Framework for Numerical Simulation of Crossflow Development and Transition to Turbulence on Swept Wings.

AE5110 Aerodynamics & Wind Energy Thesis

Florian Adam

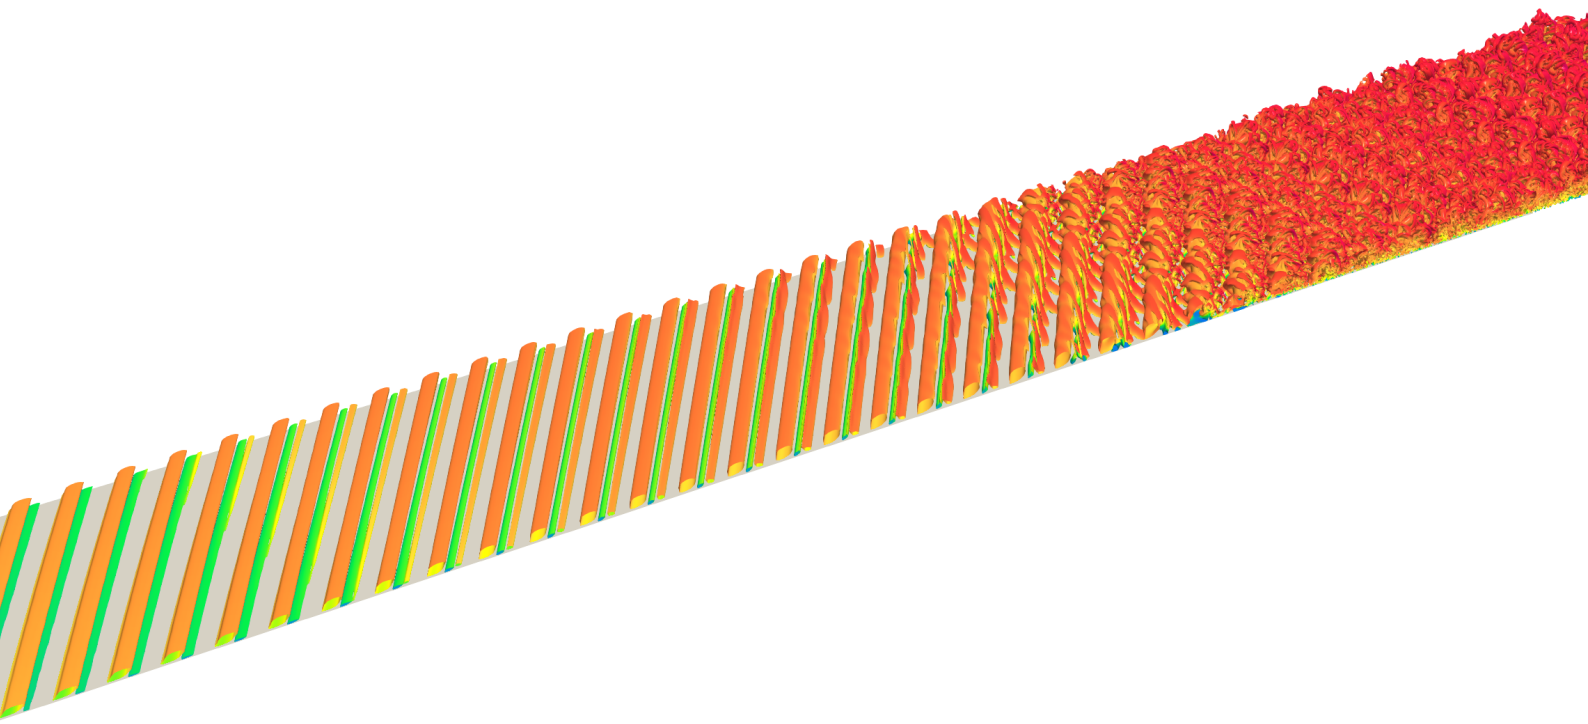

TU Delft

Framework for Numerical Simulation of Crossflow Development and Transition to Turbulence on Swept Wings.

by

Florian Adam

to obtain the degree of Master of Science
at the Delft University of Technology,
to be defended publicly on Wednesday August 23, 2023 at 10:00 AM.

Student number: 4560973

Thesis committee:

Prof. dr. ir. S. Hickel,	TU Delft
Dr. ir. M. Kotsonis,	TU Delft
Dr. ir. D. Ragni,	TU Delft
Dr. ir. A. F. Rius Vidales,	TU Delft
ir. J. Casacuberta Puig	TU Delft

TU Delft

Acknowledgements

This endeavour would not have been possible without my group of supervisors, Prof. Dr. Stefan Hickel, Dr. Marios Kotsonis, Dr. Alberto Felipe Rius-Vidales and Jordi Casacuberta Puig. Throughout this challenging project, involving both numerical and experimental components, they supported me through numerous group discussions and brainstorming sessions. Thanks to the guidance of every single one of you, I was able to meet the goals set out for this research. To Jordi, I owe a special thanks for being my first point of contact during this work. Your availability and openness to discussing new ideas at any time during the past year were a major contribution to the outcome of this thesis.

I would also like to extend my sincere thanks to my long-term roommates over the years at TU Delft. Bart, Arthur, Thomas, Mathieu and Sebastien, thank you for pulling me out of my books and from behind my computer at times, to be there to discuss ideas and brainstorm about the future, for the late-night (and sometimes early morning) therapy sessions on the roof, and just purely for the amazing friendship build during the roller-coaster that is life at TU Delft.

Lastly, I would be remiss not to mention my parents and sister, supporting me at home through all my years in Delft. Without them, I would never have grown into the man I am today, with all the tools necessary to continue life after academics and much more. Although choosing an unconventional track through university for a Belgian engineering student, they never stopped encouraging me during all the endeavours I embarked on.

Florian Adam
Delft, August 2023

Abstract

Maintaining laminar flow on large swept surfaces of subsonic transport aircraft, i.e. the wings and the stabilisers, is currently posing a considerable challenge for aerodynamic design. Improving the efficiency of aircraft by delaying or removing the laminar-to-turbulent transition process over the wing and tail parts can substantially reduce contaminant emissions. The dominant flow instability causing laminar-turbulent transition of swept-wing flow is the so-called crossflow instability (CFI). Ongoing research at TU Delft has shown potential to delay transition by use of passive mechanisms. As such, a framework has been designed to numerically compute crossflow development and transition to turbulence on swept wings. Through the use of experimental data acquired in wind-tunnel measurements at TU Delft, the CFI development and transition process on swept wings has been modelled numerically by means of Direct Numerical Simulation (DNS). Based on a DNS laminar flow field generated from the pressure distribution along the model surface, a numerical primary CFI mode in good agreement with the experiment was obtained through Non-linear Parabolized Stability Equations (NPSE). Following this steady flow field analysis, the simulation was made unsteady by the implementation of numerical free-stream turbulence. This novel method resulted in unprecedented modelling of the receptivity mechanisms of transition in three-dimensional crossflow cases, overcoming ad-hoc treatments. Both experimental and numerical flow fields indicated a Type-I dominant secondary CFI (i.e. K-H type of response in the laterally inclined shear layer of the stationary crossflow vortex), which consequently carries the formation of near-wall hairpins and ultimately turbulence. Crossflow vortex frequency content also agrees well in the low-frequency band ($450\text{Hz} \leq f \leq 3000\text{Hz}$), whilst the numerical high-frequency content ($3500\text{Hz} \leq f \leq 9000Hz$) does show a distinct delay in amplitude growth throughout the majority of the transition region. Contradicting the promising qualitative analysis of the free-stream turbulence methodology, this discrepancy in the frequency spectrum indicates the major shortcoming in the numerical setup, which was shown to be biased towards introducing more low-frequency disturbances at the inflow boundary.

