

AERODYNAMIC SHAPE OPTIMIZATION OF A
VERTICAL AXIS WIND TURBINE

by

TRAVIS JUSTIN CARRIGAN

Presented to the Faculty of the Graduate School of
The University of Texas at Arlington in Partial Fulfillment
of the Requirements
for the Degree of

MASTER OF SCIENCE IN AEROSPACE ENGINEERING

THE UNIVERSITY OF TEXAS AT ARLINGTON

December 2010

Copyright © by TRAVIS JUSTIN CARRIGAN 2010
All Rights Reserved

To my mother Linda and my father Brian who always
pushed me to work hard and pursue my goals.

ACKNOWLEDGEMENTS

I would like to start off by thanking my supervising professor Dr. Brian H. Dennis for giving me a place to work and providing advice and insight in the areas of fluid mechanics, aerodynamics, and computational fluid dynamics. I also would like to thank my academic advisors Dr. Zhen Xue Han and Dr. Bo Ping Wang for their interest in my research and words of encouragement and for taking time to serve on my thesis committee.

I would also like to thank my friends and coworkers at Pointwise for their continued support throughout my education. I wish to thank Dr. Chris Sideroff and Ms. Carolyn Dear for always pushing me to work harder and for the many discussions and invaluable pieces of advice.

I would like to thank my parents, Linda and Brian, and my sister Kayla, for their encouragement and patience throughout my education. I would also like to express my gratitude to Pete Stewart who not only trusted me enough to allow me to help restore aircraft in his hangar, but for inspiring me to pursue a degree in Aerospace Engineering.

Finally, I am extremely grateful to Richard Neves for sitting in the right seat of N8137U and teaching me how to fly.

September 30, 2010

ABSTRACT

AERODYNAMIC SHAPE OPTIMIZATION OF A VERTICAL AXIS WIND TURBINE

TRAVIS JUSTIN CARRIGAN, M.S.

The University of Texas at Arlington, 2010

Supervising Professor: Brian H. Dennis

The purpose of this study is to introduce and demonstrate a fully automated process for optimizing the airfoil cross-section of a vertical axis wind turbine (VAWT). The objective is to maximize the torque while enforcing typical wind turbine design constraints such as tip speed ratio, solidity, and blade profile. By fixing the tip speed ratio and solidity of the wind turbine, there exists an airfoil cross-section for which the torque can be maximized, requiring the development of an iterative design system.

The design system required to maximize torque incorporates rapid geometry generation and automated hybrid mesh generation tools with viscous, unsteady computational fluid dynamics (CFD) simulation software. The flexibility and automation of the modular design and simulation system allows for it to easily be coupled with a parallel differential evolution algorithm used to obtain an optimized blade design that maximizes the efficiency of the wind turbine.

TABLE OF CONTENTS

ACKNOWLEDGEMENTS	iv
ABSTRACT	v
LIST OF FIGURES	ix
LIST OF TABLES	xi
Chapter	Page
1. INTRODUCTION	1
1.1 Alternative Energy	1
1.2 Wind Turbine Types	2
1.2.1 Horizontal Axis Wind Turbines	2
1.2.2 Vertical Axis Wind Turbines	5
1.3 Computational Modeling	9
1.4 Objectives	11
2. VERTICAL AXIS WIND TURBINE PERFORMANCE	14
2.1 Ideal Performance and the Betz Limit	14
2.2 Wind Speed and Tip Speed Ratio	14
2.3 Geometry Definition	15
2.4 Performance Prediction	16
2.4.1 Angle of Attack	16
2.4.2 Average Torque	18
2.4.3 Power and Efficiency	19
3. METHODOLOGY	21
3.1 Requirements	21

3.2	Unique Modular Design	21
3.3	Directory Structure	23
3.4	Complete Automation	24
3.5	Parametric Studies and Optimization	24
4.	TOOLBOX	27
4.1	Geometry Generation	27
4.1.1	NACA 4-Series Airfoils	27
4.1.2	Airfoil Constraints	29
4.2	Grid Generation	30
4.2.1	Cell Types	31
4.2.2	Grid Considerations	31
4.2.3	Pointwise	34
4.3	Solver	35
4.3.1	Conservation Equations	36
4.3.2	Turbulence Modeling	37
4.3.3	FLUENT	39
4.3.4	Computational Considerations	41
4.4	Postprocessing	41
4.4.1	Average Torque	42
4.5	Optimization	43
4.5.1	Differential Evolution	44
4.5.2	Mutation	45
4.5.3	Crossover	47
4.5.4	Selection	47
5.	RESULTS	49
5.1	Grid Dependency Studies	49

5.1.1	Structured Grid Topology	50
5.1.2	Hybrid Grid Topology	53
5.1.3	Grid Independent Solution	55
5.2	Baseline Geometry	58
5.2.1	Baseline Performance	59
5.2.2	Published Performance Trends	61
5.3	Case 1	62
5.3.1	Optimization Results	63
5.3.2	Baseline Comparison	66
5.4	Case 2	70
5.4.1	Optimization Results	70
5.4.2	Baseline Comparison	74
6.	CONCLUSIONS	78
	REFERENCES	80
	BIOGRAPHICAL STATEMENT	85

LIST OF FIGURES

Figure	Page
1.1 Wind turbine types	3
1.2 HAWT blade root and tip airfoil sections [1]	5
1.3 Velocity and force components for a Darrieus type VAWT	7
1.4 Velocity and force components for a Savonius type VAWT	8
2.1 Force components for an airfoil	17
3.1 VAWT optimization methodology	22
3.2 Requirement for a simple directory structure	23
3.3 Directory structure for parametric study (left) and optimization (right)	25
3.4 Optimization directory details	26
4.1 NACA 4424 4-series airfoil	28
4.2 Range of airfoils: NACA 0004 (top) and NACA 9940 (bottom)	30
4.3 Different cell types used for boundary layer grids	32
4.4 Grid type accuracy vs. ease of use [2]	33
4.5 Hybrid grid for an airfoil	34
4.6 Leading edge boundary layer grid for blade geometry	34
4.7 Variation of torque as a function of time	42
4.8 Illustration showing the generation of the mutant vector \vec{v}	46
5.1 Structured grid topology	51
5.2 Structured farfield grid	52
5.3 Hybrid grid topology	54
5.4 Hybrid farfield grid	54

5.5	Structured and hybrid grid family comparison	56
5.6	Average torque comparison for grid independent solution determination	57
5.7	Hybrid grids for VAWT geometry, $\sigma = 1.5$ (left) and $\sigma = 0.4$ (right)	57
5.8	NACA 0015 baseline geometry	58
5.9	NACA 0015 performance envelope, $\sigma = 1.5$	59
5.10	NACA 0015 performance envelope, $\sigma = 0.4$	60
5.11	Simulated performance (left) and experimental performance (right)	61
5.12	Initial populations, 1 st test case	63
5.13	COP vs. generation for all population members, 1 st test case	64
5.14	Max COP vs. generation, 1 st test case	64
5.15	Optimized NACA 4-series airfoil geometry, 1 st test case	65
5.16	NACAopt performance envelope, 1 st test case	66
5.17	NACAopt vs. NACA 0015 performance, 1 st test case	67
5.18	NACAopt vs. NACA 0015 torque for $\lambda = 1$, 1 st test case	68
5.19	NACAopt (left) and NACA 0015 (right) vorticity contours, 1 st test case	69
5.20	Initial populations, 2 nd test case	71
5.21	COP vs. generation for all population members, 2 nd test case	72
5.22	Max COP vs. generation, 2 nd test case	72
5.23	Optimized NACA 4-series airfoil geometry, 2 nd test case	73
5.24	NACAopt performance envelope, 2 nd test case	74
5.25	NACAopt vs. NACA 0015 performance, 2 nd test case	75
5.26	NACAopt vs. NACA 0015 torque for $\lambda = 3$, 2 nd test case	76
5.27	NACAopt (left) and NACA 0015 (right) vorticity contours, 2 nd test case	77

LIST OF TABLES

Table		Page
5.1	Optimization Test Cases	49
5.2	Structured Grid Family	53
5.3	Hybrid Grid Family	55

CHAPTER 1

INTRODUCTION

1.1 Alternative Energy

Wind is everywhere. As long as the Earth continues to provide the right conditions, it will remain that way. All it takes is a difference in pressure to get a mass of air moving. This movement of air from areas of high pressure to areas of low pressure is what generates wind. Because this mass of air is moving, it has energy, renewable energy that has been used to provide thrust to sailboats and ships crossing the oceans, to windmills used to pump water for irrigation or grinding up grain. Even today, wind is still harnessed for much the same reason as it was thousands of years ago, but something it can provide in today's day and age is electricity. Today, only a small fraction of the world's electricity is generated by wind, however, demand for this renewable energy resource will continue to increase with the depletion of fossil fuels.

As the world continues to use up non-renewable energy resources, wind energy will continue to gain popularity. A new market in wind energy technology has emerged that has the means of efficiently transforming the energy available in the wind to a useable form of energy, such as electricity. The cornerstone of this new technology is the wind turbine.

A wind turbine is a type of turbomachine that transfers fluid energy to mechanical energy through the use of blades and a shaft and converts that form of energy to electricity through the use of a generator. Depending on whether the flow is parallel to the axis of rotation (axial flow) or perpendicular (radial flow), determines the

classification of the wind turbine. Each type of wind turbine has its strengths and weaknesses, but in the end, all wind turbines accomplish the same task.

1.2 Wind Turbine Types

Two major types of wind turbines exist based on their blade configuration and operation. The first type is the horizontal axis wind turbine (HAWT). This type of wind turbine is the most common and can often be seen littered across the landscape in areas of relatively level terrain with predictable year round wind conditions. HAWTs sit atop a large tower and have a set of blades that rotate about an axis parallel to the flow direction. These wind turbines have been the main subject of wind turbine research for decades, mainly because they share common operation and dynamics with rotary aircraft.

The second major type of wind turbine is the vertical axis wind turbine (VAWT). This type of wind turbine rotates about an axis that is perpendicular to the oncoming flow, hence, it can take wind from any direction. VAWTs consist of two major types, the Darrieus rotor and Savonius rotor. The Darrieus wind turbine is a VAWT that rotates around a central axis due to the lift produced by the rotating airfoils, whereas a Savonius rotor rotates due to the drag created by its blades. There is also a new type of VAWT emerging in the wind power industry which is a mixture between the Darrieus and Savonius designs. All the wind turbine types mentioned can be seen in Fig. 1.1.

1.2.1 Horizontal Axis Wind Turbines

The HAWT is the most popular and widely used type of wind turbine. Companies such as Vestas, Siemens, and GE develop and deploy HAWTs around the world, making them the largest and most successful wind turbine manufacturers. Backing

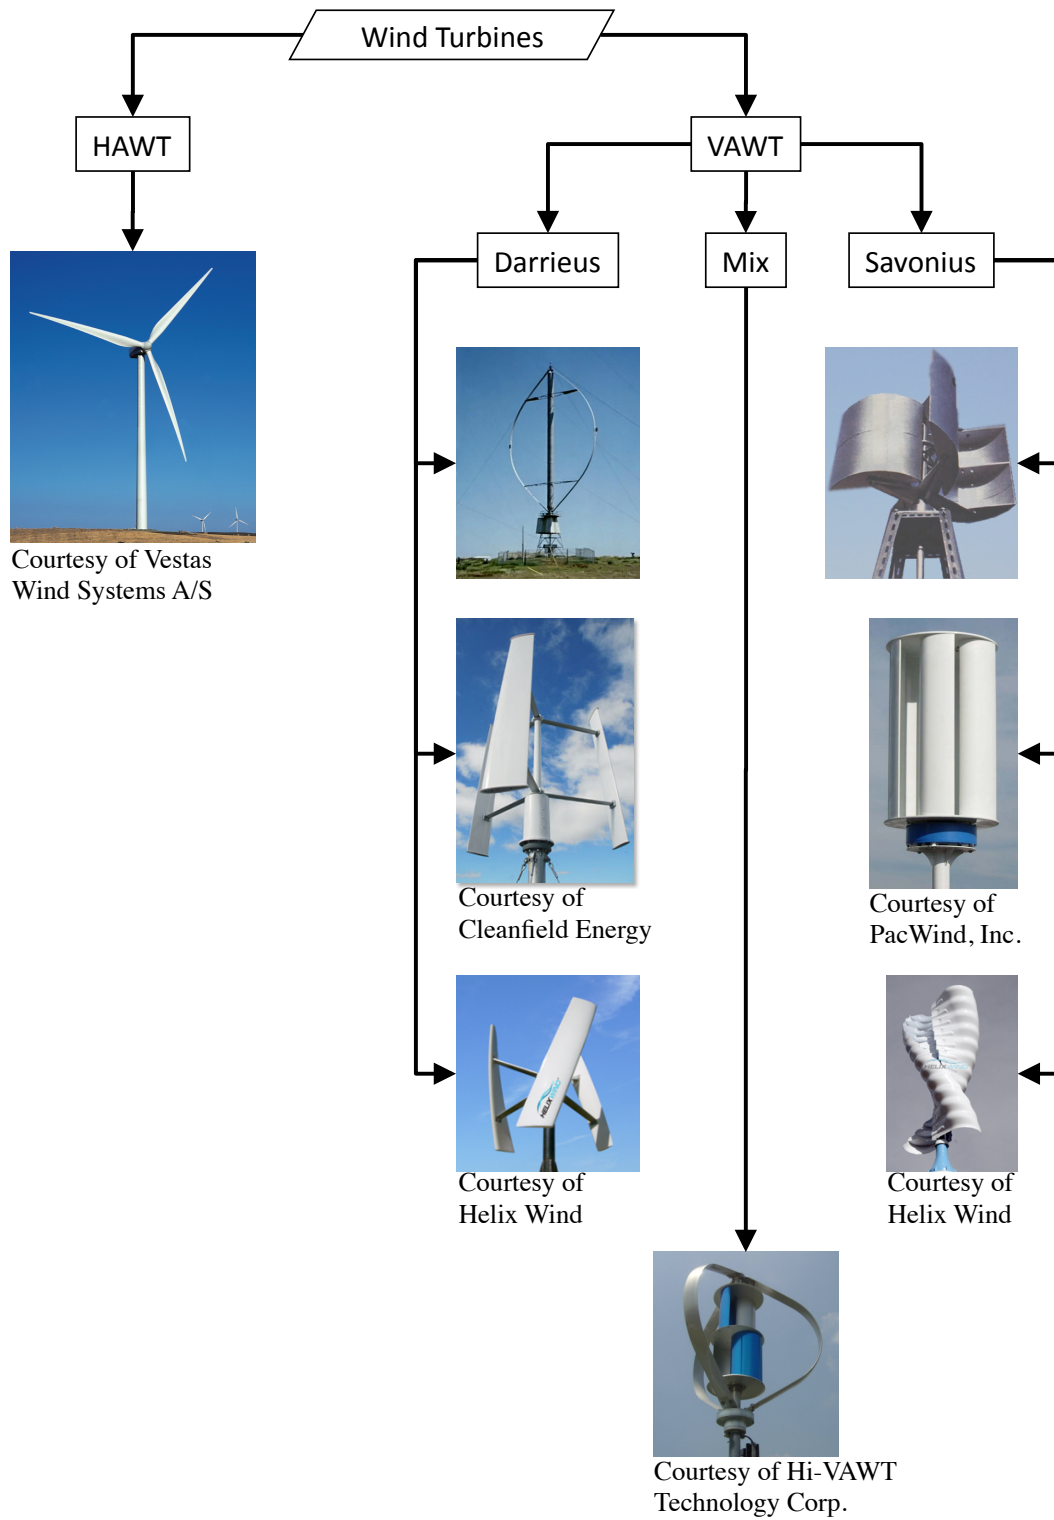


Figure 1.1. Wind turbine types.

these companies is rock solid wind turbine technology based on the foundation and experience gained from propellor and rotary-wing aircraft. All major wind farms around the world employ many HAWTs working as a team to aide in the generation of electricity for small towns and large cities.

