.1	+14.3	−36.2
3pt-Eq	+12.4	+12.9	+16.0	+13.8	−36.4

The optimized mass values, reported in Table 7.6, differ by 0.2%. The 3pt-Eq highest load increase is again at MR. Due to the lower γ_{MR}, torque is 1.1% lower than the 3pt-Wb at the same design condition. Similarly, the higher γ_{HW} leads to a 1.9% higher torque at HW for the 3pt-Eq design. The performance at the BR condition is within 0.3% for the two configurations. The γ-averaged torque is 0.5% higher for the 3pt-Wb case.

As shown in Figure 7.4, the airfoils and sectional loads show small differences between the two designs. The pressure peak for the 3pt-Wb layout is marginally lower than the 3pt-Eq case and shifted aft. Changing γ_k values has a smaller impact on the final design than the formulation differences discussed in Section 7.2. However, 1-2% torque output variation can make a large impact at industrial scale. This suggests that appropriate site-specific data needs to be taken into account for high-fidelity rotor optimization.

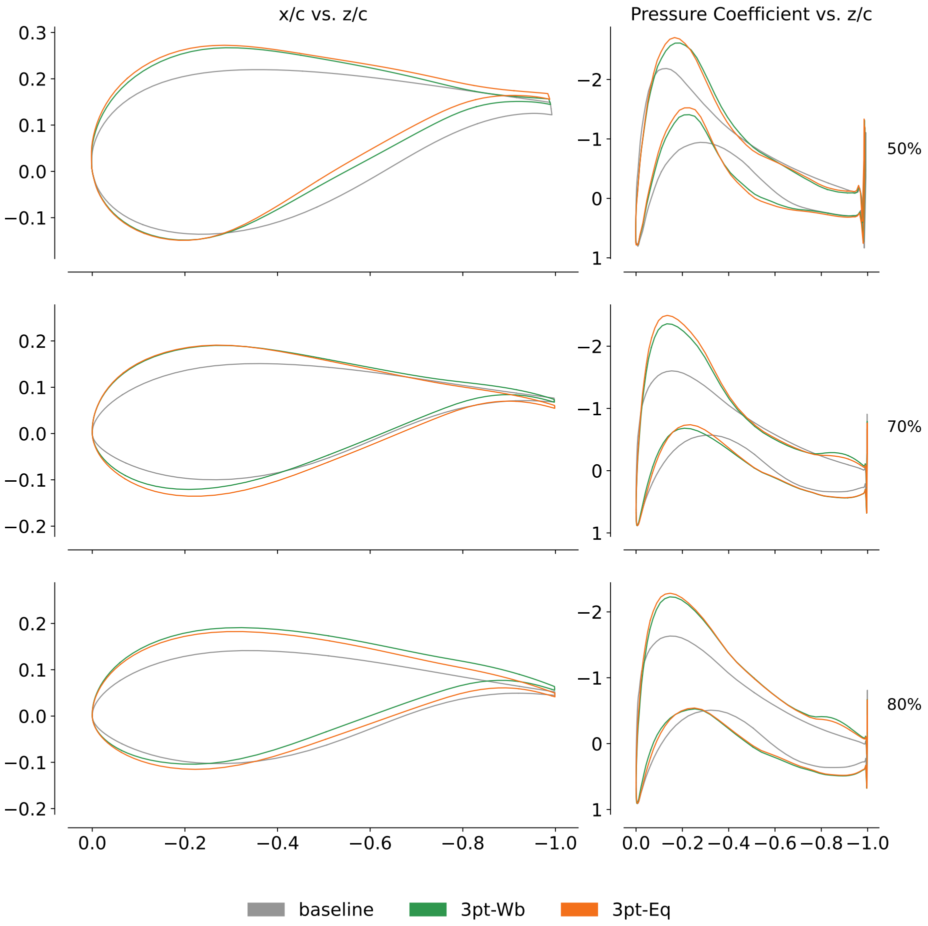

Figure 7.4: Airfoils and pressure distributions at MR for 3-point optimizations with varying design point weight γ_k.

7.4 Varying objective weight

We now discuss two optimizations with same design point weights γ_k and a different objective weight ω. The 3pt-Wb case is compared with the 3pt-WbT, which has $\omega = 0.9$.

Table 7.7: Performance and mass variation for multipoint optimizations with different objective weight, compared to baseline.

	Torque [%]			$\sum \gamma_k \Delta Q_k$ [%]	Mass [%]
	HW	BR	MR		
3pt-Wb	+10.5	+13.2	+17.1	+14.3	−36.2
3pt-WbT	+14.9	+15.8	+23.1	+18.7	−16.4

The mass decrease for 3pt-WbT is less than a half of the 3pt-Wb reference, as reported in

Table 7.7. The torque increases are between 2.6% and 4.4% higher on the 3pt-WbT rotor, with a torque 23.1% higher than baseline at MR.

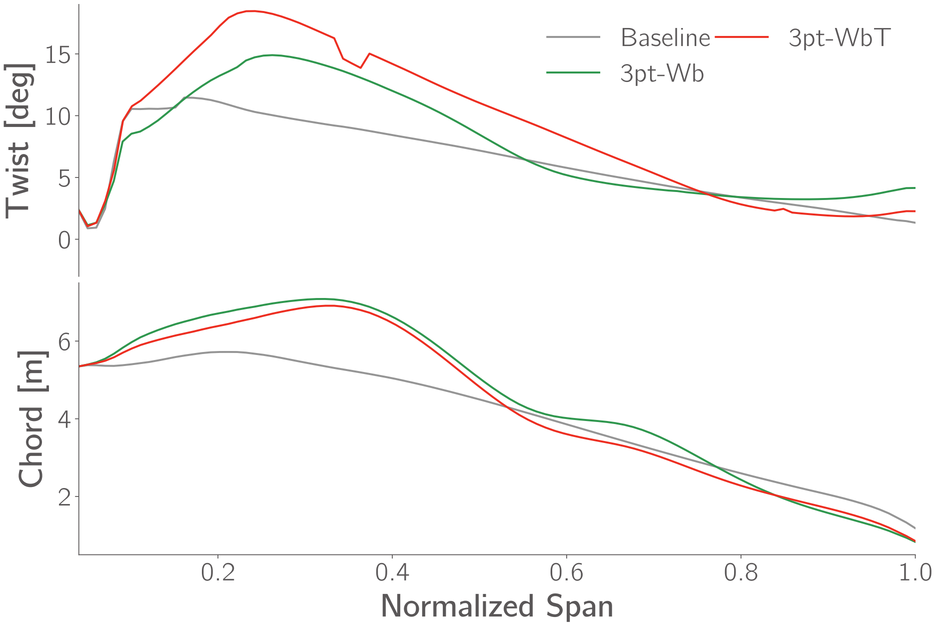

Figure 7.5: Optimized twist and chord for optimized designs with varying objective weight ω.

Ignoring the "hump" due to postprocessing artifacts around the span 40% in Figure 7.5, the 3pt-WbT case has a higher twist angle (and thus lower angle of attack) up to approximately the span 75%. Conversely, the 3pt-WbT blade incidence is approximately 2° higher than 3pt-Wb in the span outer 20%. The chord is shorter throughout the span for the 3pt-WbT design, and decreases more linearly between the span 60%-80%.

The airfoil differences in Figure 7.6 are more substantial than what observed in Section 7.3. The loads at the span 90% match closely despite the difference in twist seen in Figure 7.5. The span 80% section has a sharper LE for the 3pt-WbT case, leading to a lower C_p peak. At 60% of the span, the 3pt-WbT has a similar LE curvature but a flatter upper surface, consequence of different chordwise thickness and camber trends. The resulting C_p peak is lower than for 3pt-Wb but decreases less abruptly, indicating higher chordwise flow speed over the blade suction side.

These planform and section design differences shift the spanwise loads closer to the blade tip for the 3pt-WbT case, as shown in Figure 7.7. Both the edgewise and flatwise loads are higher than the 3pt-Eq case in the outermost 80% of the blade. The resulting higher bending moment explains both the increase in torque and the lower mass reduction seen in Table 7.7.

Figure 7.8 illustrates the structural differences between the two designs. Practical sizing based on extreme loads would likely lead to thicker structures. Nevertheless, discussing the structural design trends is relevant to understand the aerostructural tradeoffs that the optimizer identifies

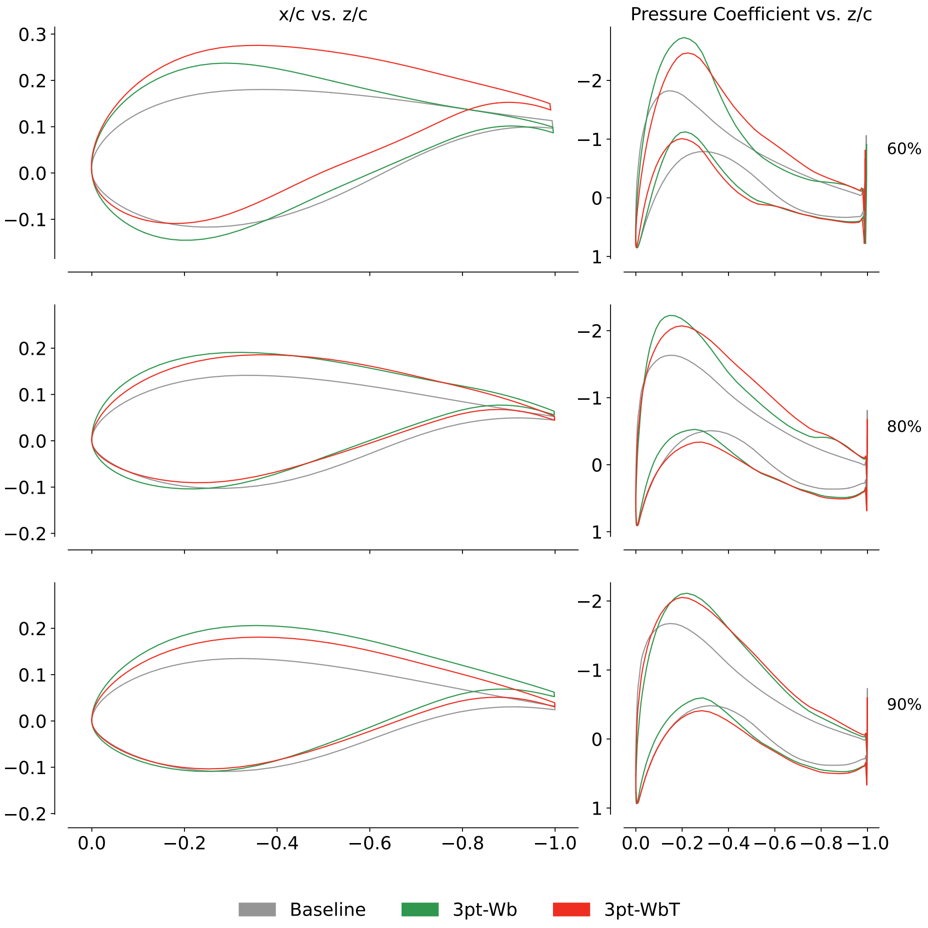

Figure 7.6: Airfoils and pressure coefficient distributions at BR for optimized designs with varying objective weight ω.

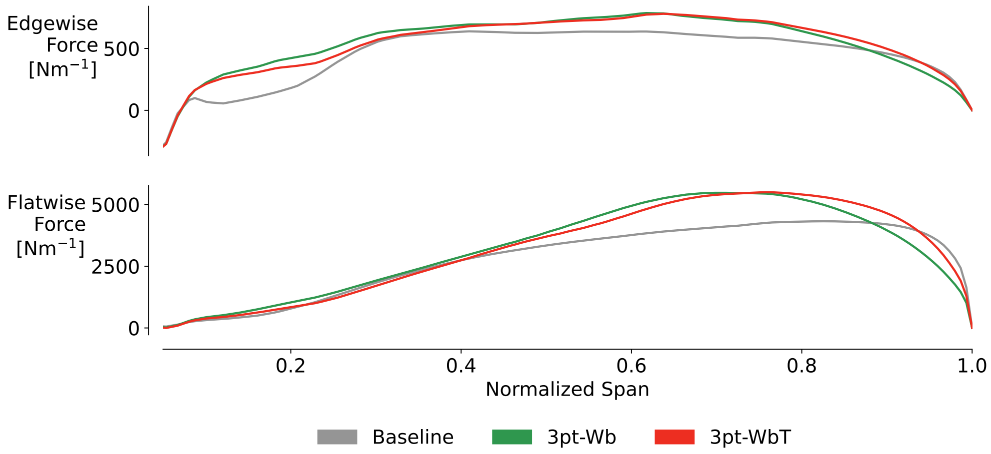

Figure 7.7: Lift distribution at BR for optimized designs with varying objective weight ω.

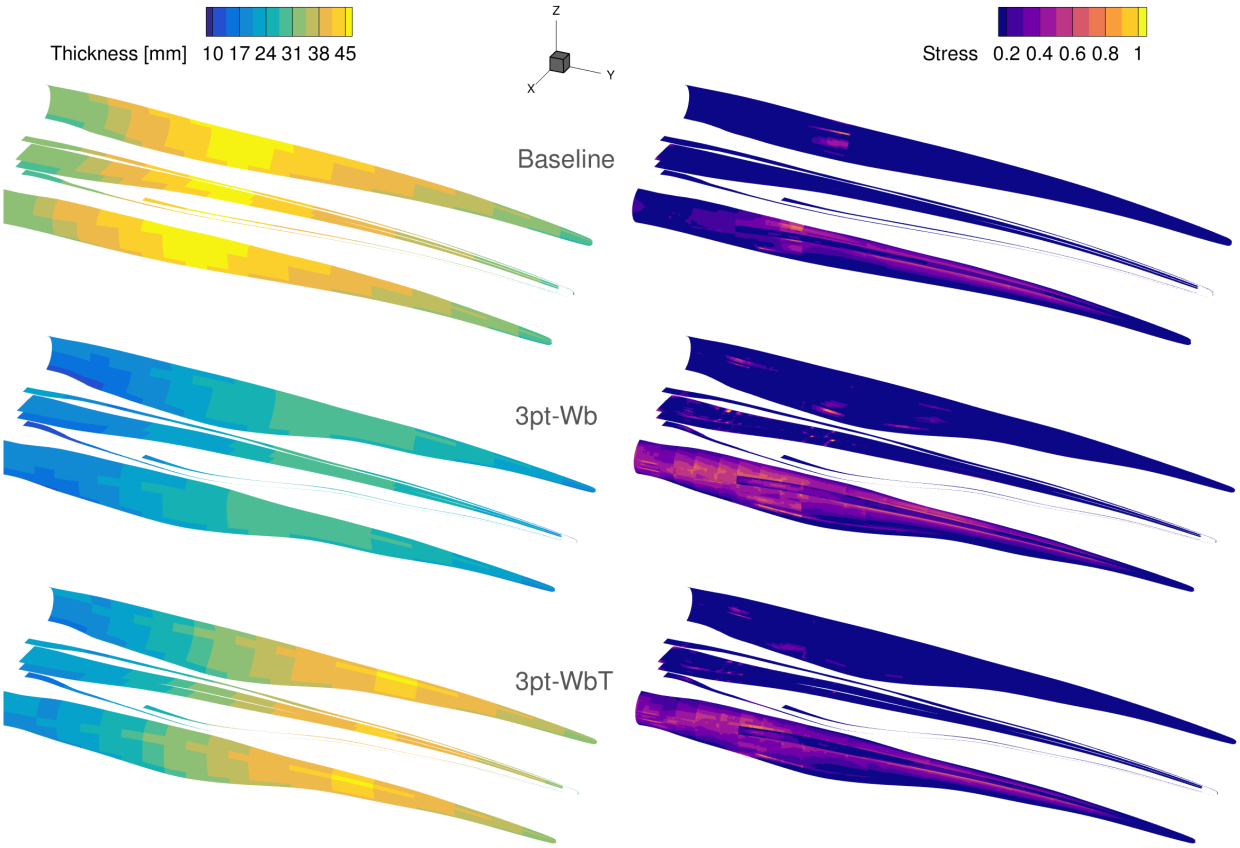

Figure 7.8: Structural thickness and stress distribution for the baseline rotor and optimized blades with varying objective weight ω.

using the problem formulation proposed in this dissertation.

In both designs, the optimizer increases the chord and thins the panels at the span 40%, where the baseline layout is the thickest. This enables mass reduction without compromising the blade structural stiffness. The 3pt-Wb rotor has a maximum thickness of approximately 30 mm at mid-span. The heavier 3pt-WbT goes up to 45 mm at the 70% section. The outboard shift of the loads seen in Figure 7.7 is consistent with these modifications.

Displacement and stress constraints are active on both designs, except for the spar stress constraint on the 3pt-WbT rotor. The stress distributions on the right of Figure 7.8 suggest that the baseline design (obtained through the same loosely coupled approach from Section 5.2) is driven by the stress constraint. The highest stress sections correspond to the most stressed area. However, the maximum thickness is located further outboard in both optimized designs, closer to the sections with the highest out-of-plane curvature. This strongly hints at the tip deflection constraint being the driving design factor for these structural layouts.

7.5 5-point optimization

In this section we investigate how including more design points in the optimization affects the rotor performance within and outside the design TSR schedule. Compared to the 3pt-Wb reference, the 5pt case includes BR-T10 and BR-T5 inflow conditions, weighted as per Table 7.1.

Table 7.8: Performance and mass variation for 3-point and 5-point optimizations, compared to baseline. Gray values are not part of the optimization formulation.

	Torque [%]					Thrust [%]	Mass [%]
	HW	BR	MR	BR-T10	BR-T5	BR	
3pt	+10.5	+13.2	+17.1	+12.7	−0.7	+14	−36.2
5pt	+12.5	+13.4	+16.3	+14.5	+18.2	+14	−34.5

The 5-point design is less than 2% heavier than the 3-point reference, as shown in Table 7.8. Due to the different formulation and higher γ_{HW}, the optimizer increases the HW torque by 2% rather than further thinning the blade structure. Conversely, the MR torque is just 0.8% lower than for 3pt-Wb.

The two rotors have a torque difference of almost 19% at the BR-T5 design point, which is not included in the 3pt-Wb optimization. The lift distribution reported in Figure 7.9 shows a loss of aerodynamic load for the 3pt-Wb layout in the region between the span 50% and 60%.

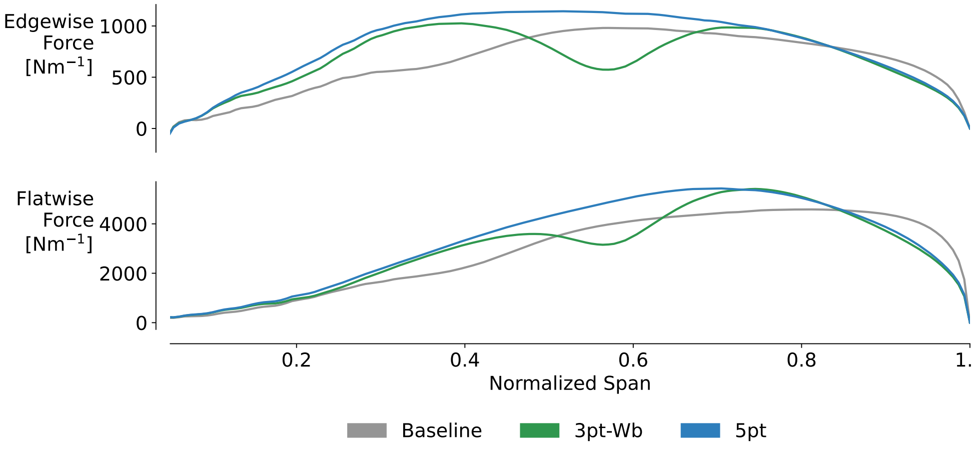

Figure 7.9: Lift distribution at BR-T5 for the 3-point and 5-point optimized designs.

Examining Figure 7.10, the 5-point case has a larger chord and lower inflow angle than the 3-point counterpart in the affected region. The lower twist of the 5-point layout in the outermost 40% of the blade is the main planform feature that contributes to the difference in torque outputs at the other inflow conditions reported in Table 7.8.

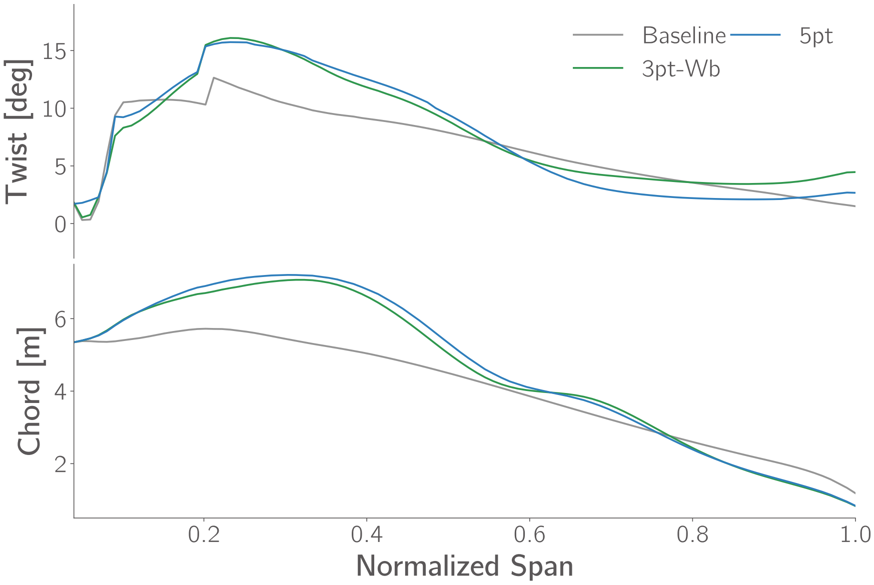

Figure 7.10: Optimized twist, chord, and relative thickness for 3-point and 5-point optimized designs.

The pressure distribution at the span 55% and 60% sections in Figure 7.11 suggests that the 3pt-Wb design is experiencing stall at this inflow condition. The in-plane velocity contours in Figure 7.12 highlight a large separation region from 40% of the chord. The combination of the relatively small planform modifications shown in Figure 7.10 and the lower thickness shown in Figure 7.11 suffices to prevent stall at BR-T5 condition.

ADflow and RANS solvers in general have limitations in predicting the stall onset, as discussed in Chapter 5. Without considering transition, surface roughness, and other turbulence models, we cannot verify that the blade would stall with the same mode in the same real world scenario. Nevertheless, this specific modeling concern does not invalidate the core findings of this comparison: the optimizer identifies the stall detrimental effect and corrects the design to prevent separation and increase the torque at this off-design (for 3pt-Wb) condition. These results highlight the effectiveness of the multipoint approach to optimize the design over a broader operational envelope.

We re-analyze the airfoils in Figure 7.11 with XFoil [154] to verify stall is captured with a conventional design tool. The C_p curves in Figure 7.13(a) are consistent with ADflow results. Separation occurs when an early turbulent transition is forced upon the boundary layer (here using tripping at 10% of the chord). However, complementary results (not shown) obtained with the

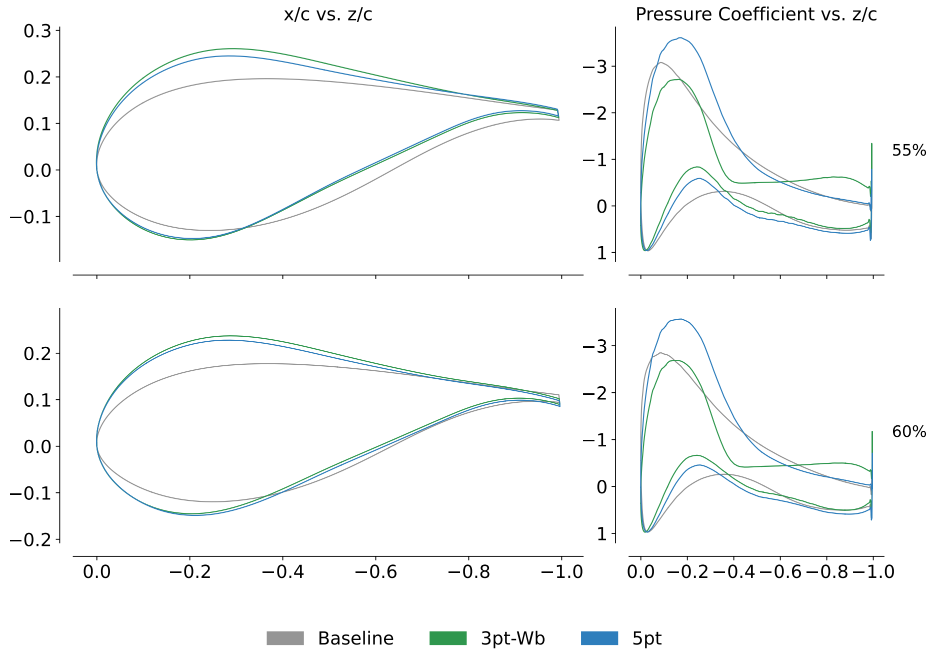

Figure 7.11: Airfoils and pressure coefficient distribution at BR-T5 for 3-point and 5-point optimized designs. The 3pt-Wb C_p highlight a large load loss due to flow separation.

