High-Fidelity Aerodynamic Shape Optimization for Natural Laminar Flow

by

Ramy Rashad

A thesis submitted in conformity with the requirements
for the degree of Doctor of Philosophy
Graduate Department of Aerospace Science and Engineering
University of Toronto

Abstract

High-Fidelity Aerodynamic Shape Optimization for Natural Laminar Flow

Ramy Rashad
<ra.rashad@gmail.com>

Doctor of Philosophy
Graduate Department of Aerospace Science and Engineering
University of Toronto
2016

To ensure the long-term sustainability of aviation, serious effort is underway to mitigate the escalating economic, environmental, and social concerns of the industry. Significant improvement to the energy efficiency of air transportation is required through the research and development of advanced and unconventional airframe and engine technologies.

In the quest to reduce airframe drag, this thesis is concerned with the development and demonstration of an effective design tool for improving the aerodynamic efficiency of subsonic and transonic airfoils. The objective is to advance the state-of-the-art in high-fidelity aerodynamic shape optimization by incorporating and exploiting the phenomenon of laminar-turbulent transition in an efficient manner. A framework for the design and optimization of Natural Laminar Flow (NLF) airfoils is developed and demonstrated with transition prediction capable of accounting for the effects of Reynolds number, freestream turbulence intensity, Mach number, and pressure gradients.

First, a two-dimensional Reynolds-averaged Navier-Stokes (RANS) flow solver has been extended to incorporate an iterative laminar-turbulent transition prediction methodology. The natural transition locations due to Tollmien-Schlichting instabilities are predicted using the simplified e^N envelope method of Drela and Giles or, alternatively, the compressible form of the Arnal-Habiballah-Delcourt criterion. The boundary-layer properties are obtained

directly from the Navier-Stokes flow solution, and the transition to turbulent flow is modeled using an intermittency function in conjunction with the Spalart-Allmaras turbulence model.

The RANS solver is subsequently employed in a gradient-based sequential quadratic programming shape optimization framework. The laminar-turbulent transition criteria are tightly coupled into the objective and gradient evaluations. The gradients are obtained using a new augmented discrete-adjoint formulation for non-local transition criteria. Using the e^N transition criterion, the proposed framework is applied to the single and multipoint optimization of subsonic and transonic airfoils, leading to robust NLF designs. The aerodynamic design requirements over a range of cruise flight conditions are cast into a multipoint optimization problem through a composite objective defined using a weighted integral of the operating points. To study and quantify off-design performance, a Pareto front is formed using a weighted objective combining free-transition and fully-turbulent operating conditions. Next we examine the sensitivity of NLF design to the freestream disturbance environment, highlighting the on- and off-design performance at different critical N-factors. Finally, we propose and demonstrate a technique to enable the design of airfoils with robust performance over a range of critical N-factors.

TO MY PARENTS

for making me believe that anything is possible
and for making everything possible

Acknowledgements

I could not have worked with a better doctoral supervisor. I am very grateful to Professor David Zingg for his guidance and support. His vast technical knowledge and creative insight have been crucial to the completion of this thesis and have contributed significantly to my own theoretical understanding of aerodynamics and numerical methods. I am most grateful to Professor Zingg for exposing me to his contagious endeavour towards sustainable aviation and his commitment to excellence in research and teaching. His professionalism and integrity make him a great leader and an excellent role model. It was an honour and pleasure to work with him.

I would like to thank the members of my doctoral examination committee. Professors Clinton Groth and Philippe Lavoie have helped steer this work through the many discussions we have had over the years. For their involvement in the final oral examination, I would like to thank Professors Markus Bussman and Siva Nadarajah. In particular, I would like to thank Professor Nadarajah for his external examination and valuable appraisal. It was a pleasure discussing my work with the entire committee.

I have had many friends at the aerospace institute who have supported and influenced this work. In particular, I would like to thank Howard Buckley, David Del Rey Fernandez, Hugo Gagnon, John Gatsis, David Brown, David Boom, and Thomas Reist for their friendship, collaboration, insight, and all the coffee and lunch breaks that have helped me survive the long arduous hours. Indeed, to all the members of the CFD groups, you have made my time at UTIAS incredibly enjoyable, fulfilling, and never lonely.

Last, but certainly not least, I would like to thank my family. To my sister Nadine, thank you for being one of my greatest supporters. Your optimism has always kept a smile on my face, even in the most challenging of times. To my parents Nahla and Abdel Rahman – to whom this thesis is dedicated – I cannot begin to thank you enough for the endless love, encouragement, and support that you have given me over the years. And finally, to my wife Jacqui, thank you for your unwavering love, support, and friendship. Words cannot express how grateful I am to have met and married you. It is your companionship that I hold dearest in this world.

Ramy Rashad

University of Toronto Institute for Aerospace Studies
November 13, 2015

Contents

Abstract **iii**

List of Tables **xi**

List of Figures **xiii**

List of Symbols and Abbreviations **xv**

1 Introduction **1**

1.1 Motivation 1
1.2 Background 9
1.3 Review of Existing Techniques and Challenges 14
1.4 Summary of Objectives 22
1.5 Thesis Outline 23

2 Numerical Solution of the Navier-Stokes Equations **25**

2.1 Governing Equations 25
2.2 Flow Solution Algorithms 30

3 Transition Modelling and Prediction in a RANS Solver **35**

3.1 Transition Region Models 35
3.2 Determination of the Boundary-Layer Edge 38
3.3 Calculation of the Boundary-Layer Properties 40
3.4 Evaluation of the Transition Prediction Criteria 42
3.5 Convergence to a Free-Transition Steady State 47
3.6 Flow Solver Efficiency 48
3.7 Flow Solver Robustness and Additional Considerations 49

4 Optimization Techniques **53**
4.1 Optimization Framework 54
4.2 Problem Definition 55
4.3 Mesh Movement 58
4.4 Sequential Quadratic Programming 59
4.5 Gradient Evaluation Techniques 62
4.6 Augmented Adjoint Formulation for Non-Local Transition Prediction 65
4.7 Solving the Augmented Adjoint System 69
4.8 Multipoint Optimization Strategy 75

5 Verification and Validation **77**
5.1 Transition Region Models 78
5.2 Ensuring a Smooth Design Space 79
5.3 Boundary-Layer Edge Finding Methods 81
5.4 Accuracy Assessment of Boundary-Layer Properties 82
5.5 Transition Prediction 84
5.6 Accuracy and Efficiency of Gradient Computations 88

6 Design Optimization Studies **93**
6.1 Cessna 172R Skyhawk 95
6.2 Bombardier Dash8-Q400 97
6.3 Boeing 737-800 97
6.4 Multipoint Optimization 99
6.5 Pareto Front Design Study 101
6.6 Robustness to Freestream Disturbance Environment 103

7 Contributions, Conclusions and Recommendations **109**
7.1 Contributions and Conclusions 109
7.2 Recommendations for Future Work 111

References **115**

Appendices **129**

A Governing Equations in Generalized Curvilinear Coordinates **131**

B B-spline Parameterization of Airfoil Geometry **135**

List of Tables

4.1 Formulation of the optimization problem 55

5.1 Comparison of objective function gradient components computed using various gradient evaluation techniques: RAE-2822, $Re = 15.7 \times 10^6$, $M = 0.6$, $AoA = 1^\circ$ 90

5.2 Comparison of lift-constraint gradient components computed using various gradient evaluation techniques: RAE-2822, $Re = 15.7 \times 10^6$, $M = 0.6$, $AoA = 1^\circ$ 91

5.3 Time required to compute drag and lift gradients for a single optimization iteration with 15 design variables: RAE-2822, $Re = 15.7 \times 10^6$, $M = 0.6$, $AoA = 1^\circ$ 91

6.1 Optimization cases 94

6.2 Case A summary of optimization results: $Re = 5.6 \times 10^6$, $M = 0.19$, $C_l^* = 0.3$ 96

6.3 Case B summary of optimization results: $Re = 15.7 \times 10^6$, $M = 0.6$, $C_l^* = 0.42$ 98

6.4 Case C summary of optimization results: $Re = 20.3 \times 10^6$, $M = 0.71$, $C_l^* = 0.5$ 99

6.5 Design points and weighting for multipoint optimization 100

6.6 Summary of multipoint optimization results 101

6.7 Summary of single-point optimizations, off-design performance, and multipoint optimization over a range of critical N-factor values 105

List of Figures

1.1 Growth of global annual air passengers 2

1.2 Typical drag breakdown for a transonic commercial transport aircraft (data by Airbus [122]) 6

1.3 Comparison of laminar and turbulent boundary-layer profiles on a flat plate at zero-incidence (exaggerated vertical scale) [20] 7

1.4 Illustration of the natural transition process on a flat plate [132] 11

1.5 Illustration of a stability diagram (top) and N-factor envelope (bottom) [8] . 14

3.1 Intermittency distributions with $x_{\text{tr}}=0.3$ and $l_{\text{tr}}=0.1$ 38

3.2 Demonstration of the application of the e^N transition criterion 44

3.3 Demonstration of the application of the AHD criterion 46

3.4 Shape factor and amplification factor extrapolation 51

4.1 Schematic of the Optimization Framework 54

4.2 B-spline parameterization of NACA 0012 airfoil 56

4.3 Illustration of zone of influence for upper surface transition point. 69

4.4 Augmented adjoint system: convergence and efficiency studies 72

5.1 Eddy-viscosity contours for different forced transition locations 79

5.2 Eddy viscosity ramp-up on a NACA-0012 with transition at 30% chord . . . 80

5.3 Design space comparison using different transition region models and convergence tolerances 81

5.4 NACA-0012 verification of edge finding methods 82

5.5 NACA-0012 off-wall grid refinement study on displacement and momentum thicknesses 83
5.6 NACA-0012 off-wall grid refinement study on shape factor 85
5.7 NACA-0012 effect of numerical dissipation on shape factor 85
5.8 NACA-0012 transition prediction verification 86
5.9 NLF-0416 transition prediction and drag polar validation 87
5.10 S809 transition prediction validation 88
5.11 RAE-2822 B-spline parameterization 90

6.1 RAE-2822 C-grid with 575×224 nodes 94
6.2 Case A optimization results: $Re = 5.6\times10^6$, $M = 0.19$, $C_l^* = 0.3$; symbols indicate transition point locations 96
6.3 Case A optimization convergence history: $Re = 5.6\times10^6$, $M = 0.19$, $C_l^* = 0.3$ 96
6.4 Case B optimization results; $Re = 15.7\times10^6$, $M = 0.6$, $C_l^* = 0.42$; symbols indicate transition point locations 98
6.5 Case C optimization results; $Re = 20.3\times10^6$, $M = 0.71$, $C_l^* = 0.5$; symbols indicate transition point locations 99
6.6 Comparison of single-point and multipoint optimization results (design point 6) 102
6.7 Pareto front study, Case B: $Re = 15.7\times10^6$, $M = 0.6$, $C_l^* = 0.42$ 103
6.8 Comparison of single-point optimizations performed at different critical N-factors values 106
6.9 Comparison of single-point and multipoint optimizations with different critical N-factor values 108

List of Symbols and Abbreviations

Roman Symbols

AoA	angle of attack
N	amplification factor of disturbance
N_{crit}	critical amplification factor
e^N	amplification factor transition criterion
A	amplitude of disturbance, area of airfoil
A_o	amplitude of disturbance at critical point
s	arc length
CO_2	carbon dioxide
x, y	Cartesian coordinates
c	chord length
C_m	moment coefficient
C_l	coefficient of lift
C_d	coefficient of drag
C_{d_p}	coefficient of pressure drag
C_{d_f}	coefficient of skin friction drag
C_l^*	target coefficient of lift
C_1	crossflow transition criterion
ϵ_{nk}	convergence tolerance for Newton-Krylov solver
ϵ_{af}	convergence tolerance for approximate-factorization solver
ϵ_{tr}	convergence tolerance for transition residual
ϵ_{adj}	convergence tolerance for linear solve of adjoint system
$\widetilde{\mathbf{c}}$	vector of equality constraints
$\widehat{\mathbf{c}}$	vector of inequality constraints
ρ	density

$\mathbf{x}$	vector of design variables
$\boldsymbol{\ell}$	vector of lower bounds on design variables
$\boldsymbol{u}$	vector of upper bounds on design variables
$\mathcal{Z}$	design-point weighting
e	internal energy
$\mathcal{F}$	feasibility
f	frequency of disturbance
$\boldsymbol{\mathcal{G}}$	gradient vector
J	metric Jacobian, flow Jacobian matrix
$\boldsymbol{\mathcal{A}_1}$	first-order flow Jacobian matrix
$\boldsymbol{\mathcal{A}_2}$	second-order flow Jacobian matrix
$\mathcal{L}$	Lagrangian
$\mathcal{J}$	objective function
M	Mach number
N_{Q}	number of discrete flow variables
N_{D}	number of design variables
N_{c}	number of constraints
N_{i}	number of inequality constraints
N_{e}	number of equality constraints
N_{tr}	number of transition points
N_{G}	number of nodes
$\mathrm{NO_X}$	nitrogen oxides
$\mathcal{P}$	optimality
p	pressure
$\mathbf{Q}$	vector of flow variables
$\boldsymbol{Q}_{\mathrm{c}}$	vector of mean flow variables
$\mathcal{T}$	quadrature weights
$\boldsymbol{\mathcal{R}}$	flow residual vector
$\boldsymbol{\mathcal{R}}_{\mathrm{tr}}$	transition residual vector
Re	Reynolds number
Re_θ	momentum thickness Reynolds number
$Re_{\theta_{\mathrm{tr}}}$	transitional momentum thickness Reynolds number
$Re_{\theta_{\mathrm{cr}}}$	critical momentum thickness Reynolds number
T	temperature

t	time, thickness
l_{tr}	transition length
$\mathbf{X}_{\mathrm{f}}$	vector of forced transition points
$\mathbf{X}_{\mathrm{p}}$	vector of predicted transition points
T_{up}	upper surface transition location
T_{lo}	lower surface transition location
T_{u}	freestream turbulence intensity
u, v	velocity components
H	shape factor
H_k	kinematic shape factor
f_{t_1}, f_{t_2}	Spalart-Allmaras trip functions
h	step-size
$\mathrm{SO_X}$	sulfur oxides
W	weight
Ω	vorticity
U	magnitude of velocity vector

Greek Symbols

$\boldsymbol{\psi}$	vector of adjoint variables
$\boldsymbol{\psi_{\mathbf{tr}}}$	vector of transition adjoint variables
α	complex wavenumber of disturbance
η, ξ	curvilinear coordinates
δ^*	displacement thickness
μ	fluid viscosity
γ	intermittency function, ratio of specific heats
$\boldsymbol{\lambda}$	vector of Lagrange multipliers
Λ_2	Pohlhausen parameter
$\overline{\Lambda_2}$	mean Pohlhausen parameter
θ	momentum thickness
τ	shear stress
μ_t	turbulent eddy-viscosity

Abbreviations

AD	adjoint method
iAD	iterative adjoint method
nAD	noniterative adjoint method
AHD	Arnal-Habiballah-Delcourt
ASO	Aerodyanmic Shape Optimization
CAEP	Committee on Aviation Environmental Protection
CFD	Computational Fluid Dynamics
ERA	Environmentally-Responsible Aviation
FD	Finite-Difference
GARDN	Green Aviation Research and Development Network
GHG	Greenhouse Gas
GMRES	Generalized Minimum Residual
ICAO	International Civil Aviation Organization
IPCC	Inter-governmental Panel on Climate Change
KKT	Karush-Kuhn-Tucker
MBM	Market Based Measure
NASA	National Aeronautics and Space Administration
NLF	Natural Laminar Flow
RANS	Reynolds-averaged Navier-Stokes
SA	Spalart-Allmaras
SN	flow sensitivity method
USD	United States Dollar
VITAS	Vettore Innovativo per il Trasporto Aereo Sostenibile

Subscripts

e	conditions at boundary-layer edge
cr	critical point
∞	freestream conditions
ft	fully-turbulent conditions
inc	incompressible boundary-layer property
lt	laminar-turbulent conditions
lo	lower surface of airfoil
tr	transition point
up	upper surface of airfoil

Over-symbols

$\overline{\mathbf{v}}$	augmented vector in adjoint formulation
$\tilde{v}$	dimensional scalar quantity
$\widehat{\mathbf{v}}$	vector in generalized curvilinear coordinates

Chapter 1

Introduction

This thesis is concerned with the development and demonstration of an effective design tool for improving the aerodynamic efficiency of subsonic and transonic airfoils. The objective is to advance the state-of-the-art in high-fidelity aerodynamic shape optimization by incorporating and exploiting the phenomenon of laminar-turbulent transition. In this chapter, the motivation for this work is first presented in the context of a larger effort towards sustainable aviation. Next, the necessary background information on the phenomenon of laminar-turbulent transition as it relates to subsonic and transonic flight is presented, along with a review of the existing design techniques and challenges. Finally, the specific objectives of this work are discussed and an outline provided.

1.1 Motivation

1.1.1 Towards Sustainable Aviation

In an attempt to ensure the sustainability of aviation, serious effort is underway to mitigate the escalating economic, environmental, and social concerns of the industry. A sustainable aviation industry must take an integral approach to strategically addressing these concerns, which often overlap and compete. From an economic perspective, the industry must maintain high levels of growth and employment while accounting for unit development costs, operational costs, growing infrastructure, and the relative affordability and competitiveness of air travel [127]. From an environmental perspective, the industry must take action to reduce the climate change impact of its technology, as well as address the issues surrounding noise pollution, local air quality, the consumption of natural resources, and end-of-life disposal [91]. Finally, from a social perspective, the industry is first-and-foremost responsible for public safety, but also public accessibility, passenger comfort, travel times, and public opinion when informing decisions on policy and infrastructure [127].

These concerns are growing alongside the industry itself. Indeed, the demand for air transportation of passengers and goods has grown tremendously since the World War II era.

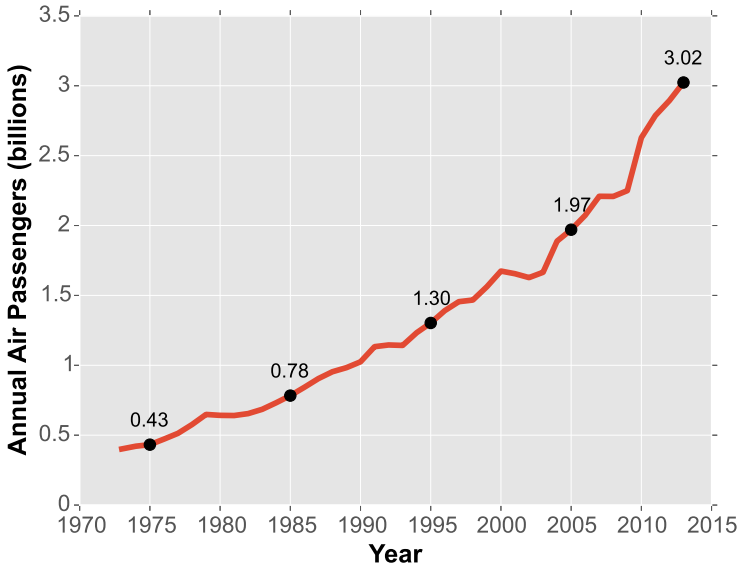

Figure 1.1: Growth of global annual air passengers

Figure 1.1 shows the growth in the annual number of revenue generating passengers over the past forty years (data from The World Bank [124]). The world's average year-on-year growth rate over this period was 5.25%, averaging 6% over the past decade alone, despite the economic recession.

Future projections by Boeing predict a 5.0% annual increase in revenue-passenger-kilometers up to 2032, with a 5.8% annual increase in freight-tonne-kilometers [12]. Almost identical numbers have been projected by Airbus [39]. Furthermore, Boeing's market outlook projects that by the year 2032, the total number of commercial transport aircraft in operation will double to 41,240 aircraft, as compared to 20,310 in operation as of 2013 [12]. In fact, airlines are expected to invest more than 1 trillion USD in fleet renewal over the next 20 years [74].

From an economic standpoint this growth directly contributes to increased revenue, job creation, and the world's gross domestic product. However, from an environmental standpoint, it implies that the total and relative impact of aviation on climate change will continue to increase as a result of the projected growth exceeding efficiency gains. Not only is aviation's impact on climate change growing, but it is growing relatively faster than other industries. Gossling and Upham have suggested that there is no other human activity pushing *individual* emission levels as fast and as high as air travel [39].

Environmental Motivation

Through the combustion of fossil fuels, aircraft emit gases and particles that alter the composition of the atmosphere. This, in turn, causes a net warming effect that contributes to climate change. The aviation sector currently contributes between 2-3% of the total greenhouse gas (GHG) emissions worldwide. When considering the total radiative forcing due to all emissions from aircraft (including non-GHGs), the aviation sector is responsible for upwards of 5% of anthropogenic climate change [68].

The direct GHG emissions from gas turbine engines are carbon dioxide (CO_2) and water vapour [91]. In addition, gas turbines emit other gases and particles that affect the composition of the atmosphere, including: nitrogen oxides (NO_X), sulfur oxides (SO_X), and soot. Under the right conditions, gas turbine engines are also responsible for the formation of condensation trails (or contrails) and also contribute to cirrus cloud formation. Detailed descriptions of these emissions, including their radiative forcing and our current levels of scientific understanding may be found in the literature [31, 39, 42, 57, 70, 77, 91, 128]. Other ways in which aviation negatively impacts the environment include issues surrounding local air quality, noise pollution, the consumption of natural resources, and end-of-life disposal.

The sources of these environmental impacts may be divided into three categories: 1) flight operations, 2) airport operations, and 3) maintenance and disposal [141]. The first category presents the greatest potential to reduce the environmental impact of the aviation sector. Significant improvements in the energy efficiency and emission levels are required through the development of advanced (and unconventional) technologies and operational procedures.

Initiatives by various national and international organizations [4, 22, 33–35, 43, 48, 49, 77, 122, 125] have resulted in detailed publications that outline the vision for the future of an environmentally responsible aviation sector. At the highest-level, these initiatives highlight the environmental concerns of the industry, and identify the various operational and technological strategies with the greatest potential to reduce aviation's impact on the environment. Many of these initiatives aim to meet specific targets on reducing emissions. For example, the Inter-governmental Panel on Climate Change (IPCC) has targeted carbon neutral growth by 2020 and a 50% reduction in CO_2 emissions by 2050 (relative to 2005 levels). NASA has also set aggressive targets through its Environmentally Responsible Aviation (ERA) program [130], which takes its guidance from the U.S. National Aeronautics Research and Development Plan [75]. The ERA targets for the 2030-2035 timeframe include a 71 dB reduction in noise, 75% reduction in NO_X, and 70% reduction in fuel burn, all relative to the Boeing 737-800 with CFM56/7B engines (representative of 1999 entry into service technology) [75]. The aviation sector is thus faced with the difficult challenge of attempting to meet these critical targets, while simultaneously meeting the growing demands for air transportation.

Economic Motivation

The overwhelming majority of the energy used for air travel is provided by fossil fuels [3]. Fuel has become one of the largest direct operating costs for airlines and is a key concern for the profitability, sustainability and competitiveness of the industry. The scarcity of fossil fuel and the volatility of its price have added motivation to the development of (i) alternative and renewable energy sources, such as biofuels, and (ii) more energy efficient flight operations. Improving the energy efficiency of aircraft and engines is particularly important and relevant to the economic sustainability of the aviation industry.

Furthermore, governments around the world are beginning to impose economic incentives by putting a price on emissions through Market-Based Measures (MBMs). Carbon taxes, for example, have been successfully deployed in several jurisdictions worldwide and have shown promising reductions in emissions without slowing economic growth [100]. Alternatively, cap-and-trade systems, which provide a means to limit the total quantity of anthropogenic emissions in a given jurisdiction, have also been deployed worldwide, albeit with greater difficulty due to the legal and administrative complexities involved. The international aviation sector faces additional challenges such as the multilateral agreements made in the 1944 Chicago Convention on International Civil Aviation [50], which prohibit the taxation of aviation jet fuel. Nonetheless, at the 38th Session of the International Civil Aviation Organization (ICAO) Assembly in 2013, more than 170 member states and international organizations pledged to develop a global MBM applicable to the aviation sector [1]. A global MBM scheme is to be developed (focusing primarily on CO_2) with the goal of incentivizing significant reductions in climate change impact. The objective is to assign an economic value to our environment and to send a signal to the aviation sector (including government, industry and academia) to focus on its climate change impact.

In the same vein, government regulations and certification standards on emissions and noise have become increasingly stringent, with the goal of incentivizing the research, development, and deployment of cleaner and quieter technologies. They provide a direct link between quantifiable noise and emission reductions and the technology goals of manufacturers. These standards stem from recommendations put forth by ICAO's Committee on Aviation Environmental Protection (CAEP). CAEP has focused primarily on ensuring sufficiently low levels of perceived noise and NO_X emissions, with the typical approach being to balance what is technologically feasible with the health impacts of people living in the vicinity of airports [141]. In 2013, the requirement for a new aircraft-level CO_2 emissions standard was approved by CAEP, with the delivery of the full standard expected at the CAEP 10 meeting in 2016 [51].

The combination of new MBMs and strict regulations has a direct impact on the economics of aviation and, as such, incentivizes the aviation industry to take more aggressive action towards more efficient flight operations. Manufacturers must strategically invest in the advanced technologies that will mitigate the financial risk of certifying and marketing new aircraft and engines that are expected to operate for decades to come.

Advanced Technologies for Sustainable Aviation

Over the past decade, the number of initiatives which have been formed to address sustainable and environmentally responsible aviation has grown tremendously. Examples include NASA's ERA [130] program, the European Union's CleanSky initiative [22], and Canada's Green Aviation Research and Development Network (GARDN) [43]. All of these have made recommendations on a host of advanced technologies and operational strategies with the greatest potential for ensuring the sustainability of aviation. The proposed strategies can be divided into four categories: 1) traffic management/operations, 2) airframe technology, 3) propulsion technology, and 4) alternative fuels. An excellent review of the proposed technologies may be found in Agarwal [2, 3].

The work presented in this thesis is specifically concerned with the development and application of a design tool for airframe technology. One of the goals in developing such tools is that they will help accelerate technology development by providing airframe designers with a new means to explore their design space and to understand the various trade-offs involved.

Arguably the two most important goals of an airframe designer are to improve aerodynamic efficiency and reduce structural weight [141]. In this work, we consider the aerodynamic efficiency of aircraft, specifically looking at how to minimize aerodynamic drag. Indeed, reducing aerodynamic drag directly reduces fuel burn. Not only does this ameliorate the economics of aviation, it also directly reduces the consumption of natural resources and the emissions of CO_2, aerosols, and NO_X per passenger-km of flight. As part of the ongoing effort to reduce drag, aerodynamicists are assessing the feasibility of *natural laminar flow* as a key enabler of sustainable aviation. Indeed, natural laminar flow ranks among the most important technologies for reducing the fuel consumption and emissions per passenger-km of flight [3, 41, 42, 133].

1.1.2 Laminar Flow: The Holy Grail of Drag Reduction

The sources of drag on an aircraft are now well understood. Figure 1.2 depicts the typical drag breakdown for a modern transonic transport aircraft. It is clear that the greatest

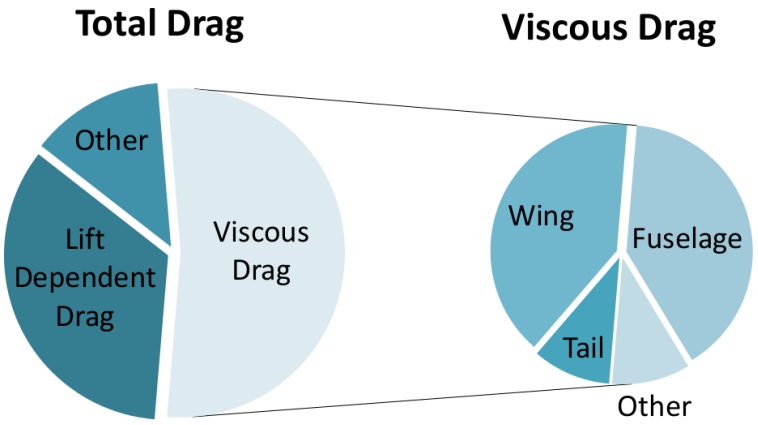

Figure 1.2: Typical drag breakdown for a transonic commercial transport aircraft (data by Airbus [122])

contributor to aircraft drag is the drag due to viscous (frictional) effects. Laminar flow technologies promise significant reductions in viscous drag by exploiting the phenomenon of laminar-turbulent transition.

In the late nineteenth and early twentieth centuries, the breakthrough work of Reynolds and Prandtl began to shed light on the theory and experimentation of boundary layers and laminar-turbulent transition [132]. It was observed that a wall-bounded flow begins as a smooth laminar flow with relatively (often significantly) lower drag, and somewhere along the way the boundary layer transitions to an apparently chaotic, random, three-dimensional turbulent flow with relatively higher drag. A physical description of the transition process is presented in Section 1.2.2.

The negative consequences of transition on viscous drag are observed as an increase in both the skin friction (tangential shearing forces) and the pressure drag (normal forces). In Figure 1.3, the average velocity distribution of a laminar and turbulent boundary layer are compared. In a turbulent boundary layer, the velocity profiles near the wall become significantly steeper in the off-wall direction. The shear force, τ, is directly proportional to the off-wall velocity gradient at the surface, such that

$$\tau = \mu \left. \frac{\partial U}{\partial y} \right|_{\text{y=wall}}, \tag{1.1}$$

where μ is the dynamic viscosity of the fluid. The observed increase in the pressure drag component, for a fully-attached flow, can be attributed to the thickening of the boundary layer, which occurs more rapidly in the turbulent region, as depicted in Figure 1.3 (with an exaggerated vertical scale) [20].

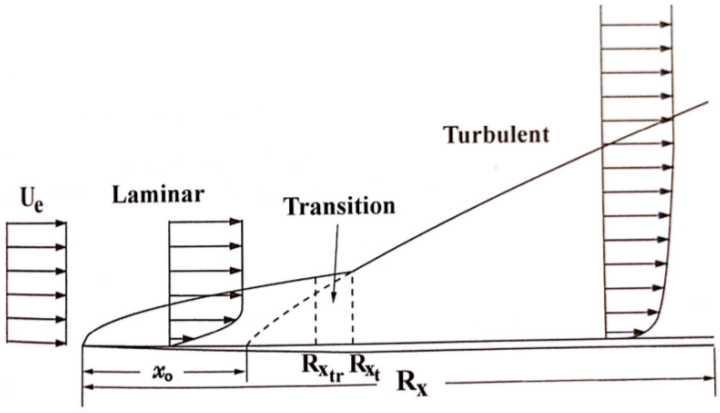

Figure 1.3: Comparison of laminar and turbulent boundary-layer profiles on a flat plate at zero-incidence (exaggerated vertical scale) [20]

The principal advantage of Natural Laminar Flow (NLF) wings is that they are characterized by extended regions of laminar flow. Depending on the flight conditions, a typical NLF airfoil could have laminar flow well beyond 50% of the chord length on both the upper and lower surfaces. In contrast to *hybrid* laminar flow control, which involves some form of active flow control (such as boundary-layer suction through slotted or perforated surfaces [20]), *natural* laminar flow refers to the passive flow control technique of appropriately shaping the wing to achieve a pressure profile more conducive to maintaining laminar flow [108].

Design for Natural Laminar Flow

Design for NLF is not a new endeavour. An excellent historical review is given by Bushnell [16, 17]. In the late 1930s the pioneering work of Jacobs led to the development and design of the first NACA laminar flow airfoils [6], and shortly after resulted in the first attempted application of NLF design on the P51-Mustang of World War II. Unfortunately, over the years, NLF has proven to be as elusive as it is desirable. Today, there remain very few NLF applications in the modern commercial fleet. The nacelles on the recent Boeing 787 and the wings on the new Honda business jet are among the first, if not the only applications to date in commercial aviation [19].

Indeed, it has been more than 75 years since Jacobs first designed a laminar-flow airfoil and more than a century since the breakthrough work of Reynolds and Prandtl. Today's airframe designers have become heavily reliant on sophisticated computational tools, such as Computational Fluid Dynamics (CFD) – the branch of fluid dynamics that uses numerical algorithms and high-speed computers to model fluid flow. Aerodynamic Shape Optimization (ASO) is also an important and growing engineering strategy in which numerical optimization

frameworks are used to alter a geometry (such as the shape of a wing) to minimize or maximize a design objective (such as drag) while simultaneously satisfying a set of design constraints (such as structural and/or lift requirements).

In a review by Sengupta [108], he discusses the DESB154 and DESB165 NLF airfoils designed by Viken in the early 1980s. Sengupta notes that the NLF designs "proceeded with hand drawing an airfoil (based on experience!) and then calculating the corresponding (i) pressure distribution, (ii) boundary layer, and (iii) its stability properties, in that sequence. This procedure was carried through iteratively until the design goals were met". Similar design methodologies were carried out by Boeing and NASA in the late 1970s [80] and early 1980s [110]. Also in the early 1980s, Zingg [139] developed an inverse design procedure for high lift to drag ratio airfoils by specifying desirable transition and pressure recovery characteristics.

These aerodynamicists developed an expertise (and patience) in shaping NLF airfoils (or specifying their characteristics) by hand. One question this thesis seeks to answer is, to what extent can numerical shape optimization be used to automatically design NLF airfoils using state-of-the-art high-fidelity flow solvers?

Over the past few decades, the conservative engineering assumption of *fully-turbulent flow*, which considers the flow to be fully turbulent right from the leading edge (over the entire wetted area), has been widely and overwhelmingly common in CFD, and especially in ASO. While this assumption has allowed for stunning advancements in aerodynamic design, the conservatism leaves something to be desired. Indeed, the development of design tools capable of exploiting laminar-turbulent transition are required to enable the robust design of NLF technology and to provide a means by which we may better understand the design trade-offs involved.

Factors Affecting Transition

While the lack of NLF applications in the fleet points to the sparsity of available design tools, it also points to the challenges in reliably realizing extended regions of laminar flow in flight. The transition to turbulence is affected by many factors, including: Reynolds number (Re), freestream turbulence intensity (T_{u}), pressure gradients, Mach number (M), surface roughness, heating, noise, vibrations, rain, hail, icing, and even insect impacts [13, 119, 132].

Although the preceding list of factors is formidable, the quest for the holy grail of drag reduction must be persistent and steadfast. On this quest, we must investigate all theoretical, experimental, and computational methods that may actualize NLF in commercial aviation. Engineers continue to explore a variety of techniques that work to promote laminar flow,

including, but not limited to: shaping the aircraft to maintain favourable pressure gradients [105], lower sweep angles and distributing surface roughness elements to stabilize crossflow instabilities [47, 103], slat-less wing configurations [10], boundary-layer suction and wall heating [73], plasma and piezoelectric actuators [45, 89], non-stick materials and coatings [135], and new manufacturing and maintenance procedures. As such, design tools capable of efficiently exploiting any of the available techniques are of immediate consequence and benefit to the present day designer.

In this work, the NLF design of airfoils is demonstrated through high-fidelity, single and multipoint aerodynamic shape optimization with transition prediction capable of accounting for the effects of Re, T_{u}, M, and the pressure gradients. A more detailed outline of the specific objectives and scope of this work is provided in Section 1.4. In the next section, we introduce the theory of laminar-turbulent transition as it relates to subsonic and transonic flight.

1.2 Background

1.2.1 Transition in Subsonic and Transonic Flight

Commercial transport aircraft flying at subsonic and transonic speeds travel through relatively still air with very low freestream turbulence intensities ($T_{\mathrm{u}} < 0.1\%$ of U_∞) [8]. While there are many factors that affect transition (as described in the previous section), we can categorize the class of transition that is expected in subsonic and transonic flight based on the given geometry and external flow conditions. The various classes of transition that can be expected in flight are as follows [8, 99]:

Natural transition: Transition that arises due to small disturbances that take the form of two-dimensional periodic instabilities, known as Tollmien-Schlichting waves [105], is referred to as natural transition. The ensuing stability analysis of these waves has been widely used in the development of several correlation-based transition prediction criteria [20]. The specific phases of the natural transition process are described in the next section.

Crossflow transition: The situation in three dimensions is significantly more complex – if the wing is swept – since there are two components of the boundary-layer velocity profile to consider. In this case, transition theory considers the additional type of instability known as crossflow waves, which make an angle of 85 to 89 degrees to the freestream direction and may be stationary or traveling in space [8].

Separated flow transition: This type of transition occurs when a laminar boundary layer separates off the surface, transitions, and reattaches as turbulent, in turn forming a transitional separation bubble. In this case, some or all of the natural transition process occurs in the region of the bubble [132].

Bypass transition: Bypass transition is a more abrupt path to transition in which three-dimensional breakdown is observed before the development of two-dimensional Tollmien-Schlichting instabilities. The various bypass mechanisms include: high freestream turbulence intensities, high surface roughness, and Görtler vortices occurring on concave surfaces [90].

Attachment line contamination: When the wing is swept, the attachment line (that is, the locus of points which would be considered the stagnation points on an unswept wing) has a flow component in the spanwise direction and thus a boundary layer develops. The state of this attachment-line boundary layer may become turbulent if large disturbances exist and emanate from the wing-fuselage junction (or due to high surface roughness). This can, in turn, lead to fully-turbulent flow (over the full chord) on both the upper and lower surfaces [92]. If the contamination exists, which can be verified by various empirical criteria [20, 92, 108], Schmitt *et al.* [106] have demonstrated the effectiveness of a so-called Gaster bump placed near the fuselage that creates a new stagnation point in the spanwise direction. Designers can also avoid this contamination by considering fairings that smooth the wing-fuselage junction.

The most commonly observed paths to transition observed in high Reynolds number flight at subsonic and transonic speeds are those of natural transition and – if the wings are swept – crossflow transition. Note that at relatively lower Reynolds numbers (typical of single-engine piston aircraft or unmanned aerial vehicles) the corresponding boundary layers are more readily separated and thus may experience separated flow transition. For internal flows (typical of the flow through a gas turbine engine) much higher levels of turbulence intensities are observed and, in turn, bypass transition is to be expected.

While separated flow transition is considered in this work, bypass transition is not. Furthermore, since this work is restricted to the analysis and optimization of airfoils in two dimensions, crossflow and attachment line transition are not considered. We make it clear that any design examples involving swept wings may require further modification (particularly in the design near the wing root and wing tip) that aims to address the crossflow instability, while not allowing any significant degradation in Tollmien-Schlichting stability [80]. Nonetheless, even for swept wings, natural transition has been observed to occur over

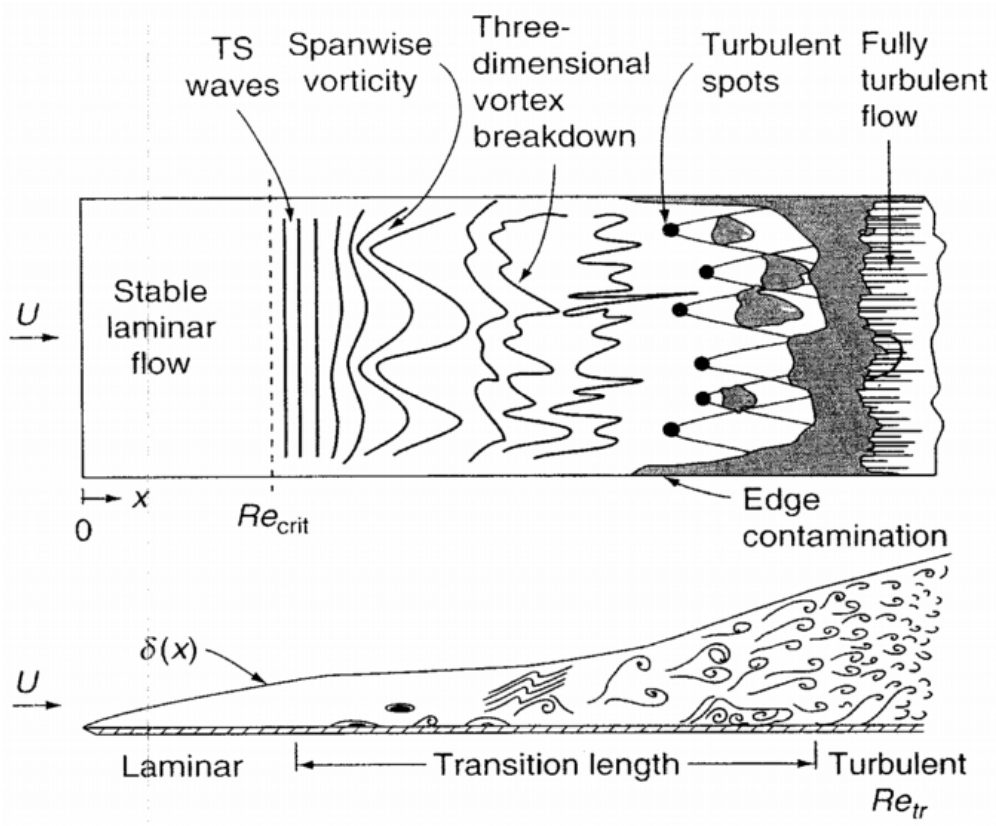

Figure 1.4: Illustration of the natural transition process on a flat plate [132]

significant portions of the wingspan [117]. Finally, we note that the transition criteria and methods used in this work have been implemented with the ability to be extended and coupled to available crossflow and attachment line criteria in three dimensions.