List of Figures

Nomenclature

Abbreviations

ASU	Arizona State University
BL	Boundary Layer
CF	Crossflow
CFI	Crossflow Instability
CFD	Computational Fluid Dynamics
DHPC	Delft High-Performance Computing Centre
DLR	Deutsche Zentrum für Luft- und Raumfahrt
DNS	Direct Numerical Simulation
DRE	Discrete Roughness Element
FFS	Forward-Facing Step
FSC	Falkner-Skan Cooke
FST	Free-Stream Turbulence
HPC	High-Performance Computing
HWA	Hot-Wire Anemometry
ILST	Incompressible Local Stability Equations
IR	Infrared
K-H	Kelvin-Helmholtz
KTH	Royal Institute of Technology
LES	Large Eddy Simulation
LFC	Laminar Flow Control
LTT	Low Turbulence Tunnel
MFD	Mean Flow Distortion
NASA	National Aeronautics and Space Administration
NPSE	Nonlinear Parabolized Stability Equations
OS	Orr-Sommerfeld
PIV	Particle Image Velocimetry
PSD	Power Spectral Density
PSE	Parabolized Stability Equations
SFD	Selective Frequency Damping
SWIFT	Swept Wing In Flight Testing
TS	Tollmien-Schlichting

Symbols

a	Unsteady disturbance amplitude
A	Steady disturbance amplitude
B_L	Low-frequency band
B_H	High-frequency band
C_f	Skin friction coefficient
c_p	Pressure Coefficient
c_X	Airfoil chord length (leading-edge-normal direction)
E_v	Kinetic Energy
f	Frequency
G	Görtler constant
L_{11}	Turbulence integral length scale
n	Amplification factor
N_p	Number of points per sphere
p	Pressure
p_q	Dynamic pressure
q	State velocity vector
Q	Velocity measured by the Hot-wire sensor
Re	Reynolds Number
R	Radius of curvature
t	Time
T	Temperature
Tu	Turbulence Intensity
u_τ	Friction velocity
u,v,w	Leading edge oriented velocity components
U,V,W	Wind tunnel oriented velocity components
x,y,z	Leading edge oriented coordinate system
X,Y,Z	Wind tunnel oriented coordinate system
α	Streamwise wavenumber
α_i	Disturbance spatial growth rate
β	Spanwise wavenumber
δ_0	Boundary layer thickness at numerical inflow boundary
δ_{99}	Boundary layer thickness
δ^*	Boundary layer displacement thickness
Δ	Time constant
ϵ	Error
γ	Wall-normal wavenumber
κ	Wavenumber magnitude
Λ	Sweep angle
λ	Wavelength
μ	Dynamic viscosity
ν	Kinematic viscosity
ω	Angular velocity
ϕ	Eigenfunction
ρ	Density
σ	Standard deviation
θ	Boundary layer momentum thickness
τ	Wall-shear stress
χ	Feedback control coefficient

Accents and Subscripts

x_B	Baseflow quantity
x_{DB}	Distorted Baseflow quantity
x_w	Quantity at the wall
$\hat{x}$	Shape function
x'	Perturbation component
$\overline{x}$	Time-average quantity
x^+	Non-dimensional half mesh cell-size
x_∞	Quantity related to the free-stream
x_e	Quantity related to the edge of the boundary layer
x_0	Quantity at the inflow boundary of the numerical domain
$\langle x \rangle_z$	Spanwise root-mean-square of quantity
x^*	Complex conjugate of quantity
$\mathbf{x}$	Vector quantity
$\dot{x}$	Time-derivative of quantity
$\tilde{x}$	Eigenfunction
x^q	Quantity related to the Q velocity component
x^u	Quantity related to the u velocity component

Contents

1 Introduction

Aerodynamics is what makes flying possible, and has always been a crucial factor in aircraft design. Focus in the past was on generating lift but has since shifted towards efficiency, implying the need for drag reduction. This drag reduction is preferably achieved without impacting other aircraft characteristics, minimising the need for compromises in aircraft design. The total drag an aircraft generates is made up of several components (i.e. pressure drag, skin-friction drag, lift-induced drag). Reducing pressure drag would result in major structural design changes, and whereas lift-induced drag can be minimised over non-lifting surfaces, the majority contribution from the main wings is directly coupled to the lift needed to fly. Skin friction drag, on the other hand, depends less on the overall aircraft configuration. With approximately half of the total drag attributed to skin friction drag, it is a potentially large improvement as well (Szodruch, 1991). The contribution of skin friction drag is directly correlated with the state of the boundary layer. Turbulent boundary layers generate more drag than their laminar counterparts, so laminar flow is favoured over most aircraft surfaces. For fully laminar flow wings to become reality, a better understanding of boundary layer transition is desired. On swept wings, the most prevalent instabilities causing transition are Tollmien-Schlichting (TS) waves and crossflow instabilities (CFI). Where the former is found more dominant at low sweep angles, in boundary layers with a zero or adverse (positive) pressure gradient, CFIs are found in combination with higher sweep angles and a favourable (negative) pressure gradient (Borodulin, Ivanov, Mischenko, et al., 2019). With increased wing sweep angle leading to drag and stability benefits in transonic and supersonic aircraft, an improved understanding of CFI-based transition can become the key to designing fully laminar wings, or laminar flow control (LFC) devices for next-generation aircraft (Donlan & Weil, 1952). Which, once achieved, can significantly reduce aircraft emissions.

1.1. Boundary-Layer Flow

The boundary layer is the region between the bulk fluid flow and the surface of an object, which both move at different rates of speed. The thickness of this layer can be defined as the wall-normal distance up to where the velocity reaches 99% of the bulk fluid flow, referred to as δ_{99}. Another definition commonly used is displacement thickness or δ^*, based on the mass-flow deficit compared to inviscid flow. It describes the wall-normal distance the wall would have to move in inviscid flow with slip boundary conditions, to obtain the same mass-flow rate as the viscid flow field (Lighthill, 1958). Prandtl (1904) was the first to describe these boundary layer concepts and introduced a set of equations to define the boundary layer mathematically. By reducing the three-dimensional Navier-Stokes equations based on the assumption of steady, incompressible flow with constant viscosity and density, the following set of equations resulted,

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} + \frac{\partial w}{\partial z} = 0, \tag{1.1}$$

$$u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} + w\frac{\partial u}{\partial z} = -\frac{1}{\rho}\frac{\partial p}{\partial x} + \nu\left(\frac{\partial^2 u}{\partial x^2} + \frac{\partial^2 u}{\partial y^2} + \frac{\partial^2 u}{\partial z^2}\right). \tag{1.2}$$

$$u\frac{\partial v}{\partial x} + v\frac{\partial v}{\partial y} + w\frac{\partial v}{\partial z} = -\frac{1}{\rho}\frac{\partial p}{\partial y} + \nu\left(\frac{\partial^2 v}{\partial x^2} + \frac{\partial^2 v}{\partial y^2} + \frac{\partial^2 v}{\partial z^2}\right), \tag{1.3}$$

$$u\frac{\partial w}{\partial x} + v\frac{\partial w}{\partial y} + w\frac{\partial w}{\partial z} = -\frac{1}{\rho}\frac{\partial p}{\partial z} + \nu\left(\frac{\partial^2 w}{\partial x^2} + \frac{\partial^2 w}{\partial y^2} + \frac{\partial^2 w}{\partial z^2}\right), \tag{1.4}$$

Referred to as the boundary layer equations, this set of partial differential equations can be simplified with several assumptions. Assuming an infinite swept wing model in incompressible flow, and performing scale analysis, reduces the boundary layer equations to,

$$\frac{\partial u}{\partial x} + \frac{\partial v}{\partial y} = 0, \tag{1.5}$$

$$u\frac{\partial u}{\partial x} + v\frac{\partial u}{\partial y} = -\frac{1}{\rho}\frac{\partial p}{\partial x} + \nu\left(\frac{\partial^2 u}{\partial y^2}\right). \tag{1.6}$$

$$\frac{\partial p}{\partial y} = 0, \tag{1.7}$$

$$u\frac{\partial w}{\partial x} + v\frac{\partial w}{\partial y} = \nu\left(\frac{\partial^2 w}{\partial y^2}\right), \tag{1.8}$$

indicating the lack of wall-normal (y_t) and spanwise (z) pressure gradients. Based on these boundary layer equations, Blasius (1908) developed the Blasius boundary layer equations, describing the steady two-dimensional boundary layer over a flat plate where $\partial p/\partial x = 0$. Falkner and Skan (1931) later extended the Blasius boundary layer to flow over wedges. These boundary layer approximations are a computationally cheap but fairly accurate way of obtaining a laminar boundary layer and allow validation of the DNS baseflow during this research.

This boundary layer theory also allows the calculation of the velocity profiles inside the boundary layer, because the difference in momentum between the surface and the fluid inevitably generates a velocity gradient that introduces shear forces in the boundary layer and at the wall. These shear forces are what quantify the skin friction drag. A measure for the shear at the wall is given by Equation 1.9, where τ_w is the shear stress at the wall, μ the dynamic viscosity of the fluid, and $\partial u/\partial y$ the velocity gradient at the wall.

$$\tau_w = \mu\left(\frac{\partial u}{\partial y}\right)_{y=0} \tag{1.9}$$

Equation 1.9 points out that skin friction drag depends almost entirely on the velocity gradient at the wall. This gradient is in turn dependent on the boundary layer state, being either laminar or turbulent. Under laminar conditions, fluid particles move quasi-parallel to each other with very little mixing, resulting in a very smooth flow and a low boundary layer thickness. In this type of flow, the velocity gradient is relatively well distributed along the entire boundary layer height, resulting in low shear forces at the wall. Turbulent boundary layers then contain excessive particle mixing, with fluid particles moving very non-parallel. The chaotic-looking motion of the particles creates a thicker boundary layer, with a peak velocity gradient close to the wall. This results in a considerably higher shear force at the wall than in a laminar boundary layer. Therefore, to reduce overall drag, aircraft designers prefer laminar boundary layers and want to avoid or delay transition to a turbulent state as much as possible. Both types of boundary layers, as well as the transition region, are visualised in Figure 1.1.

