

POLITECNICO DI TORINO

Master's Degree in Aerospace Engineering



**Politecnico
di Torino**



Development of a novel experimental setup for wing
surface morphing to delay transition on swept wings

Supervisors

Prof. Gioacchino CAFIERO
Prof. Marios KOTSONIS
Dr. Alberto Felipe RIUS VIDALES
Ir. Marco RADAELLI

Candidate

Luca COZZANI



July 2026



Development of a novel experimental setup for wing surface morphing to delay transition on swept wings

Luca Cozzani

Instructors:	Prof. Marios Kotsonis Dr. Alberto Felipe Rius Vidales Prof. Gioacchino Cafiero
Teaching Assistant:	Ir. Marco Radaelli
Project Duration:	October, 2025 - April, 2026
Faculty:	Faculty of Aerospace Engineering, Delft

Cover: Cleaning the M3J Wing model in the Low Turbulence Tunnel test section from the DREs residuals - photograph taken by M. Radaelli from the divergent section of the wind tunnel.

Abstract

Smooth surface features, or humps, were recently found to stabilise stationary Crossflow Instabilities (CFI) in swept-wing boundary layers, making them a promising technique for passive laminar flow control. The aim of this thesis is to facilitate further research in this area through the design, commissioning, and assessment of a new setup capable of morphing smooth surface features on a spanwise-invariant swept-wing model. The morphing setup is made by three main components: the pre-existing spanwise-invariant symmetrical wing model, modified with an internal cavity on one side, a motorised assembly housed within the cavity, driving twenty independent vertically moving actuators, and a silicone rubber skin connected to those actuators. The influence of the generated surface features on the laminar–turbulent transition is assessed through infrared thermography measurements, while the resulting surface geometry is verified through laser distance measurements. Their effect is further investigated under different forced-transition conditions using discrete roughness elements (DREs) placed just downstream the neutral point, following configurations previously reported in the literature.

The commissioning of the setup revealed the need for a better laser measurement system and highlighted a spanwise non-uniformity of the skin. Despite these limitations, at the midspan location, the setup reliably reproduces the surface features and yields repeatable transition-front data. The comparisons between the modified and clean sides of the symmetrical wing model, together with data acquired before the silicone skin was installed, show that the effect of the modified side is limited to a downstream shift of the transition front, with the original transition-location trends preserved. The setup allows the skin geometry to be morphed, infrared images to be acquired, and the resulting surface shape to be verified in real time within the wind tunnel, enabling fast data acquisition and troubleshooting. Furthermore, the setup has been successfully employed as a platform for live optimisation campaigns, with infrared images, laser measurements, and actuators' positions integrated into a single software environment. Several observations were made regarding the parameters that govern a hump's effect. The main trends concern the existence of a critical ratio between the hump height and the boundary layer thickness, although not yet quantified; the better performance of asymmetric humps with a steeper gradient on the upstream side; the upstream movement of the transition location produced by humps placed progressively further downstream; and the ability of wider humps to keep delaying transition at higher Reynolds numbers, whereas narrower humps trip the flow.

The most unexpected finding is the minimal influence of dimples, humps with a negative height, on the transition front. Conversely, the most promising result for future research is the persistent transition delay produced by humps in an unforced CFI environment. Overall, it can be concluded that, although the setup can still be improved, it already constitutes an efficient and versatile platform for advancing the understanding of smooth surface features and their application to passive laminar flow control.

Preface

I'm not really skilled at greeting people, especially when I have to do it on paper. Because of this, I hope that even if these words may sound a bit generic, the meaning behind them is great.

First I would like to thank every person that helped me throughout this thesis journey, starting from Marios and Marco, followed by all the professors, PhDs and WT members who maybe gave me good advice or just improved my mood of the day. Regarding Delft I cannot not mention my Italian group of friends and my crazy football team, without the low cortisol time they gifted me I couldn't have achieved any good result.

This chapter of my life in the Netherlands would never have begun if the entire world of the Formula Student and in particular my teammates in Squadra Corse hadn't been part of my life. As I write this, the summer races are taking place, Forza SC. Special thanks to my dear aerodynamics friends both brothers, sons and grandsons (as silly as it seems).

Of course I want to say something to my friends in Turin with whom, one exam after another, we grew as a group, even if now our paths are slowly diverging, thank you!. Special nomination to the two most annoying housemates in Turin, no rubber chicken or bun of bread had been mistreated in that house.

Something I feel I have to do but I really know that my words cannot carry the meaning is to thank all of my friends from my hometown Spezia and my Suvero second family. I will not say anything else but U KNOW all of it.

Finally, I want to thank my family, starting from my brother, that is not so useless, my cousins, relatives, my mum and dad. The two of them are truly the reason without which I would not have been able to develop myself, I will figure out a proper way to thank you. Finally finally I want to thank my two grandfathers that I really feel had led my way to the engineering world, each for a different reason. (Regarding my grand father Cozzani, during my thesis works I discovered with great amusement that I'm not the first generation to appear in the scientific literature as an author [4])

*Luca Cozzani
La spezia, July 2026*

Contents

Abstract	ii
Preface	iv
Nomenclature	xii
1 Introduction	1
1.1 Motivation	1
1.2 Research aim and document structure	2
2 Literature Review	3
2.1 Boundary-layer transition fundamentals	3
2.2 Linear stability and transition prediction	4
2.3 Swept-wing boundary layers and instability mechanisms	5
2.4 Crossflow instabilities	8
2.5 Smooth surface features	9
2.6 Final remarks and research motivation	12
3 Methodology	14
3.1 Introduction	14
3.2 Wing assembly	14
3.2.1 Skin	14
3.2.2 Mechanical assembly	19
3.2.3 M3J wing model and adaptation	21
3.3 Tunnel facility and measurement setup	24
3.4 Measurement procedures	27
3.4.1 Infrared thermography	27
3.4.2 Laser data handling	29
3.4.3 Setup data acquisition procedure	32
3.4.4 Preliminary optimization loop	33
4 Results	36
4.1 Silicone skin characterization	36
4.1.1 Skin geometric design	36
4.1.2 Skin sensitivity analysis	39
4.1.3 Skin span-wise measurements	40
4.2 Data validation	42
4.2.1 Preset procedure repeatability	42
4.2.2 Effect of angle of attack on the preset	42
4.2.3 Repeatability of the setup	43
4.2.4 Cross model sides validation of the IR data	44
4.2.5 Cross-campaign validation of the IR data	45
4.3 Effect of the DREs	46
4.3.1 Effect of DRE spanwise spacing	46
4.3.2 Effect of amplitude	46
4.3.3 Unforced results	47
4.4 Parametric Study Results	48
4.4.1 Height effect	49
4.4.2 Ratio effect	49
4.4.3 Skewness effect	51
4.4.4 Chord-wise location effect	51

4.4.5	Maximum delay visualization	52
4.5	Unconventional shapes	54
4.5.1	Reverse Humps	54
4.5.2	Additional unconventional shapes	54
4.6	Preliminary Optimization Loop	57
5	Conclusions and Recommendations	60
5.1	Conclusions	60
5.2	Recommendations	61
	References	63

List of Figures

2.1	The path from receptivity to laminar-turbulent transition as illustrated by Morkovin [21]	4
2.2	Inviscid and viscous instability maps	5
2.3	Inflectional instability on a rotating disk	6
2.4	Comparison of stability tools against experiment	7
2.5	Three-dimensional swept-wing boundary-layer profile	7
2.6	Görtler vortices and the velocity modulation they induce	8
2.7	Saw-tooth transition front of stationary crossflow vortices	9
2.8	Secondary instability modes of a saturated crossflow vortex	9
2.9	Hump visualization from Rius-Vidaes et al. [26]	10
2.10	Thermal maps (flow from left to right). Solid orange line and orange region indicate hump apex and width. Prefix A (high) and B (low) indicate the DRE forcing amplitude and solid grey lines their streamwise location. Suffix C indicates the reference case and H the hump case. Image from Rius-Vidaes et al. [26]	11
2.11	(a–e) Base flow components of chordwise velocity (U) (red dotted and black solid lines) and wall-normal velocity ($V \times 10$) (red dash-dotted and black dashed lines) in the hump (red) and reference (black) cases. (f) Spatial evolution of the cross-flow velocity in the hump case. Black dashed line indicates the δ_{99} . Grey dashed lines indicate loci of inflection points. The hump is indicated in orange. Image from Westerbeek et al. [38]	11
2.12	(a–f) Chordwise (solid line) and wall-normal (dashed line) perturbation velocity profiles over the hump (red) and reference case (black), at selected chordwise stations. For ease of visualisation, $ \hat{v} $ is scaled by a factor 10. Horizontal dashed blue lines indicate the location of inflection points in the cross-flow profile of the base flow in the hump case. (g) Contour map of the amplitude function of the chordwise perturbation velocity from $ \hat{u} = 0$ (white) to $ \hat{u} = 0.5$ in 10 levels. Red solid and black dashed lines indicate the location of local maxima of $ \hat{u} $ in the hump and reference case, respectively. Image from Westerbeek et al. [38]	12
3.1	On the left the codes workflow and on the right the design parameters schematic	15
3.2	On the left the Boundary Condition's code plot. The red stars are a section of the M3J shape, the green line the target curve with an hump (difference with the M3J shape) and the blue stars are the boundary conditions points (20 actuators, LE and TE strips). On the right is shown every node of the deformed mesh in output from the solver.	15
3.3	CAD model of the skin mould	16
3.4	Final skin section from the CAD model	17
3.5	On the left the problem of the silicone lip and on the right a close up on the fit of the complete assembly composed of the Inner wing box, LE strip and silicone skin. In the centre the CAD visualization of the non-parallel top and bottom surfaces of the LE aluminium strip.	17
3.6	First test of the casting procedure. On the right the visualization of a leakage	18
3.7	Casting procedure steps. From left to right: the legs assembly, composed by the three transparent bridges on the top, the twenty aluminum legs and the twenty-one transparent legs supports; the plexiglass mould fixed on a wooden board with the LE and TE aluminum strips embedded; The first step of the casting allowing the silicone to level out; dipping step of the legs assembly in the pool of silicone towards a complete skin.	18
3.8	Assembly of the skin and the mechanical actuators using an experimental green skin composed of only five legs.	19
3.9	Axonometric CAD view of one actuator unit	19
3.10	PCB on the back side of the metallic box of the mechanical assembly	20
3.11	Movement chain schematics	20

3.12 Mechanical assembly components	20
3.13 Displacement validation. On the left, the measured top surface height of 19 of 20 actuators in 8 positions. On the right, the error between the reference position and the target position (hall sensor position)	21
3.14 Actuators assembly mechanism	21
3.15 Comparison between the NACA66018 and the 66018M3J from Serpieri et al. [34]	21
3.16 Process of the M3J wing adaptation. On the left the starting Hole. In the middle the jig and the metallic supports glued to the model. On the right the image before the assembly lay down, with the three thermocouples and the engraved bottom surface	22
3.17 Wing model with the mechanical assembly on display and with the skin. Note the possibility of extracting the mechanical assembly through handles	22
3.18 On the left, the surface finish achieved at the LE between the wing model, the fixed metal strip and the aluminium strip of the silicone skin. On the right, a close-up picture of the LE of the model in the WT test section; the golden dots are the DREs.	23
3.19 Schematic drawing of the Low Turbulence Wind Tunnel. The drawing is taken from Rius-Vidales [27] where a more detailed description can be found.	24
3.20 Measurement chain schematics. Wind tunnel section room: control area on the left, wind tunnel section on the right, passage under the WT stairs on the top. Flow direction bottom-up.	24
3.21 Silicone side of the wind tunnel test section during the campaign. Equipment outside the test section: IR camera 1, halogen lamps, Initium and laser scanner. Equipment inside the test section: wing model and silicone skin delimited by the aluminium strips.	25
3.22 Dewarping procedure on the clean side of the wing model. On the left, the picture of the paper target sheet positioned on the wing model. In the middle, the original IR image taken with the IR camera. On the right, the final dewarped image with the non-dimensional chordwise coordinate system superimposed.	27
3.23 IR DIT procedure. The flow goes from right to left and the axis are in mm. In the first picture there is the dewarped image, in the second one the subtraction between the first image and a lower Reynolds one, and finally there is the result of the binarisation. The Reynolds numbers of the image and of the subtracted one are $Re = 2.81 \cdot 10^6$ and $Re = 2.42 \cdot 10^6$, respectively.	28
3.24 On the left the laser Scanner setup shooting the red beam trough the glass onto the wing model. On the right a single snapshot of raw data in its coordinate system (without equal axis)	29
3.25 Live laser scan visualization code flow chart	30
3.26 Laser raw data of 1000 consecutive snapshots in the same spot with no air flow. Full field and zoom view (the red rectangle in the full view represent the zoom)	30
3.27 Orthographic projections of the laser measurement tool. The laser beam is shot from the left side of the tool and the light sensor is on the right part of it. Based on the position of the reflected light on the sensor itself, the distance of the reflecting point can be calculated.	31
3.28 Laser post processing code flow chart	31
3.29 Bimodal filter visualisation. The raw data sample is composed of the red and black dots.	32
3.30 Error assessment through the standard deviation of the filtering output. On the left, the full view of a scan where the hump shape can be noted on the right-hand part of the picture. On the right, a zoom of the full view of a different sample of data, where three different curves are visible, as well as the final standard deviation of the black raw points.	32
3.31 PSO Code Flow chart	34
3.32 PSO reward function evaluation procedure. From left to right: the original raw image, the cropped raw image showing the location of the reference pixel, and the normalized image used as input to the reward function. The reward function is computed by summing the values of all non-NaN pixels in the normalized image.	34
3.33 Example of a PSO progress visualization plot at its last iteration.	35

4.1	Initial study on the skin design. Boundary condition parameter sensitivity. In blue, the RMS error and, in orange, the maximum difference between the target and the simulated top surface. The top row shows the effect of the gradient; the left column shows the effect of the position of the peak of the hump; the right column shows the effect of the leading edge.	37
4.2	Initial study on the skin design. Design parameter sensitivity. On the vertical axis, the RMS error in millimetres. On the left, the results of the BC <i>h3 R20 e000 x_c235</i> and on the right the BC <i>h3 R50 e000 x_c260</i> , one representing a high-gradient case and the other a lower one	37
4.3	Ranking of the geometries. On the bottom side the code names of the parameter sets. Higher the counter the better the geometry is in different BC. Blue name with red bars are the geometries with a height lower than 6.5 mm (6 mm was the design target). . . .	38
4.4	Skin sensitivity to temperature and flow velocity. In each image there are four curves describing the effect of the environmental conditions on the silicone skin surface shape. The curves in the four images are: the M3J preset, the $x_c = 0.15$ hump, the $x_c = 0.25$ hump and the $x_c = 0.35$ hump.	39
4.5	Laser line projected on the wing surface (perpendicular to the LE) and an arrow indicating the path covered by the laser line.	40
4.6	Preset spanwise non-uniformity visualisation. On the left, a 3D mapping of the laser data showing the error between the laser measurements and the M3J shape. On the right, a qualitative validation picture of the wing exposed to a grazing light directed along the surface to highlight local waviness.	40
4.7	Preset spanwise non-uniformity study. On the left, a heatmap of the top view of the laser data, showing the error between the laser measurements and the M3J shape as colour. On the right, the measurements of the skin compared with the M3J shape at three different span locations and the error of those profiles.	41
4.8	Full view of the preset consistency study. The presets are the M3J baseline, which is modified to obtain a hump. The first 55% of the chord is shown in addition to the silicone skin location (grey vertical bars).	42
4.9	Error in the silicone skin between the M3J target and the different presets. On the left, the actual errors when all the measurements were taken on current presets; on the right, the measurements taken when only the last preset was the current one. Note the difference in the y-axis between the two figures and that only one preset can be the current one. . .	42
4.10	On the left the error between the M3J target and the M3J target calculated at 3 degrees. On the right, the comparison between using the two different targets as a reference in the error evaluation for the same data.	43
4.11	Comparison of the use of the two M3J targets (M3J and M3J_3deg) in the calibration phase and in the error evaluation phase.	43
4.12	Repeatability of the transition fronts location data between different days. On the right the M3J preset data is isolated.	44
4.13	Cross sides validation, the same DREs forcing are drawn in the same colour. On the left the IR data acquired on the clean side of the wing. On the right the IR data of the silicone side with the preset M3J shape of the skin. The AoA is 3 and -3 degrees to match the same pressure distribution on the pressure side, where the IR measurement is taken. .	44
4.14	Comparison of the IR data of the present campaign (continuous line) and the IR data taken in a previous campaign (dotted line). The flow conditions are the same, but the hump manufacturing technology changes between the modified wing model and a cast hump glued in position on the wing surface.	45
4.15	Amplitude study on the IR data of the transition location. On the left, the high-amplitude case and, on the right, the low-amplitude case. The humps shaped in both cases are the same. The black line represents the M3J preset, so any point to the right of it is delaying transition.	46
4.16	Clean case with no forcing through DREs: IR transition data. The measurements are taken from the silicone side of the wing model, morphing different hump shapes with constant $R = 50$ and $x_c = 0.15$. The black line represents the M3J preset, so any data point to the right of it is a transition delay	47

4.17 Hump geometrical shape and parameters. Note the different scales of the axes.	48
4.18 Height study. In particular, these data are hump geometries with fixed $R = 50$, $e = 0$ and $x_c = 0.15$, with different heights. The black line represents the M3J preset, so any data point to the right of it is a transition-delaying geometry. As a reference, in the top part of the figure there is half of the M3J profile in grey and the location of the silicone skin pasted on it in red.	49
4.19 Ratio study. In the first three images there are sweeps of Ratios at the three constant heights of 1 mm, 1.5 mm and 2 mm. In the last image the effect of the Ratio is evaluated in different chordwise positions	50
4.20 Laser measurement Ratio. The image shows the surface measurement of the $h = 1.5$ mm, $R = 20$, $e = 0$ and $x_c = 0.15$ both in full view (on the top) and as the error between the measurement and the target hump. More Reynolds curves are shown.	50
4.21 Laser measurement Skewness. The image shows the surface measurement of the $h = 2.5$ mm, $R = 50$, $e = -0.6$ and $x_c = 0.15$ both in full view (on the top) and as the error between the measurement and the target hump. More Reynolds curves are shown.	51
4.22 Skewness study. The first three images are part of a height sweep study where the fixed parameters are $R = 50$ and $x_c = 0.15$, whereas in bottom right there is a sweep of skewness and chordwise position combined at a fixed $h = 2$ mm and $R = 50$	52
4.23 Chordwise location study. In the three images there are sweeps of chordwise location at the three constant heights of 1 mm, 1.5 mm and 2 mm.	53
4.24 Visualization of some normalized averaged dewarped full images (flow from right to left). All the images are taken with a $\lambda = 7.5$ mm DREs configuration and similar Reynolds number (written in the title). The first image is the M3J preset, the second one is the $h2.5 R50 e = -0.6 x_c = 0.15$ and the third one the $h2.0 R50 e = 0 x_c = 0.15$	53
4.25 Reverse hump IR transition location data. There are data of different reverse humps at two x_c locations. The black line represents the M3J preset, so any point to the right of it is delaying transition.	54
4.26 Laser measurement reverse Hump. The image shows the surface measurement of the $h = -3$ mm, $R = 50$, $e = 0$ and $x_c = 0.165$ both in full view (on the top) and as the error between the measurement and the target hump. More Reynolds curves are shown	55
4.27 Comparison of the unconventional shapes tested and the PSO run 3 result. The plot shows the error between the measurement and the M3J target.	55
4.28 Unconventional shapes IR transition location data. There are data of different shapes shown in Fig. 4.27. The black line represents the M3J preset, so any point to the right of it is delaying transition The purple line is the reference hump that should have been improved by the unconventional shapes.	55
4.29 PSO final update visualization, run 1	57
4.30 PSO final update visualization, run 3	57
4.31 PSO reference IR image. Normalized IR image of a M3J shape of the silicone skin under $2.2 \cdot 10^6$ Re with a $\lambda = 7.5$ mm DRE forcing	58
4.32 PSO output shape, run 1. The plot shows the error between the laser measurement and the M3J target	58
4.33 PSO output shape, run 3. The plot shows the error between the laser measurement and the M3J target	59

