



All Theses and Dissertations

2018-01-01

Investigation into Integrated Free-Form and Precomputational Approaches for Aerostructural Optimization of Wind Turbine Blades

Ryan Timothy Barrett
Brigham Young University

Follow this and additional works at: <https://scholarsarchive.byu.edu/etd>

 Part of the [Mechanical Engineering Commons](#)

BYU ScholarsArchive Citation

Barrett, Ryan Timothy, "Investigation into Integrated Free-Form and Precomputational Approaches for Aerostructural Optimization of Wind Turbine Blades" (2018). *All Theses and Dissertations*. 6673.
<https://scholarsarchive.byu.edu/etd/6673>

This Thesis is brought to you for free and open access by BYU ScholarsArchive. It has been accepted for inclusion in All Theses and Dissertations by an authorized administrator of BYU ScholarsArchive. For more information, please contact scholarsarchive@byu.edu, ellen_amatangelo@byu.edu.

Investigation into Integrated Free-Form and Precomputational Approaches for
Aerostructural Optimization of Wind Turbine Blades

Ryan Timothy Barrett

A thesis submitted to the faculty of
Brigham Young University
in partial fulfillment of the requirements for the degree of
Master of Science

S. Andrew Ning, Chair
Steven E. Gorrell
John L. Salmon

Department of Mechanical Engineering
Brigham Young University

Copyright © 2018 Ryan Timothy Barrett
All Rights Reserved

ABSTRACT

Investigation into Integrated Free-Form and Precomputational Approaches for Aerostructural Optimization of Wind Turbine Blades

Ryan Timothy Barrett
Department of Mechanical Engineering, BYU
Master of Science

A typical approach to optimize wind turbine blades separates the airfoil shape design from the blade planform design. This approach is sequential, where the airfoils along the blade span are pre-selected or optimized and then held constant during the blade planform optimization. In contrast, integrated blade design optimizes the airfoils and the blade planform concurrently and thereby has the potential to reduce cost of energy (COE) more than sequential design. Nevertheless, sequential design is commonly performed because of the ease of precomputation, or the ability to compute the airfoil analyses prior to the blade optimization. This research investigates two integrated blade design approaches, the precomputational and free-form methods, that are compared to sequential blade design.

The first approach is called the precomputational method because it maintains the ability to precompute, similar to sequential design, and allows for partially flexible airfoil shapes. This method compares three airfoil analysis methods: a panel method (XFOIL), a Reynolds-averaged Navier-Stokes computational fluid dynamics method (RANS CFD), and using wind tunnel data. For each airfoil analysis method, there are two airfoil parameterization methods: the airfoil thickness-to-chord ratio and blended airfoil family factor. The second approach is called the free-form method because it allows for fully flexible airfoil shapes, but no longer has the ease of precomputation as the airfoil analyses are performed during the blade optimization. This method compares XFOIL and RANS CFD using the class-shape-transformation (CST) method to parameterize the airfoil shapes. This study determines if the precomputational method can capture the majority of the benefit from integrated design or if there is a significant additional benefit from the free-form method.

Optimizing the NREL 5-MW reference turbine shows that integrated design reduce COE significantly more than sequential design. The precomputational method improved COE more than sequential design by 1.6%, 2.8%, and 0.7% using the airfoil thickness-to-chord ratio, and by 2.2%, 3.3%, and 1.4% using the blended airfoil family factor when using XFOIL, RANS CFD, and wind tunnel data, respectively. The free-form method improved COE more than sequential design by 2.7% and 4.0% using the CST method with XFOIL and RANS CFD, respectively. The additional flexibility in airfoil shape reduced COE primarily through an increase in annual energy production. The precomputational method captures the majority of the benefit of integrated design (about 80%) for minimal additional computational cost and complexity, but the free-form method provides modest additional benefits if the extra effort is made in computational cost and development time.

Keywords: wind turbine optimization, integrated blade design, free-form, precomputational

ACKNOWLEDGMENTS

I would like to thank my graduate advisor, Dr. Andrew Ning, for his great support, assistance, and unlimited patience. I would like to thank my other committee members, Dr. Steve Gorrell and Dr. John Salmon, for their assistance and guidance. I would like to thank my colleagues, Eric Tingey and Jared Thomas, for answering thousands of my questions as well as just listening until I would answer my own questions. I would like to acknowledge the National Renewable Energy Laboratory for helping to fund this research. And finally, I would like to acknowledge and thank my beautiful wife, Audrey, for her patience and support during this entire process.

TABLE OF CONTENTS

LIST OF TABLES	vi
LIST OF FIGURES	vii
NOMENCLATURE	ix
Chapter 1 Introduction	1
1.1 Wind Energy	1
1.2 Wind Turbines	2
1.3 Wind Turbine Blades	4
1.4 Cost of Energy	5
1.5 Blade Design	6
1.5.1 Sequential Blade Design	6
1.5.2 Integrated Blade Design	7
1.5.3 Precomputation	7
1.6 Contribution	9
1.7 Outline	14
Chapter 2 Aerostructural Blade Analysis and Optimization	15
2.1 Aerodynamics	15
2.1.1 Airfoil Analysis	15
2.1.2 Blade Element Momentum Theorem	19
2.2 Structures	20
2.3 Uncertainty	21
2.4 General Optimization Setup	22
Chapter 3 Precomputational Methods	26
3.1 Theory and Methodology	26
3.2 Parameterization Methods	27
3.2.1 Thickness-to-chord Ratio	27
3.2.2 Blended Airfoil Family Factor	28
3.3 Surrogate Model	29
3.4 Airfoil Analysis Correction (Wind Tunnel)	29
3.5 Results	31
3.6 Discussion and Conclusions	33
Chapter 4 Free-Form Methods	39
4.1 Theory and Methodology	39
4.2 Parameterization Method - Class Shape Transformation	40
4.3 Airfoil Shape Gradients	42
4.4 Results	45
4.5 Discussion and Conclusions	47

Chapter 5	Conclusions	52
5.1	Thesis Review	52
5.1.1	Research Objectives	52
5.1.2	Main Takeaways	54
5.2	Future Work	55
REFERENCES		57
Appendix A	Detailed Optimization Results	61

LIST OF TABLES

1.1	Optimization Cases Summary	14
2.1	Number of Design Variables Summary	24
2.2	Constraints Summary	25
3.1	Precomputational Method Results	32
3.2	Improvement in COE Reduction - Precomputational over Sequential	34
4.1	Free-Form Method Results	46
4.2	Improvement in COE Reduction - Free-Form over Sequential	47
5.1	Improvement in COE Reduction - Integrated Design over Sequential	52
A.1	Precomputational XFOIL Optimization Results	62
A.2	Precomputational CFD Optimization Results	63
A.3	Precomputational Wind Tunnel Optimization Results	64
A.4	Free-Form XFOIL Optimization Results	65
A.5	Free-Form CFD Optimization Results	66

LIST OF FIGURES

1.1	Wind turbines at the Smøla Wind Farm in Norway. (Bjørn Luell, “Wind turbine”, March 24, 2009 via Flickr, Creative Commons Attribution)	2
1.2	The main components that compose a wind turbine. (Hanuman Wind, “Wind turbine components”, August 11, 2009 via Wikimedia Commons, Creative Commons Attribution)	3
1.3	Visualization of several blade shape parameters that are found at 2D cross-sections along the blade span.	4
1.4	Lift and drag coefficients are dependent on the airfoil shape, the angle of attack, and the flow properties characterized by the Reynolds number.	5
1.5	Distribution of wind speed (red) and energy (blue) at the Lee Ranch facility in Colorado. The histogram shows measured data, while the curve is the Rayleigh model distribution for the same average wind speed. A combination of these two datasets calculate annual energy production for use in COE. (Gregors, “Lee Ranch Wind Speed Frequency”, September 5, 2011 via Wikimedia Commons, Creative Commons Attribution)	6
1.6	In sequential design the airfoil shapes are pre-selected and held constant and in integrated design the airfoil shapes are allowed to change during the optimization.	7
1.7	Process to precompute the airfoil drag coefficients for use during the blade analysis and optimization. a) Drag coefficients are calculated for a range of angles of attack for a set airfoil. b) Drag coefficient spline is generated from the precomputed data. c) Drag coefficient spline is used to quickly obtain the drag coefficients for the needed angles of attack (red circles) during the blade analysis.	8
1.8	Workflow comparison of integrated blade design approaches. The precomputational method performs the airfoil analyses and surrogate model generation prior to the blade optimization, while the free-form method performs the airfoil analyses during the blade optimization.	10
1.9	Number of function evaluations required to converge optimization as a function of number of design variables. Sparse Nonlinear OPTimizer is used for the gradient-based results and Augmented Lagrangian Particle Swarm Optimizer for the gradient-free results, however similar trends were observed using Sequential Least Squares Programming (gradient-based) and Non Sorting Genetic Algorithm II (gradient-free). Reference lines for linear and quadratic scaling are also shown. (Image and caption reprinted from “Integrated Design of Downwind Land-based Wind Turbines using Analytic Gradients”, by A. Ning, 2016, Wind Energy, 11, 44. Copyright 2016 John Wiley & Sons, Ltd.)	12
2.1	DU21_A17 airfoil unstructured C-mesh created with <i>SU2</i> mesh deformation tool and CFD field visualization.	17
2.2	Comparison of different airfoil analysis techniques (XFOIL, RANS CFD, and Wind Tunnel) for the DU21_A17 airfoil.	19

2.3	Parameters specifying inflow conditions of a rotating blade section used for the BEM method. (Image reprinted from “A simple solution method for the blade element momentum equations with guaranteed convergence”, by A. Ning, 2014, Wind Energy, 17, 9. Copyright 2014 John Wiley & Sons, Ltd.)	20
2.4	Structural profile of blade cross-section. (Image reprinted from “Users Guide to PreComp”, by G. Bir, 2005, Retrieved from https://nwtc.nrel.gov/PreComp . Copyright 2005 National Renewable Energy Laboratory.)	21
2.5	Visualization of several optimization design variables.	24
3.1	Thickness-to-chord ratio = thickness / chord and is used to parameterize the airfoils for the precomputational method.	28
3.2	Comparison of 50% blended airfoil at thickness-to-chord ratio of 21%. Used to parameterize the airfoils for the precomputational method.	29
3.3	Comparison of the lift and drag coefficient surrogate models for the TU-Delft airfoil family based on angle of attack and thickness-to-chord ratio.	30
3.4	Wind tunnel spline using corrected XFOIL data at $\alpha = 5.0^\circ$	31
3.5	Summary of main results from the precomputational optimizations. Increased airfoil flexibility lead to better COE, mainly through an increase in AEP.	32
3.6	Precomputational chord, twist, spar cap thickness and trailing edge thickness distributions.	36
3.7	XFOIL precomputational airfoil shape results.	37
3.8	CFD precomputational airfoil shape results.	37
3.9	Wind tunnel precomputational airfoil shape results.	38
4.1	The DU21_A17 airfoil is defined by the summation of weighted polynomial components (four on top and four on bottom).	41
4.2	Surface drag coefficient sensitivities ($\partial c_d / \partial x_1, \dots, \partial c_d / \partial x_m$) for each airfoil coordinate normal to the surface of the DU21_A17 airfoil and used to find $\partial c_d / \partial S_{af}$	44
4.3	Summary of main results from the free-form optimizations.	46
4.4	Free-form chord, twist, and spar cap thickness and trailing edge thickness distributions. The RANS CFD tended to have heavier properties than XFOIL such as larger chord, more twist, and thicker materials.	50
4.5	Free-form airfoil shape results. The precomputational method was able to achieve many of the major changes in the airfoil shapes using fewer parameters than the free-form method.	51

NOMENCLATURE

S_{af}	airfoil shape parameters
t/c	airfoil thickness-to-chord ratio
α	angle of attack
AEP	annual energy production
BEM	Blade Element Momentum
B_{af}	blended airfoil family factor
c	chord
c_{max}	chord maximum location
CST	class-shape-transformation
CFD	computational fluid dynamics
COE	cost of energy
IEC	International Electrotechnical Commission
A	Kulfan parameters (CST)
$NREL$	National Renewable Energy Laboratory
$RANS$	Reynolds-averaged Navier-Stokes
t_{spar}	spar cap thickness
λ_2	tip-speed ratio in Region 2
t_{te}	trailing edge thickness
TCC	turbine capital cost
θ	twist
c_d	two-dimensional drag coefficient
c_l	two-dimensional lift coefficient

CHAPTER 1. INTRODUCTION

As technology becomes more integrated into our daily lives, so does our reliance on the energy needed to power our modern lifestyle. While this energy can come from a variety of sources, wind energy is clean, renewable, and can be cost competitive with other energy sources. Wind turbines are the devices used to convert the power inherent in the wind to the electrical power used in our homes. The performance of a wind turbine is predominately determined from the aerodynamics and structures of its blades. Wind turbine blades greatly influence a wind turbine's ability to efficiently convert wind energy to electrical energy. Optimizing the blade shape has the potential to increase the amount of energy output while also reducing the blade costs. The blade is defined by a series of two-dimensional airfoils that constitute its three-dimensional shape. A typical blade design is performed sequentially, where the airfoils are chosen or optimized by the designer a priori and then the blade planform is optimized. Integrated blade design can achieve better performance than sequential design by optimizing the airfoil shapes and blade planform concurrently. This research investigates two integrated blade design approaches, the precomputational and free-form methods.

This chapter discusses the fundamentals of wind energy and wind turbines. This includes a description of the origin of wind energy, wind turbine components, and the wind turbine blade design and optimization process. This chapter provides a review of previous research into integrated blade design and how this research contributes to the field. An outline of the rest of the thesis is provided.

1.1 Wind Energy

Wind is a form of solar energy. The combination of the uneven heating of the earth by the sun, the earth's rotation, and the variability in the earth's surface produces wind due to areas of high and low air pressure. The energy from wind has been extracted for human use for hundreds

of years in the form of wind mills and has recently in the past few decades transitioned to electrical power production. Worldwide capacity for wind energy has increased dramatically during the last few decades. Wind energy has the potential to supply more than 40 times current worldwide consumption of electricity [2]. More efficient wind turbines are needed to take advantage of wind energy's vast potential.

1.2 Wind Turbines



Figure 1.1: Wind turbines at the Smøla Wind Farm in Norway. (Bjørn Luell, “Wind turbine”, March 24, 2009 via Flickr, Creative Commons Attribution)

Wind turbines, as pictured in Figure 1.1, convert wind energy into electrical power. A wind farm is a collection of many wind turbines located together. The wind's kinetic energy causes the wind turbine blades to rotate and generate rotational shaft energy that is used to generate electrical power with a generator. Figure 1.2 shows several of the main components of a wind turbine. The

tower holds up the rotor and the nacelle at heights where the wind is faster and steadier. The rotor is made of the blades and the hub, with the hub connecting the blades together (two or three blades is typical). The rotor rotates due to the force of the wind on the blades and spins the shaft connected to the drive train. The nacelle houses the drive train and electrical components that generate the electrical power from the spinning shaft. The yaw system points the rotor in the direction of the incoming wind. There are two types of wind turbine blades: horizontal-axis wind turbines (HAWT) and vertical-axis wind turbines (VAWT). HAWTs have the blade's axis of rotation parallel to the ground, while VAWTs have the blade's axis of rotation perpendicular to the ground. This study only considers HAWTs.




Figure 1.2: The main components that compose a wind turbine. (Hanuman Wind, “Wind turbine components”, August 11, 2009 via Wikimedia Commons, Creative Commons Attribution)

1.3 Wind Turbine Blades

The most important component to the power generation of a wind turbine is the blade. There are several parameters that compose a wind turbine blade. These parameters are defined at various 2D cross-sections along the blade span and are then interpolated in between the defined 2D sections. The main parameters that define the shape of the blade used in this research are the airfoil shapes, chord, twist, and internal structures as seen in Figure 1.3.




Figure 1.3: Visualization of several blade shape parameters that are found at 2D cross-sections along the blade span.

The chord is the distance from the leading edge of the airfoil to the trailing edge. The twist is the airfoil's rotation angle as compared to the chord line at the root of the blade. The blade twist is used to generate favorable angles of attack for the airfoils along the blade span. The internal structure of the blade is made of different composite layers and reinforced with an internal spar box.

The outer blade shape at each section is defined by an airfoil shape. Airfoils have certain curved characteristics that generate a favorable ratio of lift and drag forces when passed over by a fluid, in this case air. Lift is the force perpendicular to the fluid flow and drag is the force parallel to the fluid flow. The aerodynamics and the structures of the blade depend highly on the lift and drag forces from the 2D airfoils along the blade span. These forces are characterized by lift and drag coefficients that are functions of the airfoil shape, angle of attack, and Reynolds number as seen in Figure 1.4.




Figure 1.4: Lift and drag coefficients are dependent on the airfoil shape, the angle of attack, and the flow properties characterized by the Reynolds number.

