



# CFD-based curved tip shape design for wind turbine blades

Mads H. Aa. Madsen<sup>1</sup>, Frederik Zahle<sup>1</sup>, Sergio González Horcas<sup>1</sup>, Thanasis K. Barlas<sup>2</sup>, and  
Niels N. Sørensen<sup>1</sup>

<sup>1</sup>Aero- and Fluid Dynamics (AFD) section, DTU Wind and Energy Systems, Lyngby Campus,  
Nils Koppels Allé, Building 403, 2800 Lyngby, Denmark

<sup>2</sup>Airfoil and Rotor Design (ARD) section, DTU Wind and Energy Systems, Risø Campus,  
Frederiksborgvej 399, 4000 Roskilde, Denmark

**Correspondence:** Mads H. Aa. Madsen (mham@dtu.dk)

Received: 30 September 2021 – Discussion started: 1 November 2021

Revised: 25 March 2022 – Accepted: 4 May 2022 – Published: 18 July 2022

**Abstract.** This work presents a high-fidelity shape optimization framework based on computational fluid dynamics (CFD). The presented work is the first comprehensive curved tip shape study of a wind turbine rotor to date using a direct CFD-based approach. Preceding the study is a thorough literature survey particularly focused on wind turbine blade tips in order to place the present work in its context. Then follows a comprehensive analysis to quantify mesh dependency and to present needed mesh modifications ensuring a deep convergence of the flow field at each design iteration.

The presented modifications allow the framework to produce up to six-digit-accurate finite difference gradients which are verified using the machine-accurate Complex-Step method. The accurate gradients result in a tightly converged design optimization problem in which the studied problem is to maximize power using 12 design variables while satisfying constraints on geometry, as well as on the bending moment at 90 % blade length.

The optimized shape has about 1 %  $r/R$  blade extension, 2 %  $r/R$  flapwise displacement, and slightly below 2 %  $r/R$  edgewise displacement resulting in a 1.12 % increase in power. Importantly, the inboard part of the tip is de-loaded using twist and chord design variables as the blade is extended, ensuring that the baseline steady-state loads are not exceeded. For both analysis and optimization an industrial-scale mesh resolution of above  $14 \times 10^6$  cells is used, which underlines the maturity of the framework.

## 1 Introduction

The wind energy industry has for decades focused on minimizing the levelized cost of energy (LCoE) (Ning et al., 2014; Kalken and Ceyhan, 2017; Matheswaran et al., 2019). Innovative design has helped to keep the mass increase low as rotor size has increased to generate more annual energy production (AEP) which in turn lowers the LCoE. While manufacturing new wind turbines with diameters now exceeding 200 m helps the industry meet the present day energy demand, one should also look to already installed wind turbines which may hold promise for an increase in AEP as well. A first approach could be to completely refurbish these older

wind turbines with new rotors, although the associated costs are considerable. A less invasive operation would be to only modify the very tip of the blades with a sleeve-like solution. Indeed, using this approach design engineers have successfully met the challenge of avoiding to compromise the integrity of the original rotor and at the same time found an increase in power of up to several percent; thus, a further investigation is warranted (Zahle et al., 2018; Matheswaran et al., 2019; Barlas et al., 2021b).

In light of the above, this study investigates the promise of re-designing existing wind turbine blades by optimizing the shape of the blade tip. To maintain structural feasibility the performance gain will be brought about without exceed-

ing the load envelope. The present work should be placed in the context of the SmartTip project which in several efforts (e.g., Zahle et al., 2018; Li et al., 2018; Barlas et al., 2021b) studied innovative blade tip modeling and design, in particular how blade tips could yield a load-neutral performance gain. Zahle et al. (2018) find a 2.6 % increase in power using 12 design variables with an aerodynamic surrogate model, whereas Barlas et al. (2021b) report an up to 6 % increase in AEP using 11 design variables for an aeroelastic surrogate approach. Given that the above-mentioned works vary in included disciplines and model fidelities one should take care when comparing the final results. However, since the present study uses the exact same overall design problem which was addressed in the surrogate-based design study by Zahle et al. (2018), a comparison across flow model fidelities is given later on. The present study should in this context be seen as a contribution considering aerodynamics only but with a direct high-fidelity modeling approach using a computational fluid dynamics (CFD) solver. Thus, the aim for the present study is to test the developed CFD-based design framework using the finite difference method in an industrial-scale setting for the first time to quantify how viable this approach is and if its limitations are outweighed by the ease of implementation associated with the method. Reasons for choosing a CFD-based approach with a gradient-based optimization algorithm are given below, after which the content of the remaining paper is outlined.

Lower-fidelity methods based on the blade element momentum (BEM) theory rely on engineering assumptions and can as a result handle time-dependent simulations efficiently, which explains why they are relied on heavily in the wind energy community. High-fidelity CFD-based approaches, on the other hand, allow for investigations that do not depend on underlying engineering assumptions by solving the full Reynolds-averaged Navier–Stokes (RANS) equations on a geometrically resolved rotor configuration. Directly resolving the geometrical features at the blade tip ensures a correct modeling of the highly complex 3-D flow phenomena at the blade tip. However, also disadvantages such as an increase in computation time, as well as in implementation effort, are to be expected. Furthermore, as high-fidelity models typically are used in steady state, it is currently difficult to arrive at realistic design-driving load cases with this approach. For these reasons, one should favor a complementary use of the two approaches and only use high-fidelity methods if needed. Given that the area in question in the present study, i.e., the blade tip, is indeed an area of difficulty for lower-fidelity approaches, the use of a CFD-based model is warranted despite the increase in computation time.

To accurately describe the blade tip it was in the present study necessary to use 12 design variables. This is a considerable amount of design variables compared to many of the other tip studies mentioned below which typically use less than five. Due to the increased size of the parameter space a gradient-based optimization algorithm was preferred. More-

over, a proper step size was chosen through a gradient verification study using the Complex-Step method (Lyness, 1967; Lyness and Moler, 1967) as a machine-accurate reference gradient. The accurate gradients led to a tight optimization problem convergence.

While the presented CFD-based approach using the finite difference method is functioning well, it is not feasible to undertake a full simultaneous design of airfoils and blade planform. For this, the associated cost of computing the gradient is too high since 3-D shape optimizations of full rotor configurations involve hundreds of design variables (Nielsen and Diskin, 2012; Madsen et al., 2019). To carry out a design optimization study of a full rotor configuration using finite-difference-based gradients, one could apply a more conventional approach in which a sequential procedure in two steps should work (Barrett and Ning, 2018): first, the airfoils could be optimized with a method of choice (e.g., a panel method). Then, the planform could be optimized using the present CFD-based methodology. In practice, however, one would likely need to iterate between the two steps to arrive at a final design. See Barrett and Ning (2018) for further details and alternative approaches.

The remaining paper is structured as follows: a literature review of shape optimizations focusing on wind turbine blade tips is given in Sect. 2, whereafter the methodology used in the present study is presented in Sect. 3. Then follows a comprehensive analysis of the baseline rotor (Sect. 4) followed by a presentation of the design optimization problem (Sect. 5) before all optimization results are presented in Sect. 6. Finally, a conclusion is given (Sect. 7) in which overall findings are summarized.

## 2 Literature review

This literature review focuses exclusively on shape optimization of the blade tip from purely aerodynamic works. A literature review on general high-fidelity shape optimization of full rotor configurations within wind energy research is offered elsewhere (Madsen et al., 2019, their Sect. 2) in which an updated overview table can be found in a more recent work (Madsen, 2020, their Table 3.1). With respect to literature surveys focused on tips and winglets there are already a few present in the wind energy community (e.g., Gertz et al., 2012). However, some are several years old, and none are as comprehensive as the below-given literature survey. While some of the works mentioned below are indeed experimental, the focus is on numerical studies given that the present work is purely numerical. A literature survey including further experimental studies is presented by Mühle et al. (2020).

In order to structure the covered works the survey is split into the following sections:

- early works on wind turbine blade tips (Sect. 2.1),
- parametric studies (Sect. 2.2), and

- design optimization studies (Sect. 2.3),

before overall conclusions are presented in Sect. 2.4.

For an up front overview of all central works for the present study across the above-mentioned subsections one can consult Table 1. From this table it is possible to obtain rough estimates of, for example, what a reasonable performance increase for curved tip shapes is. It is also evident from Table 1 that only a few works on industrial-scale wind turbines exist, in which in particular actual optimization works seem to be few in number.

Care must be taken when comparing all the works listed in Table 1. Many works (Hansen and Mühle, 2018; Khalafallah et al., 2019; Khaled et al., 2019; Mourad et al., 2020; Papadopoulos et al., 2020; Aju et al., 2020) are focusing on smaller turbines, whereas only three other works (Kalken and Ceyhan, 2017; Matheswaran et al., 2019; Zahle et al., 2018) focus on larger megawatt turbines, as is the case in the present study. Furthermore, as will be evident in Sect. 2.2–2.3 there are few actual optimizations available, whereas the majority of works are parametric studies, and thus better-performing tips could very likely be found in the vicinity of the selected parameter settings for these rotor configurations. Still, Table 1 does provide the reader with an overall survey of the related available literature and should help temper the expectations with respect to the possible performance increase when optimizing the tip. Importantly, there are three works (bold rows) that manage to provide a performance-enhancing tip design that does not violate the initial load envelope. These studies can therefore not be compared directly to the studies without load constraints. Finally, it should be noted that a clear progression in level of model fidelity can be seen over time. Thus, many of the later works rely exclusively on CFD allowing researchers to analyze the finer details of the 3-D complex flow phenomena present at the blade tip regions.

## 2.1 Early works on wind turbine blade tips

The concept of a winglet can be dated more than a century back to the English engineer Frederick W. Lanchester's 1897 patent application for fixed-wing aircraft (Lanchester, 1897). However, more recent history picks up in the 1970s when Richard Travis Whitcomb (Whitcomb, 1976) further refined the idea (<https://appel.nasa.gov/2014/07/22/this-month-in-nasa-history-winglets-helped-save-an-industry/>, last access: 30 September 2021). It has been known since the pioneering days of Whitcomb that even a small winglet can limit the spanwise velocity component and reduce downwash by displacing the tip vortex which ideally diffuses or is smeared out (Whitcomb, 1976). Furthermore, with respect to the winglet orientation Whitcomb says that lower (i.e., pressure side) winglets should be as effective as upper (i.e., suction side) winglets (Whitcomb, 1976, p. 8). In aerospace, the former are typically the smaller ones due to concerns of ground clear-

ance. If anything, this is quite the opposite for wind turbines for which it is the winglet on the suction side, i.e., an “upper” winglet in aerospace terminology, which for wind turbines can cause a tower strike and therefore should be kept small. Another considerable difference is that the system of interest in wind energy is rotating, which adds some complexity to the induced velocity seen by the blade from the tip vortex. There are also other differences to study when transitioning from one sister science (aerospace) to another (wind energy). For a concise recap of essential works on winglets for fixed wings one can consult the introduction in Gaunaa et al. (2011).

About a decade after Whitcomb's 1976 study, Lissaman and Gyatt (1985) present the perhaps first comprehensive study for wing tip devices dedicated to wind turbines. As seen from their reference list a considerable amount of research on these tip devices had already been carried out in the aerospace community at this point. However, as Lissaman and Gyatt (1985) point out, no studies focus directly on wind turbines. Using both field testing and numerical analysis codes (vortex methods) they analyze three tip shapes. However, they find that none of the shapes resulted in actual performance improvement. Interestingly, they also focus on noise and report that both the standard winglet and the split winglet produce significantly more noise compared to the baseline configuration, which they attribute to added turbulent flow. For future improvements, Lissaman and Gyatt (1985) see a need for increased model fidelity in the computer code used, and, as a result, they believe the most cost-effective approach is further field or wind tunnel tests. Furthermore, means of visualizing the flow seem of interest. Not surprisingly, such techniques have more recently been heavily used, such as smoke (Mühle et al., 2020, their Fig. 6), particle image velocimetry (Aju et al., 2020, their Fig. 9), and oil (Andersen et al., 2001, their Fig. 7).

More than a decade after the earliest cited wind energy work (Lissaman and Gyatt, 1985) one will find the first significant increase in numerical model fidelity in which Madsen and Fuglsang (1997) present actual CFD-based results of novel tip designs. This is to the best of the authors' knowledge the earliest 3-D CFD-based study. Madsen and Fuglsang (1997) motivate their use of a high-fidelity CFD model by stating that the BEM-based approaches are more uncertain in the tip region due to complex flow phenomena. Using the presented methodology they are able to present a design with a non-separating tip vortex, thus lowering the tip noise. Thanks to the higher model fidelity, they can use streamlines to visualize the finer details of the flow. This allows them to align the twist angles of the winglet to match the streamlines thereby avoiding separation (Madsen and Fuglsang, 1997, their Figs. 6–11). Impressively, the final design was also tested in an experimental setup (Andersen et al., 2001) a few years later, thus lending further credibility to the results.

**Table 1.** Overview of the tip and winglet studies most relevant to the present work. Bold rows signify a reported load-neutral final design. In the case of a sequence of studies by the same author(s) a representative study was chosen. For works including both a simple straight blade extension and an actually novel tip device including a flapwise displacement, it is the latter that has been included below.

Reference	Dir. <sup>a</sup>	Fidelity	Turbine <sup>e</sup>	Mesh size <sup>d</sup>	Optimization	Design variables	Improvement	Objective
Johansen and Sørensen (2006)	++	CFD	–	$1.4 \times 10^7$	–	2	1.4 %	Power
Johansen and Sørensen (2007)	SS	CFD	–	$1.8 \times 10^7$	–	4	2.8 %	Power
Ferrer and Munduate (2007) <sup>b</sup>	÷	CFD	NREL VI	$1.2 \times 10^7$	–	2	20.1 %	Torque/thrust
Elfarrar et al. (2014) <sup>b,c</sup>	SS	Surrogate <sup>h</sup>	NREL VI	$(7.0 \times 10^5)$	✓	2	9.0 %	AEP
Aravindkumar (2014)	PS	Field test	SHWT	–	–	1	2.0 %	Power
Tobin et al. (2015)	SS	Wind tunnel	SHWT	–	–	1	8.2 %	Power
Ariffudin et al. (2016)	++	CFD	SHWT	$6.8 \times 10^6$	–	1	3.2 %	Power
Zhu et al. (2017) <sup>c</sup>	++	CFD	SHWT	$1.8 \times 10^7$	–	1	4.0 %	Power
Kalken and Ceyhan (2017)	PS	Vortex code	N80 2.5 MW	–	–	3	2%–9 %	Power
Hansen and Mühle (2018) <sup>b</sup>	SS	Surrogate <sup>g</sup>	SHWT	$(1.3 \times 10^7)$	✓	6	7.8 %	Power
<b>Zahle et al. (2018)</b>	<b>PS</b>	<b>Surrogate<sup>f</sup></b>	<b>IEA 10 MW</b>	<b><math>(6.0 \times 10^6)</math></b>	<b>✓</b>	<b>12</b>	<b>2.6 %</b>	<b>Power</b>
Reddy et al. (2019) <sup>b</sup>	++	Surrogate <sup>f</sup>	Vestas27	$(3.6 \times 10^7)$	✓	5	4.5 %	Power
<b>Matheswaran et al. (2019)</b>	<b>SS</b>	<b>Vortex code</b>	<b>NREL 5 MW</b>	–	–	<b>4</b>	<b>2.5 %</b>	<b>Power</b>
Khalafallah et al. (2019) <sup>b</sup>	++	CFD	SHWT	$5.4 \times 10^6$	–	4	4.4 %	Power
Khaled et al. (2019) <sup>b</sup>	PS	Surrogate <sup>h</sup>	SHWT	$(1.0 \times 10^7)$	✓	2	6.3 %	Power
Farhan et al. (2019)	SS	CFD	NREL VI	$1.7 \times 10^7$	–	3	9.8 %	Power
Mourad et al. (2020)	++	CFD	SHWT	$3.8 \times 10^6$	–	2	2%–6 %	Power
Sy et al. (2020) <sup>b</sup>	SS	CFD	NREL VI	$2.8 \times 10^6$	–	–	1.2 %	AEP
Papadopoulos et al. (2020) <sup>b</sup>	++	CFD	SHWT	$1.4 \times 10^7$	–	2	10.9 %	Power
<b>This study</b>	<b>PS</b>	<b>CFD</b>	<b>IEA 10 MW</b>	<b><math>1.4 \times 10^7</math></b>	<b>✓</b>	<b>12</b>	<b>1.1 %</b>	<b>Power</b>

<sup>a</sup> Whether the tip is directed towards the suction side (SS) or the pressure side (PS). When both directions were investigated, the symbol is ++, and for tips in the rotor plane without any flapwise displacement the symbol ÷ is used. <sup>b</sup> Due to their periodic boundaries the reported mesh size has been multiplied by the number of blades. <sup>c</sup> Only the amount of mesh vertices are listed, not mesh cells. <sup>d</sup> Number of cells in largest mesh used for optimization. Parentheses signify that the mesh pertains to an underlying training material. <sup>e</sup> Small horizontal axis wind turbines (SHWT) are smaller wind turbines of various configuration up to about 1 m in diameter (Ariffudin et al., 2016, report a 4 m diameter). Winglets have been reported to have effects up to 50 m inboard (Zahle et al., 2018), meaning that one should care when combining SHWT works with the remaining works. <sup>f</sup> Response surface. <sup>g</sup> Kriging model. <sup>h</sup> Artificial neural network (ANN).

Imamura et al. (1998) use a vortex lattice method with free wake modeling to compare the performance of rotors with and without winglets. Five different downstream winglets were tested, as well as a baseline blade and a pure tip extension. The numerical results (Imamura et al., 1998, their Fig. 10) show that all winglets indeed have a beneficial effect on the power coefficient,  $C_p$ , where no quantification is given. However, the pure blade extension is not far removed from the baseline performance. It should be noted that increases in bending moment are also reported (Imamura et al., 1998, Fig. 11). Again, there is no quantification. This effect is worst for straight extensions, and the bending moment increase is seen to be somewhat mitigated as the winglet increases its bending towards the downstream direction.

Already at this point, many of the key topics evident throughout the literature have been presented: it is not straightforward to gain a performance increase unless leveraging a meticulous design. Even in the cases when this is achieved, there are many other characteristics such as noise production which are important to consider. Perhaps for this very reason, many works start to focus on a select few parameters in order to narrow down the related effect for each parameter change. This trend is particularly evident in the next section.

