


Investigation of Airfoil Flow Control Using Adaptive Surface Roughness

Anders Aaboe Jensen

Delft University of Technology & Technical University of Denmark



Investigation of Airfoil Flow Control Using Adaptive Surface Roughness

by

Anders Aaboe Jensen

to obtain the dual degrees of Master of Science
in Aerospace Engineering at the Delft University of Technology,
and
in Wind Energy at the Technical University of Denmark,
to be defended publicly on Wednesday, August 26, 2026 at 11:00 AM.

Student number:	s214509 (DTU) / 6309186 (TU Delft)		
Project duration:	November 3, 2025 – July 31, 2026		
Thesis committee:	Associate Professor D. Ragni, Aerodynamics Consultant S. Hjort, Senior Researcher R. F. Mikkelsen, Assistant Professor W. Yu, Associate Professor H. Sarlak,	TU Delft Aerospace, SHJ Consult ApS, DTU Wind, TU Delft Aerospace, DTU Wind,	Chair External examiner Supervisor Supervisor Co-supervisor

Cover: Conceptual visualisation of smooth and dimpled airfoil wake behaviour, created using image generation.
Style: TU Delft Report Style, with modifications by Daan Zwaneveld



*Til min kære mor og far, tak for jeres ubetingede kærlighed, jeres utrættelige
opbakning gennem hele livet og for altid at have troet på mig.
Uden jer var jeg ikke blevet den, jeg er i dag.*

Preface

This thesis was prepared as part of the European Wind Energy Master (EWEM) dual-degree programme, leading to the degrees of Master of Science in Aerospace Engineering at Delft University of Technology and Master of Science in Wind Energy at the Technical University of Denmark. The work was carried out between November 2025 and July 2026 within the Department of Wind and Energy Systems at DTU and the Faculty of Aerospace Engineering at TU Delft.

I gratefully acknowledge Siemens Fonden (the Siemens Foundation, Denmark) for their financial support of this project. Their generous support contributed to the experimental work and the manufacture of the proof-of-concept prototype presented in this thesis.

My sincere gratitude goes to Senior Researcher Robert Flemming Mikkelsen, Assistant Professor Wei Yu, and Associate Professor Hamid Sarlak for their guidance, support, and valuable discussions throughout this project. The weekly meetings and day-to-day sparring sessions made the project not only better but also far more enjoyable. To each of you, thank you for your encouragement, your time, and for making this project such a rewarding experience.

I am also grateful to the technicians and PhD candidates at the Low Speed Lab at TU Delft, who were unfailingly friendly and generous with their time throughout my measurement campaign there. Furthermore, I would like to extend a special thanks to the DTU Construct Workshop Team 413. Across all five years of my studies at DTU, walking through the doors of that workshop to ask for advice or learn something new about manufacturing has been one of the genuine pleasures of my time there, and their help with the practical work in this thesis was no exception. Thank you to every single one of you.

To my dear friend Oliver, we met on the very first day of our bachelor's, and since then we have shared nearly every project, class, and challenge along the way. Throughout these five years, you have been my steadfast support as a colleague and, more importantly, as a friend. These years would not have been the same without you. Thank you, my dear friend.

Finally, my deepest gratitude belongs to my parents. Thank you for your unconditional love, your unwavering belief in me, and your endless encouragement. Thank you for always giving me the freedom to pursue the path that felt right for me, wherever it took me, and for supporting me every step of the way.

Anders Aaboe Jensen
Kongens Lyngby, July 2026

Abstract

Surface dimples are a well-established means of reducing drag on bluff bodies, most famously the golf ball, where tripping the boundary layer into turbulence delays the flow separation and narrows the wake, reportedly halving drag relative to a smooth sphere. Whether the same mechanism benefits a lifting surface is far less settled. An airfoil in attached flow has little separation to delay, so a dimple that helps a sphere can just as easily add drag to a wing. Existing studies on dimpled airfoils reflect this tension, reporting inconsistent benefits that depend strongly on dimple depth, size, and chordwise location, with no overwhelming consensus on where along the chord dimples should be placed. This thesis investigates how the controlled variation of dimple depth and chordwise placement affects the aerodynamic characteristics and flow-separation behaviour of a lifting surface across varying operating conditions, and develops a proof-of-concept prototype capable of mechanically varying dimple depth as a first step toward a fully adaptive surface.

A NACA 63-415 airfoil with chord $c = 200$ mm was tested with three chordwise dimple coverage families, mid (0.33c to 0.59c), aft (0.59c to 0.85c), and full (0.33c to 0.85c) suction-side coverage, each at three dimple depths, across three Reynolds numbers (1×10^5 , 2×10^5 , and 5×10^5). Measurements spanned an angle-of-attack range of -10° to 20° , with 0° to 16° forming the most substantial part of the campaign. The campaign combined wake-rake and wall-pressure measurements, surface-pressure taps, a force balance, particle image velocimetry (PIV), and tuft flow visualisation across two wind-tunnel facilities, the Red Wind Tunnel at DTU and the Small Low-Turbulence Tunnel at TU Delft, supported by RANS and LES simulations.

The results show that dimple depth, chordwise location, and Reynolds number interact strongly, and no single fixed combination performs well across the full envelope tested. At the lowest Reynolds number, mid and full coverage produce substantial glide-ratio gains, up to 132% over the smooth baseline, driven primarily by drag reduction from delayed separation. At the highest Reynolds number, this reverses entirely. Mid coverage becomes uniformly detrimental, while aft coverage becomes the only layout to show any benefit, concentrated near stall. Dimple depth shows a similarly reversing, non-monotonic effect. The depth that gives the largest drag reduction at the lowest Reynolds number is the most harmful choice at the highest, and near stall this non-monotonicity appears even within a single Reynolds number, where a depth of 2.50 mm produces a genuine drag reduction while both a shallower (1.25 mm) and a deeper (5.00 mm) variant do not. PIV, surface-pressure, and wake-rake measurements on the standout mid configuration confirm the underlying mechanism directly. The dimples energise the near-wall boundary layer, suppress separation that is otherwise present on the smooth baseline, and narrow the wake to roughly half the width of its baseline counterpart, extending to a measurable stall-delay effect at a higher Reynolds number. RANS and LES reproduce the broad direction of these trends but not the mechanism itself in detail, consistent with the RANS model's lack of a transition formulation.

These findings show that no single, fixed dimple geometry is optimal across a realistic operating envelope, motivating a surface that can adjust its own geometry in response to the flow condition. A proof-of-concept prototype was designed and built around a pneumatically actuated, depth-adjustable membrane, and structured-light scanning related each actuation setting to a physical dimple depth. Wind-tunnel testing at $Re = 1 \times 10^5$ showed the prototype reproducing the drag behaviour of the static configurations in trend, with a peak drag reduction of approximately 22% relative to its own unactuated state, and reproducing the same suction-side wake narrowing at matched depth. The correspondence is qualitative rather than exact, and no lift benefit is observed, but a single airfoil is shown to select between aerodynamic states that individually resemble distinct static geometries, establishing the feasibility of an adaptive dimpled surface.

Nomenclature

Abbreviations

Abbreviation	Definition
AoA	Angle of Attack
CAD	Computer-Aided Design
CFD	Computational Fluid Dynamics
DTU	Technical University of Denmark
EWEM	European Wind Energy Master
FOV	Field of View (PIV)
LES	Large-Eddy Simulation
LLM	Large-Language Model
LSB	Laminar Separation Bubble
NACA	National Advisory Committee for Aeronautics
Nd:YAG	Neodymium-doped Yttrium Aluminium Garnet (PIV laser)
OpenFOAM	Open-source Field Operation And Manipulation (CFD toolbox)
PIV	Particle Image Velocimetry
PTU	Programmable Timing Unit (PIV)
RA	Research Activity
RANS	Reynolds-Averaged Navier–Stokes
RWT	Red Wind Tunnel (DTU)
SGS	Sub-Grid-Scale (LES)
SLT	Small Low-Turbulence Tunnel (TU Delft)
SQ	Sub-Question
SST	Shear Stress Transport (turbulence model)
TKE	Turbulent Kinetic Energy
UAV	Unmanned Aerial Vehicle
XFoil	Viscous–inviscid panel-method airfoil solver

Symbols

Symbol	Definition	Unit
A	Area	m^2
A_{proj}	Projected frontal area of the model	m^2
A_{ts}	Wind-tunnel test-section cross-sectional area	m^2
b	Airfoil span	m
c	Airfoil chord length	m
C_a	Axial-force coefficient (body axis)	—
C_D	Drag coefficient	—
C_f	Skin-friction coefficient	—
C_L	Lift coefficient	—
C_L/C_D	Lift-to-drag (glide) ratio	—
$C_{m,c/4}$	Quarter-chord pitching-moment coefficient	—
C_n	Normal-force coefficient (body axis)	—
C_p	Pressure coefficient	—
D	Drag force per unit span	N/m
D_{total}	Total drag force	N
D_d	Dimple diameter	m
D_d/c	Dimple diameter-to-chord ratio	—
D_s	Dimple centre-to-centre spacing	m
D_s/D_d	Dimple spacing ratio	—
F_x	Streamwise force component (force balance)	N
F_y	Force component normal to the freestream (force balance)	N
G	Glide ratio, $G \equiv C_L/C_D$	—
h_{ts}	Wind-tunnel test-section height	m
H	Boundary-layer shape factor (δ^*/θ)	—
K_1	Body-shape factor (solid-blockage correction)	—
k	Dimple depth	m
k	Turbulent kinetic energy (Chapter 5)	m^2/s^2
k_∞	Freestream turbulent kinetic energy	m^2/s^2
k/c	Dimple depth-to-chord ratio	—
k/D_d	Dimple depth-to-diameter ratio	—
L	Lift force per unit span	N/m
ℓ	Turbulent length scale	m
M_v	Model volume	m^3
p	Static pressure	Pa
p_∞	Freestream static pressure	Pa
$p_{t,\infty}$	Freestream total pressure	Pa
$p_{t,w}$	Local total pressure at wake rake	Pa
p_w	Wall static pressure	Pa
q_∞	Freestream dynamic pressure	Pa
Re	Reynolds number	—
Re_k	Roughness Reynolds number	—
S_{ref}	Reference planform area, $S_{\text{ref}} = cb$	m^2
t/c	Thickness-to-chord ratio	—
Tu	Freestream turbulence intensity	%
u	Streamwise velocity component	m/s
U_e	Local velocity at boundary-layer edge	m/s
u_τ	Friction velocity	m/s
U_∞	Freestream velocity	m/s
x/c	Non-dimensional chordwise position	—
y/c	Non-dimensional position in the lift direction	—
z/c	Non-dimensional spanwise position	—
x^+	Dimensionless streamwise cell size (LES)	—
y^+	Dimensionless wall distance	—
z^+	Dimensionless spanwise cell size (LES)	—

Symbol	Definition	Unit
α	Angle of attack	$^{\circ}$
β	Blockage ratio	—
Δ	LES filter width	m
δ	Boundary-layer thickness	m
δ^*	Displacement thickness	m
ε	Turbulent dissipation rate	m^2/s^3
ε_{sb}	Solid-blockage correction factor	—
ε_{wb}	Wake-blockage correction factor	—
θ	Momentum thickness	m
μ	Dynamic viscosity	$\text{Pa}\cdot\text{s}$
μ_t	Turbulent dynamic viscosity	$\text{Pa}\cdot\text{s}$
ν	Kinematic viscosity	m^2/s
ν_t	Turbulent (eddy) kinematic viscosity	m^2/s
ν_{sgs}	Sub-grid-scale eddy viscosity	m^2/s
ρ	Fluid density	kg/m^3
σ	Streamline-curvature parameter	—
τ_w	Wall shear stress	Pa
ω	Specific dissipation rate	s^{-1}

Contents

Preface	i
Abstract	ii
Nomenclature	iii
1 Introduction	1
1.1 Research Questions and Activities	3
2 Theory and Literature Review	4
2.1 Boundary Layers	5
2.1.1 Laminar and Turbulent Boundary Layers	5
2.1.2 Wall Shear Stress and Skin Friction	6
2.1.3 Boundary-Layer Development and Transition	7
2.1.4 Reynolds Number and Critical Reynolds Number	8
2.1.5 Integral Thicknesses and the Shape Factor	9
2.2 Adverse Pressure Gradients and Flow Separation	9
2.2.1 Adverse Pressure Gradient	10
2.2.2 Flow Separation	11
2.2.3 The Laminar Separation Bubble	12
2.3 Surface Roughness and the Dimple Mechanism	13
2.3.1 Roughness-Induced Transition	14
2.3.2 Dimple-Induced Local Separation and Reattachment	14
2.4 Bluff-Body Drag Reduction: Dimples on a Golf Ball	15
2.5 Transfer of the Dimple Mechanism to Airfoils	16
2.5.1 Dimple Geometry as Design Parameters	17
2.5.2 Expected Effects Across Airfoil Operating Regimes	18
2.6 Literature Review	19
2.7 Airfoil Design Parameters	22
2.7.1 Wind-Tunnel Wall Interference and Blockage	22
2.7.2 Selection of Chord Length	23
2.7.3 Selection of Airfoil Profile	24
2.7.4 Selection of Dimple Diameter, Depth, and Spacing	25
2.7.5 Selection of Dimple Chord Location	26
2.7.6 Reynolds Number Range	28
2.8 Summary of Final Design Parameters	28
3 Methodology I - Aerodynamic Coefficient Measurements	30
3.1 Wind Tunnel Facility	30
3.2 Airfoil Models and Setup	31
3.3 Wake Rake and Wall Pressure Measurements	32
3.3.1 Experimental Setup	32
3.3.2 Wake and Wall Pressure Test Matrix	33
3.3.3 Summary of Wake and Wall Pressure Setup	34
3.4 Surface Pressure Taps	35
3.4.1 Experimental Setup	35
3.4.2 Surface Pressure Tap Test Matrix	37
3.4.3 Summary of Surface Pressure Tap Setup	37
4 Methodology II - Flow Visualisation	39
4.1 Wind Tunnel Facility	39
4.2 Airfoil Models and Setup	40
4.3 Particle Image Velocimetry	41

4.3.1	Experimental Setup	41
4.3.2	PIV Test Matrix	43
4.3.3	Summary of PIV Setup	44
4.4	Force Balance Measurements	45
4.4.1	Experimental Setup	45
4.4.2	Force Balance Test Matrix	45
4.4.3	Summary of Force Balance Setup	46
4.5	Tuft Flow Visualisation	47
4.5.1	Experimental Setup	47
4.5.2	Summary of Tuft Setup	48
5	Methodology III - Computational Fluid Dynamics	49
5.1	Reynolds Averaged Navier-Stokes	50
5.1.1	Study Design and Test Matrix	50
5.1.2	Computational Setup	52
5.1.3	Turbulence Modelling	52
5.1.4	Computational Domain and Boundary Conditions	55
5.1.5	Mesh Generation	56
5.1.6	Mesh Sensitivity Study	58
5.1.7	Summary of RANS Setup	59
5.2	Large-Eddy Simulations	61
5.2.1	Study Design and Test Matrix	61
5.2.2	Governing Equations and Sub-Grid-Scale Model	62
5.2.3	Computational Domain and Boundary Conditions	63
5.2.4	Mesh Generation and Wall-Unit Estimates	63
5.2.5	Summary of LES Setup	65
6	Experimental and Numerical Validation	66
6.1	Cross-Method Validation and Comparison	68
6.1.1	RANS Validation	68
6.1.2	LES Validation	68
6.1.3	Wall-Pressure and Wake Rake Validation	68
6.1.4	Surface Pressure Tap Validation	69
6.1.5	Force Balance Validation	69
6.2	Wake Rake Traverse Study	70
6.3	Spanwise Uniformity from Tuft Visualisation	71
6.4	PIV Validation	72
7	Results	74
7.1	Results Overview	75
7.2	Aerodynamic Screening	76
7.2.1	Reynolds Number 1×10^5	77
7.2.2	Reynolds Number 2×10^5	77
7.2.3	Reynolds Number 5×10^5	79
7.3	Effect of Chordwise Dimple Location	80
7.3.1	Reynolds Number 1×10^5	80
7.3.2	Reynolds Number 2×10^5	81
7.3.3	Reynolds Number 5×10^5	82
7.4	Effect of Dimple Depth	83
7.4.1	Reynolds Number 1×10^5	83
7.4.2	Reynolds Number 2×10^5	85
7.4.3	Reynolds Number 5×10^5	86
7.5	Investigation of Flow Behaviour of Mid-2.50	88
7.5.1	PIV Flow Field	88
7.5.2	Surface Pressure Distributions	89
7.5.3	Wake Rake Momentum Loss	91
7.5.4	Stall Delay Benefits	91
7.6	Experimental and Numerical Comparison of Mid-2.50	94
7.6.1	Flow-field Comparison	94

7.6.2	Surface Pressure Comparison	96
7.7	Summary of Results	98
7.7.1	The Reynolds Number Trend	98
7.7.2	Chordwise Placement	98
7.7.3	Dimple Depth	98
7.7.4	The Flow Behaviour and Flow Fields	98
7.7.5	Numerical Method Limitations	99
7.7.6	Connection to the Regime-Dependent Behaviour Predicted in Theory	99
8	Airfoil with Adaptive Dimple Depth	100
8.1	Concept and Motivation	100
8.2	Prototype Setup	100
8.2.1	Actuation and Control System	102
8.2.2	Depth Calibration	103
8.3	Test Matrix	103
8.4	Results	104
9	Discussion & Conclusion	107
9.1	Answering the Research Questions	107
9.2	Contributions	112
9.3	Limitations	112
9.4	Recommendations for Future Work	115
	Data and Code Availability	116
	Acknowledgements	117
	References	118
A	Supportive Figures	121
A.1	SLT Adapter for Mounting	124
A.2	Surface Pressure Tap Model - CAD Render	125
A.3	SLT Airfoil model	125
B	RWT and RANS Result Figures	126
B.1	Wake Rake and Wall Pressure C_L and C_D Screening Heatmaps	127
B.1.1	Wake Rake and Wall Pressure C_L Screening Heatmaps	127
B.1.2	Wake Rake and Wall Pressure C_D Screening Heatmaps	128
B.2	RANS C_L and C_D Screening Heatmaps	130
B.2.1	RANS C_L Screening Heatmaps	130
B.2.2	RANS C_D Screening Heatmaps	131
B.3	RANS Timeseries Results Reynolds number 200k	132
B.4	RANS Timeseries Results Reynolds number 500k	134
C	Surface Pressure Tap Result and Figures	136
C.1	Surface Pressure C_P Figures	137
C.1.1	Reynolds Number 1×10^5	137
C.1.2	Reynolds Number 2×10^5	140
C.1.3	Reynolds Number 5×10^5	143
C.2	Numerical and Experimental Surface Pressure C_P Comparison Figures	146
C.2.1	Reynolds Number 2×10^5	146
C.2.2	Reynolds Number 5×10^5	149
D	PIV, Force Balance and Tuft Results and Figures	152
D.1	SLT C_L and C_D Screening Heatmaps	152
D.1.1	SLT C_L/C_D Screening Heatmaps	152
D.1.2	SLT C_L Screening Heatmaps	153
D.1.3	SLT C_D Screening Heatmaps	154
D.2	PIV Flow Field Figures	155
D.2.1	Reynolds Number 1×10^5	155
D.2.2	Reynolds Number 2×10^5	159
D.3	Tufts Pictures	163

1

Introduction

Modern wind turbines and aircraft operate across a wide envelope of inflow conditions, and their aerodynamic efficiency depends critically on keeping the boundary layer attached to the lifting surface for as long as possible. Flow separation, whether from a gust-induced angle-of-attack excursion on a turbine blade or an aircraft manoeuvring outside its design point, causes an abrupt loss of lift and a sharp rise in drag, degrading performance and, in extreme cases, compromising structural loading and control. A large body of work in applied aerodynamics is therefore devoted to controlling where and how that separation occurs.

Broadly, flow-control devices fall into two categories. Active devices, such as synthetic jets, actuated flaps, or oscillating vortex generators, can be switched on only when needed and tailored to the instantaneous flow state, but they require power, actuation hardware, and added system complexity. Passive devices, such as fixed vortex generators, riblets, or Gurney flaps, need none of this, but by the same token cannot adapt. A passive device is shaped once, at the design stage, for a single assumed operating condition, and is then carried unchanged through every other condition the vehicle or blade will ever encounter. The ordinary aircraft wing flap illustrates why this matters. A flap does not hold one fixed position throughout a flight, it extends to generate the extra lift needed for take-off and landing, then retracts flush against the wing for efficient cruise, where the same deflection would only add drag. If only a single flap position were possible, no one choice would serve every phase of flight well. Whichever position was picked would help in one regime and cost performance in another. A fixed, passive flow-control device faces exactly the same dilemma, without the benefit of being allowed to move.

Among passive devices, the dimple occupies an unusual position. It is best known not from an aircraft or a turbine blade but from the golf ball, where a pattern of small surface indentations trips the boundary layer to turbulence and, by delaying separation on the rear of the ball, roughly halves its drag relative to a smooth sphere. In a cruder, full-scale demonstration of the same idea, the television series *MythBusters* covered a car in golf-ball-style dimples and measured a fuel-economy improvement of around 11 – 12.5% over an otherwise identical smooth surface car as baseline [22]. The appeal of transferring this mechanism to a lifting surface is obvious. A similarly small, unpowered surface modification that could delay separation and improve the lift-to-drag ratio at negligible weight or complexity cost would of course be desirable. Chapter 2 shows, however, that this transfer is far from automatic. A golf ball is a bluff body whose drag is almost entirely pressure drag, so essentially any delay of separation helps it, whereas an airfoil in attached flow already has very little separation to delay, and the same dimples can just as easily add drag as remove it. Existing studies on dimpled airfoils reflect this tension. Reported effects vary widely with dimple depth, size, and especially chordwise location, and different studies vary in their findings on where along the chord dimples should even be placed.

This inconsistency is not, on the present reading, a sign that dimples are unreliable so much as a sign that each study has tested one fixed dimple geometry at one operating condition and reported whatever that particular combination happened to produce. If the benefit of a dimple genuinely depends on the local state of the boundary layer, and that state itself depends on angle of attack, Reynolds number, and pressure gradient, then a dimple fixed at the time of manufacturing can only ever be well matched to a

single point in the operating envelope. A closely related idea has already been demonstrated, not on an airfoil but on the sphere itself: Vilumbrales-Garcia et al. [39] built a morphable sphere whose dimple depth could be actuated pneumatically in real time, and showed that varying the depth shifts the drag crisis and cuts drag by as much as 63% relative to a fixed-depth equivalent. That result, that a dimpled surface which can change its own geometry outperforms any single fixed one, is a direct inspiration for the central idea developed in this thesis: instead of fixing dimple depth and location once, allow them to be varied, so that a lifting surface, not just a sphere, can in principle be re-tuned to remain beneficial as the flow condition changes.

Investigating that idea in full, a dimple array whose depth and chordwise placement both respond in real time to the local flow, is a substantial undertaking, and this thesis approaches it in two stages of very different size. The primary and by far the larger part of the work is a broad, wind-tunnel and CFD campaign, sweeping static chordwise placement of the dimples and dimple depth across three Reynolds numbers and a full angle-of-attack range, to establish where and how deep the dimples should be placed on an airfoil at each operating condition, and to explain, through PIV, surface pressure, and tuft flow visualisation, why the resulting trends occur. The secondary, smaller part of the work uses these static findings to guide the design of a proof-of-concept prototype capable of mechanically varying its own dimple depth, establishing the feasibility of the actuation concept as a basis for future, fully adaptive systems.

Thesis Outline

The remainder of this thesis is structured as follows. Section 1.1 states the research questions that guide the work and the research activities undertaken to answer them. Chapter 2 builds the physical basis for the dimple mechanism, from boundary-layer development and separation through to the golf-ball analogy and its transfer to a lifting surface, and closes with a review of the existing dimpled-airfoil literature and the geometric design parameters adopted for the present campaign. Chapter 3 covers the aerodynamic coefficient measurements carried out at DTU, the wake-rake and wall-pressure measurements and the surface-pressure tap model. Chapter 4 covers the flow visualisation methods carried out at TU Delft, the PIV setup and the tuft flow visualisation, along with the force-balance measurements. Chapter 5 describes the CFD methodology, covering the RANS and LES setups used throughout this work. Chapter 6 validates the smooth baseline against published data, cross-checks the experimental and numerical methods against one another, and confirms the spanwise two-dimensionality of the flow. Chapter 7 presents the main results of the static campaign: the aerodynamic screening across the full geometry matrix, the effect of chordwise dimple placement and dimple depth, a detailed flow-field investigation of the standout configuration, and a comparison between the experimental and numerical methods. Chapter 8 presents the proof-of-concept adaptive dimple prototype, its mechanical and electrical design, and its preliminary results. Finally, Chapter 9 answers the research questions, discusses the contributions and limitations of the present work, and recommends directions for future research.

1.1. Research Questions and Activities

Before the experimental and numerical work, a detailed survey of existing aerodynamic studies on surface dimples was undertaken. The literature review in the following chapter establishes the current state of knowledge, summarises the most relevant findings, and identifies several shortcomings in the existing body of work. These research gaps motivated the formulation of the main research question that guides this thesis and informed the design choices for the experimental prototype.

Gaps in Literature

- No findings on studies investigating *dynamically actuated* dimples applied to lifting surfaces.
- Existing studies on lifting surfaces predominantly consider *static, single-depth* dimples; the influence of systematically *varying dimple depth* within the same flow field has received little attention.
- Prior studies exhibit highly inconsistent findings regarding the *optimal chordwise location* of static dimples on lifting surfaces.
- Few studies on lifting surfaces investigate more than one dimple location simultaneously, leading to limited insight into how *multiple chordwise rows* may interact or combine to affect performance.
- Adaptive or closed-loop surface concepts involving dimples have so far been reported only for bluff bodies, with no clear evidence of similar approaches being applied to airfoil configurations.

Research Questions

The gaps outlined above frame the key uncertainties that remain in understanding how adaptive dimple concepts may be applied to lifting surfaces. To address these gaps and guide both the numerical investigations and the experimental prototype development, the main research question is formulated:

How can the controlled variation of dimple depth and chordwise placement be exploited to alter the aerodynamic characteristics and flow-separation behaviour of a lifting surface across varying operating conditions?

To explore the main research question in a structured manner, the investigation is decomposed into the following sub-questions:

- **SQ1:** How do dimple depth, number of rows, and chordwise location influence lift, drag, stall onset, and flow-separation behaviour under varying steady flow conditions?
- **SQ2:** How does the near-wall flow topology and boundary-layer development differ between the smooth baseline and the standout dimpled configuration identified from the static campaign, as resolved using PIV, and how does this explain the observed aerodynamic gains?
- **SQ3:** How do CFD predictions compare to the experimental measurements in reproducing both the aerodynamic performance of the dimpled configurations and the near-wall flow mechanism responsible for it?
- **SQ4:** How do the aerodynamic characteristics produced by a mechanically adjustable, proof-of-concept dimple-depth prototype, at discrete depth settings, compare to those established for the equivalent static, fixed-depth configurations?

Research Activities

Each sub-question (**SQ**) corresponds to a specific set of research activities (**RA**):

- **RA1:** Conduct a comprehensive wind-tunnel test campaign varying dimple depth, row count, chord location, angle of attack, and Reynolds number.
- **RA2:** Perform PIV measurements comparing the smooth baseline to the standout dimpled configuration identified from RA1, to resolve the near-wall momentum, separation suppression, and wake development responsible for the observed aerodynamic gains.
- **RA3:** Conduct a CFD study on a subset of key configurations to provide numerical comparison against the corresponding experimental measurements.
- **RA4:** Design and manufacture a proof-of-concept prototype with actuated dimples, based on the findings of RA1 and RA2, and test it in the wind tunnel at discrete fixed-depth settings.

2

Theory and Literature Review

Modifying a surface to control the flow over it is one of the oldest ideas in applied aerodynamics. Much of it is *passive*—fixed geometry that requires no power or moving parts—and it takes many forms serving different ends: vortex generators on wings and wind-turbine blades, riblets that reduce turbulent skin friction, trip strips/wires that fix transition, and the various fins, strakes, and spoilers most commonly found on aircraft, cars and drones. Among the less common of these devices is the dimple, best known not from a vehicle, but from the golf ball.

A dimple is a form of passive flow control. It's a fixed geometric modification of the wetted surface intended to alter the development of the boundary layer, the location of laminar-to-turbulent transition, and ultimately the point at which the flow separates. Its best-known aerodynamic application is the golf ball, on which a pattern of dimples reduces the pressure drag of a bluff body by tripping the boundary layer to turbulence and so delaying separation. The transfer of this mechanism to a lifting surface is, however, not straightforward. A golf ball is a bluff body whose drag is almost entirely pressure drag, whereas an airfoil in attached flow is streamlined, and its skin friction, pressure recovery, transition, the laminar separation bubble, and stall interact in a far more delicate balance. Whether a dimple helps or harms an airfoil therefore depends on the local state of the boundary layer where it is placed, and not merely on the established fact that dimples reduce the drag of a sphere.

This chapter assembles the physical basis required to interpret that distinction, and then reviews the existing literature in its light. Section 2.1 introduces the boundary layer and its laminar, transitional, and turbulent states, together with the Reynolds-number dependence that governs them. Section 2.2 describes how an adverse pressure gradient drives flow separation and gives rise to the laminar separation bubble characteristic of the present Reynolds-number range. Section 2.3 treats surface roughness as a means of forcing transition and develops the cavity mechanism by which a single dimple re-energises the near-wall flow. Section 2.4 examines that mechanism in its cleanest setting, the dimpled sphere, before Section 2.5 considers how it transfers to a lifting surface, including the geometric parameters that govern the effect and its expected variation across the operating envelope. Section 2.6 draws these threads together into the motivation for controlling dimple depth adaptively and surveys the experimental and numerical studies of dimpled airfoils that define the research gap this thesis addresses. Section 2.7 then turns from theory to design, establishing the wind-tunnel constraints, airfoil section, chord length, and dimple geometry adopted for the present campaign, and Section 2.7.4 consolidates the resulting choices into the final set of test geometries carried forward into the experimental and numerical work.

2.1. Boundary Layers

When a viscous fluid flows over a solid surface, the no-slip condition requires the fluid in contact with the wall to be at rest relative to it. Away from the wall the flow is effectively unaffected and behaves as if inviscid, so the influence of viscosity is confined to a thin region adjacent to the surface, in which the streamwise velocity rises from zero at the wall to the local freestream velocity. This region is the *boundary layer*, and its identification by Prandtl is what allows a high-Reynolds-number flow to be treated as an inviscid outer flow coupled to a thin viscous layer near the surface [18].

The structure of the boundary layer is shown in Figure 2.1. The streamwise velocity u approaches the freestream velocity U_∞ asymptotically, so the boundary-layer thickness δ is defined by convention as the wall-normal distance at which u reaches 99% of U_∞ . Although drawn here on a flat surface, the same description applies to the curved suction surface of an airfoil. The essential addition there is a streamwise pressure gradient, which governs whether the layer remains attached and is further touched upon in Section 2.2. The shape of the near-wall velocity profile, shown in Figure 2.1, is the quantity of central interest in this work, a profile that carries a large velocity gradient at the wall is both a source of friction and an indicator of how much near-wall momentum the layer holds, and it is this momentum that determines the layer's ability to resist separation.




Figure 2.1: The boundary layer adjacent to a solid surface. The streamwise velocity rises from zero at the wall, in accordance with the no-slip condition, and approaches the external velocity U_∞ asymptotically; the thickness δ is defined as the wall-normal distance at which $u = 0.99 U_\infty$.

2.1.1. Laminar and Turbulent Boundary Layers

A boundary layer exists in one of two fundamentally different states, as illustrated in Figure 2.2. In a *laminar* boundary layer the fluid moves in smooth, ordered layers, and momentum is exchanged in the wall-normal direction only by molecular viscosity. The resulting profile is comparatively thin and carries relatively little momentum close to the wall, which makes a laminar layer efficient in terms of skin friction but poorly equipped to resist a rising pressure, which therefore separates readily once the pressure begins to recover [18].

A *turbulent* boundary layer, by contrast, contains unsteady three-dimensional velocity fluctuations superimposed on the mean flow. Decomposing the instantaneous velocity components into mean and fluctuating parts following the Reynolds decomposition,

$$u = \bar{u} + u', \quad v = \bar{v} + v', \quad (2.1)$$

where an overbar denotes a time average and a prime the fluctuation about it. These fluctuations transport fluid in the wall-normal direction, continuously exchanging low-momentum fluid near the wall for high-momentum fluid from the outer part of the layer. The result is a markedly fuller near-wall velocity profile, which is precisely why a turbulent boundary layer can withstand a steeper adverse pressure gradient before separating, and hence why it separates later than a laminar layer under the same conditions. However, this benefit is not without cost. The same near-wall mixing steepens the velocity gradient at the surface (as depicted in Figure 2.2), so a turbulent layer also produces a higher wall shear stress and greater skin-friction drag than a laminar one. The wall shear stress that quantifies this difference is explained in the following section.




Figure 2.2: Comparison of the velocity profiles of a laminar (left) and a turbulent (right) boundary layer at the same edge velocity. The turbulent profile is fuller near the wall, carrying more near-wall momentum and a steeper wall gradient. It is therefore more resistant to separation but produces a higher wall shear stress.

2.1.2. Wall Shear Stress and Skin Friction

The friction the boundary layer exerts on the surface follows directly from the shape of the near-wall velocity profile. Because the streamwise velocity rises from zero at the wall, there is a velocity gradient normal to the surface, and the viscous shear stress it produces at the wall is the *wall shear stress* (Figure 2.3)

$$\tau_w = \mu \left. \frac{\partial u}{\partial y} \right|_{y=0}, \quad (2.2)$$

where τ_w is the wall shear stress [Pa], μ is the dynamic viscosity [Pa s], u is the streamwise velocity [m/s], and y is the wall-normal distance [m]. The wall shear stress is therefore proportional to the slope of the velocity profile at the wall. The fuller turbulent profile of Figure 2.2 has the steeper wall gradient and so the larger τ_w , which is the origin of the higher skin-friction, as explained in Section 2.1.1.

Integrating τ_w over the airfoil surface and resolving in the freestream direction gives the skin-friction component of drag, while the local value is conventionally expressed through the skin-friction coefficient

$$C_f = \frac{\tau_w}{\frac{1}{2} \rho U_\infty^2}, \quad (2.3)$$

where ρ is the fluid density [kg/m³] and U_∞ is the freestream velocity [m/s]. The sign of C_f carries additional information. Where the flow is attached, the wall gradient is positive and $C_f > 0$, whereas a vanishing or negative C_f marks the points at which the flow separates and reattaches, which is further touched upon in Section 2.2.




Figure 2.3: The shear stress τ arising from the velocity gradient du/dy within the boundary layer. The streamwise velocity rises from zero at the wall to the freestream value U_∞ , and the slope of the profile sets the local shear stress; evaluated at the wall, this gradient gives the wall shear stress of Equation (2.2).

Equation (2.2) describes the purely viscous stress that acts at the wall, where the no-slip condition forces the turbulent fluctuations themselves to vanish. Away from the wall, however, the wall-normal mixing introduced in Section 2.1.1 contributes an additional apparent stress. Decomposing the instantaneous velocity components into mean and fluctuating parts following the Reynolds decomposition from Equation 2.1, the wall-normal transport of streamwise momentum gives rise to the turbulent Reynolds shear stress

$$\tau_t = -\rho \overline{u'v'}, \quad (2.4)$$

so that the total shear stress at any height y within a turbulent boundary layer is the sum of a viscous and a turbulent contribution,

$$\tau(y) = \mu \frac{\partial \bar{u}}{\partial y} - \rho \overline{u'v'}. \quad (2.5)$$

The wall shear stress of Equation 2.2 is recovered as the special case $y = 0$, where the no-slip condition forces the fluctuations, and with them the Reynolds stress of Equation 2.4, to vanish. It is this turbulent mixing that fills out the mean profile and so steepens the wall gradient in Equation (2.2). The Reynolds stress thus raises the wall shear stress indirectly, through the fuller mean profile, even though it vanishes at the wall itself [18].

2.1.3. Boundary-Layer Development and Transition

The laminar and turbulent states of Section 2.1.1 are not fixed properties of a surface but successive stages through which the boundary layer passes as it develops downstream. At the leading edge the layer is vanishingly thin and laminar. As the flow proceeds, viscosity slows an increasing quantity of fluid, and the boundary-layer thickness grows with streamwise distance as $\delta(x)$, the wall-normal velocity profile $u(y)$ also thickens as it develops. This growth is shown in Figure 2.4. As shown, beyond a critical location x_{cr} the laminar layer becomes unstable and, through a short transition region, breaks down into the turbulent state, which is thicker and, as established in Section 2.1.1, fuller near the wall. This turbulent layer also develops an inner structure of its own, comprising a thin viscous sublayer and a buffer layer beneath the fully turbulent outer region, examined in Section 2.1.4.




Figure 2.4: Development of the boundary layer through its laminar, transitional, and turbulent states. The laminar layer grows smoothly downstream of the leading edge until, beyond the critical location x_{cr} , it becomes unstable and passes through a transition region into the turbulent state. The turbulent layer is thicker and carries the fuller near-wall profile of Section 2.1.1. Here its inner region comprises a viscous sublayer and a buffer layer beneath the fully turbulent outer flow.

The change from the laminar to the turbulent state is the process of *transition*, and the location at which it occurs is not set by geometry alone. On a smooth surface it is governed by a Reynolds-number criterion. Once a critical value is exceeded, small disturbances within the laminar layer are amplified rather than damped until the flow breaks down. This criterion, and the critical Reynolds number that expresses it, is introduced in Section 2.1.4.

2.1.4. Reynolds Number and Critical Reynolds Number

Whether a boundary layer is laminar or turbulent, and where along a surface it changes from one to the other, is governed by the relative importance of inertial and viscous forces in the flow. This ratio is expressed by the Reynolds number. For an airfoil the relevant form is the chord Reynolds number, where the length used is the chord,

$$Re = \frac{U_\infty c}{\nu}, \quad (2.6)$$

where Re is the chord Reynolds number [–], U_∞ is the freestream velocity [m/s], c is the airfoil chord [m], and ν is the kinematic viscosity of the fluid [m²/s]. A large Reynolds number indicates that inertial forces dominate, so disturbances within the flow are readily amplified. A small Reynolds number indicates that viscosity dominates, damping disturbances and keeping the flow ordered. The Reynolds number thereby governs the thickness of the boundary layer, the location of transition, and the susceptibility of the flow to separation [18].

The transition introduced in Section 2.1.3 can now be further explained. As the boundary layer develops downstream, the streamwise distance x from the leading edge acts as the characteristic length, and transition sets in once the Reynolds number based on this distance reaches a *critical* value at the transition location x_{cr} ,

$$Re_{cr} = \frac{U_\infty x_{cr}}{\nu}, \quad (2.7)$$

Transition does not occur gradually everywhere but sets in once this local Reynolds number reaches a *critical* value, Re_{cr} , at which the laminar layer first becomes unstable to small disturbances. Upstream of the corresponding location x_{cr} the layer is laminar, and downstream of it the disturbances are amplified until the flow breaks down into turbulence. The critical Reynolds number is therefore the quantitative counterpart of the transition location x_{cr} marked in Figure 2.4.

The critical value is not a universal constant. It depends on the freestream turbulence level, the surface roughness, and on an airfoil, the streamwise pressure gradient, all of which can promote or delay the onset of instability and so move the transition point [18]. This sensitivity is what makes transition controllable, a deliberate surface disturbance can lower the effective critical Reynolds number and force transition earlier than it would occur on a smooth surface, a possibility further explored in Section 2.3.

Beyond governing transition, the Reynolds number also organises the internal structure of a turbulent boundary layer. Immediately adjacent to the wall, the appropriate measure is not the chord Reynolds number but a local one based on the wall-normal distance, the dimensionless wall coordinate

$$y^+ = \frac{y u_\tau}{\nu}, \quad u_\tau = \sqrt{\frac{\tau_w}{\rho}}, \quad (2.8)$$

where u_τ is the friction velocity [m/s] and τ_w is the wall shear stress of Equation (2.2). According to the value of y^+ , the turbulent boundary layer divides into an inner region governed by viscosity and an outer region governed by turbulent mixing, the inner region itself comprising the two thin sublayers below.

Viscous sublayer

The viscous sublayer is the thin layer in direct contact with the wall, occupying roughly $y^+ < 5$. Here the turbulent fluctuations are suppressed by the wall, the Reynolds stress vanishes at the surface, so viscous shear dominates, and the velocity varies almost linearly with distance from the wall. Though very thin, it is the region that sets the wall shear stress, and hence the skin friction, of the whole layer.

Buffer layer

Above the viscous sublayer lies the buffer layer, roughly $5 < y^+ < 30$, in which the viscous and turbulent shear stresses are of comparable magnitude and neither dominates. It is the area of the most intense turbulence production in the layer, and it blends the near-linear profile of the viscous sublayer into the logarithmic profile of the fully turbulent region beyond $y^+ \approx 30$.

2.1.5. Integral Thicknesses and the Shape Factor

The boundary-layer thickness δ defined in Section 2.1 is a convenient but somewhat arbitrary measure, since the velocity approaches the freestream value U_∞ asymptotically and the 99% cut-off is a convention rather than a physical boundary. Two integral measures characterise the layer more precisely by quantifying the deficits the slowed near-wall flow produces relative to an equivalent inviscid flow [18].

Displacement Thickness

The *displacement thickness* δ^* is the distance by which the wall would have to be displaced outward, in an equivalent inviscid flow, to account for the reduction in mass flow caused by the boundary layer,

$$\delta^* = \int_0^\delta \left(1 - \frac{u}{U_\infty}\right) dy, \quad (2.9)$$

where u is the streamwise velocity [m/s], U_∞ is the freestream velocity [m/s], and y is the wall-normal distance [m]. A small δ^* corresponds to a thin layer that disturbs the outer flow little. A large δ^* indicates a thicker layer that displaces a greater volume of fluid away from the surface.

Momentum Thickness

The *momentum thickness* θ measures the corresponding deficit in momentum flux, and is the thickness of an equivalent layer of freestream fluid carrying the momentum lost to wall friction,

$$\theta = \int_0^\delta \frac{u}{U_\infty} \left(1 - \frac{u}{U_\infty}\right) dy. \quad (2.10)$$

The momentum thickness is directly related to the skin-friction drag developed in Section 2.1.2, since the streamwise momentum removed from the flow is the drag exerted on the surface.

Shape Factor

The ratio of the two integral thicknesses defines the *shape factor*,

$$H = \frac{\delta^*}{\theta}, \quad (2.11)$$

which characterises the fullness of the velocity profile independently of the layer's absolute size. A full, momentum-rich turbulent profile has a low shape factor, on the order of 1.5, whereas a laminar profile has a higher value, on the order of 2.5 [18]. The shape factor is therefore a good indicator of the layer's state and, more usefully for the present work, of its proximity to separation. As the near-wall profile is drained of momentum under an adverse pressure gradient, it becomes progressively less full, H rises, and a sufficiently large value signals that separation is imminent.

2.2. Adverse Pressure Gradients and Flow Separation

The state of the boundary layer established in Section 2.1 determines how well it resists separation, but whether it actually separates depends on the pressure it must flow against. On the suction side of an airfoil the external flow is first accelerated as it passes around the leading edge and over the forward part of the chord, reaching a velocity maximum near the suction peak. Beyond this point it must decelerate as the surface pressure recovers towards the trailing edge. The forward, accelerating region is one of *favourable* pressure gradient, in which the static pressure falls in the streamwise direction, while the aft, decelerating region is one of *adverse* pressure gradient, in which it rises, as shown in Figure 2.5. It is this adverse-pressure-gradient region that governs whether the flow remains attached, and it is where boundary-layer control devices such as dimples are expected to act.




Figure 2.5: Surface-pressure distribution over an airfoil, shown as arrows normal to the surface (outward for suction, inward for positive pressure). Over the leading edge on the suction side, the flow accelerates, and the pressure falls, giving a favourable pressure gradient. Aft of the suction peak the flow decelerates, and the pressure recovers, giving rise to an adverse pressure gradient. Separation is most likely to begin in this aft recovery region.