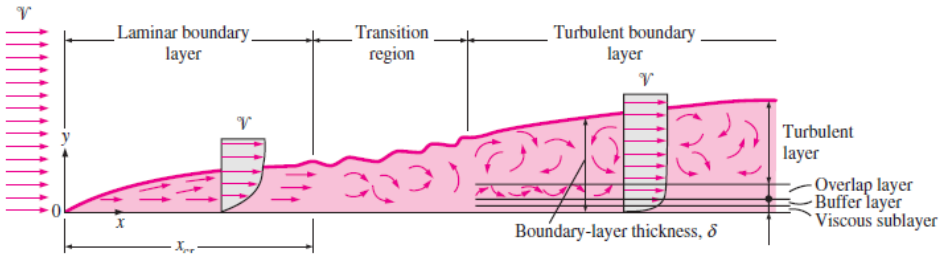

Figure 1.1: Boundary layer profiles for laminar and turbulent flow states (Çengel & Cimbala, 2006).

1.2. Hydrodynamic Stability

The behaviour of a boundary layer can be analysed by assessing its hydrodynamic stability. That is the reaction of the baseflow to superimposed external disturbances, where baseflow refers to the fully laminar steady-state solution to the Navier-Stokes equations. When disturbances in the baseflow are allowed to grow indefinitely, the flow is considered unstable, while stable flows damp out these external disturbances. The assessment of the flow stability can be performed based on two criteria. If the Reynolds number over an aircraft wing is sufficiently below the range of $1.5x10^5$ to $1.5x10^6$, the flow can be considered stable (Sforza, 2014). If this is not the case, then kinetic energy is used to evaluate stability (Joseph, 1976; Schmid & Henningson, 2001). Defining a flow asymptotically stable when the kinetic energy of the disturbance over time goes to zero,

$$\lim_{x\to\infty} \frac{E_v(t)}{E_v(0)} \to 0. \tag{1.10}$$

Or more simply, the flow can also be seen as monotonically stable when the change in kinetic energy of the disturbance over time is smaller than zero for all time steps,

$$\frac{dE_v}{dt} < 0 \quad \text{for all } t > 0. \tag{1.11}$$

Many different types of hydrodynamic instabilities exist, though only a handful are relevant for subsonic flow over swept wings (Borodulin, Ivanov, Kachanov, et al., 2019). The attachment line instability, Tollmien-Schlichting waves, and crossflow instabilities (Hall et al., 1984; Ustinov, 2013). While all of these can lead to transition, boundary layer characteristics determine which are allowed to grow (Borodulin, Ivanov, Mischenko, et al., 2019). The leading edge shape and convex airfoil surface along the pressure side, combined with the negative pressure gradient, can avoid the unsteady growth of attachment line instabilities and T-S waves (Saric et al., 2011). Although TS waves do occur frequently on wings with smaller sweep angles and adverse pressure gradients, crossflow instabilities are the dominant cause of transition on highly swept wings with favourable pressure gradients, therefore becoming a hydrodynamic instability of particular interest to the aviation industry (Borodulin, Ivanov, Mischenko, et al., 2019). Independent of the stability of the boundary layer, the disturbance growth can be simulated either through CFD, or analytic stability tools. Various levels of accuracy can be considered depending on the type of flow, from the fast but limited Incompressible Linear Stability Theory (ILST) to the more elaborate but accurate Non-linear Parabolized Stability Equations (NPSE) (Westerbeek, 2020). The range of different stability tools utilised throughout this work are outlined in detail in chapter 3.

1.2.1. Crossflow Instability

The second half of the 20^{th} century sees the emergence of many new transonic and supersonic aircraft. Accompanying this aviation boom is a rise in aircraft-related pollution. With climate change becoming a point of concern in more recent years, the research efforts to reduce emissions were sped up by the idea of laminar flow control (LFC), reducing the skin friction drag of these aircraft. To design such devices, a better understanding of the mechanisms behind transition to turbulence is required, generating a rise in research into this field. With modern aircraft increasingly making use of swept wings, the discoveries made by Gray (1952a, 1952b) re-gains focus. Gray (1952a, 1952b) found that when introducing sweep to wings, the transition location moves closer to the leading edge, and an instability referred to as crossflow appears. In crossflow-dominated boundary layers, streamlines follow three-dimensional trajectories around swept wings. These streamlines can be considered inviscid based on Prandtl (1904), establishing that viscosity effects are only evident near the wing surface (i.e. inside the boundary layer). The pressure forces along the inviscid streamlines are therefore balanced by the centrifugal forces pulling on the flow particles. This force balance along the wing chord results in a three-dimensional streamline showing a curved pattern, as visualised in Figure 1.2.

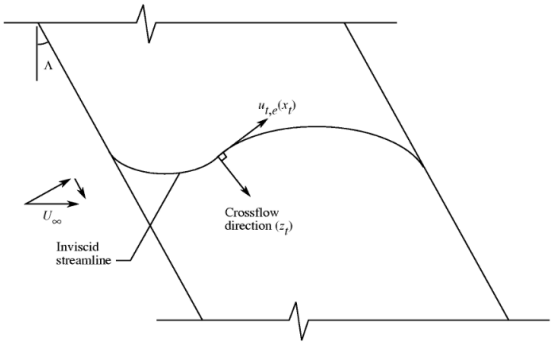

Figure 1.2: Inviscid streamline over a section of a swept wing (Dagenhart & Saric, 1999).

When viscid forces become relevant near the surface (i.e. inside the boundary layer), this equilibrium no longer holds. The momentum deficit when approaching the no-slip wall allows the forces on the flow to generate a velocity component perpendicular to the inviscid streamline. This component becomes more dominant with the increase in momentum deficit but reduces to zero at the wall due to the no-slip condition. These two velocity components are visualised in Figure 1.3, where the tangential component follows the inviscid streamline direction and the crossflow component moves perpendicular to it. The inflection point generated by the velocity vector is a prime point for instability introduction.

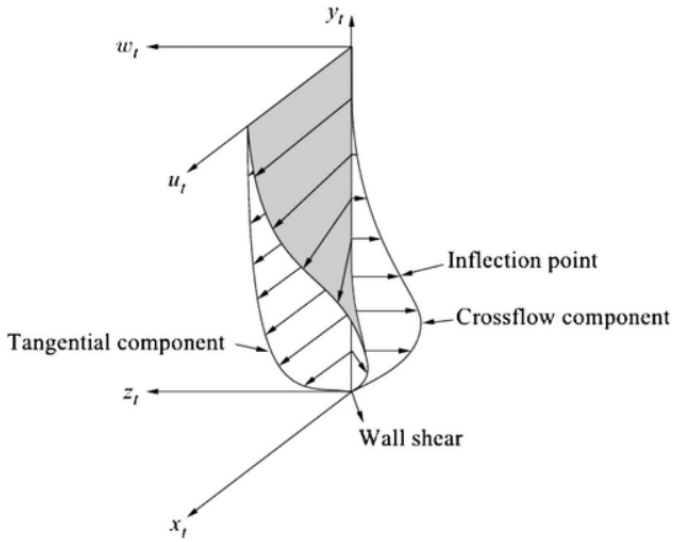

Figure 1.3: Illustration of the individual crossflow velocity components (Serpieri, 2018).

Primary Crossflow Instability

Arnal et al. (1984) and Poll (1985) showed that crossflow consists of travelling and stationary crossflow vortices, growing in amplitude in the downstream direction. Through the receptivity of steady disturbances, crossflow vortices are introduced into the boundary layer. These vortices create alternating upwelling and downwelling regions in the boundary layer, generating a wave-like boundary layer spanwise modulation. Referred to as the primary crossflow instabilities, these crossflow vortices can either remain fixed in place or travel along the span. What ratio of stationary to travelling vortices is present depends largely on the turbulence intensity and surface roughness encountered. High free-stream turbulence combined with low surface roughness tends to favour travelling crossflow vortices, while low free-stream turbulence environments produce stationary vortices with amplitudes depending on the surface roughness level (Deyhle & Bippes, 1996). The latter property can be used in research to artificially impose stationary crossflow vortices with a specific amplitude and fixed fundamental wavelength in the boundary layer, through tools such as DREs. These disturbances tend to grow when propagating downstream, at varying speeds depending on the flow conditions. Growth continues until a saturation

amplitude is reached and breakdown occurs. The breakdown into turbulence is not initiated by the primary CFI but, instead, occurs due to the development of unsteady high-frequency secondary CFI (Joslin & Streett, 1994; Kohama et al., 1991).