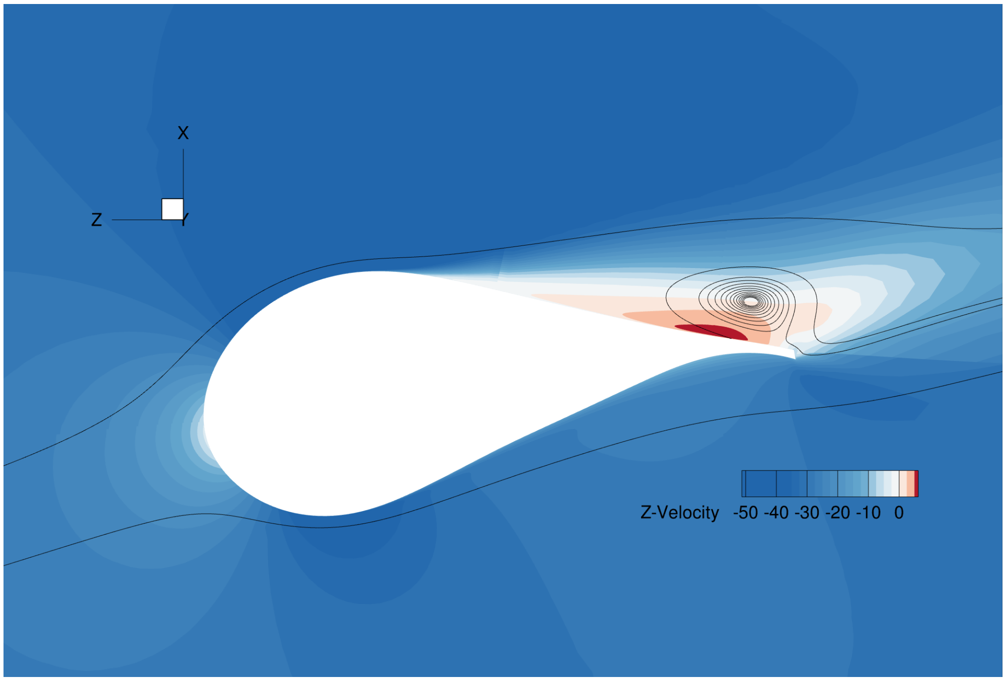

Figure 7.12: In-plane flow velocity over the blade section at the span 55% of the 3pt-Wb design at BR-T5. The streamlines highlight the flow recirculation in the aft section of the airfoil.

e^n method to compute the location of the transition indicate no stall for n as small as 1. This suggests that the fully turbulent assumption in ADflow may have caused premature stall in the BR-T5 condition. A more accurate treatment of boundary layers in RANS (such as transition models) would be necessary to optimize turbine rotors for practical applications. The C_l polar in Figure 7.13(b) shows that the optimizer effectively postpones separation onset on the 5pt design. The CFD solver does not explicitly calculate the induction factor so the operating angles of attack might be overestimated in the figure.

The goal of adding the BR-T5 condition is to obtain a design less sensitive to angle-of-attack changes over a rotor revolution or under gust loads. The turbine is not designed to specifically operate at this wind condition, so stall prediction accuracy does not affect the torque yield prediction directly.

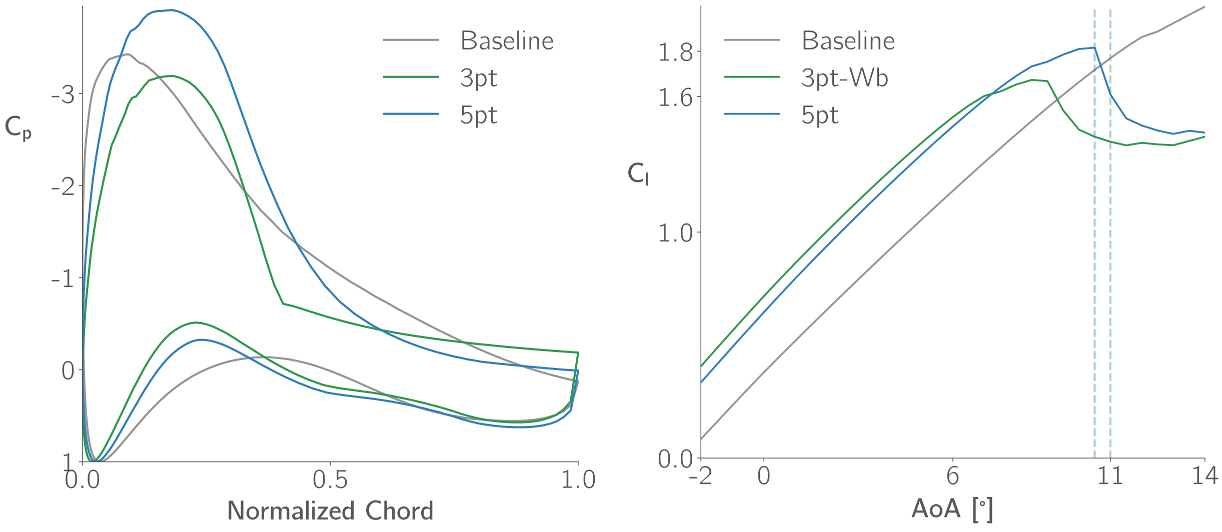

(a)

(b)

Figure 7.13: XFoil-based pressure distribution (a) and C_l polar (b) for airfoils at the span 55% at BR-T5. The vertical lines in (b) indicate the estimated operating local angle of attack.

7.6 Optimizing airfoil shape and fiber angles

Finally, in this section we combine the structural design variables from Chapter 6 with the geometrical design variables and load conditions discussed earlier in this chapter. We compare the reference 3pt-Wb case with an equivalent optimization problem that also includes “**S**ingle” fiber angle parametrization. Considering the additional 148 fiber angle variables, the resulting problem has a total of 422 structural and geometrical design variables. The constraints and objective of this case, named 3pt-WbS, are the same used for 3pt-Wb.

Table 7.9: Torque and mass variation for the optimization with shape and fiber angle variables (3pt-WbS) compared to the reference optimizations from Chapter 6 (2pt-S) and Chapter 7 (3pt-Wb).

	Torque [%]			$\sum \gamma_k \Delta Q_k$ [%]	Mass [%]
	HW	BR	MR		
2pt-S	+7.6	+9.5	-	+8.8	-9.7
3pt-Wb	+10.5	+13.2	+17.1	+14.3	-36.2
3pt-WbS	+10.3	+13.1	+16.8	+14.1	-39.0

Table 7.9 compares this new optimization case with the reference optimizations from Chapters 6 and 7, 2pt-S and 3pt-Wb respectively. 2pt-S is largely outperformed due to the small set of geometrical design variables used in that problem. Compared to 3pt-Wb, 3ptWb-S average torque increase is just 0.2% lower, without any design point showing performance gaps significantly larger than the others. However, the fiber angle tailoring of the latter case leads to a further 2.8% mass reduction – more than 2.5 t.

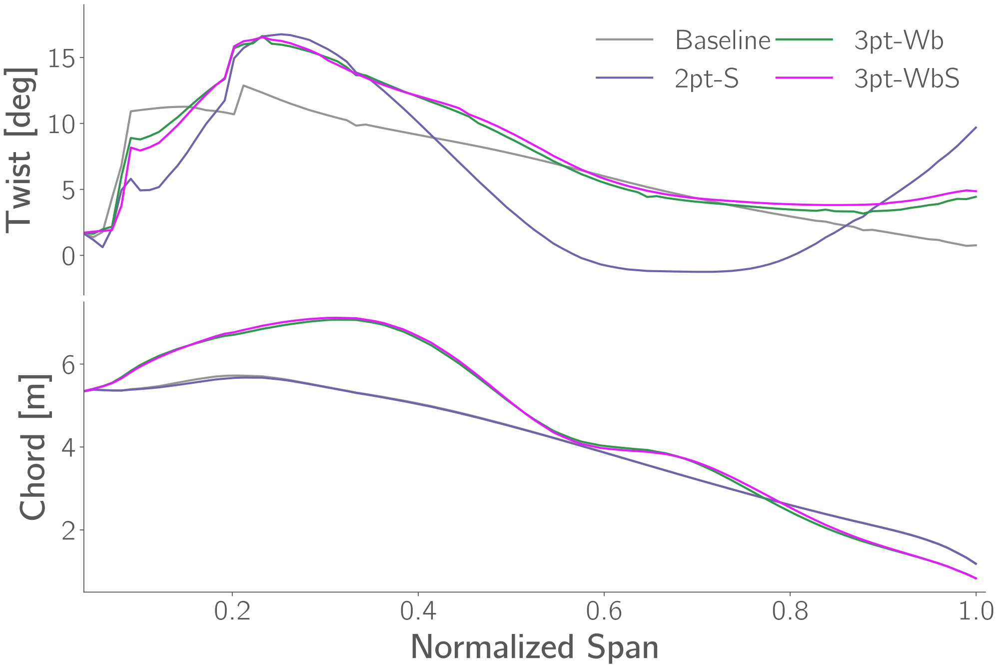

Figure 7.14: Optimized twist and chord for optimizations with different combinations of design variables.

The twist and chord trends reported in Figure 7.14 show minimal differences between 3pt-Wb and 3pt-WbS. The twist for the latter case is less than 1° higher in the outboard 25% of the span.

The chord is practically identical except for a small increase around the span 80%.

When planform and airfoil shapes are not allowed to change, seen in 2pt-S and Chapter 5, the optimizer converges to a more aggressive twist distribution to increase the aerodynamic loads. These loads have a higher peak shifted more inboard, as shown in Figure 7.15. Despite the minor planform and twist differences shown in Figure 7.14, the aerodynamic loads for 3pt-Wb and 3pt-WbS are substantially the same.

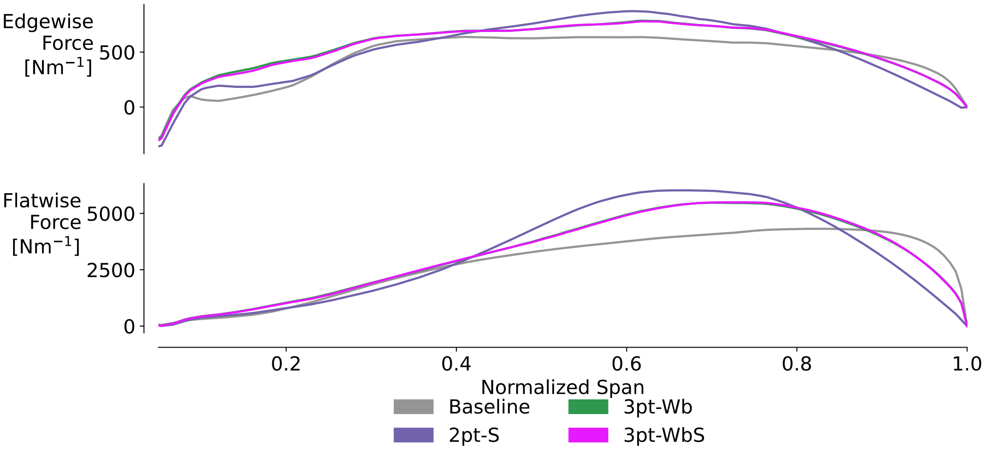

Figure 7.15: Lift distribution at BR for optimizations with different combinations of design variables.

Looking at the airfoil shapes and airfoil distributions in Figure 7.16, we observe few differences between the two 3-point optimizations. At the span 60%, 3pt-WbS is slightly thicker and flatter on the suction side. At 80% and 90% of the span, the higher twist of 3pt-WbS leads to marginally lower C_p peaks. These difference have little effect on the chordwise loads at these sections. We observe that a large increase in local angle of attack introduced by 2pt-S to leads to a C_p peak comparable to the 3-point optimization, although is shifted forward and leads to higher local loads (see Figure 7.15 at the span 60%).

Finally, the structural layouts of the different optimized layouts discussed in this sections are reported in Figure 7.17. The thickness distributions of the 3-point optimization appear almost identical. 3pt-WbS outboard half panels are less than 0.2 mm thicker than for 3pt-Wb. The largest thickness decreases, up to 5 mm, occur between the span 10% and 30% on the spars and on the leading edge panels at the span 30%.

The 3pt-WbS layout shows less aggressive fiber angle steering than the 2pt-S. The planform changes concurrently affect the aerodynamic load distribution and the local structural stiffness, so comparable stiffness can be obtained with lower fiber steering. The spar cap fibers remain once again aligned with the spanwise reference axis. Fibers are conversely rotated clockwise by almost 10° in the areas where thickness is reduced the most compared to 3pt-Wb.

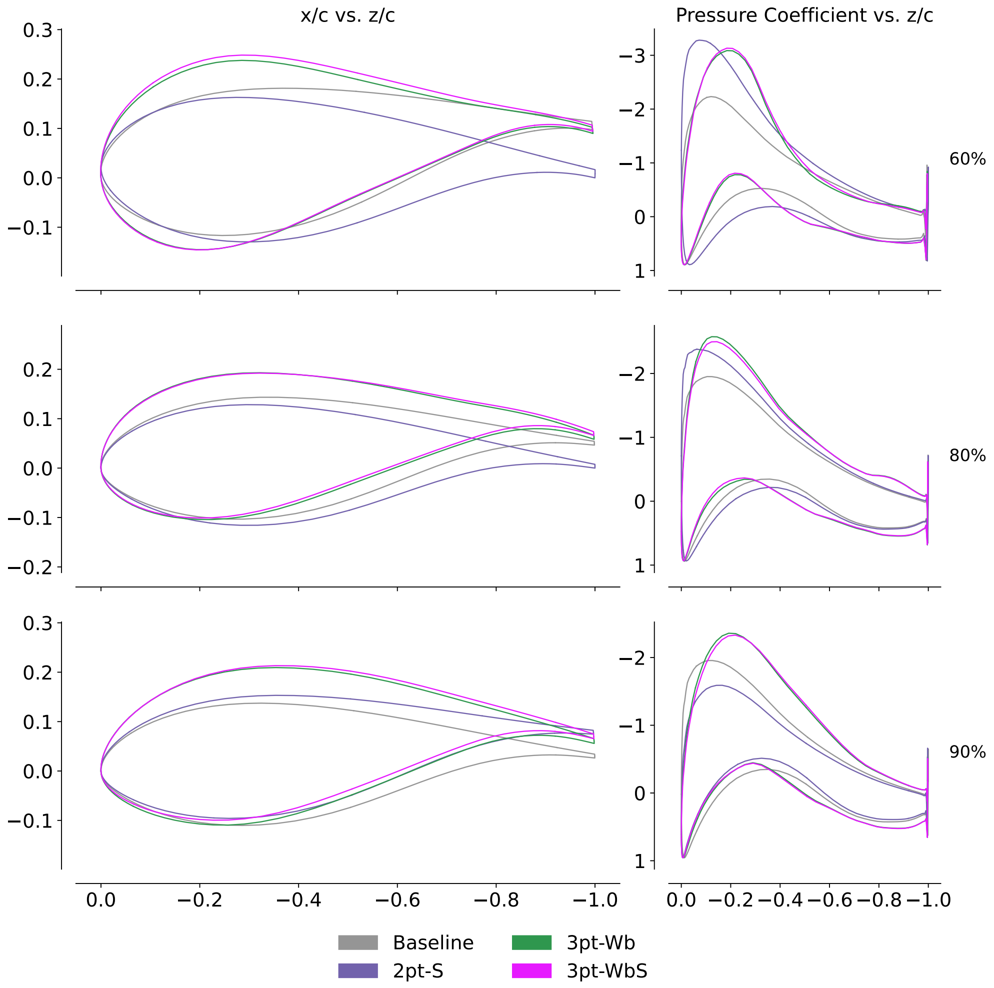

Figure 7.16: Airfoils and pressure coefficient distributions at BR for optimizations with different combinations of design variables.

Conversely to 2pt-S, there is a large area on 3pt-WbS with clockwise-rotated fibers around the span 70%. The adjacency constraints are active, limiting the rotation to slightly more than 10°. Since there are no noticeable mass gains compared to 3pt-Wb, this fiber steering must be mainly contributing the bend-twist coupling of the blade, driving the local twist angle difference shown in Figure 7.14.

Although we cannot directly extend the design trends shown in this section to more general studies, we can make some considerations to be verified in future investigations. Since 3pt-Wb and 3pt-WbS converge to practically the same outer mold line, we can conclude that there is a low chance of geometrical multimodality with this problem formulation. The same cannot be said for the structural layout, which should be investigated using more initial design points. As pointed out in Chapter 6, fiber angle tailoring enables sizable mass reductions for the same aerodynamic loads.

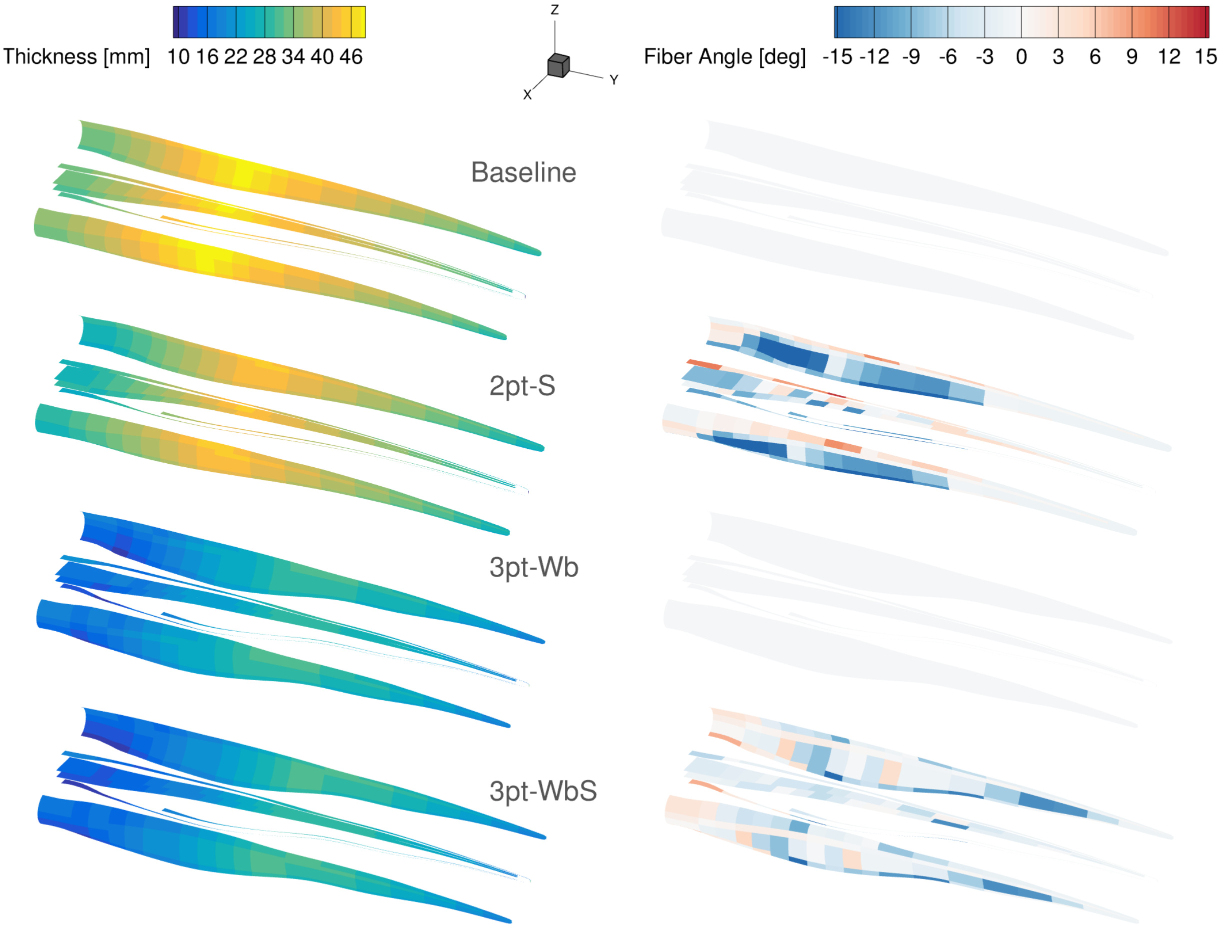

Figure 7.17: Structural thickness and fiber angle distribution for optimizations with different combinations of design variables.

Finally, we observe some loose coupling of the bend-twist characteristics of the blade with its fiber angle distribution. Future works should focus on this specific design feature to better understand how planform and structural features can be further exploited to obtain better rotor layouts.

7.7 Conclusions

In this chapter we presented a set of multipoint aerostructural wind turbine rotor optimizations using a high-fidelity framework. We use more than 250 structural and geometrical design variables, including planform and blade section shape modification. The optimized designs are up to 39% lighter and have more than 10% higher torque than the baseline rotor.

We discuss the selection of relevant wind speeds and tip speed ratios, and quantify their impact on the final layout. Considering at least 3 design points leads to less extreme designs and a more consistent performance at below-rated conditions. Different design points weights in the objective lead to differences in torque output up to 2%.

When the rotor mass has a lower weight in the objective, the mass reduction is cut in half but the torque further increases between 2.6% and 5%. The optimized loads peak closer to the blade tip, leading to a substantially different structural layout. We further discuss the driving design role of stress and displacement constraints.

Following that, we show we can prevent off-design flow separation by extending the design point stencil. The redesigned airfoil performance is verified with Xfoil. The fully turbulent boundary layer assumption triggers an early flow separation. These results highlight the necessity of including a transition model in future works.

Finally, we run a case with the same 3-point problem formulation used in this chapter and the fiber angle parametrization from Chapter 6. While the outer mold line does not show noticeable differences, the different fiber angle distribution leads to a 2.8% mass reduction for the same torque output, and a marginally improved bend-twist coupling response.

The optimizations we present are numerically tractable and can be extended to more complex studies including additional structural design variables, aerostructural constraints, and design conditions for more robust operational performance.

Chapter 8

A Mixed-Fidelity Framework for Life-Cycle Sizing

The increasingly large blade span of industrial-scale wind turbines presents design challenges beyond what has been investigated in the previous chapters of this dissertation. Steady-state performance within the operating envelope is just one of many wind systems design drivers. The higher blade flexibility and larger inertia affect the dynamic blade response and, consequently, the operation and lifespan of the rotor and other stressed components. The increased loads and size affect the amplitude and frequency of loading cycles, which also impact the operating life of the system. Modern design approaches must include holistic simulation capabilities accounting for multidisciplinary and multi-component system interactions.

One of the main shortcomings of the high-fidelity optimization approach discussed in Chapters 5 to 7 is the lack of life-cycle sizing considerations. Practical rotor designs need to account for the whole operating conditions and life of the turbine. The aerostructural simulations must account for this multidisciplinary interaction, and how environmental conditions exacerbate the risk of premature failures. At the same time, design tools must retain the complexity and design space exploration capabilities such as the high-fidelity MDAO strategy discussed in earlier chapters. Life-cycle considerations must be included in MDAO frameworks to obtain practical system layouts.

The International Energy Commission defines standards for wind turbine certification that encompass a wide range of scenarios [25]. Modern turbines must be designed considering a set of representative DLC. These DLCs include a range of extreme but statistically relevant load cases scenarios. The turbine must be sized with respect to emergency shutdowns and once-in-50-years hurricane-level wind gusts before the product can be safely deployed. Fatigue loads must also be accounted for, to ensure the system longevity and minimize maintenance downtime over the multi-decade expected operating life. These requirements are intended to guarantee the safety and integrity of the system over its planned operating life. Standards are updated over time to consider additional scenarios and keep the pace of technological developments. Analysis and design approaches must be scalable to accommodate for an ever increasing number of design points to be included in the design loop. This applies particularly to future floating systems, with additional combinations of wind and sea state that would become prohibitive to be simulated at every design iteration.

To address these design concerns and extend MACH capabilities, in this chapter we present the development and preliminary results of a mixed-fidelity framework combining MACH with a conventional turbine simulation tool. The goal is to perform high-fidelity gradient-based MDO with life-cycle sizing considerations. In practice, we use conventional tools to enforce extreme and fatigue loads constraints within the high-fidelity loop using a staggered approach.

We build on the work presented in Chapter 5. Even the high-fidelity multipoint optimization studies in Chapter 7 are limited to 5 concurrent design points. The approach hereby proposed includes a larger number of time-accurate inflow scenarios, using the approach proposed by Ingersoll and Ning [62] to integrate them within a gradient-based optimization loop. Two different coupling approaches are presented. A first, load-scaling-based approach [87] uses additional high-fidelity structural analyses with OpenFAST-informed aerodynamic loads to account for fatigue and extreme loads. A second, more consistent and robust implementation, yet unpublished, is discussed later in this chapter. Moving on from some modeling limitations of the load-scaling approach, we use a beam condensation tool, the aggregated load output from OpenFAST, and analytical strain calculations directly within the MACH loop. Results for both approaches are preliminary but showcase the potential of mixed-fidelity MDAO, leveraging high-fidelity capabilities while including life-cycle considerations. This work outlines a promising and little explored research thrust that could bridge the gap between academic tools and industrial applications, fostering the use of high-fidelity MDAO tools in the wind energy industry.

The work presented in this chapter is the result of a collaboration with Brigham Young University's FLOW Lab [1], who conceptualized the high-level framework and led the implementation effort. Early results of this work have been presented by Caprace et al. [87], building on a previous low-fidelity implementation [62]. My support focused on the integration effort with MACH, further constraint conceptualization, and support in effectively setting up and running the optimization cases.

In Section 8.1 we discuss the numerical treatment of fatigue damage and extreme loads. Their implementation in MACH is detailed in Section 8.2. Finally, preliminary results using this novel mixed-fidelity framework are presented in Section 8.3.

8.1 Assessing life cycle loading on a turbine

Section 1.3.3 discusses relevant previous optimization studies that include fatigue and or extreme load constraints. In this section we briefly introduce how we estimate extreme and fatigue loads

---

[1] FLOW Lab web page

from short-time OpenFAST simulations. The extrapolation process for once-in-50-years extreme loads is discussed in Section 8.1.1. Fatigue life estimation is a vast and challenging topic with implications on safety, costs, and reliability of mechanical systems. We redirect to more extensive reviews by Veers [155] and Ravikumar et al. [156] for an exhaustive treatment of the subject. Our fatigue damage estimation approach is discussed in Section 8.1.2.