1.2.2 The Natural Transition Process

Theoretical investigations into transition have been largely based on the assumption that laminar flows breakdown by the growth of small, unstable, disturbances [105]. The natural transition process is illustrated in a planform and side profile view of the flow over a flat plate in Figure 1.4 [132]. The three stages of natural transition applicable to both 2D and 3D flowfields are as follows [8]:

Stage 1 - Receptivity: The receptivity process describes the mechanisms by which forced disturbances enter the laminar boundary layer. These forced disturbances may be the result of a number of factors, such as surface roughness, freestream turbulence, acoustics, and vibrations [132]. Ultimately, the receptivity process provides the initial amplitude, phase and frequency of the instability waves that lead to turbulence. Unfortunately,

the physics of receptivity is difficult to study (experimentally and computationally) and the process as a whole is not yet well understood [102, 104]. This continues to drive research into the fundamental physics of the transition process, as reviewed by Saric [104].

Stage 2 - Intermediate Growth: The disturbances take the form of periodic waves (initially two-dimensional), the energy of which is convected in the streamwise direction. In the intermediate growth phase, some disturbances will decay and others will amplify and will ultimately be responsible for transition [8]. The intermediate growth of the disturbances is well described by *linear stability theory*, with successful extensions to crossflow waves observed on swept wings [116]. Linear stability theory is discussed further in the next section, as it provides the basis for many transition prediction criteria.

Stage 3 - Three-Dimensional Breakdown and Nonlinear Interactions: As the unstable waves grow to finite amplitude (on the order of 1% of U_∞), secondary instabilities lead to three-dimensional breakdown and nonlinearities lead to turbulent spots [8]. The turbulent spots, as depicted in Figure 1.4, are local intense bursts of turbulence that appear at random points in space and time and coalesce as they propagate downstream into fully developed turbulent flow. Note that while the nonlinear interactions are critical to the transition process itself, they are not critical to predicting if or where the transition process will occur. For a detailed review of the various nonlinear paths to fully-turbulent flow, the reader is referred to a review by Stuart [118].

Finally, note that while the receptivity process and nonlinear interactions differ for Tollmien-Schlichting and crossflow waves, in the intermediate growth phase, linear stability analysis may be used to asses the stability of both types of disturbances [8, 116].

1.2.3 Linear Stability Analysis

The goal of linear stability analysis is to assess the growth of the spectrum of instabilities that exist in a laminar boundary layer in order to determine if or where the disturbances grow to a level sufficient for the three-dimensional breakdown and nonlinear interactions (Stage 3 processes) to occur. It is important to note that linear theory is no longer applicable when the nonlinear effects come into play. The goal here is to elucidate the concepts and terminology relevant to the transition criteria employed in this work; a more rigorous discussion of boundary-layer stability theory may be found in the literature [105, 132].

The well-known Orr-Sommerfeld equation is the fundamental equation for the linear stability theory of boundary layers; for brevity it is not presented here, but it can be found in several textbooks [20, 105, 132]. While originally developed for 2D incompressible flows, it has successfully been extended to 3D compressible flows [8]. One of the first assumptions made in its derivation is that the 2D disturbances are periodic and take the following form:

$$q' = q(y) \exp\left[i(\alpha x - ft)\right], \tag{1.2}$$

where q' represents the disturbance, $q(y)$ is the amplitude function, α is the complex wavenumber, and f is the frequency of the given disturbance. Note that the amplitude function is not a function of the streamwise position, x, which highlights the aptly named *parallel flow assumption*. In order to account for curvature effects, such that $q(y)$ is instead $q(x, y)$, one must resort to the parabolized stability equations [5, 99]. While these equations are generally more accurate than the linear stability equations, they are also more cumbersome and expensive to employ.

In spatial amplification theory, f is a real number, and the stability of a given disturbance is thus described by the real part of the exponential, given by $\exp(-\alpha_i x)$, where α_i is the complex part of the wave number. As such, waves of a given frequency (f) and at a given local Reynolds number that have a negative α_i will decay, and vice versa, those with positive α_i will grow. In essence, solutions to the Orr-Sommerfeld equation over the entire range of unstable frequencies and Reynolds numbers informs us of the frequencies that are most unstable, and how much they have grown at each streamwise station along the airfoil surface. Indeed, if we take the locus of points $\alpha_i = 0$ we can plot a stability diagram in $f - x$ space, where x is the streamwise position, as illustrated in Figure 1.5. This thumb-shaped stability curve is such that any disturbances lying outside of the thumb are stable, while any inside the thumb are unstable. The point at which (spatially) the first disturbance becomes unstable is known as the *critical point*.

At a given streamwise station, x, the total amplification ratio of a disturbance is given as

$$\ln\left(\frac{A}{A_o}\right) = \int_{x_o}^{x} (-\alpha_i) dx\,, \tag{1.3}$$

where A is the amplitude at the given streamwise position, and A_o and x_o represent the amplitude and streamwise position at the point at which the given frequency first becomes unstable [8].

The so-called N-factor envelope, also depicted in Figure 1.5, represents the amplification

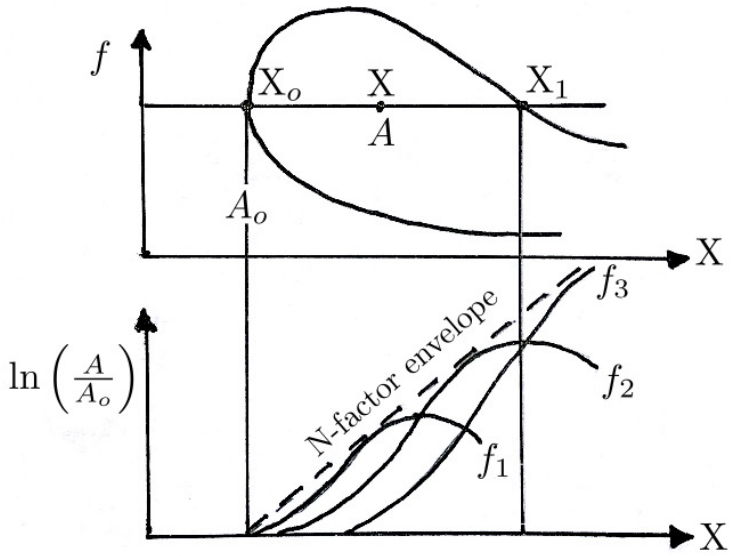

Figure 1.5: Illustration of a stability diagram (top) and N-factor envelope (bottom) [8]

of the most unstable disturbance at a given station, such that

$$N(x) = \max_{f} \left[\ln \left(\frac{A}{A_o} \right) \right] . \tag{1.4}$$

This N-factor is often used as a criterion for transition prediction, known as the e^N transition criterion first developed in 1956 by Van Ingen [129] and, independently, by Smith and Gamberoni [109]. The e^N criterion assumes that transition will occur when a critical value, denoted N_{crit}, has been reached. Typical values for N_{crit} in low freestream turbulence environments are in the range of 7 to 13. Finally, for incompressible flows, an empirical relation developed by Mack may be used to relate the freestream turbulence intensity, T_{u}, to the critical N-factor as follows [72]:

$$N_{\text{crit}} = -9.34 - 2.4 \ln(T_{\text{u}}) \text{ valid for } 0.0007 \leq T_{\text{u}} \leq 0.0298. \tag{1.5}$$

1.3 Review of Existing Techniques and Challenges

Research in the area of high-fidelity aerodynamic shape optimization for natural laminar flow is sparse. The majority of research in this field employs inviscid-viscous coupling strategies, making use of boundary-layer codes for the viscous formulation, and either a panel method or the Euler equations for the inviscid formulation. Although the inviscid-viscous coupling strategies can be computationally cheaper than higher-fidelity solvers, Reynolds-averaged

Navier-Stokes (RANS) solvers have seen great success and widespread use in recent decades. Indeed, RANS solvers provide an excellent balance between accuracy, robustness, and efficiency, while being readily extended to three-dimensional and separated flows. Furthermore, the industry's general trend toward the use of RANS solvers strongly suggests that NLF design tools should follow suit. This work makes exclusive use of RANS solvers for the analysis and optimization of aerodynamic flows with free transition. The details of the RANS solver will be presented in Chapters 2 and 3. The development of design tools that make use of high-fidelity solvers are essential to accelerating the development of technology and minimizing the risk of novel design by more accurately representing the physics (and hence the design spaces) involved. In this section we review the existing techniques and challenges involved with transition prediction in RANS solvers, and aerodynamic shape optimization frameworks that incorporate transition.

1.3.1 Review of Transition Prediction in RANS Solvers

The challenges in reliably predicting laminar-turbulent transition continue to limit our ability to compute many aerodynamic flows with accuracy [140]. In fact, the turbulence models used in RANS solvers do not have the stand-alone capability to predict the laminar-turbulent transition locations in a flow field; in order to predict transition, one must apply a transition criterion. Consequently, the development of transition prediction criteria and flow solver coupling strategies of varying complexity and fidelity are ongoing.

In recent years, several approaches for coupling transition prediction criteria into RANS solvers have been developed. An excellent review by Arnal *et al.* [8] discusses the various advantages and disadvantages of each approach in detail. The following list attempts to categorize the available strategies:

1. coupling of a RANS code with a linear or parabolized stability solver and the e^N criterion [59, 63, 83, 115, 116];
2. direct implementation of simplified e^N methods into the RANS code [13, 78, 81, 117];
3. direct implementation of correlation-based transition onset functions (criteria typically based on Re_x or Re_θ) into the RANS code [20, 23, 81, 117];
4. coupling of a RANS code with a boundary-layer code and transition prediction using the criteria of methods 1 through 3 applied to the boundary-layer solution [30, 62, 63];
5. coupling of additional transport equations to the RANS turbulence model, such as the $\gamma - \widetilde{Re}_{\theta_t}$ transition model, originally developed by Langtry and Menter with variants

developed by other researchers [7, 24, 66, 67, 79, 82]. These approaches make use of analytical or correlation-based transition onset functions built into the transport equations.

In all of the above strategies, the transition criteria employed by the RANS solvers are based on either (i) the e^N criterion or (ii) transition onset functions. To apply the e^N criterion, which was introduced in the previous section, one must first approximate the N-factor curves, representing the amplification ratios of the unstable frequencies of the disturbances in the boundary layer. Transition is assumed to occur when the maximum local N-factor has exceeded some critical value (N_{crit}). Values for N_{crit} must be specified *a priori* based on the freestream disturbance environment and/or experimental calibration. In computing the N-factor curves, there are several methods of varying fidelity and computational cost. The highest fidelity approach (for RANS solvers) is to solve the parabolized or linearized stability equations at each station to obtain the local N-factors for the unstable frequencies. Simplified approaches have been developed to alleviate the demanding computational cost of such methods, including: the use of database methods generated from linear stability studies [121], approximating the N-factor envelope through boundary-layer properties (such as the shape factor) [28], and approximating the N-factor through the pressure gradient and curve fits based on linear stability theory, without the need for boundary-layer properties [19].

Examples of transition criteria based on a transition onset function are Michel, Granville, H–Re_x, Abu-Ghannam and Shaw, Gleyzes-Habiballah and Arnal-Habiballah-Delcourt [20, 23, 66]. These methods have been classified as "simpler" or "analytical" transition criteria [8]. Each has its range of applicability and limitations. These criteria typically compare the boundary-layer properties or related quantities (such as Re_θ) to an empirically calibrated transition onset function (such as $Re_{\theta_{\text{tr}}}$). The transition onset functions are typically computed from the integrated boundary-layer properties; the exception being the local transport equation approach developed by Langtry and Menter [66, 67]. The transition point is the first point at which, for example, $Re_\theta \geq Re_{\theta_{\text{tr}}}$.

In this work, for the prediction of the natural transition locations due to Tollmien-Schlichting instabilities, we employ approaches 2 and 3 from the above list. The boundary-layer properties are computed directly from the Navier-Stokes solution [98]. While several transition criteria have been implemented, the present framework primarily makes use of Drela's e^N envelope method [28] in approach 2, and the relatively new compressible form of the AHD criterion [23] in approach 3; these criteria are presented in Sections 3.4.1 and 3.4.2, respectively. Transition prediction is implemented into the RANS solver using an iterative approach, and transition to turbulence is modelled using an explicit intermittency function

in conjunction with the SA turbulence model – both are discussed in Chapter 2.

The transport equation approach developed by Langtry and Menter [66, 67] has the advantage of a local formulation in the sense that integrated boundary layer properties are replaced by local flow properties. The local approach makes it more straightforward to parallelize and differentiate the code. While the method was originally coupled to the k-ω turbulence model, more recent works by Medida and Baeder [79] as well as Aranake *et al.* [7] have successfully coupled the $\gamma - \widetilde{Re}_{\theta_t}$ model to the Spalart-Allmaras turbulence model. Unlike the transport equation approach, the other approaches in the above list are non-local in their formulation, which has some disadvantages. However, these are being addressed; for example, approaches 1 and 2 have been successfully parallelized and extended to three-dimensional flows by Krimmelbein *et al.* [59, 60], and there is no obstacle to their use in an implicit Newton-Krylov type solver, as demonstrated in this work. There is also no required calibration specific to a particular turbulence model [79]. Furthermore, correlations for crossflow instabilities (such as the C_1 criterion) have already been successfully combined with these approaches, with experimental validation demonstrating accurate transition prediction on transonic swept wings in three dimensions [81, 117]. Finally, the modular implementation of the proposed transition prediction framework facilitates the use of higher fidelity methods (such as database methods, linear stability theory, or parabolized stability equations) if so desired.

1.3.2 Review of Aerodynamic Shape Optimization for Natural Laminar Flow

The use of ASO in the design process has several advantages. When introduced early on in the design process it provides a means of exploring the design space and unlocks the potential for the discovery of new and improved aerodynamic shapes tailored to the designer's specific application. Second, ASO helps inform design decisions by providing the means to better understand and study the various design trade-offs at play. Third, when a multipoint optimization strategy is employed, ASO can incorporate off-design performance, leading to more robust and practical designs. ASO is a design tool that adds a level of automation to the tedious cut-and-try design approach (Section 1.1.2). It lowers the risk and accelerates the development of aerodynamic systems providing optimal performance. Finally, as previously mentioned, the use of high-fidelity solvers is critical to the design process of novel configurations. As the complexity and scale of CFD flow solvers increase (with improved physical modelling), it is imperative that the optimization strategies be adapted to employ the higher-fidelity solvers.

The following discussion presents a summary of previous research in ASO incorporating laminar-turbulent transition prediction. The following cases, presented in chronological order, demonstrate the sparsity of research in the field [29]. It is important to reiterate that most of the research in this field has involved the use of low fidelity flow solvers through inviscid-viscous coupling strategies, such as the combination of the (inviscid) Euler equations coupled with a (viscous) boundary-layer code. We begin by reviewing a few selected works that make use of low fidelity solvers, and follow with a comprehensive overview of RANS-based optimizations for NLF. The areas for which there is significant room for improvement will be made clear along the way.

Optimization Strategies using Low-Fidelity Flow Solvers

Inviscid-viscous coupling strategies do not solve the Navier-Stokes equations over the entire domain of interest. Instead they iteratively couple an inviscid solver (using, for example, a vortex-panel method or the Euler equations) with a viscous solver (that solves a reduced form of the Navier-Stokes equations known as the boundary-layer equations [105], which are applicable only in the boundary layer). The following is a summary of selected research on NLF optimizations that has been carried out using inviscid-viscous coupling strategies. The list is not exhaustive.

In 1993, Dodbele developed an optimization method for the design of axisymmetric body shapes in compressible flow [26]. The objective function gradient was calculated using a simplified Granville's transition criterion. The strategies employed were restricted to axisymmetric shapes.

In 1996, Green *et al.* developed an optimization method for the constrained design of natural laminar flow airfoils [40]. Convergence of the optimization was towards a so-called *N-factor distribution* (referring to the e^N method). As noted by Driver [29], a drawback of Green's approach is that it requires a significant level of expertise and user know-how in order to specify the appropriate N-factor distributions to NLF design.

In 2003, Pralits developed an optimization method applied to natural and hybrid laminar flow control on wings using solutions of the adjoint equations to obtain the objective function gradients [94]. The composite gradient was based on three different governing equations: the steady compressible Euler equations, the steady boundary-layer equations, and the parabolized stability equations. The adjoint equations are solved separately for each of the three equations and then combined. The gradient was used in a quasi-Newton algorithm to minimize the objective function. This was the first demonstration that an efficient method to minimize boundary-layer disturbances is to formulate the adjoint based on the stability

equations themselves [29]. The approach requires significant know-how in relation to stability theory and the required initial conditions of individual disturbance frequencies.

In 2003, Sturdza *et al.* first developed an optimization method for the natural laminar flow design of wings used on supersonic aircraft [61, 119]. In this work, neither subsonic nor transonic applications are considered. As in the previous two cases, an Euler solver is coupled to a boundary-layer code. Sturdza *et al.* make use of the e^N-method to determine the location of the transition point, including modifications to incorporate cross-flow instabilities. The objective function is based on the instability envelope of the boundary layers. Sturdza's work demonstrated a significant challenge associated with incorporating a transition prediction routine into an optimization process (applied to supersonic aircraft). Driver [29] notes that, "The highly nonlinear characteristics of the laminar-turbulent transition can cause problems with effectively navigating the design space". Sturdza has continued his work over the past decade focusing on database methods and applying the inviscid-viscous solver to the optimization of subsonic and transonic airfoils [120].

In 2006, Amoignon *et al.* considered gradient based optimization aimed at delaying transition by considering the total kinetic energy of a single disturbance (the disturbance giving the largest value of kinetic energy) [5]. Amoignon notes that the results "show a reduction of the total amplification of a large number of disturbances, which is assumed to represent a delay of the transition in the boundary layer" [5]. A second objective function considered attempts to simultaneously delay the transition and reduce the pressure drag coefficient, while maintaining the coefficients of lift and pitch moments. The gradients were obtained by solving the adjoint of the Euler, boundary layer, and parabolized stability equations. Inviscid-viscous coupling between the Euler and boundary layer solvers were ignored [5]. The final geometries (from the optimizations) were analyzed using a RANS solver with prescribed transition locations. This work follows the earlier work of Pralits [94] and was continued in 2009 by Hanifi *et al.* [44].

In 2010, Cella *et al.* considered the design of a natural laminar flow wing for a medium sized transonic business jet, as part of the Italian VITAS project. An inviscid-viscous coupling strategy was first used to perform airfoil shape optimization based on target pressure distributions (inverse design). Several 2D airfoil sections were optimized and then interpolated in the spanwise direction to form a 3D wing. A post-optimization Navier-Stokes analysis was performed on the 3D geometry (with imposed transition locations) and compared to wind tunnel results with good agreement.

Optimization Strategies using High-Fidelity (RANS) Flow Solvers

Research making use of RANS solvers to optimize with transition prediction has shown promising results. The following presents a comprehensive review of the past and current research efforts.

In 2005, Driver and Zingg [29, 30] developed a two-dimensional Newton-Krylov aerodynamic shape optimization algorithm with laminar turbulent transition. At each design iteration, the transition locations were predicted using the coupled Euler and boundary-layer solver, MSES [27, 28], and then used in a compressible RANS solver with the one-equation Spalart-Allmaras turbulence model. The sensitivities of the transition locations to the design variables were obtained by centered-difference approximations, requiring multiple calls to the MSES solver. This information was used in combination with the existing discrete-adjoint formulation for the RANS solver such that the resulting gradient included the effect of the movement of the transition points. The design examples demonstrated the algorithm's ability to control the transition point locations to provide optimum performance [29]. The use of MSES provided a stop-gap approach for investigating the feasibility of transition prediction in the Newton-Krylov aerodynamic optimization algorithm [30]. Many of the recommendations found in Driver's thesis [29] have been implemented in the present work.

In 2009, Lee and Jameson [69] successfully coupled a RANS solver to a boundary-layer code and an e^N database method (making use of the Baldwin-Barth turbulence model) for NLF design in two and three dimensions. The gradient calculations in their work did not include the transition prediction, and their optimizations focused on the elimination of shock-waves for reduced wave drag.

The next five contributions described all make use of the Langtry-Menter γ–$\widetilde{Re}_{\theta_t}$ approach [66, 67] tied to the k-ω turbulence model, and all are applied to airfoils in two dimensions. First, in 2011, Khayatzadeh and Nadarajah [54–56] successfully extended the Langtry-Menter approach to an adjoint-based optimization framework in two dimensions and applied the framework to the design of low Reynolds number NLF airfoils with separation bubbles. Design objectives investigated included the minimization of turbulent kinetic energy and the maximization of the lift-to-drag ratio. In 2013, Jing *et al.* [53] used a particle swarm optimization method to demonstrate the NLF design of transonic airfoils. They were able to demonstrated that the resulting airfoils perform well over a range of uncertainty of lift coefficients. In 2014, Zahle *et al.* [136] performed wind turbine blade optimizations for NLF. Their work made use of gradient-based multi-point optimization to successfully improve on the design of a wind turbine blade that was previously optimized using XFOIL [28]. Unlike the efficient adjoint-based gradient approach employed by Khayatzadeh and Nadarajah

[54, 55], the gradients in this work were obtained by a finite-difference approximation with an optimization cycle requiring greater than 2600 calls to the CFD flow solver [136]. Also in 2014, Zhao *et al.* [137] used a non-sorting genetic algorithm in conjunction with the $\gamma-\widetilde{Re}_{\theta_t}$ approach to demonstrate the robust design of an NLF supercritical airfoil via a multi-objective evolutionary approach. Most recently, in 2015, Wang *et al.* [131] performed transonic airfoil optimization for NLF using a genetic algorithm. Their use of the so-called snapshot gradient-free algorithm also involved thousands of calls to the RANS flow solver for each design example. Note that all publications in this paragraph, with the exception of [54, 55], make use of costly optimization techniques.

1.3.3 Present Challenges and the Path Forward

There are several specific areas that require continued research and development and have significant room for improvement. The direct use of high-fidelity RANS solvers as opposed to the inviscid-viscous coupling strategies is gaining momentum (particularly in the past five years) and should continue to be developed and explored. Further work is required to explore more thoroughly the design spaces and trade-offs at higher Reynolds and Mach numbers (representative of subsonic and transonic transport aircraft). Optimizations of more realistic and practical designs are required through the use of multipoint design optimizations that incorporate the off-design performance during the optimization cycle [15, 18]. In the same vein, design objectives and constraints should aim to reflect the industry's aerodynamic design objectives. All too often, design objectives are formulated specifically or primarily to delay transition. Design objectives based on indirect functions – such as N-factor curves – run the risk of creating a gap between the design tool's capabilities and the actual objectives of manufacturers.

Next, to ensure that the design tools are practical for manufacturers, more efficient optimization strategies are required, such as the continued development and application of the adjoint-based gradient technique, in turn reducing the number of calls to the RANS flow solver (which are themselves more expensive when solving free-transitional flow as compared to fully-turbulent flow). Furthermore, given that the Spalart-Allmaras [112] turbulence model is a very popular choice for wall-bounded aerodynamics [134], there is a significant lack of optimization frameworks that employ a transition prediction model coupled to the Spalart-Allmaras model. As mentioned, for flow analysis, Medida and Baeder [79] as well as Aranake *et al.* [7] have successfully coupled the $\gamma-\widetilde{Re}_{\theta_t}$ model to the Spalart-Allmaras turbulence model, however, their approaches have yet to be used for ASO. Indeed, all optimizations listed above that make use of the $\gamma-\widetilde{Re}_{\theta_t}$ approach have been coupled to the k-ω turbulence model as developed by Langtry and Menter.

Note that all but one of the high-fidelity optimizations above have been applied to airfoils in two-dimensions. The sole exception is the direct coupling of the e^N method with a RANS solver in the work of Lee and Jameson. In order to account for three-dimensional features such as sweep angle, taper ratio, and twist [54], there is a clear need for continued development of NLF design tools applicable to full aircraft configurations in three dimensions. Ultimately, such design tools should include the prediction of transition due to crossflow instabilities and they should incorporate the transition criteria into the evaluation of the objective, constraints and gradients of the design problem. This thesis will attempt to address some, but not all, of the above challenges. The specific scope of this work is described in the next section.

1.4 Summary of Objectives

The objective of this thesis is to advance the state-of-the-art in high-fidelity aerodynamic shape optimization by incorporating and exploiting laminar-turbulent transition prediction in an accurate and efficient manner. This work considers two-dimensional flow fields involving airfoil sections, with transition prediction capable of accounting for the effects of Re, T_{u}, M, and the pressure gradients. The focus of the optimizations will be on lift-constrained drag minimization at cruise conditions. The three major research milestones required to achieve the proposed objective are as follows:

1. **Extension of a two-dimensional RANS solver making use of the Spalart-Allmaras turbulence model to incorporate laminar-turbulent transition prediction.** The aim is not to develop a new transition prediction criterion, but rather to create a flexible and efficient framework for the incorporation of existing transition prediction criteria into a RANS flow solver.

2. **Development of a new discrete-adjoint formulation for non-local correlation-based transition prediction criteria.** The goal is to compute the adjoint gradients such that they incorporate the sensitivities of the transition prediction methodology. This will allow the optimizer to exploit transition prediction during the optimization cycle. A significant effort is required to ensure high levels of accuracy and efficiency of the adjoint gradient computations.

3. **Application of the ASO framework to practical subsonic and transonic flight conditions, ranging from the Cessna 172 to the Boeing 737.** The new framework will be employed to study various design trade-offs between fully-turbulent and free-transitional design spaces, which can be used to inform conservative decision-making.

The framework will also be extended to include a weighted multipoint optimization capability that facilitates the design of NLF airfoils that perform well over a given cruise flight envelope and are robust to perturbations. Finally, we wish to examine the sensitivity and explore the robustness of NLF design to the freestream disturbance environment.

1.5 Thesis Outline

The organization of the thesis is as follows. In Chapter 2, the baseline RANS solver is briefly presented, and the key elements of the algorithm are highlighted as a precursor to Chapter 3, in which the necessary extensions to incorporate a flexible and automatic transition prediction framework are presented. In Chapter 4, we present the gradient-based optimization techniques employed, with emphasis on the modifications required to incorporate the sensitivities of the transition prediction criteria. In Chapter 5, a detailed verification and validation is carried out on both the flow solver and optimization algorithms by comparison to available numerical and experimental data. In Chapter 6, to demonstrate the capabilities of the proposed optimization framework for the purpose of NLF airfoil design, various single and multipoint optimizations are performed at conditions associated with subsonic and transonic commercial aircraft. Finally, in Chapter 7 we close with contributions, conclusions, and recommendations for future work.

Chapter 2

Numerical Solution of the Navier-Stokes Equations

In this chapter we begin with a presentation of the governing equations for aerodynamic flows, considering specifically the Reynolds-averaged Navier-Stokes equations with the Spalart-Allmaras turbulence model. Next, an overview of the implicit finite-difference algorithm used in this work for the solution of the steady, fully-turbulent, RANS equations is presented. The presentation is short and emphasis throughout is on the features that are necessary and relevant to the development in the following chapters, which are concerned with the extensions required to incorporate transition prediction into the flow solver (Chapter 3) and the optimization framework (Chapter 4).

2.1 Governing Equations

2.1.1 The Navier-Stokes Equations

The behaviour of a compressible viscous fluid is governed by the conservation of mass, momentum, and energy [71]. These well known governing equations are collectively referred to as the Navier-Stokes equations. The Navier-Stokes equations form a coupled system of nonlinear partial differential equations (PDEs) that provides a mathematical description for the manner in which the basic fluid properties relate and evolve in both space and time. The following is a presentation of the two-dimensional form of the Navier-Stokes equations, where the *Reynolds-averaging* procedure used to model turbulent flow is discussed in the next section.

The basic fluid properties are defined as the local density, $\tilde{\rho}$, velocities, $\tilde{u}$ and $\tilde{v}$, and total energy, $\tilde{e}$, all per unit of volume. In this work, the flow variables, along with the fluid's dynamic viscosity, $\tilde{\mu}$, are made nondimensional using the freestream density, $\tilde{\rho}_\infty$, sound speed,

$\tilde{a}_\infty$, and dynamic viscosity, $\tilde{\mu}_\infty$, as

$$\rho = \frac{\tilde{\rho}}{\tilde{\rho}_\infty}\,, \qquad u = \frac{\tilde{u}}{\tilde{a}_\infty}\,, \qquad v = \frac{\tilde{v}}{\tilde{a}_\infty}\,, \qquad e = \frac{\tilde{e}}{\tilde{\rho}_\infty \tilde{a}_\infty^2} \qquad \text{and} \qquad \mu = \frac{\tilde{\mu}}{\tilde{\mu}_\infty}\,. \tag{2.1}$$

The chord Reynolds number, $\mathcal{R}e_c$, is computed using $\tilde{a}_\infty$ as the reference velocity scale, along with the airfoil chord, c, as the reference length scale.

The conservative, nondimensional, partial differential form of the Navier-Stokes equations may be written in vector-notation as [96]

$$\frac{\partial Q_{\rm c}}{\partial t} + \frac{\partial E}{\partial x} + \frac{\partial F}{\partial y} = \frac{1}{\mathcal{R}e_c}\left(\frac{\partial E_v}{\partial x} + \frac{\partial F_v}{\partial y}\right)\,, \tag{2.2}$$

where the vector $Q_{\rm c}$ represents the vector of conserved, nondimensional, dependent variables given by

$$Q_{\rm c} = \begin{pmatrix} \rho & \rho u & \rho v & e \end{pmatrix}^T\,. \tag{2.3}$$

The inviscid flux vectors in (2.2) are given by

$$E = \begin{pmatrix} \rho u \\ \rho u^2 + p \\ \rho u v \\ u(e+p) \end{pmatrix} \qquad \text{and} \qquad F = \begin{pmatrix} \rho v \\ \rho v u \\ \rho v^2 + p \\ v(e+p) \end{pmatrix}\,, \tag{2.4}$$

and the viscous flux vectors are given by

$$E_v = \begin{pmatrix} 0 \\ \tau_{xx} \\ \tau_{xy} \\ \varphi_1 \end{pmatrix} \qquad \text{and} \qquad F_v = \begin{pmatrix} 0 \\ \tau_{xy} \\ \tau_{yy} \\ \varphi_2 \end{pmatrix}\,, \tag{2.5}$$

where

$$\begin{aligned} \tau_{xx} &= (\mu + \mu_t)(4u_x - 2v_y)/3\,, \\ \tau_{xy} &= (\mu + \mu_t)(u_y - v_x)\,, \\ \tau_{yy} &= (\mu + \mu_t)(-2u_x + 4v_y)/3\,, \\ \varphi_1 &= u\tau_{xx} + v\tau_{xy} + (\mu \mathcal{P}r^{-1} + \mu_t \mathcal{P}r_t^{-1})(\gamma - 1)^{-1}\partial_x(a^2)\,, \ \text{and} \\ \varphi_2 &= u\tau_{xy} + v\tau_{yy} + (\mu \mathcal{P}r^{-1} + \mu_t \mathcal{P}r_t^{-1})(\gamma - 1)^{-1}\partial_y(a^2)\,. \end{aligned} \tag{2.6}$$

In (2.6), the laminar and turbulent Prandtl numbers are taken as $\mathcal{P}r=0.72$ and $\mathcal{P}r_t=0.90$, respectively, and the ratio of specific heats is taken as $\gamma=1.4$ for air. This work considers only polytropic gases (that is, calorically perfect gases) obeying the ideal-gas law and allowing us to relate the nondimensionalized pressure, p, to the flow variables as

$$p = (\gamma - 1)\left[e - \frac{1}{2}\rho(u^2 + v^2)\right] . \tag{2.7}$$

The dynamic viscosity of the fluid, μ, is given by Sutherland's law as

$$\mu = \frac{a^3(1 + S^*/T_\infty)}{a^2 + S^*/T_\infty} , \tag{2.8}$$

where T_∞ is the freestream temperature, in °R, and S^* is Sutherland's constant taken as 198.6°R. Finally, the variable μ_t represents the *turbulent eddy viscosity*, which is nondimensionalized as

$$\mu_\mathrm{t} = \tilde{\mu}_\mathrm{t}/\tilde{\mu}_\infty . \tag{2.9}$$

In the next section, we discuss the assumptions and turbulence modelling employed to compute μ_t.

2.1.2 Turbulence Modelling

The time and length scales exhibited by turbulent flow lend themselves to a time-averaging procedure in which the chaotic turbulent fluctuations of the flow are not resolved. Instead, a time interval over which flow quantities are averaged is selected to be large enough to replace the turbulent fluctuations with their effect on the mean flow, but small enough to capture other temporal features of the flow that are not directly associated with the turbulent fluctuations [96]. This averaging procedure – when density-weighted for compressible flows – is formally referred to as *Favre-averaging* [134], but more commonly (and henceforth) referred to as *Reynolds-averaging*.

The procedure has two main consequences. First, the dependent variables in (2.1)–(2.8) no longer represent their instantaneous values, but rather their time-averaged counterparts. Second, a new tensor quantity appears in the viscous fluxes of the governing equations referred to as the *Reynolds stress tensor*, along with some additional heat flux terms [96]. In this work, the Reynolds stress tensor is modelled through a *turbulent eddy viscosity*, μ_t. The eddy-viscosity concept, first introduced by Boussinesq in 1877 [134], is based on the assumption that the Reynolds stress tensor relates to μ_t in a manner analogous to the way

the shear stress tensor relates to μ, the molecular viscosity of the fluid. By invoking the Boussinesq assumption, all that remains to close the RANS equations is a means by which to compute the turbulent eddy viscosity.

In this work, the turbulent eddy viscosity is computed using the well-known one-equation Spalart-Allmaras (SA) turbulence model [112]. The SA model has been extensively validated and optimized for aerodynamic flows [134]. This semi-empirical model requires the solution to a differential transport equation that is tightly coupled to the conservation (or mean-flow) equations. The differential transport equation is given by

$$\frac{D\tilde{\nu}}{Dt} = \frac{c_{b1}}{\mathcal{R}e}\left(1 - f_{t2}\right)\tilde{S}\tilde{\nu} + \frac{1}{\sigma\mathcal{R}e}\left\{\left(1 + c_{b2}\right)\nabla\cdot\left[\left(\nu + \tilde{\nu}\right)\nabla\tilde{\nu}\right] - c_{b2}\left(\nu + \tilde{\nu}\right)\nabla^2\tilde{\nu}\right\} - \frac{1}{\mathcal{R}e}\left(c_{w1}f_w - \frac{c_{b1}}{\kappa^2}f_{t2}\right)\left(\frac{\tilde{\nu}}{d_{\mathrm{w}}}\right)^2 + \mathcal{R}e f_{t1}\Delta U^2\,, \quad (2.10)$$

where $\tilde{\nu}$ is a nondimensional turbulence-like variable referred to herein as the *working variable*. The working variable, $\tilde{\nu}$, is related to the turbulent eddy viscosity by

$$\mu_{\mathrm{t}} = \rho\tilde{\nu}f_{v1}\,, \quad (2.11)$$

where

$$f_{v1} = \frac{\chi^3}{\chi^3 + c_{v1}^3}\,, \quad (2.12)$$

and

$$\chi = \frac{\tilde{\nu}}{\nu}\,. \quad (2.13)$$

The first term on the right-hand side of (2.10) is the production term, and it is given by

$$\tilde{S} = S\,\mathcal{R}e + \frac{\tilde{\nu}}{\kappa^2 d_{\mathrm{w}}^2}f_{v2}\,, \quad (2.14)$$

where

$$S = \left|\frac{\partial v}{\partial x} - \frac{\partial u}{\partial y}\right| \quad (2.15)$$

is the magnitude of the vorticity, d_{w} is the distance to the closest wall and

$$f_{v2} = 1 - \frac{\chi}{1 + \chi f_{v1}}\,. \quad (2.16)$$

The destruction function in (2.10) is given by

$$f_w = g\left[\frac{1+c_{w3}^3}{g^6+c_{w3}^6}\right]^{\frac{1}{6}} , \tag{2.17}$$

where

$$g = r + c_{w2}(r^6 - r) , \tag{2.18}$$

and

$$r = \frac{\tilde{\nu}}{\tilde{S}\kappa^2 d_{\mathrm{w}}^2} . \tag{2.19}$$

Note that functions denoted f_{t1} and f_{t2} in (2.10) allow the user to specify a *fixed* or *forced* transition location in the model. For fully-turbulent flow, these functions are simply set to zero. To be clear, the SA model cannot itself predict the laminar-turbulent transition locations [112]. In fact, Spalart and Allmaras state that, "On no account should the turbulence model be trusted to predict the transition location" [112]. We shall postpone the description of the trip functions for now, as they are better suited to our discussion of transition region models in Section 3.1. The remaining parameters, which are described in more detail by Spalart and Allmaras [112], are given by

$$\begin{aligned}
&c_{b1} = 0.1355 \qquad c_{b2} = 0.622 \qquad \kappa = 0.41 \qquad \sigma = \frac{2}{3} \\
&c_{w1} = \frac{c_{b1}}{\kappa^2} + \frac{(1+c_{b2})}{\sigma} \\
&c_{w2} = 0.3 \qquad c_{w3} = 2.0 \qquad c_{v1} = 7.1 \qquad c_{v2} = 5.0 .
\end{aligned}$$

Finally, we note that the addition of the SA turbulence model has increased the number of unknown dependent variables at each node by one. Henceforth, the vector of unknown dependent variables, Q, shall include the mean-flow variables as well as the SA working variable, such that $Q = [\, Q_{\mathrm{c}} \,;\, \tilde{\nu} \,]$, or

$$Q = \begin{pmatrix} \rho & \rho u & \rho v & e & \tilde{\nu} \end{pmatrix}^T . \tag{2.20}$$

2.2 Flow Solution Algorithms

2.2.1 Overview of the Implicit Finite-Difference Algorithm

This work makes use of a baseline implicit finite-difference algorithm developed and implemented by Pueyo and Zingg [95] and Nemec and Zingg [86, 87], which is heavily based on ARC2D [96]. The following list serves to highlight the key elements of the algorithm, with the details found in a recent textbook by Pulliam and Zingg [96], in combination with Nemec's PhD Thesis [85]:

1) **Structured grid topologies** are used to discretize the two-dimensional solution domain. While this work makes use of the C-type grid topology around single-element airfoils, the baseline algorithm is readily applied to other structured topologies, such as H-topology grids around multi-element airfoils [85]. In this work, a bold-faced vector denotes one that spans the set of discrete nodes in a structured grid. For example, $\mathbf{Q}$ contains $Q_{j,k}$ at each node (indexed by j and k) such that $\mathbf{Q} \in \mathbb{R}^{N_\mathrm{Q}}$, where $N_\mathrm{Q} = 5N_\mathrm{G}$, and N_G is the total number of nodes in the grid.

2) A **generalized curvilinear coordinate transformation** is used to map the physical grid (a body-fitted mesh) onto a rectilinear grid in computational space. The transformation is such that the new curvilinear coordinate directions, denoted ξ and η, follow the grid lines in the streamwise and off-wall directions, respectively. In computational space the nodes are equally spaced, such that $\Delta\xi = \Delta\eta = 1$, which is amenable to the application of finite-difference operators. For a stationary grid that does not vary with time, the curvilinear coordinates may be written as functions of the Cartesian coordinates, $\xi = \xi(x, y)$ and $\eta = \eta(x, y)$. The relationship between their spatial derivative operators is given by

$$\begin{bmatrix} \partial_x \\ \partial_y \end{bmatrix} = \begin{bmatrix} \xi_x & \eta_x \\ \xi_y & \eta_y \end{bmatrix} \begin{bmatrix} \partial_\xi \\ \partial_\eta \end{bmatrix} , \tag{2.21}$$

using subscript notation of the form, $\partial_x \equiv \frac{\partial}{\partial x}$. The metric relations of the transformation are given by

$$\xi_x = Jy_\eta \, , \; \xi_y = -Jx_\eta \, , \; \eta_x = -Jy_\eta \, , \text{ and } \; \eta_y = Jx_\xi \, , \tag{2.22}$$

where the metrics x_ξ, x_η, y_ξ, and y_η are approximated using second-order centered differences in computational space [96]. Finally, the metric Jacobian of the transformation is given by

$$J = (x_\xi y_\eta - x_\eta y_\xi)^{-1}. \tag{2.23}$$

3) A **thin-layer approximation** is invoked to simplify the Navier-Stokes equations for fully-attached or mildly-separated flows at high Reynolds numbers. The viscous fluxes in the off-wall η-direction are included in the Navier-Stokes equations, while the viscous fluxes in the ξ-direction are ignored.

4) **Second-order centered-differences** are used to approximate the spatial derivatives of the governing equations written in generalized curvilinear coordinates.