List of Tables

4.1	Initial study Boundary Condition cases, Hump shapes: hX , height in mm; RXX , Ratio; $eXXX$, skewness $\ast 10^{-2}$; x_cXXX , chord non dimensional position $\ast 10^{-3}$	36
4.2	Initial Design parameter space, 27 possible foot geometries	36
4.3	Final study Boundary Condition cases, Hump shapes: hX , height in mm; RXX , Ratio; $eXXX$, skewness $\ast 10^{-2}$; x_cXXX , chord non dimensional position $\ast 10^{-3}$	38
4.4	Final Design parameter space, 36 possible foot geometries	38
4.5	Shapes recorded in the spanwise laser measurements. The other humps parameters were $R = 50$ and $e = 0$ and the wind tunnel speed was 0 m/s.	40
4.6	Data acquired during the DREs study. The dimensions of the DREs were fixed (see Section 3.3). The other hump parameters were fixed to: $x_c = 0.15$, $R = 50$ and $e = 0$. The "-.-." line stands for one step of the sweep for each "-" or ".".	46
4.7	Data acquired during the parametric study	49

Nomenclature

Abbreviations

Abbreviation	Definition
AoA	Angle of Attack
BC	Boundary Conditions
DRE	Discrete Roughness Element
e	Skewness of the Hump shape
IR	Infra-Red
LE	Leading Edge
LL	Leg Length (skin design variable)
LT	Leg Thickness (skin design variable)
LTR	Leg Thickness Ratio (skin design variable)
M3J	Name of the wing model
R	Ratio of the Hump shape: base / height
RMS	Root Mean Square
RMSE	Root Mean Square Error
ST	Skin Thickness (skin design variable)
TE	Trailing Edge

Introduction

1.1. Motivation

Air transport is a central part of modern life and supports global commerce, tourism and human mobility. Because of its economic importance and the connectivity it provides, the sector has grown rapidly in recent years [36]. This growth, however, has an environmental cost. The combustion products of aircraft engines, such as carbon dioxide, nitrogen oxides, soot and water vapour, are harmful to the atmosphere [2], and they make the targets set by programmes such as Flightpath 2050 [15] harder to meet. Together with a growing awareness of climate change, this has motivated the research towards a more efficient and sustainable aviation sector.

Besides advances in propulsion, the reduction of aerodynamic drag is one of the most effective ways of lowering fuel consumption. For a transport aircraft, skin-friction drag is the largest drag contribution and accounts for more than 50% of the total [33]. This viscous force is caused by the shear stress between the fluid and the surface moving through it. At the same Reynolds number, a laminar boundary layer has a much lower wall shear stress, and therefore a lower skin-friction drag, than a turbulent one [33]. Preventing or delaying the formation of a turbulent boundary layer is thus an attractive goal: a 40% extension of the laminar region over the wing, tailplane, fin and nacelles has been estimated to give a total drag reduction of the order of 16% [33]. The active and passive techniques used to stabilise the boundary layer and extend the laminar region are known collectively as Laminar Flow Control (LFC).

The use of laminar flow control on modern transport aircraft is made more difficult by swept wings. Sweep is used to reduce the effective Mach number normal to the leading edge and so to delay the wave drag caused by compressibility effects in the transonic regime. As a result, the flow develops both normal and parallel to the leading edge, which produces a three-dimensional boundary layer [31]. Near the wall, the momentum deficit unbalances the pressure and centrifugal forces acting on the curved inviscid streamline. This generates a secondary crossflow, whose inflectional velocity profile is inviscidly unstable. The crossflow instabilities (CFI) that arise are among the main mechanisms driving transition on wings of moderate to high sweep [3], and a way to suppress them completely has not yet been found.

For this reason, the control of crossflow instabilities has become an active field of research. A recent and different approach uses a smooth surface feature, or hump, placed on the suction side of the wing. Instead of acting on the incoming disturbances directly, a hump of suitable height, width and chordwise location reshapes the base flow, which can in turn stabilise the developing crossflow instabilities [39]. The transition delay reported for such features makes the hump a promising device for passive laminar flow control.

The study of these features is, however, demanding. Each hump geometry has so far been made as a dedicated, individually manufactured surface. Examining the influence of parameters such as height, width and chordwise location therefore requires a separate part to be produced and installed for every configuration. A broad parametric study, and still more an optimisation campaign, is thus slow

and labour-intensive. The present work aims to ease this constraint by supporting further research into smooth surface features. To this end, a setup able to morph such features on the wing surface has been developed, such that it can offer an efficient and versatile platform for studying the effect of humps on crossflow-driven transition, and ultimately for assessing their potential as a passive laminar flow control device.

1.2. Research aim and document structure

Following the motivation outlined above, the aim of this thesis is to facilitate further research on smooth surface features through the design, commissioning and assessment of a new setup capable of morphing such features on the spanwise-invariant M3J swept-wing model. In practice, this aim translates into three objectives: to design and manufacture the morphing setup; to validate its measurements and characterise its limitations; and, since the time constraint allowed it, to collect as many observations as possible on the effect of surface features on the transition location. It is important to note that these observations are not meant to explain the underlying physics, which is deliberately left outside the scope of this work, but to guide the research questions of future physics-oriented studies.

To help the reader, the structure of this document is briefly outlined here. Chapter 2 reviews the literature on boundary-layer transition, swept-wing instability mechanisms and smooth surface features, closing with the motivation of the present work. Chapter 3 describes the wing assembly and its development, the wind tunnel facility, the measurement procedures and the preliminary optimisation loop mechanics. Chapter 4 presents the results, starting with the skin design outcome and the data validation, followed by an analysis of the effects of the forcing and the surface features on the transition front, and concluding with the optimization results. Finally, Chapter 5 gathers the conclusions drawn from this first campaign, together with the recommendations for the future development of both the setup and the research line.

2

Literature Review

2.1. Boundary-layer transition fundamentals

A body moving in a viscous fluid is subjected to the no-slip condition, that forces the flow to slow down near its surface. This thin region close to the wall, where viscous forces are dominant, was called "boundary layer" by Prandtl [22]. Its thickness is hard to fix exactly, because the velocity reaches the outer value only far from the wall. By convention, the boundary-layer thickness δ_{99} is taken as the wall-normal distance at which the flow reaches 99% of the edge velocity u_e [32]. Outside this region the viscous forces are negligible and the flow can be treated as inviscid. The boundary-layer thickness is related to the Reynolds number at which the flow is. A higher Reynolds number yields thinner boundary layers. The Reynolds number is evaluated as:

$$Re = \frac{\rho_e L u_e}{\mu} \quad (2.1)$$

where ρ_e the density outside the boundary-layer, μ is the dynamic viscosity and L a reference length.

The velocity gradient perpendicular to the wall sets up a shear stress. For a simple two-dimensional layer it is

$$\tau = \mu \frac{\partial u}{\partial y}, \quad (2.2)$$

where $\partial u / \partial y$ the wall-normal velocity gradient. The shear stress at the wall is usually written in non-dimensional form as the skin-friction coefficient,

$$C_f = \frac{\tau_{\text{wall}}}{\frac{1}{2} \rho_e u_e^2}, \quad (2.3)$$

In practice, integral quantities are preferred to δ_{99} when the size of the layer needs to be quantified, since the 99% is difficult to measure experimentally. The displacement thickness δ^* measures the loss of mass flow caused by the velocity deficit,

$$\delta^* = \int_0^{\delta_{99}} \left(1 - \frac{u}{u_e} \right) dy, \quad (2.4)$$

and the momentum thickness θ measures the loss of momentum,

$$\theta = \int_0^{\delta_{99}} \frac{u}{u_e} \left(1 - \frac{u}{u_e} \right) dy. \quad (2.5)$$

Their ratio of the two of them defines the shape factor,

$$H = \frac{\delta^*}{\theta}, \quad (2.6)$$

that is a single parameter that indicates how full the velocity profile is, and for this reason is a useful indicator of the state of the layer.

A boundary layer can be laminar or turbulent. In a turbulent layer, mixing carries high-momentum fluid towards the wall, which increase the near-wall velocity gradient and, through Eq. 2.2, the wall shear stress. A turbulent layer therefore produces a much higher skin-friction drag than a laminar one at the same Reynolds number. This is the reason for keeping the flow laminar over as much of the surface as possible, as set out in Section 1.1.

The change from a laminar to a turbulent state is not abrupt. On a wing the layer starts laminar near the leading edge and grows along the surface. Big disturbances can trip the boundary layer at the spot leading to an abrupt transition to turbulence. Conversely, small disturbances from the environment, such as free-stream turbulence, sound and surface roughness, can enter the layer and grow inside it. The process by which these external disturbances set the initial amplitude of the disturbances inside the boundary layer is called receptivity [21], and it is the first step towards transition.




Figure 2.1: The path from receptivity to laminar-turbulent transition as illustrated by Morkovin [21]

Once inside the boundary layer, a disturbance can reach turbulence along several routes. Reshotko [25] grouped these into five paths, ordered by the level of the external disturbance, see Fig 2.1. In a low-disturbance environment, such as that of an aircraft in cruise, the relevant routes are the first ones: the disturbances are weak and grow as exponential eigenmodes, a process known as modal growth, over a long streamwise distance. When their amplitude becomes large, non-linear effects and secondary instabilities take over and the layer breaks down to turbulence. At higher disturbance levels, shorter transient and bypass routes replace this slow modal growth. As the present work concerns a low-disturbance swept-wing flow, recreating the cruise flight environment, only the low disturbance paths are investigated.

2.2. Linear stability and transition prediction

Linear stability theory studies how a small disturbance added to a steady base flow grows or decays. While the disturbance stays small its behaviour is linear, and it is written as a wave whose amplitude rises or falls along the surface [17]. A common form for a spatial wave is:

$$q'(x, y, z, t) = \hat{q}(y) e^{i(\alpha x + \beta z - \omega t)} \quad (2.7)$$

where $\hat{q}(y)$ is the shape of the disturbance and α , β and ω describe its wave behaviour. The imaginary part of the streamwise wavenumber, α_i , sets the growth: the local growth rate is $-\alpha_i$, so the wave is amplified where α_i is negative and damped where it is positive.

For a two-dimensional boundary layer, two kinds of instability are distinguished. The viscous instability, where the unstable region of the wavenumber–Reynolds plane shrinks and finally disappears as the Reynolds number increases, because the viscosity is generating the instability; this is typical of flows with a zero or favourable pressure gradient. The inviscid instability, where a band of unstable wavenumbers remains unstable even at high Reynolds number; this is typical of flows with an adverse pressure gradient. The two cases are sketched in Fig. 2.2. The key feature behind an inviscid instability is an inflection point in the velocity profile: by Rayleigh’s criterion, a point where $d^2U/dy^2 = 0$ is a necessary condition for this type of instability [17].




Figure 2.2: Unstable regions in the wavenumber–Reynolds-number plane. Left: an inviscid instability keeps a band of unstable wavenumbers at high Reynolds number. Right: a viscous instability closes as the Reynolds number increases. Images from Mack [17]

A clear example of an inflectional, inviscid instability in a three dimensional boundary layer is the flow over a rotating disk. The radial velocity is zero both at the wall and far from the disk, so its profile contains an inflection point, instead the tangential velocity profile does not, see Fig. 2.3 on the left. Between these two directions there is a range of intermediate profiles that are inflectional, and therefore unstable, which gives rise to co-rotating vortices spaced by a fixed wavelength. These appear as spiral streaks on the disk surface, showed in Fig. 2.3 on the right. Because of a similar inflectional mechanism and a similar three dimensional boundary layer, the rotating disk is often used as a model problem for the crossflow instability on a swept wing [31].

Stability theory by itself gives the growth of a disturbance, not the transition location. The two are linked by the e^N , or N-factor, method, proposed by Smith and Gamberoni (1956) [37] and, independently, by van Ingen [12]. The disturbance amplitude is integrated from the first unstable position x_0 , the neutral point,



$$N = \ln \frac{A}{A_0} = - \int_{x_0}^x \alpha_i dx, \quad (2.8)$$

and transition is assumed to take place once N reaches a value set by correlation with experiments on similar flows. The method is widely used because it is cheap, although it is empiric depending on this correlation and says nothing about the physics of the breakdown itself.

Several tools of increasing cost are available to compute the disturbance growth. Local linear stability theory, based on the Orr–Sommerfeld equation, assumes a parallel base flow and is the simplest, but the hypothesis ignore the slow streamwise growth of the boundary layer. The parabolised stability equations (PSE) relax this assumption and, in their non-linear form, also include the interaction between modes, which lets them capture the amplitude saturation of the crossflow vortices [11]. Direct numerical simulation solves the full Navier Stokes equations and is the most accurate but also the order of magnitude more expensive. As Fig. 2.4 shows, the non-linear PSE follows the measured amplitude into the saturated region, where the linear methods lose accuracy [10, 24].

2.3. Swept-wing boundary layers and instability mechanisms

The swept wings are common in modern transport aircraft to fly efficiently in the transonic regime. In fact, sweeping the wing lowers the effective Mach number normal to the leading edge, which delays



Frontispiece. Boundary-layer flow induced by a rotating disk (china-clay photograph), showing laminar flow near the centre and subsequent instability in the form of spiral vortices which cause transition to turbulent flow near the rim.
(See Gregory, Stuart, and Walker 1955)

Figure 2.3: Inflectional instability on a rotating disk. On the left the radial velocity profile contains an inflection point while the tangential one does not. On the right the resulting co-rotating vortices, seen as spiral streaks on the disk surface.

the strong rise in wave drag that appears as the flow approaches the speed of sound [1]. The downside of the sweep is that makes the flow move both normal and parallel to the leading edge, developing a three dimensional boundary layer, Saric et al.[31].

The three dimensional nature of the boundary layer is best understood by following the inviscid streamline just outside it. As the flow accelerates over the front of the wing, the favourable pressure gradient curves this streamline inboard and conversely, as the flow slows down the streamline curves outboard, see Fig 2.5. Outside the boundary layer the curvature of the streamline is balanced by the pressure field on the wing surface. However, closer to the wall, the flow carries less momentum while being subjected to the same pressure gradient (which is nearly constant across the boundary layer). As a result, the balance between inertia and the adverse pressure force is progressively lost. This causes the birth of a small secondary flow perpendicular to the inviscid streamline, called crossflow. The full velocity profile can be split into a component along the inviscid streamline and the crossflow component normal to it, as shown in Fig. 2.5. Because the crossflow velocity must vanish both at the wall and at the boundary layer edge, its profile contains an inflection point. By Rayleigh's criterion this makes the profile inviscidly unstable, and it is the source of the crossflow instability [31, 23].

A swept-wing boundary layer is subject to four instability mechanisms. Which of them is dominant depends on various factors, among which the sweep angle and the surface curvature (Bippes [5] and Saric [31]). The first three are described below, while the crossflow instability is treated in detail in Section 2.4.

The first mechanism is attachment-line instability and contamination. A swept wing has no stagnation line; instead an attachment line runs along the leading edge from root to tip. Turbulence produced at the wing root, for example at the wing-body junction, can travel along this line and trip the boundary layer over the whole span. This effect grows with the leading-edge radius, so a small radius helps to avoid it, Saric et al.[23]. On current aircraft a small fairing at the wing root, known as the Gaster bump, forces a fresh stagnation point and stops the contamination, from Gaster (1965)[9].

The second mechanism is the Görtler instability. Over a concave surface, the imbalance between




Figure 2.4: Disturbance amplitude of a stationary crossflow mode computed with the Orr–Sommerfeld equation, linear PSE and non-linear PSE, compared with experiment. The non-linear PSE captures the saturation that the linear methods miss. Adapted from [10]; experimental data from [24].




Figure 2.5: On the Left the swept-wing boundary-layer velocity profile, split into a tangential component along the inviscid streamline and a crossflow component normal to it. The crossflow component vanishes at the wall and at the boundary-layer edge, which gives the inflection point from which the crossflow instability grows. Adapted from Saric et al.[31]. On the Right the development of inviscid streamline on an infinite swept wing from Rius-Vidaes [27].

the pressure and centrifugal forces produces counter-rotating streamwise vortices [30], shown on the left of Fig. 2.6. These vortices carry high-momentum fluid towards the wall and low-momentum fluid away from it. This modulates the streamwise velocity along the span and forms alternating high- and low-shear regions near the wall, shown on the right of Fig. 2.6. Unlike the crossflow instability, the Görtler instability does not need an inflection point; it needs just a concave curvature. This instability mechanism is the less influent because most swept wings have convex surfaces near the leading edge, so this mechanism is usually absent [31].

The third mechanism is the Tollmien–Schlichting (TS) wave, a viscous instability that travels downstream as a wave. On wings with low sweep, roughly below 25° , TS waves are the main route to transition [13]. They grow in regions of adverse pressure gradient, while a favourable pressure gradient stabilises them [14]. Given his characteristics, this is the most dominant instability mechanism in general boundary layers but, a wing with a long favourable pressure gradient over its surface is designed to delay TS waves as much as possible.

As been said, the first three mechanisms can all be controlled by design. A small leading-edge radius and a Gaster bump suppress the attachment-line problem, convex surfaces near the leading edge




Figure 2.6: Görtler instability on a concave wall. Left: the counter-rotating streamwise vortices, spaced by the spanwise wavelength λ . Right: the resulting spanwise modulation of the streamwise velocity, with vortex-induced upwash and downwash and the alternating high- and low-shear regions near the wall. Images from Floryan [8]

prevent Görtler vortices, and a long favourable pressure gradient holds back TS waves. The same favourable pressure gradient, however, is exactly what feeds the crossflow instability, which becomes the dominant mechanism on wings of moderate to high sweep, above roughly 30° [13, 5].

2.4. Crossflow instabilities

As described in Section 2.3, the crossflow component of a swept-wing boundary layer has an inflection point and is therefore inviscidly unstable. The instability it produces is driven by the favourable pressure gradient: the stronger this gradient, the larger the near-wall imbalance of forces, and the stronger the crossflow [5, 23]. Crossflow instabilities therefore grow in the front part of the wing, where the flow accelerates [31].

The instability appears as co-rotating vortices whose axes lie within a few degrees of the inviscid streamline. Depending on the receptivity process, these vortices develop either as travelling modes, which move along the span, or as stationary modes, which remain fixed with respect to the wing [31]. Which of the two families dominates is not decided by linear theory alone: although the travelling modes are the ones with the higher predicted growth rate, they take over only in noisy flows with high free stream turbulence, while in the quieter conditions of cruise flight the stationary modes prevail [5, 6]. The two types leave different patterns at the wall: a transition front shaped by travelling modes appears smooth, whereas one shaped by stationary modes shows a sharp saw-tooth pattern [40] visible in Fig. 2.7. Because the present work targets the low-disturbance conditions of flight, the stationary modes are the ones of interest.