8.1.1 Extreme loads

We run each relevant DLC (listed later in Section 8.2.1) for 600 seconds of time-accurate simulation time. Both the expected fatigue damage and extreme loads are respectively aggregated and extrapolated from these short-term time series, using the methods proposed by Moriarty et al. [157].

An extreme load is defined as an extreme event with a once-every-50-years occurrence. Short-time loads are sorted by their intensity to build a binned probability distribution. Every "bin" counts the occurrences of loads with a given intensity over the OpenFAST time simulations - normalized by the total number of load samples. The distributions from multiple operating conditions for a given DLCs are aggregated using a weighted sum, where the sum weights are based on the probability of a certain inflow condition occurring over the turbine life time. We do not include fault and maintenance down time in this stochastic estimation. An example of resulting probability density distribution (PDF) for in-plane (F_t) and out-of-plane (F_n) loads is illustrated in Figure 8.1.

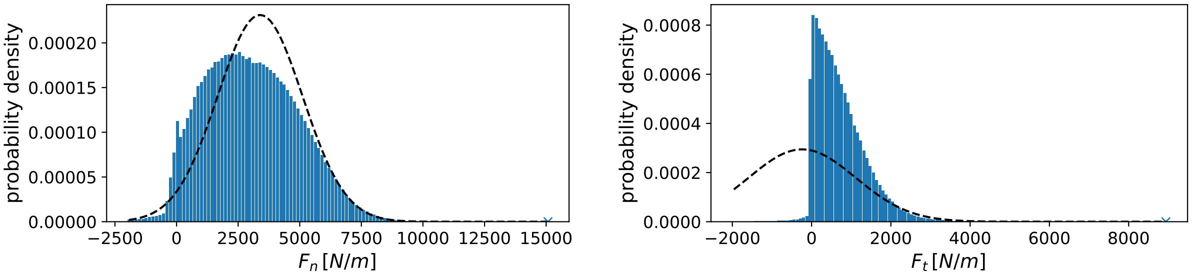

(a) Normal load

(b) Tangential load

Figure 8.1: Probability density function of mid-span aerodynamic loads measured. We fit the tail of the distribution with a normal distribution (dashed line). From Caprace et al. [87].

A Gaussian distribution is fit on the binned PDF to extrapolate the probability of extreme loads beyond the duration of our time-resolved analyses. A least-square approach is used to minimize the error and tune the Gaussian mean and variance. We use the logarithm of the probability of exceedance (1 - PDF), to match the tail of the distribution, as shown in Figure 8.1.

The probability of an extreme load is then given by:

$$P\left(L > L^{\mathrm{EXTR}}\right) = 1 - \frac{dt}{50\,\text{year}}, \tag{8.1}$$

with $dt/(50\,\text{year})$ the ratio of the simulation sampling time (defining the binned loads) over 50 years. An example of the fit with the extrapolated extreme load L^{EXTR} is shown in Figure 8.2.

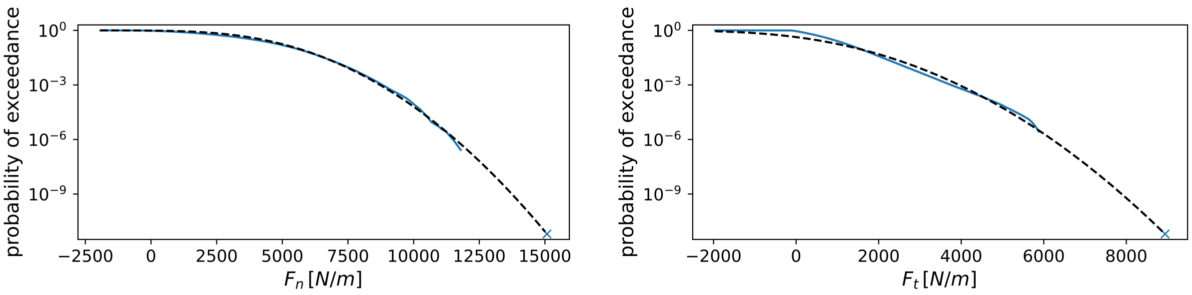

(a) Normal load

(b) Tangential load

Figure 8.2: Mid-span aerodynamic load probability of exceedance. The extrapolated L^{EXTR} is highlighted by the "×" symbol. From Caprace et al. [87].

8.1.2 Fatigue loads

The turbine structure undergoes cyclical loading due to periodical wind speed velocity variations, wind turbulent conditions, and the cyclical gravitational loads on the rotating blade. At nominal operating conditions, these loads are well below the material ultimate yield strength. Nevertheless, thousands of load cycles weaken the structure until the mechanical components fail at conditions well below nominal failure loads. This phenomenon, in a material science context, is called fatigue.

Fatigue is a complex scenario to model in numerical simulations, since a "deterministic" model would require sub-scale crack modeling. We instead rely on practical guidelines from the MLife manual [158] to estimate fatigue damage on our rotor structure. We design the rotor for 20 years of operations.

The physics of fatigue

Over their (usual) 20-year life span, turbines are exposed to a number of aerodynamic and gravitational alternating loads. A breakdown example of chordwise bending loads from Hau [5] is shown in Figure 8.3.

From a physical point of view, the cyclic loading of mechanical components triggers small crack propagation over time, leading to structural weakening and ultimately failure. Thus, all the

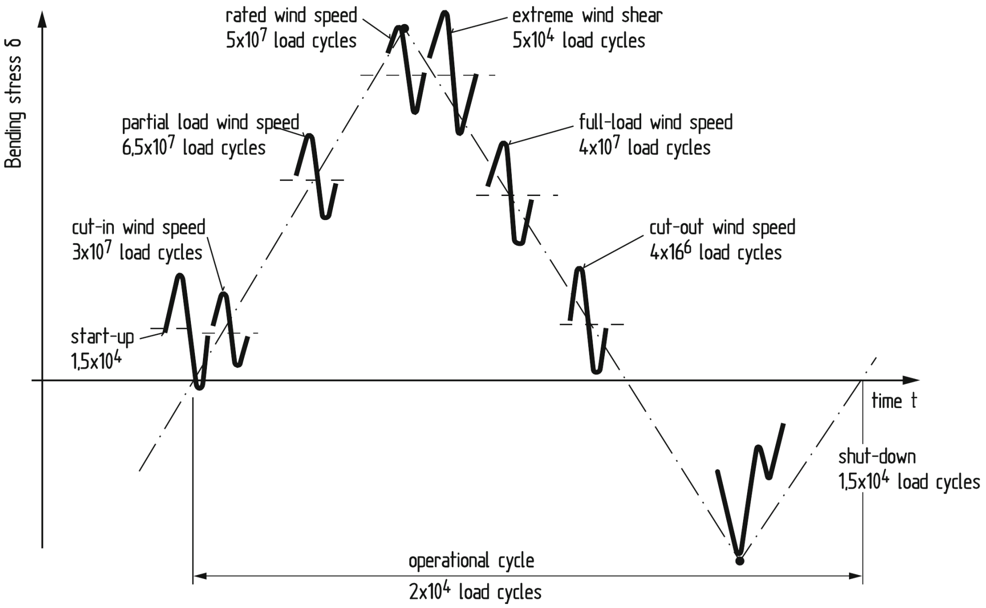

Figure 8.3: Example of fatigue load occurrence over a turbine lifetime, from Hau [5, Fig. 6.23].

turbine components exposed to cyclical loads are exposed to fatigue-induced failure over their lifetime. The turbine tower, drivetrain, and the blades in particular are exposed to large load oscillations that can shorten their operating life. Under regular operations and start-up/shut-down procedures, the blade weight has a major influence on chordwise bending moments, while wind loads primarily affect out-of-plane loading. [5] Broader operational and site-specific considerations come into play for lifetime cycle estimations. Fatigue constraints usually drive the rotor structural sizing [5, Table 6.21]

Despite it being traditionally observed on metals, a wide range of materials including carbon- and glass-fiber composite materials is affected by fatigue. Local material imperfections or stress concentrations initiate small fractures, which progressively expand driven by both elastic and plastic local stress response. Fracture mechanics modeling goes beyond the scope of this work, and is hardly practical to be implemented in coupled aerostructural optimization studies. Nevertheless, fatigue life estimation strategies (some of which have been mentioned in Section 1.3.3) have been implemented in numerical design studies.

From experimental or numerical observations, the load cycles can be sorted by mean stress value σ_m and range σ_r, which is twice the cycle amplitude $\sigma_r = 2\sigma_a$ as shown in Figure 8.4.

The fatigue life of a given material is usually characterized using an S-N curve, or Wöhler curve, that correlates the load average and range of a constant-in-time load cycle to the number of cycles expected to failure. When accounting for the full load spectrum, different load cycles contribute to fatigue damage accumulation. The power-law equation of an S-N curve like the one

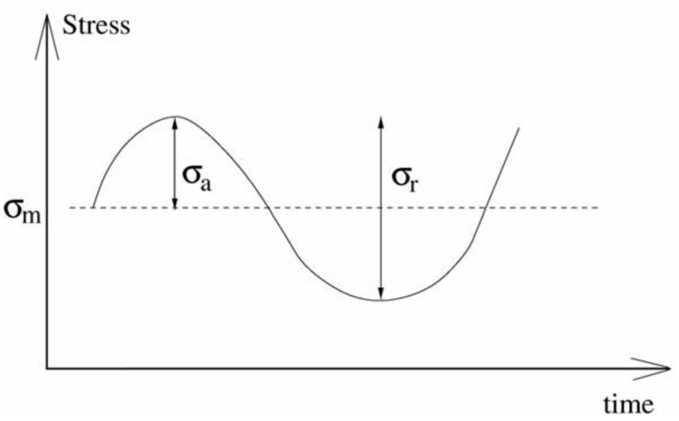

Figure 8.4: Example of oscillatory loads, from Hansen [34, Figure 15.1].

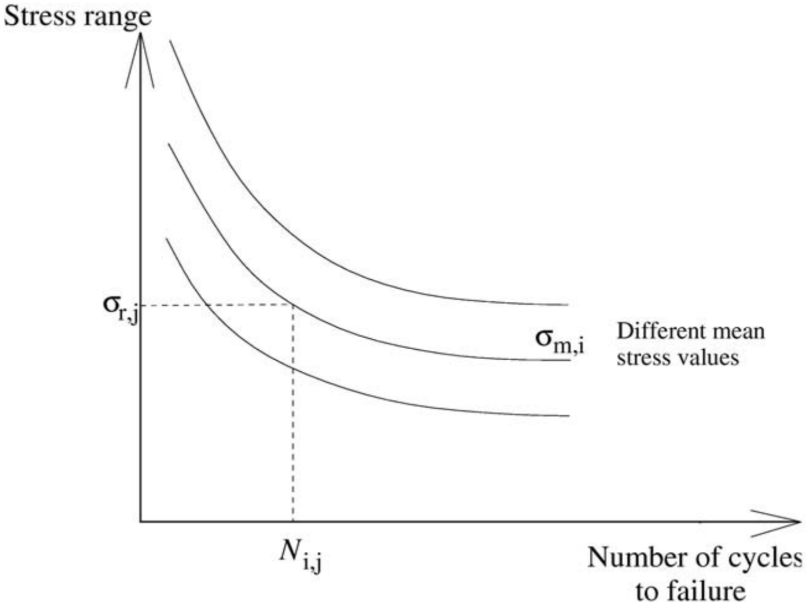

Figure 8.5: Examples of S-N fatigue curves, from Hansen [34, Figure 15.2].

shown in Figure 8.5 defines the number of cycles to failure N_{ij} for the i^{th} cycle in the j^{th} time series as:

$$N_{ij} = \left(\frac{L^{\text{ult}} - |L^M_{i,j}|}{\frac{1}{2} L^R_{i,j}} \right)^m , \tag{8.2}$$

where L^{ult} is the ultimate load of the material and $L^R_{i,j}$ and $L^M_{i,j}$ are the measured load range and average load respectively (standing for the general σ_r and σ_a in Figure 8.4).

Numerical treatment

The fatigue constraints we implement are based on the estimation of damage-equivalent loads (DEL) over the blade components. DELs are static loads which produce the same damage as the fatigue accumulated over the life time.

We decompose the load signal from each OpenFAST time-accurate simulation into a sequence of load cycles, using the Rainflow counting algorithm. Equation 8.2 is used to estimate the cycles-to-failures for each load cycle. We set the Wöhler coefficient to $m = 10$. This value reflects properties of conventional composite materials rather than the simplified isotropic model (Section 4.1.1.1) used for this study.

The contribution of each load cycle for a time series adds linearly using the Palmgren–Miner's rule:

$$D_j = \left(\sum_i \frac{1}{N_{i,j}} \right). \tag{8.3}$$

The damage over the design lifetime D is obtained through a weighted average of the time-series damage:

$$D = \sum_j f_j D_j = \sum_j f_j \left(\sum_i \frac{1}{N_{i,j}} \right) = \left(\frac{L^{\text{ult}}}{\text{DEL}^{\text{life}}} \right)^{-m}, \tag{8.4}$$

The weighting factor f_j scales the short-term damage up to operating life using:

$$f_j = \frac{p_j T^{\text{life}}}{T_j}. \tag{8.5}$$

where T^{life}/T_j is the ratio of the simulation time over the operating life and p_j is the probability of occurrence of a given time series loading. Equation 8.4 introduces the definition of lifetime damage DEL^{life}. When we use the strain-based "condensation" approach (discussed later in Section 8.2.3), we know the ultimate load L^{ult}. In this case, we can apply the Goodman correction assuming an equivalent load range $L_{i,j}^{R0}$ for each load cycle so that:

$$N_{ij} = \left(\frac{L^{\text{ult}}}{\frac{1}{2} L_{i,j}^{R0}} \right)^m \tag{8.6}$$

so that the final expression for lifetime DEL is given by:

$$\text{DEL}^{\text{life}} = \frac{1}{2} \left(\sum_j f_j \left(\sum_i (L_{ij}^{R0})^m \right) \right)^{1/m}. \tag{8.7}$$

8.2 Integrating fatigue and extreme loads in a monolithic MDO architecture

Time-accurate CFD simulations are still numerically untractable for practical large scale MDO applications. The mixed-fidelity approach we propose leverages the lower computational cost of unsteady BEMT simulations to predict the response extreme loads and operating oscillatory cycles for life estimation. At the same time, we retain the benefits of blade-resolved high-fidelity analysis codes for power estimation and detailed design manipulation.

In the studies presented in this chapter, we use the same high-fidelity setup from Chapter 5. We consider a single design point - BR - but the approach could be extended to include multiple high-fidelity design points as done in Chapters 6 and 7. The BR condition is representative of power production at below-rated conditions [15]. OpenFAST simulations are used to evaluate the unsteady response of the rotor on a consistent low-fidelity model. These analyses are used to evaluate the turbine rotor response against extreme and fatigue loads. Combining the analyses from both tools, we perform a first-of-its-kind fatigue-compliant high-fidelity optimization.

The framework we propose enables concurrent optimization of structural, twist, chord, and thickness design variables. However, for demonstration purposes, the optimizations presented in Section 8.3 only use structural design variables (and pitch, in one case) at this stage.

The mixed-fidelity framework high-level architecture is illustrated in Figure 8.6. The aerostructural kernel in the right block is the MACH framework as presented in Chapter 2. The framework setup used in Chapter 5 is extended to include two additional analysis components for yield at extreme load conditions and fatigue damage accumulation. These two components are formulated differently depending on the fidelity coupling approach.

The load-scaling approach discussed in Section 8.2.2 uses two additional instances of TACS. Static analyses using appropriately scaled aerodynamic loads are used as a proxy for extreme and fatigue damage equivalent loads.

The new approach described in Section 8.2.3 uses a newly implemented analysis tool called Beamstrain. This component wraps another NREL-developed tool, called PreComp [159], implemented in Julia by the FLOW Lab. PreComp provides span-variant beam structural properties for composite blades, to be used to set up OpenFAST analysis codes. In practice, the code returns a set of spanwise blade stiffness properties given the rotor internal topology and ply material properties. The version of PreComp we use provides analytically accurate derivatives of the stiffness properties with respect to the inputs. We calculate structural strains for given OpenFAST output loads with analytical formulas, whose derivatives can be obtained by hand. This way, we can integrate Beamstrain within the high-fidelity gradient-based framework, effectively coupling multiple fidelities within the same optimization loop.

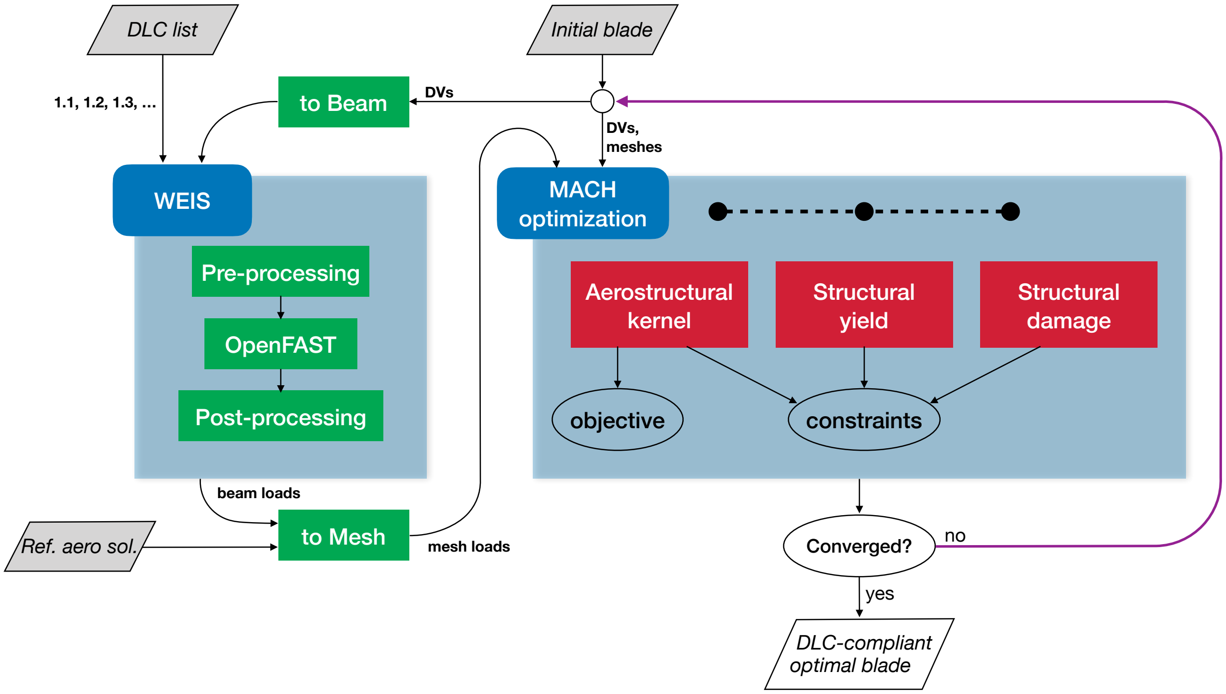

Figure 8.6: Mixed-fidelity framework architecture. From Caprace et al. [87].

Regardless of the approach used for constraint evaluation, OpenFAST time-accurate analyses are performed outside the optimization loop. One of the limitations of OpenFAST is the lack of accurate derivatives for gradient-based optimization. Moreover, despite the lower computational cost compared to CFD analyses, re-evaluating the full set of DLC at every optimization iteration would significantly lengthen the design process. To circumvent these implementation challenges, we used a staggered approach, where OpenFAST is run sequentially with the main optimization loop.

We initially provide a baseline design and the set of DLC for extreme and fatigue load evaluation (top left corner in Figure 8.6). The blade structural information is translated into a coupled BEMT-beam model, analyzed by OpenFAST(wrapped by WEIS). The output is then passed to the MACH-based optimization to enable life-cycle constraints evaluation. Different integration approaches require different data from OpenFAST, as discussed later in Section 8.2.2 and Section 8.2.3 respectively. The additional components are arranged in a multipoint formulation, similar to what done in Chapters 6 and 7. In this case, we have a single aerostructural case, used for torque, tip displacement, and detailed mass evaluation.

In both constraint formulations, the extreme and fatigue loads are frozen during the optimization loop. This is a limitation of the staggered approach we propose, since we cannot capture the feedback effect from the optimization on the OpenFAST-based loads. The optimizer modifies the structure without accounting for the impact of such changes on the rotor blade dynamic response. For this reason, once the optimization is completed, we re-run OpenFAST analyses using the optimized rotor layout. The newly obtained extreme and fatigue loads are then passed back to MACH for another optimization run. We refer to this information back-and-forth between MACH and OpenFAST as the optimization "outer loop" or "outer iteration". This procedure is repeated several times, until the loads and optimized design do not change between consecutive "outer loop" iterations.

The resulting design is potentially sub-optimal in a strictly monolithic MDO sense. Nevertheless, it is a fatigue-compliant rotor design optimized using high-fidelity analysis tools. Ingersoll and Ning [62] demonstrated on a low-fidelity framework that this approach converges in less than 10 "outer iterations" to the same solution as a fully coupled optimization. The fatigue and extreme loads are marginally affected by further structural design changes, so the frozen-load assumption is not a critical limitation. Our implementation avoids the need to compute derivatives of time accurate analysis, at the cost of needing a sequence of "outer loop" iterations to converge the final design.

8.2.1 BEMT-based analysis with WEIS/OpenFast

We rely on existing conventional analysis tools to run time-accurate turbine simulations under a range of different design load conditions (see Section 8.1). OpenFAST [160][2], developed by the National Renewable Energy Laboratory (NREL), represents the state-of-the-art for wind energy simulations. It is capable to perform aero-hydro-servo-elastic simulations of a holistic wind turbine model, both for land-based and floating offshore systems. A dynamic BEMT model is used for rotor aerodynamic simulations. The aerodynamic module is coupled with a beam structural model to simulate the blade dynamic response under a large range of unsteady inflow conditions. The structural properties are obtained combining information on the blade topology and classical laminate theory for material properties. More details on how we transfer the characteristics of a 3D FEM structural model into its beam representation are discussed later in Section 8.2.3.

We use WEIS[3], also developed by NREL, to setup and run our OpenFAST analyses. WEIS is a Python wrapper based on OpenMDAO [40]. It integrates OpenFAST and other utility software, such as ROSCO [4] controller, TurbSim [5] for turbulent wind simulation, and pCrunch [6] for load analysis and aggregation. A YAML-based ontology [161] is used to univocally define the turbine

---

[2] OpenFAST GitHub repository, v2.6.0.

[3] Wind Energy with Integrated Servo-control (WEIS) GitHub repository, v0.2.

[4] ROSCO GitHub repository.

[5] TurbSim GitHub repository.

[6] pCrunch GitHub repository.

components and analysis scenarios. We use WEIS to generate the inflow conditions, tune the ROSCO controller, and postprocess the analysis results

WEIS can use pyOptSparse as OpenMDAO driver for optimizations. In this work, we only run OpenFAST analyses outside the optimization loop and use pyOptSparse as integrated with MACH. Since we are focused on rotor design, we remove tower and generator degrees of freedom from the model. Future works could extend this approach to other components and include, for example, tower sizing in the design loop.

As mentioned in Section 1.3.3, considering the full set of DLC from the IEC [25] standard during coupled optimizations is hardly practical. For demonstration purposes, we select a subset of DLCs, namely DLC 1.2 and DLC 1.3, listed in Table 8.1.

DLC	Inflow velocities [m/s]	N seeds	Constraint
1.2	5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25	6	Fatigue
1.3	5, 7, 9, 11, 13, 15, 17, 19, 21, 23, 25	6	Extreme

Table 8.1: Load selection for OpenFAST time-accurate simulations.

At every “outer iteration” in Figure 8.6, each DLC is run at 11 wind speeds with 6 turbulent seeds each, for a total of 132 simulations runs. Future studies could extend the analyses to include more DLCs or define a reduced design load base as proposed by Pavese et al. [83] to reduce the computational burden and wall time.

8.2.2 The load-scaling approach

In this first approach, presented by Caprace et al. [87], we translate the ultimate and fatigue loads extracted from OpenFAST (see Section 8.1) into 3D equivalent aerodynamic loads to be applied on a blade-resolved TACS [117] FEM model. Following the definition of ultimate and fatigue equivalent loads, enforcing failure constraints under these conditions acts as a proxy for extreme- and fatigue-load compliant structural sizing.

Using a high-fidelity model requires little implementation effort in MACH. Two additional TACS instances, with fixed equivalent ultimate and fatigue loads respectively, are added to the setup used for Chapter 5 in a multipoint fashion. Design variables are modified consistently in each model, at each optimization iteration. The aerostructural model and the two additional structural models are evaluated in parallel, with the latter ones driving the overall sizing and thickness distribution layout.

However, transferring load information from the low-fidelity OpenFAST output into an input for a blade-resolved simulation requires intermediate processing steps. The radial load distribution from the beam model must be extrapolated along the chordwise direction, similarly to what proposed by Bottasso et al. [80]. To reconstruct equivalent 3D aerodynamic loads over the blade airfoils, we scale a reference aerodynamic load using the information extracted from the extrapolated or lifetime-aggregated aerodynamic loads. The ADflow surface solution for the baseline rotor provides the reference pressure distribution $\boldsymbol{\sigma}^{\text{ref}}(r, s)$, with r the spanwise location and s the airfoil curvilinear coordinate. The in-plane and out-of-plane components of this load distribution are defined as:

$$F_n^{\text{ref}}(r) = \int_{\partial s} \left(\boldsymbol{\sigma}^{\text{ref}}(r, s) \cdot \hat{\mathrm{e}}_n\right)\, ds, \tag{8.8}$$

$$F_t^{\text{ref}}(r) = \int_{\partial s} \left(\boldsymbol{\sigma}^{\text{ref}}(r, s) \cdot \hat{\mathrm{e}}_t\right)\, ds, \tag{8.9}$$

with ∂s the airfoil contour, and $\hat{\mathrm{e}}_t, \hat{\mathrm{e}}_n$ the in- and out-of-plane normal vectors. The integrated spanwise and surface loads are shown in Figure 8.7(a) and Figure 8.7(b) respectively.