1.4 Cost of Energy

A metric to compare the costs of various energy sources (i.e. wind, solar, coal, etc.) is what is known as the cost of energy (COE). The COE is the total cost to build and maintain the energy production system divided by the energy output as defined in Eq. 1.1 and is often measured in either dollars per kilowatt-hour or cents per kilowatt-hour. This is an important metric as the investments needed to build new energy capacity tend to follow the lowest source of COE. As the cost of wind energy decreases, there is a greater incentive to invest in and build larger scale wind turbines.

$$COE = \frac{\text{total cost (\$)}}{\text{annual energy production (kWh)}} \quad (1.1)$$

One of the main goals of wind energy research is to reduce the cost of wind energy to make it equal to or less than other energy sources. Reducing the COE occurs by decreasing the cost of the turbine and/or increasing the energy output. COE is a useful metric for aerostructural optimization because it takes into account both the aerodynamics through the energy production and the structures through the cost. Optimizing only for the energy production can produce blades that are too heavy, too costly, or fail during operation. Energy generation can be dependent on the turbine type, site conditions, and wind speed. The power generated at each wind speed is combined with the frequency of each wind speed at the chosen site to calculate the annual energy production. A comparison of the power and wind frequency at different velocities for the Lee Ranch facility in

Colorado is shown in Figure 1.5. More details on the specific implementation of COE used in this research are found in Chapter 2.




Figure 1.5: Distribution of wind speed (red) and energy (blue) at the Lee Ranch facility in Colorado. The histogram shows measured data, while the curve is the Rayleigh model distribution for the same average wind speed. A combination of these two datasets calculate annual energy production for use in COE. (Gregors, “Lee Ranch Wind Speed Frequency”, September 5, 2011 via Wikimedia Commons, Creative Commons Attribution)

1.5 Blade Design

Given an initial blade design, the blade can be optimized to minimize COE and generate better performance. The optimization provides the optimal blade design that can be further tested and analyzed before moving to manufacturing. This optimization design process is iterative, where the blade parameters such as chord, twist, etc. are changed and the blade reanalyzed until the COE is minimized. In this study, we investigate two forms of blade design, sequential and integrated blade design. The main difference between sequential and integrated blade design is how the airfoil shapes are optimized, either before or during the blade planform optimization.

1.5.1 Sequential Blade Design

A typical blade design is performed sequentially, where the airfoils are chosen or optimized by the designer a priori and then the blade planform is optimized [3–6]. However, optimal airfoil

shapes can be dependent on site conditions, turbine type, and the other blade parameters. Therefore, sequential blade design inherently generates sub-optimal blades because the design space is limited to pre-selected airfoil shapes.

1.5.2 Integrated Blade Design

Integrated blade design optimizes the airfoils and the blade planform concurrently by adding airfoil shape parameters to the optimization as design variables. Regardless of the field or discipline, integrated design performs better than, or at least equal to, sequential design by taking advantage of the interactions between design variables. Figure 1.6 shows a comparison of sequential and integrated wind turbine blade designs.

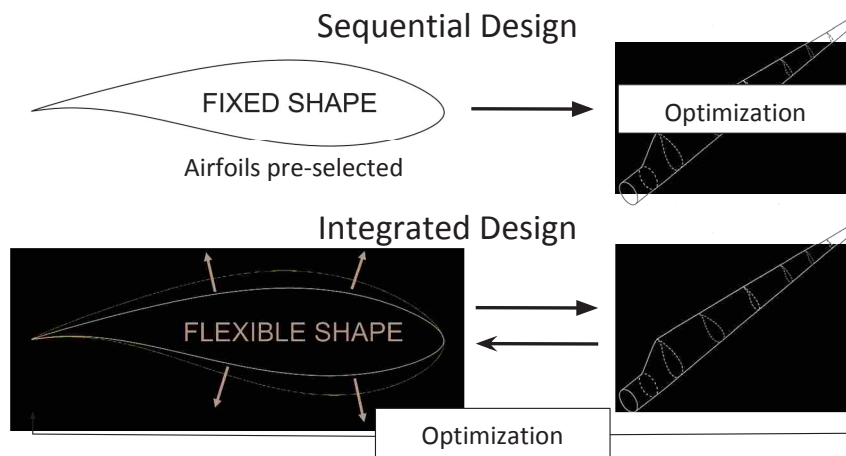


Figure 1.6: In sequential design the airfoil shapes are pre-selected and held constant and in integrated design the airfoil shapes are allowed to change during the optimization.

1.5.3 Precomputation

Although integrated design has been shown to find better performing blade designs [7–15], one of the main reasons that sequential design is commonly used is for the ability to separate the analysis of the airfoil shapes from the blade optimization, referred to in this paper as precomputation. The aerodynamics and the structures of the blade depend highly on the lift and drag forces

along the blade span. These forces are characterized by the lift and drag coefficients that are functions of the airfoil shape, angle of attack, and Reynolds number. During a blade optimization, when design variables such as chord or twist update, the lift and drag coefficients would need to be recomputed at the updated angles of attack. Depending on how the airfoils are analyzed, recomputing at each iteration can be computationally expensive. However, in a typical sequential design the airfoil shapes are pre-selected so that the lift and drag coefficients can be analyzed for a number of angles of attack before the optimization and then read from pretabulated airfoil tables during the optimization. This precomputation also makes it easier to use higher-fidelity airfoil analysis techniques because the analysis only needs to be run once at each angle of attack.

The process to precompute the airfoil analyses is demonstrated is shown in Figure 1.7 for a fixed airfoil shape. Figure 1.7a shows the drag coefficients calculated for the fixed airfoil at a number of angles of attack. Figure 1.7b shows the drag coefficient spline generated from the precomputed data. Both the lift and drag coefficients tend to follow predictable trends due to changes in the angles of attack and therefore a spline tends to capture the data well. Figure 1.7c shows the drag coefficient spline being used to quickly obtain the drag coefficients for the needed angles of attack (red circles) during the blade analysis.




Figure 1.7: Process to precompute the airfoil drag coefficients for use during the blade analysis and optimization. a) Drag coefficients are calculated for a range of angles of attack for a set airfoil. b) Drag coefficient spline is generated from the precomputed data. c) Drag coefficient spline is used to quickly obtain the drag coefficients for the needed angles of attack (red circles) during the blade analysis.

1.6 Contribution

An integrated blade design achieves better results than a sequential blade design because the airfoil shapes and the rest of the blade are optimized simultaneously. However, not precomputing the airfoil analyses makes the blade analysis difficult and computationally expensive. In this study, we develop a method that combines some of the advantage of better performance from integrated design with the advantage of precomputation from sequential design. This integrated design approach is referred to as the precomputational method: an integrated blade design that gives the airfoil shapes some limited flexibility *and* precomputes the airfoil analyses. To maintain the ability to precompute, the airfoil shapes are only allowed to vary such that the change in lift and drag coefficients is continuous and relatively smooth so that the data can be interpolated using a spline. The airfoil analyses are performed for a number of angles of attack and airfoils based on the chosen airfoil shape parameter. A surrogate model is generated from the airfoil analyses that can be used to quickly obtain the lift and drag coefficients for a number of airfoil shapes during the optimization.

Nevertheless, a disadvantage of the precomputational method is that it has limited degrees of freedom, typically only being able to vary the airfoil shape within an airfoil family. To evaluate the performance of the precomputational method, we must determine how it compares to a design approach that allows for fully flexible airfoils, called the free-form method. The objective of this comparison is to determine if most of the benefit of an integrated blade design is captured by the precomputational method or if there is a significant benefit from the additional airfoil flexibility. If the precomputational method can obtain most of the benefit compared to the free-form method, the precomputational method becomes a more attractive option. The reason that the free-form method is not as ideal, despite giving the airfoil shapes more degrees of freedom and being able to obtain better blade performance, is that precomputation is no longer possible. The airfoil analyses must be performed during the blade optimization, which can be computationally expensive. Precomputation is not a feasible option for the free-form method because a surrogate model that would work with all possible airfoil shapes would require an excessive number of airfoil analyses and negate the benefit of precomputation. Depending on the number of airfoil analyses used to generate the precomputational surrogate model, running the airfoil analyses within the optimization could require 100-200X more airfoil analyses overall than those needed to generate the precomputational

surrogate model. Although performing the airfoil analyses during the optimization is costly, it is necessary to realize and explore the full potential of an integrated blade design for comparison to the precomputational method. The workflows of these two integrated design approaches are compared in Figure 1.8.

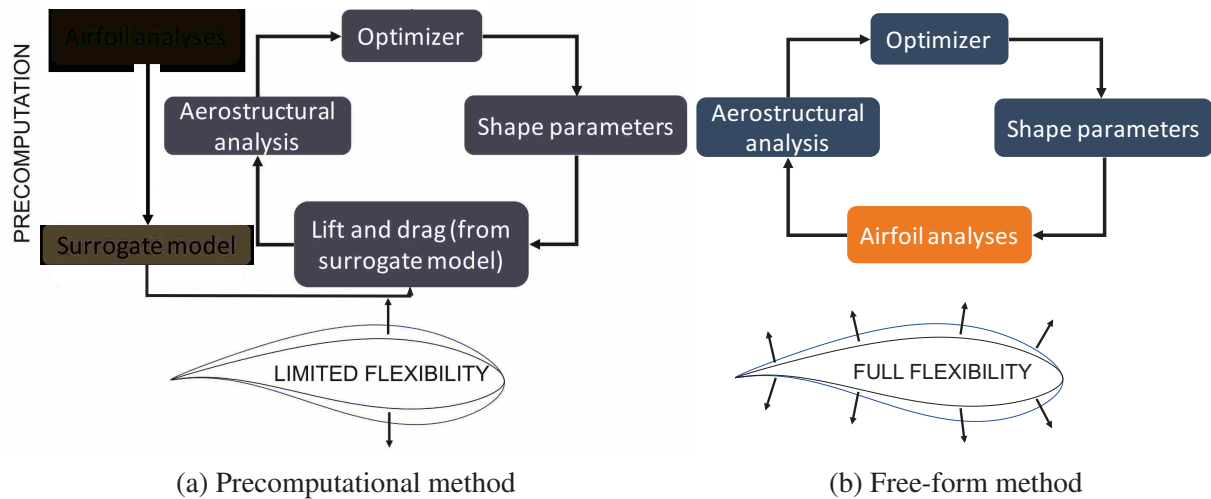


Figure 1.8: Workflow comparison of integrated blade design approaches. The precomputational method performs the airfoil analyses and surrogate model generation prior to the blade optimization, while the free-form method performs the airfoil analyses during the blade optimization.

The three main areas to consider with an integrated blade design are the airfoil analysis method (how to obtain the airfoil’s lift and drag coefficients), the airfoil parameterization method (how the airfoil coordinates translate into optimization design variables), and the optimization method.

This research extends previous investigations of integrated blade design by using Reynolds-averaged Navier-Stokes computational fluid dynamics (RANS CFD) for the airfoil analysis method and comparing it to a panel method, XFOIL. All known prior integrated designs that use 2D analysis methods have used XFOIL [7–13]; although RANS CFD have been used in 3D CFD integrated design [14, 15]. While panel methods are relatively fast, they have known problems converging in highly separated flow and are typically not as accurate as other higher-fidelity techniques such as RANS CFD or using a wind tunnel.

The CST method used in this research for the airfoil parameterization has been highly rated compared to other parameterization methods due to its ability to span the airfoil design space with a relatively low number of parameters and the assurance of a smooth curve [17, 18]. The method was developed by Kulfan for the purpose of reducing the number of parameters needed to define any 2D or 3D smooth aerodynamic shape [19, 20] and has successfully been used in aircraft wing design and optimization [21, 22].

The choice of optimization method is critical to the feasibility of the free-form method because of the increase of both the number of design variables and the computational cost per function evaluation. Previous research has focused on gradient-free methods [9–12, 14] and finite-differencing gradients [7, 8] for integrated design. Gradient-free methods in particular are well-suited for finding the global optimum. Vicina et al. demonstrated the potential of gradient-free methods for use with integrated design in wind turbine optimization, even when using CFD [14]. However, because both gradient-free and finite-differencing methods tend to scale poorly with the number of design variables the optimization problems have typically been limited in size. Analytic gradients are well-suited for larger scale problems because they are exact and scale well with the number of design variables. Methods designed for large-scale optimization problems are needed as it is likely that their size will increase as wind turbine optimization continues to develop.

Figure 1.9 shows the relationship between the number of design variables and the number of function evaluations needed to converge a relatively simple optimization for the different optimization methods of gradient-free, gradient-based finite-differencing, and gradient-based analytic gradient methods for the multidimensional generalization of the Rosenbrock function (Image and caption from Ning et al. [1]). Although the Rosenbrock function is less complex than wind turbine optimization, Lyu et al. shows that similar trends exist specifically for aerodynamic shape optimization using RANS CFD [23]. They conclude that gradient-based methods are the only viable option for large-scale aerodynamic design optimization because gradient-free methods can require up to three orders of magnitude more computational effort than gradient-based methods [23]. Rios et al. corroborates this idea when they tested 502 optimization problems with 22 gradient-free solvers and showed that, generally speaking, the performance of the optimizer suffers when there are over 30 design variables [24]. Using RANS CFD with the setup used in this research, the time difference to optimize between using gradient-free, finite-differencing, and analytic gradient

methods can be on the order of months, weeks, and days, respectively. There are some downsides to gradient-based methods such as the possibility of converging on a local optimum and the development of the gradients. Nevertheless, these risks can be mitigated through the use of analytic gradients and using a multi-start approach. In many cases local optimum exist due to numerical noise caused by poor gradients. This research uses gradient-based methods to reduce the computational cost and increase the performance of the optimization.

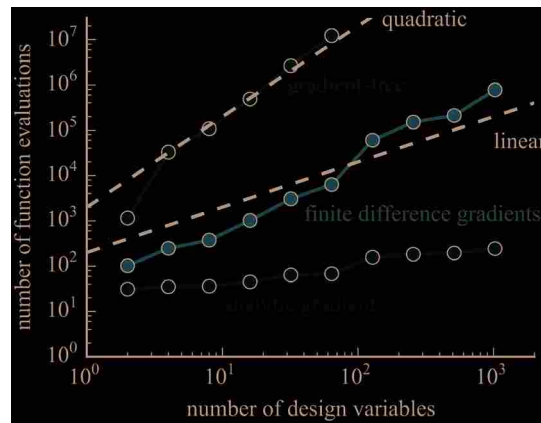


Figure 1.9: Number of function evaluations required to converge optimization as a function of number of design variables. Sparse Nonlinear OPTimizer is used for the gradient-based results and Augmented Lagrangian Particle Swarm Optimizer for the gradient-free results, however similar trends were observed using Sequential Least Squares Programming (gradient-based) and Non Sorting Genetic Algorithm II (gradient-free). Reference lines for linear and quadratic scaling are also shown. (Image and caption reprinted from “Integrated Design of Downwind Land-based Wind Turbines using Analytic Gradients”, by A. Ning, 2016, Wind Energy, 11, 44. Copyright 2016 John Wiley & Sons, Ltd.)

To summarize, this research contributes to prior investigations in five main ways and has the following objectives:

1. Determine the effectiveness of integrated blade design approaches (precomputational and free-form methods) in reducing cost of energy as compared to the sequential design. By performing a direct comparison under the same conditions between the different designs, the impact of blade design on a blade optimization can be determined. Comparing the integrated precomputational and free-form methods allows us to quantify how much of the benefit of integrated blade design can be captured while still maintaining the ease of precomputation.

2. Determine the effect of airfoil shape flexibility on blade performance by comparing three airfoil parameterization methods: thickness-to-chord ratio, blended airfoil family factor, and the CST method. While more degrees of freedom in the airfoil shapes provide additional benefits, it is important to quantify the differences between airfoil parameterization methods to determine if most of the benefit from airfoil shape flexibility can be provided by one or two airfoil shape parameters or if additional airfoil shape flexibility is needed.
3. Develop techniques for using integrated blade design with higher-fidelity airfoil analysis methods such as RANS CFD or using wind tunnel data and compare these methods to XFOIL. The optimal blade shape is dependent on the airfoil analysis method; in order to generate a blade design that more closely mimics the behavior of a real wind turbine, more accurate airfoil analysis methods are needed.
4. Develop airfoil shape parameter analytic gradients instead of using finite-differencing gradients or using a gradient-free method. By developing better gradients, the optimization can converge in a shorter amount of time. Due to the computationally expensive nature of the airfoil analysis methods, reducing the number of function evaluations as much as possible provides large reductions in the time needed to optimize the blade.
5. Provide a general tool for performing integrated blade design by adding code to NREL WISDEM open-source software. The current NREL WISDEM open-source software has the capability to do sequential blade design. The objective is to add the option to perform the integrated blade design approaches used in this research, including the precomputational and the free-form methods. Doing so will allow others to more easily implement integrated blade design for other turbine types and site conditions.