For stationary modes the initial amplitude is set mainly by the surface roughness of the wing (Bippes, 1996 [6]). This sensitivity is turned into a tool: small bumps called discrete roughness elements (DRE) can be placed in a row near the leading edge to force a selected spanwise wavelength [29, 28], see again Fig. 2.7. The spacing between the elements fixes the wavelength of the forced mode [42], which makes the DREs a simple and repeatable way to control the crossflow vortices. The elements work best when placed close to the neutral point, the streamwise location where the mode first becomes unstable, since this is where its amplitude starts to grow [24].

The growth of the stationary vortices happens in two stages [31]. In the first linear stage the weak wall-normal and spanwise disturbance velocities move low-momentum fluid away from the wall and high-momentum fluid towards it. This behaviour marks the surface with alternating low and high skin friction regions. Because the near-wall velocity gradient is large, even small displacements quickly build a strong streamwise disturbance further downstream. This growth sets the second non-linear stage, in which the boundary layer is heavily distorted and the disturbance amplitude saturates.

The saturated vortices shape strong shear layers and inflectional profiles in both the wall-normal and spanwise directions. From these, high-frequency secondary instabilities grow and cause a rapid break-



Figure 2.7: Saw-tooth transition front produced by stationary crossflow vortices. Captured by infrared thermography on the clean side of the wing model during the experimental campaign. The forcing is generated by 12 mm spaced DREs located only in the middle span, as it can be seen by the transition delay in the spanwise extremities; the flow goes from right to left.

down to turbulence, usually within a few per cent of the chord [18, 41]. Following the classification in [35], three secondary modes are distinguished according to where they sit on the instability vortex, see Fig. 2.8. These types are: the type-I (or z) mode, a high-frequency mode sit on the outer side of the upwelling region where the spanwise gradient is smallest; the type-II (or y) mode, a high-frequency mode near the top of the vortex where the wall-normal gradient is largest; and the type-III mode, a low-frequency mode placed on the inner side of the upwelling region where the spanwise gradient is largest.




Figure 2.8: Locations of the type-I, type-II and type-III secondary instability modes on the mean streamwise velocity of a saturated stationary crossflow vortex. Reproduced from [20].

The strong sensitivity of the crossflow instability to its initial conditions is in fact a way to control it. The most common control technique is upstream flow deformation: a DRE row forces a stationary mode of shorter wavelength than the naturally dominant one. Only the harmonics of this forced mode reach the non-linear stage, so the growth of the critical mode and its harmonics is delayed and transition is postponed [29, 28, 24]. A different approach changes the base flow directly so as to weaken the crossflow component, which stabilises the vortices [7]. Both routes show that a small, well-placed change to the flow or the surface can have a large effect on transition.

2.5. Smooth surface features

As the previous section noted, a small and well-placed change to the surface can have a large effect on the crossflow instability, this can clearly be a smooth hump. The study of surface features such as

smooth humps and discontinuous steps in a crossflow dominated boundary layer is recent, and the main contributions to the topic come almost entirely from Delft University of Technology research group. The transition delay effect of such features was in fact a serendipitous finding: while investigating forward- and backward-facing steps, Rius-Vidales [27] observed a delay that ran counter to the established consensus, which was that such features only advance transition. This unexpected result started further research and, in time, led to the smooth surface hump, pursued both numerically by Westerbeek et al. [39] and experimentally by Rius-Vidales et al. [26]. A visualization of such surface feature is shown in Fig. 2.9.




Figure 2.9: Hump visualization from Rius-Vidales et al. [26]

The experiments of Rius-Vidales et al. [26] were carried out in the Low Turbulence Tunnel at Delft University of Technology on the M3J swept-wing model, a spanwise-invariant wing with a 45° sweep angle at an angle of attack of 3° for which the surface carries a predominantly favourable pressure gradient. As described before, this setting is promoting stationary crossflow instabilities and is in fact the exact settings in which this thesis work had been carried out. A shallow hump, manufactured as a cast epoxy add on and verified by laser profilometry, was bonded to the wing in the favourable pressure gradient region with its apex at $x/c_x \approx 0.15$. The transition front was obtained by infrared (IR) thermography, and the crossflow vortices were characterised by particle image velocimetry (PIV). The stationary crossflow modes were forced through the use of DREs (Discrete Roughness Elements) near the leading edge selecting the spanwise wavelength.

The central finding is that the effect of the hump depends strongly on the amplitude of the incoming crossflow disturbances, which was set by the streamwise position of the DRE row relative to the neutral point of the forced mode. At a high forcing amplitude the weakly nonlinear crossflow vortices interacting with the hump produced an abrupt advancement of transition, effectively tripping the flow. At a low forcing amplitude, for which the vortices reached the hump while still growing linearly, the opposite occurred: the primary crossflow instability was stabilised and transition was delayed by about 14% of the chord with respect to the reference case without a hump. This delay proved robust to changes in the chord Reynolds number, see fig 2.10. The measured value was in fact limited downstream by the wing's own pressure recovery rather than by the cross flows instability, so that the stabilising potential of the feature may be even greater.

The velocimetry measurements traced how this stabilisation comes about, and led to a description of the interaction in terms of four chordwise regions. Between the hump leading edge and its apex the local favourable pressure gradient accelerates the boundary layer and slightly amplifies the primary instability. Downstream of the apex the curvature reverses and imposes a short adverse pressure gradient, which decelerates and thickens the boundary layer and slightly reduces the disturbance amplitude. More importantly, in this region the crossflow vortices change their wall-normal orientation, tilting against their natural direction, which points to a weakening or local reversal of the near-wall crossflow component. Farther downstream the nominal favourable pressure gradient of the wing restores the boundary layer to the reference state. Yet well beyond this recovery a pronounced and unexpected stabilisation of the primary mode and its first harmonic is measured. This stabilization is quantified by the growth rate of the instability which remains lower than in the clean case. The consequences for breakdown, and so the transition location, follow from this weakening of the primary instability. With lower amplitude in the stationary vortices, the spanwise shear that feeds the unsteady secondary instabilities (described




Figure 2.10: Thermal maps (flow from left to right). Solid orange line and orange region indicate hump apex and width. Prefix A (high) and B (low) indicate the DRE forcing amplitude and solid grey lines their streamwise location. Suffix C indicates the reference case and H the hump case. Image from Rius-Vidales et al. [26]

in 2.4) is reduced, and the velocity fluctuations associated with their growth are weaker. The rise of secondary instability, and hence of laminar breakdown, is thus postponed, which is consistent with the downstream shift of the transition front seen in the thermography.

The mechanisms behind these observations were also examined numerically. Westerbeek et al. [38] studied the same interaction with linear and nonlinear stability simulations and gave particular attention to the role of the incoming disturbance amplitude. The hump was found to distort the base flow strongly in its vicinity, with the chordwise velocity undergoing successive acceleration and deceleration, see fig. 2.11. The skin friction remained positive, so that no separation or chordwise reverse flow occurred and the changes in the instability could not be caused by a separation bubble. The defining feature of the modified base flow was a crossflow reversal just downstream of the hump apex. In fact, after a strengthening in the canonical direction of the crossflow component, enhanced by the favourable pressure gradient of the first half of the hump, the crossflow then reverses close to the wall over the short chordwise extent of the second half of the hump, to finally relax back to the base flow reference condition.




Figure 2.11: (a–e) Base flow components of chordwise velocity (U) (red dotted and black solid lines) and wall-normal velocity ($V \times 10$) (red dash-dotted and black dashed lines) in the hump (red) and reference (black) cases. (f) Spatial evolution of the cross-flow velocity in the hump case. Black dashed line indicates the δ_{99} . Grey dashed lines indicate loci of inflection points. The hump is indicated in orange. Image from Westerbeek et al. [38]

On this base flow the incoming instability was tracked through a sequence of destabilising and stabilising events, shown through the perturbation velocity in fig. 2.12. The perturbation's components are first amplified over the upstream part of the hump, then stabilised near and downstream of the apex within the region of crossflow reversal, where its shape develops a second, near-wall peak that follows the inflection points of the reversed crossflow profile. The effect persists well downstream of the hump, over roughly two hump lengths, even though the base flow itself has already recovered. Taken together, these events leave the instability with about 60% of the amplitude it would have reached in the reference case, in agreement with the transition delay measured in the experiments.




Figure 2.12: (a–f) Chordwise (solid line) and wall-normal (dashed line) perturbation velocity profiles over the hump (red) and reference case (black), at selected chordwise stations. For ease of visualisation, $|\hat{v}|$ is scaled by a factor 10. Horizontal dashed blue lines indicate the location of inflection points in the cross-flow profile of the base flow in the hump case. (g) Contour map of the amplitude function of the chordwise perturbation velocity from $|\hat{u}| = 0$ (white) to $|\hat{u}| = 0.5$ in 10 levels. Red solid and black dashed lines indicate the location of local maxima of $|\hat{u}|$ in the hump and reference case, respectively. Image from Westerbeek et al. [38]

The simulations also reproduced the strong amplitude dependence seen in the experiments. At a low incoming amplitude the interaction is essentially linear and the fundamental mode is stabilised, leaving the domain weaker than in the reference case, actually simulating a transition delay. As the incoming amplitude is raised, a localised destabilisation at the downstream end of the hump grows rapidly and pushes the instability above its reference level, generating the strong shears that favour the early breakdown observed in the high amplitude experiments. The conclusion of the analysis of Westerbeek et al. [38] attributed the stabilisation mainly to a linear mechanism, where the dominant route of extraction of energy from the base flow is weakened. This weakening is present where the crossflow vortex is deformed by the hump, but in the downstream of the hump the changes are in the perturbation shape rather than in the base flow that recovered. The instability enters the downstream region of the hump as a distorted, less unstable structure that grows less efficiently than an undisturbed crossflow vortex thanks to the spanwise shear effect applied by the hump. It can be concluded that the hump stabilises the flow by deforming the incoming instability rather than by altering the base flow it later encounters.

2.6. Final remarks and research motivation

The picture emerging from this chapter can be summarised in a few points. In swept-wing boundary layers at moderate to high sweep angles, the stationary crossflow instability is the dominant route to transition in the low-disturbance conditions of flight, and its strong sensitivity to the initial conditions makes it controllable. Smooth surface features are a recent addition to the available control techniques: the experiments of Rius-Vidales et al. [26] showed that a shallow hump placed in the favourable pressure gradient region can delay transition by about 14% of the chord when the incoming vortices are still in their linear stage, while the simulations of Westerbeek et al. [38] attributed this stabilisation to the deformation of the incoming instability rather than to a lasting change in the base flow.

This experimental evidence, however, rests on a limited hump geometry space, where the manufactur-

ing process consisted of casting epoxy add-ons, which were subsequently bonded to the model. With this approach, every new shape requires a new manufacturing and installation cycle, which makes a systematic exploration of the shape parameter space impractical. In fact, basic questions on the role of the height, width, skewness and chordwise position of the feature remain open. This is the gap the present work addresses: as stated in Section 1.2, the aim of this thesis is to facilitate further research on smooth surface features through the design, commissioning and assessment of a new setup capable of morphing such features on the same M3J swept-wing model. This aim translates into the following research questions, grouped into three themes that mirror the structure of Chapter 4.

Setup validation:

1. How accurately does the morphing setup reproduce the target surface shapes?
2. To what extent is the resulting shape affected by the flow conditions?
3. To what extent does the modified side of the wing model introduce a bias on the transition location with respect to the clean configuration?
4. How repeatable are the surface shapes and the resulting transition locations produced by the setup?

Effect of the surface features:

5. What is the effect of each geometric parameter of the hump – height, width-to-height ratio, skewness and chordwise location – on the transition-front location?
6. To what extent does the effect of the geometric parameters depend on the Reynolds number and the forcing conditions?
7. To what extent do humps with a negative height affect the transition front location?
8. To what extent do humps delay transition in an unforced CFI environment?

Optimisation:

9. Is the implementation of an automated optimization framework for the identification of transition-delaying surface shapes feasible with the proposed experimental setup?

3

Methodology

3.1. Introduction

This chapter discusses the development of the new setup as well as the tools, the facility and the post-processing methodology that led to the results. To help the reader, this chapter can be divided into the wing assembly (Section 3.2), the tunnel facility & measurement setup (Section 3.3) and the measurement procedures descriptions (Section 3.4).

3.2. Wing assembly

The wing model assembly is derived from an existing wing model, which was modified by integrating a morphing section. The morphing section consists of a silicone skin bonded to a metallic mechanical assembly housed within a cavity in the wing model. For the reader's convenience the following sections are presented in chronological order.

3.2.1. Skin

The development of the skin was the first task of this thesis work. The design inputs and constraints are expressed here: the position of the silicone skin in the wing profile, going from $x/c = 0.1$ to $x/c = 0.4$; the skin material, selected as a rubber silicone already mechanically characterized by TU Delft internally; the leading edge, trailing edge and span aluminium components; the number of legs was set to twenty, spaced 13 mm apart.

Skin design

Starting from the inputs described above, the development of the skin geometrical features was studied using a bi-dimensional Finite Element solver called *Abaqus*. The developed workflow consisted of a set of Python scripts used to generate the mesh, define the loading conditions, run the solver, and post-process the results. This enabled the implementation of an automated loop capable of testing numerous geometric configurations and identifying the optimal set of geometric parameters through a brute-force approach. This The scheme of the design workflow is reported in Fig.3.1. The load condition was imposed by defining an initial condition to all the mesh nodes, where the skin was "flat" without any stress, and the morphed condition, where a series of displacements and rotation was applied to specific nodes. This procedure reproduced the physical load case morphing the upper surface of the skin towards the target shape from a flat skin that is cast through the mould. As mentioned before, the simulations were two-dimensional so the three-dimensionality was neglected.

In this loop there are three different codes; the first one is responsible for building the two-dimensional geometry of the skin as it would be on a flat plate (step 0). This approach is realistic because it will be like this after the casting procedure. The geometrical parameters that the code can change are the *Skin Thickness*, *Leg Length*, *Leg Thickness* and *Leg Thickness Ratio* shown in Fig.3.1. The second code is responsible for applying the Boundary Conditions (load case) on the "flat" geometry. These conditions could be of two types: the vertical displacements on the legs' bottom part and the roto-translation of the




Figure 3.1: On the left the codes workflow and on the right the design parameters schematic

edge parts, both shown in red in Fig.3.1. It is important to note why the load conditions are chosen like the one described above: the direction of the displacements on the legs is only vertical because that is how the mechanical setup will handle the movement; the edge parts are angled so that the constraints are perpendicular to the wing's profile because that is how the aluminium strips will be positioned; and the reference system of the solver is not parallel to the wing's profile chord but is rotated 7.4 degrees as the real actuators' box.

The boundary condition code changes the legs height by applying a displacement from the bottom edge (see Fig. 3.1) with the aim of reproducing the target shape on the top surface of the skin. In this study the target shapes are Humps. Fig. 3.2 presents a pre-simulation plot showing the M3J shape (the name of the wing profile), the target curve and the positions of the legs in the "flat" state in addition to a solved deformed mesh.




Figure 3.2: On the left the Boundary Condition's code plot. The red stars are a section of the M3J shape, the green line the target curve with an hump (difference with the M3J shape) and the blue stars are the boundary conditions points (20 actuators, LE and TE strips). On the right is shown every node of the deformed mesh in output from the solver.

Once the simulation is run it is possible to use the third and last code that will take the deformed shape, visible on the right in Fig.3.2, find the upper nodes' curve and compare it to the original target curve. The metrics used to compare the two curves were the root mean square error, the maximum vertical error (in the *Abaqus* reference system) and the curvature root mean square error. The root mean square error is calculated using the difference of y-coordinates at the same x-coordinate of the two curves:

$$RMS_{error} = \sqrt{\frac{1}{N} \sum_{i=1}^N (y_{target,i} - y_{upper\ nodes,i})^2} \quad (3.1)$$

The result of the design process are presented in the section 4.1.1.

Skin mould manufacturing

During the skin design phase it proved useful to start the design of the parametric CAD model of the skin's mould. This model was built using the software *CatiaV5*. The final version of the mould is shown in Fig.3.3.

Some important points to note on the mould design:




Figure 3.3: CAD model of the skin mould

- The design of the mould was driven by the selected materials and manufacturing process, with a focus on cost and time. The majority of the parts are 2D plexiglass plates of different thickness stacked on each other. The only aluminium parts are the ones that remain attached to the silicone
- The skin is built upside down with the surface facing the flow facing the bottom plate of the mould
- The green volume is the silicone skin that will solidify after being poured in the void left by the mould's shape
- There are three bridges (white areas in the cross section) used to both straighten the leg supports' block in the spanwise direction and to act as a top restrain of the whole leg's assembly. These bridges need to be rigid and are 10 mm thick
- The silicone skin forms a C shape onto the leading and trailing edge parts to ensure a better adhesion
- The two thin black spacers on the spanwise side of the mould are an a posteriori design choice with the goal of increasing the skin thickness (see Sec.4.1.1)
- In the spanwise direction there are spacers designed to build a step on the skin. This step will be useful to seal the wing model cavity
- The roughness of the upper surface of the skin is ensured by the roughness of the bottom plate. By contrast, the continuity and tangency conditions at the LE and TE of the skin are ensured by the alignment of the aluminium parts (in the casting procedure)

The desired skin assembly is shown in Fig.3.4.

Casting procedure

After the manufacturing of the various parts of the mould the casting procedure needed to be tested and developed. This process was carried out using an iterative approach addressing the issues that arose one at a time.

The first test was carried out without the 21 legs supports, creating a thick sheet of silicone. This test showed the importance of an external clamping force on the mould's perimeter to prevent the liquid silicone from leaking from in between the mould layers. The most critical locations were the sharp




Figure 3.4: Final skin section from the CAD model

corners. Pictures of this test are shown in Fig. 3.6.

From the first cast, the cast skins resulted in skins progressively closer to the desired one. For brevity, the most critical aspects addressed and their solutions are listed here:

The continuity and tangency between the skin and the aluminium parts at the LE and TE were not easy to ensure, especially without having excess silicone slices on top of the aluminium, see Fig.3.5. This problem was caused by the design of the LE and TE aluminium parts, where their respective top and bottom surface were non-parallel to each other. This feature prevents the correct packing of the mould's layers leading to small leaks during the silicone pouring. The solution to this problem was achieved through some dedicated wedges that compensate for the lack of parallelism between the components thereby enabling the assembly to be tightly clamped.




Figure 3.5: On the left the problem of the silicone lip and on the right a close up on the fit of the complete assembly composed of the Inner wing box, LE strip and silicone skin. In the centre the CAD visualization of the non-parallel top and bottom surfaces of the LE aluminium strip.

The quality of the silicone material was a result of a homogeneous mix between the two components free of bubbles of air. This mixture was obtained by a dual step mixing procedure with an electrical mixer followed by a 10-minute degassing.

