



The
University
Of
Sheffield.

Department
Of
Mechanical
Engineering

PhD Mechanical Engineering Thesis

Aerodynamic Optimisation of Vertical Axis Wind Turbines Using Adjoint Methods in CFD

Harry M **DAY**

Supervisors:

Prof. Lin Ma

Prof. Derek Ingham

Prof. Mohamed Pourkashanian

A dissertation submitted to the Department of Mechanical Engineering, University of Sheffield in fulfilment of the requirements for the degree of PhD.

November 2021

Confirmation of Authorship & Publications List

Confirmation of Authorship

The work presented in this thesis is the sole effort of the author unless otherwise stated in the text. Acknowledgement and references are given where appropriate to any sources that have been used. Furthermore, no part of this work has been previously submitted for an award at this or any other institution.

In the below publications list, the co-authors are my supervisors who have acted in an advisory role.

Publications

Title: “Adjoint based optimisation for efficient VAWT blade aerodynamics using CFD”

Authors: Day, H., Ingham, D., Ma, L. & Pourkashanian, M.

Journal of Wind Engineering and Industrial Aerodynamics, Vol. 208 (2021). 104431.

Date of Publication: **November 2020**

The content of this publication relates most specifically to Chapter 5 of the thesis, but the methods in Chapters 6 and 7 are also based on this foundation.

Planned Publications

Title: “Improved Adjoint based optimisation of VAWTs using CFD”

Authors: Day, H., Ingham, D., Ma, L. & Pourkashanian, M.

Journal: TBC

Date of Publication: **Yet to be published (work in progress)**

The content of this publication relates specifically to Chapters 6 and 7.

Title: “Morphing VAWT blades according to Adjoint optimised geometries in CFD”

Authors: Day, H., Ingham, D., Ma, L. & Pourkashanian, M.

Journal: TBC

Date of Publication: **Yet to be published (work in progress)**

The content of this publication will utilise the results from Chapter 6 as part of the proposed methodology.

Acknowledgements

I would like to extend my gratitude to the Engineering & Physical Sciences Research Council (EPSRC) for funding this research, and indeed other work that pursues solutions to energy and climate change issues.

I wish to thank my supervisors Professor Ma and Professor Ingham for their constructive words of support, guidance and optimism offered to me throughout my project. I am very grateful to my colleagues and friends also working on wind energy research. In particular; Dr Nidiana Rosado Hau and Dr Mohamed El Sakka who have been an invaluable source of council and expertise, patiently helping me up the learning curve of wind turbine CFD and providing insightful discussions and suggestions. I also thank Dr Helen Theaker for her generous support towards the end of it.

Most of all, I thank my family.

Abstract

Vertical Axis Wind Turbines (VAWTs) have potential to provide a greater contribution to the deployment of wind energy around the world. At present however, progress is hindered by a massive shortfall in effectiveness of design methods associated with the complexity of VAWT aerodynamics.

Predicting and designing for VAWT flows is not only an interesting technological challenge, but also one of great societal importance. If the performance and efficiency of VAWTs are improved, then their uptake and deployment around the world will be increased. Due to some key advantages of VAWTs, their widespread use in conjunction with Horizontal Axis Wind Turbines (HAWTs) could drastically improve the cost competitiveness of wind energy overall.

This thesis describes the development of a new design/optimisation methodology that can be applied to VAWT aerodynamics. It is based on the powerful Adjoint method, which has been applied to other fields with great success. In the application of Adjoint methods to VAWT aerodynamics, this work shows an efficient new way to do VAWT blade shape optimisation.

A 'semi-transient' optimisation process is proposed, using Adjoint optimisation data from individual instances in time to improve VAWT performance. Such a method is novel in the field of VAWTs, and the use of Adjoint methods with low cost CFD models provides an efficient optimisation methodology that can be readily adopted by the VAWT design community. Throughout this work the commercial CFD software ANSYS Fluent (18.2) was used, although the proposed methodologies can be reproduced in other CFD codes which contain an Adjoint solver.

This thesis opens an exciting new avenue of research departing from conventional parameter-based design methods, and the results show that Adjoint-optimised blades can provide significant increases to a VAWTs power coefficient. Discussion of these novel blade shapes also sheds new and important light on the flow physics and optimal blade geometries throughout the VAWTs revolution. Additionally, the results show that the semi-transient Adjoint method can be used to provide blades which greatly reduce the aerodynamic loading on the VAWT (and thus fatigue damage) whilst improving the average power coefficient.

Contents

Confirmation of Authorship & Publications List	ii
Acknowledgements	iii
Abstract	iv
Contents	v
Glossary	viii
Nomenclature	x
1 Introduction	1
1.1 VAWT TECHNOLOGY (CONFIGURATIONS)	4
1.1.1 <i>Straight Bladed Darrieus</i>	4
1.1.2 <i>Egg-Beater (Phi-Rotor Darrieus)</i>	5
1.1.3 <i>Helical Rotor (Twisted Blade Darrieus)</i>	5
1.2 CHARACTERISTICS OF A VAWT	6
1.2.1 <i>Aerofoil Shape Parameters</i>	9
1.3 PROJECT AIMS	14
2 Vertical Axis Wind Turbines (VAWTs)	15
2.1 EARLY VAWT DEVELOPMENT	15
2.2 VAWT AERODYNAMICS	16
2.2.1 <i>Dynamic Stall</i>	16
2.2.2 <i>The Self-Starting Problem</i>	18
2.2.3 <i>Unsteady Wind/ Analogy with Flapping Wing</i>	20
2.2.4 <i>Flow Control Methods</i>	21
2.3 SIMPLIFIED MODELS (COMPUTATIONAL AERODYNAMICS)	24
2.4 CFD AND VAWTs	25
2.5 EXPERIMENTAL STUDY OF VAWTs	30
3 Aerodynamics Optimisation of VAWTs	31
3.1 ADJOINT SENSITIVITY ANALYSIS	40
3.2 CHOICE OF OBJECTIVE FUNCTION	42
3.3 MULTI-POINT/MULTI-OBJECTIVE	42
3.4 ADJOINT LITERATURE	43
3.5 ADJOINT BASED OPTIMISATION FRAMEWORKS IN CFD	45

4	Research Aim	47
4.1	RESEARCH NOVELTY	48
5	Single-Blade Optimisation for VAWTs	49
5.1	SINGLE BLADE APPROXIMATION MODEL	50
5.1.1	<i>Introducing the Reference Turbine</i>	52
5.1.2	<i>Single-Blade CFD Model Setup</i>	54
5.1.3	<i>Single-Blade CFD Model Validation</i>	59
5.1.4	<i>Validating the VAWT CFD Model</i>	60
5.2	SEMI-TRANSIENT ADJOINT OPTIMISATION PHILOSOPHY	61
5.2.1	<i>Objective Function</i>	61
5.2.2	<i>Snapshots</i>	62
5.3	OPTIMISATION ALGORITHM	63
5.3.1	<i>Adjoint Module Setup</i>	66
5.3.2	<i>Aggression</i>	69
5.3.3	<i>Algorithm Code</i>	70
5.3.4	<i>Algorithm Key Notes</i>	75
5.4	SINGLE-SNAPSHOT INVESTIGATION	78
5.4.1	<i>Candidate VAWT Production</i>	86
5.4.2	<i>Candidate VAWT Results – Single Snapshot Optimisation</i>	87
5.4.3	<i>Aerodynamic Analysis for Single-Blade, 1-Snapshot Optimisation</i>	94
5.5	CONCLUDING REMARKS ON 1-SNAPSHOT OPTIMISATION, SINGLE-BLADE METHOD	101
6	In-Situ VAWT (ISV) Optimisation	102
6.1	VAWT MODELLING PHILOSOPHY AND VALIDATION	103
6.2	SEMI TRANSIENT ADJOINT OPTIMISATION PHILOSOPHY	106
6.2.1	<i>Objective Function</i>	106
6.2.2	<i>Snapshots</i>	107
6.3	OPTIMISATION ALGORITHM	108
6.3.1	<i>Adjoint Module Setup</i>	109
6.3.2	<i>Aggression</i>	111
6.3.3	<i>Algorithm Code</i>	113
6.3.4	<i>Algorithm Key Notes</i>	118
6.4	SINGLE-SNAPSHOT INVESTIGATION	119
6.4.1	<i>Results – Single Snapshot Optimisation</i>	121
6.4.2	<i>Aerodynamic Analysis for ISV, 1-snapshot Optimisation</i>	137
6.5	CONCLUDING REMARKS ON 1-SNAPSHOT OPTIMISATION, ISV METHOD	145
7	Multiple-Snapshot Optimisation	147
7.1	MULTI-SNAPSHOT INVESTIGATION WITH THE ISV METHOD	147

7.1.1	<i>Philosophy of the Multi-Snapshot Investigation</i>	147
7.1.2	<i>Wave 1</i>	149
7.1.3	<i>Wave 2</i>	160
7.1.4	<i>Wave 3</i>	161
7.1.5	<i>Wave 4</i>	163
7.1.6	<i>Aerodynamic Analysis for Multi-Snapshot, ISV Optimisation</i>	166
7.2	CONCLUDING REMARKS ON MULTI-SNAPSHOT OPTIMISATION, ISV METHOD	179
8	Conclusions & Future Work	182
8.1	EVALUATION/LIMITATIONS OF METHODS	185
8.1.1	<i>Single Blade Method</i>	185
8.1.2	<i>ISV Method</i>	186
8.1.3	<i>Multiple-Snapshot ISV Method</i>	187
8.2	FUTURE WORK	188
9	References	190
10	Appendix A – Single-Blade CFD Model Validation	200
11	Appendix B – VAWT CFD Model Validation	205
12	Appendix C – Single-Blade Oscillation UDF	209
13	Appendix D – Further Detail on Algorithm Key Notes (S-B Method)	211
13.1	MESH MORPHING (CONSTRAINTS, NEGATIVE CELL VOLUMES AND MESH DISTORTION)	211
14	Appendix E – Governing Equations and Turbulence Modelling	214
14.1	RANS EQUATIONS	214
14.2	SST K-W TURBULENCE MODEL	215

Glossary

AoA	Angle of Attack
BET	Blade Element Theory
BSL	Menter Baseline (turbulence model)
CFD	Computational Fluid Dynamics
CFL	Courant-Friedrichs-Lewy (condition, number)
CST	Class/Shape Function Transformation
CPU	Central Processing Unit (computer)
DES	Detached Eddy Simulation
DW	Downwind
FFS	Free Form Scaling
HAWT	Horizontal Axis Wind Turbine
HPC	High Performance Computing
ISV	In-situ VAWT (optimisation method)
LES	Large Eddy Simulation
LEV	Leading Edge Vortex
LCOE	Levelised Cost of Energy
MAV	Micro Aerial Vehicle
NREL	National Renewable Energy Laboratory
N-S	Navier-Stokes (equations)
O&M	Operations and Maintenance
PARSEC	Parametric Section
PIV	Particle Image Velocimetry
RANS	Reynolds Averaged Navier-Stokes (equations)
RNG	Re-Normalisation Group (turbulence model)
RSM	Response Surface (optimisation) Methods
S-A	Spalart-Allmaras (turbulence model)
SI	Standard International (units)

SST	Shear-Stress Transport (k- ω turbulence model)
SPIV	Stereoscopic Particle Image Velocimetry
TSR	Tip Speed Ratio (λ)
UDF	User Defined Function
UW	Upwind
VAWT	Vertical Axis Wind Turbine
VGOT	Variable Geometry Oval Trajectory
WCT	Wall Clock Time

Nomenclature

Unless stated otherwise, the International System for units (SI) is adopted throughout this work.

c	Chord length of an aerofoil or VAWT blade, (m)
C_D	Drag force coefficient
C_L	Lift force coefficient
C_M	Moment coefficient
C_P	Power coefficient
C_T	Tangential force coefficient
k	Reduced frequency
R	Radius of turbine, (m)
U_{wind}	Velocity of oncoming wind, (m/s)
V_{rel}	Relative velocity seen by a VAWT blade, (m/s)
α	Angle of Attack (AoA), ($^{\circ}$)
θ	Azimuthal (angular) position of VAWT blade, ($^{\circ}$)
λ	Tip Speed Ratio (TSR)
ρ	Density of fluid (air), (kg/m^3)
σ	Turbine solidity
ω	Angular velocity of VAWT, (rad/s)
ω_p	Pitching frequency of an oscillating aerofoil, (rad/s)

1 INTRODUCTION

Tackling climate change is one of the key challenges of the current generation. To achieve a cleaner and more sustainable energy future a global transition towards a heavily renewable energy infrastructure is required. In developed countries this constitutes moving away from the centralised fossil fuel energy networks that have been inherited. For developing countries this means utilising existing international knowledge of sustainable technologies, to make a renewables-focused energy infrastructure from which to build upon. A more decentralised energy network utilising high proportions of renewables, would provide a future with more energy security, reduced climate impact, and a level of redundancy against oil and gas shortages and political/market pressures.

The main challenges in increasing the penetration of wind power (and other renewables) into the energy mix relate to overcoming socio-economic factors. Renewable energy technologies already exist which *could* be more widely deployed, and there are regions of abundant resource (such as wind) around the world which could be utilised. Currently however, these renewable sources are not replacing, or negating conventional power stations at a sufficient rate for climate change goals to be met. Policy makers are hampered in efforts to introduce more renewables due to the higher costs associated with these emerging technologies, and with cultural/systemic inertia favouring fossil fuel industries. The engineering challenges lie in making renewable technologies more dispatchable, and in continually reducing their Levelised Cost of Energy (LCOE) to make them as cost competitive as possible.

Focussing specifically on wind energy, the offshore wind industry is becoming well established in Europe for large scale commercial energy production. Horizontal Axis Wind Turbines (HAWTs) are used for this application, and in the past decades, HAWTs have seen a far higher rate of research and development compared with Vertical Axis Wind Turbines (VAWTs). Large scale offshore HAWT projects will continue receiving investment in European markets, and the industry is rapidly growing in other parts of the world. Despite this, the rate of improvement to HAWT aerodynamic performance is plateauing (Ferrer & Montlaur, 2015), with most innovation in the industry focussed on upscaling the size of offshore HAWTs, improving condition monitoring systems and reducing operation and maintenance (O&M) costs. This has aided a renewed interest in VAWT technology. VAWTs are sometimes regarded as an inherently poorer concept than HAWTs in terms of power output and efficiency, but the VAWT industry is not well established so the comparison is somewhat unfair. The relatively underdeveloped VAWTs do indeed have some fundamental drawbacks; most notably poor self-starting behaviour and

performance issues (particularly for simple straight bladed configurations). Furthermore, producing efficient VAWT designs proves difficult due to challenges involved in predicting the flow phenomena (see Section 2.2) which are comparatively complex next to HAWTs. It is not however correct to declare VAWTs as a poorer concept in general - It may be surprising to some to learn that VAWTs are not of intrinsically lower theoretical efficiency than HAWTs (see Table 1-1).

HAWTs and VAWTs are inherently suited to different general purposes and operating conditions. Table 1-1 provides a comparison of the two arrangements, but it should be noted that only straight bladed Darrieus type VAWTs are compared in the table. The many other variations of VAWT (some of which are discussed in Section 1.1) have been omitted from Table 1-1 for brevity. The predominating disadvantage of all VAWTs is the comparative lack of research and development which they have received. This gives HAWT technology an unfair start, thus making it more difficult to assess the true merits of the concepts alongside each other.

Early progress (1970's – 1990s) with VAWT technology was marred by a series of failures as demonstrator projects attempted to upscale the technology. Such events are described by Tjiu *et al.* (2014), and these prematurely stunted interest and faith in the VAWT concept. Most wind turbine research conducted around this time was based on quasi-2D and quasi-steady methods, thus limiting the achievable progress at the time towards gaining a good understanding of the flow physics (Vassberg, *et al.*, 2005).

With modern computational techniques now allowing deep insight into VAWT flow physics, there is an opportunity for significant further improvement to be made in VAWT aerodynamic designs. VAWTs offer potential for further increase in renewables uptake when deployed in the appropriate regions and markets. With consideration of Table 1-1, it is apparent that VAWTs are more suited to deployment in regions of unsteady winds and varying wind direction. VAWTs can be simpler and more compact systems with easier deployment and maintenance, as well as placing reduced demand on the support structure. In future, VAWTs may play an important role for offshore applications, as well as in urban environments, micro-grids, and remote regions. VAWT aerodynamics should be researched further, to improve technical feasibility and cost competitiveness of the technology, making it easier for governments, investors, and climate-conscious businesses to favour.

The inherent issues that remain with VAWT technology present themselves as an opportunity for the engineer, and more specifically the aerodynamicist. This research project aims to explore the application of advanced aerodynamic shape optimisation techniques to the problem of Darrieus-type VAWT design. Powerful

computational techniques have the potential to make significant contributions to VAWT technology, as they have done in the aerospace field.

Table 1-1 – Comparison of VAWTs (Darrieus Type) with HAWTs

Advantages are coloured green, disadvantages are coloured blue and are in italics

VAWT (Darrieus Type)	HAWT
Aerodynamics / Performance	
<i>Aerodynamic complexity means relatively challenging analysis and design methodology</i>	Simpler aerodynamic behaviour, with well researched analysis and design methods
<i>Self-starting behaviour can be poor for some conditions and configurations</i>	Generally good self-starting behaviour
Good tolerance of very high winds	<i>Need to be deactivated in very high winds</i>
Operate independently of wind direction and therefore need no yaw control mechanism	<i>Performance is highly dependent on alignment of turbine with wind direction; yaw control required</i>
Ideal efficiency up to 72% but difficult to get close to this in reality (Bhutta <i>et al.</i> , 2011)	Efficiency up to 60%, from Betz Law (Bhutta <i>et al.</i> , 2011)
Manufacturing / Maintenance Cost	
Ground access to key components: gearbox, generator, bearings	<i>Key components located in tower which reduces access for maintenance and replacement</i>
Lower manufacturing costs	
<i>“Blade length of a Darrieus VAWT is approx. twice as long as a HAWT of equivalent rotor swept area” (Sutherland <i>et al.</i>, 2012)</i>	<i>Costly blades due to continuous twist/taper along the span</i>
Structural / Mechanical Complexity	
No yaw control required	<i>Yaw control required for good performance</i>
<i>High variation in blade loads (and fatigue) due to large cyclic variation of aerodynamic forces</i>	Smaller amplitude cyclic blade loads
Lower maximum loads exerted on support structure, and simpler overall structural design	<i>High maximum loads imparted on the tower / support structure</i>
Fewer mechanical components and complexity	-
Able to use direct drive generators reducing mechanical complexity	<i>Generally not used with direct drive generators</i>
Other (Environmental)	
More compact design requiring less space	-
Easier to integrate with building architecture from a visual and structural perspective	-
Reduced bird collision rate due to visibility, height and rotational speeds	<i>Higher incidence of bird collisions</i>
Inherently quieter due to low operating speeds	<i>Can be loud, causing problems with the public for onshore installations</i>

1.1 VAWT Technology (Configurations)

Within the envelope of VAWT technology there exists a wide variety of turbine configurations, each with their own specific features and advantages/disadvantages. This section discusses these various VAWT configurations. An excellent account of the history of VAWT development is given by Tjiu *et al.* (2014) with discussion of early research and demonstration projects carried out largely in Canada in the 1970's and 1980's. The failures of these early projects contributed to the stunted development of VAWTs, even though the problems were not due to inherent flaws in the VAWT concept.

VAWT configurations are broadly grouped into two categories; "lift type / Darrieus" turbines, and "drag type / Savonius" turbines. This review will focus on Darrieus-type VAWTs, as they are generally accepted as offering greater efficiency, capacity, and lower weight compared to the Savonius-type (Hill *et al.*, 2008). The limitation on Savonius arrangements is stated by Vassberg *et al.* (2005) as the inability of the blade speed to exceed free stream velocity, and their theoretical efficiency is significantly lower than Darrieus turbines (Almohammadi *et al.*, 2015). For information on additional wind turbine configurations refer to Bhutta *et al.* (2011) where discussion is given on: Savonius, Savonius-Darrieus Hybrid, Crossflex, Two-leaf semi rotary, Zephyr Turbine, Omni-directional Guide-vane, VGOT (Variable Geometry Oval Trajectory), and the Two-tier Rotor (Darrieus-Musgrove). If the reader is interested, Shires (2013) describes even more unconventional concepts such as the Mono V-VAWT, NOVA V-VAWT and the Aerogenerator.

Within the family of Darrieus-type turbines, several concepts have been developed in order to address the prominent aerodynamic issues. A detailed discussion of these sub-configurations is provided by Bhutta *et al.* (2011). In the present work, only a very brief description and illustration of the simplest, and most widely researched Darrieus configurations is given.

1.1.1 Straight Bladed Darrieus

The most simple and typical example of a Darrieus turbine has straight blades, often of constant cross section and constant blade pitch (see Figure 1-2). Although this arrangement is simple, the connecting arms between the tower and blades can carry a significant penalty to the aerodynamic performance of the turbine. Such turbines tend to be associated with a low starting torque and in some cases require starter motors before they can achieve a rotating speed which produces significant power. Variations of this concept with variable pitch exist which can aid with generating self-starting torque and overall performance, albeit at the penalty of system complexity and cost.

1.1.2 Egg-Beater (Phi-Rotor Darrieus)

The shape of the Egg-Beater turbine blades is chosen to minimise bending moment loads in the blades. This also minimises the required size and number of the structural connecting arms (see Figure 1-1). To achieve this, the blade shape is approximated to a Tropkien profile – a rotating idealised rope which offers no bending resistance (Islam *et al.*, 2006). The complex geometry however does incur an associated cost increase due to the manufacturing of the blades. Costs however would invariably reduce as the technology/industry progresses from infancy to mass production.

1.1.3 Helical Rotor (Twisted Blade Darrieus)

The helical rotor exhibits a more constant aerodynamic loading on the tower structure which is beneficial for improving structural fatigue life. This is achieved via the perimeter-wise sweep of the blades (see Figure 1-3). Rather than concentrating a large load as a straight blade would pass through the azimuthal angle of maximum torque, for a helical blade, the instantaneous blade area passing through any azimuthal position is kept close to constant (Tjiu *et al.*, 2014).



Figure 1-2 – H-Rotor
(Tjiu *et al.*, 2014)



Figure 1-1 – FloWind
(Sutherland *et al.*, 2012)



Figure 1-3 – Helical Rotor
(Tjiu *et al.*, 2014)

1.2 Characteristics of a VAWT

“Such an unsteady aerodynamic situation where an oscillating airfoil moves from a low Reynolds number to a high Reynolds number in a single cycle leads to a kaleidoscope of nonlinear phenomena that designers have to face.” (Laneville & Vittecoq, 1986).

Predicting VAWT aerodynamics is more challenging than HAWTs owing to the added complexities of the transient flow field (Wang et al., 2010). This makes VAWT design more difficult; so regardless of the maximum theoretical efficiencies, most *current* VAWTs tend to exhibit lower operating efficiencies than HAWTs.

A steadily rotating HAWT blade sees a fairly constant flow field, whereas VAWT blades experience a constantly varying onset velocity, as well as varying Angle of Attack (AoA, α) which can be both positive and negative over the course of a revolution. Lift and drag forces (or coefficients; C_L , C_D) act perpendicular and parallel to the onset velocity at each blade, producing a resultant force which has a blade normal component (C_N), and a tangential/chordwise (C_T) component which gives the useful torque:

$$C_T = C_L \sin(\alpha) - C_D \cos(\alpha) \quad (\text{Eqn. 1.1}).$$

The AoA (and thereby the C_T) depends on the free stream wind velocity, the rotational velocity of the turbine blade, and the instantaneous angular position (azimuthal angle, θ) of the blade. Figure 1-4 shows the convention used for wind velocity (U_{wind}) and relative velocity (V_{rel}), and the aerodynamic force coefficients.

The Tip Speed Ratio (TSR) relates the wind velocity and rotational velocity:

$$\lambda = \frac{\omega \cdot R}{U_{\text{wind}}} \quad (\text{Eqn. 1.2}).$$

The relative velocity seen by a blade (with zero fixing angle) can be calculated in terms of TSR (λ) and azimuthal angle (θ) of the blade (Ferrer & Montlaur, 2015):

$$V_{\text{rel}} = U_{\text{wind}} \sqrt{\lambda^2 + 2\lambda \cos\theta + 1} \quad (\text{Eqn. 1.3}).$$

The corresponding angle of attack is defined as (Ferrer & Montlaur, 2015):

$$\tan \alpha = \frac{\sin\theta}{\lambda + \cos\theta} \quad (\text{Eqn. 1.4}).$$




Figure 1-4 – Blade Angle of Attack and Relative Velocity and Notation




Figure 1-5 - Angle of Attack vs Azimuthal Position of a VAWT Blade

As the TSR increases, the AoA variation reduces and approaches the blade fixing angle. Figure 1-5 shows this effect; several TSRs are shown as well as a sine wave, demonstrating how a sine wave is a reasonable approximation to the AoA curve at high TSR (the sine wave is scaled by a factor, B, to achieve amplitude consistency with the curve of TSR = 5). It should be noted that this AoA calculation does not account for effects of velocity induction, thus offering a good approximation but not exact values for a real VAWT AoA.

Clearly from Figure 1-5, as the TSR reduces the AoA range increases (the turbine is rotating slowly), and as a blade moves through the upwind and downwind parts of the rotation it may operate in both the stalled and un-stalled regimes. This dynamic stall comprises complex formulation of vortices, followed by their development and separation into wakes, and this becomes more prevalent with lower TSR (Wang et al., 2010). The aerodynamic forces on each blade are further complicated by the interaction with other blade wakes and vortices (Edwards, 2012). The present flow field around a VAWT blade inherits flow disturbances and effects from the previous blade having passed through the same space earlier, and this is known as ‘hysteresis’ (Vassberg, 2005).

Figure 1-6 shows the variation in relative velocity for a VAWT blade over a full revolution (from azimuthal angle 0 to 360) using Equation 1.3. The curve for a sample turbine at TSR = 4.5, and $U_{wind}=9.3$ m/s is shown. Also shown is a line corresponding to the tangential velocity of the VAWT blades which can be taken to approximate the average relative velocity that the blade sees. It should be noted that for a real turbine the tangential velocity in fact fluctuates because the combined aerodynamic torque of all the blades varies constantly.

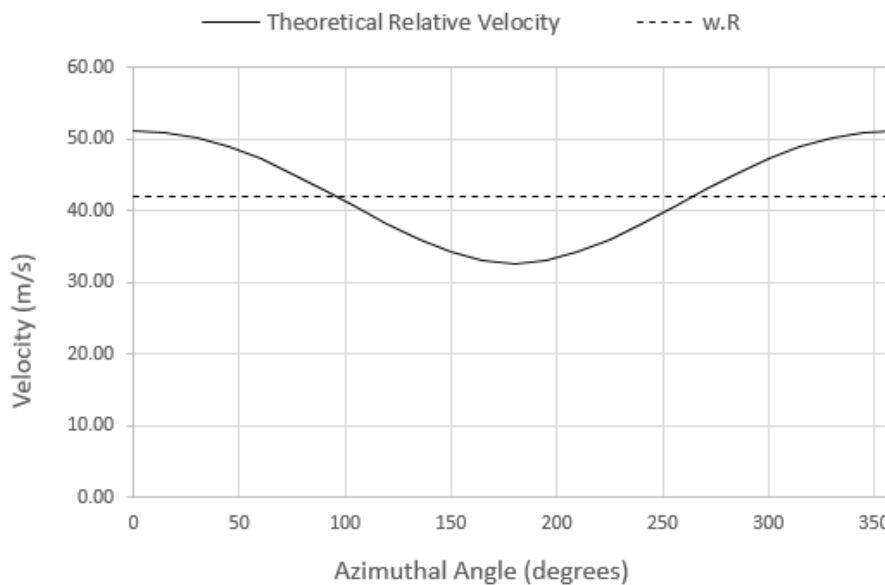


Figure 1-6 – Relative Velocity vs Azimuthal Position of a VAWT Blade

A VAWT operates successfully when a net positive torque (or moment) is achieved over a complete blade revolution; the blades must therefore generate positive torque over a range of positive and negative AoA in a complex, unsteady flow field. At low TSR there remains a portion of a blades revolution where negative torque is produced which must be overcome by the system inertia and the positive torque of the other blades. The possible range of AoA seen by the blades is magnified due to the variation in TSR that a turbine can experience. Operation with $TSR < 1.0$ (during start up or low wind speeds) tends to see high proportions of negative torque, as the blades have to operate over a greater range of sub-optimal AoA. This complex variation in flow conditions that VAWT blades experience, poses a far greater challenge to the aerodynamic designer compared to a HAWT blade. The silver lining to this is that while HAWT blades are optimised for a comparatively small range of AoA, VAWT performance is less effected by rapidly varying wind speed and direction. This enhances VAWT competitiveness for certain applications (Zhu *et al.*, 2016) where a wide envelope of wind conditions is required.

The terms of moment/torque coefficient (C_M) and power coefficient (C_P) are introduced as follows (note that the term C_T is reserved for tangential force coefficient of 1 blade):

$$C_M = \frac{M_F}{\frac{1}{2}\rho U_{wind}^2 A R} \quad (\text{Eqn 1.5}).$$

Where M_F , is the blade moment of force (N.m), T is the torque (sum of M_F for all blades in N.m), A is the swept area (m^2), and R is the turbine radius (m).

$$C_P = \frac{T \cdot \omega}{\frac{1}{2}\rho U_{wind}^3 A} \quad (\text{Eqn 1.6}).$$

The complex time dependent flow physics described in Section 1.2 have proven difficult to predict accurately, thus making aerodynamic design and optimisation of VAWTs a significant challenge.

1.2.1 Aerofoil Shape Parameters

VAWT performance is affected by a wide range of parameters regarding turbine layout and blade geometry. Historical parametric investigations into VAWT design variables such as blade geometry parameters are not comprehensive, and geometry/performance trends are still not well understood over the full operating range (Edwards, 2012). Research in VAWT development has typically been conducted using aerospace aerofoils such as the NACA series. This allowed efforts to characterise performance trends based on conventional/traditional geometrical parameters such as camber and thickness (Tjiu *et al.*, 2014). However, such

aerospace blade profiles are not necessarily suited to VAWTs. Aerospace wings are designed (especially in civil aircraft) for cruising/loitering conditions, which are large steady flow regimes and thereby alien to the highly transient nature of VAWT flow fields. Despite this, VAWT blades still largely employ typical symmetrical aerospace profiles as can be seen across the literature.

The drastic differences in application between a cruising wing and a rotating VAWT blade can be appreciated by the comparison of operating conditions given by Tjiu et al. (2014) which is recounted in Table 1-2.

Table 1-2 - Operating conditions of a Darrieus VAWT blade and an airplane blade (Table 1 from Tjiu et al., 2014)

Parameter	VAWT Blade	Aircraft Wing
Angle of Attack	Operate in unsteady conditions: oscillate between positive and negative AoA twice per revolution, which are often exceeding +/- 90°.	Operate in nearly steady conditions at near zero AoA
Stall	Encounter stall frequently, especially in strong wind.	Encounter stall only in unusual operating conditions.
Reynolds Number	Between a few hundred thousand and a few million.	Usually between three and thirty million.

The choice of aerospace aerofoils in VAWT research therefore appears somewhat arbitrary and potentially, sub-optimal. Hopefully, further research into blade profiles will see greater decoupling and independence from typical aerospace aerofoils; such research may be more likely to unlock the potential of VAWTs. The following briefly discusses the effects of classical aerodynamic shape parameters on VAWT performance. Studies to characterise and link these parameters with VAWT performance have been made in the literature, but general trends are not defined over the range of possible VAWT operating conditions (Edwards, 2012).

Camber

Research to date is inconclusive in providing generalised trends for the effects of camber on VAWT performance. Early works by Baker (1983) proposed that introducing camber should be beneficial to power output, arguing that separation will invariably occur at low Reynolds number; so camber will thereby have little detriment to separation but will aid reattachment on the upwind side improving

overall performance. More recent works fail to give convincing conclusions for performance trends on camber variations. Islam *et al.* (2007) concludes that NACA 4-digit symmetric aerofoils are not suitable for small VAWTs, while Worasinchai *et al.* (2016) calls symmetrical NACA sections a “simple and attractive choice for Darrieus rotors”.

Due to the variations across the various VAWT configurations, TSRs and turbine sizes available, it may be the case that there is no simple, convenient, and definitive trend in which the effects of camber can be generalised. Indeed, it can be seen that high camber geometries offer benefit at low TSRs due to the improved flow attachment; the converse becomes true as the TSR becomes larger and the AoA becomes smaller and more constant. The optimal camber therefore is a function of the specific turbine and operating conditions.

In this thesis, positive and negative camber are defined as shown in Figure 1-7. Positive camber means that the top surface is concave, moving away from the shaft. Negative camber means that the bottom surface is concave, moving closer to the shaft.




Figure 1-7 - Camber definition, (Left) positive, (Right) negative

Thickness

Islam *et al.* (2007) states that a large blade thickness is beneficial for self-starting improvements because thicker blades help delay stall at low Reynolds number. Thickness is shown to be less desirable at higher TSR where separation issues become less prominent and a higher profile drag is experienced as a penalty for thickness. There remains a lack of research into the effect of thickness on VAWT performance over the full range of operating conditions/turbine configurations.

Fixing Angle

Fixing angle, β is the angle made between the blade chord and the blades tangential velocity (which is at right angles to the turbine connecting arm). Figure 1-8 shows the definition/convention for fixing angle used throughout this thesis.

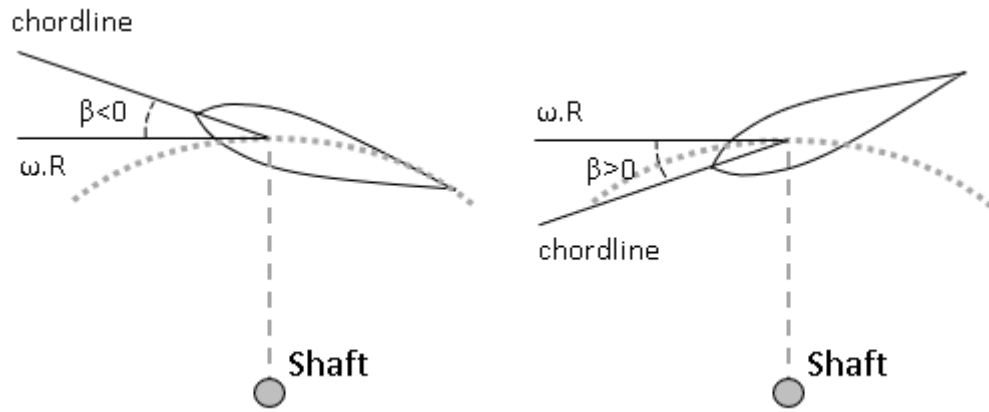


Figure 1-8 - Fixing angle definition (Left) negative, (Right) positive

Across the literature, fixing angle may also be referred to as ‘pitching angle’, or ‘mean angle’. This angle is geometrically fixed unless one is considering turbines with a variable pitch mechanism – the present work does not consider variable pitch turbines. Having a constant non-zero fixing angle causes a permanent skew to the range of AoA experienced by the blade. Klimas & Worstell (1981) investigated symmetrical aerofoils at different fixing angles and found small variations in fixing angle can exhibit great changes in cut-in TSR, efficiency and peak power coefficient. Coton *et al.* (1996) explored the effects of fixing angle on the turbine torque production around the revolution and noted delay of stall angle compared to the zero fixing angle case. Contrary to Klimas & Worstell (1981), it is concluded by Coton *et al.* that a fixing angle between +/- 6° has limited potential to improve performance. Small angles were found to reduce power output at low TSR with an improvement at high TSR, while large fixing angles were found to reduce performance across all TSRs considered. Rezaeiha *et al.* (2017) also explored the effects of changing the fixed pitch angle on VAWT performance. A -2 degree fixing angle achieved a +6.6 improvement to VAWT C_p , concluding that the change in AoA can reduce tendency of flow separation thus delaying the stall onset.

For low TSRs, the effect of blade fixing angle is small since the wind velocity predominates over blade rotational velocity, thus the onset angle diverges from the fixing angle. At high TSR the AoA converges to the fixing angle, and so the performance is more greatly influenced by the fixing angle (Hill *et al.*, 2008). Hill *et al.* claimed that choosing either a positive or negative fixing angle produces identical effects over a complete cycle when the fixing angle (magnitude) is less than the stall angle.

Solidity

Solidity, σ , of a VAWT depends on the number of blades (N), blade chord (c) and rotor radius (R) and is defined as:

$$\sigma = \frac{Nc}{R} \quad (\text{Eqn. 1.7}).$$

Solidity has a strong effect on the performance of a VAWT. High solidity turbines operate more efficiently at low TSR and exhibit a sharp loss of efficiency away from the optimum. Low solidity turbines exhibit a smoother power curve, and experience maximum efficiency at higher TSRs (Edwards, 2012).

Eboibi *et al.* (2016) investigated the effects of solidity on power output using experimentation. Particle Image Velocimetry (PIV) was used for flow visualisation. It was found that stall onset and the associated shedding of vortices occurs earlier for the lower solidity case, thus reducing efficiency. The higher solidity turbine was found to provide better average power coefficient. Howell *et al.* (2009) experimented with solidity by varying the number of blades. It was found that the 3-bladed machine drastically outperformed the lower solidity 2-blade machine over the majority of the operating range considered.

Surface Roughness

Howell *et al.* (2009) found surface roughness to have a significant effect on VAWT performance via experimentation. It was concluded that rough blades improved performance below a critical Reynolds number; above this Reynolds number smoother blades were favourable. Little research has been found on this topic, specifically regarding *wind* turbines, but the trend of roughness aiding performance only at low Reynolds numbers is demonstrated in research on *tidal* vertical axis turbines (Priegue & Stoesser, 2017).

The existing studies of VAWT blade parameters are not definitive or exhaustive but are certainly informative. Designing VAWT blades using these parameters is however to some extent, constricting. Better designs may be possible if a higher resolution in geometric variables could be used rather than adopting the conventional parameters. Optimisation studies normally need a large number of runs to examine just a few parameters, which makes only a small number of variables feasible (Bianchini *et al.*, 2014). Sophisticated methods that can efficiently handle large numbers of design variables are therefore sought after, as such methods could provide more optimal designs with less computation effort.

1.3 Project Aims

This section describes the general aims of the thesis, with more specific and detailed objectives given in Chapter 4.

Section 1.2 has described why VAWTs pose a significant design challenge for aerodynamicists. Producing accurate simulations of VAWT flows requires detailed Computational Fluid Dynamics (CFD) models which carry an associated time cost for CFD based design and optimisation methods. The general aim of this thesis is to develop efficient, CFD based aerodynamic optimisation procedures using a specific type of optimisation technique known as the Adjoint method. Adjoint solutions or Adjoint sensitivity analysis play a key part in massively improving the efficiency of gradient based optimisation processes for aerodynamics problems. The field of VAWTs has not yet benefitted from the power of Adjoints and this thesis aims to shed light on such benefits and to open new and important avenues of VAWT research.

For this thesis, the Adjoint solver in ANSYS Fluent (18.2) will be applied. ANSYS Fluent is chosen due to the wide user base and availability of the software, and also because the majority of simulations in the VAWT research field have used it (Balduzzi *et al.*, 2015). The methods developed here could be applied to other CFD codes which have an Adjoint solver, so this makes the work very accessible and repeatable for the general CFD/VAWT community, increasing the potential adoption and impact of this research. Details of the novelty in this work are discussed in Section 4.1.

2 VERTICAL AXIS WIND TURBINES (VAWTS)

This chapter begins with a short summary of the early developments in VAWT technology, followed by a discussion of the principal aerodynamic phenomena which are observed in VAWT flows. Discussion then moves onto the early aerodynamic models and experimentation used to explore VAWT flows, and the subsequent adoption of advanced CFD and experimental methods. Finally, advanced aerodynamic shape optimisation is introduced with focus on gradient based techniques such as the Adjoint method, with a view of exploring possible ways that such methods could be used with much advantage, to contribute to VAWT technology.

2.1 Early VAWT Development

The lift-type VAWTs, also known as Darrieus turbines are thus named after the French engineer who patented the concept in 1931 (Darrieus, 1931). There was little research on the topic during the following decades, but as global populations rocketed, knowledge and concerns regarding energy scarcity and security grew, and wind/renewable energies saw an advent of development during the late 20th century. The majority of the early VAWT developments (between 1970s and 1990s) were led by the National Research Council of Canada, and the US Department of Energy (via the Sandia Corporation). Sandia explored various VAWT configurations (focussing on Darrieus type) using a purpose built 34m test bed project. The test bed was used to produce comparative data for the design methods of the time (Sutherland *et al.*, 2012). Sandia produced a 17m rotor diameter model which saw deployment of more than 500 turbines by FloWind Corp (Sutherland *et al.*, 2012). Subsequent developments included the creation of turbines suitable for wider grid-tied commercial use. FloWind closed due to financing issues for the mass production of these turbines, which impeded the momentum of VAWT technology developments at the time. Subsequent research and development in wind technology was largely restricted to HAWTs for the following years.

The favour of HAWTs was also influenced by the comparative simplicity in their flow physics and thereby the ease of aerodynamic analysis. HAWTs provide a fairly constant power output and rotation speed, with less complex phenomena such as dynamic stall and hysteresis/blade wake interaction which are present in VAWTs. Modelling methods used in early wind turbine research such as Blade Element Theory (BET), were sophisticated enough for HAWTs, but inadequate at predicting VAWT flow complexities. At this critical time when the wind energy industry was in its infancy, VAWT designs were not competitive so HAWTs were chosen to fulfil the role as the work horse in large-scale wind energy production (Vassberg *et al.*, 2005).

The resurgence of VAWT research in the 2000's and 2010's employed modern computation power and sophisticated methods in attempts to address the outstanding issues with VAWT technology.

2.2 VAWT Aerodynamics

This section expands on the aforementioned complex features of VAWT flows, focussing on the following topics:

- Dynamic Stall
- The Self-Starting Problem
- Unsteady Winds/ Analogy with Flapping Wing (considered together)
- Flow Control Methods

2.2.1 Dynamic Stall

The aerodynamic effects of rapidly varying an aerofoils AoA were initially reported in 1932 by Kramer (1932) who investigated the effects of sudden gusts on increasing the maximum lift coefficient of a wing in otherwise steady flow (Wernert *et al.*, 1996). This phenomenon has since become known as dynamic stall, and much of the research on the topic has been connected with helicopter applications in the mid-late 1900's (Wang *et al.*, 2010). Dynamic stall simply refers to the more complicated stall process seen when considering a pitching aerofoil in transient flow, compared to one in a steady flow at constant AoA. The details of the flow features and mechanisms involved are complex, but simply put, the transient flow phenomena which appear when an aerofoil undergoes dynamic stall mean that lift can be maintained for longer, so that the aerofoil remains useful even after its steady flow stall angle is exceeded. Sharma & Visbal (2019) credit this to the presence of non-zero pitch rate causing a change to the effective camber of the aerofoil.

A range of research has been conducted on dynamic stall in general, providing much insight into the physical development of the flow features involved. Such research spans flow visualisation via experimentation, and detailed CFD analysis. Kim & Xie (2016) gave a comprehensive discussion on the nature of dynamic stall for HAWT blades, and a summary of the pertinent literature available at the time. Kim & Xie also used contour plots from their CFD campaign using Large Eddy Simulation (LES) to provide illustration of the flow processes such as the development of the Leading Edge Vortex (LEV), laminar separation bubble and boundary layer suppression.

For a VAWT at a high TSR, the AoA of a VAWT blade tends towards the fixing angle (constant) and the flow should be largely attached without stall. In such a case the blade velocity predominates over the wind velocity. With reducing TSR the stall point

will occur at some point as the AoA range diverges. Laneville & Vittecoq (1986) concluded that this point is equal or less than $TSR = 4$, while Ferreira *et al.* (2008) quotes $TSR < 5$. Stall means loss of desirable pressure distribution over the aerofoil and deterioration of overall turbine torque (Laneville & Vittecoq, 1986). Understanding the dynamic stall process is therefore a critical phenomenon in VAWT aerodynamics, specifically for the start-up phase of operation. Wang *et al.* (2010) showed that an oscillating aerofoil will have similar dynamic stall behaviour compared to a 2D VAWT blade. A recap will now be given of the dynamic stall sub-processes with reference to an aerofoil with sinusoidal oscillating AoA.

Wernert *et al.* (1996) characterised dynamic stall by four phases; i) Attached Flow, ii) Development of Vortex, iii) Post Stall Vortex Shedding, and iv) Flow Reattachment. For an aerofoil in a constant free stream, with sinusoidal pitching motion:

i) Attached Flow

- The pitching angle increases from zero, and the lift coefficient increases linearly with pitch. At small angles of attack the boundary layer is comparatively thin.
- At an incrementally greater AoA, the prevailing flow conditions remain for that of the previous (smaller) AoA, which has a thinner boundary layer. This inherited flow field then influences the flow to delay the thickening of the boundary layer. This is known as 'boundary layer suppression' (Kim & Xie, 2016).

ii) Development of a Vortex

- As the AoA becomes larger, the boundary layer thickens and eventually separates, but the aerofoil remains lifting even after the steady state critical AoA is exceeded. This is due to the development of a vortex at the leading edge (LEV). A recapturing of lift is observed due to the low pressure associated with the LEV on the suction side of the aerofoil (Wang *et al.*, 2010).
- The boundary layer suppression, and LEV both contribute to 'stall delay' (Kim & Xie, 2016).
- The LEV evolves and moves towards the trailing edge.

iii) Post Stall Vortex Shedding

- Upon detachment of the LEV from the trailing edge, a rapid drop of lift is seen as the aerofoil transitions into a deep stall (McAlister *et al.*, 1978).
- A second LEV in fact develops as well, but its effects are not as large as the first LEV. The second LEV can form after the first one has departed the trailing edge, and a slight recovery in lift is exhibited (Kim & Xie, 2016).
- At this point, the AoA is high and a vortex then develops from the trailing edge.
- The AoA reaches its maxima, and then begins reducing. All the vortex features are detached and convected downstream.

iv) Flow Reattachment

- The AoA continues to reduce and the flow begins to reattach, from the leading edge first, to the trailing edge (Wernert *et al.*, 1996).

The severity of the dynamic stall, i.e. the degree of its effects on the aerodynamic forces is grouped into two broad categories by Sharma & Visbal (2019). 'Light stall' means only a small area of separated flow is observed, which occurs when the maximum AoA does not greatly exceed the static stall angle. 'Deep stall' means a much larger region of separation and the presence of the large vortices described above. For light stall the separated region is said to be of size comparable to the aerofoil thickness. For deep stall the separated region is said to be of size comparable to aerofoil chord.

An important quantity considered with dynamic stall research of an oscillating blade is the Reduced Frequency, k , which is defined as (where c is the chord length and ω is the pitching frequency) (McAlister *et al.*, 1978):

$$k = \frac{\omega \cdot c}{2 \cdot U_{\infty}} \quad (\text{Eqn. 2.1}).$$

For VAWTs at high TSR, U_{∞} can be taken as the tangential velocity ($\omega \cdot R$) so that this can also be expressed as;

$$k = \frac{c}{2R} \quad (\text{Eqn. 2.2}).$$

A very small k means the wind velocity predominates over the rotational time period (Kim & Xie, 2016). This can be visualised as an aerofoil slowly increasing its AoA with the transient effects being small. The smaller the k value becomes, the more closely the stall process reflects a steady state one. Higher k values indicate higher rates of pitching, and under these conditions the unsteady flow features are more prominent and yield a greater stall delay influence. Even for relatively small reduced frequencies a noticeable stall delay can be seen (Lee & Gerontakos, 2004).

2.2.2 The Self-Starting Problem

HAWTs have a simple self-starting mechanism whereby the blades begin to rotate slowly from rest, under the predominating influence of drag. In this low TSR phase the blades are operating at a high AoA range because the blade velocity remains low compared to wind speed. The HAWT then undergoes rapid acceleration into its normal operating regime as the blades begin to generate lift effectively, and the subsequent increase in rotational speed further improves the angle of attack which generates progressively more lift etc. (which translates to turbine torque).

VAWTs do not have such a simple self-starting mechanism, and the self-starting problem has hindered progression of VAWT technology for decades. As discussed in Section 1.2, each blade produces a variation in thrust during each revolution - the challenge lies in designing an aerofoil which produces a net positive torque (per revolution) over a large range of AoA (at low TSRs). This challenge is complicated by the requirement for such an aerofoil to then operate effectively at higher TSRs once the turbine is up to speed. Wind turbines that are designed with emphasis on the maximum performance condition (generally high TSR) may neglect performance at the off-design conditions. Such an approach leads to aerofoils that struggle to achieve net positive torque when TSR is low during start-up. Some authors have declared that VAWT designs require starter motors in order to place them within a satisfactory TSR range where the turbine can become self-sustaining. Although this cannot be generalised to all VAWTs, such a possible drawback would weaken the VAWTs advantage of being a simple system compared to HAWTs. More importantly, the desirability for a VAWT and thus its route into the mass market hinges on ease of use and social acceptance of the technology.