5) Numerical stability is achieved by an **artificial dissipation model** using either the scalar dissipation scheme of Jameson *et al.* [52] or the matrix dissipation scheme of Swanson and Turkel [123]. The effect this has on the laminar boundary layer will be discussed in 5.4. To avoid the oscillations near discontinuities, a **pressure sensor** is employed to detect shock-waves (or otherwise unresolved flow features) and trigger the stabilizing second-difference dissipation term [96].

6) The **boundary conditions** are imposed at the farfeild boundaries using locally one-dimensional Riemann invariants, and at the surface by imposing the no-slip condition, adiabatic flow, and a zero wall-normal pressure gradient [85].

7) An **implicit Euler time-marching** algorithm is used to time-step the initial solution state toward the steady-state solution. This start-up algorithm makes use of approximate factorization, local-time stepping, and the diagonal form to accelerate convergence [96].

8) An **inexact Newton-Krylov method** [86, 87] is then used for efficient and deep convergence of the steady-state solution. A sufficiently converged starting iterate for the inexact Newton method is provided by the implicit Euler start-up phase.

Of the elements in the preceding list, the last two will be discussed further in the next section, as they describe the manner in which the algorithm iterates to steady-state. For completeness, the thin-layer Navier-Stokes equations in generalized curvilinear coordinates are presented in Appendix A, where the notation is such that the transformed vectors in the governing equations are denoted using an *over-hat*. For example, the vector of unknown dependent variables (including the conserved flow variables and SA working variable) is given at each node by

$$\widehat{Q} \equiv J^{-1} \left(\begin{array}{ccccc} \rho & \rho u & \rho v & e & \tilde{\nu} \end{array} \right)^T . \tag{2.24}$$

2.2.2 Convergence to a Fully-Turbulent Steady State

The spatially discretized RANS equations may be written in a concise form as

$$\frac{d\widehat{\mathbf{Q}}}{dt} = -\boldsymbol{\mathcal{R}}(\widehat{\mathbf{Q}}) \,, \tag{2.25}$$

where $\boldsymbol{\mathcal{R}}(\widehat{\mathbf{Q}}) \in \mathbb{R}^{N_\mathrm{Q}}$ is known as the discrete flow residual vector, or simply the residual, which includes the right-hand sides of the mean flow and turbulence transport equations at each node, see Appendix A. Given that we are only interested in steady-state solutions, the governing equations reduce to

$$\boldsymbol{\mathcal{R}}(\widehat{\mathbf{Q}}) = 0 \,. \tag{2.26}$$

Nemec and Zingg developed an algorithm based on Newton's method, specifically, an inexact Newton-Krylov (NK) method [86, 87] for the solution of (2.26). Newton's method is applied using the local linearization of the flow residual about time level n, such that

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \widehat{\mathbf{Q}}}^{(n)} \Delta\widehat{\mathbf{Q}}^{(n)} = -\boldsymbol{\mathcal{R}}^{(n)} \,, \tag{2.27}$$

where the superscript denotes the time level, $\boldsymbol{\mathcal{R}}^{(n)}$ is the residual evaluated using $\widehat{\mathbf{Q}}^{(n)}$, and $\Delta\widehat{\mathbf{Q}}^{(n)} = \widehat{\mathbf{Q}}^{(n+1)} - \widehat{\mathbf{Q}}^{(n)}$. The matrix on the left-hand side of (2.27) is called the exact flow Jacobian, or the $\mathcal{A}_2$-Jacobian $\in \mathbb{R}^{N_\mathrm{Q} \times N_\mathrm{Q}}$.

The NK algorithm involves two-levels of iteration, referred to as *outer* and *inner* iterations. At each outer iteration, the solution state, $\widehat{\mathbf{Q}}$, is updated by solving (2.27) for $\Delta\widehat{\mathbf{Q}}^{(n)}$. Note that (2.27) represents a large sparse system of linear algebraic equations. This system is solved to a desired tolerance using a preconditioned Generalized Minimum Residual (GMRES) method. The iterations of the GMRES solver are the *inner iterations* of the NK solver. Note that in the larger context of aerodynamic shape optimization, we will see the same GMRES method applied to the evaluation of the flow sensitivity and adjoint gradients evaluations (Chapter 4). The GMRES algorithm is a Krylov subspace projection method [101]. The algorithm consists of four main steps: initialization, Krylov subspace orthogonalization, solving a least-squares problem, and updating the solution [36]. Additional details on the GMRES method may be found in Saad [101]. Practical aspects of its implementation in the context of this work may be found in Gatsis's PhD thesis [36].

The solution procedure in GMRES requires only matrix-vector products, so it is not necessary to explicitly form the $\mathcal{A}_2$-Jacobian; the products may be formed using the Fréchet derivative approximation [85]. Furthermore, to improve conditioning, the linear system is

right-preconditioned using an incomplete LU factorization with a level of fill equal to 2. The preconditioner is obtained using an approximation to the $\mathcal{A}_2$-Jacobian, referred to as the first-order flow Jacobian matrix, or the $\mathcal{A}_1$-Jacobian $\in \mathbb{R}^{N_\mathrm{Q} \times N_\mathrm{Q}}$. Approximations are made in order to reduce the bandwidth of the preconditioner and circumvent the analytic differentiation of some complex terms, such as the pressure sensor and matrix dissipation model. This has been shown by Nemec and Zingg [86, 87] to significantly reduce memory requirements while maintaining a robust and efficient solution procedure. Note that these approximations have no effect on the final steady-state solution satisfying (2.26). Also note that at each outer iteration we solve (2.27) inexactly by reducing the L_2-norm of the algebraic system by one order of magnitude. This has the desirable effect of reducing the number of inner iterations required to compute $\Delta\widehat{\mathbf{Q}}^{(n)}$ for each Newton update while retaining a desirable convergence rate [85].

The NK method (or outer iterations) is considered to be steady-state converged when the L_2-norm of the flow residual, $||\boldsymbol{\mathcal{R}}||$, has decreased below a specified tolerance, ϵ_nk. In this work we take $\epsilon_\mathrm{nk} = 10^{-11}$ to ensure that tight convergence is achieved, as required for the adjoint gradient formulation presented in Chapter 4.

Finally, we note that the convergence of Newton's method often relies on an initial guess that is sufficiently close to the final solution [96]. This in turn necessitates a *globalization* strategy to bring the initialized solution state (which we take everywhere to be the freestream conditions) to a sufficiently accurate initial iterate for the NK solver. In this work, globalization is achieved by making use of an implicit Euler time-marching method with approximate factorization (AF) and local time-stepping to accelerate convergence. The details of the AF algorithm and local-time stepping may be found in [96]. When the L_2-norm of the flow residual has decreased below a specified tolerance of $\epsilon_\mathrm{af} = 10^{-6}$, the switch to the inexact Newton phase is made.

Chapter 3

Transition Modelling and Prediction in a RANS Solver

The new components required for laminar-turbulent transition prediction in the RANS solver described in the previous chapter are now developed. There are five such components that augment the RANS solver: (i) the transition region model; (ii) the determination of the boundary-layer edge; (iii) the calculation of the boundary-layer properties; (iv) the evaluation of existing transition prediction criteria; and (v) the iterative transition prediction procedure. Each will be presented in sequence from Sections 3.1 to 3.5, with two sections that follow to discuss matters of computational efficiency and robustness.

3.1 Transition Region Models

The goal of the transition region model is to facilitate a ramp-up of the turbulent eddy-viscosity from its value in the laminar region of the boundary layer to the turbulent region. *It is not the role of the transition region model to predict the location of transition.* Ultimately the transition region model takes its cue from either user-specified fixed-transition locations on the airfoil, or a transition prediction criterion. It should be noted that no attempt is made to capture the nonlinear interactions or any of the underlying physics of the transition process itself. Indeed this is not possible with the time-averaged approach of a steady RANS solver.

Two transition region models have been included and studied in this work. The first makes use of the SA model trip functions and the second is a newly defined S-shaped intermittency function. The former modifies the eddy-viscosity from within the turbulence model itself, while the latter scales the eddy-viscosity directly. In Chapter 5, a verification of their ability to control the eddy-viscosity is presented along with a comparison of their ramp-up behaviour and the corresponding impact on the design space.

3.1.1 Spalart-Allmaras Trip Functions

The trip functions that appear in the production, destruction and transition source terms of the SA transport equation (2.10) are defined as

$$f_{t1} = c_{t1} g_t \exp\left[-c_{t2}\frac{\omega_t^2}{\Delta U^2}\left(d^2 + g_t^2 d_t^2\right)\right] , \tag{3.1}$$

$$f_{t2} = c_{t3} \exp\left(-c_{t4}\chi^2\right) , \tag{3.2}$$

where d_t is the distance to the transition point, ω_t is the vorticity at the transition point, ΔU is the difference between the velocity magnitude at the transition point (taken as zero at the surface of the airfoil) and the velocity magnitude of the field point under consideration, and

$$g_t = \min\left(0.1, \frac{|\Delta U|}{\omega_t \Delta x}\right) , \tag{3.3}$$

where Δx is the grid spacing along the wall at the transition point [112]. The remaining constants are given by

$$c_{t1} = 5 , \qquad c_{t2} = 2 , \qquad c_{t3} = 1.2 , \qquad \text{and} \qquad c_{t4} = 0.5 .$$

The role of the f_{t1} function in the transition (or trip) source term, given as $\mathcal{R}e f_{t1}\Delta U^2$ in (2.10), is to ramp up the production of eddy-viscosity in a zone of influence near the transition point. In the definition of f_{t1}, the goal of the g_t factor is to ensure that the source term is felt over at least a few streamwise stations [112]. However, the use of Δx in (3.3) has the drawback of introducing grid-dependent solutions, which is generally undesirable. The role of the f_{t2} function is to alter the production and destruction terms to ensure stable solutions with the small values of $\tilde{\nu}$ present in the laminar region. Finally, it should be noted that the SA trip functions *do not* allow for the specification or control of a transition length, that is, the streamwise distance over which the eddy-viscosity is increased from laminar to turbulent values. This limitation can be problematic and lead to sharp increases in eddy-viscosity, as demonstrated and discussed in Section 5.1.

3.1.2 Intermittency Function

The second transition region model considered makes use of a streamwise varying *intermittency function*. The use of the term *intermittency* is not to be confused with the intermittency that is physically observed at the edge of the boundary layer as one moves in the off-wall

direction into the freestream; the off-wall intermittency is accounted for by the destruction term in the SA model, which depends on an off-wall distance calculation [112].

Following the work of Cliquet *et al.* [23], the intermittency function, γ, scales the turbulent eddy-viscosity directly from the SA model (or any other eddy-viscosity model) such that

$$\mu_{\mathrm{tr}} = \max(\, \gamma\mu_{\mathrm{t}} \,,\, \mu_{\mathrm{t}_\infty} \,) \;, \quad \text{where} \quad 0 \leq \gamma \leq 1 \,. \tag{3.4}$$

The max function in (3.4) ensures that the scaled eddy-viscosity, μ_{tr}, does not take on values less than the freestream value in the laminar regions.

In this work, a new form of the streamwise intermittency function has been designed based on a different function developed by Dhawan and Narasimha [20]. The new form of the S-shaped function is computed as a function of the local chord position, x, the specified transition point, x_{tr}, and the specified transition length, l_{tr} as

$$\gamma(x, x_{\mathrm{tr}}, l_{\mathrm{tr}}) = \exp(-5\,\xi^2) \,, \quad \text{where} \quad \xi = 1 + \frac{x_{\mathrm{tr}} - x}{l_{\mathrm{tr}}} \,, \tag{3.5}$$

whereas Dhawan and Narashimha's original intermittency function takes the form

$$\gamma(x, x_{\mathrm{tr}}, l_{\mathrm{tr}}) = 1 - \exp(0.412\,\zeta^2) \,, \quad \text{where} \quad \zeta = 3.36\,\frac{x - x_{\mathrm{tr}}}{l_{\mathrm{tr}}} \,. \tag{3.6}$$

Figure 3.1 provides a plot of the intermittency distributions with a specified transition location at 30% chord and a specified transition length of 10%. By design, the intermittency functions asymptote to a value of zero in the fully-laminar region and to unity in the fully-turbulent region. The present form of the intermittency given by (3.5) is designed to introduce a more gradual initial ramp-up of the eddy-viscosity, and is very close to a horizontally mirrored image of (3.6). The importance of the smooth initial ramp-up is two-fold. First, it helps to avoid oscillations that can appear in the flow variables just upstream of the transition point due to an abrupt increase in μ_{tr}. These oscillations are often observed in the pressure and skin friction profiles and sometimes lead to unphysical or unresolved flow separation (negative skin friction coefficients at a single node) on the airfoil. Second, the smooth ramp-up ensures a sufficiently smooth design space as required for gradient-based optimization (to be further discussed in Section 5.1).

The intermittency transition region model is particularly easy to implement since the only modification necessary is to replace the variable μ_{t} with its scaled equivalent, μ_{tr}, everywhere that it appears in the mean-flow equations (2.6). Note that it is not necessary to modify the turbulence model explicitly in any way, instead the solution to $\tilde{\nu}$ implicitly accounts for

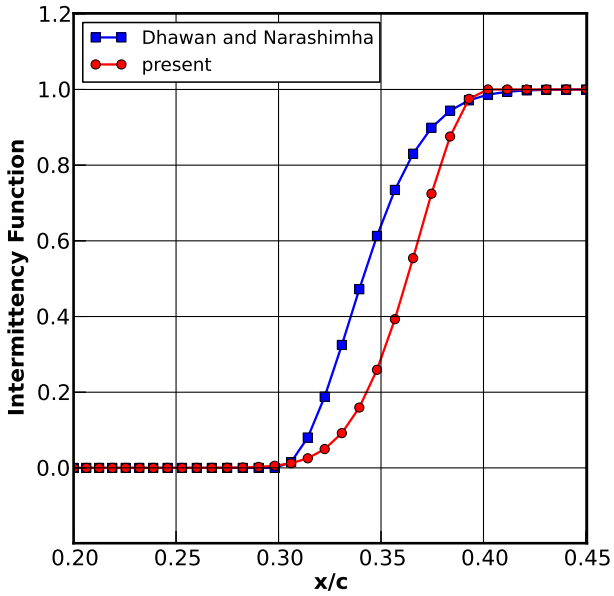

Figure 3.1: Intermittency distributions with $x_{\mathrm{tr}} = 0.3$ and $l_{\mathrm{tr}} = 0.1$

the intermittency scaling through its coupling to the mean-flow variables as convergence to steady-state is achieved. Other approaches to applying an intermittency function include scaling the production term of the SA model [62], or scaling the flow vorticity prior to its use in the SA model [78]. While both of these approaches have been implemented in the present flow solver, numerical experiments have revealed them to be generally less stable and robust. As such, the scaling approach of (3.4) shall be used in this work.

As a final note, the structured grids generated for transitional flows in this work have all been designed and smoothed for orthogonality with respect to the wall, in turn creating near-normal grid lines. As such, the numerical implementation of the intermittency scaling does not require interpolation, but instead consists of a constant γ value being applied to each η-direction grid line. In the next two sections, we discuss the analysis of the laminar boundary layer (which is formed upstream of the transition region).

3.2 Determination of the Boundary-Layer Edge

There are two main challenges when using RANS solvers to directly evaluate a transition criterion (approaches 1 to 3 discussed in Section 1.3). The first relates to defining the boundary-layer edge directly from the Navier-Stokes solution, and the second is to sufficiently resolve and accurately compute the laminar boundary-layer properties. The second challenge is addressed in Sections 3.3 and 5.4.

For arbitrary pressure distributions the boundary-layer edge velocity, U_e, is not known *a priori*. Hence, we must somehow define the boundary-layer edge based on the Navier-Stokes solution. An alternative approach is to couple the RANS solver to a boundary-layer solver that provides a boundary-layer thickness as part of the solution process [30, 64, 65, 69]. The main benefit of using a boundary-layer solver is in the accurate calculation of the boundary-layer thickness while allowing for coarser grids to be used in the RANS solver [64]. The disadvantages of the coupling approach include: (i) redundancy in solving both the Navier-Stokes and boundary-layer equations, (ii) restriction to fully-attached and mildly separated flows, (iii) implementation and convergence issues in coupling the two solvers, (iv) complexities when moving to compressible flows, wings of finite span, and parallel implementations, and (v) complexities in differentiating the boundary-layer solver when forming the Jacobian matrix (see Section 2.2). For these reasons, this work makes use of the available RANS solution, altogether avoiding the use of a boundary-layer solver.

We have implemented and compared three boundary-layer edge-finding methods. A verification, comparison and assessment of their accuracy can be found in Section 5.3. While we are only interested in computing the properties of the laminar boundary-layer, all three methods are capable of finding the boundary-layer edge in both the laminar and turbulent regions. The methods are summarized as follows:

Compressible Bernoulli Equation: Following Nebel *et al.* [83], the first method makes use of the local wall pressure, p_w, to approximate the edge velocity via the Bernoulli equation for compressible flows,

$$U_e = \sqrt{U_\infty^2 - \frac{2\gamma}{\gamma - 1}\frac{p_\infty}{\rho_\infty}\left[\left(\frac{p_w}{p_\infty}\right)^{\frac{\gamma-1}{\gamma}} - 1\right]}\,. \tag{3.7}$$

The boundary-layer thickness for the given streamwise station is then searched in the off-wall direction for the point $\delta = y$ at which $U = 0.99U_e$, where y is used here to denote the normal off-wall distance.

Baldwin-Lomax Diagnostic Function: Following Stock and Haase [114] and Nebel *et al.* [83], the second method makes use of a so-called diagnostic function derived from the Baldwin-Lomax turbulence model. The diagnostic function,

$$F = y^a \left[\frac{dU}{dy}\right]^b, \tag{3.8}$$

is first computed and its maximum value in the off-wall direction is found. The boundary-layer thickness is then computed as $\delta = \epsilon \cdot y_{max}$, where $y_{max} = y$ at which $F = F_{max}$. The values of

the constants have been determined through numerical and experimental investigation to be $a_\mathrm{l}=3.9$, $b_\mathrm{l}=1.0$, and $\epsilon_\mathrm{l}=1.294$ for laminar boundary layers, and $a_\mathrm{t}=1.0$, $b_\mathrm{t}=1.0$, and $\epsilon_\mathrm{t}=1.936$ for turbulent boundary layers [83].

Vorticity and Shear-Stress Method: Following Cliquet and Arnal [23], the third method makes use of the local vorticity, Ω, and an approximation to the total shear stress, τ_tot. The total shear stress is defined as $\tau_\mathrm{tot}=\tau_\mathrm{l}+\tau_\mathrm{t}$, where τ_l and τ_t can be expressed in the following manner [23]:

$$\tau_\mathrm{l} = \mu|\Omega| \quad \text{and} \quad \tau_\mathrm{t} = \mu_\mathrm{t}|\Omega| \,. \tag{3.9}$$

The boundary-layer thicknesses based on vorticity and shear stress are then searched in the off-wall direction, such that

$$\delta_\Omega = y \quad \text{at which} \quad |\Omega| = \epsilon_\Omega \cdot |\Omega|_\mathrm{max} \,, \tag{3.10}$$

and

$$\delta_\tau = y \quad \text{at which} \quad |\tau_\mathrm{tot}| = \epsilon_\tau \cdot |\tau_\mathrm{tot}|_\mathrm{max} \,, \tag{3.11}$$

where $\epsilon_\Omega = 0.001$ and $\epsilon_\tau = 0.015$. Finally, the boundary-layer thickness is taken as the minimum, such that $\delta=\min(\delta_\Omega, \delta_\tau)$.

The numerical implementation of the above methods includes a limited off-wall search distance of $0.2c$ (where c is the chord length of the airfoil) to avoid unnecessary computations taken to the far-field boundary, typically a distance of approximately $20c$ from the airfoil. As an additional measure of protection, all methods have a fall-back criterion that assumes δ is reached if U decreases with increasing y [23]. With the boundary-layer edge defined, we can now calculate the various boundary-layer properties.

3.3 Calculation of the Boundary-Layer Properties

Many transition prediction criteria for wall-bounded flows make use of boundary-layer properties. For example, the AHD criterion requires the calculation of the displacement thickness (δ^*), momentum thickness (θ), shape-factor (H), and Pohlhausen number (Λ_2), all of which require the calculation of the boundary-layer thickness (δ). By definition, the various boundary-layer properties are non-local, since they require the integration of flow quantities from the wall to the boundary-layer edge.

As previously mentioned, the structured grids generated in this work have the advantage of near-normal grid lines relative to the airfoil surface, allowing us to integrate directly along

the lines of constant ξ (in the off-wall direction), without the need to interpolate between nodes in the streamwise direction. The boundary-layer properties are then stored at each streamwise station along the airfoil surface. If one is considering a grid that is unstructured, or with grid lines that are non-orthogonal to the surface, Cliquet *et al.* [23] have demonstrated that a viable alternative is to define normal *computation lines* stemming from stations along the airfoil surface; these are data containers that store interpolated solution data for the integration of the boundary-layer properties. Alternatively, Krumbein [62] has made use of hybrid grids with an unstructured grid in the freestream and a structured grid in the boundary layer, facilitating the straightforward integration of the boundary-layer properties.

The various boundary-layer properties required for the transition prediction criteria employed in this work are the *incompressible* properties, given as

$$\text{Displacement Thickness:} \quad \delta^*_{\text{inc}} = \int_o^\delta \left(1 - \frac{U}{U_\text{e}}\right) dn \tag{3.12}$$

$$\text{Momentum Thickness:} \quad \theta_{\text{inc}} = \int_o^\delta \frac{U}{U_\text{e}} \left(1 - \frac{U}{U_\text{e}}\right) dn \tag{3.13}$$

$$\text{Shape Factor:} \quad H_{\text{inc}} = \frac{\delta^*_{\text{inc}}}{\theta_{\text{inc}}} \tag{3.14}$$

$$\text{Pohlhausen Number:} \quad \Lambda_{2_{\text{inc}}} = \frac{\theta^2_{\text{inc}}}{\nu_\text{e}} \frac{dU_\text{e}}{ds}\,, \tag{3.15}$$

and their *compressible* counterparts, given as

$$\text{Displacement Thickness:} \quad \delta^* = \int_o^\delta \left(1 - \frac{\rho U}{\rho_\text{e} U_\text{e}}\right) dn \tag{3.16}$$

$$\text{Momentum Thickness:} \quad \theta = \int_o^\delta \frac{\rho U}{\rho_\text{e} U_\text{e}} \left(1 - \frac{U}{U_\text{e}}\right) dn \tag{3.17}$$

$$\text{Shape Factor:} \quad H = \frac{\delta^*}{\theta} \tag{3.18}$$

$$\text{Pohlhausen Number:} \quad \Lambda_2 = \frac{\theta^2}{\nu_\text{e}} \frac{dU_\text{e}}{ds}\,. \tag{3.19}$$

In the above equations, δ is the boundary-layer thickness determined by one of the three methods discussed previously, and the subscript e denotes the linearly interpolated values at the boundary-layer edge. Geometrically, the n-direction corresponds to the local off-wall direction, and the s-direction corresponds to the local streamwise direction along the airfoil surface. Following the work of Mayda [78], the incompressible and compressible displacement and momentum thicknesses are numerically integrated using the trapezoidal integration rule

as follows:

$$\delta^* \ , \ \delta^*_{\text{inc}} \ , \ \theta \ , \text{ and } \theta_{\text{inc}} \approx \left[\frac{1}{2} \sum_{k=2}^{k_\delta} \left(I_{k-1} + I_k \right) \left(n_k - n_{k-1} \right) \right] + \frac{1}{2} \left(I_\delta + I_{k_\delta} \right) \left(\delta - n_{k_\delta} \right) \ , \qquad (3.20)$$

where I is the discrete (pointwise) value of the integrand defined for the given boundary-layer property in (3.12), (3.13), (3.16) and (3.17). The subscripts denote the nodal index at which I is evaluated, where k is the nodal index in the off-wall direction, k_δ is the nodal index of the last node inside the boundary layer, and the subscript δ indicates a linearly interpolated value at the boundary-layer edge location (interpolated between the k_δ and $k_\delta + 1$ nodes).

The calculation of the Polhausen number requires the local gradient of the edge velocity in the streamwise direction. To compute this gradient, we make use of a suggestion by Cliquet *et al.* [23] to replace dU_e/ds with dp/ds. This eliminates any complications in computing the gradient tangential to the surface at the boundary-layer edge. The relationship between the velocity and pressure gradients is established by considering first-order boundary-layer theory, as follows:

$$\frac{dU_\text{e}}{ds} = -\frac{1}{\rho U_\text{e}} \left(\frac{dp}{ds} \right) \ . \qquad (3.21)$$

As a final note, we must consider a means by which to extrapolate the laminar boundary-layer properties into the turbulent region of the flow; for this we make use of a linear extrapolation. The reasons and details of the extrapolation procedure will be discussed in Section 3.7.

3.4 Evaluation of the Transition Prediction Criteria

In this section we present the two methods employed in this work for predicting the location of transition, namely, a simplified e^N envelope method by Drela and Giles [28], and the compressible form of the Arnal-Habiballah-Delcourt criterion [9].

3.4.1 Simplified e^N Envelope Method

Drela has developed CFD tools (named XFOIL and MSES [28]) that have been widely used for airfoil analysis and design. Both make use of inviscid-viscous coupling schemes. The former makes use of a panel code, the latter solves the Euler equations. Common to both solvers is the simplified e^N envelope method used to apply the e^N transition criterion. This is the first transition criterion implemented in the present RANS solver.

The simplified e^N envelope method makes direct use of the boundary-layer properties to approximate the envelope of the spatial amplification of the disturbances (the N-factors). The approach avoids the necessity to solve the linear stability equations, which can be a computationally expensive procedure. While the envelope method does not track individual frequencies, it is significantly more efficient [78]. The correlations that relate the boundary-layer properties to the N-factor are developed based on linear stability results for the Falkner-Skan family of velocity profiles. The envelopes of the amplification factors are locally approximated as straight lines with respect to the streamwise direction, s, as follows:

$$\frac{dN}{ds} = \mathrm{fcn}(H_\mathrm{k}, \theta) = \frac{dN}{dRe_\theta} \cdot \frac{m+1}{2} \cdot l \cdot \frac{1}{\theta}\,, \tag{3.22}$$

where $\frac{dN}{dRe_\theta}$, m, and l are functions of the so-called kinematic shape factor, H_k. The kinematic shape factor is computed based on the incompressible shape factor, H_inc, and the Mach number at the boundary-layer edge, M_e, as

$$H_\mathrm{k} = \frac{H_\mathrm{inc} - 0.290 M_\mathrm{e}^2}{1 + 0.113 M_\mathrm{e}^2}\,. \tag{3.23}$$

The N-factor envelope is then obtained by integrating (3.22) in the streamwise direction, beginning at the critical point. The approach is then a two-step method: first find the critical point, then integrate from the critical point to the present station to obtain the local N-factor. The critical point is the first point at which, locally, $Re_\theta = Re_{\theta_\mathrm{cr}}$, where Re_{θ_cr} is defined by

$$\log_{10} Re_{\theta_\mathrm{cr}} = \left(\frac{1.415}{H_\mathrm{k}-1} - 0.489\right) \tanh\left(\frac{20}{H_\mathrm{k}-1} - 12.9\right) + \frac{3.295}{H_\mathrm{k}-1} + 0.44\,. \tag{3.24}$$

The functions $\frac{dN}{dRe_\theta}$, m, and l are given as follows: [28]

$$\begin{aligned}
\frac{dN}{dRe_\theta} &= 0.01\sqrt{[2.4H_\mathrm{k} - 3.7 + 2.5\tanh(1.5H_\mathrm{k} - 4.65)]^2 + 0.25} \\
m(H_\mathrm{k}) &= \left(0.058\frac{(H_\mathrm{k}-4)^2}{H_\mathrm{k}-1} - 0.068\right)\frac{1}{l(H_\mathrm{k})} \\
l(H_\mathrm{k}) &= \frac{6.54H_\mathrm{k} - 14.07}{H_\mathrm{k}^2}
\end{aligned}$$

Downstream of the critical point, the N-factor envelope begins to increase from zero. Finally, the transition point is taken as the linearly interpolated streamwise location at which $N = N_\mathrm{crit}$, where N_crit is user-specified (typically it is experimentally correlated for a given application

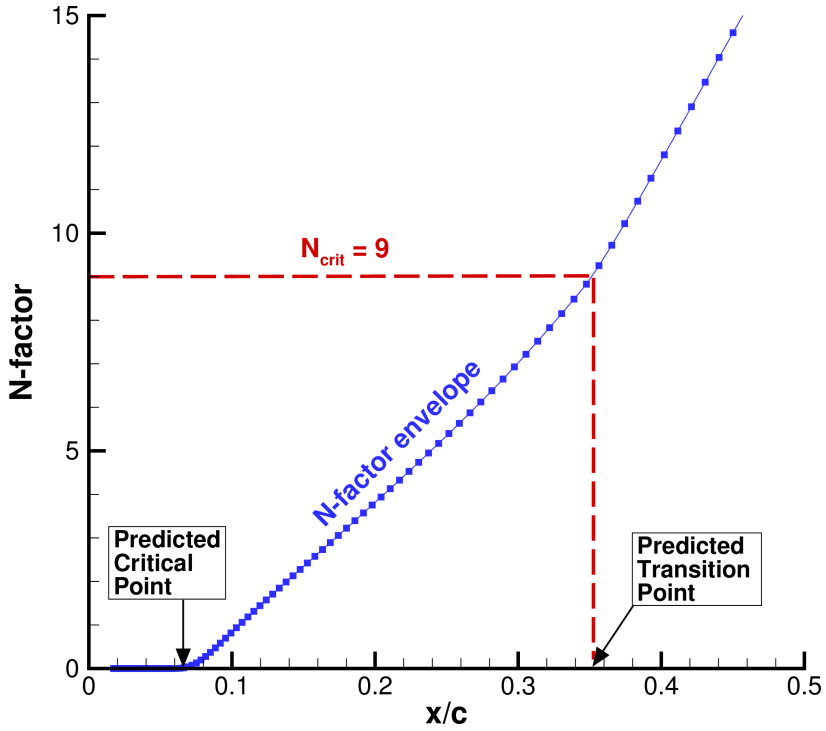

Figure 3.2: Demonstration of the application of the e^N transition criterion

and/or wind-tunnel). Figure 3.2 illustrates the application of the e^N criterion with $N_{\text{crit}} = 9$. Note that while the above envelope method does not predict crossflow transition, it has been combined with other N-factor transition criteria for such instabilities in three-dimensional flow [69].

3.4.2 Compressible Form of the AHD Criterion

A second criterion for predicting the natural transition locations due to Tollmien-Schlichting instabilities makes use of the original work by Arnal, Habiballah, and Delcourt [9, 23, 81, 117]. The Arnal-Habiballah-Delcourt (AHD) criterion is designed for low to moderate freestream turbulence intensities ($T_u \leq 0.1\%$), as typically encountered in external aerodynamic cruise conditions for transport aircraft [23]. The method has the advantage of being applicable to a wide range of pressure gradients, as well as compressible flows [23]. It does not predict transition due to crossflow instabilities; however, Arnal [9] has combined the AHD criterion with other criteria (such as the C1 criterion) for that purpose.

Beginning at the stagnation point, we march toward the trailing edge of the airfoil, treating the upper and lower surfaces independently. Our first task is to find the streamwise location of the neutral stability point, s_{cr}. Upstream of the neutral stability point, it is assumed (from

linear stability theory) that all small disturbances over all frequencies remain stable and damp out. The neutral stability point is found using the *critical* Reynolds number, calculated as a function of the incompressible shape factor, $H_{\text{inc}} = \delta^*_{\text{inc}}/\theta_{\text{inc}}$, as

$$Re_{\theta_{\text{cr}}} = \exp\left[\frac{E}{H_{\text{inc}}} - F\right], \tag{3.25}$$

such that s_{cr} is the first point at which, locally, $Re_\theta = Re_{\theta_{\text{cr}}}$. The incompressible shape factor is used even for compressible flow, as recommended by Cliquet *et al.* [23]. The functions E and F are computed as a function of the Mach number at the boundary-layer edge, as defined below. Note that $Re_{\theta_{\text{cr}}}$ typically decreases in the streamwise direction and is greater than Re_θ upstream of the critical point.

The next step is to find the streamwise location of the laminar-turbulent transition point, s_{tr}. The transition criterion is computed and checked only at points downstream of the neutral stability point, s_{cr}. The AHD criterion uses the Falkner-Skan self-similar solutions to represent the laminar boundary-layer profiles, which are characterized by the local Pohlhausen number [23]. Following the work of Granville, the necessary relationships are extended from self-similar boundary layers to actual flows by replacing Λ_2 with its mean value as follows [23]:

$$\Lambda_2 = \frac{\theta^2}{\nu}\frac{dU_{\text{e}}}{ds} \quad \Longrightarrow \quad \overline{\Lambda_2} = \frac{1}{s - s_{\text{cr}}}\int_{s_{\text{cr}}}^{s} \Lambda_2 ds\,. \tag{3.26}$$

Arnal *et al.* [23] proposed the following expression for the transitional Reynolds number, $Re_{\theta_{\text{tr}}}$:

$$Re_{\theta_{\text{tr}}} = Re_{\theta_{\text{cr}}} + A\cdot\exp(B\cdot\overline{\Lambda_2})\left[\ln(C\cdot T_{\text{u}}) - D\cdot\overline{\Lambda_2}\right], \tag{3.27}$$

where T_{u} is the freestream turbulence level, and the functions A through F are computed as a function of the Mach number at the boundary-layer edge, M_{e}, as follows: [9]

$$\begin{aligned}
A &= 98.64M_e^3 - 356.44M_e^2 + 117.13M_e - 236.69\\
B &= -13.04M_e^4 + 38.5M_e^3 - 30.07M_e^2 + 10.89M_e + 22.7\\
C &= 0.21M_e^3 + 4.79M_e^2 - 1.76M_e + 22.56\\
D &= -3.48M_e^4 + 6.26M_e^3 - 3.45M_e^2 + 0.23M_e + 12\\
E &= 0.6711M_e^3 - 0.7379M_e^2 + 0.167M_e + 51.904\\
F &= 0.3016M_e^5 - 0.7061M_e^4 + 0.3232M_e^3 - 0.0083M_e^2 - 0.1745M_e + 14.6
\end{aligned}$$

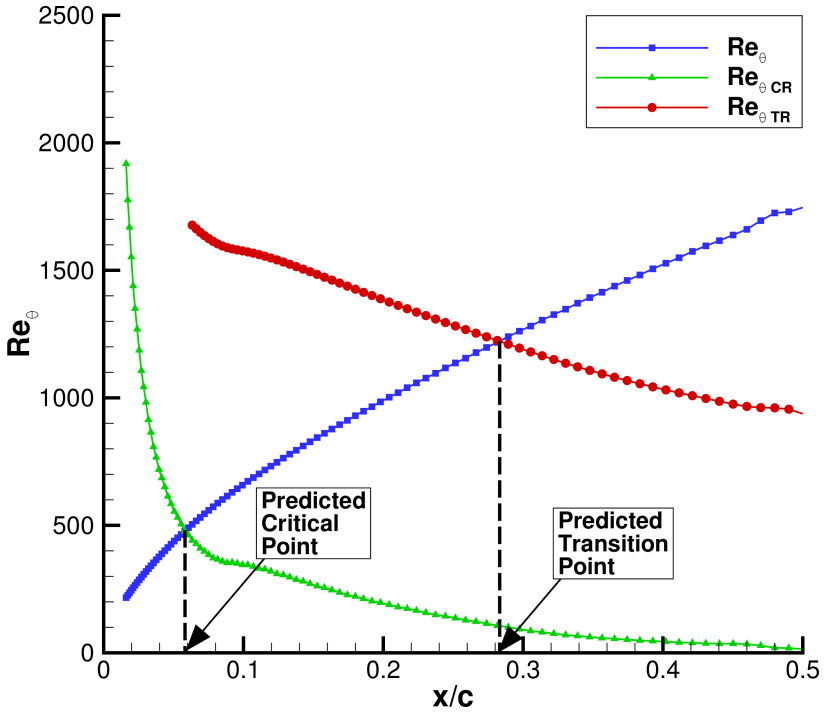

Figure 3.3: Demonstration of the application of the AHD criterion

The transition point is then taken as the streamwise location on the airfoil surface at which $Re_\theta = Re_{\theta_{tr}}$. This is computed by finding the intersection point of the two curves, using linear interpolation between streamwise stations. Figure 3.3 demonstrates the application of the AHD transition criterion. Note that $Re_{\theta_{tr}}$ typically decreases in the streamwise direction and is greater than Re_θ upstream of the transition point.

3.4.3 Additional Transition Criteria

One of the advantages of the proposed transition prediction framework is the relative ease with which different or new correlation-based transition criteria may be implemented and investigated. For example, two additional transition criteria have been implemented into the framework for testing purposes. The first is known as Granville's criterion, which is evaluated using a two-step method analogous to the AHD criterion (details may be found in Cebeci and Smith [20]). The second is referred to as the Modal Amplitude Tracking and Transition Computation (MATTC), which is a relatively new criterion developed by Campbell *et al.* [19]. MATTC has the benefit of providing N-factor envelopes without the need for boundary-layer properties.

Unfortunately, Granville's criterion is restricted to incompressible flow and is thus not suitable for the majority of applications of interest in this work. Furthermore, the numerical

verification and validation of the MATTC criterion performed in this work did not demonstrate sufficient accuracy or reliability. As such, neither of these criteria are used for the purpose of transition prediction in the RANS flow solver.

3.5 Convergence to a Free-Transition Steady State

Automatic transition prediction in the RANS solver is achieved through an iterative process. Variations of this procedure have been developed and employed by Brodeur and van Dam [13], Krumbein [62], Cliquet *et al.* [23], and Mayda [78]. The present implementation borrows ideas from all of the above works, particularly that of Krumbein.

The incorporation of free transition into the implicit solver and convergence to steady-state is achieved in the following three phases:

1) Start-up from freestream conditions with fixed transition locations (AF solver).

2) Convergence with free transition point movement (AF solver).

3) Convergence with the final fixed transition locations (NK solver).

In the first phase an initial guess of the transition locations (top and bottom surfaces) is required and typically taken at 25% chord. Transition is then fixed in the Navier-Stokes solution to occur at the initial locations using the transition region model. When the L_2-norm of the flow residual has been reduced to 5×10^{-6}, the second phase begins. Note that the tight tolerance during the start-up phase was chosen to ensure sufficiently accurate boundary-layer properties for transition prediction, which, in turn, requires that any oscillations in the density, pressure and velocities have been sufficiently damped [64, 65].

When the second phase begins, the transition prediction module is invoked to determine the new transition locations. This step includes the determination of the boundary-layer edge, the calculation of the laminar boundary-layer properties, and the evaluation of a transition prediction criterion. The forced transition points are then moved upstream or downstream as required toward the newly predicted transition points. Note that in order to predict transition downstream of a forced transition location, the boundary-layer properties must be extrapolated (as discussed in Section 3.7.1). The transition point movement is under-relaxed, such that

$$X_{\mathrm{f}}^{\mathrm{new}} = \min\left(\Delta X_{\max}\,,\; X_{\mathrm{f}}^{\mathrm{old}} - \omega\left(X_{\mathrm{f}}^{\mathrm{old}} - X_{\mathrm{p}}\right)\right)\,, \tag{3.28}$$

where ω is the under-relaxation factor, X_{f} and X_{p} represent the normalized chord locations of the forced and predicted transition points, respectively, and the **min** function limits the

transition point movement at any given update. This limit on the transition point movement is selected as $\Delta X_{\max} = 0.2$ for $M_\infty < 0.65$ and $\Delta X_{\max} = 0.05$ for $M_\infty \geq 0.65$, the reasons for which are discussed in Section 3.7. When the L_2-norm of the flow residual returns to a magnitude of 5×10^{-6}, the predicted and forced transition points are again updated. The iterative transition prediction procedure is considered converged when the absolute value of the *transition residual*, given as $|\boldsymbol{\mathcal{R}}_{\text{tr}}| = |X_{\text{f}} - X_{\text{p}}|$, has converged to a tolerance of ϵ_{tp}. From numerical experimentation over a wide range of flight conditions, an under-relaxation factor of $\omega = 0.9$ has been selected as it demonstrates a good compromise between efficiency and robustness. When the transition prediction procedure has converged, the third and final phase begins.

In the third phase, with X_{f} equal to X_{p}, the flow solver switches from the AF algorithm to the NK algorithm. The solver then continues to converge until the magnitude of the flow residual has reduced to a tolerance of ϵ_{nk}. For the purposes of gradient-based aerodynamic shape optimization, ϵ_{tp} and ϵ_{nk} are set to 10^{-8} and 10^{-11}, respectively, ensuring a sufficiently smooth design space for optimization.