Secondary Crossflow Instability
Streamwise velocity distortions generated by the primary crossflow instability can locally create doubly inflected velocity profiles. These double inflection points are seen as attachment points for secondary crossflow instabilities which are introduced by external disturbances such as free-stream turbulence, vibration or surface roughness, through the process of receptivity. Secondary crossflow instabilities are high-frequency instabilities that tend to grow exponentially when introduced into the boundary layer, often causing breakdown shortly after their first occurrence in the flow field. In general, secondary instabilities can manifest themselves in three different ways. As type-I, type-II or type-III instabilities, depending on their location in the boundary layer (Fischer & Dallmann, 1991; Malik et al., 1999; Serpieri & Kotsonis, 2016; Wassermann & Kloker, 2002). Type-I-driven transition is driven by the shear forces in spanwise direction, and the related secondary CFI grow in the upwelling region of the crossflow vortex. Type-II instabilities can be found near the top of the crossflow vortex, and are dependent on wall-normal shear. Type-III instabilities are then driven by the interaction of travelling and stationary CF vortices, so their presence is largely dependent on the ratio of these two vortex types. Type-III instabilities present themselves near the lowest part of the vortex, in the wave trough (Casacuberta, Hickel, et al., 2022).

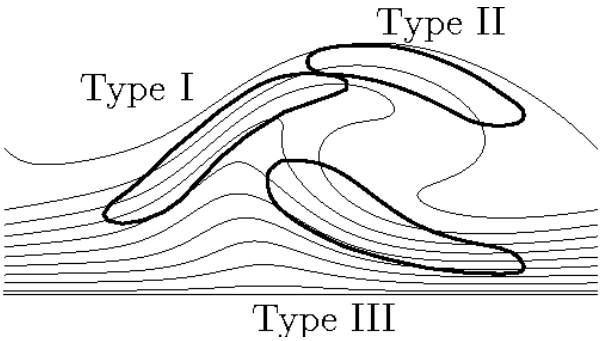

Figure 1.4: CFI type locations with respect to the CF vortex (Friederich, 2013).

These secondary instabilities complete the crossflow-dominated boundary layer, which consists of three main components,

$$q(x,y,z,t) = \underbrace{q_{\mathrm{B}}(x,y) + \dot{q}(x,y,z)}_{q_{\mathrm{DB}}} + q'(x,y,z,t) \tag{1.12}$$

where q_B refers to the unperturbed baseflow, $\dot{q}$ represents the primary crossflow instability, and $q^{'}$ the secondary instability. The baseflow and primary crossflow combined are often labelled the distorted baseflow, or q_{DB} (Casacuberta, Groot, et al., 2022).

1.2.2. Recent Scientific Advances

Crossflow Instability Characteristics
While computational power was still relatively limited in the 1990s, the main numerical improvements during this decade were linked to stability theory and achieved with the help of major experiments at institutions such as DLR and ASU. These experiments set the standard for crossflow wind tunnel testing for many years, using swept flat plates at first, to later shift to infinite swept wing models with enough flow measurement equipment to set up a numerical comparison (Deyhle et al., 1993; Ng & Crouch, 1999; Reibert et al., 1996). Several inadequacies in linear theory, and its limitations, were addressed during this time Deyhle et al. (1993), Reed et al. (1996), and Wassermann and Kloker (2002). With CFI becoming common knowledge by the end of the last century, experimental work shifted to understanding CFI interaction with irregularities such as roughness and steps. These are systematic ways to replicate any irregularities on aircraft wings encountered during flight. An in-depth understanding of this type of interaction could ultimately lead to better control over CFI growth. Deyhle and Bippes (1996) investigated the effect of roughness, finding that rough surfaces and low turbulence intensity favour stationary

crossflow vortices. Smooth surfaces and high turbulence then encourage travelling crossflow vortices. This knowledge then led to the creation of Direct Roughness Elements (DRE), where small roughness elements near the leading were used to force stationary crossflow vortices with a fixed fundamental wavelength (Radeztsky, 1994). In later years, DREs were commonly employed to enforce specific crossflow scenarios (Rius-Vidales & Kotsonis, 2022; Rizzetta et al., 2010; Saric et al., 2011). Forward- and backwards-facing steps are another common technique to investigate crossflow behaviour, creating a repeatable way to replicate discontinuities such as rough or irregular surfaces. Rius-Vidales and Kotsonis (2021) and Rius-Vidales and Kotsonis (2022) showed the impact of Forward-Facing Steps (FFS) on the flow field behaviour, and the large dependency on the step height. Although Rius-Vidales and Kotsonis (2022) found that step heights which cause a significant upstream shift of the transition front show larger temporal velocity fluctuations not linked to any of the three CFI types, the lack of spatially-resolved measurements limited the conclusions that could be drawn from this experiment to topological mechanisms. Even with the large improvements made in measurement capabilities over recent years, there is still a need for more experimental data. The development of techniques such as PIV resulted in large jumps in experimental capabilities, but the current state-of-the-art could still benefit from the addition of spatially and temporally resolved numerical data.

Laminar Flow Control

In terms of applicable laminar flow control, detailed knowledge of flow physics is still deemed inadequate to effectively design LFC devices. Out-of-the-box ideas were brought forward, like the SWIFT experiment, which performed flight tests to study the impact of DREs with a swept wing section attached underneath a small aircraft (Saric et al., 2015). And although 30 of the 63 tests resulted in usable data, the inconsistency of the results ultimately led to the tests being unable to confirm the use of DREs as a viable laminar flow control technique. Afterwards, Saric et al. (2019) pointed towards the need for more computational research into LFC devices before they can become a reality. With computational power increased, and the capability to run DNS becoming more accessible, this numerical research becomes more and more of a reality as well. However, the questions can be asked if LFC testing in CFD might be too artificial at this point in time. A structural approach to numerical research is often taken, through fundamental CFD of simple surfaces and interaction with artefacts like steps or roughness elements. Casacuberta et al. (2021) used the FFS approach to further investigate the flow mechanism over steps numerically, in a similar fashion as Rius-Vidales and Kotsonis (2021) did experimentally. Casacuberta et al. (2021) found that crossflow disturbances are amplified upstream of the step. The disturbances get lifted up and a flow reversal region is created in front of the step. For larger step cases, a second region of flow reversal is discovered downstream of the step, generating secondary inflection points. The step also generates a secondary peak in the disturbance shape function with an amplitude depending on the step height. The step shows stabilising properties if the step height is sufficiently large, but mainly increases the sub-harmonic disturbance content in the boundary layer downstream of the step. The mechanisms behind the flow over steps are discovered gradually due to structure research work of this type. However, techniques used in CFD to trigger transition and instability growth in research like (Casacuberta, Hickel, et al., 2022) are often not fully representative of what happens in flight and wind tunnel conditions. Forcing approaches that introduce disturbances at the wall, such as Blowing and Suction or random forcing techniques, are not present in wind tunnels, nor on aircraft wings (Casacuberta et al., 2021; Hader & Fasel, 2022; Högberg & Henningson, 1998). More representative ways of triggering instabilities in CFD are slowly being created, such as acoustic waves in compressible flows, or free-stream turbulence. The latter can be of interest in subsonic crossflow instability research but will need more proof to become the standard practice. So work continues towards a better understanding of crossflow-dominated boundary layers, but it is clear a crucial piece of the puzzle is missing.

Numerical Equivalent of Experimental Setups

More capabilities are desired to advance knowledge in this field even further. Stand-alone numerical or experimental research can be limited in this context. Using experimental data to validate numerical concepts and tools has been done frequently (Deyhle & Bippes, 1996; Deyhle et al., 1993; Reed et al., 1996). The other way, using numerical data as an addition to experiments much less. Most often, the numerical data is only used as a validation of what the experiments measured, and not as a complementary data set (Arnal et al., 1998; Reibert et al., 1996). With the recent improvements

in DNS capabilities, it would be useful to combine both of these approaches to increase the amount of data available for each experimental setup. By using the experimental data to set up an exact numerical replication, the numerical setup is validated. Then using the numerical time-dependent three-dimensional data as an addition to the experimental data would allow for a significantly larger data set to analyse crossflow-driven transition behaviour. Nishino and Shariff (2010) is thought of to be the first that attempted to match an experiment exactly, using DNS. The preliminary results presented by Nishino and Shariff (2010) set a benchmark for the closest match at that time, although the disturbance amplitudes seen in the DNS were still orders of magnitude removed from experimental measurements. The possible causes for this discrepancy were thought to be related to the simulation and receptivity of the roughness elements, as well as the lack of free-stream turbulence and insufficient grid resolution. Tempelmann et al. (2011) greatly improved this benchmark only a year later. Although a significant leap in results compared to Nishino and Shariff (2010), the disturbance amplitude of the steady crossflow vortices reported by Tempelmann et al. (2011) was only 40% of that seen in the series of experiments performed at ASU (Reibert et al., 1996; Saric et al., 1998; White & Saric, 2005). Numerical matching of crossflow experiments was still far from perfect at that point. Tempelmann et al. (2011) also attributed the differences to the lack of free-stream turbulence in the CFD, as well as the potential difference in roughness over the airfoil surface. Other than Nishino and Shariff (2010), Tempelmann et al. (2011) also pointed to small discrepancies in the pressure gradient near the DRE section of the surface, which could explain the difference further downstream of the DRE. Up until this point, to the author's knowledge, this is the closest match achieved through DNS. Borodulin, Ivanov, Kachanov, et al. (2019) later managed to closely replicate the stability characteristics of unsteady crossflow modes over a real airfoil, through the use of boundary layer theory and Parabolized Stability Equations (PSE). However, this research focused on using stability theory to match stability characteristics and generated a large leap forward in capabilities. Unfortunately, PSE cannot simulate flow developments close to the transition region, and the need remains for an exact experiment replication with DNS. This way, the mechanisms of complex flow fields near the transition region observed during experiments can be studied in more detail thanks to the additional numerical data.