To achieve the stated objectives and contributions, eleven optimization cases are performed, as shown in Table 1.1, using various analysis and parameterization techniques. This study contributes to the development of wind turbines by comparing free-form and precomputational integrated design approaches and providing higher-fidelity airfoil analysis techniques for the aerostructural optimization of wind turbine blades. These methods can be used for more detailed blade design, make wind turbines more cost-effective, and contribute to the world's growing energy needs by continuing to reduce COE.

Table 1.1: Optimization Cases Summary

	XFOIL	RANS CFD	Wind Tunnel
sequential	✓	✓	✓
integrated precomputational (thickness-to-chord ratio)	✓	✓	✓
integrated precomputational (blended airfoil family factor)	✓	✓	✓
integrated free-form (CST)	✓	✓	N/A

1.7 Outline

This first chapter introduced the fundamentals of wind energy and the optimization of wind turbine blades. It summarized the current body of research and how this work contributes to the field. The second chapter discusses in more depth how the blade is analyzed and optimized specifically for this research. The third chapter focuses on the methodology and results of the first integrated design approach, the precomputational method. The fourth chapter focuses on the methodology and results of the second integrated design approach, the free-form method. The fifth chapter summarizes the main conclusions from the results and the comparison between the two integrated design approaches.

CHAPTER 2. AEROSTRUCTURAL BLADE ANALYSIS AND OPTIMIZATION

To generate a practical wind turbine blade design, both the aerodynamics and the structures of the blade must be simultaneously considered. This is referred to as an aerostructural approach. Separating the blade aerodynamic and structural analysis can result in sub-optimal designs that are either too heavy, fail during operation, or have poor energy output [25]. This aerostructural approach is performed by including aerodynamic and structural design variables, constraints, and objectives.

2.1 Aerodynamics

Generating the power of a wind turbine comes from the Blade Element Momentum (BEM) theorem. An essential part of this theorem is to iteratively obtain the lift and drag coefficients of each airfoil section along the blade. The power generation is needed at a number of wind speeds to obtain the annual energy production.

2.1.1 Airfoil Analysis

The airfoils' lift and drag coefficients have a major effect on the aerostructural analysis. Three common ways to obtain these lift and drag coefficients are through a lower-fidelity panel method (XFOIL), Reynolds-averaged Navier-Stokes computational fluid dynamics (RANS CFD), or using a wind tunnel. The optimal blade shape is sensitive to the airfoil analysis method and therefore using higher-fidelity techniques such as RANS CFD or a wind tunnel is warranted. More details on the implementation of XFOIL, RANS CFD, and the wind tunnel data used in this research are described in the following sections.

Panel Method (XFOIL)

XFOIL is a software program developed by Drela that uses a linear potential (panel) method in the design and analysis of subsonic isolated airfoils [26]. Panel methods can be used to numerically solve for the lift and drag coefficients of a 2D or 3D shape, in this case for a specific airfoil shape. They are based on simplifying assumptions of the flow physics and air flow properties where compressibility and viscosity are considered negligible. XFOIL uses a simple boundary layer theory of surface flow displacement on top of the panel method to account for more viscous effects.

Given the 2D coordinates from the airfoil parameterization, XFOIL calculates the pressure distribution to obtain the lift and drag coefficients. XFOIL performs the analysis very quickly, but integral boundary layer approaches are not consistent with the physics of highly separated flows and post-stall performance can be inaccurate. The advantage of using XFOIL is that it performs the analysis very quickly. The disadvantages of using XFOIL, however, are that the computational data gathered can be inaccurate and can only be used in a limited range as XFOIL does not converge well in highly separated flows. This inaccuracy tends to increase when analyzing thicker airfoils or at higher angles of attack as these conditions tend to have more highly separated flows. Since XFOIL is based on idealized computational models, it tends to under-predict drag coefficients and slightly over-predict lift coefficients. The angle of attack and the airfoil shape parameters are tightly bounded so as to reduce the likelihood of non-convergence.

RANS CFD

RANS CFD more accurately estimates the airfoil's lift and drag coefficients by better resolving viscous effects, especially post-stall. A comparison of CFD compared to other methods is shown in Figure 2.2. The open-source CFD software *SU2* [27] is used because it provides adjoint gradients and has a Python interface. A major challenge in performing CFD analysis within an optimization is that the mesh needs to be regenerated whenever the airfoil shape changes. The *SU2* mesh deformation tool [27] can automatically deform any base mesh to any airfoil shape, a deformed mesh is visualized in Figure 2.1a and 2.1b. To calculate the cost of energy for the blade design, fourteen airfoil sections (six airfoil shapes are repeated at different sections) are analyzed

along the blade at ten wind speeds to generate the power curve and each airfoil section is also analyzed an additional four times to perform the structural analysis. The 2D airfoil meshes use 14,336 elements in an unstructured C-mesh to balance the trade-off between accuracy and time.

The RANS governing equations use the one-equation Spalart-Allmaras turbulence model with compressible flow. Although other turbulence models such as $k - \omega$ may be more accurate for wind turbine applications [29], this method was chosen because it was the best available in the *SU2* CFD package. The convergence criteria specifies a change in residual values of less than 10^{-9} with a max number of iterations of 50,000. Using 16 cores, the typical required computational time is several minutes and around 20,000 iterations to convergence. Figure 2.1c shows the CFD field visualization around one of the airfoils at an angle of attack of zero degrees. For this research,

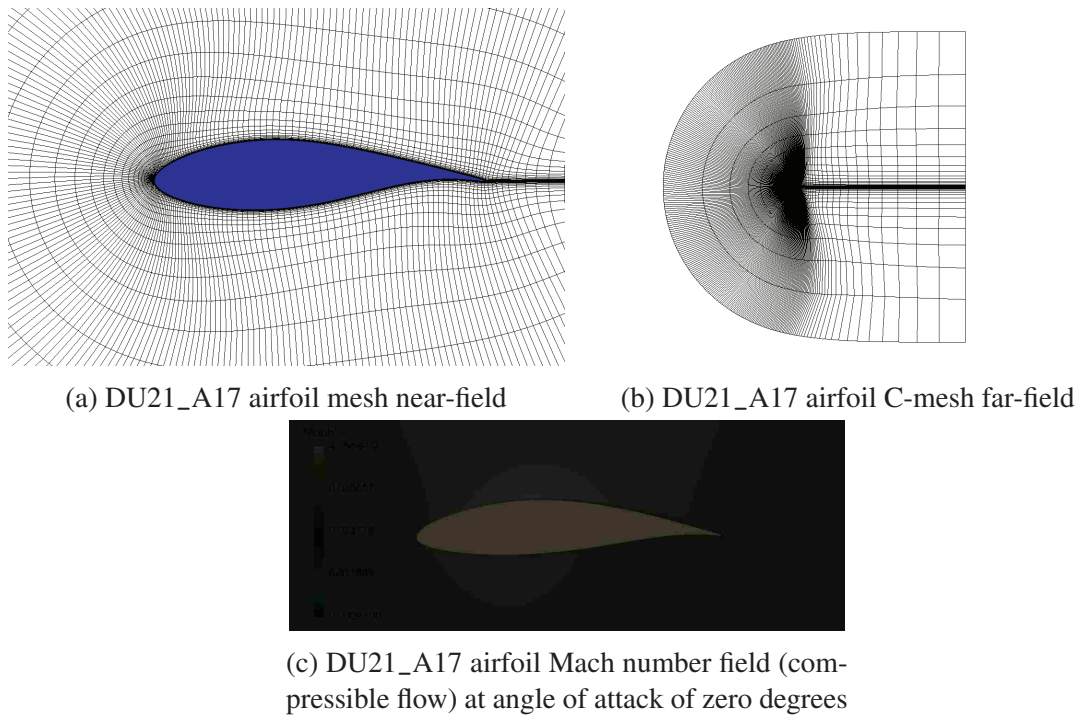


Figure 2.1: DU21_A17 airfoil unstructured C-mesh created with *SU2* mesh deformation tool and CFD field visualization.

a supercomputer was used to speed up the CFD simulations. However, for a single 2D airfoil CFD simulation, above about 16 to 32 cores the time savings from additional cores was less than the increase in time from communication overhead between cores. To take advantage of more computing resources, the analysis was setup to run multiple CFD simulations simultaneously. For

example, each blade analysis needed to run 14 airfoils along the blade span multiple times. Those 14 airfoils could be run simultaneously on the supercomputer where each airfoil is given 16 cores for a total of 224 cores being used. This led to about a 14x speedup of the CFD optimization compared to running each simulation sequentially.

Wind Tunnel

A wind tunnel is a device that can physically test different shapes in air flows of different speeds. An airfoil can be built and positioned in the wind tunnel and the lift and drag forces can be physically measured. Therefore a wind tunnel is used to experimentally obtain the lift and drag coefficients, while XFOIL and CFD are computer simulations. Although wind tunnel data is the highest-fidelity data that is most commonly used for analysis if used with accurate instrumentation, obtaining the data is difficult, expensive, and time-consuming. As each airfoil would need to be built and analyzed separately in a wind tunnel, it would be difficult (if not impossible) to use in a free-form approach. Nevertheless, the precomputational method can be used with publicly available wind tunnel data for the airfoils used in the NREL 5-MW reference turbine [30] because many of the airfoils used are from the same airfoil family.

A comparison of the aerodynamic performance of an airfoil using XFOIL, RANS CFD, and known wind tunnel data is shown in Figure 2.2. As seen in the figure, the RANS CFD data more closely matches the trend of the wind tunnel data as compared to XFOIL, especially post-stall [31]. XFOIL tended to over-predict lift and under-predict drag as compared to the wind tunnel data, while the CFD tended to slightly under-predict lift and over-predict drag for the specific mesh and configuration used. Therefore CFD, in this case, tended to be a more conservative estimate of the airfoil aerodynamic performance.

The lift and drag coefficients from the aerodynamic analyses in this study are analyzed at a Reynolds number of 10^6 and are rotational corrected and extrapolated using the NREL Airfoil-Preppy Python tool¹ to prepare for blade analysis. The three-dimensional rotational corrections are performed using Du's method [32] to augment the lift and Eggers' method [33] to modify the drag.

¹<https://github.com/WISDEM/AirfoilPreppy>

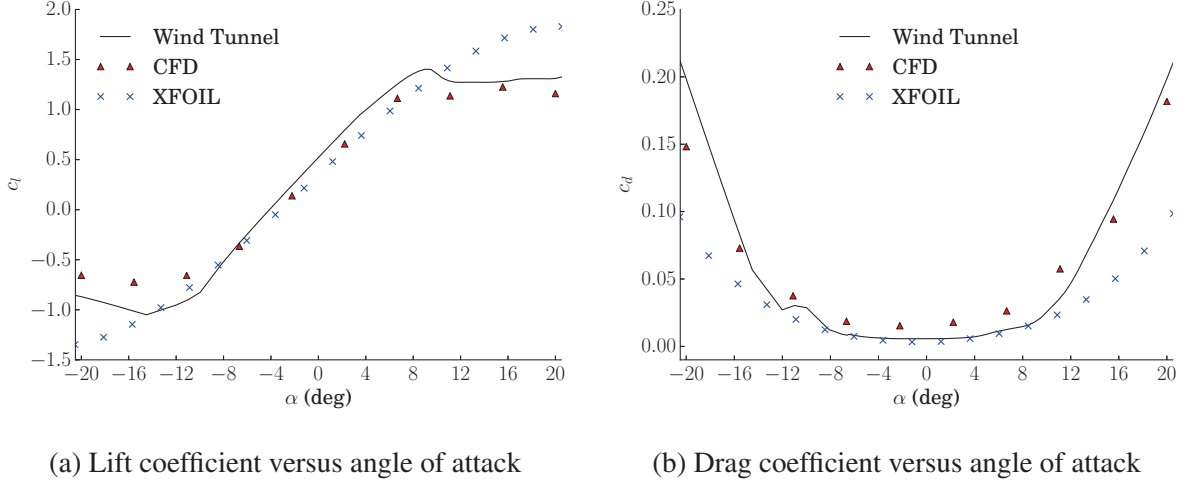


Figure 2.2: Comparison of different airfoil analysis techniques (XFOIL, RANS CFD, and Wind Tunnel) for the DU21_A17 airfoil.

2.1.2 Blade Element Momentum Theorem

Blade element momentum (BEM) theory converts the lift and drag coefficients from all the 2D airfoils into the 3D blade analysis. The chosen BEM method has guaranteed convergence (CCBlade²) [34] and is used with the wind blade analysis tool RotorSE³. The BEM theory defines a residual equation (Eq. 2.1) that must be solved iteratively to converge on the local inflow angle (ϕ) [34]. The induction factors (a and a') are then computed and the lift and drag coefficients of each airfoil section determined at the specified angle of attack.

$$R(\phi) = \frac{\sin \phi}{1 - a(\phi)} - \frac{\cos \phi}{\lambda_r(1 + a'(\phi))} = 0 \quad (2.1)$$

Due to the expensive nature of RANS CFD and the iterative behavior of the BEM method [34], the induction factors are converged using lift and drag coefficient splines from XFOIL. The splines are generated from analyzing the airfoil shapes at angles of attack from -30° to 30° , then the Viterna method is used to extrapolate the lift and drag coefficients from -180° to 180° [35]. Once the induction factors and corresponding angles of attack are known, the specified airfoil analysis method (either XFOIL or RANS CFD) generate the lift and drag coefficients.

²<https://github.com/WISDEM/CCBlade/>

³<https://github.com/WISDEM/RotorSE/>




Figure 2.3: Parameters specifying inflow conditions of a rotating blade section used for the BEM method. (Image reprinted from “A simple solution method for the blade element momentum equations with guaranteed convergence”, by A. Ning, 2014, Wind Energy, 17, 9. Copyright 2014 John Wiley & Sons, Ltd.)

Induction factors measure in a way the performance of the wind turbine, the axial induction factor (a) is defined as the difference between the wind speed far away from the turbine (U_1) and the wind speed at the turbine (U_2) divided by wind speed far away as in Eq. 2.2.

$$a \equiv \frac{U_1 - U_2}{U_1} \quad (2.2)$$

Due to the expensive nature of RANS CFD and the iterative behavior of the BEM method [34], the induction factors are converged using lift and drag coefficient splines from XFOIL. The splines are generated from analyzing the airfoil shapes at angles of attack from -30° to 30° , then the Viterna method is used to extrapolate the lift and drag coefficients from -180° to 180° [35]. Once the induction factors and corresponding angles of attack are known, the specified airfoil analysis method (either XFOIL or RANS CFD) generate the lift and drag coefficients.

2.2 Structures

While the blade’s aerodynamics are important for calculating the power conversion, the blade’s structural analysis is necessary to ensure the blade does not fail during operation. A beam finite element method is used to analyze the blade structure called pBEAM (polynomial beam element analysis module), which uses Euler-Bernoulli beam elements with twelve degrees of freedom (three translational and three rotational at each element end)⁴ [25]. For the composite materials

⁴<https://github.com/WISDEM/pBEAM>

found along the blade sections shown in Figure 2.4, a modified classical lamination theory (CLT) combined with a shear-flow approach is used called PreComp⁵. Similar structural analysis approaches are fairly common [8, 36] and the configuration used is described in more detail by Ning et al. [25]. The airfoil shapes affect the blade structures through the sizing of the sections that, in turn, affect blade mass, strain, etc. Composite panels in the spar cap, web, and trailing edge make up the majority of the structural integrity. There are a number of layers including: the Gel-Coat, glass fabrics, SNL TRIAX ($[\pm 45]_2[0]_2$), SaerTex Double-Dias (DB, $[\pm 45]_4$), carbon fabrics, generic foam, and epoxy resins [25]. These structural composite layers can adapt to changes in the chord and airfoil shapes. Trailing edge material thickness and spar cap material thickness are the optimization variables and the remaining structural elements scale accordingly.




Figure 2.4: Structural profile of blade cross-section. (Image reprinted from “Users Guide to Pre-Comp”, by G. Bir, 2005, Retrieved from <https://nwtc.nrel.gov/PreComp>. Copyright 2005 National Renewable Energy Laboratory.)

2.3 Uncertainty

As in all simulations, uncertainty plays a factor on the accuracy of the results. Other studies have been done uncertainty on blade optimization and analysis in general for both the environmen-

⁵<https://nwtc.nrel.gov/PreComp>

tal conditions and blade geometry [37,38]. In terms of this study, the purpose was not to perform uncertainty quantification, but it would still be beneficial to consider areas that integrated blade design would affect uncertainty. The main area of uncertainty from this study comes from the lift and drag coefficients. This is a difficult problem because even wind tunnel data can be inaccurate compared to the actual performance of a wind turbine blade. Additional work could be done with integrated blade design to perform robust optimization and take into account some of these uncertainties into the optimization process.