The VAWT rotational speed initially increases before entering an idling period when the rotation speed remains reasonably constant (Worasinchai *et al.*, 2016). This is the phenomenon known as the 'plateau' or 'dead-band' region as a VAWT gathers speed very slowly. The VAWT self-starting problem is when (in some cases and conditions) a turbine is unable to further increase its rotational speed and escape the dead-band phase. If and when the dead-band phase is overcome, it is followed by a phase of quick acceleration as the turbine approaches its fully operational TSR where significant power can be produced (Baker, 1983).

Research has historically given conflicting conclusions on the VAWT self-starting behaviour. Investigations for NACA aerofoil based H-rotors (2 bladed) and for steady wind conditions have, in some cases, concluded that VAWTs cannot self-start and require external assistance (Zhu *et al.*, 2015). On the other hand, Hill *et al.* (2008) showed via experimentation, that a fixed pitch H-rotor can in fact self-start. Additionally, Dominy *et al.* (2006) showed that 3 bladed VAWTs can self-start in steady winds as well as 2 bladed VAWTs albeit in a narrower band of conditions.

An insightful approach to describing and modelling VAWT self-starting behaviour is proposed by Worasinchai *et al.* (2012), where the VAWT blade is analogised to a flapping wing. It is argued that the self-starting mechanism can be split into two phases; the combined lift-drag driven phase, and the fully lift driven phase which is only realised above $TSR = 1$. Worasinchai *et al.* (2012) state that self-starting torque relies on exploitation of flow unsteadiness which is in turn dependent on the chord-diameter ratio of the rotor (as well as number of blades and aspect ratio). Ferrer &

Montlaur (2015) described how successful self-starting behaviour is coupled with complex vortex generation associated with unsteadiness of VAWT flows. Accordingly, the structure of shed vortices required to nurture self-starting occurs only in certain conditions where a critical level of unsteadiness can be reached inside the rotor. Ferrer & Montlaur (2015) defined the self-starting process in terms of transition of the flow structure rather than in terms of reaching a designated TSR (as is the definition by other authors). Zhu *et al.* (2016) explores the influence of unsteady wind fluctuations on self-starting behaviour using 2D CFD and compares the results to experimental data of other authors. Zhu *et al.* found sensitivity to wind amplitude and fluctuation frequency which in some cases can improve the self-starting behaviour, concurring with the conclusions of Worasinchai *et al.* (2012).

The self-starting problem remains one of the main barriers to VAWT technology becoming widely deployed. Despite the conflicting conclusions in early 'self-starting' research based on steady winds and traditional aerofoils, it appears that unsteadiness may be a key factor which contributes to effective self-starting and is a topic which should be further researched.

2.2.3 Unsteady Wind/ Analogy with Flapping Wing

It is widely accepted that VAWT performance is less adversely affected by unsteadiness of wind speed and direction than HAWTs. Not only are VAWTs comparatively more suited to unsteady wind conditions, but it has been shown that unsteadiness may be a key factor for their efficient operation. The power delivery of VAWTs has been found to increase during unsteady winds of specific parameters, which allows improvement to both self-starting characteristics and maximum power output. Zhu *et al.* (2016) found that high frequency, small amplitude fluctuations to the free stream velocity can produce VAWT performance improvements due to delayed stall (although frequency variation is said to play a smaller role). Large amplitude fluctuations can however reduce performance as this shifts the flow into a drag dominated regime (Danao *et al.*, Vol 116, 2013). Scheurich & Brown (2012) presented an investigation into performance and wake dynamics of a VAWT in steady and unsteady winds, covering 3 VAWT configurations. It was found that the helical blade configuration experienced less power loss due to fluctuations compared to straight blades. The straight and curved VAWTs, exhibited steeper gradients of power coefficient (C_P) to TSR in the mid-operating range. Blade curvature and twist must therefore be carefully chosen.

Danao *et al.* (Vol. 116, 2013) performed numerical simulations to investigate the effects of unsteady wind at various amplitudes and frequencies. In a separate work,

Danao *et al.* (Vol 107, 2013) conducted experimental investigations on the same theme, with results demonstrating that the performance of a VAWT in unsteady wind does not follow the steady C_p curves, calling into question simplified analyses of steady flows in VAWT research. Similar conclusions to Zhu *et al.* (2016) have been drawn (regarding effects of small amplitude and high frequency wind fluctuations on turbine performance) by Scheurich & Brown (2012), and Danao *et al.* (Vol 107, 2013). The experimental work of Danao *et al.* (Vol 107, 2013) however saw a drop in performance for unsteady winds. It can be observed therefore, that unsteady winds are only beneficial to turbine performance within certain fluctuation parameters.

An insightful way of presenting VAWT aerodynamics is by analogy to a flapping wing, as both possess periodic behaviour in onset velocity and angle of attack. Zhu *et al.* (2016) discussed works on the topic of bird flight, indicating that birds utilise the delayed stall mechanism present in fluctuating flow for advantages in thrust. Gorelov (2009) demonstrates the VAWT/flapping-wing analogy and suggests that Darrieus VAWTs should be analysed via nonlinear theory of a wing in unsteady flow. The inherently unsteady flow in a Darrieus VAWT produces thrust according to the wind fluctuations in amplitude and frequency (Bhutta *et al.*, 2011).

2.2.4 Flow Control Methods

The approaches for improving VAWT performance via flow control can be grouped into two main categories, active and passive control. Passive control methods use various insightful modifications of a blade's geometry to better deal with the wide range of onset angles. Mohamed *et al.* (2020) studied the effect of slotted aerofoil geometries on VAWT performance where the bypassing flow in the slot was hypothesised to aid delay of boundary layer separation. Leading edge serrations (inspired by marine animals) have been another novel method of flow control under investigation. Wang & Zhuang (2017) found that serrations reduced flow separation over a significant portion of the VAWT revolution and that the average power coefficient was increased, Guide-Vane Shrouds are another passive control method; similar to gas turbine engines using a ring of fixed stator vanes to guide the flow into the central rotor (Jin *et al.*, 2014). Passive control also includes the use of compliant surfaces that react only to the aerodynamic forces on them and are not actuated by a control system.

Active control methods continually enforce changes to blade geometry in some manner, whether via traditional control surfaces (ailerons, slats, flaps), blade pitching, morphing surfaces, and even boundary layer suction/blowing.

Velasco *et al.* (2017) applied synthetic jets to influence the VAWT blade flow field. They found that applying the jets to the inside surface of the blade improves the

power coefficient in the upwind region of the VAWT revolution. This is because the flow separation that would normally develop on that surface is mitigated by the energising of the flow from the jets. In the downwind region, applying jets to the outer surface was found to delay flow detachment under the same rationale. Furthermore, Velasco *et al.* (2017) concluded that the reduction of flow separation helped improve the turbine power coefficient by more than the power requirement to drive the jets. The contrasting approach to boundary layer blowing has also been shown to be successful by Rezaeiha *et al.* (2019). Rezaeiha *et al.* (2019) shows that leading edge suction can be used to prevent the laminar separation bubble from bursting and thus delaying/avoiding separation, and this approach yielded large increases in the power coefficient.

Bianchini *et al.* (2019) conducted numerical investigations into the effects that Gurney flaps can have on a rotating aerofoil where stall delay and an increase to the power coefficient could be observed. This study was extended to VAWT blades where significant improvements to the downwind power extraction could be produced, especially at low TSR. Similar observations were made by Yan *et al.* (2020) using a similar approach to implementing Gurney flaps.

Zhuang *et al.* (2020) applied a morphing trailing edge technology to VAWT blades. This approach was found to greatly reduce the severity of fluctuating aerodynamic loads (having a positive influence on structural fatigue life) by continually altering the blades camber. Similarly, Ai *et al.* (2019) found through experimental campaigns that morphing blade systems can dramatically reduce aerodynamic loading, as well as improving VAWT performance. As a brief tangent from the pure-aerodynamic discussion, it is worth further emphasising the fatigue issue since fatigue damage was responsible for the failures of the early VAWT demonstrator projects mentioned in Chapter 1. Fatigue caused by large load fluctuation continues to be a problem in VAWT design today, especially if fewer blades are used (Zhuang *et al.*, 2020). Some authors have studied the reduction of these fluctuating loads by the use of control surfaces and morphing geometries using 'smart rotor control' systems, such as Zhang *et al.* (2015), and Van-Wingerden *et al.*, (2008).

The most commonly researched active control method for VAWT applications is variable pitch blades. Variable blade pitch systems are the norm for large scale HAWT wind farms, and such systems are indeed being researched for application to VAWTs (Zhang *et al.*, 2014). Xisto *et al.* (2014) proposed a cycloidal VAWT which has blades that self-adjust their pitch via a mechanical system, which is claimed to improve self-starting performance as well as operating efficiency at lower TSRs - this is due to a reduction in the AoA range, thus mitigating dynamic stall effects. Zhao *et al.* (2018) proposed that a pitching regime which improves torque over a greater

range of azimuthal angles, will outperform one that maximises the highest torque - they reach an 18.9% increase in the maximum power coefficient but do not quote average C_P improvements. Macphee & Beyene (2012) review VAWT technology including variable pitch systems and conclude that increases in power of up to 30% can be achieved by this type of control method. A comprehensive review of the pitching VAWT literature is given by El Sakka (2020).

Variable pitching systems could offer notable improvement to the self-starting behaviour and power coefficient for VAWTs, but the robustness of the systems and associated capital costs and operations & maintenance costs will be a significant challenge to overcome for commercial turbines. In addition, mechanically controlled pitching systems require power themselves and increase the parasitic drag of the turbine due to linkages and connecting arms etc. Remotely controlled actuation is an option that removes mechanical complexity associated with control mechanisms. In the same way that fly-by-wire systems substantially improved reliability and performance of aircraft control, a similar approach could be instrumental in VAWT technology reaching maturity. Servomotors at the blades can be commanded by one central computer controller, as investigated by Zhang *et al.* (2012) who claimed significant improvements to the power coefficient using centrally controlled blade pitching.

Tjiu *et al.* (2014) noted that complicated straight blade VAWT configurations (whether actively or passively controlled) such as variable geometry (Musgrave rotor), variable pitch straight bladed, Diamond-rotor, Delta-rotor and V/Y-Rotor variations have not been heavily developed due to low economic value. At present many of these systems are not considered feasible for cost competitiveness in real world markets by other authors (Jin *et al.*, 2014). However, it is important not to discount these options from future VAWT research. Some degree of design complexity is inevitable in order to provide high performance VAWT designs. An optimal balance between a well performing technical solution and a simple/affordable one is yet to be determined.

2.3 Simplified Models (Computational Aerodynamics)

The early research into VAWTs took place before the advent of advanced computational methods. The mathematical models developed to predict VAWT aerodynamics were supplemented by experimental research, but were nevertheless limited in complexity and could not accurately predict the intricacies of the flow phenomena (complex vortex behaviour, hysteresis, unsteadiness, and dynamic stall etc.). The main modelling methods applied to VAWTs are grouped broadly into three categories; Momentum Models, Vortex Models, and Cascade Models (Jin *et al.*, 2014). These were attractive options in the 1970's and subsequent decades, due to the cost/time advantages compared to conducting experimentation. Momentum methods were accurate only for very simplified problems, i.e., inviscid flows and for lightly load turbines. There were also limitations on solidity and tip speed ratio for which the model accuracy held (Jin *et al.*, 2014). Vortex methods (which provide modelling of potential flows) tend to neglect stall effects which are prevalent in VAWT flows. Such methods were therefore not particularly useful for VAWT applications (Jin *et al.*, 2014).

The unsteady characteristics and complex dynamic vortex and wake structures are difficult to predict with accuracy using these mathematical modelling techniques (Edwards, 2012). Despite continuous research and development of these aerodynamic models, they were largely superseded by sophisticated experimentation and the evolving CFD methods which allow better quantification (via measurement or numerical calculation) of the detailed flow physics.

2.4 CFD and VAWTs

The early mathematical models formulating the traditional aerodynamics approach struggle to reflect VAWT flow complexities, and more recent sophisticated CFD has allowed improved understanding and design methods. Alongside the CFD, crucial validation data has been provided via advanced experimentation techniques such as PIV (Particle Image Velocimetry). Despite the advances in computation power, it can still be costly to compute VAWT flow fields, but computation cost may be a necessary evil in order to produce efficient designs and unlock the potential of the technology.

VAWT CFD research has predominantly consisted of 2D simulations using RANS (Reynolds Averaged Navier-Stokes) turbulence models to alleviate high computational costs (Balduzzi *et al.*, 2015), in conjunction with validation of the CFD results against experimental data. The governing equations which are resolved by the CFD software are presented in Appendix E.

The validity of 2D vs 3D models for VAWT analysis is postulated by authors across the field and it is generally agreed that 2D analyses commonly overestimate the power coefficient. Untaroiu *et al.* (2010) compared experimental data with CFD to validate 2D VAWT simulations, concluding that a transient 2D analysis can determine the operating speed with errors less than 12%. Untaroiu *et al.* (2010) suggested that 2D simulations can be preferable over 3D so long as a suitable flow model is chosen, due to the implications of degraded mesh quality and time cost associated for 3D models. Other authors have come to similar conclusions by modelling 2D flow fields and comparing the results to experimentation, such as Lanzafame *et al.* (2013).

Contrary to the above, Howell *et al.* (2009), Almohammadi *et al.* (2015), and Jin *et al.* (2014) each found that 2D simulations drastically over predicted the performance of a VAWT compared to both 3D simulations and experimental results. This discrepancy is attributed to the over tip vortices which occur in reality and cannot be predicted by 2D simulations. Tip vortices are not solely responsible for the variance; the overall performance of a turbine blade depends on the complex pressure field which varies along the span-wise direction (third dimension) in a complex fashion (Li *et al.*, 2016). The structural connecting arms also greatly influence the span-wise flow field around the blades, which are generally not modelled in 2D investigations. Regarding the H-rotor Darrieus configuration, these suffer both from large connecting arms and over-tip vortices at the ends of each blade. These factors make for greater loss of accuracy in 2D simulations when predicting overall power, but 2D can still provide a good agreement with 3D predictions at the mid-span of the blades (El Sakka, 2020). Other configurations such as the egg-beater have lower flow velocities and losses at the blade tips, and

less onerous connecting arms making their 2D analysis less affected by these aspects. As a counterpoint, the egg-beater configuration (and others) does not have a constant geometry in the vertical direction, so multiple 2D models would be required to model the various span locations.

Edwards *et al.* (2013) observed a discrepancy between the 2D and 3D results, and also found that their CFD predicted stall as a separation moving forward from the trailing edge, while PIV experimentation (as shown elsewhere) exhibits stall as a sudden leading edge vortex moving towards the trailing edge. The simulations also predicted delay in reattachment of the flow. Such disagreement between CFD and experimental data could be due to poor choice of mesh quality or other settings in the simulation. Balduzzi *et al.* (2015) set out to determine the most appropriate CFD settings to employ in order to get best reflection of reality when doing 2D simulations for H-rotors. Balduzzi *et al.* (2015) reviewed VAWT research papers across the field to identify trends in the CFD settings. Balduzzi *et al.* (2015) stated that 3D analysis is needed if one requires the truest prediction of VAWTs power output; but that 2D, if setup correctly, is able to provide comprehensive analysis of Darrieus VAWT flow phenomena. El Sakka (2020) studied the differences between 2D and 3D model accuracy and concluded that 2D simulations provide good enough agreement with experimental data to be useful in design and optimisation studies. With design/optimisation of a large numbers of variables the consideration of computation cost becomes even more critical making 3D likely to be prohibitively computationally expensive. This thesis utilises 2D simulations only, aligning with this logic and consideration of the literature.

Once choices have been made regarding 2D/3D modelling, the next key decision is the selection of a suitable turbulence model. Untaroiu *et al.* (2010) concluded that the solution is much more affected by the choice of flow model than the difference between a 2D and 3D simulation. Some turbulence models are unsuited for modelling the specific aerodynamic characteristics of VAWTs to a sufficient level of accuracy (Buchner *et al.*, 2015), but throughout the literature there is disagreement about which models can be accurate enough.

The widely used RANS turbulence models, such as the Spalart-Allmaras (S-A), $k-\omega$ models, and $k-\epsilon$ models can offer reasonable computation costs while supplying sufficiently accurate results for many applications where information on time averaged flow properties is sufficient (Versteeg & Malalsekera, 2007). The RANS turbulence models predict mean flow properties, and do not capture the turbulent fluctuations/fine motions of the turbulent eddies. The validity and accuracy of these models for application to VAWT flow physics is therefore debated across the literature.

Dhert *et al.* (2016) used the S-A model, which is a one equation model, for CFD based optimisation of a HAWT. The justification for this choice of model was that the HAWT is able to operate with largely attached flow conditions, and thereby not overstepping the limitations of the turbulence model. Hand *et al.* (2017) stated that the S-A model can predict separation and reattachment as accurately as some of the two equation models but for 80% less computational cost. On the other hand, Standish and van Dam (2003) investigated HAWTs using the S-A model with concluding remarks that it could not accurately describe the flow in the stalled regime.

The two equation models ($k-\epsilon$ and $k-\omega$) have been used with some success, with most authors choosing various augmented versions of these methods. Edwards *et al.* (2013) discussed how well various models that are available within ANSYS Fluent (version 12.1) can predict VAWT dynamic stall. By comparison with experimental data sets, it was concluded that the SST (Shear-Stress Transport) $k-\omega$ model most accurately reflected the results compared with the S-A and RNG (Re-Normalisation Group) $k-\epsilon$ models. The choice of SST $k-\omega$ model was further validated by comparison with PIV measurements. This conclusion was seconded by Balduzzi *et al.* (2015) who studied various RANS models finding that the $k-\epsilon$ models (Standard and RNG) could not provide a converged solution. Howell *et al.* (2009) however employed the RNG $k-\epsilon$ model and noted that the standard $k-\epsilon$ model did not sufficiently model flow separation. Hand *et al.* (2017) condemned ϵ -based models and favours the SST $k-\omega$ compared to the other 2-equation models. Wang *et al.* (2010) stated that across the literature, most RANS models (such as B-L, RNG $k-\epsilon$, S-A) cannot provide adequate accuracy and agreement with experimental VAWT data. Wang *et al.* described $k-\epsilon$ models as insensitive to adverse pressure gradients which makes them poor at modelling boundary layer separation such as for a stalled blade, or indeed for external flows in general.

The $k-\omega$ models are superior at modelling flows under an adverse pressure gradient, and the subsequent flow separation – the BSL (Menter Baseline) or SST variants of this model are recommended to achieve the best accuracy. This position is demonstrated by Wang *et al.* (2010) who conducted 2D simulations of an oscillating aerofoil in dynamic stall conditions, using the standard $k-\omega$ and the SST $k-\omega$. The conclusion was that the standard model gave poor data matching to the experiment, but that the SST variant could capture the complex flow structures giving good data agreement in general except at high angles of attack. Wang *et al.* (2010) focussed on low Reynolds number flows, and therefore a demanding flow regime of great relevance to VAWTs. Rezaeiha *et al.*, (2019) thoroughly investigated a range of turbulence models; S-A, realizable $k-\epsilon$, RNG $k-\epsilon$, $k-\kappa-\omega$, $k-\omega$ SST, $k-\omega$ SST with

intermittency, and Transitional SST. Various flow conditions were examined, and the SST models were found to adequately represent VAWT flow dynamics including dynamic stall.

Buchner *et al.* (2015) explored dynamic stall prediction of the Menter-SST model by comparing CFD results to experimental data obtained via Stereoscopic Particle Image Velocimetry (SPIV) for low/moderate Reynolds number flows. The findings showed that for a variety of TSRs the RANS model gave good agreement with SPIV results, including over the stalled regime. Buchner *et al.* concluded that the Menter-SST model does not capture some large scale vortex effects but still offers “*sufficient computational efficiency*” for VAWT applications. Lanzafame *et al.* (2013) experimentally validated the choice of the SST-Transition model (as opposed to the SST k- ω model) for an unsteady CFD campaign. Lanzafame *et al.* (2013) claimed that the SST Transition model is more appropriate for VAWT flows due to the advantages in modelling the laminar/turbulent transition that the model provides.

Ferreira *et al.* (2009) used PIV experiments to decide the choice of turbulence model best suited for VAWTs. In this work Ferreira stated that RANS models (such as S-A and k- ϵ) were unable to accurately model the leading edge vortices and trailing edge wake behaviour, and instead recommends the Detached Eddy Simulation (DES) model. DES is a hybrid method between RANS methods and Large Eddy Simulation (LES). DES enjoys a reduction in computation cost compared to the costly LES, while modelling the vortex behaviour more accurately than RANS methods. A similar viewpoint is taken by Kim & Xie (2016), advocating LES with the argument that RANS models are not sufficient to model transient flows. Dhert *et al.* (2016) however condemns LES as being too costly for practical use in optimisation studies, indeed LES and DES are not largely used across the literature due to cost reasons (Hand *et al.*, 2017). Wang *et al.* (2010) defends use of RANS next to DES and LES for the same reason of balancing computation cost with necessary accuracy. Balduzzi *et al.* (2015) conducted a review of VAWT literature concluding that no agreement could be made regarding choice of turbulence model, and subsequently selected the SST k- ω model.

The work performed in this thesis employed the SST k- ω model - The equations defining this model are presented in Appendix E. This is the most accurate model for VAWT applications that is also supported by ANSYS Fluent’s Adjoint solver. Interlinked with the choice of turbulence model is an appropriate meshing strategy that will allow sufficient accuracy near the wall and in areas of possible separation. In this thesis a refined wall treatment is employed to provide high near-wall

resolution; such details and further discussion of the adopted meshing strategy are deferred to Section 5.1.2,

Balduzzi *et al.* (2015) also reviewed other important 2D CFD settings with the aim of providing a summary of best practice for VAWT modelling. Across the literature, Balduzzi *et al.* (2015) found that the SIMPLE pressure-velocity coupling algorithm is the most frequently used, and that for smaller time steps the SIMPLE and Coupled schemes were similar in prediction of C_T against AoA, but that for coarser time steps the SIMPLE algorithm exhibited over prediction. Balduzzi *et al.* (2015) disregarded the PISO method as it proved to be the least accurate. Balduzzi *et al.*'s survey found the discretisation method generally preferred in the literature was the Second Order Upwind scheme.

The above review served to aid in constructing the CFD simulations. Indeed, there are many factors influencing the quality of CFD results and Table 2-1 summarises these along with referrals to chapters of the thesis where that factor is discussed in more details. Table 5-2 (for Single-Blade models) and Table 6-1 (for VAWT model) summarise many of the choices adopted in this thesis regarding CFD setup.

Table 2-1 – Summary of Key Factors Affecting CFD Simulation Results

Setup Factor	Reference section(s) within Thesis
2D/3D	2.4
Turbulence Model	2.4, 5.1.2
Mesh (type/quality/refinement)	5.1.2, 5.1.3, Appendix A, 6.1, Appendix B
Domain (size/shape)	5.1.3, Appendix A, 6.1, Appendix B
Time step size	5.1.2, 5.1.3, Appendix A, 6.1, Appendix B
Target CFL	5.1.2, 5.1.3, Appendix A, 6.1, Appendix B
Number of cycles for periodic convergence	5.1.3, Appendix A, 6.1, Appendix B
Solver settings:	
Solver type (density/pressure based)	5.1.2
Pressure-velocity coupling scheme	5.1.2
Spatial discretisation scheme	5.1.2
Temporal discretisation scheme	5.1.2
Number of iterations per time step	5.1.2
Convergence criteria (residuals)	5.1.2

2.5 Experimental Study of VAWTs

Historically, there are two main kinds of experimental method used to investigate VAWT aerodynamics; direct measurement (torque measurement or pressure/velocity measurements), and flow visualisation (smoke/thread/digital flow visualisation). Early experimental campaigns used basic versions of both of these in order to aid a better understanding of the fundamental flow phenomena involved in VAWTs. Also, parametric studies were made towards defining effects of blade geometry on performance for simple configurations such as H-rotors. Limitations in the equipment and techniques constrained the level of insight that was achievable at the time. Vittecoq & Laneville (1983) conducted an experimental campaign to determine instantaneous normal and tangential forces of a blade VAWT across a range of fixing angles and TSRs. Although such a data set is useful, approaches such as measuring surface pressures or exerted torques and quantifying turbine performance directly do not provide insight into the complexities of the flow field. As more advanced experimental instrumentation has appeared in recent years (such as PIV, and blades with engrained pressure transducers), we are gaining further understanding of the underlying, complex flow phenomena. More recent PIV works quantify the flow field around the turbine by capturing movement of seeded particles, using specialist laser and camera equipment. PIV campaigns have shed light on the detailed aerodynamics, particularly with regard to vortex structures and features of dynamic stall. Ferreira *et al.* (2008) conducted a 2D PIV investigation of dynamic stall to aid with validation of their numerical/CFD models. The study detailed the vorticity behaviour coupled with flow separation and vortex shedding. In subsequent works, Ferreira *et al.* (2009) used PIV experimentation to assess the accuracy of a range of turbulence models for modelling VAWT flows.

Fujisawa & Shibuya (2000) also investigated dynamic stall using PIV; the results form a detailed illustrated commentary on the transient flow mechanisms. Edwards *et al.* (2017) used PIV in conjunction with direct performance measurements to validate a 2D CFD model. Howell *et al.* (2009) investigated the performance variations of a VAWT under different conditions and solidities using wind tunnel experiments as well as 2D and 3D CFD to provide insight into the detailed aerodynamics. Eboibi *et al.* (2016) focussed a PIV study specifically on the effects of solidity on VAWT performance. A large portion of experimental research performed for VAWTs covers only steady wind conditions; this is coupled with the fact that most historic CFD conducted was based on steady winds which required appropriate validation using steady wind experiments. Although some examples exist there remains room for valuable further PIV research in the future for VAWTs in unsteady winds.

3 AERODYNAMICS OPTIMISATION OF VAWTS

This chapter starts with a description of the general optimisation methods, before exploring CFD based shape optimisation techniques. The Adjoint method is then highlighted for both its power in determining shape sensitivity gradients, and for its lack of use in the VAWT literature. The Adjoint method is then discussed along with how optimisation frameworks using Adjoint sensitivity gradients can be formulated in CFD. The choice of commercial CFD software ANSYS Fluent is the platform of this thesis and this is justified, as well as a discussion of the technical limitations adopted by making this choice.

The problem of VAWT optimisation can be approached in many different ways. The designer can firstly consider the blades as fixed (to the connecting arms), and study how geometrical changes can affect the performance. With this approach the significant challenge of producing an aerofoil shape that performs well in both upwind and downwind regions is unabated. Section 2.2 has described this challenge, and how upwind improvements tend to carry penalties in the downwind region and vice versa. Nevertheless, there are many techniques available to the fixed blade VAWT designer that can offer routes to improving VAWT performance, as is discussed in this section. To address this upwind/downwind design paradox however, many authors have also explored active or passive control systems to continually change the angle or shape of the VAWT blades during the revolution. In this fashion the worst negative-power parts of the revolution can be mitigated, and best advantage can be taken during the greatest power producing parts of the cycle. The key drawbacks to the control-approaches are that the mechanical systems required can be complicated and heavy, with new modes of mechanical/structural failure introduced into the VAWT. Furthermore, it is a challenge to know what geometrical profiles should be pursued by these control systems. The most commonly discussed VAWT control system for aerodynamic optimisation is the variable-pitch technique. Details of this and other techniques are given in Section 2.2.4, but such approaches have been found to improve the VAWT power coefficient by significant margins. Despite the potential benefits to aerodynamic performance, the complexity of these systems detracts from the elegance and simplicity of the VAWT concept which can operate in fluctuating winds environments without any yaw or pitch-control systems being required. Active control systems also consume power to operate as well as adding significant costs. It is important to note however that a fixed blade can only be of the optimum geometry for a single TSR. Therefore, control mechanisms can offer aerodynamic benefits not only at a single optimum TSR, but across the entire operating range of a VAWT and this includes self-starting. Nevertheless, discussion is now focussed more specifically on techniques used to

tackle the optimisation problem of fixed VAWT blade geometries. Such techniques could indeed be used in conjunction with control mechanisms to gain further increases to VAWT power coefficient.

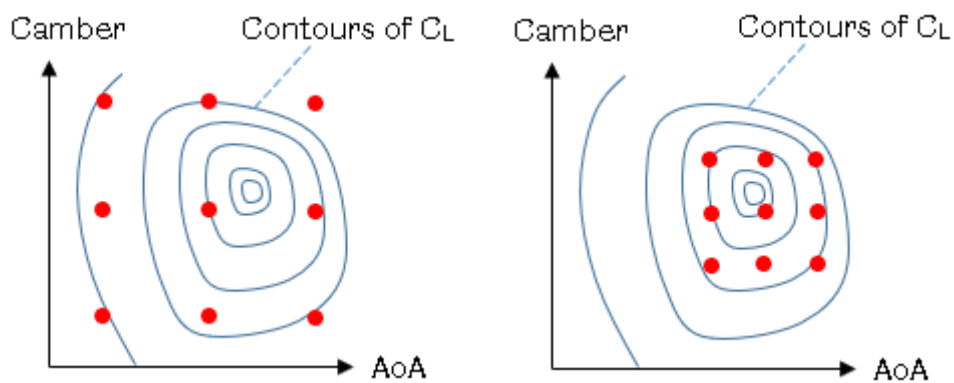
In general, optimisation problems revolve around pursuing maxima or minima (extrema) of a designated objective function (sometimes referred to as the “observable”), i.e. maximising lift or minimising drag for an aerodynamic design problem. A field of constraints is defined which the optimised solution cannot violate. This set of constraints should replicate to a reasonable degree, the real-world requirements from other engineering disciplines, such as structural limitations, manufacturing feasibility, capital costs, and operation & maintenance cost, etc. In some cases, a pure aerodynamic optimisation process may determine a solution which is unachievable in reality due to over relaxed or poorly defined constraints (perhaps an aerodynamically efficient turbine geometry that could not be manufactured for a reasonable cost). The inputs, or design variables in the VAWT aerodynamic optimisation problem are defined as the geometry of the blades.

Ning *et al.* (2014) provides a broad discussion of optimisation problems with reference to wind turbines, examining the implications of problem definition (setting the objectives and constraint functions) on the final solution. Ning *et al.* (2014) concludes that maximising annual energy production of a turbine is detrimental to the cost of energy for that system solution. Choosing to minimise the cost of energy is then discussed as an objective function, accounting for structural performance as well as system costs. Consideration of the full spectrum of constraints over the wide range of engineering disciplines involved in a wind turbine project is unfeasible for the scope of most academic research. Nevertheless, these considerations are useful to keep in mind when embarking into the literature of high-fidelity aerodynamic optimisation, which can lose sight of the basic objective – to make wind energy more cost effective and thereby more widely used. A simpler set of optimisation objectives/constraints is adopted in this thesis as will be discussed in due course.

Once the optimisation problem is constructed with the objective selected, the method of optimisation is chosen. State of the art aerodynamic optimisation methods combine the disciplines of applied mathematics, fluid dynamics and computer science. The optimisation method constitutes the mathematical formulation and algorithmic processes that conducts the necessary calculations/computations. A range of optimisation methods exist that have been applied to aerodynamics problems – different methods are suited to different types of problem. Broadly, these can be split into *gradient-free* and *gradient-based* methods. Gradient-free methods calculate the performance of the system for

various arbitrary changes to the input parameters and use the results to identify more promising configurations to be tested in the next iteration (Funke *et al.*, 2013). These methods can search for global maxima/minima of the objective function, but a large cost is associated due to the time wasted from many poor system configurations being disregarded along the way.

The most rudimentary optimisation process is the Design of Experiments (DoE) approach (which is gradient-free), whereby a selection of input parameters are chosen as variables and then the system performance is evaluated after making changes to these parameters – only a small number of input parameters can be investigated within a feasible time scale. This is illustrated in Figure 3-1 which shows a hypothetical case where maximising an aerofoil lift coefficient is pursued using two parameters, camber and AoA. By inspecting results of the initial set of tests (each red dot represents a test/CFD solution) the designer can produce a more refined set. This requires many tests in total.



**Figure 3-1 – Optimisation by Design of Experiments Approach
(Left) Initial Search, (Right) Refined Search**

Such a method is partly a trial and error process because the parameter changes are made in a semi-arbitrary fashion, without knowing what the effect will be on performance. DoE methods are simple to implement but are inefficient and the outcomes are highly subject to the choices of parameter variables. Bianchini *et al.* (2014) is a good example of how much effort is required for a DoE approach, showing a vast number of design cases being tested and very large amounts of CPU time. Such methods may not be appealing to a prospective designer looking for a fast and effective design methodology. Lots of other DoE literature is presented as 'parametric studies' where standard aerofoil parameters are explored to observe their impact on performance such as Gosselin *et al.* (2016), and much of the referenced material in Section 1.2.1 falls into this category. There are many publications using DoE/parametric-studies which spread between many different turbines with a range of operating conditions, so it is hard to quantify the efficiency

of this approach in general - but it can be observed that a large amount of computational time is required.

The discussion of literature examining shape parameters of fixed-pitch VAWT blades has been already mentioned in Section 1.2.1. This is not repeated, but such work utilises the DoE approach and makes up a significant proportion of the literature in aiming to characterise the geometrical trends and optimise VAWT blades via geometrical changes. In this thesis attention is given to alternative and more sophisticated methodologies than DoE. Two types of gradient-free optimisation method which are more sophisticated than DoE, and also commonly found in applications to aerodynamics, are Evolutionary Algorithms (including Genetic Algorithms) and Response Surface Methods (RSM) (Edwards, 2012). Genetic Algorithms have been used extensively in VAWT literature and work by producing a multitude of semi-random design variations and eliminating those that evaluate poorly against the objective function. The remaining designs are combined to produce a new subset of competing designs and the process of elimination and breeding continues. Mutations in the designs are introduced at each breeding stage, and the process ends when the performance convergence is approached. Such methods can offer global maxima/minima and can deal with discrete variables (such as the number of turbine blades) but are limited to a small number of design input variables, in order to keep the computational costs acceptable (Coppin, 2014). A further limitation is that the resulting blade shape is implicitly constrained by the way in which the geometry is parameterised. The typical parameterisation techniques used in the pertinent research are Parametric Section (PARSEC) and Class/Shape Function Transformation (CST), the details of which are not discussed in the present work. Ram *et al.* (2012) gave a brief account of various blade parameterisation methods.

Chehouri *et al.* (2015) conducted a review of optimisation techniques relating to the multi-discipline optimisation of a HAWT, highlighting the importance of cost, and structural aspects. Orman & Durmus (2016) observed improvements of ~15% to the lift/drag ratio for HAWT blades after 200 geometry generations (starting from NACA2411 blades). Similar optimisations were repeated with two variants of geometry parameterization method (namely CST and PARSEC, with CST producing the greatest improvement). He & Agarwal (2014) optimised the NREL (National Renewable Energy Laboratory) HAWT with the purpose of demonstrating a multi-objective optimisation algorithm. The work was performed using ANSYS Fluent to calculate the flow field with a k- ω turbulence model. The authors concluded that the results of the HAWT optimisation matched well with results obtained by an Adjoint

optimisation by Ritlop & Nadarajah (2009). The discussion by He and Agarwal (2014) does not however mention the computation time required to achieve this solution.

Daroczy *et al.* (2018) used genetic algorithms (with PARSEC) in conjunction with 2D CFD (with the k- ω SST turbulence model) to maximise the energy output of a VAWT using NACA0021 blades as the starting geometry. Approximately 20% improvement in C_P (at optimum TSR) was made for the optimised geometry, which was reached after 900 geometry variants and a colossal 200,000 hours of computational time. Liang & Li (2018) used genetic algorithms (CST) with 3D CFD (with the k- ω SST turbulence model) to optimise a VAWT with NACA0015 blade geometry. Across a range of TSRs the largest improvement exhibited in C_P was 7%. Similar improvements were reached by Carrigan *et al.* (2011) using a differential evolutionary algorithm; 6% gains in efficiency were made to the VAWT with NACA0015 baseline blades. Chern *et al.* (2021) used genetic algorithms to produce improvements to the tangential force coefficient of a similar order (~6%). Their findings showed that the original aerofoil shapes were already close to the optimum one delivered by the optimisation process. Ma *et al.* (2018) achieved 27% improvement to the power coefficient by applying a multi-island genetic algorithm to a high solidity 3-bladed turbine of NACA0018 baseline blades. Ferreira & Geurts (2014) conducted an aerodynamic/structural multi-discipline optimisation using genetic algorithms. Zhiqiang *et al.* (2016) used genetic algorithms to do multi-point optimisation considering multiple-AoA based design methods and applying weightings to the objective function over several AoA. Genetic algorithms have also been used for optimising the blade pitching angle of VAWTs (Li *et al.*, 2018), fouled blade performance (Ram *et al.*, 2013), and Savonius turbine blade shapes (Chan *et al.*, 2018). It is also common in genetic algorithm research for XFOIL (or similar) software to be used as a precursor step to the CFD. This approach applies geometrical changes and evaluation of the aerofoils within XFOIL, and although accuracy is lost this can greatly reduce run times compared to using CFD throughout the entire process (Ferreira & Guerts, 2014), (Liang & Li, 2018).

Vavalle & Qin (2007) proposed a Response Surface Method (RSM) rather than Genetic Algorithms or Gradient Based methods (for aerospace applications). In RSMs a range of design cases are first tested on a preliminary basis. This provides numerous evaluations of the objective function across the solution space. This data is then converted into an approximate response surface which estimates the objective function continuously across the design space. In other words, the response surface describes the sensitivity of the designs performance to variations in the input variables. This response surface therefore holds the key information needed for optimisation, but in the form of a smooth analytical surface which can be

manipulated mathematically. Classical optimisation procedures can then be applied to the analytical surface to provide a problem solution (Le Moigne, 2003). The validity of the global solution depends on the resolution of the initial data sampling, which infers a proportional computational cost. The process is repeated after each modification to the design, whereby a subsequent more refined response surface is constructed.

El Sakka (2020) applied RSM to VAWTs, using both 2D and 3D models as a basis for the optimisation. For the 2D basis a 3% increase in the power coefficient was reached, while for the 3D basis ~35% was achieved but this was largely due to changes in the blade aspect ratio. Kear *et al.* (2016) achieved a 23% improvement to the C_P using RSM however the turbine was a novel hybrid Savonius/Darrieus design which may have had a very suboptimal design to begin with.

With consideration of the above, there is clearly a huge variation in what improvements are achieved using different types of method. The potential improvements depend not only on the optimisation method selected but also on the specific case being investigated itself. Indeed, some baseline turbines might have an unsuitably poor blade geometry to begin with, thus allowing room for significant improvements to be measured after optimisation. Conversely, for some turbines the baseline aerofoil may already be close to the optimum shape, making the possible improvement margins of any optimisation process smaller.

The Gradient Free methods described so far do not explicitly use flow field data to determine the appropriate change in geometry for the next iteration. This family of methods has generally been used in the existing VAWT literature to date. On the other hand, Gradient Based methods appear more intelligent as they quantify how a change to each input parameter would affect the objective function, i.e., how a small blade shape change would affect the lift or drag. The input parameters (geometry) are then updated using this objective-gradient in a way which offers a definitive system improvement, so this minimises wasted computational effort. Indeed, Gradient Based methods have been used with much success for aerospace applications.

The sensitivity gradients are determined using the flow field data. To further describe the meaning of the term “*sensitivity gradients*”, consider a CFD flow solution of a general aerofoil in steady flow. The shape of the aerofoil (for a fixed set of boundary conditions and flow properties) determines the resulting flow field and thereby the value of the objective function (such as lift or drag) for this aerofoil. Any small change in the aerofoil shape will result in an associated change to the objective function. The proportion to which the objective function changes shows how

sensitive it is to modification of a certain shape parameter. In other words, sensitivity gradients describe how the performance will vary for a given change in the blade geometry. Gradient based methods therefore make intelligent changes to the aerofoil based on consideration of these sensitivities.

A gradient based optimisation process is performed iteratively. After the sensitivity gradients are calculated they are used to propose the next new geometry (for the subsequent iteration), the new performance is then calculated, and the sensitivity gradients are re-calculated. Figure 3-2 illustrates this using the same hypothetical case as shown in Figure 3-1.

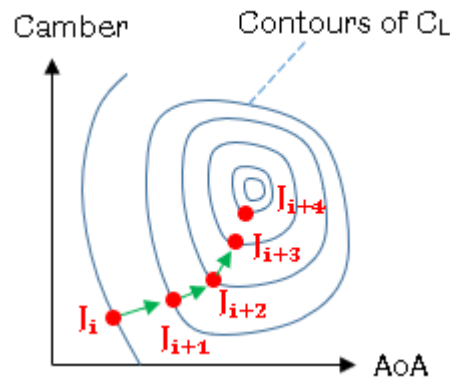


Figure 3-2 - Optimisation with a Gradient-Based Approach

For each iteration in the process, an improvement (δJ) to the objective function (J) (such as the lift coefficient) is pursued using the gradient ($\partial J / \partial X^n$) as follows:

$$\delta J = \frac{\partial J}{\partial X^n} \delta X^n \quad (\text{Eqn. 3.1})$$

Where X^n is the set of geometry variables which affect the flow field and the value of J . The changes to the geometric variables must be made according to *direction* of the sensitivity gradient as shown by the green arrows on Figure 3-2. This ensures that the proposed geometry changes are going to be beneficial.

$$\delta X^n = \lambda \frac{\partial J}{\partial X^n} \quad (\text{Eqn. 3.2})$$

Where λ is some scaling factor chosen to produce a sensibly sized geometry change, and this is represented by the *length* of the green arrows on Figure 3-2. This λ can also be positive or negative depending on whether a maxima or minima is desired.

The ‘‘Gradient Descent’’ or ‘‘Steepest Descent’’ method can be applied within this context. With an evaluation of the objective function at the i^{th} iteration (J_i) and its gradient with respect to the geometry variables ($\partial J_i / \partial X_i^n$), the estimate of J_{i+1} can be made for the subsequent iteration (Kreyszig, 2006):

$$J_{i+1} = J_i + \frac{\partial J_i}{\partial X_i^n} \delta X_i^n \quad (\text{Eqn. 3.3})$$

Once the geometry changes (δX_i^n) have been implemented the new aerofoil shape (X_{i+1}^n) is formed, and so a new CFD solution must be computed to evaluate the J_{i+1} . The new sensitivity gradients can also be computed ($\partial J_{i+1} / \partial X_{i+1}^n$) and so the iterative process continues until the difference between J_n and J_{n+1} is suitably small (converged). The question remains regarding how the sensitivity gradients are computed, but this will be addressed shortly.

Gradient based methods are therefore sophisticated because they directly pursue the *closest* extrema (of which there may be several) in the design space. This is shown by the red dots representing each iteration in Figure 3-2. The number of iterations/tests required to reach an optimum compared to gradient-free methods can be greatly reduced. As a drawback, the final solution of the optimisation is thereby heavily subject to the initial design. The global optimum solution will not necessarily be reached which is the main limitation of gradient based methods, thus suggesting their unsuitability for applications with several local optima. They also cannot handle discrete variables (such as number of blades on a turbine). Vavalle & Qin (2007) explore the other shortfalls and drawbacks associated with gradient based methods. Some authors (Peter & Dwight, 2010 and Chehouri *et al.*, 2015) suggest coupling genetic algorithms with gradient based methods to form a hybrid optimisation methodology. Such a method would use global optimisation techniques (such as evolutionary algorithms) to provide a wide search of the solution space, before using gradient based methods to efficiently refine the solution locally towards a true optimum.

Having stated the key merits of gradient based optimisation compared to other methods, it is important to note that computing the sensitivity gradients ($\partial J / \partial X^n$) *can* be costly. For a large number of design variables (X^n where n is large) this *can* mean a proportionately large number of required calculations. The choice of method used to compute the sensitivity gradients is therefore a key factor in determining the overall efficiency of the gradient based optimisation process. Consider a general aerofoil again; if we only wanted to optimise its geometry based on one variable, e.g. thickness, then a sensitivity gradient could be calculated by running a baseline CFD simulation, and then a second simulation with a small change in thickness. The change in performance divided by this geometry change would be the sensitivity gradient for thickness ($\partial J / \partial \text{Thickness}$), and this quantity shows the designer whether adding more thickness has a positive or negative effect on performance. This basic 'finite difference' technique of calculating sensitivity gradients would require an extra flow solution to be calculated for each input

variable under consideration. This would be very costly for aerodynamics problems where there are dozens or hundreds of blade geometry parameters to study (Coppin, 2014). It is possible to consider every single aerofoil geometry node (n) as a design variable (X^n) - but a much more efficient way to compute sensitivities would be required so that the optimisation process does not become prohibitively expensive. An investigation of the various methods for computing the gradients is given in Peter & Dwight (2010), but the Adjoint method is largely regarded as being the most efficient technique. This makes the Adjoint method well suited for aerodynamic design problems as these have many input parameters and few objective functions. In addition, Adjoint methods have been used successfully for aerospace applications which highlights their potential value if applied to VAWTs.

In summary, a gradient based approach using Adjoint sensitivity analysis appears to be a very good choice for VAWT blade optimisation, as the design problem lends itself to the advantages of the Adjoint approach. It is possible that significant rewards could be attained by investigating these methods for VAWTs, as such they have proved highly effective in other fields. If a strategy for applying Adjoint based optimisation to VAWTs is formulated then faster CPU times and greater efficiency of VAWT optimisation could be realised, where far less computational wastage is made.

3.1 Adjoint Sensitivity Analysis

This section outlines the principles of the Adjoint method in relation to aerodynamic optimisation and discusses the literature on the topic. Much of the literature in the realm of Adjoint methods focusses on the mathematical aspects. The present work is concerned with the applications of the Adjoint method in CFD, rather than the complex inner workings of the mathematical formulation. The principles of the Adjoint optimisation framework are however important. As such, this section describes the important concepts more verbally rather than performing a lengthy mathematical recital. There is a wealth of literature which the reader may consult for such mathematical insights, such as Errico (1997), Carpentieri (2009), Coppin (2014), Le Moigne (2003), or the reader may consult the early literature which is described at the start of Section 3.4.

Chapter 3 has mentioned that computing the sensitivity gradients directly can be very computationally expensive for problems with a large number of input variables, even if the number of objective functions (lift or drag, etc.) is low, such as for an aerofoil problem. Adjoint sensitivity analysis provides a distinct advantage to alleviate this issue of computational cost.

The Adjoint method for aerodynamic optimisation is typically presented in the literature in a simplified form, applying principles of sensitivity analysis to an abstract objective function. The objective is expressed in terms of its dependents, a vector of input variables (blade geometry) and a vector of flow field variables. Then through use of Lagrangian multipliers, the equations are manipulated in such a way that the explicit dependence of the objective on flow field perturbations is removed (Dhert *et al.*, 2016). This means that the flow field does not need to be re-computed again and again in order to determine the sensitivity gradients for each of the many input variables. The gradient of the objective function with respect to the flow field state variables is computed just once using the standard CFD solution (Ritlop & Nadarajah, 2009). In other words, instead of making additional sensitivity calculations for every input variable, the Adjoint variables are computed just once (per objective function) regardless of the number of input variables (Le Moigne, 2003). The number of variables needed to sufficiently describe an aerofoil is large; Lee & Liou (2012) state that 10's or 100's of shape variables are required for 2D transonic optimisation problems. With the Adjoint approach a turbine blade can therefore be optimised in high resolution with minimal impact on cost, rather than being constrained to use a small number of shape parameters in order to avoid prohibitive computational times.

Adjoint sensitivity analysis can therefore increase the efficiency of an aerofoil optimisation process by orders of magnitude (Funke *et al.*, 2013). It should be noted that other applications of Adjoint sensitivity analysis exist, but for the present work the combination of Adjoint sensitivity analysis along with gradient based optimisation will be referred to as 'Adjoint based optimisation'. Furthermore, since Adjoint optimised blade shapes can be decoupled from the constraints of parameterisation, non-intuitive, novel, and unexpected solutions are possible. Such possibilities enhance the likelihood of departing from the use of traditional aerospace aerofoils on VAWTs which may well be a paradigm that researchers have inherited which hampers progress in VAWT development (see Section 1.2.1).