3.6 Flow Solver Efficiency

The efficiency of the free-transition RANS solver can be compared to a fully-turbulent flow solve on identical grids. In general, the iterative transition prediction procedure typically requires two to four times the computational cost of a fully-turbulent flow solve, with no significant addition to the memory requirements. Here we are comparing the cost to converge two different solution states; nonetheless, the fully-turbulent case provides a familiar basis for comparison. It is also important to note that the optimization cases employ a tight convergence tolerance for the transition points ($\epsilon_{\text{tp}} = 10^{-8}$). This is required to achieve a smooth design space for the purposes of gradient-based shape optimization (as demonstrated in Section 5.2). If one is only interested in performing flow analysis, then this tolerance may be relaxed. For example, a relaxed tolerance of $\epsilon_{\text{tp}} = 10^{-4}$, would require approximately 40% fewer AF iterations, in turn significantly reducing the cost of the flow solver while staying well within the accuracy tolerance of the given transition criterion. Finally, the computational cost of the flow solver, along with the various gradient evaluation techniques, will be quantified and compared in Section 5.6.1.

3.7 Flow Solver Robustness and Additional Considerations

Given that the end goal is to repetitively evaluate the aerodynamic characteristics of various airfoil geometries for the purpose of shape optimization, the robustness of the flow solver is of paramount importance. Indeed, the success of the optimization depends heavily on the flow solver's ability to converge the perturbed geometries generated during the optimization cycle (see Chapter 4 for further discussion). A robust logic has been determined through extensive numerical experimentation and code verification to handle the various outcomes of the transition prediction module. The methodology employed for handling such cases and improving robustness is discussed in this section.

When augmenting the flow solver with a transition prediction capability it becomes clear that RANS solutions under fixed and free-transition conditions are, in general, less robust than RANS solutions under fully-turbulent conditions. In this context, the term *robustness* refers to the ability of the flow solver to converge to steady-state over a wide range of flight conditions and geometries. The general decrease in robustness for flows with a region of laminar flow can be understood by considering the increased tendency of a laminar boundary layer to separate from the airfoil surface compared to a turbulent boundary layer. This separation may occur due to adverse pressure gradients and/or shock-wave interaction during transition point movement (or a poorly placed fixed transition point) particularly if the transition point is too far downstream. This flow separation often leads to unsteady flow phenomena that cannot be accurately captured by the steady RANS algorithm employed in this work. In some cases, if sharp changes in the flow arise, it can also lead to unphysical fluid properties such as negative density or pressure, at which point the flow solve is terminated.

For example, when moving to higher Reynolds and Mach numbers, it became increasingly evident that successful convergence of phases 1 and 2 (as described in Section 3.5) require a more conservative approach. This is the reason $\Delta X_{\max}$ is taken as 0.05 when $M_\infty \geq 0.65$. Regardless of the flight conditions, we take the quarter-chord position as the initial transition-point locations.

3.7.1 Extrapolation of Laminar Boundary-Layer Properties

The ability to predict transition downstream of the forced transition points is critical to the efficiency and robustness of the automatic transition prediction procedure. The transition criteria, however, do not make use of turbulent boundary-layer properties, as their role is to

assess the stability of *laminar* boundary layers. Thus, when our initial guess of the transition point is too far upstream, or if we encounter the situation in which the transition criterion fails to predict a transition point (upstream of the current transition location), the new transition point must be moved aft by some appropriate step length. To approximate the step length, we make use of a linear extrapolation of the laminar boundary-layer properties, thereby allowing us to predict a transition location that is located in the turbulent region of the boundary layer.

The point of linear extrapolation is computed based on the initial ramp up of the intermittency function, specifically the first point at which $\gamma \geq 1\times10^{-4}$. When using the SA trip terms, the point of extrapolation is based on the point at which the shape factor takes its maximum value (in the upstream vicinity of the transition point). The slope of the linear extrapolation is then computed using the neighbouring properties immediately upstream of the extrapolation point. As an example, Figure 3.4 demonstrates the linear extrapolation of the shape factor and the natural extrapolation of the N-factor envelope that results. Due to the potential for small oscillations in the boundary-layer properties near the transition point (if, for example, the transition region is not sufficiently smooth), an additional check is performed to ensure that the extrapolation is taken from a monotonic region. In the case that it is not, we march in the upstream direction and extrapolate from the nearest monotonic region, where the extent of the region must be over a minimum specified distance (taken as 3% chord). This ensures a suitable extrapolation and improves the robustness of the transition point movement and convergence.

3.7.2 Flow Separation

If laminar flow separation is detected and a transitional separation bubble forms, the e^N envelope method is able to predict the transition location in the separation bubble [28]. Note that when using the AHD criterion, we select the laminar separation point as the transition point, which follows the work of Krumbein [64]. For the higher Reynolds numbers considered in this work, the separation bubbles that occur are typically very small in size (less than 5% chord), and thus the potential error incurred in choosing the separation point as the transition point is minimal and within the accuracy of a correlation-based transition criterion.

When allowing for free transition in transonic flight conditions, it is possible that a shock wave occurs upstream of the transition point, resulting in shock-induced laminar flow separation. This often leads to unsteadiness which prevents flow solver convergence. In an attempt to mitigate this issue, the limit in (3.28) applied to the maximum transition point movement at any given update is reduced to $\Delta X_{\max} = 0.05$. The more gradual

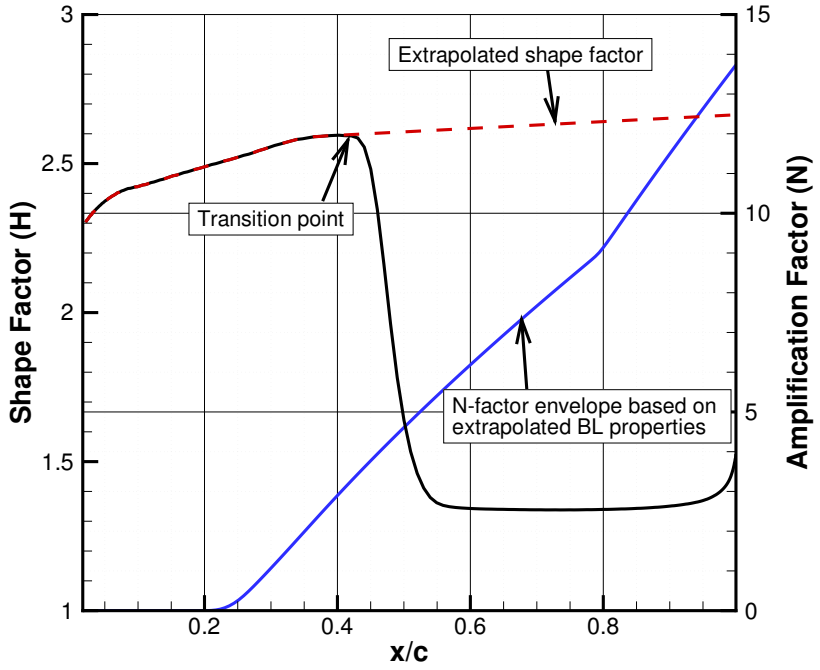

Figure 3.4: Shape factor and amplification factor extrapolation

movement of transition points is applied in general when the freestream Mach number is greater than or equal to 0.65. The gains in numerical robustness and the resulting consistency in the convergence behaviour are well worth the sacrifice in computational efficiency. When comparing cases that can converge with or without the modification, the increase in computational cost is typically less than 30%.

3.7.3 Fully-Laminar Flow

The last special case we shall consider is that of *fully-laminar* flow on a given surface. If this situation occurs, it is most likely to occur on the lower surface of an NLF airfoil, particularly at lower Reynolds numbers and with a positive angle of incidence (which has the effect of moving the lower surface transition location aft). In this work, we enforce an aft limit on the transition points of 90% chord, in order to allow sufficient chord length for the smooth ramp-up of the eddy viscosity. While this may not be a physical requirement of flow, it is imposed as an added measure of robustness. It should be noted that the transition-point location in this work defines the beginning of the ramp-up region; thus we must add the transition length of the transition region to this aft limit. The transition length is typically

specified or calculated at less than 10% chord. Thus, the imposed aft-limit is designed to ensure that the flow leaves the trailing edge of the airfoil as fully turbulent.

As a final note, the verification and validation of the above framework for automatic transition prediction in a RANS solver will be presented in Chapter 5. In the next chapter, we will discuss the optimization techniques employed to design for NLF.

Chapter 4

Optimization Techniques

Optimization is the maximization or minimization of a function subject to various constraints on its variables [88]. Aerodynamic Shape Optimization (ASO) is concerned with the development of numerical optimization techniques that enable a level of automation in the aerodynamic design process. Recall that we reviewed the various advantages of employing ASO in the design process in Section 1.3. Within the field of numerical optimization there is an ever-growing toolbox of optimization strategies that have become increasingly complex in order to match the increasing complexity and wide array of applications.

In this work we consider only continuous optimization, that is, we do not make use of discrete design variables, nor do we consider stochastic optimization strategies. Instead, we make exclusive use of gradient-based optimization, in which we must compute the sensitivity of our design objective to the changes in our design variables. There exist several gradient-free approaches that have seen wide-spread use, including genetic algorithms [88, 97], the Nelder-Mead simplex method [84], and the simulated annealing method [58]. However, gradient-based methods can be significantly more efficient than those that are gradient-free. As an example, Zingg *et al.* [138] compared genetic to gradient-based optimization algorithms and quantified the difference in computational expense. They found that the gradient-free method required significantly more calls to the flow solver than a gradient-based algorithm.

A drawback of gradient-based methods is found in the optimality conditions; when satisfied they can only guarantee that a local optimum has been found. When dealing with multi-modal design spaces, Chernukhin and Zingg [21] have shown that an effective strategy to more thoroughly explore the design space is to use a multi-start gradient-based approach, in which multiple starting geometries are each optimized independently. The approach was shown to be more computationally efficient than a gradient-free genetic algorithm [21]. Furthermore, the studies performed by Chernukhin and Zingg [21] suggest that the design space is unimodal in two-dimensions. Indeed, this work makes use of a single starting geometry for all lift-constrained drag minimizations; it has been verified that two very different starting geometries arrive at two very similar optimal designs.

In this chapter we begin by providing an overview of the optimization framework and

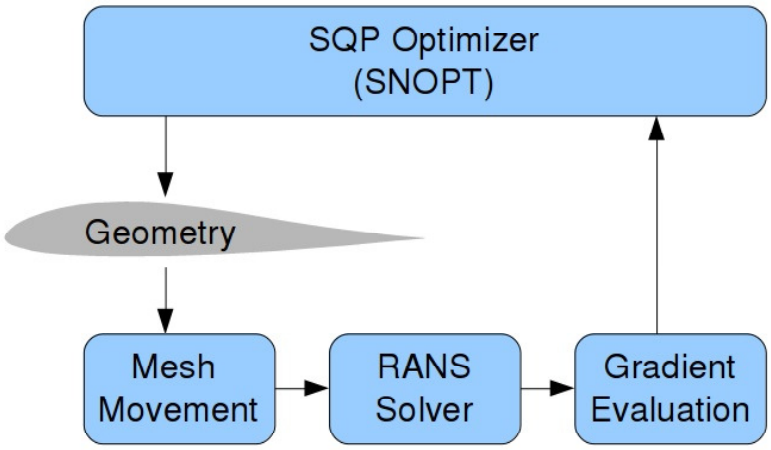

Figure 4.1: Schematic of the Optimization Framework

formally presenting the optimization problem at hand. We then discuss the sequential quadratic programming algorithm and describe the various gradient evaluation techniques investigated. Next, we propose the implementation and application of a general and flexible framework for incorporating the sensitivities of the non-local transition prediction criteria into the gradients of the objective and constraints. And finally, we discuss the multipoint optimization strategy employed for the robust design of NLF airfoils.

4.1 Optimization Framework

The present optimization framework consists of the following five constituents: (i) a geometry parametrization, (ii) a mesh movement algorithm, (iii) a two-dimensional RANS flow solver with transition prediction capabilities, (iv) various gradient evaluation techniques, and (v) a Sequential Quadratic Programming (SQP) algorithm. Figure 4.1 depicts the order in which an iteration of the optimization proceeds through each of the framework's constituents. At each iteration, the design variables are updated and a new airfoil geometry is defined. The mesh is then moved to fit the new geometry, the RANS flow solver is called to provide the objective and constraints, and the gradients of the objective and constraints (with respect to the design variables) are evaluated. Next, the objective, constraints and their gradients are passed back to the SQP optimizer to compute the next design update. The process continues to drive the geometry toward a locally optimum solution that satisfies the first-order optimality conditions (Section 4.4.1). With the exception of the RANS solver discussed in the previous two chapters, each of the constituents listed above will be discussed further in this chapter. Special attention is given to the discrete-adjoint gradient evaluation as it requires significant modifications for the incorporation of free transition.

4.2 Problem Definition

The aerodynamic shape optimization problem is defined as follows:

min	$\mathcal{J}\left(\mathbf{x}, \mathbf{Q}(\mathbf{x})\right)$	
w.r.t.	$\mathbf{x} \in \mathbb{R}^{N_\mathrm{D}}$	
s.t.	$\widetilde{c_j}\left(\mathbf{x}, \mathbf{Q}(\mathbf{x})\right) = 0$	for $j = 1 \ldots N_\mathrm{e}$
	$\widehat{c_j}\left(\mathbf{x}, \mathbf{Q}(\mathbf{x})\right) \leq 0$	for $j = 1 \ldots N_\mathrm{i}$
	$A_\mathrm{L}\mathbf{x} - \mathbf{b} = 0$	
	$\boldsymbol{\ell} \leq \mathbf{x} \leq \boldsymbol{u}$	

Table 4.1: Formulation of the optimization problem

The goal of our optimizations is thus to minimize an aerodynamic design objective, $\mathcal{J}$, with respect to the design variable vector, $\mathbf{x} \in \mathbb{R}^{N_\mathrm{D}}$, subject to linear and nonlinear constraints, where N_D is the number of design variables defined in the next section. The nonlinear constraints are divided into the equality type, $\widetilde{c_j}$, and inequality type, $\widehat{c_j}$, which are numbered up to N_e and N_i, respectively. The linear constraints may be represented by a matrix, $A_\mathrm{L} \in \mathbb{R}^{N_\mathrm{D} \times N_\mathrm{D}}$, and a bounding vector, $\mathbf{b} \in \mathbb{R}^{N_\mathrm{D}}$, which provide a means of enforcing any linear constraint relationships between the design variables. The design variables are also subject to lower and upper bounds, $\boldsymbol{\ell} \in \mathbb{R}^{N_\mathrm{D}}$ and $\boldsymbol{u} \in \mathbb{R}^{N_\mathrm{D}}$, respectively.

A few notes regarding the above formulation. We have chosen to express the objective and constraints as functions of both the design variables, $\mathbf{x}$, and the intermediate flow variables, $\mathbf{Q}(\mathbf{x})$. Other formulations are possible, such as reverse algorithmic differentiation, and the all-at-once formulation, as reviewed in Hicken's PhD thesis [46]. In Table 4.1, we have implicitly imposed that the discrete RANS flow equations are satisfied for any given $\mathbf{x}$, that is, the problem is *strictly* PDE-constrained. This is implicitly enforced through the dependence of $\mathbf{Q}$ on $\mathbf{x}$, in that $\mathbf{Q}(\mathbf{x})$ *must satisfy* $\boldsymbol{\mathcal{R}}\left(\mathbf{Q}(\mathbf{x})\right) = 0$ as per (2.26).

Unlike all other nonlinear constraints, infeasibility with respect to the PDE constraint is not permitted at any point during the iterative optimization cycle. Consequently, if the RANS solver fails to converge the flow residual for any given $\mathbf{x}$, then the optimizer receives a *failure status* (without cause or reason) and must backtrack $\mathbf{x}$ as required. This can stall the optimization process and underlies the importance of developing a robust free-transitional flow solver (Section 3.7).

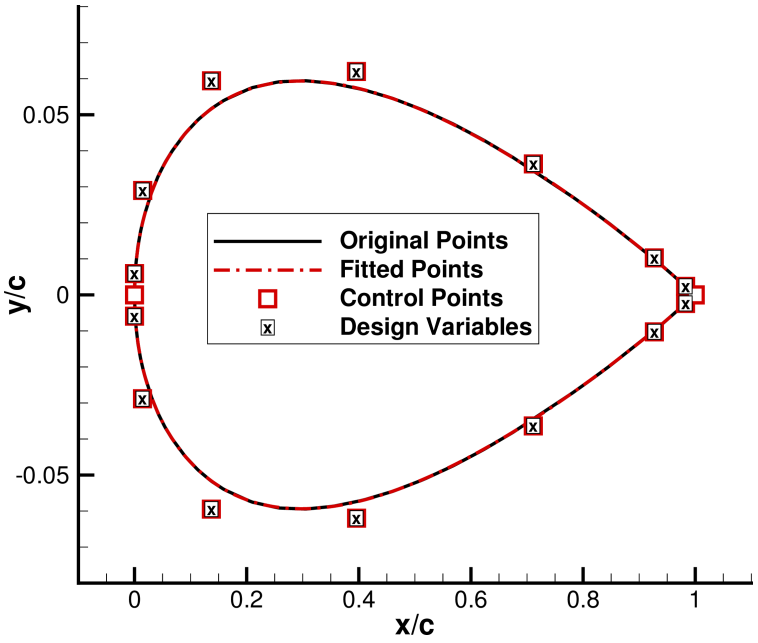

Figure 4.2: B-spline parameterization of NACA 0012 airfoil

4.2.1 Design Variables and Geometry Parameterization

The airfoil geometry is parametrized using cubic B-splines [85]. At the onset of the optimization cycle, the initial locations of the B-spline control points are determined by a least-squares fitting procedure that minimizes the discrepancy between the B-spline curve and the initial airfoil geometry. The details of the parametric representation of the airfoil using B-splines are given in Appendix B. As a typical example used in this work, Figure 4.2 shows the cubic B-spline representation of the NACA 0012 airfoil using a total of seventeen control points (two of which are coincident at the trailing edge). Note that the control points are clustered near the leading and trailing edges allowing for more geometric flexibility in those regions. The B-spline curve is also forced to pass through the leading and trailing edges in order to maintain the orientation of the chord [85].

The control points are free to move in the vertical direction to facilitate shape changes during the optimization cycle. As such, the design variables, $\mathbf{x}$, are defined as the y/c coordinates of the B-spline control points. The design variables are marked with an "x" in Figure 4.2. In order to keep a constant chord, the two coincident control points located at the sharp trailing edge are kept fixed, as is the single control point located at the leading edge.

The angle of attack of the airfoil (in degrees) is also included as a design variable (the last component of $\mathbf{x}$). Thus, for an airfoil parameterized by seventeen B-spline control points,

the design variables would include fourteen control points plus the angle of attack, giving $N_{\mathrm{D}} = 15$.

Finally, design variable scaling has been shown to have a potentially significant impact on the success and efficiency of the optimization cycle. Zingg *et al.* [142] have reviewed several different scaling options and quantified their effect on efficiency. In this work, the angle of attack (in degrees) is left unscaled, while the geometric design variables are increased in magnitude by a factor of one hundred. This scaling improves the efficiency and robustness of the optimizer by equalizing the order of magnitude of the corresponding gradient components, in turn, preventing one type of design variable from overshadowing the other.

4.2.2 Design Objectives

The optimization framework is capable of handling several design objectives, including, but not limited to: minimizing the drag coefficient, maximizing lift-to-drag ratio, and maximizing the endurance factor. See, for example, the various works of Zingg *et al.* [15, 30, 86, 87]. Recall from the review in Section 1.3 that several researchers using ASO for NLF design [5, 44, 54, 94] have chosen design objectives such as minimizing the kinetic energy over a range of disturbance frequencies, minimizing the total turbulent kinetic energy, or matching a target pressure distribution or target N-factor envelope.

This research has made a concentrated effort towards considering a design objective that is more in line with the practical aerodynamic objectives of airframers and manufacturers, and, perhaps more importantly, one that accounts for the possible drag trade-offs that might not be captured by an objective that specifically aims to delay transition. We focus on lift-constrained drag minimization in which the design objective is the total drag on the airfoil and a target lift coefficient must be achieved. Using the near-field formulation the total drag is computed by numerically integrating the component contributions from both skin friction (shearing) forces and pressure (normal) forces in the freestream direction. As a final note, we will later present design objectives that are modified for the purposes of weighted multipoint optimization and Pareto-front design studies.

4.2.3 Constraints

In this section we discuss some physical constraints imposed that are relevant to the applications at hand. The SQP algorithm works best when the problem is well constrained and appropriately bounded, as this tends to improve convergence and robustness. For example, it is good practice to impose constraints which avoid impractical or infeasible geometries [37].

The following list summarizes the linear and nonlinear constraints applied to the airfoils in this work:

1. A **nonlinear inequality area constraint** is imposed that ensures that the final area of the airfoil is greater than or equal to the initial area: $A_{\text{final}} \geq A_{\text{initial}}$.

2. A set of **nonlinear inequality thickness constraints** are enforced. In this work we typically place one thickness constraint near the leading edge and one near the trailing edge to ensure sufficient local thickness of the airfoil. Specific details will be provided in Chapter 6.

3. A **nonlinear equality lift constraint** is enforced when performing lift-constrained drag minimizations. The lift coefficient of the airfoil must match the user-specified target lift coefficient, C_l^*, such that $C_l = C_l^*$.

4. A set of **linear equality constraints** are imposed to improve the robustness of the optimizer and avoid non-physical airfoil shapes during the optimization cycle. The geometric design variables that are vertically aligned with one another (upper and lower surfaces) are paired up, and a straightforward linear constraint is imposed between them such that $x_{\text{up}} > x_{\text{low}}$. Note that the thickness and area constraints are nonlinear and, as such, may be violated during the optimization cycle [37]. Since the linear constraints are strictly enforced at each iteration [37], we can guarantee that no cross-over of the design variables will occur.

4.3 Mesh Movement

As the vertical positions of the B-spline control points are updated, the nodal locations of the grid must also be updated to fit the mesh to the new geometry. An algebraic grid-perturbation strategy [85, 126] is used to ensure that the computational grid is smoothly adjusted to conform to the changing geometric configurations. The algorithm preserves the farfield boundary and makes changes to the grid points based on their normalized distance along the off-wall (η-direction) grid lines. It has been verified that the algorithm is capable of maintaining a smooth and near-orthogonal grid even when dealing with large shape changes near the leading edge, which, as we shall see, are often encountered when designing for NLF. Specific details of the mesh movement algorithm may be found in Nemec's PhD thesis [85].

4.4 Sequential Quadratic Programming

In this section we discuss the responsibilities bestowed upon the optimizer, and the techniques it uses to drive the objective toward an optimal point. The gradient-based optimizer has two primary responsibilities. The first is to compute a *search direction* that dictates the direction in which the design variables should move, and the second is to determine an appropriate *step-length* that dictates how far to move them. For this purpose we seek to use the Sequential Quadratic Programming (SQP) technique, which is considered one of the most effective gradient-based methods for treating smooth, nonlinearly constrained, large-scale optimization problems [88]. It is well-known that implementing a robust SQP algorithm that efficiently performs the aforementioned tasks can be extremely challenging [88]. For example, the algorithm must circumvent such issues as failed function evaluations, highly infeasible design points, and poor scaling, all in a robust and efficient manner. As such, this work relies on a well-established general-purpose Sequential Nonlinear OPTimizer (SNOPT) developed by Gill, Murray, and Saunders [37].

4.4.1 SNOPT solution procedure

In this section we briefly review the SNOPT solution procedure in order to highlight its key components and the relevant measures of feasibility and optimality used to assess convergence. As input, SNOPT requires the initial design variable state, $\mathbf{x}_{k=1}$, as well as the objective, constraints, and their gradients. The algorithm itself involves *major* and *minor* iterations. At each major iteration the design variables are updated based on a computed search direction and step-length. As a pre-processing step, SNOPT reformulates the optimization problem given in Table 4.1 by performing the following modifications:

Pre-Processing Steps:

1. Convert all *inequality* constraints into *equality* constraints. This is done by introducing slack variables, s_j, to which bounds are applied, that is, $\widehat{c}_j \leq 0$ becomes $\widehat{c}_j - s_j = 0$ for $j \in [1, N_\mathrm{i}]$, with upper and lower bounds defined as $\ell_{s_j} \leq s_j \leq u_{s_j}$. The new equality constraints can now be added to the existing equality constraints to form a single constraint vector, $\mathbf{c} \in \mathbb{R}^{N_\mathrm{c}}$, where $N_\mathrm{c} = N_\mathrm{i} + N_\mathrm{e}$.

2. Convert the *constrained* objective into an *unconstrained* objective. This is done by multiplying each constraint by a positive scalar quantity. These scalars are known as the Lagrange multipliers, $\boldsymbol{\lambda} \in \mathbb{R}^{N_\mathrm{c}}$. This defines a new *Lagrangian* objective function

of the form:

$$\mathcal{L} = \mathcal{J} - \sum_{j=1}^{N_c} \lambda_j c_j \, . \tag{4.1}$$

Note that the gradient of the Lagrangian, $\nabla\mathcal{L}$, includes the gradients of the objective and constraints with respect to the design variables:

$$\nabla\mathcal{L} = \nabla\mathcal{J} - \sum_{j=1}^{N_c} \lambda_j \nabla c_j \, . \tag{4.2}$$

SNOPT then takes the following steps at each major iteration (indexed by k):

1. Strictly satisfy all linear constraints by minimizing the L_1-norm of the linear constraint violations. If the linear constraints remain infeasible, SNOPT terminates.

2. Approximate (or update) the Hessian matrix, $H_k = \nabla^2\mathcal{L}$, which contains the second derivatives of the Lagrangian, using the quasi-Newton method of Broyden, Fletcher, Goldfarb, and Shanno (BFGS) [11].

3. Form a local quadratic model of the Lagrangian, $\mathcal{L}_q$, using the gradient and Hessian:

$$\mathcal{L}_q(\mathbf{x}) = \mathcal{L}_k + \nabla\mathcal{L}^T(\mathbf{x} - \mathbf{x}_k) + \frac{1}{2}(\mathbf{x} - \mathbf{x}_k)^T H_k(\mathbf{x} - \mathbf{x}_k) \, . \tag{4.3}$$

4. Linearize the constraint functions, such that $\mathbf{c}(\mathbf{x}) = \mathbf{c}(\mathbf{x}_k) + \nabla\mathbf{c}^T(\mathbf{x} - \mathbf{x}_k)$.

5. Minimize the quadratic model (4.3) subject to the linearized constraints to find its optimal point, denoted $\mathbf{x}^*_{\mathrm{QP}}$. The optimization is carried out using an active-set method [37] which is itself an iterative process; the iterations of the active-set method are the *minor* iterations of the SQP algorithm.

6. Define the search-direction vector, $\mathbf{p}_k$, as pointing toward the optimal solution of the quadratic model, such that $\mathbf{p}_k = \mathbf{x}^*_{\mathrm{QP}} - \mathbf{x}_k$

7. Determine an appropriate step-length, $\alpha_k \in [0, 1]$, by performing a line-search. The step-length must ensure sufficient decrease in the *augmented Lagrangian merit function*, which is defined as

$$\mathcal{M}(\mathbf{x}) = \mathcal{J}(\mathbf{x}) - \boldsymbol{\lambda}^T\mathbf{c}(\mathbf{x}) + \frac{1}{2}\sum_{j=1}^{N_c} \rho_j \left(c_j(\mathbf{x})\right)^2 \, , \tag{4.4}$$

where $\boldsymbol{\rho} \in \mathbb{R}^{N_c}$ is a vector of penalty parameters that are adjusted to vary the weight of the constraints relative to the objective. If the merit function increases, then the line-search is backtracked toward $\mathbf{x}_k$ as necessary. In essence, the merit function is a relative measure between iterates that combines the objective function value with measures of constraint violations [88]. When all constraints are satisfied (to machine accuracy), it can be seen in (4.4) that the merit function equals the objective function (to machine accuracy).

8. Update the design variables as follows: $\mathbf{x}_{k+1} = \mathbf{x}_k + \alpha_k \mathbf{p}_k$.

The process is repeated until the Karush-Kuhn-Tucker (KKT) conditions are satisfied. The KKT conditions provide the first-order necessary conditions for a locally optimal point as follows:

$$\frac{\partial \mathcal{L}}{\partial x_i} = \nabla \mathcal{L}_i = \frac{\partial \mathcal{J}}{\partial x_i} - \sum_{j=1}^{N_c} \lambda_j \frac{\partial c_j}{\partial x_i} = 0 \qquad \text{for} \quad i = 1 \ldots N_D \tag{4.5}$$

$$\frac{\partial \mathcal{L}}{\partial \lambda_j} = c_j = 0 \qquad \text{for} \quad j = 1 \ldots N_c \tag{4.6}$$

A formal proof of the KKT conditions may be found in Nocedal and Wright [88]. SNOPT computes two metrics to assess the KKT conditions and to provide the convergence criteria for the optimization cycle. The first metric is termed *feasibility*, the second *optimality*. Feasibility ($\mathcal{F}$) in SNOPT is defined (based on the second KKT condition) as the maximum nonlinear constraint violation normalized by the magnitude of the design variable vector:

$$\mathcal{F} = \max_j \frac{|c_j(\mathbf{x})|}{||\mathbf{x}||} \,. \tag{4.7}$$

Optimality ($\mathcal{P}$) in SNOPT is defined based on the first KKT condition such that

$$\mathcal{P} = \max_i \frac{(\bar{x} \nabla \mathcal{L}_i)}{||\boldsymbol{\lambda}||} \,, \tag{4.8}$$

where

$$\bar{x} = \begin{cases} \min(x_i - l_i, 1) & \text{if } \nabla \mathcal{L}_i \geq 0; \\ \min(u_i - x_i, 1) & \text{if } \nabla \mathcal{L}_i < 0. \end{cases}$$

The optimization is considered converged ($\mathbf{x}_k = \mathbf{x}^*$) when $\mathcal{F}_k \leq \epsilon_f$ and $\mathcal{P}_k \leq \epsilon_p$; in this work we take the convergence tolerances to be $\epsilon_f = 1 \times 10^{-5}$ and $\epsilon_p = 1 \times 10^{-6}$. However, in practice, if convergence is not achieved we observe the merit function to see if it has flat-lined and

often use this as an indicator that the optimizer is not able to further improve the geometry, or at least not to a degree that is worthy of the extra computational expense. In that case, we check that optimality has been reduced by approximately one order of magnitude relative to the initial iterate, which in some cases may already be a near-optimal design.

Finally, we reiterate that SNOPT requires the gradients of the objective function and constraints at each major iteration. Ensuring sufficiently accurate gradients is of paramount importance to the success of the SQP algorithm. Three methods for computing accurate gradients are introduced in the next section. The extensions required to incorporate the sensitivities of the non-local transition criteria are developed in Section 4.6.

4.5 Gradient Evaluation Techniques

The first method implemented to evaluate the gradient of the objective and constraints is a parallelized finite-difference approximation, the second is a flow sensitivity (or direct) method, and the third is a discrete-adjoint method. Note that the end-goal of this research is to make exclusive use of an adjoint gradient; however, many lessons were learned and the accuracy assessment of the adjoint gradient was made possible, by the implementation of the first two methods. A brief description of each approach is presented in the following sub-sections.

4.5.1 Finite-Difference Gradient Evaluation

The gradients of the objective and constraint functions with respect to the design variables – the B-spline control points and angle of attack – may be obtained through a finite-difference approximation. In doing so, we can approximate the gradients without having to differentiate the solver. Meanwhile the resulting gradients implicitly contain the sensitivities of the given transition criterion that defines the transition locations, in turn allowing the optimizer to exploit that information.

The finite-difference approximation makes small, finite perturbations in the design variables and evaluates the sensitivity of the objective and constraint functions to those changes [88]. The components of the gradient vector, G_n, are evaluated using a second-order centered-difference approximation as follows:

$$G_i = \frac{d\mathcal{J}}{dx_i} = \frac{J\left(\mathbf{x} + h\boldsymbol{e}_i, Q\left(\mathbf{x} + h\boldsymbol{e}_i\right)\right) - J\left(\mathbf{x} - h\boldsymbol{e}_i, Q\left(\mathbf{x} - h\boldsymbol{e}_i\right)\right)}{2h} \qquad \text{for} \quad i = 1 \ldots N_{\mathrm{D}} \quad (4.9)$$

where $\boldsymbol{e}_i$ is the i^{th} unit vector, and h is a positive scalar defined as $h = \max(\epsilon_{\mathrm{fd}} \cdot |x_i|, 10^{-8})$.

The step-size, h, is computed by scaling the magnitude of the given design variable, $|x_i|$, by the constant ϵ_{fd}. A value for ϵ_{fd} and hence the step-size magnitude must be carefully selected to avoid subtractive cancellation issues associated with a finite-difference approximation; a value of $\epsilon_{\text{fd}} = 10^{-6}$ was selected based on step-size studies.

The centered-difference approximation requires two flow solves to evaluate each component of the gradient, and as such, it does not scale well with the number of design variables. In this work, the gradient evaluation has been parallelized such that multiple processors compute the necessary function evaluations – on the entire set of perturbed design variables – in parallel. The gradient evaluation is also carried out in parallel with the evaluation of the objective and constraints on the baseline geometry of the given design iteration. For example, a case with fifteen design variables would employ thirty-one processors in the parallel computation of the objective and gradients. Thus, one can obtain all of the necessary information in approximately the same turn-around time as a single serial flow solve. Although this approach is feasible in two dimensions for single-point optimization, the required computational resources become prohibitive in three dimensional and/or multipoint optimizations.

4.5.2 Flow Sensitivity Gradient Evaluation

The flow sensitivity gradient evaluation differentiates the objective with respect to each design variable in a direct fashion as follows:

$$G_i = \frac{d\mathcal{J}}{dx_i} = \frac{\partial \mathcal{J}}{\partial x_i} + \frac{\partial \mathcal{J}}{\partial \mathbf{Q}} \frac{\partial \mathbf{Q}}{\partial x_i} \qquad \text{for} \quad i = 1 \ldots N_{\text{D}} \,. \tag{4.10}$$

The partial derivative of the objective with respect to each of the design variables, $\frac{\partial \mathcal{J}}{\partial x_i}$, translates into the sensitivity of the objective to the changes in the grid or, more specifically, to the changes in the grid metrics that occur when the i^{th} B-spline control point is perturbed. This is computed by a centered-difference approximation. Since the flow variables at each node, $\mathbf{Q}$, remain unchanged, the procedure does not require any additional flow solutions and is computationally inexpensive.

The convention in (4.10) is such that $\frac{\partial \mathcal{J}}{\partial \mathbf{Q}}$ is a row vector $\in \mathbb{R}^{1 \times N_{\text{Q}}}$, and $\frac{\partial \mathbf{Q}}{\partial x_i}$ is a column vector $\in \mathbb{R}^{N_{\text{Q}} \times 1}$, making their product a scalar quantity. The partial derivative of the objective with respect to the flow variables, $\frac{\partial \mathcal{J}}{\partial \mathbf{Q}}$, is obtained by a straightforward analytical differentiation, which is carried out by hand.

The main challenge and expense associated with the flow sensitivity approach is in the computation of $\frac{\partial \mathbf{Q}}{\partial x_i}$, representing the sensitivity of the flow variables to each design variable.

The discretized flow residual of the RANS equations can be written as a function of the design variables and the flow variables as $\boldsymbol{\mathcal{R}}(\mathbf{x}, \mathbf{Q}(\mathbf{x}))$. Since the residual is satisfied for any given $\mathbf{x}$, we have

$$\frac{d\boldsymbol{\mathcal{R}}}{dx_i} = 0 = \frac{\partial \boldsymbol{\mathcal{R}}}{\partial x_i} + \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}} \frac{\partial \mathbf{Q}}{\partial x_i} \qquad \text{for} \quad i = 1 \ldots N_{\mathrm{D}} \,. \tag{4.11}$$

Thus, for each design variable we can compute $\frac{\partial \mathbf{Q}}{\partial x_i}$ by solving a large, sparse system of linear equations given by

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}} \frac{\partial \mathbf{Q}}{\partial x_i} = -\frac{\partial \boldsymbol{\mathcal{R}}}{\partial x_i} \qquad \text{for} \quad i = 1 \ldots N_{\mathrm{D}} \,. \tag{4.12}$$

The partial derivative of the flow residual with respect to the design variables (keeping $\mathbf{Q}$ constant), $\frac{\partial \boldsymbol{\mathcal{R}}}{\partial x_i}$, is computed using a centered-difference approximation requiring two residual calculations for each design variable, which is computationally inexpensive.

The system of equations given by (4.12) is solved using preconditioned GMRES as described in Section 2.2 for the Newton-Krylov flow solution. The two main differences here are that we must converge the system to a much tighter tolerance (taken as 1×10^{-8}) and that the $\mathcal{A}_2$-Jacobian (as previously described in Section 2.2) is used for $\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}$. The system is preconditioned using the $\mathcal{A}_1$-Jacobian. While no new flow solves are required to obtain $\frac{\partial \mathbf{Q}}{\partial x_i}$, we must solve (4.12) for each design variable independently, in turn accruing significant computational expense. Similar to the finite-difference gradient presented in the previous section, the flow sensitivity method does not scale well with the number of design variables. To address the scaling issue, the adjoint method, described in the next section, has emerged as a popular alternative.

4.5.3 Discrete-Adjoint Gradient Evaluation

The principal advantage of the adjoint method is that its cost does not scale with the number of design variables, but rather with the number of objectives and nonlinear constraints (not including geometry-based nonlinear constraints). Hence, the objective function gradient evaluation only requires one flow solve and one adjoint solve. For lift-constrained drag minimization, an additional adjoint solve is required for the gradient of the lift constraint.

In the discrete-adjoint approach, the gradient is evaluated using the following expression [86]:

$$G_i = \frac{d\mathcal{J}}{dx_i} = \frac{\partial \mathcal{J}}{\partial x_i} - \boldsymbol{\psi}^T \frac{\partial \boldsymbol{\mathcal{R}}}{\partial x_i} \qquad \text{for} \quad i = 1 \ldots N_{\mathrm{D}} \,, \tag{4.13}$$

where $\boldsymbol{\psi} \in \mathbb{R}^{N_{\mathrm{Q}} \times 1}$ is the *adjoint vector*. A detailed derivation of the adjoint formulation for fully-turbulent flow may be found in Nemec and Zingg [86, 87]. The adjoint vector is

obtained by solving a large, sparse system of linear equations given by

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi} = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^T , \tag{4.14}$$

The various sensitivities in (4.13) and (4.14) are computed as described in the previous section. The adjoint equations (4.14) are solved using preconditioned GMRES in the same manner as the flow sensitivity equations (4.12). Indeed, the adjoint and flow sensitivity approaches have several features in common. Their commonality can be leveraged when verifying their implementations, particularly when major modifications are made, such as those required to incorporate non-local transition prediction criteria.

4.6 Augmented Adjoint Formulation for Non-Local Transition Prediction

4.6.1 The Augmented Adjoint System

The e^N and AHD transition criteria (described in Sections 3.4.1 and 3.4.2, respectively) are non-local in their formulation in that they use flow variables in their evaluation belonging to grid points which are not physically near the transition point. As such, special consideration must be taken when evaluating and deriving an adjoint formulation capable of incorporating their sensitivities. Whenever we populate terms in the $\mathcal{A}_2$-Jacobian, $\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}$, the local formulation of the RANS transport equations produce a sparsity pattern that is amenable to iterative solvers, such as preconditioned GMRES, and may also be exploited by re-ordering techniques [36, 85] to further improve efficiency. Gatsis provides examples of such sparsity patterns in his PhD thesis [36] and details the available preconditioning techniques, including the ILU right-preconditioning used herein. Maintaining a Jacobian matrix with a desirable sparsity pattern is thus critical to the efficiency, memory, and indeed the numerical feasibility of solving the associated system of equations (4.14). However, any discretized quantity that is computed using highly off-diagonal (non-local) flow information, such as the predicted transition points, threatens the sparsity pattern of the matrix containing its differentiation with respect to the flow variables.

The proposed approach is to append a new adjoint vector, $\boldsymbol{\psi}_{\text{tr}}$, to the original adjoint vector, such that $\overline{\boldsymbol{\psi}} \Rightarrow [\boldsymbol{\psi} \; ; \; \boldsymbol{\psi}_{\text{tr}}]$. Henceforth, the overbar shall be used to indicate an augmented vector. The length of $\boldsymbol{\psi}_{\text{tr}}$ corresponds to the number of transition points, N_{tr}, which is equal to two for a single-element airfoil. The length of $\overline{\boldsymbol{\psi}}$ is then equal to $N_{\text{Q}}+N_{\text{tr}}$.

For 3D wing configurations, the transition lines may be defined by a spanwise distribution of transition points [63]; the total number of spanwise transition points on all surfaces gives N_{tr}.

To compute the new adjoint variables, we specify a corresponding number of new residual equations, such that $\overline{\boldsymbol{\mathcal{R}}} \Rightarrow [\boldsymbol{\mathcal{R}}\,;\,\boldsymbol{\mathcal{R}}_{\text{tr}}]$. The new transition residual equations represent the distance between the forced and predicted transition locations, $\boldsymbol{\mathcal{R}}_{\text{tr}} = \mathbf{X}_{\text{f}} - \mathbf{X}_{\text{p}}$, as described in Section 3.5. The transition residual vector is satisfied ($\boldsymbol{\mathcal{R}}_{\text{tr}} = 0$) when the forced transition points are in locations consistent with the given transition criterion ($\mathbf{X}_{\text{f}} = \mathbf{X}_{\text{p}}$).