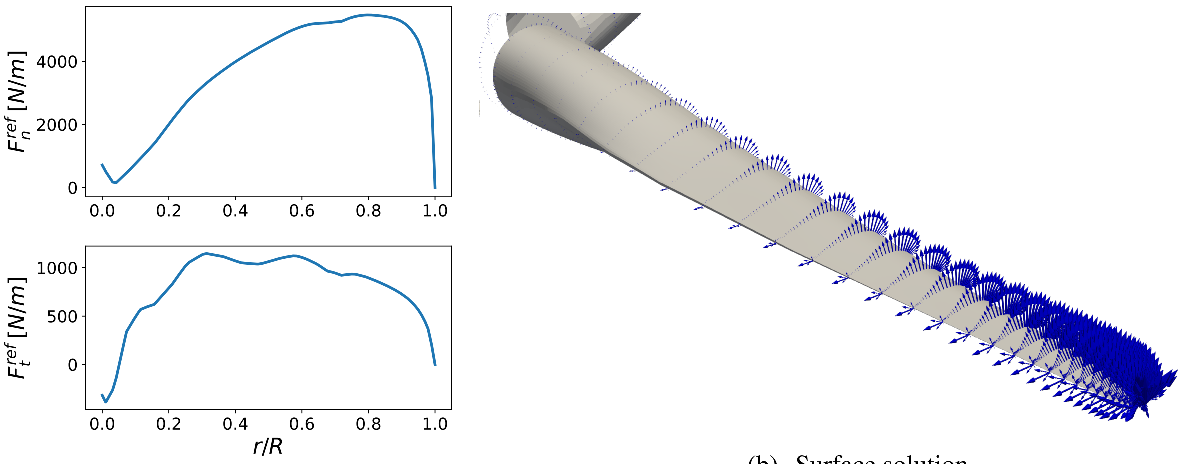

(a) Spanwise loading

(b) Surface solution

Figure 8.7: Aerodynamic loads on the reference solution. From Caprace et al. [87].

The spanwise target loads obtained from OpenFAST data postprocessing, $F_t^{\text{target}}(r)$ and $F_n^{\text{target}}(r)$ respectively, are shown as dashed lines in Figure 8.8(a). We introduce a radial scaling factor $\lambda(r) = F_n^{\text{target}}(r)/F_n^{\text{ref}}(r)$ to extrapolate the scaled-up surface loads to be used during the optimization. The surface loads shown in Figure 8.8(b) are given by:

$$\boldsymbol{\sigma}(r, s) = \lambda(r)\, \boldsymbol{\sigma}^{\text{ref}}(r, s). \tag{8.10}$$

We use a single $\lambda(r)$ scalar scaling factor at every spanwise station. As $\lambda(r)$ is defined using the ratio of target and reference loads, the out-of-plane radial load distribution matches exactly

$F_n^{\text{target}}(r)$ - as shown at the top of Figure 8.8(a). However, applying the same factor to the tangential loads:

$$F_t(r) = \lambda(r) \int_{\partial s} \left(\boldsymbol{\sigma}^{\mathbf{ref}}(r,s) \cdot \hat{\mathbf{e}}_t\right)\, ds, \tag{8.11}$$

leads to an over estimation of the tangential loads (see bottom Figure 8.8(a)). For sake of consistency, we maintain the single factor scaling approach rather than decoupling the loads. This also includes an additional inherent safety factor in the procedure. Moreover, the in-plane loads are lower in magnitude, while the out-of-plane loads drive the structural sizing.

The preliminary results presented in Section 8.3 do not use geometrical design variables aside pitch, so the blade outer mold line is fixed during the optimization. If the shape changed, the reference aerodynamic solution would be updated using the latest available surface load set from MACH. The extrapolated aerodynamic solution would then be calculated again using Eq. 8.10. The coefficient $\lambda(r)$ is fixed during the MACH-based optimization and is only updated after an "outer iteration".

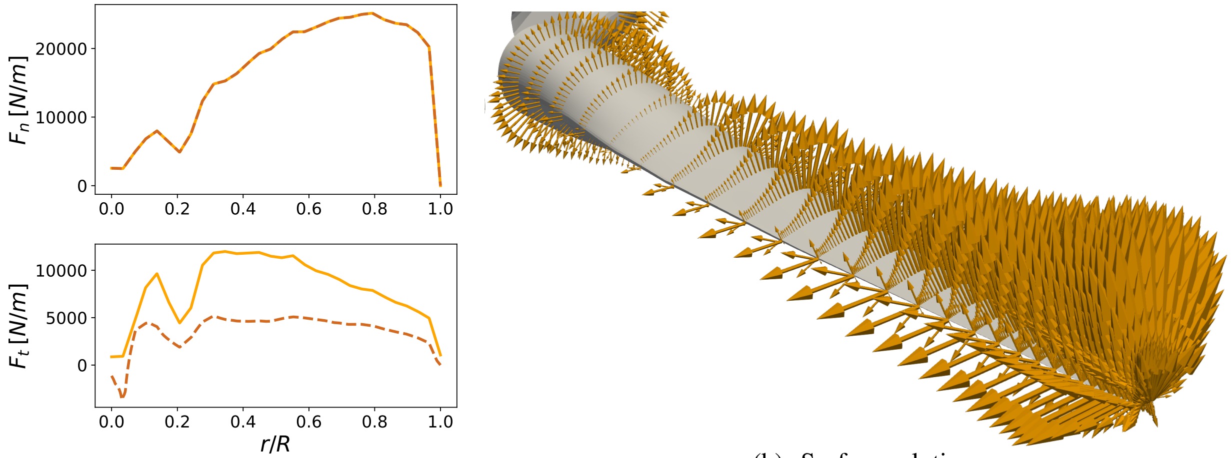

(a) Spanwise loading

(b) Surface solution

Figure 8.8: Scaled aerodynamic loads (continuous line). The target loading is indicated by the dashed line. From Caprace et al. [87].

An XDSM representation of this framework implementation is reported in Figure C.1.

8.2.3 The condensation approach

The "condensation" approach uses a more conventional treatment of extreme and fatigue loads. The damages extracted from OpenFAST simulations with the procedure outlined in Section 8.1.1 and Section 8.1.2 are passed directly to a beam model.

The local strain over the blade at a position (x, y) in the beam cross-section, given the input in-plane bending moment M_x, out-of-plane bending moment M_y, and radial load F_z, is given by the classical beam theory formula:

$$\epsilon(r) = -\left(\frac{M_x(r)}{E(r)I_{xx}(r)}y - \frac{M_y(r)}{E(r)I_{yy}(r)}x + \frac{F_z(r)}{E(r)A(r)}\right), \tag{8.12}$$

with E the cross section Young's modulus, I_{ii} the mass moment of inertia along the i axis, and A the cross sectional area.

If the extreme loads form OpenFAST are given as input to Eq. 8.12, we obtain the extreme strain $\epsilon(r) = \epsilon^{\text{EXTR}}(r)$. Conversely, if we use the fatigue loads in Eq. 8.12, we obtain a $\epsilon(r) = \epsilon^{\text{life}}(r)$ lifetime-damage equivalent strain. Assuming that the DEL load average is zero (i.e. load cycles have the same semi-amplitude and are centered around zero), using Eq. 8.2 we obtain the number of cycles to failure N^{life} as:

$$N^{\text{life}} = \left(\frac{\epsilon_{\max}}{\epsilon^{\text{life}}}\right)^m, \tag{8.13}$$

An XDSM representation of the fatigue damage integration proposed in this section is reported in Figure C.2.

Beamstrain

A newly implemented wrapper for PreComp, named Beamstrain, has been developed by FLOW Lab collaborators to calculate the damage equivalent strain for a given set of loads and blade properties. The tool takes as input the geometrical and structural design variables, the blade material properties, and the aggregated spanwise moments and loads from OpenFAST, and returns the blade damage equivalent strain. It uses the same KS constraint aggregation used in TACS, aggregating the yield $\epsilon^{EXTR}/\epsilon_{\max}$ or damage $\epsilon^{life}/\epsilon_{\max}$. This constraint aggregation was not used by Ingersoll and Ning [62].

Beamstrain wraps a version of PreComp written in Julia. The advantage of the Julia implementation is its native support for automatic differentiation. Total derivatives of the sectional stiffness and inertia are obtained by propagating design variables perturbation seeds in forward mode. Due to the low computational cost of a single PreComp evaluation - especially compared to the high fidelity solver - we neglect scalability issues in favor of ease of implementation. PreComp returns the sensitivities of the y/EI_{xx}, z/EI_{yy}, $1/EA$ with respect to the full design variables. The formula in equation 8.12 is hand derived using the chain rule to obtain the total strain sensitivity with respect to the full set of design variables. The bending moments M_x, M_y, and radial load F_z effectively act as weights for the sectional properties partial derivatives in the total sensitivity calculation.

8.2.4 Constraint definition in MACH

The different approaches outlined in Section 8.2.2 and Section 8.2.3 respectively require a different constraint definition within the optimization.

Extreme load constraint

When using the “load scaling” approach (Section 8.2.2), the extrapolated loads illustrated in Figure 8.8 are applied directly on the blade-resolved finite element model. Gravity and centrifugal forces are added as separate load sets. We use the same formulation for the failure constraint described in Section 5.1.3 - namely a von Mises stress evaluation, aggregated using the KS formulation [124]. The constraint Y can then be generally formulated as:

$$Y = \eta_Y \, \mathcal{F}(\mathbf{x}, F_n^{\mathrm{EXTR}}, F_t^{\mathrm{EXTR}}, \mathbf{\Omega}) \leq 1, \tag{8.14}$$

with x the design variables, F_n^{EXTR} and F_t^{EXTR} the normal and tangential loads respectively, and $\mathbf{\Omega}$ is the rotational speed used to calculate centrifugal loads. Conventional approaches use a linear relationship between stress and strain on the beam models. Evaluating this constraint on the high-fidelity structural model relaxes this assumption and accommodates multi-axial loads, leading to theoretically higher accuracy of stress prediction. The main assumption of this approach relies on the aerodynamic load extrapolation and aggregation discussed in Section 8.2.2.

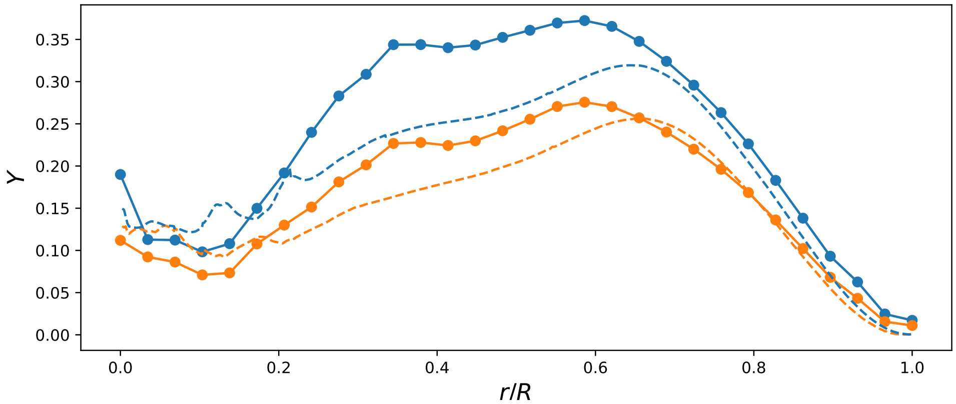

Figure 8.9: Extreme yield failure on the spar caps. Loads from the “condensation” method from Sec. 8.2.3 have solid lines (pressure side: (–●–), suction side: (–●–)), while loads from the “load-scaling” approach from Sec. 8.2.2 have dashed lines (pressure side: - - -, suction side: - - -). From Caprace et al. [87].

Conversely, the extreme load constraint for the “condensation” approach outlined in Section 8.2.3 is defined as:

$$Y = \frac{\eta_Y \, \epsilon^{\mathrm{EXTR}}}{\epsilon_{\max}} \leq 1, \tag{8.15}$$

Where $\epsilon_{\max}$ is the material yield strain and ϵ^{EXTR} is the radial strain obtained inserting the OpenFAST extreme extrapolated load into Eq. 8.12. A safety factor of $\eta_Y = 1.35$ is used for both approaches.

We compare the radial extreme load distribution using both approaches in Figure 8.9 for a wind velocity of 9 $\mathrm{m/s}$ and normal [25] turbulent conditions. Load trends, especially over the blade outermost 30%, are consistent. The absolute yield constraint values between the span 20% and 60% show discrepancies up to 13% of the maximum allowable yield on the blade pressure side. The "condensation" approach appears more conservative in this scenario, but further assumptions need to be considered when the model will be extended to include composite material properties.

Damage constraint

The canonical damage expression (discussed more in detail below) is essentially a failure criteria elevated to the Wöhler coefficient m:

$$D = \left(\eta_D \, \mathcal{F}(\mathbf{x}, F_n^{\mathrm{life}}, F_t^{\mathrm{life}}, \mathbf{g}, \boldsymbol{\Omega})\right)^m \leq 1. \tag{8.16}$$

As explained in Section 8.1.2, the damage equivalent loads in the out-of-plane and in-plane direction, F_n^{life} and F_t^{life} respectively, are static loads intended to generate the same damage as the lifetime cyclical loads that the turbine encounters in its operating life. This approach resembles the experimental procedure proposed by Musial et al. [162], where a blade prototype was tested under static damage equivalent loads. The assumptions on the damage aggregation coming from the Miner's rule and the material life model are the most relevant limitation of this approach [163]. However, using time-accurate or time-spectral dynamic simulations currently remains unpractical for MDO applications.

On the other hand, the "condensation" approach circumvents some of these assumptions and more simply combines Eq. 8.4 with Eq. 8.13 to obtain the equivalent damage:

$$D = \left(\frac{\eta_D \, \epsilon^{\mathrm{life}}}{\epsilon_{\max}}\right)^m \leq 1, \tag{8.17}$$

with safety factor $\eta_D = 1.35$ applied directly to the loads for a larger safety constraint.

Similarly to what shown in Figure 8.9, Figure 8.10 compares the spanwise damage distributions from the two different methods for the same inflow condition. We run OpenFAST for a 10 minutes simulation and then we scale the damage up to represent a year of operations. We observe the same trends discussed for the extreme yield constraint, except for the "load-scaling" approach on the suction side, which now predicts a higher damage than the conventional approach. These discrepancies might originate from the different way aerodynamic and gravitational loads are treated. More verification efforts would be required to use the “load-scaling” approach on large scale optimizations, to ensure the optimizer does not take advantage of model limitations under certain load scenarios. These limitations also motivate the use of the condensation method over the load-scaling approach, as it works with fewer assumptions.

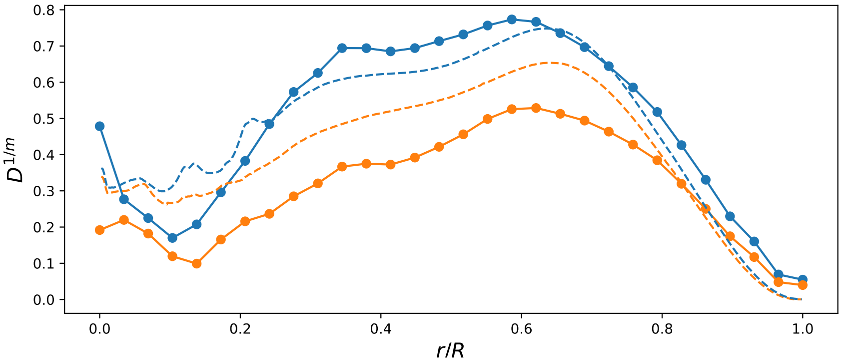

Figure 8.10: Fatigue damage failure on the spar caps. Loads from the “condensation” method from Sec. 8.2.3 have solid lines (pressure side: (—●—), suction side: (—●—)), while loads from the “load-scaling” approach from Sec. 8.2.2 have dashed lines (pressure side: - - -, suction side: - - -). From Caprace et al. [87].

8.3 Preliminary optimization with extreme loads and fatigue constraints

In this section we summarize some preliminary results using both approaches from Section 8.2.2 and Section 8.2.3. While limitations still exists in the implementation of the two different models, these studies are the first example of mixed-fidelity optimizations combining a high-fidelity framework with conventional analysis tools for life-cycle sizing considerations.

In Section 8.3.1, we first test the “load-scaling” approach on a purely structural optimization problem, to demonstrate the framework functioning and the mixed-fidelity optimization convergence on the cheapest possible problem. Following that, we extend the mass minimization problem to include a torque constraint, evaluated using the tightly coupled aerostructural model in Section 8.3.2. Finally, we introduce the first results for the “condensation” approach in Section 8.3.3

8.3.1 Structural mass minimization using the load scaling approach

We first minimize the mass of the rotor, starting from a uniform thickness distribution for the blade skins and spars. We use the sole 117 structural design variables presented in Figure 4.1. Figure 8.11 shows the convergence history of the first "inner iteration" - namely the first optimization with MACH. The optimization converges in 265 iterations.

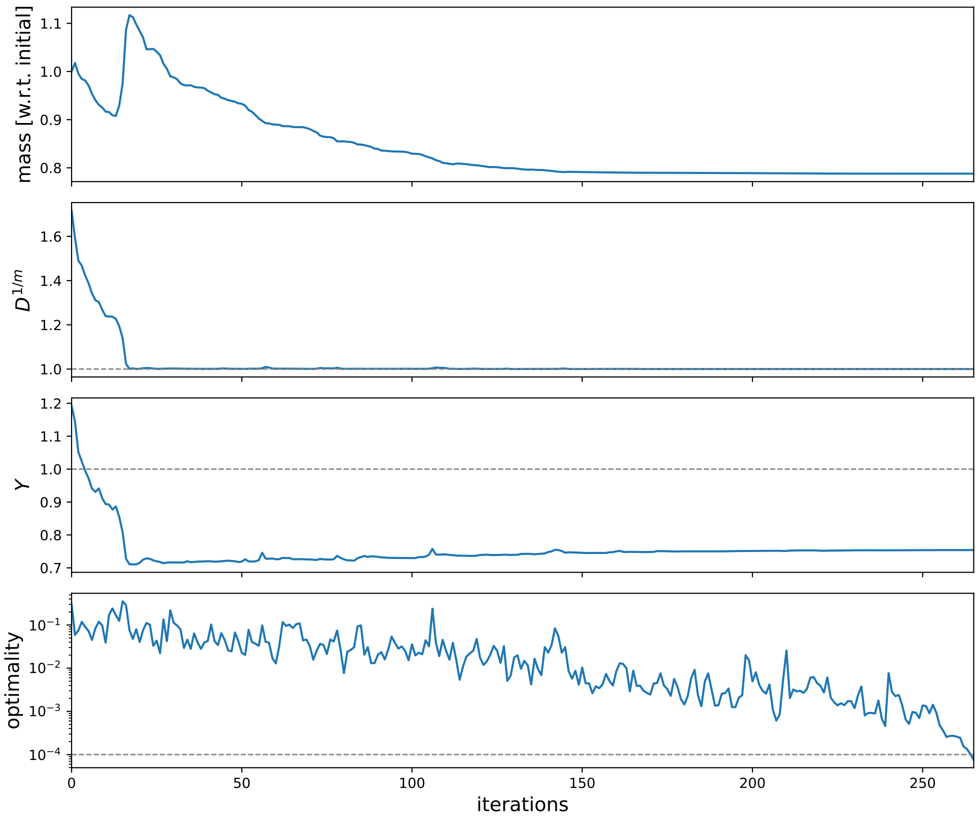

Figure 8.11: Combined-fidelity structural optimization convergence history, from Caprace et al. [87].

The total mass initially increases in order to satisfy the damage constraint. Following that, the optimizer shaves the structural thickness to obtain a design 20% lighter than the initial layout. In this load scenario, only the damage constraint is active. The structure has an approximately 30% margin on the extreme yield constraint. The spars and spar caps are the most critically loaded blade components.

The DLC selection (see Table 8.1) has a major role in the constraint definition. This optimization demonstrates the effective implementation of the load transfer between OpenFAST and MACH. However, more practical design considerations require a larger set of DLCs as prescribed by the international standards [25].

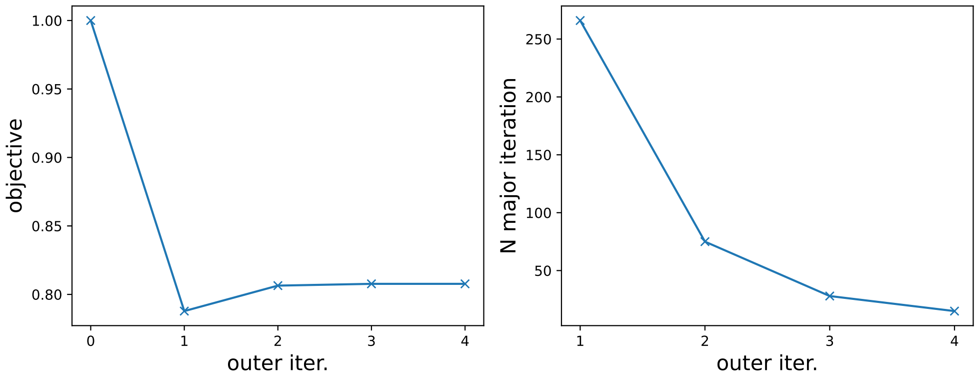

Figure 8.12: “Outer loop” convergence of normalized blade mass variation (left), and “inner loop” optimization iterations (right). From Caprace et al. [87].

The full optimization converges in 4 “outer iterations”, as shown in Figure 8.12. The first MACH “inner iteration” discussed above in Figure 8.11 is the most expensive. The second step requires less than 100 iterations to converge, and the final sub-optimization takes less than 25 iterations. Each following optimization starts from the previous iteration optimized layout, using updated loads from OpenFAST. The design at the end of the first “inner iteration” is lighter than the optimum of the full mixed-fidelity optimization. The resulting structure is thus undersized, and the damage constraint is violated at the beginning of the second “inner iteration”. This highlights the need to iteratively update the damage and extreme loads through additional OpenFAST simulations, instead of relying on a single MACH optimization with frozen loads.

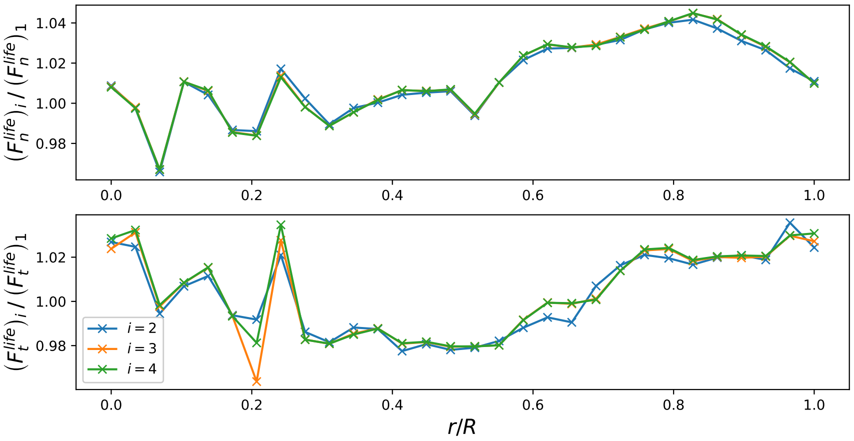

Figure 8.13: Normalized (w.r.t. $i = 1$) damage-equivalent out-of-plane (top) and in-plane (bottom). i refers to the “outer iteration”. From Caprace et al. [87].

Nevertheless, the fatigue loads are only marginally affected by the structural design updates.

The blade response and deflection is different between iterations, but with this problem formulation there are essentially no changes in the outer mold line of the blades. The damage equivalent load changes by less than 4% between the first and second "outer iteration" and then quickly converges, as shown in Figure 8.13.

The final thickness layout is shown in Figure 8.14. Compared to the initial 15 mm thickness, the optimizer increases the thickness on the the pressure side cap up to 95 mm. Despite the assumptions and limitations of this optimization case, this design feature is consistent with conventional practical designs where the shear web bears the highest loads.

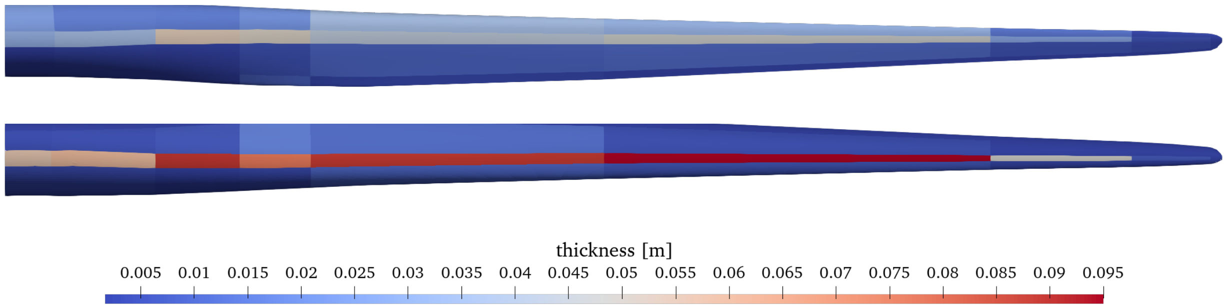

Figure 8.14: Optimized thickness distribution on the blade-resolved FEM model, on the suction side (top) and pressure side (bottom). From Caprace et al. [87].