The geometry of the legs in the silicone without trapped air pockets and an even thickness was obtained by means of a dual step procedure; first the silicone was poured into the mould without presence of the so called legs' assembly, then a period of 5 minutes was allowed for the silicone to flow and flatten, after which the legs' assembly was sunk in place. Visual explanation in Fig. 3.7

Given the procedures listed above the silicone skin started to be good enough and the last adjustment made was towards a less waviness on the surface. For this reason the Skin Thickness (ST in the previous chapter) was increased from 3 to 3.5 millimetres using a 3D-printed spacer to lift the legs' assembly.

Finally it is time to give a step by step description of the casting procedure: After cleaning and making sure to have all the components and screws necessary for the mould, the first step is to build separately the legs' assembly and the mould. The legs' assembly is composed of 21 plexiglass bars, 20 L-shaped aluminium legs, 3 plexiglass bridges and the screws that keep everything together, see Fig 3.7. The mould is composed of several layers. In the LE or TE side the layers are: bottom plate, aluminium strip, plastic spacer (just 0.3 millimetres), first top layer and second top layer. On the spanwise sides the layers are: bottom plate, first spacer (wider and the height of the ST), second spacer (that creates




Figure 3.6: First test of the casting procedure. On the right the visualization of a leakage




Figure 3.7: Casting procedure steps. From left to right: the legs assembly, composed by the three transparent bridges on the top, the twenty aluminum legs and the twenty-one transparent legs supports; the plexiglass mould fixed on a wooden board with the LE and TE aluminum strips embedded; The first step of the casting allowing the silicone to level out; dipping step of the legs assembly in the pool of silicone towards a complete skin.

a lip on the sides), First top plate, second top plate and side spacers (black 3D-printed to increase the height by 0.5 millimetres). It is important to note that this assembly was clamped together in several locations by screwing wooden blocks into the bottom wooden plank, see Fig 3.7.

After the preparation of the assemblies described above, the mixing procedure can begin. Before mixing, however, a primer must be applied to the aluminium surfaces that will be permanently bonded to the silicone, namely the surfaces of the L-shaped legs and the leading-edge (LE) and trailing-edge (TE) strips. In this step the silicone A and B components are weighed and mixed with a black dye. In the final skin the total silicone weight was close to 2 kilograms but split in 6 smaller batches for an easier pouring, mixing and degassing. The mixing was first done with an electrical helical mixer, then the mixture was scraped manually from all the container's walls and finally mixed again electrically. The mixed batch was then degassed (10 minutes) and poured into the mould.

To ensure a flat surface and an evenly spread silicone layer a period of 5 minutes was allowed before sinking the legs' assembly inside the pool of silicone. Once the assembly was in place the last steps were only to fix it to the mould with screws and pins and to check the perimeter for large air pockets, to be burst with a needle.

Final skin assembly

The final assembly of the silicone skin to the mechanical structure is achieved by means of screws. The screws fasten the L-shaped aluminum legs bonded to the silicone skin to the corresponding L-shaped actuator legs, as shown in Fig. 3.8. It is important to note the screw-tightening procedure adopted during the assembly. The screws are tightened one row at a time, starting from the trailing edge (TE).



Figure 3.8: Assembly of the skin and the mechanical actuators using an experimental green skin composed of only five legs.

If the skin is mounted with the wing model in the upright position, the weight of the skin makes this procedure difficult to control and may introduce spanwise shear stress in the silicone, as discussed in Section 4.1.3.

3.2.2. Mechanical assembly

The mechanical assembly that generates the movement of the skin was commissioned from an external company. This assembly is housed inside the wing and, as mentioned before, will be screwed to the L-shaped aluminium strips bonded to the skin, see the top L-shaped strip in Fig. 3.8. This assembly consists of a metallic box that serves as a support structure, ensuring the correct alignment of the 20 actuator units mounted on it.




Figure 3.9: Axonometric CAD view of one actuator unit

These units consist of a bottom machined metallic frame that ensure rigidity to the unit, a rod to transmit the horizontal movement, a linear DC-Servomotor that moves the rod, a 3D-printed guide, two carriages, that supports the top metallic frame, moved by the rod, some Teflon guides for the vertical movement of the top frame and a pair of springs on the sides of the unit that keep the top frame in contact with the carriages, see Fig. 3.9 and Fig. 3.14.

In more detail, the movement is produced by linear DC-Servomotors *LM 1247-020-11* by *Faulhaber*, these servomotors are equipped with analogue Hall sensors for the estimation of the movement of the rod. These servomotors are controlled by motion controllers *MC 3001 B ET* also from *Faulhaber*. The PCB containing these is located on the back side of the metallic box supporting the actuator units, see Fig. 3.10.

The motion of the L-shaped strip is generated by the servomotors as a horizontal displacement and is subsequently converted into a vertical displacement through two wedge-shaped profiles machined into the upper metallic frame. The slope of the wedge determines the accuracy and force that the unit can deliver. The sizing of the units was carried out by the external company and the testing confirmed the calculations to be correct.

The software to move the actuators assembly was also developed by the external company. The complete data flow required to move the actuators is illustrated in Fig. 3.11.




Figure 3.10: PCB on the back side of the metallic box of the mechanical assembly






Figure 3.11: Movement chain schematics

The movement of the motors is controlled by the motor controllers which communicate via CAN to a *Raspberry Pi* located outside the wing itself. This *Raspberry Pi* could be accessed through an Ethernet LAN connection and operated through its local web interface via the device IP address. The web interface has to be opened by an external PC and, most importantly, could be linked to a Python code. The data available to the final user during operation are the positions requested from the user, the actual positions read from the Hall sensors and the power consumption drawn from each motor (expressed as percentage of the maximum nominal continuous power in the motor's datasheet). During the initial setup of the motors, called Homing, additional setting are available; an example is the adjustable velocity of each motor in this phase. The movement of the motors is then controlled in position and the correct position is not reached only in case of errors in the Hall sensor measurements or an excessive drain of power.



(a) Raspberry Pi used in the test campaign



(b) Measurement setup for the validation of the mechanical assembly actuators position

Figure 3.12: Mechanical assembly components

Given the working principles and data shared by the mechanical setup, it is of primary importance to evaluate the error in the Hall sensors' readings. To do so an external measurement of the position of the actuators needed to be taken. This measurement setup was built with a *Zaber* traverse system and a *Microepsilon scanControl 30x0-50* blue light model, see Fig. 3.12b. In this type of measurement, the *Microepsilon laser* measures the relative distance between the actuators' top surface and the sensor

in a line of approximately 40 millimetres. Perpendicular to this line, the Zaber traversing system will move the line itself so that, in a sweeping motion, a 3D point cloud of the entire top surface can be reconstructed. Measuring the top surface at different actuator positions (heights) makes it possible to validate the effective displacement of the actuators against the displacement read by the Hall sensors.




Figure 3.13: Displacement validation. On the left, the measured top surface height of 19 of 20 actuators in 8 positions. On the right, the error between the reference position and the target position (hall sensor position)

In Fig. 3.13 the results of the validation can be seen. In this study the height of the actuators has been varied from -4 millimetres to 3 millimetres in 1 millimetres steps.

The measurements show that a 1 mm step is reproduced with a relative error lower than 10% of the step and that with higher steps the absolute error is lower than 0.3 mm (relative error of 4%). This result is sufficient to consider the Hall sensor error negligible compared with the errors arising from the top skin displacements. In addition, the graph shows that the zero position of the actuators is not flat; this is not an error but the a design choice. In fact, the support box was designed to facilitate the M3J shape once installed in the wing.




Figure 3.14: Actuators assembly mechanism

3.2.3. M3J wing model and adaptation

In this subsection the wing model and its adaptation to accommodate the mechanical assembly are discussed.




Figure 3.15: Comparison between the NACA66018 and the 66018M3J from Serpieri et al. [34]

The M3J wing model features a streamwise chord of $cX = 1.27$ m, a span of $b = 1.25$ m and a sweep angle of 45 deg. It was designed by Serpieri and Kotsonis [34] to study the transition induced by CF

vortices while limiting other sources of destabilization as much as possible (e.g. Tollmien-Schlichting (TS) waves, attachment line flow or Görtler vortices) and has been a valid test bench for CFI since then. In Fig. 3.15 the shape of the M3J cross-sectional profile is shown. For more detailed information about the wing model the reader is referred to [34] or [27].

The adaptation of the wing model consisted of making space in the hollow wing model for the mechan-




Figure 3.16: Process of the M3J wing adaptation. On the left the starting Hole. In the middle the jig and the metallic supports glued to the model. On the right the image before the assembly lay down, with the three thermocouples and the engraved bottom surface

ical assembly, positioning it in the right 3D location and orientation and, ensuring a structural rigidity of the wing model comparable to the original one.

This procedure started by cutting a large hole on the wing model, which entailed removing a large part of the central spar, see Fig. 3.16.

The next step was to reinforce the cutout region, restoring the structural stiffness of the wing. This was done using sandwich panels made with fibreglass fibres; the manufacturing process was a wet layup. After the panels had hardened outside the wing, they were cut and joined to the wing model through other fibreglass sheets soaked in resin.

Once the wing had a closed inner box, as can be seen in Fig. 3.16, the next step involved gluing the



Figure 3.17: Wing model with the mechanical assembly on display and with the skin. Note the possibility of extracting the mechanical assembly through handles

fixed part of the mechanical assembly to the model's inner box. A jig was developed to glue the mechanical assembly component in correct position with respect to the wing. Once the fixed components had been bonded, the inner cavity was prepared for closure by the mechanical assembly. This step re-

quired the removal of material from the cavity bottom, the installation of four thermocouples at different locations inside the enclosure, and the drilling of holes in the side panels for cable routing and venting. The outcome of these modifications is shown in the last two images of Fig. 3.16. The thermocouples are the blue wires that terminate in the top right, on the old spar, top left, on the aluminium LE part and in the bottom. The fourth thermocouple was placed on the bottom side of the mechanical assembly close to the PCBs. The resulting adapted model could host the mechanical assembly and most importantly the skin on top of it, see Fig. 3.17. This design enables the extraction of the mechanical assembly through handles, to easily perform checks or fix errors. The usage of the handles is suboptimal when the skin is fitted and, it is recommended that a new S-shaped handles to be manufactured in the future. The redesign of the handles is required due to the tight tolerances that have been achieved on the LE and TE strips; indeed even a slight bend on this components will be felt as a step once the assembly is fitted in place.

Finally, the last steps were to make sure the LE and TE transitions between the different components were as smooth as possible and to fill the screw holes with clay. The former was achieved first by filling the bigger holes with resin and then by an iterative process of filling and sanding, see Fig. 3.18. The latter was done immediately after screwing the final skin on the model directly in the test section.




Figure 3.18: On the left, the surface finish achieved at the LE between the wing model, the fixed metal strip and the aluminium strip of the silicone skin. On the right, a close-up picture of the LE of the model in the WT test section; the golden dots are the DREs.

3.3. Tunnel facility and measurement setup

The low-turbulence environment is an essential element for the development of stationary CF vortices, as described by Saric [31]. The tunnel is a closed-loop subsonic wind tunnel, capable of reaching speeds up to 117 m/s (421 km/h) by means of its six-bladed fan driven by a 525 kW DC motor. The turning vanes, which curve the flow in the corners of the tunnel, avoiding separation, are equipped with a cooling system to control the flow's properties. The test section is octagonal, and it is 1.80 m wide,




Figure 3.19: Schematic drawing of the Low Turbulence Wind Tunnel. The drawing is taken from Rius-Vidales [27] where a more detailed description can be found.

1.25 m high and 2.60 m long, with a slight divergence of the vertical walls in the streamwise direction to account for the solid blockage effect caused by the BL thickness. Due to the large contraction ratio in the convergent section of 17.9 and the capability of deploying multiple active anti-turbulence screens, the freestream turbulence intensity at the velocity of interest in this campaign is $Tu \leq 0.03\%$, which provides sufficient conditions for the investigation of stationary CF instability. A detailed description of the turbulence intensity study can be found in Serpieri and Kotsonis [35], and a detailed description of the facility can be read in [27].




Figure 3.20: Measurement chain schematics. Wind tunnel section room: control area on the left, wind tunnel section on the right, passage under the WT stairs on the top. Flow direction bottom-up.

The measurement chain of the experimental campaign was based on three main tools: the mechanical setup, the infrared cameras and the measurement laser. In addition to these three main data sources, the pressure data from the pressure port embedded in the wing model were also saved. The experi-

mental setup is represented in Fig. 3.20, where a schematic of the test section room is shown, whereas the WT facility is drawn in Fig. 3.19.

The IR thermography measurements were conducted on the pressure side of the model using IR cameras (*PI640, Optris GmbH*), whose focal-plane-array (FPA) uncooled sensor has a resolution of 640×480 px with a $17 \mu\text{m}$ pixel pitch, is sensitive over the $7.513 \mu\text{m}$ spectral range and features a noise-equivalent temperature difference (NETD) of 75 mK. A wide-angle lens of focal length $f = 10.5$ mm is mounted on the camera, yielding a full-field view that captures the entire model surface in a single frame. The pressure side is imaged assuming a surface emissivity of 0.94, and a typical acquisition consists of approximately 80 thermal images recorded at a sampling rate of $f_s \approx 3.5$ Hz.

A measurable thermal contrast between the laminar and turbulent regions is obtained by actively heating the model with four halogen lamps of 400 W. The lamps are positioned outside the test section and distributed along the chord, so as to minimise chordwise and spanwise variations of the surface temperature, and irradiate the model until thermal equilibrium is reached.

The transition front is investigated on the pressure side of both faces of the model. At an AoA of $\pm 3^\circ$ only the windward face acts as the pressure side, while the opposite suction side is fully turbulent, so each measurement is run separately by reversing the sign of the angle of attack. To this end, a dedicated IR camera and a set of halogen lamps are arranged on each side of the test section, moving some of these lamps between the two arrangements if required by the low thermal contrast of the IR images. The setup image can be seen in Fig. 3.21.




Figure 3.21: Silicone side of the wind tunnel test section during the campaign. Equipment outside the test section: IR camera 1, halogen lamps, Initium and laser scanner. Equipment inside the test section: wing model and silicone skin delimited by the aluminium strips.

The laser used to scan the silicone surface of the wing model through the test section glass was a *Microepsilon scanControl LLT30x0-600* red-light model. This model is built as a precise profile measurement platform for industrial tasks, with a resolution of 2,048 points, a profile frequency of 10,000 Hz and a measuring range up to 1050 mm. For additional information, see the *Microepsilon* website.

The Initium is the pressure transducer used in the Low Turbulence Wind Tunnel. It converts the pressure values into analogue voltage data.

The power supply of the mechanical assembly was an *EA-PS 9080-60 T* from *Elektro-Automatik*, capable of delivering DC output voltages between 0 and 500 V, output currents between 0 and 60 A and output power ratings between 0 and 1500 W. The output voltage needed by the mechanical assembly is 24 V, with less than 3 A at steady state, but the peak current measured during homing reached more than 10 A.

To condition the wavelength of the crossflow instabilities, some Discrete Roughness Elements, or DREs, are placed near the LE of the wing model (see Fig. 3.18). These elements are small circular disks cut from adhesive transfer vinyl film using a vinyl cutter¹. As well documented in the literature, the shape

¹The vinyl cutter model used is a *Jaguar V61 LX*, with an accuracy of ± 0.254 mm or $\pm 0.1\%$ on the thickness.

of these elements can be used to control the CFI: the height of the disks can increase or decrease the amplitude of the developing CFI; similarly, they can be placed downstream of the neutral point of the wing (positioned at $x_c = 0.031$), where the amplitude is higher the closer they are to the neutral point; and finally, they can control the spanwise wavelength of the CFI.

In this experimental campaign, the nominal dimensions of the disks were kept constant, with a diameter of $d_D = 2$ mm and a thickness of $k_d = 25\mu\text{m}$. The spanwise spacing, λ , and the location of the disks with respect to the neutral point were varied between tests.

3.4. Measurement procedures

This section provides a detailed description of the methods and procedures that led to the presented results. It begins by describing the infrared thermography methodology and the processing of the laser measurement data, followed by the data acquisition procedure and the preliminary optimization loop.

3.4.1. Infrared thermography

Infrared (IR) thermography is the main technique employed in this work to determine the location of the laminar–turbulent transition front on the wing model. The method is non-intrusive, since the measurement relies solely on the thermal radiation emitted by the model surface and therefore introduces no probes or surface alterations, making it an appropriate technique to study the boundary-layer state. Its robustness for the study of swept-wing transition is well established, and it has been applied extensively to crossflow-dominated boundary layers on the same wing model [27].

The detection of the transition location rests on the Reynolds analogy, which relates the skin-friction coefficient C_f to the convective heat-transfer coefficient h_c at a given Prandtl number:

$$h_c = \frac{1}{2} C_f \rho c_p U_\infty Pr^{-2/3}, \quad (3.2)$$

where ρ is the air density, c_p the specific heat at constant pressure, U_∞ the freestream velocity and Pr the Prandtl number. Equation 3.2 implies that an increase in wall shear stress, such as that produced by a turbulent boundary layer, is accompanied by an increase in surface heat transfer. When the model is actively heated above the freestream temperature, the turbulent region therefore exchanges heat with the flow more effectively and settles at a lower surface temperature, whereas the laminar region remains warmer. The resulting temperature difference between the two regions is exploited to locate the transition front.

The thermal contrast required for this purpose is established by actively heating the model above the freestream temperature with the halogen lamps described in Section 3.3, the measurements being taken once thermal equilibrium has been reached so that the surface temperature varies minimally during an acquisition. It should be noted that uniform wall heating is known to be destabilising for boundary layers dominated by Tollmien–Schlichting (TS) waves. In flows dominated by crossflow (CF) instability, however, the destabilising effect is considerably milder, and a change in wall temperature mainly influences the travelling CF modes. Lemarechal et al.[16] reported that, for a ratio of model-to-fluid temperature $T_m/T_f \leq 1.04$, no appreciable shift of the transition location was observed on a swept wing dominated by stationary CF modes. Some tests were conducted with and without external heating, detecting no difference in the transition front location.




Figure 3.22: Dewarping procedure on the clean side of the wing model. On the left, the picture of the paper target sheet positioned on the wing model. In the middle, the original IR image taken with the IR camera. On the right, the final dewarped image with the non-dimensional chordwise coordinate system superimposed.

The thermal maps are acquired in the camera image plane and must be transformed into physical coordinates before the transition front can be located; to this end a geometrical transformation, referred to as dewarping, is applied. A calibration target consisting of a rectilinear grid of black dots printed on a flexible white paper sheet is first placed on the model surface and imaged under the same illumination as the measurements. This illumination heats and highlights the black dots with respect to the white background. The known dot diameter of 10 mm and spacing of 100 mm provide a set of reference

points from which a polynomial distortion correction is computed, accounting for the curvature of the airfoil and the position of the camera relative to the surface. The same transformation is then applied to all the thermal maps of a given configuration, assuming the location of the IR camera to be fixed during the whole campaign. In addition, the dewarping converts the image coordinates from pixels to millimetres and places the origin of the coordinate system at the mid-span of the leading edge, so that the front can be tracked with respect to a physically meaningful reference. The procedure steps are shown in Fig. 3.22, where the regular arrangement of the white dots in the last image confirms the accuracy of the dewarping process.