The blades of a HAWT work to extract energy from the wind by generating lift, resulting in a net torque about the axis of rotation. As torque is produced through the slow turning of the wind turbine blades, the gearbox speeds up this rotation for the production of electricity through the use of a generator. To accomplish this task efficiently, especially for large HAWTs, active pitch controllers are used to ensure that each blade is adjusted to maintain an optimal angle of attack for maximum power extraction for a given wind speed. A yaw controller is also used to actively yaw the blades into the wind. However, these active control systems are complex and require more moving parts and effort to install than a VAWT assembly where the only moving part is the rotor, and the majority of components are located at the base of the turbine.

The design and manufacturing of a HAWT blade is complex as the blade is tapered and twisted with varying cross-sections in order to achieve optimum aerodynamic performance. A HAWT blade assembly can be seen in Fig. 1.2. The change in the cross-section and twist of the blade from the root to the tip is due to the variation of the relative velocity component w . Because the tip of the blade spins much faster than the root, the twist of the blade is shallow and the cross-section is thin. It can be seen in the figure that due to the high relative velocity at the tip, the resultant force F acting on the blade section is extremely high, but only a small portion of that force F_d is driving the rotation. However, because of the blades design, the thin blade profile with a low angle of twist at the tip produces roughly the same amount of torque as the root of the blade due to its large moment arm. Because of the slower

rotation speed at the root, the blade profile is much thicker with a higher angle of twist than at the tip. This design means a larger portion of the resultant force is in the direction of rotation. However, due to the lower moment arm at the root, the torque distribution over the entire blade is fairly uniform.

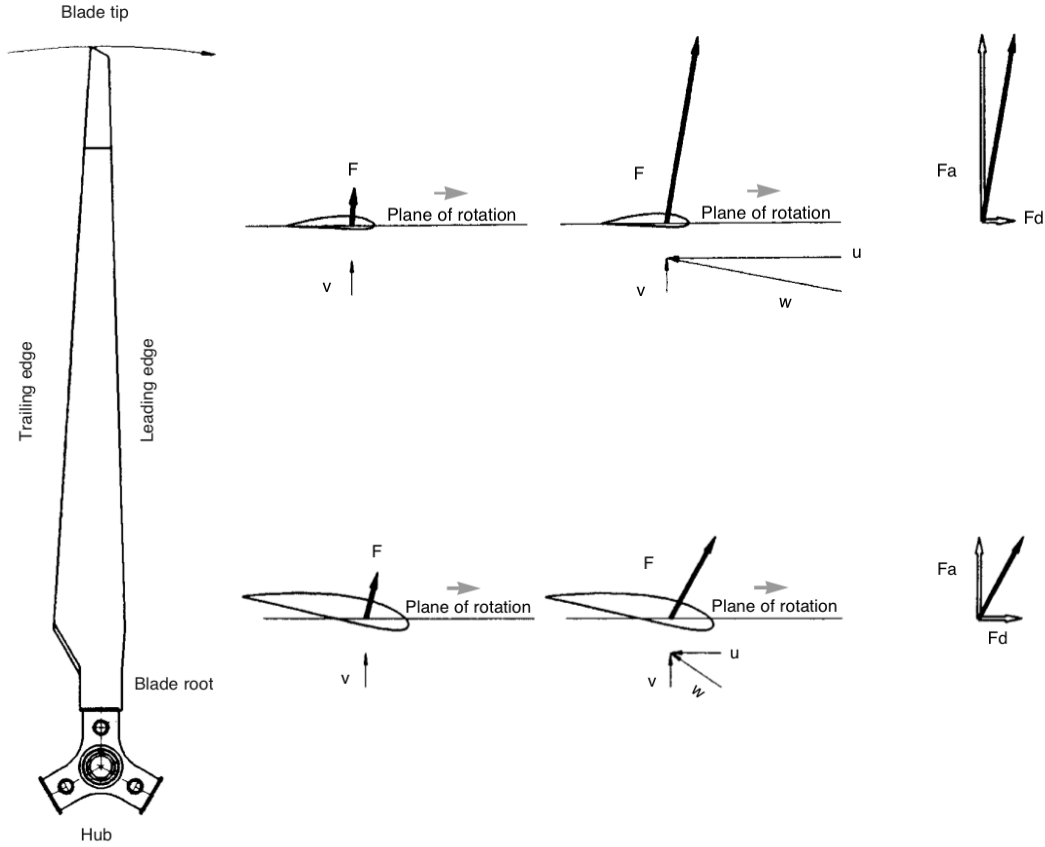


Figure 1.2. HAWT blade root and tip airfoil sections [1].

1.2.2 Vertical Axis Wind Turbines

Recently, VAWTs have been gaining popularity due to interest in personal green energy solutions. Small companies all over the world have been marketing these new devices such as Helix Wind, Urban Green Energy, and Windspire. VAWTs target

individual homes, farms, or small residential areas as a way of providing local and personal wind energy. This reduces the target individual's dependence on external energy resources and opens up a whole new market in alternative energy technology. Because VAWTs are small, quiet, easy to install, can take wind from any direction, and operate efficiently in turbulent wind conditions, a new area in wind turbine research has opened up to meet the demands of individuals willing to take control and invest in small wind energy technology.

The device itself is relatively simple. With the major moving component being the rotor, the more complex parts like the gearbox and generator are located at the base of the wind turbine. This makes installing a VAWT a painless undertaking and can be accomplished quickly. Manufacturing a VAWT is much simpler than a HAWT due to the constant cross section blades. Because of the VAWTs simple manufacturing process and installation, they are perfectly suited for residential applications.

The VAWT rotor, comprised of a number of constant cross-section blades, is designed to achieve good aerodynamic qualities at various angles of attack. Unlike the HAWT where the blades exert a constant torque about the shaft as they rotate, a VAWT rotates perpendicular to the flow, causing the blades to produce an oscillation in the torque about the axis of rotation. This is due to the fact that the local angle of attack for each blade is a function of its azimuthal location. Because each blade has a different angle of attack at any point in time, the average torque is typically sought as the objective function. Even though the HAWT blades must be designed with varying cross-sections and twist, they only have to operate at a single angle of attack throughout an entire rotation. However, VAWT blades are designed such that they exhibit good aerodynamic performance throughout an entire rotation at the various angles of attack they experience leading to high time averaged torque. The blades of

a Darrieus VAWT (D-VAWT) accomplish this through the generation of lift, while the Savonius type VAWTs (S-VAWT) produce torque through drag.

A snapshot of the cross-section of a D-VAWT blade can be seen in Fig. 1.3. As the blade rotates, the local angle of attack α for that blade changes due to the variation of the relative velocity W . The induced velocity V_i and the rotating velocity ωr of the blade govern the orientation and magnitude of the relative velocity. This in turn changes the lift L and the drag D forces acting on the blade. As the lift and drag change both their magnitude and orientation, the resultant force F_R changes. The resultant force can be decomposed into both a normal component F_N and a tan-




Figure 1.3. Velocity and force components for a Darrieus type VAWT.

gential component F_T . It is this tangential force component that drives the rotation of the wind turbine and produces the torque necessary to generate electricity. Due

to the mechanism driving the rotation of a D-VAWT, it is possible that the blades can travel faster than the speed of the freestream velocity V_∞ , or in other words, the tip speed ratio at which these wind turbines operate is often higher than 1, leading to higher efficiency.

A S-VAWT generates electricity through drag rather than lift like the D-VAWT. A cross-section of a Savonius wind turbine can be seen in Fig. 1.4. As the wind hits the concave portion of the blade (the bucket), it becomes trapped and pushes the blade around, advancing the next bucket into position. This continues as long as the wind is blowing and can overcome the friction of the shaft about which the blades rotate. A Savonius rotor typically rotates with a velocity equivalent to the speed of the freestream velocity, or a tip speed ratio of 1. Because of its lower rotation speed, Savonius rotors are associated with lower efficiencies and are not capable of




Figure 1.4. Velocity and force components for a Savonius type VAWT.

providing adequate electricity, but rather serve as a device used to reduce the overall dependence on other energy resources. However, due to the Savonius wind turbines simplicity, it is extremely easy to construct; some have even been built using large plastic blue polydrums with the capability of providing up to 10% of a household's electricity [4].

1.3 Computational Modeling

The majority of wind turbine research is focused on accurately predicting efficiency. Various computational models exist, each with their own strengths and weaknesses that attempt to accurately predict the performance of a wind turbine. Descriptions of the general set of equations that the methods solve can be found in Chap. 2. Being able to numerically predict wind turbine performance offers a tremendous benefit over classic experimental techniques, the major benefit being that computational studies are more economical than costly experiments.

A survey of aerodynamic models used for the prediction of VAWT performance was conducted by Islam [5] and Paraschivoiu [6]. While other approaches have been published, the three major models include momentum models, vortex models, and computational fluid dynamics (CFD) models. Each of the three models are based on the simple idea of being able to determine the relative velocity, and in turn, the tangential force component of the individual blades at various azimuthal locations.

The first major method used for predicting VAWT performance is the momentum model. The momentum model assumes the rotor is an actuator disk across which the forces acting on the blades of the wind turbine are equal to the change in momentum of the air. Airfoil lift, drag, and pitching moment coefficients are obtained from experimental data, making this method semi-empirical. The first momentum method developed for a VAWT was the single streamtube method introduced by Templin [7].

Templin assumed that the entire wind turbine rotor was enclosed by a single streamtube in which the induced velocity was constant. The multiple streamtube model was introduced a year later by Strickland [8]. In the multiple streamtube model, several streamtubes were used that stretched across the blade rotor. In each streamtube the induced velocity was assumed to be constant. Further modification was done by Paraschivoiu [9], Berg [10], and Templin [11] in which a double multiple streamtube model was introduced. This model assumes that two actuator disks in tandem arrangement are used along with multiple stream tube theory such that the induced velocity was different in the upstream and downstream portions of the rotor. Various other momentum models have been introduced using curvilinear streams rather than rectilinear flow approaching the wind turbine [12], and improving the prediction of local blade Reynolds number when determining the aerodynamic forces acting on the blade [13]. While these methods are relatively simple to understand and implement, they tend to break down for high tip speed ratios and high solidity where viscous effects become dominant.

Vortex models are the second major method used for predicting VAWT performance and operate under the assumption of incompressible, potential flow. The blades are replaced by a moving bound vortex. As the blades spin, the strength of the bound vortex varies as blade forces change given by the Kutta-Joukowski theorem. This can be attributed to the change in relative wind and angle of attack the blades experience as they rotate. In order to conserve total circulation, a vortex sheet is shed into the wake. At any point in the flow, the induced velocity due to a vortex sheet can be found using the the Biot-Savart law. With the freestream velocity and induced velocity known, the relative velocity can be determined for each of the blades. Once the relative velocity for the blades is known at all azimuthal locations, the lift and drag of the airfoil can be determined using empirical airfoil coefficient data tables.

While the model operates under the assumption of potential flow, the use of experimental data allows for unsteady, viscous effects to be considered. The vortex model was first introduced by Larsen [14] and Strickland [15]. But more recent advances by McIntosh [16] have been developed for swept bladed VAWTs. Vortex models have been shown to be extremely accurate, but tend to suffer convergence related issues and still rely on empirical data to approximate viscous effects.

The last method used for predicting VAWT performance is through the use of CFD where the entirety of the flow is calculated by numerically solving the Navier-Stokes equations. The Navier-Stokes equations are a set of non-linear, coupled, partial differential equations for which an exact solution still does not exist. For the simulation of a wind turbine, the equations include the continuity equation and momentum equations. Turbulence is considered using the Reynolds-averaged Navier-Stokes equations where the introduction of a new term representing the turbulent stress gives rise to a number of turbulence models. Because of its flexibility, CFD has been gaining popularity for analyzing the complex, unsteady aerodynamics involved in the study of wind turbines [17], [18], [19] and has demonstrated an ability to generate results that compare favorably with experimental data [20], [21], [22]. CFD has shown no problems predicting the performance of either high or low solidity wind turbines or for various tip speed ratios. However, it is important to note that predicting the performance of a wind turbine using CFD typically requires large computational domains with sliding interfaces and additional turbulence modeling to capture unsteady effects, therefore, CFD can be computationally expensive.

1.4 Objectives

The objective of the present work is to demonstrate a proof of concept optimization system aimed at maximizing the torque, hence, the efficiency of a VAWT for

a fixed solidity and tip speed ratio. To accomplish this, an appropriate model for predicting the performance of a VAWT is to be selected along with a robust optimization algorithm and flexible family of airfoil geometries.

Recent research has been conducted coupling models used for performance prediction with optimization algorithms. Bourguet [23] used CFD coupled with a design of experiments/response surface method approach, focusing on only symmetric blade profiles in two dimensions using a seven control point Bezier curve. Bourguet only simulated one blade with a low solidity as to avoid undesirable unsteady effects. He found that when there exists a possibility of several local optima, stochastic optimization algorithms are better suited for the job as they tend to be more efficient than gradient based algorithms. Jureczko [24] and Fuglsang [25] coupled low order performance prediction methods with optimization algorithms in a multi-criteria optimization routine. Both of their approaches focused on HAWTs rather than VAWTs. Research has also lead to patented blade designs using CFD coupled with optimization [26]. Rather than using optimization techniques, Saeed [27] was able to utilize inverse design methods to find an optimum design for a fixed tip speed ratio that satisfied the specified design performance characteristics. However, inverse design techniques require experience and intuition in order to specify desired performance, whereas optimization allows for designs to be generated that are more often than not beyond the intuition of a designer.