Finally, Figure 8.15 breaks down the computational cost of the mixed-fidelity structural optimization discussed in this section. The optimization has been performed on a node of the Marilou9 HPC from the BYU Fulton Supercomputer Lab using processors with Intel Broadwell 2.4 GHz architecture.

The total time for WEIS includes pre- and post-processing steps on top of the OpenFAST time-accurate simulations, executed in serial for each DLC. The number of DLC turbulent seeds and simulation real-time drives the computational cost of the conventional tools.

The high-fidelity optimization has the largest cost by far at the first "outer iteration" due to the number of optimization iterations (reported in Figure 8.11). This cost drastically reduces over subsequent iterations due to the shorter high-fidelity optimizations steps. The cost for the WEIS steps is practically fixed, aside for the computation of turbulent inflow conditions that is performed only once in the first "outer iteration". These inflow boundary conditions are unchanged over the full mixed-fidelity optimization.

8.3.2 Aerostructural optimization using the load scaling approach

We use the structural layout discussed in the previous section as the initial design for a new optimization problem that includes an instance of the tightly coupled aerostructural solver. The cheap

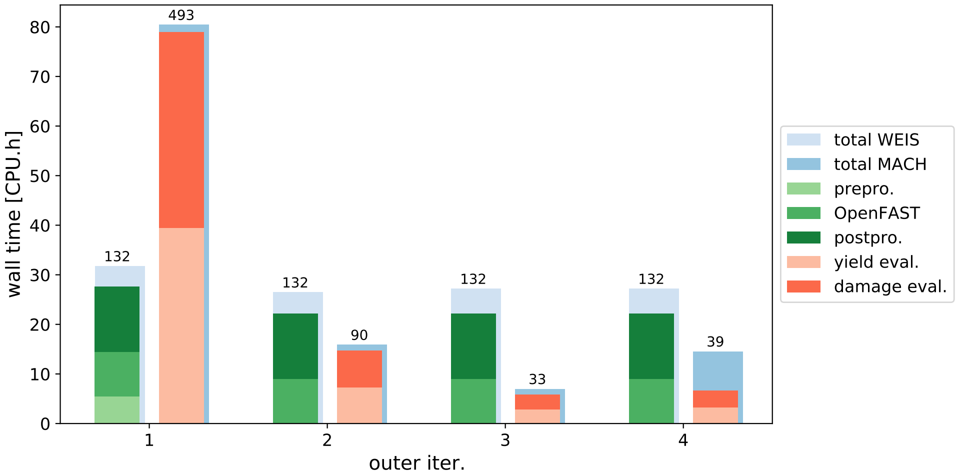

Figure 8.15: Optimization computational cost breakdown. We dissect pre- and post-processing times in addition to OpenFAST runs for WEIS. The structural solution cost for the additional TACS instances (see Figure 8.6) is decoupled from the aerostructural analysis for MACH. The numbers on top of the columns indicate the number of function evaluations. From Caprace et al. [87].

but low-accuracy L3 aerodynamic mesh is used for this demonstration.

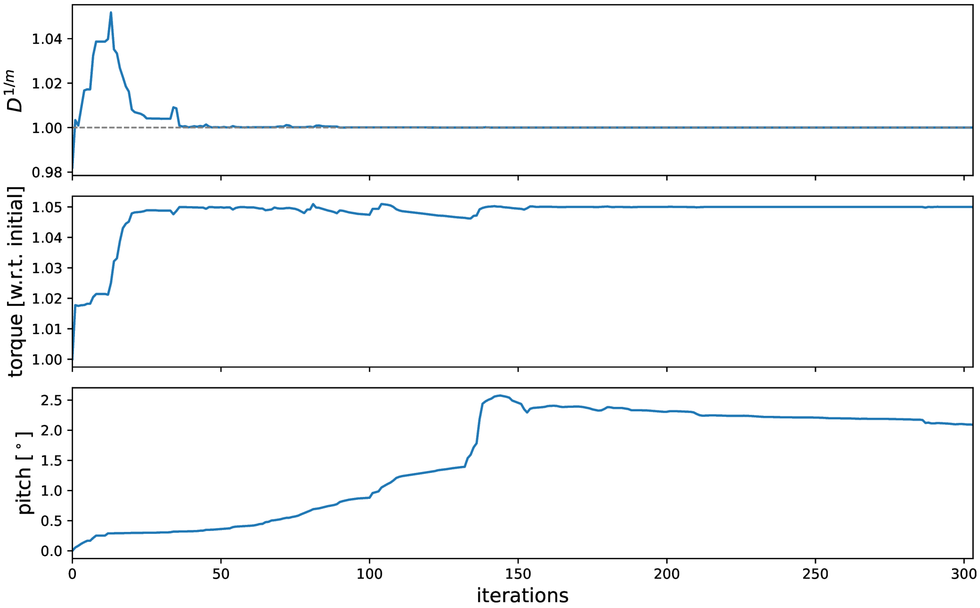

Figure 8.16: Aerostructural optimization convergence history, from Caprace et al. [87].

We use the aerostructural analysis to enforce the final torque to be 5% higher than baseline at the BR condition. To achieve this, we also include a geometrical design variable to adjust the pitch of the blade. This simplified formulation demonstrates the effective coupling of the full MACH optimization capabilities with the life-cycle constraints based on the "load-scaling" approach. The optimizer simultaneously satisfies the load constraints while increasing the torque output at nominal operating condition, opening the door to future more complex optimizations.

The optimization convergence history in Figure 8.16 shows that the optimizer increases the pitch by approximately 2 ° on the optimized design. The fatigue damage constraint is satisfied through structural layout tailoring as seen in the previous section. The torque output is affected by both geometrical and structural design variables. The torque is initially increased largely by a change in the structural thickness distribution and torsional stiffness, since the blade pitch is below 1 ° for the first 100 iterations. The constraint is satisfied by iteration 25, but at that point the damage constraint is violated. The optimizer then simultaneously updates the structural thickness and pitch value to maintain the torque increase, while satisfying the damage constraint. Constraint and design variable scaling affect the sequence of these steps during the optimization. Nevertheless, optimization in Figure 8.16 proves the successful implementation of the "load scaling" approach.

8.3.3 Structural mass minimization using the condensation approach

The "condensation" approach outlined in Section 8.2.3 is under active development at the time of writing this dissertation. In this section we present preliminary results to demonstrate the implementation of Beamstrain within MACH. We present the results of three different optimizations with three different sets of active constraints. For demonstration purposes, we limit the mixed-fidelity optimization to a single "outer iteration". We use different inflow conditions and initial thickness in this more recent case, so the optimizations illustrated in this sections are not directly comparable with the cases from Caprace et al. [87].

Once again, we perform a mass minimization study using structural design variables. Conversely to Section 8.3.1, the aerostructural analysis is included in the optimization even if there are no geometrical design variables (similarly to what done at the end of Section 5.2). Due to the absence of geometrical design variables and the shape of the design space with this new formulation, we do not include a torque constraint. Fatigue and extreme load constraints drive the design of a heavier and stiffer blade than what obtained in Chapter 5. Since the life-cycle loads are fixed, there are no concerns of the optimizer exploiting excessive operating load reductions to further reduce the rotor mass. The torque at BR conditions on the final designs shown in this section differ by a maximum of 2%

For the three cases in Table 8.2, we enforce either only the extreme constraint (Case 1), the extreme and fatigue damage constraint (case 2), and finally add a tip displacement constraint (Case 3). This latter case demonstrates the effective coupling of MACH and Beamstrain in the mixed-fidelity framework, as both tools concurrently affect the final optimized design.

Table 8.2: Total mass and tip displacement for rotors optimized with the "condensation" approach and different constraint sets.

Case	Active Constraints	Mass [t]	d_{tip} [m]
1	Extreme	93.6	4.83
2	Fatigue	160.1	2.97
3	Fatigue Displacement	167.6	2.80

For this specific set of DLC and rotor configuration, the extreme load constraint is the least "demanding" from a structural point of view. When both extreme load and fatigue damage constraints are enforced (case 2), only the latter is active at the optimum. The final structure for this design problem is more than 66 t (71%) higher than case 1. We also observe a tip displacement reduction of 38% between case 1 and 2 due to the overall higher structural thickness of the latter. Finally, in Case 3 we use the same setup of Case 2 with an additional displacement constraint - imposing a maximum tip deflection of 2.8 m. The final mass for this design is more than 79% higher than case 1 and 4.7% higher than case 2.

The thickness and stress distribution under BR condition for the high-fidelity model are illustrated in Figure 8.17. The stresses at the reference operating condition (BR), as expected, become less and less relevant as the average structural thickness increases from case 1 to 3. For case 3, stress at BR peaks at less than 25% of the critical stress around the span 65%. Case 1 shows the highest local stresses, concentrated in the root region and along the skin-spar junction on the pressure side. This layout exceeds the local failure limit using the prescribed safety factor of 4.6 at BR condition. This indicates that under these assumptions and formulation the high-fidelity stress constraint would also be active at the optimum for case 1.

The optimized layouts show different thickness distributions from previous optimization studies. In the first place, despite the blade upper and lower skin thicknesses being decoupled (conversely to what done in Chapters 6 and 7), the beam model drives the design to comparable thickness distributions over the two blade skins. This might be due to the beam model mainly increasing the thickness of the spar caps, which in turn affects the other skin panels due to the enforced adjacency constraints.

Case 1 has a thin root section (corresponding to the highest stress peaks) and a thicker outboard section. Thickness values range between and 4.7 mm and 12.7 mm

The active fatigue damage constraint in case 2 leads to higher thickness all over the span.

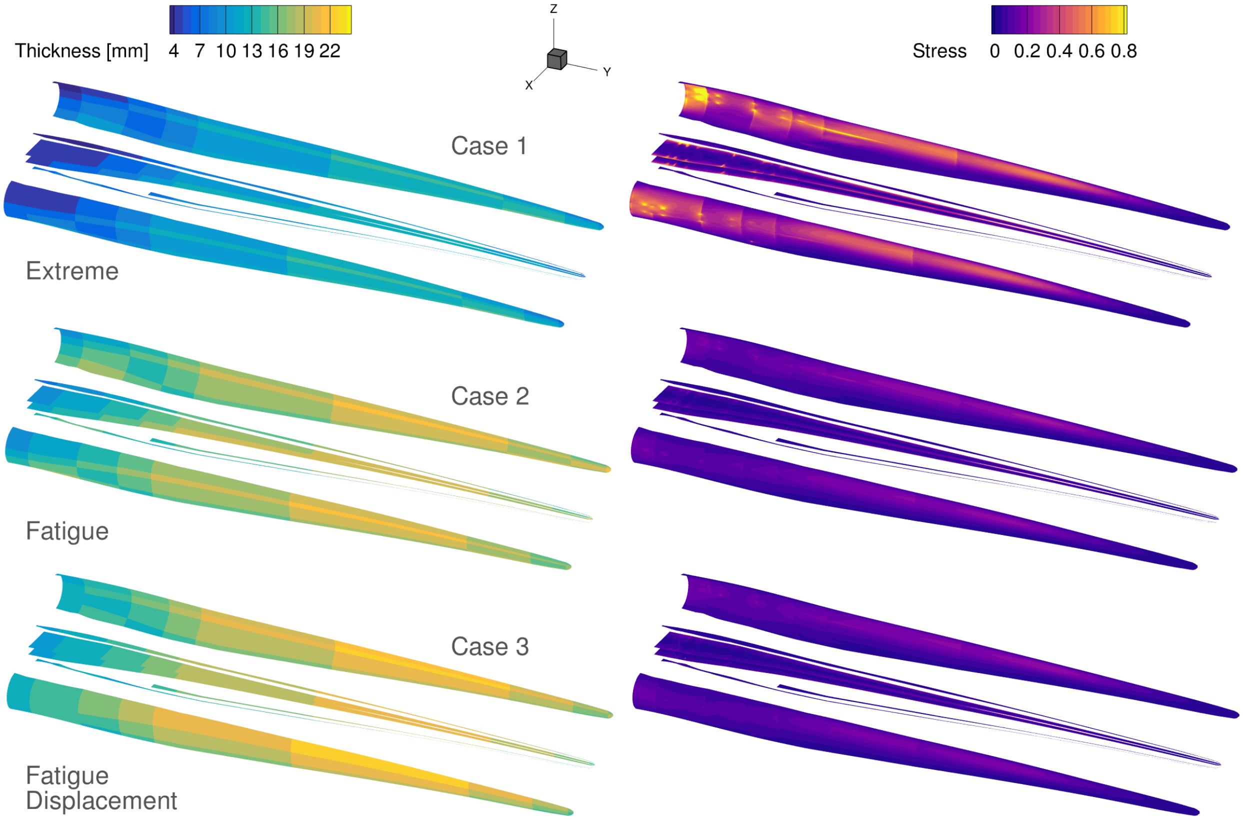

Figure 8.17: Thickness and stress distributions for the rotors in Table 8.2

The spanwise thickness trends however are comparable to case 1. Both case 2 and case 3 have the thinnest panels in the root section, with a minimum thickness of 9.4 mm and 10.1 mm respectively. The spar cap is consistently the thickest component, going up to 21.4 mm between the span 50% and 80%.

Adding a displacement constraint in case 3 leads to a small increase in maximum thickness in the same section, now 22.2 mm thick. However, the adjacent leading and trailing edge panels are thicker than for case 2 to increase out-of-plane bending stiffness and thus satisfy the displacement constraint. This occurs on a spanwise location where the thin airfoil shape privileges aerodynamic performance over the sectional structural stiffness. Conversely to the previous cases, the leading edge panel thickness is not driven by the spar cap and adjacency constraints.

Despite the results presented in this section are preliminary, they demonstrate the successful coupling of the two tools in our mixed-fidelity tool using the newly implemented “condensation” approach.

8.4 Conclusions

In this section we presented two different approaches to implement life-cycle sizing constraints within a high-fidelity optimization loop. We combine simulations from OpenFAST with MACH to obtain a fatigue- and extreme-load-compliant rotor layout optimized with a coupled CFD-CSM analysis tool. The optimizer can simultaneously increase the performance at normal operating conditions and satisfy the newly added constraints.

We first briefly introduced the physical phenomena of fatigue and the conventional numerical treatment of extreme wind scenarios and operational oscillatory loads. We extrapolate these damage equivalent loads from BEMT-based time accurate simulations. Then, we discuss the integration of these extreme load and damage evaluations in MACH.

A first approach transfers scaled static aerodynamic loads from OpenFAST to a 3D structural FEM model. A stress constraint under extrapolated load and aggregated damage is used as a proxy for extreme yield and fatigue failures. We use a spanwise scalar coefficient to consistently scale up the baseline aerodynamic load distribution and obtain steady state DEL.

The second approach directly uses the conventionally aggregated life-cycle loads on a beam model consistent with the 3D model used in MACH. We wrapped pre-existing tools in MACH to ensure the structural variables are consistently updated during the optimization loop. This approach relaxes certain assumptions on the load and model consistency when using the first "load-scaling" approach. However, this introduces more challenges in ensuring the properties of the 3D and 2D structural model are consistent.

Finally, we present three simplified optimization studies to demonstrate the successful implementation of the two approaches. For the "load scaling" approach, we discuss the results of a structural and an aerostructural optimization problem. In the first case, we show the rapid convergence of our proposed staggered approach. The design trends on the optimized layout are far from being practical, but are consistent with engineering intuition. In the second case, we highlight how the optimizer can simultaneously satisfy aerodynamic and damage fatigue constraints by manipulating the structural thickness distribution and the blade pitch.

The third and final optimization study compares three optimizations using the "condensation" approach using three different sets of active constraints. We demonstrate how both MACH and the new tool, named Beamstrain, are active simultaneously at the optimum. This opens the door for future optimization studies that include planform geometrical design variables.

The mixed-fidelity framework presented in this chapter is at a relatively early development stage. More verification efforts are needed to ensure the consistency of the life-cycle constraints between the simplified problems presented above. Moreover, we need to implement the full set of design variables available in MACH into Beamstrain, and verify the derivatives accuracy. More complex optimization studies could combine the large-scale optimization formulation presented in Chapters 6 and 7 with the life-cycle sizing constraints outlined in this chapter.

Chapter 9

Conclusions and Recommendations for Future Work

This dissertation proposes a first extensive set of high-fidelity aerostructural optimization studies for the design of wind turbine rotors. I use a coupled CFD-CSM solver with and efficient sensitivity solver to perform large-scale gradient-based optimizations. This approach leverages the enhanced accuracy and modeling capabilities of high-fidelity approaches at a tractable computational cost. The optimized layouts I obtained are lighter and more efficient over a range of operating wind conditions.

The approach I presented exploits a larger design freedom than conventional design methodologies. Using a blade-resolved model, hundreds of structural and geometrical design variables are optimized simultaneously. The tight coupling and monolithic optimization framework enables the optimizer to effectively capture aerostructural tradeoffs that would otherwise be unaccounted for. I alter the section shape, planform, structural thickness, and composite fiber angle of the rotor to increase the output torque and reduce the system mass.

The analysis tools I use, despite their cost and limitations, provide an unprecedented level of insight on the optimization process and outcome. This high-fidelity approach however cannot replace lower fidelity holistic tools, but rather complement their modeling capabilities within the design process. I demonstrate this by coupling the high-fidelity framework with conventional analysis software to include life-cycle sizing considerations in the optimization.

In Chapter 1 I outlined the broader context of this work and its potential implications. I stressed on the urgency of a rapid decarbonization and how novel design strategies, like the one discussed here, can contribute to the cause. After a brief outline of the physics, history, and components of wind turbine systems, I presented the most relevant wind turbine design optimization works from scientific literature, identifying a research niche for this work.

Chapter 2 summarizes the details of the existing high-fidelity computational framework used for the optimization studies presented in this dissertation. One of the challenges of this work was to adapt the MACH framework to optimize wind turbine rotors rather than aircraft wings. I discuss the aerostructural model verification and implementation in Chapter 3.

Despite the higher accuracy and model resolution, high-fidelity cannot capture the full operating range of a wind turbine. I dissect our optimization approach in Chapter 4, discussing how the optimization problem can be formulated to obtain practical and meaningful design insights.

In Chapter 5 I presented the first aerostructural optimization study of a wind turbine rotor using a coupled CFD-CSM solver. I optimize rotor mass, torque output, or a combination of the two by modifying the structural thickness distribution and the blade outer mold line shape. I verify consistent load trends on higher and lower fidelity models, and ensure that the local blade sections have a reasonable stall margin. A set of Pareto fronts for different sets of design variables is identified. Increasing design freedom leads to lighter turbine for comparable aerodynamic loads. Moreover, I discuss some unintuitive design trends, such as inboard load shifting as tradeoff between energy yield and blade mass reduction.

In Chapter 6 I use an improved structural model that accounts for the anisotropic properties of the composite sandwich panels. I enable the optimizer to simultaneously modify the fiber angle orientation and panel thickness. I only consider twist design variables and two design points to investigate the aerostructural bend-twist coupling properties of the turbine and focus on the effect of tailored anisotropic properties. The mass is reduced by up to 19% while increasing the aerodynamic loads over the blade, compared to the baseline layout.

In Chapter 7 I move on to more complex multipoint optimizations including airfoil shape and blade planform variables. I demonstrate the rotors obtained with this approach perform better compared to single-point optimizations while satisfying the full set of aerostructural constraints. I quantify the impact of design points and objective weights on the final design in terms of mass reduction and torque yield. Moreover, I show how flow separation can be eliminated at an off-design point through design optimization. A case with more than 400 design variables, including airfoil shape and fiber angles, is finally shown to enable an averaged 14% increase in torque and a 39% mass reduction.

An ongoing collaborative effort with the FLOW Lab aims at coupling conventional codes with MACH to inform fatigue and extreme-loads constraints with loosely coupled dynamic simulations. In Chapter 8 we present two different ways of integrating these life-cycle sizing constraints with a staggered load-estimation and optimization approach. We first detail an approach based on scaling reference aerodynamic loads to replicate extreme loads and fatigue damage through a steady-state FEM analysis. We then discuss a condensation-based approach that transfers the high-fidelity structural properties into an equivalent beam model, evaluated directly within the optimization loop. The preliminary results demonstrate the feasibility of this approach and pave the way for more extensive tool implementation and verification.

9.1 Contribution summary

Reconnecting with the work outline in Section 1.4, my dissertation brings the following novel contributions:

Developed and verified a CFD-CSM wind turbine model to perform MDO studies with the MACH framework. (Chapters 3 and 4) My research effort involved both model pre-processing and software development. I developed from scratch a set of structural models that fit the preexisting aerodynamic model and have an internal topology consistent with the benchmark rotor. I finalized and publicly released a long standing bug-fix in ADflow to enable gradient-based optimizations with moving reference frames. I verified the derivatives of both the aerodynamic and the tightly coupled aerostructural model against finite difference and complex step derivatives. The consistency of the load-displacement transfer module has also been verified. Finally, the geometry parametrization module has been updated to handle multiple rotor blades consistently.

Demonstrated the advantages of tightly coupled high-fidelity optimization by performing the first aerostructural wind turbine rotor design study of its kind. (Chapter 5) I compared tightly coupled optimizations with staggered, loosely coupled aerostructural mass minimization problems, showing how even for a simplified problem like this the optimizer can capture aerostructural tradeoffs. Then, I performed a set of aerostructural optimizations enabling concurrent blade shape and structural thickness distribution modifications with an unprecedented level of detail. I demonstrated how the coupled approach can concurrently reduce the rotor mass and increase the turbine torque output while satisfying structural stress and horizontal force constraints. The optimization convergence behavior has also been investigated. I finally performed post-optimality analyses with higher and lower fidelity approaches to verify the performance improvements.

Tailored the blade composite fiber orientation and obtained lighter and stiffer turbine rotors using high-fidelity MDO. (Chapter 6) I further refined the initial structural model to include anisotropic composite material properties and to use fiber angle as a design variable. I then investigated how material anisotropy can be leveraged to obtain lighter and stiffer turbine blades. I explored different parametrization setups, layup configurations, and load scenarios to identify the blade components that would benefit the most from fiber misalignment. I finally quantified the benefits of more granular fiber steering and demonstrated how fiber angles can be tailored to obtain even lighter turbine rotors for the same torque output.

Investigated the impact of geometrical parametrization, design objective, and site-specific conditions on the optimized rotor layout. (Chapters 5 and 7) I leveraged on the design space exploration capabilities of the high-fidelity optimization framework to perform a range of parametric optimization studies. I showed how the optimizer can exploit increasing design freedom to design lighter blades for the same torque output. I investigated how a different emphasis of mass and torque in the objective function affects the final design, identifying a set of Pareto fronts. Finally, I reformulated the turbine design optimization into a multipoint problem, enabling concurrent blade planform and local shape modification. This further extended the design space to obtain more consistent performance over the turbine operational range. The impact of different site conditions has also been assessed through another set of parametric studies.

Contributed to the development of a mixed-fidelity framework to include life-cycle sizing constraints (for fatigue damage and extreme loads) in rotor design optimization. (Chapter 8) In collaboration with the BYU's Flow Lab, I worked on the integration of the MACH framework with conventional wind turbine design tools. I supported the formulation and implementation of a framework where the high-fidelity optimization includes additional fatigue and extreme loads constraints informed by time-dependent simulations. We investigated two different approaches to ensure a consistent information flow between the low and high-fidelity models. We demonstrated the model consistency and the overall optimization convergence, performing a first of a kind mixed-fidelity optimization with CFD-CSM in the loop.

9.2 Future research opportunities

This first application of high-fidelity MDO to wind energy systems paves the way for more complex and practical design studies and novel research opportunities. The results in this dissertation show the feasibility and potential using a CFD-CSM solver early in the design process. However, there is still a large gap to bridge between these academic studies and practical industrial applications.

There are several aspects worth in-depth research effort. Some of the foreseeable opportunities are specific to the toolkit and limitations of the present work, others are more general. I list what I consider the most relevant next research steps below:

Improve flow solver accuracy and speed for rotating blades applications: ADflow enables efficient CFD-based gradient-based optimization. However, I identified and discussed some of its speed and accuracy limitations for wind turbine applications. To address this in the short term, some of the recent improvements on ADflow low-speed preconditioner should be tested on the most refined DTU 10 MW meshes used in this work. On the one hand, speed improvements could make optimizations with the L0 mesh viable. On the other hand, reducing mesh discretization sensitivity would make L1-based optimizations more robust. In the longer term, implementing either cyclic boundary conditions, to limit the analysis to a single blade, or dedicated low-speed solvers would make the solver better suited for wind turbine applications.

Replicate optimization capabilities with MPhys and test different solvers: Most of the MACH components are publicly available, but users face a steep installation and usage learning curve. MPhys, an OpenMDAO-based framework for high-fidelity optimization, replicates and extends MACH capabilities in a more flexible and modular way. Porting the capabilities and setup I developed into MPhys would enable other researchers and engineers to replicate the study with their own MPhys-compatible tools. This would not just ensure this work a legacy, but would greatly facilitate the integration of additional components for constraint evaluation, cost modeling, and more.

Explore complex and unconventional system configurations: This work focused on conventional utility-scale wind turbine layouts. However, our tool and approach can be applied to any turbine configuration - the main overhead cost is mainly related to meshing and setup. Moreover, the problem formulation could be further expanded to include more design points and a broader set of design variables. The optimizer could manipulate the blade span to exploit additional torque and weight tradeoffs. Investigating out-of-plane geometrical deformations could yield even larger benefits. Improved designs could be obtained by optimizing the blade pre-cone angle or modeling winglets at the blade tip.