In addition, the vector of conserved flow variables must be augmented to include the forced transition locations, such that $\overline{\mathbf{Q}} \Rightarrow [\mathbf{Q}\,;\,\mathbf{X}_{\text{f}}]$. The entire adjoint vector, $\overline{\boldsymbol{\psi}}$, is computed by solving the augmented linear system of equations given by

$$\frac{\partial \overline{\boldsymbol{\mathcal{R}}}}{\partial \overline{\mathbf{Q}}}^{T} \overline{\boldsymbol{\psi}} = \frac{\partial \mathcal{J}}{\partial \overline{\mathbf{Q}}}^{T} \longmapsto \left[\begin{array}{c:c} \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}} & \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_{\text{f}}} \\ \hdashline \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}} & \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{X}_{\text{f}}} \end{array}\right]^{T} \left[\begin{array}{c} \boldsymbol{\psi} \\ \hdashline \boldsymbol{\psi}_{\text{tr}} \end{array}\right] = \left[\begin{array}{c} \frac{\partial \mathcal{J}}{\partial \mathbf{Q}} \\ \hdashline \frac{\partial \mathcal{J}}{\partial \mathbf{X}_{\text{f}}} \end{array}\right]^{T} . \tag{4.15}$$

Two solution strategies have been developed to solve the new augmented system, as presented in Section 4.7. Upon solving for $\overline{\psi}$, the gradient vector may be computed by the augmented version of (4.13) given by

$$G_i = \frac{d\mathcal{J}}{dx_i} = \frac{\partial \mathcal{J}}{\partial x_i} - \overline{\boldsymbol{\psi}}^{T} \frac{\partial \overline{\boldsymbol{\mathcal{R}}}}{\partial x_i} \qquad \text{for} \quad i = 1 \ldots N_{\text{D}} \,, \tag{4.16}$$

where the last N_{tr} entries of $\frac{\partial \overline{\boldsymbol{\mathcal{R}}}}{\partial x_i}$ are given by the differentiation of the transition residual with respect to the design variables, that is, $\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial x_i}$. Keeping $\overline{\mathbf{Q}}$ constant, this differentiation is cheaply and accurately computed by a centered-difference approximation.

In the remaining subsections, prior to presenting the solution strategies, we discuss the details of how the augmented system is formed, term by term, followed by the advantages and challenges of the proposed approach.

4.6.2 Forming the Augmented Adjoint System

The various terms in the new augmented adjoint system are defined and computed as follows:

$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}$: The original flow Jacobian contains the sensitivities of the flow residuals to the flow variables and is analytically derived making extensive use of the chain rule. It is almost identical to the fully-turbulent case, with the only difference being that the turbulent eddy-viscosity is multiplied through by the intermittency function. Note that the intermittency function is treated as a constant in the differentiation with respect to $\mathbf{Q}$, since it only depends explicitly on the fixed transition points, $\mathbf{X}_\mathrm{f}$. The matrix has dimensions $N_\mathrm{Q} \times N_\mathrm{Q}$, where N_Q was defined in Section 2.2 as $N_\mathrm{Q} = 5N_\mathrm{G}$, where N_G is the total number of nodes in the grid.

$\frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial \mathbf{Q}}$: A matrix containing the sensitivities of the transition residuals to the flow variables. The matrix (pre-transpose) has dimensions $N_\mathrm{tr} \times N_\mathrm{Q}$. Note that the transition residuals depend on the predicted transition points, $\mathbf{X}_\mathrm{p}$, which, in turn, depend on the transition criteria. Due to the complexity and non-locality involved in computing the transition criteria, which requires integrations and interpolations in both the off-wall and streamwise directions, the complex-step approximation was selected to accurately and efficiently compute these sensitivities. The complex-step method [76, 113] provides accurate first-order derivatives while altogether avoiding the subtractive cancellation issues that plague the finite-difference gradient. The application requires that the transition module be re-written using complex variable declarations. We can then make complex perturbations to the flow variables to evaluate the sensitivity of $\boldsymbol{\mathcal{R}}_\mathrm{tr}$ as follows [76]:

$$\frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial Q_i} = \frac{\mathrm{Im}\left[\boldsymbol{\mathcal{R}}_\mathrm{tr}(\mathbf{Q} + ih\boldsymbol{e}_i)\right]}{h}, \quad \text{for} \quad i = 1 \ldots N_\mathrm{Q} \tag{4.17}$$

where $\boldsymbol{e}_i$ is the i^{th} unit vector, and $h = 10^{-20}$ is a scalar defining the complex-step perturbation. Further implementation details are discussed in Section 4.6.3.

$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_\mathrm{f}}$: A matrix containing the sensitivities of the flow residuals to the forced transition points. The forced transition points appear in the intermittency function which multiplies the eddy-viscosity. Since there are only N_tr transition points, this matrix is computed using centered-difference approximations that perturb $\mathbf{X}_\mathrm{f}$, while keeping $\mathbf{Q}$ constant. This requires a total of only $2 \times N_\mathrm{tr}$ residual evaluations, which is cheap and also found to be very insensitive to the step-size. The matrix (pre-transpose) has dimensions $N_\mathrm{Q} \times N_\mathrm{tr}$.

$\frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial \mathbf{X}_\mathrm{f}}$: A matrix containing the sensitivities of the transition residuals to the forced transition points. Since the transition residual is defined as $\boldsymbol{\mathcal{R}}_\mathrm{tr} = \mathbf{X}_\mathrm{f} - \mathbf{X}_\mathrm{p}$, the differentiation results in the identity matrix of dimensions $N_\mathrm{tr} \times N_\mathrm{tr}$.

$\frac{\partial \mathcal{J}}{\partial \mathbf{Q}}$: A vector containing the sensitivities of the design objective to the flow variables. Typically the design objectives involve drag and/or lift which are not explicit functions of the turbulence variable or the intermittency function. As such, this vector remains identical to the fully-turbulent case and is analytically computed. The vector has a length of N_Q.

$\frac{\partial \mathcal{J}}{\partial \mathbf{X}_\mathrm{f}}$: A vector containing the sensitivities of the design objective to the forced transition points. Typically the design objectives involve drag and/or lift which are not explicit functions of the forced transition points. For $\mathcal{J} = C_\mathrm{d}$, this vector is the null vector of length N_tr.

4.6.3 Advantages of the Augmented Formulation

The proposed approach has several advantages. First and foremost, the sensitivities of the given transition criterion with respect to the design variables are explicitly incorporated into the adjoint gradient formulation, in turn allowing the optimizer to exploit that information. Second, the non-locality in the given transition criterion is confined to the last N_tr rows and columns of the new Jacobian matrix, $\frac{\partial \overline{\boldsymbol{\mathcal{R}}}}{\partial \overline{\mathbf{Q}}}$. Note that the new adjoint system is only slightly larger than the original system, since the number of additional rows and columns in the new Jacobian is only N_tr; the number of rows and columns in the original Jacobian corresponds to the number of computational grid nodes multiplied by the number of unknowns per node. Thus, the use of the complex-step approximation in the calculation of $\frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial \mathbf{Q}}$ does not incur

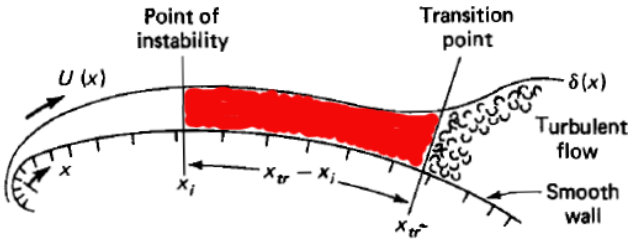

Figure 4.3: Illustration of zone of influence for upper surface transition point.

significant additional expense. Furthermore, the specific nodes involved in satisfying the transition criteria (i.e. from the critical point to the transition point, and from the airfoil surface to the boundary layer edge) represent a known zone of influence. Thus, only that subset of nodes is perturbed when using the complex-step approximation to evaluate $\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}$. This is illustrated in Figure 4.3, where the zone of influence for the upper-surface transition point has been highlighted.

Third, the iterative flow solution procedure used to determine the final forced transition locations need not be explicitly included in the gradient calculation, since the converged RANS solution satisfies the transition criterion to a sufficient tolerance, and the sensitivities of the given transition criterion are included by the addition of the new adjoint variables. Finally, the sensitivities of $\boldsymbol{\mathcal{R}}$ and $\mathcal{J}$ with respect to $\mathbf{X}_{\text{f}}$ (that is, $\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_{\text{f}}}$ and $\frac{\partial \mathcal{J}}{\partial \mathbf{X}_{\text{f}}}$) need not contain any information about the transition criterion.

Of course, no algorithm is without its disadvantages. Those of the proposed augmented formulation include the added computational expense of computing the complex-step approximation of $\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}$, which has been mitigated to a large degree as discussed above, and also the computational expense of solving the augmented system. However, with the appropriate solution strategies (developed in the next section) the method is sufficiently quick to converge.

4.7 Solving the Augmented Adjoint System

4.7.1 Iterative Solution Strategy

An iterative approach was first developed to solve the augmented adjoint system by making multiple calls to the GMRES solver. The solution procedure requires an initial guess for the transition adjoint variables, $\boldsymbol{\psi}_{\text{tr}}^{n=1}$, and is summarized as follows:

1. Use preconditioned GMRES to solve the following linear system of equations for $\boldsymbol{\psi}^{n+1}$,

where n is the iteration counter:

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi}^{n+1} = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^T - \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi}_{\text{tr}}^{n} \,. \tag{4.18}$$

2. Update the $\boldsymbol{\psi}_{\text{tr}}$ vector by solving the following linear system of equations (directly):

$$\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{X}_{\text{f}}}^T \boldsymbol{\psi}_{\text{tr}}^{n+1} = \frac{\partial \mathcal{J}}{\partial \mathbf{X}_{\text{f}}}^T - \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_{\text{f}}}^T \boldsymbol{\psi}^{n+1} \,. \tag{4.19}$$

Note that since $\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{X}_{\text{f}}}$ is the identity matrix and $\frac{\partial \mathcal{J}}{\partial \mathbf{X}_{\text{f}}}$ is the null vector, (4.19) simplifies to

$$\boldsymbol{\psi}_{\text{tr}}^{n+1} = -\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_{\text{f}}}^T \boldsymbol{\psi}^{n+1} \,. \tag{4.20}$$

3. Increment the iteration counter, $n \leftarrow n+1$, and return to step 1.

4. Stop when a given convergence criterion is satisfied (discussed in the next section).

The convergence behaviour and efficiency of the above iterative solution procedure have been studied in detail with respect to the *major* and *minor* iterations required to obtain the new adjoint vector, $\overline{\boldsymbol{\psi}}$. Herein, the *minor* iterations refer to the number of GMRES iterations required to solve the system of linear equations given by (4.18), and the *major* iterations refer to the counter denoted n above: that is, the number of iterations required to solve the overall adjoint system (or the number of updates made to $\boldsymbol{\psi}_{\text{tr}}$). The various aspects of the proposed solution strategy that have been investigated are summarized as follows:

- The selection of the **initial guess values** for $\boldsymbol{\psi}_{\text{tr}}^{n=1}$ was observed to have little to no effect on the ability of the solver to successfully converge, or on the final converged values. An initial guess of zero was initially selected for the entries of the $\boldsymbol{\psi}_{\text{tr}}$ vector, however, it was later found that setting the initial guess to an order of magnitude of 0.01 results in a marginal speed-up of $\sim 5\%$.

- There are two **convergence criteria** selected for the iterative procedure. The first makes use of the L_2-norm of the augmented adjoint system, such that

$$\left|\left| \frac{\partial \overline{\boldsymbol{\mathcal{R}}}}{\partial \overline{\mathbf{Q}}}^T \overline{\boldsymbol{\psi}} - \frac{\partial \mathcal{J}}{\partial \overline{\mathbf{Q}}}^T \right|\right|_2 \leq \epsilon_{\text{adj}} \,. \tag{4.21}$$

A value of $\epsilon_{\text{adj}} = 10^{-4}$ ensures that the final solution vector, $\overline{\boldsymbol{\psi}}$, sufficiently satisfies the augmented system of equations. This corresponds to a reduction of approximately

seven orders of magnitude relative to the initial L_2-norm. The second criterion ensures that the change in the L_2-norm of the gradient vector, $|| \left(\frac{d\mathcal{J}}{d\mathbf{x}}\right)^{n+1} - \left(\frac{d\mathcal{J}}{d\mathbf{x}}\right)^{n} ||_2$, between major iterations is less than 10^{-6}. Figure 4.4(a) shows an example of the convergence behaviour of the iterative adjoint solve with the L_2-norms of the adjoint system as well as the changes in the gradient vector plotted against the total number of minor (or GMRES) iterations. Note that each point on the plot represents a major iteration.

- Initially it was expected that the **relative convergence tolerance in GMRES** may be reduced only one or two orders of magnitude before updating the $\boldsymbol{\psi}_{\text{tr}}$ vector at each major iteration. However it has been observed that a loose convergence of each GMRES solve results in an adjoint vector that does not satisfy the system of equations. In other words, the L_2-norm given by (4.21) is not sufficiently reduced and tends to flat-line too early when GMRES is not sufficiently converged at each major iteration. When a relative reduction of 12 orders of magnitude is used as the GMRES convergence criterion, the L_2-norm of the augmented system is sufficiently reduced before it begins to flat-line, as shown in Figure 4.4(a).

- **Convergence acceleration** of the iterative procedure is achieved by under-relaxing the updates made to $\boldsymbol{\psi}_{\text{tr}}$, such that

$$\boldsymbol{\psi}_{\text{tr}}^{n+1} = \boldsymbol{\psi}_{\text{tr}}^{n} + \omega \left(\boldsymbol{\psi}_{\text{tr}}^{n+1} - \boldsymbol{\psi}_{\text{tr}}^{n} \right) \ . \tag{4.22}$$

  A value for the **under-relaxation parameter** of $\omega = 0.7$ was selected through numerical studies. Figure 4.4(b) shows the typical number of minor iterations required to satisfy the convergence criterion ($\epsilon_{\text{adj}} = 10^{-4}$) for different values of ω; it can be seen that a value of 0.7 is an appropriate choice to minimize the computational cost. Note that when the updates are not under-relaxed ($\omega = 1$) the $\boldsymbol{\psi}_{\text{tr}}$ values tend to oscillate between two solution states and the system fails to converge entirely.

4.7.2 Non-Iterative Solution Strategy

While the iterative solution strategy (discussed above) provides accurate gradient vectors (to be demonstrated in Chapter 5), the gradient computation unfortunately requires upwards of 10-20 major iterations to converge. This is in contrast to a single call to GMRES (i.e. one major iteration) required for the original fully-turbulent (or fixed-transition) adjoint system given by (4.14). Furthermore, for lift-constrained drag minimization, two adjoint gradients are required, and thus, the GMRES solver will be invoked upwards of 20-40 times

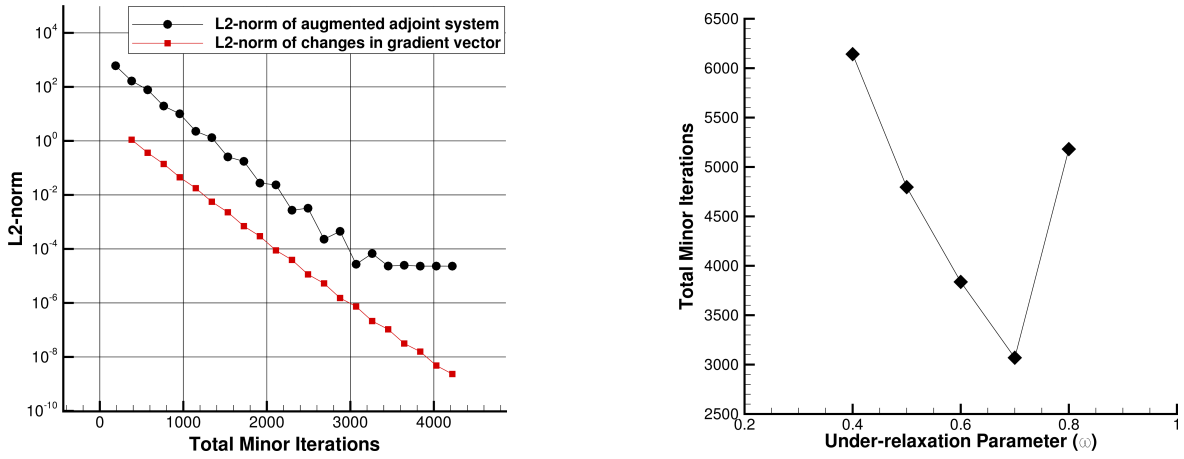

(a) Convergence of augmented adjoint system

(b) Under-relaxation parameter study

Figure 4.4: Augmented adjoint system: convergence and efficiency studies

for the convergence of two augmented systems. This motivates the development of a more efficient solution strategy for the augmented system. Such a strategy is possible by solving three intermediate problems defined by making some strategic substitutions prior to directly computing ψ and ψ_{tr}.

We begin by writing the augmented system of equations (4.15) as follows:

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi} + \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi}_{\text{tr}} = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^T , \tag{4.23}$$

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_{\text{f}}}^T \boldsymbol{\psi} + \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{X}_{\text{f}}}^T \boldsymbol{\psi}_{\text{tr}} = \frac{\partial \mathcal{J}}{\partial \mathbf{X}_{\text{f}}}^T . \tag{4.24}$$

Let us now write the identity matrix as the product of the Jacobian-transposed with its own inverse, that is,

$$\mathcal{I} = \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \left(\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \right)^{-1} \tag{4.25}$$

By pre-multiplying the second term on the left-hand side of (4.23) by (4.25) we have

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi} + \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \left(\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \right)^{-1} \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T \boldsymbol{\psi}_{\text{tr}} = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^T \tag{4.26}$$

From the quadruple-product term in (4.26), we can define an intermediate matrix, $M \in \mathbb{R}^{N_{\text{Q}} \times N_{\text{tr}}}$, as

$$M = \left(\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \right)^{-1} \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T . \tag{4.27}$$

Solving for the two columns that make up M (since $N_{\text{tr}}=2$) defines the first two intermediate problems that must be solved. To define the third intermediate problem we first substitute the matrix M into (4.26) and rearrange to get

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T (\boldsymbol{\psi} + M\boldsymbol{\psi}_{\text{tr}}) = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^T . \tag{4.28}$$

Next we define a new vector, $\widetilde{\boldsymbol{\psi}} \in \mathbb{R}^{N_{\mathbf{Q}}}$, as

$$\widetilde{\boldsymbol{\psi}} = \boldsymbol{\psi} + M\boldsymbol{\psi}_{\text{tr}} , \tag{4.29}$$

and substitute it into (4.28) to get a system of equations of the form

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T \widetilde{\boldsymbol{\psi}} = \frac{\partial \mathcal{J}}{\partial \mathbf{Q}}^T . \tag{4.30}$$

Solving for $\widetilde{\boldsymbol{\psi}}$ represents the third and final intermediate problem. Note that the large, sparse system of equations given by (4.30) is identical to the original adjoint system (4.14) without any augmentation.

To solve for the matrix M we first rewrite (4.27) as follows:

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T M = \frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T . \tag{4.31}$$

Letting the subscripts 1 and 2 denote the first and second columns of both M and $\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T$, we can then solve the following systems of equations for each column of M independently:

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T M_1 = \left(\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T \right)_1 , \tag{4.32}$$

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{Q}}^T M_2 = \left(\frac{\partial \boldsymbol{\mathcal{R}}_{\text{tr}}}{\partial \mathbf{Q}}^T \right)_2 . \tag{4.33}$$

Since (4.31), (4.32), and (4.33) all share the same Jacobian matrix on the left-hand side, we can solve each system using the same preconditioned GMRES approach previously discussed.

All that remains is to solve for $\boldsymbol{\psi}$ and $\boldsymbol{\psi}_{\text{tr}}$. From (4.29) we have

$$\boldsymbol{\psi} = \widetilde{\boldsymbol{\psi}} - M\boldsymbol{\psi}_{\text{tr}} , \tag{4.34}$$

and substituting this expression into (4.24) we have

$$\frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_\mathrm{f}}^T \left(\widetilde{\boldsymbol{\psi}} - M \boldsymbol{\psi}_\mathrm{tr} \right) + \frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial \mathbf{X}_\mathrm{f}}^T \boldsymbol{\psi}_\mathrm{tr} = \frac{\partial \mathcal{J}}{\partial \mathbf{X}_\mathrm{f}}^T \,. \tag{4.35}$$

Having previously solved for $\widetilde{\boldsymbol{\psi}}$ and M, we can solve for $\boldsymbol{\psi}_\mathrm{tr}$ directly (using Cramer's rule, for example) by rearranging (4.35) to get the following system of equations (which is of length $N_\mathrm{tr} = 2$):

$$\left(\frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial \mathbf{X}_\mathrm{f}}^T - \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_\mathrm{f}}^T M \right) \boldsymbol{\psi}_\mathrm{tr} = \frac{\partial \mathcal{J}}{\partial \mathbf{X}_\mathrm{f}}^T - \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_\mathrm{f}}^T \widetilde{\boldsymbol{\psi}} \,. \tag{4.36}$$

Note that since $\frac{\partial \boldsymbol{\mathcal{R}}_\mathrm{tr}}{\partial \mathbf{X}_\mathrm{f}}$ is the identity matrix and $\frac{\partial \mathcal{J}}{\partial \mathbf{X}_\mathrm{f}}$ is the null vector, (4.36) simplifies to

$$\left(\mathcal{I} - \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_\mathrm{f}}^T M \right) \boldsymbol{\psi}_\mathrm{tr} = - \frac{\partial \boldsymbol{\mathcal{R}}}{\partial \mathbf{X}_\mathrm{f}}^T \widetilde{\boldsymbol{\psi}} \,. \tag{4.37}$$

Finally, the vector $\boldsymbol{\psi}$ can be solved for directly from (4.34).

To summarize, the following five steps are performed to solve the augmented adjoint system:

1. Use preconditioned GMRES to solve (4.30) for $\widetilde{\boldsymbol{\psi}}$.
2. Use preconditioned GMRES to solve (4.32) for M_1.
3. Use preconditioned GMRES to solve (4.33) for M_2.
4. Solve (4.37) directly for $\boldsymbol{\psi}_\mathrm{tr}$.
5. Solve (4.34) directly for $\boldsymbol{\psi}$.

It is important to note that if the optimization problem requires any additional adjoint gradient computations for any nonlinear constraints (such as lift and/or moment constraints), it is not necessary to repeat steps 2 and 3 above since they are based solely on the left-hand side matrix of the augmented system, which does not depend on the objective or constraints under consideration.

The non-iterative solution strategy has several key advantages over the iterative solution strategy discussed in the previous section. First, it does away with the issues surrounding an iterative procedure, including the need for an initial guess on $\boldsymbol{\psi}_\mathrm{tr}$, convergence criteria, and the studies required for convergence acceleration and robustness. Furthermore, the accuracy of the resulting gradient vectors do not depend in any way on the convergence of the major iterations that exist for the iterative approach. Perhaps most important from a practical

standpoint, the computational cost of the noniterative strategy is many times cheaper than the iterative approach. The computational cost of the various gradient evaluation techniques will be quantified and compared in Section 5.6.1. For a single adjoint gradient, the non-iterative method requires only three calls to GMRES rather than 10-20 calls. For a lift-constrained drag minimization, the cost benefits are even greater, requiring a total of only four calls to GMRES rather than 20-40 calls. As a result of these advantages, the non-iterative method is selected as the solution strategy for the augmented adjoint system.

4.8 Multipoint Optimization Strategy

Multipoint optimization is used in this work to demonstrate that the resulting aerodynamic designs perform reasonably well over a given flight envelope and are robust to perturbations. This is particularly important in the design of NLF airfoils, which, in order to maximize the extent of laminar flow, tend to take the boundary layer very close to the point of separation prior to pressure recovery. As such, the off-design performance of NLF airfoils must be considered during the design process to ensure practical and robust designs. We use the methodology of Buckley and Zingg [14, 15] to perform multipoint optimization capable of handling a comprehensive set of aerodynamic design requirements.

In particular, we are interested in considering a range of Reynolds numbers, Mach numbers, and aircraft weights (W). We keep the cruise altitude constant (however it can also be included) and by specifying a range of Mach numbers and aircraft weights, we can obtain the corresponding range of Reynolds numbers and lift requirements. The optimizer then minimizes the weighted integral of the objective (in our case, the drag coefficient) subject to the lift constraints (one for each design point). Each operating point has an associated angle of attack, all of which are included as additional design variables. The weighted integral is defined as [15]:

$$\int_{W_1}^{W_2}\int_{M_1}^{M_2} C_d\left(M,W\right)\mathcal{Z}\left(M,W\right)dMdW \tag{4.38}$$

where $\mathcal{Z}$ is a weighting function to be specified by the designer. This weighting function allows the designer to specify the importance of each design point according to their own priorities or fleet mission data. The objective function, $\mathcal{J}$, is an approximation to (4.38) given by

$$\mathcal{J} = \sum_{i=1}^{N_\mathrm{W}}\sum_{j=1}^{N_\mathrm{M}} \mathcal{T}_{i,j}\, C_d\left(M_i,W_j\right)\mathcal{Z}\left(M_i,W_j\right)\Delta M\Delta W\ , \tag{4.39}$$

where N_{M} and N_{W} represent the number of quadrature points, and ΔM and ΔW are the corresponding spacings between quadrature points. The $\mathcal{T}_{i,j}$ are the associated quadrature weights used to approximate the integral. In this work, the trapezoidal quadrature rule is employed.

The above multipoint formulation requires one flow solution and two adjoint solutions (for the gradients of the objective and lift-constraint) for each operating point. Buckley and Zingg [15] have parallelized the multipoint framework such that multiple processors compute the necessary objective, constraint, and gradient information. This approach has been shown to be an effective technique for robust and efficient aerodynamic design over a range of operating conditions [14, 15]. Full details of the various operating conditions studied and their associated weights are presented in Section 6.4.

Chapter 5

Verification and Validation

In this chapter we present the verification and validation of the free-transition RANS solver and the discrete-adjoint gradient computation required for gradient-based ASO. These represent the two components of the ASO framework that have been heavily modified to incorporate and exploit transition prediction. The verification process, as defined in the AIAA guidelines [93], is the "process of determining that a model implementation accurately represents the developer's conceptual description of the model and the solution to the model" [93]. Defined differently is the validation process, which is the "process of determining the degree to which a model is an accurate representation of the real world from the perspective of the intended uses of the model"[93].

We begin by considering the verification of the transition region models in order to better understand and compare the manner in which they transition the flow from a laminar to a turbulent state. Next, the accuracy of the boundary-layer edge finding methods are verified by comparison to a boundary-layer code, which is then followed by a similar verification of the boundary-layer properties, involving grid convergence studies. The transition prediction capabilities of the RANS solver are then verified and validated by comparison to available numerical and experimental data. Finally, we consider the accuracy assessment of the discrete-adjoint gradient computation by comparison to the more costly finite-difference and flow-sensitivity gradient computation methods.

As stated in Section 1.4, it is not the goal of this work to develop a new transition criterion. Improving the accuracy and extending the applicability of correlation-based transition prediction criteria is a challenging and active area of research, but not considered in this work. To be clear, the verification and validation of the transition prediction framework presented herein is aimed primarily at ensuring the correct implementation and application of the selected criteria in the present RANS solver.

5.1 Transition Region Models

The transition region models are responsible for increasing the turbulent eddy-viscosity in the vicinity of the predicted or prescribed transition locations. As discussed in Section 3.1, no attempt is made to accurately model the unsteady, nonlinear physical processes that lead to transition; indeed, this is not possible using the present approach. In this section our goal is to characterize the transition region models in order to better understand, verify, and compare the manner in which they perform the ramp-up of eddy-viscosity. Of particular importance is the smoothness with which the eddy-viscosity is seeded into the boundary layer over the given transition length.

The first study demonstrates that the eddy-viscosity is ramped-up as expected from the laminar (freestream) state, with $\gamma = 0$, to the fully-turbulent state, with $\gamma = 1$. This study is performed using the intermittency function defined by (3.5) and by prescribing (or forcing) the upper and lower transition points at various locations on a NACA-0012 at $Re = 10 \times 10^6$, $M = 0.20$, $AoA = 0°$, and using a 575×224 C-grid. Figure 5.1 shows the eddy-viscosity contours on a log-scale for the various prescribed transition locations (upper and lower surface transition points fixed at the same chord position). It can be seen that the eddy-viscosity is indeed introduced as expected in the vicinity of the transition locations. Note that in Figures 5.1(c) and 5.1(d) the transition has been forced to occur further downstream than would be predicted by the AHD and e^N criteria (presented in Section 5.5.1).

Next, we compare the SA trip functions to the intermittency functions (all presented in Section 3.1) to better understand their behaviour. Figure 5.2 compares the ramp-up of the eddy-viscosity using the different transition region models. The eddy-viscosity is plotted at each streamwise station by following an upper-surface gridline located inside of the boundary layer, 40 nodes off the surface at an off-wall distance of approximately $10^{-4}c$. It can be observed that the SA trip functions increase the eddy-viscosity in a rather sharp manner, and slightly upstream of the specified transition location [112]. Recall that the SA trip functions do not allow for the specification of a transition length. Regardless of the values used for the constants c_{t_1} and c_{t_2} in (3.1), no change in the sharp increase of the eddy-viscosity was observed. On the other hand, the intermittency function typically provides a much smoother introduction of the eddy-viscosity in accordance with the transition length, which in this case is fixed at 10% chord. While the present intermittency function, given by (3.5), has a slight overshoot in the eddy-viscosity at the end of the transition length, this has little effect on the resulting drag value or the design space.

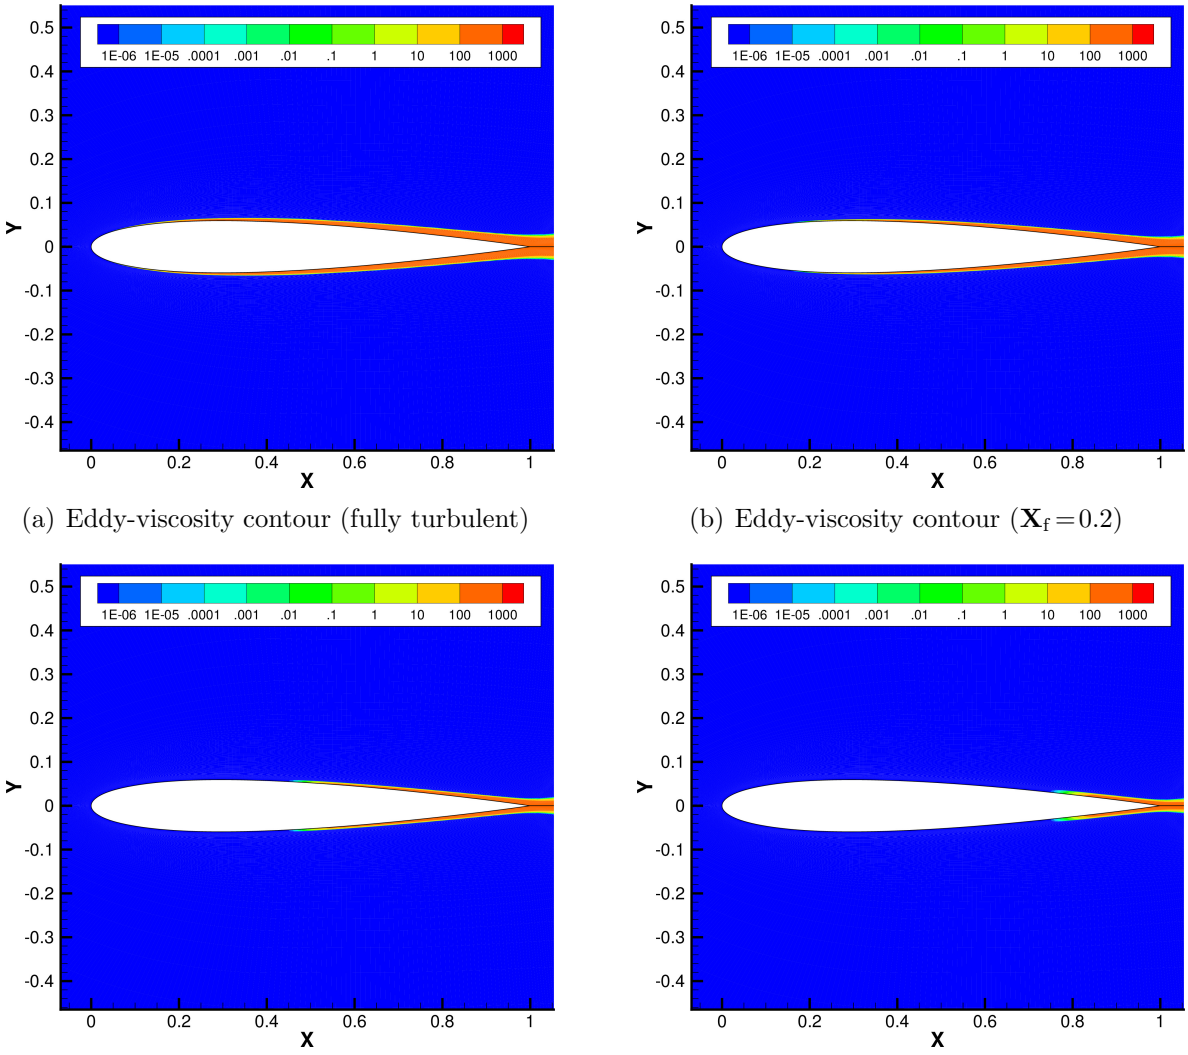

(a) Eddy-viscosity contour (fully turbulent)

(b) Eddy-viscosity contour ($\mathbf{X}_\mathrm{f} = 0.2$)

(c) Eddy-viscosity contour ($\mathbf{X}_\mathrm{f} = 0.5$)

(d) Eddy-viscosity contour ($\mathbf{X}_\mathrm{f} = 0.8$)

Figure 5.1: Eddy-viscosity contours for different forced transition locations

5.2 Ensuring a Smooth Design Space

The smoothness of the initial ramp-up of the eddy-viscosity was found to be particularly important to the smoothness of the design space. A sharp increase in the eddy-viscosity was found to cause locally non-smooth design spaces for the grids typically employed for transition prediction. This is demonstrated by making very small changes to a given design variable and calling the flow solver (with free transition) on each perturbed geometry. As an example, Figure 5.3 plots the drag values against the magnitude of the given design variable: the y-coordinate of an upper surface B-spline control point located at approximately 20% chord. For both the present intermittency function and the SA trip functions, the general

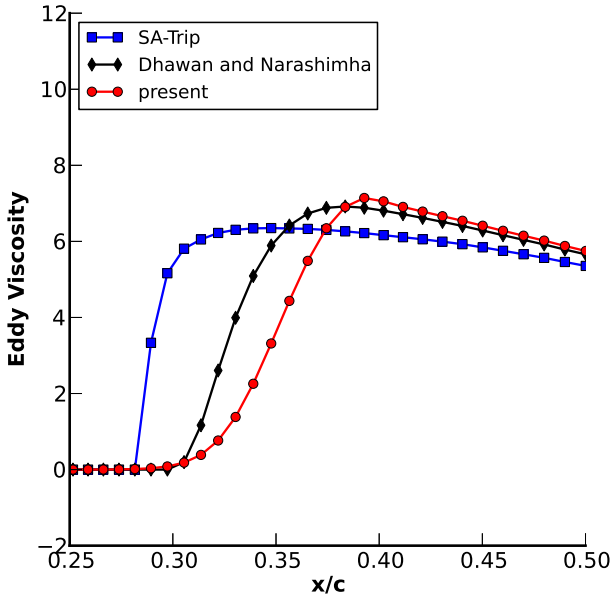

Figure 5.2: Eddy viscosity ramp-up on a NACA-0012 with transition at 30% chord

trend is a decrease in drag with an increase in the design variable magnitude. However, upon closer inspection, a step-wise pattern in the drag was observed when under-resolving the sharp ramp-up given by the SA model trip functions. This noise in the design space is typically not present when using an intermittency function, indicative of an eddy-viscosity ramp-up that is sufficiently resolved by the typical grids used for transition prediction in this work. The step-wise design space depicted in Figure 5.3 can be detrimental to the progress and convergence of gradient-based optimization, due to the local minima that arise. This behaviour was first observed when making use of a finite-difference approximation, for which small perturbations to each design variable are required to compute the gradient vector.

Figure 5.3 also demonstrates the importance of selecting a tight convergence tolerance on the transition residual, $\boldsymbol{\mathcal{R}}_{\text{tr}}$. The intermittency function with a loose convergence of $\epsilon_{\text{tp}} = 0.03\text{c}$, labelled *Intermittency (Loose)*, shows a non-smooth step in the drag values due to the imprecision of the transition prediction. When using a tight tolerance on the transition residual, $\epsilon_{\text{tp}} = 10^{-8}\text{c}$, labeled *Intermittency (Tight)*, the design space is sufficiently smooth for the purpose of gradient-based optimization. Consequently, the intermittency function along with a tight convergence tolerance on the transition residuals are selected for all optimization cases performed in this work.

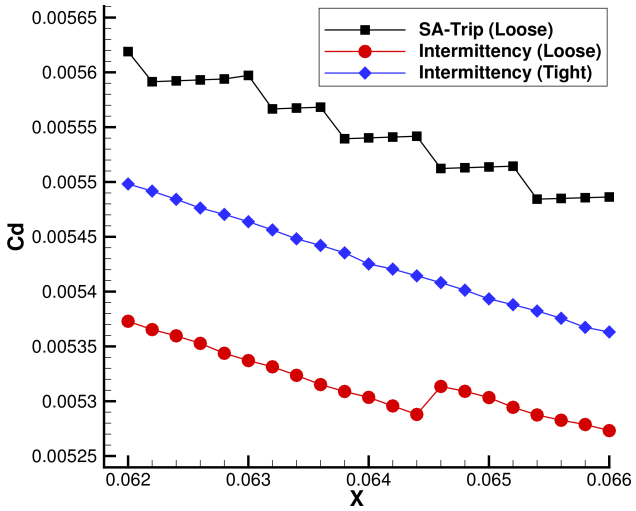

Figure 5.3: Design space comparison using different transition region models and convergence tolerances

5.3 Boundary-Layer Edge Finding Methods

In order to verify the three boundary-layer edge-finding methods described in Section 3.3, numerical flow solutions using fixed transition locations are performed. The results from Optima2D are compared to numerical results obtained from XFOIL, developed by Drela [28]. XFOIL is a two-dimensional, incompressible flow solver which couples an inviscid solver to a viscous boundary-layer code. The inviscid formulation in XFOIL is a linear vorticity-streamfunction panel method. The viscous flow in the boundary layer and wake is modeled with a two-equation lagged dissipation integral boundary-layer method [28].

Flow solutions were computed on the NACA-0012 airfoil at $Re = 1 \times 10^6$, $M = 0.20$ and zero incidence, using a C-grid with 321×384 nodes. In both solvers, transition was fixed at 50% chord on the top and bottom surfaces of the airfoil. A comparison of the three edge-finding methods in Optima2D is presented in Figure 5.4, along with the edge velocity obtained using XFOIL. The comparison verifies the ability of the various methods to define the boundary-layer edge, without the use of a boundary-layer solver. Good agreement is observed between the methods in Optima2D and XFOIL in both the laminar and turbulent regions, although our interest here is restricted to the laminar region for the purpose of predicting transition.

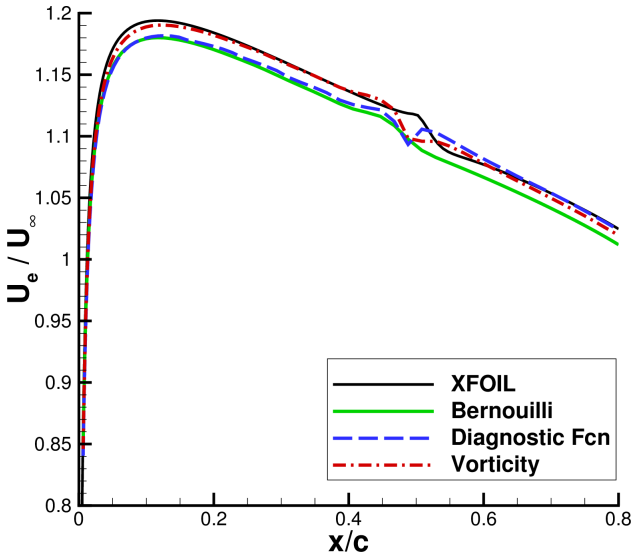

Figure 5.4: NACA-0012 verification of edge finding methods

5.4 Accuracy Assessment of Boundary-Layer Properties

For the transition prediction strategies employed in this work, accurately computing the boundary-layer properties is critical to accurately predicting transition. It has been documented by several researchers [23, 62, 63], and also found in this work, that the various transition prediction criteria are very sensitive to any inaccuracies in the laminar boundary-layer properties and/or velocity profiles. To accurately compute the boundary-layer properties, we must ensure that we are accurately resolving the laminar boundary layer in our RANS solver.