After reviewing the available models and recent research efforts, CFD was chosen as the appropriate tool for predicting the performance of a VAWT because of its flexibility and accuracy. Due to the possibility of local optima, and the requirement for floating-point optimization for geometric flexibility, a parallel stochastic differential evolution algorithm was chosen for the optimization. The NACA 4-series family of wing sections was chosen as the geometry to be parameterized for the optimization,

allowing either symmetric or cambered airfoil shapes to be generated. What separates this approach from all previous work is the consideration of both symmetric and cambered airfoil geometries, along with a full two-dimensional, unsteady simulation for a three-bladed wind turbine for various design points.

In this paper, the equations for predicting wind turbine performance will be introduced to provide insight into the iterative nature of the wind turbine design process. Next, the optimization methodology will be defined and the tools used to predict the VAWT performance and optimize the design will be described. The implementation of the design methodology will be introduced as the results of the optimization are analyzed and compared with a baseline geometry to quantify the performance gained in using this approach. Lastly, the conclusions and future recommendations for this work will be discussed.

CHAPTER 2

VERTICAL AXIS WIND TURBINE PERFORMANCE

2.1 Ideal Performance and the Betz Limit

Wind turbines are designed to extract the energy from the wind as efficiently as possible. Exactly how efficient? Early in the days of wind turbine design a man by the name of Albert Betz calculated this theoretical limit to be 59.3%, meaning that the maximum amount of kinetic energy available in the wind that a turbine can extract is only 59.3%. This is a pretty strong statement to make considering most turbomachinery exceed this limit. However, the Betz limit was derived under the assumption of an ideal wind turbine using simple one-dimensional momentum theory. An ideal wind turbine is one that is frictionless and produces no rotational component of velocity in the wake [28]. Realistically, wind turbines are not frictionless and always shed vorticity, therefore, more accurate methods of predicting wind turbine performance exist, such as those discussed in Chap. 1. The general equations used to predict the performance of wind turbines will be introduced in Sec. 2.4, but first the operating speed and geometry of the wind turbine must be defined.

2.2 Wind Speed and Tip Speed Ratio

According to the National Climatic Data Center, the average annual wind speed in the United States is approximately 4 m/s [29]. Realizing that the majority of wind turbines that have been developed to this day typically start producing power in winds as low as 3 m/s, the standard rated wind speed is still as high as 12 m/s. Determining the wind speed at which a wind turbine will operate is the most important step in

predicting its performance and even aids in defining the initial size of the wind turbine. Once the operating wind speed of the turbine has been decided upon, the first step in wind turbine design is to select a operating tip speed ratio [30] which can be expressed by

$$\lambda = \frac{\omega r}{V_{\infty}}, \quad (2.1)$$

or the ratio of the rotational velocity of the wind turbine ωr and the freestream velocity component (wind speed) V_{∞} . The selection of a λ not only depends on the wind speed, but also the generator available to convert the mechanical energy of the spinning wind turbine into electricity. Some generators are not able to operate for low values of ω .

2.3 Geometry Definition

Once λ has been chosen, the geometry of the VAWT can be defined through a dimensionless parameter known as the solidity

$$\sigma = \frac{Nc}{d}, \quad (2.2)$$

which is a function of the number of blades N , the chord length of the blades c , and the diameter of the rotor d . The solidity represents the fraction of the frontal swept area of the wind turbine that is covered by the blades. Deciding on a value for σ is not entirely random. Remember that the maximum theoretical power a wind turbine can extract from the wind is governed by the Betz limit of 59.3%. Realizing that the kinetic energy available in the wind for a given swept area is

$$P_W = \frac{1}{2} \rho d h V_{\infty}^3, \quad (2.3)$$

the maximum possible amount of power the wind turbine can generate is simply $0.59P_W$. Therefore, for a given V_{∞} and density ρ , the frontal area of the VAWT can

be approximated as dh , where d is the diameter of the rotor and h is the height of the wind turbine. This is a generous approximation based on an ideal wind turbine, so the final diameter or height of the wind turbine will be larger. A more conservative, realistic efficiency of 30%-40% could be used for initial sizing. Something to keep in mind when defining the geometry of the wind turbine is that P_W is directly related to the frontal area of the wind turbine. As a result, simply doubling either the diameter or height of the wind turbine, for instance, would lead to twice the available kinetic energy in the wind.

2.4 Performance Prediction

With λ chosen and the basic geometry of the VAWT defined, the next step is to predict the actual performance of the wind turbine. To do this, it is important to determine the forces acting on each blade. This is governed by the relative wind component W and the angle of attack α seen in Fig. 2.1. From introductory aerodynamics, an airfoil at an angle of attack to the relative wind produces lift L perpendicular to the oncoming wind. The drag D of the airfoil due to pressure and skin friction acts parallel to W . The resultant force F_R acting on the airfoil can be decomposed into a tangential force F_T and a normal force F_N . As was discussed previously in Sec. 1.2.2, F_T is the force that drives the wind turbine and is of interest when predicting the power output and efficiency.

2.4.1 Angle of Attack

As the rotor spins, the local azimuth angle for each individual blade changes, and with it, so does W and α . The angle of attack given by [5] can be expressed as

$$\alpha = \tan^{-1} \left[\frac{\sin \theta}{\lambda / (V_i / V_\infty) + \cos \theta} \right], \quad (2.4)$$




Figure 2.1. Force components for an airfoil.

where θ is the azimuth angle and V_i/V_∞ is the ratio of induced velocity to freestream velocity, a value that is unity if wake effects are neglected. However, because wake effects play a major role in calculating W , they are rarely neglected when determining α .

Close inspection of Eq. (2.4) reveals that α is governed by the tip speed ratio λ . Small tip speed ratios result in high angles of attack, whereas high tip speed ratios result in low angles of attack. Therefore, the value of λ is important as large changes in α lead to aerodynamic stall and can increase the magnitude of the cyclic loading on the wind turbine.

As a result of the changes in α associated with the decision in λ , one can conclude that thicker airfoil cross-sections would perform more reasonably for low values of λ , while thinner cross-sections would perform well for high values of λ . This is due to the fact that thicker sections can achieve a higher angle of attack before stall [31]. Additionally, assuming wake effects to be negligible, a reasonable assumption

for a low solidity rotor, the blades of a VAWT experience both positive and negative angles of attack of the same magnitude leading to the belief that symmetric airfoils should be used in these situations. However, as the solidity becomes high enough wake effects and vortex interaction cannot be neglected and cambered airfoil geometries may offer an increase in performance. It can also be understood that as all airfoils exhibit stall behavior at high angles of attack, cyclic loading can be reduced by simply increasing the operating tip speed ratio.

2.4.2 Average Torque

Once α has been determined, L and D can be found using empirical data or calculated using CFD. L and D are then non-dimensionalized by dividing through by the dynamic pressure to obtain the coefficient of lift C_l and the coefficient of drag C_d . These coefficients are used to calculate the tangential force coefficient

$$C_t = C_l \sin \alpha - C_d \cos \alpha. \quad (2.5)$$

To retrieve the actual tangential force, C_t is multiplied by the dynamic pressure

$$F_T = \frac{1}{2} C_t \rho c h W^2. \quad (2.6)$$

It is important to note that (2.6) represents the tangential force at only a single azimuthal position. Therefore, the process of determining α , C_t , and F_T must be repeated at all azimuthal locations before the torque can be calculated.

Because F_T is calculated at all azimuthal locations, it is said to be a function of θ and the average tangential force for a single rotation of one blade is

$$F_{Tavg} = \frac{1}{2\pi} \int_0^{2\pi} F_T(\theta) d\theta, \quad (2.7)$$

where the average torque for N blades located at radius r from the axis of rotation is given by

$$\tau = NF_{T_{avg}}r. \quad (2.8)$$

2.4.3 Power and Efficiency

The final step in predicting the performance of the wind turbine is determining the power it is able to extract from the wind and how efficiently it can accomplish that task. The amount of power the wind turbine is able to draw from the wind is given by

$$P_T = \tau\omega. \quad (2.9)$$

Therefore, the efficiency of the wind turbine is simply the ratio of the power produced by the wind turbine and the power available in the wind given by the expression

$$COP = \frac{P_T}{P_W} = \frac{\tau\omega}{\frac{1}{2}\rho dhV_\infty^3}. \quad (2.10)$$

Equation (2.10) is significant in this work because it represents a non-dimensional coefficient of performance (COP) that is a function of the torque to be used as the objective function for the aerodynamic shape optimization.

In Sec. 2.3 it was mentioned that increasing the frontal area of the wind turbine would increase the available wind energy for extraction. Therefore, the goal in designing a wind turbine is to do so in such a way to extract as much energy as possible. Analyzing Eq. (2.7) and Eq. (2.10) reveal that while an increase in the height of the wind turbine would increase F_T , hence τ , theoretically COP would remain unaffected. However, if an increase in P_T is all that is desired, the height of the wind turbine could be increased. In order to increase the efficiency of the wind turbine for a given λ , the blade shape and σ would have to be adjusted. Equation (2.7) is a function of C_t and c , where C_t is a function of the blade shape and c is a function of σ . Because

the shape of the blade, σ , and C_t are tightly coupled, it makes it difficult to select a geometry that maximizes the efficiency. Simplifying the problem by fixing both λ and σ reduces the problem to a single unknown, the airfoil cross-section. Therefore, it can be stated that for a fixed λ and σ there exists a single airfoil cross-section for which the torque, hence the efficiency of the wind turbine, would be maximized. Accomplishing this task is not straightforward, and requires an iterative approach.

For an iterative design process, if an initial geometry is unable to perform as well as desired, the geometry must be adjusted and reanalyzed. To achieve the optimum wind turbine that satisfies the given design constraints, the process of adjusting the geometry, whether for wind turbine size or blade shape, and reanalyzing the design may have to be executed numerous times before the wind turbine achieves desired performance. This iterative process, when conducted manually, is known as a parametric optimization. In this work we will eliminate the designers requirement to understand the wind turbine design sensitivities through numerical optimization. Therefore, it is up to the optimization algorithm to adjust the design parameters, releasing the designer from manually seeking out the design sensitivities. Obtaining the optimized design in this manner is the focus of this work and requires the implementation of a simple, automated optimization methodology.

CHAPTER 3

METHODOLOGY

3.1 Requirements

For the objectives of the current study to be feasible, the requirement for a straightforward, modular, and automated design framework became realized. Successfully taking a physical system, such as a wind turbine, and attempting to adjust, analyze, and optimize the design to satisfy an objective, or multiple objectives, required more than a few bundled files and programs requiring user input. In fact, it was this realization that sparked the idea of a simple, automated optimization methodology. The optimization methodology proposed is a unique and modular system aimed at relaxing the ties between the computer and operator and simplifying the design process so more time can be spent analyzing a solution and understanding the physics of the problem.

3.2 Unique Modular Design

A modular system is one in which entire parts of the system can be removed and replaced without compromising the process flows within the system. Therefore, for a module to be removed and replaced, it must be substituted with a module of equivalent functionality. Figure 3.1 illustrates the concept as it applies to wind turbine optimization and the methodology used in this work. The first step in the optimization process is generating the geometry. This geometry is described as a set of cartesian coordinates and is passed to the mesh generation module. This tool is used to discretize the fluid domain and output a specific file to the CFD solver module. This

module calculates a solution and passes information to the postprocessing module. The postprocessing tools manipulate the data, calculate the objective function value, and pass it to the optimizer. If the objective function is considered a maximum value, the optimization terminates. If not, the process starts over. Details pertaining to each of the modules used in this work will be discussed in the next chapter.

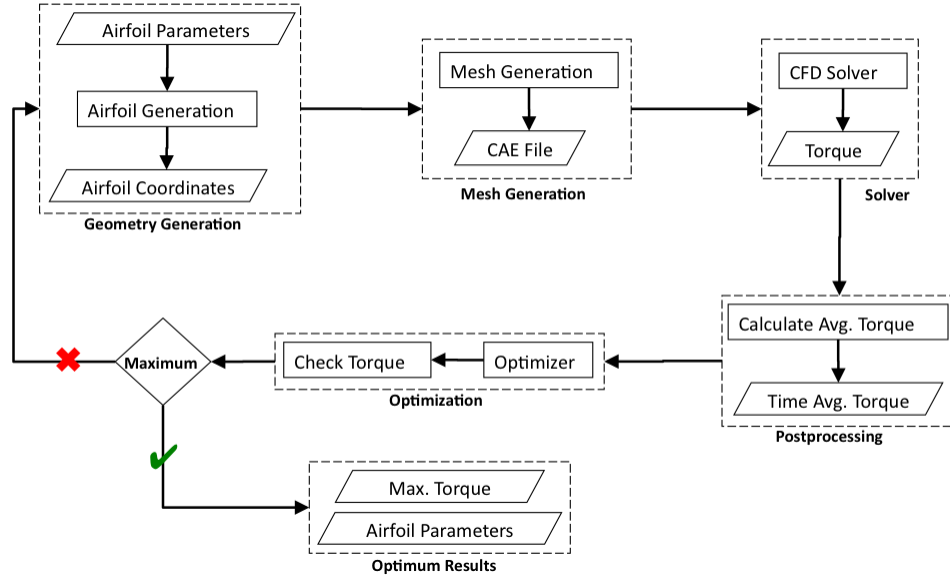


Figure 3.1. VAWT optimization methodology.

The general idea is that any of the modules can be replaced. For instance, maybe a user would like to use a different flow solver. They can easily remove the current solver module and replace it with the tool they are familiar with. For example, rather than using FLUENT, a user may prefer OpenFOAM. The user need only replace the FLUENT input files with input files for OpenFOAM to implement the new flow solver. Something that is unique to this modular design and optimization system is that no global file type or format is enforced. As it can be seen in the figure, each module

not only contains the tool, but the output of the tool. Therefore, it is up to the user as to how the modules communicate.

3.3 Directory Structure

One of the major requirements of the optimization system was to maintain a simple directory structure. Organization is not only important for allowing modules to communicate, but is also beneficial from a users point of view. File names are obvious and each one represents a different module. Therefore, all one would have to do in order to use another tool would be to simply place a new set of files within the modules directory. The directory structure for this work can be seen in Fig. 3.2 below.

```
OPTIMIZATION/  
|-- geometry  
|-- mesh  
|-- optimizer  
|-- postprocess  
|-- scripts  
`-- solver
```

Figure 3.2. Requirement for a simple directory structure.

The directory labeled *geometry* contains the executables and input files necessary to generate the geometry. The *mesh* folder contains the output from the geometry module and a script used to automatically generate the mesh used by the solver. The *solver* directory contains the final mesh that was generated and the scripts required to run the solver. Once the solver has finished running, the data gets passed to the *postprocess* directory where all the postprocessing takes place, including the calculation of the objective function value. This information gets sent to the

optimizer directory, where the optimization occurs. The *scripts* directory contains the tools used to automate the entire procedure.