## 2.2 Parametric studies

In the beginning of the 2000s, more works of medium- and high-fidelity models start to emerge. The vast majority of the works covered in this literature review are so-called parametric studies which also are known as design of experiments. For these works, the parameters are methodically varied in order to explore the underlying design space. Several works use the terms “optimizer” and “optimization” rather loosely, and thus, some works below (Farhan et al., 2019; Papadopoulos et al., 2020, etc.) will describe their work as optimizations. However, if the design parameter variations are chosen beforehand and there is no individual optimizer component that traverses the design space, they will in the present literature survey be categorized as parametric studies.

Johansen and Sørensen (2006) study five wind turbine configurations with different camber and twist distributions in a parameter study using CFD. All five configurations had the exact same chord distribution. Both upstream and downstream winglets were analyzed, and the latter type is found to outperform the former. They use cant angle and sweep angle to describe their winglet geometries, in which cant angle refers to the angle of incidence the flapwise displaced part (i.e., the out-of-plane part) of the tip makes with a straight blade reference axis (i.e., a 0° cant angle is a straight blade,



and  $90^\circ$  is a true winglet). Similarly, the sweep is the incident angle for the part of the tip that lies in the rotor plane. In this study cant angle ( $90^\circ$ ) and sweep ( $0^\circ$ ) are not altered but merely used to describe the designed winglets. The various tested twist and camber settings result in power increases of 0.6%–1.4%, whereas trust increases 1.0%–1.6% for winglets of about 1.5% height.

Not long after, Johansen and Sørensen (2007) again carry out a CFD-based parameter study, this time showing that the tested winglets bring about a 1.0%–2.8% increase in power at the cost of a 1.2%–3.6% increase in thrust. The varied winglet parameters are height, curvature radius, sweep, and twist, having the following definitions. Height refers to the distance that the winglet protrudes in the out-of-plane region (i.e., the distance to the rotor plane). Curvature radius is a related measure given in percentage of the winglet height which describes how smooth the transition from straight blade to winglet occurs (0% means a  $90^\circ$  kink and 100% means a very smoothly transitioning blade that first at the very tip of the winglet reaches the maximum projected blade length). In light of their previous results (Johansen and Sørensen, 2006) they only test downstream winglets. They test a total of 10 different rotor configurations. In relation to the present study it is relevant to mention that they find no effect on power for sweep and that only limited effect from changing the twist can be observed.

Ferrer and Munduate (2007) also use CFD and analyze three tip configurations for the NREL Phase VI wind turbine rotor. The three configurations are a square tip, a highly tapered tip where the tip ends at the pitch axis, and a highly tapered tip where the trailing edges were aligned resulting in a swept-back tip. Here, the pitch axis should be taken as the quarter-chord location. They find that both tapered tips outperform the rectangular tip which has a large tip vortex. Of the tapered tips it is the configuration with its tip on the pitch axis that has the best torque to thrust ratio. Furthermore, they specifically point to a complementary use of CFD in lower-fidelity design approaches.

Gaunaa and Johansen (2007) show that the increase in  $C_p$  resulting from a winglet is owed to a reduction in tip effects (i.e., tip loss) and not as previously thought due to a downstream shift in the wake vorticity. However, in the same work a comparison between the developed free wake lifting line model and a CFD reference was not entirely successful leading to a follow-up study (Gaunaa and Johansen, 2008) in which said comparison was improved. In both works they advocate for downstream winglets which they find to be more efficient than their upstream counterparts.

A few years later, Gaunaa et al. (2011) used computationally lighter models based on lifting line theory to analyze blades with winglets. The tested models are a free wake model and a much faster prescribed wake model. The study relates highly to previous work by the authors (Gaunaa and Johansen, 2007) in which the free wake model was used. Moreover, the validation of the developed prescribed

wake model is carried out against CFD-based results from Johansen and Sørensen (2006). By comparing results from both the free wake model and the prescribed wake model to CFD results, they conclude that these faster model types successfully can predict the effect of adding winglets to wind turbine rotors (errors: 4%–16% and 17%–28%, respectively; Gaunaa et al., 2011, their Table 2). Considering how well the wake models approximate the results from the CFD solver it would be very interesting to see these codes applied in an optimization context in future works.

Using high-fidelity models Sørensen et al. (2011) sought a deeper understanding of the underlying flow phenomena related to several tip shapes. This is, for example, a necessity when a thorough understanding of the generated noise level is needed. Furthermore, they used access to both high- and lower-fidelity models to carry out a detailed comparison across fidelities. This allowed Sørensen et al. (2011) to implement an improved tip correction for lifting line models.

Aravindkumar (2014) use experimental field tests to compare with a CFD model in order to investigate performance increase and noise reductions. They find that adding an upstream winglet increases the generator power by 2.0%, whereas the noise is reduced by 25%.

Tobin et al. (2015) use wind tunnel experiments to compare one rotor configuration without a winglet with a rotor configuration including a winglet and observe a  $C_p$  increase of 8.2%. A related thrust (i.e.,  $C_t$ ) increase of 15.0% was observed. Having studied the literature they see a need for further insight in wake performance enhancement resulting from winglet design. They find that winglets do not significantly alter tip vortex strengths.

Ariffudin et al. (2016) use a CFD model to investigate four different configurations; a straight extension, a swept extension that lies in the rotor plane, and configurations with winglets directed either upstream or downstream. Between the two first configurations it is the swept configuration that outperforms the straight configuration with a 9.1% and a 7.3%  $C_p$  increase, respectively. As for the winglets it is the downstream configuration that results in the highest performance increase at 3.2% compared to just a 1.8% increase for the upstream winglet. No quantification of similar changes in the thrust coefficient,  $C_t$ , is offered.

Zhu et al. (2017) carry out a study aiming at determining the best direction for a winglet. Upstream, downstream, and even split winglets are tested. The split winglet pointing both upstream and downstream is found to be the best performing device, and a  $C_p$  increase of up to 4.0% is observed. Comparing just the pressure side and suction side winglets the former outperforms the latter since the pressure side winglet results in an up to 3.8% increase, whereas the suction side winglet only reaches a 3.4% increase (Zhu et al., 2017, their Table 5). They also conclude that the angle of the winglet should match the incoming flow angle as much as possible (something focused on already in the very early works referenced above; Madsen and Fuglsang, 1997). Finally, they

mention that an actual optimization would further increase the performance.

Kalken and Ceyhan (2017) study three different tip concepts (turbulators, winglet, and conventional tip) and submit the final designs to experimental testing on a 2.5 MW wind turbine for a final validation. The early design phase leverages BEM and lifting line models combined with free wake models, whereas CFD is used as an analysis tool on the finalized tip design. Kalken and Ceyhan (2017) report based on measurements that a more than 4 % power increase is owed to a simple blade extension, whereas the benefit of using different blade tip shapes result in a 2 %–9 % power increase. The CFD analysis shows that a conventional extension results in a higher power to thrust ratio gain than the studied winglet. However, they point to other beneficial side effects for choosing a winglet (noise reduction, height restriction, etc.).

The work by Matheswaran et al. (2019) is particularly interesting as it is one of the rare studies with a reported load-neutral tip design. This recent work has even resulted in a patent ([https://patentscope.wipo.int/search/en/detail.jsf?docId=US242151286&recNum=7&docAn=16186876&queryString=\(IC/F03D\)&maxRec=92796](https://patentscope.wipo.int/search/en/detail.jsf?docId=US242151286&recNum=7&docAn=16186876&queryString=(IC/F03D)&maxRec=92796), last access: 30 September 2021, USPTO. Application number 16186876), stressing the industrial relevance of novel tip designs that do not compromise the load envelope for the baseline wind turbine. Matheswaran et al. (2019) motivate their study by stressing the point that most tip designs for wind turbines do not focus on, for example, the added flapwise bending moment. As a result, the initially minimally invasive operation of retrofitting the very tip of the blade may become an intractable proposition altogether. By balancing the centrifugal force with the aerodynamic forces generated by the winglet itself, they present a lightweight winglet which focuses on minimizing the bending moments. The presented design methodology is based on the vortex lattice method which, unlike traditional BEM approaches, manages to model the flow at the very tip of the blade. The model has a prescribed wake, although Matheswaran et al. (2019) do point out that a “free wake” approach would be ideal, albeit also computationally more demanding. After having validated their vortex code against both experimental and numerical results they introduce an LCoE cost model and carry out a parametric study using four parameters (height, taper ratio, and twist and cant angles). Only suction side winglet designs are investigated. Having studied various configurations they decide on a design with the best compromise between a respectable increase in  $C_p$  (2.5 %), while the force ratio between aerodynamic loads and centrifugal forces is kept manageable (1.4–1.8). They carry out a structural analysis to prove the designs can withstand the required loads. For future work they point to the need for higher model fidelity for which they specifically mention using a CFD model.

Khalafallah et al. (2019) present a parametric study of winglets’ possible effect on wind turbine power production. Both swept and straight blades are tested using a CFD model with up to 1.8 million cell meshes with periodic boundary conditions modeling only one of three blades for the rotor. Upstream and downstream winglet configurations with various cant angles and twist angles are tested, and the best result is achieved for the upstream configuration with a 4.4 % power coefficient increase. Using results from a previous study they select a few well-performing swept baseline blades and investigate the effect of adding winglets on the resulting power coefficient. A total of 18 straight blade configurations and 33 swept blade configurations were simulated. In general, they find that downstream winglets outperform upstream winglets for straight blades (Khalafallah et al., 2019, their Table 3). However, for swept blades (Khalafallah et al., 2019, their Tables 4–5) the conclusions are more ambiguous. Indeed, the best swept winglet design is an upstream-directed winglet yielding a 4.4 % increase in  $C_p$ . Overall, they conclude winglets can indeed be used to enhance rotor performance and that they may as well lead to a reduced thrust coefficient. It should be noted that a reduction in thrust coefficient is only observed in the comparison to a swept baseline blade (Khalafallah et al., 2019, Fig. 13), whereas it is quite clear that the coefficient of thrust for the winglet coefficients in general increase compared to the straight blade baseline.

Farhan et al. (2019) investigate winglet planform and airfoils’ importance for winglet design using CFD. They vary height, cant angle, and planform using two different airfoils and find that the best performance increase is owed to about a distance equivalent to 3 % of the blade length, which from now on will be written as 3 %  $d/R$  for brevity. The configuration had a rectangular winglet with a 45° cant angle. To choose the best turbulence model they start out by comparing two RANS models, i.e., the Spalart–Allmaras model and the shear stress transport (SST) model by Menter (1992), and find that the latter outperforms the former in emerging stall regimes. Finally, they test two airfoils and vary cant angle and winglet length (24 combinations) and report a –13.5 % to +9.8 % change in power depending on the wind speed and configuration (Farhan et al., 2019, their Table 3). The tested winglets were pointing in the downstream direction. Having chosen winglet height and cant angle, they then fix the airfoil shape to S809 and study planform effects (rectangular versus elliptical winglets). They find that both winglet length and cant angle are amongst the most important design variables for improving performance.

Mourad et al. (2020) use an initial literature survey covering 10 references to conclude that (i) there is no agreement on optimum winglet configuration, (ii) height is the most effective winglet parameter, and (iii) an upstream-directed winglet should be preferred over downstream configurations. Since they found no study covering toe angle, i.e., the angle of attack between tangential velocity component and winglet profile, they carry out a parameter study using height and toe

angle as their two parameters. The analysis showed that of the winglet heights ranging from 0.8 % to 8.0 %  $d/R$  it was the smaller winglet height of 0.8 %  $d/R$  that gave the largest power increase. Furthermore, Mourad et al. (2020) report that a downstream-directed winglet is not useful for wind turbine rotors. The poor performance of the winglet with 8 % height is not easily aligned with previously reported results by Gertz and Johnson (2011) and Gertz et al. (2012) stating a 5 % increase in power for that winglet height – albeit for different overall parametrization. Finally, it is found that from toe angles ranging from  $-30$  to  $+30^\circ$  it is the (upstream-directed)  $+20^\circ$  toe-angled winglets that have the most beneficial effect, in which a 2 %–6 % increase in power coefficient is observed depending on the tip speed ratio. However, a related increase in thrust coefficient of 4.6 %–9.8 % is reported, meaning that the initial load envelope is compromised.

Kulak et al. (2021) present both experimental and numerical results in a study of a small wind turbine rotor (20 cm radius) in which the aim is to raise overall power output. Wind tunnel tests of configurations with and without winglets were carried out to quantify the differences. In total, one baseline configuration, one suction side winglet (4 % height), and three pressure side winglets (3 %, 4 %, and 5 % height) were tested. The numerical results for the comparison were only generated on the pressure side winglet with 4 % height, as well as for the baseline rotor. A nice detail in this study is that they add a transitional model to increase model fidelity. In agreement with Gupta and Amano (2012) an increase in  $C_p$  for particularly pressure side winglets is observed in which an up to 6 % increase is reported. No efficiency increase is seen for suction side winglets. Instead, a decrease in  $C_p$  is observed compared to the baseline performance. Relating the experimental results to the numerical findings shows a misalignment since the numerically investigated pressure side winglet performs worse than the baseline rotor. Thus, experimental results do not agree with numerical results. They conclude that the precise winglet geometrical features must be carefully defined in order to gain the desired performance increase.

Sy et al. (2020) use CFD to study a split winglet design and its ability to lower the induced drag. In total, four different designs are analyzed (baseline, straight, suction side winglet, and split winglet). The three modified meshes all had a 1.5 %  $r/R$  extension, and winglets had a  $45^\circ$  angle offset. They observe a 1.23 % and a 2.53 % increase in power for ordinary winglet and split winglet, respectively. However, also thrust increases by 0.83 % and 2.05 % for the designs in question.

Mühle et al. (2020) investigate the promise of using winglets to enhance wake recovery using an experimental setup. These types of studies view a wind turbine not only as an isolated system but also as part of a whole, for example, a wind farm. Thus, improved wake recovery in one turbine will lead to a performance increase in the next one. The investigated rotor is a two-bladed model-scale rotor in which the wing tip can be exchanged with a downstream winglet

tip to compare performance characteristics. This is the same rotor which was investigated by Hansen and Mühle (2018). The wind tunnel measurements show that a winglet not only can be used to increase power production but also to provoke earlier tip vortex interaction resulting in a faster wake recovery.

Papadopoulos et al. (2020) perform a numerical analysis of six rotor configurations of a small rotor. They set twist and toe angle to  $0^\circ$  on the basis that available literature point to a minimal effect. The winglet height was set to 5 % blade radius. They find that the cant angles of  $\pm 45^\circ$  are better than the  $\pm 90^\circ$  angles and that sweep only has a limited effect. However, as they also document (Papadopoulos et al., 2020, Fig. 2) the corresponding projected in-plane area of the rotor is slightly increased for  $\pm 45^\circ$  angles, which should give these configurations an advantage. A nice detail in this study is the use of a four-equation transition model allowing for finer flow modeling than assuming a fully turbulent flow. The maximum observed effect was an 11 % increase in  $C_p$ . For the configurations that did not increase the projected area the increase in  $C_p$  (3.7 %) was more modest. Unfortunately, the impact on thrust is not reported in this study. With respect to the discussion on suction side versus pressure side winglets it is worth mentioning that the above-mentioned two configurations are pressure side and suction side winglets, respectively ( $+45$  and  $-90^\circ$ ). They point to future studies on additional parameters in order to further the understanding of winglet parametrization.

Aju et al. (2020) use an experimental setup to investigate the promise of using downstream winglet pitching as a means to lower turbine rotation and reduce thrust coefficient. This has relevance when, for example, wind turbines should be protected during extreme weather conditions. Given that the winglet mass only makes up 1.8 % of the blade the investigated method should provide a much faster response time compared to pitching the entire blade. Also, winglet pitching is shown to accelerate flow recovery in wake regions. Aju et al. (2020) agree with Mühle et al. (2020) in that there is great promise for winglet use in wake recovery.

Of the numerous parameter studies mentioned above many salient winglet features such as power enhancement and load mitigation have been investigated. However, several contradicting works were also found. Much can be owed to a misalignment in the studies, and one efficient way to unify the efforts would be to agree on design optimization problems to solve for. Not surprisingly, the need for actual optimizations is also brought up in some of the works (e.g., Zhu et al., 2017). In the next section these few but important works in the literature are discussed.

### 2.3 Design optimization studies

Design optimization studies often demand meticulous implementation with a great attention to detail, making them few in number in the wind energy literature. Still, some can be

found. In fact, the very first work including CFD (Madsen and Fuglsang, 1997) also includes both single-point and multipoint shape optimizations. However, most of the optimization studies are by far more recent works carried out within the last 5 years as detailed below.

From 2011 to 2015 a series of studies emerge in which CFD is used to investigate possible winglet shapes on the NREL Phase VI rotor which has a 10 m diameter (Elfarra, 2011; Elfarra et al., 2014, 2015). Given that one of the studies (Elfarra et al., 2015) does not contain an actual optimization the focus will be placed on the other two works. In both studies containing optimizations (Elfarra, 2011; Elfarra et al., 2014) a direct CFD-based approach is deemed too computationally expensive, and surrogate models in the form of artificial neural networks are favored along with a gradient-free genetic optimization algorithm. Using 24 CFD evaluations to train the artificial neural network they are able to carry out winglet optimizations using cant and twist angle as design variables and obtain about a 9% increase in power. A multipoint objective is targeted which gives a more robust design. However, as seen in Table 1 it has a rather coarse mesh resolution. While this may be due to the fact that it is indeed an early work it is fair to contemplate whether the used meshes are of adequate resolution as training material for the surrogate. After the optimization, a comparison of the flow characteristics for the baseline winglet design and the final winglet design reveal that the new design manages to attach the flow farther outboard, resulting in the reported improvement. The ensuing parameter study (Elfarra et al., 2015) focuses on power enhancement using 16 winglet configurations to study how winglet direction, sweep, cant angle, and twist influence the generated power. After having validated the numerical setup against experimental data, they test the 16 configurations, whereafter they conclude that suction side winglets generate more power than pressure side winglets. Depending on configuration and wind speed they observe an up to 10.43% power increase (Elfarra et al., 2015, their Table 6). Interestingly, this particular configuration did not have a 90° winglet but a 45° winglet. The winglet was twisted +2° (towards lower angle of attack). In general, this study sequence exemplifies how high-fidelity models such as CFD solvers can be used to visually inspect the flow and analyze the underlying flow phenomena at play.