Extend and verify the mixed-fidelity optimization framework: The prototype tool we presented has large margins of improvements. We first need to verify the consistency between high and low fidelity models more extensively, before we move on to larger scale optimization studies. The fatigue damage and extreme load extrapolation need more thorough verification with benchmark cases. The set of DLCs we include in the OpenFAST-based analyses could also be expanded to include DLC 3.x and DLC 6.x scenarios. The full capabilities of the condensation approach are yet to be implemented. In particular, we need to enable geometrical design variables within the high-fidelity/beam-model design loop. Finally, we can explore strategies for improving the robustness and speed of the sequential optimization loop, such as automatic metric-based termination criteria to switch between the low-fidelity and high-fidelity loops.

Include dynamic stability constraints (flutter, principal frequency) to optimization formulation: This work focused on steady-state turbine performance. However, MACH has been used to perform flutter-constrained optimization. Including dynamic stability considerations in the optimization would likely lead to a shape and mass distribution redesign to push a possible flutter onset outside the nominal operating envelope. Additionally, TACS includes eigenvalue analysis tools that could be used to constrain the values of the dominating structural frequencies to avoid dangerous resonance phenomena. Finally, MACH time-spectral capabilities for aerodynamic shape optimization could be extended to the aerostructural model to investigate the blade dynamic behavior over a larger operating range.

Explore coupling with control co-design frameworks: Control co-design is emerging as a promising strategy to further reduce the levelized cost of energy of wind turbine. The coupled design of control policies and physical system could identify tradeoffs early in the design process, similarly to what discussed for the monolithic integration of aerodynamic and structural analyses. Although the cost of high-fidelity simulations would still be prohibitive, their additional design freedom offers large margin of improvement. Would it be possible to co-design the controller and the blade airfoil shape? The different models could be coupled with a nested (sequential) approach. The high-fidelity model could include proxy metrics for the controller design, such as pitch actuation forces. Finally, surrogate models based on high-fidelity analyses could be included within a control co-design framework.

Appendix A

Stress Concentrations in Structural Grid Convergence

Finite element structural codes are known to introduce artificial stress concentrations as the mesh refinement is increased. Unrealistic stress accumulation would lead the optimizer to oversize the rotor in a way that is not practical or meaningful. The mesh selection discussed in Section 4.1.1 took this into account. We identified a grid size that accurately captures the displacement under distributed aerodynamic loading, without incurring in these spurious stress concentrations.

The critical areas affected by spurious stresses are mainly located at the spar-skin junction. On the upper skin we observe a high stress area where the reinforcement spar connects with the skin close to the root. On the lower skin, the stress increases with mesh refinement around 60% of the span, at the front spar connection. The affected areas are highlighted in Figure A.1.

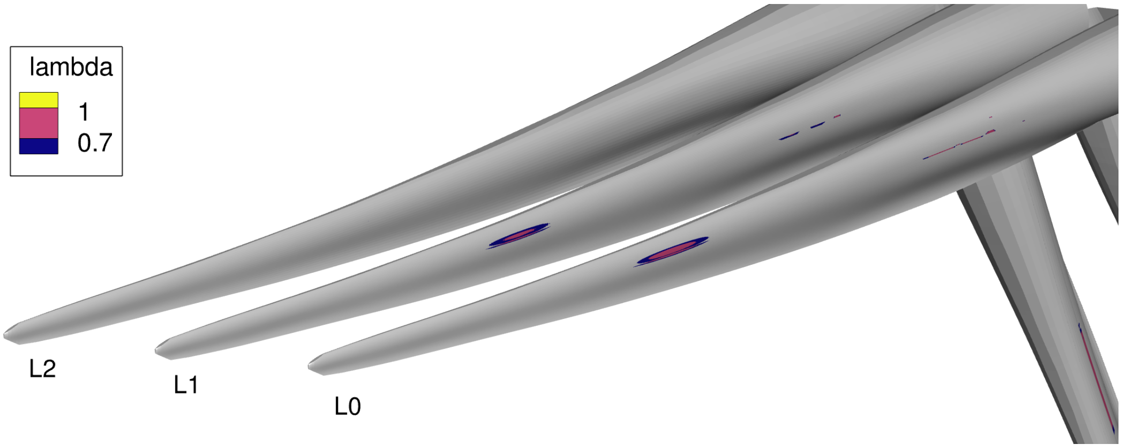

Figure A.1: “lambda” levels on the blade lower skins for the current mesh family. Areas with “lambda” $<=$ 0.69 are gray.

The convergence plots for the mesh family in Figure 3.9 are shown in Figure A.2 for the KS-aggregated and actual maximum stress values over the blade. We quantified stress-based grid convergence on both the full mesh and a subset of elements. In the latter case, we removed the critical areas highlighted in Figure A.1 from the TACS function used to evaluate the blade maximum stress.

Note that Figure A.2 has different y-axis scales between the right and left plot to account for the wider range of stress values when spurious concentrations are not removed. The convergence slopes on the left plots are more shallow than their counterpart on the right.

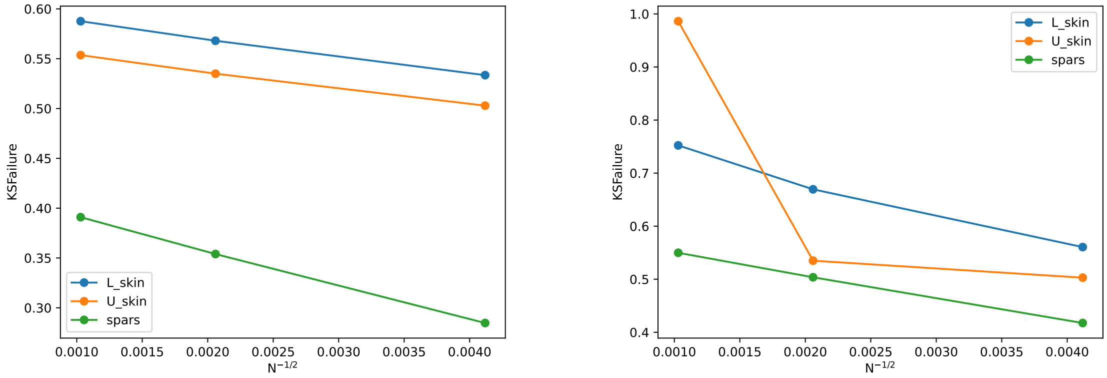

(a) KSFailure, no stress concentration

(b) KSFailure, full mesh

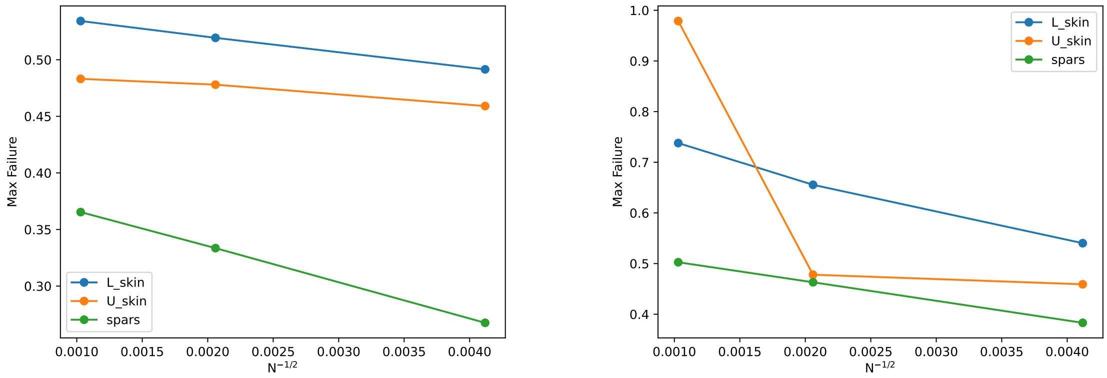

(c) MaxFailure, no stress concentration

(d) MaxFailure, full mesh

Figure A.2: Grid convergence study of KS-aggregated and maximum stress values.

Appendix B

Structural Parametrization

The parametrization used for the prototype structural model in Section 4.1.1.1 and chapter 5 has 117 design variables. It has only 9 spanwise sections but the upper and lower skins are decoupled to increase the number of design variables.

The updated model in Section 4.1.1.2 has a substantially different skin discretization. This is the result of a thorough literature review and model verification. The reference benchmark [24] defines the thickness distribution with up to a hundred spanwise sections and individual component ply layups. Such a granular definition is impractical for high-fidelity optimizations as the large design space would require several hundreds of iterations and increase the risk of incurring in local minima. The current parametrization shown in Figure 3.5 is based on Fig.4.11/4.13 of the DTU reference document [24], with the sections between 37 m and 81.8 m further split in half with respect to the original parametrization.

This discretization amounts to a total of 35 spanwise sections. At each spanwise section, the skins are split into four chordwise sections: a spar cap, a front panel, a rear panel, and an additional panel in between the reinforcement spar and the trailing edge. Each main spar, the reinforcement spar, and the leading and trailing edge panels are handled separately. This discretization leads to a high number of structural design variables that might over-complicate the problem or lead to a unrealistic structural layout.

We thus need to consider how to aggregate these panels into a practical parametrization. Appendix B.1 and Appendix B.2 discuss the spanwise and chordwise parametrization respectively. Appendix B.3 reports some structural-only optimization studies done to support the parametrization selection.

B.1 Spanwise parametrization

Debugging scripts used at this stage allowed the user to define the skins aggregation “on the fly”. Two examples of different parametrizations are shown in Figure B.1.

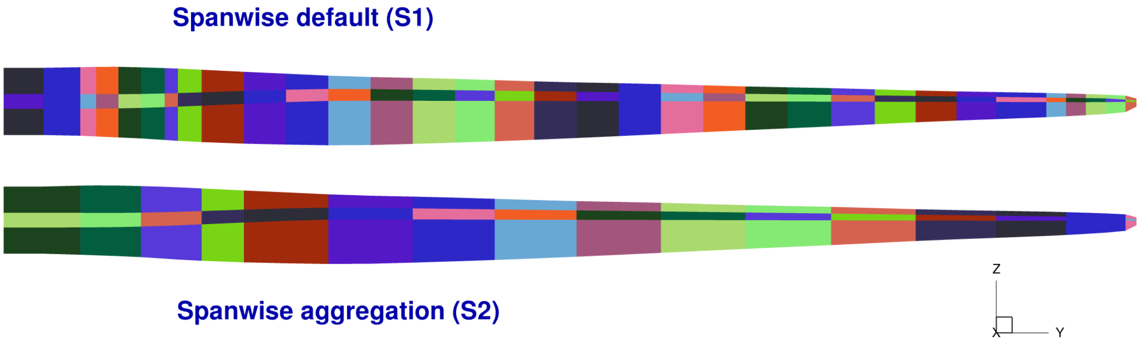

Figure B.1: Color-coded skin panel parametrizations. The most granular set of DVs is at the top, and a reduced set is at the bottom. Note that in this example there are only two chordwise variables, one for the spar cap and one for the rest of the section.

The top parametrization uses all the available spanwise sections defined in the *.bdf* file. The bottom parametrization in Figure B.1 (a debugging mesh with 33 spanwise sections) groups the spanwise sections as follows:

```
[0, 1],
[2, 3, 4],
[5, 6, 7],
[8],
[9, 10],
[11, 12],
[13, 14],
[15, 16],
[17, 18],
[19, 20],
[21, 22],
[23, 24],
[25, 26],
[27, 28],
[29, 30, 31],
[32],
```

Looking at the DTU parametrization and some preliminary optimization results reported in Appendix B.3, we found a compromise between the two parametrizations in Figure B.1. Figure B.2 is an example compromise with denser spanwise variables in the spanwise inner 30% and a more uniform panel size up to the tip. Compared to the default parametrization, has fewer DVs in the tip area and mid span.

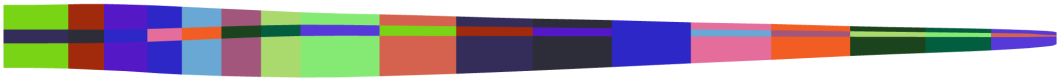

Figure B.2: One of the investigated parametrization, with updated panel aggregation at the tip and the root.

B.2 Chordwise parametrization

Panel parametrization along the chordwise required more practical and literature-informed considerations. Figure B.3 shows the set of panels that can be handled individually by TACS (and thus the optimizer) at each section.

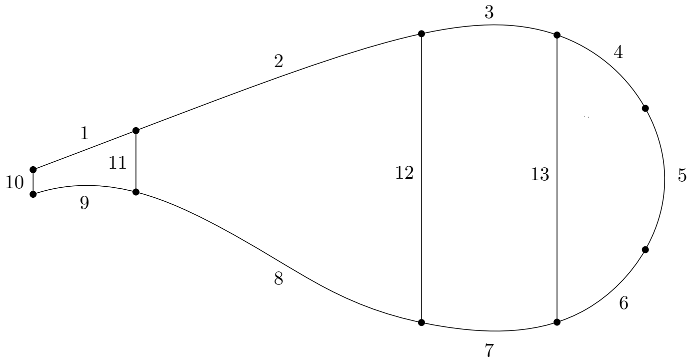

Figure B.3: Sketch of the structural parametrization of a blade section.

Starting from the full set of variables in Figure B.3 as reference, we tested several grouping options:

C0 is the parametrization shown earlier in Figure B.1, for a total of 2 variables. This grouping would likely excessively limit the design space. C1 splits C0 into having two separate variables for upper and lower skin respectively. Finally, in C3 we aggregate together panels 2-8, 3-7, 4-5, and finally 1-9-10 as a trailing edge block.

The remaining spars are independent, with panel 5 acting as LE reinforcement although effectively be a LE spar. We give the optimizer some extra design freedom by allowing the two main spars (12, 13) to be independent, but we do not generally observe large differences the optimized thickness distribution between the two spars.

We enforce some sections of the airfoil to have the same structural properties for practical purposes. As we add ply angle, the number of design variables increases by a factor between 2 and 3 depending on the ply layup. Moreover, blades are usually manufactured by layering several fiberglass plies on the upper and lower skin separately, and add the shear web as the two skins are glued together.

Table B.1: Chordwise parameterization examples.

Parametrization ID	DV Group	Segment #
C0	Skins	1,2,4,6,8,9
	Spar caps	3,7
C1	Upper Skins	1,2,4
	Lower Skins	6,8,9
	Spar caps	3,7
C2	Front Skins	4,6
	Spar caps	3,7
	Mid Skins	2,8
	Aft Skins	1,9,10
C3	No grouping	-

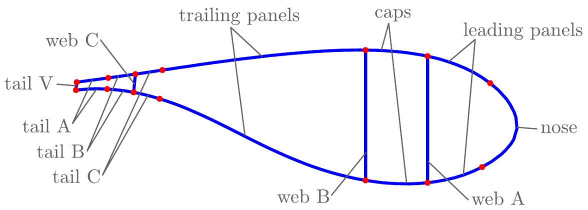

Figure B.4: Benchmark topology parametrization, from Bak et al. [24, Fig.4.3].

The benchmark parametrization from Bak et al. [24, Fig.4.3] is shown in Figure B.4.

We explored alternative approaches for the DTU 10 MW model the one proposed by Bottasso et al. [8]. Carlo Bottasso's team at TU Munich is one of the few doing comprehensive aerostructural optimization for wind turbines - with low fidelity codes, and a non-monolithic architecture. They aggregate the leading panels into a single component and have two separate design variables for the trailing panels on the suction and pressure side respectively. Figure B.5 illustrates their approach.

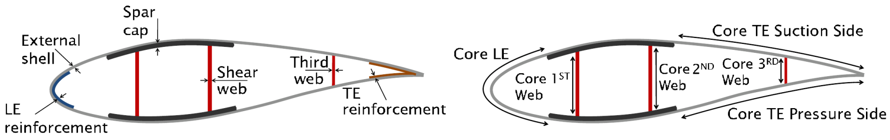

(a) Main structural components (b) Secondary core structures

Figure B.5: Parametrization example from Bottasso et al. [8].

NREL had a pretty simplified approach for the structural parametrization of the "successor" of the DTU 10MW, the IEA 15 MW [22]. The reference documentation [1] defines all the skins in the same way, with the spar caps and the trailing edge having different material definitions. More details and scripts are available on the official GitHub repository [2].

Textbook examples of wind turbine structures shown in Figure B.6 also suggest mirroring the upper and lower skins. The trailing panels have different layups than the leading panels.

(a) Figure 11.1, Hansen [34] textbook

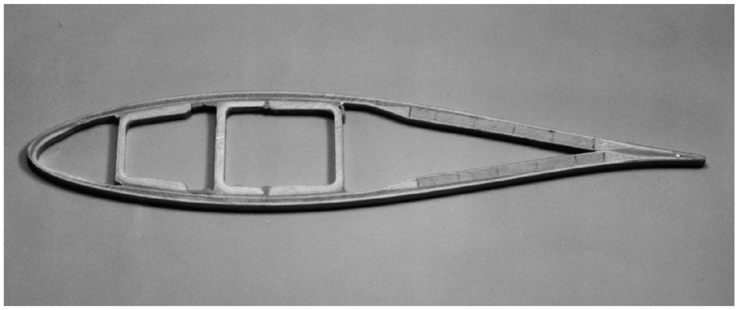

(b) Fig. 8.18, Hau [5] textbook

Figure B.6: Pictures of blade sections, from textbooks

Given these considerations, the selected parametrization is C3 in Table B.1 with independent spars (numbered 11, 12, 13, 3, and 7 in Figure B.3):

There is no need to use linear constraints to aggregate panels in the same section. The loop used for aggregating spanwise DVs to do the same on the same spanwise section. A code snippet from the TACS setup file is reported below:

```
# --- Spanwise panel aggregation ---
spanwise_set = [
    [0, 1],
    [2, 3],
    [4, 5],
    [6, 7],
```

---

[1] NREL report(check page 14)

[2] IEA 15 MW GitHub repository

```
    [8],
    [9],
    [10],
    [11],
    [12],
    [13],
    [14],
    [15, 16],
    [17, 18],
    [19, 20],
    [21, 22],
    [23, 24],
    [25, 26],
    [27, 28],
    [29, 30, 31, 32, 33, 34],
]

[....]


# --- SKINS ---

for j in range(len(spanwise_set)):
    for comp in blades:
        TEPatchName = []
        for _, k in enumerate(spanwise_set[j]):
            # adding TE group
            TEPatchName.append(f"L_SKIN{comp}/L_SKIN.{k:03}/" + "SEG.00")
            TEPatchName.append(f"U_SKIN{comp}/U_SKIN.{k:03}/" + "SEG.00")
            TEPatchName.append(f"SPARS{comp}/SPAR.00/" + f"SEG.{k:02}")
        FEASolver.addDVGroup("TE", include=TEPatchName)

    # Looping over TE
    for i in [1, 2, 3]:
        for comp in blades:
            patchName = []
```

```
        for _, k in enumerate(spanwise_set[j]):
            for skin in ["U", "L"]:
                groupName = f"{skin}_SKIN"

                patchName.append(f"{groupName}{comp}/{groupName}.{k:03}

        FEASolver.addDVGroup(f"CHORD_{i}", include=patchName)
```

B.3 Preliminary structural optimization results

In this section reports the results of a set of optimizations using 4 chordwise parametrizations (C0,C1,C2,C3) and 3 spanwise parametrizations (S2,S3,S4). On top of the practical considerations mentioned in Appendices B.1 and B.2, we tested the different parametrizations to understand their impact on the optimization convergence and final results.

Table B.2 lists the number of design variables for each combinations. Tables B.3 to B.7 report structural optimization results for different parametrizations and different combination of finite elements constitutive models and material properties.

Table B.2: Number of DVs for combination of spanwise and chordwise parametrization

Parametrization	S2	S3	S4
C0	105	114	128
C1	137	148	166
C2	121	131	147
C3	201	216	242

Table B.3: Structural optimization results using isoFSDT elements. Case highlighted in red did not converge.

Param	Runtime [s]	Maj Iters	Final mass [Kg]	Struct mesh	Aero Loads
C0/S2	496	120	39999	L2	L1
C1/S2	596	144	38519	L2	L1
C2/S2	511	114	38610	L2	L1
C3/S2	614	148	36567	L2	L1
C0/S3	492	120	39452	L2	L1
C1/S3	1123	192	38036	L2	L1
C2/S3	427	99	38091	L2	L1
C3/S3	678	163	36186	L2	L1
C0/S4	501	126	39100	L2	L1
C1/S4	694	129	37836	L2	L1
C2/S4	678	142	37757	L2	L1
C3/S4	745	166	35994	L2	L1

Table B.4: Structural optimization results using smearedFSDT elements with aluminum properties. Case highlighted in red did not converge.

Parametrization	Runtime [s]	Maj Iters	Final mass [Kg]	Struct mesh	Aero Loads
C0/S2	657	167	39629	L2	L1
C1/S2	973	214	38038	L2	L1
C2/S2	899	205	38212	L2	L1
C3/S2	1045	232	36108	L2	L1
C0/S3	915	197	39044	L2	L1
C1/S3	1004	231	37530	L2	L1
C2/S3	1019	246	37689	L2	L1
C3/S3	1318	273	35734	L2	L1
C0/S4	1008	234	38654	L2	L1
C1/S4	1210	254	37318	L2	L1
C2/S4	1076	260	37339	L2	L1
C3/S4	1378	305	35506	L2	L1

Table B.5: Structural optimization results using smearedFSDT elements with composite properties. Case highlighted in red did not converge.

Parametrization	Runtime [s]	Maj Iters	Final mass [Kg]	Struct mesh	Aero Loads
C0/S2	212	36	75627	L2	L1
C1/S2	302	50	74003	L2	L1
C2/S2	252	45	75171	L2	L1
C3/S2	368	75	72411	L2	L1
C0/S3	215	28	73566	L2	L1
C1/S3	350	61	68390	L2	L1
C2/S3	219	53	72888	L2	L1
C3/S3	400	83	66730	L2	L1
C0/S4	342	58	72906	L2	L1
C1/S4	270	48	68136	L2	L1
C2/S4	200	33	71828	L2	L1
C3/S4	385	80	66514	L2	L1

Table B.6: Structural optimization results using smearedFSDT elements with composite properties, including a single fiber angle per panel as DV. Case highlighted in red did not converge.

Parametrization	Runtime [s]	Maj Iters	Final mass [Kg]	Struct mesh	Aero Loads
C0/S2	4936	1000	67245	L2	L1
C1/S2	1790	496	65112	L2	L1
C2/S2	1700	511	66906	L2	L1
C3/S2	2586	783	63523	L2	L1
C0/S3	2693	726	65701	L2	L1
C1/S3	2409	678	62154	L2	L1
C2/S3	1898	547	65491	L2	L1
C3/S3	3900	1127	60935	L2	L1
C0/S4	2392	591	66583	L2	L1
C1/S4	2674	504	62002	L2	L1
C2/S4	2060	628	65006	L2	L1
C3/S4	3008	850	60515	L2	L1

Table B.7: Structural optimization results using smearedFSDT elements with composite properties, including fiber angle as DV. Case highlighted in red did not converge.

Parametrization	Runtime [s]	Maj Iters	Final mass [Kg]	Struct mesh	Aero Loads
C0/S2	2067	500	68459	L2	L1
C1/S2	2557	820	63349	L2	L1
C2/S2	2130	656	64418	L2	L1
C3/S2	6715	1263	60812	L2	L1
C0/S3	2069	462	67324	L2	L1
C1/S3	2724	863	60740	L2	L1
C2/S3	2948	986	62680	L2	L1
C3/S3	9552	1842	61463	L2	L1
C0/S4	3100	734	66158	L2	L1
C1/S4	3770	1257	60435	L2	L1
C2/S4	3039	1003	62335	L2	L1
C3/S4	14868	2553	63013	L2	L1

Appendix C

Mixed-Fidelity Framework XDSM Diagrams

We attach two XDSM diagrams describing the mixed-fidelity framework architecture from Chapter 8. Figure C.1 refers to the “load-scaling” approach in Section 8.2.2, while Figure C.2 refers to the “condensation” approach in Section 8.2.3.