3.4 Complete Automation

The optimization process requires that the entire system be automated. The current methodology allows for minimum user intervention as long as each individual module is automated. The *scripts* directory mentioned earlier contains scripts that automate each module. A script is a file that contains a list of commands required for executing programs and file manipulation and is interpreted by the operating system's command line interpreter. Scripting frees the user from manually typing a group of commands and reduces the need for human intervention throughout the entire process. In the end, only a single *run* script is written that calls all the scripts for each module, automating the entire process. The user need only provide the module and a single script in the *scripts* directory for automating that module.

3.5 Parametric Studies and Optimization

To test the feasibility of the proposed methodology, the idea for an automated modular design system with a simple directory structure was applied to the VAWT optimization study. The result was an automated system that allowed for parametric investigations to be conducted. This is beneficial in that multiple VAWT designs could be simulated and their performance analyzed, getting an idea of the design sensitivities, or what parameters have the largest effect on performance. The ability to carry out these types of studies is also important when conducting a grid sensitivity study, which will be discussed in more detail in Chap. 5.

Figure 3.3 illustrates the idea of the automated parametric study without the *optimizer* directory. It also goes to show how the automated parametric investigation can be placed into the final automated optimization process. This methodology accomplishes what it was originally intended to do. However, it not only provides a completely automated optimization solution, but a way to conduct trend studies, leading to a more flexible methodology then was originally conceived. The final directory structure for the optimization can be seen in Fig. 3.4. Each module is encapsulated in a black box with a top level description. The details of each module including their inputs and outputs will be discussed in Chap. 4.

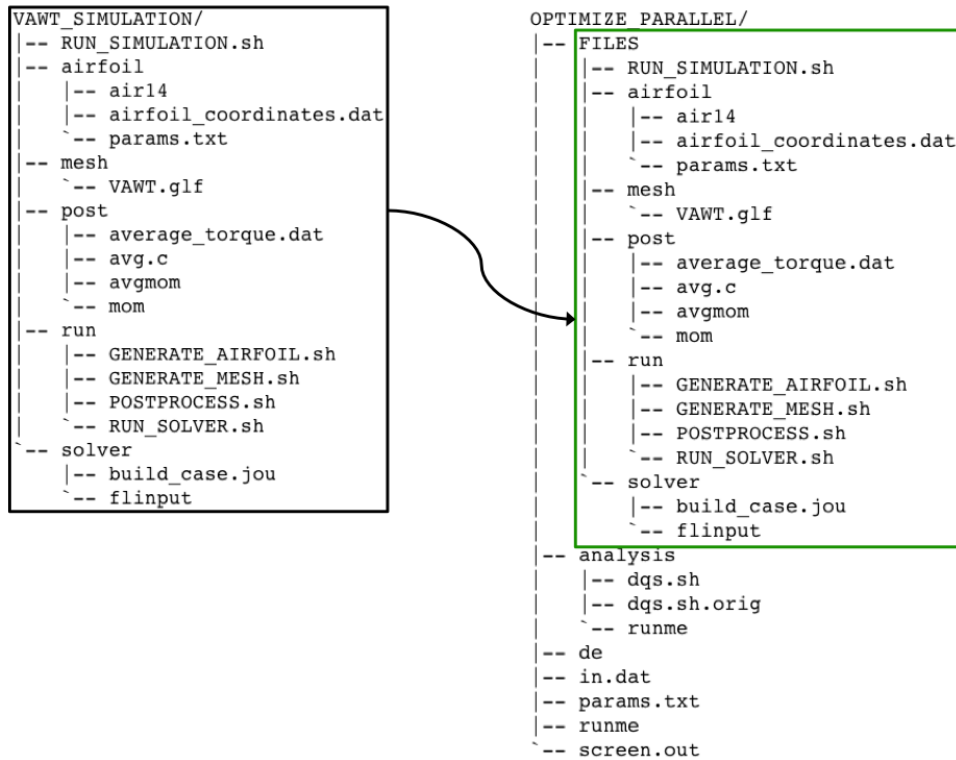


Figure 3.3. Directory structure for parametric study (left) and optimization (right).

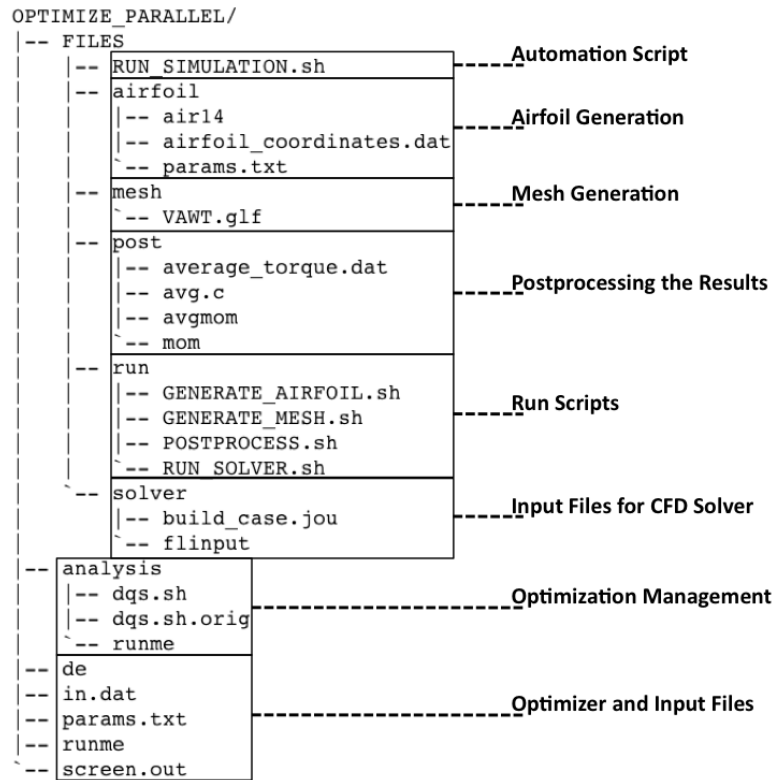


Figure 3.4. Optimization directory details.

CHAPTER 4

TOOLBOX

4.1 Geometry Generation

The objective of the optimization was to find an aerodynamic shape that for a fixed solidity and tip speed ratio would maximize the efficiency of the wind turbine. The first step was to select a suitable shape, or series of shapes, that could be adjusted throughout the optimization process. An obvious choice was the NACA 4-series airfoil. The majority of VAWTs utilize NACA airfoil sections because they are easy to manufacture and their characteristics are widely available.

4.1.1 NACA 4-Series Airfoils

The NACA 4-series airfoil family is a group of analytically designed airfoil sections expressed by a mean camber line and thickness distribution. This approach allowed airfoil designers to systematically test the effects of camber and thickness distribution. This work was published and has become one of the most successful wing section investigations of all time [31].

As it was mentioned earlier, the NACA 4-series airfoil sections are defined by a mean camber line and a thickness distribution. In Fig. 4.1, the mean camber line is the dashed line that splits the airfoil in half. The chord line is simply a straight line connecting the leading edge to the trailing edge of the airfoil whose length is defined as the chord length. The maximum thickness t is located at 30% of the chord for NACA 4-series airfoil sections. The maximum camber m , or maximum ordinate of the mean camber line, is located a distance p from the leading edge of the airfoil. The

values of m , p , and t are expressed as percentages of the chord length and represent the four digits defining the NACA 4-series airfoil. For example, the NACA 4424 shown in the figure has a maximum camber of 4% located at 40% of the chord with a maximum thickness of 24% located at 30% of the chord.




Figure 4.1. NACA 4424 4-series airfoil.

A NACA 4-series airfoil can be generated by the three parameters m , p , and t and are the parameters used in the optimization. A C-code was written to generate the shape of the airfoil based on the three parameters. Abbott and Doenhoff [31] introduced the equations used to define the shape of a NACA 4-series airfoil. The mean camber line of the airfoil was described as an analytically defined curve which was the combination of two parabolic arcs that are tangent at the point of maximum camber. For an x -coordinate, the ordinate of the mean camber line can be expressed as

$$y_c = \begin{cases} \frac{m}{p^2}(2px - x^2) & \text{forward of maximum ordinate} \\ \frac{m}{(1-p)^2}[(1-2p) + 2px - x^2] & \text{aft of maximum ordinate} \end{cases} \quad (4.1)$$

where m is the maximum camber and p is the chordwise location of the maximum camber. Once the camber line has been defined, the thickness distribution can be found by the following equation

$$\pm y_t = \frac{t}{0.20}(0.29690\sqrt{x} - 0.12600x - 0.35160x^2 + 0.28430x^3 - 0.10150x^4) \quad (4.2)$$

where t is the maximum thickness of the airfoil located at 30% of the chord. After the camber and thickness distribution have been defined for various x locations (typically ranging from 0-1), the coordinates of the upper airfoil surface can be found by

$$\begin{aligned}x_U &= x - y_t \sin \theta, \\y_U &= y_c + y_t \cos \theta\end{aligned}\tag{4.3}$$

while the coordinates of the lower surface are defined by

$$\begin{aligned}x_L &= x + y_t \sin \theta, \\y_L &= y_c - y_t \cos \theta\end{aligned}\tag{4.4}$$

and where θ is defined as

$$\theta = \arctan \frac{dy_c}{dx}.\tag{4.5}$$

4.1.2 Airfoil Constraints

For the purposes of maintaining high cell quality during the grid generation process and a converged CFD solution, constraints were placed on the parameters defining the airfoils that could be generated for the optimization. The lower limit of the constrained geometry was the NACA 0004, an extremely thin, symmetric airfoil section. The upper bound was the NACA 9940, an airfoil with a high thickness and large aft camber. The lower and upper limit of the the constrained airfoil geometries is shown in Fig. 4.2. The three parameters defining the airfoils had to be normalized from 0 to 1 for the optimization. This lead to the NACA 0004 being generated using $m = 0$, $p = 0$, and $t = 0$ and the NACA 9940 begin generated using $m = 1$, $p = 1$, and $t = 1$. Placing constraints on the airfoils to be generated lead to a possible 3700 NACA 4-series designated design options. However, since the optimization algorithm chosen

used floating-point values to define the parameters, essentially an infinite number of airfoil designs were attainable.



Figure 4.2. Range of airfoils: NACA 0004 (top) and NACA 9940 (bottom).

4.2 Grid Generation

After the airfoil geometry for the VAWT had been defined, the next step was to discretize the computational domain as a preprocessing step in the CFD process. The act of discretizing the domain is termed grid generation and is one of the most important steps in the CFD process. For simple geometries where the direction of the flow is known beforehand, creating the grid is usually straightforward. For flows such as this, high quality structured grids can be used that can accurately capture the flow physics. However, as geometry becomes complex and the flows more difficult to predict with the onset of turbulence and separation, grid generation is no longer a trivial task. For flows such as this, unstructured grids consisting of triangles and tetrahedra provide increased flexibility and are often used. Although, it should be mentioned that using unstructured grid generation techniques tend to introduce more

error into the solution than if structured grids are used due to the nature of the discretization techniques. Before moving on, it is important to discuss the different cell types used in grid generation.

4.2.1 Cell Types

The complexity of the geometry or the simulation tends to dictate the type of cells to be used and overall topology of the grid. Structured grids consist of quadrilaterals in 2D and hexahedra in 3D. These types of grids are typically used when the geometry is relatively simple and the flow path is known beforehand. For complex geometries, unstructured grid generation techniques are preferred because of their flexibility and automation and have been used extensively for aerodynamic design and optimization [32],[33],[34],[35],[36]. Unstructured grids consist of triangles in 2D and tetrahedra in 3D. Because the current work involves the optimization of a 2D section cut of a VAWT, only 2D grids will be used in this work.

4.2.2 Grid Considerations

When the flow is viscous and turbulent, a large portion of the grid generation process is spent focusing on accurately resolving the boundary layer. In 2D, the different cell types typically encountered in the boundary layer region can be seen in Fig. 4.3. The first type is a structured quadrilateral cell. For simple geometries, such as an airfoil, layers of high aspect ratio quadrilaterals can be marched off the geometry into the farfield. The disadvantage of using this approach is that in order to maintain a structured topology in the boundary layer, as soon as a single cell violates a specified quality criteria, the entire front must stop [37]. For complex geometries, this may happen sooner than desired. The second cell type is the anisotropic triangle. In 2D it is essentially a diagonalized quadrilateral. These cell types are commonly used

in the generation of boundary layers for complex geometries. Unlike the structured technique, as anisotropic triangles are marched off the surface of the geometry, if a cell violates a user specified quality constraint or becomes ideal, the boundary layer extrusion will stop locally while allowing the front to continue propagating elsewhere. The last cell type is an isotropic triangle. In order to maintain isotropy, a large amount of cells would be required in the boundary layer region, therefore, isotropic triangles are typically not desired when trying to maintain a low cell count.




Figure 4.3. Different cell types used for boundary layer grids.

Due to the number of grid types that can be used to discretize a computational domain, it is important to exercise the capabilities of the solver to determine its sensitivity to the varying cell types. A good rule of thumb proposed by Baker [2] can be seen in Fig. 4.4. From the figure, a multiblock structured grid is said to provide the highest level of viscous accuracy; yet, it also suggests that a hybrid grid topology would provide a balanced level of accuracy and automation, an important characteristic for the optimization process.

Before a final decision was to be made on the type of grid used for the VAWT simulation, a comprehensive grid dependency study was conducted. The goal of a grid dependency study is find a grid independent solution. To conduct a grid dependency study, a family of grids must be generated, each with higher global and local resolutions than the last. At some point, when the solution stops changing, the grid is said to have converged. For this work, a family of multiblock structured grids

were created alongside a family of equivalent hybrid grids. The torque was calculated for each of the grids, and a grid independent solution was found. Details regarding the grid generation and dependency studies will be discussed further in Chap. 5.