The Adjoint method can be applied to the governing equations in one of two ways; the "Discrete" and "Continuous" approaches. Le Moigne (2003) provides a good discussion of the technical/mathematical differences, and the relative merits of the two approaches, explaining that the two often yield similar results. Le Moigne concluded that the choice between the approaches is largely dictated by personal preference. Since both approaches yield similar results the difference lies in the formulation and implementation of the mathematical structure. The Discrete method is favoured by Giles & Pierce (2000) as it is more readily comprehensible, and less complicated to implement for viscous flows. An example of this is that the Discrete approach takes an intuitive set of boundary conditions whereas the Continuous approach contains the objective function explicitly within the boundary conditions, thus making it more challenging to define a physically meaningful and appropriate set of boundary conditions. Further technical insight on the different approaches is given by Nadarajah & Jameson (1999). A continuous solver is largely decoupled from the standard flow solver, and there is more freedom to solve the Adjoint equations using different discretisation/solving practices. This freedom can also hinder the solution accuracy due to variance across the solution methods between flow solution and Adjoint solution. Such problems are manifested in the loss of accuracy in computation of the sensitivity gradients, particularly for complicated flow problems - an approximate gradient is given by the continuous method, but an exact one is given by the discrete method (Nadarajah & Jameson, 1999). ANSYS Fluent uses the Discrete Adjoint method; this approach is less prone to inaccuracies in sensitivity data for complex flow problems.

The Adjoint solution in itself is not the solution to the aerodynamic design problem and these sensitivity gradients are used in an optimisation algorithm to approach the optimum design solution. Sensitivity gradients are computed in the surface normal direction over the blade surface and a gradient value is determined at each discrete point over the blades geometry (Economou *et al.*, 2016). The optimisation

algorithm and mesh morpher decide the degree to which a point on the blade surface is changed according to the sensitivity gradient at that location. The new and improved blade geometry is thus constructed for the subsequent iteration of the optimisation process.

In summary, the research in this thesis implements the Adjoint method for a number of key reasons. Firstly, Adjoints have not been applied to VAWTs as of yet within the published research, and they hold great potential for VAWTs. In addition, widespread CFD codes such as ANSYS Fluent contain an Adjoint solver module, meaning that the methods developed in this thesis are accessible to the general CFD/VAWT community thereby improving the potential adoption and impact of this research. The key obstacle in applying Adjoint methods to VAWTs is their highly unsteady flows. This suggests the use of transient Adjoint methods however these are extremely complex and time consuming although some success has been shown in turbomachinery applications (Li et al., 2011), (Walther & Nadarajah, 2015), (Luo et al., 2011). The present work develops an “engineering approach” which carefully applies a steady-state Adjoint solver to the transient problem of VAWT aerodynamics. This can be done despite the unsteady nature of the VAWT flow field.

3.2 Choice of Objective Function

In an optimisation problem, the quantity which should be maximised/minimised is referred to as the ‘objective’ or ‘objective function’. In some literature this can also be called the ‘observable’. In a typical aerodynamics context, there are usually a small number of objectives that the designer is interested in, such as lift or drag of an aerofoil. An Adjoint solution can be computed for such quantities, yielding the associated sensitivity data for use in the gradient based optimisation. With regard to VAWT research one can generally assume that the quantity of interest is the average power coefficient at a given TSR. This however is not an instantaneous quantity meaning it cannot be characterised by a single instantaneous flow field. With a steady Adjoint solver average C_P cannot be chosen as the objective. Therefore, in this work the blade tangential force coefficient or moment coefficient is chosen as the objective.

3.3 Multi-point/multi-objective

A simple case for an optimisation problem may be an aerofoil operating at a constant velocity and constant angle of attack, such as a cruising aircraft. This is a single flight condition, and a single ‘point’ in the design space. There would also only be a single objective for this case, such as the lift/drag ratio. If this were a real aircraft, the aerofoil would also need to exhibit satisfactory performance during take-off, landing

and loitering/manoeuvring conditions. Each of these conditions constitutes a different design 'point' due to the different flow speeds (and Reynolds number). Each of these flight conditions may also require a different objective, such as maximising lift rather than the lift/drag ratio. A VAWT arguably has a more complex design space than this aircraft because there are not merely a handful of flight conditions, but a continuous variation in Reynolds numbers and AoA as the TSR varies all the time. Furthermore, even at a constant TSR, each blade experiences an enormously varying flow as it completes a single rotation. Applying Adjoint methods to VAWTs clearly constitutes a great challenge in terms of capturing the full operating envelope. The present work is concerned with optimisation at a single constant TSR, although future developments of these methods will need to address multi-point optimisation.

3.4 Adjoint Literature

Adjoint optimisation has relatively little history in the application to wind turbines. However, aircraft and turbomachinery (gas turbines) have enjoyed considerable attention in this area. The aforementioned powers of the Adjoint method have been long known in concept but have only been applied in the field of computational fluids engineering across the last few decades. The Adjoint method was applied to aerofoil optimisation and popularised by Anthony Jameson in 1988 which was presented as "Aerodynamic Design via control Theory" (Jameson, 1988), having taken inspiration from Lighthill's (1945) work on aerofoil optimisation via conformal mapping. In the decades that followed, research (in large part by Jameson and his disciples) focussed on high speed flows for aerospace applications. The complexity of the methods and their applications gradually evolved from potential flow in 2D steady flows, to the optimisation of full aircraft configurations in 3D.

Li *et al.* (2011) presented 2D viscous optimisation of a turbine cascade using the Adjoint method and CFD. For this thermodynamic design problem the objective function was set to minimise the rate of entropy increase. Walther & Nadarajah (2015) built upon a 3D, single design point Adjoint optimisation method to construct a multi-point optimisation framework. The application was for a single stage axial compressor (of a gas turbine engine) and considered the performance at off-design conditions as well as peak efficiency conditions. Luo *et al.* (2011) considered the choice of objective function and conducted an optimisation to minimise the rate of entropy increase, as well as an optimisation for a desired outlet flow angle from the turbine stage.

Dhert *et al.* (2016) constructed an Adjoint optimisation framework with 3D CFD using the NREL HAWT to demonstrate the effectiveness of the optimisation method.

Dhert *et al.* claimed this to be the first CFD based wind turbine optimisation that provisioned for a suitably large set of blade shape variables (of which there were 240). This was claimed to be due to the large computation cost of such work with other methods, restricting the number of shape variables. Optimisation was carried out for a single wind speed and also for a multi-point/ multi-wind speed case. The optimisation process took place on a 256 core system and completed in around 24 hours. Significant performance improvements of around 20% (torque coefficient) were realised which provides optimism for the potential of Adjoint methods when applied to VAWTs. Lei & He (2016) presented Adjoint optimisation for a 2D aerofoil in unsteady, low Reynolds number flow with application to micro aerial vehicles. Unsteady laminar Navier-Stokes equations were used (via a flow solver developed in-house) to account for the prominent aerodynamic phenomena of viscous unsteady flow fields, and the time averaged lift/drag ratio was set for the objective function. Lee & Liou (2012) discussed the improvement of thrust for a flapping airfoil in laminar flow, via an unsteady discrete Adjoint optimisation approach. Such methods will incur a larger computational cost as well as complexity in implementation compared to steady Adjoint formulations. Lee & Liou highlighted the importance of the details in the development and shedding of vortices over the wing and discussed how varying the modes of shedding during flight (by modifying the flapping motion) can achieve flow patterns for more optimal aerodynamics. Their study included both shape optimisation and wing trajectory optimisation, but the Hick-Henne function they used defined the blade geometry in 20 shape parameters which appears to be somewhat coarse. Lee & Liou used time averaged objective functions and stressed the importance of resolving the transient flow features with a suitable time scale in order to reflect the correct value of the averaged objective function.

The literature leaves a discernible gap for application of Adjoint methods to VAWTs. Furthermore, the literature of Adjoint methods tends to be more on the theoretical side than the applied side. This project therefore aims to pursue the practical application of Adjoint methods to VAWTs, using a commercial CFD code so that these methodologies can be readily applied in the VAWT design community. Details of the project aims & objectives are discussed further in Section 4.1.

3.5 Adjoint Based Optimisation Frameworks in CFD

An Adjoint optimisation framework is a gradient based optimisation algorithm that uses Adjoint sensitivity analysis to provide the sensitivity gradients. These frameworks are composed of several very complicated elements aside from the Adjoint solver itself (Carpentieri, 2009). The following components are required:

- Standard CFD flow solver (a flow field solution is needed in order to solve the Adjoint equations)
- Shape parameterisation (defines the method by which the input variables/ geometry are expressed mathematically)
- Adjoint solver (determines the shape sensitivity gradients from the flow field)
- Optimisation algorithm (uses the shape sensitivity gradients to determine the proposed new and improved geometry for the next iteration)
- Geometry/mesh morpher (to implement the required geometry changes into the CFD model)

It is not until recently that stable Adjoint solvers have become available in readily accessible CFD codes (Peter & Dwight, 2010). Across the literature, researchers have constructed their own Adjoint Frameworks, but significant programming efforts are required in order to wrap the necessary components together into a functional framework. In this way, researchers can construct frameworks with desired features (such as a transient Adjoint solver or the use of a specific turbulence model). Such ventures require extensive knowledge of Adjoint mathematics, CFD, and computer science/programming, and these projects absorb a significant portion of time assembling the frameworks. There remains a time lag between capabilities of the latest Adjoint solvers, and their implementation within more widely available/commercial CFD software. As such, the state of the art in Adjoint solvers (such as transient Adjoint solvers) are not yet contained within ANSYS Fluent, the most widely used CFD package. ANSYS Fluent is however an attractive option for the aerodynamic designer seeking to make use of and explore the power of Adjoint methods. Fluent's main selling point in this regard is that it contains all the necessary components of an Adjoint optimisation framework configured together and ready to use.

The present project aims to maintain focus on the exploring the *application* of Adjoint methods to VAWTs, so ANSYS Fluent with its existing Adjoint framework is chosen. Because Fluent is so widely used and accessible internationally, this also means that the methods presented here could be readily deployed by the wider VAWT engineering community yet also applied within other Adjoint frameworks.

In recent versions of ANSYS Fluent the Adjoint module has been configured for compatibility with rotating reference frames. Rotating reference frame approaches are typically used for some rotating machinery where the oncoming wind is along the axis of the turbines rotation. In such unsteady problems such as HAWT and gas turbine design, the rotational symmetry of the flow means that these can be approximated to a steady flow regime within the rotating reference frame in CFD (Dhert *et al.*, 2016). VAWTs however cannot make use of this simplification due to their intrinsic unsteadiness and lack of rotational symmetry (see Section 2.2) (Balduzzi *et al.*, 2015). When aiming to perform aerodynamic optimisation it can be relatively straightforward to apply Fluent's Adjoint module for simple problems such as a 2D aerofoil. It becomes complicated to investigate the use of Fluent's Adjoint module for VAWTs, where these steady flow approximations cannot be used. The unsteadiness of VAWT flow fields therefore poses a significant challenge to the prospect of applying steady Adjoint solvers in a useful manner.

This challenge forms the basis of the present research project (see research question in Chapter 4), where an engineering approach must be constructed to make use of Fluent's Adjoint solver for the transient problem of VAWTs.

4 RESEARCH AIM

Having described VAWT aerodynamics and the literature pertaining to aerodynamic optimisation, this chapter will concisely state the objectives of the research project and the novelty of the work. The research question for this project is as follows:

“Can VAWT blade shape design be benefited from Adjoint based optimisation and how could such methods be formulated using a widely available CFD code?”

Addressing this question saw major challenges due to the lack of literature on the topic, and so a novel semi-transient Adjoint based optimisation methodology was created within the commercial CFD code ANSYS Fluent (18.2). Variations of the method were developed and tested, and these approaches differ between the types of models used to predict the VAWT blade aerodynamics, and also the amount of Adjoint sensitivity data considered in the optimisation. Progressing the research project in this step-by-step fashion meant that early feasibility of the approach could be demonstrated on simpler models, before gradually adding refinements and complexity to the method. The main chapters present this evolution of the method as follows:

- “Single-Blade” optimisation – see Chapter 5. A single oscillating blade model is used to approximate the VAWT blade flow field. The optimisation process is applied to this Single-Blade model, but the resulting blade geometry is tested in a VAWT model to measure the performance. Adjoint sensitivity data is considered from a single instance in time during the oscillation cycle.
- “In-situ VAWT” (ISV) optimisation – see Chapter 6. A standard VAWT model is used to simulate the flow field, and the optimisation is applied to this model. The performance is also measured from the same model; no intermediate models are used. Adjoint sensitivity data is considered a single instance in time during the VAWT revolution.
- “In-situ VAWT” (ISV) optimisation with multiple Adjoint snapshots – see Chapter 7. In a more refined approach, Adjoint sensitivity data is considered from multiple instances across the VAWT revolution in a combined fashion. This aims to produce aerofoils that performs well over the majority of the VAWTs revolution.

In this work these methods were developed within ANSYS Fluent, however they can be reproduced and developed in other CFD codes which have an Adjoint solver. The work is confined to straight bladed Darrieus turbines analysed in 2D CFD simulations, with constant TSR, and wind speed.

4.1 Research Novelty

The novelty of this project is composed of several main elements:

- The advantages of Adjoint based optimisation are utilised by applying Adjoint methods to VAWTs. There is an absence of literature on this topic, despite research found for the application of Adjoint methods to other technologies. This thesis contributes novel Adjoint based VAWT optimisation methods in CFD, which can improve the average power coefficient, the instantaneous power coefficient. This sets the foundation for a new and exciting avenue of VAWT research.
- Optimisation algorithms were developed from scratch and coded into executable scripts for use with commercial CFD software. These methods utilise the ANSYS Fluent's Adjoint module in a way that is outside of its intended purpose, demonstrating the value of the engineering approach of 'Semi-transient' Adjoint based optimisation.
- Novel VAWT blade geometries are discovered, which due to the Adjoint method are produced in a highly efficient manner as well as being free from the limitations of conventional aerofoil parameterisation. These resulting geometries are explored with an in-depth aerodynamic discussion, thus offering new and important insight into the optimum blade geometry. Furthermore, the optimum blade shape varies continually throughout the VAWT revolution (according to the changes in AoA and relative velocity) and this is also explored.
- The approach offers a novel solution to improving the fatigue life of VAWTs by drastically reducing the fluctuating load magnitudes while preserving the average power coefficient (without the need for a control system).
- This thesis also contributes practical "how-to" guidance on use of the ANSYS Fluent Adjoint solver for a real problem (see Sections 5.3.1 and 6.3.1). This contributes to filling a large gap in the publicly available guidance material. Such helpful practical content is lacking within the current literature and will aid the implementation of Adjoint optimisation methods (not just for VAWTs) by competent ANSYS Fluent users who are not necessarily Adjoint experts.

5 SINGLE-BLADE OPTIMISATION FOR VAWTS

The methodology of this research chapter has been described and published in Day *et al.* (2020).

Section 2.2.3 has described how a single flapping wing with an oscillating AoA can provide an approximation to the flow field of a VAWT blade. This similarity is utilised in the present chapter which adopts a single oscillating blade model to represent the VAWT blade, and the optimisation is applied to this model. This Single-Blade optimisation method therefore benefits greatly from model simplicity, instead of applying the optimisation process directly to a VAWT model (see Chapter 6). The reasons that the Single-Blade approach was developed first are as follows:

- Single-Blade optimisation was used at the beginning of the project to test the feasibility of the Adjoint method for the simpler aerodynamics problem of a fixed aerofoil. This was extended to a harmonically oscillating blade, and subsequently refined further to approximate a VAWT blade flow by creating User Defined Functions (UDFs) to model the VAWT AoA profile and relative velocity profile.
- Single-Blade CFD models can be significantly smaller and require far fewer periods of oscillation to reach a stable CFD solution. In an iterative optimisation process this can make a Single-Blade model far less computationally expensive.
- Single-Blade optimisation uses a less complex CFD model, such that the optimisation scripts required are simpler, faster and easier to validate.
- The Adjoint solutions are more stable and exhibit greater convergence with the Single-Blade model compared to the VAWT model.
- Building up from simple cases in this fashion was a pragmatic approach to develop the initial Adjoint based optimisation methodology, without the additional challenges of VAWT modelling.

The main drawbacks of the Single-Blade method are discussed in Section 8.1.

This chapter is divided up into the following sections:

- Section 5.1 – Single-Blade Approximation Model; describes and justifies the simplification of the VAWT flow field to that of a single oscillating blade, and then validates the CFD model.
- Section 5.2 – Semi-Transient Adjoint Optimisation Philosophy; describes the principle of using Adjoint optimisation data from discrete instances in time for

the transient problem of VAWT aerodynamics. This includes introduction of the concept of Adjoint ‘snapshots’ which is used throughout this thesis.

- Section 5.3 – Optimisation Algorithm; describes the architecture of the optimisation process algorithm of the Single-Blade method. This includes setup details of the Adjoint solver, optimiser aggression and mesh morphing. Practical guidance/ best practice for using the ANSYS Fluent Adjoint module in this context is also provided.
- Section 5.4 – Single-Snapshot Investigation; describes the application of the optimisation process to a VAWT blade, where 1 Adjoint snapshot is used per turbine cycle. Results are presented for a range of cases tested where the snapshot position has been varied. Note that the Multiple-Snapshot optimisation investigation is presented separately in Chapter 7.

5.1 Single Blade Approximation Model

A single, isolated aerofoil with an oscillating pitch (AoA) can be used to approximate the aerodynamic flow field of a VAWT, although vortex/wake interactions and some plunging motion components are neglected (Wang *et al.*, 2010). This approximation can be of sufficient accuracy to enable the use of a Single-Blade model instead of a full VAWT model during the optimisation process. Choosing to make this simplification drastically reduces the model size and therefore saves CPU (Central Processing Unit) time. At the end of the Single-Blade optimisation the resulting geometry can be tested on a standard VAWT model to accurately evaluate the blades’ performance.

This link between a Single-Blade model and the VAWT it represents is illustrated in Figure 5-1, where V_{rel} is the relative velocity and α is the AoA at an arbitrary Location A.




Figure 5-1 – AoA Similarity (link) between the Single-Blade model and the VAWT

The theoretical V_{rel} and AoA profiles of a VAWT blade can be applied to the Single Blade model using User Defined Functions (UDFs) in ANSYS Fluent. The V_{rel} UDF is applied to the inlet boundary where the inlet velocity varies according to Equation 1.3; further details of this UDF are given in Appendix C.

The oscillating AoA UDF is applied to a local rotating zone in the mesh which contains the blade (see Figure 5-2); the rotational frequency of the VAWT, ω , is used as the pitching frequency of the oscillating aerofoil, ω_p . The simplest formulation of this AoA UDF would be to implement Equation 1.4 exactly. However, Equation 1.4 neglects important real-turbine effects that make a significant difference to the AoA profile (Rosado-Hau, 2021). The interference/blockage factor of the VAWT blades and the extraction of energy across the rotor cause a reduction in velocity and AoA for VAWT blades compared to the geometrical AoA in Equation 1.4. The difference is most pronounced for the AoA of the blade in the downwind part of the turbine (Rosado-Hau, 2021). To improve the accuracy of the Single-Blade model flow field, a knockdown factor can be applied to the AoA during the oscillation. Gosselin et al. (2013) applied a constant AoA knockdown factor of 0.5 to the downwind part of a Single-Blade model. This approximately accounts for the transverse velocity components arising from energy extraction in the upwind and associated slow-down of the flow across the rotor (Gosselin et al., 2013). In the present work, this approach of a constant downwind knockdown factor was adopted for the Single-Blade model during the investigation presented in Section 5.4.

To implement the UDF for varying the AoA, the Single-Blade CFD model uses multiple mesh zones along with the sliding mesh technique. This technique consists of a circular non-conformal interface between the exterior mesh and the rotating subdomain, where mass/momentum exchange takes place. This practice is common within the literature (Hand et al., 2017). Figure 5-2 shows the domain and mesh zones of the model.

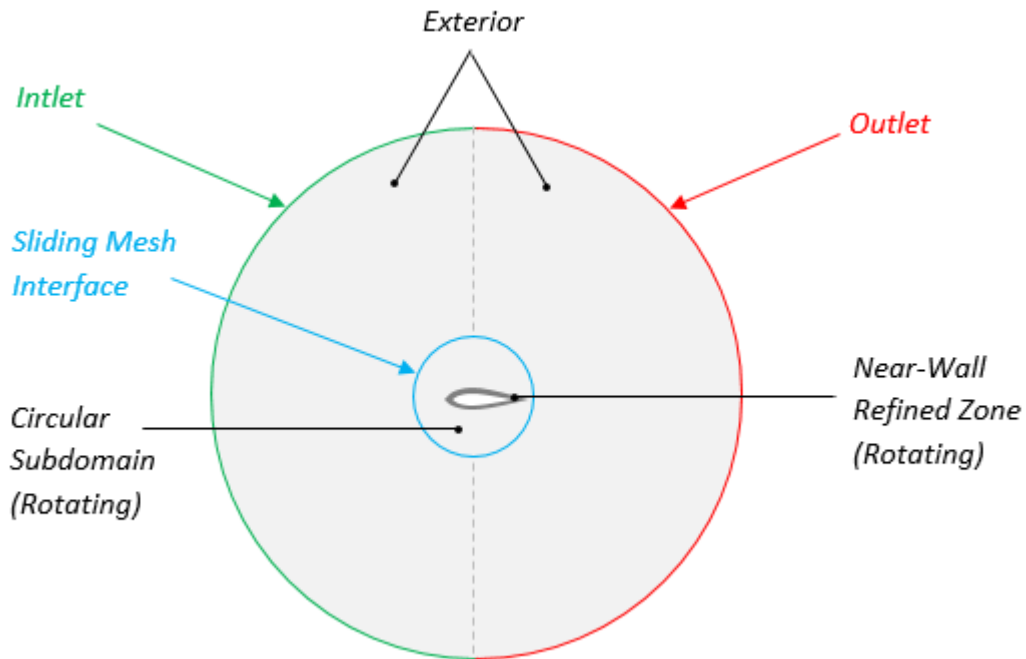


Figure 5-2 – Domain Type Used for the Single-Blade Model

Note that one cannot simply assign the VAWT wind speed as the inlet velocity of the Single-Blade model since this does not account for the blade rotational velocity component. For cases where the TSR is very high however, a constant inlet velocity equal to the tangential velocity ($\omega \cdot R$) could be used instead of the V_{rel} UDF, saving CPU time. The justification for this was given earlier, see Figure 1-6.

5.1.1 Introducing the Reference Turbine

A high TSR (4.5), 2 bladed VAWT from Rezaeiha et al. (Vol 107, 2017) was selected as the reference turbine used to develop this optimisation method with. A high TSR was chosen so that AoA variation and unsteadiness of the flow would be minimised, thereby avoiding additional flow complexities during early development of the methodology. The VAWT blades also have a fixing angle of zero and are fixed (not a variable-pitch turbine). Using this reference VAWT the agreement between a Single-Blade model and the corresponding VAWT can be seen in Figure 5-3. It should be noted that the VAWT data shown in Figure 5-3 is extracted from a VAWT model which is described in Section 6.1.




Figure 5-3 - Comparison of a Single-Blade Model Data and corresponding VAWT Model Data

The reasons for agreement/disagreement between the two models in Figure 5-3 have been discussed in Section 5.1.

This VAWT will be used as the example case to illustrate the optimisation method developed here. This VAWT is as described in Rezaeiha et al. (Vol 107, 2017) with the turbine details presented in Table 5-1. It is selected due to its high TSR and the availability of data which can be used for CFD model validation. It should be noted that the CFD work of Rezaeiha et al. (Vol 107, 2017) is based on an experimental study conducted by Tescione et al. (2014).

Table 5-1 - Details of the VAWT from Rezaeiha *et al.* (Vol 107, 2017)

Turbine Blade Profile	NACA0018
Blade Fixing Angle	0°
Number of Blades	2
Blade Chord Length	0.06m
Blade Length	1m
Turbine Diameter	1m
Rotational Velocity, ω	83.8 rad/s
Free-stream Wind Speed, U_∞	9.3 m/s
Tip Speed Ratio, λ	4.5

5.1.2 Single-Blade CFD Model Setup

Section 5.1 has described the modelling philosophy of the Single-Blade model. This section describes specific details of the CFD setup.

Before detailed discussion is given, Table 5-2 will summarise the key features of the CFD setup adopted for the Single-Blade model.

Table 5-2 - Summary of CFD Setup for Single-Blade Models

Setup Factor	Choice Adopted
2D/3D	2D
Turbulence Model	SST k- ω
Maximum mesh y^+	1.5
Number of cells around aerofoil wall	310 (see Section 5.1.3)
Cell count	54,000 (see Section 5.1.3)
Domain	(See Section 5.1.3)
Time step size	800 steps/rev (see Section 5.1.3)
Target CFL	200 (see Section 5.1.3)
Number of cycles for periodic convergence	3 (see Section 5.1.3)
Inlet boundary condition	Velocity specified by Equation 1.3 (UDF)
Outlet boundary condition	Pressure outlet
Solver type	Pressure based
Pressure-velocity coupling scheme	Coupled
Spatial discretisation scheme	Second order upwind
Temporal discretisation scheme	Bounded second order implicit
Number of iterations per time step	30
Convergence criteria	1e-5 (for all residuals)

The SST k- ω turbulence model is used for all the CFD simulations; this choice was made with consideration of the literature given in Section 2.4 as well as the following points. Numerous authors have conducted investigations of turbulence model suitability for VAWTS using CFD and experimental data. Many authors claim that the basic 2 equation models (standard k- ϵ and standard k- ω) are incapable of accurately

predicting VAWT flows, while the SST k- ω variant is favoured over other models of similar complexity/cost and can adequately reproduce the VAWT flow fields (Wang et al., 2010), (Balduzzi et al., 2015), (Hand et al., 2017). Rezaeiha et al. (2019) conducted a comprehensive study of the S-A, RNG k- ϵ , realizable k- ϵ , k- ω SST, k- ω SST with intermittency, k-k1- ω , and Transitional SST. A range of flow conditions were studied with the conclusion that the SST models can provide reasonable predictions of VAWT flows including dynamic stall.

Other models exist that could improve the solution accuracy for transitional flows, such as the k- ω SST with intermittency, Transition SST and LES. Despite this, the SST k- ω model is adopted for the present work due in part to accuracy/cost considerations. The SST k- ω model is the most accurate type supported by ANSYS Fluent's Adjoint module while also being deemed suitable for VAWTs across the literature. Consideration has been given to the k- ω SST with intermittency model; it is more much sensitive to variations in the mesh, time step and other modelling choices (El Sakka, 2020). In addition, for the high TSR case adopted here it is judged that the SST k- ω turbulence model provides adequate prediction of the unsteady flow since the dynamic stall effects are minimal. Considering the many recommendations in the literature, the k- ω SST (not intermittency) is therefore considered to most robust and appropriate choice for the purposes of this thesis. The methodology developed here can be implemented in other CFD codes if desired, such that other specific turbulence models could be used if necessary.

In the CFD simulations a value of 5% is used for the turbulence intensity, in accordance with the research paper of the VAWT (Rezaeiha et al., Vol 107, 2017). The turbulence length scale, in lieu of specified values, is set as the turbine diameter (Rezaeiha et al. Vol 156, 2018). The CFD simulations implement the Coupled Numerical Scheme to improve the solution speed. The pressure based solver is used with the second-order upwind scheme for spatial discretisation. The bounded second-order implicit scheme is adopted for the transient formulation due to its robustness (Hand et al., 2017). A limit of 30 iterations per time step is used and a minimum convergence criteria of $1e-5$ is set for all residuals. These settings are common across the VAWT CFD literature in order to achieve sufficient solution convergence (Li et al., 2018, Wang et al., 2010, Rezaeiha et al., Vol 107, 2017, and Guo et al., 2019).

The meshing strategy for the Single-Blade model is shown in Figure 5-4 and Figure 5-5, but in Figure 5-4 a highly coarsened version of the mesh is shown for brevity.

The red line in Figure 5-4 shows the (non-conformal) interface between the circular subdomain (sliding mesh zone) and the far-field zone, where mass/momentum

exchange takes place. The element spacing on each side of the interface should match closely in order to “*minimise interpolation error between the two zones*” (Hand *et al.*, 2017). This circular subdomain also allows local mesh refinement and a structured mesh with quadrilateral cells is used throughout (Almohammadi *et al.*, 2012). A further subdomain (“Near-Wall Refined Zone”) around the aerofoil surface is used to ensure extra fine mesh quality near the aerofoil wall, which is shown in Figure 5-5; the red line marks the conformal interface to the rest of the rotating region. A fine mesh in this region has the purpose of providing suitable resolution of the boundary layer complexities/velocity gradients near the blade wall. A y^+ of approximately 1 can be maintained which alleviates the need for a wall function (Wang *et al.*, 2010). The y^+ indicates the relative first cell size; it is a non-dimensional quantity linked to the time step and flow velocity. Correct specification of y^+ ensures that the first cell is located in the viscous sub-layer (Almohammadi *et al.*, 2012). Achieving the appropriate y^+ is therefore critical to the accuracy of the simulation due to the very high velocity gradients adjacent to the wall and the mesh y^+ is (along with the Courant number) the most important quantity linked to the solution quality (Balduzzi *et al.*, 2015). The y^+ value should be sufficiently small for the problem at hand.




Figure 5-4 – Example Mesh of the Rotating Circular Subdomain (Coarsened Mesh Shown)




Figure 5-5 – Example Mesh of Near-Wall Refined Zone (a) Leading Edge; (b) Trailing Edge

The mesh y^+ is defined as follows;

$$\text{Mesh } y^+ \quad y^+ = \frac{\rho U_\tau y}{\mu} \quad (\text{Eqn. 5.1}).$$

where μ is viscosity of the fluid, ρ is the density of the fluid, y is the height of the first cell adjacent to the wall and;

$$\text{Shear velocity;} \quad U_\tau = \sqrt{\frac{\tau_w}{\rho}} \quad (\text{Eqn. 5.2}).$$

τ_w is the wall shear stress which can be estimated, or alternatively the CFD code can compute this to provide accurate y^+ values for all cells at the wall.

In the literature survey of Balduzzi *et al.* (2015) it was found that mostly, a y^+ value below 5 is selected for VAWTs. Wang *et al.* (2010) stated that the first element size should be such that $y^+ \sim 1.0$. Rosado Hau (2021) maintains a maximum y^+ of 1.8. El Sakka (2020) calculates the y^+ considering the position of maximum relative velocity, which occurs at 0deg where the oncoming wind velocity and blade tangential velocity ($\omega.R$) are in maximum combination. In the present work a maximum y^+ of 1.5 is reached, at the 0 degree position, and this is in line with the literature recommendations.

The time step size needs to be considered alongside mesh size rather than independently from it because the two are strongly related (Wang *et al.* 2012). This concept is represented by the Courant number which indicates how many cells are traversed by the fluid over a time step and this is also referred to as 'CFL' (Courant-Friedrichs-Lewy) within the literature.

$$\text{Courant Number (or CFL)} = \frac{\text{Characteristic Flow Velocity} \cdot \Delta t}{\text{Typical Cell Size}} \quad (\text{Eqn. 5.3}).$$

The literature generally encourages the use of a small Courant number. Hand *et al.* (2017) recommends between 5 and 10, while Trivellato & Castelli (2014) states that less than 0.5 could even be employed.

As can be seen from Equation 5.3 the CFL will varies for each cell in the CFD model because the cell size and flow velocity are different for all cells. If a high CFL is exhibited, then numerical divergence is more likely since the fluid is effectively skipping cells during each time step. When using the Coupled solver this problem is mitigated by ANSYS Fluent making automatic corrections to time step size if divergence is detected. The user can also specify a target CFL limit, which ANSYS Fluent will satisfy by making automatic reductions to time step size. For smaller time steps this produces a smaller CFL, and the user specified CFL limit becomes irrelevant, as the specified time step can be adopted.

Possible savings to the CPU cost should be considered for simulations with heavy CPU requirements, especially in the context of optimisation where many repetitions of CFD runs are required. Therefore, using a CFL as low as 1 or 5 etc. as stated in the previous literature may not be appropriate. A low Courant number or low time-step size becomes computationally demanding and an appropriate accuracy/CPU cost balance should be pursued, so long as solution convergence is reached.

5.1.3 Single-Blade CFD Model Validation

Once a Single-Blade model has been produced with consideration to the above, the necessary validation exercises must be conducted:

- Periodic solution convergence study
- Grid independence study
- Time step independence study
- Courant number study
- Domain size study

Such validation exercises are routine within the literature, so the details are provided in Appendix A. The outcomes of this work demonstrate suitable validation for the Single-Blade CFD model when the setup is as follows:

- 3 cycles were sufficient for a periodically converged solution
- Domain size in terms of chord length, c ; Exterior diameter = $40c$, rotating zone diameter = $6c$, boundary zone thickness = $0.05c$
- Y^+ of around 1 (justification in literature review)
- Target CFL number of 200 in the Coupled solver
- Time step corresponding to 800 steps/revolution
- 310 cells around the aerofoil wall
- Total cell count of approx. 54,000
- Quad elements used throughout all subdomains (justification in the literature)

Following these validation studies, the flow field of the Single-Blade model must be checked against the VAWT flow field. Such a comparison has been shown in Figure 5-3; in this case data from a validated VAWT model (see Section 6.1) is converted into C_T vs AoA data allowing direct comparison with Single-Blade model data. Single-Blade model data can also be converted to virtual C_M vs virtual azimuthal angle (using Equation 1.5) allowing the direct comparison with VAWT data as shown in Figure 5-6.

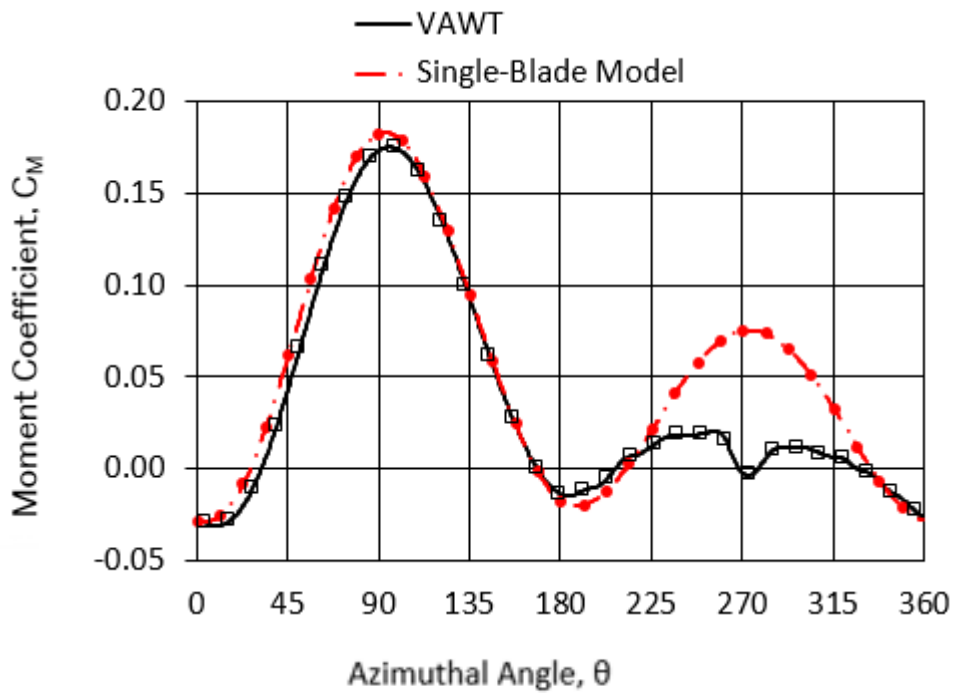


Figure 5-6 – Comparison of Single-Blade Model Data and VAWT Model Data (C_M)

Figure 5-6 demonstrates that the Single-Blade model provides an acceptable level of accuracy in predicting the flow field of the VAWT blade. The upwind region is replicated well but the downwind region (specifically between 225° and 325°) is inaccurate. This is because the Single-Blade model does not account for blockage/interference effects, curvature effects and shaft/blade wake interactions in the downwind region – further discussion has been given in Section 5.1. However, the Single-Blade model is deemed a suitable approximation for use in the optimisation process since the agreement is close for the majority of the cycle. However, owing to the discrepancies mentioned, care must be taken if considering optimisation data from the downwind region (this is further discussed in Section 5.2.2).

5.1.4 Validating the VAWT CFD Model

The VAWT model used to validate the Single-Blade model, also requires validating. These details are provided in Section 6.1.

5.2 Semi-Transient Adjoint Optimisation Philosophy

Section 5.1 has shown that a Single-Blade approximation model can be used as a basis for optimising VAWT blades. This section describes the principle of the optimisation methodology developed in the present work. The ANSYS Fluent Adjoint solver is used which is stated as being configured for steady-state flows. ANSYS also states however that *this is not a hard restriction*, and steady Adjoint solutions can be of value and of use for engineering approaches to unsteady problems (Eggenspieler, 2012). The methodology developed here is such an engineering approach for the transient problem of VAWT aerodynamics. Even though VAWT aerodynamics are highly unsteady, a steady Adjoint solver can be applied to instances in time during transient flow field, these are called Adjoint ‘snapshots’ (see Section 5.2.2). This optimisation process can therefore be considered a ‘semi-transient’ approach since the flow field solution arises from a fully transient CFD simulation, but the Adjoint solutions consider data from a series of individual instances in time.

5.2.1 Objective Function

Section 3.2 has described the role of the objective function in Adjoint based optimisation. For the oscillating Single-Blade model, the objective is chosen as the tangential coefficient, C_T (see Figure 1-4). The reason for this choice is that this quantity translates to the moment of force M_F , and hence the moment coefficient (C_M) and power coefficient (C_P) of a VAWT (see Eqns 1.5 & 1.6).

In this way, the sensitivity gradients (defined in Section 3.2) resulting from the Adjoint solver calculation describe the increase in C_T that would arise from a given change to the blade geometry (on a node-by-node basis). A blade geometry once optimised for C_T in the Single-Blade model, will therefore produce a higher torque and thereby power output when placed on a VAWT and undergoing a similar flow field.

The C_T can be programmed as the objective in the Adjoint solver via its components in the X (C_D) and Y (C_L) directions as per Equation 1.1. The objective cannot be specified as a continuous function (i.e., of time) within ANSYS fluent so the correct AoA needs to be specified via a script (or manually) whenever an Adjoint solution needs to be computed for the objective C_T ; this will be discussed more in Section 5.3.3.

5.2.2 Snapshots

In this work the term ‘snapshot’ refers to a single Adjoint solution taken at any instance in solution-time while the Single-Blade model is oscillating. More specifically, it regards the sensitivity data which arises from that Adjoint solution (see Section 3.1).

As stated in Section 2.2, the VAWT blade flow field varies greatly over the course of one revolution. In turn, the Adjoint sensitivity data will also vary depending on the flow field (AoA, relative velocity, etc.) at the time of the snapshot. The choice of snapshot position is therefore a crucial factor in influencing the resulting blade geometry and performance once the optimisation process is completed.

To illustrate this, Figure 5-7 shows a C_T vs AoA performance curve for the baseline aerofoil, with the position of an arbitrary snapshot marked in terms of AoA. In this case the snapshot position is within the upwind part of the cycle, where the AoA is approaching its maximum negative value (see Section 1.2 for sign convention). The arrows indicate the direction of the blades oscillating motion.




Figure 5-7 - Single Blade Performance Curve with the Snapshot Location Marked

Section 5.4 investigates the impact of snapshot location on the outcomes of the Single-Blade optimisation process. The equivalent investigation using direct VAWT model optimisation (rather than a Single-Blade model) is presented in Section 6.4. Finally, a multi-snapshot study is presented in Chapter 7. These investigations are deferred until later in the thesis so that details of the methodology can first be discussed.

5.3 Optimisation Algorithm

The previous sections (such as Section 3.5) have explained the building blocks of the semi-transient Adjoint based optimisation process for VAWTs. This section illustrates the Single-Blade optimisation process architecture in detail and describes the algorithm code. Also presented are details of the Adjoint module setup and key notes regarding implementation of the process philosophy into the CFD software.

This process can be applied to other CFD codes with an Adjoint solver. The level of detail provided here aligns with the ANSYS Fluent (18.2) Adjoint module; more steps may be involved for implementation in other CFD codes.

Figure 5-8 shows a generalised Adjoint optimisation process for a generic aerofoil.

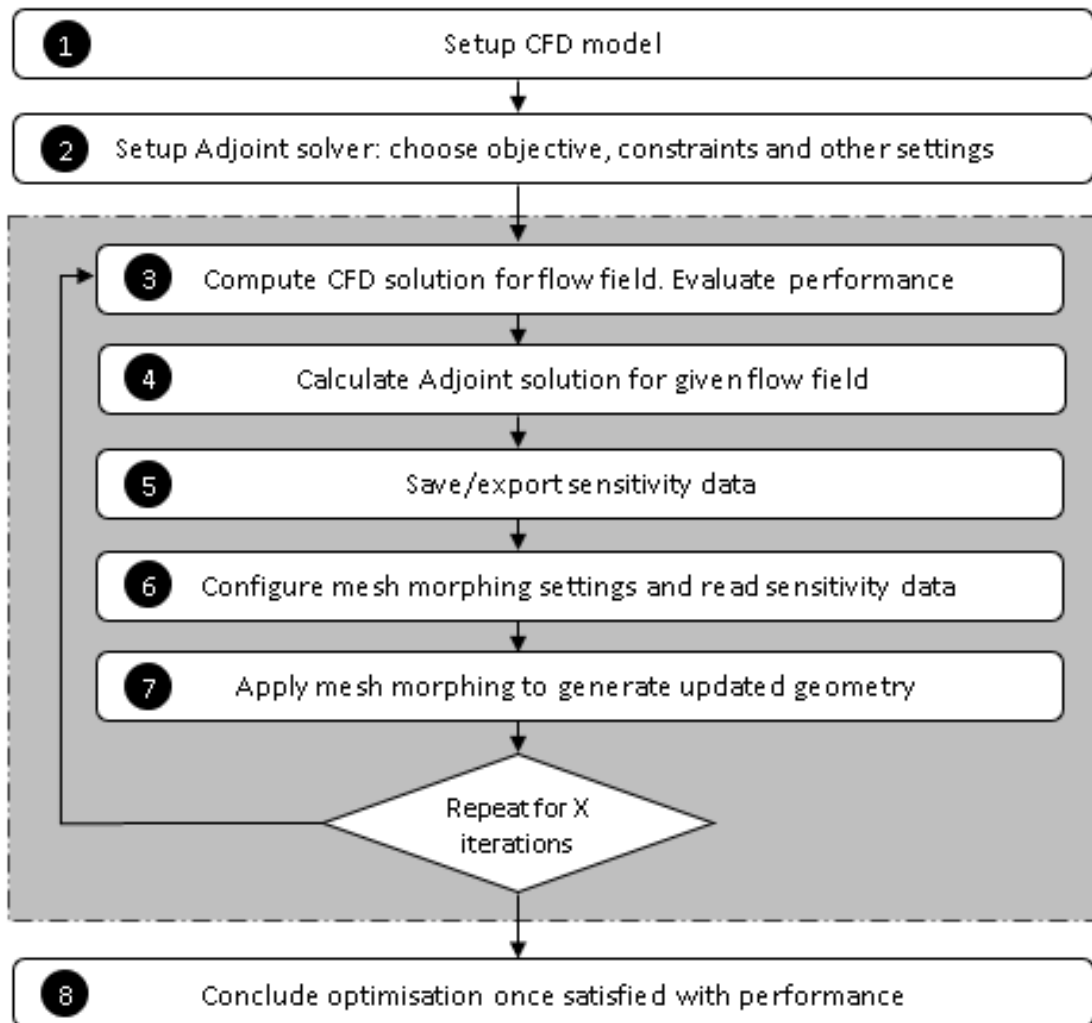


Figure 5-8 - General Overview of the Adjoint Optimisation Procedure

The Adjoint solution requires a flow field solution to have been produced from the CFD (step 3). The Adjoint solver computes the sensitivity gradients at each node on the blade surface (step 4) based on that flow field. As described in Section 3.1, the sensitivity data indicates how the performance is affected by a geometrical change at each surface point of the aerofoil. An example set of sensitivity vectors for an aerofoil can be seen in Figure 5-9. These vectors suggest which direction to make geometry changes in order to improve the objective function (performance).




Figure 5-9 - Example Vectors of Shape Sensitivity

The mesh morpher computes an appropriate deformation at each node according to the size and direction of the sensitivity vectors (step 7). The amount of deformation arising from the mesh morphing process also depends on the constraints, and 'aggression' settings (see Section 5.3.2).

Figure 5-10 shows the concept of the Adjoint based semi-transient optimisation process for VAWTs using a Single-Blade approximation model.

The aim is to produce a VAWT with increased average power coefficient, C_P compared to the original VAWT. Graphs of C_P as a function of azimuthal angle are shown for the VAWT with baseline blades (on the left), and for the VAWT with candidate blades (on the right).

The CFD solution begins, and the blade oscillates producing a transient flow field. When the blade reaches the predetermined snapshot location, the transient CFD simulation is paused, and an Adjoint solution is made generating the sensitivity data. The algorithm can produce sensitivity data at multiple snapshots but in Figure 5-10 three snapshots are shown for illustrative purposes. The Adjoint module can then combine the multiple sets of sensitivity data to produce a single set that the blade morphing process will use. The mesh morphing tool is then implemented to deform the blade geometry according to the combined sensitivity data. The new transient flow field is then produced for the updated blade geometry by running further oscillation cycles of the CFD model.

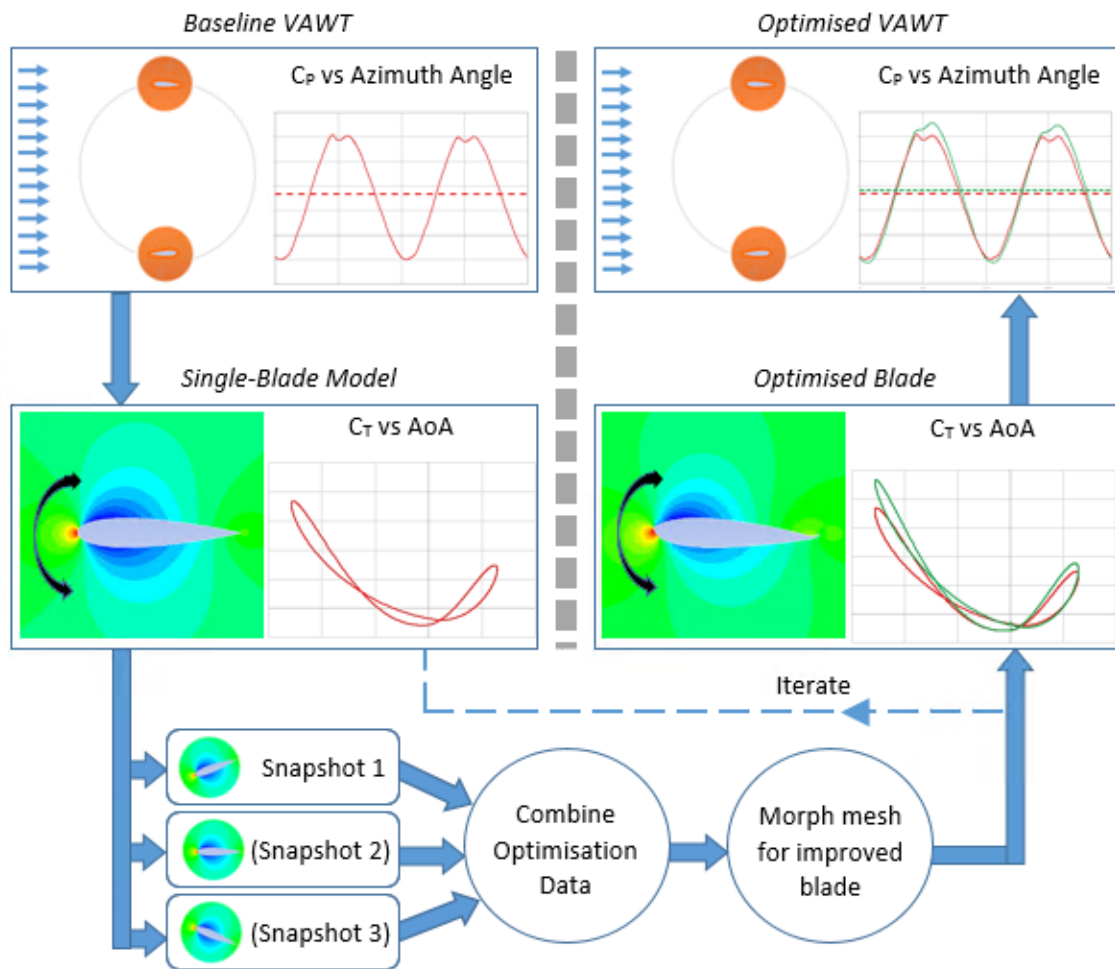


Figure 5-10 - General Schematic of the Single-Blade Optimisation Process

This process is repeated for several iterations until the C_T has improved, resulting in a 'candidate blade' which has an improved blade geometry compared to the baseline blade. At this stage the true performance of the blade geometry is not known due to the inaccuracies of the Single-Blade model (see Figure 5-6). This candidate blade geometry is then used to produce a 'candidate VAWT' model, so that the VAWT C_M (C_P) improvement can be evaluated accurately (top right of Figure 5-10).