Another sequence of relevant works is that by Hansen (2017), who applies evolutionary optimization algorithms in winglet design. The two works, Hansen (2018) and Hansen and Mühle (2018), focus on wind turbine airfoil and winglet design, respectively. The resulting optimization framework from the airfoil study (Hansen, 2018) is used in the final winglet study (Hansen and Mühle, 2018) to prepare the new airfoils needed for the ensuing wind tunnel tests. In the winglet optimization study (Hansen and Mühle, 2018), a Kriging surrogate model is trained on CFD evaluations of a model-scale wind turbine with a winglet. The optimization involves six design variables, and a total of 100 shapes

are traversed. Finally, the wind turbine performance with and without winglet is validated using a 3-D printed experimental model in the NTNU wind tunnel (<https://www.ntnu.edu/ept/laboratories/aerodynamic>, last access: 30 September 2021). Power is numerically predicted to increase 7.8%, and a subsequent wind tunnel experiment of the winglet reports a 8.9%–10.3% power increase depending on the inflow turbulence. The 7.8% power increase came at the cost of a related 6.3% increase in thrust. Hansen and Mühle (2018) point out that a better shape could be obtained by including more design variables to further refine the design parameter space.

Zahle et al. (2018) carry out a surrogate-based tip study in which the aim is to increase AEP without increasing the load envelope for the baseline blade considering only steady-state normal operating conditions. It is important to clarify that they optimize the tips for improved AEP, but the final results are reported as improvements in power. Using a gradient-based approach to efficiently manage the 12 design variables, they achieve a 2.6% and a 0.76% power increase for winglet-like and straight tip extensions, respectively. These numbers are not far off compared to the report made several decades before by Whitcomb (1976) stating that a winglet can result in a lift-to-drag ratio improvement more than twice that of a straight blade extension (Whitcomb, 1976, p. 13). The surrogate-based approach allows for a very efficient optimization process, although an added difficulty is that the final design may prove to (slightly) violate the bending moment for the underlying CFD model. This study has the same design optimization problem as is used in the present CFD-based design study which is further explained in Sect. 5. Zahle et al. (2018) also study the effect that each of the three tip-dedicated design variables have on the mechanical power and find that the achievable mechanical power improvement should increase as curvature, flapwise displacement, and edgewise displacement (i.e., sweep) of the tip are increased. Noticeably, the underlying CFD model predicts a maximum for the sweep design variable around 2%  $d/R$  (Zahle et al., 2018, their Fig. 4). Interestingly, the efficient surrogate procedure allows them to explore the Pareto curve between bending moment increase and increase in mechanical power. Their mesh resolution of 5.97 million cells combined with the 300 sampling points allow for a detailed analysis of the parameter space, and the surrogate exhibited below 2% error for both torque and bending moment. Their overall workflow includes tools for surface and volume mesh generation besides the exact same flow solver, EllipSys3D (Michelsen et al., 1992, 1994; Sørensen, 1995), as is used in the present study. All investigated tip shapes pointed upstream towards the pressure side of the blade to avoid a tower strike. Overall, Zahle et al. (2018) find that the CFD evaluations agree fairly well with the surrogate-based model. Furthermore, they conclude that a winglet-like extension should be favored over a straight extension, although the advantage is attenuated for higher wind speeds (Zahle et al., 2018, Fig. 6). Consequently, the predicted 2.6% increase in



mechanical power is valid for the lower wind speeds (6 and  $8 \text{ m s}^{-1}$ ). Finally, they explain how the tip shape in itself does not drastically increase power production but that its role instead is to diffuse and move the tip vortex further away in order to lower the induced drag.

Reddy et al. (2019) use high-fidelity CFD evaluations to train a surrogate response surface and carry out a shape optimization study resulting in a  $C_p$  increase of 4.5 % while only introducing a minor thrust force penalty. Before undertaking the actual optimization they validate their numerical setup by comparing it to experimental results. The defined objective is a compound function including coefficient of power,  $C_p$ , coefficient of thrust,  $C_t$ , and the twisting moment around the blade axis. They use five design variables: span, twist, dihedral, and sweep angle, as well as taper ratio. As stated by the authors, it seems to be the first published multi-objective optimization study for wind turbine winglets. Indeed, as seen in Table 1 only very few actual optimizations can be found in the related wind energy literature. Reddy et al. (2019) state that they use the much faster surrogate methods since an actual CFD evaluation of each configuration is too expensive. To train their surrogate 50 initial blade designs are used. In comparison, Zahle et al. (2018) use 300 CFD evaluations to train their surrogate. However, Zahle et al. (2018) also use 12 design variables and as a result would have to use more evaluations to train their surrogate since the underlying parameter space increases in size from 5 to 12 design variables. Based on their study, Reddy et al. (2019) conclude that winglets can increase power production and that their optimization framework is capable of handling the multidimensional design.

Khaled et al. (2019) use a parameter study to investigate the performance effect of winglet length and cant angle for a small wind turbine and find that both power and thrust coefficients increase when a winglet is present. Using a low-speed wind tunnel they fabricate and test six rotor configurations and are able to quantify the error ( $\sim 5\%$ ) between their computational setup and the experimental data. Subsequently, an artificial neural network is used to predict the best winglet shape which has a 6 % winglet length and a cant angle of  $48^\circ$  resulting in a 9 % increase in both power and thrust coefficient.

In summary, very few actual optimizations of wind turbine winglets have been found. While it in principle is fine to compare (i) across model fidelity and (ii) between gradient-based and gradient-free optimization procedures, one can already identify one major issue in doing so: the cited works rarely quantify how well the design optimization problem is solved, and in not doing so one can basically not know if the problem is solved at all.

## 2.4 Overall trends in the covered literature

Table 1 should mainly be seen as an overview of numerical studies within wind energy of rotors fitted with winglets, and correspondingly there are only a few studies (e.g., Tobin

et al., 2015) with experimental results, whereas the remaining works include numerical investigations. Turning to the nature of the cited studies one will find that 12 are shape analyses, whereas only 6 include actual shape optimizations. It is fair to state that a chronological trend of transitioning from parameter studies to optimization studies can be observed since by far most optimization studies are very recent. For the parameter studies it is likely that better shapes could be found, making it difficult to compare these with the optimization works. Still, it is possible to give several general statements which can be found below.

First of all there is a general consensus that winglets do indeed provide a promising means of increasing rotor performance. Some of the most studied effects are

- an increase in power production (Lissaman and Gyatt, 1985; Imamura et al., 1998; Zahle et al., 2018; Hansen and Mühle, 2018; Sy et al., 2020),
- noise reduction (Lissaman and Gyatt, 1985; Madsen and Fuglsang, 1997; Aravindkumar, 2014; Ebrahimi and Mardani, 2018), and
- accelerated wake recovery (Tobin et al., 2015; Kalken and Ceyhan, 2017; Aju et al., 2020; Mühle et al., 2020).

Starting with discussing the role of the winglet it is quite clear that there is a consensus in the literature. The underlying physical principle of a winglet is to mitigate the induced drag which is introduced on any loaded wing close to the tip where a spanwise velocity component flows from the pressure side to the suction side. This spanwise flow results in a tip vortex which reduces the lift force. By tailoring the winglet one can manipulate the spanwise flow and change how and where the tip vortex occurs. The winglets' role is therefore not necessarily in itself to increase power production locally but merely to transport the tip vortex further away to lower the induced drag which in turn raises the power production further inboard (Zahle et al., 2018; Hansen and Mühle, 2018; Sy et al., 2020; Mourad et al., 2020).

Also the direction of wind turbine winglets has been discussed in numerous studies. Several works (Zhu et al., 2017; Khalafallah et al., 2019; Kulak et al., 2021) have found the best performance increase for upstream-directed winglets. However, it has also been reported that downstream winglets are most efficient (Johansen and Sørensen, 2006; Bak et al., 2007; Gaunaa and Johansen, 2008; Ariffudin et al., 2016; Khalafallah et al., 2019) and that these should be relatively short (Bak et al., 2007). In the latter case it should be noted that even short winglets may raise concern on tower clearance (Johansen and Sørensen, 2006). Indeed, tower strike is an important aspect once transitioning from academic exercises to actual industrial applications. The most straightforward conclusion to the above apparent contradiction on winglet direction is that it can be attributed to differences in parametrization and flow model fidelity. Furthermore, it is

likely that the optimal downstream winglet and the optimal upstream winglet look different and therefore should have different parametrizations.

Indeed, taking in the whole body of work mentioned above it can be seen that the parametrization of winglets and novel tip shapes vary greatly. Parametrizations of 2 to 12 design variables have been reported. In general, the most typical considered design variables in Table 1 are winglet height, twist, sweep, cant angle, toe angle, extension, and curvature. For the higher numbers reported, for example, the 12 design variables used by Zahle et al. (2018), it is typical that the majority (9) of the design variables are used as planform-type design variables controlling, for example, the twist and chord distribution of the very tip of the blade. However, only three works (Hansen and Mühle, 2018; Zahle et al., 2018; Reddy et al., 2019) manage to take more than a few design variables into account simultaneously, and it is fair to speculate whether a mere two to four design variables indeed are sufficient to accurately model the full complexity of a winglet. That being said one can deduce several rules of thumbs for the use of design variables: one should, for example, expect the optimizer to leverage flapwise displacement to decrease bending moment (Imamura et al., 1998) or manipulation of the thrust coefficient (Khalafallah et al., 2019). Flapwise displacement can also be found to be favored slightly over sweep (Zahle et al., 2018, their Fig. 4), but both are used heavily to displace the tip vortex in numerous works (Hansen and Mühle, 2018; Sy et al., 2020; Mourad et al., 2020). Few works also include twist and chord design variables for the winglet planform where the expected trends of a reduced chord distribution towards the tip and an initial increase in twist distribution followed by a sharp decrease at the very end of the blade can be observed (Zahle et al., 2018). However, these trends for twist and chord probably are difficult to generalize and are likely to be specific to pressure side winglets.

Turning to the ability of the winglet to increase overall performance the results in Table 1 vary from an increase of 1.1 % to 20.1 % depending on objective function and parametrization. However, not all investigated shapes were reported to be load-neutral compared to the baseline, which could compromise already installed rotors. One should therefore only compare load-neutral works (bold rows in Table 1) or clearly state by how much the initial load envelope is allowed to be exceeded in order to arrive at meaningful comparisons. Based on the load-neutral works in Table 1 it is fair to state that only a few percent improvement can be expected depending on the choice of objective function.

Despite the above-mentioned salient features recent publications (Reddy et al., 2019; Papadopoulos et al., 2020) agree that still only a limited use in the wind energy industry has been found for winglet applications. Moving forward, at least two trends can be pointed out to help spread the role of the winglet and to counter remaining contradictions found in the literature.

- Firstly, action should be a general increased model fidelity when studying a feature such as the winglet. Numerous works (Matheswaran et al., 2019; Døssing, 2007, etc.) point to the need of an increase in model fidelity when studying the tip area where conventional BEM models struggle. Using CFD and other high-fidelity references also allows for improving said lower-fidelity models, which has already been reported in several studies (e.g., Sørensen et al., 2011). The increase in fidelity would also help capture the complex flow phenomena to better understand the role of the winglet.
- Secondly, one should aim for optimizations with unifying design optimization problems. By agreeing on unifying design optimization problems solved by the various frameworks one can eliminate some of the differences separating these works, and more consensus should be possible.

The novelty and contribution of the present study is to develop a CFD-based design optimization framework and to optimize a wind turbine blade tip, thereby addressing the needs identified in the presented literature review, i.e., (i) to study the tip using a high-fidelity model (CFD) and (ii) to carry out an optimization on a unifying design optimization problem. To this end, the remaining paper has sections for “Methodology” (Sect. 3), “Baseline analysis” (Sect. 4), “Design optimization problem” (Sect. 5), “Results” (Sect. 6), and “Conclusion” (Sect. 7).

### 3 Methodology

The overall aim of this study is to optimize the geometry of an academic wind turbine considering rotor aerodynamics only, and all aeroelastic effects will be disregarded. Constraints on load and geometry are used as a surrogate to maintain structural feasibility. This section describes the overall design framework used in the study, and the remaining sections will describe the optimization problem in further detail, as well as present the final results. Below, a general description of the framework is given (Sect. 3.1). Then, a description of the components – deformation library (Sect. 3.2), flow solver (Sect. 3.3), and optimizer (Sect. 3.4) – is given.

#### 3.1 FlowOpt: a high-fidelity shape optimization framework

The high-fidelity shape optimization framework called “FlowOpt” at DTU Wind Energy is built around the in-house flow solver, EllipSys3D. The framework is focused on gradient-based shape optimization, and in this study two different step-based approaches will be used to compute the gradient, namely the finite difference method and the Complex-Step method. The Complex-Step method can provide machine-accurate gradients, and up to 16 digits of gradient accuracy have been reported for the implementation in

EllipSys3D (Madsen, 2020, their Fig. 7.9). This method will therefore be used in a gradient verification of the finite difference method gradients. However, it was observed that a few meshes (of the several hundred encountered throughout an optimization) did not lead to a deep convergence, resulting in impaired Complex-Step gradient accuracy for these few meshes. For this reason it was decided to use the finite difference method in the present study during the actual optimizations.

All components in the FlowOpt framework are written in compiled low-level programming languages for maximum efficiency. However, these components also have user-friendly interfaces allowing for a lenient interaction using the interpreted high-level Python programming language. It is through these interfaces that the Python-based FlowOpt optimization framework is built. A visualization of the FlowOpt framework can be seen in Fig. 1.

The design framework visualization in Fig. 1 starts in the upper-left corner where an optimizer (0, blue) sends a set of design variables,  $\mathbf{x}_{DV}$ , to a mesh deformation library (1, red) which in turn sends a resulting deformed CFD mesh surface,  $\mathbf{x}_{surf}$ , to a flow solver (2, green) which finally can compute a flow field,  $\mathbf{w}$ , functions of interest,  $\mathbf{f}$ , and constraints,  $\mathbf{c}$ . All functions of interest and constraints are then together with the related gradients returned to the optimizer (3, blue), and the overall cycle is repeated until a final design,  $\mathbf{x}_{DV}^*$ , is identified by the optimizer.

The choice of gradient computation method in the design framework has great impact on which optimization problems can be solved in a timely manner, but the actual course an optimizer takes during an optimization should not change. Indeed, it has previously been shown (Madsen, 2020, their Fig. 11.11) that the FlowOpt framework will carry out essentially identical optimizations when using either the Complex-Step method or the adjoint method as long as the flow field is well-converged to ensure that the gradients are machine-accurate.

Returning to Fig. 1 it can be noted that there are several layers to the mesh deformation component (1, red) and flow solver component (2, green). This signifies that the optimization may involve several simultaneous flow computations. As a result, the design framework supports a nested parallelism (through OpenMDAO; Gray et al., 2019) both with respect to running multipoint optimizations but also with respect to gradient evaluation. Thus, if sufficient computational resources are available, one can dedicate a separate group of CPUs for each design variable and maintain a fixed gradient computation time. This allows for an execution time comparable to those seen for adjoint-based optimizations. However, unlike the adjoint method the cost in computational resources will in this case increase linearly with the number of design variables being evaluated simultaneously. The presented study will therefore be rather costly in terms of CPU usage to ensure a state-of-the-art computation time.

### 3.2 FFDlib: a free-form deformation library

The in-house free-form deformation library, FFDlib (Madsen, 2020, their chap. 5), has been developed as an integral part of the FlowOpt design optimization framework and has previously been used in trailing edge flapping device studies (Horcas et al., 2018), as well as in high-fidelity shape optimization studies with various gradient computation techniques (adjoint, Complex-Step, etc.) as described elsewhere (Madsen, 2020). As the name suggests the parametrization library leverages a free-form deformation (FFD) formulation to propagate changes from a few chosen design variables out to every single embedded mesh point in the computational mesh.

The FFD methodology has numerous salient features and was chosen particularly due to the following three aspects:

- exact numerical mesh representation (if the inverse search is converged to machine precision),
- mesh topology agnostic, and
- analytical gradients.

As explained in the original FFD paper (Sederberg and Parry, 1986) the basic principle in free-form deformation is to embed an object in a rubber-like material, meaning that for the present study the tip of the blade will be embedded in deformation boxes. Mathematically speaking, an inverse search is used to map the discrete mesh points to the normalized parameter space spanned by the tuple coordinates,  $(s, t, u)$ , according to the following equation:

$$\mathbf{X}_{FFD}(s, t, u) = \sum_{i=0}^l \binom{l}{i} (1-s)^{l-i} s^i \left[ \sum_{j=0}^m \binom{m}{j} (1-t)^{m-j} t^j \left[ \sum_{k=0}^n \binom{n}{k} (1-u)^{n-k} u^k \cdot \mathbf{CP}_{ijk} \right] \right], \quad (1)$$

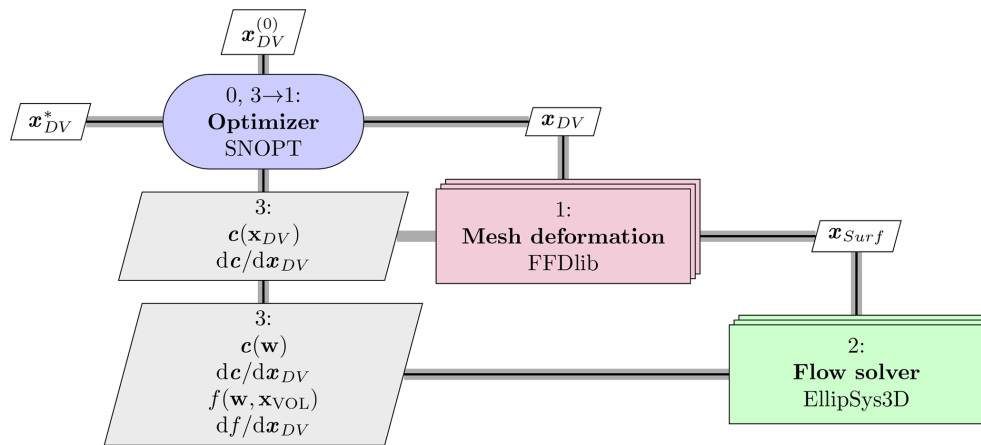
where  $\mathbf{CP}_{ijk}$  are the  $l \times m \times n$  control points of the FFD box, and  $\mathbf{X}_{FFD}(s, t, u)$  is the reconstructed 3-D point in the mesh computed from the normalized coordinates,  $(s, t, u)$ , and the control points,  $\mathbf{CP}_{ijk}$ .