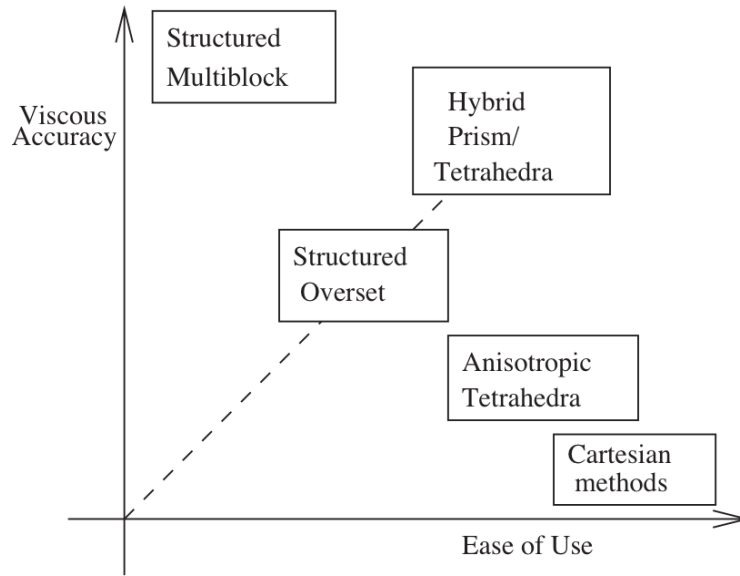


Figure 4.4. Grid type accuracy vs. ease of use [2].

As a result of the grid dependency study and the extensive work done in the areas of aerodynamic design and optimization using unstructured grid generation techniques, a hybrid grid was chosen as the most appropriate for this work. The hybrid grid consists of a structured boundary layer transitioning to isotropic tetrahedra in the farfield and can be seen for a single VAWT blade in Fig. 4.5 and Fig. 4.6. This choice provided a flexible and completely automated approach to the grid generation of numerous VAWT geometries.




Figure 4.5. Hybrid grid for an airfoil.




Figure 4.6. Leading edge boundary layer grid for blade geometry.

4.2.3 Pointwise

Pointwise V16.02 was used to generate the grids used for the VAWT study. For the grid dependency study, structured grids were constructed manually, while automated grid generation techniques were used to construct the hybrid grids. The farfield was split into a rotating and non-rotating zone and was discretized using either structured or unstructured elements. The scripting capabilities of Pointwise were exercised while constructing the hybrid grids and allowed for the grid generation of the

airfoil geometries to easily be integrated into the optimization process. A Pointwise script was written in Glyph2, based on Tcl, that imported the airfoil geometry as a list of x, y coordinates and constructed the hybrid grid based on several user defined parameters such as the initial cell height in the boundary layer and the solidity of the wind turbine. The automated script also set up the appropriate boundary conditions and exported the file for the solver.

4.3 Solver

Once the computational domain had been discretized, the equations governing fluid flow were solved using an appropriate discretization technique in order to calculate the torque. The most common technique used is called the finite volume method. The finite volume method solves the integral form of the conservation equations, specifically the conservation of mass and momentum. Approximations to the integrals are achieved using appropriate quadrature formula where the integrand containing the dependent variables, namely the convective and diffusive fluxes, are interpolated from surrounding discrete control volumes using various schemes. Algebraic equations are constructed as an outcome of the discretization that represent approximations of the original governing integral equations. The algebraic system is then linearized and dependent variables are solved for using iterative methods. At this point it is important to introduce the governing conservation equations for fluid flow to be discretized and solved.

4.3.1 Conservation Equations

If ϕ is considered to be an intensive property like mass or velocity, then the general conservation equation for a control volume can be expressed by

$$\frac{d}{dt} \int_{CM} \rho \phi dV = \frac{d}{dt} \int_{CV} \rho \phi dV + \int_{CV} \rho \phi (\vec{v} - \vec{v}_b) \cdot \vec{n} dS, \quad (4.6)$$

which states that the time rate of change of the intensive variable ϕ for a control mass is equal to the rate of change of the variable inside the control volume plus the net flux of the variable across the surfaces of the control volume. This general form of the conservation of some intensive variable ϕ is known as Reynold's transport theorem, the derivation of which can be found in most introductory fluid mechanics textbooks [38], [39].

By allowing $\phi = 1$, the conservation of mass is defined as

$$\frac{\partial}{\partial t} \int_V \rho dV + \int_S \rho \vec{v} \cdot \vec{n} dS = 0, \quad (4.7)$$

while allowing $\phi = \vec{v}$, the conservation of momentum is defined by

$$\frac{\partial}{\partial t} \int_V \rho \vec{v} dV + \int_S \rho \vec{v} \vec{v} \cdot \vec{n} dS = \int_S \vec{T} \cdot \vec{n} dS + \int_V \rho \vec{b} dV \quad (4.8)$$

where \vec{T} represents the stress tensor responsible for the surface forces acting on the control surfaces of the volume, and \vec{b} represents the body forces acting on the volume as a whole. If Gauss's theorem is used to convert the volume integrals into surface integrals,

$$\int_V (\nabla \cdot \vec{F}) dV = \int_S \vec{F} \cdot \vec{n} dS, \quad (4.9)$$

which states that the divergence of some vector $\nabla \cdot \vec{F}$ is equal to the flux of that vector across the surfaces of the volume, then the conservation of mass for an infinitesimally small control volume becomes

$$\nabla \cdot \vec{v} = 0, \quad (4.10)$$

which states that the volume dilatation is equal to zero. Using the same approach, the conservation of momentum becomes

$$\frac{\partial \vec{v}}{\partial t} + \nabla \cdot (\vec{v}\vec{v}) = -\frac{1}{\rho}\nabla p + \nabla \cdot (\nu\nabla\vec{v}) + \vec{b}, \quad (4.11)$$

where the terms on the left hand side represent the local and convective accelerations of the fluid which are equal to the forces on the fluid due to the pressure gradient, viscous diffusion, and body forces. Equations 4.10 and 4.11 represent the conservation equations for incompressible flow.

4.3.2 Turbulence Modeling

In the case of simulating the interaction of fluid with a VAWT, the calculation becomes rather complex. The rotation of the blades leads to time dependent, turbulent phenomena as the blades stall and vortices are shed into the wake.

Generally speaking, turbulence is characterized by the random motion of a fluid. The Reynolds number, or ratio of inertial to viscous effects, is often used to classify a flow as either laminar or turbulent. For high Reynolds number flows inertial effects become dominant leading to three dimensional velocity fluctuations in the flow field. These fluctuations occur for a wide range of both time and length scales. One way to resolve turbulent flows is to discretize the spatial and temporal domain as to capture even the smallest turbulent structure. However, this direct numerical simulation (DNS) as it is often termed, is computationally expensive. Therefore, other numerical solutions have evolved to approximate the nature of turbulent flows, one of which is known as the Reynolds-averaged approach.

Reynolds averaging begins by taking a variable and expressing it as the sum of the mean value of that variable over time and the fluctuation of that variable about the mean value, or

$$\phi = \bar{\phi} + \phi'. \quad (4.12)$$

Writing all the variables as the sum of the time averaged value $\bar{\phi}$ and fluctuating value ϕ' and applying this technique to the Navier-Stokes equations yields the Reynolds-averaged Navier-Stokes equations (RANS) [40]. Reynolds averaging of the Navier-Stokes equations leads to a new stress term due to the fluctuations in the velocity field often referred to as the Reynolds stress term $-\overline{u'_i u'_j}$. To provide closure to the RANS equations, several turbulence models have been developed to approximate the Reynolds stress term, one of which is known as the Spalart-Allmaras (SA) turbulence model [41].

The SA turbulence model is a one-equation turbulence model that approximates the Reynolds stresses as

$$-\overline{u'_i u'_j} = \nu_t \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right), \quad (4.13)$$

dropping the turbulent kinetic energy term k most often solved for in a two-equation turbulence model. Equation 4.13 is the product of the eddy viscosity ν_t and the strain rate tensor, leading to the argument that turbulence is found where vorticity is present. The SA model solves a transport equation for ν_t given by [42]

$$\frac{\partial}{\partial t}(\rho \tilde{\nu}) + \frac{\partial}{\partial x_i}(\rho \tilde{\nu} u_i) = G_\nu + \frac{1}{\sigma_{\tilde{\nu}}} \left[\frac{\partial}{\partial x_j} \left\{ (\mu + \rho \tilde{\nu}) \frac{\partial \tilde{\nu}}{\partial x_j} \right\} + C_{b2} \rho \left(\frac{\partial \tilde{\nu}}{\partial x_j} \right)^2 \right] - Y_\nu + S_{\tilde{\nu}} \quad (4.14)$$

where $\tilde{\nu}$ is the modified turbulent viscosity, G_ν represents the production of turbulent viscosity, and Y_ν is the destruction of turbulent viscosity. Unlike the two-equation turbulence models where often difficult to judge approximations of the turbulence intensity and length scale must be provided, the boundary conditions for the SA

turbulence model are trivial. Typically, $\tilde{\nu} = 0$ is prescribed for a wall boundary and $\tilde{\nu} < \frac{\nu}{10}$ in the freestream.

4.3.3 FLUENT

The commercial solver FLUENT v6.3 was used for this work. FLUENT uses the finite volume method to discretize the integral form of the governing equations. For the simulation, a pressure based segregated solver was chosen where the SIMPLE algorithm was used to handle the pressure-velocity coupling that exists. The segregated solver, while slower than a coupled algorithm, is more memory efficient as it solves the governing equations sequentially for the variables of interest. A 2^{nd} order interpolation scheme for pressure was used along with a 2^{nd} order upwind discretization scheme for the momentum equation and modified turbulent viscosity. The 2^{nd} order upwind scheme relies on achieving 2^{nd} order accuracy through a Taylor series expansion of the solution at a cell center about the cell centroid, requiring not only the value of the variable, but also the gradient. The gradients required for the discretization of the convective and diffusive fluxes were computed using a cell-based approach where the face value of the variable in question is obtained by averaging the values at the neighboring cell centers [42].

Turbulence modeling was accomplished through the use of the Spalart-Allmaras one-equation turbulence model described in the previous section. The y^+ for the blades varied at different azimuthal locations, but consistently placed the first cell centroid of the wall-adjacent cells inside the viscous sublayer ($y^+ < 5$) of the boundary layer. Therefore, because the grid was fine enough to resolve the viscous sublayer, the laminar stress-strain relationship $u^+ = y^+$ was used to determine the wall shear stress.

Because the simulation was time dependent, a 2^{nd} order implicit time integration was chosen for the temporal discretization. Implicit methods are often preferred over explicit methods due to their unconditional stability, whereas explicit methods are conditionally stable. However, the use of an implicit method does not imply that a large time step should be taken. Therefore, a time step was chosen small enough to reduce the number of iterations per time step and to properly model the transient phenomena.

Due to the structure of the matrices resulting from the discretization and linearization of the governing integral equations, special care must be taken when solving the system. FLUENT uses an algebraic multigrid (AMG) method coupled with a point implicit Gauss-Seidel solver. The basic idea behind an AMG method is to project the low frequency errors onto a coarser mesh turning them into high frequency errors. Because the Gauss-Seidel method is solved point by point, it is able to reduce local high frequency error rather efficiently. So in moving to a coarser grid (restriction), the newly obtained high frequency errors can be reduced using the Gauss-Seidel solver as a smoother. The solution can then be projected back onto a fine mesh (prolongation) and eventually the original mesh where hopefully the global low frequency error has been minimized [40].

The interior domain containing the wind turbine blades was considered as a moving mesh, while the outer domain was stationary. The interior sliding domain rotated with a given rotational velocity for a specified λ . The inlet to the computational domain was defined as a velocity inlet with a uniform velocity component. Typically, the rotation velocity of the sliding domain was kept constant while the inlet velocity was adjusted to perturb λ and build a performance envelope for various blade shapes. The modified turbulent viscosity $\tilde{\nu}$ at the inlet was chosen to be equal to 5ν , where ν is the molecular kinematic viscosity of air. It was found through comprehensive test-

ing that adjustments in $\tilde{\nu}$ did not tremendously affect the results of the simulation. The outlet was marked as a pressure outlet with the gauge pressure set to zero.

The simulation was set to automatically report the moment due to pressure and viscous forces about the axis of rotation every 15 time steps. This data was written to a file and then passed to the postprocessing module in order to compute the average torque for the optimizer.

4.3.4 Computational Considerations

Due to the size and unsteady nature of the problem, the Linux cluster named Grendel located in the CFD Lab at The University of Texas at Arlington was used to run the simulations. Grendel consists of 9 nodes, each of which run CentOS 5.5. The 9 computational nodes consist of 7 2.83 GHz Intel Core2Quads and 2 2.8 GHz Pentium 4 processors. Only the quad core processors were used during the simulations. Each simulation was run on its own dedicated quad core processor with 4 Gb of system memory. Overall average computation time to achieve a quasi-steady state took approximately 2.5 hours, 4 times faster than on a single processing core.

4.4 Postprocessing

Once the solution had been calculated using FLUENT and all relevant data had been written to a file, the average torque could then be determined. A small script parsed through the output file and saved only the torque values that were recorded every 15 time steps. This file then contained torque as a function of time. A graph of the torque vs. time can be seen in Fig. 4.7 below. At a certain point, the flow became quasi-steady, and the oscillations were more uniform. One single rotation of the wind turbine has been outlined in the figure. The three peaks in the torque represent the times at which each blade passed around the front of the wind turbine,

while the three valleys represent the times at which each blade moved around the back of the wind turbine. Therefore, more blades would result in a higher frequency oscillation for the same rotation speed. In order to calculate a single scalar value of the torque for the optimizer, the oscillating torque was averaged.




Figure 4.7. Variation of torque as a function of time.

4.4.1 Average Torque

In Chap. 2, the tangential force component driving the wind turbine, and also used to compute the torque, was a function of azimuthal location. During the simulation process, the torque was recorded as a function of time, therefore, it is important to introduce the average value of a function $f(t)$ over the interval $[a, b]$ as

$$f_{avg} = \frac{1}{b-a} \int_a^b f(t) dt, \quad (4.15)$$

where a would represent the time at the beginning of a single rotation and b is the time at the end of a single rotation. This equation states that the average value of the function $f(t)$ is equal to the integral of that function for a single rotation divided

by the time required to complete a single rotation. Using the trapezoidal rule, the definite integral can be expressed by

$$\int_a^b f(t)dt = \frac{b-a}{2n} \left[f(t_o) + 2 \sum_{i=1}^{n-1} f(t_i) + f(t_n) \right], \quad (4.16)$$

where n is the number of segments used to split the interval of integration. Using Eqns. 4.15 and 4.16 and replacing n by $\frac{t_n-t_o}{\Delta t}$, the average value of the torque for a single rotation of the wind turbine is defined as

$$\tau_{avg} = \frac{\Delta t}{2(t_n - t_o)} \left[\tau(t_o) + 2 \sum_{i=1}^{n-1} \tau(t_i) + \tau(t_n) \right], \quad (4.17)$$

where t_o is the time at the beginning of a rotation, t_n is the time at the end of the rotation, and Δt is the time step used when recording the torque in FLUENT.

A C-code was written to take the file containing the torque as a function of time and calculate the average torque for the final rotation of the wind turbine using Eq. 4.17. This value was then passed on to the optimization module as the objective function value to be maximized.