2.2.1. Adverse Pressure Gradient

The pressure the boundary layer experiences is imposed on it by the external, effectively inviscid outer flow at its edge. Along the surface this edge flow obeys the inviscid (Euler) momentum balance,

$$\frac{dp_e}{dx} = -\rho U_e \frac{dU_e}{dx}, \quad (2.12)$$

where p_e is the static pressure at the boundary-layer edge [Pa], x is the streamwise coordinate [m], ρ is the fluid density [kg/m³], and U_e is the local velocity at the edge of the boundary layer [m/s], which, unlike the constant freestream U_∞ , varies along the surface. Where the edge flow decelerates ($dU_e/dx < 0$) the pressure rises in the streamwise direction ($dp_e/dx > 0$), which is the defining condition of an *adverse* pressure gradient. And vice versa, where it accelerates, the gradient is *favourable*.

The distribution of surface pressure is conveniently expressed through the pressure coefficient,

$$C_p = \frac{p - p_\infty}{\frac{1}{2}\rho U_\infty^2}, \quad (2.13)$$

where C_p is the pressure coefficient [–], p is the local static pressure [Pa], p_∞ is the freestream static pressure [Pa], and $\frac{1}{2}\rho U_\infty^2$ is the freestream dynamic pressure [Pa]. A pronounced suction peak followed by a steep recovery of C_p towards the trailing edge is the signature of a strong adverse pressure gradient on the suction side, the steeper the recovery, the more severe the deceleration the boundary layer must withstand, and the more readily it will separate [4].

The effect of this rising pressure on the boundary layer itself is shown in Figure 2.6, which follows the velocity profile through four streamwise stations under an adverse pressure gradient. At station A the profile is comparatively full and the layer is firmly attached. At station B the near-wall fluid has been decelerated, the profile is less full and develops an inflection, yet the wall shear stress remains positive ($\tau_w > 0$) and the flow stays attached. At station C the near-wall velocity gradient, and with it the wall shear stress, goes to zero, identifying the point of separation. At station D the near-wall flow has reversed ($\tau_w < 0$), forming a recirculating region beneath the separation streamline. The mechanism by which this progression ends in separation is the subject of the following section.




Figure 2.6: Evolution of the boundary-layer velocity profile through four streamwise stations (A–D) under an adverse pressure gradient ($dp/dx > 0$). The profile is full and attached at A, becomes progressively less full and inflected at B ($\tau_w > 0$), reaches the separation point at C where the wall shear stress vanishes ($\tau_w = 0$), and exhibits reversed near-wall flow at D ($\tau_w < 0$). The lower panel maps the same stations onto the suction surface of an airfoil. Figure adapted from Kaushik [14].

2.2.2. Flow Separation

Separation is not caused by low velocity alone but by a deficit of near-wall momentum. The fluid within the boundary layer has already been slowed by wall friction, so it carries less streamwise momentum than the external flow. When this near-wall fluid with low-momentum is forced to move against the rising pressure of an adverse gradient it is decelerated further, and if its momentum is insufficient to overcome the gradient it is brought to rest and then reversed [18]. The point at which the near-wall velocity gradient, and with it the wall shear stress, first vanishes,


$$\left. \frac{\partial u}{\partial y} \right|_{y=0} = 0 \implies \tau_w = 0, \quad (2.14)$$

defines the point of separation, as shown at station *C* in Figure 2.6, downstream of it the wall shear stress is negative, and the near-wall flow recirculates beneath the dividing streamline as it lifts away from the surface, as traced through stations A–D in Figure 2.6.


On an airfoil, the size of the resulting separated region grows with angle of attack, as shown in Figure 2.7. At a low angle of attack, the flow remains attached over most of the chord and separates only near the trailing edge, leaving a thin wake. As the angle of attack increases, the adverse pressure gradient strengthens, the separation point moves forward, and the recirculating region enlarges. At the stall angle, the flow separates well forward on the suction surface, leaving a large region of reversed flow and a broad turbulent wake over much of the chord, with an accompanying loss of lift and rise in drag. The aft, pressure-recovery region of the suction surface is therefore the region most vulnerable to separation, and it is here that boundary-layer control devices such as dimples are intended to act.

At the low-to-moderate Reynolds numbers of the present work, however, the boundary layer is typically still laminar when it first separates. The detached shear layer may then transition to turbulence and reattach a short distance downstream, enclosing a region of recirculating fluid known as the laminar separation bubble. This is the subject of the following section.

(a) Low angle of attack



(b) High angle of attack



(c) Stall angle




Figure 2.7: Development of flow separation on an airfoil with increasing angle of attack. At a low angle of attack (a) the flow separates only near the trailing edge, leaving a thin wake. At a higher angle (b) the separation point moves forward and a region of reversed flow forms. At the stall angle (c) the flow separates well forward on the suction surface, producing large-scale separation and a broad turbulent wake.

2.2.3. The Laminar Separation Bubble

At the low-to-moderate Reynolds numbers considered here ($Re \approx \mathcal{O}(10^5)$), the boundary layer on the forward suction surface is typically laminar when it meets the adverse pressure gradient, and being poorly equipped to resist it, separates. The detached shear layer is highly unstable and once free of the wall it transitions to turbulence rapidly, carrying the fuller profile of Section 2.1.1, the now-turbulent flow then reattaches a short distance downstream. The recirculating region enclosed between the laminar separation point, the transitioning shear layer, and the turbulent reattachment is the *laminar separation bubble* (LSB) [12], illustrated in Figure 2.8. Its signature in the surface pressure is a plateau in C_p over the separated region, followed by a sharp recovery at reattachment, as shown in Figure 2.9.




Figure 2.8: Schematic of a laminar separation bubble on the suction surface of an airfoil. The laminar layer separates under the adverse pressure gradient, and the detached shear layer transitions and reattaches downstream, enclosing a region of recirculating fluid beneath the separation streamline.

LSB's are classified as *short* or *long*. A short bubble occupies a small fraction of the chord and can even be beneficial, fixing transition and leaving a reattached turbulent layer that better resists further separation [12]. A long bubble extends over much of the chord and can carry a measurable loss of lift and increase in drag (Figure 2.9). The bubble is sensitive to angle of attack and Reynolds number, as the angle of attack rises the adverse gradient strengthens until, beyond a critical condition, the bubble fails to reattach—bubble *bursting*—producing an abrupt leading-edge stall [12]. The bubble is therefore an active participant in the stall behaviour at these Reynolds numbers, not a passive feature.




Figure 2.9: A short and a long LSB and their surface-pressure signatures. The short bubble produces a localised C_p plateau, whereas the long bubble produces an extended plateau usually associated with loss of lift and increase in drag.

Two features of the bubble matter here. First, it is a marginal, sensitive flow state, exactly the kind a localised surface modification can influence. A dimple placed in or upstream of the bubble may promote earlier transition, shorten it, and move reattachment upstream. Binci et al. observed precisely this on a dimpled NACA 64-014A laminar airfoil at $Re = 1.75 \times 10^5$, where the dimples measurably reduced the bubble's extent [8]. Second, the bubble is difficult to predict numerically, resolving laminar separation, shear-layer transition, and turbulent reattachment demands a faithful transition model, which standard fully turbulent RANS closures lack. Binci et al. found that RANS overestimated the transition and reattachment locations, whereas LES agreed far more closely with experiment [8].

2.3. Surface Roughness and the Dimple Mechanism

The earlier sections established that the state of the boundary layer governs its resistance to separation, and that at low Reynolds numbers the laminar layer is prone to separation and to the formation of a laminar separation bubble. A natural strategy for influencing this behaviour is to force the boundary layer into the more separation-resistant turbulent state earlier than it would transition naturally. Surface roughness is the classical means of doing so. A sufficiently large surface disturbance destabilises the laminar layer and trips it to turbulence [27], as illustrated in Figure 2.10. A dimple may be regarded as one particular form of such roughness.




Figure 2.10: Surface roughness tripping a laminar boundary layer to turbulence. A disturbance of sufficient height causes near-wall mixing that carries the layer through transition into the fuller turbulent profile downstream.

2.3.1. Roughness-Induced Transition

The ability of a roughness element to provoke transition is conventionally characterised by the roughness Reynolds number,

$$Re_k = \frac{U_k k}{\nu}, \quad (2.15)$$

where Re_k is the roughness Reynolds number [–], k is the roughness height [m], U_k is the undisturbed local velocity at that height [m/s], and ν is the kinematic viscosity [m²/s]. Below a threshold value, the element is submerged in the near-wall flow and has little effect; above it, the element generates disturbances strong enough to trip the layer to turbulence.

This framework must be applied to dimples with care. The roughness Reynolds number was developed for *protruding* elements such as sand grains or trip wires, for which k is the height by which the element projects into the flow. For an inward dimple k is instead a cavity depth, a depression rather than a protrusion, which does not shed vorticity into the boundary layer in the manner of a discrete particle. Equation (2.15) therefore serves only as an indicative scaling and does not capture the governing physics of a cavity. A dimple must thus be large enough relative to the local viscous scale to influence the flow, yet not so large as to become a source of excessive separation or form drag.

It is worth distinguishing dimples from the passive devices they are often compared with. A vortex generator is a protruding vane that sheds a streamwise vortex to transport high-momentum fluid towards the wall. A trip strip is a protrusion whose sole purpose is to force transition. Both add a protruding element to the surface. A dimple reaches a similar effect, energising the near-wall flow, through the opposite geometric means, a cavity that induces a local separation and turbulent reattachment. The vortex-generator analogy is useful for intuition, since both ultimately produce a fuller, more separation-resistant profile. Still, the mechanisms differ and a dimple introduces no protrusion into the external flow.

2.3.2. Dimple-Induced Local Separation and Reattachment

The dimple-induced local flow separation and reattachment mechanism is best understood at the scale of a single cavity. As the boundary layer reaches a dimple, the flow separates at the upstream rim where the surface falls away, forming a free shear layer across the mouth of the cavity, beneath which the fluid recirculates. This shear layer is unstable, just as in the laminar separation bubble of Section 2.2.3. However, now triggered deliberately and locally by the dimple, the instability amplifies, the turbulence intensity rises, and the flow reattaches a short distance downstream, carrying substantially higher near-wall momentum. Each dimple thus generates a small, controlled separation bubble whose reattached flow is fuller and more energetic than the layer that entered it [9], as depicted in Figure 2.11

This is precisely the momentum-transfer mechanism of Section 2.1.1, the turbulence generated in the shear layer enhances the wall-normal exchange of momentum, producing a fuller, more separation-resistant near-wall profile. The re-energised layer leaving the dimple better withstands the downstream adverse pressure gradient, delaying the global separation point. Choi et al., in detailed measurements over a dimpled sphere, resolved exactly this sequence, local separation, a rapid rise in shear-layer fluctuations, reattachment with high near-wall momentum, and a consequent delay of the main separation, and concluded that generating such a local separation bubble is itself an effective flow-control strategy [9]. The cavity-scale flow is shown in Figure 2.11.

Crucially, this mechanism is conditional. It benefits the flow only where the downstream boundary layer is in danger of separating. Where the flow is already attached and the adverse pressure gradient is weak, for example on the forward part of an airfoil at low angle of attack, there is no separation to delay, and the dimple acts purely as a surface disturbance. It then thickens the local layer, raises skin-friction drag through the same near-wall mixing that would otherwise be beneficial, adds a cavity pressure loss, and slightly reduces the effective camber, for a net drag penalty with no compensating gain. The same dimple placed in or upstream of a near-separation region may instead delay separation and improve performance.

The mechanism is the same whether the body is a sphere or an airfoil, but the context in which it acts, and therefore the benefit it produces, differs substantially between the two. The next section examines the canonical bluff-body case, where the mechanism is most clearly established, before considering how it transfers to a lifting surface.




Figure 2.11: Schematic of the flow over a single dimple. The incoming boundary layer separates at the upstream rim and forms a free shear layer across the cavity, beneath which the flow recirculates. The shear layer becomes unstable and transitions, and the reattached flow leaves the dimple with a fuller, higher-momentum near-wall profile that better resists the downstream adverse pressure gradient. Adapted from the mechanism described by Choi et al. [9].

2.4. Bluff-Body Drag Reduction: Dimples on a Golf Ball

The clearest demonstration of the dimple mechanism is the sphere, and in particular the golf ball. A sphere is a bluff body, whose drag is dominated not by skin friction but by the pressure imbalance between a high-pressure stagnation region at the front and a low-pressure separated wake at the rear. It is expressed through the drag coefficient

$$C_D = \frac{D_{total}}{\frac{1}{2}\rho U_\infty^2 A}, \quad (2.16)$$

where C_D is the drag coefficient [–], D_{total} is the total drag force [N], ρ is the fluid density [kg/m³], U_∞ is the freestream velocity [m/s], and A is the frontal reference area [m²]; for a sphere of diameter d , $A = \pi d^2/4$.

For a smooth sphere in the subcritical regime the boundary layer is laminar when it separates, and it does so early, well before its equator, producing a wide wake and a high pressure drag (Figure 2.12, top). As the Reynolds number is raised the boundary layer eventually transitions before separation, the turbulent layer, carrying the fuller near-wall profile of Section 2.1.1, stays attached further around the rear, the wake narrows, and the drag coefficient falls abruptly. This is the *drag crisis* [1]. Dimples lower the Reynolds number at which it occurs: by tripping the boundary layer they bring the transition, and the associated delay of separation, forward to lower speeds than a smooth sphere would require (Figure 2.12, bottom). This is the basis of golf-ball aerodynamics, in which the dimpled ball achieves substantially lower drag, and hence greater range, than a smooth ball over the relevant speed range [7].

The detailed mechanism was resolved by Choi et al. and is exactly the local separation–instability–reattachment process of Section 2.3.2. The flow separates locally at the dimples, the separating shear layer becomes unstable and turbulent, and the reattached high-momentum flow overcomes the strong adverse pressure gradient over the rear, delaying the main separation to beyond roughly 110° from the front stagnation point [9]. This yields more than a 50% drag reduction relative to the smooth sphere at $Re \approx 0.9 \times 10^5$, and because the reattachment, and hence the separation angle, is thereafter fixed, the drag coefficient remains low and nearly constant across the tested range $0.5\text{--}2.8 \times 10^5$ [9]. Later morphable-sphere experiments confirm reductions of the same order. Vilumbrales-Garcia et al. reports a drag coefficient of $C_D = 0.25$ at $Re = 1.3 \times 10^5$ for a shallow fixed dimple depth, again about a 50% reduction relative to the smooth sphere [39].




Figure 2.12: Flow over a smooth sphere (top) and a dimpled golf ball (bottom) at the same Reynolds number. On the smooth sphere the laminar boundary layer separates early, leaving a thick wake and high pressure drag. On the dimpled ball the dimples trip the layer to turbulence. The fuller turbulent profile stays attached further around the rear, moving the separation point downstream and leaving a thinner wake (the same separation delay described in Section 2.3.2, here on a bluff body).

2.5. Transfer of the Dimple Mechanism to Airfoils

The bluff-body picture of Section 2.4 must be transferred to an airfoil with care, and the distinction is more than one of degree. On a sphere the dimple's task is to provoke the drag crisis, moving a single dominant separation line rearward on a body whose drag is almost entirely pressure drag, where essentially any delay of separation translates into a large drag reduction. On a streamlined airfoil at low angle of attack the flow is largely attached, the pressure drag is small, and there is little separated wake to suppress. The dimple then has no large pressure-drag penalty to recover and, as established in Section 2.3.2, may simply add drag. The regime in which a dimple can help an airfoil is therefore not the whole operating range, as it effectively is for a sphere, but a narrower set of conditions, near transition, within a laminar separation bubble, or approaching stall, where the suction-side boundary layer is genuinely close to separating.

On an airfoil, a dimple array may act through three related but distinct mechanisms. The first is *transition promotion*: where the suction-side layer would otherwise remain laminar into the adverse-pressure-gradient region, the dimples force an early transition, so the fuller turbulent profile of Section 2.1.1 reaches the pressure recovery already equipped to resist separation. The second is *modification of the laminar separation bubble*: where a bubble is present, dimples within or upstream of it destabilise the detached shear layer, promote earlier transition, and move reattachment upstream, shortening the bubble and improving pressure recovery (Section 2.2.3). The third is *momentum mixing ahead of a turbulent separation*: dimples just upstream of where a turbulent layer is about to separate enhance the near-wall momentum exchange and delay that separation. Which mechanism dominates depends on the state of the baseline boundary layer where the dimples are applied, and the three are not mutually exclusive.

This dependence on the baseline flow shows up clearly in the literature. Binci et al. studied a dimpled NACA 64-014A airfoil, chosen because it reliably forms a laminar separation bubble, and found that the dimples reduced the bubble's extent [8]. This is the second mechanism working as intended, but it only worked because the dimples were placed where the baseline flow actually had a bubble to act on. The same dimples on a section with no bubble, or at a point where the flow was firmly attached, may

have done nothing useful.

The chordwise position of the array is therefore not a free geometric choice but a flow-dependent design variable. Placed too far forward, in a favourable or mild gradient, a dimple trips the flow to turbulence too early and saddles the rest of the chord with turbulent skin friction while delaying no separation in return. Placed too far back, after the flow has already separated, it has little influence over a flow that has already left the surface. The useful position sits between these limits: within, or just upstream of, the region where the airfoil first develops a strong adverse gradient, a laminar separation bubble, or an incipient separation. That region is not fixed, though, as the airfoil is loaded up, the suction peak sharpens and the separation point moves forward, so the best place for a dimple most likely shifts with angle of attack and Reynolds number.

This has two consequences that frame the rest of the chapter. First, it explains why studies report beneficial dimple locations scattered across a wide range of the chord with little agreement. Each study fixes one geometry and tests it, and the location that works is simply the one that happens to line up with that study's separation-prone region. The disagreement is what one would expect from testing a condition-dependent device at a single condition. Second, it motivates the central idea of this work; if the depth and position of the dimples could change with the operating condition instead of being fixed at manufacture, a dimpled surface might stay effective across a range of conditions rather than only the one it was designed for.

2.5.1. Dimple Geometry as Design Parameters

The mechanism set out in the preceding sections depends not only on *where* a dimple is placed but on its geometry. Four parameters generally characterise a dimple array: the dimple diameter D_d , the dimple depth k , the centre-to-centre spacing D_s , and the orientation of the cavity relative to the surface, which may be inward (a depression) or outward (a protrusion). To allow comparison between studies that use airfoils and bodies of different sizes, these are expressed in non-dimensional form,

$$\frac{D_d}{c}, \quad \frac{k}{c}, \quad \frac{k}{D_d}, \quad \frac{D_s}{D_d}, \quad (2.17)$$

where c is the airfoil chord [m]. The geometric definitions of these quantities are given in Section 2.7.4 (Figure 2.14). The present section sets out what each ratio controls physically.

The two depth ratios describe different things. The ratio k/D_d is the depth of the cavity relative to its own opening, and governs the *local* cavity flow. It sets how steeply the surface falls away at the upstream rim and therefore how strongly the cavity provokes the local separation and shear-layer instability that drive the mechanism of Section 2.3.2. The ratio k/c , by contrast, measures the depth relative to the global scale of the airfoil and indicates how large a geometric disturbance the dimple represents to the section as a whole. The diameter ratio D_d/c fixes the in-plane size of each cavity relative to the chord, and the spacing ratio D_s/D_d determines how closely the dimples are packed. A small spacing produces a near-continuous corrugated surface in which the vortical structures of neighbouring cavities interact strongly, whereas a large spacing leaves the dimples effectively isolated.

Of these, the depth is the parameter of greatest interest to the present work, because it is the quantity that the adaptive concept varies. Its effect is non-monotonic. A dimple that is too shallow does not perturb the boundary layer strongly enough to trigger the local separation and reattachment on which the mechanism depends, and so produces little benefit. A dimple that is too deep ceases to behave as a shallow cavity that re-energises the flow and instead acts as a surface cut-out. The local separation it induces no longer reattaches cleanly, the cavity becomes a source of pressure loss in its own right, and the net effect on performance is adverse. Soundharya et al., for example, examined relatively deep semi-spherical cavities whose depth approached a significant fraction of the airfoil thickness and found that such large cut-out-like cavities degraded performance over most of the operating envelope [32]. Between these extremes there is an intermediate depth at which the benefit is greatest. The adaptive-sphere experiments make this explicit; Vilumbrales-Garcia et al. observed that increasing the dimple depth reduced drag up to a critical value, beyond which further increase in depth caused the drag to rise again, and Terwagne et al. similarly demonstrated that the drag of a morphable dimpled surface depends strongly and non-monotonically on the actuated cavity depth [39, 37]. The existence of an optimal depth, rather than a monotonic "deeper is better" relationship, is the geometric foundation of the adaptive concept developed in this thesis: if the optimum depends on the operating condition, a fixed depth cannot be optimal everywhere.

2.5.2. Expected Effects Across Airfoil Operating Regimes

The mechanism and its dependencies developed in the preceding sections can now be drawn together into a single predictive picture. Because the dimple effect is conditional, beneficial only where the boundary layer is close to separating, and detrimental where the flow is already attached, its expected influence varies systematically with the operating regime of the airfoil. Table 2.1 summarises these expectations across the regimes relevant to the present work, from fully attached flow at low incidence through to deep stall.

Table 2.1: Expected influence of suction-side dimples on an airfoil under different flow regimes, following from the mechanism developed in Sections 2.3.2–2.5.1.

Flow regime	Expected dimple effect
Attached flow at low angle of attack	Limited benefit. With no separation to delay, the dimples act as surface roughness, increasing skin-friction drag and adding a local cavity pressure loss while slightly reducing effective camber. A net drag penalty is expected.
Transitional flow or laminar separation bubble regime	Potentially beneficial. Dimples may promote transition in the detached shear layer, shorten the laminar separation bubble, and move reattachment upstream, improving pressure recovery where the baseline bubble is detrimental.
Pre-stall adverse-pressure-gradient regime	Potentially beneficial. By energising the near-wall flow ahead of an impending turbulent separation, dimples may delay that separation, increasing lift, reducing pressure drag, or improving the lift-to-drag ratio.
Deep stall or massively separated flow	Limited authority. The flow is dominated by large-scale separated structures rather than by local boundary-layer development, so small surface cavities are expected to have little influence.

This regime-dependence is the key to reading the airfoil-dimple literature surveyed next. Reported gains in lift, drag, or stall angle are not universal properties of dimples but conditional outcomes, holding only for a particular combination of Reynolds number, baseline transition state, pressure-gradient distribution, and dimple geometry and placement. A configuration that helps in one regime may be neutral or harmful in another, and a result at one operating point need not carry over to another. This is what makes the literature difficult to compare. Each study reports a single geometry at a handful of conditions, and the apparent disagreements between them are, on the present reading, the expected consequence of testing a condition-dependent device at different conditions. The following section reviews that body of experimental and numerical work with this in mind, to establish what is reliably known, to see how consistent the reported effects really are, and to identify the gaps this thesis sets out to address.

2.6. Literature Review

Extensive experimental and numerical investigations have explored how surface dimples influence the aerodynamic performance of airfoils. These studies broadly aim to understand how small, localised surface cavities modify boundary-layer behaviour, suppress laminar separation bubbles, and affect the onset and progression of stall. Reported outcomes include reductions in pressure drag, enhancements in lift, and improvements in overall lift-to-drag ratio under certain flow conditions, although the magnitude and consistency of these benefits vary considerably across the literature.

The majority of these studies implement a single row of *static* dimples placed at various chordwise locations on the suction side of the airfoil, tested under a wide range of Reynolds numbers. The underlying premise of the fundamental mechanisms is consistent across the literature: a sufficiently small cavity generates a streamwise vortex pair that energises the local boundary layer and redistributes momentum. However, the aerodynamic effect of a dimple is highly sensitive to its size, depth, spacing, and especially chordwise position, which contributes to the often conflicting and case-dependent results reported across the literature.

As mentioned, several studies have demonstrated that surface dimples can improve lift or reduce drag under specific geometric conditions. Livya et al. [19] analysed a dimpled NACA 0018 airfoil using four cavity geometries (square, hexagonal, cylindrical, and spherical) and found that inward spherical dimples provided the most favourable aerodynamic outcome, increasing lift while inducing only a modest drag penalty at moderate angles of attack. Srivastav [33] performed both 2-D and 3-D numerical studies on the same NACA 0018 profile. The 2-D simulations indicated that outward elliptical dimples reduced drag, while the 3-D study showed that outward spherical dimples enhanced lift but increased drag. Importantly, Srivastav also concluded that inward spherical dimples were the most effective overall, a finding that also aligns with the results of Kaushik et al. [15]. Taken together, these contrasting observations highlight the strong dependence of the dimple effect on both dimensionality and cavity shape.

Chordwise placement is another key parameter exhibiting substantial inconsistency across studies. Saraf et al. [26] investigated a NACA 0012 airfoil with inward spherical dimples placed at 10%, 25%, 50%, and 75% of the chord, concluding that a 75%-chord location provided the most beneficial behaviour. Singh et al. [30] experimentally studied a modified NACA 4311 using outward hexagonal dimples, semi-cylindrical dimples of three sizes, and elongated strip-type modifications. They determined that the semi-cylindrical cavity placed at 66% chord delivered the highest lift-to-drag ratio. Binci et al. [8] evaluated inward elliptical dimples on a NACA 64-014A profile and showed that dimples placed at approximately 55% chord mitigated the laminar separation bubble and reduced pressure drag. Stolt et al. [34], in contrast, examined leading-edge spherical dimples (0–20% chord) using tomographic PIV and found that the aerodynamic impact was subtle, highly Reynolds-number dependent, and did not yield consistent drag reduction.

Other studies have introduced more complex dimple configurations. Ali et al. [3] investigated a NACA 65(3)-218 wing incorporating three rows of dimples at 15%, 50%, and 85% chord simultaneously, tested in six configurations spanning suction-side arrays, pressure-side arrays, and symmetric (both sides) arrangements, with both inward and outward geometries. Their study showed that inward dimples consistently outperformed outward ones, reporting a notable improvement of 9.29% in L/D at a 12° angle of attack for the suction-side inward-dimple configuration. Kaushik et al. [15] performed a 3-D CFD study on a NACA 0018 wing with a single inward spherical dimple at 24% chord and found modest gains at low angles of attack. Soundharya et al. [32] examined relatively large semi-spherical cavities whose depths approached a significant fraction of the airfoil thickness; while the authors reported small improvements under narrow conditions, their results clearly show that such deep “cut-out-like” cavities degraded performance over most of the operating envelope. Mustak et al. [21] placed inward spherical dimples along the rear half of a NACA 4415 and observed substantial improvements in lift, stall angle, and glide ratio (L/D), demonstrating that aft-located dimples can be particularly effective.

Collectively, these studies demonstrate that the chordwise placement of static dimples remains an open research question: improvements have been reported at 15%, 24%, 35%, 50%, 55%, 66%, 75%, and 85% chord, with no consensus on a universally optimal location. A few trends are, however, broadly consistent across the literature. First, inward-facing dimples placed on the suction side generally outperform outward or pressure-side configurations. Second, the dimple geometry must remain relatively small, typically on the order of a few percent of the chord, as larger cavities behave more like surface

cut-outs and tend to degrade aerodynamic performance. Beyond these two common findings, the inconsistencies in performance stem from variations in cavity geometry, Reynolds number, turbulence model, test dimensionality (2-D vs. 3-D), and the local state of the boundary layer.

Table 2.2 summarises the underlying parameters and conclusions from the reviewed literature. The table expands on the overview originally compiled by Stolzman [35], includes small corrections to a few inaccuracies, has independently verified all reported values, and incorporates additional studies while excluding those for which the sources could not be verified.

As summarised in Table 2.2, a consistent finding is that *all* prior investigations on lifting surfaces utilise passive, fixed-shape dimples. None of them has examined the effect of dynamically adjusting dimple depth to respond to instantaneous flow conditions. In contrast, adaptive morphable-surface concepts have been successfully demonstrated on bluff bodies. Vilumbrales-Garcia et al. [39] showed that actively varying dimple depth on a sphere, through an internal pressure-driven actuation mechanism, can shift the drag crisis, alter separation location, and reduce total drag significantly, by up to 63%. Similar behaviour was reported by Terwagne et al. [37], who examined a comparable spherical configuration with controllable cavity depth and likewise observed a strong sensitivity of drag to the dynamically adjusted surface topology. Together, these findings highlight a clear opportunity for adaptive surface roughness to modulate boundary-layer behaviour.

Across the literature, almost all reported performance benefits arise under a narrow set of flow conditions, and the location at which a dimple is beneficial varies substantially between studies. This lack of consistency strongly suggests that a given dimple geometry and chordwise position is only optimal for a specific flow state. If the incoming flow changes, the optimal dimple geometry and configuration likely change with it. Notably, none of the existing airfoil investigations vary dimple *depth* or *chord location* dynamically within the same flow field, nor do they attempt to maintain performance across a changing aerodynamic environment.

This motivates the central research gap addressed in the present work: while static dimples on airfoils have been widely studied, the influence of a *dynamically actuated, multi-row dimple array* on a lifting surface has not, to the author's knowledge, been explored. In particular, where adverse pressure gradients promote separation at elevated angles of attack, have not been examined using an actively controlled surface topology, on a lifting body.

Taken together, these observations indicate that the inconsistent results reported for static dimples may derive from their inability to adapt to changing pressure-gradient conditions along the airfoil. By enabling real-time variation of dimple depth, and by distributing several rows of inward spherical dimples along the suction-side of the airfoil, as indicated by the trends in the literature and summarised in Table 2.2, an adaptive system has the potential to overcome these inconsistencies and provide a reconfigurable mechanism for robust, condition-dependent stall mitigation. In analogy to aircraft's conventional aerodynamic control surfaces, which continuously adjust their deflection to maintain optimal performance under changing flight conditions, it is reasonable to expect that surface-integrated dimples may likewise require dynamic adjustment to remain effective across varying flow states.

Table 2.2: Summary of research efforts on the dimple effect.
Expanded and refined from the overview presented in Stolzman (2021) [35].
Entries that could not be independently verified have been omitted, and inaccuracies have been corrected.

Researcher	Profile Used	Profile Length	Reynolds Number	Exp. / Num.	2D / 3D	Optimal Dimple Geometry	Optimal Dimple Location	Dimple Diameter	Dimple Depth	AoA	Adaptive Dimple Depth
Mustak et al. [21]	NACA 4415	210 mm	6.2×10^5	Exp.	3D	Inward Spherical	$> 50\%$ SS	NaN	NaN	[0, 18]	No
Livya et al. [19]	NACA 0018	NaN	NaN	Num.	2D	Inward Spherical	$< 50\%$ SS	NaN	NaN	[0, 20]	No
Kaushik et al. [15]	NACA 0018	250 mm	3.1×10^5	Num.	3D	Inward Spherical	24% SS	8 mm	NaN	4, 7, 10	No
Saraf et al. [26]	NACA 0012	NaN	NaN	Num.	2D	Inward Spherical	75% SS	0.02c	NaN	[0, 16]	No
Srivastav (2D) [33]	NACA 0018	160 mm	3.2×10^5	Num.	2D	Outward Elliptical	$> 50\%$ SS	5 mm	2 mm	[−5, 15]	No
Srivastav (3D) [33]	NACA 0018	160 mm	3.2×10^5	Num.	3D	Inward Spherical	$> 50\%$ SS	4 mm	2 mm	[0, 20]	No
Singh et al. [30]	NACA 4311	136 mm	2.1×10^5	Exp.	3D	Outward Cylindrical	66% SS	14.5 mm	7 mm	[0, 0]	No
Binci et al. (Num.) [8]	NACA 64-014A	100 mm	1.75×10^5	Num.	3D	Inward Elliptical	55% SS	0.0335c	0.00705c	0	No
Binci et al. (Exp.) [8]	NACA 64-014A	100 mm	$[1.5, 12] \times 10^5$	Exp.	3D	Inward Elliptical	55% SS	0.0335c	0.00705c	0	No
Stolt et al. [34]	NACA 0015	135 mm	1.0×10^5	Exp.	3D	Inward Spherical	0 – 20%	5 mm	2.5 mm	10, 12, 15	No
V.-Garcia et al. [39]	Sphere	∅100mm	$[0.6, 13] \times 10^5$	Exp.	3D	Inward Spherical	NaN	8.7 mm	[0, 2]mm	0	Yes
Ali et al. [3]	NACA 65(3)-218	121 mm	3.0×10^5	Num.	3D	Inward Spherical	$50\% \pm 35\%$ SS	0.02c	0.0015c	0, 4, 8, 12	No
Omer et al. [23]	NACA 0012	1000 mm	3.0×10^6	Num.	2D	Inward Spherical	75% SS	0.03c	0.015c	[0, 20]	No
Terwagne et al. [37]	Sphere	∅40mm	$[3, 12] \times 10^4$	Exp.	3D	Inward Spherical	NaN	2.2 mm	[0, 0.8]mm	0	Yes
Stolzman [35]	NACA 4414	100 mm	2.0×10^5	Num.	2D	Inward Spherical	35% SS	2.5 mm	0.5 mm	[−15, 15]	No

2.7. Airfoil Design Parameters

The theory and literature review together set out both what is known about dimpled airfoils and what remains open, and it is these findings that the present work is intending to test. Doing so requires a physical airfoil and a matching numerical model, and the choices that define that geometry—chord, span, angle-of-attack range, and the dimple parameters themselves—should not be made arbitrarily but should follow from the prior work reviewed above and the mechanisms established in the theory. This section sets out those design parameters and the constraints that fix them, beginning with the wind-tunnel envelope that constrains the model and ending with the geometric dimple choices informed by Section 2.5.1.

The design of the dimpled airfoil is constrained primarily by the available test-section dimensions of the DTU Red Wind Tunnel (RWT) and the TU Delft Small Low-Turbulence Tunnel (SLT). The chord, span, and maximum angle of attack must be chosen so that the model can be physically accommodated while keeping wall interference and blockage effects within a manageable range. Since the same geometry will be used for both the wind-tunnel model and the numerical investigations, the wind-tunnel envelope effectively sets the outer limits for the CAD model and, consequently, for the CFD setup.

From the literature overview in Table 2.2 and the research questions formulated in Section 1.1, the present work focuses on relatively high angles of attack, where flow separation and stall are expected to be strongly influenced by the dimple actuation. In this regime, blockage ratios and wall-interference corrections become non-negligible and must be explicitly considered. For a given span and tunnel cross-section, the chord length and the angle of attack are the dominant parameters governing the projected frontal area of the model and, therefore, the blockage ratio. These quantities thus become key design parameters: they must be selected such that the airfoil prototype is large enough to house the internal actuation mechanism, yet small enough that the resulting blockage remains within the range where classical correction theory can still provide reliable free-air equivalents.

2.7.1. Wind-Tunnel Wall Interference and Blockage

A closed-test-section wind tunnel constrains the flow around the model, increasing the effective freestream velocity relative to the unconfined case and thereby distorting the measured C_L , C_D , and α relative to their free-air equivalents. The severity of this interference grows as the model becomes larger relative to the test section. A simple, geometry-only measure of this is the frontal-area blockage ratio

$$\beta = \frac{A_{\text{proj}}(\alpha)}{A_{\text{ts}}}, \quad (2.18)$$

where $A_{\text{proj}}(\alpha)$ is the projected frontal area of the model at a given angle of attack and $A_{\text{ts}} = H \cdot W$ is the test-section cross-sectional area. The blockage ratio does not enter the classic correction equations directly, rather, it serves as a conservative design proxy that can be evaluated from the model geometry alone, before any aerodynamic data are available. Keeping β small ensures that the underlying interference effects; solid blockage, wake blockage, and streamline curvature, all remain small, and hence that the linearised correction theory remains valid. Blockage corrections are generally considered reliable for small blockage ratios, with acceptable thresholds in the literature typically cited in the range of 5–10% [6], beyond which the corrections progressively lose accuracy. The blockage ratio is therefore used here as the primary sizing criterion for the model. The chord and thickness ratio are selected such that β remains within this manageable range across the full angle-of-attack envelope, ensuring that the correction framework can be applied if required. Table A.1 lists the blockage ratio for a range of chord lengths, thickness ratios, and both wind-tunnel cross-sections. This table forms the basis of the chord and thickness selection discussed in Section 2.7.2.

Correction Framework

For completeness, the correction framework that would be applied to recover free-air values is summarised below. The present work adopts the formulation of Selig et al. [29], in which the lateral boundaries of a two-dimensional test section give rise to four distinct phenomena: buoyancy, solid blockage, wake blockage, and streamline curvature.

Buoyancy is an additional drag force arising from the streamwise static-pressure gradient caused by the growth of the wall boundary layers. For a constant-area test section this effect is generally negligible for airfoils and, where relevant, is absorbed directly into the freestream velocity correction.

Solid blockage accounts for the physical displacement of the flow by the model volume, which reduces the effective flow area and increases the local velocity. It is expressed as

$$\varepsilon_{sb} = \frac{K_1 M_v}{A_{ts}^{3/2}}, \quad (2.19)$$

where K_1 is a body-shape factor, M_v is the model volume, and A_{ts} is the test-section area.

Wake blockage accounts for the reduced velocity within the airfoil wake, which by continuity must be balanced by an increased velocity outside the wake. It scales with the measured drag and is given by

$$\varepsilon_{wb} = \left(\frac{c}{2 h_{ts}} \right) C_{D_u}, \quad (2.20)$$

where c is the chord, h_{ts} is the test-section height, and C_{D_u} is the uncorrected drag coefficient.

Streamline curvature arises because the tunnel walls prevent the free-air curvature of the streamlines over the lifting body, effectively increasing the airfoil camber. For a closed test section this increases the measured lift, pitching moment, and angle of attack while leaving the drag essentially unaffected. The lift and angle-of-attack corrections are

$$\Delta C_{L,sc} = \sigma C_L, \quad \Delta \alpha_{sc} = \frac{57.3 \sigma}{2\pi} (C_L + 4 C_{m,c/4}), \quad (2.21)$$

with the curvature parameter

$$\sigma = \frac{\pi^2}{48} \left(\frac{c}{h_{ts}} \right)^2. \quad (2.22)$$

Application to the Present Work

It is emphasised that **all aerodynamic coefficients reported in this thesis are presented in their uncorrected form**. No wall-interference or blockage corrections have been applied to any of the measured or simulated results. The airfoil model was deliberately sized through the blockage criterion in Table A.1 such that the blockage ratio remains within the range where the correction framework of Selig et al. [29], summarised above, remains valid. Should free-air equivalent data be required, these corrections can therefore be applied directly to the reported values without further qualification.

2.7.2. Selection of Chord Length

Based on the blockage-correction framework outlined in Section 2.7.1, the chord must be selected such that the blockage remains within a regime where solid, wake, and lift-interference corrections remain small and negligible or reliable. Table A.1 summarises the blockage ratios β for different combinations of chord length c , thickness ratio t/c , and tunnel configuration. As expected, β increases with both chord length and angle of attack, with the smaller DTU RWT, of course, showing higher blockage ratios than the larger TU Delft SLT.

A practical design target was established to keep the blockage ratio below approximately $\beta = 0.10$ throughout the majority of the operating range. Inspection of Table A.1 shows that all but one configuration gives a value of $\beta > 0.10$ in the RWT at $\alpha = 20^\circ$. Measurements at the highest angles of attack in the RWT, $\alpha \geq 15^\circ$, are therefore not ideal but remain manageable using the correction procedures described in Section 2.7.1 if needed.

From this, a chord length of $c = 200$ mm combined with a thickness ratio of $t/c = 0.15$ represents a practical upper limit. In the RWT, where the model spans the full 0.50 m tunnel width, this configuration reaches $\beta = 0.1077$ at $\alpha = 15^\circ$ and $\beta = 0.1288$ at $\alpha = 20^\circ$. The same section is tested in the SLT at a span of 600 mm, again spanning the full tunnel width, and the larger cross-section gives a lower blockage at matched incidence, $\beta = 0.0933$ at $\alpha = 16^\circ$ rising to $\beta = 0.1073$ at $\alpha = 20^\circ$. Since the SLT campaign extends only to $\alpha = 16^\circ$, every SLT measurement reported in this thesis falls within the $\beta < 0.10$ design target, whereas in the RWT the target is exceeded from $\alpha = 15^\circ$ upward.

The prototype must provide sufficient internal volume to accommodate the adaptive dimple actuation mechanism and supporting structure. Consequently, the design objective is not to minimise the blockage ratio, but rather to maximise the physical size of the model while remaining within acceptable blockage limits. A thinner section ($t/c = 0.12$) or a shorter chord would reduce blockage further but

would unnecessarily restrict the available space for the internal hardware. Conversely, increasing the chord length to 220 mm or the thickness ratio to $t/c = 0.18$ would result in excessive blockage over a larger portion of the operating envelope. The combination of $c = 200$ mm and $t/c = 0.15$ therefore represents the largest practical configuration that remains compatible with the blockage constraints identified in Table A.1, while still providing sufficient space for the adaptive dimple mechanism and associated instrumentation.

2.7.3. Selection of Airfoil Profile

With the chord and thickness ratio determined, the next design decision concerns the choice of airfoil family. Three candidate profiles are considered: a symmetric NACA 00-series section (NACA 0015), a cambered NACA 4-digit section (NACA 4415), and a NACA 6-series laminar-flow section (NACA 63-415), as illustrated in Figure 2.13. Since the thickness ratio is fixed at $t/c = 0.15$, the geometric blockage characteristics of the three profiles are essentially identical. The selection therefore rests on aerodynamic relevance, availability of reference data, and suitability for representing wind-turbine blade sections.




Figure 2.13: Airfoil geometries for NACA 0015, NACA 4415, and NACA 63-415. All plots are based on coordinate datasets obtained from AirfoilTools [2].

The NACA 00-series offers clear practical advantages. The geometric symmetry simplifies manufacturing, coordinate generation, mounting, and calibration. The relatively thick mid-section provides ample internal volume for the actuation mechanism, and the aerodynamic behaviour is straightforward to model. However, symmetric airfoils are only weakly representative of the lifting surfaces used in drones, small aircraft, or wind-turbine blades, where positive lift at zero incidence is generally required.

The NACA 4415 represents a moderately cambered and moderately thick airfoil that is widely used in aerodynamic studies and well documented in the literature. Its stall behaviour is typically characterised by progressive trailing-edge separation, which is advantageous for analysing separation control mechanisms. Nevertheless, compared to modern wind-turbine blade sections, the NACA 4415 is less representative in terms of pressure distribution and loading characteristics, particularly in the context of laminar-to-turbulent transition and mid-span blade aerodynamics.

The NACA 63-415 laminar-flow airfoil is selected as the baseline geometry for the present study. This profile is more representative of wind-turbine blade sections, particularly in the mid-span region, while remaining relevant for UAV and small-aircraft applications. Importantly, the NACA 63-415 is well documented through both numerical and experimental data. Steady aerodynamic polars are available from

international databases such as AirfoilTools [2], while high-quality experimental static polars have been obtained in wind-tunnel campaigns conducted at DTU Risø [5]. The availability of both numerical and experimental reference data enables robust validation of the CFD simulations and provides a reliable baseline for assessing the aerodynamic impact of adaptive dimple actuation.

For these reasons, the **NACA 63-415** airfoil with a chord length of $c = 200$ mm and a thickness ratio of $t/c = 0.15$ is selected as the baseline geometry for the airfoil prototype. This choice provides a suitable balance between aerodynamic relevance, data availability, and practical manufacturability, while remaining compatible with the blockage constraints established in Section 2.7.2.

2.7.4. Selection of Dimple Diameter, Depth, and Spacing

With the chord fixed at $c = 200$ mm, the next geometric design variables are the dimple diameter D_d , dimple depth k , and dimple spacing D_s . These quantities are illustrated in Figure 2.14. The dimple diameter D_d is defined as the opening diameter of the dimple, the dimple depth k is the maximum distance between the original airfoil surface and the deepest point of the cavity, and the spacing D_s is the centre-to-centre distance between neighbouring dimples.

To allow comparison between studies employing different airfoil sizes, these quantities are commonly expressed in the non-dimensional form introduced in Equation 2.17, in which the dimple diameter and depth are normalised by the airfoil chord length while the spacing is expressed relative to the dimple diameter.

The staggered dimple arrangement shown in Figure 2.14 further introduces a projected streamwise spacing

$$\Delta x = \sin\left(\frac{\pi}{3}\right) D_s, \quad (2.23)$$

which governs the distance encountered by the flow in the streamwise direction.




Figure 2.14: Definition of dimple geometry and spacing. The dimple diameter is denoted by D_d , the centre-to-centre spacing by D_s , and the dimple depth by k . The staggered arrangement defines a projected streamwise spacing $\Delta x = \sin(\pi/3) D_s$.