The inverse search can be carried out by solving the following equation (Casale and Stanton, 1985, their Eq. 7):

$$\begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix} = \mathbf{0} \Leftrightarrow \mathbf{X}_{FFD}(s, t, u) - \mathbf{P} = \mathbf{0}, \quad (2)$$

where  $\mathbf{P}$  is the point which is to be embedded. This equation is solved in FFDlib with the iterative Newton search:

$$\begin{bmatrix} s \\ t \\ u \end{bmatrix}^{n+1} = \begin{bmatrix} s \\ t \\ u \end{bmatrix}^n - \begin{bmatrix} \frac{\partial f_1}{\partial s} & \frac{\partial f_1}{\partial t} & \frac{\partial f_1}{\partial u} \\ \frac{\partial f_2}{\partial s} & \frac{\partial f_2}{\partial t} & \frac{\partial f_2}{\partial u} \\ \frac{\partial f_3}{\partial s} & \frac{\partial f_3}{\partial t} & \frac{\partial f_3}{\partial u} \end{bmatrix}^{-1} \cdot \begin{bmatrix} f_1 \\ f_2 \\ f_3 \end{bmatrix}. \quad (3)$$



**Figure 1.** The FlowOpt framework when using a step-based approach for gradient computations.

The actual implementation of the above-described FFD method is split in two parts to maximize efficiency: it consists of an underlying code base written in Fortran containing the computationally heavy operations and a user-oriented interface written in Python which is integrated in the FlowOpt shape optimization framework.

Furthermore, the basic FFD methodology can be extended in many ways such as ensuring  $C^i$ -continuity control between deforming and non-deforming interfaces or volume preservation capabilities (see Hahmann et al., 2012). While FFDlib has indeed been extended with both mentioned features, it is only the  $C^i$ -continuity control that will be used in the present study to ensure a high mesh quality is maintained as described further in Sect. 5.1.

FFDlib will in the present study only be used to deform the CFD surface mesh. The deformed surface mesh will then be propagated down to the flow solver which in turn updates the volume mesh.

### 3.3 EllipSys3D: a general purpose flow solver

For the present study the general purpose flow solver, EllipSys3D, based on the finite volume method is used to solve the steady-state incompressible RANS equations as the underlying flow model along with the  $\kappa - \omega$  shear stress transport (SST) turbulence model by Menter (1992). Furthermore, velocity and pressure variables are coupled using the semi-implicit method for pressure-linked equations (SIMPLE) (Patankar, 1980), in which the Rhie–Chow interpolation (Rhie and Chow, 1983) is used to avoid checkerboard patterns. Flow solutions are obtained with a third-order accurate discretization scheme. Finally, the internal volume mesh deformation routines will be used as explained further in Sect. 3.3.1.

At present, the EllipSys3D has been used extensively in numerous application areas ranging from blinded comparisons (Simms et al., 2001) to rotor analysis (Sørensen et al.,

2002), DES simulations (Johansen et al., 2002), large eddy simulations (LESSs) (Berg et al., 2018), and studies in vortex-induced vibrations (Horcas et al., 2018) to name but a few. EllipSys3D was also recently tested on the MareNostrum (<https://www.bsc.es/marenostrum/marenostrum>, last access: 30 September 2021) supercomputer and achieved above a 50 % scaling efficiency when using more than 16 000 CPUs. This aspect is particularly important for the present study for which a high number of CPUs must be leveraged per rotor computation in order to arrive at competitive computation timings for the entire optimization.

As seen, EllipSys3D has during the last three decades of development been thoroughly extended, including an overset grid method (Zahle, 2006), transition modeling (Sørensen, 2009), and an adjoint solver (Madsen, 2020). It is outside the scope of the present work to account for all these applications. To be clear; the transition modeling was not used in the present study, and the flow was assumed to be fully turbulent. Depending on the condition of the surface and type of airfoil, this assumption is not correct since the laminar-to-turbulent transition will not take place at the leading edge but at some point along the chord. However, one can consider it a conservative modeling choice to ensure robustness of the design under conditions in which the boundary layers are turbulent (e.g., due to surface soiling).

Alone within studies focused on tip shapes and winglets the EllipSys3D flow solver has been used in about a dozen works (Johansen and Sørensen, 2006, 2007; Johansen et al., 2008; Gaunaa and Johansen, 2007, 2008; Gaunaa et al., 2011; Sørensen et al., 2011; Kalken and Ceyhan, 2017; Zahle et al., 2018). Therefore, only a paragraph focusing on the shape optimization studies leveraging the EllipSys3D flow solver is given in the following.

The EllipSys flow solver has been used in slat design using an overset grid method (Gaunaa et al., 2013) on a 10 MW rotor configuration, in which a series of 2-D cross-sections were optimized with multi-element airfoils. Ellip-



Sys was also used to design airfoils using a gradient-based setup with finite-differenced gradients, where the optimized airfoils were subsequently validated with wind-tunnel testing (Zahle et al., 2014). More recently, EllipSys3D figured in a surrogate-based design study (Zahle et al., 2018) in which an optimization problem similar to the one seen in the present study was solved. They showed there was great promise in coupling the CFD solver to lower-fidelity methods in the interest of saving computation time while maintaining the essential flow physics. Finally, the EllipSys3D flow solver has been integrated in the FlowOpt design optimization framework and applied in high-fidelity shape optimization. The chosen design optimization problem (Madsen, 2020, their Eq. 11.3) was a drag minimization of a 3-D wing subject to a lift constraint using five twist design variables. The problem was chosen to evaluate the newly developed framework on one of the few aerodynamic shape optimization problems that actually has an analytical optimal solution: an elliptic lift distribution. Mesh sizes up to  $664 \times 10^3$  and  $83 \times 10^3$  cells were used in the optimization using the Complex-Step and adjoint methods, respectively. It was found that the analytical elliptic lift distribution was well-approximated (Madsen, 2020, their Figs. 11.6 and 11.9), and most optimizations were tightly converged below a  $10^{-4}$  threshold. Finally, it was demonstrated that essentially identical optimizations (Madsen, 2020, their Fig. 11.11) occur when computing gradients either with the Complex-Step method or with the adjoint method for well-converged flows.

### 3.3.1 Mesh deformation

In this study the EllipSys3D flow solver will receive modified surface meshes as the optimization progresses. EllipSys3D will then subsequently propagate the deformation change between the original surface mesh and the deformed surface mesh out through the volume mesh using internal mesh deformation propagation routines. This deformation method in EllipSys is based on an analytical approach in which both the translational and overall re-orientations (i.e., rotations) of the surface are propagated and attenuated in the volume mesh using a hyperbolic tangent function, blended into the original volume mesh based on the distance to the blade surface along the given grid line. The consideration of the rotation of the surface greatly improves the quality of the deformed meshes for certain configurations compared to attenuating only the displacement resulting from the deformation. For instance this update has been instrumental for the present work as considerable local rotations were identified for some of the explored tip shapes. Without the consideration of rotations in the mesh deformation routines, this led to mesh folding in the boundary layer region.

## 3.4 Optimizer

In this work the Sparse Nonlinear OPTimizer (SNOPT) version 7.2-10 is used (Gill et al., 2018, 2005) (<https://web.stanford.edu/group/SOL/guides/sndoc7.pdf>, last access: 30 September 2021). SNOPT is based on a sequential quadratic programming (SQP) algorithm which allows for infeasible steps to be taken during the optimization. At convergence all optimizations were feasible within the requested tolerance.

The SNOPT optimizer is accessed in the FlowOpt framework through the open-source Python wrapper called pyOptSparse (Wu et al., 2020) by choosing the pyOptSparseDriver in OpenMDAO. pyOptSparse is leveraged extensively (<https://mdolab-pyoptsparse.readthedocs-hosted.com/en/latest/publishedWorks.html>, last access: 30 September 2021) throughout the broader numerical optimization community and is as the name suggests dedicated to constrained nonlinear optimization of large sparse problems.

## 4 Baseline analysis

The present section contains a description of baseline planform, surface mesh, and volume mesh. Mesh modifications were made to the original geometry to enable a deep convergence. These changes will also be presented, and the effects of the geometry changes will be assessed. Finally, a scaling study is also included to discuss the necessary computational resources needed to carry out direct CFD-based optimizations.


### 4.1 Computational mesh

The baseline geometry is the IEA 10 MW reference wind turbine (<https://www.nrel.gov/docs/fy19osti/73492.pdf>, last access: 30 September 2021; <https://github.com/ieawindtask37/iea-10.0-198-rwt>, last access: 30 September 2021) (Bortolotti et al., 2019, Appendix B). The chord and twist planform distributions can be inspected in Fig. 2.

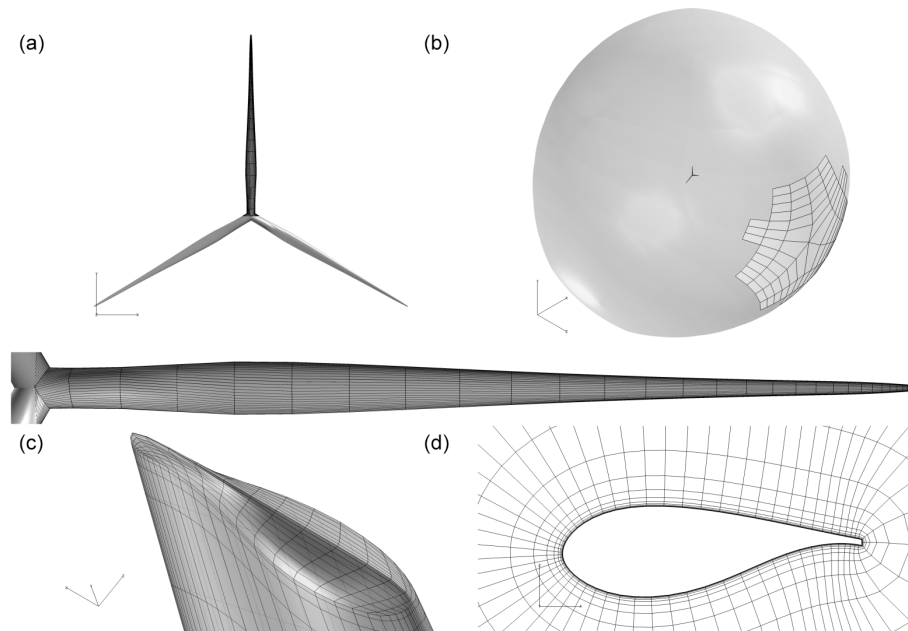
Now follows a description of how the structured surface and volume meshes are generated. Both surface mesh and volume mesh can be inspected in Fig. 3.

The rotor surface mesh has been generated from the planform data and the FFA-W3 airfoil family used for the IEA 10 MW with the in-house Parametric Geometry Library (PGL) tool. The surface mesh on each blade has 256 cells in the chordwise direction and 128 cells in the spanwise direction, partitioned into blocks of  $32 \times 32$  cells. Four blocks of  $32 \times 32$  cells form the tip cap, resulting in a total of  $3 \cdot 36 \cdot 32 \cdot 32 = 110592$  mesh cells for the surface mesh.

The baseline volume mesh is prepared with the in-house hyperbolic mesh generator, HypGrid3D (Sørensen, 1998). The volume mesh has an O–O topology in which 128 layers are grown from the surface mesh resulting in 14.16 million



**Figure 2.** The IEA 10 MW wind turbine (baseline) planform.



**Figure 3.** Surface (a) and volume (b) for which a close-up of one of the blade meshes is given below. Visualization of tip cap mesh (c) and a view of the airfoil discretization (d) are also offered. Notice that for clarity only one out of four mesh lines is shown. Boundary condition zones for inflow (transparent gray) and outflow (gray with black mesh lines) are visualized in the upper-right plot.

cells. By setting the first boundary layer cell below  $10^{-6}$  m, a  $y^+$  below 1.0 is ensured given the operational conditions seen in Table 2.


The above description is for the baseline mesh at the very start of the optimization. All subsequent volume meshes throughout the optimization are computed using the internal mesh deformation routines in EllipSys3D. As a result, they are all likely to exhibit a slight reduction in mesh quality due to impaired orthogonality of the mesh as the tip is created from a straight blade planform.

Finally, with respect to boundary conditions the rotor surface is a no-slip boundary, whereas the far-field zone is split into two sections: an approximately circular area behind the

rotor is an outflow-scaling zone, whereas the rest of the far-field region is an (uniform) inflow zone.

#### 4.2 Geometrical modifications

Initially, the described surface mesh of the IEA 10 MW reference wind turbine was used in the optimizations without any additional modifications. However, as finer grid levels were taken into use a seriously impaired gradient quality was observed. The explanation proved to be a lack of flow convergence on the finer grid levels due to complex swirling 3-D flow phenomena occurring in the root area of the rotor. Such a massive blunt object will inherently cause complex flow phenomena resulting in a lack of convergence for steady-



**Figure 4.** The IEA 10 MW reference wind turbine baseline geometry (a) and mesh modifications needed (b), resulting in an improved flow solver convergence.

state incompressible CFD solvers based on the SIMPLE algorithm.

There are several remedies suggested in wind energy research on shape optimization for the described convergence issues. Nielsen and Diskin (2012) were able to carry out shape optimization not just of the rotor itself<sup>1</sup> but also including nacelle and tower (Nielsen and Diskin, 2012, their Fig. 4) by using an unsteady RANS formulation. It is very likely that the unsteady RANS formulation in EllipSys3D would somewhat mitigate the observed impaired convergence. However, the computation time would drastically increase, which is why this option was ruled out.

Yet another option found in the literature is to exclude some mesh regions: Dhert et al. (2017) cut out the root of the rotor configuration to improve convergence, whereas Vorspel et al. (2018) exclude both root and tip regions from being deformed due to a local inferior gradient quality.

Of the above-mentioned alternatives the mesh (root) modification option is favored since the present shape optimization study is focused on tip shapes. However, instead of removing the root section altogether it was decided to simply re-shape it aerodynamically to reduce separation and in turn improve convergence. The modified baseline mesh can be inspected in Fig. 4.

To assess the quality of the modified mesh a full grid sequence is run on both the original baseline mesh and the modified baseline mesh using the operational conditions seen in Table 2.

The resulting convergence behavior, as well as spanwise forces, can be inspected in Fig. 5. As is evident from the figure it is only the innermost part of the spanwise forces where one can discern a minor change. The design optimization problem of optimizing the tip is in other words practically unaltered, and the convergence issues have been addressed without changing the task at hand.

<sup>1</sup>Interested readers will notice that the actual root of the NREL VI rotor was excluded in this study. Still, the inclusion of nacelle and tower introduce a massive blunt object in the flow.

### 4.3 Mesh convergence study

A mesh convergence study is carried out using the modified rotor geometry to investigate mesh dependence of the flow solution. To generate the coarse mesh levels to be used in the mesh convergence study a grid coarsening sequence inside the flow solver is used: the above-described computational mesh is the finest mesh called L1. The next mesh level, L2, is generated by removing every second mesh point in all directions. Similarly, the L3 mesh level, which is the coarsest mesh level used in the present study, is generated by removing every second grid mesh point from L2. All three mesh levels, L1, L2, and L3, are listed in Table 3.

It has previously been shown (Madsen et al., 2019, their Table 4) that a flow solver with a noticeable grid dependency may even suggest incorrect design trends (Madsen et al., 2019, their Sect. 6.2) on the coarser grid levels. Therefore, one should ensure that all mesh levels exhibit low error percentages compared to the Richardson extrapolation. As listed in Table 3 the error percentage on both mesh levels L1 and L2 is well below 10%, and it should be reasonable to expect relevant design optimization results at least for these two mesh levels.

### 4.4 Computational resources

As a final section in the baseline analysis the computational resources needed to carry out shape optimizations on all three grid levels are assessed. A visualization of flow solver scaling on the high performance computing (HPC) cluster<sup>2</sup> at DTU Wind Energy can be inspected in Fig. 6.

The most important aspect to consider when inspecting scaling results as shown in Fig. 6 is the finest mesh level, L1, since it by far is the most time consuming and since it will determine the final results. As seen, the scaling is indeed close to ideal on L1 for EllipSys3D, and if computational resources are available, one can advantageously use as many CPUs as possible on the given rotor mesh. In the present case that number is 432 – one CPU for each block in the mesh.


## 5 Design optimization problem

The single-point design optimization problem used in the present study is to optimize the power production using 12 design variables while satisfying constraints on smoothness of the geometry and on the flapwise bending moment computed at 90% span, which from now on will be written as  $0.9 r/R$  for brevity. Mathematically, the design optimization

<sup>2</sup>The Sophia HPC cluster at DTU Wind Energy comprises 516 computational nodes where each of these is a x86-64 computer with 32 cores. More information is available at <https://windenergy.dtu.dk/nyheder/2019/12/ny-computer-cluster-paa-risoe-campus?id=a495d2e5-a7f1-4133-9fb2-488d150f7c01>, last access: 30 September 2021.

**Table 2.** Operational conditions for the simulations used in both analysis and optimization. Density is set to the density of air at sea level and 15 °C,  $\rho = 1.225 \text{ kg m}^{-3}$ , and dynamic viscosity is set to  $\mu = 1.784 \times 10^{-5} \text{ kg m}^{-1} \text{ s}^{-1}$ .

Run	Wind speed [m s <sup>-1</sup> ]	RPM [-]	Rotation rate, $\omega$ [rad s <sup>-1</sup> ]	Pitch [°]	TSR [-]
wsp08	8.0	8.164590	0.8681	0.0	10.58



**Figure 5.** A comparison of the baseline mesh before (gray) and after (black) root modifications reveal that the modified baseline mesh ensures a deep convergence (a) on all the grid levels: L3 (it = 0–4000), L2 (it = 4000–12 000), and L1 (it = 12 000–30 000). Mesh levels are listed in Table 3. Also the normal force (b) and driving force (c) are visualized. As seen, the spanwise forces only differ noticeably within the 30 m farthest inboard.

problem can be formulated as follows:

$$\begin{aligned} &\text{minimize: } -\frac{P(\mathbf{x})}{P(\mathbf{0})}, \\ &\text{with respect to } \begin{cases} \text{twist: } \theta_1, \theta_2, \theta_3, \theta_4 \\ \text{chord: } c_1, c_2, c_3, c_4 \\ \text{tip shape: } c_{wl}, h_{wl}, s_{wl}, R_{ext} \end{cases}, \\ &\text{subject to } \frac{M_{\text{bending}}|_{0.9r/R}^{\text{initial}}}{M_{\text{bending}}|_{0.9r/R}^{\text{initial}}} \leq 1, \frac{d\text{chord}}{dS} \leq 0. \end{aligned} \quad (4)$$

Above, the objective function is mechanical power,  $P = \omega \cdot Q$ , found from rotational rate,  $\omega$ , and torque,  $Q$ , where the operational conditions used in this study can be found in Table 2.