4.5 Optimization

In order to maximize the average torque of the wind turbine given the NACA 4-series airfoil design parameters (Sec. 4.1) and the solidity and tip speed ratio design constraints, a simple and robust optimization algorithm was required. This act of searching for the minimum or maximum value of a function while varying the parameters, or values of that function, and incorporating any constraints is called optimization. In optimization, the function is often termed the objective function or cost function, and it is the goal of the optimization algorithm to find the true minimum or maximum of that objective function as efficiently as possible. However, in design the objective function may be a rather complex, nonlinear, or non-differentiable

function that is under the influence of many parameters and design constraints. This possibility rules out any simple, gradient-based optimization algorithms such as the method of steepest descent or Newton’s method, as these algorithms require the objective function to be differentiable and are only efficient at finding local minimum or maximum values. Therefore, global optimization algorithms are favored for design optimization.

Evolutionary algorithms (EA) are global, direct search algorithms that are able to overcome the limitations of local, gradient-based methods. Two the of the most popular EAs are the genetic algorithm (GA) and differential evolution (DE). The underlying idea behind EAs and other conventional direct search methods is to generate populations of design parameters that are thought of as vectors. Variations in the design parameters are created and it is decided whether to accept or reject the new parameter vector. The majority of direct search algorithms use the greedy search method where the new parameter vector is accepted, in the case of a minimization problem, if the objective function value has been reduced. However, greedy search methods tend to get trapped by local minimum and require special techniques such as annealing to help a vector climb away from local minimum. This disadvantage, while common for most EAs and undesirable for design optimization, is eliminated by the DE algorithm as proposed by Storn and Price [43], [44], [45].

4.5.1 Differential Evolution

The DE algorithm is a stochastic direct search method aimed at minimizing an objective function based on constraints that are represented by floating-point values rather than binary strings like most EAs. The DE algorithm is robust, fast, simple, and easy to use as it requires very little user input. These traits lead to the choice of DE as the algorithm used in the current study. However, for this work the focus

was on maximizing the average torque while allowing the airfoil geometry to change, therefore, a change in sign was all that was required to adjust the DE algorithm from a minimization to a maximization method.

In order to determine the maximum value of the objective function, the DE algorithm starts with an randomly populated initial generation of NP D-dimensional parameter vectors, where NP is the number of parents in a population and D is the number of parameters. For this work, $D = 3$ and $NP = 14$. Each member of the population is perturbed, generating a new vector through a mutation operation. The new vector's parameters are then mixed with the parameters of another vector called the target vector, to create a trial vector through an operation known as crossover. In the last step of the optimization scheme, the trial vector is compared with the target vector and if it so happens that the trial vector outperforms the target vector, it replaces the target vector in the next generation. However, if the target vector is found to yield a higher objective function value, it is retained for the next generation. The optimization is said to have converged when either the maximum objective function value is found or the maximum number of user specified generations is reached, which for this study was 11 generations.

4.5.2 Mutation

After initializing the population, each target vector $\vec{x}_{i,G}$ in that generation undergoes a mutation operation given by

$$\vec{v}_{i,G+1} = \vec{x}_{r1,G} + F(\vec{x}_{r2,G} - \vec{x}_{r3,G}), \quad (4.18)$$

where the index r represents a random population member in the current generation, F is a scaling factor $\in [0, 2]$ dictating the amplification of the difference vector $(\vec{x}_{r2,G} - \vec{x}_{r3,G})$, and the result $\vec{v}_{i,G+1}$ is called the mutant vector. A scaling factor $F = 0.8$ was

selected for this work. This mutant operation is a characteristic of a variant of DE that utilizes a single difference operation, therefore, $NP \geq 4$ such that the index i is different than the randomly chosen values of $r1, r2, r3$. This operation can be seen graphically in Fig. 4.8.




Figure 4.8. Illustration showing the generation of the mutant vector \vec{v} .

Other variants of DE exist that utilize more difference operations to determine the mutant vector. The DE strategy used in this work *DE/best/2/exp* utilizes two difference vectors. The idea is that by using two difference vectors the diversity of large populations can be improved, increasing the possibility that members of a population span the entire solution space and reducing the risk of premature convergence. The mutation operation for the DE variant used in this work is given by

$$\vec{v}_{i,G+1} = \vec{x}_{best,G} + F(\vec{x}_{r1,G} + \vec{x}_{r2,G} - \vec{x}_{r3,G} - \vec{x}_{r4,G}), \quad (4.19)$$

where $\vec{x}_{best,G}$ represents the best performing parameter vector from the current population. This is different than the previous strategy that utilized a random population

member to perform the mutation operation. The hope is that by using the best parameter vector in the population, the number of generations required for convergence will decrease.

4.5.3 Crossover

The crossover operation generates a trial vector by selecting pieces of the target vector and mutant vector. The trial vector is determined by

$$\vec{u}_{ji,G+1} = \begin{cases} \vec{v}_{ji,G+1} & \text{if } (\text{randb}(j) \leq \text{CR}) \text{ or } j=\text{rnbr}(i) \\ \vec{x}_{ji,G} & \text{if } (\text{randb}(j) > \text{CR}) \text{ and } j \neq \text{rnbr}(i) \end{cases} \quad (4.20)$$

where CR is the crossover constant, or crossover probability $\in [0, 1]$, and $\text{randb}(j)$ is a randomly chosen number $\in [0, 1]$ evaluated during the j^{th} evaluation where $j=1,2,3,\dots,D$. If the value of $\text{randb}(j)$ happens to be less than or equal to CR, the trial vector gets populated with a parameter from the mutant vector. However, if the random number that has been generated happens to be larger than CR, the trial vector gets a parameter from the target vector. To ensure that at least one parameter value is chosen from the mutant vector, a random value $\in 1, 2, 3, \dots, D$ is chosen as $\text{rnbr}(i)$. If $\text{CR}=1$, all trial vector parameters will come from the mutant vector. This illustrates that the choice in CR works to control the crossover probability. In this work, $\text{CR} = 0.6$.

4.5.4 Selection

The last step in the DE algorithm is selection. Once the trial vector has been formed, it must be decided whether or not it should move to the next generation. Therefore, in the selection process, if the trial vector performs better than the target vector, resulting in a larger objective function value, the trial vector moves on to the next generation. However, if the newly generated trial vector is outperformed by the

original target vector, the target vector remains a population member in the next generation.

In this work, the DE code generated new airfoil parameters each generation through mutation, crossover, and selection operations. Each new set of parameters was used to generate the airfoils of the VAWT for which the torque could then be calculated using the FLUENT solver. The torque was then averaged by the postprocessing module and used as the objective function value driving the DE algorithm.

CHAPTER 5

RESULTS

The overall objective of the work was to successfully demonstrate a proof of concept optimization system capable of maximizing the efficiency of a three-bladed VAWT for a fixed solidity and tip speed ratio. Two test cases were conducted to demonstrate the robustness of the optimization system for a wide range of design constraints and are defined in Table 5.1. Before the final results of the optimization are presented, a detailed overview of the grid dependency studies will be introduced. Next, the performance of a baseline geometry will be presented as well as experimental VAWT performance trends for comparison. Finally, the results of the two optimization test cases will be introduced and compared with the performance of the baseline geometry.

Table 5.1. Optimization Test Cases

Case	σ	λ	V_∞ (m/s)	ω (rad/s)
1	1.5	1	10	10
2	0.4	3	4	12

5.1 Grid Dependency Studies

The goal of a grid dependency study is to find a grid independent solution. When conducting a new study, it is highly unlikely that the first grid generated provides the necessary resolution to adequately capture the flow physics, therefore,

typically a family of grids is created and compared. The hope is that as the number of cells in the grids increase, there will be a point at which the solution will not change, and the solution is said to be grid independent. For this work, a family of structured and equivalent hybrid grids were created in hopes to find a grid that provided adequate resolution of the unsteady phenomena while the construction of the grid would remain highly automated for the optimization process.

To begin, a simple blade shape was used for the grid study. The VAWT consisted of three 60 degree semi-circular blades with a constant thickness of 0.025 m giving the rotor a solidity $\sigma = 1.5$. Each blade was separated by 60 degrees at a radius of 1 m from the axis of rotation, providing a simple geometry with which to define the initial topology. Because the blades had to spin in the simulation, an interior sliding domain was to be constructed with a radius of 25 m. The outer stationary domain, and the extent of the farfield, was defined to be 50 m. The farfield domain was large enough such that the unsteady flow characteristics would develop and dissipate inside the domain, eliminating the concern for reverse flow. The construction of the structured and hybrid grid topologies chosen for the grid study will be discussed below.

5.1.1 Structured Grid Topology

Due to the nature of structured grids, the topology must be well defined before construction begins. A structured grid consists of quadrilateral elements, therefore, opposing grid lines must contain the same number of points to construct a domain consisting of purely quadrilateral cells. The final topology of the structured grid can be seen in Fig. 5.1. In the figure, the blades are outlined in red, while the grid lines are highlighted in blue. The topology chosen consisted of a square domain inside the rotor. The square was then used to construct the 5 structured interior domains as

shown in the figure. Grid boundary lines were allowed to float to increase the quality of the cells in the mesh.




Figure 5.1. Structured grid topology.

The quality measure for all grids in the study was the maximum included angle α_{in} . This is the maximum interior angle for the quadrilateral elements, or the triangular elements for the hybrid grids. While the optimum interior angle for quads is 90 degrees and 60 degrees for triangles, best practices suggest maintaining a maximum $\alpha_{in} < 170$ degrees.

A characteristic, and even a disadvantage of the structured grid topology, is that the local blade resolution used to resolve the boundary layer is propagated into the farfield. It can be seen that as the leading and trailing edge of the blades in Fig. 5.1 are resolved locally, the grid spacing and the number of grid points propagates

both into the interior and into the farfield. This tends to lead to larger cell counts when using structured grids. Figure 5.2 shows the farfield grid for the structured topology with the inner rotational domain and outer stationary domain. The number of grid points in the circumferential direction was governed by the number of points used to resolve the blade geometry.




Figure 5.2. Structured farfield grid.

Once the structured topology had been defined, a family of grids was constructed for the grid dependency study. Table 5.2 below provides a summary of the structured grid family. The cell count for each grid is provided, along with the initial spacing in the boundary layer Δs_{init} , the growth rate of the cells away from the

blades GR , the number of points used to resolve the blade geometry $Points_{af}$, and the maximum included angle α_{in} .

Table 5.2. Structured Grid Family

Grid	Cell Count	Δs_{init} (m)	GR	$Points_{af}$	Max α_{in} (deg)
1	50,000	0.0025	1.2	129	130
2	100,000	0.001	1.1	177	122
3	400,000	0.0005	1.1	355	122

5.1.2 Hybrid Grid Topology

The hybrid grid topology used for this study consisted of a structured boundary layer transitioning to unstructured triangles. Unlike the structured grid, the hybrid grid topology was easy to construct and automate. The final hybrid grid topology can be seen in Fig. 5.3. The boundary layer was constructed using a normal hyperbolic extrusion technique in Pointwise [46]. The user need only specify the initial cell height, growth rate, and the number of layers for the extrusion.

A benefit from using the hybrid topology was that appropriate boundary layer resolution was obtained while maintaining a low cell count in the farfield. The farfield grid for the hybrid topology can be seen in Fig. 5.4. Unlike the structured grids, the farfield and boundary layer were almost entirely decoupled, resulting in a much lower cell count.

Table 5.3 provides a summary of the hybrid grid family. Each of the hybrid grids was constructed to be equivalent to the structured grids in terms of local blade resolution. While α_{in} was higher for the hybrid grids, it still remained below the acceptable limit.




Figure 5.3. Hybrid grid topology.




Figure 5.4. Hybrid farfield grid.

Table 5.3. Hybrid Grid Family

Grid	Cell Count	Δs_{init} (m)	GR	$Points_{af}$	Max α_{in} (deg)
1	50,000	0.0025	1.2	129	140
2	55,000	0.001	1.1	177	136
3	150,000	0.0005	1.1	355	135

A comparison of the structured and hybrid grid families can be seen in Fig. 5.5. The local blade resolution was preserved using the normal hyperbolic extrusion technique during the construction of the hybrid grids. As was mentioned earlier, the resolution and spacing constraints placed on the structured grid near the blade can be seen propagating away from the blade itself, whereas for the hybrid case there is a clear boundary that separates the boundary layer grid from the rest of the farfield. For instance, while the coarse grids contain 50,000 cells for the same local blade resolution, once the spacing was adjusted to achieve higher resolution, the structured grid contained 100,000 cells as opposed to the hybrid grid containing only 55,000 cells. In this case, the local blade resolution that was enforced for the structured grid propagated into the farfield. However, while the local blade resolution changed for the hybrid grid, the farfield remained unaffected.

5.1.3 Grid Independent Solution

The torque was calculated for each grid using the FLUENT solver, the settings of which were discussed in Sec. 4.3.3. A time step of $\Delta t = \frac{2\pi}{\omega N}$ was used where ω is the rotation rate, and N is the number of cells in the circumferential direction. This was calculated for the 100,000 cell structured mesh with $\omega = 10$ rad/s to be approximately 0.001 s, the time step that was used for all simulations. This represents the amount




Figure 5.5. Structured and hybrid grid family comparison.

of time for the sliding domain to move one grid point in the circumferential direction and was found adequate for the simulation.

Convergence was monitored by observing the residuals as well as the torque. In the ideal case, the residuals should converge to true zero. However, a more relaxed convergence criteria of $1e-5$ was enforced for continuity, momentum, and modified turbulent viscosity. The residuals were monitored every time step, while the torque was recorded every 15 time steps. The solution consistently became quasi-steady after 5 rotations, approximately 3150 time steps. The time step used for the simulation allowed the solution to converge after 30 iterations per time step, resulting in nearly 100,000 iterations to achieve a quasi-steady state.

The average torque was calculated as described in Sec. 4.4.1 for the last rotation of the wind turbine for each of the grids. Figure 5.6 shows the average torque for all the grids. For each of the hybrid grids, the resolution of the farfield domain remained relatively constant, while the local blade resolution was increased to match the structured grids. From this grid dependency study, it is obvious that the farfield




Figure 5.6. Average torque comparison for grid independent solution determination.

domain resolution had a minimal effect on the results, as the hybrid grids compared well with their structured equivalents. The 100,000 cell structured grid and 55,000 cell hybrid grid seem to exhibit grid convergence. However, due to the complexity of the topology required to construct the structured grid and the difficulty of applying this topology to varying geometries, the 55,000 cell hybrid grid topology was chosen for the optimization.




Figure 5.7. Hybrid grids for VAWT geometry, $\sigma = 1.5$ (left) and $\sigma = 0.4$ (right).