The transition front is then extracted following the differential infrared thermography (DIT) technique, which was introduced for transition detection by Raffel and Merz [19] and adopted on the present model by Rius-Vidales [27]. The technique consists of: first, averaging in time each set of infrared images acquired at a fixed Reynolds number, to reduce the random measurement noise; the DIT itself consists of subtracting two such time-averaged maps obtained at two different Reynolds numbers, chosen so that the corresponding transition locations are clearly separated. Because both maps are recorded under identical settings (optical and illumination arrangement), the stationary background pattern is common and is largely cancelled by the subtraction; the differential map therefore isolates the band swept by the front between the two conditions, with a markedly improved signal-to-noise ratio and a reduced likelihood of a spurious identification. The front is located within this differential map by a sequence of image-processing operations. A global intensity threshold is first determined from the histogram of the differential map by means of Otsu's method, which selects the threshold that minimises the intra-class variance of the two intensity populations; the bimodal character of the histogram is essential for this step, as it is the separation between the two populations (the two hills of the histogram) that determines the binarisation. Once the map is binarised, the intensity gradient of the binarised map is computed and its peaks are interpreted as the front locations, the downstream peak corresponding to the lower Reynolds number and the upstream peak to the higher one. A straight line and a confidence gap are finally fitted to the identified peaks along the span, and the reference transition position is extracted at the spanwise mid-domain location. Some steps of the procedure in addition to the final front location are shown in Fig. 3.23; these images are reporting data based on the results of the $\lambda = 7.5$ mm DREs configuration on the clean side of the wing model for the highest Reynolds case and a lower one used for the subtraction. The images show the transition point locations (green), the transition line (purple, dashed), and the confidence bounds (purple, solid). The blue rectangle indicates the region of interest, where the image subtraction is performed. The light blue dashed lines represent the maximum forward and backward limits, while the large green dot marks the centroid of the clear region identified by the subtraction process.




Figure 3.23: IR DIT procedure. The flow goes from right to left and the axis are in mm. In the first picture there is the dewarped image, in the second one the subtraction between the first image and a lower Reynolds one, and finally there is the result of the binarisation. The Reynolds numbers of the image and of the subtracted one are $Re = 2.81 \cdot 10^6$ and $Re = 2.42 \cdot 10^6$, respectively.

The automatic identification described above is not infallible. Its reliability rests on the bimodality of the differential histogram, which can be degraded by a low thermal contrast; in such cases the Otsu threshold and the subsequent gradient step may either fail or yield a spurious front. To retain the affected measurements, an arbitrary coefficient to increase the thermal contrast based on the Reynolds number and a manual selection of the transition front were implemented as a final fallback, whenever the automatic procedure did not provide a satisfactory result.

The uncertainty of the transition location is governed mainly by systematic errors. The principal contributions arise from turbulent wedges generated by imperfectly assembled components or by particles

deposited on the wing surface, from a non-uniform heating of the model prior to acquisition and from a misalignment of the calibration target during dewarping; the first two degrade the quality of the differential map, whereas the last introduces a geometrical bias in the mapped coordinates. A rigorous quantification of these contributions would require a large number of repeated tests under varied conditions, such as repeated calibrations, which could not be carried out within the time available for the present work. The uncertainty reported for the transition location is taken as the standard deviation of the error of the linear fit of the front. As this quantity is larger than the NETD of the camera, it is the dispersion of the fit rather than the sensor resolution that governs the overall uncertainty.

3.4.2. Laser data handling

It should first be noted that the first use of the laser coincided with the first day of the WT slot; for this reason, the know-how presented was built during the experiments.

The processing of the laser data can be divided into two parts: the processing performed during the test campaign, using the laser as live feedback on the skin surface, and the post-processing performed as a check of the surface shape once the data had been acquired.





Figure 3.24: On the left the laser Scanner setup shooting the red beam trough the glass onto the wing model. On the right a single snapshot of raw data in its coordinate system (without equal axis)

The live laser data visualisation was obtained by means of an SDK extension from *Microepsilon* and a Python script accessing the raw data from the laser in real time. These data were the z-x coordinates of the detected points on the red laser line projected onto the model surface (see Fig. 3.24). As can be seen, the raw data were noisy, with vertical oscillations of the order of 0.3 mm on very smooth surfaces such as the old M3J model surface as well as on other surfaces like a flat metallic ruler. The same noisy behaviour could also be observed at the same location over time, where a single point could oscillate between different positions. This behaviour is caused by the inner mechanics of the laser, which will be discussed later.

Because the live visualisation of the skin was crucial to the campaign, the code had to overcome these problems. The flow chart of the code implemented can be seen in Fig. 3.25. The first line of the flow chart is the initialisation; in this phase the code sets up the target curve based on the model angle of attack (the M3J section), sets the calibration offsets, and initialises the communication with the laser and the figures. The calibration offsets are the coefficients of the *roto-translation* needed to fit the raw data to the M3J target curve, which is set with the chord parallel to the x-axis starting at (0, 0). These coefficients are the X and Z translations and an angle of rotation. To evaluate these numbers, the root mean square error between the raw data points and the target shape was minimised. This procedure was improved with the weighted RMS error minimisation, in which more accuracy was imposed on the LE region, where the higher curvature ensured better results.

Once the initialisation was completed, the live loop started to show the laser data overlaid on the M3J shape with a refresh rate close to 1 Hz. In addition, the capability of pushing a button and saving n snapshots of the raw data was implemented in the loop. The solution proposed for the noise problem is a polynomial fit of the raw points. In particular, the fitting was performed only in the silicone area and the order of the polynomial is high (25–30). This is done because the aim of the laser data is to reconstruct the surface of the skin, and a polynomial fit ensured a more physical approach than a sliding average.




Figure 3.25: Live laser scan visualization code flow chart

This approach yielded better results during the experiments, increasing the sensitivity to the actuators' displacement (even 0.002 mm was noticeable) as well as a qualitatively smoother surface to the touch (once the M3J preset had been done). By means of this fitting, the skin shape could be seen in real time during the campaign, providing crucial feedback.




Figure 3.26: Laser raw data of 1000 consecutive snapshots in the same spot with no air flow. Full field and zoom view (the red rectangle in the full view represent the zoom)

It should be noted that the post-processing of the laser data has been done with little time available and can be extended to yield more information. In post-processing, the temporal noise problem became evident after the acquisition of one thousand consecutive snapshots (see Fig. 3.26). The raw data oscillate in time between different shapes, even some millimetres apart. The causes of this behaviour are still not clear, but some reasonable guesses are small vibrations in the system that became critical due to the wear of the sensor's glass, or an interference phenomenon. To understand the possible cause of the noise, it is first useful to understand the laser measurement mechanics. The laser measures the distance as a triangulation of a laser beam shot from one side of the tool and collected again after the reflection caused by the surface. A visual explanation is presented in Fig. 3.27.

The wear of the sensor's glass can be the cause of the noisiness because of the variable refractions that it yields on the laser sensor in the presence of micro-vibrations. This effect can change the detected location of the laser on the sensor plane based on the glass defects, and the resulting distance no longer matches the calibration of the tool. These defects were clearly visible on the laser tool. The interference phenomenon is instead called *speckle* and is the manifestation of the constructive or destructive interference of the red laser light reflecting from the model surface. The reflections on the model are provided by the micro-roughness on the surface that reflects the incoming laser light with different phases. Since the light source is coherent, all the reflected light beams have the same wavenumber but different phases, so an interference pattern can arise. If the fringes on the laser sensor are larger than the pixel size, this will artificially change the detected position of the surface. Because of this phenomenon, the surface position is highly susceptible to slight changes in the air density or micro-vibrations of the setup. Note that in the data collected during the campaign this behaviour appears as two or more *families* of parallel curves.




Figure 3.27: Orthographic projections of the laser measurement tool. The laser beam is shot from the left side of the tool and the light sensor is on the right part of it. Based on the position of the reflected light on the sensor itself, the distance of the reflecting point can be calculated.




Figure 3.28: Laser post processing code flow chart

The logical flow chart of the post-processing is shown in Fig. 3.28 and, as can be seen, it is composed of different steps. The first step of the process applies a filter to the raw data in order to be able to do, in the second step, a spatial bin averaging of the different raw profiles captured. The filtering action aims to delete from the raw data the profiles with reflections, detecting them with a gradient-based filter, and the different *families* of curves described above, selecting only profiles close to the most recurrent one. A visualisation of this filtering action can be seen in Fig. 3.29, where the red dots' curve is discarded. It is important to note that the noise amplitude of the selected data is higher than in the data acquired during the parametric study, as shown in Fig. 3.30.

Once the first two steps are done, the data of a single geometry are expressed as a single profile in the raw data coordinate system. The next step is to create the correct target curve, which means adding the correct hump to the M3J profile if needed, and to match the mean profile with the target curve through the best roto-translation. This procedure is done by minimising the root mean square error between the two profiles, similarly to the live laser data visualisation, but improving the accuracy of the result by adding more steps (for example a more suitable minimisation method named Nelder-Mead refinement). Once the data are in the target coordinate system, they can be analysed and plotted to show the difference between the skin surface shape and the target one.

The error estimation of the laser data has not been done as an independent analysis, but some information about it can be assessed from the output of the filtering action. In particular, the standard




Figure 3.29: Bimodal filter visualisation. The raw data sample is composed of the red and black dots.




Figure 3.30: Error assessment through the standard deviation of the filtering output. On the left, the full view of a scan where the hump shape can be noted on the right-hand part of the picture. On the right, a zoom of the full view of a different sample of data, where three different curves are visible, as well as the final standard deviation of the black raw points.

deviation of the filtered data can be used to calculate the magnitude of the errors, as no other external validation measurements are available. The standard deviation of the laser data measured together with the IR images was lower than 0.15 mm; Fig. 3.30 presents two representative cases, each shown either a full or zoomed view, to illustrate the measurement noise level and the magnitude of the corresponding standard deviation. In the first case the full shape is visible and the writer assures that the standard deviation is close to the latter image. The resulting information on the error estimation is that a first statistical estimate predicts that the measurement is within a ± 0.15 mm range from the plotted result, ensuring an error on the hump height lower than 15% (with humps higher than 1 mm), which is acceptable.

3.4.3. Setup data acquisition procedure

Because of the complexity of the setup, the campaign data acquisition developed as a multistep procedure. The steps into which it is divided are: the homing procedure, the preset making, the wind tunnel initialisation, the wind tunnel run, the IR image capturing and the laser measurements.

The homing is the first action to do in order to control the mechanical assembly, which implies morphing the silicone skin surface. This process aims to initialise each Hall sensor in every actuator assembly in order to locate the "zero" position. This procedure works by pushing every motor as much as possible (150% or 200% of the nominal power of the motor) in one direction until the movement stops; at that point the Hall sensor knows it is at the maximum location and, going backwards by a precise amount, can position itself at the "zero" location. As can be easily understood, the problems of this method are the correct calibration of the precise amount to find the "zero" and the false detections of the dead stops. The calibration is not a problem, as it has been done and tested by the external company and validated by some tests outside the wing. The erroneous dead stop problem is instead of primary importance; in fact, a bigger friction or a broken chunk could get — and actually got — the actuator's carriage stuck in a location different from the limit. In this campaign, various methods, such as cleaning every sliding part and making sure not to bend any component, were tried but, contrary to the tests on a flat surface, as soon as the complete skin was installed this problem persisted.

The solution to this unreliable homing procedure was found through the preset procedure. The preset is the name of the set of actuator positions that makes the skin follow the M3J shape; these values are calculated after the homing and, because of this, are not susceptible to the homing unreliability. In this way, the reliability of matching the M3J shape is shifted from the Hall sensor readings (and correct homing) to the laser measurements. The evaluation of the preset values is an iterative procedure where

the M3J target is matched manually thanks to the live laser feedback. Being an iterative procedure, a fixed standard termination of the procedure has to be set; this criterion is that the maximum error of the skin surface with respect to the M3J target must be lower than ± 0.15 mm along the entire skin length. This step has to be done just once at every power-up of the system and could be automated in the future, but in this campaign it could last as much as 45 minutes.

The wind tunnel initialisation involves all of the wind tunnel tools, such as the wind fan, the pressure transducer, the motor controlling the AoA, the IR cameras and the halogen lamps. Once this step was completed, the wind tunnel could be used to acquire data on different skin-generated geometries. In this campaign multiple Reynolds runs had to be run on the same geometry, so the measurement procedure was designed to automate this.

The wind tunnel runs were set to automatically change between five Reynolds numbers every ~ 40 seconds; in this time interval the measurements were made. As soon as the Reynolds (the velocity) stabilised, the WT software automatically captured ~ 80 IR images, the WT environment values such as the velocity and the temperature, and the pressure port values; then the WT operator had just to switch off the halogen lamps for a few seconds and trigger the laser profile measurements, which saved 50 scans.

The whole procedure has to be repeated for each geometry, and once the setup is initialised, it can achieve a turnaround time between different geometries of 6 minutes (in a 5 Reynolds study), which is an incremental gain of time compared to the previous methods.

3.4.4. Preliminary optimization loop

The last challenge of this test campaign was to show that the setup could be used as a live optimiser. This meant that the setup had to be able to do three things without human intervention: taking and saving some images from the IR camera, processing the images and computing a reward function, and moving the actuators to the next position to test.

To prove this capability of the setup, a PSO (Particle Swarm Optimisation) loop was implemented. The choice of the optimisation algorithm was based both on the physical setup (20 degrees of freedom, one per actuator) and on the simplicity of the implementation. This algorithm is well known among optimisation methods because of its simplicity. In fact, the mechanism of the model consists of a swarm of individual particles (much like a swarm of birds) that can explore the solution space. In this approach, at each iteration an individual (bird) knows three fundamental quantities: its last direction (V_i^t), its best guess ($P_{Best,i}$) and the best guess of the swarm ($P_{BestGlobal}$). From these three pieces of information the next position is calculated using Eq. 3.3:

$$V_i^{t+1} = wV_i^t + c_1r_1(P_{Best,i} - P_i^t) + c_2r_2(P_{BestGlobal} - P_i^t) \quad (3.3)$$

where the r_1 and r_2 coefficients are randomly generated numbers in $[0, 1)$, calculated at every evaluation of Eq. 3.3. The V are velocity vectors, which represent the difference between two positions (e.g. $V^{t+1} = P^{t+1} - P^t$), and the P are the position vectors, which are the vertical displacements of the 20 actuators.

As can be seen from Eq. 3.3, the exploration of the solution space is driven by the w, c_1, c_2 coefficients. The w coefficient is the inertial component, keeping the direction of the previous iteration, which tries to avoid falling into local minima of the solution. The c_1 coefficient is the cognitive (personal) component, storing in memory the best actuators' position of the individual particle and steering the exploration towards it. Finally, the c_2 coefficient is the social (global) component, saving in memory the best actuators' position of the whole swarm across different iterations and steering the direction towards it. The choice of the values of these coefficients is well described in the literature and is crucial for the convergence of the algorithm. Because the aim of the experiment is not the result of the optimisation loop but the demonstration of the setup capability, no detailed study on the influence of the parameters has been carried out.

After the brief explanation of the PSO algorithm, the code flow chart can be presented (see Fig. 3.31). In the code there are two nested loops: one for the PSO iterations (outer loop) and one for each particle of the swarm (inner loop). In this way, at each inner loop iteration the position and velocity of a single particle are updated, and at each outer loop iteration the global best and convergence plot are updated as well.




Figure 3.31: PSO Code Flow chart

The implementation of an algorithm with steps involving random coefficients must include some checks to prevent the system from drifting or breaking. In this case there were limits both on the actuators' position and on the maximum torque read by the *Raspberry Pi*. In particular, there were limits on the maximum top and bottom displacements tailored to each actuator, the maximum gradient between two neighbouring actuators and the maximum sum of the actuators (to prevent an all-out or all-in situation that could draw too much current). These checks are part of the feasibility check section of the code.

One important point that deserves further discussion is the reward function. In these tests the aim of the optimiser is to change the positions of the twenty actuators in order to delay the transition line as much as possible. The transition location is detectable through the temperature differences on the wing model surface, which depend on the flow-field state. In fact, in the turbulent regions the flow cools the wing surface more, whereas in the laminar regions the surface temperature is higher. Based on




Figure 3.32: PSO reward function evaluation procedure. From left to right: the original raw image, the cropped raw image showing the location of the reference pixel, and the normalized image used as input to the reward function. The reward function is computed by summing the values of all non-NaN pixels in the normalized image.

this principle, the IR images displaying the surface temperature were cropped to a specific region (see Fig. 3.32) and, after averaging over n snapshots ($n = 25$ here), the pixel temperature values were normalised by a reference pixel at the LE (located in an always laminar region) and, finally, the reward value was the sum of the pixel values. The higher the sum, the larger the surface area in the laminar state. Despite the simplicity of the reward function, it is able to lead the algorithm to move the transition line downstream. The limitations given by its simplicity were several, such as the increase of the

sum value over time, or the fact that a tripping hump in the downstream part of the skin had a higher reward than a clean M3J surface, which was a crucial aspect in the starting iterations. An alternative reward function approach could be to detect the laminar and turbulent zones through a binarisation of the images, thereby identifying a transition region.

During the optimisation iterations, the code performs two tasks: displaying the optimisation progress and saving the iteration data. The progress is shown through six images, of which three regard the external iterations and three the internal ones (see Fig. 3.33). In particular, the top-left image tracks




Figure 3.33: Example of a PSO progress visualization plot at its last iteration.

the best reward function value among the individuals of each external iteration; its rising characteristic implies a progressive delay of the transition front. The top-right image is instead the normalised field of the best IR image, and its respective actuator shape is shown in the middle image. The bottom row is important to update the user on the optimisation loop state, as the histograms represent the individuals' reward function values calculated in each internal iteration, while the actuator disposition (in the middle) and the normalised IR image describe the physical state of the last measurement, beyond the little information carried by the reward function.

The saved data of the optimisation loop are of two types: the mean of the IR images captured in each internal iteration and a log .csv file. This log file contains the information of the internal and external iterations about the reward function value, the actuators' disposition and the reference point location (even if it is constant).

4

Results

The aim of this chapter is to present the processing and analysis of the data acquired during the first experimental campaign using the new setup. Although the primary objective of the campaign was the commissioning and validation of the experimental apparatus, part of the collected dataset already provides useful observations for the research on surface humps.

The chapter is organised into several sections. It begins with the characterisation and assessment of the experimental setup and progressively transitions to the presentation and discussion of the first research results obtained with the system, ending with the description of the preliminary optimisation loop runs.

4.1. Silicone skin characterization

In this section the silicone skin is fully described, beginning with the choice of the skin geometrical parameters and continuing with the first tests addressing the behaviour of the skin in the wind tunnel environment.

4.1.1. Skin geometric design

The first step of the design process was the assessment of the sensitivity of the simulation results to both the design parameters and the hump shape parameters. The parameter space for this initial study is shown in Table 4.2:

Table 4.1: Initial study Boundary Condition cases, Hump shapes: hX, height in mm; RXX, Ratio; eXXX, skewness $\ast 10^{-2}$; x_c XXX, chord non dimensional position $\ast 10^{-3}$.