The literature review presented in Table 2.2 was therefore re-analysed using these non-dimensional parameters, and the resulting values are summarised in Table 2.3. Most airfoil studies cluster in a relatively narrow range of dimple diameters, with $D_d/c \approx 0.02$ – 0.05 , while significantly larger dimples such as those used by Singh et al. [30] ($D_d/c \approx 0.11$) are the exception rather than the rule. Reported depth ratios typically lie in the range $k/c \approx 0.005$ – 0.02 for airfoils and adaptive spheres. Where spacing information is available, the spacing ratio generally falls within $D_s/D_d \approx 1.5$ – 2.25 , indicating moderately separated dimples rather than a nearly continuous corrugated surface.

Table 2.3: Overview of dimple diameter and depth normalised by chord length c for selected studies from Table 2.2. “–” indicates that the quantity was not specified. Dimple Depth k ranges for adaptive or morphable concepts are indicated in brackets. The spacing ratio D_s/D_d denotes the centre-to-centre spacing D_s normalised by dimple diameter D_d , as depicted in Figure 2.14.

Researcher	Profile used	c [mm]	D_d/c [–]	k/c [–]	D_s/D_d [–]
Saraf et al. [26]	NACA 0012	–	0.020	–	–
Kaushik et al. [15]	NACA 0018	250	0.032	–	2.25
Srivastav (2D) [33]	NACA 0018	160	0.031	0.013	–
Srivastav (3D) [33]	NACA 0018	160	0.025	0.013	–
Singh et al. [30]	NACA 4311	136	0.107	0.051	–
Binci et al. (Num.) [8]	NACA 64-014A	100	0.034	0.007	2.0
Binci et al. (Exp.) [8]	NACA 64-014A	100	0.034	0.007	2.0
Stolt et al. [34]	NACA 0015	135	0.037	0.019	1.6
V.-Garcia et al. [39]	Sphere	∅100	0.087	[0.00, 0.02]	–
Ali et al. [3]	NACA 65-218	121	0.020	0.002	1.5
Omer et al. [23]	NACA 0012	1000	0.030	0.015	–
Terwagne et al. [37]	Sphere	∅40	0.055	[0.00, 0.02]	–
Stolzman [35]	NACA 4414	100	0.025	0.005	–

On this basis, a target dimple diameter of $D_d/c = 0.05$ is selected for the present airfoil prototype. This choice places the design at the upper end of the “small-dimple” regime suggested by Table 2.3, providing geometric resolution for manufacturing and actuation while remaining consistent with the ranges where beneficial effects on separation control have been reported. For an idealised, fully developed inward spherical dimple with $D_d/c = 0.05$, the corresponding depth would be $k/c \approx 0.025$. In practice, the adaptive concept considered here allows the depth to be varied continuously up to this value, so that both shallow and relatively deep configurations within the literature-supported range can be explored.

For the initial prototype, the dimple spacing is set to a nominal value of $D_s/D_d = 2.0$, corresponding to a centre-to-centre distance of $D_s = 2D_d$. This choice is consistent with the spacing ratios reported for lifting-surface studies in Table 2.3, and provides a compromise between strong local interaction and excessive overlap of the induced vortical structures. Overall, the selected combinations of D_d/c , k/c , and D_s conform to the prevailing conclusion from prior work that dimples must remain relatively small and moderately spaced in order to promote local mixing and delay separation without introducing excessive form drag or large-scale geometric disturbance.

2.7.5. Selection of Dimple Chord Location

The final geometric design parameter is the chordwise location of the dimples. Unlike the dimple diameter and depth, which can be scaled directly from the literature using the non-dimensional ratios D_d/c and k/c , the optimal dimple location depends strongly on the local flow physics and airfoil geometry.

The literature review presented in Table 2.2 reveals a clear trend. With the exception of the leading-edge study by Stolt et al. [34], most authors report optimum dimple locations on the mid- to aft-suction surface. Reported locations include approximately 35% chord by Stolzman [35], 50–55% chord by Ali et al. [3] and Binci et al. [8], 66% chord by Singh et al. [30], and 75% chord by Saraf et al. [26]. Several studies further conclude that dimples located downstream of approximately 50% chord are most effective at influencing separation and stall behaviour.

Consequently, the present work focuses on the region between approximately $0.33c$ and $0.85c$, which encompasses nearly all optimum locations reported in the literature while also covering the expected separation region of the selected NACA 63-415 airfoil.

The chordwise locations of the seven candidate dimple rows are defined about the mid-chord position using the spacing parameters introduced in Figure 2.14. For a dimple diameter of $D_d = 10$ mm and a centre-to-centre spacing of $D_s = 20$ mm, the projected streamwise spacing becomes $\Delta x = \sin\left(\frac{\pi}{3}\right) D_s = 17.3$ mm. The resulting row locations are listed in Table 2.4.

Table 2.4: Chordwise locations of the seven dimple rows for $c = 200$ mm, with $D_s = 2D_d = 20$ mm.

Row	Expression for x [mm]	x [mm]	x/c [%]
1	$\frac{c}{2} - 2 \sin(\pi/3) D_s$	65.4	32.7
2	$\frac{c}{2} - \sin(\pi/3) D_s$	82.7	41.3
3	$\frac{c}{2}$	100.0	50.0
4	$\frac{c}{2} + \sin(\pi/3) D_s$	117.3	58.7
5	$\frac{c}{2} + 2 \sin(\pi/3) D_s$	134.6	67.3
6	$\frac{c}{2} + 3 \sin(\pi/3) D_s$	152.0	76.0
7	$\frac{c}{2} + 4 \sin(\pi/3) D_s$	169.3	84.7

The seven rows were subsequently combined into four geometry families, shown in Table 2.5. The Smooth geometry is the undimpled baseline configuration, while Mid, Aft, and Full cover the mid-chord, aft-chord, and full extent of the dimple region, respectively.

The selected layouts are intentionally simple and were chosen to balance physical relevance with the practical limitations of the project. The literature review suggests that aft-chord dimples are generally most effective, but it remains unclear how sensitive the performance is to the exact extent of the dimpled region. The three dimpled configurations are therefore intended to provide a systematic comparison between different partial chordwise coverage while keeping the total number of test cases manageable.

Table 2.5: Geometry families used for the experimental and numerical campaigns.

ID	Rows	x/c range	D_d/c	k/c
Smooth	—	—	—	—
Mid	1–4 (4/7)	0.33–0.59	0.05	0–0.025
Aft	4–7 (4/7)	0.59–0.85	0.05	0–0.025
Full	1–7 (7/7)	0.33–0.85	0.05	0–0.025

Many additional combinations could have been investigated, including isolated dimple rows, forward-shifted layouts, and varying coverage lengths. However, such an expanded parameter study was considered beyond the scope of the present project.

2.7.6. Reynolds Number Range

With the chord fixed at $c = 200$ mm, the accessible Reynolds number range is set by the usable velocity envelope of the DTU RWT and the TU Delft SLT. In practice, both facilities provide accurate low-speed measurements from $U_{\min} \approx 5$ m/s, while the upper limits for the present campaign are taken as $U_{\max} \approx 40$ m/s for the RWT and $U_{\max} \approx 35$ m/s for the SLT, consistent with their documented operating ranges for low-speed aerodynamic testing [11, 25]. Air is assumed incompressible ($\text{Mach} < 0.3$) with properties corresponding to $T \approx 20^\circ\text{C}$ and $p \approx 1$ atm, for which the kinematic viscosity is taken as $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$, the value used to set the tunnel speeds throughout the campaign.

The chord-based Reynolds number follows from basic fluid mechanics [40]:

$$Re = \frac{U_\infty c}{\nu}, \quad (2.24)$$

Using $c = 0.20$ m and $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$, the resulting Reynolds number ranges for the two tunnels are summarised in Table 2.6.

Table 2.6: Velocity and corresponding chord Reynolds number ranges for the DTU Red Wind Tunnel (RWT) and the TU Delft Small Low-Turbulence Tunnel (SLT), assuming $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$.

Wind tunnel	Chord [mm]	$U_{\min}$ [m/s]	$U_{\max}$ [m/s]	$Re_{\min}$ [–]	$Re_{\max}$ [–]
DTU RWT	200	5	40	6.7×10^4	5.3×10^5
TU Delft SLT		5	35	6.7×10^4	4.7×10^5

As shown in Table 2.2, most dimpled-airfoil studies in the literature operate in Reynolds number ranges of $\mathcal{O}(10^4\text{--}10^6)$. The present test window, approximately $Re \in [6 \times 10^4, 5 \times 10^5]$, lies well within this interval and overlaps the lower to mid portion of the ranges reported for comparable NACA profiles. The chosen operating conditions thereby provide Reynolds numbers that are directly comparable to prior work, while still allowing systematic exploration of stall and separation behaviour under adaptive dimple actuation.

Table 2.7: Selected operating conditions for the experimental campaign with $c = 200$ mm and $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$.

Wind tunnel	U_∞ [m/s]	Re [–]
DTU Red Wind Tunnel	7.5	1.0×10^5
	15.0	2.0×10^5
	37.5	5.0×10^5
TU Delft Small Low-Turbulence Tunnel	7.5	1.0×10^5
	15.0	2.0×10^5

The values were chosen to lie within the operating envelope of both wind tunnels while spanning a representative portion of the Reynolds number range identified in the literature review (Table 2.2). The corresponding freestream velocities are listed in Table 2.7. Reynolds numbers of $Re = 1 \times 10^5$ and 2×10^5 were tested in both facilities, while the higher Reynolds number of $Re = 5 \times 10^5$ was only achievable in the RWT.

2.8. Summary of Final Design Parameters

The airfoil prototype configuration adopted in this work follows directly from the blockage constraints, the manufacturing requirements, and the geometric sizes and ranges supported by the literature on dimpled lifting surfaces. The chord and thickness ratio were selected to ensure sufficient internal volume for the actuation system while keeping wind-tunnel blockage within the regime where classical wall-interference is negligible or correction theory remains valid. The choice of dimple geometry is based on the dimensional ranges reported in Table 2.3, emphasising small, inward spherical cavities with moderate spacing. Finally, the NACA 63-415 profile is selected for its aerodynamic relevance

to wind-turbine blade sections and its extensive numerical and experimental documentation, enabling reliable validation of both CFD simulations and wind-tunnel measurements.

Table 2.8: Summary of selected design parameters for the prototype airfoil.

Parameter	Selected Value
Chord length c	200 mm
Thickness ratio t/c	0.15
Airfoil profile	NACA 63-415
Dimple diameter ratio D_d/c	0.05
Maximum dimple depth ratio k/c	≈ 0.025 (continuously variable)
Dimple spacing ratio D_s/D_d	2.0
Dimple geometry	Inward spherical
Dimple placement region	$x/c \in [0.33, 0.85]$, Suction side
Number of chordwise rows	4 to 7
Operating Reynolds numbers	$Re \in 1 \times 10^5, 2 \times 10^5, 5 \times 10^5$
Wind tunnels used	DTU Red Wind Tunnel, TU Delft Small Low-Turbulence Tunnel

The parameters in Table 2.8 collectively define the baseline configuration for the numerical simulations, the experimental screening tests, and the subsequent dynamically actuated prototype used in the wind-tunnel and PIV measurements. The chosen ranges maintain consistency with documented dimple-airfoil behaviour while providing sufficient flexibility for systematic variation during the investigation of adaptive flow-control concepts.




Figure 2.15: CAD geometries of the four airfoil dimple layout configurations from Table 2.5. Smooth is the undimpled NACA 63-415 baseline. Mid covers the mid-chord region, Aft the aft region, and Full the full extent. All chordwise positions of rows follow from Table 2.4. All dimpled configurations use inward spherical dimples.

The CAD models shown in Figure 2.15 represent the final geometries derived from the design process summarised in Table 2.8. Together, the four configurations define the baseline for the test matrix used throughout the numerical and experimental investigations. Smooth serves as the undimpled reference case, while Mid, Aft, and Full introduce progressively different chordwise dimple coverage patterns within the selected design space. These geometries form the basis for the CFD studies, wind-tunnel measurements, and subsequent evaluation of the adaptive dimple concept.

Methodology I - Aerodynamic Coefficient Measurements

3.1. Wind Tunnel Facility

All experiments in this chapter were conducted in the Red Wind Tunnel (RWT), a low-turbulence wind tunnel at DTU Fluid Lab operated in open-loop configuration [25], depicted in Figure 3.1. The tunnel is powered by a 34 kW axial fan and incorporates four anti-turbulence screens in the settling plenum, followed by a contraction with an area ratio of 8.2:1. The contraction outlet has a rectangular cross-section of $0.50 \times 0.75 \text{ m}^2$, and the test section downstream is 2 m long. The test section is constructed from 18 mm thick transparent polycarbonate panels. The freestream turbulence intensity has been characterised as $Tu < 0.1 \%$ across the operating velocity range [25], making the facility well suited to aerodynamic measurements at the low Reynolds numbers considered here.




Figure 3.1: The DTU Red Wind Tunnel (RWT) used for the low-speed wind-tunnel measurements. The red contraction section upstream of the test section conditions and accelerates the incoming flow, while the transparent rectangular test section provides optical and mechanical access for mounting the airfoil model and measurement equipment. The downstream red section guides the flow out of the test section into the axial fan, which is out of picture.

The freestream speed U_∞ was derived from measurements of the dynamic pressure using a Pitot-static probe connected to a high-accuracy differential pressure transducer. The dynamic pressure is related to the total and static pressures by

$$q_\infty = p_{t,\infty} - p_\infty = \frac{1}{2} \rho U_\infty^2$$

using a Pitot-static probe connected to a high-accuracy differential transducer. Solving for velocity under the incompressible assumption gives

$$U_\infty = \sqrt{\frac{2(p_{t,\infty} - p_\infty)}{\rho}}.$$

The air density ρ was estimated from laboratory temperature and barometric pressure, assuming ideal-gas behaviour.

3.2. Airfoil Models and Setup

The test models, summarised in Table 3.1 and illustrated in Figure 3.2, consist of a smooth baseline airfoil and a set of modified geometries featuring surface dimples of varying depth. The baseline corresponds to the smooth NACA 63-415 profile (denoted Smooth), while the Mid, Aft, and Full configurations carry dimples over the mid-chord, aft-chord, and full suction-side regions, respectively, each tested at three dimple depths: $k = 1.25$ mm, 2.50 mm, and 5.00 mm. This yields ten model variants in total, as listed in Table 3.1. All models share a chord length of $c = 200$ mm and a span of 500 mm, corresponding to the full width of the RWT test section. All models were 3D-printed, allowing rapid manufacturing and straightforward replacement between test runs.

Each model is mounted vertically in the test section via a stainless steel support rod passing through the airfoil at the quarter-chord point. The lower end of the rod is seated in a force balance hub, which provides the primary means of setting the angle of attack where the model is rotated about the rod axis to the desired α and held in place. The force balance itself is not used for load measurement in this campaign as aerodynamic forces are instead obtained from the wake rake and wall pressure taps described in Section 3.3. To accommodate the higher freestream velocities tested in the RWT compared to the SLT campaign, a second support rod is attached at the upper end of the airfoil span, bracing the model against the top wall of the test section and preventing aeroelastic deflection at the highest Reynolds numbers, the reason for which this force balance sensor is not utilised in the RWT. The full mounting arrangement, including both support rods and the wake rake position, is shown in Figure 3.3.

Table 3.1: Geometry configurations used in the RWT campaign. The numeric suffix denotes the fixed dimple depth in millimetres for $c = 200$ mm.

ID	Rows	x/c range	D_d/c	k/c	b/c
Smooth	–	–	–	–	2.5
Mid-1.25	1–4 (4/7)	0.33–0.59	0.05	0.0063	2.5
Mid-2.50	1–4 (4/7)	0.33–0.59	0.05	0.0125	2.5
Mid-5.00	1–4 (4/7)	0.33–0.59	0.05	0.0250	2.5
Aft-1.25	4–7 (4/7)	0.59–0.85	0.05	0.0063	2.5
Aft-2.50	4–7 (4/7)	0.59–0.85	0.05	0.0125	2.5
Aft-5.00	4–7 (4/7)	0.59–0.85	0.05	0.0250	2.5
Full-1.25	1–7 (7/7)	0.33–0.85	0.05	0.0063	2.5
Full-2.50	1–7 (7/7)	0.33–0.85	0.05	0.0125	2.5
Full-5.00	1–7 (7/7)	0.33–0.85	0.05	0.0250	2.5




Figure 3.2: CAD geometries of the four airfoil dimple layouts as manufactured for the RWT campaign, corresponding to the design configurations of Figure 2.15. Smooth is the undimpled NACA 63-415 baseline. Mid covers the mid-chord region, Aft the aft region, and Full the full extent from $x/c = 0.33$ to 0.85. All dimpled configurations use inward spherical dimples.

3.3. Wake Rake and Wall Pressure Measurements

This section describes the wake rake and wall pressure measurement campaign conducted on all ten airfoil configurations in the Red Wind Tunnel at DTU. The objective was to obtain aerodynamic force polars, lift coefficient C_L from the tunnel wall pressure taps and the drag coefficient C_D from the wake momentum deficit, across the full angle-of-attack and Reynolds number ranges defined in the test matrix. Because both quantities are derived from simple pressure measurements acquired simultaneously on the same model runs, this approach provided the broadest parametric coverage of the campaign at low per-case cost, forming the primary screening dataset of the present work.

3.3.1. Experimental Setup

The experimental arrangement is shown in Figure 3.3. Aerodynamic forces were obtained indirectly from two sets of pressure measurements acquired simultaneously on each model run. A wake rake was positioned downstream of the trailing edge, from which the drag coefficient was derived via the momentum deficit and Static pressure taps flush-mounted in the tunnel side walls, from which the lift coefficient was derived via the wall-pressure method. No force balance was used for load measurement in this campaign.




Figure 3.3: CAD render of the RWT experimental setup.

Wake Rake and Drag Coefficient

The wake rake consists of 16 total-pressure tubes arranged in a horizontal comb at mid-span, positioned 195 mm ($0.975c$) downstream of the airfoil trailing edge. The tubes are non-uniformly spaced to provide higher resolution in the wake core centred around the wake centreline. The rake spans 168 mm in total, covering the full expected wake width at the tested conditions.

The drag coefficient is obtained from the streamwise momentum deficit in the wake using the classical wake-integral method [6]. For incompressible flow, C_D is given by

$$C_D = \frac{2}{c} \int_{-\infty}^{\infty} \sqrt{\frac{p_{t,w} - p_{\infty}}{p_{t,\infty} - p_{\infty}}} \left(1 - \sqrt{\frac{p_{t,w} - p_{\infty}}{p_{t,\infty} - p_{\infty}}} \right) dy, \quad (3.1)$$

where $p_{t,w}$ is the local total pressure measured at each wake rake tube, $p_{t,\infty}$ is the freestream total pressure, p_{∞} is the freestream static pressure, and the integration is performed across the wake in the direction normal to the freestream y . In practice, the integral is evaluated numerically using the trapezoidal rule over the 16 discrete rake positions, with the dynamic pressure $q_{\infty} = p_{t,\infty} - p_{\infty}$ supplied by the Pitot-static probe described in Section 3.1. This form assumes that the static pressure across the wake has recovered to the freestream value at the rake station.

Wall Pressure Taps and Lift Coefficient

Static pressure taps are flush-mounted in both side walls of the test section, running the entire length of the 2-meter test section. The full array consists of 16 taps per side wall, of which 8 per side are active and connected to the data acquisition system. The active taps span from $x = -620$ mm to $x = 1140$ mm relative to the airfoil quarter-chord, covering the region from $-3.1c$ upstream to $+5.7c$ downstream and capturing the full wall-pressure signature induced by the airfoil circulation. The suction-side wall taps (positive y) and pressure-side wall taps (negative y) are mirror-symmetric in position.

The lift coefficient is derived from the measured wall static pressures using the wall-pressure integration method. The pressure coefficient at each tap location is defined as

$$C_p(x) = \frac{p_w(x) - p_\infty}{q_\infty}, \quad (3.2)$$

where $p_w(x)$ is the measured wall static pressure at streamwise position x , p_∞ is the freestream static pressure, and $q_\infty = \frac{1}{2}\rho U_\infty^2$ is the freestream dynamic pressure. The lift per unit span is then obtained by integrating the pressure difference between the two walls across all streamwise tap positions,

$$L = \int_{x_{\text{up}}}^{x_{\text{down}}} [p_w^{\text{PS}}(x) - p_w^{\text{SS}}(x)] dx, \quad (3.3)$$

where p_w^{SS} and p_w^{PS} are the static pressures on the suction-side and pressure-side walls, respectively. The pressure-side wall carries the higher static pressure when the airfoil is positively loaded, so the difference is taken in this order to give a positive lift for a positive angle of attack. The integral is evaluated numerically using the trapezoidal rule over the eight active tap positions per side. The lift coefficient follows as

$$C_L = \frac{L}{q_\infty c}. \quad (3.4)$$

All pressure signals — wake rake, wall taps, and the Pitot reference — were sampled simultaneously by the data acquisition system and time-averaged over each measurement run to obtain mean force coefficients at each operating condition.

3.3.2. Wake and Wall Pressure Test Matrix

The RWT campaign covered the ten model variants summarised in Table 3.1 and illustrated in Figure 3.2: the undimpled Smooth baseline and all three dimpled geometry families Mid, Aft, and Full, each tested at three dimple depths of 1.25, 2.50, and 5.00 mm. This campaign constitutes the primary parametric screening of the present work, enabled by the relative ease of manufacturing the 3D-printed airfoil models and the low per-case cost of wake rake and wall pressure measurements compared to optical or numerical methods.

The three Reynolds numbers span the full operating envelope of the RWT and overlap with the RANS and PIV datasets, enabling direct cross-method comparison at $Re = 1.0 \times 10^5$ and 2.0×10^5 . The angle-of-attack range was extended to cover the full polar from pre-stall through deep stall, as summarised in Table 3.2. Wake rake and wall pressure data were acquired simultaneously on the same model runs, so no additional cases were required to obtain both datasets.

Table 3.2: Operating conditions used in the wake and wall pressure measurement campaign. The three Reynolds numbers correspond to inlet velocities of $U_\infty = 7.5$ m/s, 15 m/s, and 37.5 m/s, using $\nu = 1.5 \times 10^{-5}$ m²/s and $c = 200$ mm.

Parameter	Values
Angle of attack α	$[-10^\circ, 20^\circ]$ in steps of 2°
Reynolds number Re	1.0×10^5 , 2.0×10^5 , 5.0×10^5
Inlet velocity U_∞	7.5 m/s, 15 m/s, 37.5 m/s
Flow regime	Incompressible

Combining ten geometry configurations, sixteen angles of attack, and three Reynolds numbers gives a total of 480 measurement runs, as summarised in Table 3.3.

Table 3.3: Summary of wake and wall pressure case count.

Item	Dimpled cases	Reference cases
Geometry configurations	3	1
Dimple depth levels	3	—
Angles of attack	16	16
Reynolds numbers	3	3
Number of cases	432	48
Total cases		480

3.3.3. Summary of Wake and Wall Pressure Setup

The wake and wall pressure campaign consisted of 480 measurement runs across ten model variants at three Reynolds numbers and sixteen angles of attack, constituting the primary parametric screening dataset of the present work. Measurements were performed in the Red Wind Tunnel at DTU using a 16-tube wake rake for drag and 8 active wall pressure taps per side for lift, acquired simultaneously on each run. The key instrumentation and campaign parameters are consolidated in Table 3.4. The results provide aerodynamic force polars across the full operating envelope of the RWT for all geometry families and dimple depths investigated.

Table 3.4: Summary of wake rake and wall pressure measurement setup and campaign parameters.

Parameter	Value
<i>Facility and technique</i>	
Wind tunnel	Red Wind Tunnel (RWT), DTU
Technique	Wake rake (drag) + wall pressure taps (lift)
Freestream turbulence	$Tu < 0.1\%$
<i>Operating conditions</i>	
Chord length c	200 mm
Reynolds numbers Re	1.0×10^5 , 2.0×10^5 , 5.0×10^5
Inlet velocities U_∞	7.5 m/s, 15 m/s, 37.5 m/s
Angles of attack α	$[-10^\circ, 20^\circ]$ in steps of 2°
Kinematic viscosity ν	$1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$
<i>Wake rake</i>	
Rake position	195 mm ($0.975c$) downstream of trailing edge
Active tubes	16 total-pressure tubes + 1 reference
Rake span	168 mm
Drag method	Momentum-deficit integral [6]
<i>Wall pressure taps</i>	
Active taps per side	8 (suction-side wall and pressure-side wall)
Streamwise coverage	$-3.1c$ to $+5.7c$ relative to quarter-chord
Lift method	Wall-pressure integration [6]
<i>Study scope</i>	
Geometry configurations	10 (Smooth, plus Mid, Aft, and Full at 1.25, 2.50, and 5.00 mm)
Total measurement runs	480

3.4. Surface Pressure Taps

This section describes the surface pressure tap measurements conducted on two selected airfoil configurations in the Red Wind Tunnel at DTU. The objective was to obtain chordwise pressure distributions $C_p(x/c)$ on both the suction and pressure surfaces, providing direct comparison with the RANS surface pressure predictions and supplementing the PIV flow field and integrated force coefficients obtained from the wake rake and wall pressure campaign.

3.4.1. Experimental Setup

Two dedicated pressure tap models were manufactured for this campaign: the smooth baseline Smooth PT and the mid-chord dimpled configuration Mid-2.50 PT. These correspond to the same geometries as Smooth and Mid-2.50 used throughout the RANS, LES, and PIV campaigns, with the single addition of 32 flush surface pressure taps — 20 on the suction side and 12 on the pressure side — at the chordwise locations listed in Table 3.5 and illustrated in Figure 3.4 and Figure A.2. Each tap consists of a 0.7 mm diameter hole drilled normal to the local surface, connected to an internal pressure channel routed through the model.




Figure 3.4: Transparent side-view CAD renders of the Smooth PT (top, blue) and Mid-2.50 PT (bottom, green) pressure tap models, revealing the internal pressure channels and tap locations. Both models share identical tap positions, with the distribution densest near the leading edge where the steepest pressure gradients are expected. The dimple cavities of Mid-2.50 PT are visible on the suction surface.

The models are 3D-printed in three parts: two hollow outer end caps and a fully solid centre section. The pressure channels are printed directly into the solid centre piece, keeping all internal passages airtight. Flexible rubber hoses of $\varnothing 1$ mm inner diameter connect the channel outlets to the pressure measurement system. The hoses exit the model through the hollow end caps alongside the stainless steel mounting rod, as shown in Figure 3.5. After assembly, all three parts were bonded with adhesive and the outer surface sanded to a smooth finish consistent with the other airfoil models in the campaign.

Table 3.5: Chordwise positions of the surface pressure taps on the NACA 63-415 pressure tap models. Positions are given as x/c fractions of the chord length $c = 200$ mm.

Side	x/c locations
Suction (upper)	0.000, 0.005, 0.022, 0.034, 0.047, 0.059, 0.071, 0.096, 0.147, 0.197, 0.248, 0.298, 0.349, 0.399, 0.450, 0.500, 0.601, 0.701, 0.801, 0.901
Pressure (lower)	0.010, 0.028, 0.053, 0.079, 0.104, 0.153, 0.203, 0.253, 0.351, 0.500, 0.649, 0.800

The pressure tap models are mounted in the RWT using the same support rod arrangement described in Section 3.2. The 32 surface tap hoses are connected directly to the same pressure measurement system used for the wake rake campaign, replacing the 32 wall and wake rake tubes. The data acquisition format, sampling rate, and zero-wind reference correction procedure are therefore identical to those described in Section 3.3.




Figure 3.5: Physical pressure tap models prior to final assembly, showing Mid-2.50 PT (top row) and Smooth PT (bottom row).

Each model consists of two hollow outer end caps and a fully solid 3D-printed centre section containing the airtight internal pressure channels. The rubber tubes emerging from the solid centre piece are routed through one of the hollow end caps and exits alongside the stainless steel mounting rod, visible on the right-hand side of each assembly. The three parts were bonded and sanded after this photograph.

Pressure Coefficients

The surface pressure coefficient at each tap is defined as

$$C_p\left(\frac{x}{c}\right) = \frac{p_{\text{tap}} - p_{\text{ref}}}{q_{\infty}}, \quad (3.5)$$

where p_{tap} is the time-averaged gauge pressure at the tap, p_{ref} is the zero reference pressure recorded with no airfoil model present in the wind tunnel, and $q_{\infty} = \frac{1}{2}\rho U_{\infty}^2$ is the freestream dynamic pressure from the Pitot-static probe. Results are presented as $-C_p$ versus x/c following the standard aerodynamic convention, with suction peaks appearing as positive values.

Lift and Drag Coefficients

The lift and drag coefficients are derived from the measured surface pressure distribution by numerical integration following the approach of Anderson [4]. The normal and axial force coefficients in the body axis system are first obtained by integrating the pressure coefficient over the airfoil contour,

$$C_n = \int_0^1 [C_{p,l}(x/c) - C_{p,u}(x/c)] d(x/c), \quad (3.6)$$

$$C_a = \int_0^1 C_{p,u} d(y_u/c) - \int_0^1 C_{p,l} d(y_l/c), \quad (3.7)$$

where $C_{p,u}$ and $C_{p,l}$ are the pressure coefficients on the suction and pressure surfaces respectively, y_u/c and y_l/c are the local surface ordinates, and the integrals are evaluated numerically using the trapezoidal rule over the discrete tap positions. The lift and drag coefficients in the wind axis system follow from the standard body-to-wind axis rotation,

$$C_L = C_n \cos \alpha - C_a \sin \alpha, \quad C_D = C_n \sin \alpha + C_a \cos \alpha, \quad (3.8)$$

where α is the angle of attack. It should be noted that this approach yields the pressure drag component only and does not include skin-friction drag. The resulting C_D values are therefore not directly comparable to the total drag obtained from the wake rake, and are used here primarily as a consistency check alongside the C_p distributions.

3.4.2. Surface Pressure Tap Test Matrix

The surface pressure tap campaign covers the same two configurations as the PIV, RANS, and LES campaigns, Smooth and Mid-2.50, tested across the full angle-of-attack and Reynolds number ranges of the RWT, as summarised in Table 3.6. This allows direct comparison of the measured C_p distributions with the numerical surface pressure predictions at matching conditions.

Table 3.6: Operating conditions used in the surface pressure tap campaign. The three Reynolds numbers correspond to inlet velocities of $U_\infty = 7.5 \text{ m/s}$, 15 m/s , and 37.5 m/s , using $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$ and $c = 200 \text{ mm}$.

Parameter	Values
Angle of attack α	$[-10^\circ, 20^\circ]$ in steps of 2°
Reynolds number Re	1.0×10^5 , 2.0×10^5 , 5.0×10^5
Inlet velocity U_∞	7.5 m/s , 15 m/s , 37.5 m/s
Flow regime	Incompressible

Combining two geometry configurations, sixteen angles of attack, and three Reynolds numbers gives a total of 96 measurement runs, as summarised in Table 3.7.

Table 3.7: Summary of surface pressure tap case count.

Item	Dimpled cases	Reference cases
Geometry configurations	1	1
Dimple depth levels	1	—
Angles of attack	16	16
Reynolds numbers	3	3
Number of cases	48	48
Total cases		96

3.4.3. Summary of Surface Pressure Tap Setup

The surface pressure tap campaign consisted of 96 measurement runs across two model variants at three Reynolds numbers and sixteen angles of attack. Measurements were performed in the Red Wind Tunnel at DTU using the same pressure acquisition system as the wake and wall pressure campaign, with the 32 wake rake connections replaced by the 32 surface tap hoses. The key parameters are consolidated in Table 3.8.

Table 3.8: Summary of surface pressure tap measurement setup and campaign parameters.

Parameter	Value
<i>Facility and technique</i>	
Wind tunnel	Red Wind Tunnel (RWT), DTU
Technique	Surface pressure taps
Freestream turbulence	$Tu < 0.1 \%$
<i>Operating conditions</i>	
Chord length c	200 mm
Reynolds numbers Re	1.0×10^5 , 2.0×10^5 , 5.0×10^5
Inlet velocities U_∞	7.5 m/s, 15 m/s, 37.5 m/s
Angles of attack α	$[-10^\circ, 20^\circ]$ in steps of 2°
Kinematic viscosity ν	$1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$
<i>Instrumentation</i>	
Tap diameter	0.7 mm, normal to surface
Suction-side taps	20 (leading edge to $x/c = 0.901$)
Pressure-side taps	12 ($x/c = 0.010$ to 0.800)
Hose inner diameter	$\varnothing 1 \text{ mm}$ rubber
Pressure system	Shared with wake rake campaign
<i>Study scope</i>	
Geometry configurations	2 (Smooth PT, Mid-2.50 PT)
Total measurement runs	96

Methodology II - Flow Visualisation

4.1. Wind Tunnel Facility

All experiments in this chapter were conducted in the Small Low-Turbulence Tunnel (SLT) at the Low-Speed Laboratory of TU Delft, a low-turbulence wind tunnel operated in open-loop configuration [16], depicted in Figure 4.1. The tunnel is powered by a 75 kW axial fan and incorporates seven anti-turbulence screens in the settling plenum, followed by a contraction with an area ratio of 11:1. The contraction outlet has a rectangular cross-section of $0.9 \times 0.6 \text{ m}^2$, and the working section downstream is constructed from 12 mm thick transparent polycarbonate panels, providing full optical access for laser-sheet illumination and PIV imaging. For the present experiments, a single 1.8 m-long test section module was used. The freestream turbulence intensity has been characterised as $Tu \approx 0.03\text{--}0.07\%$ across the operating velocity range [16], making the facility well suited to measurements of boundary-layer flows sensitive to inflow disturbances.



Figure 4.1: The Small Low-Turbulence Tunnel (SLT) in the Low Speed Lab at TU Delft. The axial fan and motor assembly are visible on the right, followed by the flow conditioning section housing the anti-turbulence screens. The contraction section (wooden structure, centre) accelerates the flow into the rectangular test section. The tunnel is shown in its open-loop configuration, without the test section present, before the measurement campaigns.

The freestream velocity U_∞ was determined from the pressure difference measured across the tunnel contraction using a differential pressure transducer. Applying the incompressible Bernoulli equation between the settling chamber (station 1) and the contraction outlet (station 2), combined with continuity, gives

$$U_\infty = \sqrt{\frac{2 \Delta p}{\rho \left[1 - \left(\frac{A_2}{A_1} \right)^2 \right]}}, \quad (4.1)$$

where $\Delta p = p_1 - p_2$ is the measured static pressure difference, ρ is the ambient air density, and A_1 , A_2 are the cross-sectional areas of the settling chamber and contraction outlet, respectively. Given the

contraction area ratio $A_1/A_2 = 11$, the correction term $(A_2/A_1)^2 = 1/121 \approx 0.008$ is negligible, and Equation 4.1 reduces in practice to $U_\infty = \sqrt{2 \Delta p / \rho}$. The air density was estimated from laboratory temperature and barometric pressure, assuming ideal-gas behaviour.

4.2. Airfoil Models and Setup

The test models, presented in Figure 4.2 and Table 4.1, consist of a smooth baseline airfoil and a set of modified geometries featuring surface dimples of varying depth. The baseline corresponds to the smooth NACA 63-415 profile (denoted Smooth), while the modified configuration (denoted Mid) incorporates dimples at three depths: $k = 1.25$ mm, 2.50 mm, and 5.00 mm, yielding four model variants in total: Smooth, Mid-1.25, Mid-2.50, and Mid-5.00. All models share a chord length of $c = 200$ mm. The models are the same physical airfoils used in the DTU Red Wind Tunnel campaign described in Chapter 3, but extended to a total span of 600 mm using smooth NACA 63-415 end extensions to match the full width of the SLT test section, as depicted in Figure A.3.

Table 4.1: Geometry configurations used in the SLT campaign. The numeric suffix denotes the fixed dimple depth in millimetres for $c = 200$ mm.

ID	Rows	x/c range	D_d/c	k/c	b/c
Smooth	–	–	–	–	3
Mid-1.25	1–4 (4/7)	0.33–0.59	0.05	0.0063	3
Mid-2.50	1–4 (4/7)	0.33–0.59	0.05	0.0125	3
Mid-5.00	1–4 (4/7)	0.33–0.59	0.05	0.0250	3

Each model is mounted vertically in the test section via a central support rod connected to a three-component force balance through a custom-designed adapter (see Appendix Figure A.1). The force balance serves a dual purpose: it is the primary structural support for the airfoil and simultaneously measures the aerodynamic loads acting on the model. Continuous rotation of the force balance adapter allows accurate and repeatable setting of the angle of attack α throughout the measurement campaign. Because the force balance is the common mounting hardware for all experiments in this chapter — PIV, force measurements, and tuft visualisation — it is described here once and referenced in the subsequent sections.




Figure 4.2: CAD geometries of the four airfoil configurations. All dimpled configurations use inward spherical dimples with differing dimple depths.

4.3. Particle Image Velocimetry

This section describes the Particle Image Velocimetry (PIV) investigation conducted on the reference airfoil and selected configurations with surface dimples. The objective is to resolve the flow structures associated with boundary-layer development, separation, and reattachment, and to assess how these change with the adaptive surface geometry at the angles of attack of interest.

4.3.1. Experimental Setup

A planar PIV system was assembled to measure the two-component velocity field in the mid-span plane of the NACA 63-415 airfoil and its dimpled variants. The main components are: a seeding generator, a double-pulse Nd:YAG laser with sheet-forming optics, a high-resolution sCMOS camera with lens, and DaVis 8 software for image acquisition and cross-correlation processing.

The laser sheet was oriented horizontally at the mid-span height of the vertically mounted airfoil, illuminating the suction surface from leading to trailing edge. The camera was positioned directly above the test section, oriented vertically downward and normal to the laser sheet, imaging the suction-side flow field in plan view. This arrangement ensures that the measurement plane coincides with the airfoil mid-span, where the flow is expected to be most two-dimensional.




Figure 4.3: CAD render of the PIV experimental setup. (a) Full assembly inside the test section. (b) Measurement assembly with the test section removed, showing the force balance adapter and the relative geometry of the laser sheet and camera field of view.

Seeding

Flow seeding was provided by a SAFEX fog generator producing water–glycol droplets with a mean diameter of approximately $1\ \mu\text{m}$. The seeding concentration was adjusted to yield sufficient particle image density in the interrogation windows without introducing any meaningful aerodynamic disturbance to the boundary layer under investigation.

Laser

The illumination was delivered by a Quantel Evergreen 200 double-pulsed Nd:YAG laser generating a laser with wavelength $\lambda = 532 \text{ nm}$ through a second-harmonic conversion of the 1064 nm main source. Each pulse carries up to 200 mJ of energy with an 8 ns duration and a maximum repetition rate of 15 Hz . The beam was expanded by a set of cylindrical lenses into a thin light sheet oriented horizontally at the mid-span location of the model, so that it illuminates the suction surface of the airfoil from leading to trailing edge. Synchronisation of the laser pulses and camera exposures was managed by a programmable timing unit (PTU) controlled through the DaVis 8 interface.

Camera and Lens

Particle images were recorded with a LaVision Imager sCMOS CLHS camera, featuring 2560×2160 pixels at a pixel pitch of $p_{px} = 6.5 \mu\text{m}$ and 16-bit dynamic range. A Nikon AF Micro-Nikkor 60 mm focal-length lens provided the required magnification and depth-of-field control. DaVis 8 (LaVision) managed hardware control, image acquisition, and cross-correlation processing; interrogation window size, overlap, iteration scheme, and vector-validation settings were all configured through its interface.

Optical Setup Parameters

To capture the full chord length with margin for the near-wake, the field of view was set to 1.5 times the chord length:

$$\text{FOV} = 1.5 c = 1.5 \times 200 \text{ mm} = 300 \text{ mm}. \quad (4.2)$$

The system magnification M is the ratio of the physical sensor size to the field of view:

$$M = \frac{\text{Sensor size}}{\text{FOV}} = \frac{N_{\text{px}} \times p_{\text{px}}}{\text{FOV}} = \frac{2560 \times 6.5 \mu\text{m}}{300 \text{ mm}} = \frac{16.640 \text{ mm}}{300 \text{ mm}} \approx 0.0555. \quad (4.3)$$

The object distance d_o and image distance d_i follow from the thin lens equation $1/f = 1/d_o + 1/d_i$, with the substitution $d_i = M d_o$:

$$d_o = \frac{f(1+M)}{M} = \frac{60(1+0.0555)}{0.0555} \approx 1141 \text{ mm}, \quad d_i = M d_o \approx 63.3 \text{ mm}. \quad (4.4)$$

Aperture Selection

To avoid peak-locking errors, the particle image diameter d_τ should lie above approximately 1.5 pixels, with the range 2–3 pixels regarded as optimal [28]. The total particle image diameter combines a diffraction contribution and a geometric contribution in quadrature:

$$d_\tau = \frac{\sqrt{d_{\text{diff}}^2 + d_{\text{geo}}^2}}{p}, \quad d_{\text{diff}} = 2.44 \lambda f_\# (1+M), \quad d_{\text{geo}} = M d_p, \quad (4.5)$$

with $\lambda = 532 \text{ nm}$ and seeding particle diameter $d_p = 1 \mu\text{m}$. Rearranging for the $f_\#$ that yields the target diameter of 2 px (i.e. $\Delta d_\tau = 2 \times 6.5 \mu\text{m} = 13 \mu\text{m}$):

$$f_\# = \frac{\sqrt{\Delta d_\tau^2 - (M d_p)^2}}{2.44 \lambda (1+M)} = \frac{\sqrt{(13.0 \times 10^{-6})^2 - (0.0555 \times 10^{-6})^2}}{2.44 \times 532 \times 10^{-9} \times 1.0555} \approx 9.5. \quad (4.6)$$

The nearest available stops are $f_\# = 8$ and $f_\# = 11$, neither of which reaches the optimal band. Of the two, $f_\# = 8$ admits roughly twice the light and was selected on that basis, since the available laser energy and seeding density set the binding constraint on image quality in this setup. It yields a particle image diameter of

$$d_\tau \approx 1.69 \text{ px}. \quad (4.7)$$

This value remains above the 1.5 px lower bound recommended by Sciacchitano [28] as the threshold below which peak-locking bias becomes significant. The wider aperture was preferred to maximise the

light admitted to the sensor, which was necessary given the large object distance and the low scattering cross-section of the $1\text{ }\mu\text{m}$ seeding particles.

Pulse Separation

According to the one-quarter rule [28], the particle displacement between two laser pulses must not exceed $w/4$, where w is the interrogation window width, to ensure a sufficient fraction of particle pairs are retained within each window. For the 32×32 pixel window used here, this gives a maximum displacement of 8 px. The required pulse separation follows from

$$\Delta t = \frac{(w/4) p_{px}}{M U_\infty}, \quad (4.8)$$

giving

$$\Delta t|_{U_\infty=7.5\text{ m/s}} = \frac{8 \times 6.5 \times 10^{-6}}{0.0555 \times 7.5} \approx 125\text{ }\mu\text{s}, \quad \Delta t|_{U_\infty=15\text{ m/s}} = \frac{8 \times 6.5 \times 10^{-6}}{0.0555 \times 15} \approx 62.5\text{ }\mu\text{s}. \quad (4.9)$$

In practice, pulse separations of $\Delta t = 124\text{ }\mu\text{s}$ and $62\text{ }\mu\text{s}$ were used for $U_\infty = 7.5$ and 15 m/s respectively, both yielding a maximum particle displacement of approximately 8 px, consistent with the one-quarter rule. A summary of all optical and acquisition parameters is provided in Table 4.4.

4.3.2. PIV Test Matrix

The PIV campaign was restricted to the four model variants summarised in Table 4.1 and illustrated in Figure 4.2: the undimpled Smooth baseline and the mid-chord dimpled configuration Mid at three dimple depths, 1.25, 2.50, and 5.00 mm. The Mid-2.50 configuration was selected to match the LES and RANS reference cases, while the additional depths of 1.25 and 5.00 mm allow the effect of dimple depth to be isolated at constant chordwise placement.

The operating condition was selected to match the available LES and RWT datasets. Both Reynolds numbers were tested across five angles of attack covering attached flow, the onset of separation, and separated conditions, as summarised in Table 4.2.