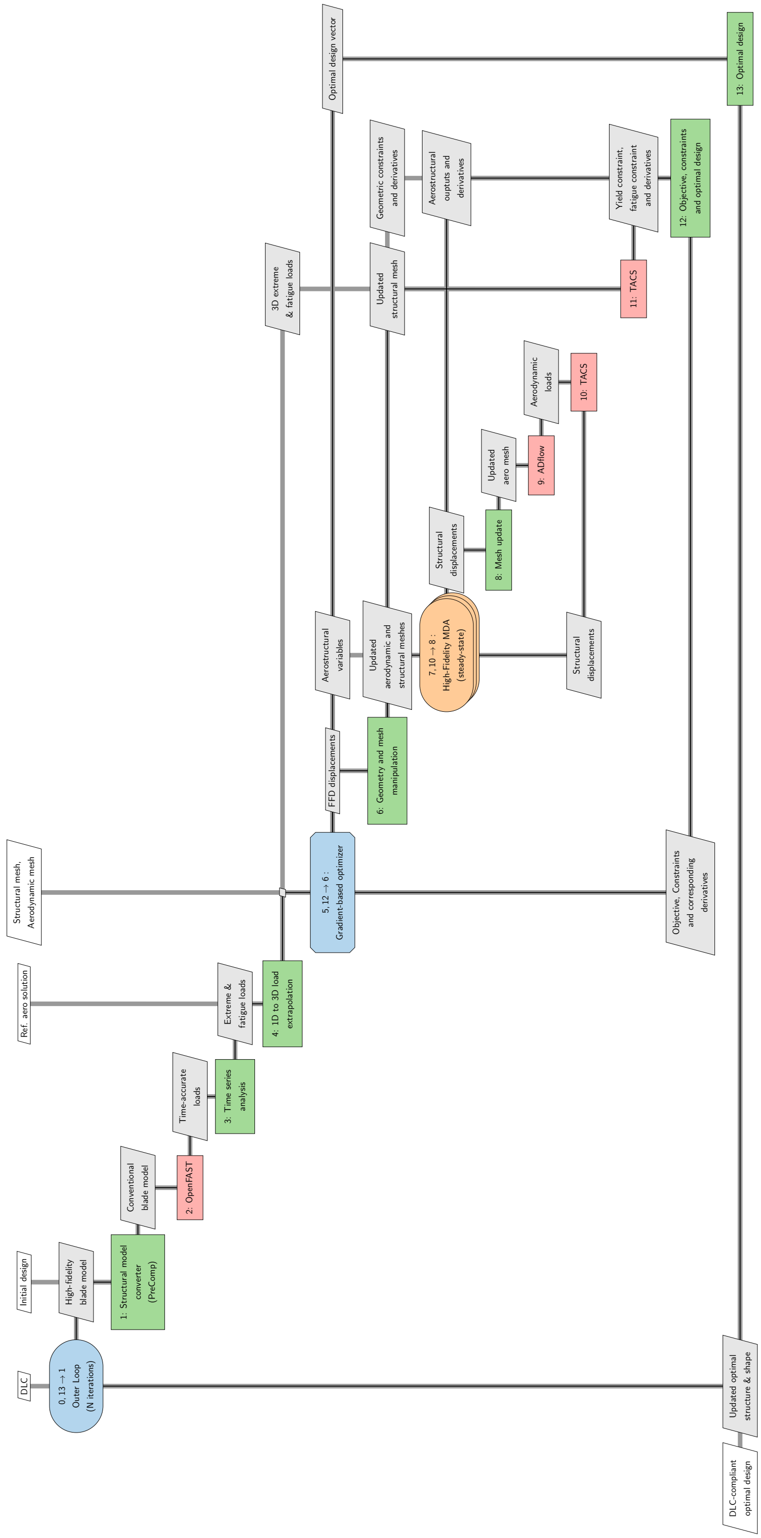

Figure C.1: XDSM diagram for the mixed-fidelity framework implemented in Caprace et al. [87].

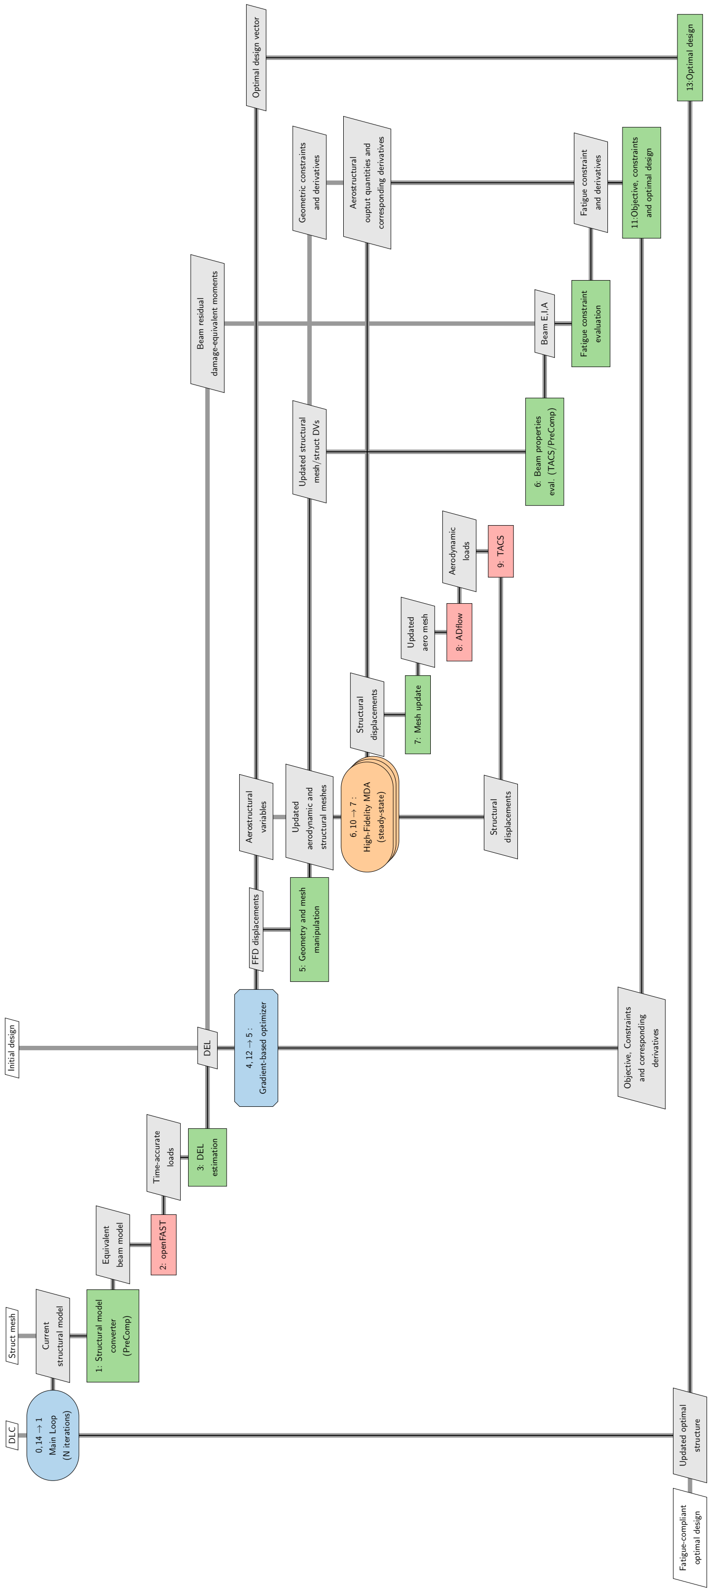

Figure C.2: XDSM diagram for the fatigue damage integration from Sec. 8.2.3.

Bibliography

[1] Valérie Masson-Delmotte, Panmao Zhai, Hans-Otto Pörtner, Debra Roberts, Jim Skea, Priyadarshi R Shukla, Anna Pirani, W Moufouma-Okia, C Péan, R Pidcock, et al. Global warming of 1.5 c. an ipcc special report on the impacts of global warming of 1.5 °c above pre-industrial levels and related global greenhouse gas emission pathways, in the context of strengthening the global response to the threat of climate change, sustainable development, and efforts to eradicate poverty. Technical report, Intergovernmental Panel on Climate Change (IPCC), 2018. URL `https://www.ipcc.ch/site/assets/uploads/sites/2/2019/06/SR15_Full_Report_High_Res.pdf`.

[2] Laura Cozzi, Timur Gül, and Stéphanie Bouckaert. Net zero by 2050 —a roadmap for the global energy sector. Technical report, IEA, May 2021.

[3] Paul Veers, Katherine Dykes, Eric Lantz, Stephan Barth, Carlo L Bottasso, Ola Carlson, Andrew Clifton, Johney Green, Peter Green, Hannele Holttinen, et al. Grand challenges in the science of wind energy. *Science*, 366(6464):eaau2027, 2019.

[4] P. Veers, C. Bottasso, L. Manuel, J. Naughton, L. Pao, J. Paquette, A. Robertson, M. Robinson, S. Ananthan, A. Barlas, A. Bianchini, H. Bredmose, S. G. Horcas, J. Keller, H. A. Madsen, J. Manwell, P. Moriarty, S. Nolet, and J. Rinker. Grand challenges in the design, manufacture, and operation of future wind turbine systems. *Wind Energy Science Discussions*, 2022:1–102, 2022. doi:10.5194/wes-2022-32. URL `https://wes.copernicus.org/preprints/wes-2022-32/`.

[5] Erich Hau. *Wind Turbines — Fundamentals, Technologies, Application, Economics*. Springer, 2nd edition, 2006.

[6] W. Musial, D. Heimiller, P. Beiter, G. Scott, and C. Draxl. 2016 offshore wind energy resource assessment for the united states. Technical Report NREL/TP-5000-66599, National Renewable Energy Laboratory, 2016. URL `http://www.nrel.gov/docs/fy16osti/66599.pdf`.

[7] Jason M. Jonkman. Dynamics of offshore floating wind turbines—model development and verification. *Wind Energy*, 12(5):459–492, 2009. doi:https://doi.org/10.1002/we.347. URL `https://onlinelibrary.wiley.com/doi/abs/10.1002/we.347`.

[8] Carlo Luigi Bottasso, P Bortolotti, Alessandro Croce, and Federico Gualdoni. Integrated aero-structural optimization of wind turbines. *Multibody System Dynamics*, 38(4):317–344, 2016. doi:10.1007/s11044-015-9488-1.

[9] R. Barrett and A. Ning. Integrated free-form method for aerostructural optimization of wind turbine blades. *Wind Energy*, 21(8):663–675, 2018. doi:10.1002/we.2186.

[10] Joaquim R. R. A. Martins and Andrew B. Lambe. Multidisciplinary design optimization: A survey of architectures. *AIAA Journal*, 51(9):2049–2075, September 2013. doi:10.2514/1.J051895.

[11] Mario Garcia-Sanz. A metric space with lcoe isolines for research guidance in wind and hydrokinetic energy systems. *Wind Energy*, 23(2):291–311, 2020. doi:10.1002/we.2429.

[12] Gaetan K. W. Kenway, Graeme J. Kennedy, and Joaquim R. R. A. Martins. Scalable parallel approach for high-fidelity steady-state aeroelastic analysis and adjoint derivative computations. *AIAA Journal*, 52(5):935–951, May 2014. doi:10.2514/1.J052255.

[13] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multipoint high-fidelity aerostructural optimization of a transport aircraft configuration. *Journal of Aircraft*, 51(1):144–160, January 2014. doi:10.2514/1.C032150.

[14] Tristan Dhert, Turaj Ashuri, and Joaquim R. R. A. Martins. Aerodynamic shape optimization of wind turbine blades using a Reynolds-averaged Navier–Stokes model and an adjoint method. *Wind Energy*, 20(5):909–926, May 2017. doi:10.1002/we.2070.

[15] Mads H. Aa. Madsen, Frederik Zahle, Niels N. Sørensen, and Joaquim R. R. A. Martins. Multipoint high-fidelity CFD-based aerodynamic shape optimization of a 10 MW wind turbine. *Wind Energy Sciences*, 4:163–192, April 2019. doi:10.5194/wes-4-163-2019.

[16] Hoesung Lee, Katherine Calvin, Dipak Dasgupta, Gerhard Krinner, Aditi Mukherji, Peter Thorne, Christopher Trisos, José Romero, Paulina Aldunce, and Alexander C Ruane. Climate change 2023: Synthesis report. Technical report, Intergovernmental Panel on Climate Change, 2024.

[17] Julia Kirch Kirkegaard, David Philipp Rudolph, Sophie Nyborg, Helena Solman, Elizabeth Gill, Tom Cronin, and Mary Hallisey. Tackling grand challenges in wind energy through a socio-technical perspective. *Nature Energy*, pages 1–10, 2023.

[18] IRENA. World energy transitions outlook 2023: 1.5° c pathway, volume 1, 2023.

[19] Yulizza Henao, Russell Gentry, Tristan Al-Haddad, Lawrence C Bank, and John E Taylor. Construction assessment framework of electrical transmission structures from decommissioned wind turbine blades. In *Construction Research Congress 2022*, pages 588–598, 2022.

[20] Alejandra S Escalera Mendoza, Shulong Yao, Mayank Chetan, and Daniel Todd Griffith. Design and analysis of a segmented blade for a 50 mw wind turbine rotor. *Wind Engineering*, 46(4):1146–1172, 2022.

[21] Venkata Yaramasu and Bin Wu. *Model predictive control of wind energy conversion systems*. John Wiley & Sons, 2016.

[22] Evan Gaertner, Jennifer Rinker, Latha Sethuraman, Frederik Zahle, Benjamin Anderson, Garrett Barter, Nikhar Abbas, Fanzhong Meng, Pietro Bortolotti, Witold Skrzypinski, George Scott, Roland Feil, Henrik Bredmose, Katherine Dykes, Matt Sheilds, Christopher Allen, and Anthony Viselli. Definition of the IEA 15-megawatt offshore reference wind turbine. Technical report, International Energy Agency, 2020. URL `https://www.nrel.gov/docs/fy20osti/75698.pdf`.

[23] Andrew Scholbrock, Paul Fleming, David Schlipf, Alan Wright, Kathryn Johnson, and Na Wang. Lidar-enhanced wind turbine control: Past, present, and future. In *2016 American Control Conference (ACC)*, pages 1399–1406. IEEE, 2016.

[24] Christian Bak, Frederik Zahle, Robert Bitsche, Taeseong Kim, Anders Yde, Lars Christian Henriksen, Anand Natarajan, and Morten Hansen. Description of the dtu 10 mw reference wind turbine. Technical Report DTU Wind Energy Report-I-0092, Danish Technical University, DTU Wind Energy Roskilde, DK, 2013. URL `https://rwt.windenergy.dtu.dk/dtu10mw/dtu-10mw-rwt/-/blob/master/docs/DTU_Wind_Energy_Report-I-0092.pdf`.

[25] International Electrotechnical Commission. Wind turbines – part 3: Design requirements for offshore wind turbines. Technical Report IEC 61400-3, International Electrotechnical Commission, Geneva, Switzerland, 2009.

[26] C. L. Johnson and M. Smith. *Kelly: More than my share of it all*. Smithsonian Books, 1985.

[27] M Baudino Bessone, M Zaaijer, D Von Terzi, K Dykes, E Jump, and A Viré. Including installation logistics costs in the optimal sizing of semi-submersibles for floating wind farms. In *Journal of Physics: Conference Series*, volume 2265, page 042018. IOP Publishing, 2022.

[28] Anders Myhr, Catho Bjerkseter, Anders Ågotnes, and Tor A Nygaard. Levelised cost of energy for offshore floating wind turbines in a life cycle perspective. *Renewable energy*, 66: 714–728, 2014. doi:10.1016/j.renene.2014.01.017.

[29] Katherine Dykes. Optimization of wind farm design for objectives beyond lcoe. In *Journal of Physics: Conference Series*, volume 1618, page 042039. IOP Publishing, 2020.

[30] Juliet Simpson, Eric Loth, and Katherine Dykes. Cost of valued energy for design of renewable energy systems. *Renewable Energy*, 153:290–300, 2020. ISSN 0960-1481. doi:https://doi.org/10.1016/j.renene.2020.01.131. URL `https://www.sciencedirect.com/science/article/pii/S0960148120301531`.

[31] Katherine Dykes, Jennifer King, Nicholas DiOrio, Ryan King, Vahan Gevorgian, David Corbus, Nate Blair, Kate Anderson, Greg Stark, Craig Turchi, et al. Opportunities for research and development of hybrid power plants. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2020.

[32] Jeffrey M Loiter and Vicki Norberg-Bohm. Technology policy and renewable energy: public roles in the development of new energy technologies. *Energy Policy*, 27(2):85–97, 1999.

[33] A. Betz. Das maximum der theoretisch möglichen ausnötzung des windes durch windmotoren,. *Zeitschrift für das gesamte Turbinewesen*, 26:307–309, 1920.

[34] Martin O.L. Hansen. *Aerodynamics of Wind Turbines*. James & James (Science Publishers) Ltd., 35-37 William Road, London, 2000.

[35] Patrick J Moriarty and A Craig Hansen. Aerodyn theory manual. Technical report, National Renewable Energy Lab., Golden, CO (US), 2005.

[36] Joaquim R. R. A. Martins and Andrew Ning. *Engineering Design Optimization*. Cambridge University Press, Cambridge, UK, 2021. ISBN 9781108833417. doi:10.1017/9781108980647. URL `https://mdobook.github.io`.

[37] Olivier Pironneau. On optimum profiles in Stokes flow. *Journal of Fluid Mechanics*, 59 (01):117–128, 1973. doi:10.1017/S002211207300145X.

[38] Raphael T. Haftka. Optimization of flexible wing structures subject to strength and induced drag constraints. *AIAA Journal*, 15(8):1101–1106, 1977. doi:10.2514/3.7400.

[39] A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3(3): 233–260, September 1988. doi:10.1007/BF01061285.

[40] Justin S. Gray, John T. Hwang, Joaquim R. R. A. Martins, Kenneth T. Moore, and Bret A. Naylor. OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization. *Structural and Multidisciplinary Optimization*, 59(4):1075–1104, April 2019. doi:10.1007/s00158-019-02211-z.

[41] François Gallard, Rémi Lafage, Charlie Vanaret, Benoit Pauwels, Damien Guénot, Pierre-Jean Barjhoux, Vincent Gachelin, and Anne Gazaix. GEMS: A Python library for automation of multidisciplinary design optimization process generation. In *18th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, June 2018. doi:10.2514/6.2018-0657.

[42] NREL. Openfast. open-source wind turbine simulation tool, 2018. `https://github.com/OpenFAST/openfast` [Accessed: 2022].

[43] Frederik Zahle, Carlo Tibaldi, Christian Pavese, Michael K. McWilliam, Jose P. A. A. Blasques, and Morten H. Hansen. Design of an aeroelastically tailored 10 MW wind turbine rotor. *Journal of Physics: Conference Series*, 753(6):062008, 2016.

[44] Samuel Scott, Terence Macquart, Carlos Rodriguez, Peter Greaves, Paul McKeever, Paul Weaver, and Alberto Pirrera. Preliminary validation of atom: an aero-servo-elastic design tool for next generation wind turbines. In *Journal of Physics: Conference Series*, volume 1222, page 012012. IOP Publishing, 2019. doi:10.1088/1742-6596/1222/1/012012.

[45] Alfonso del Carre, Arturo Muñoz-Simón, Norberto Goizueta, and Rafael Palacios. Sharpy: A dynamic aeroelastic simulation toolbox for very flexible aircraft and wind turbines. *Journal of Open Source Software*, 4(44):1885, 2019. doi:10.21105/joss.01885.

[46] Dennis Marten, Johannes Wendler, George Pechlivanoglou, Christian Navid Nayeri, and Christian Oliver Paschereit. Qblade: an open source tool for design and simulation of horizontal and vertical axis wind turbines. volume 3, pages 264–269, 2013.

[47] Mareike Leimeister, Athanasios Kolios, and Maurizio Collu. Development of a framework for wind turbine design and optimization. *Modelling*, 2(1):105–128, 2021. ISSN 2673-3951. doi:10.3390/modelling2010006.

[48] Sergio González Horcas, Néstor Ramos-García, Ang Li, Georg Pirrung, and Thanasis Barlas. Comparison of aerodynamic models for horizontal axis wind turbine blades accounting for curved tip shapes. *Wind Energy*, 26(1):5–22, 2022. doi:10.1002/we.2780. URL `https://onlinelibrary.wiley.com/doi/10.1002/we.2780`.

[49] K. Boorsma, G. Schepers, H. Aagard Madsen, G. Pirrung, N. Sørensen, G. Bangga, M. Imiela, C. Grinderslev, A. Meyer Forsting, W. Z. Shen, A. Croce, S. Cacciola, A. P. Schaffarczyk, B. Lobo, F. Blondel, P. Gilbert, R. Boisard, L. Höning, L. Greco, C. Testa, E. Branlard, J. Jonkman, and G. Vijayakumar. Progress in the validation of rotor aerodynamic codes using field data. *Wind Energy Science*, 8(2):211–230, 2023. doi:10.5194/wes-8-211-2023. URL `https://wes.copernicus.org/articles/8/211/2023/`.

[50] JG Schepers, Koen Boorsma, H Aa Madsen, GR Pirrung, G Bangga, Giorgia Guma, Thorsten Lutz, Thomas Potentier, Caroline Braud, Emmanuel Guilmineau, et al. Iea wind tcp task 29, phase iv: Detailed aerodynamics of wind turbines. Technical report, International Energy Agency, 2021.

[51] Néstor Ramos-García, Matias Sessarego, and Sergio González Horcas. Aero-hydro-servo-elastic coupling of a multi-body finite-element solver and a multi-fidelity vortex method. *Wind Energy*, 24(5):481–501, 2021. doi:10.1002/we.2584.

[52] Joachim C. Heinz, Niels N. Sørensen, and Frederik Zahle. Fluid-structure interaction computations for geometrically resolved rotor simulations using cfd. *Wind Energy*, 19(12): 2205–2221, 2016. doi:10.1002/we.1976.

[53] Matias Sessarego, Néstor ía, Hua Yang, and Wen Zhong Shen. Aerodynamic wind-turbine rotor design using surrogate modeling and three-dimensional viscous-inviscid interaction technique. *Renewable Energy*, 93:620–635, 2016. ISSN 0960-1481. doi:https://doi.org/10.1016/j.renene.2016.03.027. URL `https://www.sciencedirect.com/science/article/pii/S0960148116302117`.

[54] Lorenzo Cozzi, Juri Bellucci, Matteo Giovannini, Francesco Papi, and Alessandro Bianchini. Towards the development of an advanced wind turbine rotor design tool integrating full cfd and fem. In *Journal of Physics: Conference Series*, volume 2265, page 042050. IOP Publishing, 5 2022. doi:10.1088/1742-6596/2265/4/042050. URL `https://iopscience.iop.org/article/10.1088/1742-6596/2265/4/042050`.

[55] Jichao Li, Xiaosong Du, and Joaquim R. R. A. Martins. Machine learning in aerodynamic shape optimization. *Progress in Aerospace Sciences*, 134:100849, October 2022. doi:10.1016/j.paerosci.2022.100849.

[56] Kyoungsoo Lee, Ziaul Huque, Raghava Kommalapati, and Sang-Eul Han. Fluid-structure interaction analysis of nrel phase vi wind turbine: Aerodynamic force evaluation and structural analysis using fsi analysis. *Renewable Energy*, 113:512–531, 2017. ISSN 0960-1481. doi:10.1016/j.renene.2017.02.071.

[57] Thomas R. Wainwright, Daniel J. Poole, Christian B. Allen, Jamil Appa, and Oliver Darbyshire. High fidelity aero-structural simulation of occluded wind turbine blades. In *AIAA Scitech Forum*, 2021. doi:10.2514/6.2021-0950.

[58] Ping Cheng, Yang Huang, and Decheng Wan. A numerical model for fully coupled aero-hydrodynamic analysis of floating offshore wind turbine. *Ocean Engineering*, 173:183–196, 2019. ISSN 0029-8018. doi:10.1016/j.oceaneng.2018.12.021.

[59] Andrew Ning, Rick Damiani, and Pat Moriarty. Objectives and constraints for wind turbine optimization. *Journal of Solar Energy Engineering*, 136(4):041010–1–12, November 2014. doi:10.1115/1.4027693.

[60] Justin S. Gray, Tristan Hearn, Kenneth Moore, John T. Hwang, Joaquim R. R. A. Martins, and Andrew Ning. Automatic evaluation of multidisciplinary derivatives using a graph-based problem formulation in OpenMDAO. In *Proceedings of the 15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Atlanta, GA, June 2014. doi:10.2514/6.2014-2042.

[61] Andrew Ning and Derek Petch. Integrated design of downwind land-based wind turbines using analytic gradients. *Wind Energy*, 19(12):2137–2152, 2016. doi:10.1002/we.1972.

[62] Bryce Ingersoll and Andrew Ning. Efficient incorporation of fatigue damage constraints in wind turbine blade optimization. *Wind Energy*, 23(4):1063–1076, 2020. doi:10.1002/we.2473.

[63] Pietro Bortolotti, Kristian Dixon, Evan Gaertner, Megan Rotondo, and Garrett Barter. An efficient approach to explore the solution space of a wind turbine rotor design process. In *Journal of Physics: Conference Series*, volume 1618, page 042016. IOP Publishing, 9 2020. doi:10.1088/1742-6596/1618/4/042016. URL `https://dx.doi.org/10.1088/1742-6596/1618/4/042016`.

[64] Carlo Luigi Bottasso, Filippo Campagnolo, and Alessandro Croce. Multi-disciplinary constrained optimization of wind turbines. *Multibody System Dynamics*, 27(1):21–53, 2012. doi:10.1007/s11044-011-9271-x.

[65] P. Bortolotti, C. L. Bottasso, and A. Croce. Combined preliminary–detailed design of wind turbines. *Wind Energy Science*, 1(1):71–88, 2016. doi:10.5194/wes-1-71-2016.

[66] S Scott, Peter Greaves, Paul M Weaver, Alberto Pirrera, and Terence Macquart. Efficient structural optimisation of a 20 mw wind turbine blade. In *Journal of Physics: Conference Series*, volume 1618, page 042025. IOP Publishing, 2020. doi:10.1088/1742-6596/1618/4/042025.

[67] Michael K McWilliam, Frederik Zahle, Antariksh Dicholkar, David Verelst, and Taeseong Kim. Optimal aero-elastic design of a rotor with bend-twist coupling. In *Journal of Physics: Conference Series*, volume 1037, page 042009. IOP Publishing, 2018. doi:10.1088/1742-6596/1037/4/042009.

[68] Sergio González Horcas, François Debrabandere, Benoît Tartinville, Charles Hirsch, and Grégory Coussement. Cfd study of dtu 10 mw rwt aeroelasticity and rotor-tower interactions. In *MARE-WINT*, pages 309–334. Springer, Cham, 2016. doi:10.1007/978-3-319-39095-6_18.

[69] P. Fuglsang, C. Bak, J. G. Schepers, B. Bulder, T. T. Cockerill, P. Claiden, A. Olesen, and R. van Rossen. Site-specific design optimization of wind turbines. *Wind Energy*, 5:261–279, 2002.

[70] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Aerostructural shape optimization of wind turbine blades considering site-specific winds. In *Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, Victoria, BC, September 2008. doi:10.2514/6.2008-6025. AIAA 2008-6025.