As seen, the constraint on bending moment ensures that the bending moment at 0.9  $r/R$  span does not increase compared to the baseline value. The geometric chord design variable constraint ensures that the optimization does not increase the chord towards the very tip of the blade. There are no constraints on the twist design variables. In the following section (Sect. 5.1) the 12 design variables are further described.

### 5.1 Parameterization

In the design optimization problem there are 12 design variables: four twist variables ( $\theta_{1,2,3,4}$ ), four chord variables ( $c_{1,2,3,4}$ ), one extension variable ( $R_{ext}$ ), one flapwise tip displacement variable ( $h_{wl}$ ), one edgewise tip displacement variable ( $s_{wl}$ ), and one tip curvature variable ( $c_{wl}$ ). We have adopted all design variable nomenclature from a related

surrogate-based optimization study (Zahle et al., 2018) to allow for an easy comparison. The twist, chord, and extension design variables change the planform at the blade tip region, whereas  $c_{wl}$ ,  $h_{wl}$ , and  $s_{wl}$  are used to shape the tip towards winglet-like shapes.

All 12 design variables are imposed by manipulating the three FFDlib tip boxes (blue) seen in Fig. 7. As seen, the deforming tip region (red mesh lines) embedded in the FFD boxes only makes up 10 % of the total blade length of the baseline rotor (gray). The twist and chord variables are imposed section-wise on FFD box sections  $S_3$ ,  $S_4$ ,  $S_5$ , and  $S_9$  in Fig. 7. Sections  $S_{6-8}$  are given from interpolation once sections  $S_5$  and  $S_9$  are set. Also seen are two darker blue areas on the FFD boxes signifying that the related FFD spanwise sections have a particular functionality: the darker blue part further inboard visualizes the two FFD box sections,  $S_{1-2}$ , which are locked to ensure a  $C^1$  continuity with the remaining surface mesh. The darker blue part towards the very tip of the blade visualizes three FFD sections,  $S_{6-8}$ , that have been inserted as a tip-cap protection since it is crucial that this part of the mesh retains a high mesh quality. Notice that the tip can indeed still be deformed by moving the outermost FFD box section, in which case the three tip-cap protection sections are interpolated to their correct position.

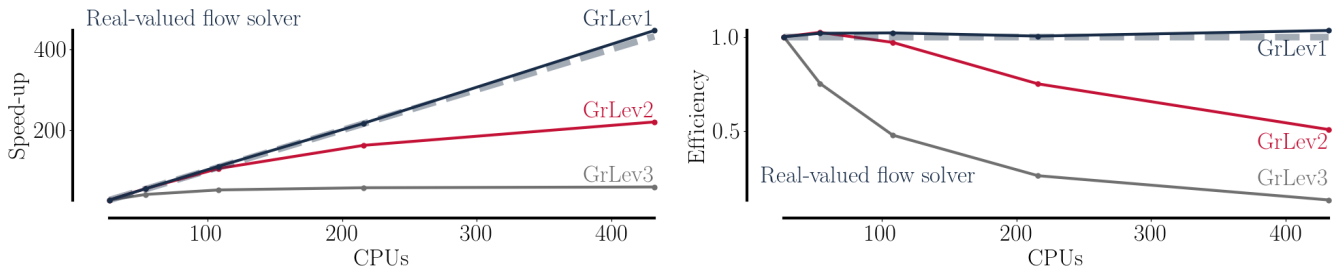
#### 5.1.1 Tip design variables

The three tip design variables – flapwise tip displacement ( $h_{wl}$ ), edgewise tip displacement ( $s_{wl}$ ), and curvature of the tip ( $c_{wl}$ ) – can be inspected in Fig. 8 where the FFD box visualization has been omitted in order to more easily inspect the



**Table 3.** Mesh convergence study with error percentages obtained from Richardson extrapolations. The quadratic upstream interpolation for convection kinematics (QUICK) discretization scheme is used for all cases.

Mesh	Cells [million]	QUICK: third-order stencil			
		Torque $\times 10^6$ [Nm]	Error [%]	Thrust $\times 10^6$ [N]	Error [%]
L3	0.221	5.676	10.5	1.169	6.2
L2	1.769	5.394	5.0	1.124	2.1
L1	14.155	5.200	1.3	1.106	0.5
Extrapolation	$\infty$	5.136	0.00	1.101	0.0



**Figure 6.** Measured scaling on the Sophia cluster for the EllipSys3D flow solver on the modified baseline mesh using 27, 54, 108, 216, and 432 CPUs. For the L2 mesh it seems up to 108 CPUs will result in very efficient CPU usage. For L1, which is by far the most interesting and time-consuming mesh level, the scaling is ideal, and up to 432 CPUs can advantageously be used. On L3 the mesh is so coarse that the computational task is modest, making scaling on this mesh level less important.

deformed geometry. Notice that the curvature is an interpolation variable from 0.0 (maximum curvature) to 1.0 (straight tip). It is also relevant to point out that the role of the three locked FFD sections (dark blue in Fig. 7) next to the outermost tip section is to protect the volume cells at the tip, which at the same time results in a reduced maximal curvature for the parametrization.

## 5.2 Gradient verification

Given that the present study uses the finite difference method an initial step size study is carried out to identify a suitable step size. Table 4 shows how the finite difference gradient accuracy correlates to the chosen finite difference step size. The machine-accurate reference gradient is computed using the Complex-Step method. Gradient accuracy for all constraints has also been verified (not shown) and exhibits similar accuracy.

By inspecting Table 4 one learns that for most design variables (e.g., twist, edge, extension) the best finite difference step size is  $10^{-4}$ . However, for chord design variables a step size of  $h = 10^{-5}$  seems better suited. Up to six significant digits can be seen, which is more than plenty to carry out gradient-based design optimization. However, it is also evident that the step size does not have to be much off before the gradient precision is worsened. This may factor in for longer optimizations for which many new rotor shapes are intro-

duced, meaning that also the optimal step size might change slightly throughout the course of the optimization.

## 6 Results

The main results of this study fall in two parts. The first part (Sect. 6.1) is carried out solely on mesh level L3 and is a study in finding the best settings that balance the most accurate gradient computation on the one hand with more robust settings on the other hand that result in properly converged optimization problems. The second part (Sect. 6.2) is a shape optimization study using three grid levels in which focus is laid on analyzing the shape of the resulting blade tip.

### 6.1 Step size study for shape optimizations based on finite difference

To identify the best functioning step size a series of six shape optimizations were run using step sizes  $h = 10^{-1}$ ,  $h = 10^{-2}$ ,  $h = 10^{-3}$ ,  $h = 10^{-4}$ ,  $h = 10^{-5}$ , and  $h = 10^{-6}$ . All optimizations were allowed to run until either the upper limit of 100 major iterations (i.e., design steps) was reached or the optimizer exited since it could not converge the design problem further<sup>3</sup>. The lower and upper design variable bounds for this

<sup>3</sup>The related SNOPT message is as follows: SNOPTC EXIT 40 - terminated after numerical difficulties. SNOPTC INFO 41 - current point cannot be improved.




Figure 7. Overview of baseline rotor (gray), deformable tip region (red), and FFD boxes (blue). Grid level 3 is used to visualize mesh lines.

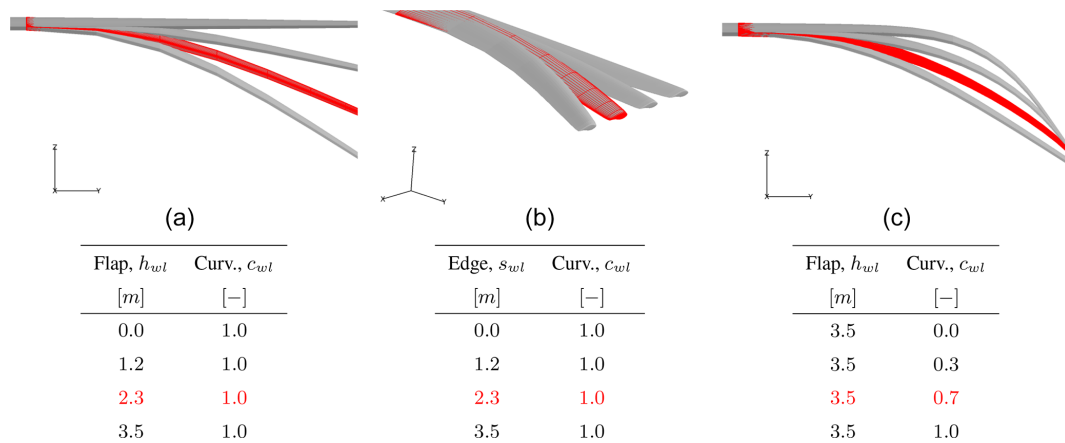


Figure 8. Tip design variables: flapwise tip displacement (a), edgewise tip displacement (b), and curvature of the tip (c). The design variables used are listed below each visualization. The rows colored in red show the design variable settings used to generate the red meshes. The flapwise and edgewise displacement visualizations are made with no curvature, forcing the tip to be as straight as possible. The curvature visualization is made with a maximally flapped tip. Grid level 3 is used to visualize mesh lines.

step size study are

$$\begin{bmatrix} -10.0, -10.0, -10.0, -10.0 \\ \cdot 0.5, \cdot 0.5, \cdot 0.5, \cdot 0.5 \\ \cdot 0.0, +0.0, +0.0, \cdot 1.0 \end{bmatrix} \leq \begin{bmatrix} \theta_1, \theta_2, \theta_3, \theta_4 \\ c_1, c_2, c_3, c_4 \\ c_{wl}, h_{wl}, s_{wl}, R_{ext} \end{bmatrix} \leq \begin{bmatrix} [l] + 5.0, +5.0, +5.0, +5.0 \\ \cdot 1.0, \cdot 1.0, \cdot 1.0, \cdot 1.0 \\ \cdot 1.0, +3.5, +3.5, \cdot 2.0 \end{bmatrix}. \quad (5)$$

Above, the units for the design variables are the following. Twist variables,  $\theta$ , are in degrees. Chord variables,  $c$ , use a unitless scaling factor. Curvature is likewise a unitless interpolation factor from 0.0 to 1.0, where 0.0 represents the maximally allowed curvature for the parametrization, and 1.0 represents a straight tip (see Fig. 8c). Flapwise displacement,  $h_{wl}$ , and edgewise displacement,  $s_{wl}$ , are in meters, and the

extension scaling variable,  $R_{ext}$ , is a unitless scaling variable which stretches the entire FFD box.

The shape optimization results for the six different step sizes are listed in Table 5, and their optimization histories are visualized in Fig. 9. Before discussing the results in Table 5 it should be noted that the step size is but one of several important settings that one must fine-tune to arrive at a properly functioning optimization framework. Another important consideration to mention is that these step size studies ideally should be done on each grid level. However, that is extremely expensive on the finest grid level for 12 design variables. Based on the present study it is the authors' experience that for flow solvers as consistent across grid levels as seen in Table 3 it will suffice to carry out the step size study on mesh level L3. For flow solvers less consistent across grid levels one might have to redo the step size study on each grid level.

**Table 4.** Table showing how the finite difference gradient accuracy depends on the chosen step size. The machine-accurate reference gradient is computed with the Complex-Step method.

Step size, $h$	Significant digits on L3:			
	Twist, $\theta_1$ ( $\cdot 10^{-4}$ )	Twist, $\theta_2$ ( $\cdot 10^{-4}$ )	Twist, $\theta_3$ ( $\cdot 10^{-4}$ )	Twist, $\theta_4$ ( $\cdot 10^{-4}$ )
$10^{-1}$	-5.85049867	-4.43009974	-3.34042748	-6.96420069
$10^{-2}$	-5.86822272	-4.44143106	-3.34801787	-7.04257508
$10^{-3}$	-5.87003860	-4.44256502	-3.34877921	-7.05052147
$10^{-4}$	-5.87025903	-4.44271548	-3.34889474	-7.05135166
$10^{-5}$	-5.87069215	-4.44313697	-3.34930883	-7.05186998
$10^{-6}$	-5.87511817	-4.44770443	-3.35383055	-7.05632885
(reference)	-5.87055619	-4.44270932	-3.34887356	-7.05133069
	Chord, $c_1$ ( $\cdot 10^{-3}$ )	Chord, $c_2$ ( $\cdot 10^{-3}$ )	Chord, $c_3$ ( $\cdot 10^{-3}$ )	Chord, $c_4$ ( $\cdot 10^{-2}$ )
$10^{-1}$	-8.79500858	-6.08880387	-4.32113105	-1.03168014
$10^{-2}$	-8.12767663	-5.70326167	-4.06860284	-0.80373667
$10^{-3}$	-8.06665832	-5.66956278	-4.04640816	-0.78776642
$10^{-4}$	-8.06085604	-5.66622427	-4.04407679	-0.78611997
$10^{-5}$	-8.06021495	-5.66583677	-4.04381366	-0.78595010
$10^{-6}$	-8.05994610	-5.66558572	-4.04358425	-0.78591230
(reference)	-8.06008832	-5.66584359	-4.04382519	-0.78592965
	Curv., $c_{wt}$ ( $\cdot 10^{-3}$ )	Flap, $h_{wt}$ ( $\cdot 10^{-4}$ )	Edge, $s_{wt}$ ( $\cdot 10^{-4}$ )	Ext., $R_{ext}$ ( $\cdot 10^{-1}$ )
$10^{-1}$	-3.74767833	+3.28947962	-3.19416366	-1.41397308
$10^{-2}$	-3.86239738	+3.64583550	-3.71786342	-1.41788435
$10^{-3}$	-3.87348172	+3.67988163	-3.77135090	-1.41846530
$10^{-4}$	-3.87460335	+3.68325048	-3.77628964	-1.41852803
$10^{-5}$	-3.87475825	+3.68314301	-3.77729514	-1.41853454
$10^{-6}$	-3.87520704	+3.67891606	-3.78163278	-1.41853715
(reference)	-3.87632872	+3.68048836	-3.77670622	-1.41852678

Red digits (i.e.,  $\theta_4[h = 10^{-1}]$ ,  $c_1[h = 10^{-6}]$ , and  $h_{wt}[h = 10^{-3}]$ ) are used to indicate that the error is not on the indicated digit but one order of magnitude below.

The first thing to note when inspecting Table 5 is that one should take care not to use finite difference step sizes that are too small. Indeed, the most accurate gradient step sizes identified in Sect. 5.2 (i.e.,  $h = 10^{-4}$  and  $h = 10^{-5}$ ) do not result in very successful optimizations. Although  $h = 10^{-4}$  and  $h = 10^{-5}$  indeed where the most accurate step sizes on the baseline mesh, they do not (in the present study) seem to be a robust choice over the course of an entire optimization. One potential explanation could be that cancellation errors due to step sizes that are too small are likely to occur for at least a few of the several hundred rotor shapes generated by the optimizer during an optimization. Even if this happens for just a few shapes, it might result in the optimizer having to reset the Hessian which could impair the optimization problem convergence. One could implement individual step sizes

for each design variable in the attempt to gain better gradient accuracy and robustness, but in general it seems advisable to use a step size 1–2 orders of magnitude larger than  $h = 10^{-4}$  to deeply converge optimization problems.

Judging from Table 5 the main benefit from finding a good compromise between gradient accuracy and overall robustness is that a sound procedure for ending the optimization is available: optimizations with a functioning step size (SPL3e1c, SPL3e2c, and SPL3e3c) all have around 50 major iterations when terminating and result in approximately the same improvement, whereas the other optimizations (SPL3e4c, SPL3e5c, SPL3e6c) differ both in number of major iterations and in the final improvement. Noticeably, these optimizations (e.g., SPL3e4c) may still result in approximately as much improvement as the better-performing

**Table 5.** Overview of six CFD-based shape optimizations of wind turbine blade tips using different finite difference step sizes. Operational conditions are found in Table 2. The naming convention explanation using “SPL3e2c” as an example is as follows. SP: a single-point optimization; L3: on mesh level L3; e2: using a finite difference step size of  $10^{-2}$ ; and c: (cold-)started from a straight baseline configuration. All optimizations were carried out using 648 CPUs split into 12 groups – one group of 54 CPUs for each design variable.

ID	Mesh level (see Table 3)	Wall clock* $10^{-7}$ threshold/full [h]	Maj. iter.	Step size ( $h =$ )	Convergence (orders of magn.)	Mech. power (improvement)
SPL3e6c	L3	-/28.2	88	$10^{-6}$	1	0.37 %
SPL3e5c	L3	-/14.1	12	$10^{-5}$	1	0.37 %
SPL3e4c	L3	-/11.6	40	$10^{-4}$	2	0.39 %
SPL3e3c	L3	6.0/11.6	54	$10^{-3}$	8	0.39 %
SPL3e2c	L3	5.3/10.5	46	$10^{-2}$	8	0.39 %
SPL3e1c	L3	5.7/9.2	50	$10^{-1}$	9	0.39 %

\* Given that some groups of CPUs on the HPC cluster will be faster than other CPU groups a representative computation speed (in  $s^{-1}$ ) for each grid level has been computed as an average over an entire optimization on a given grid level. This allows for a fair comparison between optimizations carried out on the same grid level although they have not been computed using the exact same CPUs.

optimizations, but the optimization problem is only converged by 1–2 orders of magnitude, and one cannot be sure whether the optimization problem is actually solved.

Optimizations SPL3e5c and SPL3e6 in Table 5 resemble two typical end-case scenarios for optimizations with inefficient settings. Either the optimization finishes prematurely (e.g., SPL3e5c) or the optimizer keeps trying to converge the optimization problem using excessive iterations (e.g., SPL3e6c). In both cases, the optimizer exits due to numerical difficulties resulting from the inaccurate gradient. While the SPL3e5c optimization finishes much sooner than the SPL3e6c optimization, thus saving considerable computation time, it actually results in a shape which performs slightly worse (not visible in Table 5). Thus, SPL3e5c represents the least successful optimization in Table 5.