2.4 General Optimization Setup

The different optimization cases vary the airfoil parameterization or analysis methods, but the underlying optimization method remains the same.

Objective Function The optimization objective is to minimize the cost of energy, the total cost of the turbine divided by its energy production. The analysis uses a simplified model that assumes that the other aspects of the turbine (hub, nacelle, and tower) remain constant because the rotor thrust is constrained to not exceed its initial thrust and the rated power is held constant. This means that the only effect from the blade is on the AEP and the TCC. In this case, financing aspects are ignored and the cost of energy is found with Eq. 2.3 [39].

$$COE = \frac{FR(TCC + BOS) + (1 - T)OPEX}{AEP} \quad (2.3)$$

In this equation, COE is the project levelized cost of energy, TCC is the total turbine capital costs for the project, BOS is the total balance of station costs for the project, AEP is the annual energy production, OPEX is the overall project operational expenditures, FR is the financing rate, and T is the tax deduction rate on OPEX. The TCC is the sum of the cost of the tower, nacelle, and rotor. In this case, we assume the tower cost and the nacelle cost remain constant. The TCC is the sum of the tower, nacelle, and rotor as in Eq. 2.4.

$$TCC = \text{tower cost} + \text{nacelle cost} + \text{rotor cost} \quad (2.4)$$

Rotor cost is the sum of the hub cost and the blades cost where the blades costs are estimated to be linearly proportional to the blade mass. An AEP loss factor of 0.885 and the turbine capital cost multiplier of 1.56 are used with a FR of 0.095 and T of 0.4 [39]. Standard International Electrotechnical Commission (IEC) specifications for a land-based high-wind-speed site (IEC Class IB) are used corresponding to a mean wind speed of 10.0 m/s [40]. The wind conditions follow a Weibull distribution with a shape parameter of 2.0.

Design Variables The relevant design variables in each case are summarized in Table 2.1. In this study, fourteen airfoil sections are used with three non-airfoil sections. The chord distribution is an array of four control points that define a spline that defines the chord along the blade span. The max chord location defines the point along the blade span where the maximum chord occurs. The twist distribution defines four control points for the entire blade span. The tip speed ratio defines the ratio of speed of the tip of the blade over the speed of the incoming wind. The trailing edge and spar cap thickness distribution defines five control points that define the thicknesses of the composite materials along the blade span. Depending on the blade design there are different airfoil shape parameters. In all cases, there are six airfoils used along the blade that are defined by the airfoil shape parameters either thickness-to-chord ratio, blended airfoil family, or the CST method. Both thickness-to-chord ratio and blended airfoil family factor are one variable each per airfoil while the CST method uses eight variables for each airfoil. The design variables are explained in more detail in [25] and later on in this study. Figure 2.5 shows a cross-section of the blade for a better understanding of various design variables and constraints.

Constraints There are a number of constraints on the optimization to ensure the blade is structurally sound. The categories include constraints on the strain and buckling of the spar cap and trailing edge, the flap-wise and edge-wise frequency, and the rotor thrust. The strain is constrained for extreme load conditions according to IEC standards. The buckling is constrained for maximum operating conditions. All blade natural frequencies are to be above the rotor rotation speed with an added margin to avoid resonance. The rotor thrust is constrained to not exceed its initial thrust to ensure that the same tower and drivetrain can be used. The rated power is kept constant at 5-MW for a similar reason. In total, there are 33 constraints on the optimization that are described in more

Table 2.1: Number of Design Variables Summary

		Sequential	t/c	B_{af}	CST
chord distribution	c	4	4	4	4
max chord location	c_{max}	1	1	1	1
twist distribution	θ	4	4	4	4
tip-speed ratio in Region 2	λ_2	1	1	1	1
trailing edge thickness distribution	t_{te}	5	5	5	5
spar cap thickness distribution	t_{spar}	5	5	5	5
thickness-to-chord ratio distribution	t/c	-	6	6	-
blended airfoil family factor	B_{af}	-	-	6	-
Kulfan parameters (CST)	A	-	-	-	48
total #		20	26	32	68




Figure 2.5: Visualization of several optimization design variables.

detail by Ning et al. [25]. A overview of the constraints on the optimization is summarized in Table 2.2.

The optimization is performed using a gradient-based sequential quadratic programming method called the Sparse Nonlinear OPTimizer (SNOPT) [41] optimization package. This optimization method generally finds the optimum quickly and robustly. Since there are 34 outputs (33 constraints and one objective), the sequential, precomputational t/c , and precomputational B_{af} use the direct method within the OpenMDAO [42] framework because the number of design variables is fewer than the number of outputs at 20, 26, and 32, respectively. The free-form method uses the adjoint method because the number of design variables is 68. All of the design variables are scaled

Table 2.2: Constraints Summary

Description
spar cap strain \leq ultimate strain at 7 stations along blade
trailing edge strain \leq ultimate strain at 8 stations along blade
spar cap buckling \leq critical buckling at 8 stations along blade
trailing edge buckling \leq critical buckling at 7 stations along blade
flap-wise/edge-wise frequency \geq blade passing frequency
rotor thrust at rated power \leq initial rotor thrust

so that the gradients are of a similar magnitude and the initial airfoil shapes and blade parameters are taken from the NREL 5-MW reference turbine [30]. The formulation of the optimization is summarized below:

$$\begin{array}{ll}
 \text{minimize} & COE(x) \\
 \text{with respect to} & x = S_{af} \text{ (either } t/c \text{ and } B_{af}, \text{ or } A), c, c_{max}, \theta, \lambda_2, t_{spar}, t_{te} \\
 \text{subject to} & c_{set}(x) < 0 \text{ (buckling, strain, frequency, rotor thrust)}
 \end{array}$$

All of the optimization results are performed using the same general setup. The only differences in terms of the optimization setup between the various optimization cases are the airfoil shape parameters and the airfoil analysis method.

CHAPTER 3. PRECOMPUTATIONAL METHODS

3.1 Theory and Methodology

The objective of the precomputational method is to allow the airfoil shapes to change during the blade optimization while still computing the airfoils' lift and drag coefficients beforehand. The precomputational method combines the performance advantage of integrated blade design with the ability to precompute the airfoil analysis as in sequential design. For this to be possible, the parameter that changes the airfoil shapes needs to vary such that the design space for the lift and drag coefficients is continuous and relatively smooth so that a surrogate model can define the design space well. Airfoil families tend to exhibit this behavior, so for this research we decided to limit the airfoil shape parameters to change the airfoil shapes within specific airfoils families. A surrogate model of both the lift and drag coefficients is generated by running the airfoil analyses for many airfoils within the chosen airfoil family. Surrogate models closely emulate the behavior of the more complicated analysis method, while being computationally inexpensive to evaluate compared to the real analysis method. The lift and drag coefficients can then be inexpensively obtained from the surrogate model during the optimization. In summary, a typical sequential design generates splines for each airfoil to make the lift and drag coefficients a function of angle of attack. In the precomputational method, 2D splines are generated for each airfoil family to make the lift and drag coefficients a function of both angle of attack and an airfoil shape parameter. The process to perform the precomputational method is as follows:

1. choose airfoil shape parameters that can generate continuous and relatively smooth lift and drag coefficients throughout the design space
2. compute the airfoil analyses for a number of airfoils by varying the chosen airfoil shape parameter and angle of attack

3. create surrogate lift and drag coefficient models from the precomputed airfoil analyses using 2D splines
4. optimize to reduce COE using the surrogate model to obtain the lift and drag coefficients as needed by the optimization with the airfoil shape parameter and the angle of attack as inputs

The creation of the surrogate model enables the workflow to be very similar to that of the sequential design as seen in Figure 1.8. Instead of looking up the lift and drag coefficients from airfoil tables, they are generated by the surrogate model. Through the use of the surrogate model, we are able to achieve the objective of obtaining some of the benefit of an integrated design *and* still precompute.

3.2 Parameterization Methods

The choice of airfoil shape parameter is an important consideration. One of the most common ways to classify a group of similar type airfoils is through the airfoil family. A good way to give the airfoil some flexibility within a specific airfoil family is to change the airfoil thickness-to-chord ratio. Another degree of flexibility would be to perform a blend between airfoil families, called in this paper the blended airfoil family factor. For this case, the fixed airfoil families match those used in the NREL 5-MW reference turbine: the TU Delft and NACA 64-series airfoil families. The TU Delft airfoil family is used for the first two-thirds of the blade and the NACA 64-series for the last third as the starting conditions to match the NREL 5-MW reference turbine blade. The thickness-to-chord ratio and the blended airfoil family factor were chosen specifically because these are common approaches used in determining the airfoils to use in sequential design. A goal for the development of the precomputational method is to make it as easy as possible to implement from existing sequential design methods.

3.2.1 Thickness-to-chord Ratio

A main contributor to an airfoil's aerodynamic performance and structural integrity is its thickness. Aerodynamic performance tends to improve with thinner airfoils while the blade's bending stiffness tends to improve with thicker airfoils. Choosing another airfoil shape parameter could have had an impact, but likely not as large as the thickness-to-chord ratio. This thickness

can be controlled through the thickness-to-chord ratio (t/c) for a fixed airfoil family as shown in Figure 3.1. The thickness-to-chord ratio is simply the maximum thickness of an airfoil normalized the airfoil chord.




Figure 3.1: Thickness-to-chord ratio = thickness / chord and is used to parameterize the airfoils for the precomputational method.

3.2.2 Blended Airfoil Family Factor

The thickness-to-chord ratio allows for the airfoil to vary within a single airfoil family, while a blended airfoil family factor (B_{af}) allows for an additional degree of freedom by blending between two airfoil families at the same t/c . The lift and drag coefficients and the airfoil coordinates are linearly blended based on a factor that varies continuously between 0.0 and 1.0 that indicates the degree to which the second airfoil family (NACA 64-series) is blended into the first (TU-Delft). For example, a factor of 0.3 would mean that that airfoil is 70% TU-Delft and 30% NACA 64-series. The aerodynamic data is linearly blended using the AirfoilPreppy blend tool¹. The underlying equations are a simple linear blend shown in Eq. 3.1 and Eq. 3.2, where 0 refers to the first airfoil family and 1 refers to the second airfoil family.

$$c_l = c_{l_0} + B_{af}(c_{l_1} - c_{l_0}) \quad (3.1)$$

$$c_d = c_{d_0} + B_{af}(c_{d_1} - c_{d_0}) \quad (3.2)$$

The airfoil coordinates for the structures are also blended linearly. Performing this type of airfoil blending is common practice as some analysis tools, such as FAST [43] through the AirfoilPrep

¹<https://github.com/WISDEM/AirfoilPreppy>

tool, use a similar airfoil blending between sections. This research applies the blending instead to each airfoil section rather than between sections. An example of this blending between airfoil families can be seen in Figure 3.2.





Figure 3.2: Comparison of 50% blended airfoil at thickness-to-chord ratio of 21%. Used to parameterize the airfoils for the precomputational method.

3.3 Surrogate Model

For both airfoil families, ten airfoils with t/c ranging from 13% to 42% are analyzed and the lift and drag coefficients are extracted at various angles of attack. Both the lift and drag coefficients are 2D splined across angles of attack and thickness-to-chord ratios as seen in Figure 3.3. A smoothing factor is applied to the bivariate spline that resulted in maximum error values of 0.01 and 0.005 for the lift and drag coefficients splines, respectively. Using the surrogate models, the lift and drag coefficients can be estimated for any angle of attack and thickness-to-chord ratio within either airfoil family. The accuracy was verified by adding additional thickness-to-chord ratios until there were only small changes to the surrogate model. This made it relatively accurate for each chosen thickness-to-chord ratio and angle of attack within the provided limits.

3.4 Airfoil Analysis Correction (Wind Tunnel)

While wind tunnel data is publicly available for most thickness-to-chord ratios of the airfoil families used, when the wind tunnel data is not available, a correction is applied to XFOIL data to mimic the wind tunnel data. The XFOIL data is used for the correction instead of the RANS CFD data because of its speed and because only a few corrections are needed. Wind tunnel data



(a) Lift coefficient surrogate model

(b) Drag coefficient surrogate model

Figure 3.3: Comparison of the lift and drag coefficient surrogate models for the TU-Delft airfoil family based on angle of attack and thickness-to-chord ratio.

is known for five thickness-to-chord ratios ranging from 21% to 40% [30]. XFOIL is calculated at those same five thickness-to-chord ratios and the difference between the wind tunnel and XFOIL lift and drag coefficients are taken and averaged at each angle of attack. For the lower thickness-to-chord ratios where wind tunnel data is not available, XFOIL is used and then the correction applied so that the data more closely matches the wind tunnel data. The XFOIL correction is therefore a drag and lift offset based on the known difference between the wind tunnel and computational data under the same conditions. The Reynolds number is matched between the XFOIL correction data and the wind tunnel data at 10^6 . The surrogate model is created this way for both airfoil families so it affects both the thickness-to-chord and blended airfoil family factor results. Since the wind tunnel data is not available for the corrected data it is difficult to quantify the accuracy of the correction without testing the airfoils in a wind tunnel. This correction is demonstrated in Figure 3.4 that shows a combined lift and drag correction for lift over drag (although the correction is applied to each separately) at an angle of attack of 5.0° for the TU Delft airfoil family.

The wind tunnel data is used for most of the analysis and the corrected XFOIL data is used for the smaller thickness-to-chord ratios where wind tunnel data is not available. Therefore, the lift and drag coefficients are anchored with wind tunnel data and augmented with corrected XFOIL

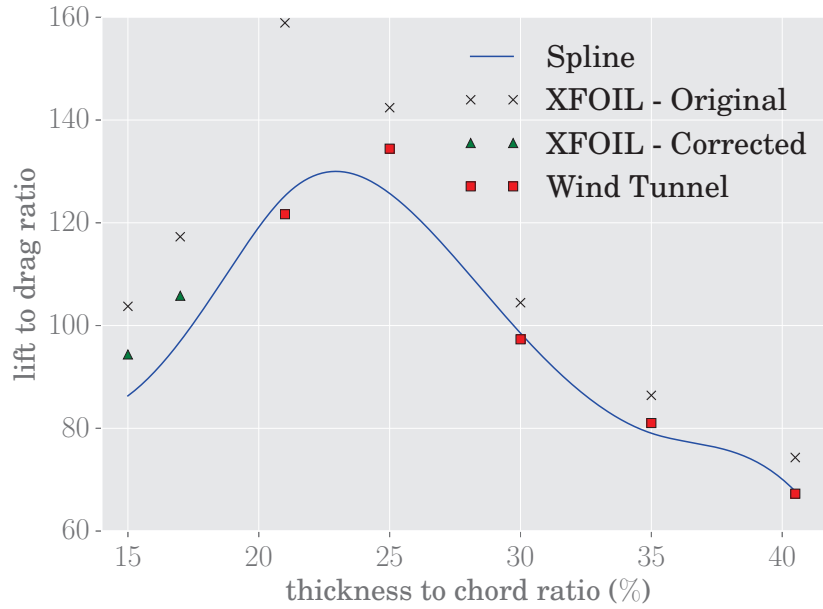


Figure 3.4: Wind tunnel spline using corrected XFOIL data at $\alpha = 5.0^\circ$.

data. This gives a more accurate estimate of the lift and drag coefficients than just XFOIL alone. A similar technique could be employed in any case where the airfoil analysis is particularly expensive such as it is with a wind tunnel.

3.5 Results

The full results from the precomputational method optimization cases are shown in Appendix Tables A.1, A.2, and A.3 for using XFOIL, RANS CFD, and wind tunnel data, respectively. Each compares the results from the sequential design, the precomputational method with thickness-to-chord ratio, and the precomputational method with thickness-to-chord ratio and the blended family airfoil family factor. The optimization results are all compared to the NREL 5-MW reference turbine evaluated with that airfoil analysis technique (i.e., the XFOIL reference blade is evaluated with XFOIL, the RANS CFD reference blade with RANS CFD, etc.). A summary of the major results are shown in Figure 3.5 and Table 3.1, including the changes in COE, the annual energy production (AEP), and the turbine capital costs (TCC) compared to the reference blade. A comparison of the chord, twist, spar cap thickness and trailing edge thickness distributions along the blade for each airfoil analysis is shown in Figure 3.6. A graphical representation of the airfoil

shapes along the blade span is shown in Figure 3.7, 3.8, and 3.9 for using XFOIL, RANS CFD, and wind tunnel data, respectively. For this research, the finally obtained solution from the precomputational method is run again with the free-form method to directly compare the performance.