Hump shape	$\lambda = \frac{base_{Hump}}{2Spacing_{act}}$	Position - on actuator
h3 R50 e000 x_c 150	5.77	Front - ON
h3 R50 e000 x_c 260	5.77	Middle - OFF
h3 R20 e000 x_c 235	2.31	Middle - ON
h3 R20 e000 x_c 244	2.31	Middle - OFF
h3 R20 e000 x_c 260	2.31	Middle - OFF
h1 R50 e000 x_c 150	1.92	Front - ON
h1 R50 e000 x_c 235	1.92	Middle - ON
h1 R60 e000 x_c 260	2.31	Middle - OFF
h05 R50 e000 x_c 235	0.96	Middle - ON
h05 R50 e000 x_c 260	0.96	Middle - OFF

Table 4.2: Initial Design parameter space, 27 possible foot geometries

Parameter	Values
ST	2.5 – 3.5 – 4.5
LT	2.0 – 3.5 – 5.0
LL	4.0 – 6.0 – 7.5
LTR	1.0

After this first round of simulations the results were promising; in fact, the RMS errors were lower than 0.1 mm. Some trends were observed both between different hump target curves and between design parameters.

Among the different hump shapes, the most influential characteristics were: the gradient present on the shape, where an increase in the steepness of the curve corresponded to an increase in the error — in particular, the increase of $\lambda = \frac{base_{Hump}}{2Spacing_{act}}$, the ratio between the hump's base (width) and double the spacing of the actuators, led to bigger RMS errors; the relative position between the peak of the hump shape and the underlying actuator, with a lower error for better alignment; and the error arising from the cutting effect of the skin LE, which limits the physical representation of humps on the skin to a certain upstream chord location. These results are shown in Fig. 4.1.




Figure 4.1: Initial study on the skin design. Boundary condition parameter sensitivity. In blue, the RMS error and, in orange, the maximum difference between the target and the simulated top surface. The top row shows the effect of the gradient; the left column shows the effect of the position of the peak of the hump; the right column shows the effect of the leading edge.

More importantly, regarding the effect of the design parameters, some results are shown in Fig. 4.2. In these images there are two graphs, taken from the results of a steeper hump BC and a smoother one. The trends are similar but show how a smoother hump shape performs better with a thicker skin, unlike the steeper case. The similar trends are on the LT, where thinner is better, and on the LL, where shorter is better, even if barely noticeable in the smoother case.




Figure 4.2: Initial study on the skin design. Design parameter sensitivity. On the vertical axis, the RMS error in millimetres. On the left, the results of the BC h3 R20 e000 x_{c235} and on the right the BC h3 R50 e000 x_{c260} , one representing a high-gradient case and the other a lower one

Following this first study, the parameter space was changed and a final set of simulations was solved.

In this case the hump shapes were all on top of the actuators and, to further increase the gradient of the shape, the skewness was investigated. Learning from the previous study, the choice of the design parameters focused more on the full range of the LL and LTR and on shorter ST and LT (see Table 4.4).

Table 4.3: Final study Boundary Condition cases, Hump shapes: hX, height in mm; RXX, Ratio; eXXX, skewness $\times 10^{-2}$; x_c XXX, chord non dimensional position $\times 10^{-3}$.

Hump shape	$\lambda = \frac{baseHump}{2spacing_{act}}$	Position - on actuator
h3 R50 e000 x_c 15	5.77	Front - ON
h3 R50 e000 x_c 235	5.77	Middle - ON
h3 R20 e000 x_c 235	2.31	Middle - ON
h3 R50 e065 x_c 235	5.77	Middle - ON
h3 R50 e-65 x_c 235	5.77	Middle - ON
h1 R50 e000 x_c 235	1.92	Middle - ON
h1 R50 e065 x_c 235	1.92	Middle - ON
h05 R50 e000 x_c 235	0.96	Middle - ON

Table 4.4: Final Design parameter space, 36 possible foot geometries

Parameter	Values
ST	2.0 – 3.0 – 3.5
LT	2.0 – 3.0
LL	3.0 – 4.0 – 6.5 – 7.5
LTR	0.5 – 1.0

To finally arrive at the choice of the design geometry of the skin, the ranking of the RMSE results of the parameter space is shown in Fig. 4.3. This ranking is nothing more than a counter of the good-enough results of every geometry in the different BC cases. This is done by counting for how many different BCs the geometries appear within a certain threshold of the best RMS geometry in every BC (the BCs are eight in total, so the maximum rank can be eight).

The threshold used is 0.015 mm, so every geometry that yields an $RMSE > RMSE_{bestBC} + threshold$, i.e. higher than the best RMSE result in every BC, has been discarded. In addition, the remaining geometries are classified as physically doable if the total height is lower than 6.5 mm and are highlighted with a blue name and a red bar.




Figure 4.3: Ranking of the geometries. On the bottom side the code names of the parameter sets. Higher the counter the better the geometry is in different BC. Blue name with red bars are the geometries with a height lower than 6.5 mm (6 mm was the design target).

The final choice of the geometry of the skin was made after some practical considerations: the vertical space left for the skin by the preliminary design was 6 mm, so more than 6.5 mm would have limited the movement of the actuators; the LT of 2 mm is dangerous from the production point of view, because of the 1.5 mm thickness of the aluminium leg. The choice of an LT of 2 mm would have resulted in 0.25 mm slices of silicone bonded on both sides of the aluminium, which is doable but with higher risk; also, the LTR was limited to 1, given the resulting simplification of the support legs' manufacturing; finally,

the choice of the *ST30-LT30-LL30-LTR10* corresponded to a 10 mm support leg, which again simplified the manufacturing.

Even though the final choice of this study was the *ST30-LT30-LL30-LTR10* geometry, it was changed to an *ST35-LT30-LL30-LTR10* during the casting phase to be more resilient against a waviness observed on the skin top surface. This new choice is in fact in the ranking, so it was considered a good choice, trading 0.5 mm of the actuators' range for a better surface finish.

4.1.2. Skin sensitivity analysis

In this section the sensitivity of the skin to temperature and flow velocity is tested through some laser measurements. The data were collected by measuring the skin surface under different conditions and with four different shapes. The four shapes, divided into the four images of Fig. 4.4, are three $h = 1.5$ mm, $R = 50$ humps located at $x_c = 0.15, 0.25, 0.35$, and the current M3J preset. Each shape was measured under four conditions, in this order: without the wind and without the halogen lamps; without the wind and after a 7 minute warming time of the skin, which made the skin temperature increase by 7°C (much more than in the experiments); with the wind on and the halogen lamps off, but only after a cooling time for the wing to return to the flow temperature; finally, with the wind and the halogen lamps on, as in the experiments (turning off the lamps during the actual laser measurements). The freestream velocity adopted in this study was selected to match the average Reynolds number used in previous studies, resulting in a test Reynolds number of $2.3 \cdot 10^6$. However, a posteriori, it appears that a higher freestream velocity would have been more appropriate getting closer to the maximum velocity tested in this campaign (equivalent Reynolds number of $2.9 \cdot 10^6$).




Figure 4.4: Skin sensitivity to temperature and flow velocity. In each image there are four curves describing the effect of the environmental conditions on the silicone skin surface shape. The curves in the four images are: the M3J preset, the $x_c = 0.15$ hump, the $x_c = 0.25$ hump and the $x_c = 0.35$ hump.

The first observations on the shapes are that the preset itself appears to slope upward in the chordwise direction and that the hump shape is not significantly modified by the different conditions. Something that is influenced by the different forcings is the waviness of the surface, which changes in magnitude and phase. No clear correlation was observed between the waviness and the different forcings; this could be attributed to the low amplitude of the waviness itself, 0.15 mm (10% of the hump height), and to the small uncertainty added to the data by the roto-translation applied to the measured laser data.

It can be concluded that the main factors that modify the skin surface are not the wind velocity or the heating of the halogen lamps, but rather the silicone skin inner stresses and the actuators' displacements.

4.1.3. Skin span-wise measurements

The aim of the data collected and presented in this section was to characterise the spanwise variability of the silicone skin. The need for these measurements was due to the initial IR transition front images, where the transition location shifted a lot in the chordwise direction along the span. The data collected in this investigation were some laser measurements, in particular the profiles of some skin surface humps at different spanwise locations. The data collection consisted of measuring seven different skin surface shapes (see Table 4.5) at a given spanwise location, after which the measurement position was moved by about 20 mm in the spanwise direction and the procedure was repeated. The area covered by this procedure is shown in Fig. 4.5 by the red line perpendicular to the LE (the recorded laser profile) and by the yellow arrow that defines the spanned region.

Hump shapes measured
hump $h = 0.0$ mm - preset M3J
hump $h = 1.5$ mm - $x_c = 0.15$
hump $h = 1.5$ mm - $x_c = 0.25$
hump $h = 1.5$ mm - $x_c = 0.35$
hump $h = 2.5$ mm - $x_c = 0.15$
hump $h = 2.5$ mm - $x_c = 0.25$
hump $h = 2.5$ mm - $x_c = 0.35$

Table 4.5: Shapes recorded in the spanwise laser measurements. The other humps parameters were $R = 50$ and $e = 0$ and the wind tunnel speed was 0 m/s.




Figure 4.5: Laser line projected on the wing surface (perpendicular to the LE) and an arrow indicating the path covered by the laser line.

The first conclusion drawn from the analysis of the preset configuration (M3J shape without humps) is the presence of a spanwise non-uniformity of the skin. This result did not come as a surprise given the non-uniformity of the transition fronts in the IR images. The visualisation of the non-uniformity can be seen in Fig. 4.6 and a more detailed characterisation of it in Fig. 4.7.




Figure 4.6: Preset spanwise non-uniformity visualisation. On the left, a 3D mapping of the laser data showing the error between the laser measurements and the M3J shape. On the right, a qualitative validation picture of the wing exposed to a grazing light directed along the surface to highlight local waviness.

The data collected show how the middle locations are the most accurate, where the errors between the target and the laser profile are less than 0.2 mm. This observation is the consequence of the preset procedure, which ensures that the skin is flat in the middle span location, where the laser is kept during the operation. The upper part of the span is characterised by a higher spot towards the LE followed by a valley. Instead, the lower part of the span is characterised by a higher spot towards the TE and three wide valleys in the central streamwise location. The sudden appearance of these valleys in the span measurements suggests that they may not correspond to the actual skin surface. This aspect needs further investigation.




Figure 4.7: Preset spanwise non-uniformity study. On the left, a heatmap of the top view of the laser data, showing the error between the laser measurements and the M3J shape as colour. On the right, the measurements of the skin compared with the M3J shape at three different span locations and the error of those profiles.

The most important result of these data is the selection of the spanwise range in which the silicone skin is close enough to the desired shape. This range is 1520 cm wide and will be used as the area in which the IR transition location will be detected.

Similar observations can be made in the other hump shape measured.

The origin of this spanwise non-uniformity remains unclear. A possible explanation is related to a non-uniform stress state within the silicone skin, in which shear stresses may play a dominant role. The elevated regions observed at opposite corners of the skin (top left and bottom right in Fig. 4.7) could be explained by a shear deformation induced between the LE and TE. One possible source of this shear is the assembly procedure of the silicone skin onto the wing. During the installation process, a non-uniform spanwise force may have been applied while tightening the screws of the 20 actuator assemblies. Variations in this force among different actuator rows could have produced a slight spanwise misalignment of the actuator structure, resulting in a shear deformation of the skin. This hypothesis is further supported by an additional potential source of shear. During the manufacturing of the wing model, the actual model hole was shifted by 3 mm in the spanwise direction at the LE, whereas the silicone skin was manufactured according to the nominal design geometry. This discrepancy may have introduced a residual shear strain when mounting the skin onto the wing. Confirmation of this hypothesis requires further investigation. In particular, measurements of the actuator assembly without the skin installed would allow verification of the alignment of the mechanical structure. In addition, a new skin installation procedure, with increased attention to spanwise alignment during screw tightening, could help assess the influence of assembly-induced shear. If neither of these actions eliminates the observed non-uniformity, the manufacture of a skin incorporating the same 3 mm spanwise offset as the model hole is suggested.

4.2. Data validation

The aim of this section is to present the data validation phase of the experimental campaign. This phase was carried out before the experimental tests and was intended to determine when the acquired data could be considered sufficient to properly characterise the experimental setup. Defining this point was important for balancing the effort devoted to the setup characterisation with the time available for its experimental testing during the campaign.

4.2.1. Preset procedure repeatability

For clarity, it is worth recalling that the preset refers to the actuator configuration used to reproduce the surface shape of the M3J model. The first validation test was about the consistency of the preset making procedure. The measurements consisted of making six different presets and testing the consistency of the resulting measurements. Every time a new preset had to be done, the whole assembly was powered down to emulate the end-of-the-day procedure; the last preset shown was actually taken on a different day.




Figure 4.8: Full view of the preset consistency study. The presets are the M3J baseline, which is modified to obtain a hump. The first 55% of the chord is shown in addition to the silicone skin location (grey vertical bars).

As shown in Fig. 4.8, the errors between the preset M3J shape and the target one are minimal. To actually measure the error between the curves, the data handling procedure described in Section 3.4.2 was used. The errors shown in Fig. 4.9 are in most cases within a ± 0.1 mm range from the target, and within an even smaller range between different presets. The preset made on a different day also verifies the consistency of the procedure even with different settings of the skin and setup (mainly temperature). In the figure on the right, the measurements show how the skin behaves when a current preset is compared with the actuators' displacements of a different preset. This shows why there is the need for a preset procedure after a power-off. Even on the same day, the presets are really different




Figure 4.9: Error in the silicone skin between the M3J target and the different presets. On the left, the actual errors when all the measurements were taken on current presets; on the right, the measurements taken when only the last preset was the current one. Note the difference in the y-axis between the two figures and that only one preset can be the current one.

4.2.2. Effect of angle of attack on the preset

The next test was on the influence of the angle of attack on the preset procedure. This parameter is important not because of its values, which can change only between 0° and 3° , but because of the slight changes that it has on the preset procedure. In fact, a different AoA meant that the wing's cross section changed and that both the calibration and the error evaluation had to be done with a different target curve. The cross section variation is not a simple rotation of the cross section, as it might seem. Because the axis of rotation is not parallel to the LE and the laser beam does not rotate with the wing model, the beam measures a slightly different span location of the wing model and the resulting cross section is not perpendicular to the LE any more. The new shape was found after designing the wing

3D surface in CAD, applying the 3° rotation and cutting the cross section with the original 45° cutting plane, actually reproducing the physical space. From this new target shape arose the choice of using the new curve as a target for either the preset or the error evaluation, or both of them. The possible preset procedures are: doing the preset at an AoA of 0° and then using the new target only for the 3° AoA measurements, or doing the preset directly at an AoA of 3° and using the new target, but in this case the comparison is on whether or not to use the different M3J target curve as the reference for the error evaluation or the calibration (roto-translation).

Fig. 4.10 shows, on the left, the error between the M3J targets calculated with the 0° and 3° AoA, restricted to the silicone section, and, on the right, the results of using the two targets on the same calibrated data (calibrated with the zero-degree shape).




Figure 4.10: On the left the error between the M3J target and the M3J target calculated at 3 degrees. On the right, the comparison between using the two different targets as a reference in the error evaluation for the same data.

The effect of using the original M3J target to evaluate the error on a rotated wing model affects the overall error plot with an increasing difference in the x-coordinate. The magnitude of the error is 0.03 mm, increasing in x to 0.055 mm, which, in the present study, can be neglected. In this section the target shape used for the error evaluation will be written in the legend of the figures, but in the other sections the standard will be the use of the original M3J shape.

The use of the two different targets in the process has been tested a posteriori and shown in Fig. 4.11. The results show that there is a negligible difference when using the two targets, provided that the same target is consistently used during both the calibration and the error evaluation phases. Regarding the choice of acquiring the preset at an AoA of either 0° or 3° , qualitative observations suggest that this parameter influences the results. However, its effect has not been quantitatively characterised in the present study. A dedicated investigation is therefore recommended in future work.




Figure 4.11: Comparison of the use of the two M3J targets (M3J and M3J_3deg) in the calibration phase and in the error evaluation phase.

4.2.3. Repeatability of the setup

In this section, the repeatability of the IR transition location is assessed through the repetition of hump shapes on different days. These shapes are a height sweep of humps with constant $R = 50$ and $x_c = 0.15$. The DREs configuration and the AoA are set as follows: $\lambda = 7.5$ mm low amplitude for the DREs and AoA = 3° . In particular, the preset comparison is isolated on the right in Fig. 4.12.

The first observation is that the repeatability of the transition location is not always assured. In particular,




Figure 4.12: Repeatability of the transition fronts location data between different days. On the right the M3J preset data is isolated.

with higher features the repeatability is much worse. Conversely, if the humps are short, or the skin is not even morphed, the results are good. Another observation is that the dotted lines of the higher humps, where the bigger error lies, are constantly advancing transition with respect to the continuous line. At present, this effect may be attributed to something similar to a systematic bias, as the perturbation amplitude was slightly higher upstream in one of the two cases. Nevertheless, this hypothesis remains speculative and needs further investigation.

4.2.4. Cross model sides validation of the IR data

In this section the data acquired on the silicone side and on the clean side are compared. This comparison is valuable because the wing model is symmetric and, by mirroring the AoA, effectively going to 3° , the flow conditions on the clean side of the wing model surface are the same as on the silicone side. Thus these data can be used to validate the modified wing model side. The results of the comparison are shown in Fig. 4.13.




Figure 4.13: Cross sides validation, the same DREs forcing are drawn in the same colour. On the left the IR data acquired on the clean side of the wing. On the right the IR data of the silicone side with the preset M3J shape of the skin. The AoA is 3 and -3 degrees to match the same pressure distribution on the pressure side, where the IR measurement is taken.

The most evident feature is the general advancement of the transition fronts on the silicone side, not influenced by the Reynolds number or the spacing of the DREs (CFI wavelength value). In fact, a closer look at the trends highlights a clear resemblance between the two sides, implying a sort of constant advancement effect of the modified side, as if the amplitude of the perturbations were higher from the start. This observation is challenged only by the $\lambda = 7.5$ mm low amplitude case, where this behaviour is not present and the trend differs between the two sides.

4.2.5. Cross-campaign validation of the IR data

The analysis in this section is a comparison between the data acquired in previous experiments and those of the present campaign. In particular, there is a comparison of the data taken months prior to the wing model modification and after it; that is, a comparison between humps made with a cast system, treated as an external component glued on the wing surface, and the silicone skin method. The data were reproduced with matching flow conditions and DREs configuration. The humps tested are of different heights but the same $R = 50$ and $x_c = 0.15$. It is important to note that the hump surface shapes are slightly different because of the flaws of the two manufacturing methods. The skin hump, as described in the sections above, is not spanwise invariant; instead, the cast hump is not in continuity with the wing surface: there is a small step both at the LE and TE of the hump, caused by the manufacturing procedure.




Figure 4.14: Comparison of the IR data of the present campaign (continuous line) and the IR data taken in a previous campaign (dotted line). The flow conditions are the same, but the hump manufacturing technology changes between the modified wing model and a cast hump glued in position on the wing surface.

The results of this analysis are similar to those obtained in the cross-sides validation. The IR transition locations follow the same trends but, as in the comparison with the clean side of the wing model, the silicone-modified side advances the transition. The advancement of the transition in the modified wing model seems to be lower than in the previous data at higher Reynolds numbers (see Fig. 4.14). Another observation is the absence of strong transition jumps with increasing Reynolds in the present campaign data. This could be explained by a sudden tripping of the flow by the backward-facing step in the trailing part of the cast hump.

4.3. Effect of the DREs

The focus of this section is on the data collected with different DREs configurations. The choice of the configurations shown in Table 4.6 was made to validate the setup along with acquiring novel knowledge. The validation analysis of the data is reported in Section 4.2.