Turning to Fig. 9 it is easy to see how the optimizations progress. All optimization problems start with an optimality around  $10^{-3}$  (see right y axis in Fig. 9 shown in red). The  $10^{-7}$  threshold has been chosen to ensure that all design optimization problems are converged by about 4 orders of magnitude. To relate this to relevant literature it is noted that a wind turbine design problem with a single design variable (pitch) has been well-converged with an optimality reduction of only 3 orders of magnitude on similar mesh sizes ( $325 \times 10^3$  cells in Dhert et al., 2017, their Fig. 5, and  $221 \times 10^3$  cells in Madsen et al., 2019, their Fig. 8). Indeed, Madsen et al. (2019) showed that wind energy design optimization problems with 1, 14, and 154 design variables were all well-converged with an optimality reduction of about 2–4 orders of magnitude (Madsen et al., 2019, their Figs. 8, 9, 11, 12, 14, and 17). Therefore, a chosen threshold of about 4 orders of magnitude reduction in optimality for the present study seems reasonable. As seen in Fig. 9 the chosen threshold is only reached by the optimizations which converge properly, and the threshold could even have been 3–4 orders of magnitude stricter at the limited expense of about 10 major iterations. However, it is clear from the merit

function shown in black that it would lead to no added value since the objective function for all practical purposes is fully converged. In order to accommodate reproducibility, Fig. 9 shows a “merit function” instead of the actual design optimization problem objective since it is the former metric that SNOPT exports. However, these two metrics are highly related as detailed in SNOPT’s manual, and the merit function will converge to the objective function value as the solution is approached. The optimization history for SPL3e6c clearly visualizes what may happen if said threshold is not met: the optimizer uses a large amount of excessive iterations in the attempt to further converge the optimization problem. Evidently, it is crucial to have a well-defined way to end optimizations.


An important point when discussing the cost of optimizations is to discern an optimizer’s major iterations from wall clock computation time: although SPL3e5c is the optimization with by far the fewest major iterations, it is the second most time-consuming wall-clock optimization, as also indicated in Table 5. The optimizer simply takes very few actual steps in this optimization due to impaired gradient precision.

In summary, the best functioning step sizes seem to be  $h = 10^{-1}$ ,  $h = 10^{-2}$ , and  $h = 10^{-3}$ . Given that  $h = 10^{-3}$  of the three resulted in the most accurate gradient computation when comparing to machine-accurate reference gradients (Table 4), it is the  $h = 10^{-3}$  step size that will be used in the ensuing section.

## 6.2 Aerodynamic shape optimization of wind turbine blade tips

With  $h = 10^{-3}$ , identified as a promising finite difference step size in Sect. 6.1, a shape optimization study across three grid levels has been carried out. All optimizations were again allowed to run until either the upper limit of 100 major iterations (i.e., design steps) were reached or the optimizer exited





**Figure 9.** Merit function (black, left y axis) and optimality (red, right y axis) for shape optimizations on mesh level L3 using step sizes  $h = 10^{-1}$ ,  $10^{-2}$ ,  $10^{-3}$ ,  $10^{-4}$ ,  $10^{-5}$ , and  $10^{-6}$ . The simulations are further described in Table 3 under the names SPL3e1c, SPL3e2c, SPL3e3c, SPL3e4c, SPL3e5c, and SPL3e6c, respectively. As seen, the three optimizations with the largest finite difference gradient step sizes converge deeply, which is not true for the three optimizations with the smallest step sizes.

since it could not converge the design problem further<sup>4</sup>. For large flapwise and edgewise displacements negative cell volumes were encountered on the finest mesh level as a result of the deformation of the mesh during the optimization. In order to make sure that the exact same optimization could be carried out on the various grid levels, it was therefore necessary to limit the upper design variable bounds for  $h_{wl}$  and  $s_{wl}$  to 2.0 m. All other settings from Sect. 6.1 were kept the same. The results from the final shape optimization study are listed in Table 6, and their optimization histories are visualized in Fig. 10.

As seen from Table 6 a total of five optimizations where run: three optimizations (SPL3e3cb, SPL2e3cb, and SPL1e3cb) were (cold-)started from a straight baseline blade, whereas two optimizations (SPL2e3hb, SPL1e3hb) used the optimization result from a coarser mesh as a starting point.

Starting from the left in Table 6 one can, after “ID” and “Mesh level”, find two columns describing the computational resources used in this study (i.e., “CPUs” and “Wall clock”). When discussing the amount of CPUs used in the study it is good to remember that the CPUs are split into 12 groups: one group for each design variable to reduce the gradient computation time. The reason for using only  $12 \cdot 54 = 648$  CPUs on mesh level L3 is that the efficiency study (see Fig. 6) clearly showed that very little speed-up could be gained by increasing the number of CPUs on this mesh level. Similarly, for L2 there is a drop in efficiency in Fig. 6 after 108 CPUs, which is why  $12 \times 108 = 1296$  CPUs were used for SPL2e3cb. However, for L1 the efficiency in Fig. 6 for the EllipSys3D flow solver is at the upper possible limit, and more CPUs could advantageously have been used. The reason for limiting SPL1e3cb and SPL1e3hb to 1296 CPUs has to do with the computational resources available for the present work.

<sup>4</sup>The related SNOPT message is as follows: SNOPTC EXIT 40 - terminated after numerical difficulties. SNOPTC INFO 41 - current point cannot be improved.

This also means that the computation time could be much improved: simply by raising the number of CPUs from 108 to 432 one could gain a factor of 4 in speed-up, meaning that SPL1e3hb could be carried out in only  $204.3/24/4 \approx 2$  d. For a well-converged high-fidelity optimization 2 d is certainly an acceptable computation time. One may find more optimistic computation time consumptions reported in the literature but typically with a correspondingly poor convergence of the design optimization problem. As also indicated in Fig. 10 the computation time is very dependent on the chosen threshold. Therefore, one should not discuss one without considering the other.

While on the topic of computation time it may be relevant to mention the  $\tau$  time unit which is a non-dimensional work unit to compare across HPC systems (<https://cfd.ku.edu/hioicfd/>, last access: 12 July 2022). The  $\tau$  code was downloaded, and 10 runs were carried out on the Sophia cluster resulting in an average execution time of 3.89 s. Using this result one can now compute a unitless normalized version of the reported wall clock times in Table 6. As an example, the SPL2e3cb wall clock time to reach the threshold (129.2 h) is in normalized  $\tau$  units:  $129.2 \cdot 60 \cdot 60 [s] / 3.89 [s] = 119.67 \times 10^3$ .


It is difficult to relate the reported timings to relevant literature from the wind energy community since very few high-fidelity CFD-based shape optimization studies of that magnitude exist (Madsen, 2020, Table 3.1). Elfarra et al. (2014) do mention an approximate wall clock timing of 240 h for their gradient-free optimization, but since the timings are not in normalized  $\tau$  units, it is difficult to compare across HPC platforms. Furthermore, Elfarra et al. (2014) do not mention how well their optimization problem is converged in terms of optimality reduction, which makes a comparison very difficult to carry out. Turning to the study by Dhert et al. (2017) one can, however, find both wall clock and optimality reduction reported: they spend 8.25 h on a wind turbine design problem with a  $2.6 \times 10^6$  cell mesh resolution. The reported mesh resolution is most easily compared with the SPL2e3cb results

**Table 6.** Overview of optimized shapes obtained from CFD-based shape optimizations of wind turbine blade tips. Operational conditions are found in Table 2. The naming convention explanation using “SPL1e3hb” as an example is as follows. SP: a single-point optimization; L1: on mesh level L1; e3: using a finite difference step size of  $10^{-3}$ ; h: (hot-)started from the optimized shape on L2; and b: with limited design variable bounds that work on all grid levels. All optimizations were carried out using either 648 (L3) or 1296 (L2, L1) CPUs split into 12 groups – one group of either 54 (L3) or 108 (L2, L1) CPUs for each design variable.

ID	Mesh level (See Tab. 3)	CPUs	Wall clock* $10^{-7}$ threshold / full [h]	Design variables †	$M_{\text{bending}} _{0.9r/R}$ (baseline fraction)	Mech.Power (improvement)
				$\begin{bmatrix} \theta_1, & \theta_2, & \theta_3, & \theta_4 \\ c_1, & c_2, & c_3, & c_4 \\ c_{wl}, h_{wl}, s_{wl}, R_{ext} \end{bmatrix}$		
SPL3e3cb	L3	648	11.2 / 12.3	$\begin{bmatrix} +3.8 & -10.0 & +5.0 & -5.4 \\ -1.0 & -1.0 & -1.0 & -1.0 \\ -0.0 & +1.0 & +2.0 & -1.1 \end{bmatrix}$	1.0000	0.33 %
SPL2e3cb	L2	1296	129.2 / 299.3	$\begin{bmatrix} +4.7 & -10.0 & +5.0 & -4.5 \\ -1.0 & -1.0 & -1.0 & -0.8 \\ -0.0 & +2.0 & +2.0 & -1.1 \end{bmatrix}$	1.0000	0.42 %
SPL2e3hb	L2	1296	110.7 / 246.5	$\begin{bmatrix} +4.7 & -10.0 & +5.0 & -4.5 \\ -1.0 & -1.0 & -1.0 & -0.8 \\ -0.0 & +2.0 & +2.0 & -1.1 \end{bmatrix}$	1.0000	0.42 %
SPL1e3cb	L1	1296	384.1 / 547.1	$\begin{bmatrix} +4.1 & -10.0 & +5.0 & -4.0 \\ -1.0 & -1.0 & -1.0 & -0.7 \\ -0.0 & +2.0 & +1.9 & -1.1 \end{bmatrix}$	1.0000	0.44 %
Evaluation of optimization result using a regenerated volume mesh:					0.9992	1.12 %
SPL1e3hb	L1	1296	204.3 / 332.4	$\begin{bmatrix} +4.1 & -10.0 & +5.0 & -4.0 \\ -1.0 & -1.0 & -1.0 & -0.7 \\ -0.0 & +2.0 & +1.9 & -1.1 \end{bmatrix}$	1.0000	0.44 %
Evaluation of optimization result using a regenerated volume mesh:					0.9992	1.12 %

\* Given that some groups of CPUs on the HPC cluster will be faster than other CPU groups a representative computation speed (it/[sec]) for each grid level has been computed as an average over an entire optimization on a given grid level. This allows for a fair comparison between optimizations carried out on the same grid level although they have not been computed using the exact same CPUs.

† The lower and upper design variable bounds are:  $\begin{bmatrix} -10.0, -10.0, -10.0, -10.0 \\ -0.5, -0.5, -0.5, -0.5 \\ -0.0, +0.0, +0.0, -1.0 \end{bmatrix} \leq \begin{bmatrix} \theta_1, \theta_2, \theta_3, \theta_4 \\ c_1, c_2, c_3, c_4 \\ c_{wl}, h_{wl}, s_{wl}, R_{ext} \end{bmatrix} \leq \begin{bmatrix} +5.0, +5.0, +5.0, +5.0 \\ -1.0, -1.0, -1.0, -1.0 \\ -1.0, +2.0, +2.0, -2.0 \end{bmatrix}$



**Figure 10.** Merit function (black, left y axis) and optimality (red, right y axis) for shape optimizations (Table 6) on mesh levels L3, L2, and L1 using finite difference step size  $h = 10^{-3}$ . The SPL1e3hb is a hot-started optimization on L1 where the result from L2 is used as the starting point. The SPL2e3hb correspondingly uses the L3 result as the starting point. All other optimizations start from a straight blade.

in Table 6. However, it is difficult from the reported final optimality (Dhert et al., 2017, Table 2) to learn how many orders of magnitude it has been converged, making it difficult to compare with the presented timings in this study. Madsen et al. (2019) report CPU timings (Madsen et al., 2019, Table 6) for adjoint-based high-fidelity shape optimizations using a single design variable of close to 60h for a mesh with a resolution equal to L1 in Table 3. Again, these computations were carried out on a different HPC platform than the present study, so these timings are difficult to compare, but as optimizations tend to become more difficult to converge as more design variables are included, the 60h should certainly be seen as a lower possible bound in that study. In summary,

a realistic lower bound for a high-fidelity optimization that is well converged seems to be around 2–3 d for very efficient frameworks.

Returning to Table 6 one can in the fifth column find the final design variables for each optimization. For readability all design variables ending at the upper/lower limit have been colored accordingly. Overall, it can be said that the final shape trends favored by the optimizer are an increase in sweep and an even greater increase in flapwise displacement. Both design variables are used to mitigate an increase in bending moment as the blade is extended. Of particular interest is the change in final sweep design variable from L2 to L1, where the L1 result,  $s_{wl} = 1.9$ , is not on the limit any-

more, and the optimizer seems to have found a maximum for this design variable. This aligns very well with findings in the related surrogate-based study, in which a maximum for the sweep design variable is found at 2%  $d/R$  (Zahle et al., 2018, their Fig. 4). Turning to the twist and chord design variables the optimizer favors the twist design variable over the chord design variable to shed loads. For all optimizations the optimizer favors curvature, and for most optimizations it is at the very limit of what is allowed for the parametrization. All design trends agree with the surrogate-based design study by Zahle et al. (2018).

The resulting shape is visualized in Fig. 11, and an analysis of planform and spanwise forces is given in the following subsection (Sect. 6.2.1). As seen, the produced novel curved tip shape can extend the blade in an efficient manner, and it effectively mitigates the downwash by displacing the tip vortex. The vortex is somewhat diffused and smeared out along the extended tip structure protruding into the out-of-plane region as one would expect to see for a functioning optimization.

With respect to the final two columns in Table 6 showing bending moment constraint and gained improvement, they agree well with expectations: the bending moment constraint is indeed a design-driving load constraint and will to a great extent dictate how far the blade can be extended. In practice, the load level constraint would depend on the structural capacity in the blade, which may well allow for increases in the flapwise moment. A relevant investigation would therefore be to explore the effect of varying this constraint, which however, is beyond the scope of this work. To give the optimizer more freedom to operate future studies will involve relaxing this constraint to explore the resulting shapes produced by the optimizer.

With respect to the gained improvement in mechanical power it can be said that the optimization results from the finer grid levels (SPL2e3cb, SPL2e3hb, SPL1e3cb, and SPL1e3hb) agree very well at 0.42%–0.44% improvement, in which the result from L3 is lower (0.33%). This observation agrees with earlier findings (Madsen et al., 2019) stating that very coarse mesh levels should be used with care in shape optimizations. Furthermore, it should be pointed out that the SPL3e3c improvement percentage from Table 5 is higher than the SPL3e3cb improvement percentage from Table 6 because the design variable bounds for  $h_{wl}$  and  $s_{wl}$  were changed to 3.5 from 2.0. Thus, one cannot directly compare SPL3e3c and SPL3e3cb.

Finally, while on the topic of improvement it is well to note the importance of regenerating the mesh: as seen, the final shape designs from SPL1e3cb and SPL1e3hb are for completeness evaluated with a regenerated volume mesh to ensure as accurate a result as possible. As a result, the final improvement in mechanical power changes from 0.44% to 1.12%. Importantly, the bending moment fraction constraint only changes from 1.000 to 0.9992, signifying that the final


shape is still feasible and at the upper limit of the constraint as expected.

The 1.12% improvement in mechanical power aligns very well with previous findings from the literature (see, e.g., Table 1): Matheswaran et al. (2019) report a 2.5% power increase for a load-neutral optimization, and Zahle et al. (2018) report an improvement in power of 0.76% for straight blade extensions and up to 2.6% improvement for optimized winglet shapes. The surrogate-based approach by Zahle et al. (2018) used the exact same mesh generator and flow solver, meaning that the results should align fairly well. Given that the developed FFD-based parametrization in the present study does not produce true 90° winglet shapes (see Fig. 15, left, for a comparison) it is reasonable that the expected improvement from the present study should lie somewhere between 0.76% and 2.6%, which is also the case.

The reason that mechanical power changes much more than the bending moment when the mesh is regenerated is that the bending moment computation is driven by pressure and friction, whereas the mechanical power computation is based on drag which is much more mesh-quality-dependent. The same phenomenon is seen in standard 2-D CFD airfoil computations in which lift mainly is a projection of pressure forces, whereas the viscous forces are very important for the drag. A small change in force vector will therefore mean much more for the drag than for the final lift.

It should be clearly stated that the change in final mechanical power in Table 6 from 0.44% to 0.12% deserves further investigation in future studies. Below, some of the planned investigations are described in further detail.

One could opt to restart the optimization on a regenerated mesh around the final L1 shape to obtain a result that is more independent of mesh quality. Alternatively, one could look into other popular methods such as radial basis functions or the inverse distance method to investigate whether these methods produce meshes that are closer in quality to an actual regenerated mesh. Yet another option is to regenerate the mesh after every optimization step. However, given that the aim is to arrive at a high-fidelity optimization framework that also utilizes adjoint solvers, that is not a desirable avenue to pursue since it would further complicate the gradient computation. A final option worth mentioning is to make a dedicated L3 mesh that is not a result of the grid coarsening described in Sect. 4.3. A better way to generate the L3 mesh would be to grow the mesh directly on L3 using the hyperbolic mesh generator. This approach has proved to be very efficient in the past. The dedicated meshes could be combined with grid sequencing to save further time. Table 6 gives two examples (SPL2e3hb, SPL1e3hb) of how optimizations may save time by starting from a coarser mesh level's result: as an example one could start an optimization on L2 (SPL2e3cb) and then proceed on L1 (SPL1e3hb) using the L2 result as seen in Fig. 12. Importantly, the final shape on a given grid level is the same when starting from a straight blade (SPL2e3cb, SPL1e3cb) or when starting from a result



**Figure 11.** Tip vortex visualization using the shape optimization result from grid level L1. Wind direction is along the  $z$  axis. The full rotor image (a) is a superimposed image showing both the baseline rotor (gray) with straight blades and the final rotor shape (red) with the optimized tip shapes. The tip vortex is visualized both for the straight baseline (b) and for the final optimized tip shape (c) to show that it is smeared out and moved away from the rotor plane.

from a coarser level (SPL2e3hb, SPL1e3hb) as one would expect for a well-functioning setup. Given that a speed-up of about a factor of 2 is observed this grid sequencing approach in shape optimization seems very advantageous and should be further investigated.

### 6.2.1 Analysis of the optimized shape

To inspect the optimized blade shape the planform is visualized in Fig. 13, whereafter the resulting spanwise forces are visualized in Fig. 14.