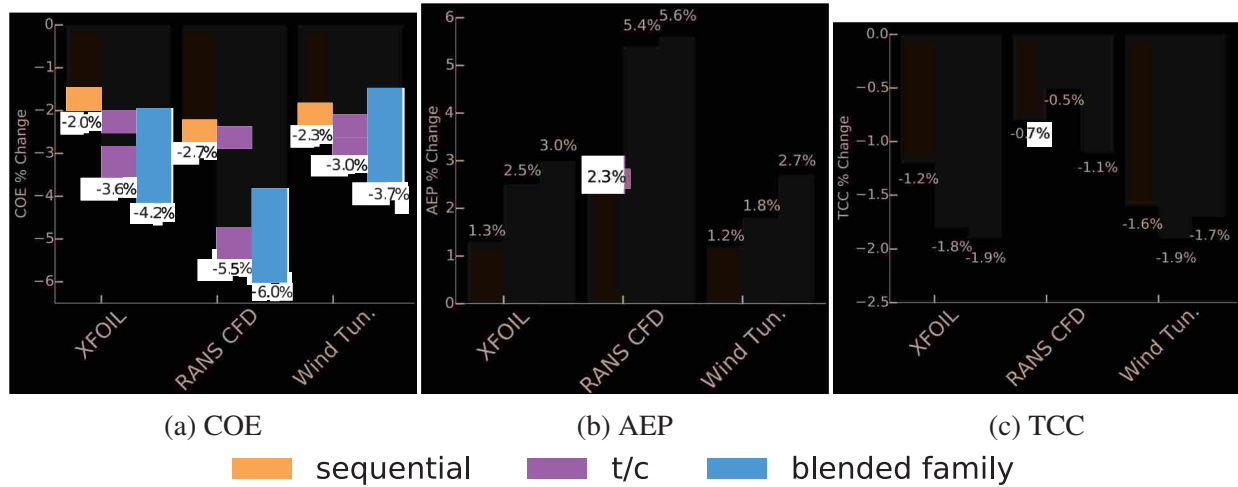


Figure 3.5: Summary of main results from the precomputational optimizations. Increased airfoil flexibility lead to better COE, mainly through an increase in AEP.

Table 3.1: Precomputational Method Results

		XFOIL	RANS CFD	Wind Tunnel
reference	COE (ϕ /kWh)	6.283	7.512	6.212
	AEP (MWh)	23,232	19,433	23,500
	TCC (\$)	9,207,436	9,207,436	9,207,436
sequential	COE (ϕ /kWh)	6.155 (-2.0%)	7.311 (-2.7%)	6.072 (-2.3%)
	AEP (MWh)	23,539 (+1.3%)	19,872 (+2.3%)	23,782 (+1.2%)
	TCC (\$)	9,100,633 (-1.2%)	9,143,561 (-0.7%)	9,060,217 (-1.6%)
integrated (t/c)	COE (ϕ /kWh)	6.056 (-3.6%)	7.102 (-5.5%)	6.023 (-3.0%)
	AEP (MWh)	23,805 (+2.5%)	20,482 (+5.4%)	23,925 (+1.8%)
	TCC (\$)	9,038,830 (-1.8%)	9,160,248 (-0.5%)	9,033,509 (-1.9%)
integrated (B_{af})	COE (ϕ /kWh)	6.020 (-4.2%)	7.058 (-6.0%)	5.980 (-3.7%)
	AEP (MWh)	23,932 (+3.0%)	20,524 (+5.6%)	24,132 (+2.7%)
	TCC (\$)	9,029,051 (-1.9%)	9,103,081 (-1.1%)	9,053,384 (-1.7%)

The results show significant reductions in COE through the integrated designs over the sequential design. For XFOIL, the COE reduction was 2.0%, 3.6%, and 4.2% for the sequential,

precomputational t/c , and precomputational B_{af} , respectively. For RANS CFD, the COE reduction was 2.7%, 5.5%, and 6.0% for the sequential, precomputational t/c , and precomputational B_{af} , respectively. For the wind tunnel, the COE reduction was 2.3%, 3.0%, and 3.7% for the sequential, precomputational t/c , and precomputational B_{af} , respectively. In every case and across all airfoil analysis methods, additional airfoil shape flexibility resulted in greater COE reductions. In Figure 3.5 we see how the thickness-to-chord ratio is able to capture the majority of the COE benefit, but the blended airfoil family factor does provide some additional benefit. The major source of COE reduction was a result of an increase in energy production rather than a reduction in turbine cost. The reduction in TCC varied widely between blade designs. Integrated design has a large effect on AEP and only a relatively small effect on TCC.

The changes in chord, twist, etc., were less affected by the airfoil parameterization method than by the airfoil analysis method. A bigger change occurred, on average, by switching from XFOIL to RANS CFD or to the wind tunnel data than by changing the design method. As can be seen in Figure 3.6, the blades from RANS CFD tended to have larger chord, more twist, and thicker materials than those from using XFOIL. This is because, in general, the drag coefficients were higher for RANS CFD than XFOIL and therefore the blade was forced to have a larger size to withstand the increased loads. The airfoil shape changes as seen in Figure 3.7, 3.8, and 3.9, in most cases, were relatively subtle. But subtle changes in airfoil shape can still have a large impact the COE. The most dramatic changes from the original airfoil shapes occurred from changes in the thickness-to-chord ratio. The thickness-to-chord ratio captured 73%, 85%, and 50% of the benefit of integrated design as compared to adding the blended airfoil family factor when using XFOIL, RANS CFD, and wind tunnel data, respectively. However, the addition of the blended airfoil family factor allowed the airfoil shapes to vary in more subtle but still important ways. The additional degree of freedom from the blended family was able to obtain about an average of about 35% improvement over the cases that only used thickness-to-chord ratio.

3.6 Discussion and Conclusions

Integrated precomputational blade designs show a significant improvement in COE over sequential designs. Optimizing the airfoil shape simultaneously with the rest of the blade allows the trade-offs between the aerodynamic performance of the blade and the needed structural thickness to

be better explored. Since the blade cost is only 17.7% of the entire rotor [44], the COE reductions of the entire turbine are actually substantial improvements to the blade cost and energy production. The reduction in COE holds relatively consistent for each of the airfoil analysis methods. The improvement in COE over the sequential design is shown in Table 3.2.

Table 3.2: Improvement in COE Reduction - Precomputational over Sequential

	XFOIL	RANS CFD	Wind Tunnel
integrated using t/c	1.4%	2.5%	0.7%
integrated using B_{af}	2.0%	3.0%	1.3%

As expected, as the flexibility of the airfoils increases so does the COE reduction. The COE reduction for the integrated design with t/c was an average of 1.7% more than the sequential design and 2.3% for the integrated design with B_{af} . The thickness-to-chord ratio captures the majority of the benefit of the precomputational method, 73%, 85%, and 50% for XFOIL, RANS CFD, and wind tunnel, respectively. The additional benefit from the blend between airfoil families had a modest but important effect of about an average of 35%. This is because the airfoil shapes could change more and better balance the lift and drag forces to both increase the turbine’s energy product as well as reduce costs.

The optimal blade design was quite different with the various airfoil analysis methods. While the overall reduction in COE between the various aerodynamic analysis cases were similar, the overall blade shape changed significantly in terms of airfoil shape, chord, twist, etc. This shows that the actual optimal blade design is quite sensitive to the aerodynamic performance of the airfoils that make up the blade. This alludes to the idea that high-fidelity data is perhaps even more important in an integrated blade optimization than it is in other scenarios. The optimal blade design is not as robust to changes in the airfoil analysis method as it is in airfoil parameterization method. While obtaining high-fidelity data is almost always recommended if possible, it is more so in this case.


The implication that the blade design is quite robust to changes in airfoil parameterization alludes to the idea that perhaps an extended sequential blade design might be adequate in some cases. The airfoils could be chosen and the blade optimized and *then the airfoils could be re-optimized* once the other design variables of the blade have been optimized. Additional work is

needed for this idea to be validated. Nevertheless, this research shows that there are changes in optimal blade design and performance with different airfoil parameterization techniques.


The main takeaway is that a significant benefit can be gained through precomputational parameterization methods *with minimal additional computational cost and complexity*. The average COE improvement over the sequential method is 2.3% and the optimization workflow is similar except for the creation of the surrogate model creation and connecting it with the structural and aerodynamic codes.

Additional airfoil shape flexibility allows for larger increases in AEP and is the major source of COE reduction. The average increase in AEP with the integrated design over the sequential design is 1.7% for t/c and 2.3% for the B_{af} . The average decrease in turbine capital cost (TCC) with the integrated design over the sequential design is 0.2% for t/c and 0.4% for the B_{af} . This shows that the major benefit from an integrated design is through an increase in AEP. In a few cases the integrated design increased cost slightly more than the sequential design, but made up for it with a larger increase in energy production.


It is recommended that the airfoil shape be added to aerostructural blade optimization with high-fidelity aerodynamic analysis tools and as much flexibility as possible while still maintaining the ability to precompute. Both thickness-to-chord ratio and blended family factor had an important impact on increasing the overall blade performance. Other precomputational parameterization methods could include blending more than two airfoil families, changing the airfoil camber, etc. This precomputational method maintained the ability to precompute and allowed the blade optimization to be faster than it would be otherwise. The free-form method is performed in the next chapter to determine whether or not this precomputational method is able to capture the majority of the benefit from an integrated design or if there is still a significant benefit to be attained from additional airfoil shape flexibility.




(a) XFOIL - chord distribution




(b) CFD - chord distribution




(c) Wind T. - chord distribution




(d) XFOIL - twist distribution




(e) CFD - twist distribution



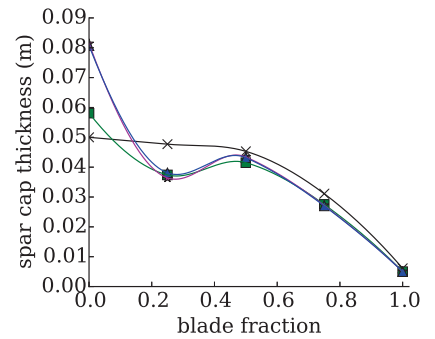
(f) Wind T. - twist distribution




(g) XFOIL - spar cap distribution




(h) CFD - spar cap distribution




(i) Wind T. - spar cap distribution



(j) XFOIL - trailing edge distribution



(k) CFD - trailing edge distribution



(l) Wind T. - trailing edge distribution

×× Reference ■■ Sequential ★★ t/c ▲▲ blended family

Figure 3.6: Precomputational chord, twist, spar cap thickness and trailing edge thickness distributions.




Figure 3.7: XFOIL precomputational airfoil shape results.

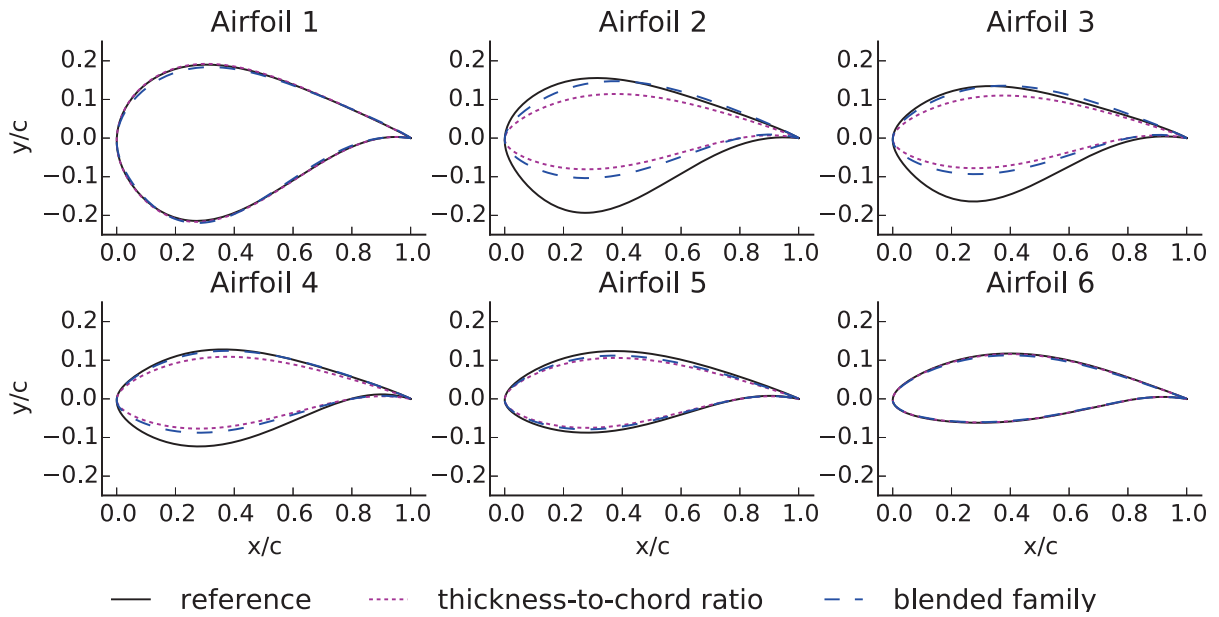


Figure 3.8: CFD precomputational airfoil shape results.




Figure 3.9: Wind tunnel precomputational airfoil shape results.

CHAPTER 4. FREE-FORM METHODS

4.1 Theory and Methodology

The precomputational method is able to achieve substantial reductions in COE for minimal additional computational cost and complexity. However, to adhere to the criteria of precomputation, only some flexibility within a specific airfoil family could be given to the airfoil shapes. Giving the airfoil shapes additional flexibility in the free-form method has the potential for additional benefit, but it is unknown how much additional benefit could be gained. Nevertheless, this additional flexibility forces us to relax the criteria of precomputing the airfoil analyses. This is because it would be essentially impossible to generate a surrogate model for the entire design space of airfoil shapes. A general model that would work with all airfoil shapes would require a very large number of airfoil analyses and negate the benefit of precomputing. The free-form method relaxes the criteria of precomputation and now performs the airfoil analyses within the optimization. The steps to the free-form method are as follows:

1. parameterize the airfoil shape with a free-form airfoil parameterization method
2. optimize to reduce COE while running the airfoil analyses directly inside of the optimization loop

Without precomputation in the free-form method, the standard workflow has changed substantially as is seen in Figure 1.8. Whereas with the sequential or precomputational method, the lift and drag coefficients could be looked up from airfoil tables or the surrogate model during the optimization, they now have to be computed through each iteration of the optimization. This can be quite computationally expensive, but allows us to have full flexibility in the airfoil shapes for the integrated blade design.

In the precomputational method, we used one or two design variables per airfoil (the thickness-to-chord ratio and/or the blended airfoil family factor). To give the airfoil shapes ad-

ditional freedom, more design variables are required (in this case we use eight parameters in the CST method explained in the next section). As discussed in the introduction, we are faced with two challenges: the amount of time per function call has increased and the number of design variables has also increased. For a free-form method with high-fidelity CFD techniques to be feasible for large-scale problems, analytic gradients are necessary even though they are more difficult to develop. Analytic gradients are obtained and used for both the free-form and the precomputational methods. More details on how these gradients are obtained are found in the Airfoil Shape Gradients section.

4.2 Parameterization Method - Class Shape Transformation

While there exists several methods to parametrize a 2D airfoil shape, including Hicks-Henne bumps, PARSEC technique, and B-spline curves [45], the CST method was chosen due to its simplicity, robustness, and the relatively low number of design parameters needed. The CST method was developed by Kulfan for the purpose of reducing the number of parameters needed to define any 2D or 3D smooth aerodynamic shape [19]. This method is described in the following equations [19,20]:

$$\zeta(\psi) = C_{N_2}^{N_1}(\psi) \sum_{n=0}^n A_n S_n + \psi \Delta \zeta_{TE} \quad (4.1)$$

Eq. 4.1 describes the main equation in the CST method. By defining certain parameters, the airfoil coordinates can be calculated for both the top and bottom airfoil surfaces. In this equation $\zeta = y/c$, $\psi = x/c$ (the normalized x- and y-coordinates of the airfoil shape), and ζ_{TE} defines the trailing edge thickness.

$$C_{N_2}^{N_1}(\psi) = (\psi)^{N_1} (1 - \psi)^{N_2} \quad (4.2)$$

Eq. 4.2 describes the class function that based on the parameters N_1 and N_2 determine the geometry class. The geometry class can range from 2D airfoils to 3D aircraft fuselage or any smooth aerodynamic shape.

$$S_i = K_{i,n} \psi^i (1 - \psi^i) \quad (4.3)$$

Eq. 4.3 describes the shape function, S_i , which is the summation of component shape functions that when multiplied by the weighing factors from Eq. 4.1, A_i , define a Bernstein polynomial.

$$K_{i,n} = \frac{n!}{i!(n-i)!} \quad (4.4)$$

Eq. 4.4 is the binomial coefficient of order n that defines the Bernstein polynomial order and the number of component shape functions ($n + 1$). The CST method is the summation of component shape functions that when multiplied by weighting factors, known as Kulfan parameters, A_1, \dots, A_n , define a Bernstein polynomial. Figure 4.1 shows two third-order Bernstein polynomials composed of the summation of four weighted component shape functions that define the top and bottom airfoil surfaces. By supplying the weighting or Kulfan parameters, A_i , where $i = 0, 1, \dots, n + 1$, the




Figure 4.1: The DU21_A17 airfoil is defined by the summation of weighted polynomial components (four on top and four on bottom).

y-coordinate can be calculated for any x-coordinate for the desired 2D or 3D smooth aerodynamic shape. For this research, only 2D airfoil shapes are needed, so by fixing $N_1 = 0.5$ and $N_2 = 1.0$ in Eq. 4.2 we limited the output coordinates to a 2D NACA type round nose and pointed aft end airfoils [19]. The class function in this case becomes $C_{1,0}^{0.5}(\psi) = \sqrt{\psi}(1 - \psi)$ and is performed for both the top and bottom surfaces of the airfoil. In summary, the CST method allows us to perform our airfoil shape parameterization through the specification of various Kulfan parameters ($A = [A_0, A_1, \dots, A_{n+1}]$) as design variables. While there is no restriction on the range of the Kulfan parameters, realistic airfoils tend to have Kulfan parameters that range from 1.0 to -1.0 based on whether it is for the top or bottom airfoil surface. Using a Python script, a function is defined with the eight Kulfan parameters as the input and the x- and y-airfoil coordinates as the output. The

Kulfan parameters for the bottom surface are defined first followed by the parameters for the top surface.