1.3. Laminar-Turbulent Transition Prediction

Laminar-to-turbulent transition describes the process of a laminar boundary layer turning turbulent. Without a fixed path to turbulence, there are several types of transition possible, depending on the flow field type and parameters. Figure 1.5 shows the most prevalent options, where path **a** indicates the transition type of interest in this work. Through receptivity of external disturbances, primary instabilities start to grow inside the boundary layer. Receptivity is the amount of influence an external disturbance has on the boundary layer, and how much it affects its stability over time. External disturbances can be presented in the form of free-stream turbulence, vibrations, roughness, particles, or sound waves. Through this receptivity process, instabilities can grow, which in this case are the primary and secondary CFI. The primary instability is allowed to grow until saturated, after which the secondary CFI trigger breakdown and transition to turbulence. As the primary CFI does not directly lead to breakdown by itself, a high-frequency secondary CFI is required to complete this step. This high-frequency disturbance tends to grow exponentially and quickly leads to transition. Predicting where this transition will occur is most accurate through experimental measurements, or numerically with DNS. Because this is an expensive process in terms of cost and time, several metrics of a flow field can aid in predicting this transition location more easily. Most fundamental is the spatial growth rate of the disturbance α_i, derived from the complex part of the streamwise wavenumber. This disturbance growth rate can be linked to the disturbance amplitude a as,

$$\frac{1}{a^q_{(k,j)}} \frac{da^q_{(k,j)}}{dx} = -\alpha^q_{i,(k,j)}, \tag{1.13}$$

with subscripts k and j referring to temporal frequency and spanwise wavenumber disturbance modes, respectively. q denotes the velocity vector, with its three components u, v, and w (Casacuberta, Groot, et al., 2022). To the disturbance growth rate, both the amplitude a and the amplification factor n can be linked as,

$$n = \ln \frac{a}{a_0} = \int_{x_0}^{x} -\alpha_i dx, \tag{1.14}$$

where a_0 and x_0 refer to the first known amplitude and its respective location along the x-axis. The n-factor can be referred to in context of a location, or in a frequency domain. It is a measure used to keep track of disturbance growth for individual frequencies and structures.

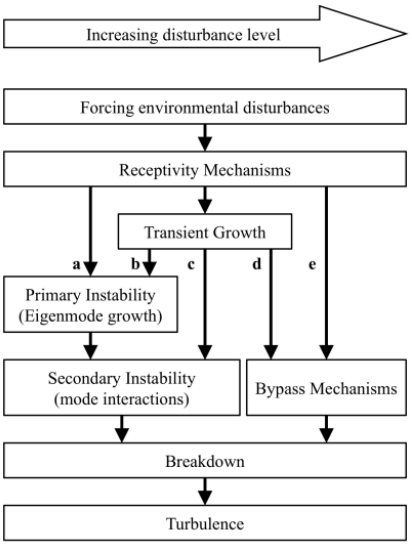

Figure 1.5: Different pathways from laminar to turbulent boundary layers, where path **a** indicates the transition type of interest for this research (Meersman et al., 2018).

Even though the transition location can be estimated using linear stability tools, such as through the e^N-method developed by Van Ingen (2008), it remains a highly non-linear process. As the name suggests, it uses the amplification factor n to calculate the amplification ratio e^N. When the n-factor reaches a predetermined value, transition is expected to occur. Correctly applied, the results of this method are remarkably similar to transition locations observed in experiments, while still being based on linear stability theory. This implies a large number of assumptions are applied (i.e. parallel flow, small disturbances, streamwise invariant velocity distribution). Although many extensions of the method have been developed, it remains hard to keep up with the latest CFD developments because of the simplifications involved. So while very computationally expensive, DNS still remains the most accurate way of estimating transition.

1.4. Motivation

Separate experimental and numerical research has contributed significantly to the understanding of crossflow-driven laminar-to-turbulent transition on swept wings. Despite being achieved before (see for instance Serpieri and Kotsonis (2016)), experimental measurement capabilities of three-dimensional volumes inside this type of boundary layer remain limited with respect to numerical simulations. In spite of the increasing availability of measurement tools such as PIV on the boundary layer scale and time-resolved HWA measurements, the lack of inter-dependency between measurement points often found in experiments tends to make it difficult to extract all the desired information from the flow field. Within numerical simulations, on the other hand, temporally and spatially correlated three-dimensional flow fields can be extracted much more easily. However, there remains a large dependency on the simulation setup. Wind tunnels are designed to develop certain flow environments throughout their test sections, hence numerical models should reproduce these expected mechanisms to allow for direct comparison with the wind tunnel data. For this to occur, boundary conditions representative of the wind tunnel environment needs to be generated numerically. Stand-alone numerical research, therefore, requires careful consideration of the applicability of the conclusions to different flow environments. When a framework would now exist to build a numerical simulation representative of the experimental CFI development and transition, it would give more insight into the driving mechanisms by combining both numerical and experimental data sets. The experimental research would have additional three-dimensional and time-dependent data to draw additional conclusions from their experiments, and the numerical setup can be validated by the experimental measurements, allowing for more trustworthy

results. Although previous attempts have not been fully successful in reaching replication (see for instance Nishino and Shariff (2010) and Tempelmann et al. (2011)), new and upcoming simulation techniques are extending the numerical capabilities. A prime example is the increased knowledge of the numerical generation of free-stream turbulence (FST). Wind tunnel environments inevitably generate a certain level of FST, depending on their design. Through receptivity mechanisms between the boundary layer and the free-stream sections of the flow, disturbances are excited in the boundary layer. By now numerically replicating this FST level measured in the wind tunnel, this work attempts to introduce this receptivity mechanism in the simulation. When implemented, this would remove the need for conventional forcing techniques such as Blowing & Suction, bringing the numerical simulation one step closer to the experimental equivalent. To test the feasibility of implementing FST numerically, and study its capabilities for generating a framework for future numerical setups, this work focuses on attempting to replicate a low-speed crossflow experiment over a 45-degree swept wing performed at TU Delft. All with the goal of establishing the capabilities of FST in this scenario, and setting up a framework for future numerical simulation of CFI development and transition to turbulence on swept wings.

1.5. Project Scope

This research aims to answer the question: ”How to generate a numerical framework to simulate the development of crossflow instabilities and transition to turbulence over a swept wing?”. Which in turn can be split into several sub-questions this research will try to answer:

- Is it possible to accurately replicate experimental CFI development and transition on a swept wing with a global free-stream method that does not a priori assume information on transition itself?
- Which components are required to generate this numerical framework?
- Does using free-stream turbulence in unsteady simulations result in a flow field that represents experiments more closely than conventional forcing techniques (e.g. Blowing and suction)?

In an attempt to answer these questions, a numerical setup will be generated based on an experiment over a smooth swept-wing performed at TU Delft. Replication of the experiment will be done in several steps, aiming to first match baseflow, then the primary crossflow instability (CFI), and finally the secondary CFI to result in an unsteady flow field. This framework will ultimately be assessed based on its capabilities, and flexibility towards future application. With the fundamentals of boundary layer transition to turbulence, crossflow instabilities and some historical context outlined before, the experimental data set will be discussed first in the next section. After which the numerical tools available are touched upon, before diving into the numerical methodology behind the simulation set-up. The results section will then discuss three main pillars: the outcome of the steady flow field matching procedure, the analysis of the unsteady flow field based on transition front, type and frequency content, and a review of the impact of the free-stream turbulence method. Conclusions will be drawn from these results, whether this framework is capable of reliable reproduction of crossflow-dominated experimental flow fields on swept wing surfaces, and if it can be used for more complex setups in the future.

2

Reference Experimental Data

Reference data from a wind tunnel experiment performed by Rius-Vidales and Kotsonis (2022), measuring crossflow-driven transition over an infinite swept wing model, is used as the basis for the numerical work. The following chapter will outline the experimental setup and the data available. Important remarks regarding the processing and use of this experimental data in the numerical setup will serve as a conclusion to this chapter, before moving on to numerical methodology.

2.1. Experimental Layout

The subsonic experiment measures the flow over a swept wing, specifically designed to study CFI development and transition by inhibiting the growth of T-S waves and attachment line instabilities (Serpieri & Kotsonis, 2015). Generating a framework to replicate this experiment numerically requires knowledge of the experimental layout and the flow parameters. Therefore an overview of the wind tunnel and model is provided here.