Inspecting the final twist distribution (Fig. 13, upper) a slightly jagged curve is seen where the optimizer tries to continue the upward rising twist curve at the start of the tip, as well as tries to introduce more negative twist towards the very tip of the blade. Reassuringly, these exact same trends were observed in a recent surrogate-based study (Zahle et al., 2018, their Fig. 8c) albeit with a smoother curve. The present study's parametrization has mainly been designed with robustness and mesh quality in mind, focusing on avoiding negative cells on the finest mesh level. Future work will entail experimenting with the parametrization in the attempt to try to arrive at a smoother twist curve. That being said, the resulting spanwise forces are very smooth even for the present setup as can be seen in Fig. 14.

The chord distribution only changes at the very tip of the blade (Fig. 13, lower) where the outermost chord design variable is used to slim the blade. One could attempt to move the second outermost FFD box section towards the tip to activate the remaining chord design variables. This would indeed also give the optimizer more freedom to shape the blade tip. How-

ever, that FFD section has been purposefully placed at a distance to the actual blade tip to protect the cell volumes at the very tip, meaning one should use caution not to compromise the mesh quality. Overall, the trend of slimming the blade using the chord design variable as the blade is extended is to be expected for these optimizations.

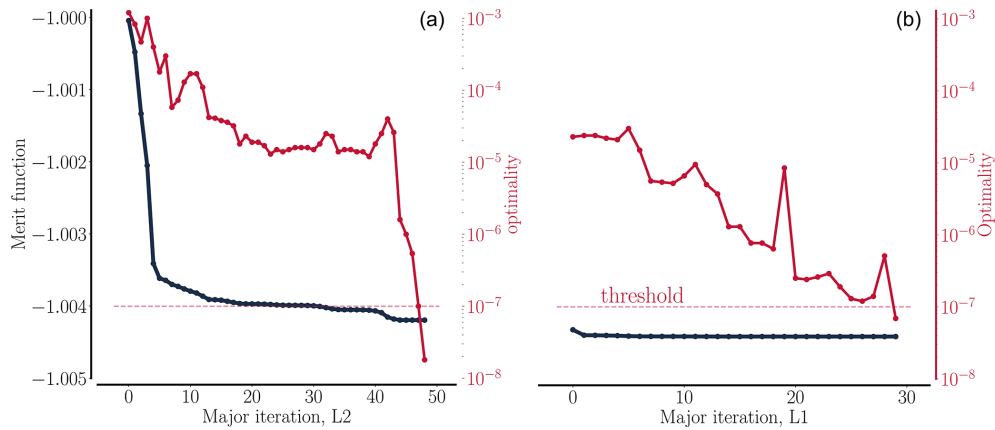
Finally, the resulting spanwise forces are shown in Fig. 14 where the driving force is placed above the normal force. The entire span is shown to the left, and a zoom-in of the outermost part of the blade is shown to the right.

The driving force is exactly as expected for this type of tip optimization: as visible in the zoomed-in plot (right) the optimizer sacrifices a small portion of power at the start of the tip (90–95 m) but generates more power towards the very tip of the extended blade, resulting in a net increase in torque.

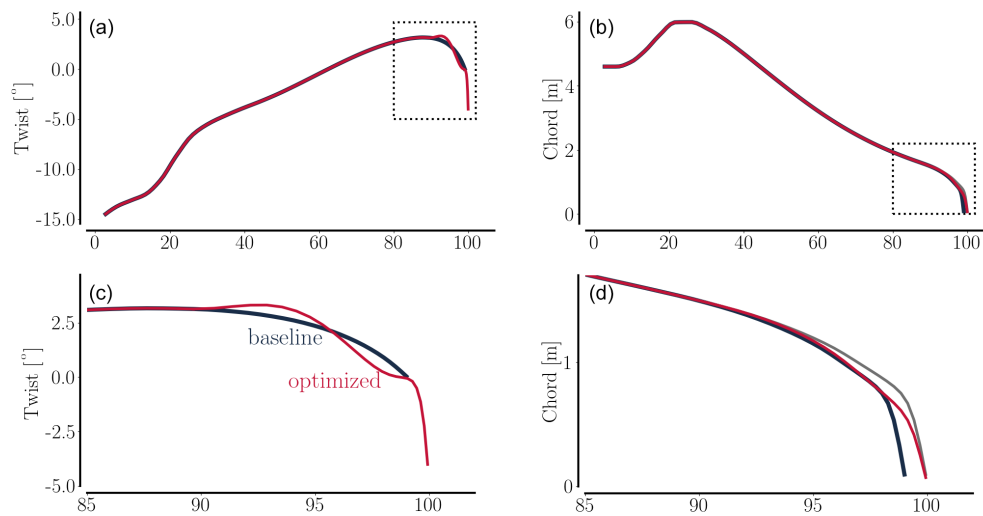
Turning to the normal force for the optimized shape the tip is effectively de-loaded mid-tip ( $\sim 97$  m span) after an initial slight increase at the beginning of the tip. The de-loading allows the optimizer to extend the blade, thus incurring new loads at the very tip.

In sum, a novel curved tip shape has been designed. The results show that with approximately 1% blade extension, 2% flapwise displacement, and slightly below 2% edgewise displacement, one can obtain a 1.12% increase in power. The design favors as much curvature as is possible with the present parameterization, and it is likely that parameterizations allowing for a curvature closer to a  $90^\circ$  winglet-like shape would find an even greater power increase. Indeed, the novel shape is able to extend the blade efficiently in the following manner: using a combination of the twist and chord design variables, it effectively sheds loads at the beginning





**Figure 12.** Example of optimization procedure using grid sequencing: the result from grid level L2 (a) is used as a starting point on grid level L1 (b) to save time. The optimizations are called SPL2e3cb and SPL1e3hb and also figure both in Table 6 and in Fig. 10.



**Figure 13.** Planform distribution for twist and chord for the shape optimization result from grid level L1. Gray chord distribution signifies a stretched blade without any chord scaling in order to better discern where the reduced chord is taking place on the stretched blade.

of the tip (Fig. 14) where also a reduced driving force is observed in the process, while extending the blade with a slender chord distribution, minimizing the load impact from the extension. Flow visualizations (Fig. 11) showed that compared to the original tip, the curved tip shape results in a more smeared out tip vortex, thus reducing tip loss. However, further investigations are needed to fully understand how, for example, the addition of sweep at the tip is favorable to a non-swept tip from an aerodynamic point of view.

### 6.3 A comparison across fidelities


Given that the study by Zahle et al. (2018) is closely related to the present work it is rewarding to compare the parametrizations and in particular the difference in maximally allowed curvature by superimposing a tip shape result from the present work (red) on to the final design (gray) from

the surrogate-based study by Zahle et al. (2018) as seen in Fig. 15. Inspecting the upper-left plot in Fig. 15 it is evident that the present study cannot produce true  $90^\circ$  winglet shapes due to concerns of the mesh quality, which is why one should expect to find final improvements slightly lower than in the study by Zahle et al. (2018). Reassuringly, a 1.12% increase in power is reported in the present study, whereas a 2.6% increase in power is reported in Zahle et al. (2018).


### 6.4 Future work

As a first step one should further increase the robustness and efficiency of the presented FlowOpt framework. This includes adding enhanced convergence methods and an adjoint method.

An enhanced convergence method would, apart from a general increase in optimization robustness, also ensure that



**Figure 14.** Driving (a, c) and normal (b, d) forces for the shape optimization result from grid level L1. Dashed black rectangles (a, b) indicate limits for the zoomed-in plots (c, d).




**Figure 15.** Resulting tip shape from the present study (red) compared to the surrogate-based design by Zahle et al. (2018) (gray). Both results exhibit maximally allowed curvature by their respective parametrization. As seen, the curvature from the present work (red) is not a true  $90^\circ$  winglet as in the study by Zahle et al. (2018), signifying that one should expect a difference in final possible improvement (final results are 1.12% and 2.6%, respectively). The upper three figure panels show different views in the rotor plane, whereas the lower panel is in the upstream direction.

aerodynamically challenging shapes due to, for example, stall, could be handled (see, e.g., Fig. 16), meaning that a significant increase in the design space could be gained as it would only be due to negative cells that one would have to limit the design variables.

The adjoint method allows for a gradient computation time that is independent of the number of design variables, which would be highly relevant for future work. Notice that the adjoint method will most likely not result in a speed-up of what can already be achieved for the presented optimizations with the current framework when using the parallelization techniques in OpenMDAO to compute perturbations for various

design variables simultaneously. In fact, given that not all adjoint solvers are as fast as their flow solver counterparts one might see a slight slowdown. However, given that the present study uses the finite difference method which clearly depends on a preceding step size study, the final run times are difficult to predict and discuss without a concrete comparison being done.

Leaving the discussion of possible lower bounds for run times aside, one thing that is certain is that the framework with an adjoint solver will be able to take on new optimization problems altogether since the gradient computation will be independent of the number of design variables. Thus, full



**Figure 16.** A sequence of three different deforming curved tip shapes (red) imposed on the baseline blade (gray) where surface-restricted streamlines can be seen as thin black lines. The flow visualization shows an emerging stall region leading to the impaired flow convergence. Shedding of vorticity is introduced as the blade is further twisted. These shapes were created by the optimizer during an early exploratory optimization. They are easier to solve with, for example, unsteady methods, and the design space was therefore limited enough to avoid the stall region.

shape optimizations using free-form techniques will be manageable – something one can never hope to do with the finite difference method.

Finally, it would be very relevant to validate the numerical results against experimental data in future studies (see, e.g., Barlas et al., 2021a).

### 6.5 Learning outcomes

Based on the above-described detailed high-fidelity shape optimization study, these are the overall findings.

- A thorough literature review showed that there is a lack of high-fidelity shape optimization studies within wind energy in which most works simply are parameter studies.
- Robust mesh deformation is an absolute key feature for high-fidelity shape optimization with  $y^+ \sim 1$  and  $\mathcal{O}(10^7)$  cells.
- In order to explore a larger design space (e.g., stall regions) there is a need for enhanced convergence methods which would also bring about an increase in robustness.
- Meticulous setup of the finite difference method will allow for deeply converged design optimization problems even without machine-accurate gradients.
- Although the finite difference method is a viable approach high-fidelity shape design, the authors can conclude based on experience with both direct CFD-based optimizations and surrogate-based optimizations that due to ease of use and a much lower computational cost, one should prefer the surrogate-based approach for optimizations of up to about a dozen design variables if one can accept the drop in model fidelity.

## 7 Conclusions


In this study a novel curved tip shape was aerodynamically designed for maximum power using a CFD solver on a 10 MW reference wind turbine with the constraint that the initial steady-state loads should not be compromised. The study showed that a 1.12% increase in power was possible while satisfying the imposed constraints on loads and geometry. The final curved tip results in a 1% blade length extension, a 2% flapwise tip displacement, just below a 2% edgewise tip displacement, and as much tip curvature as possible within the developed parameterization. Using twist and chord design variables to reduce loads and slim the outermost part of the blade, the novel curved tip shape efficiently extends the blade, and a flow analysis visualized how the final design effectively displaces the tip vortex to mitigate induced drag. A tip design as the one presented could be mounted on already installed wind turbines as a sleeve-like solution or be conceived as part of a modular blade with tips designed for site-specific conditions. Importantly, this study was not aeroelastic but aerodynamic only. Only steady-state flow conditions and normal operation were considered, and a detailed unsteady load analysis based on the IEC (International Electrotechnical Commission) standard was thus not carried out, which would be needed to design the structural geometry of the blade tip and ensure that it is indeed load-neutral.

Turning to the numerical aspect of the study it can be concluded that it is indeed possible to tightly converge direct CFD-based design optimization problems using the finite difference method as long as a meticulous step size study is carried out. However, the finite difference method is found to be extremely expensive on industrial-scale cases (above 14 million cell meshes), and a surrogate-based approach should be favored due to ease of use and implementation as long as a drop in model fidelity can be accepted. Furthermore,

the study revealed that robust mesh deformation routines are very important for a successful optimization framework.

A comprehensive literature review on blade tips preceding the optimizations revealed many overall favorable design trends could be identified. Furthermore, it was found that up to a 2.5%–2.6% increase in power should be possible for winglet-like load-neutral tip shapes. However, as also pointed out in the literature review there is a void of high-fidelity shape optimization results, and this study shows at least one way to set up the essential components in a framework for CFD-based optimization and how to tightly converge the design optimization problems through a meticulous fine tuning of the setup.

### Appendix A: Visualization in the rotor plane of the optimized tip shape



**Figure A1.** Baseline (gray) and optimized (red) blade shape from different angles.

**Data availability.** Data are available upon request to the corresponding author.

**Author contributions.** MHAM carried out the literature review, wrote the FFDlib, ran all optimizations, and wrote the bulk of the paper. FZ developed PGL and assisted in work related to the parameterization and optimization. Also, FZ assisted in postprocessing and analysis of all results and in the comparison to surrogate-based

optimizations. SGH provided meshes and helped guide the study – particularly in the early stage (develop parameterization, ensuring cross-platform compatibility, etc.) – and assisted in the analysis of results. Furthermore, SGH ensured the final parameterization is aligned with the approach used in the DTU Wind Energy FSI framework for future work. TKB guided the project, assisted with expert knowledge on blade tip design, and aligned the present study with the surrounding SmartTip study. NNS, as one of the two original EllipSys3D developers, advised on computational setup and parallelization and implemented the enhanced mesh deformation algorithm used in the present study. He also worked on the EllipSys3D coarse grid solver complexification. All authors took part in writing and editing the paper.

**Competing interests.** The contact author has declared that neither they nor their co-authors have any competing interests.

**Disclaimer.** Publisher's note: Copernicus Publications remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

**Acknowledgements.** A special mention of the head of section (ARD), Flemming Rasmussen, DTU Wind Energy, is also warranted. He has through several years supported the development of a high-fidelity design framework without which this study would not have been possible.

Finally, the authors thank PhD student Neil Wu (MDOLab, University of Michigan) and OpenMDAO project lead Justin S. Gray (NASA, Glenn) for assistance with pyOptSparse and OpenMDAO, respectively.

**Financial support.** This research has been supported by the Innovation Fund Denmark (grant no. 7046-00023B).

**Review statement.** This paper was edited by Horia Hangan and reviewed by two anonymous referees.

### References

- Aju, E., Suresh, D., and Jin, Y.: The influence of winglet pitching on the performance of a model wind turbine: Aerodynamic loads, rotating speed, and wake statistics, *Energies*, 13, 5199, <https://doi.org/10.3390/en13195199>, 2020.
- Andersen, B., Bak, C., Aagaard Madsen, H., and Søndergård, B.: Måling og beregning af aerodynamisk støj samt optimering af vingetipper på vindmøller, Tech. rep., <https://findit.dtu.dk/en/catalog/58945f798040e5ab4503f864> (last access: 3 July 2022), 2001.
- Aravindkumar, N: Analysis of the Small Wind Turbine Blade with and Without Winglet, *International Journal of Engineering Research and Applications* 4, 239–243, 2014.
- Ariffudin, M., Zawawi, F., Kamar, H., and Kamsah, N.: Effectiveness of blade tip on low speed horizontal axis