An optimisation process would typically be run for as many iterations as required to reach convergence of the objective function. As mentioned before, with the Single-Blade method the true performance is not known until the candidate VAWT is made and tested. Time costs would however be prohibitive to make a candidate VAWT after every iteration of the Single-Blade optimisation process. To mitigate this, preliminary studies were conducted to decide an appropriate number of iterations to run the Single-Blade optimisation for. 10 iterations were chosen and used in the present work, as this approximated the optimum number of iterations for the range of cases tested.

5.3.1 Adjoint Module Setup

The Adjoint solver in ANSYS Fluent was implemented in this work which has various settings that require configuring for the optimisation process to function properly. There is a severe lack of guidance in the literature regarding these Adjoint module settings and how to apply it in practice, so for the present work the Adjoint module setup has largely been determined by performing numerous preliminary studies.

Table 5-3 summarises the Adjoint settings used, the details of which are discussed in the following sub sections.

Table 5-3 - Adjoint Module Settings Summary (Single-Blade Method)

Objective function:	Blade tangential force coefficient (C_T)
Target performance change:	+3% (of the objective function)
Adjoint solution iteration limit	1000 iterations
Adjoint solution stability scheme	Automatic
Geometric constraint	Constant chord length
Size of mesh morphing zone as a multiple of chord length	1.8c (x), 1.1c (y)
Number of control points in mesh morphing zone	100 (x), 100 (y)
Freeform Scaling Scheme	Objective reference change
Freeform Scale Factor	1

5.3.1.1 Adjoint Solver Settings

A 1000 iteration limit is implemented for computing the Adjoint solution. This value is a compromise between solution convergence, and CPU cost. Stabilisation scheme options are available with the Adjoint solver when the standard advancement scheme is unstable. The 'auto-assign' option is used for the Single-Blade model; if the solution divergence is detected during the calculation of the Adjoint solution then the most appropriate stabilisation scheme is applied automatically. The Adjoint solutions for upwind snapshots have generally reached convergence well within the 1000 iteration limit, and thus do not require stabilisation. The Adjoint solutions for downwind snapshots have generally converged or plateaued with residuals settling around $1e-4$, which is considered acceptable. The Adjoint solution convergence criteria for the Single-Blade model were $1e-5$ for Adjoint continuity and velocity, and $1e-3$ for Adjoint local flow rate (these are default values). Figure 5-15 (Section 5.4)

shows examples of Adjoint solution convergence. The Adjoint solution method was chosen as the 'standard' (for pressure) and 'first order' (for momentum). These settings are the defaults which offer the most stability. A more accurate (when converged) and less stable choice is the second-order method, but in this work stability is prioritised.

The Adjoint solution produces the sensitivity data as discussed in Section 3.1, and is itself not the solution to the aerodynamic optimisation problem. The sensitivity data must be applied to the aerofoil geometry via mesh morphing (see Section 5.3.1.2), and Figure 5-16 (Sections 5.4) shows examples of Adjoint sensitivity data.

5.3.1.2 Mesh Morphing Settings & Constraints

The mesh morphing tool within the ANSYS Fluent Adjoint module has two main elements that place limits/constraints onto the mesh morphing process. Firstly, the user can specify the size and position of the morphing region – the area where mesh morphing is permitted. This relates to the surrounding mesh that must morph to accommodate the changes to the blade wall geometry. The user can also specify constraints, which control the deformation of the blade wall to within prescribed limits. These two components are now discussed further with regard to the present optimisation problem.

When specifying the size and position of the morphing region it should surround the blade and a portion of the surrounding mesh. This allows the geometrical changes of the blade wall to be accommodated along with other necessary deformations in the surrounding mesh. The region used in the present work surrounds the blade wall and no other features/boundaries/interfaces of the CFD model and is shown in Figure 5-11(b). The dimensions of the morphing region used in this work are semi-arbitrary and were determined via preliminary studies (see Table 5-3).

Typical geometrical parameterisation does not take place nor is it required here. The sensitivity gradients are computed on a node by node basis, and the mesh morphing is performed on a similar basis via the set of Control Points. The user can specify up to 200 control points in the X and Y directions and these are automatically distributed across the defined morphing region.

To give the VAWT optimisation process a more real-world applicability, constraints should be included that represent the requirements from other engineering disciplines, outside those of pure aerodynamics. For the present works where the main focus is producing a functioning optimisation process, the constraints have been approached simply and only a chord constraint is implemented. This is because even if the blade profile was preserved during optimisation, a chord change would produce a performance change as it has altered the turbine solidity. To avoid this

kind of ‘false’ optimisation, the process here should operate at a constant solidity (i.e., chord) in order to produce a valuable outcome. Numerous VAWTs have a fixing angle of zero, but the literature has explored variations in non-zero fixing angle and seen how it can change to performance. For the present work it was decided that a constant fixing angle should *not* be imposed as a constraint. The optimisation process can provide the optimum fixing angle implicitly after the geometry changes have been made, and this could be beneficial to the performance.

To implement a chord constraint whilst allowing freedom of fixing angle and other geometry changes, a circular boundary is used which is shown in Figure 5-11 (a). The constraint must be imported into the Adjoint module as a separate file, and this constraint file (.cas) is actually a doughnut shaped surface.

Consideration must be given to whether the constraint is designated as “strict” or not. The present work uses non-strict conditions because the use of strict conditions can cause severe issues with mesh morphing – see Section 5.3.4.1.




Figure 5-11 – (a) Constraint Boundary (Left) and (b) Mesh Morphing Zone (Right)

The mesh morpher considers the sensitivity data and aggression settings (see Section 5.3.2) along with any constraints and calculates the new node locations accordingly.

The ‘Freeform scaling scheme’ is chosen as ‘objective reference change’. This means that the mesh deformations are based on the specified aggression (see Section 5.3.2). If the alternative option is chosen, then the mesh deformation is based on the number of control points. A ‘Freeform scaling factor’ of 1 is used.

The smoothness function of the geometry change can be defined by the user or if left by default is automatically determined. This smoothness corresponds to the permissible relative displacement of adjacent control points during mesh morphing. The sensitivity data itself may exhibit sharper variations in direction and magnitude

than would be desirable to see in the aerofoil wall, so this smoothness factor ensures no undue discontinuities occur in the geometry during the morphing stage. During the preliminary studies, negligible differences in blade geometry were found between using 'auto-smoothness' and smoothness factors of 5 and 10.

Numerical settings are available relating to the calculation of the mesh morphing deformations. The following values were arrived at during the preliminary constraints studies and are used throughout this work. These settings aim to improve the convergence of the mesh morphing solution and infer only a tiny increase in CPU cost:

Prescribed motion and Freeform motion settings:

- Max iterations set to 30. No impact on overall CPU time, the default was 5.
- Constraint relation set to 0.1. The value should be between 0 and 1, 1 being the most aggressive. The default is 1, but preliminary studies showed this proved difficult to reach a converged solution for the mesh morpher.
- Preconditioning set to 10. A value of 0 is most aggressive but 1 is typical. The defaults were 0 (prescribed motion preconditioning) and 0.5 (freeform motion preconditioning).
- Freeform motions also has a parameter relaxation setting, this is also set to 0.1.
- Tolerances are set to $1e-5$ while the defaults are $1e-3$, this aims to tighten up the gaps between the constraint zone boundary and the control points, achieving less violation of the constraint (Note that a small amount of chord constraint violation is permissible, see Section 13.1).

5.3.2 Aggression

The sensitivity data is used in the mesh morpher to deform the revised blade geometry in a way that pursues the specified target in the objective function (performance). This magnitude of this target effects the scaling of the geometry/mesh deformation that is implemented. In ANSYS Fluent this is referred to as the 'observable objective' and is effectively the level of aggression in the optimisation process. With reference to the discussion in Section 3, this aggression is akin to the scaling factor λ .

The objective can be specified in a few ways; 'increase value', 'decrease value' or 'target change in value'. In this work the desire is for the C_T (see Section 5.2.1) to be increased so 'increase value' is selected. 'Target change in value' is not used because this does not allow the prescription of weighting factors to be applied to each

snapshot. The functionality for user-defined weightings was desired for possible future investigations.

The magnitude of the 'target/reference change' can be specified as an absolute value, or as a percentage. For the Single-Blade, 1-snapshot investigation (results in Section 5.4) the 'target/reference change' is specified as +3% C_T . The mesh morpher attempts to implement a 3% increase in C_T by deforming the geometry, moving each wall surface node according to the sensitivity vectors considered - Figure 5-9 shows an example of these sensitivity vectors. The mesh morpher calculates the geometry deformations required to achieve the target using a first order approximation. Only small geometry changes should therefore be made across each iteration because of the associated loss of accuracy in the first order prediction for larger changes. Larger aggressions can lead to difficulties in the optimisation process due to 'overshooting' effects, which is explained in Section 6.3.2 .

It is important to note that the 3% increase should not be expected to manifest as an average C_P increase of the turbine. This target is relevant to the instantaneous flow field of the snapshot(s) and not a time averaged quantity. For a 1 snapshot case the user can expect a localised instantaneous increase in objective function (C_T) around the location of the snapshot.

The value of 3% aggression was determined via an investigation where various levels of aggression were tested. Several 1-snapshot optimisation tests were carried out starting from the same baseline model. Variations of aggression 2%, 3%, 4% were tested where candidate VAWTs were produced to measure the turbine performance improvements. Smaller aggression levels needed more iterations of the optimisation process to improve the performance, but the process would be more stable. A balance is ideal and 3% was found to produce the best improvement to average C_P across the range of cases tested, whilst avoiding 'overshoot' in the process (see Section 6.3.2 for discussion of 'overshoot'). This study was therefore not exhaustive, and 3% may not be the best value – but it provided a starting point for use in the Single-Blade method. A more thorough aggression study was then conducted during the work on the ISV (In-Situ VAWT optimisation) method (see Section 6.3.2).

5.3.3 Algorithm Code

The previous sections have described the setting up of the Single-Blade CFD model (Section 5.1), and the setting up of the Adjoint module (Section 5.3). This section outlines the algorithmic process used which envelopes CFD model and Adjoint module in order to form the Single-Blade optimisation process. This is a more detailed and specific version of the generalised optimisation process shown earlier

in Figure 5-8/Figure 5-10. This algorithm script takes the form of a journal file, which Fluent reads line by line as text input commands; this is written in ANSYS Fluents' internal programming language called Scheme (Dybvig, 1996).

Once the template for a correctly functioning algorithm script was determined, a spreadsheet tool was produced to automatically generate the text lines of new scripts based on user defined parameters. These parameters mostly relate to the desired number of snapshots and their positions. In this manner the scripts for a range of test cases could be produced quickly, after the initial time investment of producing the spreadsheet tool.

Figure 5-12 shows a flow diagram of the single-blade optimisation algorithm script; it is a less abstract version of Figure 5-8. The following discussion refers to stages in Figure 5-12 by the number identifier, **X**. The light and dark grey shaded boxes correspond to the main loops in the algorithm.

Precursor items required before running the optimisation algorithm/script include the following steps:

- Stage **1**. The Single-Blade CFD model should be produced according to Section 5.1.
- Stage **2**. After the baseline Single-Blade model has completed 3 cycles, the C_T as a function of AoA curve should be plotted for the baseline blade geometry. The 3rd revolution is sufficient to provide a periodically converged solution as has been shown in Section 5.1.3. The flow field data of this revolution forms the baseline blade performance. After each iteration the performance will be compared to this. This curve can be used to make choices about the number of snapshots, position of the snapshots, and relative weighting of snapshots (see Figure 5-7). Once chosen, the spreadsheet tool produces the specific script file accordingly.
- Stage **3**. The Adjoint module should be configured according to Section 5.3.1.

Script/Algorithm

- Stage **4**. The transient solution is continued by N iterations, until the snapshot position is reached. N depends on the desired position (see Figure 5-7) of the first snapshot (a spreadsheet is used to convert snapshot azimuthal angle or AoA into a value for N).
- Stage **5** & Stage **6**. When the snapshot position is reached, the transient solution is stopped temporarily. Then the steady flow solver is selected. This is because

the ANSYS Fluent Adjoint module is configured for use with the steady flow solver. Some Adjoint module functions do not work while the transient solver is engaged. Doing this does not invalidate the flow field data or preclude its use in an Adjoint calculation (as explained in Section 5.2). Having the steady solver active at this point in time allows the use of the Adjoint module, but switching between steady and transient flow solvers requires additional considerations which are discussed in Section 5.3.4.2.

- Stage **7**. Before the Adjoint solution is calculated, the mesh and flow field within the rotating subdomain is rotated. It is rotated from the AoA of the current snapshot to AoA=zero (but will be rotated back later at Stage **10**). This has the effect of aligning the blade with the global X, Y coordinate system. This is a necessary step because the sensitivity data produced from an Adjoint solution is prescribed in the global coordinate system. To combine sensitivity data in a correct way it must all be orientated in a consistent fashion (see Section 5.3.4.3), This method maintains consistency/alignment of the orientation of sensitivity data for all snapshot positions, so that combination of sensitivity data is correct. It should be noted that when rotating the domain in this way, specifying the objective function C_T in the Adjoint module is simplified. When the blade is orientated with 0 degrees, the C_T is -1 times the force coefficient in the global X direction.
- Stage **8** and Stage **9**. The Adjoint solution is calculated with the setup as described in Section 5.3.1. When this calculation is complete, the resulting sensitivity data is exported for later use.
- Stage **10**. The counterpart step to Stage **7**. Now that the snapshot is complete and the sensitivity data is obtained, the CFD solution can resume. Before the transient solution can continue marching towards the next snapshot position, the mesh and flow field are rotated back (from AoA=0) to the correct AoA (see Section 5.3.4.3).
- Stage **11**. The transient solver must be re-engaged before resuming the transient simulation (see Section 5.3.4.2).
- Stages **4** to **11** are repeated for however many snapshots are specified.

- Stage **12**. Once all snapshots have been taken, the transient solution continues until the blade is back at the AoA=zero position. At this position, the blade is aligned with the global X and Y coordinate system. All the sensitivity data calculated previously is therefore aligned with the blade at this position, ready for mesh morphing.
- Stage **13**. The steady solver is engaged as before, to allow full functionality of the Adjoint module. Sensitivity data from all snapshots is imported. Weightings are applied to each snapshot if desired.
- Stage **14**. The mesh morphing settings are implemented according to Section 5.3.1.2. The mesh is morphed according to the combined sensitivity data and aggression, providing an updated blade geometry. The sensitivity data is then cleared from the Adjoint module.
- Stage **15**. The transient solver is then re-engaged as before and a further 3 revolutions are computed for the new geometry. The flow field data from the final revolution is used to evaluate the new Single-Blade performance (C_T vs AoA curve).
- Stages **4** to **15** are repeated for X iterations. X is the number of Adjoint iterations specified by the user. Note that for the Single-Blade method the 'best' number of iterations cannot be known at this stage because candidate blade VAWT models must be created to evaluate the true performance of a candidate blade when affixed to a VAWT. A fixed value of 10 iterations is used in this work, explanation of this is given in Sections 5.3 and 5.4.1. Creation of the candidate VAWT is discussed separately because it is not contained *within* the algorithm for the Single-Blade optimisation.
- Stage **16**. The performance data for the optimisation so far is plotted showing how the C_T as a function of AoA curve has improved with each iteration.

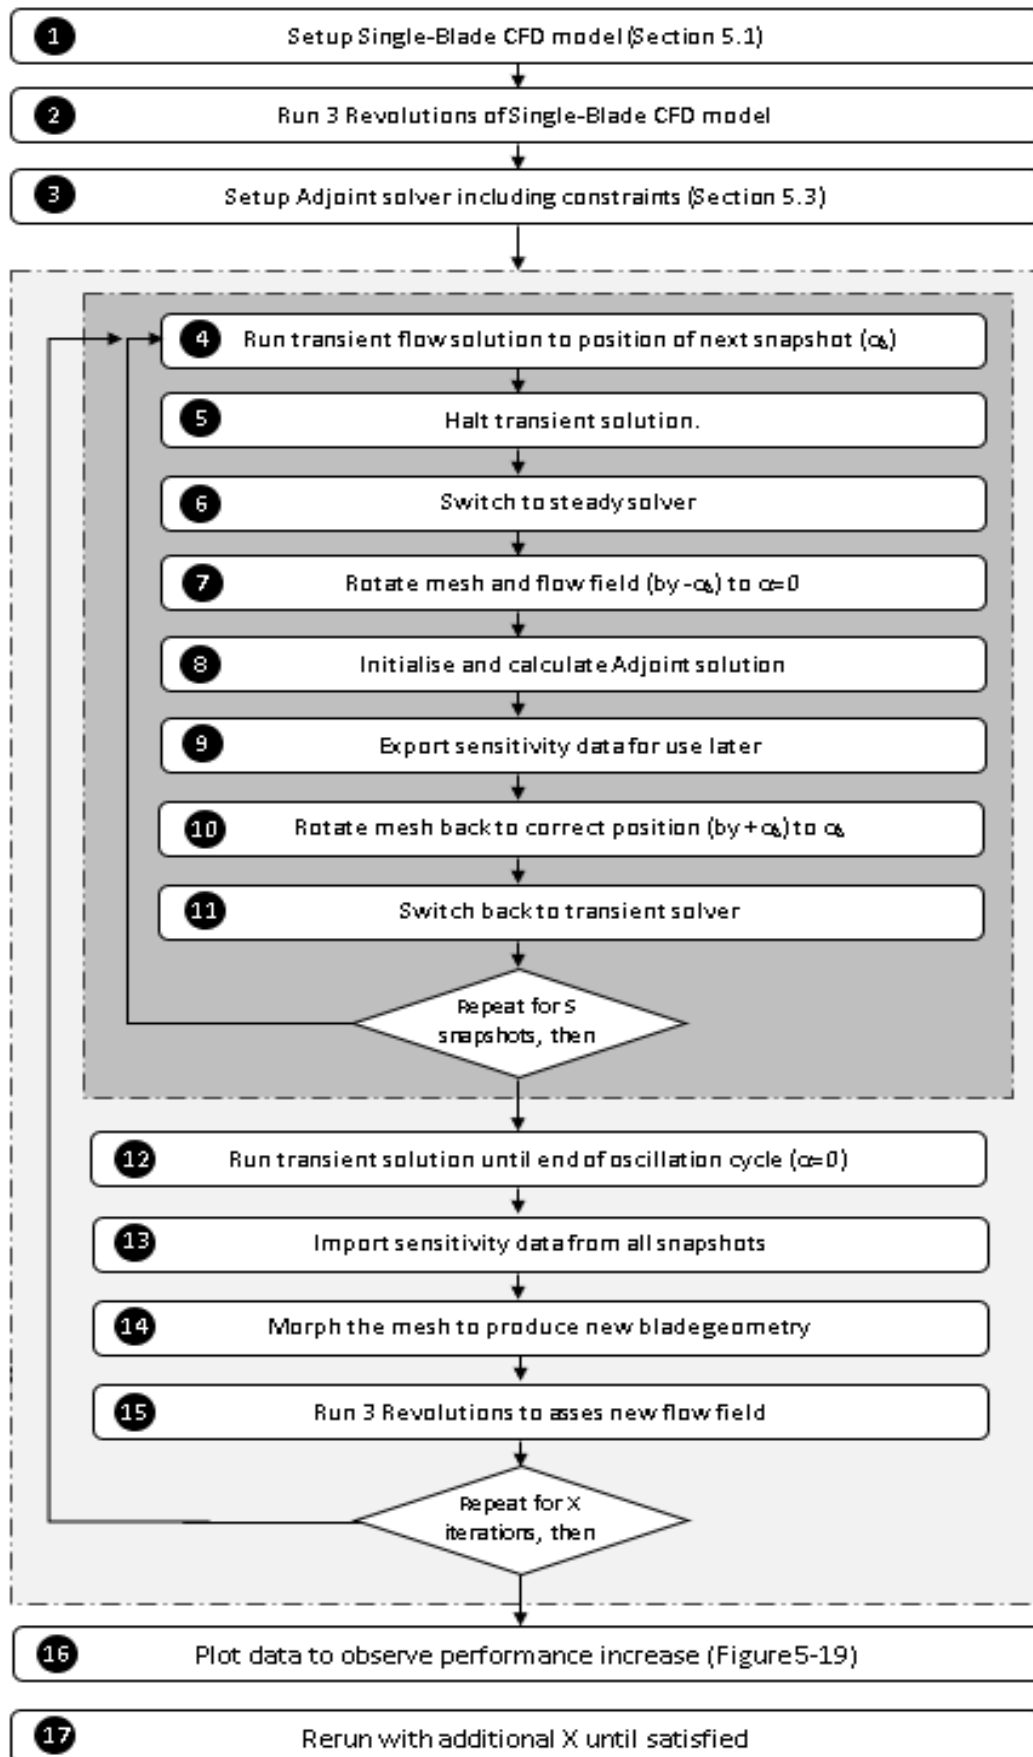


Figure 5-12 - Optimisation Process Flow Diagram

5.3.4 Algorithm Key Notes

This section expands upon some of the key elements in the algorithm which ensure the robustness and correct functionality of the optimisation process. Also discussed here are some of the ways that the algorithm was validated. A similar discussion is given later in Section 6.3.4 which includes specific considerations for the ISV method algorithm.

5.3.4.1 Mesh Morphing (Constraints/Negative Cell Volumes/Mesh Distortion)

To ensure that the mesh morphing functions correctly, tests were carried out to examine the mesh deformation in the regions where the blade wall is near to the constraint boundary (described in Section 5.3.1.2). This is at the leading and trailing edges of the blade, but the trailing edge was the main region of interest because of the highly refined mesh at this sharp geometrical feature. Investigations showed that correct mesh morphing is achieved by allowing the constraints to be implemented without imposing 'Strict' conditions to the constraint. If 'Strict' conditions are imposed this can cause unacceptable distortion of the mesh cells, specifically at the trailing edge. This deteriorates the solution quality and can often lead to the error 'Negative Cell Volumes' which halts the CFD process altogether. Further details on this point are presented in Appendix D.

5.3.4.2 Steady/Transient Solver Switching and Frame/Mesh Motion

As stated in Section 5.3.3, the CFD solution method is switched from transient to steady whenever an Adjoint solution is calculated (and also at other points in the algorithm). To ensure that the correct mesh motion settings are implemented within every mesh zone at all times, some additional commands are added to the code to specify 'frame motion' when the steady solver is used, and 'mesh motion' when the transient solver is used. This mitigates the possibility of incorrect motion settings being applied automatically when changing between transient/steady solution methods.

5.3.4.3 Adjoint Solution Correction for Global X-Y Alignment

The sensitivity data arising from an Adjoint solution is relative to the global X-Y coordinate system regardless of what orientation the blade is at. For CFD models with moving walls/mesh motion such as oscillating blades or VAWTs, this becomes a critical factor that must be considered. With the oscillating Single-Blade model, the optimisation algorithm must correct for the misalignment of the instantaneous blade orientation with the global X-Y orientation.

Figure 5-13 shows this potential misalignment issue if not mitigated, when optimising with 2 snapshots of different blade angles.

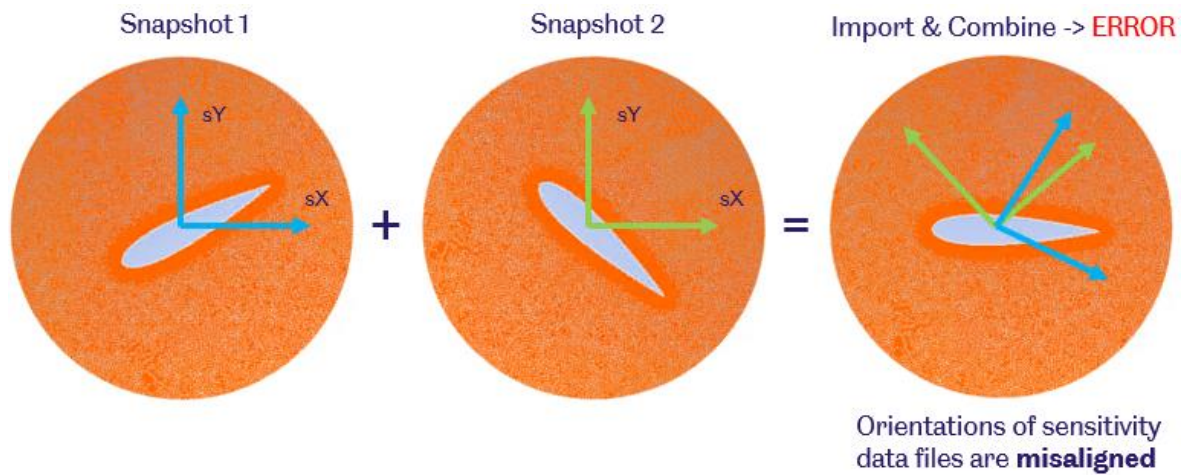


Figure 5-13 - Misalignment of Sensitivity Data

As stated in Section 5.3.3, this misalignment of sensitivity data is addressed in the algorithm by rotating the CFD model to the zero AoA position whenever an Adjoint solution is needed. This means that the entire mesh and flow field is rotated, and by doing so all the aerodynamic forces and sensitivity data are aligned in a consistent fashion, with the global X-Y axes. After a snapshot is complete, the model is rotated back to its correct position for the transient solution to continue. When the transient solution completes a revolution, the blade is returned naturally to the zero AoA position. At this point all the sensitivity data can be imported and combined correctly as per Figure 5-14.

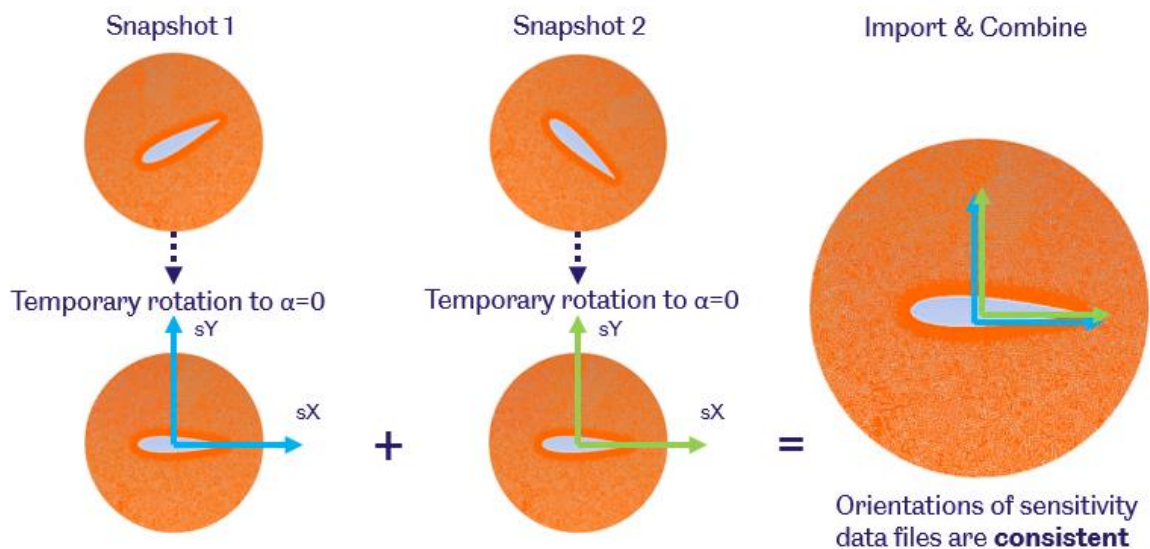


Figure 5-14 - Alignment of Sensitivity Data

The same principle is used with a slightly different implementation for the direct VAWT optimisation method, as discussed in Section 6.3.4. An important note here is that the angle of rotation to move the blade by (to the zero AoA position) must be known by the optimisation script. These angles are calculated via a spreadsheet when the script text file is generated, so each rotation angle is known for each of the

snapshots. These values depend on the UDF chosen for the oscillating AoA profile (see Section 5.1) and much care must be taken to calculate them correctly so that the correct X-Y alignment is achieved.

Algorithm validation exercises were conducted with the Adjoint module disabled, purely to ensure that the algorithm rotates the model by the correct angles, regardless of the desired snapshot positions or number of snapshots chosen. This ensured robustness of the script generator in adapting to the user inputs. Further validation investigations were conducted for the elements of the algorithm handling the Adjoint solutions, importing/exporting/combining of sensitivity data, and subsequent mesh morphing. For this, optimisation processes were conducted in a manual fashion on a desktop PC so that the critical steps were compared and validated to that produced via the algorithms executed on the HPC.

5.3.4.4 Snapshot Weightings Implementation / 'Observable Objective'

If functionality for user specified snapshot weightings is required, then within the Adjoint Design Tool the 'observable objective' type must be set to 'increase value'. Section 5.3.2 mentions that in the Adjoint Design Tool, different modes can be used to implement the 'observable objective'. This means the performance target that the optimisation process is pursuing. These modes include: 'increase value', 'decrease value' or 'target change in value' to the objective.

To ensure that user defined weightings are applied correctly to multiple snapshots, the 'increase value' mode must be selected. Testing showed that when using 'target change in value', the specification of weightings was ignored by the Adjoint module but this behaviour was not displayed in any warning messages, etc. Validation exercises were conducted to ensure that user specified weightings were implemented correctly. An optimisation was run with 2-snapshots, one upwind and one downwind snapshot, but with the weightings applied in a 99:1 fashion. The results matched closely with a single snapshot case of the same upwind snapshot.

5.3.4.5 Objective Sense-Check

Testing was carried out at numerous points during the algorithm development process to ensure that the optimised blade shapes were correctly pursuing the desired goal. Initially this entailed the use of static blade models (no mesh motion) and the trial of several different objectives (lift, drag, tangential force, etc.) to observe the variation of optimised blade performance under each objective. The performance as predicted by the Adjoint solver was then checked via manual calculation using standard CFD monitor/output data. Additional sense-checks were made to ensure the sensitivity data produced by the Adjoint solutions was consistent with the optimised geometries resulting from the process.

5.4 Single-Snapshot Investigation

This section uses the Single-Blade optimisation method (setup according to Sections 5.1, 5.2 and 5.3) with 1 Adjoint snapshot per oscillation cycle to investigate the effects of snapshot location on the outcome of the optimisation process. A multiple-snapshot investigation is presented in Chapter 7. The present 1 snapshot investigation allows a clear illustration of the cause and effect of snapshot choice, and performance changes. A range of cases are investigated where the location of the snapshot is varied. These 12 cases are listed in Table 5-4. Note that in the following discussion reference is made to the cases by virtual azimuthal angle – the Single-Blade model itself has no azimuthal angle but rotates in a way that approximates to the azimuthal angle variation of a VAWT.

As described in Section 5.1, a knockdown factor is applied to the AoA over the downwind region. This knockdown factor is why the snapshot AoA values are of lower magnitude in the downwind region (between 180 and 360 virtual azimuthal angle) than the upwind.

Table 5-4 - List of Test Cases for Single-Blade Optimisation

Case Name /Azimuthal Angle (degrees)	AoA of Blade at Snapshot (degrees)	Virtual Azimuthal Angle (for corresponding VAWT) (degrees)
0	0	0
30	-5.4	30
60	-9.8	60
90	-12.5	90
120	-12.2	120
150	-7.9	150
180	0	180
210	+3.9	210
240	+6.1	240
270	+6.3	270
300	+4.9	300
330	+2.7	330

Across these cases, the snapshot location differs by 1/12th of an oscillation cycle which corresponds to a gap of 30 degrees azimuthal angle. Although 360 cases (1 for each degree) or even more cases could be studied over the cycle, it is judged that 12 cases offers a reasonable balance between time cost of this investigation, and a sufficient resolution in studying the solution space.

Section 5.3.1.1 has mentioned the Adjoint module setup, and how convergence of the Adjoint solution is generally acceptable. Figure 5-15 shows the Adjoint solution convergence plots for a typical upwind and a typical downwind case. As will be seen in due course, there is a valid bases for grouping sets of results together to some extent, as the upwind and the downwind cases each share typical behaviours.

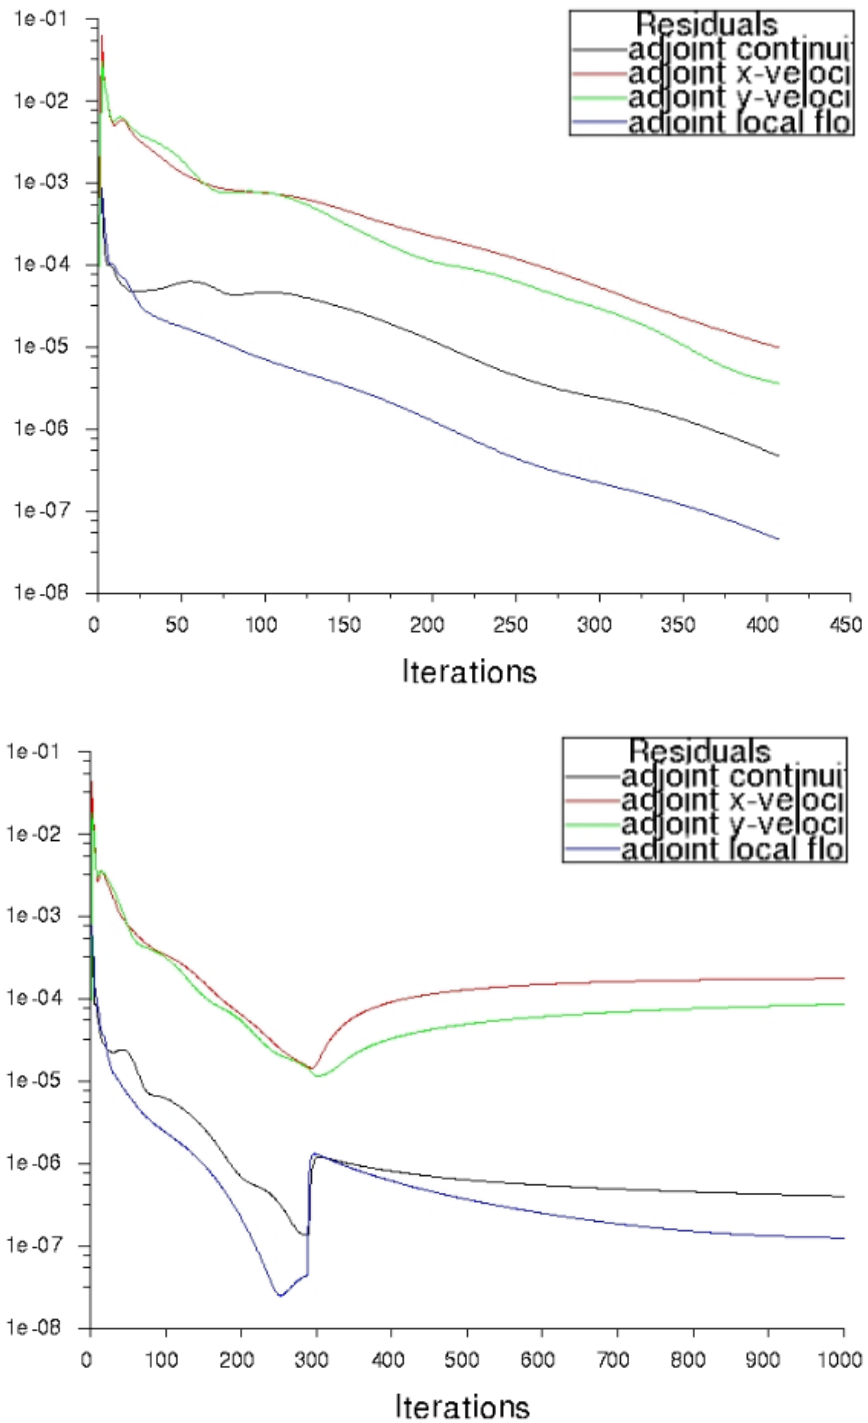


Figure 5-15 - Single-Blade Optimisation Adjoint Solution Convergence (Top) Case 90, and (Bottom) Case 270

To correspond with the convergence plots shown in Figure 5-15, Figure 5-16 displays the fields of sensitivity data generated from those Adjoint solutions. Vectors of shape sensitivity magnitude are shown, and this indicates the direction in which small geometry changes would provide improvements to C_T .

The mesh morpher applies smoothing to the geometry deformation so that the highly varying magnitudes of the sensitivity vectors do not produce discontinuities in the geometry.

After the iterations of the optimisation process are completed, the process is halted and the performance data can be processed. Table 5-5 shows the changes to the tangential force coefficient (C_T) of the Single-Blade model arising from the optimisation process. These values are maximum and average C_T of the Single-oscillating Blade. For negative values of “max C_T improvement (%)” this means that the upwind max C_T has deteriorated - such cases tend to show an improvement in the downwind C_T . An example of this is shown by Case 210 in Figure 5-17(h).




Figure 5-16 - Sensitivity Vectors, (Left) Case 90 and (Right) Case 270


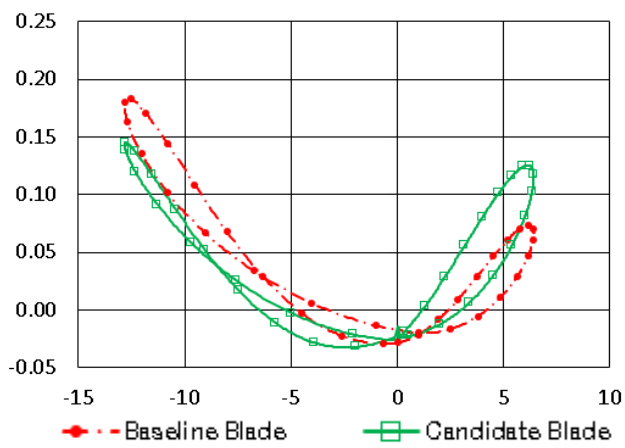
Table 5-5 – Results of 1-snapshot Single-Blade Optimisation

Case Name /Azimuthal Angle (degrees)	Max C_T improvement	Average C_T Improvement
0	-21.3 %	+6.0 %
30	-12.5 %	+2.6 %
60	+4.9 %	+8.1 %
90	+6.9 %	+9.5 %
120	+14.6 %	+11.9 %
150	+4.5 %	+7.5 %
180	-1.1 %	+8.1 %
210	-28.6 %	+1.2 %
240	-17.7 %	+2.2 %
270	-22.3 %	+2.3 %
300	-20.1 %	+2.3 %
330	-16.4 %	+2.1 %



Further details can be presented for each of the cases. Figure 5-17 shows the C_T vs AoA performance curves (left), and the candidate blade geometry (right).

Throughout Figure 5-17 (a) to (h) the red and green curves represent the baseline blade and candidate blade performance, respectively. The Y and X axes are C_T and AoA, respectively. This AoA is the geometric AoA with the downwind knockdown factor applied (described in Section 5.1), and it does not account for fixing angle changes due to the optimisation i.e., the AoA is relative to the tangential velocity vector ($\omega.R$). The contour plots illustrate the blade geometries and are shown at the 0degree azimuthal position for all cases. The non-dimensional static pressure field is calculated by normalising the local pressure with an approximate but consistent maximum pressure value of 1700 Pa. The top surface corresponds to the outside of the turbine, and the bottom surface corresponds to the inside of the turbine (towards the shaft). For the baseline blade these contour plots would show a zero fixing angle, but the optimised blades tend to have non-zero fixing angles.


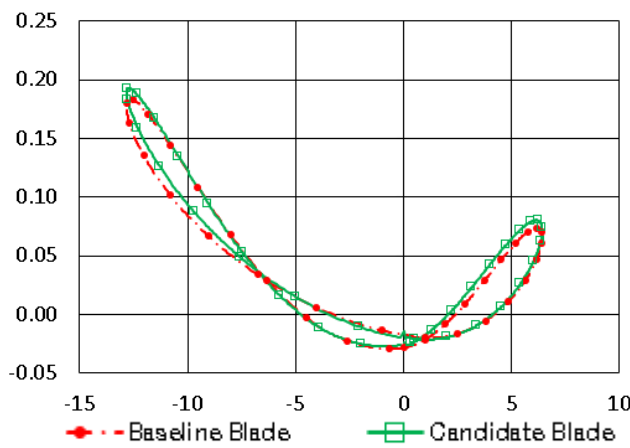
(a, $\theta = 0^\circ$, $\alpha = 0^\circ$) C_T against AoA




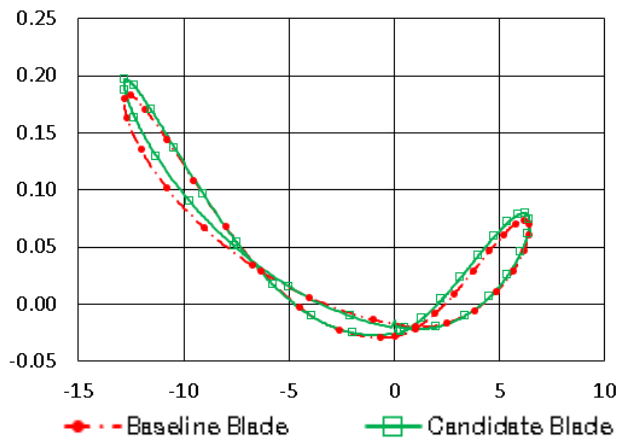
(b, $\theta = 30^\circ$, $\alpha = -5.4^\circ$) C_T against AoA





(c, $\theta = 60^\circ$, $\alpha = -9.8^\circ$) C_T against AoA





(d, $\theta = 90^\circ$, $\alpha = -12.5^\circ$) C_T against AoA





(e, $\theta = 120^\circ$, $\alpha = -12.2^\circ$) C_T against AoA





(f, $\theta = 150^\circ$, $\alpha = -7.9^\circ$) C_T against AoA





(g, $\theta = 180^\circ$, $\alpha = 0^\circ$) C_T against AoA




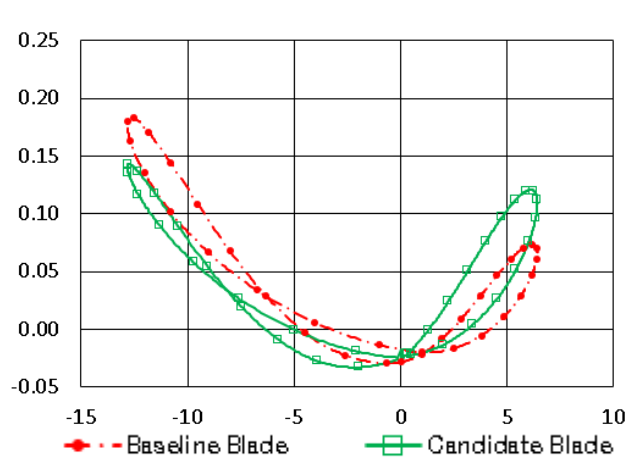
(h, $\theta = 210^\circ$, $\alpha = 3.9^\circ$) C_T against AoA





(i, $\theta = 240^\circ$, $\alpha = 6.1^\circ$) C_T against AoA



(j, $\theta = 270^\circ$, $\alpha = 6.3^\circ$) C_T against AoA



(k, $\theta = 300^\circ$, $\alpha = 4.9^\circ$) C_T against AoA



(l, $\theta = 330^\circ$, $\alpha = 2.7^\circ$) C_T against AoA





Figure 5-17 (a) to (l) – 1-Snapshot Investigation Results, Single-Blade Optimisation

These results can be considered to form two major groups, namely the cases with upwind (UW) snapshot positions, and the cases with downwind (DW) snapshot positions.

- UW snapshot cases appear to be characterised by small improvement to the C_T curve in the UW part of the cycle and there is a minor effect on the DW part of the cycle. Cases 60-150 exhibit this trend and the blades have a prominent negative camber (see Figure 1-7 for camber definition).
- DW snapshot cases appear to be characterised by a large improvement to the C_T curve in the DW part of the cycle and a deterioration in the UW part. Cases 180-330 exhibit this trend, and the blades have a small positive camber.
- Cases 0 and 30 appear to behave like the DW cases. These show a slight negative camber like the UW cases, but suffer a reduced UW performance, and this is possibly due to hysteresis.
- All cases improve the C_T at the position of the snapshot and in the range of azimuthal angles surrounding that position.
- All cases show a tendency to alter the fixing angle of the blade making it negative/toe-out (see Figure 1-8).
- The emphasis of the camber appears to be at the trailing edge, suggesting that the performance is most sensitive to trailing edge geometry changes.

With this single snapshot optimisation, it was found that an improvement to the tangential force coefficient (C_T) was observed at a range of AoA near the snapshot location. Elsewhere in the cycle, the C_T improved/deteriorated slightly at various points. This is because no optimisation data was considered from such regions away from the snapshot position. In essence, applying 1-snapshot optimisation produces a series of blade shapes showing, for each azimuthal position, what geometric features will improve the instantaneous blade torque.

Following the single-blade optimisation process, candidate VAWTs can be produced from the optimised blade geometries for each case, and thereby measuring the improvements in turbine performance.

5.4.1 Candidate VAWT Production

Candidate VAWT models are needed because although the Single-Blade model provides a reasonable approximation, it does not accurately reflect the VAWT flow field specifically in the downwind part of the cycle. A Candidate VAWT model

provides the true evaluation of average C_P (C_M) for the candidate blades. As described in Section 5.3, each candidate VAWT was made after 10 iterations of the single blade optimisation process.

The candidate VAWT model is made in a consistent fashion to the baseline VAWT model (see Section 6.1). Since the baseline VAWT model is fully validated this provides the only possible means of validating the candidate VAWT, in lieu of experimental campaigns. Since the candidate blade geometry in this work is entirely novel, there does not exist any experimental or computational data to validate with. The only difference between the baseline VAWT and candidate VAWT models is the way the blade geometry is constructed. The candidate blade geometry is made by opening ANSYS Fluent's 'post-view' module and importing the '.dat' output file which is produced by the Single-Blade optimisation process. The coordinates of the candidate blade can then be exported to .csv format from post-view and transformed into a new input geometry file.

The results of Candidate VAWT performance for the cases described within Table 5-5, are presented in Section 5.4.2.

5.4.2 Candidate VAWT Results – Single Snapshot Optimisation

Table 5-6 presents the results for the candidate VAWTs which were created according to the candidate blade geometries from each Single-Blade optimisation case. Values are given as a percentage, compared to the baseline VAWT with NACA0018 blades.

Table 5-6 – Candidate VAWT, List of Test Cases and Results

Case Name /Azimuthal Angle (degrees)	AoA of Snapshot (degrees)	Max VAWT C_M change	Average VAWT C_P Improvement
0	0	-5.1 %	+3.5 %
30	-5.3	-3.2 %	+1.4 %
60	-9.8	+2.2 %	+2.4 %
90	-12.5	+2.4 %	+2.1 %
120	-12.2	+6.0 %	+2.3 %
150	-7.9	+1.8 %	+2.2 %
180	0	+0.3 %	+3.0 %
210	+3.9	-7.6 %	+1.5 %
240	+6.1	-4.7 %	+1.3 %
270	+6.3	-5.7 %	+1.3 %
300	+4.9	-5.2 %	+1.3 %
330	+2.7	-4.3 %	+1.3 %


It should be noted that C_P is correlated to C_M , by calculating the moment (total contribution of both blades) as per Equation 1.5, and then using this value as the torque in Equation 1.6.

From these results it can be seen that Case 0 produces the greatest improvement to the average C_P , of 3.5%. The results however show a general increase in the average C_P occurring for all the cases but there is significant variation on how this is achieved.


Figure 5-18 shows performance curves of 1 blade C_M and VAWT (2-bladed) C_M for each candidate VAWT. The corresponding blade geometries are shown in previously Figure 5-17. It should be noted that since C_T was optimised in the Single-Blade model, it is expected that this will manifest as a localised C_M improvement in the region around the snapshot positions (azimuthal angle) in the 1-blade C_M graphs. The combination of the second blades aerodynamic forces concurrently cannot be considered by the single-blade model. The candidate VAWT data allows us to explore cases that result in average (combined forces of all blades) performance improvements, thus offering much insight into the most appropriate blade shapes.

Throughout Figure 5-18 (a) to (l) the red and green curves represent baseline and candidate blade performance, respectively. The Y and X axes are C_M and Azimuthal Angle, respectively.