Based on previous uses of the Kulfan parameters in optimization and in an effort to minimize the number of design variables needed, a third-order Bernstein polynomial is used to define both the top and bottom airfoil surfaces [20]. This required the use of four Kulfan parameters for both the top and bottom surfaces of the airfoil. Therefore, a total of eight design variables are needed to define one airfoil shape along the wind turbine blade. For the entire wind turbine blade, six airfoils are used as in the NREL 5-MW reference turbine. As such, a total of 48 design variables (six airfoils with eight Kulfan parameters) are used in this free-form approach for the entire blade. This is a small number compared to the possibility of having hundreds or thousands of design variables by using other methods.

4.3 Airfoil Shape Gradients

Obtaining analytic gradients can substantially reduce the computational cost due to fewer function calls and faster convergence. A convenient way to obtain gradients is to split the analysis into smaller sections and then combine the partial derivatives of each section using a coupled adjoint solve to obtain total system derivatives [46]. Many of the analytic gradients in the blade analysis code, RotorSE, have already been developed using a combination of automatic differentiation and the adjoint method [1]. Automatic differentiation breaks up a derivative into a sequence of the basic arithmetic operations and functions used in the computer on an elementary level and are repeatably combined with the chain rule. For this research, we add the gradients for the airfoil shape parameters (S_{af}). The equation for the total load gradients is found through a combination of partial derivatives shown in Eq. 4.5.

$$\frac{dloads}{dS_{af}} = \frac{\partial loads}{\partial S_{af}} - \frac{\partial loads}{\partial \phi} \frac{\partial R}{\partial S_{af}} / \frac{\partial R}{\partial \phi} \quad (4.5)$$

Where R is the residual equation from the BEM method (Eq. 2.1), ϕ is the local inflow angle, and S_{af} are the airfoil shape parameters. $\partial R / \partial \phi$ is obtained using the Tapenade automatic differentiation tool [47].

The values for $\partial loads/\partial S_{af}$ and $\partial R/\partial S_{af}$ are dependent through the chain rule on the gradients of lift and drag coefficients with respect to the airfoil shape parameters ($\partial c_l/\partial S_{af}$ and $\partial c_d/\partial S_{af}$) as seen in Eq. 4.6.

$$\frac{\partial R}{\partial S_{af}} = \frac{\partial R}{\partial c_l} \frac{\partial c_l}{\partial S_{af}} + \frac{\partial R}{\partial c_d} \frac{\partial c_d}{\partial S_{af}} \quad (4.6)$$

For the precomputational method, since we have the surrogate model for the lift and drag coefficients we can easily obtain the gradients for the parameterization, in this case $\partial c_l/\partial(t/c)$ and $\partial c_d/\partial(t/c)$. The 2D spline method used in Python provides analytic gradients. For the blended family factor the gradients are calculated by hand from Eq. 3.1 and Eq. 3.2.

$$\frac{dc_l}{dB_{af}} = c_{l_1} - c_{l_0} \quad (4.7)$$

$$\frac{dc_d}{dB_{af}} = c_{d_1} - c_{d_0} \quad (4.8)$$

The derivatives of the lift and drag coefficients with respect to the blended family factor are calculated in Eq. 4.7 and Eq. 4.8. This method to obtain the gradients for the precomputational method is the same regardless of whether XFOIL, RANS CFD, or wind tunnel data is used.

For the free-form XFOIL, a version is used that can handle complex numbers¹. This allows the use of a gradient method called complex step to find $\partial c_l/\partial S_{af}$ and $\partial c_d/\partial S_{af}$. Complex step is similar to finite-differencing except that we can obtain exact gradients by using complex numbers to avoid the subtraction cancellation error. The equation for complex step gradients is shown in Eq. 4.9. A step size (h) of 10^{-20} is used.

$$\frac{df}{dx} = \frac{Im(f(x+ih))}{h} \quad (4.9)$$

For CFD, it is more complicated to obtain $\partial c_l/\partial S_{af}$ and $\partial c_d/\partial S_{af}$. Surface sensitivities for the lift and drag coefficients at the airfoil coordinates normal to the airfoil surface ($\partial c_l/\partial x_1, \dots, \partial c_l/\partial x_m$) are generated from *SU2*'s automatic differentiation tool [48, 49]. The direction of the unit normal vector of a point on the surface is the average of the normal vectors of the lines to its adjacent

¹<https://bitbucket.org/mdolab/pyxlight>

points. An example of these surface sensitivities on an airfoil is shown in Figure 4.2. These surface sensitivities help the optimizer to determine what effect a change in the airfoil shape at that airfoil coordinate will have on that airfoil's lift and drag coefficient.




Figure 4.2: Surface drag coefficient sensitivities ($\partial c_d/\partial x_1, \dots, \partial c_d/\partial x_m$) for each airfoil coordinate normal to the surface of the DU21_A17 airfoil and used to find $\partial c_d/\partial S_{af}$.

The chain rule is used in Eq. 4.10 to convert the lift and drag coefficient gradients with respect to each airfoil coordinate to the lift and drag coefficient gradients with respect to each Kulfan parameter in the CST method [20]. The gradients of the Kulfan parameters with respect to each airfoil coordinate ($\partial x_1/\partial A_1, \dots, \partial x_m/\partial A_n$) are found using automatic differentiation of the CST equations.

$$\frac{\partial c_l}{\partial S_{af}} = \begin{bmatrix} \frac{\partial c_l}{\partial A_1} \\ \frac{\partial c_l}{\partial A_2} \\ \vdots \\ \frac{\partial c_l}{\partial A_n} \end{bmatrix} = \begin{bmatrix} \frac{\partial x_1}{\partial A_1} & \cdots & \frac{\partial x_m}{\partial A_1} \\ \frac{\partial x_1}{\partial A_2} & \cdots & \frac{\partial x_m}{\partial A_2} \\ \vdots & \ddots & \vdots \\ \frac{\partial x_1}{\partial A_n} & \cdots & \frac{\partial x_m}{\partial A_n} \end{bmatrix} \begin{bmatrix} \frac{\partial c_l}{\partial x_1} \\ \frac{\partial c_l}{\partial x_2} \\ \vdots \\ \frac{\partial c_l}{\partial x_m} \end{bmatrix} \quad (4.10)$$

The values for $\partial c_l/\partial S_{af}$ and $\partial c_d/\partial S_{af}$ are used to obtain the total loads gradients with respect to the airfoil shape parameters ($dloads/dS_{af}$) from Eq. 4.5. The load gradients are prop-

agated to the objective function and the constraints gradients using OpenMDAO [42], which uses either the direct method or the adjoint method depending on the number of inputs to outputs.

4.4 Results

The main results from the optimization of the NREL 5-MW reference turbine are summarized in Table 4.1 and Figure 4.3. The precomputational results in this section refer to the results from the blended airfoil family factor for comparison purposes. The full optimization results are shown in Appendix Tables A.4 and A.5. As previously explained, wind tunnel data was not used for the free-form method because it would be infeasible to run the wind tunnel for every new airfoil shape needed by the optimization.

Results show significant reductions in COE through the integrated designs over the sequential design. For XFOIL, the COE reduction was 2.0%, 4.2%, and 4.7% for the sequential, precomputational, and free-form designs, respectively. For RANS CFD, the COE reduction was 2.7%, 6.0%, and 6.7% for the sequential, precomputational, and free-form designs, respectively. As expected, additional airfoil shape flexibility results in a larger COE reduction. The major source of COE reduction was from an increase in energy production rather than a reduction in turbine cost. The turbine cost remained fairly constant or increased between the sequential and the free-form design. The energy production increased significantly with additional airfoil shapes flexibility. The increase in energy production is a result of the change in the airfoil shapes. The flow characteristics around the airfoil shapes changed to allow the lift force along the blade span to increase, leading to higher rotations of the blade and therefore higher energy production.

A comparison of the chord, twist, spar cap thickness, and trailing edge thickness for each optimization case along the blade span is shown in Figure 4.4. A bigger change occurred, on average, by switching the airfoil analysis method from XFOIL to RANS CFD than by changing the airfoil shape parameters. The blades when analyzed with RANS CFD tended to have larger chord, more twist, and thicker materials than those from using XFOIL. This is because, in general, the drag coefficients were higher for RANS CFD than XFOIL and therefore the blade was forced to have a larger size to withstand the increased loads. The RANS CFD is able to better resolve the boundary layer than XFOIL and resulted in generally higher drag forces, especially in the stalling region. The fundamental differences in the physics of how the lift and drag coefficients are

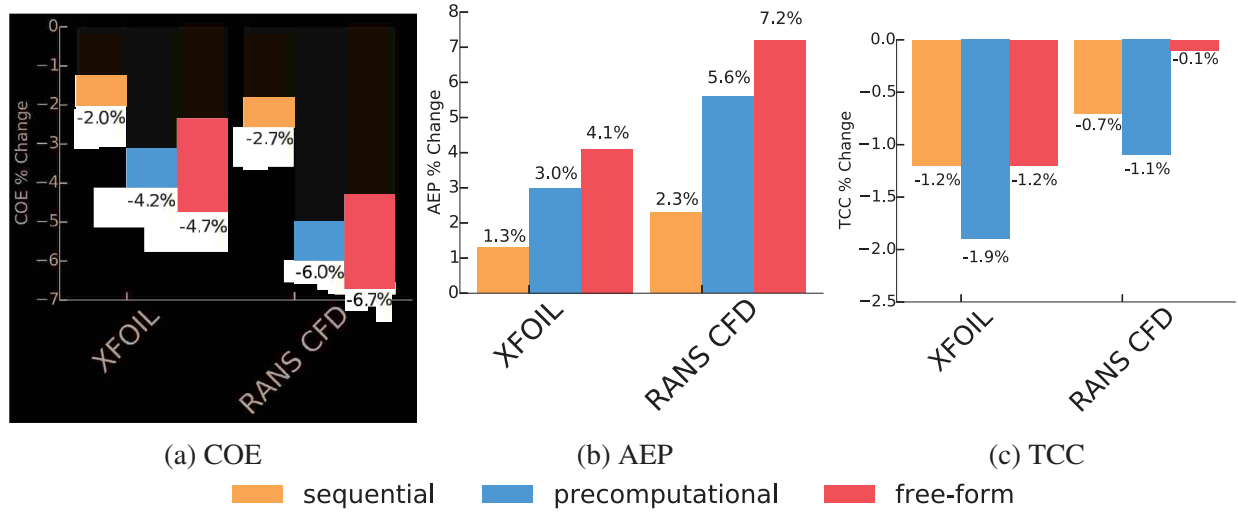


Figure 4.3: Summary of main results from the free-form optimizations.

Table 4.1: Free-Form Method Results

		XFOIL	RANS CFD
reference	COE (¢/kWh)	6.283	7.512
	AEP (MWh)	23,232	19,433
	TCC (\$)	9,207,436	9,207,436
sequential	COE (¢/kWh)	6.155 (-2.0%)	7.311 (-2.7%)
	AEP (MWh)	23,529 (+1.3%)	19,872 (+2.3%)
	TCC (\$)	9,100,633 (-1.2%)	9,143,561 (-0.7%)
precomputational (B_{af})	COE (¢/kWh)	6.020 (-4.2%)	7.058 (-6.0%)
	AEP (MWh)	23,932 (+3.0%)	20,524 (+5.6%)
	TCC (\$)	9,029,051 (-1.9%)	9,103,081 (-1.1%)
free-form (CST)	COE (¢/kWh)	5.990 (-4.7%)	7.002 (-6.7%)
	AEP (MWh)	24,176 (+4.1%)	20,835 (+7.2%)
	TCC (\$)	9,098,397 (-1.2%)	9,198,973 (-0.1%)

calculated in XFOIL and RANS CFD are directly shown in the resultant optimized blade. When comparing the blade parameters between the precomputational and free-form methods, the free-form method also tended to have heavier properties (i.e. larger chord, thickness distributions, etc.). The free-form method was able to generate airfoil shapes that had higher overall loads that resulted in increased energy production, but these increased loads had to be balanced with a larger blade. The change in chord, twist, etc. between the different design methods show that the optimal blade planform is dependent on the airfoils and vice versa.

A graphical representation of the airfoil shapes cross-sections taken at different points along the blade span is shown in Figure 4.5, where Airfoil 1 starts near the blade root and progresses to Airfoil 6 near the blade tip. The airfoil shape changes, in most cases, were relatively subtle, but still had a large impact on COE. The optimal airfoil shapes changed more from the original NREL 5-MW reference turbine airfoils with the RANS CFD cases than the XFOIL cases. This is why the benefit of integrated design was greater for RANS CFD, because the original airfoils when evaluated using RAN CFD were worse to begin with and the integrated design was thereby able to extract a greater benefit. In many cases, the airfoil shapes changes between the precomputational and the free-form methods were relatively similar. This shows that the precomputational method can converge on similar airfoil shapes despite having fewer parameters.

4.5 Discussion and Conclusions

The improvements in the COE reduction from the integrated designs over the sequential design are shown in Table 5.1. Results show that the integrated free-form method reduced COE by 2.7% and 4.0% more than the sequential design and 0.5% and 0.7% more than integrated precomputational method when using XFOIL and RANS CFD, respectfully. Precomputational methods can capture the majority of the benefit of integrated design for minimal additional computational cost and complexity. These results demonstrate the following main takeaways from this study.

Table 4.2: Improvement in COE Reduction - Free-Form over Sequential

	Precomputational	Free-Form
XFOIL	2.2%	2.7%
RANS CFD	3.3%	4.0%

Precomputational methods can capture the majority of the benefit of integrated design for minimal additional computational cost and complexity. This study quantified the differences in blade performance between two integrated blade design approaches: the precomputational and the free-form methods. The precomputational method captured 81.5% and 82.5% of the benefit found by the free-form method over the sequential design for XFOIL and RANS CFD, respectively. This method has a similar workflow to the sequential design due to its ability to precompute, but


manages to add some partial flexibility to the airfoil shapes. These results are promising for the adoption of the precomputational method because it was able to capture the majority of the benefit of integrated design (about 80%) without sacrificing the ease of precomputation. These results are promising for the adoption of the precomputational method because it was able to capture the majority of the benefit of integrated design (about 80%) without sacrificing the ease of precomputation. In many situations, the additional cost and complexity of the free-form method for only about 20% additional benefit may not be worth it. By carefully choosing the airfoil shape parameters, many of the major changes in the airfoil shapes can be obtained using fewer parameters. The choice between the two integrated design approaches balance the trade-off between computational cost and performance. If the extra effort is made, in both development time and computational cost, to obtain analytic gradients and perform the airfoil analyses during the optimization, the free-form method can reduce the COE more than the precomputational method. As wind energy continues to advance and computational capabilities increase, the extra benefit from a free-form method may be worth the additional cost and effort.

Nevertheless, it is important to note that the precomputational method cannot capture all of the benefit of integrated blade design. The free-form method was able to substantially reduce COE over sequential designs. However, when compared to the precomputational method, the free-form method achieved modest additional reductions of 25.0% and 21.2% in COE for XFOIL and RANS CFD, respectively. This about 20% improvement came as a result of the additional airfoil flexibility that allow the blade to obtain more favorable lift-to-drag ratios. The choice between the two integrated design approaches becomes a matter of balance between cost and performance. If the extra effort is made, in both development time and computational cost, to obtain analytic gradients and perform the airfoil analyses during the optimization, the free-form method can reduce the COE more than the precomputational method. As wind energy continues to advance and computational capabilities increase, the extra benefit from a free-form method may be worth the additional cost and effort.