Table 4.2: Operating conditions used in the PIV campaign. The two Reynolds numbers correspond to inlet velocities of $U_\infty = 7.5\text{ m/s}$ and 15 m/s , using $\nu = 1.5 \times 10^{-5}\text{ m}^2/\text{s}$ and $c = 200\text{ mm}$.

Parameter	Values
Angle of attack α	$0^\circ, 4^\circ, 6^\circ, 10^\circ, 14^\circ$
Reynolds number Re	$1.0 \times 10^5, 2.0 \times 10^5$
Inlet velocity U_∞	$7.5\text{ m/s}, 15\text{ m/s}$
Flow regime	Incompressible

The two freestream velocities $U_\infty = 7.5\text{ m/s}$ and 15 m/s correspond to $Re = 1.0 \times 10^5$ and 2.0×10^5 respectively, matching conditions available in both the RWT and LES databases, and allowing direct comparison with the RANS results at $Re = 2.0 \times 10^5$. The resulting PIV case count is 40 measurement sets in total as shown in Table 4.3 below.

Table 4.3: Summary of PIV case count.

Item	Dimpled cases	Reference cases
Geometry configurations	1	1
Dimple depth levels	3	—
Angles of attack	5	5
Reynolds numbers	2	2
Number of cases	30	10
Total PIV cases		40

4.3.3. Summary of PIV Setup

The PIV campaign consisted of 40 measurement sets across four model variants at two Reynolds numbers and five angles of attack. Measurements were performed in the SLT wind tunnel at TU Delft using a planar PIV system with a LaVision sCMOS camera and a Quantel Evergreen 200 Nd:YAG laser. The key instrumentation and acquisition parameters are consolidated in Table 4.4. The PIV results are used to resolve the suction-side flow structures and compare boundary-layer development, separation onset, and reattachment behaviour across the smooth and dimpled configurations.

The optical setup parameters summarised above, field of view, magnification, aperture, and pulse separation, were derived using a PIV setup calculator developed as part of this thesis to help carry out the PIV setup. The tool is made available online for others planning similar PIV campaigns, at <https://aaboe97.github.io/piv-setup-calculator/> (source: <https://github.com/Aaboe97/piv-setup-calculator>).

Table 4.4: Summary of PIV measurement setup and campaign parameters.

Parameter	Value
<i>Facility and technique</i>	
Wind tunnel	SLT, TU Delft
Technique	Planar PIV
Measurement plane	Mid-span suction side, horizontal
Freestream turbulence	$Tu \approx 0.03\text{--}0.07\%$
<i>Operating conditions</i>	
Chord length c	200 mm
Reynolds numbers Re	1.0×10^5 , 2.0×10^5
Inlet velocities U_∞	7.5 m/s, 15 m/s
Angles of attack α	0° , 4° , 6° , 10° , 14°
Kinematic viscosity ν	$1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$
<i>Optical setup</i>	
Camera	LaVision Imager sCMOS CLHS
Lens	Nikon AF Micro-Nikkor 60 mm
Sensor resolution	$2560 \times 2160 \text{ px}$
Pixel pitch p_{px}	$6.5 \mu\text{m}$
Field of view	300 mm ($1.5c$)
Magnification M	0.0555
Object distance d_o	$\approx 1141 \text{ mm}$
Image distance d_i	$\approx 63.3 \text{ mm}$
Aperture $f_\#$	8
Particle image diameter d_τ	$\approx 1.69 \text{ px}$
<i>Acquisition</i>	
Laser	Quantel Evergreen 200 Nd:YAG, 532 nm
Pulse energy	200 mJ
Interrogation window w	$32 \times 32 \text{ px}$
Pulse separation (7.5 m/s)	$124 \mu\text{s}$
Pulse separation (15 m/s)	$62 \mu\text{s}$
Seeding	SAFEX fog generator, $\approx 1 \mu\text{m}$ droplets
<i>Study scope</i>	
Geometry configurations	4 (Smooth, Mid-1.25, Mid-2.50, Mid-5.00)
Total measurement sets	40

4.4. Force Balance Measurements

This section describes the dedicated force balance measurement campaign conducted on all four model variants in the SLT wind tunnel. The objective was to obtain aerodynamic force polars — lift coefficient C_L and drag coefficient C_D as a function of angle of attack — across the two Reynolds numbers tested in the PIV campaign, enabling direct comparison of integrated aerodynamic performance between the smooth baseline and the three dimpled depth configurations.

4.4.1. Experimental Setup

The airfoil models, mounting arrangement, and span extensions are identical to those described in Section 4.2 and used in the PIV campaign. The force balance serves as both the structural mount and the load measurement device. The airfoil is rotated automatically to each angle of attack by the rotary stage of the force balance assembly, as illustrated in Figure 4.4. The three-component sensor records forces in three directions simultaneously, of which the streamwise component F_x (drag) and the component normal to the freestream F_y (lift) are used in the present analysis.

The lift and drag coefficients are defined as

$$C_L = \frac{F_y}{q_\infty S_{\text{ref}}}, \quad (4.10)$$

$$C_D = \frac{F_x}{q_\infty S_{\text{ref}}}, \quad (4.11)$$

where F_y and F_x are the time-averaged normal and streamwise force components measured by the balance, $q_\infty = \frac{1}{2}\rho U_\infty^2$ is the freestream dynamic pressure determined from the tunnel contraction pressure drop as described in Section 4.1, and $S_{\text{ref}} = c \times b = 0.200 \times 0.600 = 0.12 \text{ m}^2$ is the reference planform area based on the full model span including the end extensions.




Figure 4.4: CAD render of the Mid-2.50 airfoil mounted on the three-component force balance assembly at TU Delft. The cross-shaped base plate is bolted to the tunnel floor, and the rotary stage directly below the airfoil sets the angle of attack automatically. The stainless steel support rod connects the airfoil to the force balance sensor at the quarter-chord point. The dimple pattern is visible on the suction surface of the model.

4.4.2. Force Balance Test Matrix

The force balance campaign covers the same four model variants as the PIV campaign; Smooth, Mid-1.25, Mid-2.50, and Mid-5.00, at the same two Reynolds numbers, but sweeps a wider angle-of-attack range to capture the full polar from attached flow through stall, as summarised in Table 4.5.

Table 4.5: Operating conditions used in the force balance campaign.

Parameter	Values
Angle of attack α	$[0^\circ, 16^\circ]$ in steps of 2°
Reynolds number Re	1.0×10^5 , 2.0×10^5
Inlet velocity U_∞	7.5 m/s, 15 m/s
Flow regime	Incompressible

The angle-of-attack range extends from 0° to 16° in steps of 2° , covering the attached, transitional, and fully separated flow regimes. Combining four model variants, two Reynolds numbers, and nine angles of attack gives a total of 72 measurement runs, as summarised in Table 4.6.

Table 4.6: Summary of force balance case count.

Item	Dimpled cases	Reference cases
Geometry configurations	1	1
Dimple depth levels	3	–
Angles of attack	9	9
Reynolds numbers	2	2
Number of cases	54	18
Total force balance cases		72

4.4.3. Summary of Force Balance Setup

The force balance campaign consisted of 72 measurement runs across four model variants at two Reynolds numbers and nine angles of attack. Lift and drag were obtained directly from the streamwise and normal force components of the three-component balance, normalised by the freestream dynamic pressure and the full planform reference area. The key parameters are consolidated in Table 4.7.

Table 4.7: Summary of force balance measurement setup and campaign parameters.

Parameter	Value
<i>Facility and technique</i>	
Wind tunnel	SLT, TU Delft
Technique	Three-component force balance
AoA setting	Automatic rotary stage
Freestream turbulence	$Tu \approx 0.03\text{--}0.07\%$
<i>Operating conditions</i>	
Chord length c	200 mm
Span b	600 mm (incl. end extensions)
Reference area S_{ref}	0.12 m^2
Reynolds numbers Re	1.0×10^5 , 2.0×10^5
Inlet velocities U_∞	7.5 m/s, 15 m/s
Angles of attack α	0° to 16° in steps of 2°
Kinematic viscosity ν	$1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$
<i>Data reduction</i>	
Lift coefficient	$C_L = F_y / (q_\infty S_{\text{ref}})$
Drag coefficient	$C_D = F_x / (q_\infty S_{\text{ref}})$
Dynamic pressure q_∞	From tunnel contraction Δp
<i>Study scope</i>	
Geometry configurations	4 (Smooth, Mid-1.25, Mid-2.50, Mid-5.00)
Total measurement runs	72

4.5. Tuft Flow Visualisation

Following the force balance campaign, a qualitative tuft study was conducted to obtain a global picture of the surface flow topology on two selected configurations. Tufts provide a simple and effective means of identifying regions of attached flow, separation, and spanwise non-uniformities that would otherwise require a full PIV survey to detect.

4.5.1. Experimental Setup

Tufts were applied to the suction surface of the Smooth and Mid-2.50 models, selected to match the configurations available in the LES, PIV, RANS, and wake rake datasets, and tested at $U_\infty = 15 \text{ m/s}$ ($Re = 2.0 \times 10^5$) in the SLT wind tunnel immediately following the force balance measurements. Individual tufts were cut from yellow yarn to a length of approximately $0.2c = 40 \text{ mm}$ and fixed to the surface using thin specialised aero tape at three chordwise rows located at $x/c \approx 0.15, 0.45$, and 0.75 . Spanwise, seven columns were distributed symmetrically about the mid-span at positions $z = 0, \pm 60, \pm 120$, and $\pm 180 \text{ mm}$, supplemented by one additional column of three tufts at each span end on the surface of the end extensions, to monitor wall-proximity effects and end-condition disturbances. This yields a total of $(7 + 2) \times 3 = 27$ tufts per model, applied at identical locations on both Smooth and Mid-2.50. The tuft layout on the Mid-2.50 model prior to testing is shown in Figure 4.5.




Figure 4.5: Suction surface of the Mid-2.50 airfoil with tufts applied prior to wind tunnel testing. Yellow yarn tufts of length $\approx 0.2c = 40 \text{ mm}$ are arranged in three chordwise rows at $x/c \approx 0.15, 0.45$, and 0.75 , and nine spanwise columns including the single columns at each span end. Leading edge and trailing edge are denoted *LE* and *TE* respectively.

Two angle-of-attack sweeps were performed for each model. First, a continuous pitch from 0° to 20° and back to 0° was conducted to capture the dynamic evolution of the surface flow topology through stall onset and recovery. Second, a quasi-static sweep from 0° to 16° in steps of 2° was performed, with the model held at each angle for a sufficient duration to allow the flow to reach a steady state. Both sweeps were recorded using an iPhone camera mounted outside the test section, imaging through the transparent polycarbonate side walls. The resulting video footage and still images provide a qualitative record of the chordwise and spanwise extent of attached and separated flow regions at each condition.

4.5.2. Summary of Tuft Setup

Table 4.8: Summary of tuft flow visualisation setup and campaign parameters.

Parameter	Value
<i>Facility and technique</i>	
Wind tunnel	SLT, TU Delft
Technique	Surface tuft visualisation
Imaging	iPhone camera, external view through polycarbonate wall
Freestream turbulence	$T_u \approx 0.03\text{--}0.07\%$
<i>Operating conditions</i>	
Reynolds number Re	2.0×10^5
Inlet velocity U_∞	15 m/s
Continuous sweep	$0^\circ \rightarrow 20^\circ \rightarrow 0^\circ$
Step sweep	0° to 16° in steps of 2°
<i>Tuft arrangement</i>	
Tuft material	Yellow yarn, ≈ 1 mm diameter
Tuft length	≈ 40 mm ($0.2c$)
Chordwise rows	3 (at $x/c \approx 0.15, 0.45, 0.75$)
Spanwise columns	9 (7 mid-span + 1 per span end)
Total tufts per model	27
<i>Study scope</i>	
Geometry configurations	2 (Smooth, Mid-2.50)

5

Methodology III - Computational Fluid Dynamics

The numerical investigation was conducted to complement the experimental campaign and provide additional insight into the flow physics associated with the dimpled airfoil configurations. Numerical simulations allow systematic variation of geometric parameters while providing insight into flow quantities that are difficult to measure experimentally, including boundary-layer development, separation behaviour, surface pressure distributions, and wake structures.

Two different numerical approaches were employed. A Reynolds-Averaged Navier–Stokes (RANS) study was performed to investigate a broad range of dimple configurations and operating conditions, enabling efficient comparison of geometric trends across the design space defined in Section 2.8. The primary objective was to identify trends in aerodynamic performance and select representative cases for further investigation.

Subsequently, Large-Eddy Simulations (LES) were conducted for a limited number of selected configurations. Unlike the RANS study, the LES calculations explicitly resolved the dominant unsteady flow structures associated with boundary-layer separation, wake development, and the interaction between the dimple cavities and the outer flow. This provided a higher-fidelity representation of the flow physics governing separation and stall.

The following sections describe the numerical methods, computational setups and domains, mesh generation procedures, boundary conditions, and simulation matrices used for the RANS and LES investigations.

5.1. Reynolds Averaged Navier-Stokes

This section presents the Reynolds-Averaged Navier–Stokes screening study conducted to evaluate the aerodynamic influence of chordwise dimple placement on the NACA 63–415 airfoil. The simulations used the final design parameters established in Section 2.8, including the airfoil profile, chord length, dimple diameter, dimple spacing, and selected dimple placement region. The purpose of the RANS campaign was therefore not to redefine the geometry, but to compare the aerodynamic response of the geometry families defined in Table 2.5 under consistent numerical conditions.

The study evaluates the four airfoil geometries laid out in Section 2.8, a smooth baseline and three dimpled layouts, across a range of angles of attack and Reynolds numbers representative of the wind-tunnel test conditions. The section first defines the RANS test matrix and the simulated geometry configurations. It then describes the computational setup, turbulence modelling approach, mesh generation strategy, and mesh verification procedure used for the screening campaign.

5.1.1. Study Design and Test Matrix

A structured RANS test matrix was defined to isolate the influence of chordwise dimple placement on the aerodynamic performance of the NACA 63–415 airfoil. The baseline airfoil geometry, chord length, dimple diameter, dimple spacing, and row locations were established in Section 2.8. The present RANS study therefore does not vary the individual dimple dimensions or the dimple depth. Instead, it compares the aerodynamic response of different chordwise coverage patterns at one fixed dimple depth.

The four simulated configurations are summarised in Table 5.1 and illustrated in Figure 5.1. The Smooth geometry represents the undimpled NACA 63–415 reference case and provides the baseline against which all dimpled configurations are compared. The dimpled cases are denoted Mid-2.50, Aft-2.50, and Full-2.50, where the numeric suffix indicates a fixed dimple depth of $k = 2.50$ mm, corresponding to $k/c = 0.0125$ for the selected chord length of $c = 200$ mm.

Table 5.1: Geometry configurations used in the RANS screening study. The numeric suffix denotes a fixed dimple depth of $k = 2.50$ mm, corresponding to $k/c = 0.0125$ for $c = 200$ mm.

ID	Rows	x/c range	D_d/c	k/c	b/c
Smooth	–	–	–	–	0.2
Mid-2.50	1–4 (4/7)	0.33–0.59	0.05	0.0125	0.2
Aft-2.50	4–7 (4/7)	0.59–0.85	0.05	0.0125	0.2
Full-2.50	1–7 (7/7)	0.33–0.85	0.05	0.0125	0.2

Geometry Mid-2.50 contains rows 1–4 and represents a mid-chord coverage region from $x/c = 0.33$ to 0.59. Geometry Aft-2.50 contains rows 4–7 and isolates the aft-chord region from $x/c = 0.59$ to 0.85, where separation-control effects are expected to be most relevant. Geometry Full-2.50 combines all seven rows and covers the full investigated suction-side region from $x/c = 0.33$ to 0.85.




Figure 5.1: CAD geometries of the four airfoil configurations used in the RANS screening study. Smooth is the undimpled NACA 63–415 baseline. Mid-2.50 covers the mid-chord region, Aft-2.50 covers the aft region, and Full-2.50 covers the full investigated suction-side region from $x/c = 0.33$ to 0.85. All dimpled configurations use inward spherical dimples with $D_d/c = 0.05$, $k/c = 0.0125$, and staggered spacing $D_s/D_d = 2.0$. All computational geometries have a span of $b/c = 0.2$.

All dimpled configurations use the same inward spherical dimple geometry, with $D_d/c = 0.05$, $k/c = 0.0125$, and staggered spacing $D_s/D_d = 2.0$, as defined in Section 2.7.4. The computational models are extruded to a span of $0.2c$. By holding dimple size, depth, spacing, and span constant, the test matrix isolates chordwise dimple placement as the primary geometric variable.

The operating conditions were selected to match the design Reynolds numbers established in Section 2.7.6 and to cover the progression from attached flow to post-stall conditions. The Reynolds numbers $Re = 2.0 \times 10^5$ and $Re = 5.0 \times 10^5$ correspond to two of the primary experimental operating conditions and remain within the Reynolds-number range reported for comparable dimpled-airfoil studies in Table 2.2. The angle-of-attack range was selected to include low-angle attached flow, the onset of separation, and the near-stall regime, where the influence of the dimple configurations is expected to be most pronounced.

Table 5.2: Operating conditions used in the RANS screening study. The two Reynolds numbers correspond to inlet velocities of $U_\infty = 15.0 \text{ m/s}$ and 37.5 m/s , respectively, using $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$ and $c = 200 \text{ mm}$.

Parameter	Values
Angle of attack α	$0^\circ, 6^\circ, 10^\circ, 12^\circ, 14^\circ, 16^\circ$
Reynolds number Re	$2.0 \times 10^5, 5.0 \times 10^5$
Flow regime	Incompressible

Two further cases, Smooth and Mid-2.50 at $Re = 1 \times 10^5$ and $\alpha = 6^\circ$, were run subsequently using an identical numerical setup. These were carried out solely to provide a matching RANS flow field for the experimental and numerical comparison presented in Section 7.6, and are additional to the screening matrix defined here.

The resulting operating conditions are summarised in Table 5.2. Combining four geometry configurations, six angles of attack, and two Reynolds numbers gives a total of 48 RANS simulations, as summarised in Table 5.3. The matrix provides sufficient coverage of the primary operating parameters while maintaining a computational cost appropriate for a screening study.

Table 5.3: Summary of RANS case count.

Item	Dimpled cases	Reference cases
Geometry configurations	3	1
Dimple depth levels	1	–
Angles of attack	6	6
Reynolds numbers	2	2
Number of cases	36	12
Total RANS cases		48

5.1.2. Computational Setup

The aerodynamic performance of the investigated airfoil configurations is evaluated using Reynolds-Averaged Navier–Stokes (RANS) simulations. The simulations are performed using the finite-volume solver `pimpleFoam` in `OpenFOAM`. A pseudo-transient approach is employed, where the governing equations are advanced in time until a steady-state solution is obtained.

Governing Equations and Solver

The flow is assumed incompressible, and the instantaneous velocity and pressure fields are decomposed into mean and fluctuating parts according to the Reynolds decomposition,

$$u_i = \bar{u}_i + u'_i, \quad p = \bar{p} + p', \quad (5.1)$$

where $\bar{u}_i$ and $\bar{p}$ denote time-averaged quantities, and u'_i and p' represent the turbulent fluctuations. Substituting Equation 5.1 into the incompressible Navier–Stokes equations and applying Reynolds averaging yields the incompressible Reynolds-averaged Navier–Stokes (RANS) equations,

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0, \quad (5.2)$$

$$\frac{\partial \bar{u}_i}{\partial t} + \bar{u}_j \frac{\partial \bar{u}_i}{\partial x_j} = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \nu \frac{\partial^2 \bar{u}_i}{\partial x_j \partial x_j} - \frac{\partial \overline{u'_i u'_j}}{\partial x_j}. \quad (5.3)$$

The additional term $-\partial \overline{u'_i u'_j} / \partial x_j$ in Equation 5.3 is the divergence of the Reynolds stress tensor, representing the momentum transport induced by turbulent fluctuations. Its appearance introduces the classical *closure problem*, since the number of unknowns now exceeds the number of governing equations.

To close the system, the Reynolds stresses are modelled using the Boussinesq eddy-viscosity hypothesis. For incompressible flow this takes the form

$$-\overline{u'_i u'_j} = \nu_t \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right) - \frac{2}{3} k \delta_{ij}, \quad (5.4)$$

where ν_t is the turbulent eddy viscosity, $k = \frac{1}{2} \overline{u'_i u'_i}$ is the turbulent kinetic energy, and δ_{ij} is the Kronecker delta. The component $-\rho \overline{u'v'}$ introduced in Equation 2.4 is the $i = 1, j = 2$ entry of this tensor, which is the dominant one in the two-dimensional boundary layer considered there. In this chapter k denotes the turbulent kinetic energy, following standard notation for two-equation turbulence models. It is distinct from the dimple depth k defined in Section 2.5.1, which appears elsewhere in this thesis as a physical length, and in the normalised forms k/c and k/D_a . The two quantities never occur within the same expression. The isotropic term $-\frac{2}{3} k \delta_{ij}$ ensures that the trace of the modelled stress tensor is consistent with the definition of k . The closure problem then reduces to determining ν_t , which is the task of the turbulence model described in Section 5.1.3.

The mean-flow equations are solved using the finite-volume solver `pimpleFoam` in `OpenFOAM`. Although formulated as a transient solver, it is employed here in a pseudo-transient manner: the solution is advanced iteratively until the residuals and integral force coefficients have converged to a steady state.

5.1.3. Turbulence Modelling

To close the Reynolds-averaged momentum equations, a two-equation eddy-viscosity model is employed. The purpose of a two-equation model is to represent the turbulent properties of the flow by solving additional transport equations for a velocity scale and a turbulence length- or time-scale. This is an important improvement over algebraic models, since turbulence is not fixed locally but is convected and diffused throughout the domain. In the present work, turbulence is modelled using the k – ω shear stress transport (SST) model of Menter [20], which is based on the Boussinesq hypothesis introduced in Equation 5.4, relating the Reynolds stresses to the mean strain rate through a turbulent eddy viscosity. To motivate this choice, the standard k – ω and k – ε models are first introduced, as the SST model draws directly on both formulations.

Turbulent Scales

The first transported quantity common to all two-equation models is the turbulent kinetic energy,

$$k = \frac{1}{2} \overline{u'_i u'_i}, \quad (5.5)$$

which represents the kinetic energy contained in the turbulent velocity fluctuations and serves as the turbulence velocity scale.

The second transported quantity differs between models. In the k - ε model it is the turbulent dissipation rate ε , defined as the rate at which turbulent kinetic energy is converted into thermal internal energy ($\text{m}^2 \text{s}^{-3}$). In the k - ω model it is the specific dissipation rate ω , interpreted as the mean frequency of the turbulence and therefore as a measure of its decay rate. The two quantities are related through

$$\omega = \frac{\varepsilon}{C_\mu k}, \quad (5.6)$$

where C_μ is a model constant. From dimensional analysis, k and ω define a turbulence time-scale and length scale,

$$T \sim \frac{1}{\omega}, \quad \ell \sim \frac{\sqrt{k}}{\omega}, \quad (5.7)$$

and consequently a turbulent eddy viscosity,

$$\nu_t \sim \frac{k}{\omega}. \quad (5.8)$$

Turbulence Model k - ε

The standard k - ε model, from Launder and Sharma [17], describes turbulence through transport equations for k and the dissipation rate ε . This allows turbulence to be convected, diffused, produced, and dissipated throughout the domain, rather than being prescribed algebraically.

The transport equation for k reads

$$\frac{\partial(\rho k)}{\partial t} + \frac{\partial(\rho \bar{u}_i k)}{\partial x_i} = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_k} \right) \frac{\partial k}{\partial x_j} \right] + P_k - \rho \varepsilon, \quad (5.9)$$

where ρ is the fluid density, μ is the molecular dynamic viscosity, μ_t is the turbulent dynamic viscosity, σ_k is the turbulent Prandtl number for k , and P_k is the production of turbulent kinetic energy due to mean velocity gradients. On the right-hand side of Equation 5.9, the first term is the combined molecular and turbulent diffusion of k , the second is its production, and $-\rho \varepsilon$ is the viscous dissipation.

The corresponding transport equation for ε is

$$\frac{\partial(\rho \varepsilon)}{\partial t} + \frac{\partial(\rho \bar{u}_i \varepsilon)}{\partial x_i} = \frac{\partial}{\partial x_j} \left[\left(\mu + \frac{\mu_t}{\sigma_\varepsilon} \right) \frac{\partial \varepsilon}{\partial x_j} \right] + C_{1\varepsilon} \frac{\varepsilon}{k} P_k - C_{2\varepsilon} \rho \frac{\varepsilon^2}{k}, \quad (5.10)$$

where σ_ε is the turbulent Prandtl number for ε , and $C_{1\varepsilon}$ and $C_{2\varepsilon}$ are model constants. The production term $C_{1\varepsilon}(\varepsilon/k)P_k$ links the generation of dissipation to the local turbulence production, while $-C_{2\varepsilon}\rho \varepsilon^2/k$ represents its self-destruction.

The turbulent dynamic viscosity is recovered from

$$\mu_t = C_\mu \frac{\rho k^2}{\varepsilon}, \quad (5.11)$$

and the production term in Equation 5.9 and Equation 5.10 is evaluated as

$$P_k = \mu_t S^2, \quad S \equiv \sqrt{2 S_{ij} S_{ij}}, \quad (5.12)$$

where $S_{ij} = \frac{1}{2}(\partial \bar{u}_i / \partial x_j + \partial \bar{u}_j / \partial x_i)$ is the mean strain-rate tensor. The standard model constants are

$$C_{1\varepsilon} = 1.44, \quad C_{2\varepsilon} = 1.92, \quad C_\mu = 0.09, \quad \sigma_k = 1.0, \quad \sigma_\varepsilon = 1.3. \quad (5.13)$$

The k - ε model is robust and converges reliably for a wide range of flows. However, it performs poorly in flows with strong adverse pressure gradients and separation, due in part to its isotropic eddy-viscosity assumption and its highly modelled ε -equation. These shortcomings are particularly relevant for the present airfoil problem, and constitute a primary motivation for using the k - ω SST formulation described in section 5.1.3.

Turbulence Model $k-\omega$

The standard $k-\omega$ model of Wilcox [41] replaces the ε -equation with a transport equation for the specific dissipation rate ω . As noted in Section 5.1.3, ω represents a natural time-scale for turbulence decay, and the eddy viscosity follows directly from Equation 5.8.

The transport equation for k reads

$$\frac{\partial k}{\partial t} + \bar{u}_j \frac{\partial k}{\partial x_j} = \tau_{ij} \frac{\partial \bar{u}_i}{\partial x_j} - \beta^* k \omega + \frac{\partial}{\partial x_j} \left[(\nu + \sigma^* \nu_t) \frac{\partial k}{\partial x_j} \right], \quad (5.14)$$

where ν is the molecular kinematic viscosity, ν_t is the turbulent kinematic viscosity, $\tau_{ij} = -\overline{u'_i u'_j}$ is the specific Reynolds stress tensor, and σ^* is the diffusion coefficient for k . On the right-hand side of Equation 5.14, the first term is the production of k due to mean velocity gradients, $-\beta^* k \omega$ is its destruction, and the final term accounts for molecular and turbulent diffusion.

The transport equation for ω is

$$\frac{\partial \omega}{\partial t} + \bar{u}_j \frac{\partial \omega}{\partial x_j} = \alpha \frac{\omega}{k} \tau_{ij} \frac{\partial \bar{u}_i}{\partial x_j} - \beta \omega^2 + \frac{\partial}{\partial x_j} \left[(\nu + \sigma \nu_t) \frac{\partial \omega}{\partial x_j} \right], \quad (5.15)$$

where α , β , and σ are model constants. The production term links the generation of ω to the local turbulence production, while $-\beta \omega^2$ represents its self-destruction. Closing the system, the turbulent kinematic viscosity is given by

$$\nu_t = \frac{k}{\omega}, \quad (5.16)$$

consistent with the scaling in Equation 5.8. The standard model constants are

$$\alpha = \frac{5}{9}, \quad \beta = \frac{3}{40}, \quad \beta^* = 0.09, \quad \sigma = \frac{1}{2}, \quad \sigma^* = \frac{1}{2}. \quad (5.17)$$

A key advantage of the $k-\omega$ model over the $k-\varepsilon$ model is its good near-wall behaviour: no additional damping functions are required, making it well suited for low-Reynolds-number wall-bounded flows. It also performs better in flows with adverse pressure gradients and incipient separation. Its principal drawback, however, is a pronounced sensitivity to the prescribed free-stream value of ω , which provided the primary motivation for the SST formulation described next.

Turbulence Model $k-\omega$ SST

The $k-\omega$ SST model of Menter [20] combines the complementary strengths of the two models above: the near-wall accuracy of the $k-\omega$ formulation and the free-stream robustness of the $k-\varepsilon$ formulation. This is achieved by blending the two sets of model equations using a smooth function F_1 , which equals unity inside the boundary layer and tends to zero in the outer flow, as illustrated in Figure 5.2. All blended coefficients take the form

$$\phi = F_1 \phi_1 + (1 - F_1) \phi_2, \quad (5.18)$$

where ϕ_1 denotes a coefficient of the inner $k-\omega$ model and ϕ_2 the corresponding coefficient of the outer transformed $k-\varepsilon$ model.

The transport equation for k is

$$\frac{\partial k}{\partial t} + \bar{u}_j \frac{\partial k}{\partial x_j} = \tilde{P}_k - \beta^* k \omega + \frac{\partial}{\partial x_j} \left[(\nu + \sigma_k \nu_t) \frac{\partial k}{\partial x_j} \right], \quad (5.19)$$

and the transport equation for ω is

$$\frac{\partial \omega}{\partial t} + \bar{u}_j \frac{\partial \omega}{\partial x_j} = \alpha S^2 - \beta \omega^2 + \frac{\partial}{\partial x_j} \left[(\nu + \sigma_\omega \nu_t) \frac{\partial \omega}{\partial x_j} \right] + 2(1 - F_1) \sigma_\omega \frac{1}{\omega} \frac{\partial k}{\partial x_i} \frac{\partial \omega}{\partial x_i}, \quad (5.20)$$

where $\tilde{P}_k = \min(P_k, 10\beta^* k \omega)$ is a production limiter that prevents the build-up of excessive turbulent kinetic energy in stagnation regions, with $P_k = \nu_t S^2$ expressed here in kinematic form consistent with Equation 5.19, and S the invariant measure of the mean strain rate defined in Equation 5.12. β^* , β , α , σ_k , σ_ω , and $\sigma_{\omega 2}$ are blended model coefficients following Equation 5.18. The final term in Equation 5.20 is the cross-diffusion term, which arises from the transformation of the $k-\varepsilon$ model into $k-\omega$ form. It is

multiplied by $(1 - F_1)$ and is therefore active only in the outer flow, where the model transitions away from the standard $k-\omega$ behaviour.

The turbulent kinematic viscosity is given by

$$\nu_t = \frac{a_1 k}{\max(a_1 \omega, SF_2)}, \quad (5.21)$$

where $a_1 = 0.31$ is a model constant and F_2 is a second blending function that is active inside the boundary layer. The denominator of Equation 5.21 introduces a shear-stress limiter: in regions of strong shear where $SF_2 > a_1 \omega$, the effective eddy viscosity is reduced relative to the standard k/ω ratio. This is the principal reason why the SST model outperforms the standard $k-\omega$ and $k-\epsilon$ formulations in separated flows and flows with adverse pressure gradients.




Figure 5.2: Blending behaviour of the $k-\omega$ SST model, showing the transition from the near-wall $k-\omega$ region through the blending region to the outer $k-\epsilon$ region. Adapted from [13].

For the present study, the SST model is selected because of its suitability for flows with adverse pressure gradients and separation, which are expected in the aft region of the airfoil where the dimples are located. The mesh is designed accordingly to resolve the near-wall region with $y^+ \approx 1$, satisfying the low-Reynolds-number requirement of the model. It should be noted, however, that no laminar-to-turbulent transition model is included. At the Reynolds numbers considered here, this limits the predictive accuracy of the absolute results; the RANS simulations are therefore interpreted primarily as a comparative screening tool for evaluating relative performance differences between dimple configurations.

5.1.4. Computational Domain and Boundary Conditions

The computational domain is defined to represent a free-field configuration rather than a direct wind-tunnel replication. Its extents, summarised in Table 5.4, are chosen to be sufficiently large that the outer boundaries have a negligible influence on the solution in the vicinity of the airfoil. The streamwise extent runs from $5c$ upstream to $10c$ downstream of the airfoil quarter-chord point, while the vertical extent spans $\pm 2.5c$. In the spanwise direction, the domain is extruded to a width of $0.2c$, yielding a thin three-dimensional slab that suppresses spanwise flow variations and approximates a quasi-two-dimensional flow field. This assumption is appropriate for the present screening study, whose primary objective is to evaluate relative aerodynamic performance differences between dimple layouts rather than to resolve three-dimensional flow structures.

Table 5.4: Computational domain for RANS campaign. Computational domain extents relative to the airfoil quarter-chord point, expressed in chord lengths $c = 200$ mm.

Region	Streamwise	Spanwise	Vertical
Full domain	$-5c$ to $+10c$	$0.2c$	$\pm 2.5c$
Far-field refinement	$-1c$ to $+2c$	$0.2c$	$\pm 0.5c$
Near-field refinement	$-\frac{3}{8}c$ to $+\frac{7}{8}c$	$0.2c$	$\pm 0.25c$

The angle of attack is varied by rotating the airfoil geometry about its quarter-chord point and regenerating the mesh for each case, while the inlet velocity vector remains fixed in the x -direction. This

ensures that the domain boundaries are always aligned with the Cartesian grid, avoiding interpolation errors that would arise from rotating the domain itself.

The boundary conditions applied to each domain face are summarised in Table 5.5 and described below.

Table 5.5: Summary of boundary conditions applied in the RANS simulations. The kinematic pressure p is normalised by density ($\text{m}^2 \text{s}^{-2}$).

Boundary	U	p	k	ω	ν_t
Inlet	Fixed value	Zero gradient	Fixed value	Fixed value	Calculated
Outlet	Inlet-outlet	Fixed value (0)	Inlet-outlet	Inlet-outlet	Calculated
Airfoil wall	No-slip (0)	Zero gradient	k -wall function	ω -wall function	Spalding
Symmetry (all)	Symmetry	Symmetry	Symmetry	Symmetry	Symmetry

Inlet. A uniform velocity U_∞ is prescribed at the upstream boundary, directed normal to the inlet face. The two Reynolds numbers considered correspond to inlet velocities of $U_\infty = 15.0 \text{ m/s}$ and $U_\infty = 37.5 \text{ m/s}$, using air properties $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$ and $\rho = 1.197 \text{ kg/m}^3$, consistent with standard conditions at $T \approx 20^\circ \text{C}$. The turbulent quantities at the inlet are prescribed as

$$k_\infty = \frac{3}{2}(Tu U_\infty)^2, \quad \omega_\infty = \frac{\sqrt{k_\infty}}{C_\mu^{1/4} \ell}, \quad (5.22)$$

where Tu is the freestream turbulence intensity and ℓ is the turbulent length scale. The prescribed values $k_\infty = 0.02344 \text{ m}^2 \text{s}^{-2}$ and $\omega_\infty = 8.438 \text{ s}^{-1}$ correspond to a turbulent length scale of $\ell \approx 33 \text{ mm} \approx 0.17c$, and imply turbulence intensities of $Tu \approx 0.83\%$ at $Re = 2 \times 10^5$ and $Tu \approx 0.33\%$ at $Re = 5 \times 10^5$. These values are intentionally kept low to represent a clean, low-turbulence free-field environment.

Outlet. A fixed kinematic pressure of $p = 0 \text{ m}^2 \text{s}^{-2}$ is applied at the downstream boundary. Velocity and turbulent quantities use an `inletOutlet` condition, which reverts to the inlet values if backflow is detected, preventing numerical divergence in separated or unsteady cases.

Airfoil wall. A no-slip condition is enforced on the airfoil surface, with zero velocity and a zero-gradient pressure condition. The turbulent kinetic energy uses a `kqRWallFunction`, and ω uses a stepwise-blended `omegaWallFunction` consistent with the low- Re SST formulation. The turbulent viscosity is recovered using the Spalding wall function (`nutUSpaldingWallFunction`), which provides a smooth and continuous treatment across the viscous sublayer, buffer layer, and log-law region. This is appropriate given the target $y^+ \approx 1$ described in Section 5.1.5.

Symmetry planes. The four remaining boundaries — top, bottom, and the two spanwise faces — are treated as symmetry planes, enforcing zero normal velocity and zero normal gradients of all flow variables. This avoids the formation of artificial boundary layers at the domain limits and is consistent with the quasi-two-dimensional modelling assumption.

5.1.5. Mesh Generation

High-quality spatial discretisation is a prerequisite for reliable RANS predictions, particularly in wall-bounded flows where the boundary-layer structure governs separation onset and drag. This section describes the mesh strategy, the near-wall resolution requirements, and the sensitivity study used to select the mesh for all production runs.

Mesh Strategy

All meshes are generated using the automatic meshing functionality in SimScale, which employs a Cartesian cut-cell approach with region-based refinement and prismatic boundary-layer inflation. The domain is subdivided into three concentric refinement regions, whose extents are defined in Table 5.4 and illustrated in Figure 5.3.

The outermost region covers the full domain and is meshed at a coarse resolution sufficient to propagate the far-field boundary conditions without significant numerical dissipation. A first refinement box

encloses the near-wake and leading-edge regions, capturing the accelerated flow around the airfoil and the developing wake. A second, finer refinement box is placed immediately around the airfoil to resolve the pressure gradients on the suction surface and the separated shear layer near the trailing edge. Finally, a prismatic boundary-layer inflation is applied directly on the airfoil surface, producing a structured layer of cells whose thickness is controlled to achieve the target near-wall resolution. The specific parameters for each mesh level are listed in Table 5.6.




Figure 5.3: Computational mesh for the NACA63-415 airfoil (Mesh 4). The upper view shows the full domain with the two refinement regions visible as concentric boxes around the airfoil. The lower view provides a close-up of the near-field region, illustrating the transition from the coarse background mesh to the refined near-wall cells and prismatic boundary-layer inflation on the airfoil surface.

Near-Wall Resolution

The non-dimensional wall distance y^+ is defined as

$$y^+ = \frac{u_\tau y}{\nu}, \quad (5.23)$$

where $u_\tau = \sqrt{\tau_w/\rho}$ is the friction velocity, y is the distance from the wall to the centre of the first cell, and ν is the kinematic viscosity. The value of y^+ quantifies how deeply the first computational cell penetrates into the boundary layer: $y^+ \lesssim 5$ corresponds to the viscous sublayer, while $y^+ \sim 30\text{--}300$ corresponds to the log-law region.

The $k\text{--}\omega$ SST model, as described in section 5.1.3, is formulated as a low-Reynolds-number model that integrates the governing equations through the viscous sublayer without requiring empirical wall functions for the mean flow. To exploit this capability, the first cell centre must be placed within the viscous sublayer, requiring $y^+ \approx 1$ across the airfoil surface. Achieving this consistently is the primary driver of mesh resolution in the present study, as confirmed by the sensitivity analysis in Section 5.1.6.

An additional challenge arises in the dimpled configurations, where the surface curvature at the rim of each cavity introduces local flow acceleration and elevated wall shear stress. The y^+ distribution is therefore not uniform across the surface, and the target $y^+ \approx 1$ must be satisfied at these locations as well. This is addressed by applying the boundary-layer inflation uniformly across the entire airfoil

surface, including the dimple geometry, with a sufficiently small first-layer thickness to accommodate the locally elevated u_τ .

5.1.6. Mesh Sensitivity Study

A mesh sensitivity study was performed on the Smooth baseline airfoil at $\alpha = 0^\circ$ and $Re = 2 \times 10^5$ to identify the mesh resolution at which the aerodynamic coefficients are insensitive to further refinement. Five meshes of increasing resolution were evaluated, as detailed in Table 5.6. The key parameters varied between levels are the number of prismatic boundary-layer inflation layers, the expansion rate, and the relative first-layer thickness, while the far- and near-field cell sizes are held constant between Meshes 1–4 and only refined for Mesh 5.

Table 5.6: Mesh sensitivity study for the NACA 63-415 baseline at $\alpha = 0^\circ$, $Re = 2 \times 10^5$. Mesh parameters and computational cost are listed for each resolution level.

Parameter	Mesh 1	Mesh 2	Mesh 3	Mesh 4	Mesh 5
Cells [million]	1.6	2.0	2.2	2.5	4.7
Nodes [million]	1.9	2.4	2.6	2.9	5.1
Core hours (mesh)	2.73	2.79	3.40	6.44	11.04
Core hours (simulation)	5.54	8.79	13.76	35.97	55.34
Far-field max cell size	5 mm	5 mm	5 mm	5 mm	2 mm
Near-field max cell size	–	2 mm	2 mm	2 mm	2 mm
Boundary-layer layers	5	5	10	15	15
Expansion rate	1.3	1.3	1.2	1.2	1.2
BL min. thickness	auto	auto	2 mm	2 mm	2 mm
Relative first-layer thickness	0.055	0.055	0.050	0.025	0.025
$y_{\max}^+$	9.41	9.33	9.57	15.23	15.18
y_{avg}^+	4.29	4.26	2.07	1.07	1.07
$y_{\min}^+$	0.07	0.07	0.05	0.04	0.05

The aerodynamic performance predicted by each mesh was evaluated against Mesh 5, the finest resolution considered, which serves as the reference for the grid-convergence assessment. Results are summarised in Table 5.7, where the glide ratio $G = C_L/C_D$ is used as the primary convergence metric since it combines the errors in both lift and drag into a single figure of merit. RANS results from Villalpando et al. [38] and experimental data from Filho et al. [10] are included in the table for further comparison, though both are at a higher Reynolds number ($Re = 3 \times 10^5$ and $Re = 3.2 \times 10^5$, respectively) than the $Re = 2 \times 10^5$ considered in the present study, and are therefore not used as the convergence reference.

It should be noted that $y_{\max}^+$ in Table 5.7 does not decrease monotonically with first-layer refinement. This maximum is set by a small number of individual cells at geometric edges and boundaries, where the automatic mesher produces locally elevated skewness, and its exact value therefore varies from one mesh to the next rather than following the refinement level. All such cells remain within the quality limits enforced by the mesh generator, and the area-averaged value y_{avg}^+ is the meaningful measure of near-wall resolution here.

The results show a monotonic reduction in the glide-ratio deviation with increasing mesh resolution. The dominant driver is the near-wall resolution rather than the background cell size: Meshes 1 and 2 differ only in the addition of a near-field refinement box yet show almost no improvement (-24.7% versus -22.6%), while the step from Mesh 2 to Mesh 3, which increases the number of boundary-layer inflation layers from 5 to 10 and halves the first-layer thickness, reduces the deviation to -11.1% . A further halving of the relative first-layer thickness in Mesh 4, combined with 15 inflation layers, brings y_{avg}^+ to 1.07 and the deviation to -5.1% . Mesh 5 then reduces the far-field cell size by a factor of 2.5, from 5 mm to 2 mm, while keeping the boundary-layer setup identical, raising the glide ratio from $G = 22.06$ to $G = 23.24$.

This final step changes the glide ratio by 5.1 percentage points, which is comparable to the 6.0 percentage points gained between Mesh 3 and Mesh 4. The sequence has therefore not yet flattened, and grid independence in the strict sense is not demonstrated by the present study. The practical consequence is that the production mesh carries a residual discretisation uncertainty in G of the order of 5%. Two considerations make this acceptable for the purpose at hand. It is small relative to the modelling error introduced by the fully turbulent formulation, which is shown in Section 7.6 to be the dominant source of disagreement with experiment. It is also systematic, since the same meshing strategy is applied to all four geometries, so it largely cancels in the geometry-to-geometry comparison that the RANS screening is used for.

Table 5.7: Aerodynamic coefficients and glide ratio for the five mesh levels at $\alpha = 0^\circ$, $Re = 2 \times 10^5$. The relative deviation in glide ratio is defined as $\Delta G/G_{\text{ref}} = (G - G_{\text{ref}})/G_{\text{ref}}$, where $G_{\text{ref}} = 23.24$ is the glide ratio obtained on the finest mesh (Mesh 5), used as the reference for the grid-convergence assessment. RANS [38] and experimental [10] literature values are included for comparison but correspond to a higher Reynolds number ($Re = 3 \times 10^5$ and $Re = 3.2 \times 10^5$, respectively).