(a, $\theta = 0^\circ$)





(b, $\theta = 30^\circ$)





(c, $\theta = 60^\circ$)





(d, $\theta = 90^\circ$)





(e, $\theta = 120^\circ$)





(f, $\theta = 150^\circ$)





(g, $\theta = 180^\circ$)





(h, $\theta = 210^\circ$)





(i, $\theta = 240^\circ$)



(j, $\theta = 270^\circ$)



(k, $\theta = 300^\circ$)



(l, $\theta = 330^\circ$)





Figure 5-18 (a) to (x) – 1-Snapshot Investigation Results, Candidate VAWTs

On considering these results, the UW snapshots tend to improve the upwind performance slightly, and parts of the DW also remain similar or improve slightly. DW snapshots produce a more severe effect on the performance curve, where the DW performance improves significantly, but the UW performance deteriorates significantly. DW snapshots therefore produce a more even generation of power over the cycle which would offer significantly reduced demand on the electrical generator and lower fatigue loading on the structure. These conclusions are consistent with those stated in Section 5.4 with regard to the Single-Blade C_T vs AoA curves of Figure 5-17.

Some additional observations can be made about the results which will be of use in further developments of the method:

- All cases have resulted in an improvement to the VAWT average C_P .
- In cases where a large max C_T increase is predicted (in the upwind) with the Single-Blade model, this corresponds to an increased max C_M of the candidate VAWT - but not of the same magnitude.
- The average C_T improvements seen in the Single-Blade model do not translate to similar improvements in the average C_M (or C_P) in the VAWT model. This is because the Single-Blade model offers a VAWT approximation but has inaccuracies mainly in the downwind region (see Figure 5-3/Figure 5-6).
- Upwind snapshot cases generally improve the C_M over the majority of the revolution but only by a small amount.
- Downwind snapshot cases generally improve the C_M in the downwind, while deteriorating in the upwind. This reduces the peaks/troughs magnitude of the power curve.
- For the cases shown, a C_P improvement more evenly spread across the entire cycle (both upwind and downwind) can prove to be more beneficial in realising a greater average C_P improvement.

The results of the 1-snapshot investigation using the Single-Blade method show that the semi-transient Adjoint based optimisation process can be successful in improving the average power coefficient of a VAWT. The results provide insight into how the snapshot position effects the VAWT blade performance. Discussion of the blade geometries and aerodynamics behind these performance effects are provided in Section 5.4.3.

Some key limitations exist which should be noted regarding the method presented in this chapter which are discussed in Section 8.1.

5.4.3 Aerodynamic Analysis for Single-Blade, 1-Snapshot Optimisation

As seen in Section 5.4.3, the cases with upwind snapshot locations and downwind snapshot locations fall into two distinct groups. The following discussion considers the aerodynamics in detail, of a representative case from each of these groups.

Figure 5-18 (d) (shown previously) is the graph of the C_M as a function of the azimuthal angle for Case 90 which is the representative upwind case. The contribution from only one of the two blades is shown on the left and the combined effect of both blades is shown on the right. A 2.1% increase to the average VAWT C_P was made after 10 iterations of the optimisation process and the candidate blade geometry that gave this improvement has been presented in Figure 5-17 (d). The candidate blade has a toe-out fixing angle of 2° , a maximum camber of 2.0% chord located at 80% chord (towards the trailing edge), and a maximum thickness increase of 1% compared to the baseline NACA0018 blade.

To analyse the aerodynamics implications of this blade geometry, Figure 5-19 shows the streamlines (which are coloured by non-dimensional static pressure) for the Case 90 candidate blade, as well as the baseline NACA0018 blade when at the snapshot position of 90° azimuthal angle. The corresponding surface pressure coefficients plots are presented in Figure 5-20. The surface pressures plots for the candidate blade show a weaker negative pressure on the suction side, at the leading edge. The increased fixing angle reduces the AoA at this position, meaning less curvature is required to pass around the leading edge. This decreases the leading edge suction but allows a greater suction to be maintained along the mid-span and further aft approaching the trailing edge. The negative camber helps recover torque lost by the relaxed AoA from the fixing angle, and this is conveyed by the close matching of pressures along the mid chord. Towards the aft region, the positive pressure on the top surface is greater than the baseline blade because of the camber. A high pressure zone can be observed as the flow is slowed down by the cambered tail on the top surface of the candidate blade around the position of greatest camber. This is coupled with a greater suction magnitude (for the candidate blade) towards the trailing edge, such that a more favourable magnitude and direction of pressure gradient is achieved. In addition, the trailing edge geometry slightly changes the size and shape of a small recirculating region but the effect on the surface pressures is minor.

Figure 5-21 shows the streamlines (which are coloured by non-dimensional static pressure), and Figure 5-22 shows the surface pressures for the Case 90 candidate

blade, at the position of 270° azimuthal angle. Note: please do not confuse this with information relating to Case 270 – the analysis at positions of 90° and 270° azimuthal angle will be repeated in due course for the candidate blade resulting from Case 270.




Figure 5-19 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Candidate blade from Case 90. (Bottom) Baseline blade




Figure 5-20 - Case 90 candidate blade surface pressure coefficient at 90° azimuthal angle




Figure 5-21- VAWT blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Candidate blade from Case 90. (Bottom) Baseline blade




Figure 5-22 - Case 90 candidate blade surface pressure coefficient at 270° azimuthal angle

At the position of 270 azimuthal angle, the candidate blade (Case 90) has a larger suction peak on the top surface of the blade at the leading edge (Figure 5-22). This is due to the fixing angle of the candidate blade; at this position in the downwind region the fixing angle demands more curvature around the leading edge. In the mid chord, despite the negative camber being disadvantageous at this AoA, the fixing angle change still manifests a comparable amount of suction, and a greater amount of positive pressure compared to the baseline blade. The large suction and pressure region located at 0.8 chord (Figure 5-22) corresponds to the position of maximum camber. The camber effect therefore produces greater pressure gradients in the candidate blade, moving from the mid-chord towards the trailing edge. These exert a greater 'pushing' component onto the rear face of the blade which adds additional torque. The candidate blade geometry of Case 90 is therefore aerodynamically advantageous when considering the effects over a complete revolution, producing a greater average C_P compared to the VAWT with baseline (NACA0018) blade.

Figure 5-18 (j) (shown previously) is the graph of the C_M as a function of the azimuthal angle for Case 270 which is the representative downwind case. The contribution from only one of the two blades is shown on the left and the combined effect of both blades is shown on the right. A 1.3% increase in the average VAWT C_P was made after 10 iterations of the optimisation process and the candidate blade geometry that gave this improvement has been presented in Figure 5-17 (j). The candidate blade has a toe-out fixing angle of 2.5° , a maximum camber of 0.9% chord located at 77% chord (towards the trailing edge), and a maximum thickness increase of 0.5% compared to the baseline NACA0018 blade.

To analyse the aerodynamics implications of this blade geometry, Figure 5-23 shows the streamlines (which are coloured by non-dimensional static pressure) for the Case 270 candidate blade, as well as the baseline NACA0018 blade when at the snapshot position of 90° azimuthal angle. The corresponding surface pressure coefficients plots are presented in Figure 5-24. The surface pressures plots for the candidate blade show a weaker negative pressure on the suction side, at the leading edge. The increased fixing angle means relaxing the AoA in this position so that less curvature is required to pass around the leading edge. This decreases the leading edge suction. A weaker positive pressure on the upper surface is also observed through the mid-span and further aft approaching the trailing edge. At the trailing edge the size of the small recirculating region is reduced for the candidate blade, also due to the relaxed AoA provided by the fixing angle. Overall, at 90 degrees, and over the upwind in general, the Case 270 candidate blade provides less torque than the baseline blade. This is expected at the opposite part of the cycle to where the snapshot is taken.




Figure 5-23 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Candidate blade from Case 270. (Bottom) Baseline blade




Figure 5-24 - Case 270 candidate blade surface pressure coefficient at 90° azimuthal angle




Figure 5-25 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Candidate blade from Case 270. (Bottom) Baseline blade




Figure 5-26 - Case 270 candidate blade surface pressure coefficient at 270° azimuthal angle

Figure 5-25 shows the streamlines (which are coloured by non-dimensional static pressure), and Figure 5-26 shows the surface pressures for the Case 270 candidate blade, at the position of 270° azimuthal angle.

At the position of 270 azimuthal angle, the candidate blade (Case 270) has a much larger suction peak on the top surface of the blade at the leading edge (Figure 5-26). Also, this is clearly visible in the pressure contours (Figure 5-25). This is due to the fixing angle of the candidate blade; at this position in the downwind region the fixing angle demands more curvature around the leading edge. The camber of the blade is minimal, and as such, there is no large suction and pressure region located towards the rear of the blade. The effects on the pressures towards the aft region are minor in comparison with the baseline blade. The candidate blade geometry of Case 270 is therefore aerodynamically advantageous when considering the effects over a complete revolution, producing a greater average C_P compared to the VAWT with the baseline (NACA0018) blade.

5.5 Concluding Remarks on 1-Snapshot Optimisation, Single-Blade Method

The results presented are significant because they demonstrate the successful use of Adjoint based optimisation for VAWT aerodynamics. Furthermore, using commercial CFD software, feasibility has been shown for a promising semi-transient optimisation process by implementing the crudest form of the method (1 Adjoint snapshot per revolution). Therefore, great opportunity exists for improving and refining the method further, leading to associated increases to VAWT performance. Novel VAWT blade geometries have been generated and discussion has been made linking the performance characteristics to features of the blade geometry.

The semi-transient Adjoint based optimisation method with 1-snapshot per revolution is highly effective at increasing the instantaneous tangential coefficient for the oscillating blade model. This holds value for aerospace and other applications where engineers can easily optimise surface geometries for operation in unsteady flow conditions, and this includes both internal and external flows. The results also show that the method can increase the average C_M (C_P) and can reduce the aerodynamic loading while maintaining the VAWT average power coefficient. This carries a large value since fatigue loading of VAWT structures is reduced for more even aerodynamic blade loading over the revolution, and the significance of this structural challenge is highlighted in the literature (MacPhee & Beyene, 2018).

Considering the range of results from the Candidate VAWTs, an 'upwind/downwind paradox' is observed. This means that improvements to upwind performance tend to decrease downwind performance, and vice versa. This challenge is a fundamental drawback of the fixed-blade VAWT concept, and in this case of 1-snapshot optimisation meant that the best improvement to average C_P was 3.5%. While the effects of this design paradox reduce with increasing TSR, it cannot be directly avoided without use of a control system. With this in mind the 1-snapshot optimisation results hold value for potential control system applications, since this semi-transient Adjoint based method is very efficient at determining the set of desired geometries across the various positions in the VAWT revolution. This set of resulting blade geometries provides the basis for an active control geometry profile that aims to maximise the C_T at each azimuthal angle of the turbine. Investigation of this is recommended as future work (see Chapter 8).

This work has made important progress by setting the foundation for taking advantage of the powerful Adjoint method in VAWT design and optimisation. These methods could help unlock new knowledge and design efficiency of VAWTs which other types of method cannot provide.

6 IN-SITU VAWT (ISV) OPTIMISATION

The In-situ VAWT (ISV) optimisation method applies directly to a VAWT CFD model, as opposed to the Single-Blade method which has been presented in Chapter 5. The same semi-transient Adjoint based optimisation philosophy is used, but the implementation is different for the ISV method. Despite the relative complexity and cost of ISV optimisation, this method provides a more accurate flow field and hence the more accurate Adjoint sensitivity data.

This Chapter is divided up into the following sections:

- Section 6.1 – The modelling philosophy relating to the VAWT is described briefly.
- Section 6.2 – The implementation of Semi-Transient Adjoint Optimisation when applied directly to a VAWT model is described. Differences between the ISV and Single-Blade method are highlighted.
- Section 6.3 – Optimisation Algorithm; describes the architecture of the optimisation process algorithm, for the ISV method. This includes setup details of the Adjoint solver, optimiser aggression and mesh morphing. Practical guidance/ best practice for using the ANSYS Fluent Adjoint module in this context is also provided. Differences between the ISV and Single-Blade method are highlighted.
- Section 6.4 – Single-Snapshot Investigation; describes the application of the ISV method where 1 snapshot is used per turbine cycle. Results are presented for a range of cases based on varying the snapshot position. Note that the Multi-Snapshot Optimisation investigation is presented separately in Chapter 7.

6.1 VAWT Modelling Philosophy and Validation

A VAWT CFD model with baseline NACA0018 blades is constructed and validated before the ISV optimisation process is applied to it. Note that the data from this VAWT simulation is also used to validate the Single-Blade model described earlier in Chapter 5 (see Figure 5-3 and Figure 5-6). The validated VAWT modelling method is also used to generate the candidate blade VAWT models used for the Single-Blade method.

Table 6-1 will summarise the key features of the CFD setup adopted for the VAWT models. These are similar to those used for Single-Blade model (Table 5-2), but with some changes.

Table 6-1 - Summary of CFD Setup for VAWT Models

Setup Factor	Choice Adopted
2D/3D	2D
Turbulence Model	SST k- ω
Maximum mesh y^+	1.5
Number of cells around aerofoil wall	475
Cell count	259,000
Domain	(See Section 6.1)
Time step size	800 steps/rev
Target CFL	5
Number of cycles for periodic convergence	10
Inlet boundary condition	Wind velocity 9.3 m/s
Outlet boundary condition	Pressure outlet
Solver type	Pressure based
Pressure-velocity coupling scheme	Coupled
Spatial discretisation scheme	Second order upwind
Temporal discretisation scheme	Bounded second order implicit
Number of iterations per time step	30
Convergence criteria	1e-5 (for all residuals)

The domain of the VAWT model and dimensions used are shown in Figure 6-1 and Table 6-2, respectively.

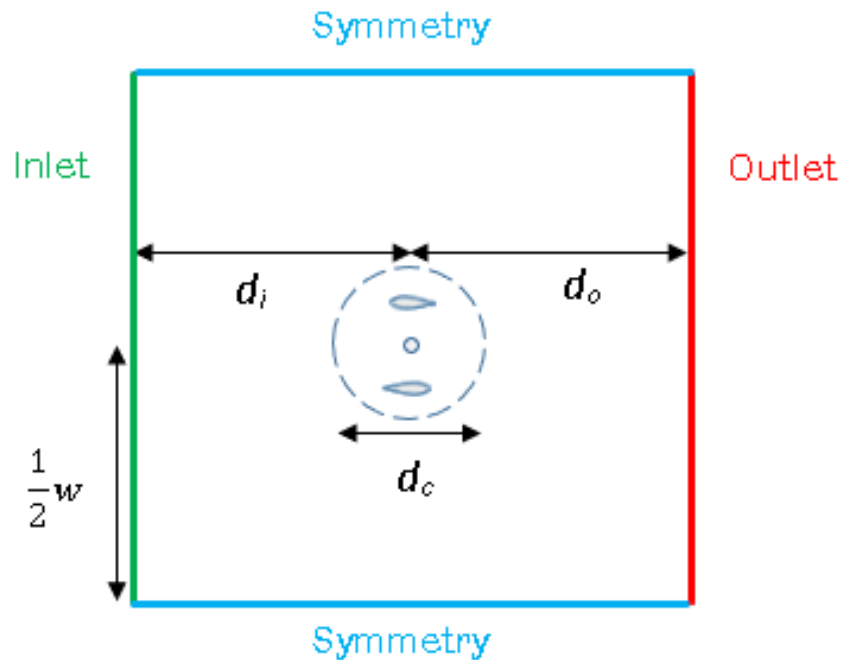


Figure 6-1 - Domain Dimensions

Table 6-2 - Domain Dimensions

Dimension	*
d_c , Rotating subdomain diameter	1.5D
d_i , Distance from turbine centre to inlet	10D
d_o , Distance from turbine centre to outlet	10D
$\frac{1}{2}w$, Half width of domain	10D

*Dimensions are presented as a multiple of the turbine diameter (x1.0m)

This domain is meshed in a similar way to that described for the Single-Blade. The same size of near-wall boundary zone is used and the same y^+ is achieved. For consistency between the two models, the VAWT model also uses the $k-\omega$ SST turbulence model. The detailed discussion on the general model setup is not recounted here except for specific areas regarding the ISV method.

Once a VAWT model has been produced with consideration to the above, necessary validation investigations must be conducted:

- Periodic solution convergence study
- Grid independence study
- Time step independence study
- Courant number study
- Domain size study

Independence studies of mesh and time-step were conducted. For the example case of Rezaeiha *et al.* (2017) used presently, the reference paper contains a thorough domain size study exploring each of the domain dimensions independently in order to provide appropriate recommendations. This provided the dimensions for the VAWT model domain. Therefore, it is deemed unnecessary for the present work to recount or reconstruct this domain independence study; the final dimensions used, as recommended by Rezaeiha *et al.* (2017), and are presented in Table 6-2 and Figure 6-1.

The other validation studies are routine within the literature, so the details are provided in Appendix B. The outcomes of this work demonstrate suitable validation for the Single-Blade CFD model when setup as follows:

- 10 revolutions were sufficient for a periodically converged solution
- Domain, 20D wide, 20D length, 1.5D rotating subdomain, circular refinement zone around blade = 6c, boundary zone thickness = 0.05c
- Y^+ of around 1 (justification in literature review)
- CFL number of 5 in the Coupled solver
- Time step corresponding to 800 steps/revolution
- 475 cells around the whole aerofoil surface
- Total cell count of approx. 259,000
- Quad elements used throughout all subdomains (justification in literature)

Having demonstrated mesh, time-step and domain size convergence of the solution, it is then necessary to check the performance prediction against that of the reference paper Rezaeiha *et al.*, (Vol 107, 2017), this is shown in Figure 6-2.




Figure 6-2 - Baseline VAWT CFD Results and Those Obtained by Rezaieha et al., (Vol 107, 2017)

The agreement between the CFD model in the present work and the reference paper is judged to be satisfactory. Note that the candidate VAWT CFD models produced as part of the Single-Blade method (see Section 5.4.1) are constructed in the same way as the validated baseline VAWT model described above.

6.2 Semi transient Adjoint Optimisation Philosophy

Section 6.1 described the VAWT CFD model used for the ISV optimisation process. The ISV method again uses instantaneous snapshots of the transient flow field to produce Adjoint solutions, in a similar fashion to the Single-Blade method (see Section 5.2). The use of a VAWT model however means that there are differences in how the method is implemented.

6.2.1 Objective Function

Section 3.2 described the role of the objective function in the Adjoint based optimisation process. An instantaneous property must be chosen in order to use a steady Adjoint solver, so the average C_M or C_P cannot be used. Furthermore, with a turbine of multiple blades, useful sensitivity data can only be taken from one blade at a time. For the ISV method, the instantaneous moment coefficient C_M of one blade is chosen as the objective function, and when evaluating the VAWT performance after optimisation, C_P calculations are made according to Equation 1.6.

It should be noted that the blades of a VAWT don't experience the exact same flow field as their sibling blades, although they are similar. This is due to hysteresis and the complex time history of the transient flow field. One could therefore optimise each blade separately using its own sensitivity data. In this work however, the sensitivity data obtained from the flow field around 1 blade is used to morph all the sibling blades identically.

6.2.2 Snapshots

Snapshots have been defined in Section 5.2.2, and this definition is maintained in the ISV method. With the ISV method however the snapshots are considered with reference to a C_M vs Azimuthal angle curve – an arbitrary snapshot is shown in Figure 6-3.




Figure 6-3 – VAWT Performance Curve with a Snapshot Location Marked

Section 6.4 presents the results of an investigation into the effect of snapshot location on the outcomes of the ISV optimisation process. A multi-snapshot study is presented in Chapter 7. These investigations are deferred until later in the thesis so that details of the methodology can first be discussed.

6.3 Optimisation algorithm

The previous sections (such as Section 3.5) have explained the building blocks of the semi-transient Adjoint based optimisation process for VAWTs. This section illustrates the ISV optimisation process architecture and describes the algorithm code. Also presented are details of the Adjoint module setup and the major implementation problems encountered during development.

This process can be applied to other CFD codes with an Adjoint solver. The level of detail provided here aligns with the ANSYS Fluent (18.2) Adjoint module; more steps may be involved for implementation in other CFD codes.

Figure 5-8 (in Section 5.3) shows a flow chart of a general Adjoint optimisation procedure that is not specific for VAWTs. Figure 6-4 shows a conceptual view of the ISV Adjoint based semi-transient optimisation process for VAWTs.

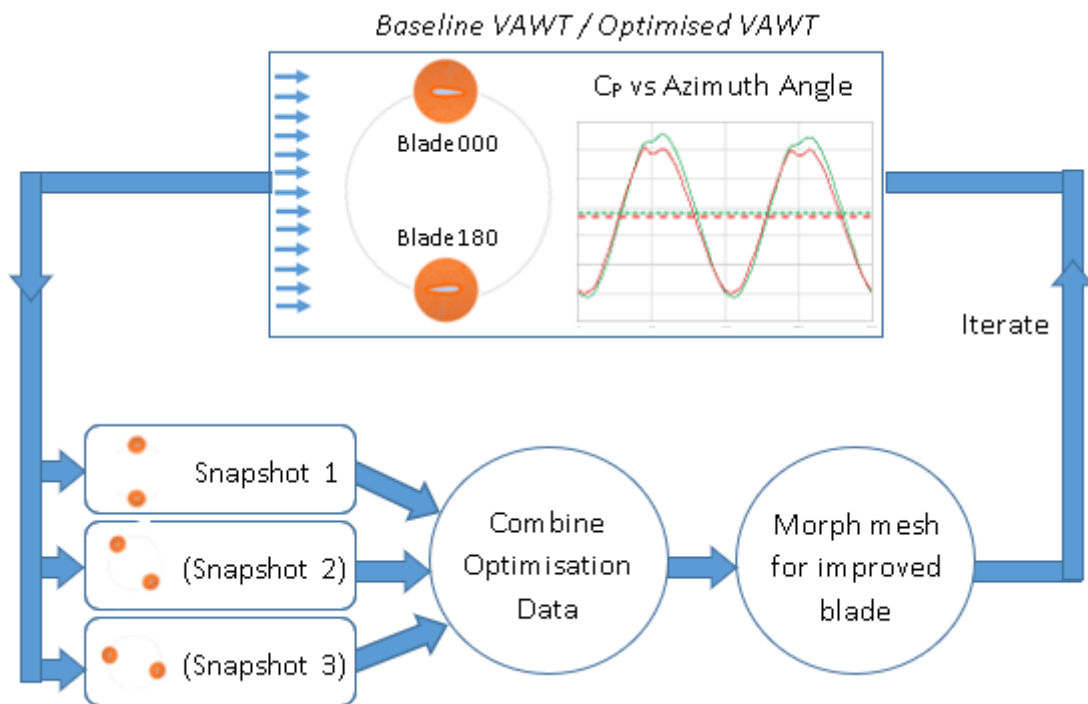


Figure 6-4 - General Schematic of the ISV Optimisation Process

The goal remains to produce a VAWT with improved average power coefficient, but for the ISV method no additional models are required (as for the Single-Blade method). The model itself is more complex, as both blades are contained, such that the algorithm is slightly more complicated. The top blade, at the 0 degree azimuthal position shall be referred to as Blade000, while the second blade beginning at 180 degrees is called the Blade180.

The transient CFD solution progresses until Blade000 reaches the snapshot location. The transient CFD simulation is then paused, while an Adjoint solution is

calculated which generates the sensitivity data. This sensitivity data describes how the C_M is affected by geometry changes of the blade. The ISV optimisation algorithm can handle several snapshots, but in this discussion a 1 snapshot case will be discussed for brevity. After the snapshot position is reached and the Adjoint solution is taken, the remainder of the revolution is completed so that Blade000 returns to 0 azimuthal angle.

At this stage, the sensitivity data can be imported and used to morph Blade000. Then the CFD solution progresses until Blade180 reaches the 0 azimuth position. The mesh morphing must take place here so that the coordinate system of the Blade180 is consistent with the sensitivity data (global X-Y). Blade180 is morphed according to the same sensitivity data generated for Blade000. The solution progresses until Blade000 once again returns to the 0 azimuthal position.

10 revolutions (see Section 6.1) are then computed to allow a periodically converged CFD solution to be reached so that the performance can be evaluated with the modified geometry. Several iterations of this Adjoint optimisation process are applied such that a new blade geometry is produced.

An optimisation process typically runs until convergence of the objective function is reached. In the present work, concern is for the average C_M (C_P) improvement of the VAWT it is in fact the local (at snapshot position) instantaneous C_M that is being optimised. The process does not need to be run until the instantaneous C_M -at-snapshot is converged. This is because the average VAWT C_M tends to be adversely affected by large geometry changes. For this work, the process is run for at least 10 iterations, which envelopes the point of maximum average C_M . While C_M -at-snapshot could increase with further iterations, it is not of use or value to pursue it. To illustrate this point Figure 6-5 shows the progression of the ISV optimisation process for two cases, and up to 10 process iterations. These show a continual increase to the C_M -at-snapshot, long after the maximum average VAWT C_M is reached. Perhaps 20 or 30 iterations would show convergence for the C_M -at-snapshot in some cases, but the associated blade geometry would be of little value in the present context, since the average VAWT C_M would become excessively poor.

6.3.1 Adjoint Module Setup

The Adjoint setup for the ISV method is similar to that of the Single-Blade method (see Section 5.3.1) but some implementation differences are necessary. Specifically, the objective function used and the Adjoint solution iteration limit are modified. Table 6-3 shows a summary of the Adjoint settings for the ISV method, the details of which are discussed in the following subsections.

Table 6-3 - Adjoint Module Settings Summary (ISV Method)

Objective function:	Blade moment coefficient (C_M)
Target performance change:	+0.0033 (to the objective function)
Adjoint solution iteration limit	3000 iterations
Adjoint solution stability scheme	Automatic
Geometric constraint	Constant chord length
Size of mesh morphing zone as a multiple of chord length	1.8c (x), 1.1c (y)
Number of control points in mesh morphing zone	100 (x), 100 (y)
Freeform Scaling Scheme	Objective reference change
Freeform Scale Factor	1

6.3.1.1 Adjoint Solver Settings

The VAWT model includes the complications of the real flow field such as curvature effects, blade/shaft interference and wake interactions, thus making the Adjoint solution less stable. An investigation was made that arrived at 3000 being a suitable number that balances solution progress with CPU cost. Stabilisation scheme options are offered for the Adjoint solver when the standard advancement scheme is unstable, and these options are the Modal scheme, Spatial scheme, and Dissipation scheme. The current work uses the 'auto-assign' option which chooses the appropriate scheme automatically if numerical divergence is detected during the calculation of the Adjoint solution. This led to automatic implementation of the Dissipation Scheme when computing the Adjoint solution.

The Adjoint solution convergence criteria for the Single-Blade model were $1e-4$ for Adjoint continuity and velocity, and $1e-3$ for Adjoint local flow rate. These are relaxed by 1 order or magnitude for the VAWT model optimisation (Section 5.3.1.1). This was required because in preliminary studies the default convergence criteria were never met regardless of the iteration limit in the Adjoint solution. Choosing to relax the convergence criteria was necessary but does mean a loss of final 'accuracy' in the Adjoint sensitivity data. This however is a small effect and does not manifest any issues in the optimisation process, as can be seen by the results in Section 6.4.

The Adjoint solutions for upwind snapshots have generally converged or plateaued around $1e-3$. The Adjoint solutions for downwind snapshots have also generally

converged or plateaued with residuals around $1e-3$, which is considered acceptable. Figure 6-8 (Section 6.4) shows examples of Adjoint solution convergence.

The Adjoint solution method was again chosen as the standard (for pressure) and first order (for momentum) to improve solution stability. The Adjoint solution produces the sensitivity data as discussed in Section 3.1, which is itself not the solution to the aerodynamic optimisation problem. The sensitivity data must be applied to the aerofoil geometry via mesh morphing (see Section 5.3.1.2). Figure 5-16 (Section 6.4) shows examples of Adjoint sensitivity data.

6.3.1.2 Mesh Morphing Settings & Constraints

The mesh morphing settings and constraints are similar for the ISV method as were discussed in Section 5.3.1.2 for the Single-Blade method.

The location of the constraint area and the morphing region must be around the starting location of Blade000, at the 0 azimuthal position. Note also that the constraint does not need to follow the blades while they rotate and can remain in the same location. This is because the both blades are both morphed in this 0 azimuth location regardless of the snapshot locations (see Section 6.3.4.3).

6.3.2 Aggression

As described in Section 5.3.2 the term ‘aggression’ used in this thesis refers to how the objective of the optimisation process is specified. During development of the ISV method, the chosen method of specifying aggression changed from that used in the Single-Blade optimisation chapters (see Section 5.3.2). This change arose as a result of additional investigations which took place following the Single-Blade method development, which aimed to improve the implementation of aggression. Instead of specifying a target increase *as a percentage* of the current value of C_M (or C_T , etc.), the ISV method specifies a target *as an absolute value*. This means that the aggression is more consistent because the same amount of change in the objective will be targeted regardless of the snapshot position.

An investigation was made with a preliminary level of aggression and this target was $+0.0033 C_M$. This value was chosen because it is 10% of the average C_M of the baseline VAWT, offering an appropriate order of magnitude target for the optimisation. The optimisation process then attempts to improve the instantaneous C_M by 0.0033 at the snapshot location, for every iteration.

Figure 6-5 shows two cases, where a single snapshot was used in the optimisation. It can be seen that the peak average VAWT C_M is reached early in the process (iteration 1 for Case 345, iteration 4 for Case 60) before it deteriorates; meanwhile

the C_M -at-snapshot steadily increases. This can be seen for many cases in Table 6-5 where the best iteration for average C_M is early on.





Figure 6-5 – Progression of Optimisation Processes (Left) Case 60, (Right) Case 345

This shows that the process has a reasonably appropriate level of aggression, but that a smaller value could improve the ‘resolution’ around the optimum. In addition, another series of tests were made using a lower value of $+0.00165 C_M$ as the objective, corresponding to 5% of the average C_M of the baseline VAWT. These results are not presented but showed that the best average C_M was typically reached around the mid-point of the process (often around iteration 4 or 5), and a marginally higher value was achieved compared to when using $+0.0033 C_M$.

Other preliminary studies showed that high aggression levels can lead to an ‘overshoot’ effect. This term is coined to describe undesirable behaviour in an optimisation process where excessive geometry changes are made at each iteration. The local optimum is overshoot each time, rather than steadily approaching it – this should be avoided. Figure 6-6 shows an example case with ‘overshoot’ caused by excessive aggression.

If aggression is too low then the maximum average VAWT C_M may not be reached even after large numbers of process iterations, thus carrying a high computational cost. Correct aggression is therefore a balance of computing cost and avoiding ‘overshoot’.



*Given as a % of the average C_M for 1 blade

Figure 6-6 - Example of the 'Overshoot' Effect in an Overly Aggressive Optimisation Process

All cases tested (except Case 15 and Case 30) exhibited smooth steady increases to the 1-blade- C_M -at-snapshot, thus demonstrating that aggression of $+0.0033 C_M$ is not too high. Case 15 and Case 30 results did exhibit some behaviour of excessive aggression, where the resulting blade shapes (see Figure 6-11 (b) and (c)) are distorted compared to the rest of the cases. $+0.0033 C_M$ is used as the aggression/objective target for the cases tested here, as stated in Table 6-3, and this is because it is widely suitable offering a good balance of stability and CPU cost.

6.3.3 Algorithm Code

The previous sections have described the setting up of the VAWT model (Section 6.1), and the setting up of the Adjoint module (Sections 6.3). This section outlines the algorithmic process used that envelopes the operation of the VAWT model and Adjoint module in order to form the ISV optimisation process. This algorithm takes the form of a journal file, which Fluent reads line by line as text input commands; this is written in ANSYS Fluents' internal programming language called Scheme (Dybvig, 1996).

In a similar fashion to the Single-Blade method, a spreadsheet tool was developed to quickly produce algorithm scripts based on the number of snapshots (user defined) and their positions.

Figure 6-7 shows a flow diagram of the ISV optimisation algorithm; it is a less abstract version of Figure 5-8/Figure 6-4. The following discussion refers to stages in Figure 6-7 by the number identifier, **X**. The light and dark grey shaded boxes correspond to the main loops in the algorithm.

Precursor items required before running the optimisation algorithm/script include the following:

- Stage **1**. The VAWT CFD model should be produced according to Section 6.1.
- Stage **2**. After the baseline VAWT model with NACA0018 blades has completed 10 cycles, the C_M as a function of azimuthal angle should be plotted for the baseline blade geometry. The 10th revolution is sufficient (in high TSR cases) to provide a periodically converged solution as has been shown in Section 6.1. The flow field data in this revolution forms the baseline blade performance. After each iteration the performance is compared to this. We can use this curve to make choices about the number of snapshots, position of the snapshots, and relative weighting of snapshots (see Figure 6-3). Discussion of how to choose these details can be found in Section 7.2. Once chosen, the spreadsheet tool produces the specific script file accordingly.
- Stage **3**. The Adjoint module should be configured according to Section 6.3.1

Script/Algorithm:

- Stage **4**. The transient solution is continued by N iterations, until the next snapshot position is reached by Blade000. N depends on the desired position (see Figure 6-3) of the first snapshot (a spreadsheet is used to convert azimuthal angle into a value for N).
- Stage **5** & Stage **6**. When the snapshot position is reached, the transient solution is stopped temporarily. Then the steady flow solver is selected. This is because the ANSYS Fluent Adjoint solver is configured for use with the steady flow solver (some Adjoint module functions do not work while the transient solver is engaged). This approach does not invalidate the flow field data or preclude its use in an Adjoint calculation (see Section 5.2). Having the steady solver active at this point in time allows the use of the Adjoint module, but switching between steady and transient flow solvers requires additional considerations which are discussed in Section 6.3.4.2.
- Stage **7**. Before the Adjoint solution is calculated, the mesh and flow field within the rotating subdomain is rotated. It is rotated from the azimuthal position θ of the current snapshot to $\theta = \text{zero}$ (but will be rotated back later at Stage **10**). This has the effect of aligning the blade with the global X-Y coordinate system. This is

a necessary step because the sensitivity data produced from an Adjoint solution is prescribed in the global coordinate system. To combine sensitivity data in a correct way it must all be orientated in a consistent fashion (see Section 5.3.4.3). This method maintains consistency/alignment in orientation of sensitivity data for all snapshot positions, so that combination of sensitivity data is correct.

- Stage **8** and Stage **9**. The Adjoint solution is calculated with the setup as described in Section 6.3.1. When this calculation is complete, the resulting sensitivity data is exported for later use.
- Stage **10**. The counterpart step to Stage **7**. Now that the snapshot is complete, and the sensitivity data is obtained the CFD model can resume. Before the transient solution can continue marching towards the next snapshot position, the mesh and flow field are rotated back (from $\theta=0$) to the correct θ (see Section 5.3.4.3).
- Stage **11**. The transient solver must be re-engaged before resuming the transient simulation (see Section 6.3.4.2).
- Stages **4** to **11** are repeated for however many snapshots are specified.
- Stage **12**. Once all snapshots have been taken, the transient solution continues until the blade is back at the θ =zero position. At this position, the blade is aligned with the global X and Y coordinate system. All the sensitivity data calculated previously is aligned with the blade at this position, ready for mesh morphing. Note the '#' at the end of this item. This corresponds to a 'break' in the algorithm where the current Fluent case & data files are saved and are then closed/reopened. This was found to be a necessity due to memory issues as described in Section 6.3.4.1. Breaks appear elsewhere in the algorithm and these are also denoted by a'#'.
- Stage **13**. The steady solver is engaged, as before, to allow full functionality of the Adjoint module. Sensitivity data from all snapshots is imported. Weightings are applied to each snapshot if desired (see Section 5.3.4.4).
- Stage **14**. The mesh morphing settings are implemented according to Section 6.3.1.2. The mesh is morphed according to the combined sensitivity data and aggression, providing an updated blade geometry. The sensitivity data is then cleared from the Adjoint module.

- Stage **15**. The transient solver is then re-engaged as before, and the solution is progressed until Blade180 reaches the $\theta=0$ position for morphing.
- Stage **16**. The steps **13** and **14** are repeated to apply the same sensitivity data in the morphing of Blade180. When morphing, the algorithm must ensure to select the correct blade wall (Blade000 or Blade180) to morph depending on which one is currently at the 0 position.
- Stage **17**. The transient solver is re-engaged, and the solution is progressed until Blade000 returns to the $\theta=0$ position. Following this, 10 cycles of the turbine are run; the flow field data from the final revolution is used to evaluate the performance of the new blade geometry.
- Stages **4** to **17** are repeated for X iterations where X is the number of Adjoint iterations specified by the user.
- Stage **18**. The performance data for the optimisation so far is plotted showing how the C_M as a function of θ curve has improved with each iteration.

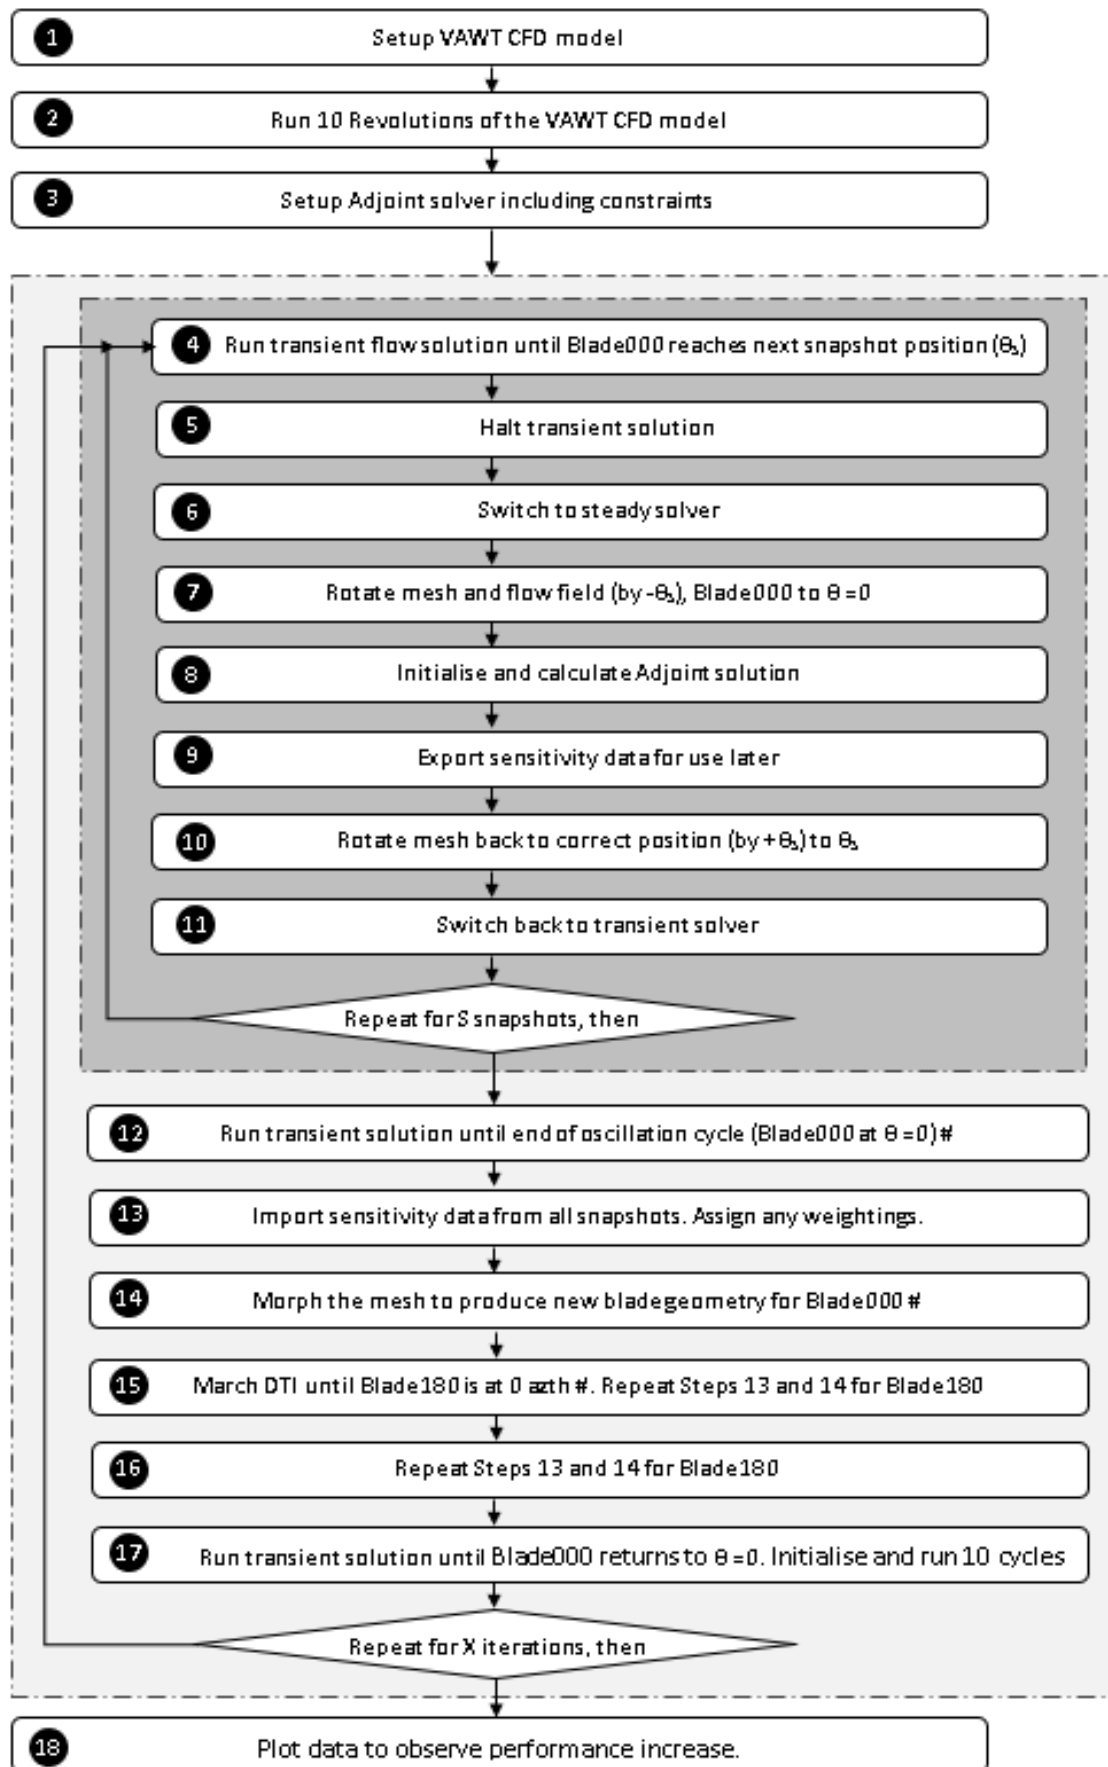


Figure 6-7 - ISV Optimisation Process Flow Diagram

6.3.4 Algorithm Key Notes

This section expands upon some of the key elements in the algorithm which ensure robustness and correct functionality of the optimisation process. Also discussed here are some of the ways that the algorithms operation was validated. A similar discussion is given in Section 5.3.4 which includes specific considerations for the Single-Blade method algorithm.

6.3.4.1 Memory & Code 'Breaks'

To ensure robust functioning of the optimisation process, code 'breaks' were introduced which effectively close and reopen the working CFD file during the process. By doing so, temporary memory caches are cleared which can become full owing to the memory demands of the process, which could cause computing faults if not mitigated. Note that error codes and log files generated may not provide specific information of the root cause of these problems. After each code-break, some of the algorithm variables and CFD settings must be redefined in the script. This approach was found to be necessary while running the process on parallel cores on an HPC. Breaks were not required for the Single-Blade method, possibly due to the smaller model size. Sufficient memory should also be requested when running on an HPC.

6.3.4.2 Steady/Transient Switching and Frame/Mesh Motion

This is similar to the consideration described in Section 5.3.4.2 for the Single-Blade method, although additional zones exist in the ISV model. This means that extra lines of code are required for correcting the frame motion/mesh motion settings, whenever the solver is switched from steady/transient.

6.3.4.3 Adjoint Solution Correction for Global XY Alignment

The discussion in Section 5.3.4.3 applies also to the ISV method. However, the implementation of the corrections for alignment of the Adjoint solution differs slightly. The Single-Blade method uses AoA corrections while the ISV method uses azimuthal angle corrections.

For the ISV method, when an Adjoint solution is taken the transient CFD solution is paused. At this point the mesh and flow field is temporarily rotated such that Blade000 is located at the zero azimuthal angle position. This places Blade000 in alignment with the global X-Y coordinate system when the Adjoint solution is taken. Afterwards the model is rotated back to the correct azimuthal angle so that the transient solution can be resumed.

An important note here is that the angles of rotation to move Blade000 to the zero-azimuth position must be known by the optimisation script. These are calculated via

a spreadsheet when the script text file is generated, so each rotation angle is known for each of the specified snapshots.

6.4 Single-snapshot investigation

This section uses the In-situ VAWT (ISV) optimisation method (setup according to Sections 6.1, 6.2 and 6.3) with 1 Adjoint snapshot per revolution to investigate the effects of snapshot location on the outcome of the optimisation process. A multiple-snapshot investigation is presented in Chapter 7. This 1 snapshot investigation allows a clear illustration of the cause and effect of snapshot choice, and performance changes. A range of cases are investigated where the location of the snapshot is varied; these are listed in Table 6-4.

Table 6-4 – List of Test Cases for ISV Optimisation

Case Name /Azimuthal Angle (degrees)	Azimuthal Angle of VAWT at Snapshot (degrees)
0	0
15	15
30	30
45	45
60	60
75	75
90	90
105	105
120	120
135	135
150	150
165	165
180	180
195	195
210	210
225	225
240	240
255	255
270	270
285	285
300	300
315	315
330	330
345	345

Across these cases, the snapshot location differs by $1/24^{\text{th}}$ of an oscillation cycle which corresponds to a gap of 15 degrees azimuthal angle. Although testing 360 cases (1 for each degree) or even more could be done, it is judged that 24 cases offer a reasonable balance between time cost, and a sufficient resolution in the studying of the solution space.

Section 6.3.1.1 mentions the Adjoint module setup, and how convergence of the Adjoint solution is generally acceptable. Figure 6-8 shows the Adjoint solution convergence plots for a typical upwind and a typical downwind case. As will be seen in due course, there is a valid basis for grouping snapshots together to some extent, as the upwind and the downwind cases each share typical behaviours respectively.




Figure 6-8 - Optimisation Adjoint Solution Convergence (Left) Case 90 (Right) Case 270

To correspond with these convergence plots Figure 6-9 displays the fields of sensitivity data. Vectors of shape sensitivity magnitude are shown, and this indicates the direction in which small geometrical changes provide improvements to C_M . It should be noted that there are some large sensitivity vectors that are not shown which appear at only a few nodes around the geometry. These correspond to inflections in pressure at the leading edge, and the sharp geometry of the trailing edge. The mesh morpher applies smoothing to the geometry deformation so that any large variations in sensitivity vectors (which may arise naturally) do not produce discontinuities in the geometry when morphing takes place.




Figure 6-9 - Sensitivity Vectors (Left) Case 90 (Right) Case 270

By comparing Figure 6-9 with Figure 5-16, a high level of consistency is observed for the sensitivity vectors of Case 90. In this upwind region, the Single-Blade model gives a good prediction of the VAWT flow field, and therefore similar sensitivity vectors. Case 270 shows a level of similarity but also some disagreement, which is due to the challenges of representing the downwind flow field with a Single-Blade model (as discussed in Section 5.1).

6.4.1 Results – Single Snapshot Optimisation

After the iterations of the optimisation process are completed, the performance data can be processed. Unlike the Single-Blade method, the ISV method has the advantage that the performance can be observed directly from the optimisation model without the necessity of creating a separate candidate VAWT model.

Section 6.3 explained that the ISV optimisation process optimises the blade geometry for C_M , according to the instantaneous flow field around one of the blades. As such, the process is not capable of considering the summation of every blade's flow field concurrently. In other words, the optimiser can only 'see' the C_M of 1 blade at a time when the Adjoint solution is taken. Therefore, the appropriate way to measure if the process is working effectively is the 1-bladed C_M at the snapshot location, rather than the turbine C_M . Of course, the aim is to improve the combined average C_M of all turbine blades acting together (i.e., the average C_P), but these key quantities are in essence a by-product of the optimisation process which improves the 1-blade C_M in the region of the snapshot(s).