The optimal blade planform is more dependent on the airfoil analysis method than on the airfoil parameterization method. There was a greater difference in the optimal blade planform by changing the airfoil analysis method than from changing the airfoil parameterization method. This demonstrates the need for higher-fidelity airfoil analysis methods. This study demonstrates tech-

niques to use higher-fidelity airfoil analysis methods for both the precomputational and the free-form methods so that blade designs can more accurately model real wind turbine blade behavior. While this study showed the results from the optimization of the NREL 5-MW reference turbine, a similar behavior is expected with other turbine types and site conditions when using integrated blade design. The plan is to add the code to perform the integrated designs to the open-source NREL WISDEM wind turbine analysis framework².


²<https://github.com/WISDEM/>




(a) XFOIL - Chord distribution




(b) CFD - Chord distribution




(c) XFOIL - Twist distribution




(d) CFD - Twist distribution




(e) XFOIL - Spar cap thickness distribution



(f) CFD - Spar cap thickness distribution



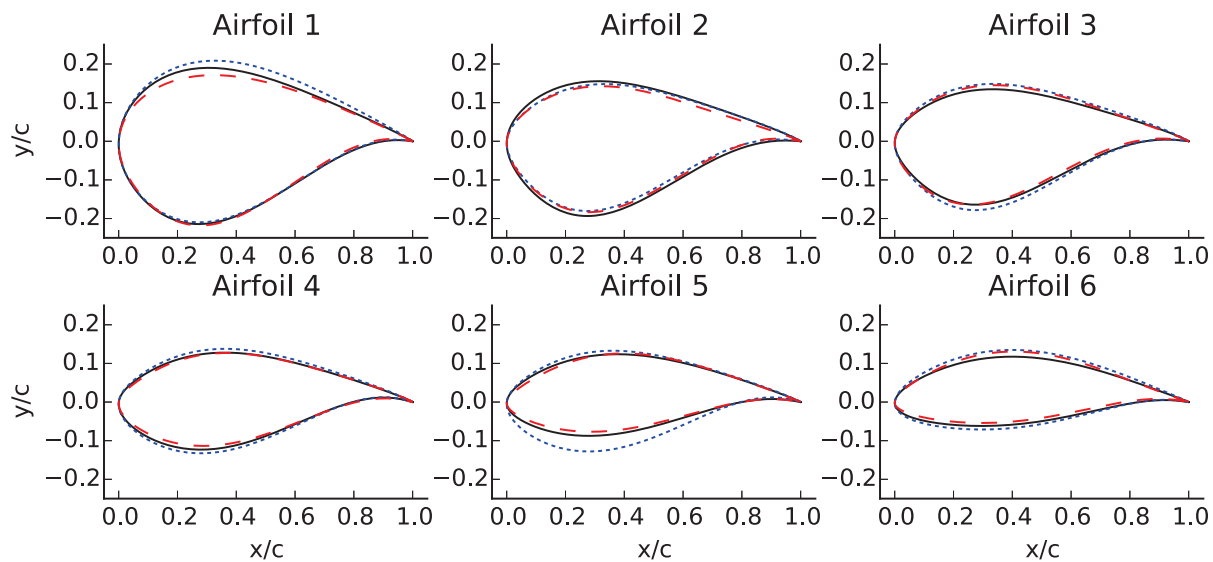
(g) XFOIL - Trailing edge thickness distribution




(h) CFD - Trailing edge thickness distribution

×—× reference ■—■ sequential ★—★ precomputational ▲—▲ free-form

Figure 4.4: Free-form chord, twist, and spar cap thickness and trailing edge thickness distributions. The RANS CFD tended to have heavier properties than XFOIL such as larger chord, more twist, and thicker materials.



(a) XFOIL



(b) RANS CFD

— reference ····· precomputational - - - free-form

Figure 4.5: Free-form airfoil shape results. The precomputational method was able to achieve many of the major changes in the airfoil shapes using fewer parameters than the free-form method.

CHAPTER 5. CONCLUSIONS

5.1 Thesis Review

The main results from all of the integrated design optimization cases as compared to sequential design is seen in Table 5.1. Integrated precomputational methods improved COE more than the sequential design by 1.6%, 2.8%, and 0.7% using thickness-to-chord ratio, and by 2.2%, 3.3%, and 1.4% using a blended airfoil family factor with XFOIL, RANS CFD, and wind tunnel data respectively. Integrated free-form methods improved COE more than the sequential design by 2.7% and 4.0% using the CST method with XFOIL and RANS CFD respectively. Giving the airfoil shapes additional flexibility reduced COE primarily by increasing annual energy production. The precomputational method can capture a majority of the benefit of integrated design for minimal additional computational cost and complexity, but the free-form method can provide modest additional benefits if the extra effort is made in computational cost and development time.

Table 5.1: Improvement in COE Reduction - Integrated Design over Sequential

	XFOIL	RANS CFD	Wind Tunnel
precomputational t/c	1.6%	2.8%	0.7%
precomputational B_{af}	2.2%	3.3%	1.4%
free-form CST	2.7%	4.0%	N/A

5.1.1 Research Objectives

The research objectives were accomplished and led to the following important results and conclusions.

1. Directly comparing the integrated design approaches to the sequential design show significant reductions in COE. The airfoil shape flexibility in the integrated blade design allow

for higher increases in the energy production of the wind turbine. The precomputational method is able to combine the advantage of precomputation with some of the benefit of integrated design. By comparing directly the precomputational method to the free-form method, it was determined that the precomputation method can capture a majority of the benefit of integrated design, in this case about 80%, for minimal additional effort and complexity. Nevertheless, the free-form method can provide some additional benefit, about 20%, over the precomputational method, if the additional effort is made in development time and computational cost.

2. A comparison of the different airfoil parameterization methods showed that additional degrees of freedom in the airfoil shapes allows for the optimization to find blade designs with lower COE. Nevertheless, each degree of freedom is not equal and there are diminishing returns with each additional degree of freedom. The thickness-to-chord ratio was able to achieve about 59% and 70% of the benefit compared to the free-form method for XFOIL and RANS CFD, respectively. Adding the additional degree of freedom of the blended airfoil family factor moved that to 81.5% and 82.5%, or an additional 22.5% and 12.5%, for XFOIL and RANS CFD, respectively. The additional 6 design variables from the CST method provided the additional 18.5% and 17.5% improvement for XFOIL and RANS CFD, respectively. The choice of airfoil shape parameters is very important and if done correctly can capture much of the benefit without needing a large number of design variables.
3. The improvements in the COE of each blade design were relatively consistent across the different airfoil analysis methods. Nevertheless, the airfoil analysis method caused large variations in the final blade shape. This is why in this study techniques to use higher-fidelity methods such as RANS CFD and using wind tunnel data were developed, so that integrated design can progress beyond only using XFOIL. This will allow for more accurate blade designs to be generated. This study demonstrated techniques for how RANS CFD and wind tunnel data can be used in the precomputational method and RANS CFD for the free-form method.
4. Due to the expensive nature of RANS CFD, analytic gradients were developed in order to reduce the number of function evaluations needed to converge an optimization. The use

of analytic gradients becomes particularly important for large-scale optimization problems. The total gradients are based on the lift and drag coefficients gradients with respect to the airfoil shape parameters ($\partial c_l / \partial S_{af}$ and $\partial c_d / \partial S_{af}$). These gradients are propagated to the total cost of energy gradients using the chain rule and either the adjoint or direct method. The development of these gradients allowed for higher-fidelity analysis methods to be used and to more accurately and quickly find the optimal blade design. These gradients allow the optimization to run more accurately and faster.

5. The code to perform integrated blade design was added to the NREL WISDEM open-source software. This involved changes to the codes RotorSE, CCBlade, and AirfoilPrep. The code was written in Python and interfaced to the XFOIL code written in Fortran and SU2 CFD code written in C++. This study was applied to NREL 5-MW reference turbine, but by making the code for integrated design open-source it can be more easily applied to other turbine types and site conditions¹.

5.1.2 Main Takeaways

By accomplishing the research objectives, the following three main takeaways were determined.

Precomputational methods can capture the majority of the benefit of integrated design for minimal additional computational cost and complexity. The precomputational method captured 81.5% and 82.5% of the benefit found by the free-form method over the sequential design for XFOIL and RANS CFD, respectively. Precomputational methods can capture the majority of the benefit (about 80% in this case) of integrated design for minimal additional computational cost and complexity. By maintaining the ease of precomputation, this method has a similar workflow to the sequential design, but manages to add some airfoil shape flexibility. The integrated precomputational method has the potential to replace the sequential design in many cases. The additional cost and complexity of the free-form method for only 20% additional benefit may not be worth it in many situations.

¹<https://github.com/WISDEM/>

Some modest improvements can be gained from the free-form method if the extra effort is made to obtain analytic gradients and run the airfoil analyses directly. In most situations, the precomputational method can provide most of the benefit needed, but cannot provide all of the benefit in blade performance from integrated blade design. The free-form method was able to achieve an additional 25.0% and 21.2% improvement over the precomputational method as a result of the additional airfoil shape flexibility for XFOIL and RANS CFD, respectively. This improvement comes through the additional effort in both development time and computational cost to obtain analytic gradients and perform the airfoil analysis during the optimization. There is a trade-off between cost and improved blade performance between these two integrated design approaches. The additional cost and effort in performing the free-form method may be worth it as wind energy continues to develop and progress and as computational costs continue to decrease.

The optimal blade planform is more dependent on the airfoil analysis method than on the airfoil parameterization method. The importance of using higher-fidelity airfoil analysis methods, such as RANS CFD, has been demonstrated by this study. The airfoil analysis method plays a large role on the optimal blade planform while the choice in airfoil shape parameters has a more moderate influence. Through techniques developed in this study, such as though the use of analytic gradients, higher-fidelity analysis methods can be used for both the precomputational and the free-form methods so that blade designs can more accurately model real wind turbine blade behavior.

5.2 Future Work

Future work is needed to continue to investigate integrated blade design. This research used the BEM method to convert the 2D airfoil shapes analysis into the 3D blade. Although the BEM method is fairly accurate, performing the 3D analysis would better capture the 3D effects and interactions that exist between the airfoil sections. As stated in the introduction, research into 3D CFD integrated design has begun, but needs further development.

The trend toward analytic gradients is also an important trend because giving the blade additional degrees of freedom will continue to generate better blade designs. As wind turbine opti-

mization continues to mature, analytic gradients will be needed to handle large-scale optimization problems. Developing these gradients for more design variables would be beneficial.

This research increased the flexibility of the airfoil shapes and additional freedom could be given to the structures. This additional flexibility in the structures could involve the choice of materials in the composite layers, the thicknesses of each individual layup instead of the entire composite, or the ply orientation and order.

REFERENCES

- [1] Ning, A., and Petch, D., 2016. “Integrated design of downwind land-based wind turbines using analytic gradients.” *Wind Energy*, Feb. 11, 42
- [2] Lu, X., McElroy, M. B., and Kiviluoma, J., 2009. “Global potential for wind-generated electricity.” In *Proceedings of the National Academy of Sciences of the United States of America*, pp. 10933–10938. 2
- [3] Chehouri, A., Younes, R., Ilinca, A., and Perron, J., 2015. “Review of performance optimization techniques applied to wind turbines.” *Applied Energy*, **142**, pp. 361–388. 6
- [4] Xudong, W., Shen, W. Z., Zhu, W. J., Sørensen, J. N., and Jin, C., 2009. “Shape optimization of wind turbine blades.” *Wind Energy*, **12**(8), pp. 781–803. 6
- [5] Méndez, J., and Greiner, D., 2006. “Wind blade chord and twist angle optimization by using genetic algorithms.” In *Proceedings of the Fifth International Conference on Engineering Computational Technology*, Vol. 6, pp. 12–15. 6
- [6] Tangler, J., and Kocurek, J. D., 2005. “Wind turbine post-stall airfoil performance characteristics guidelines for blade-element momentum methods.” In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, pp. 10–13. 6
- [7] Bottasso, C., Croce, A., Sartori, L., and Grasso, F., 2014. “Free-form design of rotor blades.” In *Journal of Physics: Conference Series*, Vol. 524, IOP Publishing, p. 012041. 7, 10, 11
- [8] Kenway, G., and Martins, J., 2008. “Aerostructural shape optimization of wind turbine blades considering site-specific winds.” In *Proceedings of the 12th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Sept*, pp. 10–12. 7, 10, 11, 21
- [9] Polat, O., and Tuncer, I. H., 2013. “Aerodynamic shape optimization of wind turbine blades using a parallel genetic algorithm.” *Procedia Engineering*, **61**, pp. 28 – 31 25th International Conference on Parallel Computational Fluid Dynamics. 7, 10, 11
- [10] Herbert-Acero, J. F., Martínez-Lauranchet, J., Probst, O., Méndez-Díaz, S., Valenzuela-Rendón, M., Réthoré, P.-E., and Castillo-Villar, K. K. “Aerodynamic Optimization of Small Wind Turbine Rotors Based on NACA 4-Digit Airfoils through Computational Intelligence.”. 7, 10, 11
- [11] Fiore, G., and Selig, M. S., 2015. “Optimization of wind turbine airfoils subject to particle erosion.” In *33rd AIAA Applied Aerodynamics Conference*, p. 3393. 7, 10, 11
- [12] Fischer, G. R., Kipouros, T., and Savill, A. M., 2014. “Multi-objective optimisation of horizontal axis wind turbine structure and energy production using aerofoil and blade properties as design variables.” *Renewable Energy*, **62**, pp. 506 – 515. 7, 10, 11

- [13] Zhu, W. J., Shen, W. Z., and Sørensen, J. N., 2014. “Integrated airfoil and blade design method for large wind turbines.” *Renewable Energy*, **70**, pp. 172–183. 7, 10
- [14] Vucina, D., Marinic-Kragic, I., and Milas, Z., 2016. “Numerical models for robust shape optimization of wind turbine blades.” *Renewable Energy*, **87, Part 2**, pp. 849 – 862. 7, 10, 11
- [15] Dhert, T., Ashuri, T., and Martins, J. R. R. A., 2016. “Aerodynamic shape optimization of wind turbine blades using a Reynolds-averaged Navier-Stokes model and an adjoint method.” *Wind Energy*. 7, 10
- [16] Barrett, R., and Ning, A., 2016. “Comparison of airfoil precomputational analysis methods for optimization of wind turbine blades.” *IEEE Transactions on Sustainable Energy*, **7(3)**, July, pp. 1081–1088.
- [17] Kulfan, B. M., and Bussoletti, J. E., 2006. “Fundamental parametric geometry representations for aircraft component shapes.” In *11th AIAA/ISSMO multidisciplinary analysis and optimization conference*, Vol. 6948, sn. 11
- [18] Sripawadkul, V., Padulo, M., and Guenov, M., 2010. “A comparison of airfoil shape parameterization techniques for early design optimization.” In *13th AIAA/ISSMO multidisciplinary analysis optimization conference*, pp. 13–15. 11
- [19] Kulfan, B., 2007. “CST universal parametric geometry representation method with applications to supersonic aircraft.” In *4th International Conference on Flow Dynamics*, pp. 26–28. 11, 40, 41
- [20] Hewitt, P., and Marques, S., 2014. “Aerofoil Optimisation Using CST Parameterisation in SU2.” In *Royal Aeronautical Society Biennial Applied Aerodynamics Research Conference, 2014*. 11, 40, 42, 44
- [21] Straathof, M. H., and L. Van Tooren, M. J., 2011. “Extension to the class-shape-transformation method based on B-splines.” *AIAA journal*, **49(4)**, pp. 780–790. 11
- [22] Ciampa, P. D., Zill, T., and Nagel, B., 2010. “CST parametrization for unconventional aircraft design optimization.” In *27th Congress of the International Council of the Aeronautical Sciences (ICAS)*. 11
- [23] Lyu, Z., Xu, Z., and Martins, J., 2014. “Benchmarking optimization algorithms for wing aerodynamic design optimization.” In *Proceedings of the 8th International Conference on Computational Fluid Dynamics, Chengdu, Sichuan, China*. 11
- [24] Rios, L. M., and Sahinidis, N. V., 2013. “Derivative-free optimization: a review of algorithms and comparison of software implementations.” *Journal of Global Optimization*, **56(3)**, Jul, pp. 1247–1293. 11
- [25] Ning, A., Damiani, R., and Moriarty, P. J., 2014. “Objectives and constraints for wind turbine optimization.” *Journal of Solar Energy Engineering*, **136(4)**, June, pp. 04010–041010. 15, 20, 21, 23, 24