To demonstrate the automation and quality of the hybrid grid topology for arbitrary blade geometries, a hybrid grid was constructed for a high solidity VAWT geometry and a low solidity geometry seen in Fig. 5.7. The normal hyperbolic extrusion created layers of quads that marched smoothly away from the blade, transitioning to isotropic triangles. Utilizing this technique resulted in high quality, automated hybrid grid generation for all airfoil geometries analyzed throughout the optimization process.

5.2 Baseline Geometry

A baseline VAWT geometry was selected with which the results of the optimization could be compared. The idea was to select a typical VAWT airfoil cross-section. Therefore, the NACA 0015 was selected as the baseline airfoil cross-section simply due to the fact that a number of researchers attribute this geometry with good overall aerodynamic performance [18], [19], [22]. The NACA 0015 airfoil cross-section can be seen in Fig. 5.8.




Figure 5.8. NACA 0015 baseline geometry.

5.2.1 Baseline Performance

The performance of the baseline three-bladed VAWT utilizing NACA 0015 airfoil cross-sections was evaluated for $\sigma = 1.5$ and $\lambda = [0.5, 1.5]$, as well as for $\sigma = 0.4$ and $\lambda = [2.5, 3.5]$. By keeping ω constant at 10 rad/s and 12 rad/s for $\sigma = 1.5$ and $\sigma = 0.4$ respectively, V_∞ was adjusted to control λ . This provided relevant performance data surrounding $\lambda = 1$ and $\lambda = 3$, the tip speed ratio design constraints for the optimization. A total of 5 simulations were run for each of the solidities tested to build up a performance envelope for the baseline geometry. The average torque was calculated for each simulation and was used to determine the coefficient of performance. The results of the analysis can be seen in Fig. 5.9 and Fig. 5.10.




Figure 5.9. NACA 0015 performance envelope, $\sigma = 1.5$.




Figure 5.10. NACA 0015 performance envelope, $\sigma = 0.4$.

Figure 5.9 defines the performance envelope for a VAWT utilizing the NACA 0015 airfoil for $\sigma = 1.5$. It can be seen that there exists a point at which the efficiency is highest ($\lambda \approx 1.2$), and can be described as the optimum tip speed ratio for the geometry. As expected, as the wind speed changes, driving the tip speed ratio away from the optimum, the efficiency decreases. From this figure, it can be deduced that in order for the baseline wind turbine design to perform optimally at $\lambda = 1$, the solidity of the rotor would have to be increased beyond $\sigma = 1.5$ while retaining the NACA 0015 airfoil cross-section. Increasing the solidity of the rotor tends to result in lower optimum tip speed ratios, whereas lower solidities increase the optimum tip speed ratio. However, in order to compare the results of the NACA 0015 with the optimization test cases, the solidity of the baseline geometry was not adjusted. The

performance envelope of the NACA 0015 for $\sigma = 0.4$ did not provide much insight, as only one side of the curve was captured during the simulation. The fact that the maximum efficiency was not captured during the simulation reveals that $\sigma = 0.4$ is much higher than necessary to achieve maximum efficiency at $\lambda = 3$ using the NACA 0015 airfoil.

5.2.2 Published Performance Trends

The results of the baseline geometry were compared with published trends for validation. By comparing with published trends, the goal was to ensure that the results obtained from the 2D simulation represented realistic physical characteristics of a full 3D VAWT. In Ch. 2 it was mentioned that while the height of the rotor is the only 3D factor that effects the torque, theoretically it does not effect the efficiency. Therefore, while a more detailed 3D flow analysis would be required to predict all flow field phenomena, a 2D simulation was found adequate enough for the optimization.




Figure 5.11. Simulated performance (left) and experimental performance (right) [3].

Figure 5.11 shows the performance envelope of the NACA 0015 baseline geometry ($\sigma = 1.5$) compared with published experimental performance data for a VAWT utilizing a NACA 6-series airfoil cross-section ($\sigma = 0.3$). While the geometry and operating tip speed ratio may be different, the important thing to note is the trend of the performance envelope. The simulation was able to predict the same trend as the experimental data, building confidence in the capabilities of the 2D numerical model. Also, observing the range in tip speed ratios for the experimental data in Fig. 5.11 goes to show that low solidity VAWTs tend to have a flatter trend, meaning that they perform better over a higher range of tip speed ratios than a high solidity design. This was observed in Fig. 5.9 ($\sigma = 1.5$) where the entirety of the performance envelope can be seen, and Fig. 5.10 ($\sigma = 0.4$) where only a fraction of the performance envelope is observed, a characteristic that was shown experimentally by Brulle [47].

5.3 Case 1

Two test cases were run through the optimization system, the details of which are shown in Table 5.1. Each case ran for approximately 1 week on the cluster described in Sec. 4.3.4, after which the maximum number of user specified generations was reached ($G = 11$). Because there was no guarantee that the optimization algorithm would find the optimum design, the goal was to obtain an improved design that was able to achieve a higher efficiency than the baseline geometry.

To demonstrate the capabilities of the optimization system, and show that 1 week was enough time to achieve an optimized geometry, 2 unique initial populations that were randomly generated by the DE code were run through the optimization system. The results of the 2 unique runs will be presented and compared with the baseline geometry. The optimization is said to have been successful if the VAWT uti-

lizing the optimized airfoil cross-section achieved a higher efficiency than the baseline geometry at the design tip speed ratio ($\lambda = 1$).

5.3.1 Optimization Results

The initial population of the 2 runs for the 1st test case can be seen in Fig. 5.12. Due to the nature of the DE algorithm, the initial population is random, and completely different for the 2 runs conducted. The reason for starting with 2 different initial populations was to ensure that 11 generations was a sufficient amount time to find an optimized design while avoiding premature convergence or stagnation.




Figure 5.12. Initial populations, 1st test case.

The diversity of the population for all generations can be seen in Fig. 5.13. This figure illustrates how the COP varies with generation. NACAopt-RUN1 and NACAopt-RUN2 refer to the 2 unique runs conducted for the 1st test case. It can be seen that the populations for each of the first 5 generations are quite diverse. However, after the 5th generation the populations begin to converge while still remaining somewhat diverse, a characteristic of the stochastic nature of the DE algorithm.




Figure 5.13. COP vs. generation for all population members, 1st test case.




Figure 5.14. Max COP vs. generation, 1st test case.

While Fig. 5.13 illustrates the diversity of the population for each generation, Fig. 5.14 provides a history of the best overall objective function value throughout the optimization. If the COP at the current generation happens to be higher than the previous maximum, it is replaced and the new airfoil design parameters are used to generate the next population. After 11 generations NACAopt-RUN1 and NACAopt-RUN2 were able to achieve a maximum COP of 0.373 and 0.374 respectively, despite the fact that each run was initialized from a different initial population.

The optimized geometry for both runs can be seen in Fig. 5.15. The NACAopt-RUN1 airfoil has a maximum camber of $0.0094c$, a maximum camber location of $0.599c$, and a maximum thickness of $0.177c$, where c is the chord length of the airfoil. The choice in a cambered airfoil geometry over a symmetric cross-section could be an indication that slight camber increases the efficiency of high solidity rotors that experience undesirable blade vortex interactions. Due to the fact that the maximum COP and the optimized airfoil cross-section for both runs is indistinguishable indicates that 11 generations is sufficient. Therefore, it is unnecessary to discuss the performance of both VAWTs and only the performance of the NACAopt-RUN1 will be presented.




Figure 5.15. Optimized NACA 4-series airfoil geometry, 1st test case.

The performance envelope for the VAWT using the NACAopt airfoil cross-section can be seen in Fig. 5.16. Because the optimization was run for $\lambda = 1$, the blade shape was tailored to perform as best as possible at this value. Therefore, the optimum tip speed ratio is much closer to 1, signifying that the solidity of the rotor is only slightly lower than necessary to achieve maximum efficiency at $\lambda = 1$.




Figure 5.16. NACAopt performance envelope, 1st test case.

5.3.2 Baseline Comparison

While the optimization algorithm was able to find an optimized NACA 4-series geometry for $\sigma = 1.5$ and $\lambda = 1$ with very little user input and little or no designer intuition or experience, it had to be compared with the baseline geometry to quantify

the performance gained by using such approach. The performance envelopes for the NACAopt and NACA 0015 VAWT designs are shown in Fig. 5.17. For the design tip speed ratio $\lambda = 1$, the NACAopt design has a $COP = 0.373$, 2.4% higher than the NACA 0015 baseline geometry, which over the lifetime of the VAWT is considered a significant improvement.

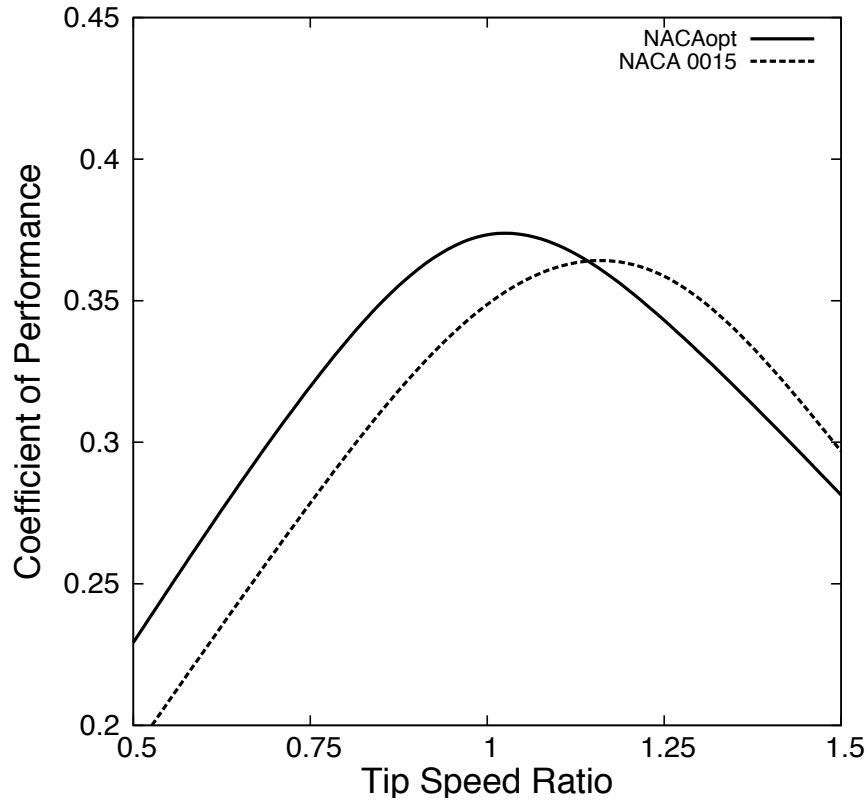


Figure 5.17. NACAopt vs. NACA 0015 performance, 1st test case.

In order to understand the mechanism for improved efficiency over the baseline geometry, the torque for a single rotation was observed, seen in Fig. 5.18. While the frequency of the oscillation in the torque is the same for both geometries, the peak-to-peak amplitude for the NACA 0015 is higher than that of the NACAopt geometry. The higher thickness of the NACAopt geometry allows for such a cross-section to

achieve a slightly higher angle of attack before stall than that of the NACA 0015 airfoil. Therefore, due to the increase in drag associated with the dynamic stall of the NACA 0015, a higher cyclic loading is observed. Not only did the NACAopt obtain a higher efficiency, but also a reduction in cyclic loading which could lead to a longer lifespan than the NACA 0015 geometry.




Figure 5.18. NACAopt vs. NACA 0015 torque for $\lambda = 1$, 1st test case.

Interesting flow field phenomena were captured when visualizing the vorticity seen in Fig. 5.19. The image clearly reveals that the wake of one blade actually interacts with the trailing blade, a characteristic typical of high solidity rotors. This interaction disturbs the flow, altering the velocity field around the trailing blade and is most likely the reason a cambered airfoil was chosen for this high solidity geometry. In the pair of images labeled *b*, a leading edge separation bubble can be seen forming on the lower left blade for the NACA 0015 geometry. However, the




Figure 5.19. NACAopt (left) and NACA 0015 (right) vorticity contours, 1st test case.

same phenomena is not observed for the NACAopt geometry. In c and d as the blade continues to rotate counterclockwise the separation bubble becomes larger and eventually separates, contributing to the trailing vortex and increasing its strength. The increase in the efficiency of the NACAopt geometry can be attributed to the airfoil cross-section's favorable characteristics at higher angles of attack, leading to the elimination of the leading edge separation bubble and a reduction in cyclic loading.

5.4 Case 2

This section focuses on the 2nd test case. Similar to the previous case, 2 initial populations were generated and their results presented and compared with the baseline geometry. Again, the optimization is said to have been successful if the optimized geometry achieved a higher efficiency than the baseline at the design tip speed ratio ($\lambda = 3$).

5.4.1 Optimization Results

The random initial populations generated by the DE code can be seen in Fig. 5.20. Two initial populations were created solely to ensure that 11 generations was a sufficient amount of time to find an optimized design that outperformed the baseline geometry while avoiding premature convergence or stagnation.

The objective function value, COP, can be seen for each generation in Fig. 5.21 for the NACAopt-RUN1 and NACAopt-RUN2, the 2 unique runs conducted for the 2nd test case. Unlike the 1st test case, several population members have a negative COP. A negative COP indicates that the airfoil geometry simply will not perform at the design tip speed ratio. Rather than integrating the pressures around the airfoils to achieve the appropriate rotation speed of the VAWT, a fixed tip speed ratio was enforced. Therefore, a negative efficiency indicates that the VAWT would




Figure 5.20. Initial populations, 2nd test case.

be unable to produce power at the specified tip speed ratio. This occurs for several population members in the first 4 generations. From the 5th generation onwards, the efficiency remains positive. It was found from the simulations conducted for the baseline geometry that $\sigma = 0.4$ was too high for the design tip speed ratio and is most likely the reason that several population members have a low COP throughout the optimization.

The history of the maximum COP throughout the optimization is shown in Fig. 5.22. After the 5th generation, the DE algorithm seems to have found the maximum COP, signifying the possibility of premature convergence. However, looking back at Fig. 5.21, the population after the 5th generation has not lost its diversity, an indication that the algorithm did not converge prematurely. After 11 generations it was observed that the NACAopt-RUN1 and NACAopt-RUN2 geometry achieved a COP of 0.340 and 0.341 respectively, despite being initialized from two different initial populations.

The optimized geometry for the 2 runs can be seen in Fig. 5.23. The NACAopt-RUN1 airfoil is symmetric with a maximum thickness of 0.109c. The choice in a




Figure 5.21. COP vs. generation for all population members, 2nd test case.




Figure 5.22. Max COP vs. generation, 2nd test case.

symmetric airfoil is significant. Because low solidity rotors do not experience strong blade vortex interactions, the positive and negative angles of attack that the blades experience are of the same magnitude, therefore, symmetric airfoils are typically used. Because the COP and airfoil cross-sections are essentially the same, an indication that 11 generations was sufficient, the NACAopt-RUN1 results will be presented.