Several authors have performed grid convergence studies in the process of verifying their free-transition RANS solvers, making use of either local or non-local transition prediction. For example, Mayda found that C-grids of approximately 576×164 were required to accurately predict transition using a non-local approach [78] with a matrix dissipation scheme. Grid convergence studies performed by Khayatzadeh and Nadarajah using the local $\gamma - \widetilde{Re}_{\theta_t}$ approach found that – even when boundary-layer properties are not explicitly computed – a resolved boundary layer was still necessary for accurate transition prediction; as such, their typical C-grid was of a size equal to 514×256 [54].

Grid convergence studies are thus a critical step in the verification and validation of any free-transition RANS solver. Here we perform grid convergence studies and compare the computed boundary-layer properties to numerical results obtained using the integral boundary-layer code found in MSES. An ICEM C-grid blocking strategy was developed that allows for grid refinement in the boundary layer, while keeping constant the number of nodes

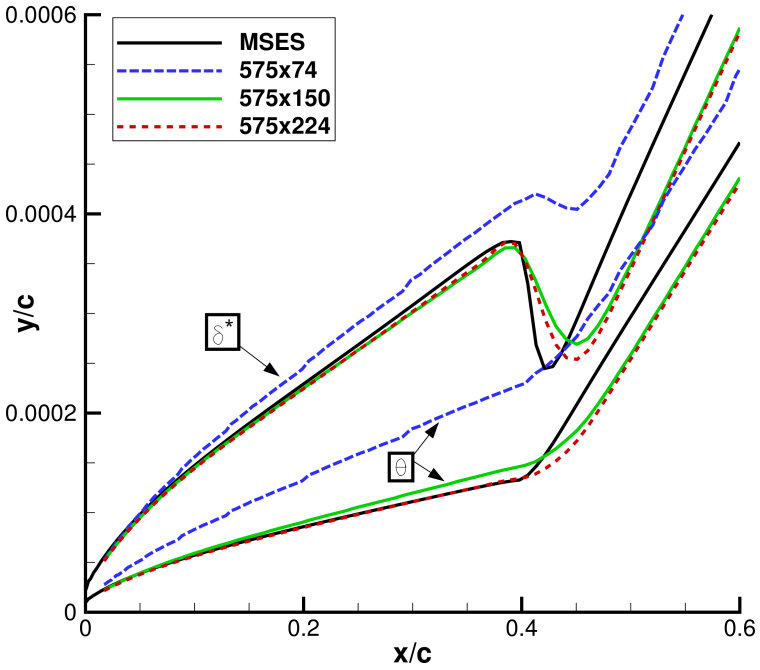

Figure 5.5: NACA-0012 off-wall grid refinement study on displacement and momentum thicknesses

in the outer region (residing entirely outside of the boundary layer). All results are obtained using scalar dissipation; however, the effects of scalar vs. matrix dissipation are discussed later in this section.

We begin with an off-wall grid refinement study that considers C-grids consisting of 575×74, 575×150, and 575×224 nodes, with the distance to the first grid node from the airfoil surface kept constant at 1×10^{-6}. Figure 5.5 presents a grid convergence study of the boundary-layer displacement thickness (top) and momentum thickness (bottom), for the NACA-0012 airfoil at $Re = 10\times10^6$, $M = 0.20$, $AoA = 0°$. The results shown here were obtained using the compressible Bernoulli edge-finding method; however, similar results were found using the other edge-finding methods. At first glance it appears as though δ^* and θ show reasonably good agreement for both the intermediate and fine grids. However, the differences in their accuracy are made more evident by comparing their ratio, that is, the shape factor, $H = \frac{\delta^*}{\theta}$.

Figure 5.6 presents a grid convergence study of the boundary-layer shape factor (H) for the same NACA-0012 test case and demonstrates the importance of sufficient grid resolution in the off-wall direction. The results also verify the current implementation and demonstrate that with reasonable grid density, sufficient accuracy of the boundary-layer properties can be computed directly from the Navier-Stokes solution, confirming similar results found by Brodeur and van Dam [13].

Grid convergence in the streamwise direction was also studied. In general, it was observed that the accuracy of the shape factor for fully-attached flow is only affected in the vicinity of the transition point. A streamwise nodal distribution of approximately 500 nodes around the airfoil was found to be adequate. The high streamwise resolution is also advantageous for resolving the transition region intermittency functions and providing a smooth design space, as well as more accurately capturing the small separation bubbles that may be encountered during the optimization cycle for the subsonic and transonic applications of interest.

Effect of Numerical Dissipation

All results presented herein (with the exception of Figure 5.7) have been obtained using a scalar dissipation scheme. Investigations into the accuracy of the boundary-layer properties revealed that the scalar dissipation scheme can be overly dissipative. A clear improvement in the accuracy of the boundary-layer properties was observed when using the matrix dissipation scheme. Similar numerical dissipation effects on the boundary-layer properties have been reported by Mayda [78]. The improvement can be seen in Figure 5.7, which compares the shape factor obtained by scalar and matrix dissipation for the 575×150 grid, for the same NACA-0012 test case.

Unfortunately, the matrix dissipation algorithm has not yet been differentiated in the second-order flow Jacobian, $\mathcal{A}_2$, and is thus not employed for the purposes of adjoint-based optimization. The task of differentiating the matrix dissipation model is left to future work. As mentioned in Section 7.2, the robustness of the flow solver using matrix dissipation for free-transitional flows should be further evaluated, as reducing dissipation typically reduces robustness. For all optimization cases presented herein, the more robust scalar dissipation model is employed in conjunction with the fine 575×224 grid to ensure a sufficiently resolved boundary layer.

5.5 Transition Prediction

5.5.1 NACA-0012 Airfoil

The predictive capability of the transition prediction framework is first verified by comparison to numerical results generated using XFOIL on the NACA-0012 airfoil. The transition prediction in XFOIL makes use of the simplified e^N envelope method developed by Drela [28] and thus provides a good benchmark for verification.

In Figure 5.8, transition prediction results are shown for a 575×224 C-grid around the

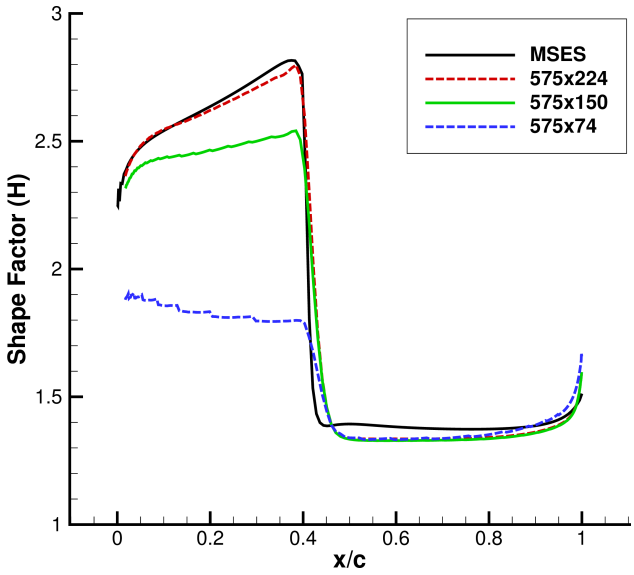

Figure 5.6: NACA-0012 off-wall grid refinement study on shape factor

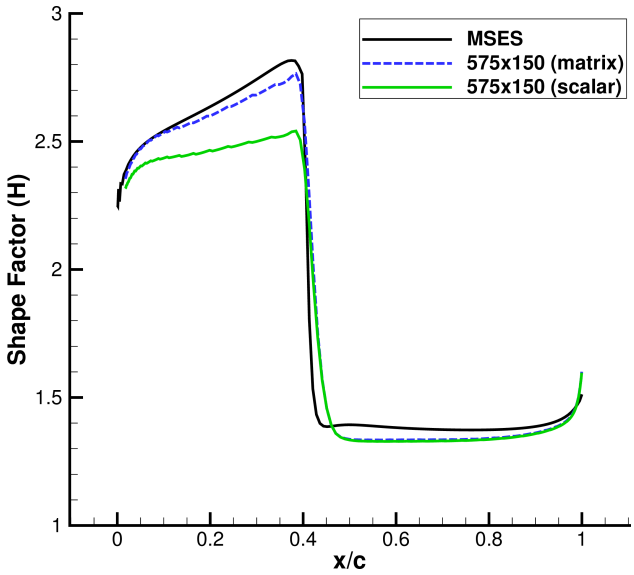

Figure 5.7: NACA-0012 effect of numerical dissipation on shape factor

NACA-0012 airfoil at $Re = 10 \times 10^6$, $M = 0.2$, and a freestream turbulence intensity (T_u) of 0.1% (corresponding to an N-factor of 8 for the e^N criterion). The Optima2D (O2D) points represent the final converged transition locations using the AHD and e^N criteria as indicated. Good agreement between XFOIL and both transition criteria in Optima2D is observed. The AHD criterion predicts transition upstream of the e^N envelope method. Similar trends have been found by Cliquet and Arnal [23] and Streit *et al.* [117].

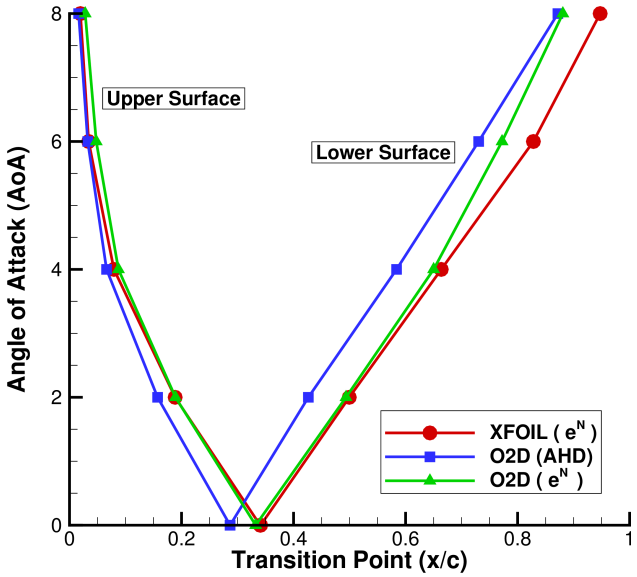

Figure 5.8: NACA-0012 transition prediction verification

5.5.2 NLF-0416 Airfoil

Validation of the transition prediction framework and the transition criteria has been carried out by comparison to available experimental transition data for the NLF-0416 airfoil developed by Somers [110]. The experimental results were obtained in the Langley Low Turbulence Pressure Tunnel (LTPT) using microphoned pressure taps [110]. The resolution of the experiments corresponds to the physical spacing of the microphoned taps along the chord of the airfoil.

Presented herein are the test case results for a 575×224 C-grid around the NLF-0416 airfoil at $Re = 4\times10^6$, $M = 0.2$, and $T_u = 0.1\%$ (and $N_{crit} = 8$ for the e^N criterion in both Optima2D and XFOIL). The transition points predicted by both Optima2D and XFOIL are presented in Figure 5.9, along with the wind tunnel experimental data. The results of this test case show that the predictive capabilities of Optima2D match closely with the published experimental results over a range of lift coefficients.

Figure 5.9(c) presents the drag polar for the NLF-0416 airfoil using both Optima2D and XFOIL. Good agreement is observed between the experimental results and the drag polars computed using both Optima2D and XFOIL. In Somers' report [110], the experimental results were obtained at $Re = 4\times10^6$ and $M = 0.1$. The freestream turbulence intensity, T_u, was unfortunately not published for the NLF-0416 experiments. It is possible that the wind tunnel may have had lower or higher T_u than the 0.1% used for the computations. The slightly higher Mach number was used for the computational results to improve the convergence of the flow solver, while the effect on the results is negligible.

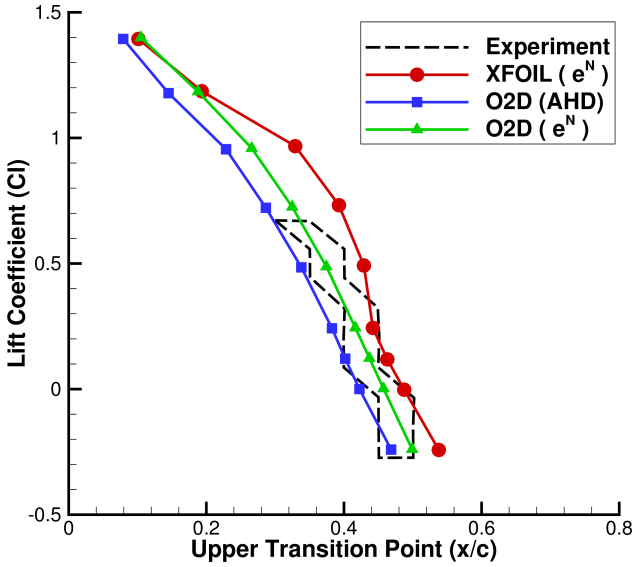

(a) Upper surface transition prediction

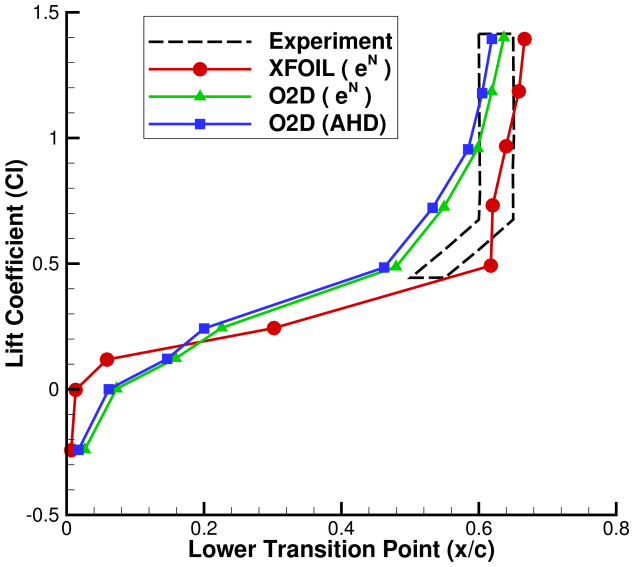

(b) Lower surface transition prediction

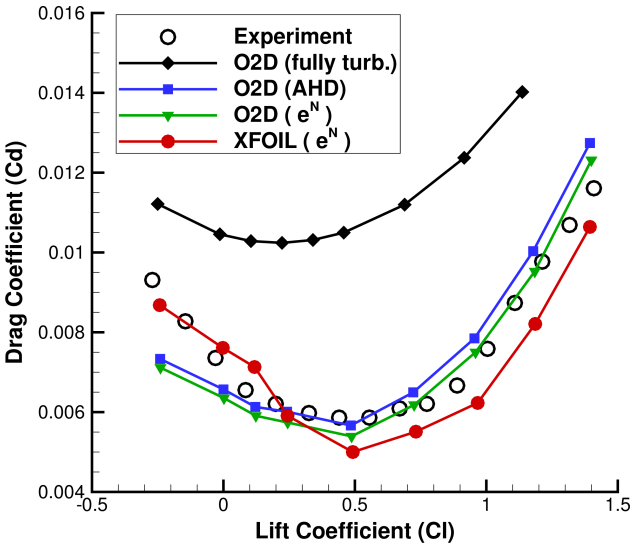

(c) Drag polars

Figure 5.9: NLF-0416 transition prediction and drag polar validation

5.5.3 S809 Airfoil

The S809 wind turbine blade was also experimentally investigated by Somers [111] in the Delft University low-turbulence wind tunnel at a Reynolds number of 3×10^6. Given the wind tunnel's freestream turbulence intensity level of approximately 0.03% [111], Mack's relationship (1.5) is used to compute a critical N-factor of 11 in both XFOIL and Optima2D.

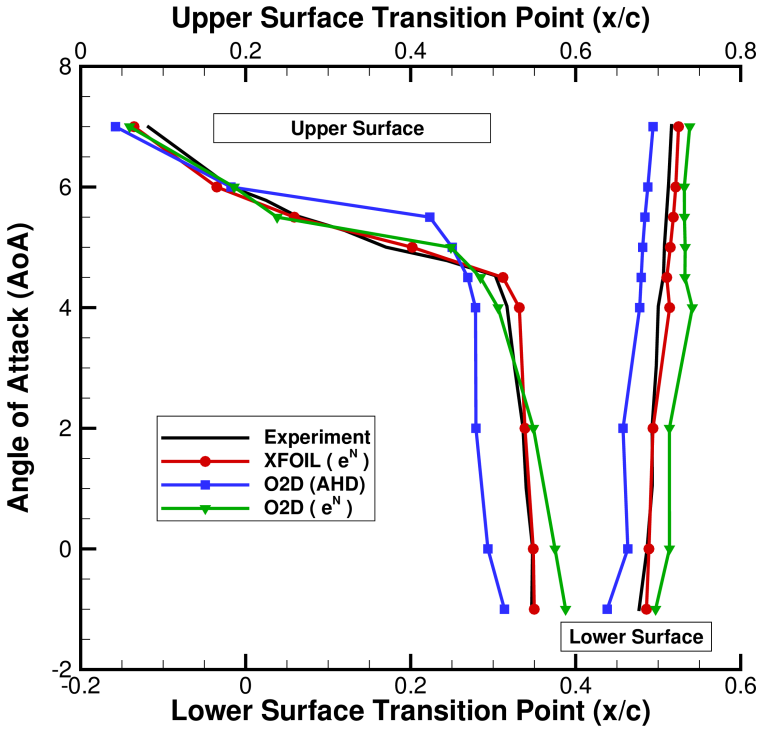

Figure 5.10: S809 transition prediction validation

Similar to the NLF-0416 investigation, transition locations were determined experimentally using microphone measurements [111].

Figure 5.10 compares the experimental data to XFOIL and the two transition criteria in Optima2D. Excellent agreement is again observed between XFOIL and Optima2D with the e^N criterion. On the upper surface, at an angle of attack of 5.5°, the AHD criterion shows some significant deviation from the experimental and e^N transition location. The mechanism for transition on the lower surface is that of a separation bubble, which is reported by Somers [111] to move slightly downstream as the angle of attack increases. The same trend is captured by XFOIL and by both criteria in Optima2D, with the AHD criterion using the separation point to approximate the point of transition.

5.6 Accuracy and Efficiency of Gradient Computations

To verify the accuracy of the gradient evaluation and, in particular, the augmented adjoint formulation for transition prediction (as presented in Sections 4.6 and 4.7) we compare the gradient vectors obtained using the following evaluation techniques (with the following labels):

- **Finite-difference gradient method (FD):** Computes a centered-difference approximation to the gradient vector, given by (4.9).
- **Flow (or direct) sensitivity method (SN):** Solves (4.10) and (4.12) making use of the same augmented Jacobian matrix as used in the discrete-adjoint formulation. An analogous iterative solution procedure to that described in Section 4.7.1 is employed to solve (4.12) for each design variable.
- **Iterative augmented adjoint method (iAD):** Solves (4.15) making use of the iterative solution strategy described in 4.7.1.
- **Noniterative augmented adjoint method (nAD):** Solves (4.15) making use of the noniterative solution strategy described in 4.7.2

The accuracy assessment is carried out by performing a single iteration of the optimization, using each of the above listed gradient evaluation techniques. As shown in Figure 5.11, the RAE-2822 airfoil geometry is parameterized by seventeen B-spline control points, fourteen of which are selected as the design variables in addition to the angle of attack. The flight conditions are $Re = 15.7\times10^6$, $M = 0.6$, $AoA = 1^\circ$, and the e^N transition prediction criterion is employed with $N_{\text{crit}} = 9$. A step-size of 1×10^{-6} is used for the finite-difference gradient, resulting from a step-size study. Tables 5.1 and 5.2 compare the resulting objective function and lift-constraint gradients, respectively. The first 14 components are the geometric design variables, the last is the angle of attack. The results demonstrate excellent agreement between all methods for computing the gradient vector. The percent difference between the FD and nAD gradients is less than 1% for all components. Since the finite-difference method does not require differentiating the code, it serves to verify that the differentiation was performed accurately in the remaining gradient evaluation techniques. The augmented Jacobian matrix is also verified by the excellent agreement found between the flow sensitivity and adjoint gradient results; these methods employ the same augmented Jacobian matrix, but use entirely different approaches for computing the gradient.

5.6.1 Comparison of Computational Cost

The computational cost of each gradient evaluation technique is now compared. For 15 design variables, a centered-difference approximation requires 30 flow solutions, the flow sensitivity method requires 15 solutions to the augmented flow sensitivity system given by (4.12), and the discrete-adjoint method requires one solution to the augmented adjoint system given by (4.15). Table 5.3 provides the computational cost in terms of total wall-clock time to compute

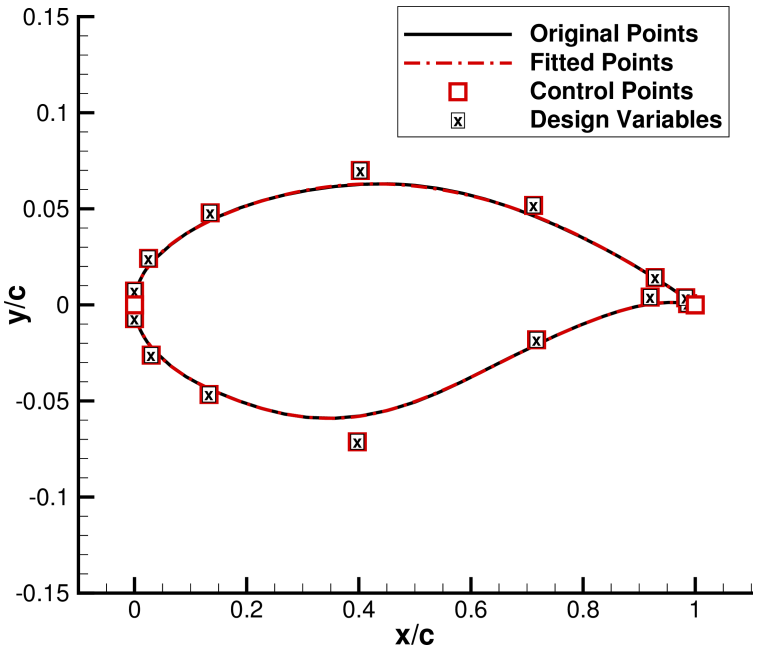

Figure 5.11: RAE-2822 B-spline parameterization

Table 5.1: Comparison of objective function gradient components computed using various gradient evaluation techniques: RAE-2822, $Re = 15.7 \times 10^6$, $M = 0.6$, $AoA = 1°$

Design	nAD	FD$-n$AD		SN$-n$AD		iAD$-n$AD	
Variable	Gradient	Diff.	% Diff.	Diff.	% Diff.	Diff.	% Diff.
1	0.0702067611	-2.2821E-06	-3.2506E-03	1.1238E-11	1.6007E-08	-3.1358E-11	-4.4665E-08
2	0.0245812187	2.7873E-06	1.1339E-02	1.3699E-12	5.5730E-09	-5.4497E-12	-2.2170E-08
3	0.0070351926	1.3651E-05	1.9404E-01	-1.7622E-12	-2.5048E-08	1.8640E-11	2.6495E-07
4	0.0198719593	-9.9423E-06	-5.0032E-02	-1.4039E-11	-7.0647E-08	-4.0467E-10	-2.0364E-06
5	-0.0113855483	-7.8927E-06	6.9322E-02	9.3495E-12	-8.2117E-08	2.3594E-10	-2.0723E-06
6	-0.0285137171	-3.8429E-06	1.3477E-02	1.4446E-11	-5.0662E-08	5.0821E-10	-1.7823E-06
7	0.0005086987	1.3408E-06	2.6358E-01	-7.0467E-12	-1.3852E-06	-1.1293E-10	-2.2199E-05
8	0.1427562391	6.9009E-06	4.8340E-03	-3.6426E-11	-2.5516E-08	-1.2300E-09	-8.6163E-07
9	-0.0254696647	4.7337E-06	-1.8586E-02	6.9483E-12	-2.7281E-08	2.0961E-10	-8.2298E-07
10	-0.0467547417	-5.8803E-06	1.2577E-02	1.2354E-11	-2.6424E-08	4.5291E-10	-9.6870E-07
11	0.0182288605	-4.1825E-06	-2.2944E-02	-1.9162E-12	-1.0512E-08	-2.4871E-11	-1.3644E-07
12	0.0235721164	-1.4554E-06	-6.1742E-03	3.2766E-12	1.3900E-08	-6.3460E-12	-2.6922E-08
13	0.0302285093	-3.6853E-06	-1.2191E-02	1.1256E-11	3.7237E-08	4.2186E-12	1.3956E-08
14	0.0363683761	-1.6034E-05	-4.4088E-02	1.1913E-11	3.2757E-08	-3.0724E-12	-8.4480E-09
AoA	0.0008542099	-4.1651E-06	-4.8760E-01	-3.9028E-13	-4.5689E-08	-6.0492E-13	-7.0816E-08

the gradient for a single optimization iteration. It also compares the cost of computing an additional gradient vector associated with the lift-constraint gradient, since both gradient vectors are required for the lift-constrained drag minimizations presented in the next chapter.

All comparisons are performed on a single CPU. Thus, the FD gradient is computed

Table 5.2: Comparison of lift-constraint gradient components computed using various gradient evaluation techniques: RAE-2822, $Re = 15.7 \times 10^6$, $M = 0.6$, $AoA = 1^\circ$

Design	nAD	FD$-n$AD		SN$-n$AD		iAD$-n$AD	
Variable	Gradient	Diff.	% Diff.	Diff.	% Diff.	Diff.	% Diff.
1	16.776029011	1.6599E-04	9.8944E-04	2.5800E-11	1.5379E-10	5.1610E-10	3.0764E-09
2	6.2471822492	-8.0749E-05	-1.2926E-03	1.8670E-11	2.9886E-10	1.8097E-10	2.8968E-09
3	3.2566403756	-4.8176E-05	-1.4793E-03	2.5397E-12	7.7987E-11	6.5170E-11	2.0011E-09
4	1.7371341558	5.0544E-05	2.9096E-03	-3.0500E-12	-1.7558E-10	7.0330E-11	4.0486E-09
5	0.9604662869	-7.2617E-05	-7.5606E-03	-1.2865E-11	-1.3395E-09	-1.3654E-10	-1.4217E-08
6	0.5277657350	1.2555E-05	2.3789E-03	-2.7312E-11	-5.1750E-09	-2.5444E-10	-4.8211E-08
7	0.2255987326	-2.9693E-05	-1.3162E-02	-6.3210E-12	-2.8019E-09	-9.8158E-11	-4.3510E-08
8	-1.6960180276	-1.0497E-04	6.1893E-03	3.1013E-10	-1.8286E-08	3.4210E-09	-2.0171E-07
9	1.0110455993	-2.3590E-04	-2.3332E-02	-6.3170E-11	-6.2480E-09	-6.8601E-10	-6.7852E-08
10	1.7778487185	1.5768E-04	8.8692E-03	-1.3330E-10	-7.4978E-09	-1.3513E-09	-7.6006E-08
11	1.7834824690	3.0993E-04	1.7378E-02	3.0100E-12	1.6877E-10	6.0390E-11	3.3861E-09
12	3.1369587960	1.6100E-04	5.1325E-03	-7.8200E-12	-2.4928E-10	8.1100E-11	2.5853E-09
13	4.4576732209	-1.4952E-04	-3.3542E-03	5.2047E-13	1.1676E-11	1.2207E-10	2.7384E-09
14	6.3471594806	2.3702E-04	3.7343E-03	4.2390E-11	6.6786E-10	1.9013E-10	2.9955E-09
AoA	0.1391872901	-1.9787E-04	-1.4216E-01	5.5803E-13	4.0092E-10	9.8980E-12	7.1113E-09

Table 5.3: Time required to compute drag and lift gradients for a single optimization iteration with 15 design variables: RAE-2822, $Re = 15.7 \times 10^6$, $M = 0.6$, $AoA = 1^\circ$

	Wall-clock time (hh:mm:ss)			
Gradients	FD	SN	iAD	nAD
Drag Only	14:30:31	03:27:23	00:16:02	**00:04:56**
Drag + Lift	14:30:31	03:27:24	00:32:57	**00:05:57**

in serial and incurs the highest wall-clock time to compute the gradient vectors, requiring upwards of 14 to 15 hours. When the FD gradient is computed in parallel for this case, it requires only 29 minutes and 32 seconds, which is approximately the same cost required to compute the flow solution. However, in this case an additional 30 CPUs are required for the FD gradient. Both the FD and SN gradients scale linearly with the number of design variables; however the SN method is typically more efficient than FD. Both the FD and SN methods have the advantage that they incur virtually no additional cost in order to compute the lift-constraint gradient. For the FD method, both the lift and drag are computed at each perturbed geometry state, and for the SN method, the linear systems given by (4.12) do not need to be solved again.

The two discrete adjoint approaches show significant improvement in efficiency and their cost is relatively independent of the number of design variables. The noniterative (nAD)

method incurs significantly less cost than all other methods, requiring approximately five minutes for the objective gradient and an additional one minute for the lift-constraint gradient. As explained in Section 4.7.2, the cost of the iAD approach approximately doubles when required to compute the lift-constraint. The nAD approach does not, since it is able to reuse information stored during the objective function gradient evaluation. For these reasons, the nAD method is selected as the preferred gradient evaluation technique for free-transitional flows. In the next chapter, we turn our attention to the design of NLF airfoils using the new gradient-based ASO framework.

Chapter 6

Design Optimization Studies

To demonstrate the capabilities of the proposed optimization framework for the purpose of NLF airfoil design, single and multipoint optimizations are performed at conditions associated with subsonic and transonic commercial aircraft. In particular, we will be considering lift-constrained drag minimizations for which the general optimization problem (previously defined in Table 4.1) is specialized as follows:

$$
\begin{aligned}
\text{min} \quad & C_d\left(\mathbf{x}, \mathbf{Q}(\mathbf{x})\right) \\
\text{w.r.t.} \quad & \mathbf{x} \in \mathbb{R}^{N_{\mathrm{D}}} \\
\text{s.t.} \quad & C_l = C_l^* \\
& A_{\text{final}} \geq A_{\text{initial}} \\
& t_{0.03c} \geq 0.025c \\
& t_{0.98c} \geq 0.002c
\end{aligned}
$$

The objective is to minimize the total drag of the airfoil constrained by a user-specified lift target, C_l^*. As mentioned in Section 1.3.3, we have selected a design objective and a set of constraints that aim to reflect the industry's aerodynamic design goals. We intentionally avoid the use of indirect objectives, such as those that focus specifically on delaying transition. In turn, the optimizer can better account for any design trade-offs that may exist (for example, between wave drag and friction drag). Additional inequality constraints are included for structural considerations; minimum thickness constraints near the leading and trailing edges are enforced, as well as an area constraint that ensures the final area of the airfoil is greater than or equal to the initial area.

The initial geometry (unless otherwise specified) is the RAE-2822 airfoil parametrized by seventeen B-spline control points, as shown in Figure 5.11. The control point located at the leading edge of the airfoil as well as the two coincident control points at the trailing edge are kept fixed throughout the optimization. The y-coordinates of the remaining 14 control points are used as the geometric design variables (shaded in blue). The angle of

Table 6.1: Optimization cases

Case	Aircraft	Reynolds Number (Re)	Mach Number (M)	Lift Coefficient (C_l^*)
A	Cessna 172R	5.6×10^6	0.19	0.30
B	Dash-8 Q400	15.7×10^6	0.60	0.42
C	Boeing 737-8	20.3×10^6	0.71	0.50

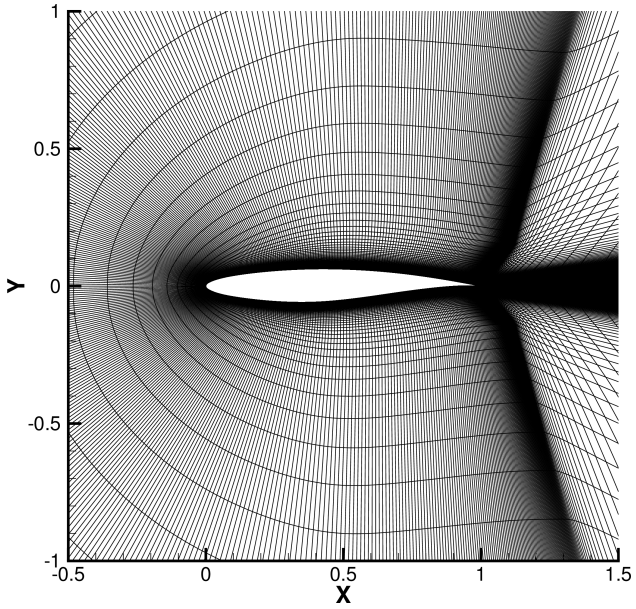

Figure 6.1: RAE-2822 C-grid with 575×224 nodes

attack is also included as an additional design variable. The computational grid consists of a 575×224 C-grid, shown in Figure 6.1, resulting from the grid convergence studies on the boundary-layer properties. All results were obtained using the scalar dissipation scheme, the compressible Bernoulli edge-finding method, the smooth intermittency function given by (3.5), and the noniterative augmented adjoint gradient computation for free-transitional flows. Furthermore, it has been observed that the AHD criterion is generally less robust and less accurate than the e^N envelope criterion. Due to the importance of having robust, accurate, and deep convergence of the free-transition flow solver (as discussed in Section 3.7) over a wide range of geometries, the e^N envelope criterion has been selected for the optimization results presented herein.

We begin by considering single-point optimizations for the subsonic and transonic cases outlined in Table 6.1. Cases A, B, and C were selected to loosely approximate the cruise flight conditions of the Cessna 172R, the Bombardier Dash-8 Q400 turboprop, and the Boeing 737-800 turbofan, respectively. Next we consider a multipoint optimization performed on

the Dash-8 Q400 (Section 6.4), followed by a Pareto front study that considers the design trade-offs between optimizing under fully-turbulent and free-transition conditions (Section 6.5). Finally, we study the effect of the freestream disturbance environment by varying the critical N-factor (Section 6.6). The study includes single-point and multipoint optimizations, with the goal of demonstrating NLF airfoil design that is more robust to the uncertainty inherent in the selection of N_{crit}.

The airfoils in this work have been investigated and optimized strictly under cruise flight conditions. More work is required to assess the low-speed, high-lift characteristics of the optimized geometries (see Section 7.2 for further discussion). The multipoint optimization presented in Section 6.4 considers the cruise segment of the Q400 flight envelope; however, the off-design performance at take-off and landing is not included. In practice, the protrusions created by high-lift devices, as well as engine pylons and flap-track fairings, are typically detrimental to maintaining laminar flow. This motivates much of the ongoing research and development of slatless and morphing wing design [10], as well as highlighting an aerodynamic incentive to rear-mounted engines (further facilitating NLF on the lower surface) [133]. For practical NLF design, in addition to including a low-speed requirement, one should also consider multipoint optimization with varying critical N-factors (for reasons made clear in Section 6.6). Nonetheless, the various design studies presented in this work aim to quantify the potential benefits of optimization for NLF employing clean wings in cruise conditions. Ultimately the following studies serve as a proof-of-concept for the design procedure in general; they demonstrate that the methodology works and is capable of being used in a variety of contexts.

6.1 Cessna 172R Skyhawk

The Cessna 172R Skyhawk is assumed to be cruising at 6000 ft, a speed of 120 knots, and a weight of 2200 lbs, corresponding to $Re = 5.6 \times 10^6$, $M = 0.19$, and $C_l^* = 0.3$. In this section we consider single-point optimization with results obtained using the e^N envelope transition criterion with $N_{\text{crit}} = 9$. In Table 6.2, a summary of the results comparing the initial and optimized airfoils is presented. Figure 6.2(a) compares the initial and optimized geometries and Figure 6.2(b) compares the pressure profiles. Note that in Figure 6.2(a), the geometries are rotated about the mid-chord position based on the AoA of the airfoil; the same is true for all such figures in this chapter. The transition locations are indicated by the solid circles. The angle of attack decreased from an initial value of 0.75° to 0.19°, the lift constraint is satisfied, and the total drag is reduced by 35.1 drag counts, or 56%. Figure 6.3 shows the convergence

Table 6.2: Case A summary of optimization results: $Re = 5.6\times10^6$, $M = 0.19$, $C_l^* = 0.3$

	C_d	C_{d_p}	C_{d_f}	C_l	C_m	$T_{\text{up}}(x/c)$	$T_{\text{lo}}(x/c)$	AoA
Initial	0.00556	0.00102	0.00454	0.3000	-0.06687	0.3349	0.5352	0.7498°
Optimized	0.00241	0.00061	0.00181	0.2999	-0.07514	0.8387	0.8669	0.1883°

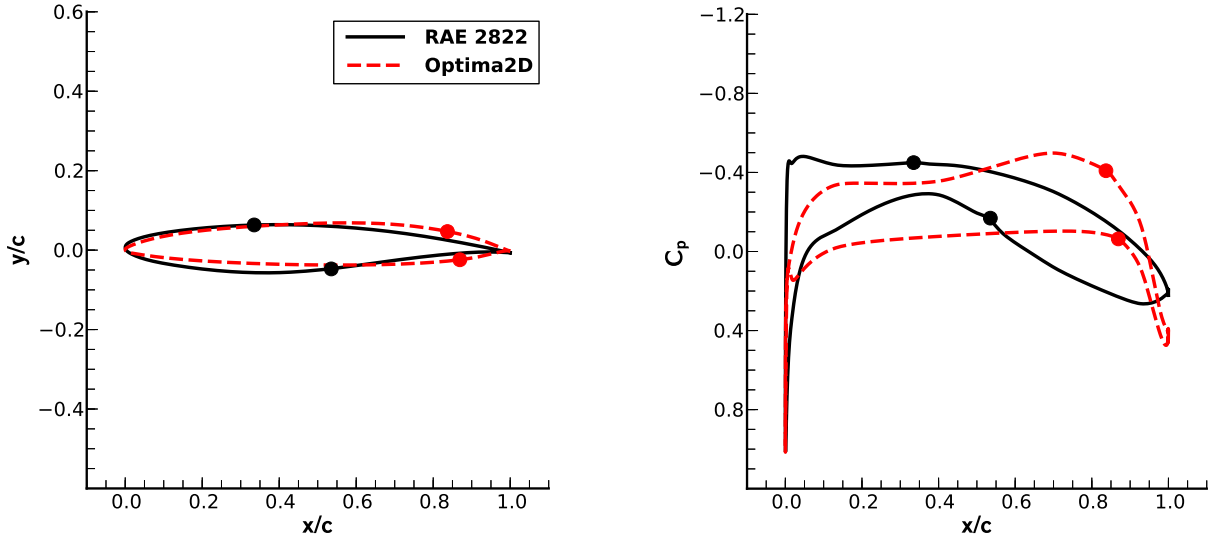

(a) Initial and optimized airfoils

(b) Initial and optimized pressure distributions

Figure 6.2: Case A optimization results: $Re = 5.6\times10^6$, $M = 0.19$, $C_l^* = 0.3$; symbols indicate transition point locations

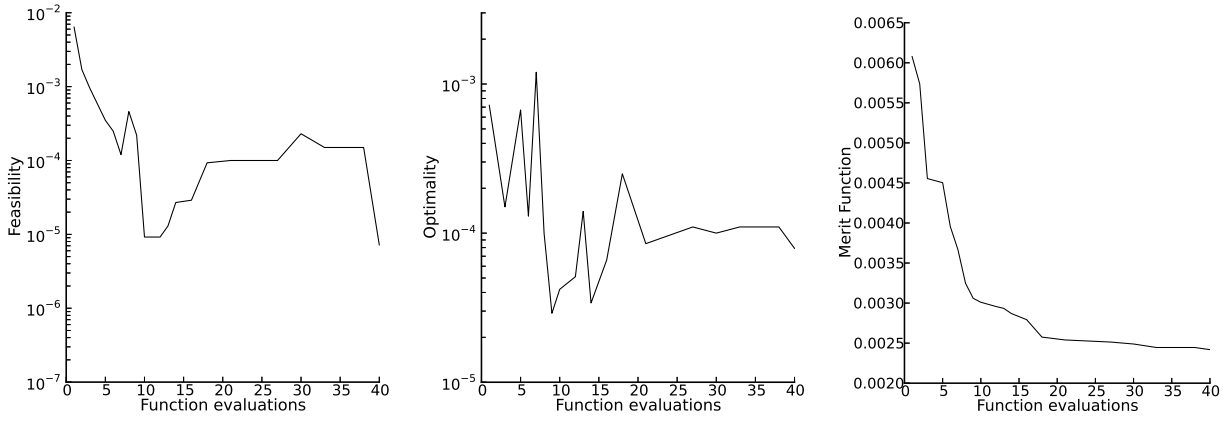

(a) Feasibility

(b) Optimality

(c) Merit function

Figure 6.3: Case A optimization convergence history: $Re = 5.6\times10^6$, $M = 0.19$, $C_l^* = 0.3$

history of the SNOPT feasibility, optimality, and merit functions. As discussed in Section 4.4 (and demonstrated for this case in Figure 6.3), we typically observe that feasibility is satisfied, optimality is reduced by approximately one order of magnitude, and the merit function levels off over a number of iterations.