With this in mind, we consider the results both in terms of changes to; a) 1-blade C_M -at-snapshot position, and b) average VAWT C_M (or C_P). Table 6-5 shows the results of the cases studied using 1-snapshot ISV optimisation. Note that the final 10th iteration provides the best 1-blade- C_M -at-snapshot for all cases, but not necessarily the best average VAWT C_M . For this reason, another column is given that shows the best average VAWT C_M that was reached during the process. Values of the improvement to 1-blade- C_M -at-snapshot are given as percentages of the average C_M ($=0.033 C_M$) of one blade over one revolution of the baseline VAWT. This provides a

consistent basis for comparison since the baseline *instantaneous* C_M varies for each azimuthal position. Average C_M values are simply given as a percentage of the baseline VAWT average C_M .


Figure 6-10 (a) to (x) shows the resulting performance curves of the 24 cases of 1-snapshot optimisation studied. Curves are shown for the 1-bladed C_M (left), along with the VAWT C_P curve (all blades) (right). The corresponding blade geometries are shown in Figure 6-11 (a) to (x). The data shown corresponds to the 10th process iteration. As shown in Figure 6-5 this means that the iteration for the best VAWT average C_P has been surpassed, but the 1-blade- C_M -at-snapshot continues to increase. Using more iterations to pursue the true maximum value of 1-blade- C_M -at-snapshot is not deemed useful as the highly deformed blade would carry significant performance deterioration away from the snapshot position and would not be viable for maintaining the average VAWT C_M . The 10th iteration results therefore achieve an appropriate level of maturity in the optimisation process, presenting a shape that will improve instantaneous blade torque, for each case (azimuthal angle). The VAWT average C_P results are presented in Table 6-5.

Table 6-5 - ISV 1 snap results After 10 Process Iterations


Case Name /Azimuthal Angle (deg)	Percentage improvements made by optimisation process		
	*1-blade C_M at snapshot (10 th Iteration)	Average VAWT C_M (10 th Iteration)	Average VAWT C_M (Best Iteration)
0	+94.2 %	-5.3 %	+0.9 % (1 st Iteration)
15	+39.0 %	-26.6 %	+3.6 % (6 th iteration)
30	+33.5 %	-4.9 %	+1.2 % (1 st iteration)
45	+84.3	-4.5	+1.4 % (1 st iteration)
60	+84.6 %	+0.3 %	+1.9 % (4 th iteration)
75	+71.9	+2.2 %	+2.5 % (7 th iteration)
90	+46.1 %	+2.8 %	+2.8 % (10 th iteration)
105	+24.3	+1.9	+1.9 % (10 th iteration)
120	+11.7 %	+1.3 %	+1.3 % (10 th iteration)
135	+13.0	+1.7 %	+1.7 % (10 th iteration)
150	+42.5 %	+1.4 %	+2.9 % (6 th iteration)
165	+43.3	-13.6	+2.5% (1 st iteration)
180	+35.4 %	-63.0 %	-0.8 % (1 st iteration)
195	+97.0 %	-14.4 %	+1.1 % (1 st iteration)
210	+68.0 %	-2.1 %	+1.2 % (2 nd iteration)
225	+80.2 %	-1.7 %	+1.2 % (2 nd iteration)
240	+83.5 %	-1.3 %	+1.2 % (2 nd iteration)
255	+40.2 %	+0.7 %	+1.2 % (4 th iteration)
270	+85.1 %	-2.8 %	+1.1 % (2 nd iteration)
285	+95.6 %	-1.3 %	+1.1 % (2 nd iteration)
300	+96.2 %	-0.7 %	+1.1 % (2 nd iteration)
315	+96.7 %	-0.6 %	+1.1 % (2 nd iteration)
330	+95.4 %	-0.6 %	+1.1 % (2 nd iteration)
345	+93.2 %	-1.0 %	+1.0 % (2 nd iteration)
*given as a percentage of $C_M=0.033$ (the average C_M of one blade of the baseline VAWT)			

Throughout Figure 6-10 (a) to (x) the red and green curves represent baseline and optimised blade performance respectively. The value of θ given is the azimuthal position of the snapshot used for that particular case. The change in 1-blade C_M -at-snapshot position has been marked in yellow so it can be readily observed.


(a, $\theta = 0^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(b, $\theta = 15^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(c, $\theta = 30^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(d, $\theta = 45^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(e, $\theta = 60^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(f, $\theta = 75^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(g, $\theta = 90^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(h, $\theta = 105^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(i, $\theta = 120^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(j, $\theta = 135^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(k, $\theta = 150^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(l, $\theta = 165^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(m, $\theta = 180^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(n, $\theta = 195^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(o, $\theta = 210^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(p, $\theta = 225^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(q, $\theta = 240^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(r, $\theta = 255^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(s, $\theta = 270^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(t, $\theta = 285^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




(u, $\theta = 300^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle



(v, $\theta = 315^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle



(w, $\theta = 330^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle



(x, $\theta = 345^\circ$), (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle


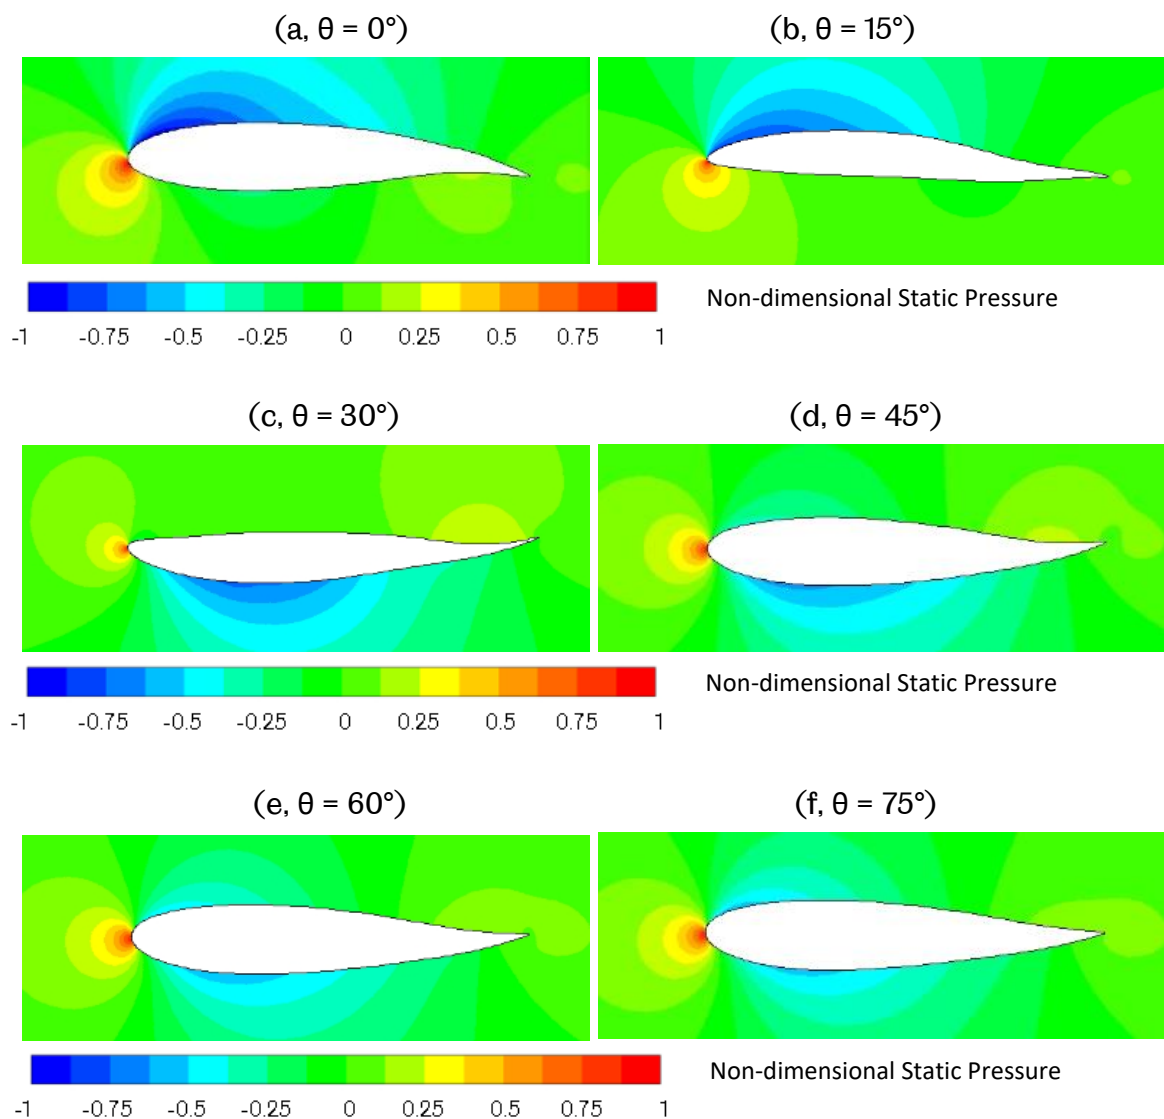







Figure 6-10 (a) to (x) - 1 Snapshot ISV Optimisation, (Left) 1-Blade C_M vs Azimuthal Angle, (Right) VAWT C_M vs Azimuthal Angle

The blade geometries corresponding to Figure 6-10 are shown in Figure 6-11 (a) to (x).

Throughout Figure 6-11 (a) to (x) the colour contours are of non-dimensional static pressure. The value of θ given is the azimuthal position of the snapshot used for that particular case. The blade geometries are shown at the 0 degree azimuthal position for all cases. The top surface corresponds to the outside of the turbine, and the bottom surface corresponds to the inside of the turbine (towards the shaft). For the baseline blade these contour plots would show a zero fixing angle, but the optimised blades tend to have slight non-zero fixing angles.





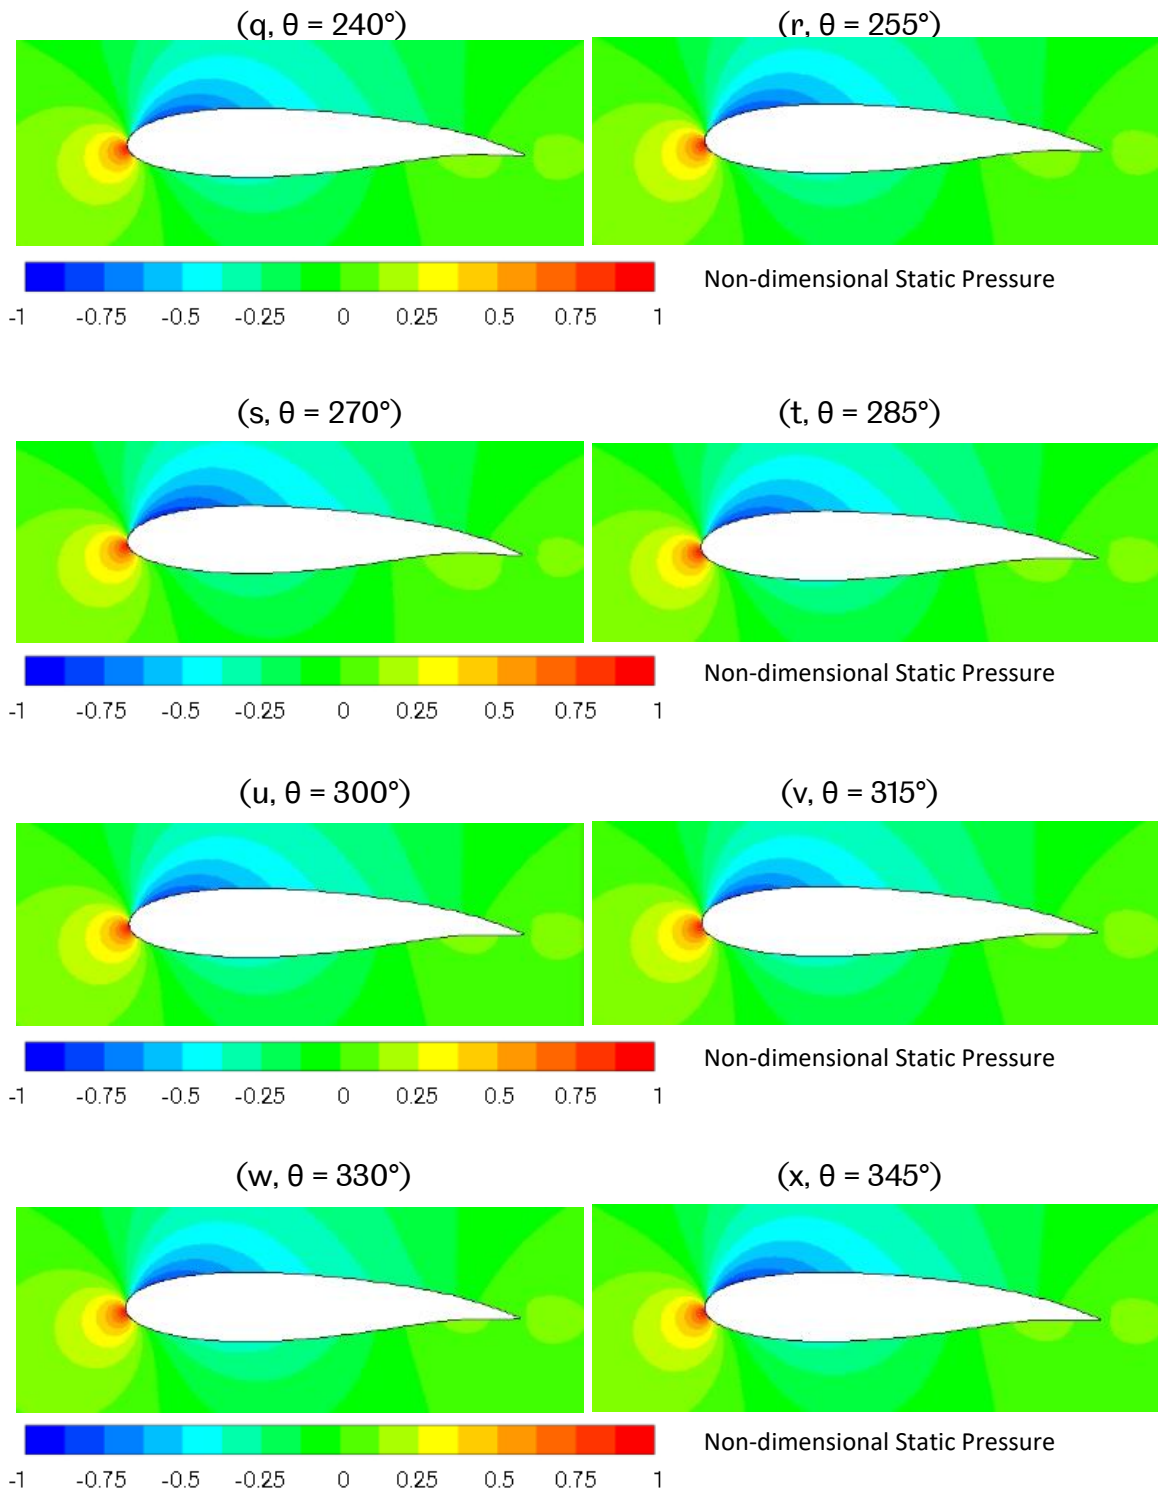


Figure 6-11 (a) to (x) - 1 Snapshot ISV Optimisation, Resulting Blade Geometries

Some major observations about the results are as follows:

- Most of the cases reached and surpassed the maximum VAWT average C_M within 10 process iterations, but 4 upwind cases needed additional iterations (suggesting higher aggression could have been useful). Case 15 and Case 30

results exhibited behaviour of excessive aggression, where the resulting blade shapes are distorted (see Figure 6-11 (b) and (c)) compared to the rest of the cases. Elsewhere the aggression levels appeared suitable. This indicates the aggression level used (see Section 6.3.2) was generally appropriate and a good compromise/balance on considering all cases.

For cases where maximum VAWT average C_M was reached early on (within 1 to 4 iterations), some were re-run at lower aggression levels to find out if the peak performance (for that case) was being missed due to over-aggression. These re-runs showed that only a very small increase to the peak performance would be seen by reducing aggression for those cases.

- In all the cases significant increases to 1-blade- C_M -at-snapshot (and the surrounding region) are observed. This is significant and for the majority of cases the instantaneous C_M increases by more than 70% or 80% of the baseline average C_M value for one blade. This demonstrates that the optimiser is correctly pursuing the goal that it is programmed to do in a highly effective manner. Regions away from the snapshot position tend not to be improved, and this is because no optimisation data is considered away from the position of the single snapshot.
- A certain increase to 1-blade- C_M -at-snapshot does not necessarily translate to an improved average VAWT C_M (or C_P) of the same magnitude, because the localised improvements (at the snapshot position) tend to carry a penalty on the opposing side of the turbines revolution
- Modest improvements to the average VAWT C_M are seen, but a positive improvement is in fact achieved for all cases except for Case 180. The best average VAWT C_M achieved is around 3.6%, and this is similar to that achieved in the Single-Blade method.
- A major conclusion about this could be that *single-snapshot* optimisation may not be likely to yield large improvements to the average VAWT C_M , and a multi-snapshot approach may be necessary. This of course heavily depends on how sub-optimal the starting design is; perhaps for the VAWT in the present example the NACA0018 is already relatively close to the optimum (at the TSR studied).

- A significant benefit of reducing the cyclic load fluctuation is seen for most of the downwind snapshot cases, whilst the average C_M is also improved slightly. This offers great value for designing to reduce fatigue loads and can be achieved very quickly and without the need for any control systems (active or passive).
- Some peculiar cases occur such as Case 0 and Case 180. These are when the AoA is rapidly changing from positive/negative, and where the instantaneous C_M is very low. It appears that in these locations the sensitivity data is not valuable, and that an increase to the C_M -at-snapshot is largely detrimental to the VAWT average C_M . This is because the flow regimes at these azimuthal positions are peculiar, so optimising for them is not beneficial for the majority of the revolution.
- Another interesting result is for Case 0, which has a blade geometry resembling the downwind snapshot cases. This is likely to be due to hysteresis, where the flow field experienced at 0 degrees azimuthal is largely dominated by typical downwind effects.
- Upwind snapshot cases generally improve the C_M over the upwind but deteriorate in the downwind. The resulting geometries for these cases tend to exhibit a negative camber.
- For upwind cases (Cases 45, 60, 75, 90), before the torque peak (at ~ 100 degrees), significant increases to the peak C_M magnitude are seen, and also for the negative peak magnitude in the downwind that follows. For upwind cases (Cases 105, 120, 135) after the torque peak (at ~ 100 degrees), only a small change is seen to these peak/trough magnitudes.
- There is a very large consistency across the downwind snapshot cases exhibiting almost the same average C_M improvement and similar C_M -at-snapshot improvements. The resulting geometries for these cases tend to exhibit a positive camber.
- Downwind cases (Cases 195-345) all improve the C_M in the downwind, while deterioration is seen in the upwind. This reduction of peak magnitudes is responsible for the reduction in load fluctuations (potential fatigue damage) already mentioned.

- For all snapshots (upwind and downwind), the camber of the resulting blade geometries appears to be focused toward the trailing edge in the aft quarter of the blades chord length.
- The 1-snapshot results suggest that morphing blade technologies could offer an appropriate means of pursuing the optimum blade shapes continuously through the cycle. If the blade shape can be changed in accordance with the resulting geometries shown above, then the large improvements to C_M -at-snapshot could potentially be realised without suffering from the upwind/downwind performance paradox (see Section 5.4.3) which is experienced for fixed blades.

The results of the 1-snapshot investigation using the ISV method show that the semi-transient Adjoint based optimisation process can be successful in improving the average power coefficient of a VAWT. The results provide insight into how the snapshot position effects the VAWT blade performance, which is of value in informing the multi-snapshot investigation of Chapter 7. Discussion of the blade geometries and aerodynamics behind these performance effects are now provided in Section 6.4.2.

Some key limitations exist which should be noted regarding the method presented here and these are discussed in Section 8.1.

6.4.2 Aerodynamic Analysis for ISV, 1-snapshot Optimisation

As has been seen in Section 6.4.1, the results from cases with upwind snapshot locations and downwind snapshot locations form two distinct groups. The following discussion considers the aerodynamics in detail, of a representative case from each of these groups.

Figure 6-10 (g) (shown previously) is the graph of the C_M as a function of the azimuthal angle for Case 90 which is the representative upwind case. The contribution from only one of the two blades is shown on the left and the combined effect of both blades is shown on the right. After 10 iterations a 46% increase was made to the 1-blade- C_M -at-snapshot (as a % of 0.033, the average C_M of one baseline blade), and a 2.8% increase to the average VAWT C_M (C_P) was made (see Table 6-5). The blade geometry that gave this improvement has been presented in Figure 6-11 (g). The blade has a toe-out fixing angle of 0.6° , a maximum negative camber of (-) 0.8% chord located at 76% chord (towards the trailing edge), and a thickness reduction of 0.5% compared to the baseline NACA0018 blade. If the optimisation process is continued further after the best average C_M was passed, the 1-blade- C_M -at-snapshot kept increasing to 87% after 15 iterations.

To analyse the aerodynamics implications of this blade geometry, Figure 6-12 shows the streamlines (which are coloured by the non-dimensional static pressure) for the resulting optimised blade (Case 90), as well as the baseline NACA0018 blade when at the snapshot position of 90° azimuthal angle. The corresponding surface pressure coefficient plots are presented in Figure 6-13. The surface pressures plots for the Case 90 blade show a weaker negative pressure on the suction side, at the leading edge. The increased fixing angle means that less curvature is required to pass around the leading edge. This decreases the leading edge suction but allows a greater suction to be maintained along the mid-span and further aft approaching the trailing edge. Towards the aft region, the positive pressure on the top surface is greater than the baseline blade because of the camber. A high pressure zone can be observed as the flow is slowed down by the cambered tail on the top surface of the Case 90 blade. This is coupled with a slightly greater suction magnitude that persists towards the trailing edge, such that a more favourable magnitude and direction of the pressure gradient is achieved. In addition, the trailing edge geometry slightly changes the size and shape of a small recirculating region but the effect on the surface pressures is minor.

Figure 6-14 shows the streamlines (which are coloured according to the non-dimensional static pressure), and Figure 6-15 shows the surface pressures for the resulting optimised blade for Case 90, at the position of 270° azimuthal angle. This is the opposite location in the cycle to where the snapshot was taken. Note: this should not be confused with information relating to the Case 270 resulting blade – the analysis at positions of 90° and 270° azimuthal angle will be repeated in due course for Case 270 which is the representative downwind case.

At the position of 270 azimuthal angle, the optimised blade (Case 90) negative camber is not beneficial because in this position the lower surface is the pressure surface. The Case 90 blade has a weaker negative pressure/suction peak at the leading edge than the baseline blade, and this remains weaker towards the mid-span which is due to the effect the fixing angle has on the AoA at this position. The larger suction and pressure region located at 0.8 chord (Figure 6-15) corresponds to the position of maximum camber. The net effect is that the Case 90 blade produces slightly less torque at 270 degrees azimuthal angle, and indeed across the downwind region.




Figure 6-12 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Resulting blade from Case 90. (Bottom) Baseline blade




Figure 6-13 - Case 90 VAWT blade surface pressure coefficient at 90° azimuthal angle




Figure 6-14 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Resulting blade from Case 90. (Bottom) Baseline blade




Figure 6-15 - VAWT blade surface pressure coefficient at 270° azimuthal angle

The above discussion regarded the aerodynamics of the resulting blade from Case 90, at two azimuthal angles (90 and 270). Case 90 was the representative upwind case. Similarly, Case 270, the representative downwind case is now discussed.

Figure 6-10 (s) (shown previously) is the graph of the C_M as a function of the azimuthal angle for Case 270 which is the representative downwind case. The contribution from only one of the two blades is shown on the left and the combined effect of both blades is shown on the right. The blade geometry after 10 iterations has been presented in Figure 6-11 (s). The blade has a toe-out fixing angle of 2.2° , a maximum camber of 3.4% chord located at 0.75% chord (towards the trailing edge), and a thickness reduction of 3.8% compared to the baseline NACA0018 blade. As per Table 6-5, Case 270 provides a significant increase to the 1-blade- C_M -at-snapshot after 10 process iterations - an 85% increase (as a % of 0.033, the average C_M of one baseline blade) was made. If the optimisation process is continued further after the best average C_M is passed, the 1-blade- C_M -at snapshot kept rising to 130% by 15 iterations. At 10 iterations the average C_M has deteriorated past its maximum for this case which was just 1.1% at iteration 2.

To analyse the aerodynamics implications of this blade geometry, Figure 6-16 shows the streamlines (which are coloured according to the non-dimensional static pressure) for the resulting optimised blade (Case 270), as well as the baseline NACA0018 blade when at the snapshot position of 270° azimuthal angle. The corresponding surface pressure coefficients plots are presented in Figure 6-17. The surface pressures plots for the Case 270 blade show much greater magnitudes of both the positive pressure and negative pressure along the respective surfaces, in comparison to the baseline blade. Even at the leading edge, where the difference in geometry is minor there is a large difference in the pressure coefficient - the substantial mid/trailing edge camber has a significant impact on the flow elsewhere. The increased AoA at this position also means that greater curvature is required to pass around the leading edge which contributes to the greater suction. The camber then allows suction to persist further down the chord. Towards the aft region, the positive pressure bump on the top surface is greater than the baseline blade due of the camber concentrated around 0.8 chord length. Note that at this position (270° azimuthal angle), the lower surface is the pressure surface such that the camber can provide this effect. The trailing edge geometry enlarges the small recirculating region which slightly effects the surface pressures, but this effect is small compared to overall aerodynamic benefits of the geometry.




Figure 6-16 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Resulting blade from Case 270. (Bottom) Baseline blade




Figure 6-17 - Case 270 VAWT blade surface pressure coefficient at 270° azimuthal angle




Figure 6-18 - VAWT blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Resulting blade from Case 270. (Bottom) Baseline blade




Figure 6-19 - Case 270 VAWT blade surface pressure coefficient at 90° azimuthal angle

Figure 6-18 shows the streamlines (which are coloured according to the non-dimensional static pressure), and Figure 6-19 shows the surface pressures for the resulting optimised blade for Case 270, at the position of 90° azimuthal angle. This is the opposite location in the cycle to where the snapshot was taken.

At the position of 90 azimuthal angle, the optimised blade (Case 270) positive camber is not beneficial. The Case 270 blade has a weaker negative pressure/suction peak and a weaker positive pressure at the leading edge than the baseline blade. This reduction in pressure gradient persists well past the mid chord until the cambered trailing edge is approached. A deterioration in torque is found at 90 degrees azimuthal angle and across the upwind region in general.

The conclusions that can be drawn from the discussion are as follows:

- The upwind and downwind representative blades both have a slight toe out fixing angle (see Figure 1-8), despite having opposing cambers (negative camber and positive camber respectively, see Figure 1-7). This reduces the upwind AoA which means that the blade does not see such a large maximum AoA and thus decreasing the tendency of flow separation. This toe-out fixing angle increases the downwind AoA which increases the lift/tangential force in the downwind region. Due to the lower nominal AoA range in the downwind region this does not risk the blade reaching its stall angle there. This conclusion is consistent with the study of Gosselin et al. (2013).
- Downwind cases cause significantly reduced fluctuations in the aerodynamic loads. This is due to the blades positive camber, which reduces the upwind peak whilst the average C_M is maintained via downwind improvements. These cases therefore offer a route to optimising a VAWT for reduced fatigue damage.
- Relatively small blade geometry changes have a large effect on the instantaneous performance. Case 90 for instance has a very small fixing angle change (-0.6°), camber (-0.8% chord), and thickness change (-0.5%), yet this can significantly increase the C_M at the snapshot position.

6.5 Concluding Remarks on 1-Snapshot Optimisation, ISV Method

The results show that the semi-transient Adjoint based optimisation method can be applied directly to a VAWT model with success at improving instantaneous C_M , and average C_M (C_P), and also reducing fluctuating loads. The algorithm/script required is more complex than for the Single-Blade method, but an intermediate Single-Blade model is not required which aids simplicity in another way. Optimisation run times are significantly longer for the ISV method due to model size, and because more cycles are required for a periodically converged solution.

The ISV method with 1 Adjoint snapshot per revolution has produced blade geometries which significantly increase the instantaneous C_M around the snapshot location, across the range of the 24 cases studied. The results also show that the method can increase the average C_M (C_P) and can greatly reduce the aerodynamic loading while maintaining the VAWT average power coefficient. This carries a large value for designers since the fatigue loading of VAWT structures is reduced when there are smaller peaks/troughs in the aerodynamic load variation. This method can therefore be used to address the significant structural challenge of fatigue, which has been highlighted in the literature (MacPhee & Beyene, 2018).

Similar conclusions can be drawn to those presented in Section 5.5 regarding the geometry trends found in upwind snapshot cases (negative camber) and downwind snapshot cases (positive camber). In the upwind region, the Single-Blade model gives a good prediction of the VAWT flow field, and therefore similar sensitivity vectors and resulting blade geometries are seen for upwind cases (such as Case 90) with the ISV method. Case 270 shows a level of similarity but also some disagreement (between the ISV and Single-Blade methods), which is due to the challenges of representing the downwind flow field with a Single-Blade model (as discussed in Section 5.1). The upwind/downwind design paradox is observed again in the ISV results, because benefits to performance in the upwind region tended to carry detriment to the downwind region, and vice versa.

The resulting novel VAWT blades show, for each azimuthal position, what geometric features will improve the instantaneous blade torque. These blade geometries have been discussed regarding how their specific shapes prove to be favourable, aerodynamically. These results therefore could be directly useful for control system applications, providing the basis for an active control geometry profile that maximises the C_M continually at each azimuthal angle of the turbine. Furthermore, this set of desired geometries can be very efficiently determined using this semi-transient Adjoint based method.

The objective for ISV optimisation was specified as $+0.033 C_M$. This appears to be more appropriate than the approach of specifying the objective as a percentage of the instantaneous C_M value (discussed in Section 5.3.2) as was done in the Single-Blade optimisation. Section 6.3.2 has described the rationale for specifying the objective in this fashion for the ISV method.

The necessary details of setting up the Adjoint module have been determined and presented, including an aggression study which is a critical factor in making a correctly functioning optimisation process. This guidance provides a basis for the methodology to be readily adopted by the CFD community.

Importantly, this chapter has shown the effectiveness of Adjoint methods at producing novel blade geometries that provide significant aerodynamic benefits to VAWTs. This provides confidence and curiosity about the future benefits that could be seen by employing Adjoint methods in this way. The main drawbacks of the ISV method are discussed in Section 8.1.

7 MULTIPLE-APSHOT OPTIMISATION

This chapter explores how the optimisation process results can be influenced by combining more than one snapshot per VAWT revolution. In this investigation, both the number of snapshots per revolution, and the position of those snapshots are varied. The ISV method will be used as the platform of this investigation as it provides the more accurate flow field (compared to the Single-Blade method) and therefore the most valid sensitivity data at each snapshot.

7.1 Multi-Snapshot Investigation with the ISV Method

7.1.1 Philosophy of the Multi-Snapshot Investigation

This multi-snapshot section builds on the material of Section 6.4 where a 1-snapshot investigation was presented (studying the effect of snapshot *position* on the outcome of the optimisation process). The goal of the present study is to achieve greater improvements to the average C_M (C_P) of the turbine, than what was achieved using just 1-snapshot per revolution.

The number of possible cases that could be studied with multiple snapshots is very large because there are many different combinations of snapshot number and snapshot location. It is therefore desirable to devise a strategy for economising on the number of cases to reduce CPU cost of the study. The results and conclusions of the 1-snapshot investigation in Section 6.4.1 are used to help form this strategy and make it less arbitrary. The 1-snapshot results showed that using Adjoint sensitivity data from some locations did not lead to benefits in the VAWT performance, and so the multi-snapshot study will not use Adjoint snapshots from the areas listed in Table 7-1. Discounting snapshots within these ranges helps to reduce the number of possible cases to test. The next consideration to help streamline the investigation is the number of permutations to include for inter-snapshot spacing. To help illustrate this, Figure 7-1 shows three arbitrary snapshot positions marked on a VAWT blade performance curve.

An arbitrary spacing of 45 degrees is shown in Figure 7-1 which is wide enough such that each snapshot has a significantly varied flow field. If several snapshots were taken with a much smaller spacing this would offer little benefit compared to just 1 snapshot in that region. This is because the sensitivity data will be similar for each snapshot if all of them are close together. In the present study, an initial minimum snapshot spacing is therefore chosen as 15 degrees as this is judged to be sufficiently large to capture significant variations in the flow field between snapshot positions. This ensures the combination of sensitivity data from multiple snapshots is worthwhile.

Table 7-1 - Excluded Snapshot Positions from Multi-snapshot Study

Azimuthal Angle Range (degrees)	Reason for excluding this region from multi-snapshot cases
170-190	Case 180 (Table 6-5) showed a deteriorated average C_M . The flow field at, and near this position is peculiar compared to the majority of the cycle. This is because the AoA is around zero and is changing rapidly. This region only lasts for a short while and also bears only a small magnitude of C_M and therefore makes only a small contribution to value of average C_M . Optimising for this flow field would therefore not be valuable while detracting from performance in the more important regions of the cycle.
350-10	This region also bears only a small magnitude of C_M and therefore a small contribution to the value of average C_M . Optimising here would also detract from the high C_M regions elsewhere in the cycle.




Figure 7-1 – VAWT Performance Curve with Multiple Snapshots Marked

Another key choice is how many snapshots to consider - a range of cases up to a maximum of 12 snapshots per revolution are explored in this thesis. The optimisation algorithms could be configured for more snapshots, but in this thesis the scope is restricted to a maximum of 12 to make the case load more manageable.

Having made these initial decisions to reduce the study size, a large number of possible cases still remain. To tackle this, the multi-snapshot study is divided into 'Waves' whereby each subsequent wave explores a new route of investigation so that subsequent waves can choose more appropriate cases to study. Each wave builds on the progress of previous waves, and comprises a subsection of this chapter as follows:

- Wave 1: A selection of 2, 4, 6 and 12 snapshot cases are run as an initial search for promising cases. There are many cases divided into 'initial' and 'follow-up' cases.
- Wave 2: Exploring cases with snapshots focussed on the downwind region.
- Wave 3: Exploring cases with snapshots focussed on the upwind region.
- Wave 4: A refined selection of cases derived from the results of previous waves.

7.1.2 Wave 1

The first wave explores both the number of snapshots, as well as the snapshot position. The snapshot positions are referred to by their azimuthal angles around the turbines revolution. Table 7-2 shows the range of initial cases used; these provide a coarse scan of the solution space so that the results can later guide a set of more refined cases with parameters closer to the most successful ones.

The case IDs represent the wave (W#), the number of snapshots used (N#), and an identifying letter for each case. The 2-snapshot cases in Wave 1 maintain a constant inter-snapshot spacing of 180 degrees, while incrementing the snapshot positions through the cycle. Figure 7-2 illustrates this by showing the snapshot locations for the first four 2-snapshot cases. The same approach is used for the 4, 6 and 12-snapshot cases with the pairs of upwind/downwind snapshots being spaced 180 degrees apart.

The ISV optimisation process was run in the same way as that described in Chapter 6, except of course for the snapshot permutations. The details of the optimisation process itself will therefore not be recounted here and the focus of the present chapter is to explore the results of the cases rather than the methodology. The optimisation cases were run until a maximum value of the average VAWT C_M was reached. Table 7-2 shows the resulting improvements for each case alongside the case details (snapshot positions).

Table 7-2 -Wave 1 Initial Cases

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W1N2w	2	15, 195	1.7 % (4 iterations)
W1N2x	2	30, 210	2.4 % (4 iterations)
W1N2a	2	45, 225	5.8 % (33 iterations)
W1N2b	2	60, 240	3.7 % (16 iterations)
W1N2c	2	75, 255	3.3 % (13 iterations)
W1N2d	2	90, 270	2.1 % (8 iterations)
W1N2e	2	105, 285	1.3 % (8 iterations)
W1N2f	2	120, 300	1.3 % (6 iterations)
W1N2g	2	135, 315	1.4 % (7 iterations)
W1N2y	2	150, 330	6.5% (28 iterations)
W1N2z	2	165, 345	3.5 % (4 iterations)
W1N4e	4	30, 90, 210, 270	5.1 % (20 iterations)
W1N4a	4	45, 105, 225, 285	2.6 % (14 iterations)
W1N4b	4	60, 120, 240, 300	2.1 % (10 iterations)
W1N4c	4	75, 135, 255, 315	2.2 % (11 iterations)
W1N4d	4	90, 150, 270, 330	3.2 % (16 iterations)
W1N4f	4	105, 165, 285, 345	4.9% (13 iterations)
W1N6a	6	45, 75, 105, 225, 255, 285	2.6 % (13 iterations)
W1N6b	6	60, 90, 120, 240, 270, 300	2.0 % (8 iterations)
W1N6c	6	75, 105, 135, 255, 285, 315	1.7 % (11 iterations)
W1N6d	6	90, 120, 150, 270, 300, 330	1.8 % (9 iterations)

Table 7-2 continues on the next page

Table 7-2 Continued

WIN12a	12	15, 45, 75, 105, 135, 165, 195, 225, 255, 285, 315, 345	2.3 % (6 iterations)
WIN12b	12	10, 40, 70, 100, 130, 160, 190, 220, 250, 280, 310, 340	2.0 % (6 iterations)
WIN12c	12	20, 50, 80, 110, 140, 170, 200, 230, 260, 290, 320, 350	10.3% (30 iterations)




Figure 7-2 – Examples of snapshot spacing/pairing, for the first four 2-snapshot cases

The results of the Wave 1 initial cases showed that the greatest improvement was reached with 12 snapshots, although some of the 2-snapshot cases also performed well. The improvements are a comparison of the VAWT average C_M to the baseline VAWT which used NACA0018 blades (Rezaeiha et al. Vol 107, 2017). Discussion of these results will first focus on the initial 2-snapshot cases.

Case W1N2a reached a 5.8% increase to the average C_M with one upwind (45 degrees) and one downwind (225 degrees) snapshot. This could mean that the sensitivity data from these specific snapshot locations is very effective for improving the average C_M on this turbine. Since the 4-snapshot case W1N4a (45, 105, 225, 285) contained both the 45 and 225 degree snapshots while reaching only 2.6 % improvement, this suggests the 105 and 285 degree snapshots are less beneficial. Case W1N6a (45, 75, 105, 225, 255, 285) also contained both the 45 and 225 degree snapshots but also reached only 2.6 % improvement. This also suggests that the inclusion of additional snapshots did not provide benefit and may even 'dilute' the positive effect of the 45 and 225 degree snapshots.

Similarly, case W1N2y reached a 6.5% increase to the VAWT average C_M with 150/330 degree snapshots while cases W1N4d and W1N6d produced smaller improvements despite also including the 150/330 snapshots. This agrees with the argument that some regions of the cycle provide very beneficial sensitivity data while other regions do not.

The initial 4, and 6 snapshot cases generally had a resulting blade geometry which was characterised by a slight positive camber and a slight improvement to the average VAWT C_M . These characteristics are synonymous with the single-snapshot cases where the snapshot position is located in the downwind region (see Figure 6-11). Figure 7-3 shows the resulting geometries from a representative 4 snapshot case, and a representative 6 snapshot case. The resulting geometries from all 4 snapshot cases tested are relatively similar, and the same is true for the 6 snapshot cases. The resulting camber is less pronounced than that seen for the single downwind snapshot cases (see Figure 6-11) due to the presence of the upwind snapshots which acts as a counterbalance. Table 7-2 shows that in these 4 and 6 snapshot cases, the number of upwind and downwind snapshots are equal. These considerations indicate that the downwind snapshots have predominated over the upwind snapshots to produce this slight positive camber effect.





Figure 7-3 - (Left) Resulting geometry from Case W1N4a; (Right) Resulting geometry from Case W1N6a (contours are of non-dimensional static pressure)

There were many initial cases using 2-snapshots, most of which gave similar results to the 4 and 6 snapshot cases that have just been discussed. This again suggests there are large ranges of snapshot positions that do not provide much benefit to improving the average VAWT C_M . A few 2 snapshot cases however provided a significant increase to the performance. This means that a smaller range of snapshot positions tended to be highly effective at contributing to the increase of average VAWT C_M .

7.1.2.1 Early 2-Snapshot Cases

As shown in Table 7-2, Case W1N2a with 45/225 degree snapshots, provided an increase of 5.8% to the average VAWT C_M . This is a reasonable improvement upon the best case that was achieved using single-snapshot optimisation (namely 3.6%, see Table 6-5). Interestingly, the geometry produced by this case (W1N2a) exhibits a slight reflex camber as can be seen in Figure 7-4; this means that a positive camber is exhibited around the mid-chord and towards the 3/4 chord region, before then transitioning into a negative camber at the trailing edge.

It is desirable to determine which of the snapshots (45 or 225 degrees) has the most positive influence on the optimisation for Case W1N2a. Therefore, some additional 2-snapshot follow-up cases are made within Wave 1, recalling that in Wave 1 the snapshot positions are varied in 15 degree increments. Table 7-3 lists these follow-up cases, firstly with cases preserving the 45 degree snapshot while varying the downwind one; and secondly preserving the 225 degree snapshot while varying the upwind one.



**Figure 7-4 - Resulting geometry from Case W1N2a
(contours are of non-dimensional static pressure)**

Table 7-3 - Wave 1 Follow-up Cases (i)

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W1N2j	2	45, 210	4.4 % (16 iterations)
W1N2a*	2	45, 225	5.8% (33 iterations)
W1N2k	2	45, 240	6.8 % (42 iterations)
W1N2l	2	45, 255	7.2% (48 iterations)
W1N2m	2	225, 30	2.1 % (4 iterations)
W1N2n	2	225, 60	3.5 % (14 iterations)
W1N2o	2	225, 75	3.2 % (12 iterations)

**Wave 1 case shown previously but the result is repeated here for reference*

The results in Table 7-3 show a few things worth noting. Cases 'm', 'n' and 'o', all vary the upwind snapshot position away from 45 degrees, and all of these provide less improvement to average the VAWT C_M compared to Cases 'j', 'k' and 'l' which preserve the 45 degree snapshot. This demonstrates that the 45 degree position is the best choice for the upwind snapshot out of these cases. The best choice for the downwind snapshot was then shown to be 255 degrees. Case W1N2l has increased the average VAWT C_M (C_P) by 7.2% by combining 45/255 degree snapshots. This is closely followed by Case W1N2k, with 6.8% and snapshots at 45/240 degrees; both cases showed a sustained gradual improvement over many process iterations. Most cases so far have reached their peak C_M after around 10 or 15 iterations, but these cases required over 40 iterations. The multi-snapshot approach tends to require more process iterations (than the single-snapshot approach), and this is because the upwind and downwind sensitivity data tend to oppose each other to some extent (see Figure 6-9), thus creating a cancelling-out effect so that shape morphing deformations are smaller in each process iteration. Also, a more greatly improved blade (i.e., cases W12Nk and W12NI) will intrinsically require more process iterations to reach the maximum value, as only a small increase is made during a single iteration.

Cases, W1N2k and W1N2l, both exhibited a reflex camber effect in their resulting blade geometries (see Figure 7-5), in a similar fashion to Case W1N2a (see Figure 7-4). There is a subtle difference between the two blade shapes, and this key difference should be highlighted as it causes different aerodynamic behaviour. Blade W1N2k has a predominantly positive camber (0.8% at its greatest), starting from the leading edge and reaching all the way up to the position of 90% chord, before a reflex is seen so that the remaining 10% of the chord length has a slight negative camber (-0.2% at its greatest). On the other hand, for blade W1N2l only a small positive camber (less than 0.1% at its greatest) is present between the leading edge up to the reflex point around 72% chord. The negative camber predominates and occupies the trailing 28% of the blade, with its maximum (-0.7%) positioned around 91% chord. The key difference between these cases is whether the reflex camber is predominantly characterised by positive or negative camber.




Figure 7-5 - (Left) Resulting geometry from Case W1N2k. (Right) Resulting geometry from Case W1N2l (contours are of non-dimensional static pressure)

Previously, Section 6.4 described the effects that camber has on the shape of the performance curve of a turbine blade. These effects can again be observed in Figure 7-6 and Figure 7-7 which show the C_M as a function of azimuthal angle for blades W1N2k and W1N2l. The blade with predominating positive camber (W1N2k) shows increased C_M over the downwind region, and the blade with more predominating negative camber (W1N2l) shows increased C_M over the upwind region. While both of these blades produced a good increase to average VAWT C_M via a reflex camber geometry, it is interesting to observe this key difference. Furthermore, this difference is a product only of changing the downwind snapshot position from 240 degrees to 255 degrees.




Figure 7-6 - Performance of Case W1N2k with Adjoint-Optimised Geometries. (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




Figure 7-7 - Performance of Case W1N2I with Adjoint-Optimised Geometries. (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle

A more detailed aerodynamic discussion of similar aerofoils to these cases (namely, Case W4N2d and Case W4N12a) is deferred to Section 7.1.6, so that all the results of the multiple snapshot cases can first be discussed.


7.1.2.2 Late 2-Snapshot Cases

As can be seen in Table 7-2, Case W1N2y with 150/330 degree snapshots, provided an increase of 6.5% to the average VAWT C_M . This is a reasonable improvement on the best case that was achieved using single-snapshot optimisation (namely 3.6%, see Table 6-5). The geometry produced by this case (W1N2y) again exhibits a reflex camber as can be seen in Figure 7-8; in this case a slight positive camber (0.5% at its greatest) is seen between 0.25c and 0.75c, which then reflexes into a negative camber (-0.9% at its greatest) over the trailing edge from 75% chord onwards.

It is desirable to determine which of the snapshots (150 or 330 degrees) has the most positive influence on the optimisation for Case W1N2y. Therefore, some additional 2-snapshot follow-up cases are made within Wave 1. Table 7-4 lists these follow-up cases, firstly with cases preserving the 150 degree snapshot while varying the downwind one; and secondly preserving the 330 degree snapshot while varying the upwind one.

The results in Table 7-4 show that if the upwind snapshot is moved away from 150 degrees, the increase to the average C_M becomes much smaller. Also, it is found that the improvements can be greater if the downwind snapshot is moved to 300 degrees.

The pairing of 150/300 degree snapshots was less effective than the pairing of 45/225 degree snapshots and was also far more sensitive to slight changes in snapshot location. These investigations have therefore shown the value of doing a 'course search' of the solution space, followed by some additional refined cases.



**Figure 7-8 – Resulting geometry from Case W1N2y
(contours are of non-dimensional static pressure)**

Table 7-4 – Wave 1 Follow-up Cases (ii)

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W1N2p	2	150, 300	6.9% (31 iterations)
W1N2q	2	150, 315	6.8% (31 iterations)
W1N2y*	2	150, 330	6.5% (28 iterations)
W1N2r	2	150, 345	6.4% (32 iterations)
W1N2t	2	135, 330	1.4% (7 iterations)
W1N2u	2	165, 330	3.4 % (4 iterations)
* Wave 1 case shown previously but the result is repeated here for clarity.			

7.1.2.3 12-Snapshot Cases

Regarding the 12-snapshot cases, Case WIN12c reached a 10.3% increase to the average C_M , whereas the other two 12-snapshot cases (WIN12a and WIN12b) produced an increase of around just 2%. These three 12-snapshot cases are the only cases that could be run whilst maintaining a constant inter-snapshot spacing of 15 degrees and also avoiding the excluded regions set out in Table 7-1. Figure 7-9 shows the C_M as a function of azimuthal angle for the WIN12c blade. A typical behaviour of a downwind snapshot case is observed here; the downwind performance is enhanced while the upwind performance deteriorates. The geometry of the blade resulting from WIN12c is shown in Figure 7-10 and the positive camber seen here aligns with these trends in the upwind/downwind performance from Figure 7-9.

This geometry is interesting because it is unlike any of the blades seen so far from either the single-snapshot cases, or the 2, 4 and 6 snapshot cases. While the previously discussed blades showed a small thickness change (compared to the baseline blade), a slight fixing angle change, and a large (positive or negative) camber, this WIN12c blade is mainly characterised by a significant thickness reduction. It is 36% thinner than the NACA0018 baseline blade, and this is accompanied by a toe-out fixing angle (see Figure 1-8) of 1.7 degrees. A positive camber is seen reaching its maximum value (1.7%) at 53% chord. This camber is also unusual compared to previous blades where the camber tends to be focussed in the trailing quarter of the chord length. For the other 12 snapshot cases (WIN12a and WIN12b) the resulting blades look like typical 1-downwind-snapshot cases, exhibiting some positive camber and only a small thickness change, and not yielding a very large performance increase.

Further consideration of additional 12 snapshot cases is given in Wave 4 (Section 7.1.5), and more detailed aerodynamic discussion of this blade is deferred to Section 7.1.6.