Figure 5.23. Optimized NACA 4-series airfoil geometry, 2nd test case.

The performance envelope for the NACAopt geometry can be seen in Fig. 5.24. Due to the fact that the solidity was fixed as a design constraint, the optimum tip speed ratio is lower than the design tip speed ratio $\lambda = 3$. This signifies that the solidity of the rotor is larger than necessary to achieve maximum performance at $\lambda = 3$. While the optimization was able to select an airfoil cross-section with a decent level of performance given the design constraints, allowing the solidity to become a design parameter would provide complete geometric flexibility. Such a design would be able to achieve maximum performance at the design tip speed ratio and has been considered a case for future work.




Figure 5.24. NACAopt performance envelope, 2nd test case.

5.4.2 Baseline Comparison

The optimization algorithm was able to find an optimized NACA 4-series airfoil cross-section for $\sigma = 0.4$ and $\lambda = 3$ in 11 generations, the performance of which was compared with the NACA 0015 baseline geometry shown in Fig. 5.25. The NACAopt design was able to achieve a $COP = 0.340$ at $\lambda = 3$, 1.1% more efficient than the baseline design. While the solidity was too high to obtain maximum performance at $\lambda = 3$, the NACAopt geometry was found to exhibit a higher efficiency than the baseline design at and around the design tip speed ratio.

In an attempt to determine the reason for the higher efficiency associated with the NACAopt geometry, the torque for a single rotation of the optimized design was compared with the NACA 0015, shown in Fig. 5.26. Although the difference is




Figure 5.25. NACAopt vs. NACA 0015 performance, 2nd test case.

subtle, the peak-to-peak amplitude of the NACA 0015 geometry is larger than that of the NACAopt geometry. For each valley, the torque generated by the NACA 0015 rotor is slightly lower than the NACAopt geometry and is a result of the increase in drag associated with the thicker cross-section. On the other hand, the peaks in the torque for both geometries are similar because while the drag component is lower for the NACAopt geometry, this difference is offset with the slight reduction in lift associated with the NACAopt airfoil's thinner cross-section.

Vorticity contours were captured at various instances during a rotation and can be seen in Fig. 5.27. Unlike the 1st test case, no major differences between the NACAopt and NACA 0015 flow fields were observed. In the pair of images labeled *b*, the wake from the blade on lower left can be seen interacting with the




Figure 5.26. NACAopt vs. NACA 0015 torque for $\lambda = 3$, 2nd test case.

blade moving around the top of the wind turbine. The DE algorithm's selection of a symmetric airfoil over a cambered geometry is most likely a result of the weaker blade vortex interaction associated with the low solidity geometry when compared with the high solidity wind turbine in the 1st test case. The increase in the efficiency of the NACAopt geometry can be attributed to the lower drag associated with the thinner cross-section blade, leading to a slightly lower peak-to-peak amplitude than the NACA 0015 and higher overall average torque.




Figure 5.27. NACAopt (left) and NACA 0015 (right) vorticity contours, 2nd test case.

CHAPTER 6

CONCLUSIONS

This work successfully demonstrated a fully automated process for optimizing the airfoil cross-section of a VAWT. The generation of NACA airfoil geometries, hybrid mesh generation, and unsteady CFD were coupled with the DE algorithm subject to tip speed ratio, solidity, and blade profile design constraints. The optimization system was then used to obtain an optimized blade cross-section for 2 test cases, resulting in designs that achieved higher efficiency than the baseline geometry. The optimized design for the 1st test case achieved an efficiency 2.4% higher than the baseline geometry. The increase in efficiency of the optimized geometry was attributed to the elimination of a leading edge separation bubble that was causing a reduction in efficiency and an increase in cyclic loading. For the 2nd test case, the optimized design outperformed the baseline geometry by 1.1% at the design tip speed ratio. This increase in efficiency was a result of the optimized airfoil's ability to achieve a slightly lower cyclic loading and higher average torque. While this study is significant, it represents an initial step towards the development of an operational VAWT utilizing an optimized blade cross-section and requires further research and development.

The decision to fix the solidity resulted in designs where there was a noticeable difference between the optimum tip speed ratio and the design tip speed ratio. This effect could be mitigated by allowing the solidity to become a parameter in the optimization. Therefore, not only would the blade shape be allowed to change, but also the size of the blade itself. The optimization algorithm would then have complete

control over the geometry and the maximum efficiency would occur at the design tip speed ratio.

While the tip speed ratio, solidity, and blade shape effect the torque and efficiency of the wind turbine, theoretically the height of the wind turbine only effects the torque, assuming any 3D effects to be negligible. Therefore, it is important to carry the simulation into 3D to capture any unsteady phenomena, like tip leakage, that can deteriorate the performance of the wind turbine. Trend studies could then be conducted to determine the effect that the height of the wind turbine has on the torque and efficiency. The 3D simulations could also be used to calculate the increase in performance obtained while using the optimized blade cross-section and see how well it compares with the 2D results.

This study focused on the aerodynamic optimization of a VAWT, however, due to the oscillations in the torque, the wind turbine experiences a cyclic loading. Therefore, a structural analysis would be required to determine whether this loading would damage or destroy the wind turbine, and may even introduce additional constraints for the optimization.

REFERENCES

- [1] H. Stiesda, “The wind turbine components and operation,” Bonus Energy A/S, 1999.
- [2] T. Baker, “Mesh generation: Art of science?” *Progress in Aerospace Sciences*, pp. 29–63, 2005.
- [3] W. Moran, “Giromill wind tunnel test and analysis volume i - executive summary,” McDonnell Aircraft Company, St. Louis, MO, Tech. Rep., 1977.
- [4] M. Percival, P. Leung, and P. Datta, “The development of a vertical turbine for domestic electricity generation,” University of Northumbria, School of Engineering, Ellison Building, Ellison Place, Newcastle Upon Tyne, Tyne and Wear, UK, NE1 8ST., Tech. Rep.
- [5] M. Islam, D.-K. Ting, and A. Fartaj, “Aerodynamic models for darrieus-type straight-bladed vertical axis wind turbines,” *Renewable and Sustainable Energy Reviews*, pp. 1087–1109, 2008.
- [6] I. Paraschivoiu, F. Saeed, and V. Desobry, “Prediction capabilities in vertical-axis wind turbine aerodynamics,” in *The World Wind Energy Conference and Exhibition*, Berlin, Germany, 2002.
- [7] R. Templin, “Aerodynamic performance theory for the nrc vertical axis wind turbine,” National Research Council Canada, Tech. Rep. LTR-LA-10, 1974.
- [8] J. Strickland, “The darrieus turbine: A performnace prediction model using multiple streamtubes,” Sandia National Laboratories Report, Tech. Rep. SAND 75-041, 1975.

- [9] I. Paraschivoiu, “Aerodynamic loads and performance of the darrieus rotor,” *Journal of Energy*, pp. 406–412, 1981.
- [10] D. Berg, “An improved double-multiple streamtube model for the darrieus-type vertical axis wind turbine,” in *6th Biennial Wind Energy Conference and Workshop*, Minneapolis, MN, 1983, pp. 231–233.
- [11] R. Templin, “The nrc double multiple stream tube theory for vertical axis wind turbines,” National Research Council Canada, Tech. Rep. TM WE 023, 1985.
- [12] A. Gorban, A. Gorlov, and V. Silantyev, “Limits of the turbine efficiency for free fluid flow,” *ASME Journal of Energy Resources Technology*, vol. 123, 2001.
- [13] D. Lian, J.-H. Lee, and Y.-C. Kim, “Power prediction of darrieus type wind turbine considering real air velocity on the wind turbine blade,” School of Mechanical Engineering, Kunsan National University, Korea, Tech. Rep.
- [14] H. Larsen, “Summary of a vortex theory for the cyclogiro,” in *Proceedings of the second US national conferences on wind engineering research*, Colorado State University, 1975.
- [15] J. Strickland, B. Webster, and T. Nguyen, “A vortex model of the darrieus turbine: An analytical and experimental study,” Sandia National Laboratories Report, Tech. Rep. SAND 79-7058, 1980.
- [16] S. McIntosh, H. Babinsky, and T. Bertenyi, “Aerodynamic modeling of swept bladed vertical axis wind turbines,” in *47th AIAA Aerospace Sciences Meeting Including The New Horizons Forum and Aerospace Exposition*, Orlando, Florida, 2009.
- [17] S. Tullis, A. Fielder, K. McLaren, and S. Ziada, “Medium-solidity vertical axis wind turbines for use in urban environments,” Department of Mechanical Engineering, McMaster University, Hamilton, Canada, Tech. Rep.

- [18] J. Vassberg, A. Gopinath, and A. Jameson, “Revisiting the vertical-axis wind-turbine design using advanced computational fluid dynamics,” in *43rd AIAA ASM*, Reno, Nevada, 2005.
- [19] C. S. Ferreira, H. Bijl, G. van Bussel, and G. van Kuik, “Simulating dynamic stall in a 2d vawt: Modeling strategy, verification and validation with particle image velocimetry data,” *Journal of Physics*, 2007.
- [20] M. Torresi, S. Camporeale, P. Strippoli, and G. Pascazio, “Accurate numerical simulation of a high solidity wells turbine,” *Renewable Energy*, pp. 735–747, 2008.
- [21] J. Menet and N. Bourabaa, “Increase in the savonius rotors efficiency via a parametric investigation,” ENSIAME, France.
- [22] J. Edwards, N. Durrani, R. Howell, and N. Qin, “Wind tunnel and numerical study of a small vertical axis wind turbine,” in *46th AIAA Aerospace Sciences Meeting and Exhibit*, Reno, Nevada, 2008.
- [23] R. Bourguet, G. Martinat, G. Harran, and M. Braza, “Aerodynamic multi-criteria shape optimization of vawt blade profile by viscous approach,” *Wind Energy*, pp. 215–219, 2007.
- [24] M. Jureczko, M. Pawlak, and A. Mezyk, “Optimisation of wind turbine blades,” *Journal of Materials Processing Technology*, pp. 463–471, 2005.
- [25] P. Fuglsang and H. Madsen, “Optimization method for wind turbine rotors,” *Journal of Wind Engineering and Industrial Aerodynamics*, pp. 191–206, 1999.
- [26] H. Rahai and H. Hefazi, “Vertical axis wind turbine with optimized blade profile,” U.S. Patent 7,393,177 B2, 1, 2008.
- [27] F. Saeed, I. Paraschivoiu, O. Trifu, M. Hess, and C. Gabrys, “Inverse airfoil design method for low-speed straight-bladed darrieus-type vawt applications.”
- [28] M. Hansen, *Aerodynamics of Wind Turbines, 2nd Ed.* London, UK: Earthscan, 2008.

- [29] D. Dellinger. (2008) Wind - average wind speed. National Climatic Data Center. [Online]. Available: <http://www.ncdc.noaa.gov/oa/climate/online/ccd/avgwind.html>
- [30] N. Cetin, M. Yurdusev, R. Ata, and A. Ozdemir, “Assessment of optimum tip speed ratio of wind turbines,” *Mathematical and Computational Applications*, vol. 10, no. 1, pp. 147–154, 2005.
- [31] I. Abbott and A. von Doenhoff, *Theory of Wing Sections Including a Summary of Airfoil Data*. New York: Dover Publications, Inc., 1959.
- [32] W. Anderson and D. Bonhaus, “Aerodynamic design on unstructured grids for turbulent flows,” National Aeronautics and Space Administration Langley Research Center, Hampton, Virginia, Tech. Rep.
- [33] M. Giles, “Aerodynamic design optimisation for complex geometries using unstructured grids,” Oxford University Computing Laboratory Numerical Analysis Group, Oxford, England, Tech. Rep.
- [34] E. Nielsen and W. Anderson, “Aerodynamic design optimization on unstructured meshes using the navier-stokes equations,” *AIAA*, 1998.
- [35] J. Elliott and J. Peraire, “Aerodynamic design using unstructured meshes,” *AIAA*, 1996.
- [36] Elliott and Peraire, “Practical 3d aerodynamic design and optimization using unstructured meshes,” Department of Aeronautics and Astronautics Massachusetts Institute of Technology, Cambridge, MA, Tech. Rep.
- [37] J. Steinbrenner and J. Abelanet, “Anisotropic tetrahedral meshing based on surface deformation techniques,” *AIAA*.
- [38] B. Munson, D. Young, T. Okiishi, and W. Huebsch, *Fundamentals of Fluid Mechanics, 6th Edition*. John Wiley and Sons, Inc., 2009.
- [39] F. White, *Fluid Mechanics, 6th Edition*. McGraw-Hill, 2006.

- [40] J. Ferziger and M. Peric, *Computational Methods for Fluid Dynamics, 3rd Edition*. New York: Springer-Verlag, 2002.
- [41] P. Spalart and S. Allmaras, “A one-equation turbulence model for aerodynamic flows,” in *30th Aerospace Sciences Meeting and Exhibit*, Reno, NV, 1992.
- [42] Fluent, “Fluent 6.3 user’s guide,” Fluent Inc., Lebanon, NH, Tech. Rep., 2006.
- [43] R. Storn and K. Price, “Differential evolution - a simple and efficient adaptive scheme for global optimization over continuous spcaes,” ICSI, Tech. Rep. TR-95-012, 1995.
- [44] Storn and Price, “Differential evolution - a simple and efficient heuristic for global optimization over continuous spaces,” *Journal of Global Optimization*, vol. 11, pp. 341–359, 1997.
- [45] K. Price, “Differential evolution: A fast and simple numerical optimizer,” *IEEE*, 1996.
- [46] J. Steinbrenner, N. Wyman, and J. Chawner, “Development and implementation of gridgen’s hyperbolic pde and extrusion methods,” *AIAA*, 2000.
- [47] R. Brulle, “Feasibility investigation of the giromill for generation of electrical power volume ii - technical discussion,” McDonnell Aircraft Company, St. Louis, MO, Tech. Rep., 1977.

BIOGRAPHICAL STATEMENT

Travis J. Carrigan was born in Glenrock, Wyoming in 1987 and moved to Blackfoot, Idaho shortly thereafter. He received both his B.S. and M.S. degree in Aerospace Engineering from The University of Texas at Arlington in 2009 and 2011, respectively. Early in his undergraduate education, he worked at Vought Aircraft as a Quality Assurance Engineer on the Boeing 787 Dreamliner program. He currently works for Pointwise, a company focused on mesh generation for CFD. His current research interests include, CFD, mesh generation, wind turbine aerodynamics, numerical optimization, and aircraft design. He is a member of AIAA and ASME. He is also a licensed private pilot and a member of AOPA and EAA.