Mesh	Cells [M]	C_D	C_L	G	$\Delta G/G_{\text{ref}}$ [%]
Mesh 1	1.6	0.0173	0.303	17.51	-24.7
Mesh 2	2.0	0.0169	0.304	17.99	-22.6
Mesh 3	2.2	0.0154	0.318	20.65	-11.1
Mesh 4	2.5	0.0146	0.322	22.06	-5.1
Mesh 5	4.7	0.0139	0.323	23.24	0.0
RANS, Villalpando et al. [38]	—	0.0127	0.3044	23.97	—
Experimental, Filho et al. [10]	—	0.0171	0.2095	12.25	—




Figure 5.4: Wall-unit distribution (y^+) on the suction side of the smooth NACA63-415 airfoil at $\alpha = 0^\circ$ for the five mesh resolutions. Leading edge to the left, trailing edge to the right. The transition from elevated y^+ values on Meshes 1–2 to a uniformly low distribution on Meshes 4–5 reflects the improvement in boundary-layer inflation described in Table 5.6.

Based on this analysis, **Mesh 4** was selected for all 48 production runs. It achieves $y_{\text{avg}}^+ \approx 1$, satisfying the low-Reynolds-number requirement of the SST model, at a simulation cost of approximately 36 core hours per case. Mesh 5 reduces the glide-ratio deviation by a further 5.1 percentage points but requires 55 core hours per simulation — a 54% increase in cost, which is prohibitive across the 48-case production matrix. Mesh 4 therefore represents the best available balance between accuracy and computational cost for the present screening campaign, at the price of the residual discretisation uncertainty noted above.

5.1.7. Summary of RANS Setup

This section has described the complete numerical setup for the RANS screening campaign. The study evaluates four airfoil configurations — a smooth baseline and three dimpled layouts with varying chordwise coverage — across six angles of attack and two Reynolds numbers, yielding a total of 48 simulations.

The governing equations are solved using the finite-volume solver `pimpleFoam` in `OpenFOAM`, operated in pseudo-transient mode until steady-state convergence is achieved. Turbulence is modelled using the $k-\omega$ SST formulation of Menter [20], selected for its established accuracy in flows with adverse pressure gradients and incipient separation. The computational domain is configured as a free-field environment with extents of $-5c$ to $+10c$ in the streamwise direction and $\pm 2.5c$ in the vertical direction, extruded to a spanwise width of $0.2c$ to enforce a quasi-two-dimensional flow field.

The mesh is generated using a Cartesian cut-cell approach with prismatic boundary-layer inflation, selected through a five-level sensitivity study on the smooth baseline airfoil. Mesh 4 was identified as the production mesh, achieving $y_{\text{avg}}^+ \approx 1$ across the airfoil surface at a computational cost of approximately 36 core hours per simulation. The key numerical settings and domain parameters are consolidated in Table 5.8.

Table 5.8: Summary of RANS simulation settings and domain parameters used across all 48 production runs.

Parameter	Value
Solver	
Solver	pimpleFoam (OpenFOAM)
Approach	Pseudo-transient, steady-state
Turbulence model	$k-\omega$ SST [20]
Flow regime	Incompressible
Operating conditions	
Chord length c	200 mm
Reynolds numbers Re	2.0×10^5 , 5.0×10^5
Inlet velocities U_∞	15.0 m/s, 37.5 m/s
Angles of attack α	0° , 6° , 10° , 12° , 14° , 16°
Kinematic viscosity ν	$1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$
Air density ρ	1.197 kg m^{-3}
Domain	
Streamwise extent	$-5c$ to $+10c$
Vertical extent	$\pm 2.5c$
Spanwise width	$0.2c$
Mesh	
Total cells	$\approx 2.5 \times 10^6$
Boundary-layer layers	15
Expansion rate	1.2
Relative first-layer thickness	0.025
y_{avg}^+	≈ 1.07
y_{max}^+	15.2
Inlet turbulence	
Freestream TKE k_∞	$0.02344 \text{ m}^2 \text{ s}^{-2}$
Specific dissipation ω_∞	8.438 s^{-1}
Turbulence intensity Tu	0.83% ($Re = 2 \times 10^5$), 0.33% ($Re = 5 \times 10^5$)
Study scope	
Geometry configurations	4 (Smooth, Mid-2.50, Aft-2.50, Full-2.50)
Total simulations	48

It should be noted that no laminar-to-turbulent transition model is included in the present setup. At the Reynolds numbers considered, natural transition occurs on the airfoil surface, and its location influences the extent of the laminar separation bubble and the onset of stall. The absence of a transition model therefore limits the absolute predictive accuracy of the RANS results. The simulations are consequently interpreted as a comparative screening tool, with the primary objective of identifying relative aerodynamic performance differences between dimple configurations rather than reproducing absolute experimental values.

5.2. Large-Eddy Simulations

A limited Large-Eddy Simulation (LES) campaign was conducted as a supplement to the RANS screening study and the experimental measurements. The objective was not to perform a full numerical parameter study, but to obtain a higher-fidelity unsteady flow-field comparison for two selected configurations: the smooth baseline airfoil and one dimpled configuration. The LES results were therefore only used for qualitative comparison with the PIV and RANS flow fields at one operating condition.

Compared with the RANS simulations described in Section 5.1, LES resolves the larger turbulent motions directly and models only the unresolved sub-grid-scale motions. This makes LES better suited for visualising separated shear layers, wake development, and the local flow response around the dimple cavities. However, because only two cases were simulated, the results are interpreted as qualitative flow-field support rather than quantitative validation data.

5.2.1. Study Design and Test Matrix

The LES campaign was restricted to two configurations: the undimpled baseline geometry Smooth and the mid-chord dimpled configuration Mid-2.50. The numeric suffix denotes a fixed dimple depth of $k = 2.50$ mm, corresponding to $k/c = 0.0125$ for the selected chord length of $c = 200$ mm. The selected cases are summarised in Table 5.9 and illustrated in Figure 5.5.

The Mid-2.50 configuration was selected because it was also tested experimentally using PIV, force measurements, and tuft visualisation. This allowed the same geometry to be compared across numerical and experimental methods. The computational span was reduced to $b/c = 0.1$ to keep the LES cost manageable.

Table 5.9: Geometry configurations used in the LES study. The numeric suffix denotes a fixed dimple depth of $k = 2.50$ mm, corresponding to $k/c = 0.0125$ for $c = 200$ mm.

ID	Rows	x/c range	D_d/c	k/c	b/c
Smooth	–	–	–	–	0.1
Mid-2.50	1–4 (4/7)	0.33–0.59	0.05	0.0125	0.1




Figure 5.5: CAD geometries of the two airfoil configurations used in the LES study. Smooth is the undimpled NACA 63-415 baseline. Mid-2.50 covers the mid-chord region from $x/c = 0.33$ to 0.59 . The dimpled configuration uses inward spherical dimples with $D_d/c = 0.05$, $k/c = 0.0125$, and staggered spacing $D_s/D_d = 2.0$. Both computational geometries have a span of $b/c = 0.1$.

The operating condition was selected to match one of the PIV measurement cases. Both simulations were performed at $Re = 1.0 \times 10^5$ and $\alpha = 6^\circ$, corresponding to an inlet velocity of $U_\infty = 7.5$ m/s for $c = 200$ mm and $\nu = 1.5 \times 10^{-5}$ m²/s. The resulting LES case count is therefore only two.

Table 5.10: Operating conditions used in the LES study.

Parameter	Values
Angle of attack α	6°
Reynolds number Re	1.0×10^5
Inlet velocity U_∞	7.5 m/s
Flow regime	Incompressible

5.2.2. Governing Equations and Sub-Grid-Scale Model

The LES calculations were performed using the transient incompressible solver `pimpleFoam` in OpenFOAM v2406. In LES, the Navier–Stokes equations are spatially filtered with a filter width Δ tied to the local mesh size. Large energy-carrying eddies are resolved directly, while the effect of unresolved sub-grid-scale (SGS) motions is modelled.

Filtered Continuity and Momentum Equations

The spatially filtered incompressible continuity equation

$$\frac{\partial \bar{u}_i}{\partial x_i} = 0, \quad (5.24)$$

where $\bar{u}_i$ is the resolved, filtered velocity component in direction i . The filtered momentum equation can be written as

$$\frac{\partial \bar{u}_i}{\partial t} + \frac{\partial}{\partial x_j} (\bar{u}_j \bar{u}_i) = -\frac{1}{\rho} \frac{\partial \bar{p}}{\partial x_i} + \frac{\partial}{\partial x_j} \left[(\nu + \nu_{\text{sgs}}) \frac{\partial \bar{u}_i}{\partial x_j} \right], \quad (5.25)$$

where an overbar $(\bar{\cdot})$ denotes a spatially filtered quantity, $\bar{u}_i$ is the resolved velocity, $\bar{p}$ is the filtered kinematic pressure (normalised by density), ν is the molecular kinematic viscosity, and ν_{sgs} is the sub-grid-scale eddy viscosity provided by the SGS model.

The SGS stress tensor

$$\tau_{ij} = \overline{u_i u_j} - \bar{u}_i \bar{u}_j$$

is closed through the Boussinesq hypothesis:

$$\tau_{ij} - \frac{1}{3} \tau_{kk} \delta_{ij} = -2\nu_{\text{sgs}} \bar{S}_{ij} \quad \text{where} \quad \bar{S}_{ij} = \frac{1}{2} \left(\frac{\partial \bar{u}_i}{\partial x_j} + \frac{\partial \bar{u}_j}{\partial x_i} \right)$$

is the resolved strain-rate tensor.

The term ν_{sgs} represents the additional momentum diffusion caused by turbulent structures that are smaller than the local grid scale. Thus, the LES closure problem is reduced to estimating ν_{sgs} .

Smagorinsky SGS Model

The SGS eddy viscosity is computed algebraically using the Smagorinsky model [31]:

$$\nu_{\text{sgs}} = (C_s \Delta)^2 |\bar{S}|, \quad |\bar{S}| = \sqrt{2 \bar{S}_{ij} \bar{S}_{ij}}, \quad (5.26)$$

where $C_s = 0.158$ is the Smagorinsky constant [24], Δ is the local filter width defined in Equation 5.27, and $|\bar{S}|$ is the magnitude of the resolved strain-rate tensor. In the OpenFOAM implementation, ν_{sgs} is derived from an algebraic estimate of the SGS kinetic energy k_{sgs} [24], but the resulting eddy viscosity is physically equivalent to Equation 5.26 for incompressible flow.

Van Driest Near-Wall Damping

The LES filter width was specified using the `vanDriest` delta model,

$$\Delta = f_{\text{vD}} \Delta_g, \quad (5.27)$$

where Δ_g is the underlying geometric filter width and f_{vD} is the van Driest damping function,

$$f_{\text{vD}} = 1 - \exp\left(-\frac{y^+}{A^+}\right), \quad (5.28)$$

with y^+ the dimensionless wall distance and $A^+ = 26$ the van Driest constant. In the present setup, Δ_g was computed using the `cubeRootVol` option,

$$\Delta_g = V_{\text{cell}}^{1/3}, \quad (5.29)$$

where V_{cell} is the local cell volume. The van Driest damping reduces the effective filter width near solid walls, ensuring $\nu_{\text{sgs}} \rightarrow 0$ as $y^+ \rightarrow 0$ and preventing the undamped Smagorinsky model from overpredicting eddy viscosity in the viscous sublayer.

Time Discretisation

Time integration was performed using the Crank–Nicolson scheme with a blending coefficient of 0.9. The coefficient 0.9 indicates that the time integration is predominantly Crank–Nicolson, with a small amount of implicit Euler off-centring for numerical damping. This provides higher temporal accuracy than a purely first-order Euler scheme while retaining near-second-order temporal accuracy appropriate for resolving unsteady flow structures in LES. This improves stability compared with a fully centred Crank–Nicolson discretisation which is beneficial when the objective is to resolve unsteady flow structures while maintaining numerical robustness.

5.2.3. Computational Domain and Boundary Conditions

The LES domain follows the same free-field layout as the RANS simulations, but with a reduced spanwise width of $b/c = 0.1$ to reduce computational cost. The streamwise domain extends from $-5c$ upstream to $+10c$ downstream of the airfoil quarter-chord point, while the vertical extent is $\pm 2.5c$. The domain and refinement regions are summarised in Table 5.11.

Table 5.11: Computational domain used for the LES campaign. Domain extents are given relative to the airfoil quarter-chord point and expressed in chord lengths, with $c = 200$ mm.

Region	Streamwise	Spanwise	Vertical
Full domain	$-5c$ to $+10c$	$0.1c$	$\pm 2.5c$
Far-field refinement	$-1c$ to $+2c$	$0.1c$	$\pm 0.5c$

The boundary conditions are summarised in Table 5.12. A fixed inlet velocity of $U_\infty = 7.5$ m/s is prescribed, while the outlet uses a fixed reference pressure. The airfoil surface is treated as a no-slip wall. Symmetry conditions are applied at the top, bottom, and spanwise boundaries.

Table 5.12: Summary of boundary conditions applied in the LES simulations. The kinematic pressure p is normalised by density ($\text{m}^2 \text{s}^{-2}$). The Smagorinsky SGS model computes ν_{sgs} algebraically; no transport equations for k or ω are solved.

Boundary	U	p	ν_{sgs}
Inlet	Fixed value (7.5)	Zero gradient	Fixed value (0)
Outlet	Inlet-outlet	Fixed value (0)	Inlet-outlet
Airfoil wall	No-slip (0)	Zero gradient	Fixed value (0)
Symmetry (all)	Symmetry	Symmetry	Symmetry

5.2.4. Mesh Generation and Wall-Unit Estimates

The LES mesh was generated using the same general Cartesian cut-cell strategy as the RANS meshes, with local refinement around the airfoil and prismatic boundary-layer inflation on the wall. The resulting mesh is shown in Figure 5.6, and the main mesh settings are summarised in Table 5.13.

For LES, the relevant wall-unit measures are

$$y^+ = \frac{u_\tau y}{\nu}, \quad x^+ = \frac{u_\tau \Delta x}{\nu}, \quad z^+ = \frac{u_\tau \Delta z}{\nu}, \quad (5.30)$$

where $u_\tau = \sqrt{\tau_w/\rho}$ is the friction velocity, τ_w is the wall shear stress, y is the wall-normal distance from the wall to the first cell centre, Δx and Δz are the near-wall streamwise and spanwise cell sizes, and ν is the kinematic viscosity. A preliminary estimate of u_τ was obtained from the turbulent flat-plate skin-friction relation

$$C_f \approx 0.0576 Re^{-1/5} \quad (5.31) \quad \text{combined with} \quad u_\tau = U_\infty \sqrt{\frac{C_f}{2}} \quad (5.32)$$

For $U_\infty = 7.5$ m/s, $Re = 1.0 \times 10^5$, and $\nu = 1.5 \times 10^{-5} \text{ m}^2/\text{s}$, this gives $C_f \approx 0.00576$ and $u_\tau \approx 0.40$ m/s. The corresponding viscous length scale is

$$\ell_* = \frac{\nu}{u_\tau} \approx \frac{1.5 \times 10^{-5}}{0.40} \approx 3.7 \times 10^{-5} \text{ m}. \quad (5.33)$$

Thus, $y^+ = 1$ corresponds to $y \approx 0.037$ mm, $x^+ = 100$ corresponds to $\Delta x \approx 3.7$ mm, and $z^+ = 50$ corresponds to $\Delta z \approx 1.9$ mm. The outer domain was therefore meshed with a maximum cell size of 4 mm, while the refined region around the airfoil used a maximum cell size of 2 mm. With a relative first-layer thickness of 0.025, leading to a maximum first-layer height of 0.05 mm, placing the near-wall resolution close to the intended $y^+ \approx 1$ target. It should be emphasised that this mesh was selected as a practical compromise between resolution and computational cost.

Table 5.13: Mesh settings for the NACA 63-415 airfoil at $\alpha = 6^\circ$ and $Re = 1.0 \times 10^5$ for the LES simulations.

Parameter	LES Mesh
Cells [million]	9.9
Nodes [million]	13.6
Core hours (mesh)	15.1
Core hours (simulation)	3417.6
Full domain max cell size	4 mm
Far-field max cell size	2 mm
Boundary-layer layers	15
Expansion rate	1.2
BL min. thickness	2 mm
Relative first-layer thickness	0.025




Figure 5.6: Computational mesh for the NACA 63-415 airfoil used in the LES campaign. The upper view shows the full domain and the refinement box around the airfoil. The lower view shows the near-field region, including the refined mesh and prismatic boundary-layer inflation on the airfoil surface.

5.2.5. Summary of LES Setup

The LES campaign consisted of two simulations: the undimpled Smooth baseline and the dimpled Mid-2.50 configuration at $Re = 1.0 \times 10^5$ and $\alpha = 6^\circ$. The simulations were performed using `pimpleFoam` in OpenFOAM with the Smagorinsky SGS model. The domain followed the same free-field layout as the RANS simulations but used a reduced span of $b/c = 0.1$ to keep the computational cost manageable.

Both simulations were advanced with a fixed time step of $\Delta t = 5 \times 10^{-4}$ s to a total simulated time of 1.0 s, corresponding to 2000 time steps and approximately 37 flow-throughs based on the convective time $c/U_\infty = 2.67 \times 10^{-2}$ s.

The key numerical settings are summarised in Table 5.14. The LES results are used only to support the comparison of instantaneous and mean flow-field structures between PIV, RANS, and LES. They are not used as a standalone aerodynamic performance dataset.

Table 5.14: Summary of LES simulation settings and domain parameters.

Parameter	Value
<i>Solver</i>	
Solver	<code>pimpleFoam</code> (OpenFOAM)
Approach	Transient LES
SGS model	Smagorinsky [31]
Flow regime	Incompressible
<i>Operating conditions</i>	
Chord length c	200 mm
Reynolds number Re	1.0×10^5
Inlet velocity U_∞	7.5 m/s
Angle of attack α	6°
Kinematic viscosity ν	$1.5 \times 10^{-5} \text{ m}^2 \text{ s}^{-1}$
<i>Domain</i>	
Streamwise extent	$-5c$ to $+10c$
Vertical extent	$\pm 2.5c$
Spanwise width	$0.1c$
<i>Mesh</i>	
Total cells	$\approx 9.9 \times 10^6$
Boundary-layer layers	15
Expansion rate	1.2
Relative first-layer thickness	0.025
<i>Temporal discretisation and averaging</i>	
Time step Δt	5×10^{-4} s
Total simulated time	1.0 s (2000 time steps)
Convective time c/U_∞	2.67×10^{-2} s
Flow-throughs simulated	≈ 37
<i>Study scope</i>	
Geometry configurations	2 (Smooth, Mid-2.50)
Total simulations	2

6

Experimental and Numerical Validation

The present campaign uses seven distinct experimental and numerical methods to characterise the aerodynamic performance of the NACA 63-415 airfoil. Of these, five are quantitative and subject to formal validation; the RWT wall-pressure and wake rake measurements (Section 3.3), the RWT surface pressure tap models (Section 3.4), the SLT force balance (Section 4.4), the RANS simulations (Section 5.1), and the LES (Section 5.2). The remaining two methods, PIV and tuft flow visualisation, are used in a qualitative nature and are not validated against any public data. Before interpreting the dimpled configurations, it is necessary to establish that each quantitative method produces physically consistent results for the smooth NACA 63-415 baseline airfoil, denoted Smooth, and that the methods agree with one another and with published reference data for the NACA 63-415.

This chapter presents that validation. Figure 6.1, Figure 6.2, and Figure 6.3 show the lift and drag polars for all available Smooth datasets at $Re = 1.0 \times 10^5$, 2.0×10^5 , and 5.0×10^5 respectively, alongside XFOil predictions [2] and literature references for the same airfoil at comparable Reynolds numbers. All references were carried out on the same NACA 63-415 airfoil profile. Since none of these references exactly match the Reynolds numbers of the present campaign, comparisons should be interpreted in terms of trend agreement and qualitative behaviour rather than point-by-point accuracy.

The experimental reference at $Re = 3.2 \times 10^5$ is Filho et al. [10], who conducted wind-tunnel measurements in an open-return tunnel with a 460×460 mm rectangular test section operated at 20 m/s. The NACA 63-415 model had a 250 mm chord and 450 mm span. Forces and moments were measured via a three-degree-of-freedom external aerodynamic balance, similar to the method used in the present campaign via the SLT force balance (Section 4.4).

The experimental reference at $Re = 1.6 \times 10^6$ is the Risø wind-tunnel campaign carried out by Bak et al. [5], conducted in the VELUX open-jet tunnel at 40 m/s. The NACA 63-415 model had a 0.600 m chord and 1.9 m span, with end plates at both tips to suppress three-dimensional effects. Aerodynamic coefficients were derived from 62 mid-span surface pressure taps and a downstream wake rake of 53 total-pressure probes rather than a force balance, analogous to the methods employed in the present campaign and described in Section 3.3 and Section 3.4.

The numerical reference is Villalpando et al. [38], who performed two-dimensional RANS simulations using the $k-\omega$ SST turbulence model on a structured mesh of approximately 120,000 cells with $y^+ \approx 1$, at $Re = 3 \times 10^5$ and 5.7×10^5 . No transition model was employed and the boundary layer was treated as fully turbulent throughout, consistent with the approach adopted for the RANS simulations in the present campaign (Section 5.1).

The remaining numerical reference is XFOil [2], a coupled viscous-inviscid solver that combines a potential-flow panel method for the outer inviscid region with a two-equation integral boundary-layer formulation to account for viscous effects. XFOil results are reliable up to the onset of stall but do not capture post-stall behaviour.




Figure 6.1: Validation of the Smooth baseline NACA 63-415 airfoil, at $Re = 1.0 \times 10^5$. Lift coefficient C_L (left) and drag coefficient C_D (right) as a function of angle of attack. Datasets shown: RWT wall-pressure (C_L) and wake rake (C_D), RWT surface pressure taps (integrated C_L and pressure C_D), SLT force balance, RANS, LES, and XFOIL at $Re = 1.0 \times 10^5$ [2].




Figure 6.2: Validation of the Smooth baseline NACA 63-415 airfoil, at $Re = 2.0 \times 10^5$. Lift coefficient C_L (left) and drag coefficient C_D (right) as a function of angle of attack. Additional references: Villalpando et al. [38] (numerical, $Re = 3 \times 10^5$) and Filho [10] (experimental, $Re = 3.2 \times 10^5$).




Figure 6.3: Validation of the Smooth baseline NACA 63-415 airfoil, at $Re = 5.0 \times 10^5$. Lift coefficient C_L (left) and drag coefficient C_D (right) as a function of angle of attack. Additional references: Villalpando et al. [38] (numerical, $Re = 5.7 \times 10^5$) and Risø [5] (experimental, $Re = 1.6 \times 10^6$).

6.1. Cross-Method Validation and Comparison

The validation polars for all three Reynolds numbers are shown in Figure 6.1, Figure 6.2, and Figure 6.3. Each quantitative method is discussed in turn below, drawing on all three figures.

6.1.1. RANS Validation

The RANS results are shown as yellow circles, where the marker denotes the mean value over the final 25 % of the iteration history and the error bars indicate the minimum and maximum values recorded over that same interval. This representation is used because the steady RANS solver cannot resolve inherently unsteady post-stall flow structures, and at higher angles of attack the solution oscillates periodically rather than converging to a fixed point, as can be seen in Figure B.11, Figure B.12, Figure B.13, and Figure B.14 in the Appendix. The error bars therefore reflect the amplitude of this periodic behaviour rather than measurement uncertainty. This uncertainty becomes most apparent at high angles of attack where the flow separates and the solver oscillates periodically around a mean value instead of reaching a single steady solution.

Across all three Reynolds numbers, the RANS results show good agreement with the XFoil and numerical references throughout the attached and transitional regime (low to moderate angles of attack), with the trends similarity deteriorating at higher angles once significant flow separation occurs. In general, the RANS, like the other numerical references, overestimates the lift and underestimates the drag relative to the experimental results across all three Reynolds numbers. This is consistent with the fully turbulent boundary-layer assumption and the absence of a laminar-to-turbulent transition model in the simulations, leading to a natural over-prediction in lift and under-prediction of drag relative to the experimental values.

A notable exception is the drag comparison at low angles of attack in Figure 6.2 and Figure 6.3. The RANS C_D and the RWT wake rake C_D are in close agreement at low angles of attack. At $Re = 5.0 \times 10^5$ the RANS drag and lift at low angles of attack also align well with the experimental Risø reference [5], providing further confidence in the numerical setup in the attached regime. This agreement between the momentum-deficit wake measurement, the CFD prediction, and the experimental literature provides confidence in both the experimental setup and the numerical estimation.

6.1.2. LES Validation

The LES result (Figure 6.1) is shown as a single yellow triangle at $\alpha = 6^\circ$, representing the C_L and C_D values at the end of the simulation, which was stopped upon reaching a converged periodic state. Only one LES data point for the baseline is available, shown in the $Re = 1.0 \times 10^5$ figure.

At this condition, the LES captures the real-flow behaviour notably better than the RANS. Whereas the RANS prediction aligns with the XFoil reference, overpredicting lift and underpredicting drag, the LES $\alpha = 6^\circ$ point falls in close agreement with all three experimental methods for both C_L and C_D . Because the LES resolves the larger turbulent structures directly rather than modelling them, it captures more of the real-flow physics, including the transitional boundary-layer behaviour that the steady RANS cannot represent. This agreement is consistent with improved fidelity of LES over RANS at this low Reynolds number, although it rests on a single operating point and a single geometry, and close agreement in two integrated coefficients does not by itself establish that the underlying flow field is correct. It is therefore treated as supporting confidence in the LES flow-field results used for comparison to experimental data in the results chapter, rather than as independent proof of their accuracy.

6.1.3. Wall-Pressure and Wake Rake Validation

The RWT wall-pressure C_L and the wake rake C_D are shown as the mean of two identical independent measurement sessions conducted on separate days, with error bars indicating one standard deviation between sessions. The small spread between sessions confirms good repeatability of the measurements. Each individual session value is itself a time-averaged measurement recorded over a fixed sampling window at each angle of attack and Reynolds number, following the same time-averaging procedure used for the SLT force balance data.

The wall-pressure C_L is the most complete lift dataset, spanning $\alpha = -10^\circ$ to 20° at all three Reynolds numbers. It follows the XFoil and literature lift curves trends closely throughout the attached-flow regime, exhibiting the same slope and a gradual stall consistent with the expected behaviour of the NACA 63-415 at these Reynolds numbers. As is typical for real wind-tunnel measurements, the mea-

sured lift sits slightly below the idealised Xfoil prediction. In general, the wake rake C_D shows excellent agreement with both Xfoil and the RANS drag across the attached regime, including a clear drag bucket, and lies systematically above the Xfoil curve only by the small margin expected given Xfoil's known tendency to underpredict drag at low Reynolds numbers. The wake rake provides the sectional, two-dimensional drag at mid-span and is therefore the most directly comparable experimental drag measure to the CFD and literature references.

Comparing the lift to the experimental references, at $Re = 5.0 \times 10^5$ the wake rake C_D shows strong agreement with the experimental Risø reference [5] from approximately $\alpha = -10^\circ$ to 10° , further confirming the validity of the momentum-deficit drag measurement in the attached regime. Additionally, the wall-pressure C_L and the Filho et al. [10] data follow one another very closely in Figure 6.2, being near identical in the attached region before slowly diverging, with Filho showing higher lift at separated, high-angle conditions. The drag is less closely matched: Filho shows higher C_D than the wake rake, which is expected since Filho's data are obtained from a force balance and therefore include the unavoidable three-dimensional effects of a finite wing in a tunnel. This is reinforced when comparing with the SLT C_D , which is also a force balance measurement. The Filho and SLT drag values are higher and in close mutual agreement, while the wake rake drag sits lower, consistent with it being a sectional two-dimensional measurement.

6.1.4. Surface Pressure Tap Validation

The surface pressure tap results follow the same statistical treatment as the wall-pressure and wake rake data: each point is the mean of two independent measurement sessions conducted on separate days, with error bars indicating one standard deviation between sessions, and each session value is itself time-averaged over a fixed sampling window at each angle of attack and Reynolds number.

The surface pressure taps yield an integrated C_L and a pressure-only C_D , available for Smooth at all three Reynolds numbers. The integrated C_L follows the wall-pressure and force balance lift curves closely throughout the attached-flow regime, providing mutual confirmation between three independent experimental techniques. A small offset is visible, the surface tap lift is slightly lower than the other two methods through most of the range, crossing over to slightly higher lift near stall. This is attributable to the finite tap coverage described in Section 3.4: manufacturing and geometric constraints meant taps could not be placed at the very trailing edge, leaving small unsampled regions that introduce error in the pressure-integrated normal force coefficient.

The surface tap C_D is substantially higher than the wake rake at moderate to high angles of attack and at $Re = 5.0 \times 10^5$. This is expected and not anomalous: the pressure integration yields pressure drag only and does not include skin-friction drag, which becomes an increasingly large fraction of the total drag as the boundary layer thickens. The agreement is nonetheless good at low and negative angles, where attached flow makes skin friction a smaller relative contribution. The surface tap drag is therefore included to confirm the correct sign and trend of the pressure drag rather than as a quantitative total drag estimate.

6.1.5. Force Balance Validation

The SLT force balance results are shown as single time-averaged values per operating point without error bars, as only one measurement session was conducted. Force balance data are available at $Re = 1.0 \times 10^5$ and 2.0×10^5 over $\alpha = 0^\circ$ to 16° .

The force balance C_L is in close agreement with the RWT wall-pressure lift across the full measured range at both Reynolds numbers, capturing the same lift-curve slope, stall onset, and post-stall behaviour despite being obtained in a different facility using a fundamentally different measurement principle. This cross-facility agreement provides strong mutual validation of both the DTU and TU Delft datasets. The force balance C_D agrees well with the wake rake at low angles of attack but gives higher drag at moderate angles. This divergence is physically consistent with the difference in measurement principle. The force balance integrates the total aerodynamic load over the entire three-dimensional model, including wall-interference and end effects, whereas the wake rake captures only the sectional two-dimensional drag at mid-span. The force balance drag therefore reflects the true three-dimensional drag of the physical model, and shows strong agreement with the Filho et al. [10] reference, which makes good physical sense since both are force balance measurements exhibiting the same elevated drag from three-dimensional effects.

6.2. Wake Rake Traverse Study

The RWT wall-pressure and wake rake campaign is by a wide margin the largest dataset of this thesis, covering all ten geometry configurations across three Reynolds numbers and the full angle-of-attack range. Its results therefore carry significant weight in the comparative analysis of the dimpled configurations, which makes it important to verify the central assumption underlying the method; that a single, fixed-position wake rake measurement at mid-span, as described in Section 3.3, is representative of the sectional aerodynamic performance. This assumption hinges on the spanwise two-dimensionality of the flow. While this is well established for a smooth airfoil, it is less obvious for the dimpled configurations, where the staggered dimple array deliberately trips the boundary layer and introduces periodic spanwise surface perturbations that could, in principle, drive three-dimensional wake structure. This section presents a dedicated traverse study to test whether the wake rake remains a valid measurement approach once surface roughness elements are introduced.

To assess the spanwise uniformity of the wake, the wake rake was traversed to three additional spanwise positions by raising it 5 mm, 7 mm, and 10 mm from the mid-span reference position using flat shim plates. These offsets correspond to $0.5D_d$, $0.7D_d$, and $1.0D_d$ respectively, where $D_d = 10$ mm is the dimple diameter. The study was performed on the Full-5.00 configuration, the airfoil with the fullest chordwise dimple coverage and the deepest dimple depth, as this is intended to represent the worst-case scenario for three-dimensional spanwise effects. If the wake is two-dimensional for Full-5.00, it can be assumed two-dimensional for all other configurations. The measurements were conducted at all three Reynolds numbers tested in the main campaign.

The resulting force polars at each traverse position are shown in Figure 6.4, Figure 6.5, and Figure 6.6 for $Re = 1.0 \times 10^5$, 2.0×10^5 , and 5.0×10^5 respectively. Across all three Reynolds numbers, the lift and drag polars obtained at the four spanwise positions collapse closely onto one another throughout the low to moderate angles of attack. Some scatter is visible in the drag coefficient near stall and at the most negative angles of attack, where the wake is widest and the momentum deficit is most sensitive to spanwise position. Nevertheless, the overall agreement is strong, showing that the wake is invariant with spanwise position over the traversed range at all tested conditions. The fixed mid-span wake rake position used throughout the main campaign is therefore considered representative. Whether the flow is two-dimensional over the full span is a separate question, addressed independently by the tuft visualisation in Section 6.3.




Figure 6.4: Wake rake traverse study for Full-5.00 at $Re = 1.0 \times 10^5$. Lift coefficient C_L (left) and drag coefficient C_D (right) are shown as a function of angle of attack for four spanwise wake rake positions: the mid-span reference position (0 mm) and offsets of 5 mm, 7 mm, and 10 mm, corresponding to $0.5D_d$, $0.7D_d$, and $1.0D_d$.

These results also provide indirect validation for the spanwise extents used in the numerical simulations, although the relevant comparison is not simply one of length. The dimples are arranged in a staggered array with a spanwise centre-to-centre spacing of $D_s = 20$ mm, and adjacent rows are offset by half of that, 10 mm. The traverse range therefore spans the full range of spanwise positions relative to the dimple pattern, from directly behind a dimple to directly behind the gap between two neighbouring dimples, rather than an arbitrary 10 mm sample of the span. That the force polars remain essentially

unchanged across this range, even for the most aggressively dimpled configuration Full-5.00, indicates that the wake does not vary significantly with spanwise phase within the dimple pattern.

On this basis the thin-slab approach is consistent with the present experimental evidence: the RANS slab of $0.2c = 40$ mm contains two complete spanwise periods of the dimple array and the LES slab of $0.1c = 20$ mm contains one, so both resolve the periodic unit that the traverse shows to be aerodynamically uniform.




Figure 6.5: Wake rake traverse study for Full-5.00 at $Re = 2.0 \times 10^5$. Lift coefficient C_L (left) and drag coefficient C_D (right) are shown as a function of angle of attack for four spanwise wake rake positions: the mid-span reference position (0 mm) and offsets of 5 mm, 7 mm, and 10 mm, corresponding to $0.5D_d$, $0.7D_d$, and $1.0D_d$.




Figure 6.6: Wake rake traverse study for Full-5.00 at $Re = 5.0 \times 10^5$. Lift coefficient C_L (left) and drag coefficient C_D (right) are shown as a function of angle of attack for four spanwise wake rake positions: the mid-span reference position (0 mm) and offsets of 5 mm, 7 mm, and 10 mm, corresponding to $0.5D_d$, $0.7D_d$, and $1.0D_d$.

6.3. Spanwise Uniformity from Tuft Visualisation

The traverse study above tests spanwise uniformity over a 10 mm range at a single spanwise reference position, stepped outward in small increments, and on one configuration, Full-5.00. A complementary check is available from the tuft flow visualisation campaign, which instead observes the surface flow across the *entire* span of the models suction side simultaneously, at three chordwise rows of tufts. One representative case, $\alpha = 6^\circ$, is shown in Figure 6.7 below. The remaining angles of the sweep, $\alpha = 0^\circ, 2^\circ, 4^\circ, 8^\circ, 10^\circ, 12^\circ$, and 14° , for both the undimpled Smooth baseline and the dimpled Mid-2.50 configuration at $Re = 2.0 \times 10^5$, are given in full in Section D.3. The angle $\alpha = 16^\circ$, which corresponds

to the stall condition central to the geometry comparison, is presented and discussed separately in Section 7.5.4 rather than as part of the present spanwise-uniformity check.




Figure 6.7: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 6^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15 \text{ m/s}$). The leading edge is to the left. Tufts of yellow yarn are visible across the entire suction surface of the airfoil.

At $\alpha = 6^\circ$, shown above, and at every other angle covered in Section D.3, and for both airfoils, the tufts lie straight and aligned with the freestream across the full span, with no visible cellular structure, no localised regions of deflection, and no asymmetry between one side of the span and the other. This is consistent with the spanwise-invariant force polars established by the traverse study, but extends that check in two ways the traverse study cannot. It covers the full span of the model rather than a 10 mm window, and it covers a second geometry, Mid-2.50, in addition to the worst-case Full-5.00 configuration used in the traverse. Taken together with the traverse results, this provides independent, whole-span confirmation that the flow remains effectively two-dimensional under the conditions tested, and further supports the thin-slab spanwise extents used in the RANS and LES simulations, as already discussed in Section 6.2.

6.4. PIV Validation

Before the PIV system is used to measure the flow over the airfoils, its ability to recover the correct freestream must be confirmed. An empty-tunnel measurement was therefore taken with no model installed, so that the flow in the field of view should be a uniform stream aligned with the tunnel axis at the set freestream velocity. Two quantities establish this. The streamwise component u_x should be uniform across the field of view and equal to the target freestream velocity, confirming that the calibration and processing recover the correct speed. The transverse component u_y should be close to zero everywhere, confirming that the flow is aligned with the tunnel axis and free of any significant cross-flow or bias in the vector fields. The measurement was repeated at the two Reynolds numbers used in the PIV campaign, $Re = 1 \times 10^5$ and 2×10^5 , corresponding to $U_\infty = 7.5 \text{ m/s}$ and 15 m/s .

The two measurements are shown in Figures 6.8 and 6.9. The transverse component u_y (left panels) is close to zero and uniform across the entire field of view at both Reynolds numbers, with no sign of any cross-flow, confirming that the freestream is aligned with the tunnel axis and that the vector fields carry no systematic transverse bias.

The streamwise component u_x (right panels) matches the target freestream velocity over most of the field of view, but shows two localised regions where it departs from it, around $y \approx [-50, 0] \text{ mm}$, $x \approx [100, 150] \text{ mm}$ and $y \approx [-75, -25] \text{ mm}$, $x \approx [-75, 0] \text{ mm}$. These are not flow features but optical artefacts caused by scratches and laser-light reflections in the test-section window, which locally cor-

rupt the correlation and produce dead regions where the recovered velocity does not represent the true streamwise flow. The same imperfections are faintly visible as small noisy patches in the same locations of the $Re = 2 \times 10^5$ transverse field (left panel).




Figure 6.8: Transverse (u_y , left) and streamwise (u_x , right) velocity fields measured by PIV in the empty test section at $Re = 1 \times 10^5$, prior to airfoil installation. The transverse component is uniformly near zero, while the streamwise component matches the target freestream velocity except in two localised regions (around $y \approx -50$ to 0 mm, $x \approx 100$ to 150 mm, and $y \approx -75$ to -25 mm, $x \approx -75$ to 0 mm) where scratches and laser-light reflections in the test-section window produce optical artefacts; velocity data in these regions should be disregarded.

Outside these regions, the streamwise component is uniform and close to the target value at both Reynolds numbers, which is the check required before the airfoil measurements can be interpreted. The scratches and reflections could in principle have been mitigated, but the short and tightly scheduled wind-tunnel campaign left no time to do so within the allocated slot. Because the affected regions are fixed in location, they are simply disregarded when interpreting the flow fields over the airfoil.




Figure 6.9: Transverse (u_y , left) and streamwise (u_x , right) velocity fields measured by PIV in the empty test section at $Re = 2 \times 10^5$, prior to airfoil installation. As in Figure 6.8, the streamwise component is uniform and close to the target freestream velocity apart from the same scratch- and reflection-induced artefact regions; here the imperfections are also faintly visible as small noisy patches in the transverse field. These regions are not physical flow features and should be disregarded when interpreting the freestream uniformity.

7

Results

The preceding chapter validated the smooth baseline against published data, establishing that the measurement and simulation methods reproduce the known aerodynamics of the NACA 63-415 before any dimples are introduced. This chapter turns to the results of the dimpled configurations. The aim is to determine how the aerodynamic performance changes with the chordwise placement of the dimples and with their depth, and to explain why those changes arise in the flow.

The investigation covers two design axes. The first is the chordwise coverage, spanning the Mid, Aft, and Full layouts. The second is the dimple depth, tested at $k = [1.25, 2.50, 5.00]$ mm. Each is examined across three Reynolds numbers and a full angle-of-attack polar. Only a single method covers this whole space, so the campaign is layered instead. A broad and low-cost force-coefficient screening covers the entire geometry matrix and forms the backbone of the chapter, while a set of richer but more detailed methods, namely surface-pressure taps, PIV, tuft visualisation, and CFD simulation, is applied to a smaller subset of configurations to explain the trends the screening reveals. The chapter follows this order. Section 7.1 maps out the available data. Section 7.2 presents the aerodynamic screening across the full geometry matrix, and Sections 7.3 and 7.4 read that screening along its two design axes in turn, chordwise dimple location and dimple depth. Section 7.5 examines the flow field of the standout cases, and Section 7.6 compares the experimental and numerical flow fields against one another at the single condition where all methods overlap. Section 7.7 then relates the findings to the regime-dependent behaviour predicted in Chapter 2.

7.1. Results Overview

The full set of experimental and numerical datasets is summarised in Table 7.1, which lists every method against the ten geometry variants (the undimpled baseline Smooth and the three coverage families Mid, Aft, and Full, each at depths $k = [1.25, 2.50, 5.00]$ mm) and the three Reynolds numbers. The table is the reference map for the rest of the chapter, and each result presented below draws on one or more of its rows.

The campaign is built around a broad screening backed by a set of focused follow-up studies. Two methods make up the screening. On the experimental side, the wall-pressure and wake-rake measurements are the only dataset complete across the entire matrix, covering all ten geometries, all three Reynolds numbers, and the full -10° to 20° polar. Their completeness follows from the low per-case cost of the 3D-printed models and the simultaneous wake-and-wall acquisition. On the numerical side, the RANS study serves the same screening role, covering all four geometries, Smooth, Mid, Aft, and Full, at the reference depth $k = 2.50$ mm and at the two higher Reynolds numbers, which allows the coverage trend seen in the experiment to be checked numerically. Together, these two methods establish what the dimples do to the integrated forces across the geometry space.

The remaining methods are more expensive per case and were therefore reserved for the configurations that the screening flagged as most interesting, concentrated on the baseline and the mid-chord family, Smooth and Mid. The surface-pressure taps resolve the chordwise C_p distribution, and with it the separation and reattachment behaviour that integrated forces alone cannot show. PIV resolves the velocity field and the near-wall momentum that the dimple mechanism is expected to modify. The tuft models give a direct qualitative picture of the surface flow and the onset of separation. The LES provides a scale-resolved flow field for the dimpled case of interest, giving a higher-fidelity reference than RANS can against which the separation and reattachment can be judged. These focused datasets are what establish why the forces behave as the screening shows, on the specific configurations where the effect is strongest.

Table 7.1: Overview of all experimental and numerical datasets in the present campaign, listed by method against the ten geometry variants, baseline Smooth and coverage families Mid, Aft, and Full at dimple depths $k = [1.25, 2.50, 5.00]$ mm. Row shading indicates Reynolds number: blue 1.0×10^5 , green 2.0×10^5 , orange 5.0×10^5 . Angle-of-attack ranges are given in the final column. Completed cases are marked $\checkmark$; a dash denotes a case not run.

Method	Re	Smooth	Mid			Aft			Full			AoA range
			1.25	2.50	5.00	1.25	2.50	5.00	1.25	2.50	5.00	
<i>Experimental</i>												
Wake and Wall pressure	1.0×10^5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	−10° to 20° (step 2°)
Wake and Wall pressure	2.0×10^5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	−10° to 20° (step 2°)
Wake and Wall pressure	5.0×10^5	✓	✓	✓	✓	✓	✓	✓	✓	✓	✓	−10° to 20° (step 2°)
Pressure tap model	1.0×10^5	✓	—	✓	—	—	—	—	—	—	—	−10° to 20° (step 2°)
Pressure tap model	2.0×10^5	✓	—	✓	—	—	—	—	—	—	—	−10° to 20° (step 2°)
Pressure tap model	5.0×10^5	✓	—	✓	—	—	—	—	—	—	—	−10° to 20° (step 2°)
PIV	1.0×10^5	✓	✓	✓	✓	—	—	—	—	—	—	0°, 4°, 6°, 10°, 14°
PIV	2.0×10^5	✓	✓	✓	✓	—	—	—	—	—	—	0°, 4°, 6°, 10°, 14°
Force balance	1.0×10^5	✓	✓	✓	✓	—	—	—	—	—	—	0° to 16° (step 2°)
Force balance	2.0×10^5	✓	✓	✓	✓	—	—	—	—	—	—	0° to 16° (step 2°)
Tuft models	2.0×10^5	✓	—	✓	—	—	—	—	—	—	—	0° to 16° (step 2°)
<i>Numerical</i>												
LES	1.0×10^5	✓	—	✓	—	—	—	—	—	—	—	6°
RANS	1.0×10^5	✓	—	✓	—	—	—	—	—	—	—	6°
RANS	2.0×10^5	✓	—	✓	—	—	✓	—	—	✓	—	0°, 6°, 10°, 12°, 14°, 16°
RANS	5.0×10^5	✓	—	✓	—	—	✓	—	—	✓	—	0°, 6°, 10°, 12°, 14°, 16°

7.2. Aerodynamic Screening

The screening forms the backbone of the results chapter. It covers the full geometry matrix and quantifies how the aerodynamic performance of each dimpled configuration compares with the undimpled baseline Smooth, across the full range of angle of attack and Reynolds number tested. Performance is expressed throughout as the glide ratio C_L/C_D , and every result is reported as a percentage change relative to the Smooth airfoil at the same operating point. Thus, a positive value denotes an improvement over the Smooth airfoil, a negative value a degradation.