- wind turbine performance, *Jurnal Teknologi*, 78, 31–39, <https://doi.org/10.11113/jt.v78.9582>, 2016.
- Bak, D., Andersen, P., Bertagnolio, F., Buhl, T., Gaunaa, M., Hansen, A., Hansen, M., Johansen, J., Larsen, G., Larsen, T., Madsen, H., Markou, H., Rasmussen, F., Sørensen, N., Thomsen, K., Hansen, K., Hansen, M., Mikkelsen, R., Shen, W., Sørensen, J., Trolborg, N., and Øye, S.: Research in aeroelasticity EFP-2006, Risø National Laboratory, Denmark, Forskningscenter Risoe, Risoe-R No. 1611(EN), [https://backend.orbit.dtu.dk/ws/files/7703181/ris\\_r\\_1611.pdf](https://backend.orbit.dtu.dk/ws/files/7703181/ris_r_1611.pdf) (last access: 3 July 2022), 2007.
- Barlas, T., Pirrung, G. R., Ramos-García, N., Horcas, S. G., Mikkelsen, R. F., Olsen, A. S., and Gaunaa, M.: Wind tunnel testing of a swept tip shape and comparison with multi-fidelity aerodynamic simulations, *Wind Energ. Sci.*, 6, 1311–1324, <https://doi.org/10.5194/wes-6-1311-2021>, 2021a.
- Barlas, T., Ramos-García, N., Pirrung, G. R., and González Horcas, S.: Surrogate-based aeroelastic design optimization of tip extensions on a modern 10MW wind turbine, *Wind Energ. Sci.*, 6, 491–504, <https://doi.org/10.5194/wes-6-491-2021>, 2021b.
- Barrett, R. and Ning, A.: Integrated free-form method for aerostuctural optimization of wind turbine blades, *Wind Energy*, 21, 663–675, <https://doi.org/10.1002/we.2186>, 2018.
- Berg, J., Trolborg, N., Menke, R., Patton, E. G., Sullivan, P. P., Mann, J., and Sørensen, N.: Flow in complex terrain – a Large Eddy Simulation comparison study, *J. Phys.-Conf. Ser.*, 1037, 072015, <https://doi.org/10.1088/1742-6596/1037/7/072015>, 2018.
- Bortolotti, P., Tarrés, H. C., Dykes, K., Merz, K., Sethuraman, L., Verelst, D., and Zahle, F.: IEA Wind TCP Task 37: Systems Engineering in Wind Energy-WP2.1 Reference Wind Turbines, Tech. rep., <https://doi.org/10.2172/1529216>, 2019.
- Casale, M. and Stanton, E.: An overview of analytic solid modeling, *IEEE Comput. Graph.*, 5, 45–56, <https://doi.org/10.1109/MCG.1985.276402>, 1985.
- Dhert, T., Ashuri, T., and Martins, J. R. R. A.: Aerodynamic Shape Optimization of Wind Turbine Blades Using a Reynolds-Averaged Navier–Stokes Model and an Adjoint Method, *Wind Energy*, 20, 909–926, <https://doi.org/10.1002/we.2070>, 2017.
- Døssing, M.: Vortex lattice modelling of winglets on wind turbine blades, Danmarks Tekniske Universitet, Risø Nationallaboratoriet for Bæredygtig Energi, Denmark, Forskningscenter Risoe, Risoe-R No. 1621(EN), [https://backend.orbit.dtu.dk/ws/files/151422033/ris\\_r\\_1621.pdf](https://backend.orbit.dtu.dk/ws/files/151422033/ris_r_1621.pdf) (last access: 3 July 2022), 2007.
- Ebrahimi, A. and Mardani, R.: Tip-Vortex Noise Reduction of a Wind Turbine Using a Winglet, *J. Energ. Eng.*, 144, 04017076, [https://doi.org/10.1061/\(ASCE\)EY.1943-7897.0000517](https://doi.org/10.1061/(ASCE)EY.1943-7897.0000517), 2018.
- Elfarrá, M., Sezer-Uzol, N., and Akmandor, I.: NREL VI rotor blade: Numerical investigation and winglet design and optimization using CFD, *Wind Energy*, 17, 605–626, <https://doi.org/10.1002/we.1593>, 2014.
- Elfarrá, M., Sezer-Uzol, N., and Akmandor, I. Sinan: Investigations on blade tip tilting for hawt rotor blades using CFD, *Int. J. Green Energy*, 12, 125–138, <https://doi.org/10.1080/15435075.2014.889007>, 2015.
- Elfarrá, M. A.: Horizontal axis wind turbine rotor blade: winglet and twist aerodynamic design and optimization using CFD, PhD thesis, <https://open.metu.edu.tr/handle/11511/21223> (last access: 3 July 2022), 2011.
- Farhan, A., Hassanpour, A., Burns, A., and Motlagh, Y.: Numerical study of effect of winglet planform and airfoil on a horizontal axis wind turbine performance, *Renew. Energ.*, 131, 1255–1273, <https://doi.org/10.1016/j.renene.2018.08.017>, 2019.
- Ferrer, E. and Munduate, X.: Wind turbine blade tip comparison using CFD, *J. Phys.-Conf. Ser.*, 75, 012005, <https://doi.org/10.1088/1742-6596/75/1/012005>, 2007.
- Gaunaa, M. and Johansen, J.: Determination of the maximum aerodynamic efficiency of wind turbine rotors with winglets, *J. Phys.-Conf. Ser.*, 75, 012006, <https://doi.org/10.1088/1742-6596/75/1/012006>, 2007.
- Gaunaa, M. and Johansen, J.: Can CB be Increased by the Use of Winglets? – or – A Theoretical and Numerical Investigation of the Maximum Aerodynamic Efficiency of Wind Turbine Rotors with Winglets, AIAA 2008-1314, 46th AIAA Aerospace Sciences Meeting and Exhibit, <https://doi.org/10.2514/6.2008-1314>, 2008.
- Gaunaa, M., Réthoré, P.-E. M., Sørensen, N. N., and Døssing, M.: A computational efficient algorithm for the aerodynamic response of non-straight blades, in: Proceedings European Wind Energy Association (EWEA), [https://backend.orbit.dtu.dk/ws/portalfiles/portal/6372784/Rethore\\_paper\\_ewea2011\\_paper.pdf](https://backend.orbit.dtu.dk/ws/portalfiles/portal/6372784/Rethore_paper_ewea2011_paper.pdf) (last access: 3 July 2022), 2011.
- Gaunaa, M., Zahle, F., Sørensen, N. N., Bak, C., and Réthoré, P.-E.: Rotor Performance Enhancement Using Slats on the Inner Part of a 10MW Rotor, *Proceedings Og Ewea 2013*, 2, 1293–1302, 2013.
- Gertz, D. and Johnson, D. A.: An evaluation testbed for wind turbine blade tip designs-baseline case, *Int. J. Energ. Res.*, 35, 1360–1370, <https://doi.org/10.1002/er.1897>, 2011.
- Gertz, D., Johnson, D., and Swytink-Binnema, N.: An evaluation testbed for wind turbine blade tip designs – Winglet results, *Wind Engineering*, 36, 389–410, <https://doi.org/10.1260/0309-524X.36.4.389>, 2012.
- Gill, P. E., Murray, W., and Saunders, M. A.: SNOPT: An SQP algorithm for large-scale constrained optimization, *SIAM Rev.*, 47, 99–131, 2005.
- Gill, P. E., Murray, W., Saunders, M. A., and Wong, E.: User’s Guide for SNOPT 7.7: Software for Large-Scale Nonlinear Programming, Center for Computational Mathematics Report CCoM 18-1, Department of Mathematics, University of California, San Diego, La Jolla, CA, <https://ccom.ucsd.edu/~optimizers/static/pdfs/snopt7-7.pdf> (last access: 3 July 2022), 2018.
- Gray, J. S., Hwang, J. T., Martins, J. R. R. A., Moore, K. T., and Naylor, B. A.: OpenMDAO: An open-source framework for multidisciplinary design, analysis, and optimization, *Struct. Multidiscip. O.*, 59, 1075–1104, <https://doi.org/10.1007/s00158-019-02211-z>, 2019.
- Gupta, A. and Amano, R.: CFD analysis of wind turbine blade with winglets, *Proceedings of the Asme Design Engineering Technical Conference*, 5, 843–849, <https://doi.org/10.1115/DETC2012-70679>, 2012.
- Hahmann, S., Bonneau, G.-P., Barbier, S., Elber, G., and Hagen, H.: Volume-preserving FFD for programmable graphics hardware, *Visual Comput.*, 28, 231–245, <https://doi.org/10.1007/s00371-011-0608-5>, 2012.
- Hansen, T.: Aerodynamic Optimisation of Airfoils and Winglets for Wind Turbine Application, PhD thesis, Norway, <https://doi.org/10.1007/s00158-019-02211-z>, 2011.

- [/ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/2463515](https://ntnuopen.ntnu.no/ntnu-xmlui/handle/11250/2463515) (last access: 3 July 2022), 2017.
- Hansen, T. and Mühle, F.: Winglet optimization for a model-scale wind turbine, *Wind Energy*, 21, 634–649, <https://doi.org/10.1002/we.2183>, 2018.
- Hansen, T. H.: Airfoil optimization for wind turbine application, *Wind Energy*, 21, 502–514, <https://doi.org/10.1002/we.2174>, 2018.
- Horcas, S., Madsen, M., Sørensen, N., and Zahle, F.: Suppressing Vortex Induced Vibrations of Wind Turbine blades with flaps, in: *Recent Advances in CFD for Wind and Tidal Offshore Turbines*, Springer Tracts in Mechanical Engineering, edited by: Ferrer, E. and Montlaur, A., European Community on Computational Methods in Applied Sciences, Springer, Cham, [https://doi.org/10.1007/978-3-030-11887-7\\_2](https://doi.org/10.1007/978-3-030-11887-7_2), 2018.
- Imamura, H., Hasegawa, Y., and Kikuyama, K.: Numerical analysis of the horizontal axis wind turbine with winglets, *Jsmc Int. J. B-Fluids T.*, 41, 170–176, <https://doi.org/10.1299/jsmeb.41.170>, 1998.
- Johansen, J. and Sørensen, N. N.: Aerodynamic investigation of winglets on wind turbine blades using CFD, Denmark, Forskningscenter Risoe, Risoe-R No. 1543(EN), [https://backend.orbit.dtu.dk/ws/files/7703268/tris\\_r\\_1543.pdf](https://backend.orbit.dtu.dk/ws/files/7703268/tris_r_1543.pdf) (last access: 3 July 2022), 2006.
- Johansen, J. and Sørensen, N. N.: Numerical analysis of winglets on wind turbine blades using CFD, *Conference Proceedings*, 2, 1184–1189, 2007.
- Johansen, J., Sørensen, N. N., Michelsen, J. A., and Schreck, S.: Detached-Eddy simulation of flow around the NREL phase-VI blade, *Asme 2002 Wind Energy Symposium*, Wind2002, 106–114, <https://doi.org/10.1115/wind2002-32>, 2002.
- Johansen, J., Gaunaa, M., and Sørensen, N. N.: Increased Aerodynamic Efficiency on Wind Turbine Rotors using Winglets, *Technical Papers*, <https://doi.org/10.2514/6.2008-6728>, 2008.
- Kalken, J. and Ceyhan, O.: InnoTIP end report rev 2, Tech. Rep. ECN-O-17-011, <https://publications.ecn.nl/WIN/2017/ECN-O--17-011> (last access: 3 July 2022), 2017.
- Khalafallah, M. G., Ahmed, A., and Emam, M.: The effect of using winglets to enhance the performance of swept blades of a horizontal axis wind turbine, *Adv. Mech. Eng.*, 11, 1–10, <https://doi.org/10.1177/1687814019878312>, 2019.
- Khaled, M., Ibrahim, M., Abdel Hamed, H., and AbdelGwad, A.: Investigation of a small Horizontal-Axis wind turbine performance with and without winglet, *Energy*, 187, 115921, <https://doi.org/10.1016/j.energy.2019.115921>, 2019.
- Kulak, M., Lipian, M., and Zawadzki, K.: Investigation of performance of small wind turbine blades with winglets, *Int. J. Numer. Method. H.*, 31, 629–640, <https://doi.org/10.1108/HFF-09-2019-0695>, 2021.
- Lanchester, F.: Provisional specification: Improvements in and relating to Aerial Machines, Patent nr. 3608 (LAN/6/34), <https://catalogue.lanchesterinteractive.org/records/LAN/6/34> (last access: 3 July 2022), 1897.
- Li, A., Pirrung, G., Madsen, H. A., Gaunaa, M., and Zahle, F.: Fast trailed and bound vorticity modeling of swept wind turbine blades, *J. Phys.-Conf. Ser.*, 1037, 062012, <https://doi.org/10.1088/1742-6596/1037/6/062012>, 2018.
- Lissaman, P. and Gyatt, G.: Development and testing of tip devices for horizontal axis wind turbines, Tech. Rep. NASA-CR-174991, NATIONAL AERONAUTICS AND SPACE ADMINISTRATION (NASA), Lewis Research Center, <https://www.osti.gov/servlets/purl/5988121> (last access: 3 July 2022), 1985.
- Lyness, J. N.: Numerical algorithms based on the theory of complex variable, *Proceedings of the 1967 22nd National Conference*, 125–133, <https://doi.org/10.1145/800196.805983>, 1967.
- Lyness, J. N. and Moler, C. B.: Numerical Differentiation of Analytic Functions, *Siam J. Numer. Anal.*, 4, 202–210, <https://doi.org/10.1137/0704019>, 1967.
- Madsen, H. and Fuglsang, P.: Numerical investigation of different tip shapes for wind turbine blades. Aerodynamic and aeroacoustic aspects, Tech. Rep. No. 891, Risø National Laboratory, [https://backend.orbit.dtu.dk/ws/files/7753502/RIS\\_R\\_891.pdf](https://backend.orbit.dtu.dk/ws/files/7753502/RIS_R_891.pdf) (last access: 3 July 2022), 1997.
- Madsen, M.: High-Fidelity CFD-based Shape Optimization of Wind Turbine Blades, PhD thesis, Denmark, <https://doi.org/10.11581/dtu:00000068>, 2020.
- Madsen, M. H. Aa., Zahle, F., Sørensen, N. N., and Martins, J. R. R. A.: Multipoint high-fidelity CFD-based aerodynamic shape optimization of a 10 MW wind turbine, *Wind Energ. Sci.*, 4, 163–192, <https://doi.org/10.5194/wes-4-163-2019>, 2019.
- Matheswaran, V., Miller, L., and Moriarty, P.: Retrofit winglets for wind turbines, *Aiaa Scitech 2019 Forum*, <https://doi.org/10.2514/6.2019-1297>, 2019.
- Menter, F.: Improved two-equation  $k$ - $\omega$  turbulence models for aerodynamic flows, <https://ntrs.nasa.gov/citations/19930013620> (last access: 3 July 2022), 1992.
- Michelsen, J. A.: Basis3D – a Platform for Development of Multi-block PDE Solvers, Technical Report AFM 92-05, Technical University of Denmark, 1992.
- Michelsen, J. A.: Block structured Multigrid solution of 2D and 3D elliptic PDE's, Technical Report AFM 94-06, Technical University of Denmark, 1994.
- Mourad, M. G., Shahin, I., Ayad, S. S., Abdellatif, O. E., and Mekhaïl, T. A.: Effect of winglet geometry on horizontal axis wind turbine performance, *Engineering Reports*, 2, e12101, <https://doi.org/10.1002/eng.2.12101>, 2020.
- Mühle, F., Bartl, J., Hansen, T., Adaramola, M., and Sætran, L.: An experimental study on the effects of winglets on the tip vortex interaction in the near wake of a model wind turbine, *Wind Energy*, 23, 1286–1300, <https://doi.org/10.1002/we.2486>, 2020.
- Nielsen, E. J. and Diskin, B.: Discrete adjoint-based design for unsteady turbulent flows on dynamic overset unstructured grids, 50th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition, Nashville, Tennessee, 9–12 January 2012, AIAA 2012-0554, <https://doi.org/10.2514/6.2012-554>, 2012.
- Ning, A., Damiani, R., and Moriarty, P.: Objectives and constraints for wind turbine optimization, *J. Sol. Energ.-T. Asme*, 136, 041010, <https://doi.org/10.1115/1.4027693>, 2014.
- Papadopoulos, C., Schmid, M., Kaparos, P., Misirlis, D., and Vlachostergios, Z.: Numerical analysis and optimization of a winglet for a small horizontal wind turbine blade, *Chem. Engineer. Trans.*, 81, 1321–1326, <https://doi.org/10.3303/CET2081221>, 2020.
- Patankar, S.: Numerical heat transfer and fluid flow, Taylor & Francis, ISBN 978-0-89116-522-4, 1980.
- Reddy, S. R., Dulikravich, G. S., Sobieczky, H., and Gonzalez, M.: Bladelets-Winglets on Blades of Wind Turbines: A Multiobjec-

- tive Design Optimization Study, *J. Sol. Energ.-T. Asme*, 141, 061003, <https://doi.org/10.1115/1.4043657>, 2019.
- Rhie, C. and Chow, W.: Numerical study of the turbulent flow past an airfoil with trailing edge separation, *AIAA J.*, 21, 1525–1532, 1983.
- Sederberg, T. W. and Parry, S. R.: Free-form deformation of solid geometric models, *ACM Siggraph Computer Graphics*, 20, 151–160, <https://doi.org/10.1145/15922.15903>, 1986.
- Simms, D., Schreck, S., Hand, M., and Fingersh, L.: NREL Unsteady Aerodynamics Experiment in the NASA-Ames Wind Tunnel: A Comparison of Predictions to Measurements, technical report, <https://doi.org/10.2172/783409>, 2001.
- Sørensen, J. N., Shen, W. Z., Zhu, W. J., Borbye, J., Okulov, V., Mikkelsen, R. F., Gaunaa, M., Réthoré, P.-E. M., and Sørensen, N. N.: Design og optimering af vingetipper for vindmøller: Slutrapport, DTU Mekanik, Tech. rep., <https://backend.orbit.dtu.dk/ws/files/5515584/NEI-DK-5519.pdf> (last access: 3 July 2022), 2011.
- Sørensen, N.: General purpose flow solver applied to flow over hills, PhD thesis, [https://backend.orbit.dtu.dk/ws/files/12280331/Ris\\_R\\_827.pdf](https://backend.orbit.dtu.dk/ws/files/12280331/Ris_R_827.pdf) (last access: 3 July 2022), 1995.
- Sørensen, N.: HypGrid2D a 2-D mesh generator, Tech. rep., [https://backend.orbit.dtu.dk/ws/files/7750949/RIS\\_R\\_1035.pdf](https://backend.orbit.dtu.dk/ws/files/7750949/RIS_R_1035.pdf) (last access: 3 July 2022), 1998.
- Sørensen, N., Michelsen, J., and Schreck, S.: Navier-Stokes predictions of the NREL phase VI rotor in the NASA Ames 80-by-120 wind tunnel, 2002 Asme Wind Energy Symposium; 40. AIAA Aerospace Sciences Meeting and Exhibit, Reno, NV, 14–17 January 2002, 94–105, 2002.
- Sørensen, N. N.: CFD modelling of laminar-turbulent transition for airfoils and rotors using the  $\gamma - \tilde{R}e_{\theta t}$ , *Wind Energy*, 12, 715–733, <https://doi.org/10.1002/we.325>, 2009.
- Sy, M., Abuan, B., and Macapili Danao, L.: Aerodynamic investigation of a horizontal axis wind turbine with split winglet using computational fluid dynamics, *Energies*, 13, 4983, <https://doi.org/10.3390/en13184983>, 2020.
- Tobin, N., Hamed, A., and Chamorro, L.: An Experimental Study on the Effects of Winglets on the Wake and Performance of a Model Wind Turbine, *Energies*, 8, 11955–11972, <https://doi.org/10.3390/en81011955>, 2015.
- Vorspel, L., Stoevesandt, B., and Peinke, J.: Optimize rotating wind energy rotor blades using the adjoint approach, *Appl. Sci. (Switzerland)*, 8, 1112, <https://doi.org/10.3390/app8071112>, 2018.
- Whitcomb, R.: A design approach and selected wind tunnel results at high subsonic speeds for wing-tip mounted winglets, <https://ntrs.nasa.gov/citations/19760019075> (last access: 3 July 2022), 1976.
- Wu, N., Kenway, G., Mader, C. A., Jasa, J., and Martins, J. R. R. A.: pyOptSparse: A Python framework for large-scale constrained nonlinear optimization of sparse systems, *Journal of Open Source Software*, 5, 2564, <https://doi.org/10.21105/joss.02564>, 2020.
- Zahle, F.: Wind turbine aerodynamics using an incompressible overset grid method, Imperial College of Science, Technology and Medicine thesis, [https://backend.orbit.dtu.dk/ws/files/204995349/thesis\\_zahle\\_final\\_published.pdf](https://backend.orbit.dtu.dk/ws/files/204995349/thesis_zahle_final_published.pdf) (last access: 3 July 2022), 2006.
- Zahle, F., Bak, C., Sørensen, N. N., Vronsky, T., and Gaudern, N.: Design of the LRP airfoil series using 2D CFD, *J. Phys.-Conf. Ser.*, 524, 012020, <https://doi.org/10.1088/1742-6596/524/1/012020>, 2014.
- Zahle, F., Sørensen, N. N., McWilliam, M. K., and Barlas, A.: Computational fluid dynamics-based surrogate optimization of a wind turbine blade tip extension for maximising energy production, *J. Phys.-Conf. Ser.*, 1037, 042013, <https://doi.org/10.1088/1742-6596/1037/4/042013>, 2018.
- Zhu, B., Sun, X., Wang, Y., and Huang, D.: Performance characteristics of a horizontal axis turbine with fusion winglet, *Energy*, 120, 431–440, <https://doi.org/10.1016/j.energy.2016.11.094>, 2017.