The ability of the optimizer to reduce drag by exploiting laminar-turbulent transition prediction is made evident by the aft movement of the transition points from 34% to 84% chord on the upper surface and 54% to 87% chord on the lower surface. The leading edge radius has decreased and the point of maximum thickness has been pushed significantly aft. As a result, the favourable pressure gradients are extended aft, in turn delaying transition.

6.2 Bombardier Dash8-Q400

The design point for the Dash-8 Q400 is taken as point 6 from the multipoint optimization case (to be presented in Section 6.4). The Q400 is assumed to be cruising at a Mach number of 0.6, at an altitude of 23,000 ft, and a weight of 60755 lbs, corresponding to $Re = 15.7 \times 10^6$, $M = 0.6$, and $C_l^* = 0.42$. The results are obtained using the e^N envelope transition criterion with $N_{\text{crit}} = 9$. Table 6.3 provides a summary of the results comparing the initial and optimized airfoils. In this case, the angle of attack is decreased from an initial value of 1.14° to 0.66°, the lift constraint is again satisfied, and the total drag is reduced by 32.6 drag counts, or 54%. The transition point on the upper surface has been moved aft by over 50% chord, while the lower surface transition point has moved aft approximately 20% chord.

Figure 6.4(a) compares the initial and optimized geometries; Figure 6.4(b) compares the pressure profiles. It can be observed that the optimizer was again successful in designing an airfoil with an extended favourable pressure gradient on both the upper and lower surfaces. As in the previous case, the optimized geometry has a smaller leading edge radius and a point of maximum thickness that is further aft. The results demonstrate the ability of the optimizer to generate new NLF airfoils which would typically require considerable aerodynamic experience to design.

6.3 Boeing 737-800

The Boeing 737-800 has a wing sweep angle of 25°and is assumed to be cruising at 35000 ft, and a Mach number of 0.785, which corresponds to an effective Mach number of 0.71. The sectional lift coefficient is approximated as 0.5. Results are obtained using the e^N envelope transition criterion with $N_{\text{crit}} = 9$. Due to the transonic flight conditions, the optimization

Table 6.3: Case B summary of optimization results: $Re = 15.7\times10^6$, $M = 0.6$, $C_l^* = 0.42$

	C_d	C_{d_p}	C_{d_f}	C_l	C_m	$T_{\text{up}}(x/c)$	$T_{\text{lo}}(x/c)$	AoA
Initial	0.00598	0.00194	0.00405	0.4200	-0.08129	0.1480	0.4912	1.1424°
Optimized	0.00272	0.00095	0.00178	0.4199	-0.08247	0.7232	0.7685	0.6613°

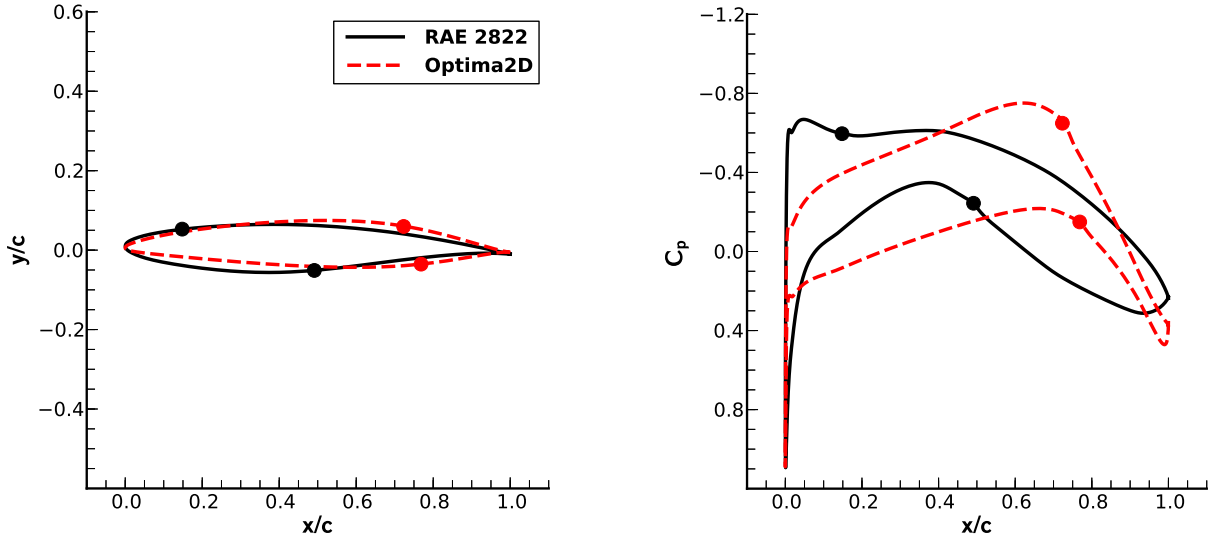

(a) Initial and optimized airfoils

(b) Initial and optimized pressure distributions

Figure 6.4: Case B optimization results; $Re = 15.7\times10^6$, $M = 0.6$, $C_l^* = 0.42$; symbols indicate transition point locations

in this case is less robust. The flow solver may fail to converge if the transition locations are moved aft of a shockwave during the transition prediction procedure, in turn causing unsteady flow separation. Modifications to the transition prediction algorithm for increased robustness have been discussed in Section 3.7.

Table 6.4 provides a summary of the results comparing the initial and optimized airfoils. The angle of attack in this case is decreased from an initial value of 1.13° to −0.25°, the lift constraint is satisfied, and the total drag is reduced by 25.8 drag counts, or 42%. The transition point on the upper surface was moved aft from 21% to 74% chord, and from 47% to 50% chord on the lower surface. Figure 6.5(a) compares the initial and optimized geometries; Figure 6.5(b) compares the pressure profiles. In this case, the optimizer is successful in designing a shock-free transonic NLF airfoil, in turn, significantly reducing the total drag. While the optimizer is able to delay transition significantly on the upper surface, the lower surface transition location remains near the mid-chord position. With the favourable pressure

Table 6.4: Case C summary of optimization results: $Re = 20.3\times10^6$, $M = 0.71$, $C_l^* = 0.5$

	C_d	C_{d_p}	C_{d_f}	C_l	C_m	$T_{\text{up}}(x/c)$	$T_{\text{lo}}(x/c)$	AoA
Initial	0.00617	0.00259	0.00358	0.4995	-0.09427	0.2090	0.4740	1.1266°
Optimized	0.00359	0.00146	0.00214	0.5014	-0.16261	0.7441	0.5006	-0.2487°

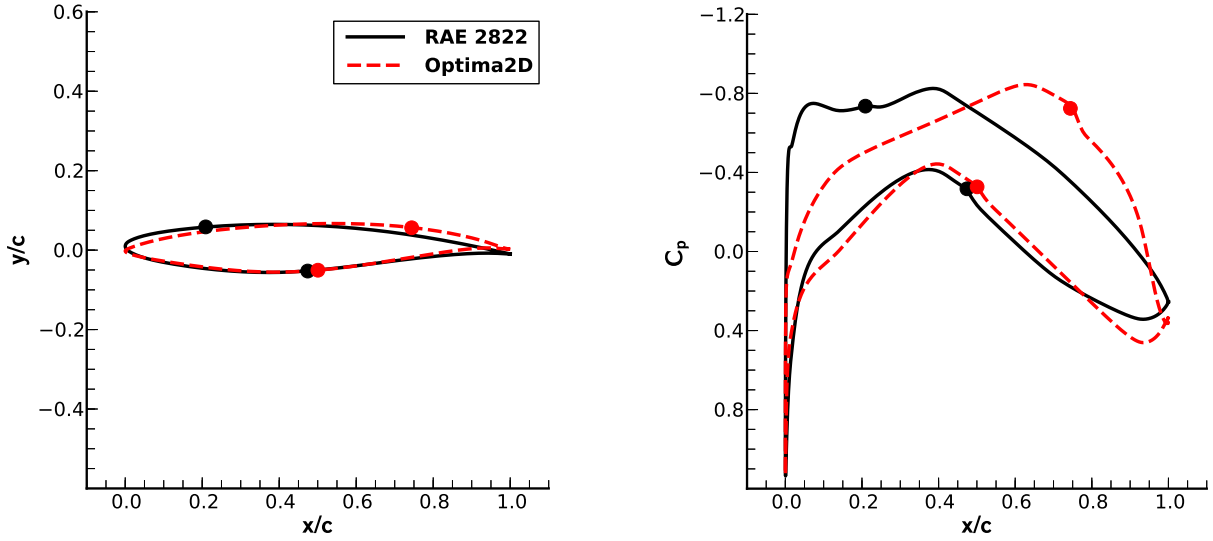

(a) Initial and optimized airfoils

(b) Initial and optimized pressure distributions

Figure 6.5: Case C optimization results; $Re = 20.3\times10^6$, $M = 0.71$, $C_l^* = 0.5$; symbols indicate transition point locations

gradient on the top surface extended further aft than that of the lower surface, the airfoil is more heavily aft-loaded. The aft-loading, in turn, results in a higher negative pitching moment and, as discussed in Section 7.2, future work will consider the addition of a pitching moment constraint for transonic NLF applications.

6.4 Multipoint Optimization

Here we consider a multipoint optimization at a range of cruise conditions associated with the Dash-8 Q400 aircraft. A nine-point stencil, presented in Table 6.5, is defined by varying the aircraft weight and Mach number. This is done to reduce the sensitivity of the final optimized shape to any perturbations in the flight conditions encountered during cruise and to enable efficient operation within this envelope. The aircraft is assumed to have a take-off weight equal to the Q400's maximum take-off weight of 64500 lbs. Given a typical payload,

Table 6.5: Design points and weighting for multipoint optimization

Design Pt.	Quadrature Weight ($\mathcal{T}$)	Aircraft Weight (W) [lbs]	Mach No. (M)	Reynolds No. (Re)	Lift Coefficient (C_l^*)
1	1	63757	0.48	12.5×10^6	0.68
2	2	63757	0.54	14.1×10^6	0.54
3	1	63757	0.60	15.7×10^6	0.44
4	2	60754	0.48	12.5×10^6	0.65
5	4	60754	0.54	14.1×10^6	0.51
6	2	60754	0.60	15.7×10^6	0.42
7	1	57751	0.48	12.5×10^6	0.62
8	2	57751	0.54	14.1×10^6	0.49
9	1	57751	0.60	15.7×10^6	0.40

the usable fuel on board (at take-off) is approximated to be 7500 lbs. The three aircraft weights considered in the multipoint stencil are calculated from a 10%, 50% and 90% fuel burn, which loosely approximate the beginning, middle, and end of cruise. The three Mach numbers considered are 0.6, 0.54, and 0.48, which roughly correspond to the high-speed, intermediate, and long-range design speeds of the Q400, respectively. Given the range of weights and Mach numbers, and assuming a constant cruising altitude of 23000 ft, we can then compute the corresponding range of Reynolds numbers and lift constraints presented in Table 6.5. The results are obtained using the e^N envelope transition criterion with $N_{\text{crit}}=9$.

Recall from Section 4.8 that the design objective given by (4.39) is an approximation to the weighted integral given by (4.38). Although any design priority weighting may be selected as desired, here we make the assumption that all design points are of equal importance, that is, $\mathcal{Z}(W_i, M_j)=1$ for all i and j. Table 6.5 also presents the quadrature weights $\mathcal{T}$ used to approximate (4.38) using the trapezoidal quadrature rule.

Table 6.6 provides a summary of the results comparing the initial and optimized airfoils, along with the various angles of attack. The lift constraint has been satisfied and the drag reduced at each operating point. The flight conditions and lift constraint of design point 6 correspond to the Q400 single-point optimization presented in Section 6.2. Figure 6.6(a) compares the single-point and multipoint optimized geometries for design point 6, and Fig 6.6(b) compares the pressure distributions. Comparing the optimized designs, it is clear that the single and multipoint results differ. The multipoint optimization has transition points that are approximately 5% further upstream when compared to the single-point optimization of Case B. Furthermore, while the total drag was reduced by 55% in the single-

Table 6.6: Summary of multipoint optimization results

Design Pt.		C_d	C_{d_p}	C_{d_f}	C_l	C_m	$T_{\text{up}}(x/c)$	$T_{\text{lo}}(x/c)$	AoA
1	Initial	0.00803	0.00345	0.00458	0.6795	-0.07315	0.0172	0.5393	3.5050°
	Optimized	0.00453	0.00211	0.00241	0.6799	-0.08278	0.5509	0.7767	2.5651°
2	Initial	0.00692	0.00253	0.00439	0.5400	-0.07691	0.0620	0.5134	2.2177°
	Optimized	0.00328	0.00128	0.00200	0.5399	-0.08434	0.6656	0.7583	1.2498°
3	Initial	0.00613	0.00205	0.00408	0.4400	-0.08123	0.1328	0.4932	1.2852°
	Optimized	0.00304	0.00105	0.00199	0.4400	-0.08730	0.6664	0.7297	0.3184°
4	Initial	0.00779	0.00323	0.00456	0.6494	-0.07328	0.0199	0.5356	3.2717°
	Optimized	0.00393	0.00174	0.00219	0.6502	-0.08252	0.6194	0.7741	2.3059°
5	Initial	0.00670	0.00234	0.00436	0.5100	-0.07704	0.0763	0.5103	1.9920°
	Optimized	0.00318	0.00117	0.00200	0.5100	-0.08347	0.6685	0.7558	1.0383°
6	Initial	0.00598	0.00194	0.00405	0.4201	-0.08129	0.1479	0.4912	1.1430°
	Optimized	0.00300	0.00100	0.00200	0.4202	-0.08662	0.6683	0.7226	0.1875°
7	Initial	0.00761	0.00298	0.00463	0.6200	-0.07355	0.0232	0.5321	3.0357°
	Optimized	0.00368	0.00155	0.00213	0.6201	-0.08191	0.6421	0.7717	2.0713°
8	Initial	0.00656	0.00222	0.00434	0.4900	-0.07710	0.0860	0.5082	1.8419°
	Optimized	0.00311	0.00111	0.00200	0.4899	-0.08287	0.6705	0.7543	0.8963°
9	Initial	0.00584	0.00183	0.00401	0.4000	-0.08133	0.1634	0.4889	1.0000°
	Optimized	0.00301	0.00095	0.00206	0.4001	-0.08590	0.6701	0.6977	0.0527°

point optimization, it was reduced by 50% in the multipoint optimization. This illustrates that the added robustness in the design (now optimized over a range of conditions) incurs a penalty relative to the single-point optimization. It also exemplifies the importance of the designer's role in carefully selecting and weighting the design points appropriately. For example, if the Q400 normally cruises at the high-speed Mach number of 0.60, then the designer might choose to place more importance on those operating points.

6.5 Pareto Front Design Study

A Pareto front can provide useful insight into the trade-offs involved in the design of NLF airfoils. While the multipoint optimization case (presented in Section 6.4) is useful for designing an airfoil that performs well over a range of cruise flight conditions, here we consider the aerodynamic performance of NLF airfoils when or if transition occurs inadvertently at the leading edge of the airfoil.

Following the work of Driver and Zingg [30], a Pareto front may be formed by minimizing

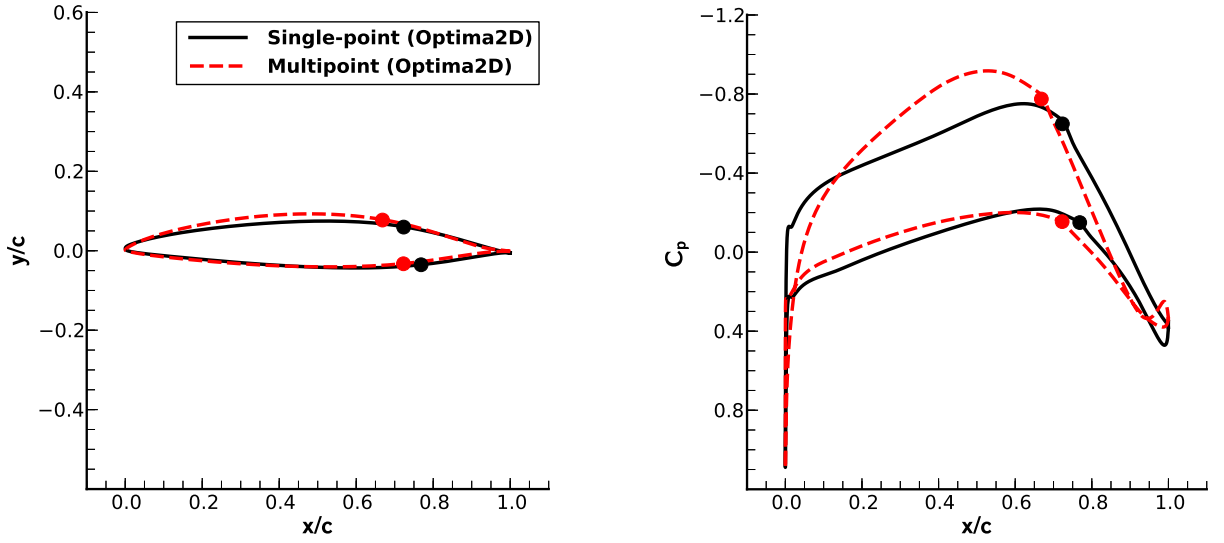

(a) Initial and optimized airfoils

(b) Point 6: Initial and optimized pressure distributions

Figure 6.6: Comparison of single-point and multipoint optimization results (design point 6)

a weighted sum objective, $\mathcal{J}$, defined as

$$\mathcal{J} = \omega_{\text{ft}}\mathcal{J}_{\text{ft}} + (1-\omega_{\text{ft}})\mathcal{J}_{\text{lt}} \,, \tag{6.1}$$

where $\mathcal{J}_{\text{ft}}$ and $\mathcal{J}_{\text{lt}}$ represent the drag coefficients under fully-turbulent and laminar-turbulent (i.e. free-transition) conditions, respectively. Each point on the Pareto front represents a two-point design problem in which we minimize $\mathcal{J}$ for a given weighting factor, ω_{ft}, where $0 \leq \omega_{\text{ft}} \leq 1$. The calculation of the two operating conditions ($\mathcal{J}_{\text{ft}}$ and $\mathcal{J}_{\text{lt}}$) in turn requires two flow solutions, each at their respective angle of attack. Furthermore, both operating conditions are constrained to meet the same lift target (set to $C_l^* = 0.42$ for Case B) to ensure sufficient lift generation at both conditions for every optimal point.

The computed Pareto front is shown in Figure 6.7 and clearly captures the advantages of favouring one operating condition over the other. As expected, the drag count values under laminar-turbulent conditions are significantly less than the fully-turbulent operating conditions. The Pareto front demonstrates that when an airfoil designed strictly for laminar-turbulent conditions (that is, $\omega_{\text{ft}} = 0$) is operating under fully-turbulent conditions, it has a drag count of approximately 84, as compared to 79 counts for an airfoil designed and operated under fully-turbulent conditions ($\omega_{\text{ft}} = 1$); a relative drag penalty of approximately 6%. On the other hand, when an airfoil designed strictly for fully-turbulent conditions is operating

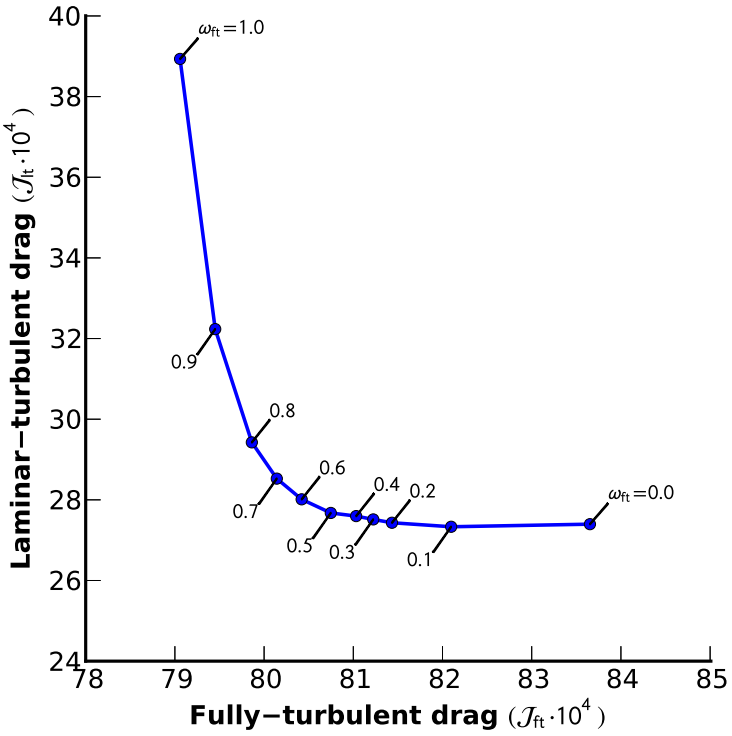

Figure 6.7: Pareto front study, Case B: $Re = 15.7\times10^6$, $M = 0.6$, $C_l^* = 0.42$

under laminar-turbulent conditions it has a drag count of approximately 39, as compared to 27 counts for an airfoil designed and operated under laminar-turbulent conditions; a relative drag penalty of approximately 44%. The remaining points on the Pareto front allow the designer to select an appropriate optimal geometry depending on their needs and conservatism. Airfoils optimized using ω_{ft} values in the range of $0.2 \leq \omega_{\text{ft}} \leq 0.7$ represent a good compromise in performance between the two operating conditions.

6.6 Robustness to Freestream Disturbance Environment

The design of NLF airfoils should be robust to a variation in the freestream disturbance environment. It is not clear that an airfoil optimized under one set of conditions, reflected in the critical N-factor, will perform well under other conditions, i.e. with a different critical N-factor. Recall that the critical N-factor selected for the application of the e^N criterion is specified by the designer and typically has some uncertainty. The uncertainty stems from discrepancies between wind-tunnel and in-flight freestream turbulence intensities (T_u), as well as many other factors that can affect the transition process, such as acoustic disturbances, vibrations, and surface roughness (see Section 1.1.2). Indeed, the process of accurately

correlating analytical transition criteria to experimental results is a difficult challenge. In this section, we first examine the sensitivity to the critical N-factor and then propose a technique to enable the design of airfoils with robust performance over a range of critical N-factors.

6.6.1 Sensitivity to the Critical N-factor

To examine the sensitivity to the critical N-factor, we repeat the single-point optimization performed in Section 6.2 by specifying different values for N_{crit}. Recall that the previously presented single-point optimization of the Q400 (Section 6.2) assumed a critical N-factor of 9. In the first case, we decrease the critical N-factor to a value of 7, which implies that transition will occur further upstream for the same geometry and flight conditions (see Figure 3.2). In the second scenario, we increase the critical N-factor to a value of 11, which implies the opposite trend. For critical N-factors of 7, 9, and 11 the freestream turbulence intensities approximated using Mack's relationship (1.5) are 0.16%, 0.07%, and 0.03%, respectively. However, it is important to note that Mack's relationship is only valid for incompressible flows and a specific range of Tu, thus, these values are not necessarily representative.

The results of the optimizations are presented in Table 6.7. The labelling is such that SP7 represents the optimized geometry from the single-point optimization performed at $N_{\text{crit}} = 7$. The off-design performance (at different N_{crit} values) and multipoint results will be discussed in the coming sections. For the single-point cases, when operating on-design, it can be seen that the lift target is met and the optimal drag coefficients decrease with increasing N_{crit}; the difference in drag between $N_{\text{crit}} = 7$ and $N_{\text{crit}} = 11$ is approximately 1.2 counts. It can also be observed that the optimal geometries have transition points on the upper surface that are further aft as we increase N_{crit} (as expected). For the $N_{\text{crit}} = 11$ case, the optimizer takes a slightly lower AoA, resulting in an optimal geometry with an upper surface transition point that is 5-7% chord further downstream and a lower surface transition point that is approximately 1-3% further upstream, as compared to the other two cases.

A comparison of the optimized geometries and pressure profiles is presented in Figure 6.8. In an attempt to maintain laminar flow, the optimizer produces steeper favourable pressure gradients as the critical N-factor decreases. Qualitatively, the optimized geometries and $AoAs$ are very similar; however, the variation in N_{crit} produces important changes to the geometry that alter the point at which pressure recovery begins and transition occurs.

Off-Design Performance of Single-point Optimizations

Table 6.6.2 also presents the results of a post-optimality study on the off-design performance of the three single-point cases discussed above; here, off-design refers to the uncertainty in

Table 6.7: Summary of single-point optimizations, off-design performance, and multipoint optimization over a range of critical N-factor values

Geometry	N_{crit}	C_d	C_{d_p}	C_{d_f}	C_l	C_m	$T_{\text{up}}(x/c)$	$T_{\text{lo}}(x/c)$	AoA
Initial	7	0.00612	0.00198	0.00414	0.4200	-0.08101	0.1206	0.4816	1.1531°
SP7	7	0.00276	0.00092	0.00184	0.4200	-0.08480	0.7124	0.7439	0.5953°
SP7	9	— did not converge due to flow separation —							
SP7	11	— did not converge due to flow separation —							
Initial	9	0.00598	0.00194	0.00405	0.4200	-0.08129	0.1480	0.4912	1.1424°
SP9	9	0.00272	0.00095	0.00178	0.4199	-0.08247	0.7232	0.7685	0.6613°
SP9	7	0.00281	0.00095	0.00186	0.4185	-0.08216	0.7111	0.7546	0.6613°
SP9	11	— did not converge due to flow separation —							
Initial	11	0.00583	0.00189	0.00394	0.4200	-0.08159	0.1783	0.4985	1.1300°
SP11	11	0.00264	0.00095	0.00168	0.4200	-0.12140	0.7780	0.7405	-0.0422°
SP11	7	0.00388	0.00122	0.00265	0.4062	-0.11163	0.7441	0.3956	-0.0422°
SP11	9	0.00272	0.00095	0.00178	0.4174	-0.12077	0.7661	0.7275	-0.0422°
MP	7	0.00297	0.00091	0.00207	0.4203	-0.08559	0.6762	0.7045	0.5707°
MP	9	0.00286	0.00089	0.00197	0.4204	-0.08581	0.6927	0.7194	0.5631°
MP	11	0.00281	0.00090	0.00191	0.4203	-0.08591	0.7010	0.7280	0.5594°

*SP7 = geometry from single-point optimization at $N_{\text{crit}} = 7$
*SP9 = geometry from single-point optimization at $N_{\text{crit}} = 9$
*SP11 = geometry from single-point optimization at $N_{\text{crit}} = 11$
*MP = geometry from multipoint optimization (Section 6.6.2)

the freestream disturbance environment. Flow solutions are computed for each combination of N_{crit} with each optimized airfoil and AoA. Typically, one would expect that optimizing using a lower critical N-factor is assumed to be more conservative and robust to factors that may prematurely trip the flow. However, optimizations performed at too low a value of N_{crit} can result in an airfoil that is not robust to the opposite scenario; that is, one in which a lower freestream turbulence intensity is encountered. For example, if we take the SP7 airfoil and AoA (optimized with $N_{\text{crit}} = 7$) and attempt a flow solution at $N_{\text{crit}} = 9$ or 11, the laminar boundary layer does not transition soon enough to avoid flow separation. The same behaviour is observed when we take the SN9 airfoil and attempt a flow solution at $N_{\text{crit}} = 11$. As indicated in Table 6.6.2, for these cases our steady flow solver is not able to converge due to the unsteadiness of the separated flow. Needless to say, this type of flow separation should be avoided in NLF design due to the significantly increased pressure drag that results.

If, on the other hand, we take the airfoils designed at the higher values of N_{crit} and perform flow analyses at lower values of N_{crit}, then flow separation is not an issue. Depending on the optimized pressure profiles, the off-design performance degradation when operating

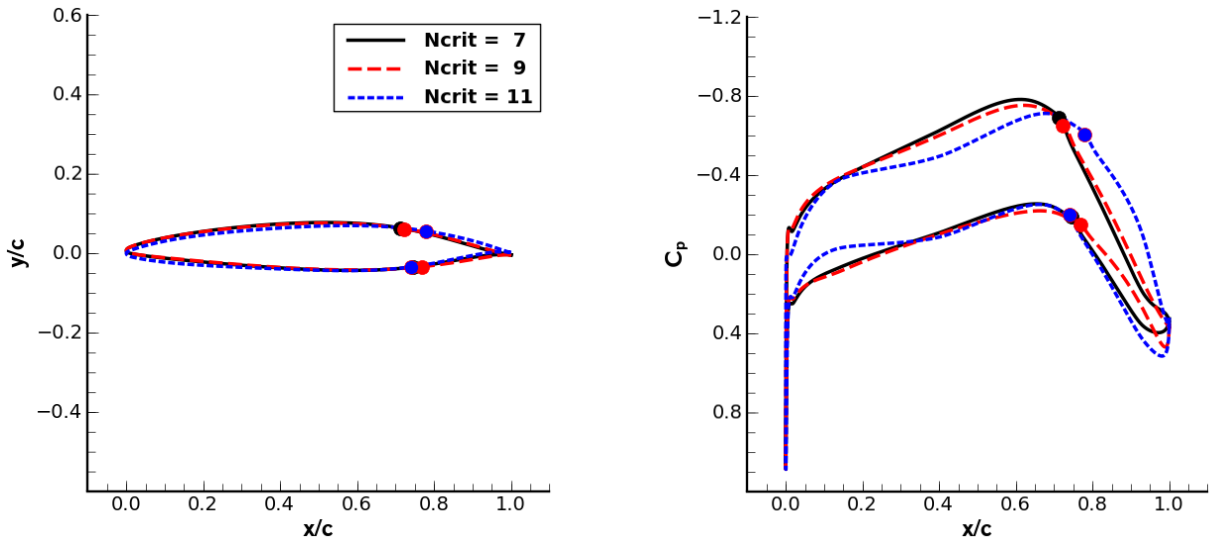

(a) Optimized airfoil geometries

(b) Optimized pressure distributions

Figure 6.8: Comparison of single-point optimizations performed at different critical N-factors values

at a lower than expected N_{crit} may or may not be significant. For example, when the SP9 airfoil is analyzed at $N_{\text{crit}}=7$, the transition points move slightly further upstream, but the airfoil continues to perform well. The same is true for the SP11 airfoil when analyzed at $N_{\text{crit}}=9$. However, when the SP11 airfoil is analyzed at $N_{\text{crit}}=7$, a noticeable drag increase of 12.4 counts results due to the lower surface transition point being much further upstream: 39.6% chord for $N_{\text{crit}}=7$ as compared to the on-design transition location of 74.4% chord. The reason for this can be observed in the lower surface pressure profile, which contains a dip just upstream of the mid-chord region. This dip presents a short region of mild adverse pressure gradient that is sufficient to trigger transition at the lower critical N-factor but not at the higher values. In general, these results imply that the single-point optimizations (while qualitatively similar) are not particularly robust to the uncertainty in the specification of N_{crit}. In the next section we propose an NLF design strategy that aims to reduce this sensitivity.

6.6.2 Robust Design Over a Range of Critical N-factors

To address the concerns raised in the previous section, we propose the use of multipoint optimization over a range of critical N-factors. The objective is to reduce the sensitivity of the final NLF airfoil to the uncertainty in the freestream disturbance environment, and hence the transition locations (*viz.* the specified N_{crit} values). The same weighted integral approach

presented in Section 4.8 is employed, with the objective defined as

$$\mathcal{J} = \int_{N_1}^{N_2} C_d\,(N)\,\mathcal{Z}\,(N)\,dN \quad \longrightarrow \quad \mathcal{J} = \sum_{i=1}^{p} \mathcal{T}_i\,C_d\,(N_i)\,\mathcal{Z}\,(N_i)\,\Delta N\,, \tag{6.2}$$

where N is used to denote N_{crit} (for clarity), and p is the number of operating points, each with a different N.

As a demonstration, we shall consider a combination of the three single-point cases ($p=3$) presented in the previous section. The objective is evaluated using the Q400 flight conditions with $N=\{7,9,11\}$, and the corresponding trapezoidal quadrature weighting of $\mathcal{T}=\{\frac{1}{2},1,\frac{1}{2}\}$. All points are assumed to be equally important and given a design-point weighting of $\mathcal{Z}(N_i)=1$. Note that the designer is free to adjust the design-point weighting function as desired.

At the bottom of Table 6.7, a summary of the multipoint optimization is presented (labelled MP). It can be observed that the optimizer is able to design an airfoil that minimizes the weighted drag objective over all three N_{crit} conditions, while meeting the lift-constraint and keeping the flow attached. The penalty in designing a multipoint airfoil over the single-point airfoils is also quantified for each value of N_{crit}; the $N_{\text{crit}}=7$ case incurs the largest drag penalty of 2.1 counts. However, recall from the previous section that when the SP7 and SP9 geometries were run at the higher N_{crit} values, the laminar boundary layer was observed to separate (and the flow solver failed to converge).

Figure 6.9(a) compares the optimal geometry from the multipoint optimization to the optimal geometry of the $N_{\text{crit}}=9$ single-point case. Figures 6.9(b)–6.9(d) compare the single-point and multipoint pressure profiles for each value of N_{crit}. The geometry and pressure profiles that result from the multipoint optimization are such that the transition locations do not vary significantly with a change in N_{crit}. For the multipoint case, the optimizer produces a geometry and angle of attack for which (i) the point of pressure recovery occurs further upstream, (ii) the initial pressure recovery occurs more gradually (particularly on the lower surface), in turn reducing the steepness of the adverse pressure gradient, and (iii) there are no dips present in the pressure profiles (as observed in the $N_{\text{crit}}=11$ case), which could prematurely transition the flow. These differences, while subtle, help to design an NLF airfoil that avoids flow separation and performs well at all three values of N_{crit}. Indeed, the more gradual pressure recovery is strikingly similar to the so-called *transition ramps* used by experienced designers as a means of controlling how and where transition occurs [32, 38, 107]. These results are very promising and demonstrate the feasibility and importance

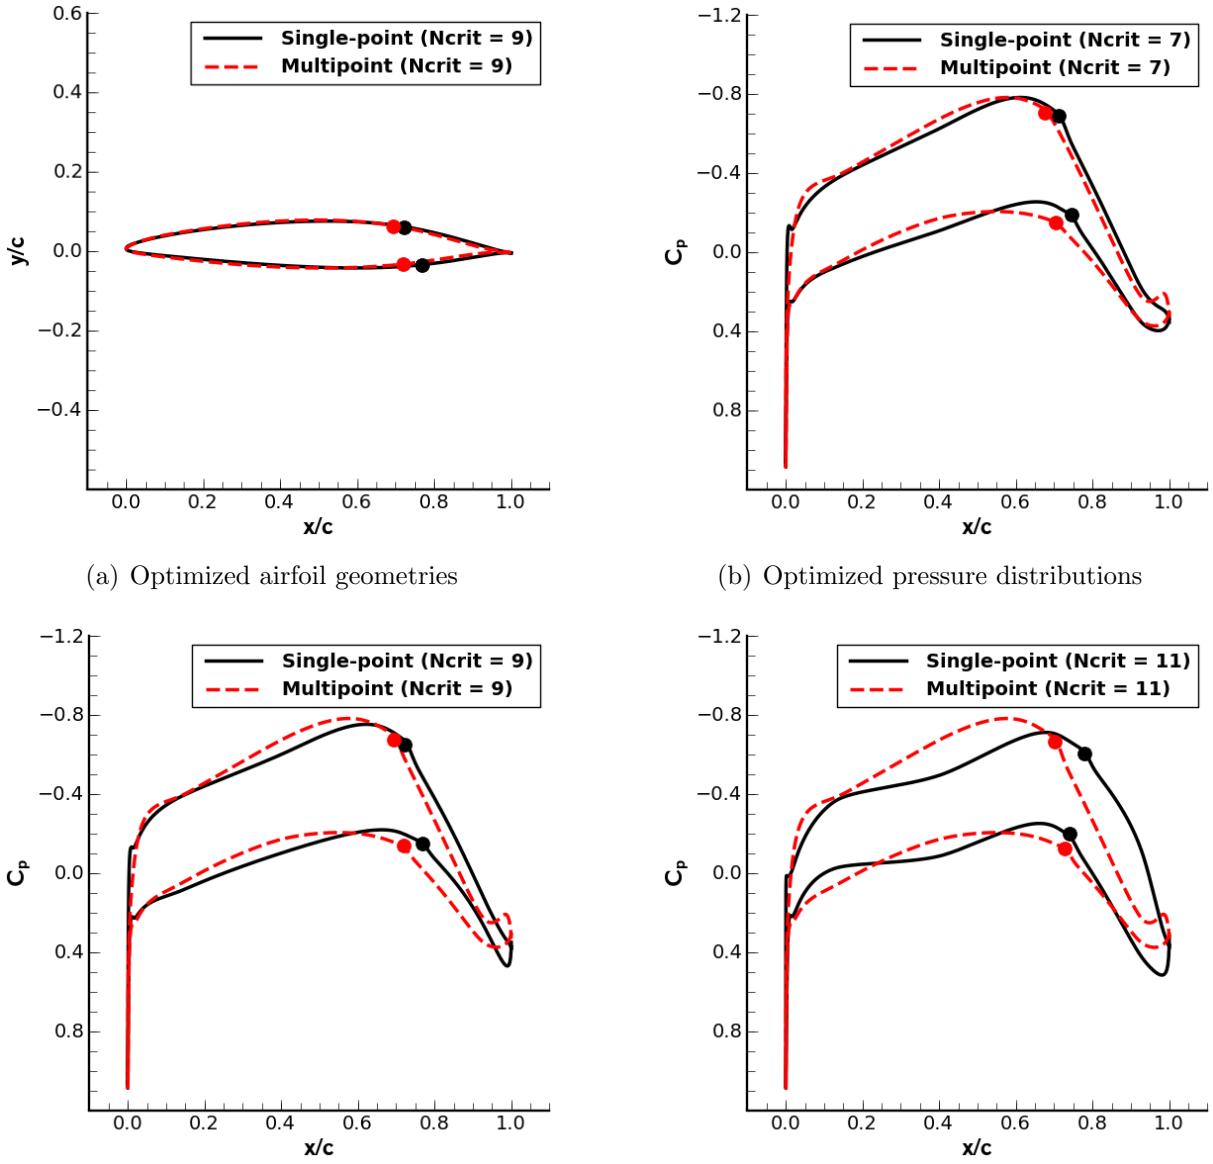

(a) Optimized airfoil geometries

(b) Optimized pressure distributions

(c) Optimized pressure distributions

(d) Optimized pressure distributions

Figure 6.9: Comparison of single-point and multipoint optimizations with different critical N-factor values

of the multipoint N-factor approach in designing NLF airfoils that are more robust to the uncertainty in the freestream disturbance environment.

Chapter 7

Contributions, Conclusions and Recommendations

7.1 Contributions and Conclusions

This research has advanced the state-of-the-art in high-fidelity aerodynamic shape optimization incorporating laminar-turbulent transition prediction by meeting the three research objectives outlined in Section 1.4 (restated below). In so doing, this research has yielded a novel optimization framework, suggested new ways of incorporating non-local transition criteria into the high-fidelity design process, and provided valuable insight into the design of NLF airfoils in the subsonic and transonic flight regime. The following presents a summary of the author's contributions and conclusions for each of the three thesis objectives:

1. Extension of a two-dimensional RANS solver making use of the Spalart-Allmaras turbulence model to incorporate a laminar-turbulent transition prediction methodology.

    - Investigation into transition region modelling has led to a new definition of an intermittency function that provides a sufficiently smooth ramp-up of the eddy-viscosity and avoids a non-smooth design space. It was observed that the SA trip term produces a sharp increase in eddy-viscosity and results in a noisy design space that is not well-suited for gradient-based optimization.
    - Implementation and investigation of several boundary-layer edge-finding methods has verified that their accuracies are in line with results from a boundary-layer code.
    - An assessment of the boundary-layer properties computed directly from the Navier-Stokes equations has demonstrated that sufficient accuracy is achievable with reasonable grid densities capable of resolving the boundary-layer profiles.
    - The compressible AHD criterion and Drela's simplified e^N envelope method have been implemented, verified, and validated in the RANS solver. The e^N method

has been selected for reasons of accuracy and robustness.

- An automatic transition prediction procedure was implemented into the RANS solver. The manner in which the transition points are updated along with the convergence criterion of the iterative procedure have been investigated and optimized to achieve consistently robust, efficient, and precise flow solves, resulting in accurate adjoint gradients and robust optimizations.

2. Development of a discrete-adjoint formulation for non-local correlation-based transition prediction.

   - An augmented adjoint system has been developed to allow for the inclusion of non-local transition criteria without destroying the sparse matrix pattern of the flow Jacobian.
   - The system has been accurately formed in an efficient and flexible manner, exploiting the complex-step approximation and the zone of influence.
   - Two strategies have been developed for solving the augmented system: an iterative solution strategy and a noniterative solution strategy. The effects of various parameters in the iterative solution procedure, such as the initial guess, under-relaxation parameter and convergence criterion have been studied.
   - The computational costs of the various gradient evaluation techniques have been quantified and compared. As a result of the demonstrated efficiency gains observed when employing the noniterative adjoint solution strategy, this approach has been selected for gradient-based optimizations.
   - An accuracy assessment of the adjoint gradient computations has demonstrated excellent agreement when compared to both the finite-difference approximation and flow sensitivity gradients, in turn, verifying the implementation and solution procedures of the proposed augmented adjoint system.