Two datasets make up the screening. The experimental data are drawn from the RWT wake-rake and wall-pressure campaign, in which the lift coefficient is obtained from the wall static pressures and the drag coefficient from the wake rake within the same run. It is the only dataset complete across the full matrix in Table 7.1, spanning all ten geometries, the full angle-of-attack polar, and all three Reynolds numbers, and therefore carries the primary weight in the comparison. The numerical screening comes from the RANS simulations, which cover the three coverage families at the reference dimple depth of 2.50 mm and provide an independent check of the experimental trend at the two higher Reynolds numbers.

Results are presented as heatmaps of the percentage change in C_L/C_D , one Reynolds number at a time. Each row is a geometry and depth combination, and each column an angle of attack, so a row reads as the performance of a single dimpled layout and depth across the polar, and a column as the response of every layout at a fixed incidence. Where RANS data exist, the numerical heatmap is placed alongside the experimental one using the same axes, colour scale, and layout, so the two can be compared cell by cell.

For completeness, the appendix Section B.1 and Section B.2 decompose the same screening data into the individual lift and drag coefficients, so that the separate contributions of C_L and C_D to the glide-ratio (C_L/C_D) change can be inspected. The figures share the same axes, colour scale, and row/column layout as the heatmaps in this section. Green cells mark an improvement over the undimpled baseline Smooth (an increase in C_L or a reduction in C_D), red cells a degradation.

Red cells are attributed to a degradation of the boundary layer where the dimples upset the flow locally, lowering lift and/or raising drag. This is a straightforward and expected failure mode and is not analysed case by case in what follows. The discussion below therefore concentrates on the green regions, where the dimples produce a genuine aerodynamic benefit, since these are the outcomes of interest to the present investigation.

As a reminder for the reader, Figure 7.1 repeats the three dimple layouts under comparison: Mid, Aft, and Full, each tested at the three dimple depths 1.25 mm, 2.50 mm, and 5.00 mm, in addition to the undimpled baseline Smooth.




Figure 7.1: The four dimple coverage layouts, reiterated from Figure 3.2.

7.2.1. Reynolds Number 1×10^5

Figure 7.2 shows the RWT screening results at the lowest Reynolds number tested. The matrix is mostly green from 0° to 8° . A pronounced peak occurs around 4° , where several configurations more than double the baseline glide ratio and the best case reaches $+132\%$. This low-incidence window is followed by a sharp and near-total reversal at 10° to 12° , where the entire matrix turns red. The polar then settles into a mixed low-amplitude pattern of small gains and losses from 14° to 16° .

This is the strongest and most consistent improvement seen anywhere in the screening matrix. Its extent is not uniform across configurations. Some remain green across nearly the full 0° to 8° window while others are only intermittently or never positive within it. Section 7.3 examines this variation in terms of the chordwise placement of the dimples.

The individual C_L and C_D decompositions are given in Figure B.1 and Figure B.4. No RANS screening was performed at this Reynolds number. The only numerical data available at $Re = 1 \times 10^5$ is the single-point RANS and LES comparison against Smooth and Mid-2.50 at 6° , presented separately in Section 7.6.




Figure 7.2: Percentage change in C_L/C_D relative to the untreated baseline geometry denoted Smooth, averaged over the two RWT Wake rake and wall pressure test campaigns, as a function of angle of attack for Reynolds numbers $Re = 1 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) paired with three dimple depth variants for each, 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate an improvement over Smooth, and red cells a degradation.

7.2.2. Reynolds Number 2×10^5

Figure 7.3 shows the RWT screening results at $Re = 2 \times 10^5$. Raising the Reynolds number sharply reduces both the extent and the magnitude of the improvement seen at 1×10^5 . The positive region contracts to a narrow band around 2° to 4° , and even there the largest gain is $+22\%$, well below the peak observed at the lower Reynolds number. Across most of the polar the matrix is now red. A particularly pronounced dip occurs at 12° , where several configurations fall below -30% and the worst case reaches -52% . A recovery towards green appears close to stall. By 16° every configuration in the matrix returns a positive value although generally a modest one, while at 14° the picture is still mixed and most configurations remain red. The individual C_L and C_D decompositions underlying this heatmap are given in Figure B.2 and Figure B.5.




Figure 7.3: Percentage change in C_L/C_D relative to the untreated baseline geometry Smooth, averaged over the two RWT Wake rake and wall pressure test campaigns, as a function of angle of attack for Reynolds numbers $Re = 2 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) paired with three dimple depth variants for each, 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate an improvement over Smooth, and red cells a degradation.

Figure 7.4 shows the corresponding RANS screening, available for Mid, Aft, and Full at 2.50 mm. It reproduces the same broad shape as the experiment. Degradation dominates through the low to mid angle of attack range and a partial recovery appears approaching stall. The two datasets do not agree closely in magnitude, however. At 12° , where the experimental screening shows its deepest dip, the RANS degradation is comparatively mild. Conversely, at 14° the RANS screening turns positive for two of the three geometries while the experimental screening remains negative at the same angle for most configurations. The qualitative trend, degradation through the attached regime followed by a recovery near stall, is therefore corroborated by RANS, but the two methods disagree on the angle at which that recovery actually sets in. The individual C_L and C_D decompositions underlying the RANS heatmap are given in Figure B.7 and Figure B.9.




Figure 7.4: Percentage change in C_L/C_D relative to the untreated baseline geometry Smooth, obtained from RANS simulations, as a function of angle of attack at a Reynolds number of $Re = 2 \times 10^5$. Rows correspond to geometry families Mid, Aft, and Full at the 2.50 mm dimple depth, the only configuration evaluated in RANS. The axes, colour scale, and row/column layout match Figure 7.3 to allow direct comparison with the wind-tunnel results; cells marked NaN denote angles of attack or geometry variants for which no RANS data exists.

7.2.3. Reynolds Number 5×10^5

Figure 7.5 shows the RWT screening results at $Re = 5 \times 10^5$. The trend of degradation through the attached regime followed by a recovery near stall at higher Reynolds numbers continues and intensifies further. The matrix is red across almost the entire polar from 0° to 10° , and the losses are larger than at either lower Reynolds number, reaching beyond -60% for some configurations. It is only within the 12° to 16° range, near stall, that a subset of configurations turns positive again, in some cases substantially so with a maximum of $+42\%$, though the majority of the matrix remains red even in this region. The individual C_L and C_D decompositions underlying this heatmap are given in Figure B.3 and Figure B.6.




Figure 7.5: Percentage change in C_L/C_D relative to the untreated baseline geometry Smooth, averaged over the two RWT Wake rake and wall pressure test campaigns, as a function of angle of attack for Reynolds numbers $Re = 5 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) paired with three dimple depth variants for each, 1.25 mm, and 5.00 mm. Green cells indicate an improvement over Smooth, and red cells a degradation.




Figure 7.6: Percentage change in C_L/C_D relative to the untreated baseline geometry Smooth, obtained from RANS simulations, as a function of angle of attack at a Reynolds number of $Re = 5 \times 10^5$. Rows correspond to geometry families Mid, Aft, and Full at the 2.50 mm dimple depth, the only configuration evaluated in RANS. The axes, colour scale, and row/column layout match Figure 7.5 to allow direct comparison with the wind-tunnel results; cells marked NaN denote angles of attack or geometry variants for which no RANS data exists.

Figure 7.6 shows the corresponding RANS screening, available for Mid, Aft, and Full at 2.50 mm. Unlike at 2×10^5 , the RANS degradation here is consistently deeper than the experimental one throughout the low to moderate angles of attack, at almost every angle where the two can be compared from 0° to 12° . Near stall, the two methods diverge differently. For Mid-2.50, the RANS screening turns sharply positive at 14° , well before the experimental screening turns positive for the same configuration, while for Aft-2.50 and Full-2.50 the two methods agree more closely on both the onset and the magnitude of the recovery at 14° and 16° . The qualitative pattern, deeper degradation through the attached regime followed by a recovery approaching stall, in higher Reynolds numbers, is therefore again reproduced by RANS, but the strength of that agreement depends on which coverage family is considered. The individual C_L and C_D decompositions underlying the RANS heatmap are given in Figure B.8 and Figure B.10.

7.3. Effect of Chordwise Dimple Location

The screening in Section 7.2 established a clear Reynolds number trend that holds independently of which specific geometry is considered. The benefit of dimples, where it exists, is concentrated at low Reynolds number and low to moderate angle of attack, while at higher Reynolds numbers the dimples are mostly detrimental across the attached flow regime, with any benefit confined to the immediate vicinity of stall. This trend was established without distinguishing between the three coverage families, and not every geometry follows it to the same degree. This section revisits the same datasets to isolate the effect of chordwise dimple placement, comparing the Mid, Aft, and Full layouts, as introduced in Figure 7.1, at each Reynolds number in turn.

7.3.1. Reynolds Number 1×10^5

Figure 7.2 shows a clear coverage trend at this Reynolds number. Mid and Full perform well throughout the 0° to 8° window, while Aft is red at nearly every angle except 0° , 14° , and 16° . This indicates that the dimples are most effective when placed forward on the airfoil, and that the Mid rows sit at the onset of transition, where the flow is still close to the surface and can be re-energised, whereas the Aft rows most likely sit within a region where the flow has already departed the surface, and there is little left to influence. Full, which combines the Mid rows with the Aft rows, largely mirrors the Mid trend but with generally lower gains, indicating that the rearward dimples added by Aft, and by extension the rear half of Full, degrade rather than aid the performance in this regime.

The individual C_L and C_D decompositions in Figure B.1 and Figure B.4 help distinguish the Mid and Full behaviour. Mid shows a strong drag reduction across the same 0° to 8° window, accompanied by a smaller lift gain, so its glide ratio improvement is driven primarily by C_D . At 0° this relationship reverses, where the lift gain exceeds the drag reduction, at roughly 30% against 8 – 15%. Full shows an even larger drag reduction than Mid, reaching nearly 50% at some angles, but shows essentially no lift gain except a small increase at 4° . Adding the rear dimples that turn Mid into Full therefore removes the lift benefit rather than simply reducing it. Aft, in contrast, which consists only of the rear dimples, despite its poor glide ratio performance, shows a lift gain across most of the polar, its degradation instead comes from a substantial drag increase across nearly the entire 0° to 16° range. Mid and Aft individually, therefore both show mostly lift gains, but combining them into Full removes rather than compounds this benefit. But the drag benefits only seem to arise when forward dimples are present, as on Mid and Full.

Figure 7.7 sums the per-angle contributions across the full polar to give an overall picture of each coverage family. Mid achieves the largest net improvement at the 1.25 and 5.00 mm depths, while Full slightly exceeds Mid at the 2.50 mm depth. In every case, both Mid and Full remain net positive across the full 0° to 16° envelope, meaning their gains outweigh their losses and both configurations outperform the baseline overall, when measured across the entire envelope. Aft shows the opposite pattern. Each depth variant returns a positive sum of gains alone, the solid bars, but the net sum including losses is negative at every depth, confirming that its occasional positive angles are outweighed by its losses elsewhere in the polar.




Figure 7.7: Net improvement in C_L/C_D over Smooth, summed across 0° to 16° at $Re = 1 \times 10^5$, from Figure 7.2, by geometry family and dimple depth. Solid bars sum only positive per-angle contributions; hatched bars give the true net sum.

7.3.2. Reynolds Number 2×10^5

At $Re = 2 \times 10^5$ the picture in Figure 7.3 is markedly different from the lower Reynolds number. The broad 0° to 8° window of improvement seen for Mid and Full at 1×10^5 contracts sharply, surviving only in a narrow band around 2° to 4° , and the exact extent of that band depends on depth: Mid-1.25 remains positive from 2° to 6° , Mid-2.50 from 2° to 4° , and Mid-5.00 from 0° to 4° . Outside this band the matrix is red for nearly every configuration, with a partial recovery reappearing close to stall at 14° and, more consistently, at 16° . Aft remains predominantly red across the polar as before, its only clear exception again confined to the angles nearest stall.

The decomposition into C_L and C_D in Figure B.2 and Figure B.5 shows that the lift gains present at 1×10^5 have essentially disappeared. C_L is flat or slightly negative across nearly the whole polar, with consistent small gains appearing only at 16° and, for a few configurations, at 14° . One exception is Full-2.50, whose lift gain at 16° reaches $+13\%$, a genuine contribution rather than a negligible one. The improvement pattern visible in the glide ratio at 2° to 4° instead closely tracks the C_D decomposition, where the same cells that are green in the C_L/C_D heatmap show a drag reduction of comparable magnitude. The benefit at this Reynolds number is therefore driven almost entirely by drag reduction on the forward-dimpled geometries, rather than by any meaningful lift gain.

$Re = 2 \times 10^5$ lies in the transitional regime for this airfoil, and this is consistent with what the heatmaps show, no lift benefit remains, and only a modest, narrowly confined drag reduction survives. The gains that do remain are confined to the forward-dimpled configurations in the attached-flow region around 2° to 4° , again pointing to the forward rows as the location where the dimples retain some, if much reduced, effect.

Figure 7.8 sums the per-angle contributions across the full polar and confirms that none of the configurations achieves a sustained improvement at this Reynolds number. Considering only the positive per-angle contributions, the solid bars, Mid-1.25 and Mid-5.00, show the largest raw gains. Once the losses across the rest of the polar are included, the hatched bars, every configuration returns a net negative value, in contrast to $Re = 1 \times 10^5$ where Mid and Full exceeded the baseline overall. Among these, Full-1.25 comes closest to break-even with the smallest net loss of any configuration tested, while Full-5.00 and Aft-2.50 show the largest net losses, and Aft shows much the same pattern as in Figure 7.7. The RANS screening at this Reynolds number (Figure 7.4) reproduces this coverage split reasonably well, with Mid and Full again showing the least degradation and the same partial recovery approaching stall. The main disagreement with the experiment remains concentrated at 14° and 16° , consistent with what was already noted for the un-split screening comparison.




Figure 7.8: Net improvement in C_L/C_D over Smooth, summed across 0° to 16° at $Re = 2 \times 10^5$, from Figure 7.3, by geometry family and dimple depth. Solid bars sum only positive per-angle contributions; hatched bars give the true net sum.

7.3.3. Reynolds Number 5×10^5

Figure 7.5 shows the RWT results at the highest Reynolds number tested, and the picture changes substantially again. Every benefit previously found at low to moderate angle of attack is gone, with a single exception at 2° for Aft-1.25. Mid shows no improvement whatsoever at any angle or depth. Full recovers a handful of gains near stall, at 14° and 16° for the 1.25 and 2.50 mm depths, though 5.00 mm remains entirely red. Aft, by contrast, now shows a genuine and consistent improvement at its deepest variant, 5.00 mm, from 12° to 16° , with a comparable though smaller recovery for 2.50 mm at 14° and 16° .

The C_L and C_D decompositions in Figure B.3 and Figure B.6 show a pattern similar to 2×10^5 . Lift contributions remain marginal across the board, with only a few percentage points of gain or loss anywhere in the matrix. The most consistent lift improvement is now found on Aft, across all three depths, at low to moderate angle of attack, occupying much the same 0° to 6° region where Mid and Full produced their lift gains at $Re = 1 \times 10^5$. The glide ratio pattern again tracks the C_D map closely. Mid and Full show very large drag penalties at low to moderate angle of attack, most severe for the deepest variants, where the forward dimples now provide no drag benefit at all and only a negligible, isolated lift gain. This is the reverse of the picture found at the lowest Reynolds number, where the same forward placement produced the strongest drag reduction in the whole campaign.

Figure 7.9 confirms this reversal across the full envelope. Mid is the worst-performing family at every depth, with a solid, positive-only sum of exactly 0%, meaning no angle offers any improvement, and a net sum reaching as low as -394% once losses are included. Full is similarly poor overall, though less extreme, with its 2.50 mm variant showing the family's best result and still a net loss of -130% . Aft is the only family to show a genuine positive contribution, its solid bars reaching as high as $+72.5\%$ at 5.00 mm, the largest positive sum of any configuration here, though even this is outweighed by losses elsewhere in the polar, leaving every Aft depth net negative too.

The RANS screening at this Reynolds number (Figure 7.6) behaves much as it did at 2×10^5 . Good qualitative agreement with the experiment, generally at larger magnitude in both directions. One clear disagreement stands out. RANS predicts an improvement for Mid at 14° and 16° that is entirely absent from the experimental screening, where Mid is red at every angle tested. This is not simply a difference of degree, as seen elsewhere between the two methods, but a disagreement in sign, and suggests the RANS simulation is not capturing the flow behaviour correctly for this configuration in this range.




Figure 7.9: Net improvement in C_L/C_D over Smooth, summed across 0° to 16° at $Re = 5 \times 10^5$, from Figure 7.5, by geometry family and dimple depth. Solid bars sum only positive per-angle contributions; hatched bars give the true net sum.

7.4. Effect of Dimple Depth

The previous section showed a consistent pattern across the three Reynolds numbers. The Mid layout is beneficial at low Reynolds number and detrimental at high Reynolds number, while Aft shows the reverse trend, and Full only partially combines the two. Of the three families, Mid stood out as the strongest performer at $Re = 1 \times 10^5$ and 2×10^5 , and as the weakest at $Re = 5 \times 10^5$. This makes Mid the natural candidate for a closer look at how dimple depth shapes the gains already identified, and the remainder of this section focuses on its three depth variants, 1.25 mm, 2.50 mm, and 5.00 mm. The heatmaps introduced earlier form the basis of this analysis, supported here by a closer look at the wake rake measurements and by the SLT force balance campaign.

7.4.1. Reynolds Number 1×10^5

At $Re = 1 \times 10^5$, Mid showed large gains in the glide ratio (Figure 7.2), but no clear relationship between the size of the gain and dimple depth is visible in C_L/C_D itself, since the underlying C_L trend is scattered rather than systematic (Figure B.1). The C_D decomposition is more informative (Figure B.4). From 0° to 8° , in the low-incidence regime, deeper dimples give progressively greater drag reduction. From 10° onward the dimples instead incur a drag penalty, and here the ordering is less clean, though the deepest variant, 5.00 mm, is consistently the worst performer at every angle in this band.

Figure 7.10 isolates this C_D trend for Mid directly. It confirms the monotonic ordering from 0° to 8° , with drag reduction increasing steadily with depth. From 10° the sign reverses for all three depths, and the ordering between the 1.25 and 2.50 mm depths is no longer consistent, though 5.00 mm remains the clear outlier with the largest penalty throughout. This breaks down further at 14° , where 2.50 mm alone returns to a drag reduction while 1.25 and 5.00 mm both remain penalised, pointing to a practical limit on how deep a dimple can be made before it begins to degrade rather than aid the flow, a limit also visible in Aft and Full.

The wake profile at 6° (Figure 7.11) shows this depth dependence directly. The wake narrows progressively as dimple depth increases, and this narrowing is confined entirely to the suction side of the wake profile (right-hand side). The pressure side (left-hand side), unaffected by the dimples, collapses onto the same curve for every geometry, as expected. The minimum wake velocity is essentially unchanged across all four cases, so the drag reduction with depth comes from a narrower wake rather than a shallower one.




Figure 7.10: Percentage change in C_D from the undimpled baseline geometry Smooth, as a function of angle of attack at $\text{Re} = 1 \times 10^5$, for the three Mid dimple-depth variants (1.25 mm, 2.50 mm, 5.00 mm) from the RWT test campaign.

This is the same mechanism described for the golf ball in Section 2.4. There, tripping the boundary layer produces a turbulent layer that resists the adverse pressure gradient for longer, delaying separation and narrows the wake behind the body (Figure 2.12). The behaviour observed here on the dimpled side of the airfoil is consistent with the same process, a deeper dimple trips the boundary layer more strongly, separation is delayed further downstream, and the wake narrows accordingly, exactly as illustrated for the sphere in Figure 2.12.




Figure 7.11: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $\text{Re} = 1 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

The wake profile at 14° (Figure 7.12) illustrates the departure from this simple picture noted above. Here, Mid-1.25 closely tracks the baseline wake, giving essentially no benefit, while Mid-5.00 produces an even deeper and wider wake than the baseline, a clear degradation. Mid-2.50 alone breaks from both, showing a narrower wake than any of the other three cases and confirming the drag reduction already seen in the C_D heatmap. Depth and performance are therefore not related in a simple monotonic way near stall, and this non-monotonic behaviour is one of the underlying motivations for making dimple depth an adjustable parameter, as a clear optimal dimple depth exists.




Figure 7.12: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $Re = 1 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

The SLT campaign, in which Mid-1.25, Mid-2.50, and Mid-5.00 were also tested on the force balance, provides an independent check on these trends. The corresponding heatmap and its C_L and C_D decompositions are given in Figure D.1, Figure D.3, and Figure D.5. The SLT results largely support the RWT findings, with good agreement on where gains and penalties occur across 0° to 12° , and some disagreement appearing only near stall, from 14° to 16° . The exact values differ between the two campaigns, which is expected since the SLT force balance captures three-dimensional effects that the RWT wake rake does not. That the same trends and the same regions of benefit appear in two independent facilities using two different measurement principles indicates that the effect of the dimples is not specific to a single wind tunnel setup, further validating the findings of the effects of varying dimple depth.

7.4.2. Reynolds Number 2×10^5

At $Re = 2 \times 10^5$ the picture for Mid is more nuanced than at the lower Reynolds number. As already established in Section 7.3.2, the gains here are smaller overall and are driven almost entirely by C_D , since the C_L contribution stays within one or two percent at nearly every angle (Figure B.2).

The C_D decomposition (Figure B.5) does not show a single trend that holds across the whole polar, but a clear depth correlation does appear at specific angles, most notably at 0° , 6° , 8° , and 14° , though the direction of the correlation is not consistent between them. The wake profiles at 6° and 14° illustrate this directly.




Figure 7.13: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $Re = 2 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

At 6° (Figure 7.13) the wake for Mid-1.25 shows a narrower but deeper wake compared to the baseline and gives a small drag reduction, while Mid-2.50 and Mid-5.00 both show a deeper and wider wake than the baseline. The degradation increases with depth, with Mid-5.00 producing the widest and deepest wake of the three, consistent with a monotonic worsening of drag with dimple depth at this angle.

At 14° (Figure 7.14) the trend reverses. All three depths now show a deeper and wider wake than the baseline, so all incur a drag penalty, but the penalty is largest for Mid-1.25 and smallest for Mid-5.00. Depth and drag penalty are therefore still monotonically related at this angle, only in the opposite order, compared to 6° .




Figure 7.14: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $Re = 2 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

The summary bar chart introduced in Section 7.3.2 (Figure 7.8) reflects the same depth dependence at the level of the full envelope. For Mid, and similarly for Full, the net loss across the polar grows with dimple depth, so the deepest variant carries the largest overall penalty in both families. Aft shows a comparable but distinct pattern restricted to its positive contributions alone, where the solid, gains-only sum grows with depth, though this remains modest at every depth tested.

The SLT results for Mid at this Reynolds number, given in Figure D.2 together with its C_L and C_D decompositions in Figure D.4 and Figure D.6, show a broadly similar picture, though the agreement with RWT is weaker here than at $Re = 1 \times 10^5$. The two campaigns agree on whether a given angle is a gain or a penalty across roughly 2° to 14° , but the exact magnitudes differ throughout. As at the lower Reynolds number, the two independent methods agree on the underlying trend even where their absolute values do not coincide.

7.4.3. Reynolds Number 5×10^5

Neither the RANS nor the SLT campaign covers $Re = 5 \times 10^5$, so the RWT measurements are the only basis available for assessing dimple depth at this Reynolds number.

As established in Section 7.3.3, the gains previously found for Mid are absent at this Reynolds number (Figure 7.5). As at $Re = 2 \times 10^5$, the lift contribution is negligible and incoherent, remaining within a few percent at almost every angle (Figure B.3), so C_D again drives the glide ratio trend (Figure B.6). A clear depth correlation appears from 0° to 8° , where the drag penalty grows monotonically with dimple depth at every angle in this band. From 10° to 16° the ordering between depths becomes scattered, with one notable exception discussed below. The same 0° to 8° depth correlation, and the same scatter beyond it, is also present for Aft and Full, with the rearward dimples on these two families recovering some gains in the 10° to 16° region, as already discussed in Section 7.3.3.

The wake profile at 6° (Figure 7.15) illustrates this attached-regime trend directly. The wake widens and deepens progressively as dimple depth increases, the reverse of the $Re = 1 \times 10^5$ behaviour, where deeper dimples narrowed the wake relative to the shallower variant. This indicates that at this Reynolds number the dimples increase mixing in the boundary layer in an unfavourable way, with deeper dimples producing more mixing and a correspondingly larger wake.




Figure 7.15: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $\text{Re} = 5 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

The wake profile at 14° (Figure 7.16) exposes a further detail not visible in the C_D heatmap alone. Here the depth ordering is genuinely monotonic, with drag penalties of 37%, 43%, and 48% for the 1.25, 2.50, and 5.00 mm depths respectively, a modest eleven percentage point spread across the full depth range. Despite the similarity in penalty magnitude, the position of the wake deficit shifts noticeably with depth. The deepest variant, Mid-5.00, sits closer to the chord centreline than the baseline Smooth, while the Smooth wake is deflected furthest away from it. Increasing dimple depth at this angle therefore changes not only the size of the wake but also its position relative to the baseline, a distinction that the C_D heatmap and the summary bar chart cannot capture on their own.




Figure 7.16: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $\text{Re} = 5 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

The summary bar chart (Figure 7.9) reflects the same depth dependence at the level of the full envelope already noted in Section 7.3.3. For Mid the net loss grows with dimple depth and no depth variant returns a positive contribution anywhere in the polar. Aft shows the opposite relationship for its positive contributions alone, where deeper dimples yield a larger gains-only sum, though this remains outweighed by losses elsewhere at every depth tested.

7.5. Investigation of Flow Behaviour of Mid-2.50

The chordwise placement and dimple depth investigations identified Mid-2.50 as the strongest performing configuration at $Re = 1 \times 10^5$. This section examines that configuration in detail to establish what happens at the airfoil surface and within the boundary layer to produce the gains already seen in the screening. The PIV campaign, the surface pressure tap model, and the tuft flow visualisation are used in turn to build this picture, focusing on the specific angles of attack and Reynolds numbers most relevant to the trends already identified. The full set of results for each method, across the entire test matrix, is given in the appendix: the pressure tap results in Appendix C, the PIV visualisations in Section D.2, and the tuft flow visualisations in Section D.3.

7.5.1. PIV Flow Field

Figure 7.17 and Figure 7.18 compare the time-averaged PIV velocity field of Smooth and Mid-2.50 at $Re = 1 \times 10^5$, at 0° and 4° respectively. These two angles were chosen because Mid-2.50 shows some of its largest glide ratio gains here, +15% at 0° and +87% at 4° (Figure 7.2).

At 0° (Figure 7.17) the two cases already show a clear contrast in the near-wall flow. Smooth has a pronounced low-velocity region hugging the suction surface, with velocities falling to as low as 0 to 2 m/s in the boundary layer. Mid-2.50 shows a markedly thinner and higher-momentum layer in the same region, with velocities closer to 1 to 3 m/s at the wall and a noticeably reduced wake width immediately downstream of the trailing edge. This is consistent with the drag reduction already noted at this angle (Figure B.4), and shows that the dimples are already energising the near-wall flow at 0° , even before the more dramatic separation-suppressing effect seen at 4° below.




Figure 7.17: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ at angle of attack 0° , comparing the undimpled baseline geometry Smooth (a) to the Mid-2.50 dimple-depth variant (b). The airfoil is masked in white. The shadow region where the laser sheet is blocked by the airfoil and no valid PIV data exists is masked in grey. The hatched region marks an area affected by excessive laser-light reflections and scratches in the test-section window and should be disregarded. For Mid-2.50 (b), the four grey half-circles on the suction-side surface mark the chordwise locations of the four dimple rows.

At 4° (Figure 7.18) the difference is much more pronounced. Smooth shows the classic signature of boundary layer separation, a thin attached layer near the leading edge that thickens rapidly from roughly a quarter chord onward into a wide region of low momentum extending well into the wake. This separated shear layer is the source of the large wake and elevated drag associated with Smooth at this angle. Mid-2.50 shows no equivalent region. The high-velocity flow remains attached to the surface along the full length of the suction side, similar to Figure D.9c, with no thickening low-momentum layer forming downstream. The dimples evidently mix sufficient high-momentum fluid down to the wall to prevent separation from occurring at all at this angle, and the wake behind the airfoil is correspondingly reduced to close to nothing. This is entirely consistent with the 42% drag reduction recorded for Mid-2.50 at 4° in Figure B.4.




Figure 7.18: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ at angle of attack 4° , comparing the undimpled baseline geometry Smooth (a) to the Mid-2.50 dimple-depth variant (b). The airfoil is masked in white. The shadow region where the laser sheet is blocked by the airfoil and no valid PIV data exists is masked in grey. The hatched region marks an area affected by excessive laser-light reflections and scratches in the test-section window and should be disregarded. For Mid-2.50 (b), the four grey half-circles on the suction-side surface mark the chordwise locations of the four dimple rows.

The lift gain of 10% recorded for Mid-2.50 at this angle (Figure B.1) is also visible in the flow field. A small region of elevated velocity, reaching up to 10 m/s, appears over the forward half of the suction surface in the Mid-2.50 case, a level of local acceleration not present anywhere on the Smooth surface at the same angle.

7.5.2. Surface Pressure Distributions

Looking further into this change in near-surface momentum, the surface pressure tap results are presented in Figure 7.19, taken from the RWT surface pressure campaign, where the Smooth and Mid-2.50 surface pressure at $Re = 1 \times 10^5$ and 4° is visualised. Blue vectors denote a negative C_p and orange a positive C_p .

The two vector fields are broadly similar in shape and the difference between them is subtle to read directly from this plot alone. The clearer comparison is the difference plot in Figure 7.20, which isolates the change in C_p between the two cases at each tap.

The forward three-quarters of the chord on the suction side shows a consistent outward vector, indicating an increase in suction magnitude for Mid-2.50 relative to Smooth. This is consistent with the higher near-wall momentum already identified in the PIV field at this angle (Figure 7.18), and with the associated gain in lift. Over the final quarter of the chord, this trend reverses, with the vectors pointing inward, indicating a reduction in suction magnitude. This is the pressure recovery signature of an airfoil whose flow stays attached, the pressure at the rear of the suction side moves closer to ambient because the boundary layer has not separated, in contrast to Smooth, where the separated wake keeps the pressure at the rear taps further from ambient. A smaller difference is also visible on the pressure side, where Mid-2.50 shows a slightly smaller C_p magnitude than Smooth throughout.

The tap nearest the leading edge, at the stagnation point, is absent from the difference plot. Since $C_p \approx 1$ at the stagnation point irrespective of incidence, this shows only that the stagnation point falls on the same tap for both models, and is therefore consistent with the two models having been mounted at the same effective angle of attack. Note also that the reference vector in Figure 7.20 is drawn at a larger scale than in Figure 7.19, since the changes being shown here are considerably smaller than the absolute pressures in the previous figure.




Figure 7.19: Surface pressure coefficient C_p represented as vectors normal to the airfoil surface at each pressure tap, for the undimpled baseline geometry Smooth (top) and the Mid-2.50 dimple-depth variant (bottom) at $\text{Re} = 1 \times 10^5$ and angle of attack 4° . Vectors point outward from the surface with magnitude $|C_p|$. The colour orange indicates $C_p > 0$ and blue indicates $C_p < 0$. Dimple locations are marked as grey inward cavities on the Mid-2.50 suction surface.




Figure 7.20: Change in surface pressure coefficient ΔC_p (Mid-2.50 – Smooth) at each pressure tap location, represented as vectors normal to the airfoil surface, at $\text{Re} = 1 \times 10^5$ and angle of attack 4° . Vector direction indicates whether the magnitude of C_p increased (outward) or decreased (inward) relative to Smooth. The colour orange indicates $C_p > 0$ and blue indicates $C_p < 0$. Dimple locations are marked as grey inward cavities on the Mid-2.50 suction surface.

7.5.3. Wake Rake Momentum Loss

The wake rake measurement at the same condition, $Re = 1 \times 10^5$ and 4° , gives a further view of this same change (Figure 7.21). The Mid-1.25 and Mid-5.00 curves are also included in this figure for completeness but are not the focus here. The Smooth wake is close to twice the width of the Mid-2.50 wake, confirming at the wake rake the same difference already seen directly in the PIV field (Figure 7.18). As in the earlier depth study, the pressure side of the wake profile is near identical across all four geometries, and the narrowing is confined entirely to the suction side, consistent with the dimples acting only on that surface. This is the same delay of separation already discussed in the context of the golf ball mechanism in Figure 2.12. The dimples trip the boundary layer, the resulting turbulent layer resists separation for longer, and the wake narrows as a direct result, exactly as observed here in the PIV field, the surface pressure taps, and the wake rake profile together.




Figure 7.21: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $Re = 1 \times 10^5$ and angle of attack 4° , comparing the undimpled baseline geometry Smooth to the Mid dimple geometry at all three depth variants: 1.25 mm, 2.50 mm, and 5.00 mm.

7.5.4. Stall Delay Benefits

A further beneficial effect observed for Mid-2.50 is a stall delay at $Re = 2 \times 10^5$. The baseline Smooth stalls around $\alpha = 16^\circ$, whereas Mid-2.50 remains unstalled at the same angle, and this is corroborated independently by three measurement techniques. The effect is already visible in the screening heatmap (Figure 7.3), where Mid-2.50 shows a small rise in C_L/C_D around stall driven by a modest lift gain and drag reduction. The surface pressure, force balance, and tuft measurements below examine this condition in detail and together give a fuller picture of the stall delay.

Surface Pressure Distribution

Figure 7.22 shows the surface pressure coefficient from the RWT pressure tap model at $\alpha = 16^\circ$. The pressure sides of the two airfoils are nearly identical, as expected since the dimples are confined to the suction side, so the differences of interest are entirely on the suction side.

The clearest indication that Mid-2.50 remains attached where Smooth has begun to break down is the scatter in the measurement itself. The leading-edge suction-side taps on Smooth carry a markedly larger standard deviation than the corresponding taps on Mid-2.50, the unsteadiness expected of a separating or intermittently stalling flow, whereas Mid-2.50 shows no comparable scatter anywhere on the suction side, consistent with a steady attached boundary layer. The mean pressure curves themselves support the same picture, if more subtly. Mid-2.50 sustains a marginally stronger suction peak over the first $0.1c$, consistent with the leading-edge flow staying attached for longer before separating, and recovers to a pressure slightly closer to ambient over the aft chord from roughly $0.2c$ to $0.9c$, consistent with a fuller pressure recovery and reduced pressure drag relative to Smooth.




Figure 7.22: Surface pressure coefficient $-C_p$ against chordwise position x/c at $\text{Re} = 2 \times 10^5$ and $\alpha = 16^\circ$, comparing the baseline Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.

Force Balance

The SLT force balance measurements (Figure 7.23) show the same stall delay in the integrated forces. Smooth stalls sharply at $\alpha = 16^\circ$, with the lift coefficient dropping abruptly and the drag coefficient rising steeply. Mid-2.50 shows neither. Its lift remains and is roughly 30% above Smooth at 16° , and its drag at the same angle is close to half that of the baseline. The continuous C_L and C_D curves of Mid-2.50, against the sharp drop in lift and sharp rise in drag for Smooth, indicate that the dimpled airfoil remains within its operating range at an angle where the baseline has already stalled.

The magnitude of this improvement is larger than the gain recorded at the same condition in the RWT screening of Section 7.3.2, where the rise in C_L/C_D near stall is present but modest. A clear difference between the campaigns accounts for the discrepancy. The RWT wake-rake and wall-pressure measurements return a sectional result taken at mid-span. In contrast, the force balance integrates the load over the full model span, including the regions closest to the tunnel walls. The tuft pictures (Section D.3) show that these regions are not two-dimensional. The rows of tufts adjacent to the side walls are deflected and unsteady at the trailing edge at every angle of attack tested, indicating a wall-junction disturbance that contributes additional drag independently of the dimples. These effects act differently on the two models once the baseline stalls, since a stalled section is considerably more sensitive to spanwise non-uniformity than an attached one. The two campaigns should therefore be read as agreeing on the existence and the sign of the stall delay, which all three methods confirm independently, while the absolute magnitude is more reliably taken from the sectional RWT measurement.





Figure 7.23: Lift coefficient C_L (left) and drag coefficient C_D (right) against angle of attack at $\text{Re} = 2 \times 10^5$, measured with the SLT force balance, comparing the baseline Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid).

Tuft Visualisation

Figure 7.24 shows the tuft response on Smooth and Mid-2.50 at $\alpha = 16^\circ$ and $Re = 2.0 \times 10^5$, viewed through the test-section side wall with the leading edge to the left. It provides a direct visual confirmation of the stall delay inferred from Figure 7.22 and Figure 7.23.

On Smooth the tufts across the entire suction surface deflect in multiple directions with no coherent streamwise alignment, the signature of a large, fully separated region. Only one or two tufts at the leading edge near mid-span retain any streamwise tendency, and the trailing-edge tufts are especially erratic, consistent with reversed flow. The airfoil is fully stalled.



(a) Smooth baseline airfoil at $\alpha = 16^\circ$. Tufts across the entire suction surface deflect in multiple directions, indicating full stall. Only one or two tufts near the leading edge at mid-span show any streamwise alignment.



(b) Dimpled airfoil Mid-2.50 at $\alpha = 16^\circ$. Tufts at mid-span are nearly straight from leading to trailing edge, indicating attached flow. Separated cells are visible near each span end between $x/c \approx 0.50$ and the trailing edge, consistent with wall-junction end-effects.

Figure 7.24: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 16^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15 \text{ m/s}$). The leading edge is to the left. The tufts demonstrate that the dimpled configuration maintains substantially more attached flow compared to the fully stalled baseline at this condition.

On Mid-2.50 the picture is markedly different. The mid-span tufts across all chordwise rows lie nearly straight from leading to trailing edge, indicating attached flow over the central span, and the leading-edge tufts are uniformly straight across the full span, in sharp contrast to Smooth. Two separated cells are visible, one near each span end, centred between the mid-chord and the trailing edge. Because these cells appear symmetrically at both ends and in the same location, they are attributed to wall-junction separation where the airfoil meets the tunnel side walls, a three-dimensional end effect of testing a finite model in a tunnel, rather than to the airfoil aerodynamics themselves. Disregarding these two end regions, the leading edge and the entire mid-span show attached flow, confirming that the dimples maintain attachment and delay stall at this condition.

7.6. Experimental and Numerical Comparison of Mid-2.50

To compare the numerical and experimental methods directly, a single case at $Re = 1 \times 10^5$ and 6° was examined across PIV, LES, and RANS. The purpose is to assess how well each numerical method captures the flow physics already established for Mid-2.50 in the previous sections. A comparison of surface pressure is also available at this condition, drawing on the pressure tap campaign alongside the same RANS and LES cases.

7.6.1. Flow-field Comparison

Figure 7.25, Figure 7.26, and Figure 7.27 show the flow field for PIV, LES, and RANS, respectively, all at $Re = 1 \times 10^5$ and 6° .

The PIV field (Figure 7.25) shows a picture very similar to the 0° comparison discussed earlier (Figure 7.17). Smooth again carries a visibly thick, low-momentum region along the suction surface, while Mid-2.50 again shows a much thinner near-wall layer with higher velocity sustained close to the surface, consistent with the dimples energising the boundary layer and mixing higher-momentum fluid down to the wall. As at 0° , the wake also differs between the two, with Mid-2.50 showing a slightly smaller and narrower wake than Smooth. The high-velocity core and the surrounding high-velocity region on the suction side appear more uniform and better developed on Mid-2.50 than on Smooth. Despite this, the RWT screening shows the lift difference between the two at this angle to be negligible, with the glide ratio gain driven almost entirely by drag reduction (Figure B.1, Figure 7.2). This suggests that the patchy, non-uniform regions visible away from the surface in the Smooth field, outlined by the black dashed ellipses in Figure 7.25(a), are more likely an artefact of laser-light scattering and reflections than genuine pockets of reduced velocity.




Figure 7.25: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (a) to the Mid-2.50 dimple-depth variant (b). The airfoil is masked in white. The shadow region where the laser sheet is blocked by the airfoil and no valid PIV data exists is masked in grey. The hatched region marks an area affected by excessive laser-light reflections and scratches in the test-section window and should be disregarded. For the Mid-2.50 case (b), the four grey half-circles on the suction-side surface mark the chordwise locations of the four dimple rows. The dashed ellipses in (a) mark the localised non-uniform regions discussed in the text, attributed to laser-light scattering rather than to genuine velocity deficits.

The LES field (Figure 7.26) is, as expected for a simulated result, considerably more uniform than the PIV measurement. Unlike the PIV comparison, LES shows no clear difference in the near-wall flow between Smooth and Mid-2.50, the boundary layer region looks essentially the same for both geometries. The wake behind the airfoil does appear marginally narrower for Mid-2.50 than for Smooth, in the same direction as the PIV result, and this is consistent with Table 7.2, where LES, together with every experimental method, shows a reduction in C_D from Smooth to Mid-2.50. However, the LES lift coefficient shows a lift penalty for Mid-2.50 relative to Smooth, in contrast to the experimental result,

where the lift difference at this angle is negligible. Given that LES shows almost no visible change in the near-wall boundary layer between the two geometries, this suggests that LES is reproducing the correct trend in drag but is not resolving the dimple-induced mixing mechanism itself with the same fidelity as the PIV measurement.




Figure 7.26: Time-averaged velocity magnitude field obtained from LES at $Re = 1 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (a) to the Mid-2.50 dimple-depth variant (b). The airfoil is shown in white. For the Mid-2.50 case (b), the four grey half-circles on the suction-side surface mark the chordwise locations of the four dimple rows.




Figure 7.27: Velocity magnitude field obtained from RANS simulation at $Re = 1 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (a) to the Mid-2.50 dimple-depth variant (b). The airfoil is shown in white. For the Mid-2.50 case (b), the four grey half-circles on the suction-side surface mark the chordwise locations of the four dimple rows.

The RANS field (Figure 7.27) presents a two-sided picture. Visually, it resembles the PIV result more closely than LES does, matching the size and location of the high velocity region, the suction bubble, the separation, and the wake more closely than LES manages. On closer inspection, however, the

steady-state nature of the solver becomes apparent, and it cannot capture the dynamic behaviour of an airfoil, particularly with dimples present.

As with LES, the difference between the two RANS cases is small. The suction bubble and high velocity core appear nearly identical for Smooth and Mid-2.50, with no visible change introduced by the dimples. Compared to PIV, the length of this bubble is a better match than the one obtained from LES, though its width is considerably thicker than in either the PIV or LES fields. This is reflected in Table 7.2, where RANS overestimates C_L for both geometries relative to their experimental counterparts and to LES.

The wake tells a similarly mixed story. RANS again shows little difference between the two geometries, where LES captures the dimple-induced narrowing more convincingly. Yet the RANS C_D for Mid-2.50 sits closer to the experimental values than the LES prediction does, even though the underlying trend is wrong: RANS predicts a drag increase for Mid-2.50 relative to Smooth, the opposite of every experimental method and of LES. This also exposes that the RANS C_D for Smooth itself is considerably off compared to both LES and the experimental methods. RANS therefore appears to land closer to the correct drag level for Mid-2.50 specifically, while getting both the Smooth baseline and the trend between the two geometries wrong, consistent with the known limitations of a steady solver without a transition model in capturing dimple-induced boundary layer mixing.