2.1.1. Wind Tunnel

Experiments were performed in the low-turbulence wind tunnel (LTT) at the Delft University of Technology. This closed-loop atmospheric wind tunnel with a contraction ratio of 17.9:1 keeps turbulence low through the use of 7 anti-turbulence screens. The 2.6m long test section has a cross-sectional area of 1.80m x 1.25m (w x h). Based on free-stream hot-wire data, the measured turbulence intensity level in the test section was 0.023% at a wind speed of 24.7 m/s and Re_{c_X} of 2.17×10^6 (Rius-Vidales & Kotsonis, 2022; Serpieri, 2018). Wind speed in the test section is calculated by subtracting the static pressure at the end of the wind tunnel contraction region from the total pressure between the turbulence grids in the settling chamber. From this pressure differential, the reference dynamic pressure (p_q) in the test section is determined based on an empty test-section calibration curve, resulting in the test section velocity $U_\infty = \sqrt{2p_q/\rho}$ (Rius Vidales, 2022). A pitot tube in the test section provides a second velocity measurement which considers the blockage effect from the swept wing. Streamwise pressure taps at the model surface, as well as access from the back of the model to extract hotwire measurements, allow for detailed characterisation of the flow field.

2.1.2. Model

The model used consists of a constant $c_X = 1.27m$ streamwise chord wing at $45°$ sweep, spanning $1.25m$ and intersecting with both wind tunnel walls. To capture the fundamental features of crossflow-driven transition, the infinite swept wing condition is required (Saric et al., 2003; Serpieri & Kotsonis, 2015). Through a sufficiently high wing aspect ratio, therefore, a spanwise-invariant pressure distribution and flow field has been generated without the need for special wall liners (Serpieri & Kotsonis, 2016). The M3J airfoil used is a variant of the NACA66018, developed by Serpieri and Kotsonis (2015) to favour crossflow instability growth while suppressing Tollmien-Schlichting waves and Görtler instabilities at low angles of attack (≈ 3 °). Compared to the NACA66018 as seen in Figure 2.1, the point of maximum thickness is moved rearward to lengthen the region of favourable pressure gradient.

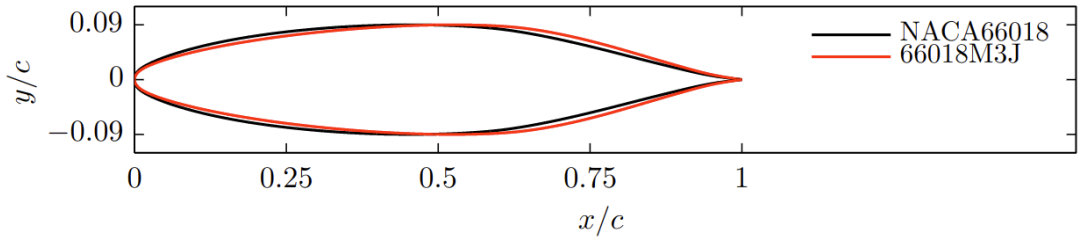

Figure 2.1: M3J airfoil cross-section compared to the NACA66018. Reprinted from Serpieri and Kotsonis (2015).

Made out of fibreglass, the model is polished to a roughness standard deviation of $R = 0.20\mu m$ (Serpieri & Kotsonis, 2016). Parallel with the wind tunnel walls, 92 static pressure taps distributed over two rows allow for measurements along the full streamwise chord of the wing (Rius-Vidales & Kotsonis, 2022). Figure 2.2 shows the location of both rows, which are 40% span apart to validate the infinite swept wing assumption. Discrete Roughness Elements (DRE) are placed $8mm$ apart near the leading edge. The DREs condition the spanwise wavelength of the CF mode. By choosing a height of $200\mu m$, strong stationary CFI growth is encouraged and transition is ensured within the crossflow dominant region. The coordinate system used on the model is aligned with the leading edge of the wing (x, y, z), while the wind tunnel itself uses a free-stream aligned coordinate system, indicated by (X, Y, Z) in Figure 2.2.

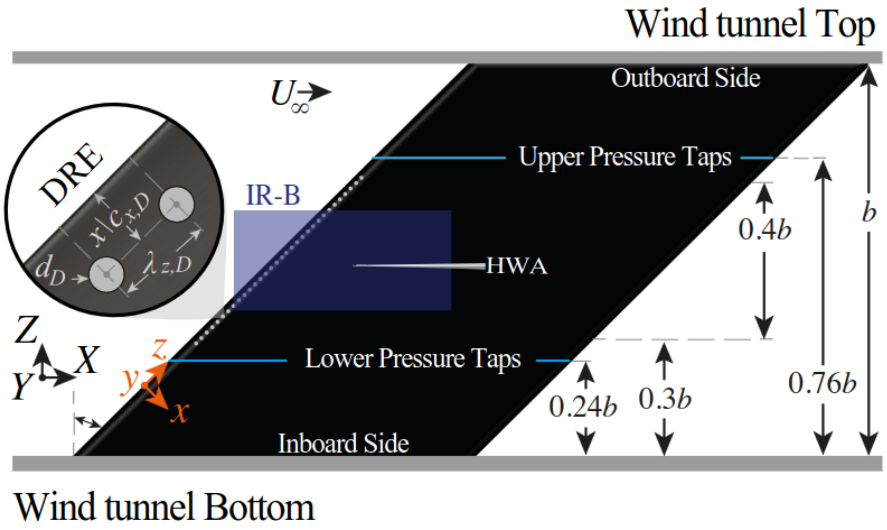

Figure 2.2: Wing model layout and positioning in the wind tunnel, indicating the position of both pressure tap rows, the area of HWA measurements and the DREs placed near the wing leading edge. Re-adapted from Rius-Vidales and Kotsonis (2022)

2.2. Available Data

The available experimental data dictates which approaches for numerical replication are viable. For this reason, recommendations on experimental data desired for future numerical replication are also given at the end of this work. The data extracted from Rius-Vidales and Kotsonis (2022) only considers the experiment without Forward Facing Step (FFS), which is referred to as the clean case. Related to this specific experiment, three separate data sets are available: wind tunnel parameters, surface static pressure port data, and hot-wire anemometry (HWA) planes. Ambient conditions such as temperature, density, and ambient pressure are recorded during testing. Since the hot-wire anemometry planes were measured over the span of multiple days, ambient conditions for each plane varied in temperature T $(+/-1K)$, pressure p $(+/-10hPa)$, density ρ $(+/-0.016kg/m^3)$, and velocity U_∞ $(+/-0.35m/s)$. Test-section U_∞, therefore, shows some variation since the wind tunnel operates at a constant Reynolds number. As previously mentioned in subsection 2.1.1, the test-section free-stream velocity is instead determined indirectly based on wind tunnel calibration data (Rius Vidales, 2022). Additionally, a pitot tube is available in the test section itself just upstream of the model, to allow for velocity corrections

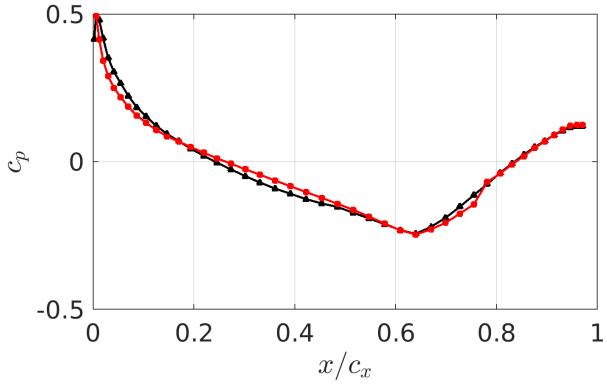

Figure 2.3: Pressure coefficient distribution in streamwise (x-coordinate) direction on the pressure side of the wing at $\alpha = 3°$ and $Re_{c_x} = 2.17 \times 10^6$ for the lower (red) and upper (black) pressure tap rows.

due to blockage effects. With the total pressure known from the ambient pressure measurements, the pressure coefficient (c_p) over the wing can be calculated using the two rows of static pressure ports along the wing chord. The c_p-distribution that emanates is shown in Figure 2.3 for both pressure tap rows. Off-surface measurements are provided by hotwire anemometry. The system consists of a movable steel sting with a single hotwire aligned vertically in Z-direction, measuring a Euclidean sum of the velocity components. The resulting magnitude of this sum equals $Q = \sqrt{(ucos\Lambda + wsin\Lambda)^2 + v^2}$, where Λ is the wing sweep angle. The hotwire is operated using a constant temperature bridge. Measuring at 51.2kHz for 2 seconds, 8 planes of 60 by 40 measurement points are extracted parallel to the leading edge. The result is a measurement resolution of $533\mu m$ in the spanwise z-direction, and between $60\mu m$ and $90\mu m$ in wall-normal (y) direction. Plane locations are expressed in a percentage of the chord and spread between $x/c_x = 0.22$ and $x/c_x = 0.28$ (Rius-Vidales & Kotsonis, 2022). How this data will be used in the numerical setup will be outlined in the methodology section, along with the processing required to generate suitable input parameters from it. Aspects to take into account when processing the experimental data will be touched upon in the next section.