3. Application of the ASO framework to subsonic and transonic case studies, including the Cessna 172R, Bombardier Q-400, and Boeing 737-8.

   - A range of design studies presented have demonstrated the validity of the new aerodynamic shape optimization framework and the various ways in which it can be used to study and design NLF airfoils.

- The ability of the optimizer to exploit the transition point information in order to significantly reduce drag while satisfying all design constraints is demonstrated through single and multipoint optimizations.
- NLF airfoils have been designed for all applications and typically include a reduced leading-edge radius, with the point of maximum thickness shifted downstream, allowing for extended regions of favourable pressure gradients and delayed transition. Further work is required to study and optimize the designs for low-speed requirements.
- A multipoint optimization procedure allows for NLF design over a range of cruise operating conditions, and it has been demonstrated that this is achievable without the loss of laminar flow by considering a wide range of Mach numbers and aircraft weights over the cruise flight envelope.
- A Pareto front generated using the new design framework provides the designer with a range of optimized shapes depending on the conservatism required for the design of the given application. As compared to an airfoil optimized under fully-turbulent conditions, it was demonstrated that airfoils optimized for NLF do not have significant performance degradation at cruise conditions when operating under the fully-turbulent off-design condition.
- The sensitivity and robustness to the freestream disturbance environment was also studied. Upon analysis of the single-point optimization results, it was found that the optimal geometries and *AoAs* are not particularly robust to the inherent uncertainty in the specification of the critical N-factor. To address this concern, the use of multipoint optimization with varying N_{crit} values was proposed and demonstrated. The results show that robust NLF airfoils can be designed to perform well over a range of N_{crit} values, in turn demonstrating the feasibility and importance of the multipoint N-factor approach.

7.2 Recommendations for Future Work

High-lift and low-speed characteristics: In designing airfoils capable of generating sufficient lift at low-speeds (typical of take-off and landing scenarios), the multipoint optimization framework can be extended to incorporate the associated off-design lift constraints. This would allow the optimizer to explore the design trade-offs involved with the smaller leading-edge radii typical of the NLF airfoils designed in this work. It

should be noted, however, that at high angles of attack the potential for unsteady flow separation increases. This may require some modification to improve the robustness of the free-transition, steady RANS solver employed in this work. For example, it may be prudent to choose an initial guess that is further upstream for the transition point on the upper surface.

Pitching moment constraints for transonic applications: Future work should also consider the application of a pitching moment constraint to the design of transonic NLF airfoils. It was found that the optimized 737 airfoil was significantly aft-loaded, which is typical of supercritical airfoils. The moment constraint would present a compromise between the aft-loading, the leading-edge radius, and the point at which pressure recovery begins on both the upper and lower surfaces, all the while attempting to maintain a shock-free and favourable pressure distribution.

Matrix dissipation: It has been demonstrated that a matrix dissipation model requires far fewer nodes in the off-wall direction to accurately compute the laminar boundary-layer properties, as compared to the more dissipative scalar dissipation scheme. The matrix dissipation model should be differentiated in the $\mathcal{A}_2$-Jacobian matrix for accurate adjoint-based gradient computations. The robustness of the flow solver using matrix dissipation for free-transitional flows should be further evaluated, as reducing dissipation typically reduces robustness.

Natural transition prediction in three-dimensional flows: The present framework may be extended to three-dimensional flows for ASO of NLF wings. A set of transition points along the span of each wing surface would form the transition lines. Following the work of Krumbein *et al.* [60, 64, 65], the non-local boundary layer properties and Tollmien-Schlichting transition criteria used in this work can be parallelized, interpolated, and applied following the external streamlines of the flow.

Crossflow transition prediction in three-dimensional flows: The extensions to three-dimensional flows should also include the evaluation of a crossflow transition criterion, particularly for aircraft with swept wings. For example, the C_1 crossflow criterion has been successfully coupled to the e^N and AHD transition criteria for three-dimensional flows [64, 114].

ASO for NLF in three-dimensions: The proposed augmented adjoint formulation for transition prediction may also be extended to three-dimensional flows. The number of additional rows and columns in the augmented adjoint formulation would increase

such that it corresponds to the total number of transition points that have been selected to form the transition lines. For both conventional and unconventional aircraft configurations, interesting design trade-offs between wave drag and viscous drag could then be explored, particularly with the inclusion of a crossflow transition criterion, in order to better understand the potential benefits of NLF for transonic flight. In particular, there is potential to design shock-free wings at lower sweep angles which, in turn, may stabilize the crossflow instabilities.

Morphing wings: High-lift devices such as slats and flaps create joints and protrusions along the wing surfaces that can be detrimental to achieving laminar flow. ASO frameworks for NLF may be of great benefit to future designers of morphing wing technology, as it has the potential to help determine the optimal NLF shapes required to significantly minimize drag while achieving the necessary lift over different regions of the flight envelope.

Active and hybrid laminar flow control: Active flow control techniques such as boundary-layer suction and piezoelectric actuators are being investigated to achieve extended regions of laminar flow in flight. These techniques may be modelled in combination with passive flow control techniques (such as those investigated in this work) to design hybrid laminar flow technologies. Indeed, given all of the factors that affect transition, hybrid flow control may be critical to reliably realizing laminar flow in flight.

References

[1] R. Abeyratne. Aviation and climate change: In search of a global market based measure. *SpringerBriefs in Law*, 2014.

[2] R. Agarwal. Environmentally responsible air and ground transportation. In *49th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, number AIAA–2011-965, Orlando, Florida, Jan 2011.

[3] R. K. Agarwal. *Recent Advances in Aircraft Technology*. InTech, Feb 2012.

[4] Air Transport Action Group. Towards sustainable aviation. 2012. Signed on Thursday 22 March 2012, at Geneva, Switzerland.

[5] O. Amoignon, M. Berggren, A. Hanifi, D. Henningson, and J. Pralits. Shape optimization for delay of laminar-turbulent transition. *AIAA Journal*, 44:AIAA–1009–1024, May 2006.

[6] J. D. Anderson. *A History of Aerodynamics*. Cambridge University Press, 1997.

[7] A. Aranake, V. Lakshminarayan, and K. Duraisamy. Assessment of transition model and CFD methodology for wind turbine flows. In *42nd AIAA Fluid Dynamics Conference and Exhibit*, number AIAA–2012–2720, New Orleans, Louisiana, June 2012.

[8] D. Arnal, G. Casalis, and R. Houdeville. Practical transition prediciton methods: Subsonic and transonic flows. In *VKI Lecture Series: Advances in Laminar-Turbulent Transition Modeling*, January 2009.

[9] D. Arnal, R. Houdeville, A. Seraudie, and O. Vermeersh. Overview of laminar- turbulent transition investigations at ONERA toulouse. In *41st AIAA Fluid Dynamics Conference and Exhibit*, number AIAA–2011–3074, Honolulu, Hawaii, June 2011.

[10] S. Barbarino, O. Bilgen, R. M. Ajaj, M. I. Friswell, and D. J. Inman. A review of morphing aircraft. *Journal of Intelligent Material Systems and Structures*, 22(9):823–877, Jun 2011.

[11] A. D. Belegundu and T. R. Chandrupatla. *Optimization Concepts and Applications in Engineering*. Cambridge University Press, New York, NY, 2 edition, 2011.

[12] Boeing. *Current Market Outlook 2013–2032*. 2013. Available at `http://www.boeing.com/assets/pdf/commercial/cmo/pdf/Boeing_Current_Market_Outlook_2013.pdf`. Accessed March 2015.

[13] R. R. Brodeur and C. P. van Dam. Transition prediction for a two-dimensional reynolds-averaged navier-stokes method applied to wind turbine airfoils. *Wind Energy*, 4(2):61–75, October 2001.

[14] H. Buckley, B. Y. Zhou, and D. W. Zingg. Airfoil optimization using practical aerodynamic design requirements. *Journal of Aircraft*, 47(5):1707–1719, September 2010.

[15] H. P. Buckley and D. W. Zingg. An approach to aerodynamic design through numerical optimization. *AIAA Journal*, 51(8):1972–1981, August 2013.

[16] D. M. Bushnell. Aircraft drag reduction – a review. In *Proceedings of the Institution of Mechanical Engineers*, volume 217, 2003.

[17] D. M. Bushnell and M. H. Tuttle. Survey and bibliography on attainment of laminar flow control in air using pressure gradient and suction – volume I. Technical report, NASA Reference Publication 1035, 1979.

[18] L. Cameron, J. Early, and R. McRoberts. Metamodel assisted multi-objective global optimisation of natural laminar flow aerofoils. In *29th AIAA Applied Aerodynamics Conference*, number AIAA–2011–3001, Honolulu, Hawaii, June 2011.

[19] R. Campbell, M. Campbell, and T. Streit. Progress toward efficient laminar flow analysis and design. In *29th AIAA Applied Aerodynamics Conference*, number AIAA–2011–3527, Honolulu, Hawaii, Jun 2011.

[20] T. Cebeci and J. Cousteix. *Modeling and Computation of Boundary-Layer Flows*. Springer, Long Beach, 2nd edition, 2005.

[21] O. Chernukhin and D. Zingg. An investigation of multi-modality in aerodynamic shape optimization. In *20th AIAA Computational Fluid Dynamics Conference*, number AIAA–2011–3070, Honolulu, Hawaii, June 2011.

[22] Clean Sky Joint Undertaking. Clean Sky Joint Technology Initiative. Available at `http://www.cleansky.eu/content/homepage/about-us`. Accessed March 2015.

[23] J. Cliquet, R. Houdeville, and D. Arnal. Application of laminar-turbulent transition criteria in Navier-Stokes computations. *AIAA Journal*, 46(5):1182–1190, May 2008.

[24] J. Coder and M. Maughmer. A CFD-compatible transition model using an amplification factor transport equation. *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, (AIAA–2013–253), Jan 2013.

[25] C. de Boor. *A Practical Guide to Splines*, volume 27 of *Applied Mathematical Sciences*. Springer-Verlag, New York, January 1978.

[26] S. S. Dodbele. Design optimization of natural laminar flow bodies in compressible flow. *Journal of Aircraft*, 29(3):343–347, May 1992.

[27] M. Drela. *A User's Guide to MSES 2.95*. 1996.

[28] M. Drela and M. B. Giles. Viscous-inviscid analysis of transonic and low reynolds number airfoils. *AIAA Journal*, 25(10):1347–1355, Oct 1987.

[29] J. Driver. Optimal aerodynamic shape design with transition prediction. Master's thesis, University of Toronto, 2005.

[30] J. Driver and D. W. Zingg. Numerical aerodynamic optimization incorporating laminar-turbulent transition prediction. *AIAA Journal*, 45(8):1810–1818, August 2007.

[31] O. Edenhofer, R. Pichs-Madruga, Y. Sokona, E. Farahani, S. Kadner, K. Seyboth, A. Adler, I. Baum, S. Brunner, P. Eickemeier, B. Kriemann, J. Savolainen, S. Schlomer, C. von Stechow, T. Zwickel, and J. Minx. *Climate Change 2014: Mitigation of Climate Change*. Cambridge University Press, Cambridge, United Kingdom and New York, NY, USA, 2014. Contribution of Working Group III to the Fifth Assessment Report of the Intergovernmental Panel on Climate Change.

[32] R. Eppler. *Airfoil Design and Data*. Springer-Verlag, New York, 1990.

[33] European Commission. Horizon 2020. Available at `http://ec.europa.eu/programmes/horizon2020/en/what-horizon-2020`. Accessed March 2015.

[34] European Commission. *Aeronautics and Air Transport Research in the 7th Framework Programme (2007-2013)*, volume 3. Luxembourg, 2013.

[35] G. J. Follen, R. Del Rosario, R. Wahls, and N. Madavan. NASA's Fundamental Aeronautics Subsonic Fixed Wing Project: Generation N+3 Technology Portfolio. In *SAE International*, number 2011-01-2521, Oct 2011.

[36] J. G. Gatsis. *Preconditioning Techniques for a Newton-Krylov Algorithm for the Compressible Navier-Stokes Equations*. PhD thesis, University of Toronto, 2014.

[37] P. E. Gill, W. Murray, and M. A. Saunders. SNOPT: An SQP algorithm for large-scale constrained optimization. *SIAM Review*, 47(1):99–131, Jan 2005.

[38] A. Gopalarathnam, B. A. Broughton, B. D. McGranahan, and M. S. Selig. Design of low Reynolds number airfoils with trips. *Journal of Aircraft*, 40(4):768–775, July 2003.

[39] S. Gossling and P. Upham. *Climate Change and Aviation: Issues, Challenges, and Solutions*. Earthscan Publications Ltd, Sterling, VA, 2009.

[40] B. E. Green, J. L. Whitesides, R. L. Campbell, and R. E. Mineck. A method for the constrained design of natural laminar flow airfoils. In *14th AIAA Applied Aerodynamics Conference*, number AIAA–96–2502, New Orleans, Louisiana, June 1996.

[41] J. E. Green. Laminar flow control - back to the future? In *38th Fluid Dynamics Conference and Exhibit*, number AIAA–2008–3738, Seattle, Washington, Jun 2008.

[42] J. E. Green. Aviation and climate - where do we go from here? In *4th UTIAS International Workshop on Aviation and Climate Change*, Toronto, Canada, May 2014. Available at `http://www.utias.utoronto.ca/sustainable-aviation/international-workshop`. Accessed March 2015.

[43] Green Aviation Research & Development Network. *2013–2014 Annual Report: Reaching Greener Heights*. 2014.

[44] A. Hanifi, O. Amoignon, J. O. Pralits, and M. Chevalier. A gradient-based optimization method for natural laminar flow design. In *Seventh IUTAM Symposium on Laminar-Turbulent Transition*, pages 3–10. Springer Science + Business Media, Nov 2009.

[45] R. E. Hanson, K. M. Bade, B. A. Belson, P. Lavoie, A. M. Naguib, and C. W. Rowley. Feedback control of slowly-varying transient growth by an array of plasma actuators. *Physics of Fluids*, 26(024102), Feb 2014.

[46] J. Hicken. *Efficient Algorithms for Future Aircraft Design: Contributions to Aerodynamic Shape Optimization*. PhD thesis, University of Toronto, 2009.

[47] S. M. Hosseini, D. Tempelmann, A. Hanifi, and D. S. Henningson. Stabilization of a swept-wing boundary layer by distributed roughness elements. *J. Fluid Mech.*, 718, Feb 2013.

[48] International Air Transport Association. Improving environmental performance. Available at `http://www.iata.org/whatwedo/environment/Pages/index.aspx`. Accessed March 2015.

[49] International Civil Aviation Organization. Committee on Aviation Environmental Protection (CAEP). Available at `http://www.icao.int/environmental-protection/pages/CAEP.aspx`. Accessed March 2015.

[50] International Civil Aviation Organization. *Convention on International Civil Aviation*. Chicago, December 1944.

[51] International Civil Aviation Organization. ICAO's CO_2 standard as part of a basket of measures to meet emission reduction goals. Technical Report A38-WP/297, September 2013.

[52] A. Jameson, W. Schmidt, and E. Turkel. Numerical solution of the Euler equations by finite volume methods using Runge-Kutta time-stepping schemes. In *14th AIAA Fluid and Plasma Dynamics Conference*, number AIAA–81–1259, Palo Alto, California, June 1981.

[53] L. Jing, G. Zhenghong, H. Jiagntao, and Z. Ke. Robust design of NLF airfoils. *Chinese Journal of Aeronautics*, 26(2):309–318, April 2013.

[54] P. Khayatzadeh and S. K. Nadarajah. Aerodynamic shape optimization via discrete viscous adjoint equations for the k-ω SST turbulence and $\gamma-\widetilde{Re}_{\theta_t}$ transition models. In *49th AIAA Aerospace Sciences Meeting*, number AIAA–2011–1247, Orlando, Florida, 4–7 January 2011.

[55] P. Khayatzadeh and S. K. Nadarajah. Aerodynamic shape optimization of natural laminar flow (NLF) airfoils. In *50th AIAA Aerospace Sciences Meeting*, number AIAA–2012–0061, Nashville, Tennessee, 9–12 January 2012.

[56] P. Khayatzadeh and S. K. Nadarajah. Laminar-turbulent flow simulation for wind turbine profiles using the $\gamma-\widetilde{Re}_{\theta_t}$ transition model. *Wind Energy*, 17(6):901–918, 2013.

[57] B. Y. Kim, G. G. Fleming, J. J. Lee, I. A. Waitz, J. P. Clarke, S. Balasubramanian, A. Malwitz, K. Klima, M. Locke, C. A. Holsclaw, and et al. System for assessing aviation's global emissions (SAGE), part 1: Model description and inventory results. *Transportation Research Part D: Transport and Environment*, 12(5):325–346, Jul 2007.

[58] S. Kirkpatrick, C. D. Gelatt, and M. P. Vecchi. Optimization by simulated annealing. *Science*, 220(4598):671–680, May 1983.

[59] N. Krimmelbein, C. Nebel, and R. Radespiel. Numerical aspects of transition prediction for three-dimensional configurations. In *35th AIAA Fluid Dynamics Conference & Exhibit*, number AIAA–2005–4764, Toronto, Ontario, June 2005.

[60] N. Krimmelbein and R. Radespiel. Transition prediction for three-dimensional flows using parallel computation. *Computers & Fluids*, 38(1):121–136, 2009.

[61] I. Kroo and P. Sturdza. Design-oriented aerodynamic analysis for supersonic laminar flow wings. In *41st AIAA Aerospace Sciences Meeting & Exhibit*, number AIAA–2003–0774, Reno, Nevada, 6–9 January 2003.

[62] A. Krumbein. On modeling of transitional flow and its application on a high lift multi-element airfoil configuration. In *41st Aerospace Sciences Meeting and Exhibit*, number AIAA–2003–724, Reno, Nevada, January 2003.

[63] A. Krumbein. e^N transition prediction for 3D wing configurations using database methods and a local, linear stability code. *Aerospace Science and Technology*, 12(8):592–598, December 2008.

[64] A. Krumbein, N. Krimmelbein, and G. Schrauf. Automatic transition prediction in hybrid flow solver, part 1: Methodology and sensitivities. *Journal of Aircraft*, 46(4):1176–1190, July 2009.

[65] A. Krumbein, N. Krimmelbein, and G. Schrauf. Automatic transition prediction in hybrid flow solver, part 2: Practical application. *Journal of Aircraft*, 46(4):1191–1199, July 2009.

[66] R. B. Langtry and F. R. Menter. Transition modeling for general CFD applications in aeronautics. In *43rd AIAA Aerospace Sciences Meeting and Exhibit*, number AIAA–2005–522, Reno, Nevada, 10–13 January 2005.

[67] R. B. Langtry and F. R. Menter. Correlation-based transition modeling for unstructured parallelized computational fluid dynamics codes. *AIAA Journal*, 47(12):2894–2906, December 2009.

[68] D. S. Lee, D. W. Fahey, P. M. Forster, P. J. Newton, R. C. Wit, L. L. Lim, B. Owen, and R. Sausen. Aviation and global climate change in the 21st century. *Atmospheric Environment*, 43(22-23):3520–3537, Jul 2009.

[69] J. Lee and A. Jameson. Natural-laminar-flow airfoil and wing design by adjoint method and automatic transition prediction. In *47th AIAA Aerospace Sciences Meeting and Exhibit*, number AIAA–2009–897, Orlando, Florida, 5–8 January 2009.

[70] J. J. Lee, I. A. Waitz, B. Y. Kim, G. G. Fleming, L. Maurice, and C. A. Holsclaw. System for assessing aviation's global emissions (sage), part 2: Uncertainty assessment. *Transportation Research Part D: Transport and Environment*, 12(6):381–395, Aug 2007.

[71] H. Lomax, T. H. Pulliam, and D. W. Zingg. *Fundamentals of Compuational Fluid Dynamics*. Springer, Berlin, 2001.

[72] L. Mack. Transition prediction and linear stability theory. In *AGARD*, number CP-224, Lyngby, Denmark, May 1977.

[73] D. Maddalon, J. Collier, F.S., L. Montoya, and R. Putnam. Transition flight experiments on a swept wing with suction. In *Laminar-Turbulent Transition*, IUTAM Symposium, pages 53–62, Toulouse, 1990.

[74] M. Manzoor. Sustainable development - a major challenge to the aviation industry. *Messier-Dowty Inc, Canada*, 2008.

[75] J. H. Marburger. National aeronautics research and development plan. Technical report, National Science and Technology Council, February 2010.

[76] J. R. R. A. Martins, P. Sturdza, and J. J. Alonso. The complex-step derivative approximation. *Transactions on Mathematical Software*, 29(3):245–262, Sep 2003.

[77] L. Maurice and D. Lee. Assessing current scientific knowledge, uncertainties and gaps in quantifying climate change, noise and air quality aviation impacts. Technical report, Final Report of the International Civil Aviation Organization (ICAO) Committee on Aviation and Environmental Protection (CAEP) Workshop, 2009.

[78] E. Mayda. *Boundary Layer Transition Prediction for Reynolds-Averaged Navier-Stokes Methods*. PhD thesis, University of California, Davis, 2007.

[79] S. Medida and J. D. Baeder. Application of the correlation based γ–Re_{θ_t} transition model to the Spalart-Allmaras turbulence model. In *20th AIAA Computational Fluid Dynamics Conference*, number AIAA–2011–3979, Honolulu, Hawaii, June 2011.

[80] D. B. Middleton, G. W. Hanks, C. W.Clay, G. R. Swinford, T. C. Versteegh, and R. L. Sullivan. Natural laminar flow airfoils analysis and trade studies. Technical Paper 159029, National Aeronautics and Space Administration, June 1979.

[81] F. Moens, J. Perraud, A. Krumbein, T. Toulorge, P. Iannelli, E. P., and A. Hanifi. Transition prediction and impact on a three-dimensional high-lift-wing configuration. *Journal of Aircraft*, 45(5):1751–1766, September 2008.

[82] A. Mosahebi and E. Laurendeau. Convergence characteristics of fully and loosely coupled numerical approaches for transition models. *AIAA Journal*, 53(5):1399–1404, May 2015.

[83] C. Nebel, R. Radespiel, and T. Wolf. Transition prediction for 3D flows using a Reynolds-averaged Navier-Stokes code and N-factor methods. In *33rd AIAA Fluid Dynamics Conference and Exhibit*, number AIAA–2003–3593, Jun 2003.

[84] J. A. Nelder and R. Mead. A simplex method for function minimization. *The Computer Journal*, 7(4):308–313, January 1965.

[85] M. Nemec. *Optimal Shape Design Of Aerodynamic Configurations: A Newton-Krylov Approach*. PhD thesis, University of Toronto, 2003.

[86] M. Nemec and D. W. Zingg. Newton-Krylov algorithm for aerodynamic design using the Navier-Stokes equations. *AIAA Journal*, 40(6):1146–1154, June 2002.

[87] M. Nemec, D. W. Zingg, and T. H. Pulliam. Multipoint and multi-objective aerodynamic shape optimization. *AIAA Journal*, 42(6):1057–1065, June 2004.

[88] J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, second edition, 2006.

[89] L. P. Osmokrovic, R. E. Hanson, and P. Lavoie. Laminar boundary-layer response to spanwise periodic forcing by dielectric-barrier-discharge plasma-actuator arrays. *AIAA Journal*, 53(3):617–628, March 2015.

[90] S. Özgen. Linear stability theory and laminar-turbulent transition. In *Middle East Technical University Lecture Series*, 2014.

[91] J. E. Penner, D. H. Lister, D. J. Griggs, D. J. Dokken, and M. McFarland. *Aviation and the Global Atmosphere*. Cambridge University Press, Cambridge, 1999. Published for the Intergovernmental Panel on Climate Change.

[92] D. I. A. Poll and M. Danks. Relaminarisation of the swept wing attachment-line by surface suction. In *Laminar-Turbulent Transition IUTAM Symposium*, pages 137–144, Sendai, Japan, September 1995. Springer Berlin Heidelberg.

[93] J. L. Porter. Guide for the verification and validation of computational fluid dynamics simulations. *AIAA Guide*, (AIAA G-077-1998), Jan 1998.

[94] J. Pralits. *Optimal Design of Natural and Hybrid Laminar Flow Control on Wings*. PhD thesis, Royal Institute of Technology Department of Mechanics, Stockholm, Sweden, October 2003.

[95] A. Pueyo. *An Efficient Newton-Krylov Method for the Euler and Navier-Stokes Equations*. PhD thesis, University of Toronto, February 1998.

[96] T. H. Pulliam and D. W. Zingg. *Fundamental Algorithms in Computational Fluid Dynamics*. Springer International Publishing, 2014.

[97] D. Quagliarella and A. Della Cioppa. Genetic algorithms applied to the aerodynamic design of transonic airfoils. *Journal of Aircraft*, 32(4):889–891, July 1995.

[98] R. Rashad and D. W. Zingg. Toward high-fidelity aerodynamic shape optimization for natural laminar flow. In *21st AIAA Computational Fluid Dynamics Conference*, number AIAA–2013–2583, San Diego, CA, June 2013.

[99] H. L. Reed and W. S. Saric. Transition mechanisms for transport aircraft. In *38th Fluid Dynamics Conference and Exhibit*, number AIAA–2008–3743, Seattle, Washington, June 2008.

[100] N. Rivers. *The Case for a Carbon Tax in Canada*. Canada 2020, 2014.

[101] Y. Saad. *Iterative Methods for Large Sparse Systems*. PWS Pub Co, second edition, January 2000.

[102] W. Saric, H. Reed, and E. Kershcen. Leading-edge receptivity to sound: Experiments, DNS, theory. In *25th AIAA Fluid Dynamics Conference*, number AIAA–94–2222, Colorado Springs, CO, June 1994.

[103] W. S. Saric, D. E. West, M. W. Tufts, and H. L. Reed. Flight test experiments on discrete roughness element technology for laminar flow control. In *53rd AIAA Aerospace Sciences Meeting*, number 2015–0539. American Institute of Aeronautics and Astronautics (AIAA), Jan 2015.

[104] W. S. Saric, E. B. White, and H. L. Reed. Boundary-layer receptivity to freestream disturbances and its role in transition. In *30th AIAA Fluid Dynamics Conference*, number AIAA–99–3788, Norfolk, VA, June 1999.

[105] H. Schlichting. *Boundary-Layer Theory*. McGraw-Hill, Toronto, 7th edition, 1979.

[106] V. Schmitt, J. Archambaud, K. Horstmann, and A. Quast. Hybrid laminar fin investigations. In *NATO/RTO/AVT Symposium on Active Control Technology*, number ADPO11101, Braunschweig, Germany, May 2000.

[107] M. S. Selig. Low reynolds number airfoil design. In *VKI Lecture Series: Applied Vehicle Technology (AVT) Panel*, November 2003.

[108] T. K. Sengupta. *Theoretical and Compuational Aerodynamics*. John Wiley & Sons, Ltd, 2015.

[109] A. Smith and N. Gamberoni. Transition, pressure gradient and stability theory. Technical Report ES–26388, Douglas Aircraft Company, 1956.

[110] D. M. Somers. Design and experimental results for a natural-laminar-flow airfoil for general aviation applications. Technical Paper 1861, National Aeronautics and Space Administration, June 1981.

[111] D. M. Somers. Design and experimental results for the s809 airfoil. Technical Paper NREL/SR-440-6918, National Renewable Energy Lab, January 1997.

[112] P. Spalart and S. Allmaras. A one-equation turbulence model for aerodynamic flows. In *30th AIAA Aerospace Sciences Meeting & Exhibit*, number AIAA–092–0439, Reno, NV, January 1992.

[113] W. Squire and G. Trapp. Using complex variables to estimate derivatives of real functions. *SIAM Review*, 40(1):110–112, January 1998.

[114] H. W. Stock and W. Haase. Feasibility study of e^n transition prediction in Navier-Stokes methods for airfoils. *AIAA Journal*, 37(10):1187–1196, October 1999.

[115] H. W. Stock and W. Haase. Navier-Stokes airfoil computations with e^n transition prediction including transitional flow regions. *AIAA Journal*, 38(11):2059–2066, November 2000.

[116] H. W. Stock and A. Seitz. Crossflow-induced transition prediction using coupled Navier–Stokes and e^n method computations. *AIAA Journal*, 42(9):1746–1754, September 2004.

[117] T. Streit, H. Horstmann, G. Shraut, S. Hein, U. Fey, Y. Egmai, J. Perraud, I. Salah El Din, U. Cella, and J. Quest. Complementary numerical and experimental data analysis of the ETW Telfona Pathfinder wing transition tests. In *49th Aerospace Sciences Meeting and Exhibit*, number AIAA–2011–881, Orlando, Florida, January 2011.

[118] J. Stuart. Perspectives in turbulence studies. 1987.

[119] P. Sturdza. *An Aerodynamic Desgin Method for Supersonic Natural Laminar Flow Aircraft*. PhD thesis, Stanford University, December 2003.

[120] P. Sturdza. Extensive supersonic natural laminar flow on the Aerion business jet. In *45th AIAA Aerospace Sciences Meeting and Exhibit*, number AIAA-2007-685, Reno, NV, January 2007. Desktop Aeronautics Inc.

[121] P. Sturdza and D. G. Rajnarayan. Design and validation of a supersonic natural laminar flow test article. In *15th AIAA/ISSMO Multidisciplinary Analysis and Optimization Conference*, number AIAA–2014–2176, June 2014.

[122] Sustainable Aviation. *Sustainable Aviation CO_2 Road Map*. London, UK, 2012.

[123] R. Swanson and E. Turkel. On central-difference and upwind schemes. *Journal of Computational Physics*, 101(2):292–306, Aug 1992.

[124] The World Bank. World development indicators: Air transport, passengers carried. Available at `http://data.worldbank.org/indicator/IS.AIR.PSGR`. Accessed March 2015.

[125] Transport Canada. *Canada's Action Plan to Reduce Greenhouse Gas Emissions from Aviation*. 2012. TP 15187.

[126] A. H. Truong. *Development of a Computer-Aided-Desing-Based Geometry and Mesh Movement Algorithm for 3D Aerodynamic Shape Optimization*. PhD thesis, University of Toronto, 2014.

[127] P. Upham, J. Maugham, D. Raper, and C. Thomas. *Towards Sustainable Aviation*. Earthscan Publications Ltd, London, 2003.

[128] N. van Dorland, F. van der Zwan, S. Ghijs, S. Santema, and R. Curran. Aviation and the environment: Rating airlines on their CO_2 efficiency. In *9th AIAA Aviation Technology, Integration, and Operations Conference (ATIO)*, number AIAA–2009-7030. American Institute of Aeronautics and Astronautics, Sep 2009.

[129] J. van Ingen. A suggested semi-empirical method for the calculation of the boundary layer transition region. Technical Report VTH–74, Delft University of Technology, 1956.

[130] R. Wahls. NRC meeting of experts on NASA's plans for system-level research in environmental mitigation. National Harbor, MD: http://www.aeronautics.nasa.gov/calendar/20090514.htm, 14 May 2009.

[131] X. Wang, J. Cai, C. Liu, and Z. Hu. Airfoil optimization based on rapid transition prediction. In *56th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*, number AIAA–2015–0400, Jan 2015.

[132] F. M. White. *Viscous Fluid Flow*. McGraw-Hill, Inc., third edition, 2005.

[133] K. Wicke, M. Kruse, and F. Linke. Mission and economic analysis of aircraft with natural laminar flow. In *28th International Congress of the Aeronautical Sciences*, number 2012-1.7.4, Brisbane, Australia, September 2012.

[134] D. C. Wilcox. *Turbulence Modeling for CFD*. DCW Industries, La Cañada, 1993.

[135] C. Wohl, J. Smith, J. Connell, E. Siochi, R. Penner, and J. Gardner. Engineered surfaces for mitigation of insect residue adhesion. In *51st AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*, number AIAA–2013–413, January 2013.

[136] F. Zahle, C. Bak, N. N. Srensen, T. Vronsky, and N. Gaudern. Design of the LRP airfoil series using 2D CFD. *Journal of Physics: Conference Series*, 524(012020), Jun 2014.

[137] K. Zhao, Z. Gao, and J. Huang. Robust design of natural laminar flow supercritical airfoil by multi-objective evolution method. *Appl. Math. Mech.-Engl. Ed.*, 35(2):191202, Jan 2014.

[138] D. Zingg, M. Nemec, and T. Pulliam. A comparative evaluation of genetic and gradient-based algorithms applied to aerodynamic optimization. *REMN*, 17:103–126, 2008.

(invited paper in a special issue on adjoint-based shape design in Revue Europeenne de Mecanique Numerique - European Journal of Computational Mechanics).

[139] D. W. Zingg. An approach to the design of airfoil with high lift to drag ratios. Master's thesis, University of Toronto, 1981.

[140] D. W. Zingg and P. Godin. A perspective on turbulence models for aerodynamic flows. *International Journal of Computational Fluid Dynamics*, 23(4):327–335, April 2009.

[141] D. W. Zingg and O. L. Gülder. Technology developments and renewable fuels for sustainable aviation. Chapter to appear in Sustainable Development in International Aviation Law.

[142] D. W. Zingg, T. Leung, L. Diosady, A. Truong, S. Elias, and M. Nemec. Improvements to a Newton-Krylov adjoint algorithm for aerodynamic optimization. In *17th AIAA Computational Fluid Dynamics Conference*, number AIAA–2005–4857, Jun 2005.

APPENDICES

Appendix A

Governing Equations in Generalized Curvilinear Coordinates

The governing equations are transformed into a generalized curvilinear coordinate system. Section A.1 presents the transformed Navier-Stokes equations, followed by the Spalart-Allmaras turbulence model in Section A.2.

A.1 Thin-Layer Navier-Stokes Equations

A thin-layer approximation is invoked to simplify the Navier-Stokes equations for fully-attached or mildly-separated flows at high Reynolds numbers. The viscous fluxes in the off-wall η-direction are included in the Navier-Stokes equations, while the viscous fluxes in the ξ-direction are ignored. The resulting equations in generalized curvilinear coordinates are given by [96]

$$\frac{\partial \widehat{Q}}{\partial \tau} + \frac{\partial \widehat{E}}{\partial \xi} + \frac{\partial \widehat{F}}{\partial \eta} = \mathcal{R}e^{-1}\frac{\partial \widehat{S}}{\partial \eta}\,, \tag{A.1}$$

where the convective flux vectors are

$$\widehat{E} = J^{-1}\begin{pmatrix} \rho U \\ \rho U u + \xi_x p \\ \rho U v + \xi_y p \\ (e+p)U - \xi_t p \end{pmatrix}, \quad \widehat{F} = J^{-1}\begin{pmatrix} \rho V \\ \rho V u + \eta_x p \\ \rho V v + \eta_y p \\ (e+p)V - \eta_t p \end{pmatrix}, \tag{A.2}$$

the contravariant velocities are

$$U = \xi_t + \xi_x u + \xi_y v \tag{A.3}$$
$$V = \eta_t + \eta_x u + \eta_y v \tag{A.4}$$

and the viscous flux vector is

$$\widehat{S} = J^{-1} \begin{pmatrix} 0 \\ \eta_x m_1 + \eta_y m_2 \\ \eta_x m_2 + \eta_y m_3 \\ \eta_x(u m_1 + v m_3 + m_4) + \eta_y(u m_2 + v m_3 + m_5) \end{pmatrix}, \tag{A.5}$$

where

$$\begin{aligned} m_1 &= (\mu + \mu_t)(4\eta_x u_\eta - 2\eta_y v_\eta)/3 \\ m_2 &= (\mu + \mu_t)(\eta_y u_\eta + \eta_x v_\eta) \\ m_3 &= (\mu + \mu_t)(-2\eta_x u_\eta + 4\eta_y v_\eta)/3 \\ m_4 &= (\mu \mathcal{P}r^{-1} + \mu_t \mathcal{P}r_t^{-1})(\gamma - 1)^{-1} \eta_x \partial_\eta(a^2) \\ m_5 &= (\mu \mathcal{P}r^{-1} + \mu_t \mathcal{P}r_t^{-1})(\gamma - 1)^{-1} \eta_y \partial_\eta(a^2)\,. \end{aligned} \tag{A.6}$$

A.2 Spalart-Allmaras Model in Generalized Curvilinear Coordinates

The generalized curvilinear transformation is also applied to the Spalart-Allmaras turbulence model given by (2.10). Neglecting any terms containing mixed derivatives, the Spalart-Allmaras turbulence model becomes [36]

$$\frac{\partial \tilde{\nu}}{\partial \tau} + U\frac{\partial \tilde{\nu}}{\partial \xi} + V\frac{\partial \tilde{\nu}}{\partial \eta} = \frac{1}{\mathcal{R}e}\left\{ c_{b1}\tilde{S}\tilde{\nu} - c_{w1} f_w \left(\frac{\tilde{\nu}}{d_{\mathrm{w}}}\right)^2 + \frac{1}{\sigma}\left[(1 + c_{b2})\, T_1 - c_{b2} T_2\right] \right\}, \tag{A.7}$$

where

$$\begin{aligned} T_1 = \xi_x \frac{\partial}{\partial \xi}\left[(\nu + \tilde{\nu})\, \xi_x \frac{\partial \tilde{\nu}}{\partial \xi}\right] + \eta_x \frac{\partial}{\partial \eta}\left[(\nu + \tilde{\nu})\, \eta_x \frac{\partial \tilde{\nu}}{\partial \eta}\right] \\ + \xi_y \frac{\partial}{\partial \xi}\left[(\nu + \tilde{\nu})\, \xi_y \frac{\partial \tilde{\nu}}{\partial \xi}\right] + \eta_y \frac{\partial}{\partial \eta}\left[(\nu + \tilde{\nu})\, \eta_y \frac{\partial \tilde{\nu}}{\partial \eta}\right] \end{aligned} \tag{A.8}$$

and

$$T_2 = (\nu + \tilde{\nu}) \times \left[\xi_x \frac{\partial}{\partial \xi} \left(\xi_x \frac{\partial \tilde{\nu}}{\partial \xi} \right) + \eta_x \frac{\partial}{\partial \eta} \left(\eta_x \frac{\partial \tilde{\nu}}{\partial \eta} \right) + \xi_y \frac{\partial}{\partial \xi} \left(\xi_y \frac{\partial \tilde{\nu}}{\partial \xi} \right) + \eta_y \frac{\partial}{\partial \eta} \left(\eta_y \frac{\partial \tilde{\nu}}{\partial \eta} \right) \right] \ . \quad \text{(A.9)}$$

Appendix B

B-spline Parameterization of Airfoil Geometry

Following the work of Nemec and Zingg [85, 86], B-splines are used to parameterize the airfoil geometries, as depicted in Figures 4.2 and 5.11. The airfoil representation based on B-splines is given by [85]

$$x_a(w_j) = \sum_{i=1}^{n+1} X_i^c B_{i,k}(w_j) \tag{B.1}$$

$$y_a(w_j) = \sum_{i=1}^{n+1} Y_i^c B_{i,k}(w_j)\,, \tag{B.2}$$

where x_a, y_a are the Cartesian coordinates of the airfoil surface, X_i^c, Y_i^c are the coordinates of the B-spline coordinates, and $n+1$ is the total number of control points. In this work, the B-spline basis functions, $B_{i,k}$, are taken to be cubic with order $k = 4$, and degree $k-1$.

For any point (indexed by j) along the surface of the airfoil, the distance along the B-spline curve is given by [85]

$$w_j = (n-k+2)\left(\sum_{m=1}^{N-1}\sqrt{L_m}\right)^{-1}\left(\sum_{m=1}^{j-1}\sqrt{L_m}\right) \qquad \text{for} \quad j = 2\ldots N\,, \tag{B.3}$$

where N is the total number of airfoil surface points, $w_1 = 0$, and L_m is the segment length between surface points. Finally, the B-spline basis functions are given by the Cox-de Boor recurrence relation [25] as

$$B_{i,k}(w_j) = \begin{cases} 1 & \text{if } d_i \le w_j < d_{i+1} \\ 0 & \text{otherwise}\,, \end{cases}$$

where d is a nonuniform knot sequence (with clustering near the leading and trailing edges)

defined as

$$d_i = \begin{cases} 0 & \text{for} \quad 1 \leq i \leq k \\ \frac{n-k+2}{2}\left[1 - \cos\left(\frac{i-k}{n-k+2}\pi\right)\right] & \text{for} \quad k+1 \leq i \leq n+1 \\ n-k+2 & \text{for} \quad n+2 \leq i \leq n+k+1\,, \end{cases}$$

The initial locations of the B-spline control points that best approximate the initial shape of the airfoil must be determined as a pre-processing step to the optimization. A description of the least-squares algorithm used to minimize the fitting error may be found in Nemec's PhD thesis [85].