Table 7.2: Lift and drag coefficients at $Re = 1 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry (Smooth) to the Mid-2.50 dimple-depth variant across all available experimental and numerical methods. RWT Pressure Taps are obtained by integrating the surface pressure distribution from the wind-tunnel pressure taps and include pressure drag only (no skin-friction contribution), so C_D is expected to undershoot the RWT Wake Survey and SLT Force Balance values, which directly capture total drag.

Method	Smooth		Mid-2.50	
	C_L	C_D	C_L	C_D
RWT Pressure Taps	0.5560	0.0535	0.6606	0.0389
RWT Wake and Wall	0.7401	0.0375	0.7352	0.0288
SLT Force Balance	0.7134	0.0436	0.7708	0.0323
LES	0.7327	0.0448	0.6644	0.0406
RANS	0.9102	0.0247	0.8329	0.0279

7.6.2. Surface Pressure Comparison

A further comparison is made between the surface pressure recorded by the RWT pressure tap model and the corresponding RANS and LES results under the same condition. Figure 7.28 compares the three methods for Smooth, and Figure 7.29 makes the same comparison for Mid-2.50.

The Smooth comparison (Figure 7.28) explains much of what was already observed in the flow-field comparison. The wider, better developed high-velocity region and the higher lift coefficient seen for RANS are both apparent here, with RANS showing a substantially higher suction-side $-C_p$ than either LES or the RWT pressure taps throughout the forward half of the chord. The longer bubble associated with RANS is likewise visible, its suction-side curve falls off more gradually than the LES curve, remaining elevated further aft.

Both CFD methods deviate from the RWT measurement in the same way: the suction-side pressure is over-predicted over the forward half of the chord and under-predicted over the rear half, which is consistent with both models' inability to accurately reproduce the wake width and boundary layer seen in the PIV. The pressure side, by contrast, tracks closely across all three methods and shows good general agreement, which is expected since this side of the airfoil is far more passive and its flow structure is easier to capture. It is also the only region where a comparison with PIV itself is not possible, given the nature of that measurement. Taken together, this confirms that most of the disagreement between the two CFD methods and experiment originates on the suction side, unsurprising given that this is the side carrying essentially all of the dynamic flow behaviour, even for a plain, undimpled airfoil.




Figure 7.28: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c for the undimpled baseline geometry **Smooth** at $Re = 1 \times 10^5$ and angle of attack 6° , comparing RWT pressure-tap measurement (blue), RANS (green), and LES (orange). This is the only AoA/Reynolds-number combination where surface pressure data from all three methods is available.

For Mid-2.50 (Figure 7.29) the overall picture is much the same. RANS again shows the highest suction-side $-C_p$ of the three methods, consistent with its overshoot already established for Smooth. Here, however, the RANS curve tracks the RWT measurement more closely than it did for Smooth, particularly around mid-chord, where the two nearly coincide. LES for Mid-2.50 agrees closely with RWT over roughly the first 30% of the chord before falling below both RWT and RANS further aft. The pressure side again shows the same close agreement across all three methods observed for Smooth.

One further feature stands out in the RWT curve itself. Between roughly $0.35c$ and $0.6c$ the RWT suction-side pressure holds a shallow plateau before dropping off sharply around $x/c = 0.6$ to 0.7 . This plateau is not present in the Smooth comparison, so it is most likely a direct signature of the dimples, consistent with the fuller pressure recovery over the aft chord already identified elsewhere for Mid-2.50. It is worth distinguishing this feature from the plateau associated with a laminar separation bubble in Section 2.2.3. The chordwise extent here, $0.35c$ to $0.59c$, coincides with the dimpled band rather than with the natural separation location on the smooth baseline, and the feature is absent from the Smooth curve at the same condition, so it is attributed to the localised recirculation within and immediately downstream of the dimples rather than to a bubble forming on an otherwise smooth surface. Neither RANS nor LES reproduces this feature, both fall away smoothly through this region with no equivalent plateau or sharp drop, a further indication that neither CFD method is capturing the specific mechanism by which the dimples modify the suction-side pressure distribution.




Figure 7.29: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c for the **Mid-2.50** dimple-depth variant at $Re = 1 \times 10^5$ and angle of attack 6° , comparing RWT pressure-tap measurements (blue), RANS (green), and LES (orange). This is the only AoA/Reynolds-number combination where surface pressure data from all three methods is available.

7.7. Summary of Results

The preceding sections examined the chordwise coverage, the dimple depth, the flow field, and the agreement between experimental and numerical methods largely in isolation. This section draws these threads together into a single picture of how the dimples behave across the test envelope, without repeating the case-by-case evidence already presented, before Chapter 9 turns to what this means for the design of an adaptive dimpled surface.

7.7.1. The Reynolds Number Trend

The screening in Section 7.2 established the dominant trend that everything else in this chapter sits inside. At $Re = 1 \times 10^5$ the dimples are broadly beneficial across the low-incidence regime, from 0° to 8° , with gains large enough in places to more than double the baseline glide ratio. At $Re = 2 \times 10^5$ this benefit collapses to a narrow band around 2° to 4° and shrinks in magnitude, driven almost entirely by drag reduction rather than any lift gain. At $Re = 5 \times 10^5$ the attached regime turns uniformly detrimental for the mid-chord and full coverage families, and any benefit that remains is confined to the immediate vicinity of stall. The dimples therefore do not have a single effect that simply weakens with Reynolds number, they have a genuinely different effect at each Reynolds number tested, beneficial and widespread at the lowest, marginal and narrowly confined at the middle, and actively harmful across most of the polar at the highest.

7.7.2. Chordwise Placement

Overlaying the coverage families onto this trend, established in Section 7.3, shows that no single family dominates across the full Reynolds number range. The Mid layout produces the strongest gains at $Re = 1 \times 10^5$ and still shows some benefit at $Re = 2 \times 10^5$, but becomes the worst performing family at $Re = 5 \times 10^5$, with no depth variant showing any improvement at any angle. The Aft layout shows close to the opposite pattern, negligible or poor at the two lower Reynolds numbers but the only family to show a genuine contribution at $Re = 5 \times 10^5$ at near stall angles. The Full layout does not simply average the two. At $Re = 1 \times 10^5$ it tracks Mid closely, but without recovering Mid's lift gain, and at $Re = 5 \times 10^5$ it inherits some of Aft's recovery near stall without matching Aft's magnitude. Combining mid-chord and aft dimples therefore does not add their individual benefits, and in some conditions removes a benefit that either family shows on its own.

7.7.3. Dimple Depth

The depth investigation in Section 7.4, focused on the Mid dimple layout, adds a further layer to this picture. In the low-incidence regime at $Re = 1 \times 10^5$, deeper dimples give progressively greater drag reduction, in agreement with the mechanism established for the golf ball in Section 2.4, where a deeper dimple trips the boundary layer more strongly and delays separation further downstream. This same trend reverses sign entirely at $Re = 5 \times 10^5$, where deeper dimples instead produce a larger drag penalty in the same angle of attack range, and the same golf-ball framing still applies, at this Reynolds number, the additional mixing introduced by deeper dimples degrades an already favourable boundary layer rather than delaying a separation that would otherwise occur. Near stall the relationship between depth and performance is not monotonic at all. At $Re = 1 \times 10^5$ and 14° , only the mid-depth variant, Mid-2.50, returns a genuine drag reduction, while both the shallowest and deepest variants underperform it. This non-monotonic behaviour, together with the sign reversal between Reynolds numbers, is the central motivation for treating dimple depth as an adjustable rather than fixed parameter, since no single depth is optimal across the full range of conditions tested.

7.7.4. The Flow Behaviour and Flow Fields

The flow-field investigation in Section 7.5 identified the physical mechanism underlying these trends for the standout case, Mid-2.50 at $Re = 1 \times 10^5$. The PIV field shows the dimples visibly energising the near-wall flow and suppressing separation that is otherwise present on the smooth baseline at the same angle. The surface pressure taps corroborate this with a measurable increase in suction magnitude over the forward chord and a fuller pressure recovery over the aft chord, and the wake rake measurement shows the resulting wake narrowed to roughly half the baseline width. These three independent measurements agree with one another and with the golf-ball mechanism from the theory chapter, delayed separation through boundary-layer tripping, narrowing the wake and reducing drag. The same configuration was also shown to delay stall at $Re = 2 \times 10^5$, corroborated independently by surface pressure, force balance, and tuft measurements, extending this mechanism from a drag-

reduction effect in the low-incidence regime to a genuine stall-delay effect closer to the edge of the operating envelope.

7.7.5. Numerical Method Limitations

The cross-method comparison in Section 7.6 showed that neither RANS nor LES fully reproduces the boundary layer mixing mechanism, despite both broadly reproducing the drag trend seen experimentally. LES showed almost no visible change in the near-wall boundary layer between Smooth and Mid-2.50, yet still predicted the correct trend of the change in drag, while introducing a lift penalty not seen experimentally. RANS showed no meaningful difference in the suction bubble between the two geometries at all, and disagreed with experiments on the sign of the drag change at 6° and on the presence of any benefit for Mid at 14° and 16° at $Re = 5 \times 10^5$. This is consistent with the absence of a transition model in the RANS simulations, a limitation already flagged in the theory chapter, and means that the RANS screening used throughout this chapter should be read as a check on the direction and rough magnitude of the experimental trend, not as an independent quantitative confirmation of it.

7.7.6. Connection to the Regime-Dependent Behaviour Predicted in Theory

Taken together, these results reproduce the regime-dependent behaviour anticipated in Section 2.5.2. Dimples are expected to help where the boundary layer is close to separating and to hurt where the flow is otherwise comfortably attached, since in the latter case there is no separation to delay and the dimples only add unnecessary mixing and drag. This is exactly the pattern observed here. At $Re = 1 \times 10^5$ the smooth baseline separates readily across the low-incidence regime, and forward dimples delay that separation to good effect. At $Re = 5 \times 10^5$ the baseline boundary layer is far more resistant to separation in the same angle of attack range, so the same forward dimples have nothing useful to delay and instead degrade the flow, while aft dimples, operating closer to the point where the higher-Reynolds-number boundary layer does eventually struggle, recover some of the benefit instead. The dependence on both dimple depth and dimple location observed across this chapter is therefore not accidental but a direct consequence of where, and how strongly, the local boundary layer sits relative to separation at each operating condition.

Airfoil with Adaptive Dimple Depth

The static campaign in the preceding chapters identified Mid-2.50 as a standout configuration and established that no single fixed dimple depth performs well across the full range of Reynolds numbers and angles of attack tested. This motivates a prototype capable of varying dimple depth in service, rather than committing to one depth at the design stage. This chapter presents a proof-of-concept airfoil built around this idea. It covers the underlying concept, the mechanical and electrical setup used to actuate the dimples, and the results obtained from the prototype

8.1. Concept and Motivation

The dimples on the static models investigated so far are fixed geometric features, printed directly into the airfoil surface at a single depth. The results in Section 7.4 show that this is a limitation rather than a convenience, since the optimal dimple depth changes with Reynolds number and angle of attack. Further, near stall, it does not vary monotonically at all. A depth that gives the largest drag reduction at $Re = 1 \times 10^5$ is actively harmful at $Re = 5 \times 10^5$, and no single static depth is optimal across the full envelope.

An adaptive dimple concept addresses this by replacing the fixed recess with an elastic membrane whose depth can be actuated on demand. Rather than committing to one geometry at manufacture, the depth can in principle be set to whatever value the static campaign has shown to be beneficial for the current operating condition. The present prototype demonstrates the mechanical feasibility of this approach. It does not implement real-time, closed-loop control based on flow measurements, that remains future work, but instead shows that a physical mechanism can reliably actuate dimple depth to a set of discrete, repeatable settings, so that its performance at each setting can be compared directly against the corresponding static configuration from the earlier chapters.

8.2. Prototype Setup

Figure 8.1 shows the overall system. A latex membrane forms the deformable dimple surface at the required chordwise location on the airfoil. Each dimple is connected to a shared cavity inside the airfoil, which in turn is connected via a tube to a vacuum pump, with two solenoid valves and a pressure transducer in the line between the two. Drawing a vacuum inside the cavity pulls the membrane of each dimple inward, forming the dimple recess. Once the desired depth is reached, the sealing valve closes and holds the vacuum, keeping the membrane at that depth without requiring the pump to run continuously. The whole system is controlled from an Arduino Uno.

Figure 8.2 shows the airfoil mounted in the tunnel for this campaign, with the yellow glued latex membranes visible on the suction surface, which act as the dimples, and the pneumatic supply line running from the base of the model out of the test section to the actuation system described in Section 8.2.1.

The dimple surface itself is a thin latex sheet, chosen for its ability to deform elastically and repeatedly under vacuum without tearing or taking a permanent set. Several sheet thicknesses were tested to find a suitable balance between the ease of actuation and the mechanical durability of the membrane under repeated cycling. A sheet of thickness 0.30 mm was selected for the prototype presented here.




Figure 8.1: Overview of the adaptive dimple prototype system. The latex membrane forming the dimple surface is connected via a tube to a pressure transducer, a pair of solenoid valves, and a vacuum pump, all controlled by an Arduino.

Each membrane patch is bonded to the airfoil surface by hand, and as a result the unactuated, zero-pressure surface is not perfectly flush with the surrounding printed geometry. Individual patches sit proud or recessed by up to approximately 0.5 mm in places, depending on the assembly. The prototype's rest state is therefore not as smooth as the printed Smooth baseline used elsewhere in this thesis, a distinction that matters when interpreting the results in Section 8.4.




Figure 8.2: The adaptive dimple prototype mounted in the DTU Red Wind Tunnel. The yellow latex membranes forming the actuated dimples are visible on the suction surface. The pneumatic fitting and supply tube, running from the base of the model to the vacuum pump outside the test section, are visible in the lower left of the frame.

8.2.1. Actuation and Control System

Dimple depth is set by controlling the pressure inside the cavity of the airfoil. Once the membrane has deformed to the desired depth, an Arduino-controlled solenoid valve closes and seals the vacuum at the dimple cavity, holding that depth without further pump operation. A second solenoid valve vents the cavity back to ambient pressure to return the membrane to its undeformed, flush state. The pressure transducer provides a continuous reading of the cavity pressure to the Arduino, giving a repeatable reference point for a given actuation rather than relying on activation time alone. Figure 8.3 shows a single dimple in both its unactuated and actuated state.




Figure 8.3: Close-up of the prototype's dimples, showing the yellow latex membrane in its unactuated (a) and actuated (b) state. The actuated state reaches a dimple depth of $k = 1.5$ mm.

The vacuum pump used is a Kamoer KLVP6, the same model used by Vilumbrales-Garcia et al. [39], rated to a maximum vacuum of 85 kPa. Figure 8.4 shows the full electrical wiring of the system, including the relay switching the pump, the two solenoid valves, and the voltage divider used to bring the transducer's 0 to 10 V output down to the Arduino's 0 to 5 V analogue input range. The divider follows directly from Kirchhoff's voltage law: two equal resistors, R_1 and R_2 , of 10 k Ω in series across the transducer output divide the signal in the ratio $R_2/(R_1 + R_2) = 0.5$, halving it exactly and mapping the full transducer range onto the full input range of the Arduino.





Figure 8.4: Electrical wiring diagram of the prototype. The Arduino switches the vacuum pump and the two solenoid valves via a relay, and reads the cavity pressure from the transducer through a resistive voltage divider.

8.2.2. Depth Calibration


Dimple depth is set and held by the cavity pressure and read continuously from the pressure transducer, rather than by pump activation time directly. A given pressure setpoint therefore corresponds to a repeatable membrane deformation regardless of how long the pump took to reach it.

To relate pressure setpoint to physical dimple depth, the actuated membrane was measured with an ATOS III structured-light 3D scanner from GOM [36], shown together with the prototype in Figure 8.5a. The scanner projects a blue-light fringe pattern onto the suction surface and reconstructs it as a point cloud, from which a triangulated surface of the dimpled region is generated. Dimple depth was extracted from this reconstructed surface as the maximum normal deflection of the membrane relative to the surrounding airfoil contour, and the procedure was repeated at each of the pressure setpoints used in the aerodynamic campaign.

Figure 8.5b shows the resulting relationship between cavity pressure and dimple depth. The error bars represent the measurement uncertainty of the reconstructed point cloud itself. Because the latex membrane is elastic and remains under load throughout a scan, any residual motion during acquisition degrades the agreement between individual fringe projections, so the scatter in the reconstructed surface grows with the degree of unsteadiness in the membrane during deflection. The measurement uncertainty of the instrument itself is 0.005 mm [36], far smaller than the scatter introduced by the membrane itself.



(a) Prototype airfoil set up for scanning



(b) Measured depth against cavity pressure

Figure 8.5: Depth calibration of the adaptive dimple. (a) The ATOS III structured-light scanner and the prototype, with the yellow membranes and pneumatic supply tube visible. (b) Measured dimple depth as a function of cavity pressure setpoint inside the prototype. Error bars give the point-cloud reconstruction uncertainty, which grows with the membrane's unsteadiness due to pressure variations.

8.3. Test Matrix

The aerodynamic campaign was conducted at $Re = 1 \times 10^5$ in the DTU Red Wind Tunnel, following the same wake-rake and wall-pressure procedure described in Section 3.3. This Reynolds number was chosen because it is where the Mid layout showed the strongest benefit in the static campaign, making it the most informative single condition for a first aerodynamic test of the prototype. Three actuated pump pressure setpoints were tested, equivalent to dimple depth of 0.75 mm, 1.25 mm, 1.50 mm denoted Mid-0.75 (Adaptive), Mid-1.25 (Adaptive) and Mid-1.50 (Adaptive). Each is compared against the same prototype airfoil in its own unactuated 0.00 mm state, denoted Mid-0.00 (Adaptive). This unactuated state, rather than the separately manufactured Smooth model, serves as the baseline for all comparisons in Section 8.4, so that the surface-finish imperfection noted in Section 8.2 is present in both the baseline and the actuated cases alike, isolating the effect of actuation itself. Measurements were taken across an angle-of-attack range of -10° to 16° as shown in Table 8.1.

Table 8.1: Operating conditions used for the airfoil with adaptive dimple depth.

Parameter	Values
Angle of attack α	$[-10^\circ, 16^\circ]$ in steps of 2°
Reynolds number Re	1.0×10^5
Inlet velocity U_∞	7.5 m/s
Flow regime	Incompressible
<i>Study scope</i>	
Depth settings	4 (Mid-0.00, Mid-0.75, Mid-1.25, Mid-1.50)
Total measurement runs	56 (4 depth settings $\times$ 14 angles)

8.4. Results

This section first establishes what actuation does to the prototype on its own terms, comparing each actuated depth against the same airfoil in its unactuated state, and then compares those actuated states against the corresponding static configurations from Section 7.4 at matched depth.

Figure 8.6 shows the resulting lift and drag polars across the test matrix described in Section 8.3. The lift curves for all four depth settings collapse more or less onto one another across the entire polar, indicating that, unlike the static configuration Mid-1.25 (Figure B.1), actuation of the present prototype does not produce a measurable lift benefit at this Reynolds number. The drag curves are where the effect of actuation appears, most clearly in the two regions already familiar from the static campaign. A modest reduction below the baseline through the low-incidence region can be seen, and a growing penalty as the polar approaches stall.




Figure 8.6: Lift coefficient C_L (left) and drag coefficient C_D (right) against angle of attack at $Re = 1 \times 10^5$, comparing the adaptive prototype in its unactuated state (Mid-0.00) to the three actuated depths Mid-0.75, Mid-1.25, and Mid-1.50.

Figure 8.7 isolates the drag change directly. From 0° to 8° all three actuated depths reduce drag relative to the unactuated baseline, and the reduction peaks sharply at 6° , reaching approximately 22% for Mid-0.75, 20% for Mid-1.25, and 15% for Mid-1.50. Within this window, the ordering is consistent, with the shallowest actuation giving the largest drag reduction throughout and the deepest the smallest. This is the reverse of the equivalent static result: in the dimple depth study at $Re = 1 \times 10^5$ for the same Mid-configuration (Figure 7.10), where the deepest variant, Mid-5.00, produced the greatest drag reduction over the same range of angles.

Between 8° and 10° the sign reverses, and all three settings move into a drag penalty that peaks around 12° . The ordering here is less clean. Mid-0.75 carries the smallest penalty at every angle beyond the crossover, reaching at most 7% at 12° , but the two deeper settings exchange places with angle: Mid-1.50 is the worst at 10° at roughly 10%, while Mid-1.25 is the worst at both 12° and 14° , peaking at approximately 15%. All three remain above the baseline at the largest angle shown, so the penalty has

not closed by the end of the measured range. The overall shape of the curve, a pronounced reduction through the low-incidence regime giving way to a penalty near stall, reproduces the static depth study closely again, but the ordering within each still differs.




Figure 8.7: Percentage change in C_D relative to the unactuated baseline (Mid-0.00) as a function of angle of attack at $Re = 1 \times 10^5$, for the three actuated depths.

This reversal in depth ordering, shallow performing best in the low-incidence regime rather than deep, is the opposite of the static Mid result at the same Reynolds number. The most likely explanation is the membrane's imperfect "rest-state" surface finish, noted in Section 8.2. Where the static models present a smooth, uniform printed surface outside the dimple recesses themselves, the prototype's glued membrane introduces a degree of surface irregularity that is not present on the static models. This explanation is plausible but not confirmed by the present data, and is offered here as the most likely candidate rather than an established finding.




Figure 8.8: Normalised streamwise wake velocity u/U_∞ across the wake rake position y/c at $Re = 1 \times 10^5$ and angle of attack 6° , comparing the unactuated Mid-0.00 and Mid-1.25 states of the adaptive prototype (solid lines) against the corresponding static Smooth and Mid-1.25 airfoils (dashed lines).

With the depth calibration of Section 8.2.2 in place, the actuated settings can be compared against the static configurations at matched dimple depth. Figure 8.8 makes this comparison at the wake, placing the unactuated and Mid-1.25 states of the prototype alongside the static Smooth and Mid-1.25 models at the same angle. Both pairs reproduce the behaviour already established throughout

the static campaign. The dimpled case narrows the wake relative to its own baseline. It does so on the suction side of the profile, consistent with the delayed-separation mechanism identified for the sphere in Section 2.4. The two adaptive wakes recover to the freestream more rapidly than their static counterparts, giving a narrower wake overall. This is consistent with the additional surface roughness carried by the assembled membrane surface outside the dimples themselves, which would promote further near-wall mixing on top of the dimple effect, although the present data cannot separate that contribution from other differences between the two models.

Figure 8.9 extends the same matched-depth comparison to the full polar, and the two coefficients behave differently. The drag curves show the adaptive model sitting above its static counterpart across most of the low-incidence regime, by a margin that is largest from -6° to -4° , and 10° to 12° , and narrows towards stall. This offset is in the direction expected from the additional surface roughness of the assembled membrane, and it is present in the unactuated state as well as the actuated one, consistent with it being a property of the model rather than of the actuation. Within each model, the dimpled case still reduces drag relative to its own baseline over the low to moderate angles of attack, so the trend introduced by the dimples survives the offset, though it appears shifted by a few degrees in angle of attack between the two models.

The lift curves differ more substantially. The adaptive model produces a noticeably steeper lift curve than the static model, lying below it at negative incidence and above it from roughly 2° onward, and it also reaches a higher maximum lift before dropping more sharply at 16° . The present campaign does not isolate a cause for this. The elastic membrane, the hand-assembled surface finish, and the separate manufacture of the two models are all plausible contributors, and it is reported here as an observed difference rather than an explained one.




Figure 8.9: Lift coefficient C_L (left) and drag coefficient C_D (right) against angle of attack at $Re = 1 \times 10^5$, comparing the unactuated Mid-0.00 and Mid-1.25 states of the adaptive prototype (solid lines) against the corresponding static Smooth and Mid-1.25 airfoils (dashed lines).

Across the four settings, the prototype therefore behaves as a single airfoil with a selectable drag characteristic. At rest it is aerodynamically close to an undimpled surface; actuated, it reproduces the low-incidence drag reduction and near-stall penalty established for the static dimpled models, with the magnitude of both depending on the depth selected. The transition between these states is made pneumatically, without any change of hardware, which is the mechanical capability the adaptive concept requires.

Discussion & Conclusion

This thesis set out to answer how the controlled variation of dimple depth and chordwise placement could be exploited to alter the aerodynamic characteristics and flow-separation behaviour of a lifting surface across varying operating conditions. To that end, a broad static campaign swept chordwise coverage, dimple depth, angle of attack, and Reynolds number across wind-tunnel, force-balance, PIV, and CFD methods, followed by a focused investigation of the flow mechanism behind the standout configuration, and the design of a proof-of-concept prototype capable of mechanically varying its own dimple depth. This chapter draws the answers to the four sub-questions together, states the contributions and limitations of the present work, and closes with recommendations for future research.

9.1. Answering the Research Questions

SQ1: Coverage, Depth, and Reynolds Number

The screening, coverage, and depth investigations in Chapter 7 together form the largest body of evidence in this thesis. SQ1 asks how chordwise location, dimple depth, and Reynolds number jointly shape aerodynamic performance, addressed below in the order the results themselves were built: first the Reynolds-number trend that holds regardless of geometry, then how location and depth each explain part of that trend mechanistically, before all three are drawn together.

Conclusion on the Impact of Reynolds Number

The screening in Section 7.2, ahead of coverage or depth, establishes a single Reynolds-number trend that the following two subsections build on. At $Re = 1 \times 10^5$ the full geometry matrix was overwhelmingly beneficial across the low-incidence regime, the 0° to 8° range in which the baseline at this Reynolds number is already separating rather than attached, with the single best result in the entire campaign, a glide-ratio improvement of 132%, occurring near 4° (Figure 7.2). At $Re = 2 \times 10^5$ the positive region contracted to a narrow band around 2° to 4° with a much-reduced peak of 22%, and the remainder of the attached regime turned red, with losses reaching -52% near 12° (Figure 7.3). At $Re = 5 \times 10^5$ the attached regime was red almost everywhere from 0° to 10° , with losses exceeding -60% in places, and any recovery was confined to the immediate vicinity of stall, and even there inconsistent across the matrix (Figure 7.5).

This is not three unrelated datasets but one continuous trend: the benefit of dimples is concentrated at low Reynolds number and contracts, then inverts, as Reynolds number rises. The drag decomposition makes this especially clear. At $Re = 1 \times 10^5$ the C_D heatmap shows widespread drag reduction across the same 0° to 8° low-incidence window (Figure B.4). At $Re = 2 \times 10^5$ this same heatmap turns mostly red, with only a handful of green cells remaining, now confined to the narrower 2° to 4° band (Figure B.5). At $Re = 5 \times 10^5$ almost no drag improvement remains anywhere in the matrix, and the penalties themselves grow far larger, with several configurations exceeding a 100% drag increase (Figure B.6). The same Reynolds-number reversal already seen in the glide ratio is therefore driven, at every stage, primarily by drag.

The RANS screening (Figure 7.4, Figure 7.6) reproduced this same qualitative shape at both higher Reynolds numbers, offering independent numerical support for the broad trend, though RANS was

consistently more optimistic in some regions and more pessimistic in others. The experimental trend itself, however, is corroborated by the SLT campaign wherever the two overlap (Figure D.1, Figure D.2) and is the single most consistent finding of the screening exercise.

This confirms, across the full test matrix rather than one isolated condition, the regime-dependent behaviour anticipated in Section 2.5.2: dimples help where the boundary layer is already close to separating and hurt where it is not, and which regime a given operating point falls into changes with Reynolds number.

Conclusion on the Impact of Chordwise Location of Dimples

The coverage investigation in Section 7.3 showed that chordwise placement is not a minor tuning parameter but the dominant factor separating success from failure at a given Reynolds number. At $Re = 1 \times 10^5$, the Mid and Full layouts both performed well across the 0° to 8° low-incidence window, while Aft performed poorly at nearly every angle in that same window (Figure 7.2). The mechanism behind Mid's gain was shown to be drag-dominated. Mid produced a strong drag reduction across this range, with a smaller accompanying lift gain, except at 0° where the lift gain briefly exceeded the drag reduction (Figure B.1, Figure B.4). Full pushed the drag reduction even further, but its lift gain effectively disappeared, showing that adding the rearward dimples that turn Mid into Full removes a benefit rather than simply diluting it. Aft told a different story, its lift was a genuine gain across most of the polar, but this was outweighed by a substantial drag increase spanning almost the entire 0° to 16° range, leaving it net negative overall despite the lift benefit. Summed across the full envelope, only Mid and Full exceeded the smooth baseline overall at this Reynolds number, Aft never did (Figure 7.7).

This pattern follows directly from where separation occurs at this Reynolds number. The boundary layer separates comparatively early on the chord, so only dimples placed upstream of that point can energise it in time, exactly what Mid and Full do and Aft does not. Placed aft of an already-separated flow, Aft's dimples have nothing left to trip, they merely disturb a region that has already broken down, consistent with the drag penalty dominating despite the lift gain. The PIV flow field confirms this directly, Figure D.7c shows the baseline separating in the same chordwise location where the Mid dimples sit, while Figure D.9c shows the dimpled case remaining fully attached along the same stretch, removing the wake altogether rather than merely shrinking it.

At $Re = 2 \times 10^5$ this picture contracted sharply. The wide benefits seen for Mid and Full narrowed to a small band around 2° to 4° , and for the first time no configuration, at any depth, returned a net positive result over the full envelope (Figure 7.3). The best any of the three managed was to come close to break-even between losses and gains across the polar (Figure 7.8).

At $Re = 5 \times 10^5$ the coverage story reversed outright. Mid lost every trace of benefit, with no angle or depth returning an improvement at all, making it the worst-performing family precisely where it had been the strongest at the lowest Reynolds number. Full recovered a modest benefit only very close to stall. Aft, by contrast, became the only family to show a genuine positive contribution at this Reynolds number, concentrated in the same 12° to 16° region where Mid and Full were failing, though even this gain remained outweighed by Aft's own losses elsewhere in the polar (Figure 7.5, Figure 7.9). This is the mirror image of the low-Reynolds-number case. Separation is now delayed much further aft, so the Mid dimples no longer act near a separation point, they instead trip a boundary layer already comfortably attached, adding drag rather than delaying anything. Aft's placement, by contrast, now sits closer to where separation actually begins, and picks up something close to the role Mid played at the lowest Reynolds number. Full sits awkwardly between the two, never the best performer at either extreme but never entirely without effect either.

The conclusion on chordwise location is therefore that forward and aft placement are not simply better or worse versions of one another. They are suited to opposite ends of the Reynolds number range tested, and full coverage does not average their behaviour. Combining the two tends to inherit the weaknesses of both locations rather than the strengths of either. The underlying reason is that rising Reynolds number pushes the natural separation point rearward, where the boundary layer becomes fuller and more resistant to the adverse pressure gradient, and needs correspondingly less help to stay attached over the forward chord. What determines whether a dimple helps or hurts is therefore not simply "forward is better" or "aft is better", but whether the dimples sit close to wherever transition and separation actually occur at that Reynolds number. Placed too far upstream of the separation and they merely trip an already-healthy boundary layer, adding parasitic drag with no separation left to delay.

Placed too far downstream of the separation and the intended effects arrive too late to influence a shear layer that has already detached.

Conclusion on the Impact of Dimple Depth

The depth investigation in Section 7.4, focused on Mid, showed that depth does not simply scale the coverage effect described above, its influence reverses with Reynolds number in much the same way location does. At $Re = 1 \times 10^5$, within the low-incidence regime from 0° to 8° , drag reduction increased monotonically with dimple depth (Figure 7.10). The wake-rake measurement at 6° confirms this. The wake narrowed progressively with depth, entirely on the suction side, while the minimum wake velocity stayed essentially unchanged (Figure 7.11). The benefit therefore came from a narrower wake rather than a shallower one, essentially mimicking the dimple effect seen on golf balls, as shown in Figure 2.12. However, this clean ordering broke down near stall. At 14° , only the mid-depth variant, Mid-2.50, returned an actual drag reduction. Mid-1.25 tracked the baseline wake with no benefit, and Mid-5.00 produced an even deeper and wider wake than the untreated airfoil (Figure 7.12). Depth has a practical optimum near stall that neither "shallower is safer" nor "deeper is stronger" describes correctly. The SLT force-balance campaign corroborated this depth ordering independently across 0° to 12° , with disagreement confined to the same near-stall region (Figure D.1), giving cross-facility support to the trend rather than resting on a single dataset and test facility.

At $Re = 2 \times 10^5$ the effect of dimple depth became angle-specific rather than following one rule across the polar, and its direction was already reversing. At 6° , deeper dimples now produced a wider, deeper wake than the baseline, the opposite of the $Re = 1 \times 10^5$ behaviour at the same angle (Figure 7.13). At 14° the ordering reversed again relative to 6° , with the shallowest depth now carrying the largest penalty (Figure 7.14). Summed across the full envelope, the net loss for Mid grew monotonically with depth (Figure 7.8).

At $Re = 5 \times 10^5$ this reversal was complete. The wake at 6° widened and deepened progressively with depth, the mirror image of the narrowing seen at $Re = 1 \times 10^5$ (Figure 7.15). The same mechanism that suppressed separation at low Reynolds number was, at this higher Reynolds number, adding unwanted mixing to a boundary layer that did not need it. At 14° the depth ordering was genuinely monotonic again, with drag penalties of 37%, 43%, and 48% for 1.25 mm, 2.50 mm, and 5.00 mm respectively. The wake profiles (Figure 7.16) revealed a detail the drag numbers alone could not: the wake's position, not just its size, shifted measurably with depth, a distinction invisible to the C_L/C_D heatmap or summary bar chart alone.

The conclusion on dimple depth is that its effect is genuinely non-monotonic and, like chordwise location, reverses sign with Reynolds number. No single depth is safe to fix at the design stage, the depth that gives the strongest benefit at $Re = 1 \times 10^5$ is actively the most harmful choice at $Re = 5 \times 10^5$, and even within a single Reynolds number the optimal depth can differ from one angle of attack to the next. This is the most direct evidence in the thesis for the motivation of treating dimple depth as an actuated variable rather than a fixed choice.

Taken together, these findings answer SQ1: location, depth, and Reynolds number interact strongly rather than independently, and no single fixed combination performs well across the full envelope tested. This is precisely why allowing dimple depth, and in principle chordwise placement, to be varied in service, rather than fixed once at manufacture, motivates the proof-of-concept prototype introduced in Chapter 8.

SQ2: The Flow Mechanism

The flow-field investigation of the standout configuration, Mid-2.50 at $Re = 1 \times 10^5$ (Section 7.5), identified the physical mechanism underlying the gains established above, using PIV, surface pressure, and wake rake measurements together rather than any one method alone. At 0° , where Mid-2.50 already returns a 15% glide-ratio gain, the PIV field shows a markedly thinner, higher-momentum layer at the wall compared to the baseline (Figure 7.17). At 4° , where the gain rises to 87%, the contrast becomes dramatic rather than incremental. The baseline shows a clearly separated shear layer thickening from roughly a quarter-chord onward, while Mid-2.50 shows high-momentum flow held attached along the entire suction side, with no equivalent low-momentum region forming at all (Figure 7.18). The dimples are visibly energising the near-wall flow and suppressing a separation that is otherwise present on the smooth baseline at the same angle.

The surface pressure taps corroborate this at the same condition. The forward three-quarters of the suction-side chord shows a consistent increase in suction magnitude for Mid-2.50 relative to Smooth, consistent with the higher near-wall momentum seen in PIV, while the final quarter of the chord shows the reverse, a reduction in suction magnitude consistent with a fuller pressure recovery on an airfoil that has stayed attached rather than separated (Figure 7.19, Figure 7.20). The wake rake measurement at the same angle shows the resulting wake narrowed to roughly half the baseline width, entirely on the suction side (Figure 7.21), consistent with the 42% drag reduction and 10% lift gain recorded at this condition. All three independent measurements, PIV, surface pressure, and wake rake, point to the same process already established for the golf ball in Section 2.4. The dimples trip the boundary layer, the resulting turbulent layer resists the adverse pressure gradient for longer, and separation is delayed, and the wake is narrowed as a direct result.

The same configuration was also shown to delay stall at $Re = 2 \times 10^5$, corroborated independently by three further measurements. The surface pressure taps at 16° show a stronger, more prolonged leading-edge suction peak for Mid-2.50 and markedly less scatter across the suction-side taps than the baseline, the signature of a steady, attached boundary layer rather than an intermittently separating one (Figure 7.22). The SLT force balance shows the same effect in the integrated forces: Smooth stalls sharply at 16° , while Mid-2.50's lift continues to rise smoothly to roughly 30% above the baseline at the same angle, with drag close to half (Figure 7.23). The tuft visualisation confirms this directly by eye, Smooth shows fully separated, incoherent surface flow across the suction side at 16° , while Mid-2.50 shows attached, streamwise-aligned flow across the mid-span (Figure 7.24). Together these extend the mechanism from a drag-reduction effect in the low-incidence regime to a genuine stall-delay effect nearer the edge of the operating envelope. This directly answers SQ2: the near-wall flow topology changes identified by PIV, surface pressure, force balance, and tuft measurements together explain the observed aerodynamic gains, corroborated independently across two facilities and four measurement techniques.

SQ3: Numerical and Experimental Agreement

The cross-method comparison in Section 7.6 examined a single reference condition, Smooth against Mid-2.50 at $Re = 1 \times 10^5$ and 6° , resolved simultaneously by PIV, surface pressure taps, and wake rake alongside the corresponding RANS and LES solutions. This condition was selected because it is the only point in the campaign where a resolved experimental flow field, a surface pressure distribution, and both numerical methods are available together, allowing the integrated forces and the flow field responsible for them to be assessed against one another rather than in isolation.

The two numerical methods diverge from experiment in different ways. LES reproduced the direction of the drag change correctly, showing a reduction in C_D from Smooth to Mid-2.50 in agreement with every experimental method (Table 7.2), and its wake was marginally narrower for the dimpled geometry, in the same direction as the PIV result. The near-wall flow, however, showed almost no difference between the two geometries (Figure 7.26), and LES additionally predicted a lift penalty for Mid-2.50 that the experimental campaign does not show, where the lift difference at this angle is negligible. The correct integrated trend is therefore recovered without the flow-field change that produces it experimentally. RANS presented the opposite balance. Its field resembled the PIV measurement more closely than LES did in the size and location of the high-velocity region, yet it showed no meaningful difference in the suction bubble between the two geometries (Figure 7.27), and it disagreed with experiment on the sign of the drag change at this condition, as well as on the presence of any benefit for Mid at 14° and 16° at $Re = 5 \times 10^5$. The visual resemblance is therefore not accompanied by agreement in the quantity that matters.

What neither method reproduced is the mechanism itself. The thinner, higher-momentum near-wall layer that PIV resolves for Mid-2.50 (Figure 7.25), and which Section 7.5 establishes as the origin of the aerodynamic gains, is absent from both numerical fields. This is consistent with the fully turbulent formulation adopted in the RANS simulations. A boundary layer that is turbulent everywhere by its very construction cannot be tripped, so the tripping mechanism the dimples rely on has no representation in the model. Any drag change the simulation returns must arise from some other cause. The agreement in trend is therefore not evidence that the same physics is being captured.

Away from this single condition, the picture is more favourable, and the distinction matters for how the numerical results are read elsewhere in this thesis. The RANS screening reproduced the qualitative Reynolds-number trends established experimentally at both higher Reynolds numbers (Figure 7.4, Fig-

ure 7.6), while running consistently more optimistic in some regions of the matrix and more pessimistic in others, and the surface pressure comparison at $Re = 2 \times 10^5$ and 5×10^5 shows the numerical distributions following the experimental ones in form while differing in magnitude. Both numerical methods also passed the baseline validation of Chapter 6 against literature and against the experimental campaign.

Taken together, these findings answer SQ3. Both numerical methods reproduce the aerodynamic performance of the dimpled configurations at the level of direction and rough magnitude, and neither reproduces the near-wall flow mechanism responsible for it, with LES recovering the correct drag trend without the corresponding flow-field change and RANS matching the flow field visually while disagreeing on the sign of the force. The consequence for this thesis is that the RANS screening is used throughout Chapter 7 as a check on the direction and rough magnitude of the experimental trend rather than as independent quantitative confirmation of it, and the mechanism itself rests on the experimental evidence of SQ2 rather than on either simulation.

SQ4: Performance comparison of Adjustable and Fixed-Depth Dimples

A proof-of-concept prototype based on the Mid dimple layout was designed and built around a latex membrane whose depth was controlled pneumatically by a vacuum pump, solenoid valves, and a pressure transducer, all managed by an Arduino (Chapter 8). Scanning of the actuated membrane (Section 8.2.2) related each cavity pressure setpoint to a physical dimple depth, giving four repeatable states, an unactuated rest condition and three actuated depths of 0.75, 1.25, and 1.50 mm, and allowed the prototype to be compared against the static configurations of Section 7.4 at matched depth rather than by actuation setting alone.

Compared against itself, the prototype reproduced the drag behaviour of the static campaign in form. All three actuated depths reduced drag through the low-incidence regime, peaking at approximately 22% at 6° , before crossing into a penalty that grows towards stall (Figure 8.7). This is the same two-region shape established for the static Mid family at the same Reynolds number (Figure 7.10). The ordering within that shape, however, is inverted. The shallowest actuation gives the largest drag reduction and the deepest the smallest, where the static study found the opposite. The most likely explanation is the surface finish of the hand-assembled membrane, which introduces an irregularity outside the dimple recesses that the printed static models do not carry and that plausibly may worsen with actuation depth. This is offered as the leading candidate rather than an established finding.

At matched depth, the correspondence is closer in drag than in lift. The wake at 6° (Figure 8.8) shows both the static and the adaptive pair narrowing on the suction side when dimpled, the same signature identified for the sphere in Section 2.4, with the adaptive wakes recovering to the freestream more rapidly than their static counterparts. Across the full polar (Figure 8.9) the adaptive drag sits above the static drag through most of the low-incidence regime, in both the unactuated and the actuated state, which is consistent with the increased surface roughness of the model rather than of the actuation. Within each model, the dimples still reduce drag relative to that model's own baseline, so the trend survives the offset. The lift curves diverge more substantially, the adaptive model producing a steeper lift curve and a higher maximum lift than the static model, a difference the present campaign does not isolate a cause for. The prototype also shows no noticeable lift benefit from actuation itself, unlike the static Mid configurations (Figure B.1).

Taken together, these findings answer SQ4. The proof-of-concept prototype was manufactured based on the findings of RA1 and RA2. The prototype reproduces the drag behaviour of the equivalent static configurations in trend, both in the two-region shape of the drag change with angle of attack and in the suction-side wake narrowing that produces it, while differing in the ordering of that behaviour with depth, in absolute drag level, and in the absence of any lift benefit. A single physical airfoil is therefore shown to select between aerodynamic states that individually resemble distinct static geometries, which is the mechanical capability the adaptive concept requires, without yet matching the quantitative performance of a precisely manufactured static dimple. The demonstration is confined to one Reynolds number and one matched depth, and the correspondence established here is thus qualitative rather than exact. It should also be noted that the actuated range reached only 1.50 mm, some 30 % of the 5.00 mm design target, so the deeper configurations that gave the largest static gains at this Reynolds number remain untested in adaptive form.

9.2. Contributions

This work makes four contributions relative to the gaps and sub-questions identified in Section 1.1. First, it provides a single, internally consistent dataset spanning all three chordwise coverage families, each built from multiple dimple rows rather than a single row, and all three dimple depths, across the same three Reynolds numbers and the same angle-of-attack polar. The row placements themselves were not chosen arbitrarily but informed by the chordwise locations reported across the existing literature, as laid out in the design parameters of Chapter 2, so that the matrix tests the range of placements prior studies have individually proposed within one consistent campaign rather than replicating any single one of them in isolation. This directly addresses the fragmentary and mutually inconsistent picture that individual single-geometry, single-row studies have produced in the existing literature.

Second, it identifies the physical mechanism connecting dimple geometry to aerodynamic outcome directly, through PIV, surface pressure, and wake rake measurements on the same configuration, rather than relying on integrated forces alone. The PIV comparison of Smooth and Mid-2.50 at $Re = 1 \times 10^5$ and 4° in particular (Figure 7.18) is not a novel measurement technique or a first-of-its-kind result, but it offers an unusually clear, almost textbook-quality visualisation of dimple-induced separation suppression on an airfoil, the kind of direct, visual confirmation that force coefficients alone cannot provide.