Figure 7-9 - Performance of Case W1N12c with Adjoint-Optimised Geometries. (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle




Figure 7-10 – Resulting Blade Geometry from W1N12c (contours are of non-dimensional static pressure)

7.1.3 Wave 2

The Wave 1 results showed that a good improvement to the average C_M could be reached by using snapshots equally distributed across the upwind and downwind regions. Wave 2 has the purpose of testing the effect of focussing optimisation data in the *downwind* region. To do this a series of 4-snapshot cases were used with either 3 or all 4 of the snapshots located in the downwind region. These results will also show whether a multi-snapshot approach covering a range of azimuthal positions in the downwind is more advantageous than just 1 single downwind snapshot (i.e., a comparison to the results of Cases 195-345 from Section 6.4.1).

Wave 3 (see Section 7.1.4) will follow the same kind of investigation as Wave 2 but for cases focussed on the upwind region.

The results presented in Table 7-5 show that in comparison to the initial upwind/downwind 4-snapshot cases of Wave 1 (W1N4a - W1N4d in Table 7-2), no added benefit was gained by focussing the optimisation on the downwind region.

A small benefit can be observed however compared to the single-snapshot results of the downwind cases (in Table 6-5). With 1-snapshot the biggest average C_M improvement reached by a downwind case was 1.2%, while the multi-snapshot downwind cases W2N4a and W2N4c produced a 1.8% increase. This shows that there can be some value in using multiple snapshots to consider sensitivity data from a range of azimuthal angles.

These Wave 2 cases are not exhaustive, and only a small range of cases were considered. Nevertheless, downwind focussed cases will no longer be investigated in this study due to the smallness of the improvements found here.

Table 7-5 - Wave 2 Downwind Focussed Cases

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W2N4a	4	60, 210, 255, 300	1.8 % (5 iterations)
W2N4b	4	90, 210, 255, 300	1.5 % (3 iterations)
W2N4c	4	60, 240, 285, 330	1.8 % (5 iterations)
W2N4d	4	90, 240, 285, 330	1.4 % (4 iterations)
W2N4e	4	195, 240, 285, 330	1.1 % (2 iterations)
W2N4f	4	210, 255, 300, 345	1.1 % (2 iterations)

7.1.4 Wave 3

The Wave 1 results showed that a good improvement to the average C_M could be reached by using snapshots equally distributed across the upwind and downwind regions. Wave 3 has the purpose of testing the effect of focussing optimisation data in the *upwind* region. To do this a series of 4-snapshot cases were used with either 3 or all 4 of the snapshots located in the upwind region. These results will also show whether a multi-snapshot approach covering a range of azimuthal positions in the

upwind region is more advantageous than just 1 single upwind snapshot (i.e., comparison to the results of Cases 15-165 from Section 6.4.1).

The results in Table 7-6 show that in comparison to the initial upwind/downwind 4-snapshot cases of Wave 1 (W1N4a - W1N4d in Table 7-2), some benefit was gained by focussing the optimisation on the upwind region.

Some additional benefit can also be observed compared to the single-snapshot results of the downwind cases (in Table 6-5). Table 6-5 presented results after 10 process iterations, and Case 90, 105, 120, 135 had not yet converged to a maximum value of average VAWT C_M . To make the present comparison fair, these cases were continued, with the most successful being Case 90 providing a 2.9% increase in average VAWT C_M after 11 iterations. This was the greatest improvement reached by a single upwind snapshot case while the multi-snapshot upwind case W3N4d produced a 5.0% increase. This shows that there can be value in using multiple snapshots to consider sensitivity data from a range of azimuthal angles; the same conclusion was made in Wave 2.

These Wave 3 cases are not exhaustive, and only a small range of cases could be considered. Nevertheless, the upwind focussed cases will no longer be investigated in this study, as they are surpassed by the equally balanced upwind/downwind cases (from 2-snapshots and 12-snapshots).

Table 7-6 - Wave 3 Upwind Focussed Cases

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W3N4a	4	30, 75, 120, 300	2.0 % (6 iterations)
W3N4b	4	30, 75, 120, 255	2.1 % (6 iterations)
W3N4c	4	60, 105, 150, 300	4.7% (26 iterations)
W3N4d	4	60, 105, 150, 255	5.0% (23 iterations)
W3N4e	4	15, 60, 105, 150	3.6% (15 iterations)
W3N4f	4	30, 75, 120, 165	2.0 % (3 iterations)

7.1.5 Wave 4

The Wave 4 cases will use a finer resolution of 5 degrees (instead of 15 degrees in the previous waves) to explore the permutations of snapshot position in more detail. 2, 4, 6 and 12 snapshot cases will be studied and a clustering strategy will be implemented to explore another avenue for defining multi-snapshot cases.

7.1.5.1 2-Snapshot Cases (Wave 4)

The first set of Wave 4 cases are presented in Table 7-7 and these use the finer resolution (5 degrees) to improve upon the results of the best previous 2-snapshot case, namely W1N2I which produced an increase of 7.2% to the VAWT average C_M using snapshots at 45 and 255 degrees.

Table 7-7 shows that initially the upwind snapshot was held at 45 degrees, while the downwind snapshot position was varied. Once the best downwind position was found (260 degrees, case W4N2d), variations in the upwind snapshot location were varied. 45 degrees was found to remain the most successful so W4N2d (45/260 degrees) was the most effective 2-snapshot case resulting in an increase of 7.4% to the average C_M .

The performance of the optimised Case W4N2d VAWT can be seen in Figure 7-11. The W4N2d blade geometry is presented in Figure 7-16. Compared to the baseline NACA0018 blade, W4N2d is 7.1% thinner, has a toe-out (negative) 0.6 degree fixing angle (see Figure 1-8), and a reflex camber. This reflex camber consists of a slight positive camber between 0.25c-0.75c which then becomes zero at around 0.8c before reaching its greatest negative camber (-0.76%), at the 95% chord position. This geometry is similar to Case W1N2I (45, 255) shown in Figure 7-5 and Figure 7-7 (Section 7.1.6). The finer resolution of 5 degrees for Wave 4 has therefore helped improve the result slightly (from 7.2% to 7.4%). It is judged unnecessary to further pursue improvements via examining cases with a 1 degree resolution in snapshot location. This would cost significant computational time but would not provide a significant benefit to the VAWT performance.

For case W4N2d when considering the C_M of 1 blade, a larger upwind C_M peak is seen as well as increased C_M in the latter part of the downwind region. This blade has consistently improved the torque of the VAWT over the vast majority of the cycle. The aerodynamic discussion of this blade important blade is presented in Section 7.1.6.

Table 7-7 - Wave 4, 2-Snapshot Cases

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W1N2a*	2	45, 225	5.8 % (33 iterations)
W4N2a	2	45, 235	6.4% (42 iterations)
W1N2k*	2	45, 240	6.8 % (42 iterations)
W4N2b	2	45, 245	7.1% (44 iterations)
W4N2c	2	45, 250	7.3% (47 iterations)
W1N2l*	2	45, 255	7.2 % (48 iterations)
W4N2d	2	45, 260	7.4% (49 iterations)
W4N2e	2	40, 250	4.4% (19 iterations)
W4N2f	2	40, 255	4.2% (19 iterations)
W4N2g	2	40, 260	4.3% (19 iterations)
W4N2h	2	50, 250	4.9% (28 iterations)
W4N2i	2	50, 255	5.1% (30 iterations)
W4N2j	2	50, 260	5.0% (28 iterations)
*Wave 1 case shown previously but the result is repeated here for reference.			




Figure 7-11 - Performance of Case W4N2d with Adjoint-Optimised Geometries. (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle

7.1.5.2 4, 6 and 12-Snapshot Clustered Cases (Wave 4)

In Waves 1, 2 and 3 it was found that the 4, 6 and 12 snapshot cases generally provided only small/moderate improvements to the average C_M (between 1-5%). The exception being case W1N12c which has been discussed. In the present study, some additional 4, 6 and 12-snapshot cases are tested which cluster the snapshots around the key positions of 45 and 260 degrees, these being the most effective locations found in any 2-snapshot case. Table 7-8 shows the range of clustered cases. Producing these extra 4, 6 and 12-snapshot cases acknowledges the fact that only a limited number of cases were tested in Waves 1, 2 and 3.

Recalling the conclusions arising from the 2-snapshot case results, it was found that a small range of azimuthal positions provide favourable sensitivity data, while a large range of positions provide unfavourable sensitivity data. A clustered-snapshot approach was therefore chosen to pursue better results (than if equal inter-snapshot spacing is enforced) because the sensitivity data can be gathered from the key most favourable regions of the revolution while avoiding the unfavourable regions.

Table 7-8 - Clustered Cases with 4, 6 and 12 Snapshots

Case ID	Number of Snapshots, N	Snapshot Positions (degrees)	Improvement to VAWT average C_M (Process Iterations)
W4N4a	4	40, 50, 255, 265	7.5% (51 iterations)
W4N4b	4	35, 45, 250, 260	3.8% (15 iterations)
W4N4c	4	45, 55, 260, 270	4.5% (21 iterations)
W4N6a	6	40, 45, 55, 250, 260, 270	6.3% (37 iterations)
W4N6b	6	40, 45, 50, 255, 260, 265	7.4% (48 iterations)
W4N6c	6	30, 45, 60, 240, 260, 275	3.7% (14 iterations)
W4N12a	12	35, 40, 45, 50, 55, 60 250, 255, 260, 265, 270, 275	6.9% (42 iterations)
W4N12b	12	30, 35, 40, 45, 50, 55 245, 250, 255, 260, 265, 270	3.8% (15 iterations)

Considering these results alongside those of Waves 1, 2 and 3, it is readily seen that for 4 and 6-snapshot cases, the clustering approach is significantly more effective than the approach of maintaining a constant inter-snapshot spacing. This is demonstrated as the clustered cases W4N4a (+7.5% average C_M) and W4N6b (+7.4%) have exceeded cases W1N4e (+5.1%) and W3N4d (+5.0), which were the best 4 or 6 snapshot cases previously seen.

W4N4a produced a 7.5% increase to average C_M with 4 snapshots clustered around the 45/260 degree positions. The previously highlighted case W4N2d with two snapshots at these 45 and 260 degree positions produced a 7.4% increase, so this demonstrates that taking additional snapshots can be more effective than just 2 snapshots as long as they are around the beneficial range of azimuthal positions. Similarly, the clustered 6-snapshot case W4N6b (7.4%) matched the improvement produced by W4N2d.

The resulting blade from Case W4N4a has a similar shape and performance curve to that of Case W4N2d. The resulting blade from Case W4N6b also has a similar blade shape and performance curve to that of Case W4N2d – see Section 7.1.6.2 for the aerodynamic discussion of the Case W4N2d resulting blade.

Case W4N12a (+6.9%) provides a significant improvement to W1N12a and W1N12b (around 2%). This remains lower than W1N12c (+10.3%) however the W4N12a blade is far more typical (than W1N12c) and similar to the rest of the high performing blades produced. W4N12a has a reflex camber with a positive predominant camber, and other similar features (fixing angle and thickness) to that of W1N2k which has been discussed in Section 7.1.2. Aerodynamic discussion of W4N12a is presented in Section 7.1.6.3

7.1.6 Aerodynamic Analysis for Multi-Snapshot, ISV Optimisation

Aerodynamic analysis is given here for some key resulting blade geometries. Because many of the resulting blades exhibit geometrical similarities, aerodynamic analysis is not repeated. For several of the key blades which have been verbally discussed, the following list redirects the reader to a suitable Section on similar aerodynamic analysis:

- W1N2k, see similar blade W4N12a – Section 7.1.6.3.
- W1N2l, see similar blade W4N4d – Section 7.1.6.2.
- W4N4a, see similar blade W4N4d – Section 7.1.6.2.
- W4N6b, see similar blade W4N4d – Section 7.1.6.2.

7.1.6.1 Blade from Case W1N12c

Sections 7.1.2 - 7.1.5 showed that the most effective aerofoil produced for increasing the average VAWT C_M exhibited was W1N12c.

As shown in Figure 7-10 and mentioned in Section 7.1.2.3, the resulting blade from Case W1N12c has a toe-out (negative) fixing angle of 1.7 degrees (Figure 1-8), a positive camber (maximum of 1.7% chord), and a 36% reduction in thickness. These geometrical changes produce the performance curves shown in Figure 7-9, the aerodynamics of which will now be discussed.

Figure 7-12 shows the streamlines and Figure 7-13 shows the surface pressure coefficients for the Case W1N12c resulting blade, when the blade is at the 90 degree azimuthal position (see Figure 1-4). At this point in the revolution, this upwind blade produces less torque than the baseline blade (see Figure 7-9, Left).

Despite both blades having similar positive pressures on the pressure-surface, the negative (suction) pressures are significantly different for the Case W1N12c blade. A far greater leading edge suction peak is followed by weaker negative pressures along the mid span. Because the suction remains far lower than the baseline blade for most of the chord length, this results in a reduced C_M at this upwind position. The negative fixing angle of the W1N12c blade reduces the magnitude of the AoA at this position, which reduces the curvature demand on the flow, such that less flow acceleration is required around the suction surface. The positive camber also acts to reduce the lift coefficient and thereby tangential force coefficient (Barnard & Philpott, 2004), and so the two effects together are responsible for this reduction of pressure gradient and loss of torque (C_M). This reduced C_M of the upwind blade (90 degrees) is counterbalanced somewhat by the benefits taking place at the downwind blade at the same time (270 degrees), the net result is still a decrease of C_M for the VAWT overall at this azimuthal position (see Figure 7-9, Right). It is also worth noting that the small recirculating region at the trailing edge is eliminated with the Case W1N12c blade, which is due to the thickness reduction as well as the fixing angle change acting so that as the flow moves further aft, it can stay attached more easily in the region of adverse pressure gradient.

While the blade at 90 degrees azimuthal angle is in the upwind region, the downwind blade is simultaneously positioned at 270 degrees. In the downwind region, the real benefits of the Case W1N12c blade geometry take hold. Figure 7-14 shows the streamlines and Figure 7-15 shows the surface pressure coefficients for the Case W1N12c blade, when the blade is at the 270 degree azimuthal position. At this point, the W1N12c blade produces more torque than the baseline blade (see Figure 7-9, Right).




Figure 7-12 - VAWT Blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Resulting blade from Case W1N12c, (Bottom) Baseline Blade

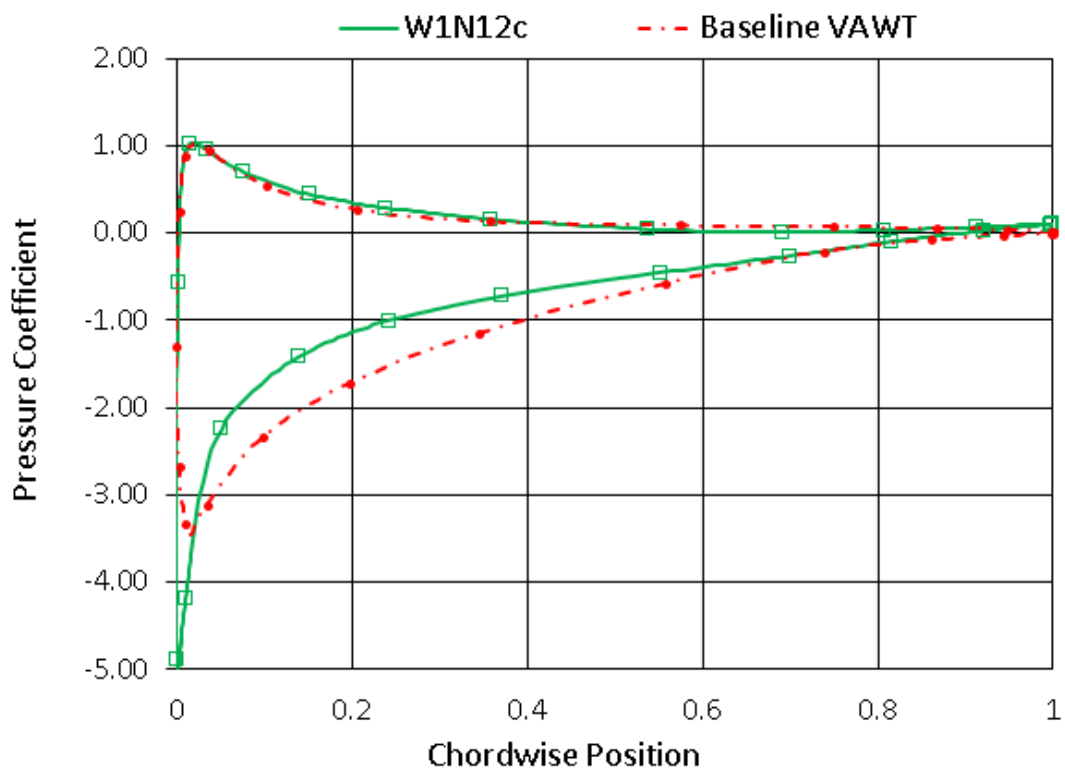


Figure 7-13 – Surface pressure coefficients at 90° azimuthal angle for the Case W1N12c resulting blade




Figure 7-14 - VAWT Blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Resulting blade from Case WIN12c, (Bottom) Baseline Blade




Figure 7-15 - Surface pressure coefficients at 270° azimuthal angle for the Case WIN12c resulting blade


The positive pressures are significantly higher for the Case W1N12c blade along the majority of the chord length. The leading edge suction peak is also much higher. In this position, the negative fixing angle increases the AoA of the blade which means greater flow acceleration is required to pass around the suction surface. The positive camber also acts to increase the lift of the blade thus improving the torque (C_M) generated.

The Case W1N12c blade has higher C_M over the majority of the cycle, with deterioration only in the range 30-170 (see Figure 7-9, Left). Considering the combination of both blades together (see Figure 7-9, Right), the Case W1N12c blade has significantly reduced the negative torque troughs while a smaller penalty in the torque peaks is observed. The net result is a significant (>10%) increase to the VAWT average C_M .

7.1.6.2 Blade from Case W4N2d

The interesting reflex camber geometry from the most successful 2-snapshot case W4N2d will now be discussed aerodynamically. Case W4N2d uses an upwind snapshot (45 degrees) and a downwind snapshot (260 degrees).

The resulting geometry from this case is shown in Figure 7-16 and this is similar to the resulting aerofoils of other cases such as W1N2a, W1N2l, (see Figure 7-4 and Figure 7-5) and W4N6a, and so the following discussion has a large amount of crossover to those aerofoils.



**Figure 7-16 - Resulting geometry from Case W4N2d
(contours are of non-dimensional static pressure)**

As stated in Section 7.1.5.1, the Case W4N2d blade has a toe-out (negative) fixing angle of 0.6 degrees (Figure 1-8), a 7.1% reduction in thickness, and a reflex camber which is predominantly negative. The positive camber between 0.25c-0.75c becomes zero around 0.8c, and then negative over the trailing edge (the greatest negative camber being -0.76% chord, positioned at 95% chord). These geometrical changes produce the performance curves shown in Figure 7-11, the aerodynamics of which will now be discussed.

Figure 7-17 shows the streamlines and Figure 7-18 shows the surface pressure coefficients for the Case W4N2d resulting blade when it is at the 90 degree azimuthal position. At this point, the blade produces more torque than the baseline blade (see Figure 7-11, Left).

Both blades have similar surface pressures on both upper and lower surfaces, spanning from the leading edge and over the mid chord. At this azimuthal position (90 degrees) the disparity in pressures is observed mainly at the trailing edge. Due to the negative camber element of the reflex camber, a slight increase in positive pressure is observed between 0.8c and 1.0c. This means that the pressure gradient between the leading/trailing faces of the aerofoil is stronger, exerting a greater force component in the blade tangential direction, and thereby improving the torque (C_M). This improved C_M of the upwind blade (90 degrees) is acting concurrently with the downwind blade positioned at 270 degrees. The downwind blade maintains a similar C_M to the baseline blade, thus the net result is an increase to the VAWT C_M overall (see Figure 7-11, Right). It is also worth noting that the small recirculating region at the trailing edge is reduced in size with the W4N2d blade, which is due to the thickness reduction as well as the fixing angle change causing a relaxation to the AoA.

Considering the resulting W4N2d blade in the downwind region (at 270 degrees), a deterioration to the performance would normally be expected since the blade geometry has improved the upwind performance. This expectation stems from the trends observed in literature and also within the blades produced in previous chapters of this thesis. It is fundamentally challenging to find a turbine blade that performs well in both the upwind/downwind regions owing to the VAWT 'design paradox' discussed in Chapter 3. However, Figure 7-19 shows the streamlines and Figure 7-20 shows the surface pressure coefficients for the Case W4N2d blade when the blade is at the 270 degree azimuthal position. At this point, this downwind blade produces a very similar torque to the baseline blade (see Figure 7-11, Left).




Figure 7-17 - VAWT Blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Resulting blade from Case W4N2d, (Bottom) Baseline Blade




Figure 7-18 - Surface pressure coefficients at 90° azimuthal angle for the Case W4N2d resulting blade




Figure 7-19 - VAWT Blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Resulting blade from Case W4N2d, (Bottom) Baseline Blade




Figure 7-20 - Surface pressure coefficients at 270° azimuthal angle for the Case W4N2d resulting blade

Figure 7-20 shows that the W4N2d blade has negative pressures slightly weaker than the baseline blade for the majority of the chord, up until the reflex camber near the trailing edge. A slightly greater positive pressure is observed until around the mid-chord.

The trailing edge region exhibits greater pressure gradients again, due to the reflex geometry producing a concentration of negative camber at the tail. Since the suction/pressure surfaces have swapped over (at around 0.75c), this means that the lower surface of the trailing edge (as depicted in Figure 7-19) carries the negative pressure.

These points considered together mean that the W4N2d resulting blade maintains a similar C_M to the baseline blade at 270 degrees, and indeed over the downwind region in general (see Figure 7-11, Left). While the downwind performance is maintained, the upwind advantages of the geometry act concurrently such that for the Case W4N2D VAWT, a similar or slightly higher C_M is produced across almost the entire revolution, with only minor deterioration in the ranges 30-75 and 200-260 degrees (see Figure 7-11, Right).

This is a very interesting result as it means that blade geometries can be produced which improve performance in some regions of the cycle, *without* large penalties elsewhere in the cycle. The net result is a significant (>7%) increase to the average VAWT C_M .

7.1.6.3 Blade from Case W4N12a

The resulting blade from Case W4N12a will now be discussed aerodynamically. Case W4N12a uses upwind snapshots clustered around 45 degrees and downwind snapshot clustered around 260 degrees. The resulting geometry is shown in Figure 7-21 and this is similar to the resulting aerofoil of other cases such as W1N2k (Figure 7-5), so the following discussion has a large amount of crossover.

The resulting blade from Case W4N12a has a toe-out (negative) fixing angle of 1.0 degrees (Figure 1-8), a 6.8% reduction in thickness, and a reflex camber which is predominantly positive. The positive camber (maximum of 0.8% chord) occupies most of the chord length before becoming zero at 88% chord and then negative at the trailing edge (the greatest negative camber being 0.4% positioned at 95% chord).

These geometrical changes produce the performance curves shown in Figure 7-22, the aerodynamics of which will now be discussed.




Figure 7-21 - Resulting geometry from Case W4N12a (contours are of non-dimensional static pressure)




Figure 7-22 - Performance of Case W4N12a with Adjoint-Optimised Geometries. (Left) C_M vs Azimuthal Angle of 1-blade, (Right) VAWT C_P vs Azimuthal angle

Figure 7-23 shows the streamlines and Figure 7-24 shows the surface pressure coefficients for the Case W4N12a resulting blade, when the blade is at the 90 degree azimuthal position. At this point in the revolution, this upwind blade produces less torque than the baseline blade (see Figure 7-22, Left).

The leading edge and mid-span positive pressures are very similar for the W4N12a blade compared to the baseline blade. Towards the trailing edge the high pressure can be seen which is associated with the negative camber after the reflex point. This provides some contribution to the tangential force component due to a stronger pressure gradient between the leading/trailing faces of the blade. The negative pressures are reduced for much of the forward part of the chord length. This is due to the reduction of AoA produced by the fixing angle combined with the positive camber such that less curvature/acceleration is demanded from the flow. The net result is that the 1 blade at 90 degrees provides slightly less C_M than the baseline blade (see Figure 7-22, Left).

Due to the fact that this blade has a predominantly positive camber, the real aerodynamic benefits are seen in the downwind region. Figure 7-25 shows the streamlines and Figure 7-26 shows the surface pressure coefficients for the W4N12a blade when the blade is at the 270 degree azimuthal position. At this point, this downwind blade produces more torque than the baseline blade (see Figure 7-22, Left).

The positive pressure peak at the leading edge is slightly greater for the W4N12a blade, followed by similar positive pressures along the rest of the chord length. On the suction side, despite a similar leading edge peak, the W4N12a blade maintains a greater suction all the way to around 75% chord which results in greater C_M . This is a result of the fixing angle increasing the AoA, and the positive camber also increasing lift of the blade in the downwind region.

The W4N12a blade at 270 degrees and indeed throughout the downwind region recovers the VAWTs overall torque so that the instantaneous total C_M is greater than the baseline VAWT (Figure 7-22, Right). The W4N12a VAWT thereby achieves an increased instantaneous C_M over the majority of the cycle with just a small range (between 40-90 degrees) having a marginal reduction in the instantaneous C_M .




Figure 7-23 – VAWT Blade streamlines (coloured by the non-dimensional static pressure) at 90° azimuthal angle. (Top) Resulting blade from Case W4N12a, (Bottom) Baseline Blade




Figure 7-24 – Surface pressure coefficients at 90° azimuthal angle for the Case W4N12a resulting blade




Figure 7-25 – VAWT Blade streamlines (coloured by the non-dimensional static pressure) at 270° azimuthal angle. (Top) Resulting blade from Case W4N12a, (Bottom) Baseline Blade




Figure 7-26 - Surface pressure coefficients at 270° azimuthal angle for the Case W4N12a resulting blade

7.2 Concluding Remarks on Multi-Snapshot Optimisation, ISV Method

The multiple-snapshot investigation has explored how the effectiveness of the semi-transient Adjoint based optimisation process can be affected by using multiple Adjoint snapshots across the VAWT revolution. The volume of possible multiple-snapshot cases to test is very large. This study has therefore not been exhaustive but constitutes a thorough investigation which allowed some key insightful conclusions to be made.

The major conclusion is that in comparison to the single-snapshot cases explored in Chapter 6, a multi-snapshot approach can be significantly more effective at increasing the average VAWT C_M (~5% and ~10% increases respectively). In many cases the multi-snapshot approach produced blade geometries that performed well in both the upwind and the downwind regions of the VAWT revolution, addressing to an extent, the VAWT design paradox described in Chapter 3. This means that benefits to performance in the one region did not carry a significant penalty to the opposing region. This is a very important point as the geometries which achieved this are novel and further investigation of them could offer a very valuable topic of future research for VAWT designers. Such aerodynamic behaviour is owed to the reflex camber geometries discovered here. These all exhibited a positive camber in the leading/mid-chord part of the blade, followed by a negative camber at the trailing edge. Within this set of geometries, some had a slight predominance of negative camber, and some a slight predominance of positive camber. For all of them a negative fixing angle was observed. In the upwind region, the fixing angle meant a relaxation to the AoA, but the aft-ward negative camber provided a localised high pressure at the trailing edge which improved the pressure gradient (tangential force component) between the leading/trailing faces of the aerofoil (Barnard & Philpott, 2004). In the downwind region, the negative fixing angle meant an increased AoA, and the positive camber element also contributed towards an increased lift and tangential force component. These blades also typically exhibited around a 6-7% reduction in thickness allowing a decrease in profile drag whilst not reaching a level of thinness that risks the introduction of separation and stall.

The discovery of these VAWT blades is due to the Adjoint based optimisation approach since it is decoupled from both the limitations of conventional blade parameterisation, and from the computational costs of non-gradient based optimisation methods. This allows these non-intuitive and highly effective blade geometries to be reached within reasonable timescales.

This study firstly identified 'beneficial regions' in the VAWT revolution via a series of 2-snapshot cases. Taking and combining Adjoint snapshots from these regions proved effective at improving VAWT average C_M . Namely these were around the 45 degree and 260 degree positions, and cases with snapshots away from these beneficial regions tended to perform less well. This trend held true when using 2, 4 or 6-snapshots. Additionally, taking a balanced number of snapshots from the upwind and downwind regions yielded better results than skewing more snapshots to either the upwind or downwind. For 4 and 6-snapshot cases, clustering the snapshots around the beneficial regions proved to be effective and thereby significant increases were realised with 2, 4, 6 and 12-snapshots using this method.

The resulting blades from the clustered multi-snapshot cases tended to be similar to the best performing 2-snapshot cases (around which the clustered cases are based). This is not surprising, but it should be noted that the most successful 2-snapshot case W4N2d (+7.4%) was matched or improved upon slightly by the clustered 4 and 6-snapshot cases W4N6b (7.4%) and W4N4a (+7.5%). This allows a conclusion to be drawn that collects the knowledge gained across the entire multi-snapshot investigation. Taking additional sensitivity data from a range of positions in the revolution can be beneficial, but only if those additional snapshots remain within the beneficial regions identified.

The result of the optimisation process can be very sensitive to the snapshot position and even a 5 degrees change to one of the snapshot positions can significantly deteriorate the percentage increase to the average C_M that is achieved.

A recommendation for applying this semi-transient Adjoint based optimisation method to another VAWT, would be to begin by testing a range of 2-snapshot cases (as per the Wave 1 Cases W1N2#) to identify the beneficial regions (for the particular turbine and operating conditions in question). Honing in on these beneficial regions could yield a resulting optimised blade that is satisfactory to the designer without the need to explore additional clustered cases with more snapshots. Such an approach is readily implementable using the guidance provided in this thesis.

The multi-snapshot cases took considerably longer to run than the 1-snapshot cases of (Chapter 6) and the reasons for this are described in Section 7.1.2.1. The following details give a summary of the computational cost for the cases that proved most effective (i.e., cases with at least 7% improvement to the VAWT average C_M):

- 16 cores were used on an HPC at the University of Sheffield, and this is a modest number of cores that some modern desktop PCs can match. Faster run times could be achieved by using more cores if desired. 16 cores were used because many cases were being run concurrently during the project.

- Average Wall-Clock-Time (WCT) for the best/longest running cases was between 20-25 days.
- Most of this CPU time is due to the 10 VAWT revolutions which take place after each iteration, to allow the flow field to stabilise.
- Significant variation was seen in the running speed of cases and this depended on the node quality within the HPC. The running speed was between 8hrs-15hrs per iteration of the optimisation process.
- The multi-snapshot study was conducted with the aggression levels determined from a 1-snapshot optimisation study (see Section 6.3.2). It was mentioned in Section 7.1.2.1 that with multiple snapshots there is a cancelling out effect when sensitivity data is combined from the upwind/downwind regions and this reduces the effective aggression of the process. It is therefore envisaged that significant WCT savings could be reached by determining a revised multi-snapshot aggression ('objective target value', see Section 6.3.2).

Importantly, this chapter has shown that a multi-snapshot approach allows the semi-transient Adjoint method to produce even more effective novel blade geometries which give significant aerodynamic benefits to VAWTs. This provides further confidence about the future possibilities of employing Adjoint methods in this way. The main drawbacks of the ISV method are discussed in Section 8.1.

8 CONCLUSIONS & FUTURE WORK

The aim of this research was to develop for the first time, methods for applying a steady Adjoint solver to the unsteady aerodynamics problem of VAWT optimisation using CFD (see Chapter 4). The functionality of these new methods was to be demonstrated by their application to a sample VAWT via 2D CFD simulations.

To address this aim suitable VAWT modelling approaches were first researched. This included a standard VAWT CFD model constructed according to the recommendations in the literature, and also a low-cost representation of VAWT blade flow fields via an oscillating Single-Blade model. Subsequently, the key elements of an Adjoint optimisation framework were identified from the literature. ANSYS Fluent was then chosen as a suitable platform for the work, providing all the necessary elements in a single software package. Preliminary feasibility studies implementing Adjoint based optimisation were then applied to simple steady flow CFD models before progressing onto the oscillating Single-Blade model and then finally the full VAWT model. Incrementing the model complexity in this way aided the systematic progression of the method and its validation. The methods were developed in conjunction with a high TSR sample VAWT. This minimised the unsteady effects of the transient flow such as dynamic stall, thus avoiding additional complications in the process of developing the methods.

Robust algorithms/scripts were produced for applying semi-transient Adjoint based optimisation to both the Single-Blade (VAWT approximation) model and the VAWT model. Points of best practice for implementation of the method have been documented to assist in addressing the significant vacuum of publicly available guidance on this topic. Automatic script generating tools were created to rapidly produce permutations of the scripts in order to study some of the key variables in the optimisation process. Investigations were then performed to study the effect of altering the position of the 'Adjoint snapshots' (see Section 5.2.2) on the outcomes of the optimisation process. This was done for both model types (Single-Blade and VAWT) before a multi-snapshot investigation was made using the full VAWT model. The VAWT blade improvements of the optimisation process were evaluated for a large variety of test cases.

The results of these investigations demonstrated that a steady Adjoint solver can be successfully applied to the transient aerodynamics problem of VAWTs. The proposed semi-transient Adjoint based optimisation method departs from conventional parameter-based design methods, and the results were novel blade geometries which exhibited significant increases to the VAWT instantaneous power coefficient. Improvements can be seen by applying the optimisation methods with

just 1 Adjoint snapshot per revolution. For the VAWT optimisation method, most of the 1-snapshot cases increased the local moment coefficient of a blade (at the snapshot position), by significant margins. Furthermore, the resulting novel VAWT blade geometries from 1-snapshot optimisation provided new insight into the most desirable blade shapes across the range of azimuthal angles in the revolution. This led to the observation that a large range of upwind cases share similarity, and a large range of downwind cases share similarity, thus forming two key groups. Upwind snapshot optimisation tended to result in a toe-out fixing angle with negative camber, while downwind snapshot optimisation resulted in a positive camber. The two groups are largely converse to each other and improving one region be it upwind or downwind, tends to detract from the other. The resulting performance curves clearly showed, across 1-snapshot cases, that improvements to the upwind carry a penalty in the downwind, and vice versa – this is the design paradox of VAWT aerodynamics. A natural conclusion of this is that active control techniques could be a necessity in achieving high VAWT efficiencies that are competitive with HAWTs. The results of the 1-snapshot optimisations also showed that the semi-transient Adjoint method can greatly reduce the aerodynamic loading (and thus fatigue damage) whilst preserving or even improving the VAWT average power coefficient. The implications of this for structural design make this method very important and valuable, as the fatigue damage problem of VAWTs has been emphasised in the literature (MacPhee & Beyene, 2012).

Considering multiple Adjoint snapshots across a revolution aimed to address the upwind/downwind design paradox. By considering the combined sensitivity data taken from different azimuthal angles, a blade could be formed with improved performance over the majority of the revolution. With this increased ‘resolution’ of the method, even greater increases to the average power coefficient were produced. Using a multi-snapshot approach, the proposed method was able to improve the average VAWT power coefficient by over 10% for the particular VAWT tested. Furthermore, the most effective blades discovered tended to exhibit a reflex camber geometry which is very interesting. Upon examining the aerodynamic behaviour of these novel blade geometries, valuable new insight was given on how certain geometrical features can maintain a higher C_M over the majority of the cycle.

The novelty and value of this work is that new methods have been proposed for applying Adjoint based optimisation to the problem of VAWT aerodynamics. Adjoint based methods hold several advantages over other approaches (which have been discussed in Chapter 3), and this work provides the foundation for an exciting new avenue of VAWT research. The aims of the research project have therefore been fulfilled. Such methods have not been applied to VAWTs in the literature so far, and

as such, novel algorithms for Adjoint optimisation of VAWTs using CFD were created. In addition, Adjoints have been used in other fields with great success, so this research area could prove to be very beneficial to the progression of VAWT technology in the future. The viability of semi-transient Adjoint based optimisation for VAWTs has been shown using ANSYS Fluent but these methods can be implemented in alternative CFD codes which contain an Adjoint solver. A low-cost modelling approach has also been implemented, which in conjunction with an Adjoint sensitivity analysis provides an efficient overall optimisation method.

The results in Section 6.4.1 present a set of novel blade geometries which are un-hampered by conventional aerofoil parameterisation due to the Adjoint method. These geometries show the aerodynamicist, for each azimuthal position, which geometries will improve the instantaneous blade torque - this carries a large value for designers. Aerodynamic discussion of the novel blades has also been provided, explaining the significant improvements to the instantaneous moment coefficient that were achieved. A similar set of ideal instantaneous geometries can be readily generated for other VAWTs using this method, and this is due to the practical information provided. Such detailed guidance is a novel contribution in itself, thus helping to bridge a massive shortfall in publicly available information on the topic. In addition, although this guidance information is specifically for VAWT applications, it can be of much value in general for users looking to apply the ANSYS Fluent Adjoint module.

Final Remarks

This thesis has shown that VAWT blade shape design can be significantly benefitted by the application of Adjoint based optimisation. Opening this avenue of new research is an important and valuable contribution to the field of VAWT technology because both industrial & academic engineers can readily adopt, implement, and build upon these Adjoint based methods to produce more efficient VAWT designs. This thesis showed that a 10% increase can be made to the VAWTs average power coefficient within just a short space of time. Efficiency improvements of this order can substantially increase the annual energy yield of an individual turbine, thereby making VAWT projects more commercially viable. It will be exciting to see what can be achieved in future developments of VAWT technology, as the VAWT becomes a more powerful, efficient, and economically competitive energy source.

8.1 Evaluation/Limitations of Methods

The main limitation of the semi-transient Adjoint based optimisation methods proposed here is that the average VAWT C_P (C_M) cannot be optimised directly. This is because an Adjoint solver for instantaneous/steady flow fields is used in conjunction with transient simulations. A transient Adjoint solver would be required to consider direct optimisation of time averaged quantities. This applies to all the method variations in this thesis (Chapters 5, 6, and 7). Despite this, the results have shown that a semi-transient method can be highly effective at improving the VAWT average power coefficient, C_P . Considering the high associated cost of calculating transient Adjoint solutions, this could also be looked on as a benefit of the semi-transient method since calculating the steady Adjoint solution is relatively fast.

The SST $k-\omega$ turbulence model was used for the CFD modelling rather than a more accurate approach such as LES, which also effects the CPU cost. The CFD simulations have been 2D instead of 3D which massively reduces computation cost but also tends to carry an associated error in performance estimation. This modelling approach is however aligned with much of the existing literature where low computation cost is needed. Since 2D CFD most accurately reflects 3D behaviour at the mid-span of the blade (for an H-rotor), the resulting blade geometries proposed here are most optimised for that position, away from the blade tips where the spanwise flow effects are minimal (El Sakka, 2020). Despite the modelling assumptions, the effectiveness of the semi-transient Adjoint optimisation method has been shown.

The VAWT used throughout the development of the methods had a TSR of 4.5. While this choice reduced the unsteady effects of the flow, it also means that the conclusions drawn so far could be limited to high TSR cases. Furthermore, a single flow condition has been considered in this optimisation. The method can readily be used to improve a VAWTs performance at a single TSR, but in its current form cannot consider multiple operating conditions. A constant oncoming wind speed and direction has also been assumed.

8.1.1 Single Blade Method

The Single-Blade method produced improvements to the average VAWT C_M (+3.5%). It was also found that this method could reduce aerodynamic fluctuating loads whilst maintaining the average C_M . Cases with the snapshot positioned in the upwind region tended to exhibit a negative camber (see Figure 1-7 for notation). This increases the lift coefficient and tangential force component in the upwind region where the AoA is negative (see Figure 1-8). Cases with the snapshot positioned in

the downwind region tended to exhibit a positive camber, and this improved the tangential coefficient in the downwind where the AoA is positive.

A limitation of the Single-Blade method (Chapter 5) is that the Single-Blade approximation model has some inaccuracy in predicting the VAWT blade flow field, especially in the downwind region. This means that the true performance of the optimised blades is only known after a 'candidate VAWT' model is made using the altered Single-Blade geometry. Producing the candidate VAWTs costs much time and effort when many cases are to be studied.

The Single-Blade method was demonstrated via a series of 1-snapshot optimisations. This poses an inherent limitation since optimisation data was only considered from a single point in the VAWT cycle. Nevertheless, even in this basic formulation positive results were shown where VAWT performance can be quickly improved.

Both these limitations were addressed via the investigation of the ISV method using multiple snapshots (Chapter 7).

8.1.2 ISV Method

The ISV method with 1-snapshot provided improvements to the average VAWT C_M (+3.6%) as well as significant improvements to the instantaneous VAWT C_M (more than +95%, given as a percentage of the average C_M for one blade). These large, localised improvements were located at the region of the VAWT revolution where the snapshot was positioned for each case. Greater improvements were found compared to the Single-Blade method because a VAWT model was used directly in the optimisation, providing more accurate sensitivity data. For some cases, large benefits in the reduction of fluctuating aerodynamic loads were also produced whilst maintaining average VAWT C_M . In a similar fashion to the Single-Blade method, cases with the snapshot positioned in the upwind region tended to exhibit a negative camber (see Figure 1-7 for notation). This increases the lift coefficient and tangential force component in the upwind region where the AoA is negative (see Figure 1-8). Cases with the snapshot positioned in the downwind region tended to exhibit a positive camber, and this improved the tangential coefficient in the downwind where the AoA is positive.

As for the Single-Blade method, the ISV (in Chapter 6) method was demonstrated via a series of 1-snapshot optimisations. This poses an inherent limitation since optimisation data was only considered from a single point in the VAWT cycle. Nevertheless, even in this basic formulation positive results were shown where VAWT performance can be quickly improved. This limitation was addressed via the investigation of the ISV method using multiple snapshots (Chapter 7).

The advantage in flow field accuracy of the ISV method compared to Single-Blade method comes at the cost of run-time. Considering a range of cases it was found that the ISV method takes approximately twice as long as the Single-Blade method for running a 1-snapshot optimisation process for 10 iterations. This is due to the model size and the necessity for more cycles of the turbine model to reach a periodically converged solution.

8.1.3 Multiple-Snapshot ISV Method

The ISV method with multiple-snapshots provided significant improvements to the average VAWT C_M (+10.3%). This was due to the fact that sensitivity data was combined from various regions of the VAWT revolution, which produced aerofoils that perform well over the majority of the cycle. The most successful cases tended to exhibit a reflex camber geometry, where a positive camber was found over the mid-chord, before changing to a negative camber at the trailing edge. This allowed aerodynamic benefit to be taken in both the upwind and downwind regions. With the instantaneous VAWT C_M improved over the majority of the cycle, significant increases to the average VAWT C_M could be reached.

The multiple-snapshot ISV method (Chapter 7) aimed to address the limitation of using a steady solver to optimise average C_M . By considering multiple snapshots around the VAWT's revolution, significant improvements to average VAWT C_M were observed. Although the average VAWT C_M is not optimised directly, these results showed that considering optimisation data from multiple Adjoint snapshots allows VAWT blades to be produced which perform well over the majority of the turbine revolution.

8.2 Future Work

Future work could improve and refine this type of method, so that even greater improvements to VAWT performance could be reached. More generally applicable incarnations of the method could also be found. The foundation for applying Adjoint based optimisation to VAWTs has been set by this thesis, and there are many possible topics that could be studied to provide added value. The recommended topics for future work envisaged by the author are as follows:

- The results show the method is highly effective at improving the instantaneous C_M at the snapshot position. A natural route to make best use of these particular benefits is to couple results from a 1-snapshot study with morphing blade active control techniques. In this fashion the optimum instantaneous blade geometries could be pursued while the blade moves around the revolution. It is expected that significant gains to the average C_P could be realised in this way, since the upwind/downwind design paradox can be alleviated directly.
- Refinements could be made to the Single-Blade oscillation profile such that it more accurately represents the VAWT flow field. The sensitivity data arising from the Adjoint snapshots would therefore be more appropriate for realising geometry benefits to the candidate VAWT blade. Another possible approach could be to apply AoA sampling techniques to a VAWT model, such as that proposed by El Sakka (2020), and then to implement this AoA profile into the Single-Blade model. This could improve the prediction of the Single-Blade model while retaining the CPU benefits of the smaller model.
- The multi-snapshot investigation could be extended to cover more cases. The range of possible case combinations is enormous when considering the possible variations in number of snapshots, snapshot position, and snapshot weighting. Only a certain number of cases were able to be studied in this work and it was not exhaustive. A study of applying non-equal weightings to snapshots could also be undertaken.
- More powerful turbulence models could be applied to both ISV and Single-Blade methods. There is a potential avenue for doing this in ANSYS Fluent but the use of alternative CFD codes may be required. This would improve the accuracy in predicting the VAWT flow field and enhance the quality of the sensitivity data.

This would also likely improve the matching between any experimental data that could be produced for validation purposes.

- The method could be applied to other turbines of different sizes and operating conditions, etc. This would then inform the development of a generalised approach for choosing the appropriate number/location of snapshots for optimising a new VAWT.
- Testing of the method at a range of TSRs would be beneficial to understand the impact that the TSR has on the effectiveness of the method. Discoveries could be made about the modifications required in the method to facilitate its use with low TSR cases.
- The method could also be extended to 3D simulations, allowing novel spanwise geometry variations of the VAWT blades to be discovered. This could also inform the optimisation of winglets or even connecting arm geometries. Crude blade tip and connecting arm geometries are known to play a significant role in blade performance losses. Well-designed winglets can reduce flow leakage at the tip maintaining a useful pressure gradient (El Sakka, 2020) and minimising the induced drag.
- A study could be made to select a more optimum level of aggression to be used in the multi-snapshot ISV optimisation. As stated in Section 8.1.3, the multi-snapshot cases were run with a level of aggression selected for 1-snapshot optimisation. Exploring this could significantly reduce run times of multi-snapshot optimisation.
- Because the blades produced by these methods are novel, there exists no numerical or experimental data about these aerofoils. Conducting experimental campaigns would be a valuable means to validate the performance of these novel blades and thereby the optimisation process.

9 REFERENCES

1. ANSYS Inc. ANSYS Help, Release 18.2. (2017).
2. ANSYS Inc. Introduction to ANSYS Fluent, Module 07: Turbulence. Release 17.0. (2016)
3. Ai, Q., Weaver, P. M., Barlas, T. K., Olsen, A. S., Madsen, H. A. & Anderson, T. L. Field testing of morphing flaps on a wind turbine blade using an outdoor rotating rig. *Renewable Energy* Vol 133. (2019), pp. 53-65.
4. Almohammadi, K. M., Ingham, D. B., Ma, L. & Pourkashanian, M. CFD Sensitivity of a Straight-Blade Vertical Axis Wind Turbine. *Wind Engineering* Vol 36. (2012), pp. 571-588.
5. Almohammadi, K. M., Ingham, D. B., Ma, L. & Pourkashanian, M. 2D CFD Analysis of the Effect of Trailing Edge Shape on the Performance of a Straight-Blade Vertical Axis Wind Turbine. *IEEE Transactions on Sustainable Energy* Vol 6. (2015), pp. 228-235.
6. Baker, J. R. Features to aid or enable self starting of fixed pitch low solidity vertical axis wind turbines. *Journal of Wind Engineering and Industrial Aerodynamics* Vol 15. (1983), pp. 369-380.
7. Balduzzi, F., Bianchini, A., Maleci, R., Ferrara, G. & Ferrari, L. Critical issues in the CFD simulation of Darrieus wind turbines. *Renewable Energy* Vol 85. (2015), pp. 419-435.
8. Barnard, R., H. & Philpott, D., R. *Aircraft Flight*, Third Edition. Pearson, Prentice-Hall (2004).
9. Bhutta, M. M. A., Hayat, N., Farooq, U.H., Ali, Z., Jamil, S.R. & Hussain, Z. et al. Vertical axis wind turbine – A review of various configurations and design techniques. *Renewable and Sustainable Energy Reviews* Vol 16. (2011), pp. 1926-1939.
10. Bianchini, A., Ferrara, G. & Ferrari, L. Design guidelines for H-Darrieus wind turbines: Optimization of the annual energy yield. *Energy Conversion and Management* Vol 89. (2014), pp. 690-707.
11. Bianchini, A., Balduzzi, F., Di Rosa, D & Ferrera, G. On the use of Gurney flaps for aerodynamic performance augmentation of Darrieus wind turbines. *Energy Conversion and Management* Vol 184. (2019), pp. 402-415.
12. Buchner, A.J., Lohrya, M. W., Martinelli, L., Soria, J. & Smits, A. J. Dynamic stall in vertical axis wind turbines: Comparing experiments and computations. *Journal of Wind Engineering and Industrial Aerodynamics* Vol 146. (2015), pp. 163-171.