[71] Mohammad Rezaei Mirghaed and Ramin Roshandel. Site specific optimization of wind turbines energy cost: Iterative approach. *Energy Conversion and Management*, 73:167–175, 2013. ISSN 0196-8904. doi:https://doi.org/10.1016/j.enconman.2013.04.016. URL `https://www.sciencedirect.com/science/article/pii/S0196890413002136`.

[72] Turaj Ashuri, Michiel B. Zaaijer, Joaquim R. R. A. Martins, Gerard J. W. van Bussel, and Gijs A. M. van Kuik. Multidisciplinary design optimization of offshore wind turbines for minimum levelized cost of energy. *Renewable Energy*, 68:893–905, August 2014. doi:10.1016/j.renene.2014.02.045.

[73] V. Cognet, S. Courrech du Pont, and B. Thiria. Material optimization of flexible blades for wind turbines. *Renewable Energy*, 160:1373–1384, 2020. ISSN 0960-1481. doi:https://doi.org/10.1016/j.renene.2020.05.188. URL `https://www.sciencedirect.com/science/article/pii/S0960148120308934`.

[74] Jie Zhu, Zhong Zhou, and Xin Cai. Multi-objective aerodynamic and structural integrated optimization design of wind turbines at the system level through a coupled blade-tower model. *Renewable Energy*, 150:523–537, 2020. ISSN 0960-1481. doi:https://doi.org/10.1016/j.renene.2020.01.013. URL `https://www.sciencedirect.com/science/article/pii/S0960148120300148`.

[75] Debiao Meng, Shiyuan Yang, Abílio M.P. de Jesus, and Shun-Peng Zhu. A novel kriging-model-assisted reliability-based multidisciplinary design optimization strategy and its application in the offshore wind turbine tower. *Renewable Energy*, 203:407–420, 2023. ISSN 0960-1481. doi:https://doi.org/10.1016/j.renene.2022.12.062. URL `https://www.sciencedirect.com/science/article/pii/S096014812201847X`.

[76] NM Karaolis, PJ Mussgrove, and G Jeronimidis. Active and passive aeroelastic power control using asymmetric fibre reinforced laminates for wind turbine blades. In *Proceedings of the 10th British Wind Energy Conference, London*, pages 446–458, 1988.

[77] Paul Veers, Donald Lobitz, and Gunjit Bir. Aeroelastic tailoring in wind-turbine blade applications. Technical report, Sandia National Lab.(SNL-NM), Albuquerque, NM (United States), 1998.

[78] M Hagenbeek, SJ van den Boom, NPM Werter, F Talagani, M van Roermund, and BH Bulder. The blade of the future: wind turbine blades in 2040. Technical report, TNO, Delft, The Netherlands, 2022.

[79] Carlo Luigi Bottasso, FILIPPO Campagnolo, ALESSANDRO Croce, and Carlo Tibaldi. Optimization-based study of bend–twist coupled rotor blades for passive and integrated passive/active load alleviation. *Wind Energy*, 16(8):1149–1166, 2013.

[80] CARLO LUIGI Bottasso, FILIPPO Campagnolo, ALESSANDRO Croce, S Dilli, FEDERICO Gualdoni, and Martin Bjerre Nielsen. Structural optimization of wind turbine rotor blades by multilevel sectional/multibody/3d-fem analysis. *Multibody System Dynamics*, 32: 87–116, 2014.

[81] Alejandro Albanesi, Facundo Bre, Victor Fachinotti, and Cristian Gebhardt. Simultaneous ply-order, ply-number and ply-drop optimization of laminate wind turbine blades using the inverse finite element method. *Composite Structures*, 184:894–903, January 2018. doi:10.1016/j.compstruct.2017.10.051.

[82] Evan M Anderson, Faisal Hasan Bhuiyan, Dimitri J Mavriplis, and Ray S Fertig. Adjoint-based high-fidelity structural optimization of wind-turbine blade for load stress minimization. *AIAA Journal*, 57(9), 2019. doi:10.2514/1.J057756.

[83] Christian Pavese, Carlo Tibaldi, Torben J Larsen, Taeseong Kim, and Kenneth Thomsen. Reduced design load basis for ultimate blade loads estimation in multidisciplinary design optimization frameworks. In *Journal of Physics: Conference Series*, volume 753, page 062005. IOP Publishing, 2016.

[84] Eirikur Jonsson, Cristina Riso, Christopher A. Lupp, Carlos E. S. Cesnik, Joaquim R. R. A. Martins, and Bogdan I. Epureanu. Flutter and post-flutter constraints in aircraft design optimization. *Progress in Aerospace Sciences*, 109:100537, August 2019. doi:10.1016/j.paerosci.2019.04.001.

[85] Sicheng He, Eirikur Jonsson, Charles A. Mader, and Joaquim R. R. A. Martins. Coupled Newton–Krylov time-spectral solver for flutter and limit cycle oscillation prediction. *AIAA Journal*, 59(6):2214–2232, June 2021. doi:10.2514/1.J059224.

[86] Sicheng He, Eirikur Jonsson, and Joaquim R. R. A. Martins. Wing aerodynanmic shape optimization with time spectral limit-cycle oscillation adjoint. In *AIAA Aviation Forum*, June 2022. doi:10.2514/6.2022-3357.

[87] Denis-Gabriel Caprace, Adam Cardoza, Andrew Ning, Marco Mangano, Sicheng He, and Joaquim R. R. A. Martins. Incorporating high-fidelity aerostructural analyses in wind turbine rotor optimization. In *AIAA SciTech Forum*, January 2022. doi:10.2514/6.2022-1290.

[88] Zhoujie Lyu and Joaquim R. R. A. Martins. Aerodynamic design optimization studies of a blended-wing-body aircraft. *Journal of Aircraft*, 51(5):1604–1617, September 2014. doi:10.2514/1.C032491.

[89] Ney R. Secco and Joaquim R. R. A. Martins. RANS-based aerodynamic shape optimization of a strut-braced wing with overset meshes. *Journal of Aircraft*, 56(1):217–227, January 2019. doi:10.2514/1.C034934.

[90] Anil Yildirim, Justin S. Gray, Charles A. Mader, and Joaquim R. R. A. Martins. Performance analysis of optimized STARC-ABL designs across the entire mission profile. In *Proceedings of the AIAA SciTech Forum*, January 2021. doi:10.2514/6.2021-0891.

[91] Yingqian Liao, Anil Yildirim, Joaquim R. R. A. Martins, and Yin Lu Young. RANS-based optimization of a T-shaped hydrofoil considering junction design. *Ocean Engineering*, 262: 112051, October 2022. doi:10.1016/j.oceaneng.2022.112051.

[92] Rhea P. Liem, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Multimission aircraft fuel burn minimization via multipoint aerostructural optimization. *AIAA Journal*, 53(1): 104–122, January 2015. doi:10.2514/1.J052940.

[93] David A. Burdette and Joaquim R. R. A. Martins. Design of a transonic wing with an adaptive morphing trailing edge via aerostructural optimization. *Aerospace Science and Technology*, 81:192–203, October 2018. doi:10.1016/j.ast.2018.08.004.

[94] Timothy R. Brooks, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Benchmark aerostructural models for the study of transonic aircraft wings. *AIAA Journal*, 56(7):2840–2855, July 2018. doi:10.2514/1.J056603.

[95] Timothy R. Brooks, Joaquim R. R. A. Martins, and Graeme J. Kennedy. High-fidelity aerostructural optimization of tow-steered composite wings. *Journal of Fluids and Structures*, 88:122–147, July 2019. doi:10.1016/j.jfluidstructs.2019.04.005.

[96] Nicolas Bons, Joaquim R. R. A. Martins, Felipe Odaguil, and Ana Paula Curty Cuco. Aerostructural wing optimization of a regional jet considering mission fuel burn. *ASME Open Journal of Engineering*, 1:011046, October 2022. doi:10.1115/1.4055630.

[97] Timothy R. Brooks, Joaquim R. R. A. Martins, and Graeme J. Kennedy. Aerostructural trade-offs for tow-steered composite wings. *Journal of Aircraft*, 57(5):787–799, September 2020. doi:10.2514/1.C035699.

[98] Nicolas P. Bons and Joaquim R. R. A. Martins. Aerostructural design exploration of a wing in transonic flow. *Aerospace*, 7(8):118, August 2020. doi:10.3390/aerospace7080118.

[99] Nitin Garg, Gaetan K. W. Kenway, Joaquim R. R. A. Martins, and Yin Lu Young. High-fidelity multipoint hydrostructural optimization of a 3-D hydrofoil. *Journal of Fluids and Structures*, 71:15–39, May 2017. doi:10.1016/j.jfluidstructs.2017.02.001.

[100] Nitin Garg, Bryce W. Pearce, Paul A. Brandner, Andrew W. Phillips, Joaquim R. R. A. Martins, and Yin L. Young. Experimental investigation of a hydrofoil designed via hydrostructural optimization. *Journal of Fluids and Structures*, 84:243–262, January 2019. doi:10.1016/j.jfluidstructs.2018.10.010.

[101] Yingqian Liao, Joaquim R. R. A. Martins, and Yin Lu Young. 3-D high-fidelity hydrostructural optimization of cavitation-free composite lifting surfaces. *Composite Structures*, 268: 113937, July 2021. doi:10.1016/j.compstruct.2021.113937.

[102] Yingqian Liao, Joaquim R. R. A. Martins, and Yin Lu Young. Hydrostructural optimization of single-layer and multi-layer composite lifting surfaces. *Composite Structures*, 307: 116650, March 2023. doi:10.1016/j.compstruct.2022.116650.

[103] Andrew B. Lambe and Joaquim R. R. A. Martins. Extensions to the design structure matrix for the description of multidisciplinary design, analysis, and optimization processes. *Structural and Multidisciplinary Optimization*, 46(2):273–284, August 2012. doi:10.1007/s00158-012-0763-y.

[104] Hannah M. Hajdik, Anil Yildirim, Neil Wu, Benjamin J. Brelje, Sabet Seraj, Marco Mangano, Joshua L. Anibal, Eirikur Jonsson, Eytan J. Adler, Charles A. Mader, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. pyGeo: A geometry package for multidisciplinary design optimization. *Journal of Open Source Software*, 8(87):5319, July 2023. doi:10.21105/joss.05319.

[105] Thomas W. Sederberg and Scott R. Parry. Free-form deformation of solid geometric models. *SIGGRAPH Comput. Graph.*, 20(4):151–160, August 1986. ISSN 0097-8930. doi:10.1145/15886.15903.

[106] Gaetan K.W. Kenway, Graeme. J. Kennedy, and Joaquim R. R. A. Martins. A CAD-free approach to high-fidelity aerostructural optimization. In *Proceedings of the 13th AIAA/ISSMO Multidisciplinary Analysis Optimization Conference*, number AIAA 2010-9231, Fort Worth, TX, September 2010. doi:10.2514/6.2010-9231.

[107] Edward Luke, Eric Collins, and Eric Blades. A fast mesh deformation method using explicit interpolation. *Journal of Computational Physics*, 231(2):586–601, January 2012. ISSN 0021-9991. doi:10.1016/j.jcp.2011.09.021.

[108] Mads Holst Aagaard Madsen. *High-Fidelity CFD-based Shape Optimization of Wind Turbine Blades*. PhD thesis, DTU Wind Energy, 2020.

[109] Charles A. Mader, Gaetan K. W. Kenway, Anil Yildirim, and Joaquim R. R. A. Martins. ADflow: An open-source computational fluid dynamics solver for aerodynamic and multidisciplinary optimization. *Journal of Aerospace Information Systems*, 17(9):508–527, September 2020. doi:10.2514/1.I010796.

[110] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th Aerospace Sciences Meeting and Exhibit*, June 1992. doi:10.2514/6.1992-439.

[111] Zhoujie Lyu, Gaetan K. Kenway, Cody Paige, and Joaquim R. R. A. Martins. Automatic differentiation adjoint of the Reynolds-averaged Navier–Stokes equations with a turbulence model. In *21st AIAA Computational Fluid Dynamics Conference*, San Diego, CA, July 2013. doi:10.2514/6.2013-2581.

[112] Anil Yildirim, Gaetan K. W. Kenway, Charles A. Mader, and Joaquim R. R. A. Martins. A Jacobian-free approximate Newton–Krylov startup strategy for RANS simulations. *Journal of Computational Physics*, 397:108741, November 2019. ISSN 0021-9991. doi:10.1016/j.jcp.2019.06.018.

[113] Shamsheer Chauhan and Joaquim R. R. A. Martins. RANS-based aerodynamic shape optimization of a wing considering propeller-wing interaction. *Journal of Aircraft*, 58(3): 497–513, May 2021. doi:10.2514/1.C035991.

[114] Yayun Shi, Charles A. Mader, Sicheng He, Gustavo L. O. Halila, and Joaquim R. R. A. Martins. Natural laminar-flow airfoil optimization design using a discrete adjoint approach. *AIAA Journal*, 58(11):4702–4722, November 2020. doi:10.2514/1.J058944.

[115] Yayun Shi, Charles A. Mader, and Joaquim R. R. A. Martins. Natural laminar flow wing optimization using a discrete adjoint approach. *Structural and Multidisciplinary Optimization*, 64:541–562, August 2021. doi:10.1007/s00158-021-02936-w.

[116] Gaetan K. W. Kenway, Charles A. Mader, Ping He, and Joaquim R. R. A. Martins. Effective adjoint approaches for computational fluid dynamics. *Progress in Aerospace Sciences*, 110: 100542, October 2019. doi:10.1016/j.paerosci.2019.05.002.

[117] Graeme J. Kennedy and Joaquim R. R. A. Martins. A parallel finite-element framework for large-scale gradient-based design optimization of high-performance structures. *Finite Elements in Analysis and Design*, 87:56–73, September 2014. doi:10.1016/j.finel.2014.04.011.

[118] Alasdair C. Gray and Joaquim R. R. A. Martins. Geometrically nonlinear high-fidelity aerostructural optimization for highly flexible wings. In *Proceedings of the AIAA SciTech Forum*, January 2021. doi:10.2514/6.2021-0283.

[119] Eirikur Jonsson, Charles A. Mader, Graeme J. Kennedy, and Joaquim R. R. A. Martins. Computational modeling of flutter constraint for high-fidelity aerostructural optimization. In *2019 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, San Diego, CA, January 2019. American Institute of Aeronautics and Astronautics. doi:10.2514/6.2019-2354.

[120] Jason E Hicken and David W Zingg. A simplified and flexible variant of GCROT for solving nonsymmetric linear systems. *SIAM Journal on Scientific Computing*, 32(3):1672–1694, 2010. doi:10.1137/090754674.

[121] Graeme J. Kennedy. A full-space barrier method for stress-constrained discrete material design optimization. *Structural and Multidisciplinary Optimization*, 54(3):619–639, April 2016. doi:10.1007/s00158-016-1428-z.

[122] G. Kreisselmeier and R. Steinhauser. Systematic control design by optimizing a vector performance index. In *International Federation of Active Controls Symposium on Computer-Aided Design of Control Systems, Zurich, Switzerland*, 1979. doi:10.1016/S1474-6670(17)65584-8.

[123] Andrew B. Lambe, Joaquim R. R. A. Martins, and Graeme J. Kennedy. An evaluation of constraint aggregation strategies for wing box mass minimization. *Structural and Multidisciplinary Optimization*, 55(1):257–277, January 2017. doi:10.1007/s00158-016-1495-1.

[124] Graeme J. Kennedy and Jason E. Hicken. Improved constraint-aggregation methods. *Computer Methods in Applied Mechanics and Engineering*, 289:332–354, 2015. ISSN 0045-7825. doi:10.1016/j.cma.2015.02.017.

[125] Joaquim R. R. A. Martins, Juan J. Alonso, and James J. Reuther. High-fidelity aerostructural design optimization of a supersonic business jet. *Journal of Aircraft*, 41(3):523–530, May 2004. doi:10.2514/1.11478.

[126] Neil Wu, Gaetan Kenway, Charles A. Mader, John Jasa, and Joaquim R. R. A. Martins. pyOptSparse: a Python framework for large-scale constrained nonlinear optimization of sparse systems. *Journal of Open Source Software*, 5(54):2564, October 2020. doi:10.21105/joss.02564.

[127] Philip E. Gill, Walter Murray, and Michael A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, 2005. doi:10.1137/S0036144504446096.

[128] Ella Wu. *Practical Appropriate Fidelity Optimization for Large-scale Multidisciplinary Aircraft Design*. PhD thesis, University of Michigan, 2022.

[129] Sabet Seraj, Anil Yildirim, Joshua L. Anibal, and Joaquim R. R. A. Martins. Dissipation and time step scaling strategies for low and high Mach number flows. *Journal of Computational Physics*, 491(4):112358, October 2023. doi:10.1016/j.jcp.2023.112358.

[130] P. J. Roache. Perspective: A method for uniform reporting of grid refinement studies. *Journal of Fluids Engineering*, 116(3):405–413, 1994. doi:10.1115/1.2910291.

[131] James G. Coder, David Hue, Gaetan Kenway, Thomas H. Pulliam, Anthony J. Sclafani, Leonel Serrano, and John C. Vassberg. Contributions to the Sixth Drag Prediction Workshop using structured, overset grid methods. *Journal of Aircraft*, 55(4):1406–1419, August 2018. doi:10.2514/1.C034486.

[132] Marco Mangano and Joaquim R. R. A. Martins. Multipoint aerodynamic shape optimization for subsonic and supersonic regimes. *Journal of Aircraft*, 58(3):650–662, May 2021. doi:10.2514/1.C036216.

[133] David Simms, Scott Schreck, Maureen Hand, and Lee Jay Fingersh. Nrel unsteady aerodynamics experiment in the nasa-ames wind tunnel: a comparison of predictions to measurements. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2001.

[134] Zhoujie Lyu, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. Aerodynamic shape optimization investigations of the Common Research Model wing benchmark. *AIAA Journal*, 53(4):968–985, April 2015. doi:10.2514/1.J053318.

[135] Neil Wu, Charles A. Mader, and Joaquim R. R. A. Martins. A gradient-based sequential multifidelity approach to multidisciplinary design optimization. *Structural and Multidisciplinary Optimization*, 65:131–151, April 2022. doi:10.1007/s00158-022-03204-1.

[136] M. H. Aa. Madsen, F. Zahle, S. G. Horcas, T. K. Barlas, and N. N. Sørensen. Cfd-based curved tip shape design for wind turbine blades. *Wind Energy Science*, 7(4):1471–1501, 2022. doi:10.5194/wes-7-1471-2022. URL `https://wes.copernicus.org/articles/7/1471/2022/`.

[137] M Maureen Hand, DA Simms, LJ Fingersh, DW Jager, JR Cotrell, S Schreck, and SM Larwood. Unsteady aerodynamics experiment phase vi: wind tunnel test configurations and available data campaigns. Technical report, National Renewable Energy Lab., Golden, CO.(US), 2001.

[138] Niels N Sørensen, JA Michelsen, and Scott Schreck. Navier–stokes predictions of the nrel phase vi rotor in the nasa ames 80 ft× 120 ft wind tunnel. *Wind Energy: An International Journal for Progress and Applications in Wind Power Conversion Technology*, 5(2-3):151–169, 2002.

[139] Niels N Sørensen, B Méndez, A Munoz, G Sieros, E Jost, T Lutz, G Papadakis, S Voutsinas, GN Barakos, S Colonia, et al. Cfd code comparison for 2d airfoil flows. In *Journal of Physics: Conference Series*, volume 753, page 082019. IOP Publishing, 2016.

[140] Joaquim R. R. A. Martins, Peter Sturdza, and Juan J. Alonso. The complex-step derivative approximation. *ACM Transactions on Mathematical Software*, 29(3):245–262, September 2003. doi:10.1145/838250.838251.

[141] L. Hascoët. Tapenade: A tool for automatic differentiation of programs. In *Proceedings of 4^{th} European Congress on Computational Methods, ECCOMAS'2004, Jyvaskyla, Finland*, 2004.

[142] Marco Mangano, Sicheng He, Yingqian Liao, Denis-Gabriel Caprace, and Joaquim R. R. A. Martins. Towards passive aeroelastic tailoring of large wind turbines using high-fidelity multidisciplinary design optimization. In *AIAA SciTech Forum*, January 2022. doi:10.2514/6.2022-1289.

[143] Graeme J. Kennedy, Gaetan K. W. Kenway, and Joaquim R. R. A. Martins. High aspect ratio wing design: Optimal aerostructural tradeoffs for the next generation of materials. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, National Harbor, MD, January 2014. doi:10.2514/6.2014-0596.

[144] Alasdair C. Gray, Graeme J. Kennedy, and Joaquim R. R. A. Martins. Geometrically nonlinear high-fidelity aerostructural optimization including geometric design variables. In *AIAA Aviation Forum*, San Diego, CA, June 2023. doi:10.2514/6.2023-3316.

[145] Timothy R. Brooks. *Design Optimization of Flexible Aircraft Wings Using Tow-steered Composites*. PhD thesis, University of Michigan, Ann Arbor, MI, 2018.

[146] Sabet Seraj and Joaquim R. R. A. Martins. Predicting the high-angle-of-attack characteristics of a delta wing at low speed. *Journal of Aircraft*, 59(4):1071–1081, July 2022. doi:10.2514/1.C036618.

[147] Ping He, Charles A. Mader, Joaquim R. R. A. Martins, and Kevin J. Maki. DAFoam: An open-source adjoint framework for multidisciplinary design optimization with OpenFOAM. *AIAA Journal*, 58(3), March 2020. doi:10.2514/1.J058853.

[148] Jess A Michelsen. Block structured multigrid solution of 2d and 3d elliptic pde's. Technical report, Technical University of Denmark, 1994.

[149] Bernardo Pacini, Malhar Prajapati, Karthik Duraisamy, Joaquim R. R. A. Martins, and Ping He. Towards mixed-fidelity aero-structural-acoustic optimization for urban air mobility vehicle design. In *AIAA Aviation Forum*, June 2023. doi:10.2514/6.2023-3905.

[150] Timothy R. Brooks, Graeme J. Kennedy, and Joaquim R. R. A. Martins. High-fidelity multipoint aerostructural optimization of a high aspect ratio tow-steered composite wing. In *Proceedings of the 58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, AIAA SciTech Forum*, Grapevine, TX, January 2017. doi:10.2514/6.2017-1350.

[151] Gaetan K. W. Kenway and Joaquim R. R. A. Martins. Multipoint aerodynamic shape optimization investigations of the Common Research Model wing. In *Proceedings of the AIAA Science and Technology Forum and Exposition (SciTech)*, Kissimmee, FL, January 2015. doi:10.2514/6.2015-0264.

[152] Katherine Dykes, Andrew Ning, Ryan King, Peter Graf, George Scott, and Paul S. Veers. Sensitivity analysis of wind plant performance to key turbine design parameters: A systems engineering approach. In *32nd ASME Wind Energy Symposium*, 2014. doi:10.2514/6.2014-1087.

[153] Mark Drela. *Frontiers of Computational Fluid Dynamics*, chapter Pros and Cons of Airfoil Optimization, pages 363–381. World Scientific, Singapore, November 1998. ISBN 978-981-02-3707-3. doi:10.1142/9789812815774_0019.

[154] H. Youngren M. Drela. *XFOIL 6.9 User Primer*. MIT Department of Aeronautics and Astronautics, `http://web.mit.edu/aeroutil_v1.0/xfoil_doc.txt`, Nov. 2001.

[155] P.S. Veers. Fatigue loading of wind turbines. In John D. Sørensen and Jens N. Sørensen, editors, *Wind Energy Systems*, Woodhead Publishing Series in Energy, chapter 5, pages 130–158. Woodhead Publishing, 2011. ISBN 978-1-84569-580-4. doi:10.1533/9780857090638.1.130. URL `https://www.sciencedirect.com/science/article/pii/B9781845695804500053`.

[156] Karthikeyan Ravikumar, Rajkumar Subbiah, Nalini Ranganathan, Joseph Bensingh, Abdul Kader, and Sanjay K Nayak. A review on fatigue damages in the wind turbines: Challenges in determining and reducing fatigue failures in wind turbine blades. *Wind Engineering*, 44 (4):434–451, 2020. doi:10.1177/0309524X19849851.

[157] Patrick J Moriarty, WE Holley, and SP Butterfield. Extrapolation of extreme and fatigue loads using probabilistic methods. Technical Report NREL/TP-500-34421, National Renewable Energy Lab., Golden, CO (US), 2004.

[158] G Hayman. Mlife theory manual for version 1.00. *National Renewable Energy Laboratory (NREL)*, 2012.

[159] Gunjit S Bir. User's guide to precomp, technical report no. nrel/tp-500-38929. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2005.

[160] Openfast v2.6.0, May 14, 2021. URL `https://github.com/OpenFAST/openfast`.

[161] Pietro Bortolotti, Christopher Bay, Garrett Barter, Evan Gaertner, Katherine Dykes, Michael McWilliam, Mikkel Friis-Moller, Mads Molgaard Pedersen, and Frederik Zahle. System modeling frameworks for wind turbines and plants: Review and requirements specifications. Technical report, National Renewable Energy Lab.(NREL), Golden, CO (United States), 2022.

[162] Walter Musial, Melissa Clark, Nathaniel Egging, Michael Zuteck, Walter Musial, Melissa Clark, Nathaniel Egging, and Michael Zuteck. Comparison of strength and load-based methods for testing wind turbine blades. In *35th Aerospace Sciences Meeting and Exhibit*, pages 1–13, 1997. doi:10.2514/6.1997-957. URL `https://arc.aiaa.org/doi/abs/10.2514/6.1997-957`.

[163] Herbert J Sutherland. On the fatigue analysis of wind turbines. Technical report, Sandia National Laboratory, United States, 1999. URL `https://www.osti.gov/biblio/9460`.