2.3. Additional Remarks

Experimental research using this data has already been performed by Rius-Vidales and Kotsonis (2022). Detailed flow physics analysis of the experimental data is thus out of the scope of this thesis. The focus instead is on the data that can be extracted and used to set up the numerical simulation. Based on the experimental research, it is known that the transition is driven by Type-I CFI. The location of transition is determined using an IR thermographic system, monitoring the pressure side of the wing. Through changes in model temperature, the transition location is determined to be at $x/c_x \approx 0.32$ from Figure 2.4b (Rius-Vidales & Kotsonis, 2022).

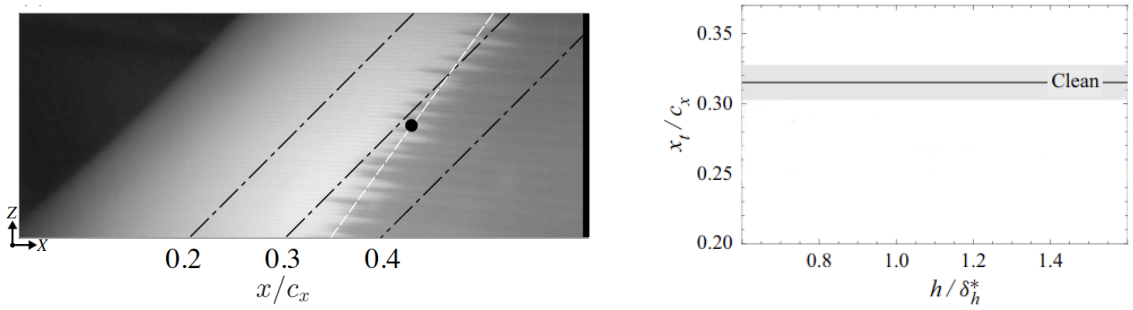

(a) (b)

Figure 2.4: (a) Thermal map from camera IR-B. Marker (•) indicates the projection of the transition linear fit(dashed white line) to the centre of the domain. Reprinted from Rius-Vidales and Kotsonis (2022). (b) Transition location at $Re_{c_x} = 2.17 \times 10^6$ and $\alpha = 3°$for the experimental model, re-adapted from Rius-Vidales and Kotsonis (2022).

To allow for successful replication of this transition behaviour, all flow characteristics need to be extracted correctly. For the experimental data, the inflow velocity is crucial. The blockage effect of the model on the calibrated test-section velocity will thus need to be taken into account. Impacts of the wind tunnel walls, although kept to a minimum, can also play a role in the inflow angle. The infinite span wing assumption, while commonly accepted, does also need to be carefully checked to exclude the possibility of large spanwise differences throughout the experimental data. The hotwire is also located on a long metal sting, which will inevitably introduce noise into the measurement data. To extract the turbulence intensity from this data it is, therefore, band-pass filtered between 2 and $5000Hz$ to remove both the mechanical vibrations of the supporting arm and the high-frequency spectrum traced to the inherent characteristics of the HWA bridge (Rius-Vidales & Kotsonis, 2022). Spectral analysis of the boundary layer in Rius-Vidales and Kotsonis (2022) is still performed on the unfiltered measurements, which has to be kept in mind when comparing experimental and numerical data. With this knowledge of the potential experimental error sources and mitigation strategies, a numerical methodology can be introduced. The different coordinate systems introduced in this section need to be kept track of, and care needs to be taken when transferring experimental data to the numerical framework. Often different coordinate and reference systems are used in both research fields, and data might be represented as dimensional or non-dimensional depending on the circumstances.

3

Governing Equations

Multiple tools are used to efficiently set up the numerical flow field. These tools are all based on fundamental fluid dynamics equations, with some accompanied by a set of assumptions to simplify implementation. This section will outline the Navier-Stokes equations the DNS solver is based on, as well as the different equations used by stability tools. The latter allows for a less computationally expensive assessment of the flow field.

3.1. Navier-Stokes Equations

The Navier-Stokes equations are the basis of fluid dynamics. Describing the behaviour of fluid flows, they consist of a set of three time-dependent partial differential equations expressing the fluid momentum conservation in each of the three dimensions. To fully characterise the flow field, the continuity equation for mass conservation is also required, along with the conservation of energy equation (Anderson, 2016). Although no general solution has been constructed for these equations, these can be solved in iterative ways. Because this is very computationally expensive, a simplified version of the Navier-Stokes equations can be used for the purpose of this work. As this concerns low-speed aerodynamics, compressibility and temperature effects can be neglected. The three momentum conservation equations can be simplified for incompressible flow. Neglecting the conservation of energy equation through the assumption of constant temperature turns the Navier-Stokes equations into Equation 3.1.

$$\rho \frac{\partial q}{\partial t} = -\rho(q \cdot \nabla)q - \nabla p + \mu \nabla^2 q \tag{3.1a}$$

$$\nabla \cdot q = 0 \tag{3.1b}$$

Here q is the velocity vector $[u, v, w]$, in x-, y-, and z-direction, respectively. μ represents the dynamic viscosity, t indicates the time and p the static pressure (Pope, 2000). DNS uses these equations to solve the crossflow solution inside the numerical domain. Because it describes the behaviour of incompressible viscous flow precisely, it is the ideal choice for experimental replication.

3.2. Linear Stability Theory

Throughout the first stage of transition, infinitesimally small disturbances show linear behaviour. Using incompressible local stability equations (ILST), the growth rate at this initial perturbation stage can be estimated. Derivation of the ILST equations starts with Navier-Stokes for incompressible flows given by Equation 3.1. The velocity vector can be split up into a steady baseflow component $\bar{q}$ and a perturbation component q'. Pressure terms can be decomposed in the same way. Then replacing the velocity vector with the three individual momentum equations results in a set of four equations, representing the extended Navier-Stokes equations. Subtracting the steady components from the equations, based on the assumption that these form a stand-alone solution when the flow is time-invariant, results in the perturbation equations. In this set, given by Equation 3.2, only the perturbation terms u', v', w', and p' are left.

$$\frac{\partial u'}{\partial x} + \frac{\partial v'}{\partial y} + \frac{\partial w'}{\partial z} = 0, \tag{3.2a}$$

$$\begin{aligned}\frac{\partial u'}{\partial t} + U\frac{\partial u'}{\partial x} + u'\frac{\partial u'}{\partial x} + u'\frac{\partial U}{\partial x} + V\frac{\partial u'}{\partial y} + v'\frac{\partial u'}{\partial y} + v'\frac{\partial U}{\partial y} + W\frac{\partial u'}{\partial z} + w'\frac{\partial u'}{\partial z} + w'\frac{\partial U}{\partial z} \\ + \frac{1}{\rho}\frac{\partial p'}{\partial x} - \nu\left[\frac{\partial^2 u'}{\partial x^2} + \frac{\partial^2 u'}{\partial y^2} + \frac{\partial^2 u^r}{\partial z^2}\right] = 0,\end{aligned} \tag{3.2b}$$

$$\begin{aligned}\frac{\partial v'}{\partial t} + U\frac{\partial v'}{\partial x} + u'\frac{\partial v'}{\partial x} + u'\frac{\partial V}{\partial x} + V\frac{\partial v'}{\partial y} + v'\frac{\partial v'}{\partial y} + v'\frac{\partial V}{\partial y} + W\frac{\partial v'}{\partial z} + w'\frac{\partial v'}{\partial z} + w'\frac{\partial V}{\partial z} \\ + \frac{1}{\rho}\frac{\partial p'}{\partial y} - \nu\left[\frac{\partial^2 v'}{\partial x^2} + \frac{\partial^2 v'}{\partial y^2} + \frac{\partial^2 v'}{\partial z^2}\right] = 0,\end{aligned} \tag{3.2c}$$

$$\begin{aligned}\frac{\partial w'}{\partial t} + U\frac{\partial w'}{\partial x} + u'\frac{\partial w'}{\partial x} + u'\frac{\partial W}{\partial x} + V\frac{\partial w'}{\partial y} + v'\frac{\partial w'}{\partial y} + v'\frac{\partial W}{\partial y} + w'\frac{\partial w'}{\partial z} + W\frac{\partial w'}{\partial z} + w'\frac{\partial W}{\partial z} \\ + \frac{1}{\rho}\frac{\partial p'}{\partial z} - \nu\left[\frac{\partial^2 w'}{\partial x^2} + \frac{\partial^2 w'}{\partial y^2} + \frac{\partial^2 w'}{\partial z^2}\right] = 0,\end{aligned} \tag{3.2d}$$