Table 4.6: Data acquired during the DREs study. The dimensions of the DREs were fixed (see Section 3.3). The other hump parameters were fixed to: $x_c = 0.15$, $R = 50$ and $e = 0$. The “-.-” line stands for one step of the sweep for each “-” or “.”.

DREs configuration	Sweepled parameters
Clean configuration	h 0.5-.-.-3.5
$\lambda = 6$ mm, $x_c = 0.05$	h 0.5-.-.-3.5
$\lambda = 9$ mm, $x_c = 0.05$	h 0.5-.-.-3.5
$\lambda = 12$ mm, $x_c = 0.05$	h 0.5-.-.-3.5
$\lambda = 7.5$ mm, $x_c = 0.04$	h 0.5-.-.-3.5
$\lambda = 7.5$ mm, $x_c = 0.05$	h 0.5-.-.-3.5

4.3.1. Effect of DRE spanwise spacing

The direct effect of the DRE spanwise spacing, λ , is to selectively amplify a specific wavelength component relative to the others. Although this mechanism increases the amplitude of the instabilities, it may also delay transition with respect to the unforced case, as the selection of a more slowly growing perturbation can suppress the development of faster-growing modes. This behaviour can be observed in Fig. 4.13, where the $\lambda = 6$ mm configuration delays transition compared to the unforced case. Conversely, the $\lambda = 7.5$ mm configuration promotes an earlier transition compared to the other wavelengths, owing to its higher growth rate.

4.3.2. Effect of amplitude

The effect of the amplitude has been described both numerically and experimentally, as reported in Section 2.5. The aim of this section is to assess the behaviour of the wing model humps under two different forcing environments. In this comparison, the IR data of the two configurations, described in Table 4.6, are shown in Fig. 4.15.




Figure 4.15: Amplitude study on the IR data of the transition location. On the left, the high-amplitude case and, on the right, the low-amplitude case. The humps shaped in both cases are the same. The black line represents the M3J preset, so any point to the right of it is delaying transition.

The influence of amplitude remains unclear. Nevertheless, the trends observed for the different hump heights appear to be qualitatively similar apart from a small transition advancement of the higher-amplitude case. These results are not as clear as the literature would have predicted and need more detailed studies on the effective amplitudes reaching the humps.

4.3.3. Unforced results

In this section the results of a hump height sweep are shown in absence of a DREs configuration, this means that the amplitude of the instability perturbations impacting with the hump is generally lower and without a defined spanwise wavelength. The results are shown in the Fig. 4.16.

In this section the results of a hump height sweep are shown in the absence of a DREs configuration; this means that the amplitude of the instability perturbations impacting the hump is generally lower and without a defined spanwise wavelength. The results are shown in Fig. 4.16.




Figure 4.16: Clean case with no forcing through DREs: IR transition data. The measurements are taken from the silicone side of the wing model, morphing different hump shapes with constant $R = 50$ and $x_c = 0.15$. The black line represents the M3J preset, so any data point to the right of it is a transition delay

The first notable and most important observation is that, at every Reynolds and hump height, the transition location is delayed. This behaviour is never reproduced by the forced cases at any height or Reynolds magnitude. This is a promising result for the future of this research and confirms the thesis that the hump shape can delay transition if the incoming perturbation wave has a lower amplitude. Although this is a great conclusion, the way to emulate the commercial free-flight cruise conditions is still long and there is the need of increasing the Reynolds number by an order of magnitude and to assess the Mach effects.

4.4. Parametric Study Results

In this section the analysis of the majority of the IR data collected is described. The aim of this set of data is to explore the design parameter space of the hump geometrical shape. It is important to emphasise once again that the primary objective of this experimental campaign was to commission and validate the experimental setup. Although some trends could already be identified, the main purpose was just to detect them, leaving the investigation of the underlying physics to future studies. A future rework of these data is suggested for two reasons: to do a more physics-oriented analysis and to do a more accurate correlation work between the IR and laser data, characterising each surface shape based on the laser measurements instead of using the actuators' displacements. This latter aspect has been found to be crucial in this parameter study.

The hump geometrical shape is governed by four parameters: height, Ratio, skewness and position. The Ratio is the fraction of the base over the height of the hump; the skewness is a parameter that can vary between negative one and positive one and determines the asymmetry of the shape (zero is symmetrical); finally, the position is the location of the hump's peak on the chord of the profile in non-dimensional coordinates. See Fig. 4.17 for a visualisation of the parameters.




Figure 4.17: Hump geometrical shape and parameters. Note the different scales of the axes.

The data in this section of the test campaign were acquired with a fixed AoA of 3° , a fixed DREs configuration¹ and a constant sweep of Reynolds². These data are a series of IR images (around 75 samples) and laser measurements (50 samples) for every geometry. The collected data can be summarised in Table 4.7.

¹Nominal DREs configuration measures: spacing of $\lambda = 7.5$ mm, position of $x_c = 0.05$ and a diameter $d = 2$ mm.

²Sweep of Reynolds of 5 points with variable values between $2.1 \cdot 10^6$ to $2.9 \cdot 10^6$

Table 4.7: Data acquired during the parametric study

Fixed main parameter	Sweepled parameters range	Other fixed parameters
pos 0.15	h 0.5-.-.3.5	R 50, e 0
pos 0.165	h 0.5-.-.3.5	R 50, e 0
pos 0.18	h 0.5-.-.3.5	R 50, e 0
pos 0.194	h 0.5-.-.3.5	R 50, e 0
pos 0.222	h 0.5-.-.3.5	R 50, e 0
e 0.3	h 0.5-.-.3.5	R 50, pos 0.15
e -0.3	h 0.5-.-.3.5	R 50, pos 0.15
e 0.6	h 1-.-.3.5	R 50, pos 0.15
e -0.6	h 1-.-.3.5	R 50, pos 0.15
R 20	h 0.5-.-.3.5	e 0, pos 0.15
R 100	h 0.5-.-.2	e 0, pos 0.15
R 100	h 0.5-.-.3	e 0, pos 0.194
pos 0.15	h 1-2, R 50-100, e -0.3-0.3	
pos 0.165	h 1-2, R 50-100, e -0.3-0.3	
pos 0.18	h 1-2, R 50-100, e -0.3-0.3	
pos 0.194	h 1-2, R 50-100, e -0.3-0.3	

4.4.1. Height effect

The height of the hump is a crucial parameter. In fact, as can be seen in Fig. 4.18 and is confirmed by a lot of other cases, the hump height greatly shifts the transition location.




Figure 4.18: Height study. In particular, these data are hump geometries with fixed $R = 50$, $e = 0$ and $x_c = 0.15$, with different heights. The black line represents the M3J preset, so any data point to the right of it is a transition-delaying geometry. As a reference, in the top part of the figure there is half of the M3J profile in grey and the location of the silicone skin pasted on it in red.

In particular, in Fig. 4.18, showing the combined effect of the height and Reynolds, it can be observed that, at a fixed Reynolds, the height of the hump increasingly delays the transition front until a certain height, where it starts to advance it quite abruptly. Moreover, increasing the Reynolds leads to a reduction of this critical height. This behaviour can be understood thanks to the BL characteristics. In fact, it seems that there exists a critical ratio between the height of the BL and the height of the hump, above which the hump trips the flow and right below which the delay is the maximum. Given the fact that the BL thickness decreases with increasing Reynolds, the critical height gets lower with higher Reynolds numbers.

4.4.2. Ratio effect

The effect of the Ratio is less clear and less critical compared to the height. Fig. 4.19 can be used to make some observations. At low heights (1 mm) the different Ratios give really close results in terms of transition location, as if the Ratio could be neglected. With increasing heights the effect of the Ratio




Figure 4.19: Ratio study. In the first three images there are sweeps of Ratios at the three constant heights of 1 mm, 1.5 mm and 2 mm. In the last image the effect of the Ratio is evaluated in different chordwise positions

becomes more visible, starting from the top and bottom Reynolds limits. In fact, in the middle range of




Figure 4.20: Laser measurement Ratio. The image shows the surface measurement of the $h = 1.5$ mm, $R = 20$, $e = 0$ and $x_c = 0.15$ both in full view (on the top) and as the error between the measurement and the target. More Reynolds curves are shown.

Reynolds there is less difference between different Ratios. The peculiar observation is that at low and high Reynolds the best-performing Ratio changes and at first sight it seems to flip (from best 50-100-20 to 20-100-50). A more accurate analysis can challenge this observation thanks to two facts: the actual hump geometry has a skewness given by the impossibility of reproducing the wide base of some data

points caused by the bounds of the skin surface (see the upper grey M3J reference); in addition, a specific focus on the actual shape of the hump $h = 1.5$ mm - $R20$ showed that the true Ratio is not $R20$ but higher (see Fig. 4.20).

After the latter consideration on the $R20$ data, there is actually a trend that is still emerging from the last plot of Fig. 4.19. There seems to be a critical Reynolds (probably again a critical height of the boundary layer with respect to the hump) where the best-performing Ratio changes from $R50$ to $R100$. It is as if the $R50$, i.e. more surface gradient, which is equivalent to a bigger pressure gradient in the stream direction, can yield better delays but is much more prone to trip the flow if the height of the BL decreases. A milder surface gradient, instead, can still delay the transition compared to the no-hump surface even at higher Reynolds.

4.4.3. Skewness effect




Figure 4.21: Laser measurement Skewness. The image shows the surface measurement of the $h = 2.5$ mm, $R = 50$, $e = -0.6$ and $x_c = 0.15$ both in full view (on the top) and as the error between the measurement and the target hump. More Reynolds curves are shown.

The effect of the skewness of the hump shape is more difficult to assess than it seems. In fact, with this type of shape modification there are more critical factors. The main one is the feasibility of the shape itself, which, given the strong chordwise gradients and the hump base streamwise shift, makes the shapes developed on the silicone skin different from the target ones. A clear example of these errors is shown in Fig. 4.21, where a negative skewness is introduced by the physical skin when it was not in the target, compromising this analysis. In this light, the results of this section have to be carefully weighed, and a further characterisation of the hump shape parameters has to be made with the aid of each laser measurement to properly evaluate the results. From the available processed data, shown in Fig. 4.22, it seems that a negatively skewed hump delays the transition front more than a positive one and, in particular, it performs better than an $e = 0$ skewness hump, which is close to or better than a positively skewed hump. This behaviour is particularly visible in humps with higher heights. It could be argued that a reason for the lower performance of the positive skewness is the cutting effect on the shape in the vicinity of the LE of the skin chord range. This reason is refuted by Fig. 4.22, where the negative skewness is consistently better than the positive one at different x_c positions, which confirms the previous observation.

An additional observation made during the data analysis is on the spanwise behaviour of the transition front: in fact, in these cases the fronts were less uniform. The stronger chordwise gradients imposed by the actuators could be the reason for a more complex stress state in the silicone skin, thus creating a more complex and non-uniform surface shape.

4.4.4. Chord-wise location effect

The chordwise location of the hump is very important because, as described by the literature in Section 2.5, the amplitude of the perturbation is key to the delay mechanism of the hump and the amplitude




Figure 4.22: Skewness study. The first three images are part of a height sweep study where the fixed parameters are $R = 50$ and $x_c = 0.15$, whereas in bottom right there is a sweep of skewness and chordwise position combined at a fixed $h = 2$ mm and $R = 50$

of the perturbation grows in x . The effect of the location is shown in Fig. 4.23, where three different humps are compared at different positions.

The first observation confirms that increasing the chordwise position of the hump advances transition. This is clearly visible from the figures of the higher humps, where the lower the chordwise position, the more downstream the transition location is. In addition, the humps located more downstream advance transition at lower Reynolds numbers compared to the upstream ones, which seems to indicate that a linear CFI interaction with humps is more important than a suitable height. Something different is happening in the furthest location tested: the $x_c = 0.31$ location delays transition, unlike the $x_c = 0.22$ location, which advances it (at selected Reynolds numbers). This is a peculiar result to be investigated with more x_c locations.

4.4.5. Maximum delay visualization

This section is intended to reconnect the reader with the physical phenomena observed during the experimental campaign. The figures show a dewarped full-span view of the wing model, with the non-dimensional chordwise coordinate superimposed on the images. The normalised infrared thermal images detect the transition front through a change in colour (black turbulent, white laminar), which makes it possible to see the jagged transition front characteristic of a stationary crossflow instability. The shadows covering up to $x_c = 0.15$ visible in the second and third images are a leftover of the previous data point that tripped the flow right after the hump location (definitely a higher Reynolds data point), suggesting that a longer settling time would have been beneficial to the quality of the data.

The choice of the data points in Fig. 4.24 comes from the desire to see the effect of the hump on the flow field. In the first image there is no hump, and in the second and third ones there is one. As can be clearly seen, the effect of the hump is to delay the transition front up to the $x_c = 0.7$ location. The




Figure 4.23: Chordwise location study. In the three images there are sweeps of chordwise location at the three constant heights of 1 mm, 1.5 mm and 2 mm.

transition front shape in the second and third images is not a jagged one because of the "cutting" effect of the change of curvature of the wing model at that location, implying that it could have delayed the transition even more. The last observation is that some wedges are visible in the last image; these were probably caused by a faulty DRE or a local dust particle on the model, which increased the amplitude of the CFI.




Figure 4.24: Visualization of some normalized averaged dewarped full images (flow from right to left). All the images are taken with a $\lambda = 7.5$ mm DREs configuration and similar Reynolds number (written in the title). The first image is the M3J preset, the second one is the $h2.5 R50 e = -0.6 x_c = 0.15$ and the third one the $h2.0 R50 e = 0 x_c = 0.15$

4.5. Unconventional shapes

The unconventional shapes are shapes different from humps. The choice of testing these shapes was made to explore a little further the parameter space of shapes achievable with the morphing setup and, admittedly, to satisfy some curiosities of the thesis writer.

4.5.1. Reverse Humps

The choice of testing reverse humps came from the direct curiosity of exploring the setup capabilities; in fact, the actuators can move both outside the M3J target and inside it. A reverse hump is a shape parametrised as a positive hump with the only difference that every calculated displacement is in the inward direction.




Figure 4.25: Reverse hump IR transition location data. There are data of different reverse humps at two x_c locations. The black line represents the M3J preset, so any point to the right of it is delaying transition.

The results of the IR measurements are very interesting. Fig. 4.25 displays humps of different heights at two chord locations and, even though the heights are over 2 mm, all the different measurements seem to collapse over the entire Reynolds range. This result seemed strange; focusing on the bigger reverse hump and checking the laser-measured surface, the hump shape is in fact smaller than the target one. A suction effect from the flow is supposed to have played a role in this skin deformation. Nevertheless, the results are more valuable and novel because of the laser measurements that verify the feature (see Fig. 4.26). In this case it can be stated that a reverse hump of 2 mm does not affect the transition location. Focusing even more on the latter hump, it behaves differently in the high Reynolds range. In fact, contrary to the usual transition advancement, it delays transition more than it does at lower Reynolds. This preliminary observation, which needs to be investigated in more detail, shows significant potential for future developments.

4.5.2. Additional unconventional shapes

The choice of testing these shapes comes from the curiosity of the writer to reproduce the results obtained in the preliminary optimization loop in the final minutes available for the wind tunnel campaign. The output shape of the optimizing loop, given all of its flows and limitations, was not a single hump shape, it was similar to a composition of a big reverse hump followed downstream by another hump. This preliminary PSO optimized configuration is shown in Fig. 4.27 in addition to the tested shapes. The colours of the shapes are consistent between the Fig. 4.27 and Fig. 4.28.

The choice of testing these shapes comes from the curiosity of the writer to reproduce the results obtained in the preliminary optimisation loop in the final minutes available for the wind tunnel campaign. The output shape of the optimisation loop, given all of its flaws and limitations, was not a single hump shape; it was similar to a composition of a big reverse hump followed downstream by another hump. This preliminary PSO-optimised configuration is shown in Fig. 4.27 in addition to the tested shapes. The colours of the shapes are consistent between Fig. 4.27 and Fig. 4.28.




Figure 4.26: Laser measurement reverse Hump. The image shows the surface measurement of the $h = -3$ mm, $R = 50$, $e = 0$ and $x_c = 0.165$ both in full view (on the top) and as the error between the measurement and the target hump. More Reynolds curves are shown




Figure 4.27: Comparison of the unconventional shapes tested and the PSO run 3 result. The plot shows the error between the measurement and the M3J target.

As can be seen, the tested shapes have the same conventional hump positioned at $x_c = 0.31$ ($x = 279$ mm) but a different upstream shape. The upstream reverse hump was linked to the downstream hump in two different ways, shaping a plateau in between the two peaks or connecting them directly through a third-order polynomial curve. The resulting shapes, however, are similar to the PSO output, but the height and the location of the reverse hump are not exactly the same. This was caused by the little time remaining to double-check the shapes.




Figure 4.28: Unconventional shapes IR transition location data. There are data of different shapes shown in Fig. 4.27. The black line represents the M3J preset, so any point to the right of it is delaying transition. The purple line is the reference hump that should have been improved by the unconventional shapes.

The IR results in Fig. 4.28 show how these shapes advance the transition location for most of the

Reynolds range, and it can be stated that they are not a good choice compared to the conventional hump reference. Despite this clear conclusion, the curiosity of the writer has not fully vanished because of the different location of the humps in the preliminary PSO result; further studies could be carried out to finally declare these shapes suboptimal.

4.6. Preliminary Optimization Loop

The best result regarding the optimisation loop is not the optimised output shape but rather the fact that it worked as an optimisation tool. The main limitation of this preliminary study was the time constraint; in fact, more time would have been needed to actually think of more clever ways to initialise the loop, keep the actuators' shape inside some physical constraints without wasting an internal iteration, build a better-suited reward function and leave the PSO running for a longer time.

Even in the presence of these limitations, the PSO actually found a transition-delaying shape in three out of five PSO runs. The problem with the two faulty runs was the initialisation shape; in fact, the optimisation loop did not optimise: the best actuators' displacement guess in the first iteration remained the best overall for all the next iterations (in terms of reward function). It can be noted that a higher external iteration number could have improved the guess, but the loop was stopped after 30 minutes to change some settings and try again.




Figure 4.29: PSO final update visualization, run 1




Figure 4.30: PSO final update visualization, run 3

The three optimization runs, which demonstrated the feasibility of delaying the onset of transition through the optimization procedure, were performed at a constant Reynolds number of $2.2 \cdot 10^6$ with

DREs forcing of $\lambda = 7.5$ mm. Under these conditions, the transition front corresponding to the preset M3J silicone shape is located on the downstream aluminum strip of the silicone skin (see Fig. 4.31).

The three optimisation runs, which demonstrated the feasibility of delaying the onset of transition through the optimisation procedure, were performed at a constant Reynolds number of $2.2 \cdot 10^6$ with DREs forcing of $\lambda = 7.5$ mm. Under these conditions, the transition front corresponding to the preset M3J silicone shape is located on the downstream aluminium strip of the silicone skin (see Fig. 4.31).




Figure 4.31: PSO reference IR image. Normalized IR image of a M3J shape of the silicone skin under $2.2 \cdot 10^6$ Re with a $\lambda = 7.5$ mm DRE forcing

The results of run 1 and run 3 are shown in Fig. 4.29 and Fig. 4.30, respectively. As can be seen from the external iteration plot, both runs progressively increased the reward function value to the point of delaying transition beyond the M3J preset limit. In both PSO loops there were five inner particles, and the number of external iterations was 20 and 25, respectively. Two different sets of PSO parameters were used in the two optimization runs. In the first run, the parameters were fixed at ($w = 0.6$), ($c_1 = 1.5$), and ($c_2 = 1.5$). In the third run, ($c_1 = 1.5$) and ($c_2 = 1.3$) were kept constant, while the inertia weight (w) was linearly decreased from 0.6 to 0.4 as the outer iteration progressed. The normalised IR images show a delaying effect of the best skin geometry configuration. To further understand the output skin geometry, the final shape was laser-measured, and the scans are in Fig. 4.32 and Fig. 4.33.