13. Carpentieri, G. PhD Thesis: An Adjoint-Based Shape-Optimization Method for Aerodynamic Design. Delft University of Technology. (2009).
14. Carrigan, T. J., Dennis, B. H., Han, Z. X. & Wang, B. P. Aerodynamic Shape Optimization of a Vertical-Axis Wind Turbine Using Differential Evolution. *International Scholarly Research Network Renewable Energy Vol 2012*. (2011), pp. 1-16.
15. Chan, C. M., Bai, H. L. & He, D. Q. Blade shape optimization of the Savonius wind turbine using a genetic algorithm. *Applied Energy Vol 213*. (2018), pp. 148-157.
16. Chehour, A., Younes, R., Ilinca, A. & Perron, J. HAWT Review of performance optimization techniques applied to wind turbines. *Applied Energy Vol 142*. (2015), pp. 361-388.
17. Coppin, J. PhD Thesis: Aerodynamics, Stability and Shape Optimisation of Unmanned Combat Air Vehicles. University of Sheffield. (2014).
18. Coton, F. N., Galbraith, R. A. M. & Jiang, D. The Influence of Detailed Blade Design on the Aerodynamic Performance of Straight-Bladed Vertical Axis Wind Turbines. *IMechE Part A: Power and Energy Vol 210*. (1996), pp. 65-74.
19. Danao, L. A., Eboibi, O. & Howell, R. J. An Experimental Investigation into the Influence of Unsteady Wind on the Performance of a Vertical Axis Wind Turbine. *Applied Energy Vol 107*. (2013), pp. 403-411.
20. Danao, L. A., Edwards, J., Eboibi, O. & Howell, R. J. A numerical investigation into the influence of unsteady wind on the performance and aerodynamics of a vertical axis wind turbine. *Applied Energy Vol 116*. (2013), pp. 111-124.
21. Daróczy, L., Janiga, G. & Thévenin, D. Computational fluid dynamics based shape optimization of airfoil geometry for an H-rotor using a genetic algorithm. *Engineering Optimization*, (2018), pp. 1-17.
22. Darrieus, G. J. M. Turbine Having its Rotating Shaft Transverse to the Flow of the Current. United States Patent Office. Pat No: 1,835,018 (1931).
23. Day, H., Ingham, D., Ma, L., Pourkashanian, M. Adjoint based optimisation for efficient VAWT blade aerodynamics using CFD. *Journal of Wing Engineering & Industrial Aerodynamics Vol 208*. (2021), 104431.
24. Dhert, T., Ashuri, T. & Martins, J. Aerodynamic shape optimization of wind turbine blades using a Reynolds-averaged Navier–Stokes model and an adjoint method. *Wind Energy Vol 20*. (2016), pp. 909-926.
25. Dominy, R., Lunt, P., Bickerdyke, A. & Dominy, J. Self-starting Capability of a Darrieus Turbine. *IMechE Part A: Power and Energy Vol 221*. (2006), pp 111-120.

26. Dybvig, R. K. *The Scheme Programming Language, Second Edition*. Available online at: <https://scheme.com/tspl2d>, Prentice Hall (1996).
27. Eboibi, O., Danao, L. A. M. & Howell, R. J. Experimental investigation of the influence of solidity on the performance and flow field aerodynamics of vertical axis wind turbines at low Reynolds numbers. *Renewable Energy* Vol 92. (2016), pp. 474-483.
28. Economon, T. D., Palacios, F., Copeland, S. R., Lukaczyk, T. W. & Alonso, J. J. SU2: An Open-Source Suite for Multiphysics Simulation and Design. *AIAA* Vol 54. (2016), pp. 828-846.
29. Edwards, J. M. PhD Thesis: The influence of aerodynamic stall on the performance of vertical axis wind turbines. University of Sheffield. (2012).
30. Edwards, J. M., Danao, L. A. & Howell, R. J. PIV Measurements and CFD Simulation of the Performance and Flow Physics and of a Small-Scale Vertical Axis Wind Turbine. *Wind Energy* Vol 18. (2013), pp. 201-217.
31. Eggensteiner, G. Product Presentation: ANSYS Fluent Adjoint Solver. ANSYS Inc. (2012).
32. El Sakka, M. PhD Thesis: The Aerodynamics of Fixed and Variable Pitch Vertical Axis Wind Turbines. University of Sheffield. (2020).
33. Errico, R. M. What is an adjoint model? *Bulletin of the American Meteorological Society* Vol 78. (1997), pp. 2577-2591.
34. Ferreira, C. S. & Geurts, B. Aerofoil optimization for vertical-axis wind turbines. *Wind Energy* Vol 18. (2014), pp. 1371-1385.
35. Ferreira, C. S., van-Kuik, G., van-Bussel, G. & Scarano, F. Visualization by PIV of dynamic stall on a vertical axis wind turbine. *Exp Fluids* Vol 46. (2008), pp. 97-108.
36. Ferreira, C. J. S., van Zuijlen, A., Bijl, H., van Bussel, G. & van Kuik, G. Simulating dynamic stall in a two-dimensional vertical-axis wind turbine: verification and validation with particle image velocimetry data. *Wind Energy* Vol 13. (2009), pp. 1-17.
37. Ferrer, E. & Montlaur, A. *CFD for Wind and Tidal Offshore Turbines*. Springer. (2015).
38. Fujisawa, N. & Shibuya, S. Observations of dynamic stall on Darrieus wind turbine blades. *Journal of Wind Engineering and Industrial Aerodynamics* Vol 89. (2000), pp. 201-214.
39. Funke, S. W., Farrell, P. E. & Piggott, M. D. Tidal turbine array optimisation using the Adjoint approach. *Renewable Energy* Vol 63. (2013), pp. 658-673.

40. Giles, M. B. & Pierce, N. A. An Introduction to the Adjoint Approach to Design. *Flow, Turbulence and Combustion* Vol 65. (2000), pp. 393-415.
41. Gorelov, D. N. Analogy Between a Flapping Wing and a Wind Turbine with a Vertical Axis of Revolution. *Journal of Applied Mechanics and Technical Physics* Vol 50. (2009), pp. 292-299.
42. Gosselin, R., Dumas, G. & Boudreau, M. Parametric study of H-Darrieus vertical-axis turbines using uRANS simulations, 21st Annual conference of the CFD society of Canada #178, Session 13-6. (2013).
43. Gosselin Gosselin, R., Dumas, G. & Boudreau, M. Parametric study of H-Darrieus vertical-axis turbines using CFD simulations. *Journal of Renewable Sustainable Energy* Vol 8. (2016), 053301.
44. Hand, B., Kelly, G. & Cashman, A. Numerical simulation of a vertical axis wind turbine airfoil experiencing dynamic stall at high Reynolds numbers. *Computers & Fluids* Vol 149. (2017), pp. 12-30.
45. Hill, N., Dominy, R., Ingram, G. & Dominy, J. Darrieus turbines: The physics of self-starting. *IMEchE Part A: Power and Energy* Vol 223. (2008), pp. 21-29.
46. Howell, R., Qin, N., Edwards, J. & Durrani, N. Wind tunnel and numerical study of a small vertical axis wind turbine. *Renewable Energy* Vol 35. (2009), pp. 412-422.
47. Islam, M., Ting, D. S. K. & Fartaj, A. Aerodynamic models for Darrieus-type straight-bladed vertical axis wind turbines. *Renewable and Sustainable Energy Reviews* Vol 12. (2006), pp. 1087-1109.
48. Islam, M., Ting, D. S. & Fartaj, A. Desirable Airfoil Features for Smaller-Capacity Straight-Bladed VAWT. *Wind Engineering* Vol 31. (2007), pp. 165-196.
49. Jameson, A. Aerodynamic Design via Control Theory. *Journal of Scientific Computing* Vol 3. (1988), pp. 233-260.
50. Jawahar, H. K., Ai, Q. & Azarpeyvand, M. Experimental and numerical investigation of aerodynamic performance for airfoils with morphed trailing edges. *Renewable Energy* Vol 127. (2018), pp. 355-367.
51. Jin, X., Zhao, G., Gao, K. & Ju, W. Darrieus vertical axis wind turbine: Basic research methods. *Renewable and Sustainable Energy Reviews* Vol 42. (2014), pp. 212-225.
52. Kear, M., Evans, B., Ellis, R. & Rolland, S. Computational aerodynamic optimisation of vertical axis wind turbine blades. *Applied Mathematical Modelling* Vol 40. (2016), pp. 1038-1051.

53. Kim, Y. & Xie, Z-T. Modelling the effect of freestream turbulence on dynamic stall of wind turbine blades. *Computers & Fluids* Vol 129. (2016), pp. 53-66.
54. Klimas, P. C. & Worstell, M. H. Effects of Blade Preset Pitch/ Offset on Curved-Blade Darrieus Vertical Axis Wind Turbine Performance, SAND-81-1762. Sandia Corporation. (1981).
55. Kramer, M. Increase in the maximum lift of an airplane wing due to a sudden increase in its effective angle of attack resulting from a gust. Technical Memorandum No 678. National Advisory Committee for Aeronautics (1932).
56. Kreyszig, E. *Advanced Engineering Mathematics*, 9th Ed. Wiley. (2006).
57. Lachenal, X., Daynes, S. & Weaver, P. M. Review of morphing concepts and materials for wind turbine blade applications. *Wind Energy* Vol 16. (2013), pp. 283-307.
58. Laneville, A. & Vittecoq, P. Dynamic Stall: The Case of the Vertical Axis Wind Turbine. *Journal of Solar Energy Engineering* Vol 108. (1986), pp. 140-145.
59. Lanzafame, R., Mauro, S. & Messina, M. 2D CFD Modelling of H-Darrieus Wind Turbines Using a Transition Turbulence Model. *Energy Procedia* Vol 45. (2013), pp. 131-140.
60. Lee, T. & Gerontakos, P. Investigation of flow over an oscillating airfoil. *Fluid Mechanics* Vol 512. (2004), pp. 313-341.
61. Lee, B. J. & Liou, M. Unsteady Adjoint Approach for Design Optimization of Flapping Airfoils. *AIAA* Vol 50. (2012), pp. 2460-2475.
62. Lei, J. & He, J. Adjoint-Based Aerodynamic Shape Optimization for Low Reynolds Number Airfoils. *Journal of Fluids Engineering* Vol 138. (2016), pp. (021401)1-6.
63. Le Moigne, A. PhD Thesis: A Discrete Navier-Stokes Adjoint Method for Aerodynamic Optimisation of Blended Wing-Body Configurations. University of Sheffield. (2003).
64. Li, Q., Maeda, T., Kamada, Y., Murata, J., Kawabata, T., Shimizu, K., Ogasawara, T., Nakai, A & Kasuya, T. *et al.* Wind tunnel and numerical study of a straight-bladed vertical axis wind turbine in three-dimensional analysis (Part I: For predicting aerodynamic loads and performance). *Energy* Vol 106. (2016), pp. 443-452.
65. Li, H., Song, L., Li, Y. & Feng, Z. 2D Viscous Aerodynamic Shape Design Optimization for Turbine Blades Based on Adjoint Method. *Journal of Turbomachinery* Vol 133. (2011), pp. (031014)1-8.

66. Li, C., Xiao, Y., Xu, Y., Peng, Y., Hu, G. & Zhu, S. Optimization of blade pitch in H-rotor vertical axis wind turbines through computational fluid dynamics simulations. *Applied Energy* Vol 212. (2018), pp. 1107-1125.
67. Liang, C. & Li, H. Effects of optimized airfoil on vertical axis wind turbine aerodynamic performance. *Brazilian Society of Mechanical Sciences and Engineering*. (2018), pp. 1-9.
68. Lighthill, M. J. A New Method of Two-dimensional Aerodynamic Design. *Aeronautical Research Council Reports and Memoranda*, No 2112. (1945).
69. Luo, J., Xiong, J., Liu, F. & McBean, I. Three-Dimensional Aerodynamic Design Optimization of a Turbine Blade by Using an Adjoint Method. *Journal of Turbomachinery* Vol 133. (2011), pp. (011026)1-11.
70. MacPhee, D. & Beyene, A. Recent Advances In Rotor Design of Vertical Axis Wind Turbines. *Wind Engineering* Vol 36. (2012), pp. 647-666.
71. McAlister, K. W., Carr, L. W. & McCroskey, W. J. Dynamic stall experiments on the NACA 0012 airfoil. *NASA, Technical Paper* 1100. (1978).
72. Menter, F. R. Two-Equation Eddy-Viscosity Turbulence Models for Engineering Applications. *AIAA No. 8* Vol. 32. (1994), pp. 1598-1605.
73. Mohamed, O. S., Ibrahim, A. A., Etman, A. K., Abdelfatah, A. A., & Elbaz, A. M. R. Numerical Investigation of Darrieus wind turbine with slotted airfoil blades. *Energy Conservation and Management: X* Vol 5. (2020), 100026.
74. Nadarajah, S. K. & Jameson, A. A Comparison of the Continuous and Discrete Adjoint Approach to Automatic Aerodynamic Optimization *AIAA* Vol 667. (1999), pp. 1-20.
75. NASA (Christopher Rumsey). Turbulence Modelling Resource. Langley Research Centre. Available online at: <https://turbmodels.larc.nasa.gov/sst.html>. (2021).
76. Ning, S. A., Damiani, R. & Moriarty, P. J. Objectives and Constraints for Wind Turbine Optimization. *Journal of Solar Energy Engineering* Vol 136. (2014), pp. (041010)1-12.
77. Orman, E. & Durmus, G. Comparison of shape optimization techniques coupled with genetic algorithm for a wind turbine airfoil - *IEEE Conference Publication*. (2016), pp. 1-7.
78. Peter, J. E. V. & Dwight, R. P. Numerical sensitivity analysis for aerodynamic optimization: A survey of approaches. *Computers & Fluids* Vol 39. (2010), pp. 373-391.

79. Priegue, L. & Stoesser, T. The influence of blade roughness on the performance of a vertical axis tidal turbine. *International Journal of Marine Energy* Vol 17. (2017), pp. 136-146.
80. Ram, K. R., Lal, S. & Ahmed, M. R. Low Reynolds number airfoil optimization for wind turbine applications using genetic algorithm. *Journal of Renewable and Sustainable Energy* Vol 5. (2009), pp. (052007)1-15.
81. Ram, K. R., Lal, S. & Ahmed, M. R. Airfoil Optimization for Small Wind Turbines Using Multi Objective Genetic Algorithm. ASME, IMECE2012-88788. (2012), pp. 1-9.
82. Rezaeiha, A., Kalkman, I., Blocken, B. CFD simulation of a vertical axis wind turbine operating at a moderate tip speed ratio: Guidelines for minimum domain size and azimuthal increment. *Renewable Energy* Vol 107. (2017), pp. 373-385.
83. Rezaeiha, A., Kalkman, I., Blocken, B. Effect of pitch angle on power performance and aerodynamics of a vertical axis wind turbine. *Renewable Energy* Vol 197. (2017), pp. 132-150.
84. Rezaeiha, A., Montazeri, H., Blocken, B. Towards accurate CFD simulations of vertical axis wind turbines at different tip speed ratios and solidities: Guidelines for azimuthal increment, domain size and convergence. *Energy Conversion and Management* Vol 156. (2018), pp. 301-316.
85. Rezaeiha, A., Montazeri, H., Blocken, B. Characterization of aerodynamic performance of vertical axis wind turbines: Impact of operational parameters. *Energy Conversion and Management* Vol 169. (2018), pp. 45-77.
86. Rezaeiha, A., Montazeri, H., Blocken, B. On the accuracy of turbulence models for CFD simulations of vertical axis wind turbines. *Energy* Vol. 180. (2019), pp. 838–857.
87. Rezaeiha, A., Montazeri, H. & Blocken, B. Active flow control for power enhancement of vertical axis wind turbines: Leading-edge slot suction. *Energy* Vol 189. (2019), 116131.
88. Ritlop, R. M. & Nadarajah, S. K. Thesis, Master of Engineering: Toward the aerodynamic shape optimization of wind turbine profiles. McGill University (2009).
89. Rosado Hau, N. PhD Thesis: Unsteady Aerodynamics of Vertical Axis Wind Turbines. University of Sheffield. (2021).
90. Scheurich, F. & Brown, R. E. Modelling the aerodynamics of vertical-axis wind turbines in unsteady wind conditions. *Wind Energy* Vol 16. (2012), pp. 91-107.
91. Sharma, A. & Visbal, M. Numerical investigation of the effect of airfoil thickness on onset of dynamics stall. *Journal of Fluid Mechanics* Vol 870. (2019), pp 870-900.

92. Shires, A. Design optimisation of an offshore vertical axis wind turbine. *ICE Energy* Vol 166. (2013), pp. 7-18.
93. Standish, K. J. & van Dam, C. P. v. Aerodynamic Analysis of Blunt Trailing Edge Airfoils. *Journal of Solar Energy Engineering* Vol 125. (2003), pp. 479-487.
94. Sutherland, H. J., Berg, D. E. & Ashwill, T. D. A Retrospective of VAWT Technology, SAND-2012-0304, Sandia Corporation. (2012).
95. Tjiu, W., Marnoto, T., Mat, T., Ruslan, M. H. & Sopian, K. Darrieus vertical axis wind turbine for power generation I: Assessment of Darrieus VAWT configurations. *Renewable Energy* Vol 75. (2014), pp. 50-67.
96. Trivellato, F. & Castelli, M. R. On the Courant-Friedrichs-Lewy criterion of rotating grids in 2D vertical-axis wind turbine analysis. *Renewable Energy* Vol 62. (2014), pp. 53-62.
97. Tzanakis, A. Masters Thesis: Duct optimization using CFD software ANSYS Fluent Adjoint Solver. Chalmers University of Technology. (2014).
98. Untaroiu, A., Wood, H. G., Allaire, P. E. & Ribando, R. J. Investigation of Self-Starting Capability of Vertical Axis Wind Turbines Using a Computational Fluid Dynamics Approach. *Journal of Solar Energy Engineering* Vol 133. (2010), pp. (041010)1-8.
99. Vassberg, J. C., Gopinath, A. K. & Jameson, A. Revisiting the Vertical-Axis Wind-Turbine Design using Advanced Computational Fluid Dynamics. *AIAA ASM* Vol 43. (2005), pp. 1-23.
100. Vavalle, A. & Qin, N. Iterative Response Surface Based Optimization Scheme for Transonic Airfoil Design. *Journal of Aircraft* Vol 4. (2007), pp. 365-376.
101. Velasco, D., Lopez Mejia, O. & Lain, S. Numerical simulations of active flow control with synthetic jets in a Darrieus turbine. *Renewable Energy* Vol 113. (2017), pp. 129-140.
102. Versteeg, H. K. & Malalasekera, W. An Introduction to Computational Fluid Dynamics – The Finite Volume Method, 2nd Ed. Pearson. (2007).
103. Walther, B. & Nadarajah, S. An Adjoint-Based Multi-Point Optimization Method for Robust Turbomachinery Design. *ASME Turbo Expo 2015, GT2015-44142*. (2015), pp. 1-12.
104. Wang, S., Ingham, D. B., Ma, L., Pourkashanian, M. & Tao, Z. Numerical investigations on dynamic stall of low Reynolds number flow around oscillating airfoils. *Computers & Fluids* Vol 39. (2010), pp. 1529-1541.

105. Wang, S., Ingham, D. B., Ma, L., Pourkashanian, M. & Tao, Z. Turbulence Modelling of Deep Dynamics Stall at Relatively Low Reynolds Number. *Journal of Fluids and Structures* Vol 33. (2012), pp. 191-209.
106. Wang, Z., & Zhuang, M. Leading-edge serrations for performance improvement on a vertical-axis wind turbine at low tip-speed-ratios. *Applied Energy* Vol 208. (2017), pp 1184-1197.
107. Wernert, P., Geissler, W., Raffel, M. & Kompenhans, J. Experimental and numerical investigations of dynamic stall on a pitching airfoil. *AIAA No. 5* Vol 34. (1996), pp. 982-989.
108. Van-Wingerden, J. W., Hulskamp, A. W., Barlas, T., Marrant, B., Van-Kuik, G. A. M., Molenaar, D. P & Verhaegen, M. On the proof of concept of a 'Smart' wind turbine rotor blade for load alleviation.
109. Worasinchai, S., Ingram, G. L. & Dominy, R. G. The Physics of H-Darrieus Turbines Self-Starting Capability: Flapping-Wing Perspective. *ASME: Turbo Expo 2012, GT2012-69075*. (2012), pp. 1-10.
110. Worasinchai, S., Ingram, G. L. & Dominy, R. G. The Physics of H-Darrieus Turbine Starting Behavior. *Journal of Engineering for Gas Turbines and Power* Vol 138. (2016), pp. (062605)1-11.
111. Xisto, C. M., Pascoa, J. C. & Leger, J. A. Wind energy production using an optimized variable pitch vertical axis rotor. *Proceedings of the ASME 2014 International Mechanical Engineering Congress and Exposition*. (2014).
112. Yan, Y., Avital, E., Williams, J. & Cui, J. Performance Improvements for a vertical axis wind turbine by means of Gurney flaps. *Journal of Fluids Engineering* Vol 142. (2020), 021205-1.
113. Zhang, L., Liang, Y., Li, E., Zhang, S. & Guo, J. Vertical axis wind turbine with individual active blade pitch control. *Asia-Pacific Power and Energy Engineering Conference, APPEEC*. (2012), pp. 1-4.
114. Zhang, L., Liang, Y., Liu, X. & Guo, J. Effect of blade pitch angle on aerodynamic performance of straight-bladed vertical axis wind turbine | Springer for Research & Development. *Central South University* Vol 21. (2014), pp. 1417-1427.
115. Zhang, M., Tan, B. & Xu, J. Parameter study of sizing and placement of deformable trailing flap on blade fatigue load reduction. *Renewable Energy* Vol 77, (2015), pp. 217-226.

116. Zhao, Z., Wang, R., Shen, W., Wang, T., Xu, B., Zheng, Y. & Qian, S. Variable pitch approach for performance Improving of straight-bladed VAWT at rated tip speed ratio. *Applied Sciences* Vol 8, 957. (2018), 103390.
117. Zhuang, C., Yang, G., Zhu, Y. & Hu, D. Effect of morphed trailing-edge flap on aerodynamic load control for a wind turbine section. *Renewable Energy* Vol 148. (2020), pp. 954-974.
118. Zhiqiang, Y., Minghui, Y., Xiaoyang, C., Zaiyu, C. & Yun, Z. Multi-AOA optimization of variable-speed wind turbine airfoils. *IEEE*. (2016), pp. 1674-1677.
119. Zhu, J., Huang, H. & Shen, H. Self-starting aerodynamics analysis of vertical axis wind turbine. *Advances in Mechanical Engineering* Vol 7. (2015), pp. 1-12.
120. Zhu, J., Jiang, L. & Zhao, H. Effect of wind fluctuating on self-starting aerodynamics characteristics of VAWT. *Central South University* Vol 23. (2016), pp. 2075-2082.

10 APPENDIX A – SINGLE-BLADE CFD MODEL VALIDATION

This appendix presents the validation studies that supplement the production of the Single-Blade CFD model, as described in Section 5.1.3.

Throughout the following studies, the settings described in Section 5.1.2 were held constant while variations were made to test solution sensitivity. An initial study was made to decide on the number of oscillation cycles necessary to achieve a periodically stable solution. From the 3rd cycle an error of less than 0.2% (average C_T) was found for a baseline model, and so the following studies were conducted considering the solution at the 3rd cycle. Such a rapid periodic convergence is observed because the Single-Blade model has no wake interactions and operates with almost no separation effects.

- Mesh fineness
- Time step size
- Courant number
- Domain size

A ‘mesh-independent’ and ‘time-step-independent’ solution changes by a negligible amount when the resolution of mesh and/or time-step is made finer. Thus, a point is reached where the improved accuracy of the solution is not worth the additional computational cost of the additional mesh/time-step fineness. This study details how a sufficiently accurate, mesh and time-step independent model was reached

The details of each mesh are presented in Table 10-1. The range of time-steps tested are represented in Table 10-2.

Throughout the figures which will follow, a C_t vs AoA curve is used to compare data for each case.

Table 10-1 - Range of Meshes Used in Validation Study

Mesh I.D.	M1	M2	M3	M4	M5
Total cells (k = 1000 elements)	54 k	102 k	150 k	245 k	345 k
Boundary zone Normal Resolution	25	35	50	70	90
Subdomain Normal Resolution	100	165	175	250	300
Aerofoil Chordwise Resolution*	310	380	475	570	680
Far-field Normal Resolution	50	70	90	110	120

*Total number of elements around the whole aerofoil wall.

Table 10-2 – Range of Time-Steps in Validation Study

Time Step I.D.	T1	T2	T3	T4	T5	T6	T7
Number of steps per cycle	100	200	400	800	1200	1600	2000

For the mesh and time step studies, the domain size was held constant. A constant Courant number of 40 was chosen according to a preliminary CFL study which is not recorded here (the CFL study is presented after the mesh study/time-step study).

Figure 10-1 compares results for the range of time-steps, with a constant mesh (M3) of medium fineness.




Figure 10-1 – C_T vs AoA for a range of Time-Steps

The coarse/medium time steps (T1, T2, T3) show progressive improvement towards the time-step converged solution. The medium-fine time-steps (T4, T5) showed near perfect agreement with each other and are therefore represented as a single line. The two finest time-steps were in near perfect agreement and are also represented as a single line. T4 and T5 have good agreement with the finest time-steps (T6, T7). The conclusion of this when considering the balance between accuracy and CPU cost, is that time-step T4 (800 time-steps/rev) is suitable.

Figure 10-2 compares results for the range of meshes, for the chosen time-step (T4).

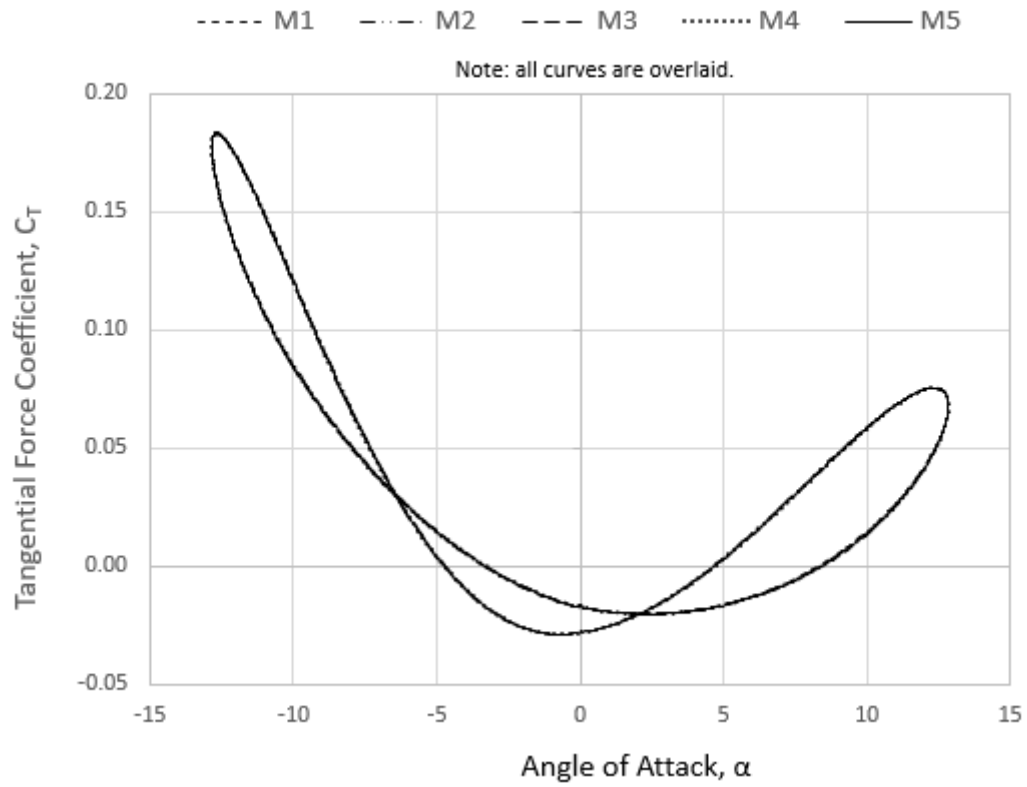


Figure 10-2 - C_T vs AoA for a range of Meshes




Figure 10-3 - C_T vs AoA for a range of Courant Numbers

The coarsest mesh and the finest mesh have near perfect agreement with each other along with the other meshes in between. The finest mesh has around 7 times the number of elements than the coarsest but offers no improvement to accuracy. The conclusion of this is that the coarsest mesh (M1) is suitable.

With the mesh and time-step selected (M1, T4), additional tests were carried out to test the sensitivity to the target CFL used in the Coupled solver. Figure 10-3 compares results for the range of CFLs tested. The largest and smallest CFL have near perfect agreement with each other, along with the other meshes in between. The conclusion of this is that the largest CFL (200) is suitable.

This CFL and time-step are larger than the values recommended in some of the literature for this type of problem. The reason for this is that the literature recommendations tend to consider low TSR cases for full VAWT models, where far more prominent unsteadiness including dynamic stall features are present. The choices here are therefore only valid for this case, and higher resolution CFL/mesh/time-step may indeed be required for lower TSR cases.

A range of domain sizes were tested to validate the model in terms of domain size convergence. The results for the various domains are presented in Figure 10-4. For each of the 5 models, the size of the rotating subdomain and the far field zone were modified as shown in Table 10-3. The size of the boundary zone subdomain and the number of nodes around the aerofoil surface were held constant. The mesh densities of the rotating and far-field zones were held approximately constant while their dimensions were changed.

The curves for the 5 domains tested are close together demonstrating good domain size convergence. D1 is a slight outlier over most of the cycle. D2 shows very minor disagreement with the rest especially at the downwind peak (at $+12^\circ$). The remaining models (D3, D4, D5) share very good agreement. D3 and D4 match closely in the downwind part when the AoA is approaching its maximum (0° to $+12^\circ$), with D5 being slightly offset. However, D4 and D5 match closely in the downwind part when the AoA is reducing from its maximum ($+12^\circ$ to 0°). Observing the minor disagreement between D3, D4 and D5, and considering that D4 and D5 have 19% and 37% more cells compared to D3 respectively, it is judged that D3 is the most appropriate domain size to use.

Table 10-3 – Range of Domains Used in Validation Study

Domain I.D.	D1	D2	D3	D4	D5
Total cells (k = 1000 elements)	33 k	43 k	54 k	64 k	73 k
Diameter of rotating subdomain*	2	4	6	8	10
Diameter of far-field zone*	15	25	40	65	100

*Diameters are presented as a multiple of the aerofoil chord length (x 0.06m).




Figure 10-4 - C_T vs AoA for a range of Domain Sizes

11 APPENDIX B – VAWT CFD MODEL VALIDATION

This appendix presents the validation studies which supplement the production of VAWT CFD models, as described in Section 5.1.4 and Section 6.1.

Throughout the following studies, the settings described in Section 6.1 were held constant while variations were made to test solution sensitivity. An initial study was made to decide the number of revolutions necessary to achieve a periodically stable solution. From the 10th revolution an error of less than 1% (average upwind C_M) was found for a baseline model, and so the following studies were conducted using the solution at the 10th revolution.

- Mesh fineness
- Time step size
- Courant number

With the domain described in Section 5.1.4, Rezaeiha *et al.* (2017) used a mesh with around 400,000 cells. The present mesh independence study therefore uses a range of cases with around 400,000 cells as a medium/fine model (see Table 11-1). To study the range of meshes, a constant time step of 800 steps/rev was chosen since this was recommended by the Single-Blade validation studies (Appendix A).

Table 11-1 – Range of Meshes Used in Validation Study

Mesh I.D.	M1	M2	M3	M4	M5
Total cells (k = 1000 elements)	99 k	174 k	259 k	461 k	689 k
Boundary zone Normal Resolution	25	35	50	70	90
Subdomain Normal Resolution	80	130	140	200	240
Aerofoil Chordwise Resolution*	310	380	475	570	680
Element size - Rotating Zone (m)	0.018	0.013	0.009	0.055	0.005
Element size - Exterior (m)	0.2	0.14	0.1	0.07	0.05

*Total number of elements around the whole aerofoil wall

Figure 11-1 shows a graph of the various meshes for the constant time step.




Figure 11-1 - C_T vs AoA for a range of Meshes

All meshes have relatively near perfect agreement, except for a small area at the downwind (max +ve AoA). Here, the coarsest meshes (M1, M2) have close agreement with each other but not with the finer meshes. Although it is difficult to see from the figure, M3, M4 and M5 are all overlaid demonstrating mesh convergence. The conclusion of this is that the coarsest mesh in this converged group (M3) is suitable. The range of time-steps studied are presented in Table 11-2. Each of them was run with the chosen mesh (M3).

Table 11-2 – Range of Time-Steps Used in Validation Study

Time Step I.D.	T1	T2	T3	T4	T5	T6	T7
Number of steps per revolution	100	200	400	800	1200	1600	2000

Figure 11-2 shows a graph of the various time steps for the chosen mesh.




Figure 11-2 - C_T vs AoA for a range of Time-Steps




Figure 11-3 - C_T vs AoA for a range of Courant Numbers

The coarse time steps (T1, T2) are outliers from the finer time steps. T3 has some minor disagreement with the finer time-steps in the downwind area (+ve AoA). T4, T5, T6 and T7 show near perfect agreement with each other and are therefore represented as a single line. The conclusion of this is that time-step T4 (800 time-steps/rev) is suitable.

With the mesh and time-step selected (M3, T4), additional tests were carried out to test the sensitivity to the target CFL used in the Coupled solver. Figure 11-3 compares results for the range of CFL (for M1, T4).

The curves for the three coarsest CFL's are overlaid. The curve for CFL=5 has a slight variation from the coarser curves. The curve for CFL=10 is in between the curves shown for CFL=5 and CFL=40/100/200, but has been omitted for brevity. There is not a small of difference between run times between the coarsest and finest CFL. Since the smaller CFL generally corresponds to a more accurate solution (see Section 5.1.2), CFL=5 is deemed appropriate.

Now that the VAWT model details have been confirmed, the results must be compared to some reference data in order to provide validation. Figure 6-2 (Section 6.1) shows how the VAWT model data agrees with the CFD of the reference paper (Rezaeiha *et al.* 2017). The curves do not overlay perfectly, but it is judged that overall there is good agreement providing confidence in the accuracy of the model.

12 APPENDIX C – SINGLE-BLADE OSCILLATION UDF

The UDF includes two parts; the AoA oscillation profile, and the relative velocity (V_{rel}) profile.

AoA Profile

ANSYS Fluent requires the UDF to specify the rotation profile as a continuous function of rotational velocity, rather than as a continuous time function of α (AoA) as per Equation (1.3). This means that the equation used in the UDF is the derivative of Equation (1.3). This derivation of Equation (1.3) is coded as follows in the UDF:

```
omega =w*(1+TSR*cos(w*time))/(1+2*TSR*cos(w*time)+(TSR*TSR));
```

The rotational frequency of the VAWT, ω , is used as the pitching frequency of oscillating aerofoil ω_p . This appears as 'w' in the UDF code.

The UDF applies a knockdown factor of 0.5 to this value of omega during the downwind, which is explained in Section 5.1.

It is important to note that this causes an initial counter-clock-wise rotation of the blade. This is in accordance with the sign convention for AoA defined in Figure 1-4 (Section 1.2). In this case, the initial counter-clockwise motion of the blade corresponds to the geometric AoA being experienced as a VAWT blade moves into the upwind part of the cycle, from a starting point of 0 degrees (azimuthal angle). For the present work the baseline blade starts with a fixing angle of zero, but if a VAWT is used where the fixing angle is non-zero then this must be implemented in the model before the UDF is applied.

Vrel Profile

To prescribe the relative velocity profile Equation 1.3 is implemented directly.

Complete UDF

The complete UDF code incorporating AoA profile and V_{rel} profile is as follows:

```
#include "udf.h"
#define TSR 4.5 /* constants , tip speed ratio*/
#define velocity_free 9.3 /*free flow velocity*/
#define thetamax 0.22393
#define w 83.7

DEFINE_TRANSIENT_PROFILE(angular_velocity,time)
{
    real omega, theta;
    theta = (w*time) - 6.283185307*floor((w*time)/6.283185307);
    if (theta > 3.141592654 && theta < 6.283185307)
        omega = 0.5*w*(1+TSR*cos(w*time))/(1+2*TSR*cos(w*time)+(TSR*TSR));
    else
        omega = w*(1+TSR*cos(w*time))/(1+2*TSR*cos(w*time)+(TSR*TSR));
    return omega;
}

DEFINE_PROFILE(unsteady_velocity, thread, position)
{
    face_t f;
    real t = CURRENT_TIME;
    real theta = w*t;
    real alpha = atan(sin(theta)/(TSR+cos(theta)));
    begin_f_loop(f, thread)
    {
        F_PROFILE(f, thread, position) =
        velocity_free*sqrt(1+pow(TSR,2)+2*TSR*cos(theta));
    }
    end_f_loop(f, thread)
}
```

13 APPENDIX D – FURTHER DETAIL ON ALGORITHM KEY NOTES (SINGLE-BLADE METHOD)

13.1 Mesh Morphing (Constraints, Negative Cell Volumes and Mesh Distortion)

See Section 5.3.4.1 for context relating to the following discussion.

Enforcing 'Strict' constraint conditions means that no amount of penetration of the blade wall through the constraint boundary will occur (see Figure 5-11(a))Figure 5-11. Not using Strict conditions alleviates this issue but permits some non-conformance at the constraint boundary (i.e., the blade geometry may partially enter the boundary), this approach is recommended.

Figure 13-1 shows an example of trailing edge mesh distortion if 'Strict' constraints are imposed when the mesh morpher is used during the optimisation process. It can be observed that the mesh morpher attempts to move trailing edge mesh across the constraint boundary (shaded yellow), The first row of cells is highly distorted as the mesh morpher tries to satisfy both the desired geometry change and the constraint. The wall nodes of these cells are prevented from moving across the constraint boundary, and so cannot move consistently with their counterpart nodes, elongating the cells.




Figure 13-1 - Mesh Distortion at the Trailing Edge

Along with the choice of Strict/Un-strict constraints, other settings within the ANSYS Fluent Adjoint 'design tool' can contribute to the mesh morphing/distortion behavior:

- Number of control points. The mesh morpher uses 'control points' as a means of discretising the mesh inside of the mesh morphing region Figure 5-11(b). The user can select up to 200 control points in each direction (X and Y). These are equally spaced over the length/height of the specified morphing region shown. The control points themselves cannot be displayed graphically in ANSYS 18.2. Studies linked to the above discussion on strict/un-conditions found that 100 control points in both X and Y was an appropriate number to achieve stable behaviour during mesh morphing.

Using un-strict constraints allows some non-conformance to the constraint. This non-conformance can be limited by using the appropriate number of control points. Figure 13-2 shows (left) the control points at the trailing edge before morphing, and (right) then after morphing. A small area of the TE penetrates the constraint area (grey), but the mesh quality is preserved.




Figure 13-2 - Non-strict Constraints & Partial Non-conformance

- Region type (coordinate system). Cartesian or Polar coordinates can be chosen; this affects the region shape (rectangular vs circular) and the axes in which the control points are distributed. It was found in preliminary studies that using the polar system tended to result in distorted blade shapes for the present problem. This is due to the distribution of control points within the chosen coordinate system. In addition, it was found to be easier to achieve a more even distribution, and more stable mesh morphing behaviour using the Cartesian system which is adopted throughout this work.

- Free Form Scaling (FFS). This determines the limit of deformation applied to each control point by the mesh morpher. It is thereby a measure of aggression used in the morphing process. The FFS value can be user defined; a value of FFS = 1.0 was used for the current work.
- FFS scheme. This alters the way that the sensitivity data is translated into a geometry change. The 'Control-point spacing' option means that the geometry deformations are based on the control point spacing, which is a function of the morphing region size and number of control points. The 'Objective reference change' option means that geometry deformations are based on the value used as the target value (objective), mentioned in 5.3.2. 'Objective Reference Change' is chosen for the present work as it allows changes to be made to the number of control points without effecting aggression of the morphing process.

This selection of settings was reached to allow robust function of the optimisation process. This method allows some minor inflation to the blade chord length as a concession required to alleviate the mesh distortion at the blade wall/constraint interface. CFD error messages associated with this issue if not mitigated, may be 'Negative Cell Volumes', but with no further information of what is causing the issue. If using a High Performance Computing (HPC) cluster, error messages such as 'SIGSEGV' may be encountered (relating to memory segmentation faults) which also do not elaborate on the root cause of the issue.

14 APPENDIX E – GOVERNING EQUATIONS AND TURBULENCE MODELLING

The Navier-Stokes equations which govern the flow of Newtonian fluids can be written in various forms but always consist of the continuity equation and momentum equations.

The continuity equation can be written as:

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} = 0 \quad (\text{Eqn. A.1})$$

The momentum equations describe the movement of a fluid parcel through space and time. Equations for momentum in the x, y & z directions can be written individually – the following form assumes an incompressible Newtonian fluid, where body forces (gravity) are neglected:

$$\frac{\partial(\rho u)}{\partial t} + \frac{\partial(\rho u^2)}{\partial x} + \frac{\partial(\rho uv)}{\partial y} + \frac{\partial(\rho uw)}{\partial z} = -\frac{\partial P}{\partial x} + \mu \left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2} \right) \quad (\text{Eqn. A.2-i})$$

$$\frac{\partial(\rho v)}{\partial t} + \frac{\partial(\rho uv)}{\partial x} + \frac{\partial(\rho v^2)}{\partial y} + \frac{\partial(\rho vw)}{\partial z} = -\frac{\partial P}{\partial y} + \mu \left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} + \frac{\partial^2 v}{\partial z^2} \right) \quad (\text{Eqn. A.2-ii})$$

$$\frac{\partial(\rho w)}{\partial t} + \frac{\partial(\rho uw)}{\partial x} + \frac{\partial(\rho vw)}{\partial y} + \frac{\partial(\rho w^2)}{\partial z} = -\frac{\partial P}{\partial z} + \mu \left(\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2} + \frac{\partial^2 w}{\partial z^2} \right) \quad (\text{Eqn. A.2-iii})$$

Total acceleration (per unit volume)

Total force (per unit volume)

These equations are analogous to Newton's law, but instead of an expression in the form $F=MA$ (force=mass x acceleration), we consider forces per unit volume $F/V=(MA)/V$, leading to the form $\rho A=F/V$ suitable for application to a continuum (a fluid).

14.1 RANS Equations

The Navier-Stokes equations can be presented in a way that allows analysis of the influence of turbulence on the flow. This approach, called Reynolds decomposition is applied to the velocity (u) in the x-momentum equation, rewriting the instantaneous velocity as the sum of the mean (U) and fluctuating velocity components (u') due to turbulence:

$$u = U + u' \quad (\text{Eqn. A.3})$$

If time averaging is applied following this substitution, the resulting x-momentum equation presents as (where the bar operator $\bar{\phi}$ indicates a time averaged quantity):

$$\rho \frac{\partial U}{\partial t} + U \frac{\partial U}{\partial x} + V \frac{\partial U}{\partial y} + W \frac{\partial U}{\partial z} + \frac{\partial \overline{u'u'}}{\partial x} + \frac{\partial \overline{u'v'}}{\partial y} + \frac{\partial \overline{u'w'}}{\partial z} = -\frac{\partial P}{\partial x} + \mu \left(\frac{\partial^2 U}{\partial x^2} + \frac{\partial^2 U}{\partial y^2} + \frac{\partial^2 U}{\partial z^2} \right) \quad (\text{Eqn. A.4})$$

This is now the Reynolds Averaged Navier Stokes (RANS) momentum equation in x. The y and z equations can be derived in a similar way. This equation now contains the $\overline{u'u'}$, $\overline{u'v'}$ and $\overline{u'w'}$ terms which are called the Reynolds stresses. The following form is commonly given in the literature to present the x, y & z directions concisely, for incompressible flows neglecting gravity (El Sakka, 2020):

$$\rho \frac{\partial u_i}{\partial t} + \rho u_j \frac{\partial u_i}{\partial x_j} = -\frac{\partial P}{\partial x_i} + \mu \frac{\partial}{\partial x_j} \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) + \frac{\partial}{\partial x_j} (-\rho \overline{u'_i u'_j}) \quad (\text{Eqn. A.5})$$

In this form of the RANS equations, the term $-\rho \overline{u'_i u'_j}$ now represents the Reynolds stress tensor (τ_{ij}), and the u_i and x_i terms provide the 3 directional velocities and directions respectively.

These Reynolds stresses represent of the effect that turbulent fluctuations have on momentum transfer in the flow. This is because turbulent eddies convect slow moving fluid to faster moving areas and vice-versa, which leads to a momentum exchange. They include interactions between all the velocity directions because turbulence is a 3D phenomenon (Versteeg & Malalsekera, 2007).

The time averaged continuity equations are then:

$$\frac{\partial U}{\partial x} + \frac{\partial V}{\partial y} + \frac{\partial W}{\partial z} = \frac{\partial \overline{u}}{\partial x} + \frac{\partial \overline{v}}{\partial y} + \frac{\partial \overline{w}}{\partial z} = 0 \quad (\text{Eqn. A.6})$$

The RANS equations are solved numerically by the CFD code. However, to find the value of these additional Reynolds stress terms, some strategy of estimating them is required – this is achieved via turbulence modelling.

14.2 SST k- ω Turbulence Model

The SST k- ω turbulence model (adopted in this work) sits within the family of “eddy viscosity” turbulence models. This means that an artificial quantity called the eddy viscosity (μ_t) is calculated by relating it in some fashion to the mean flow variables. This eddy viscosity provides a measure of how strongly momentum is transferred by the turbulence. Once calculated μ_t is then used to estimate the Reynolds stresses, thereby closing the RANS equations and allowing an approximate solution to be determined by CFD solvers.

The SST k- ω model is one such eddy viscosity method, and its suitability for VAWTs has been described in Section 2.4 (it provides a good balance of accuracy and CPU

cost). The model relates the eddy viscosity to the flows turbulent kinetic energy (k) and the specific dissipation rate (ω) of the turbulent kinetic energy (also called the turbulence frequency). The dissipation rate is the action of the smallest eddies dissipating energy to heat due to viscosity.

The SST k - ω model is a development upon earlier models which exhibit limitations/inaccuracies when predicting separated flows. SST k - ω combines two such models in a way that provides better overall accuracy; the k - ϵ is used away from the wall and the k - ω is used near the wall, with a blending function in between to transition behaviour across the two. By adopting this approach, improved accuracy is achieved near the boundary layer due to the k - ω model, while using k - ϵ away from the wall to avoid high sensitivity to freestream turbulence associated with k - ω . This description of a blended k - ϵ / k - ω model also befits the BST k - ω model, but the SST variant employs an additional correction/limiter function to the turbulent viscosity near the wall. This correction accounts for transport of turbulent shear stress, which is where the model gets its name 'SST' (Menter, 1994).

Since Menter's paper announcing the SST k - ω model, CFD codes have implemented various updates to the original one, for specific purposes. The equations defining the SST k - ω model transport equations for k and ω are as follows (NASA, 2021):

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho u_j k)}{\partial x_j} = P - \beta^* \rho \omega k + \frac{\partial}{\partial x_j} \left[(\mu + \sigma_k \mu_t) \frac{\partial k}{\partial x_j} \right] \quad (\text{Eqn. A.7})$$

$$\frac{\partial(\rho \omega)}{\partial t} + \frac{\partial(\rho u_j \omega)}{\partial x_j} = \frac{\gamma}{v_t} P - \beta^* \rho \omega^2 + \frac{\partial}{\partial x_j} \left[(\mu + \sigma_\omega \mu_t) \frac{\partial \omega}{\partial x_j} \right] + 2(1 - F_2) \frac{\rho \sigma_{\omega 2}}{\omega} \frac{\partial k}{\partial x_j} \frac{\partial \omega}{\partial x_j} \quad (\text{Eqn. A.8})$$

Menter (1994), NASA (2021), or the readers CFD source code manual can be consulted for definitions of the empirical coefficients and damping functions denoted in Equations A.7 & A.8.

Once the transport equations are solved, the SST k - ω model defines the eddy viscosity (μ_t), amongst a coefficient a_1 , the shear strain magnitude S , and a smoothing function F_2 (Menter, 1994);

$$\mu_t = \frac{a_1 \rho k}{\max(a_1 \omega, SF_2)} \quad (\text{Eqn. A.9})$$

The dependence of the Reynolds stresses on the eddy viscosity is defined using the Boussinesq viscosity hypothesis (El Sakka, 2020) where δ_{ij} is the Kronecker delta function:

$$-\rho \overline{u'_i u'_j} = \tau_{ij} = \mu_t \left(\frac{\partial u_i}{\partial x_j} + \frac{\partial u_j}{\partial x_i} \right) - \frac{2}{3} \rho k \delta_{ij} \quad (\text{Eqn. A.10})$$

With the Reynolds stresses thereby determined, a solution to the RANS equations can be reached by the CFD solver.