Third, it documents what the RANS and LES formulations used here do and do not capture, when applied to dimpled airfoils at low Reynolds number. Both reproduced the direction and rough magnitude of the drag change, yet neither reproduced the near-wall mixing that the experimental campaign identifies as its cause, with LES returning the correct drag trend without any corresponding change in the boundary layer and RANS matching the measured flow field visually while disagreeing on the sign of the force. Because the same condition was resolved simultaneously by PIV, surface pressure taps, and wake rake, the disagreement can be located in the flow field itself rather than inferred from integrated coefficients alone. This is a negative result, but a useful one for anyone intending to screen dimple geometries numerically, since agreement in trend is shown here not to imply that the underlying mechanism is being modelled.

Fourth, it translates the resulting design guidance into a physical, mechanically adjustable proof-of-concept prototype, described in Chapter 8. This is not a groundbreaking device in itself, but it moves the adaptive-dimple concept, previously demonstrated only on bluff bodies such as the sphere, one concrete step closer to a working lifting-surface implementation.

9.3. Limitations

Several limitations qualify the conclusions drawn above. None of them is a reason to distrust the data presented in this thesis, every method used here is a standard established technique in experimental and computational aerodynamics, but each carries its own sources of uncertainty, and a fair reading of the results should keep these in view.

Scope of the campaign. The screening campaign covers a single airfoil profile, the NACA 63-415, at three discrete Reynolds numbers, and the results should not be assumed to transfer directly to other profiles or to the continuous range of Reynolds numbers an operational blade or wing would encounter.

Numerical methods. The RANS simulations use a fully turbulent formulation without a transition model, which was shown in Section 7.6 to introduce real disagreement with experiment in both magnitude and, in at least one case, sign. The RANS screening should accordingly be read as qualitative support rather than independent confirmation. In addition, both the RANS and LES meshes were generated with an automatic mesher rather than a manually defined, structured mesh. This was a deliberate choice made for turnaround speed across the large number of cases required, but automatic meshing is known to produce less uniform cell distributions, particularly in regions such as the dimple recesses and the near-wall boundary layer, than a carefully hand-built mesh would, and this is a plausible contributor to some of the numerical disagreement discussed above. A third point concerns the inflow condition. The RANS inlet turbulence was prescribed through fixed values of k_∞ and ω_∞ rather than a fixed turbulence intensity, which means the effective Tu differs between the two simulated Reynolds numbers and matches the measured background turbulence of neither facility. Finally, the mesh sensitivity study in Section 5.1.6 did not reach a flattened convergence sequence, so the production mesh carries a residual discretisation uncertainty in the glide ratio of the order of 5%. Both effects are sys-

tematic across the four geometries and therefore weigh more heavily on the absolute values than on the relative comparison between them.

Wind-tunnel flow conditions. The two wind-tunnel facilities used in this thesis determine freestream velocity by different measurement principles, and set to nominally the same Reynolds number, the actual flow condition in each tunnel still carries the measurement uncertainty of its own reference instrumentation, principally the pressure transducers used to infer velocity from a reference pressure difference. Beyond this facility-to-facility difference, the achieved Reynolds number can also drift modestly from run to run within a single facility, since the reference air density and viscosity used to set the tunnel speed depend on the ambient temperature and atmospheric pressure at the time of testing, both of which vary during longer sessions rather than being held at a single fixed laboratory condition throughout the entire campaign.

Model installation. Each airfoil model was mounted using a dedicated alignment tool intended to improve repeatability between installations, but the final angular alignment was still set by hand, reading the angle from a graduated scale rather than through an automated or feedback-controlled system. A residual misalignment of under one degree between installations, in either tunnel, is therefore plausible and cannot be fully excluded from the reported angle-of-attack values.

Model manufacture. All models were 3D printed using a 0.4 mm nozzle, which sets a practical lower bound on cross-sectional accuracy of roughly ± 0.4 mm, or about $0.002c$ for the present chord length and thickness. This is small in absolute terms but not negligible, and it also governs the printed surface roughness of the models before finishing. Every model was subsequently sanded by hand to improve surface finish, which is itself a manual process and a further, if likely minor, source of geometric variation between nominally identical models.

Wake rake and wall-pressure measurements. These measurements carry the instrument uncertainty of the pressure transducers used in both the wake rake and the wall-pressure taps. They are also subject to a specific geometric constraint. At angles of attack above 16° , the separated wake behind the dimpled configurations became wider than the wake rake itself, so that part of the wake falls outside the region the rake can actually resolve. This is the reason the bulk of the reported polar is confined to 0° to 16° , with results beyond this range treated with correspondingly less confidence.

Surface pressure tap model. In addition to the transducer uncertainty shared with the wake-rake measurement, the surface pressure tap model carries the uncertainties inherent to any discrete pressure-tap surface, finite tap size, the small manufacturing tolerance of each drilled tap hole, and the limited number and spacing of taps relative to the true continuous pressure distribution, all of which were already noted where relevant in Chapter 7 for their effect on the integrated lift and drag estimates.

Particle image velocimetry. The PIV campaign was affected by laser-light reflections and scratches on the test-section window, arising from both the smooth and dimpled airfoil surfaces, which locally corrupt the cross-correlation and produce dead or erroneous velocity regions in the processed fields, as discussed where these appear in Section 7.5 and Section 7.6. Beyond this, the measurement carries the usual sources of PIV uncertainty common to any such system, including small residual misalignment between the camera and the laser sheet, and the general particle tracking and correlation uncertainty inherent to the technique.

Force balance. The SLT force-balance measurements carry the instrument uncertainty of the balance itself, in addition to the three-dimensional effects already noted in Chapter 6. Because the balance measures the total load on a finite-span model rather than a sectional slice at mid-span, its results include wall-junction and tip-region effects, such as the wall-junction separation identified in the tuft visualisation of Section 7.5.4, that the two-dimensional wake-rake measurement does not capture. The aerodynamic coefficients reported throughout this thesis are also presented uncorrected for wind-tunnel blockage, a deliberate choice discussed in Chapter 2, which means the force-balance results in particular should be read as representative of the finite model tested rather than as blockage-free, fully two-dimensional sectional data.

Adaptive prototype. The prototype demonstration is narrower than the concept it is intended to test. The design target set out in Section 2.7.4 was a continuously variable depth reaching $k/c = 0.025$, or 5.00 mm at the present chord, whereas the membrane achieved a maximum of 1.50 mm, about 30 % of that target, and overlapped the static campaign at only a single depth, 1.25 mm. The comparison against the static models is therefore restricted to the shallow end of the design space, and the deeper configurations that produced the largest static drag reductions at this Reynolds number were never reproduced in actuated form. Depth was also set open-loop from a calibrated cavity pressure rather than by closed-loop control on a measured depth or aerodynamic signal, so the prototype demonstrates the mechanical capability the adaptive concept requires without yet demonstrating the control that would be needed to exploit it.

Large-eddy simulations. Beyond the mesh and inflow points noted above, the LES campaign used the standard Smagorinsky sub-grid-scale model, which is purely dissipative and does not switch itself off in laminar or transitional regions. This is a recognised characteristic of the model at these Reynolds numbers and is a plausible contributor to the weak dimple sensitivity the LES returned, since the mechanism being resolved is precisely a transition process. No comparison against a transition-sensitive closure was run, so this remains a candidate explanation rather than an established one. The campaign was also limited to two simulations at a single angle of attack and a spanwise width of $0.1c$, which contains exactly one spanwise period of the dimple array and therefore cannot represent any spanwise structure of larger scale.

Summary. Taken together, these limitations do not point to any single flaw undermining the results, but rather reflect the ordinary combination of instrument uncertainty, hand-built and hand-aligned physical models, and the specific measurement principle of each technique used. The cross-method agreement demonstrated throughout Chapter 7, wherever two or more independent methods examine the same condition, indicates that the underlying trends reported in this thesis are robust to these individual sources of uncertainty, even where the exact numerical values differ from one method to the next.

9.4. Recommendations for Future Work

Beyond the static campaign presented here, several extensions follow naturally from the results and would be well worth pursuing.

Individually Actuated, Row-Level Prototype

The present prototype actuates a single dimple depth across the whole array at once. A natural extension is independent actuation of individual rows along the chord, combined with real-time, pressure-based feedback control, closing the loop between the local pressure state and the actuated depth, moving the concept from a mechanism proven capable of discrete, fixed settings toward a genuinely adaptive surface. The main practical obstacle to a single such model is that its un-actuated state must closely resemble the smooth baseline, since any surface irregularity in a neutral state would undermine the very zero-depth condition the comparison depends on. If that can be achieved with an acceptable surface finish, a single reconfigurable model would remove the inter-model alignment and manufacturing uncertainties discussed in the previous section, at the cost of a harder manufacturing problem than the fixed-depth models used here. The alternative, extending the present approach with a larger matrix of static models covering nearly every plausible row combination and depth, remains a viable fallback, at the price of reintroducing exactly those inter-model uncertainties.

Isolating the Dimple-Transition Interaction

A complementary, more fundamental experiment would isolate the dimple-transition interaction directly, independent of any specific airfoil geometry. A narrow boundary-layer wind tunnel with a flat, adjustable test surface, incorporating a single actuatable dimple capable of being translated upstream and downstream, would allow the proximity between a dimple and the natural transition or separation point to be swept continuously, rather than sampled at the handful of discrete chordwise rows used in the present campaign. This would test directly the proximity-dependent mechanism identified in this chapter for the aerodynamic gains observed on the airfoil.

Extending the Flow-Field and Numerical Investigation

The flow-field investigation in this work focused on a single Reynolds number and a single standout geometry. Extending the PIV and surface pressure campaigns to the higher Reynolds numbers, where the dimple effect was shown to reverse sign, would help confirm whether the same golf-ball-type mechanism is responsible for the detrimental behaviour observed there, or whether a distinct process is at play. Repeating the screening with a transition-sensitive RANS formulation, or extending the LES campaign to additional angles of attack and Reynolds numbers, would similarly help establish whether the numerical disagreements identified in Section 7.6 reflect a fundamental limitation of steady, fully turbulent CFD for dimpled surfaces, or whether they can be resolved with a more appropriate turbulence and transition model, an open question this thesis was not intended to resolve.

Broader Applications

Finally, the low-Reynolds-number gains identified here point toward applications beyond wind turbine blades. Fixed-wing UAVs operate largely in the Reynolds number range where the benefit was strongest, and within wind turbines the effect may still be relevant on the lower-Reynolds-number inboard sections of a blade. The same mechanism could plausibly extend beyond lifting surfaces entirely, to road vehicles with speed-dependent surface morphing, echoing the crude but real fuel-economy gain demonstrated for a dimpled car body discussed in the introduction, to aircraft fuselages, or more speculatively to slender axisymmetric bodies, where independently actuated dimple zones around the circumference could in principle generate asymmetric drag and thus a steering potential, building on the adaptive-sphere concept of Vilumbras-Garcia et al. [39]. The hull of a container ship, operating at a near-constant draught and speed for most of a voyage, could similarly benefit from a surface with drag-reducing properties, such as dimples. Each of these represents a distinct research direction in its own right, well beyond the scope of the present work, but each follows naturally from the regime-dependent mechanism this thesis has established.

Data and Code Availability

The full body of code, data, and supplementary material generated over the course of this thesis is collected in a dedicated repository:

<http://github.com/Aaboe97/MSc-Thesis-2026/>

Much of this material could not be included in the printed thesis itself, either because of length or because it consists of raw data, source code, or media rather than text. The repository contains code and data of the following categories:

- **Setup calculators** — the scripts and tools used to derive experimental configuration parameters ahead of testing, including the PIV optical setup calculator introduced in Section 4.3.3 and equivalent tools used to size other parts of the campaign.
- **CAD design files** — the parametric models used to generate the airfoil geometries, dimple layouts, mounting hardware, and the prototype assembly.
- **Post-processing and plotting code** — the scripts used to reduce raw measurement and simulation output into the coefficients, heatmaps, and comparison figures presented throughout this thesis.
- **Videos and images** — supplementary photographic and video material from the wind-tunnel campaigns, including footage not reproduced in the report, such as extended tuft flow visualisation and PIV acquisition sequences.
- **Extra results** — additional figures, tables, and configurations beyond those selected for inclusion in the main text and appendices, covering conditions and geometries examined during the campaign but not central to the reported findings.
- **Raw experimental data** — the unprocessed wake-rake, wall-pressure, surface-pressure tap, force-balance, and PIV measurement data underlying the results in Chapter 7.
- **CFD cases and OpenFOAM code** — the full RANS and LES case setups, including mesh generation scripts, boundary condition definitions, and any custom OpenFOAM modifications used in Chapter 5.
- **Arduino code** — the firmware and control logic used to actuate the adaptive dimple prototype described in Chapter 8.

The repository is not publicly listed. Access is granted on request, and readers wishing to consult its contents may request access directly from the author.

Acknowledgements

In the interest of transparency, this section declares the ways in which generative LLM tools were used during the preparation of this thesis. In every case described below, the tool functioned as an assistant carrying out a task defined and directed by the author. No LLM tool was used to generate research ideas, make engineering or scientific decisions, interpret results, or draw conclusions. Hence, every judgement, interpretation, and conclusion presented in this thesis is the author's own.

Software development and simulation workflows. LLM assistance was used for debugging and writing Python and C++ code used in the post-processing, analysis, and control software developed for this project, for debugging $\text{\LaTeX}$ formatting and compilation issues throughout the preparation of the document, and for assistance with OpenFOAM, particularly in handling large numbers of simulation cases and meshes, such as swapping mesh variants between cases and setting up new case directories. In all instances, the resulting code was reviewed, tested, and verified by the author before use. The tools have thus served to accelerate implementation and troubleshooting of code and workflows that the author specified.

Automated, repetitive data handling. LLM agent tools were used to automate mechanical, repetitive tasks across large numbers of files and datasets, for example, executing a post-processing script across hundreds of data files and collecting the resulting output. This use was purely executional, the processing logic, the choice of what to compute, and the interpretation of the resulting output were defined by the author, with the agent tool performing the repetitive execution itself.

Illustrative figures. LLM image-generation tools were used to draw a number of the illustrative, non-data figures used throughout this thesis to explain underlying concepts and theory, such as boundary-layer behaviour, dimple geometry, and related theoretical schematics. These figures are conceptual illustrations rather than scientific data or measurement results.

Literature search and summarisation. LLM tools were used to help locate papers relevant to this project and to summarise longer reports and articles in order to identify which sections or publications warranted closer reading and possible inclusion. The selection of which sources to ultimately cite and rely upon, and the interpretation of their content, remained the author's own.


Language and structural editing. LLM assistance was used to improve the language, clarity, and structure of the written text throughout this thesis. In every case, the underlying content, the scientific reasoning, the data interpretation, and the conclusions were first written by the author. LLM tools were used only to refine how that existing, human-authored material was phrased and organised, not to originate it.

References

- [1] Elmar Achenbach. "Experiments on the Flow Past Spheres at Very High Reynolds Numbers". In: *Journal of Fluid Mechanics* 54.3 (1972), pp. 565–575. DOI: 10.1017/S0022112072000874.
- [2] *AirfoilTools: Airfoil Database and Analysis Tool*. <http://airfoiltools.com/>. Accessed: 2025-02-28.
- [3] Haris Ali et al. "Passive flow-field control using dimples for improved aerodynamic flow over a wing". In: *Scientific Reports* 14 (2024), p. 12918. DOI: 10.1038/s41598-024-63638-z.
- [4] John D. Anderson. *Fundamentals of Aerodynamics*. 6th. New York: McGraw-Hill Education, 2017. ISBN: 978-1-259-12991-9.
- [5] Christian Bak et al. *Wind Tunnel Tests of the NACA 63-415 and a Modified NACA 63-415 Airfoil*. Tech. rep. Risø-R-1193(EN). Roskilde, Denmark: Risø National Laboratory, Dec. 2000, p. 107. URL: <https://orbit.dtu.dk/en/publications/wind-tunnel-tests-of-the-naca-63-415-and-a-modified-naca-63-415-a/>.
- [6] Jewel B. Barlow, William H. Rae, and Alan Pope. *Low-Speed Wind Tunnel Testing*. 3rd ed. New York: John Wiley & Sons, 1999. ISBN: 0-471-55774-9.
- [7] P W Bearman and J K Harvey. "Golf Ball Aerodynamics". In: *Aeronautical Quarterly* 27.2 (1976), pp. 112–122. DOI: 10.1017/S0001925900007617.
- [8] L. Binci et al. "Study of the Flow Field Past Dimpled Aerodynamic Surfaces: Numerical Simulation and Experimental Verification". In: *Journal of Physics: Conference Series* 923 (2017), p. 012030. DOI: 10.1088/1742-6596/923/1/012030.
- [9] Jin Choi, Woo-Pyung Jeon, and Haecheon Choi. "Mechanism of Drag Reduction by Dimples on a Sphere". In: *Physics of Fluids* 18.4 (2006), p. 041702. DOI: 10.1063/1.2191848.
- [10] Antonio Carlos Daud Filho, Hernán Cerón-Muñoz, and Fernando Catalano. "Experimental Study of the Influence of Vortex Generators on Airfoils for Wind Turbines". In: *VI Congreso Internacional de Ingeniería Mecánica y IV de Ingeniería Mecatrónica*. Bogotá, Colombia, May 2013.
- [11] Delft University of Technology. *W-Tunnel: Low-Speed Wind Tunnel Facility*. <https://www.tudelft.nl/lr/organisatie/afdelingen/flow-physics-and-technology/facilities/low-speed-wind-tunnels/w-tunnel>. Accessed: 2026-01-14. 2024.
- [12] M. Gaster. *The Structure and Behaviour of Laminar Separation Bubbles*. Reports and Memoranda 3595. London, UK: Aeronautical Research Council, 1967.
- [13] SimScale GmbH. *k-omega SST Turbulence Model*. Accessed: 2026-03-17. 2024. URL: <https://www.simscale.com/docs/simulation-setup/global-settings/k-omega-sst/>.
- [14] Mrinal Kaushik. "Boundary Layers". In: *Theoretical and Experimental Aerodynamics*. Singapore: Springer, 2019, pp. 251–284. DOI: 10.1007/978-981-13-1678-4_11.
- [15] V. Kaushik, M. Mahore, and S. Patil. "Analysis of Dimpled Wing of an Aircraft". In: *International Journal of Engineering Development and Research* 6.3 (2018), pp. —.
- [16] Max W. Knoop, Abdelrahman Hassanein, and Woutijn J. Baars. "Development and characterisation of a turbulent boundary layer facility at the Delft University of Technology". In: *Aerospace Science and Technology* 168 (2026), p. 110972. ISSN: 1270-9638. DOI: <https://doi.org/10.1016/j.ast.2025.110972>. URL: <https://www.sciencedirect.com/science/article/pii/S1270963825010351>.
- [17] Brian E. Launder and B. I. Sharma. "Application of the energy-dissipation model of turbulence to the calculation of flow near a spinning disc". In: *Letters in Heat and Mass Transfer* 1.2 (1974), pp. 131–137. DOI: 10.1016/0094-4548(74)90150-7.
- [18] J. Gordon Leishman. *Boundary Layer Flows*. Embry-Riddle Aeronautical University, 2023. DOI: 10.15394/eaglepub.2022.1066.n22.

- [19] E. Livya, G. Anitha, and P. Valli. "Aerodynamic Analysis of Dimple Effect on Aircraft Wing". In: *International Journal of Aerospace and Mechanical Engineering* 9.2 (2015). World Academy of Science, Engineering and Technology (WASET), pp. 350–353. DOI: 10.5281/zenodo.1099926.
- [20] Florian R. Menter. "Two-equation eddy-viscosity turbulence models for engineering applications". In: *AIAA Journal* 32.8 (1994), pp. 1598–1605. DOI: 10.2514/3.12149.
- [21] R. Mustak, Md. N. Uddin, and Md. Mashud. "Effect of Different Shaped Dimples on Airfoils". In: *Proceedings of the International Conference on Mechanical Engineering and Renewable Energy (ICMERE2015), Chittagong, Bangladesh*. 2015, pp. –.
- [22] MythBusters. *Dirty vs. Clean Car*. YouTube. Accessed: 2026-07-23. 2009. URL: <https://www.youtube.com/watch?v=haRtzAfNMdI>.
- [23] El hassen Omer, Hana Momen, and Rayan Yahia. "Aerodynamic Characteristics of Dimple Effect on Airfoil". In: 5 (Apr. 2023).
- [24] OpenCFD Ltd. *Smagorinsky LES Model*. <https://doc.openfoam.com/2306/tools/processing/models/turbulence/les/rtm/Smagorinsky/>. Accessed: 2026-06-15. 2023.
- [25] Özge Özçakmak, Andreas Fischer, and Helge Madsen. "Characterization of a Fast Response Probe for Inflow Turbulence Measurements on Wind Turbine Blades". In: *Journal of Physics: Conference Series* 2265 (May 2022), p. 032015. DOI: 10.1088/1742-6596/2265/3/032015.
- [26] Amit Saraf, Mahendra Singh, and T. S. Chouhan. "Effect of Dimple on Aerodynamic Behaviour of Airfoil". In: *International Journal of Engineering and Technology* 9 (June 2017), pp. 2268–2277. DOI: 10.21817/ijet/2017/v9i3/1709030335.
- [27] Hermann Schlichting and Klaus Gersten. *Boundary-Layer Theory*. 9th. Berlin, Heidelberg: Springer, 2017. ISBN: 978-3-662-52917-1. DOI: 10.1007/978-3-662-52919-5.
- [28] A. Sciacchitano, D. Ragni, and R. Zamponi. *AE4-180 Flow Measurement Techniques: Student Manual for the Laboratory Exercise*. Instructors: A. Grille Guerra, L. Hendriksen, B. Dsouza, S. Wang. Apr. 2025.
- [29] Michael S. Selig et al. *Summary of Low-Speed Airfoil Data: Volume 1*. Virginia Beach, VA: SoarTech Publications, 1995. ISBN: 978-0-9646747-1-4.
- [30] Diksha Singh, Radheshyam Gajghat, and Mrinal Manik. "Experimental Investigation To Examine The Effect Of Shape And Size Of Dimples At Suction Surface Of Aerofoil". In: *International Journal of Scientific Technology Research* 8 (Dec. 2019), pp. 521–534.
- [31] Joseph Smagorinsky. "General Circulation Experiments with the Primitive Equations: I. The Basic Experiment". In: *Monthly Weather Review* 91.3 (1963), pp. 99–164. DOI: 10.1175/1520-0493(1963)091<0099:GCEWTP>2.3.CO;2.
- [32] V. Soundharya et al. "Aerodynamic Analysis of Dimple Effect on Airfoil". In: *International Journal of Engineering and Advanced Technology (IJEAT)* 5.4 (2016), pp. 207–211. ISSN: 2249-8958.
- [33] D. Srivastav. "Flow Control over Airfoils using Different Shaped Dimples". In: *International Conference on Fluid Dynamics and Thermodynamics Technologies (FDTT), Singapore*. IPCSIT Vol.33, IACSIT Press. 2012, pp. 92–97.
- [34] Adam J. Stolt, Al Habib Ullah, and Jordi Estevadeordal. "Study of Leading-Edge Dimple Effects on Airfoil Flow Using Tomographic PIV and Temperature Sensitive Paint". In: *Fluids* 4.4 (2019), p. 184. DOI: 10.3390/fluids4040184.
- [35] J. E. Stolzman and D. D. Manoharan. "The Effectiveness of Dimples on a NACA Airfoil: A Numerical Investigation Conducted via an Independent Study". In: *ASEE Annual Conference Exposition*. Paper ID: 36575, ASEE PEER. 2021. URL: <https://peer.asee.org/the-effectiveness-of-dimples-on-a-naca-airfoil-a-numerical-investigation-conducted-via-an-independent-study.pdf> (visited on 11/14/2025).
- [36] Technical University of Denmark. *ATOS III 3D Scanner*. DTU Equipment Information System. DOI: 10.57735/6350. URL: <https://doi.org/10.57735/6350> (visited on 07/28/2026).
- [37] Denis Terwagne, Miha Brojan, and Pedro M. Reis. "Smart Morphable Surfaces for Aerodynamic Drag Control". In: *Advanced Materials* 26.32 (2014), pp. 5718–5724. DOI: 10.1002/adma.201401403.

- [38] Fernando Villalpando, Marcelo Reggio, and Adrian Ilinca. “Numerical Study of Flow Around Iced Wind Turbine Airfoil”. In: *Engineering Applications of Computational Fluid Mechanics* 6.1 (2012), pp. 39–45. DOI: 10.1080/19942060.2012.11015401.
- [39] R. Vilumbrales-Garcia, P. B. Sudarsana, and A. Sareen. “Adaptive drag reduction of a sphere using smart morphable skin”. In: *Flow* 5 (2025). Cambridge University Press, E17. DOI: 10.1017/flow.2025.7.
- [40] Frank M. White. *Fluid Mechanics*. 8th. McGraw-Hill Education, 2016.
- [41] David C. Wilcox. *Turbulence Modeling for CFD*. 2nd. La Cañada, California: DCW Industries, 1998. ISBN: 0963605100.



Supportive Figures

Table A.1: Blockage ratios β for different chord lengths c and thickness ratios t/c with angles of attack $\alpha = 0^\circ, 5^\circ, 10^\circ, 15^\circ$, and 20° in the DTU Red Wind Tunnel ($0.50 \text{ m} \times 0.75 \text{ m}$) and the TU Delft SLT tunnel ($0.60 \text{ m} \times 0.90 \text{ m}$). The blockage ratio is defined as $\beta = A_{\text{proj}}/A_{\text{ts}}$ where A_{proj} is the projected frontal area of the airfoil and A_{ts} is the wind-tunnel cross-sectional area. In each facility, the model is taken to span the full tunnel width, giving $b = 500 \text{ mm}$ in the RWT and $b = 600 \text{ mm}$ in the SLT. Cell background colours indicate blockage severity: green (0–5%), yellow (5–8%), light orange (8–10%), orange (10–15%). Blockage corrections are considered reliable below the 5–10% range discussed in Section 2.7.1.

c [mm]	t/c	Tunnel	$\beta(0^\circ)$	$\beta(5^\circ)$	$\beta(10^\circ)$	$\beta(15^\circ)$	$\beta(20^\circ)$
160	0.12	SLT Tunnel	0.0213	0.0367	0.0519	0.0666	0.0809
		Red Wind Tunnel	0.0256	0.0441	0.0623	0.0799	0.0970
	0.15	SLT Tunnel	0.0267	0.0421	0.0571	0.0718	0.0859
		Red Wind Tunnel	0.0320	0.0505	0.0686	0.0861	0.1030
	0.18	SLT Tunnel	0.0320	0.0474	0.0624	0.0769	0.0909
		Red Wind Tunnel	0.0384	0.0568	0.0749	0.0923	0.1090
180	0.12	SLT Tunnel	0.0240	0.0413	0.0584	0.0749	0.0910
		Red Wind Tunnel	0.0288	0.0496	0.0700	0.0899	0.1091
	0.15	SLT Tunnel	0.0300	0.0473	0.0643	0.0807	0.0966
		Red Wind Tunnel	0.0360	0.0568	0.0771	0.0969	0.1159
	0.18	SLT Tunnel	0.0360	0.0533	0.0702	0.0865	0.1022
		Red Wind Tunnel	0.0432	0.0640	0.0842	0.1038	0.1227
200	0.12	SLT Tunnel	0.0267	0.0459	0.0649	0.0833	0.1011
		Red Wind Tunnel	0.0320	0.0551	0.0778	0.0999	0.1213
	0.15	SLT Tunnel	0.0333	0.0526	0.0714	0.0897	0.1073
		Red Wind Tunnel	0.0400	0.0631	0.0857	0.1077	0.1288
	0.18	SLT Tunnel	0.0400	0.0592	0.0780	0.0962	0.1136
		Red Wind Tunnel	0.0480	0.0711	0.0936	0.1154	0.1363
220	0.12	SLT Tunnel	0.0293	0.0505	0.0713	0.0916	0.1112
		Red Wind Tunnel	0.0352	0.0606	0.0856	0.1099	0.1334
	0.15	SLT Tunnel	0.0367	0.0578	0.0786	0.0987	0.1181
		Red Wind Tunnel	0.0440	0.0694	0.0943	0.1184	0.1417
	0.18	SLT Tunnel	0.0440	0.0651	0.0858	0.1058	0.1250
		Red Wind Tunnel	0.0528	0.0782	0.1029	0.1269	0.1499

A.1. SLT Adapter for Mounting




Figure A.1: Technical drawing of the aluminium adapter used to mount the airfoil support rod to the three-component force balance in the SLT, as referred to in Section 4.2. All dimensions in millimetres.

A.2. Surface Pressure Tap Model - CAD Render




Figure A.2: Transparent top-view CAD render of the Smooth PT airfoil model's centre section, with the leading edge at the bottom and the trailing edge at the top. The internal pressure channels are routed spanwise away from one another to prevent adjacent taps from interfering downstream before being collected and guided out through the hollow end caps.

A.3. SLT Airfoil model




Figure A.3: Suction surface of the Mid-2.50 airfoil with spanwise extensions mounted before wind tunnel testing. Thin tape was applied to the seams between the extensions and the airfoil ends. The stainless steel rod for mounting is seen protruding on the left-hand side of the airfoil. Leading edge and trailing edge are denoted *LE* and *TE* respectively.

B

RWT and RANS Result Figures


B.1. Wake Rake and Wall Pressure C_L and C_D Screening HeatmapsB.1.1. Wake Rake and Wall Pressure C_L Screening Heatmaps

Figure B.1: Percentage change in C_L relative to the untreated baseline geometry Smooth, as a function of angle of attack at $Re = 1 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) at dimple depths 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate an increase in C_L over Smooth, red cells a decrease.




Figure B.2: Percentage change in C_L relative to the untreated baseline geometry Smooth, as a function of angle of attack at $Re = 2 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) at dimple depths 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate an increase in C_L over Smooth, red cells a decrease.




Figure B.3: Percentage change in C_L relative to the untreated baseline geometry Smooth, as a function of angle of attack at $Re = 5 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) at dimple depths 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate an increase in C_L over Smooth, red cells a decrease.

B.1.2. Wake Rake and Wall Pressure C_D Screening Heatmaps




Figure B.4: Percentage change in C_D relative to the untreated baseline geometry Smooth, as a function of angle of attack at $Re = 1 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) at dimple depths 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate a reduction in C_D over Smooth, red cells an increase.




Figure B.5: Percentage change in C_D relative to the untreated baseline geometry Smooth, as a function of angle of attack at $Re = 2 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) at dimple depths 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate a reduction in C_D over Smooth, red cells an increase.




Figure B.6: Percentage change in C_D relative to the untreated baseline geometry Smooth, as a function of angle of attack at $Re = 5 \times 10^5$. Each row corresponds to a geometry family (Mid, Aft, Full) at dimple depths 1.25 mm, 2.50 mm, and 5.00 mm. Green cells indicate a reduction in C_D over Smooth, red cells an increase.


B.2. RANS C_L and C_D Screening HeatmapsB.2.1. RANS C_L Screening Heatmaps

Figure B.7: Percentage change in C_L relative to the untreated baseline geometry Smooth, obtained from RANS simulations, as a function of angle of attack at $Re = 2 \times 10^5$. Rows correspond to geometry families Mid, Aft, and Full at the 2.50 mm dimple depth, the only configuration evaluated in RANS. The axes, colour scale, and layout match Figure B.2 to allow direct comparison with the wind-tunnel results. Green cells indicate an increase in C_L over Smooth, red cells a decrease; cells marked *NaN* denote angles of attack for which no RANS data exists.




Figure B.8: Percentage change in C_L relative to the untreated baseline geometry Smooth, obtained from RANS simulations, as a function of angle of attack at $Re = 5 \times 10^5$. Rows correspond to geometry families Mid, Aft, and Full at the 2.50 mm dimple depth, the only configuration evaluated in RANS. The axes, colour scale, and layout match Figure B.3 to allow direct comparison with the wind-tunnel results. Green cells indicate an increase in C_L over Smooth, red cells a decrease; cells marked *NaN* denote angles of attack for which no RANS data exists.


B.2.2. RANS C_D Screening Heatmaps

Figure B.9: Percentage change in C_D relative to the untreated baseline geometry Smooth, obtained from RANS simulations, as a function of angle of attack at $Re = 2 \times 10^5$. Rows correspond to geometry families Mid, Aft, and Full at the 2.50 mm dimple depth, the only configuration evaluated in RANS. The axes, colour scale, and layout match Figure B.5 to allow direct comparison with the wind-tunnel results. Green cells indicate a reduction in C_D over Smooth, red cells an increase; cells marked *NaN* denote angles of attack for which no RANS data exists.




Figure B.10: Percentage change in C_D relative to the untreated baseline geometry Smooth, obtained from RANS simulations, as a function of angle of attack at $Re = 5 \times 10^5$. Rows correspond to geometry families Mid, Aft, and Full at the 2.50 mm dimple depth, the only configuration evaluated in RANS. The axes, colour scale, and layout match Figure B.6 to allow direct comparison with the wind-tunnel results. Green cells indicate a reduction in C_D over Smooth, red cells an increase; cells marked *NaN* denote angles of attack for which no RANS data exists.

B.3. RANS Timeseries Results Reynolds number 200k




Figure B.11: Time histories of the lift coefficient (top row) and drag coefficient (bottom row) for geometry families Smooth (left) and Mid-2.50 (right) at $Re = 2 \times 10^5$, shown for all six angles of attack. Dashed horizontal lines indicate the corresponding XFOIL reference values [2] at the same Reynolds number. The first 500 s are excluded to omit the initial transient. Steady-state cases converge to a constant value, while cases near and beyond stall exhibit persistent oscillations; time-averaged values over the final 25 % of the record are used as the reported RANS result in figure 7.4.




Figure B.12: Time histories of the lift coefficient (top row) and drag coefficient (bottom row) for geometry families Aft-2.50 (left) and Full-2.50 (right) at $Re = 2 \times 10^5$, shown for all six angles of attack. Dashed horizontal lines indicate the corresponding XFOIL reference values [2] at the same Reynolds number. The first 500 s are excluded to omit the initial transient. Steady-state cases converge to a constant value, while cases near and beyond stall exhibit persistent oscillations; time-averaged values over the final 25 % of the record are used as the reported RANS result in figure 7.4.

B.4. RANS Timeseries Results Reynolds number 500k




Figure B.13: Time histories of the lift coefficient (top row) and drag coefficient (bottom row) for geometry families Smooth (left) and Mid-2.50 (right) at $Re = 5 \times 10^5$, shown for all six angles of attack. Dashed horizontal lines indicate the corresponding XFOIL reference values [2] at the same Reynolds number. The first 500 s are excluded to omit the initial transient. Steady-state cases converge to a constant value, while cases near and beyond stall exhibit persistent oscillations; time-averaged values over the final 25 % of the record are used as the reported RANS result in figure 7.6.




Figure B.14: Time histories of the lift coefficient (top row) and drag coefficient (bottom row) for geometry families Aft-2.50 (left) and Full-2.50 (right) at $Re = 5 \times 10^5$, shown for all six angles of attack. Dashed horizontal lines indicate the corresponding XFOIL reference values [2] at the same Reynolds number. The first 500 s are excluded to omit the initial transient. Steady-state cases converge to a constant value, while cases near and beyond stall exhibit persistent oscillations; time-averaged values over the final 25 % of the record are used as the reported RANS result in figure 7.6.

C

Surface Pressure Tap Result and Figures


C.1. Surface Pressure C_P FiguresC.1.1. Reynolds Number 1×10^5 

Figure C.1: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 0° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.2: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 2° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.3: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 4° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.4: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.5: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 8° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.6: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 10° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.7: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 12° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.8: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.9: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 1 \times 10^5$ and angle of attack 16° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.


C.1.2. Reynolds Number 2×10^5 

Figure C.10: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 0° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.11: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 2° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.12: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 4° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.13: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.14: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 8° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.15: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 10° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.16: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 12° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.17: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.18: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 16° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.


C.1.3. Reynolds Number 5×10^5 

Figure C.19: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 5 \times 10^5$ and angle of attack 0° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.20: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 5 \times 10^5$ and angle of attack 2° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.21: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 5 \times 10^5$ and angle of attack 4° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.22: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.23: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 8° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.24: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 10° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.25: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 5 \times 10^5$ and angle of attack 12° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.26: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 5 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.




Figure C.27: Surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 5 \times 10^5$ and angle of attack 16° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement.

C.2. Numerical and Experimental Surface Pressure C_P Comparison Figures

C.2.1. Reynolds Number 2×10^5




Figure C.28: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 0° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.29: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.30: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 10° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.31: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 12° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.32: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.33: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $Re = 2 \times 10^5$ and angle of attack 16° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).


C.2.2. Reynolds Number 5×10^5 

Figure C.34: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 0° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.35: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 6° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.36: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 10° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.37: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 12° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.38: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 14° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).




Figure C.39: RWT surface pressure coefficient $-C_p$ as a function of chordwise position x/c at $\text{Re} = 5 \times 10^5$ and angle of attack 16° , comparing the undimpled baseline geometry Smooth (blue, dashed) to the Mid-2.50 dimple-depth variant (orange, solid). Pressure-side taps (1–13) are shown as circles and suction-side taps (13–32) as squares; error bars indicate the standard deviation across the measurement. RANS surface pressures for the same Reynolds number and angle of attack are included for the same baseline geometry Smooth (green, solid) and the Mid-2.50 dimple-depth variant (purple, solid).

D

PIV, Force Balance and Tuft Results and Figures

D.1. SLT C_L and C_D Screening Heatmaps

D.1.1. SLT C_L/C_D Screening Heatmaps




Figure D.1: Percentage change in C_L/C_D relative to the undimpled baseline geometry Smooth, from the SLT force balance test campaign, as a function of angle of attack for Reynolds numbers $Re = 1 \times 10^5$. Each row corresponds to one of three dimple depth variants, 1.25 mm, 2.50 mm, and 5.00 mm, all with the Mid dimple chord position. Green cells indicate an improvement over Smooth, and red cells a degradation.




Figure D.2: Percentage change in C_L/C_D relative to the undimpled baseline geometry Smooth, from the SLT force balance test campaign, as a function of angle of attack for Reynolds numbers $Re = 2 \times 10^5$. Each row corresponds to one of three dimple depth variants, 1.25 mm, 2.50 mm, and 5.00 mm, all with the Mid dimple chord position. Green cells indicate an improvement over Smooth, and red cells a degradation.


D.1.2. SLT C_L Screening Heatmaps

Figure D.3: Percentage change in C_L relative to the undimpled baseline geometry Smooth, from the SLT force balance test campaign, as a function of angle of attack for Reynolds numbers $Re = 1 \times 10^5$. Each row corresponds to one of three dimple depth variants, 1.25 mm, 2.50 mm, and 5.00 mm, all with the Mid dimple chord position. Green cells indicate an improvement over Smooth, and red cells a degradation.




Figure D.4: Percentage change in C_L relative to the undimpled baseline geometry Smooth, from the SLT force balance test campaign, as a function of angle of attack for Reynolds numbers $Re = 2 \times 10^5$. Each row corresponds to one of three dimple depth variants, 1.25 mm, 2.50 mm, and 5.00 mm, all with the Mid dimple chord position. Green cells indicate an improvement over Smooth, and red cells a degradation.


D.1.3. SLT C_D Screening Heatmaps

Figure D.5: Percentage change in C_D relative to the undimple baseline geometry Smooth, from the SLT force balance test campaign, as a function of angle of attack for Reynolds numbers $Re = 1 \times 10^5$. Each row corresponds to one of three dimple depth variants, 1.25 mm, 2.50 mm, and 5.00 mm, all with the Mid dimple chord position. Green cells indicate an improvement over Smooth, and red cells a degradation.




Figure D.6: Percentage change in C_D relative to the undimple baseline geometry Smooth, from the SLT force balance test campaign, as a function of angle of attack for Reynolds numbers $Re = 2 \times 10^5$. Each row corresponds to one of three dimple depth variants, 1.25 mm, 2.50 mm, and 5.00 mm, all with the Mid dimple chord position. Green cells indicate an improvement over Smooth, and red cells a degradation.

D.2. PIV Flow Field Figures

All results presented in this section are the time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ and 2×10^5 at angle of attack 0° , 4° , 6° , 10° , 14° , comparing the undimpled baseline geometry Smooth to the Mid-1.25, Mid-2.50, and Mid-5.00 dimple-depth variants. The airfoils are masked in white. The shadow region where the laser sheet is blocked by the airfoils and no valid PIV data exists is masked in grey. The hatched region marks areas affected by excessive laser-light reflections and scratches in the test-section window and should be disregarded. For the dimpled cases, the four grey half-circles on the suction-side surface mark the chordwise locations of the four dimple rows.

D.2.1. Reynolds Number 1×10^5




Figure D.7: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ at angle of attack 0° and 4° , comparing the undimpled baseline geometry Smooth (a,c) to the Mid-1.25 dimple-depth variant (b,d).




Figure D.8: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ at angle of attack 6° , 10° , and 14° , comparing the undimpled baseline geometry Smooth (a,c,e) to the Mid-1.25 dimple-depth variant (b,d,f).




Figure D.9: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ at angle of attack 0° and 4° , comparing the Mid-2.50 (a,c) and Mid-5.00 (b,d) dimple-depth variants.




Figure D.10: Time-averaged velocity magnitude field measured by PIV at $Re = 1 \times 10^5$ at angle of attack 6° , 10° , and 14° , comparing the Mid-2.50 (a,c,e) and Mid-5.00 (b,d,f) dimple-depth variants.


D.2.2. Reynolds Number 2×10^5 

Figure D.11: Time-averaged velocity magnitude field measured by PIV at $Re = 2 \times 10^5$ at angle of attack 0° and 4° , comparing the undimpled baseline geometry Smooth (a,c) to the Mid-1.25 dimple-depth variant (b,d).




Figure D.12: Time-averaged velocity magnitude field measured by PIV at $Re = 2 \times 10^5$ at angle of attack 6° , 10° , and 14° , comparing the undimpled baseline geometry Smooth (a,c,e) to the Mid-1.25 dimple-depth variant (b,d,f).




Figure D.13: Time-averaged velocity magnitude field measured by PIV at $Re = 2 \times 10^5$ at angle of attack 0° and 4° , comparing the Mid-2.50 (a,c) and Mid-5.00 (b,d) dimple-depth variants.





Figure D.14: Time-averaged velocity magnitude field measured by PIV at $Re = 2 \times 10^5$ at angle of attack 6° , 10° , and 14° , comparing the Mid-2.50 (a,c,e) and Mid-5.00 (b,d,f) dimple-depth variants.

D.3. Tufts Pictures



(a) Smooth baseline airfoil at $\alpha = 0^\circ$.




(b) Dimpled airfoil Mid-2.50 at $\alpha = 0^\circ$.

Figure D.15: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 0^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.



(a) Smooth baseline airfoil at $\alpha = 2^\circ$.




(b) Dimpled airfoil Mid-2.50 at $\alpha = 2^\circ$.

Figure D.16: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 2^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.



(a) Smooth baseline airfoil at $\alpha = 4^\circ$.




(b) Dimpled airfoil Mid-2.50 at $\alpha = 4^\circ$.

Figure D.17: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 4^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.



(a) Smooth baseline airfoil at $\alpha = 8^\circ$.



(b) Dimpled airfoil Mid-2.50 at $\alpha = 8^\circ$.

Figure D.18: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 8^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.



(a) Smooth baseline airfoil at $\alpha = 10^\circ$.



(b) Dimpled airfoil Mid-2.50 at $\alpha = 10^\circ$.

Figure D.19: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 10^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.



(a) Smooth baseline airfoil at $\alpha = 12^\circ$.



(b) Dimpled airfoil Mid-2.50 at $\alpha = 12^\circ$.

Figure D.20: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 12^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.



(a) Smooth baseline airfoil at $\alpha = 14^\circ$.



(b) Dimpled airfoil Mid-2.50 at $\alpha = 14^\circ$.

Figure D.21: Tuft flow visualisation on the suction surface of (a) Smooth and (b) Mid-2.50 at $\alpha = 14^\circ$, $Re = 2.0 \times 10^5$ ($U_\infty = 15$ m/s). The leading edge is to the left. Tufts made of yellow yarn are visible across the entire suction surface of the airfoil.