- [26] Drela, M., 1989. “XFOIL: An Analysis and Design System for Low Reynolds Number Airfoils.” In *Low Reynolds Number Aerodynamics*, T. Mueller, ed., Vol. 54 of *Lecture Notes in Engineering*. Springer Berlin Heidelberg, pp. 1–12. 16
- [27] Palacios, F., Colonno, M. R., Aranake, A. C., Campos, A., Copeland, S. R., Economon, T. D., Lonkar, A. K., Lukaczyk, T. W., Taylor, T. W., and Alonso, J. J., 2013. “Stanford University Unstructured (SU2): An open-source integrated computational environment for multi-physics simulation and design.” *AIAA Paper*, **287**, p. 2013. 16
- [28] Palacios, F., Economon, T. D., Aranake, A. C., Copeland, S. R., Lonkar, A. K., Lukaczyk, T. W., Manosalvas, D. E., Naik, K. R., Padrón, A. S., Tracey, B., Variyar, A., and Alonso, J., 2014. “Stanford University Unstructured (SU2): Open-source analysis and design technology for turbulent flows.” *AIAA paper*, **243**, pp. 13–17.
- [29] Villalpando, F., Reggio, M., and Ilinca, A., 2011. “Assessment of turbulence models for flow simulation around a wind turbine airfoil.” *Modelling and simulation in engineering*, **2011**, p. 6. 17
- [30] Jonkman, J. M., Butterfield, S., Musial, W., and Scott, G., 2009. *Definition of a 5-MW reference wind turbine for offshore system development*. National Renewable Energy Laboratory Golden, CO. 18, 25, 30
- [31] Sande, B., Pijl, S., and Koren, B., 2011. “Review of computational fluid dynamics for wind turbine wake aerodynamics.” *Wind Energy*, **14**(7), pp. 799–819. 18
- [32] Du, Z., and Selig, M. S., 1998. “A 3-D stall-delay model for horizontal axis wind turbine performance prediction.” *AIAA Paper*, **21**. 18
- [33] Eggers, A., Chaney, K., and Digumarthi, R., 2003. “An assessment of approximate modeling of aerodynamic loads on the uae rotor.” In *ASME 2003 Wind Energy Symposium*, American Society of Mechanical Engineers, pp. 283–292. 18
- [34] Ning, A., 2014. “A simple solution method for the blade element momentum equations with guaranteed convergence.” *Wind Energy*, **17**(9), pp. 1327–1345. 19, 20
- [35] Viterna, L. A., and Janetzke, D. C., 1982. Theoretical and experimental power from large horizontal-axis wind turbines Tech. rep., National Aeronautics and Space Administration, Cleveland, OH (USA). Lewis Research Center. 19, 20
- [36] Kong, C., Bang, J., and Sugiyama, Y., 2005. “Structural investigation of composite wind turbine blade considering various load cases and fatigue life.” *Energy*, **30**(11), pp. 2101 – 2114. 21
- [37] Petrone, G., De Nicola, C., Quagliarella, D., Witteveen, J., Axerio-Cilies, J., and Iaccarino, G., 2011. “Wind turbine optimization under uncertainty with high performance computing.” *AIAA paper*(2011-3806). 22
- [38] Petrone, G., Nicola, C., Quagliarella, D., Witteveen, J., and Iaccarino, G., 2011. “Wind turbine performance analysis under uncertainty.” *AIAA Paper No. AIAA*, **544**. 22

- [39] Ning, A., and Dykes, K., 2014. “Understanding the benefits and limitations of increasing maximum rotor tip speed for utility-scale wind turbines.” *Journal of Physics: Conference Series*, **524**(012087), Jun. 22, 23
- [40] European Committee for Standardisation, 1993. Eurocode 3: Design of steel structures—part 1-6: General rules—supplementary rules for the shell structures EN 1993-1-6: 20xx, European Committee for Standardisation. 23
- [41] Gill, P. E., Murray, W., and Saunders, M. A., 2002. “SNOPT: An SQP algorithm for large-scale constrained optimization.” *SIAM Journal on Optimization*, **12**(4), pp. 979–1006. 24
- [42] Gray, J. S., Moore, K. T., and Naylor, B. A., 2010. “OpenMDAO: An Open Source Framework for Multidisciplinary Analysis and Optimization.” In *13th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference, Fort Worth, TX, AIAA, AIAA-2010-9101*, AIAA. 24, 45
- [43] Jonkman, J. M., and Buhl Jr, M. L., 2005. FAST users guide Tech. rep., National Renewable Energy Laboratory. 28
- [44] Jamieson, P., 2011. *Innovation in wind turbine design*. John Wiley & Sons. 34
- [45] Castonguay, P., and Nadarajah, S., 2007. *Effect of Shape Parameterization on Aerodynamic Shape Optimization*. American Institute of Aeronautics and Astronautics, 2016/05/31. 40
- [46] Martins, J. R., and Hwang, J. T., 2013. “Review and unification of methods for computing derivatives of multidisciplinary computational models.” *AIAA journal*, **51**(11), pp. 2582–2599. 42
- [47] Hascoet, L., and Pascual, V., 2013. “The tapenade automatic differentiation tool: principles, model, and specification.” *ACM Transactions on Mathematical Software (TOMS)*, **39**(3), p. 20. 42
- [48] Jameson, A., Shankaran, S., and Martinelli, L., 2008. “Continuous adjoint method for unstructured grids.” *AIAA journal*, **46**(5), pp. 1226–1239. 43
- [49] Albring, T., 2013. “Gradient calculation using algorithmic differentiation.” *RAE*, **2822**, p. 10. 43

APPENDIX A. DETAILED OPTIMIZATION RESULTS

The following tables provide more detailed information about the specific changes in the design parameters and objective from each optimization. They are listed in the following order:

1. Precomputational XFOIL (Table A.1)
2. Precomputational CFD (Table A.2)
3. Precomputational Wind Tunnel (Table A.3)
4. Free-Form XFOIL (Table A.4)
5. Free-Form CFD (Table A.5)

Table A.1: Precomputational XFOIL Optimization Results

	units	Reference	Sequential	Integrated t/c	Integrated B_{af}
c	m	[3.26, 4.57, 3.32, 1.46]	[2.64, 4.69, 2.93, 1.47]	[2.59, 3.85, 2.54, 1.3]	[2.50, 3.47, 2.50, 1.30]
c_{max}		0.236	0.306	0.319	0.397
θ	degree	[13.28, 7.46, 2.89, -0.09]	[11.23, 4.07, 2.07, -1.26]	[12.83, 2.63, -0.176, -3.26]	[10.01, 2.43, -0.453, -3.45]
λ_2		7.55	7.56	7.36	7.17
t_{spar}	m	[0.05, 0.047754, 0.0454, 0.031085, 0.0061398]	[0.0466, 0.0407, 0.0424 , 0.0306, 0.005]	[0.0946, 0.0388, 0.0405, 0.0249, 0.005]	[0.1314, 0.0437, 0.0414, 0.0253, 0.005]
t_{te}	m	[0.1, 0.09569, 0.06569, 0.02569, 0.00569]	[0.0658, 0.0306 , 0.0177, 0.0085 , 0.005]	[0.0122, 0.0313, 0.0121, 0.0101, 0.005]	[0.1026, 0.0334, 0.0188, 0.0081, 0.005]
t/c		[0.404458, 0.349012, 0.29892, 0.251105, 0.211299, 0.179338]	-	[0.411, 0.318, 0.310, 0.272, 0.261, 0.207]	[0.42, 0.3285, 0.3269, 0.2704, 0.2607, 0.2060]
bf		[0.0, 0.0, 0.0, 0.0, 0.0, 1.0]	-	-	[0.1621, 0.00002, 0.0154, 0.00002, 0.0, 1.0]
TCC	\$	9,207,436	9,100,633 (-1.2%)	9,038,830 (-1.8%)	9,029,051 (-1.9%)
AEP	MWh	23,232	23,539 (+1.3%)	23,805 (+2.5%)	23,932 (+3.0%)
COE	cents/kWh	6.283	6.155 (-2.0%)	6.056 (-3.6%)	6.020 (-4.2%)

Table A.2: Precomputational CFD Optimization Results

	units	Reference	Sequential	Integrated t/c	Integrated B_{af}
c	m	[3.26, 4.57, 3.32, 1.46]	[2.72, 4.82, 3.71, 1.43]	[2.83, 5.30, 3.22, 1.60]	[2.31, 4.38, 3.18, 1.3]
c_{max}		0.236	0.307	0.351	0.337
θ	degree	[13.28, 7.46, 2.89, -0.09]	[18.35, 6.49, 3.91, 2.22]	[19.0, 5.69, 2.32, 3.28]	[15.92, 5.70, 1.26, -1.76]
λ_2		7.55	6.54	6.79	6.80
t_{spar}	m	[0.05, 0.047754, 0.045376, 0.031085, 0.0061398]	[0.0743, 0.0431, 0.0498, 0.0307, 0.005]	[0.0267, 0.0547, 0.0506, 0.0310, 0.005]	[0.1120, 0.0479, 0.0474, 0.0318, 0.005]
t_{te}	m	[0.1, 0.09569, 0.06569, 0.02569, 0.00569]	[0.0439, 0.0339, 0.0173, 0.0105, 0.005]	[0.0707, 0.0289, 0.0274, 0.0094, 0.005]	[0.1026, 0.0334, 0.0188, 0.0081, 0.005]
t/c		[0.404458, 0.349012, 0.29892, 0.251105, 0.211299, 0.179338]	-	[0.409, 0.195, 0.188, 0.186, 0.181, 0.177]	[0.404, 0.251, 0.229, 0.212, 0.191, 0.173]
bfi		[0.0, 0.0, 0.0, 0.0, 0.0, 1.0]	-	-	[0.05, 0.02, 0.01, 0.004, 0.033, 0.982]
TCC	\$	9,207,436	9,143,561 (-0.7%)	9,160,248 (-0.5%)	9,103,081 (-1.1%)
AEP	MWh	19,433	19,872 (+2.3%)	20,482 (+5.4%)	20,524 (+5.6%)
COE	cents/kWh	7.512	7.311 (-2.7%)	7.102 (-5.5%)	7.058 (-6.0%)

Table A.3: Precomputational Wind Tunnel Optimization Results

	units	Reference	Sequential	Integrated t/c	Integrated B_{af}
c	m	[3.26, 4.57, 3.32, 1.46]	[2.71, 3.40, 2.34, 1.30]	[2.33, 3.93, 2.22, 1.30]	[2.48, 4.19, 2.38, 1.31]
c_{max}		0.236	0.456	0.359	0.346
θ	degree	[13.28, 7.46, 2.89, -0.09]	[9.54, 6.15, 2.44, -0.49]	[13.10, 5.78, 0.76, -1.17]	[10.11, 6.26, 1.67, -1.67]
λ_2		7.55	8.86	8.15	8.98
t_{spar}	m	[0.05, 0.047754, 0.045376, 0.031085, 0.0061398]	[0.1334, 0.0468, 0.0479, 0.0302, 0.005]	[0.0814, 0.0364, 0.0434, 0.0272, 0.005]	[0.08067, 0.03810, 0.04316, 0.02703, 0.005]
t_{te}	m	[0.1, 0.09569, 0.06569, 0.02569, 0.00569]	[0.0122, 0.0311, 0.0157, 0.0109, 0.005]	[0.0597, 0.0256, 0.0165, 0.0098, 0.005]	[0.04958, 0.02778, 0.01572, 0.00866, 0.005]
t/c		[0.404458, 0.349012, 0.29892, 0.251105, 0.211299, 0.179338]	-	[0.290, 0.215, 0.212, 0.275, 0.277, 0.220]	[0.4023, 0.3545, 0.3029, 0.2517, 0.2096, 0.1789]
bf		[0.0, 0.0, 0.0, 0.0, 0.0, 1.0]	-	-	0.0001, 0.0, 0.0, 0.0, 0.0008, 0.97]
TCC	\$	9,207,436	9,060,217 (-1.6%)	9,033,509 (-1.9%)	9,053,384 (-1.7%)
AEP	MWh	23,500	23,782 (+1.2%)	23,925 (+1.8%)	24,132 (+2.7%)
COE	cents/kWh	6.212	6.072 (-2.3%)	6.023 (-3.0%)	5.980 (-3.7%)

Table A.4: Free-Form XFOIL Optimization Results

	units	Reference	Sequential	Integrated CST
c	m	[3.26, 4.57, 3.32, 1.46]	[2.65, 4.70, 2.91, 1.43]	[2.62, 4.65, 2.88, 1.42]
c_{max}		0.236	0.306	0.302
θ	degree	[13.28, 7.46, 2.89, -0.09]	[11.16, 4.09, 2.03, -1.59]	[11.05, 4.02, 2.02, -1.66]
λ_2		7.55	7.57	7.58
t_{spar}	m	[0.05, 0.0477, 0.0453, 0.0311, 0.0061]	[0.0435, 0.0407, 0.0417, 0.0305, 0.005]	[0.0475, 0.0425, 0.0437, 0.0293, 0.0052]
t_{te}	m	[0.1, 0.0957, 0.0657, 0.0257, 0.0057]	[0.0651, 0.0311, 0.0181, 0.0090, 0.005]	[0.0679, 0.0314, 0.0183, 0.0094, 0.005]
A_1		[-0.492, -0.729, -0.382, 0.137, 0.504, 0.548, 0.376, 0.370]	-	[-0.484, -0.740, -0.375, 0.144, 0.493, 0.533, 0.377, 0.364]
A_2		[-0.380, -0.759, -0.218, 0.086, 0.384, 0.484, 0.267, 0.347]	-	[-0.373, -0.755, -0.225, 0.129, 0.355, 0.486, 0.248, 0.320]
A_3		[-0.298, -0.679, -0.157, 0.128, 0.284, 0.460, 0.217, 0.338]	-	[-0.301, -0.673, -0.153, 0.134, 0.285, 0.470, 0.223, 0.333]
A_4		[-0.274, -0.407, -0.292, 0.279, 0.236, 0.437, 0.254, 0.311]	-	[-0.268, -0.402, -0.291, 0.266, 0.226, 0.447, 0.248, 0.319]
A_5		[-0.196, -0.289, -0.206, 0.191, 0.229, 0.399, 0.289, 0.295]	-	[-0.185, -0.285, -0.203, 0.190, 0.198, 0.410, 0.291, 0.292]
A_6		[-0.172, -0.137, -0.243, 0.151, 0.206, 0.355, 0.328, 0.259]	-	[-0.141, -0.139, -0.223, 0.172, 0.197, 0.392, 0.342, 0.249]
TCC	\$	9,207,436	9,100,633 (-1.2%)	9,098,397 (-1.2%)
AEP	MWh	23,232	23,529 (+1.3%)	24,176 (+4.1%)
COE	cents/kWh	6.283	6.155 (-2.0%)	5.990 (-4.7%)

Table A.5: Free-Form CFD Optimization Results

	units	Reference	Sequential	Integrated CST
c	m	[3.26, 4.57, 3.32, 1.46]	[2.72, 4.82, 3.71, 1.43]	[2.80, 5.13, 4.09, 1.30]
c_{max}		0.236	0.307	0.347
θ	degree	[13.28, 7.46, 2.89, -0.09]	[18.35, 6.49, 3.91, 2.22]	[12.99, 7.39, 2.96, 0.90]
λ_2		7.55	6.54	5.52
t_{spar}	m	[0.05, 0.0477, 0.0453, 0.0311, 0.0061]	[0.0743, 0.0431, 0.0498, 0.0307, 0.005]	[0.0941, 0.0558, 0.0539, 0.0424, 0.0063]
t_{te}	m	[0.1, 0.0957, 0.0657, 0.0257, 0.0057]	[0.0439, 0.0339, 0.0173, 0.0105, 0.005]	[0.0554, 0.0376, 0.0219, 0.0114, 0.0058]
A_1		[-0.492, -0.729, -0.382, 0.137, 0.504, 0.548, 0.376, 0.370]	-	[-0.355, -0.289, -0.399, 0.150, 0.345, 0.549, 0.400, 0.395]
A_2		[-0.380, -0.759, -0.218, 0.086, 0.384, 0.484, 0.267, 0.347]		[-0.326, -0.433, -0.277, 0.150, 0.225, 0.496, 0.267, 0.381]
A_3		[-0.298, -0.679, -0.157, 0.128, 0.284, 0.460, 0.217, 0.338]		[-0.246, -0.323, -0.206, 0.150, 0.231, 0.476, 0.320, 0.385]
A_4		[-0.274, -0.407, -0.292, 0.279, 0.236, 0.437, 0.254, 0.311]		[-0.215, -0.283, -0.181, 0.189, 0.203, 0.417, 0.280, 0.337]
A_5		[-0.196, -0.289, -0.206, 0.191, 0.229, 0.399, 0.289, 0.295]		[-0.187, -0.160, -0.220, 0.159, 0.183, 0.370, 0.310, 0.302]
A_6		[-0.172, -0.137, -0.243, 0.151, 0.206, 0.355, 0.328, 0.259]		[-0.153, -0.200, -0.128, 0.134, 0.144, 0.296, 0.199, 0.239]
TCC	\$	9,207,436	9,143,561 (-0.7%)	9,198,973 (-0.1%)
AEP	MWh	19,433	19,872 (+2.3%)	20,835 (+7.2%)
COE	cents/kWh	7.512	7.311 (-2.7%)	7.002 (-6.7%)