Figure 4.32: PSO output shape, run 1. The plot shows the error between the laser measurement and the M3J target

A first observation on the shapes is their waviness, which confirms the random inner mechanics of the PSO algorithm. The second and most important observation is on the shape itself. It seems that a very downstream hump of 1.3 mm can delay transition if preceded by a small dimple. This mechanism




Figure 4.33: PSO output shape, run 3. The plot shows the error between the laser measurement and the M3J target

needs additional studies, as the ones presented in Section 4.5.2 are insufficient. It is crucial to point out the reward function bias of yielding a higher value in the case of a downstream tripping hump compared with the initial guesses. This bias actually encourages the geometry search to start with a downstream hump and then start an optimisation. This could be the reason for such output shapes.

This preliminary optimisation platform yielded interesting results with a view to a more complete study. The main suggestions coming out of this campaign are: a more resilient reward function that avoids the bias towards downstream humps; a longer waiting time between IR images, which allows the inner particles' IR images not to be influenced by each other; and a convergence study of the PSO algorithm in terms of inner and outer iterations and PSO coefficients.

Conclusions and Recommendations

5.1. Conclusions

The aim of this thesis was to facilitate further research on smooth surface features through the design, commissioning and assessment of a new setup capable of morphing such features on the spanwise-invariant M3J swept-wing model. This chapter gathers the conclusions drawn from the design process and from the first experimental campaign, and closes with the recommendations for the future development of both the setup and the research line.

Regarding the design of the setup, the two-dimensional finite element loop proved to be an effective tool for the selection of the skin geometry. The brute-force exploration of the design parameter space, made possible by the automated mesh generation, solving and post-processing codes, led to the final *ST35-LT30-LL30-LTR10* geometry, which reproduces the target hump shapes with RMS errors lower than 0.1 mm in the simulations (see sec 4.1.1). The casting procedure was developed iteratively, addressing one issue at a time, and now yields skins with a good surface finish and a clean continuity at the leading and trailing edges. As described in Sec. 4.1.2 the sensitivity tests showed that the skin surface is not significantly modified by the flow velocity or by the heating of the halogen lamps; its shape is governed by the inner stresses of the silicone and by the actuators' displacements.

The commissioning of the setup revealed its current limitations (more details in Sec. 4.1.3 and Sec. 4.2). The most important limitation is the spanwise non-uniformity of the skin, which restricts the usable measurement region to a 15–20 cm wide range around the midspan location, where the error with respect to the target shape stays below 0.2 mm. The homing procedure of the actuators proved unreliable once the skin was installed; this problem was solved by the preset procedure, which shifts the reliability of the M3J shape from the Hall sensor readings to the laser measurements and guarantees a maximum surface error of ± 0.15 mm. The laser measurement system itself, however, showed a temporal noise attributed to the wear of the sensor glass or to speckle interference, and it is the component of the setup that most needs to be replaced and the component whose replacement would most improve the setup. Finally, the repeatability of the transition location is good for the preset and for low surface features, whereas it is less repeatable for higher features (sec 4.2.5).

Despite these limitations, the data validation gave confidence in the measurements taken at the midspan location. The preset making procedure is consistent within ± 0.1 mm, even between different days, and the influence of the angle of attack on the error evaluation is negligible for the purposes of this study. More importantly, the cross-sides and cross-campaign comparisons showed that the effect of the modified side of the wing model is limited to a constant upstream shift of the transition front, as if the amplitude of the incoming perturbations were higher from the start, with the original transition-location trends preserved (Sec. 4.2.4 and 4.2.5).

On the silicone side, it can be concluded that the setup fulfils its main purpose. The skin geometry can be morphed, the IR images acquired and the resulting surface shape verified in real time within the wind tunnel, with a turnaround time of six minutes between different geometries in a five-Reynolds

study as in this campaign. Furthermore, the setup was successfully employed as a platform for live optimisation campaigns: the preliminary PSO loop, with IR images, laser measurements and actuator positions integrated in a single software environment, found a transition-delaying shape in three out of five runs.

This first campaign also provided a set of observations on the effect of the hump shape on the transition location. It is important to emphasise once again that the primary objective of the campaign was the commissioning and validation of the setup: the aim of these observations is not to explain the underlying physics, which is deliberately left outside the scope of this work, but to detect as many trends as possible so that they can guide the research questions of future physics-oriented studies. With this in mind, the main observations are the following. The height of the hump is the most critical parameter: at a fixed Reynolds number, increasing the height progressively delays the transition front until a certain height, above which the flow is tripped quite abruptly. Since this critical height decreases with increasing Reynolds number, the existence of a critical ratio between the hump height and the boundary-layer thickness is suggested, although not yet quantified. Regarding the width of humps: wider humps keep delaying transition at higher Reynolds numbers, whereas narrower ones trip the flow. The asymmetry of humps is an important parameter as with a steeper gradient on the upstream side (negative skewness) it perform better than the symmetric and the positively skewed ones. The first take on humps placed progressively further downstream is an advancement of the transition location, as previously suggested in the literature, which seems to confirm that an interaction with instabilities still in the linear regime is more important than a suitable height. However, the hump located furthest downstream, at ($x_c = 0.31$), does not follow this trend, suggesting that further investigation is required. The most unexpected observation is the minimal influence of the reverse humps, or dimples, on the transition front described in Sec. 4.5.1. Conversely, the most promising one is the persistent transition delay produced by humps in the unforced CFI environment, at every height and Reynolds number tested, with observed delays reaching $x_c \approx 0.7$ and limited by the pressure recovery of the wing rather than by the crossflow instability itself (Sec. 4.3.3).

Overall, it can be concluded that, although the setup can still be improved, it already constitutes an efficient and versatile platform for the study of smooth surface features and their application to passive laminar flow control. The observations collected during this first campaign do not answer the physical questions, but they sharpen them, and this is precisely the role the experimental campaign was designed for.

5.2. Recommendations

The recommendations arising from the sections of this work are grouped into four areas: the hardware of the setup, the measurement procedures and data processing, the research directions on the surface features, and the optimisation loop.

Regarding the hardware, the first recommendation is the replacement of the laser measurement system, whose temporal noise was the main limitation of the surface verification during the campaign. The spanwise non-uniformity of the skin deserves a dedicated investigation: measurements of the actuator assembly without the skin installed would verify the alignment of the mechanical structure, and a new skin installation procedure, with increased attention to the spanwise alignment during the screw tightening, would assess the influence of the assembly-induced shear. In case neither of these actions eliminates the non-uniformity, the manufacture of a new skin incorporating the same 3 mm spanwise offset as the model hole is suggested. Two further practical improvements are the manufacture of new S-shaped handles for the extraction of the mechanical assembly and the automation of the preset procedure, which in this campaign had to be repeated at every power up of the system and could last as much as 45 minutes.

Regarding the procedures and the data, a dedicated study on the influence of the angle of attack at which the preset is acquired is recommended, since qualitative observations suggest that this parameter affects the results. A rework of the parametric study data is also suggested, for two reasons: to perform a more physics-oriented analysis and to correlate the IR and laser data more accurately, characterising each surface shape from the laser measurements instead of the actuators' displacements; this latter aspect proved crucial in the ratio and skewness studies, where the shapes developed on the

skin differed from the target ones. In addition, the degraded repeatability observed with higher features needs further investigation and longer settling times between consecutive data points would improve the quality of the IR images. Finally, the unclear influence of the DRE forcing amplitude calls for a more detailed study of the effective amplitudes reaching the humps.

Regarding the research directions, the quantification of the critical ratio between the hump height and the boundary-layer thickness is the natural next question raised by this campaign. The effect of the chordwise location should be extended with more positions, since the delaying behaviour observed at $x_c = 0.31$, in contrast with the advancement at $x_c = 0.22$, remains unexplained. The behaviour of the reverse hump at the higher Reynolds numbers, where it delays transition contrary to the usual advancement, deserves attention, as do the unconventional dimple–hump combinations suggested by the optimisation output, whose testing in this campaign was too limited to declare them suboptimal. In the longer term, once the humps optimal shape and position will be characterized, the steps that the research needs to approach the commercial free-flight cruise conditions are an increase of the Reynolds number by an order of magnitude and the assessment of the Mach effects.

Concerning the optimisation loop, the main recommendation is a more resilient reward function that avoids the bias towards downstream tripping humps but that promotes transition delay; an alternative approach could detect the laminar and turbulent zones through a binarisation of the images, thereby identifying the transition region directly that would solve the progressive heating problem of the laminar part of the wing model. Longer waiting times between the IR images would prevent consecutive inner particles from influencing each other, and a convergence study of the PSO algorithm, in terms of inner and outer iterations and coefficients, together with a more careful initialisation and feasibility checks that do not waste internal iterations, would turn the preliminary loop into a proper optimisation tool.

References

- [1] J. D. Anderson. *Fundamentals of Aerodynamics*. 5th ed. New York: McGraw-Hill, 2011.
- [2] A. Anger. "Including aviation in the European emissions trading scheme: Impacts on the industry, CO₂ emissions and macroeconomic activity in the EU". In: *Journal of Air Transport Management* 16 (2010), pp. 100–105. issn: 0969-6997. doi: 10.1016/j.jairtraman.2009.10.009.
- [3] A. Anscombe and N. L. Illingsworth. *Wind-tunnel observations of boundary layer transition on a wing at various angles of sweep back*. Tech. rep. ARC RM 2968. Aeronautical Research Council, 1952. url: <https://reports.aerade.cranfield.ac.uk/handle/1826.2/3531>.
- [4] Franco Baldi, Ermanno Cozzani, and Marco Filippelli. "Gas chromatography/Fourier transform infrared spectroscopy for determining traces of methane from biodegradation of methylmercury". In: *Environmental Science & Technology* 22.(7) (1988), pp. 836–839. doi: 10.1021/es00172a015.
- [5] H. Bippes. "Basic experiments on transition in three-dimensional boundary layers dominated by crossflow instability". In: *Progress in Aerospace Sciences* 35.4 (1999), pp. 363–412. doi: 10.1016/S0376-0421(99)00002-0.
- [6] H. Deyhle and H. Bippes. "Disturbance growth in an unstable three-dimensional boundary layer and its dependence on environmental conditions". In: *Journal of Fluid Mechanics* 316 (1996), pp. 73–113. doi: 10.1017/S0022112096000456.
- [7] P. C. Dörr and M. J. Kloker. "Stabilisation of a three-dimensional boundary layer by base-flow manipulation using plasma actuators". In: *Journal of Physics D: Applied Physics* 48.28 (2015), p. 285205. doi: 10.1088/0022-3727/48/28/285205.
- [8] J. M. Floryan. "On the Görtler instability of boundary layers". In: *Prog. Aerospace Sci.* 28 (1991), pp. 235–271.
- [9] M. Gaster. "A Simple Device for Preventing Turbulent Contamination on Swept Leading Edges". In: *The Aeronautical Journal* 69.659 (1965), pp. 788–789. doi: 10.1017/s0368393100081748.
- [10] T. S. Haynes and H. L. Reed. "Simulation of swept-wing vortices using nonlinear parabolized stability equations". In: *Journal of Fluid Mechanics* 405 (2000), pp. 325–349. doi: 10.1017/S0022112099007260.
- [11] T. Herbert. "Parabolized stability equations". In: *Annual Review of Fluid Mechanics* 29.1 (1997), pp. 245–283. doi: 10.1146/annurev.fluid.29.1.245.
- [12] J. L. van Ingen. "The e^N Method for Transition Prediction. Historical Review of Work at TU Delft". In: *38th AIAA Fluid Dynamics Conference and Exhibit*. AIAA Paper 2008-3830. 2008. doi: 10.2514/6.2008-3830.
- [13] R. D. Joslin. "Aircraft laminar flow control". In: *Annual Review of Fluid Mechanics* 30.1 (1998), pp. 1–29. doi: 10.1146/annurev.fluid.30.1.1.
- [14] Y. Kachanov. "Physical Mechanisms of Laminar-Boundary-Layer Transition". In: *Annual Review of Fluid Mechanics* 26.1 (1994), pp. 411–482. doi: 10.1146/annurev.fl.26.010194.002211.
- [15] D. Knörzer, J. Szodruch, and European Commission, eds. *Innovation for Sustainable Aviation in a Global Environment: Proceedings of the Sixth European Aeronautics Days, Madrid, 30 March–1 April 2011*. IOS Press, 2012. doi: 10.2777/69176.
- [16] Jonathan Lemarechal et al. "Investigation of stationary-crossflow-instability induced transition with the temperature-sensitive paint method". In: *Experimental Thermal and Fluid Science* 109 (2019), p. 109848. issn: 0894-1777. doi: <https://doi.org/10.1016/j.expthermflusci.2019.109848>. url: <https://www.sciencedirect.com/science/article/pii/S0894177718317692>.
- [17] L. Mack. "Boundary layer linear stability theory". In: *Special Course on Stability and Transition of Laminar Flows*. AGARD-R-709. 1984.

- [18] M. R. Malik, F. Li, and C.-L. Chang. "Crossflow disturbances in three-dimensional boundary layers: nonlinear development, wave interaction and secondary instability". In: *Journal of Fluid Mechanics* 268 (1994), pp. 1–36. doi: 10.1017/S0022112094001242.
- [19] Christoph Merz, Kai Richter, and M. Raffel. "Differential Infrared Thermography for Unsteady Boundary-Layer Transition Measurements". In: vol. 52. May 2014. doi: 10.2514/1.J053235.
- [20] Vinan I. Mistry. "Simulation and control of stationary crossflow vortices". Loughborough University. PhD thesis. 2014. url: <https://hdl.handle.net/2134/16323>.
- [21] Mark V. Morkovin. "On the Many Faces of Transition". In: *Viscous Drag Reduction*. Ed. by C. Sinclair Wells. Boston, MA: Springer US, 1969, pp. 1–31. isbn: 978-1-4899-5579-1.
- [22] L. Prandtl. "Über Flüssigkeitsbewegung bei kleiner Reibung". In: *Verh. III. Int. Math.-Kongr., Heidelberg*. English translation: NACA TM-452. Teubner, Leipzig, 1904, pp. 484–491.
- [23] H. L. Reed and W. S. Saric. "Stability of three-dimensional boundary layers". In: *Annual Review of Fluid Mechanics* 21.1 (1989), pp. 235–284. doi: 10.1146/annurev.fl.21.010189.001315.
- [24] M. Reibert et al. "Experiments in nonlinear saturation of stationary crossflow vortices in a swept-wing boundary layer". In: *34th AIAA Aerospace Sciences Meeting and Exhibit*. AIAA Paper 1996-184. 1996. doi: 10.2514/6.1996-184.
- [25] E. Reshotko. "Transition issues for atmospheric entry". In: *Journal of Spacecraft and Rockets* 45.2 (2008), pp. 161–164. doi: 10.2514/1.29777.
- [26] A. F. Rius-Vidales et al. "Delay of swept-wing transition using a surface hump". In: *Journal of Fluid Mechanics* 1014 (2025), A35. doi: 10.1017/jfm.2025.403.
- [27] A.F. Rius-Vidales. "Influence of a forward-facing step on crossflow instability and transition: An experimental study in a swept wing boundary-layer". English. Dissertation (TU Delft). Delft University of Technology, 2022. isbn: 978-94-6366-544-5. doi: 10.4233/uuid:6fd8a152-ab7a-4ecd-a817-61945d431bef.
- [28] W. Saric, A. Carpenter, and H. Reed. "Passive control of transition in three-dimensional boundary layers, with emphasis on discrete roughness elements". In: *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences* 369.1940 (2011), pp. 1352–1364. doi: 10.1098/rsta.2010.0368.
- [29] W. Saric, R. Carrillo Jr., and M. Reibert. "Leading-edge roughness as a transition control mechanism". In: *36th AIAA Aerospace Sciences Meeting and Exhibit*. AIAA Paper 1998-781. 1998. doi: 10.2514/6.1998-781.
- [30] W. S. Saric. "Görtler Vortices". In: *Annual Review of Fluid Mechanics* 26.1 (1994), pp. 379–409. doi: 10.1146/annurev.fl.26.010194.002115.
- [31] W. S. Saric, H. L. Reed, and E. B. White. "Stability and transition of three-dimensional boundary layers". In: *Annual Review of Fluid Mechanics* 35.1 (2003), pp. 413–440. doi: 10.1146/annurev.fluid.35.101101.161045.
- [32] H. Schlichting and K. Gersten. *Boundary-Layer Theory*. 9th ed. Berlin: Springer-Verlag, 2017.
- [33] G. Schrauf. "Status and perspectives of laminar flow". In: *The Aeronautical Journal* 109.1102 (2005), pp. 639–644. doi: 10.1017/S000192400000097X.
- [34] J. Serpieri and M. Kotsonis. "Design of a swept wing wind tunnel model for study of cross-flow instability". In: *American Institute of Aeronautics and Astronautics* (2015). doi: 10.2514/6.2015-2576.
- [35] J. Serpieri and M. Kotsonis. "Three-dimensional organisation of primary and secondary crossflow instability". In: *Journal of Fluid Mechanics* 799 (2016), pp. 200–245. doi: 10.1017/jfm.2016.379.
- [36] S. Sgouridis, P. A. Bonnefoy, and R. J. Hansman. "Air transportation in a carbon constrained world: Long-term dynamics of policies and strategies for mitigating the carbon footprint of commercial aviation". In: *Transportation Research Part A: Policy and Practice* 45 (2011), pp. 1077–1091. issn: 0965-8564. doi: 10.1016/j.tra.2010.03.019.
- [37] A. M. O. Smith and N. Gamberoni. *Transition, Pressure Gradient and Stability Theory*. Tech. rep. El Segundo Division: Douglas Aircraft Company, Aeronautical Research Council, 1956.

- [38] S. Westerbeek, J. Casacuberta, and M. Kotsonis. "Linear and nonlinear interactions between stationary cross-flow instabilities and a smooth surface hump". In: *Journal of Fluid Mechanics* 1026 (2026), A43. doi: 10.1017/jfm.2025.11042.
- [39] S. Westerbeek et al. "Linear and Nonlinear Stability Analysis of a Three-Dimensional Boundary Layer over a Hump". In: *AIAA SCITECH 2023 Forum*. AIAA Paper 2023-0678. 2023. doi: 10.2514/6.2023-0678.
- [40] E. White et al. "Stages of swept-wing transition". In: *39th AIAA Aerospace Sciences Meeting and Exhibit*. AIAA Paper 2001-271. 2001. doi: 10.2514/6.2001-271.
- [41] E. B. White and W. S. Saric. "Secondary instability of crossflow vortices". In: *Journal of Fluid Mechanics* 525 (2005), pp. 275–308. doi: 10.1017/S002211200400268X.
- [42] G. Zoppini et al. "The near wake of discrete roughness elements on swept wings". In: *Journal of Fluid Mechanics* 960 (2023). doi: 10.1017/jfm.2023.141.