This set of perturbation equations is the basis for all stability tools discussed in this chapter, including the linear and nonlinear parabolized stability equations derived later in this section. To arrive at the ILST equations, the following Fourier-ansatz is required,

$$q'_{LST} = \hat{q}(y)\mathrm{e}^{i(\alpha x + \beta z - \omega t)} + c.c, \tag{3.3}$$

where $\tilde{q}(y)$ is the eigenfunction velocity vector, α the streamwise wavenumber, β the spanwise wavenumber, and ω the angular frequency (Westerbeek, 2020). Introducing this equation into the perturbation equations and assuming only small perturbations are present, allows for the linearisation of the equations. This entails disregarding terms with multiple disturbances multiplied, improving computational cost and time. Removal does come with the cost of losing flow information as transition is ultimately triggered by nonlinear disturbances. Other processes such as vortex breakdown are also not simulated because of this assumption. The ILST equations can be further simplified by assuming inhomogeneous disturbances in only one direction. Combine this with the parallel flow assumptions, and a one-dimensional set of equations remain, with the flow only being variable in the y-direction. In Cartesian coordinates, ILST consists of the following four equations,

$$i\beta\tilde{\omega} + \tilde{v}_y + i\alpha\tilde{u} = 0 \tag{3.4}$$

$$i\beta\tilde{u}\bar{W} + \tilde{v}\bar{U}_y + i\alpha\tilde{u}\bar{U} - i\omega\tilde{u} = -i\alpha\tilde{p} + \frac{\tilde{u}_{yy}}{Re} - \frac{\beta^2\tilde{u}}{Re} - \frac{\alpha^2\tilde{u}}{Re} \tag{3.5}$$

$$i\beta\tilde{v}\bar{W} + i\alpha\tilde{v}\bar{U} - i\omega\tilde{v} = -\tilde{p}_y + \frac{\tilde{v}_{yy}}{Re} - \frac{\beta^2\tilde{v}}{Re} - \frac{\alpha^2\tilde{v}}{Re} \tag{3.6}$$

$$\tilde{v}\bar{W}_y + i\beta\tilde{w}\bar{W} + i\alpha\tilde{w}\bar{W} - i\omega\tilde{w} = -i\beta\tilde{p} + \frac{\tilde{w}_{yy}}{Re} - \frac{\beta^2\tilde{w}}{Re} - \frac{\alpha^2\tilde{w}}{Re} \tag{3.7}$$

This set of eigenvalue problems needs to be solved to find the eigenvalue α, and the eigenfunctions $\hat{u}$, $\hat{v}$, $\hat{w}$, and $\hat{p}$. A range of solutions is possible for these equations, and thus filtering is required to find the required solution set. The possible solutions can be divided into two different groups, the discrete and continuous spectra. Solutions to the discrete spectrum induce a disturbance inside the boundary layer, which is damped exponentially outside of the boundary layer. Continuous solutions produce oscillatory eigenfunctions in the free-stream, which get damped out quickly inside the boundary layer (Jacobs & Durbin, 1998). From which spectrum the filtered solution is chosen thus depends on the desired use case. For generating boundary conditions in the free-stream region, continuous eigenfunctions are required. In the case of free-stream turbulence, a large number of these continuous eigenfunctions are superimposed, generating a randomised oscillatory behaviour. To initialise a boundary layer through a boundary condition, discrete eigenfunctions are required, leaving the free-stream untouched. Similar to transition estimation, the discrete spectrum of the eigenvalue problem is desired.

3.2.1. Orr-Sommerfeld Equation

The Orr-Sommerfeld equation is essentially a form of the ILST equations, written differently. Discovered by William McFadden Orr and Arnold Sommerfeld, it also describes linear disturbance modes in parallel flow (Eckert, 2010). Having become common knowledge in hydrodynamic stability theory, it is more often referenced than the ILST form. No further assumptions are required to switch between the ILST and Orr-Sommerfeld forms. The partial differential form of the Orr-Sommerfeld equation looks like,

$$\left[(i\omega - i\alpha U - i\beta W)\left(\frac{\partial^2}{\partial y^2} - \alpha^2 - \beta^2\right) + \alpha i\frac{d^2U}{dy^2} + \beta i\frac{d^2W}{dy^2} + v\left(\frac{\partial^2}{\partial y^2} - \alpha^2 - \beta^2\right)^2\right]\hat{v} = 0 \tag{3.8}$$

The eigenvalue problem can be solved to find the eigenfunction v and the eigenvalue α when β and the inflow velocity components U and W are known. In the numerical study of transition from laminar to turbulent flow over swept wings, ILST or Orr-Sommerfeld can be used in the early stages of the transition process. Allowing the generation of inflow boundary conditions for the fluid domain, it also acts as an initial predictor of disturbance growth. Though flow quickly tends to become non-parallel and more accurate techniques will be required there. Next to disturbance growth, the Orr-Sommerfeld equation is also the bases of the e^N transition prediction method. This method uses an ILST-based approach to allow for relatively accurate prediction of the transition location (Van Ingen, 2008).

3.3. Linear Parabolized Stability Equations

Linear Parabolized Stability Equations, also called PSE or LPSE, is a more detailed version of ILST. The major change is dropping the parallel flow assumption, implying the boundary layer shape is now dependent on the position along the wall, next to the wall-normal y-coordinate. This x-coordinate dependency is introduced through the Fourier-ansatz equation given by Equation 3.9, where the shape function q now becomes dependent on both x and y. α is integrated over the streamwise distance travelled, to provide additional information on upstream disturbance growth.

$$q'(x,y,z,t) = \hat{q}(x,y)e^{i\left(\int_{x_0}^{x}\alpha(x)dx + \beta z - \omega t\right)} \tag{3.9}$$

In Equation 3.9, q is only allowed to vary slowly and is called the shape function. The exponential term is referred to as the wave function, which is allowed to vary rapidly when required. This equation can be introduced into the perturbation equations derived before in section 3.2 and given by Equation 3.2, therefore its derivatives need to be calculated first. The shape function can be expanded through the chain rule for each of its components. An example expansion is given for the chordwise component $\hat{u}$,

$$\frac{\partial\hat{u}}{\partial x} = \frac{\partial\xi}{\partial x}\frac{\partial\hat{u}}{\partial\xi} = \frac{1}{Re}\frac{\partial\hat{u}}{\partial\xi}. \tag{3.10}$$

Using this expansion to obtain the first-order derivative of the Fourier-ansatz becomes,

$$\frac{\partial u'}{\partial x} = \left(\frac{1}{Re}\frac{\partial\hat{u}}{\partial\xi} + i\alpha\hat{u}\right)e^{i\left(\int_{\tilde{x}_0}^{\tilde{x}}\alpha(x)\partial x + \beta z - \omega t\right)}, \tag{3.11}$$

where the second-order derivative results in

$$\frac{\partial^2 u'}{\partial x^2} = \left(\frac{1}{\mathrm{Re}^2}\frac{\partial^2\hat{u}}{\partial\xi^2} + \frac{i2\alpha}{\mathrm{Re}}\frac{\partial\hat{u}}{\partial\xi} + \frac{i\hat{u}}{\mathrm{Re}}\frac{\partial\alpha}{\partial\xi} - \alpha^2\hat{u}\right)e^{i\left(\int_{\tilde{x}_0}^{\tilde{u}}\alpha(x)\partial x + \beta z - \omega t\right)}. \tag{3.12}$$

To make the equations parabolic, the elliptic term $\frac{1}{Re^2}\frac{\partial^2\hat{u}}{\partial\xi^2}$ in the second-order derivative is removed as it is negligible in magnitude compared to the remaining terms. Introducing these derivatives into the perturbation equations results in the following set of linear partial differential equations as derived by Westerbeek (2020),

$$i\alpha\hat{u} + \frac{\partial\hat{u}}{\partial x} + \frac{\partial\hat{v}}{\partial y} + i\beta\hat{w} = 0, \tag{3.13a}$$

$$-i\omega\hat{u} + Ui\alpha\hat{u} + U\frac{\partial\hat{u}}{\partial x} + \hat{u}\frac{\partial U}{\partial x} + V\frac{d\hat{u}}{dy} + \hat{v}\frac{\partial U}{\partial y} + i\beta W\hat{u} + \frac{i\alpha}{\rho}\hat{p} + \frac{1}{\rho}\frac{\partial\hat{p}}{\partial x} - v\left[i2\alpha\frac{\partial\hat{u}}{\partial x} + i\hat{u}\frac{\partial\alpha}{\partial x} + \frac{\partial^2\hat{u}}{\partial y^2} - \left(\alpha^2+\beta^2\right)\hat{u}\right] = 0, \tag{3.13b}$$

$$-i\omega\hat{v} + Ui\alpha\hat{v} + U\frac{\partial\hat{v}}{\partial x} + u'\frac{\partial V}{\partial x} + V\frac{\partial\hat{v}}{\partial y} + \hat{v}\frac{\partial V}{\partial y} + i\beta W\hat{v} + \frac{1}{\rho}\frac{\partial\hat{p}}{\partial y} - v\left[i2\alpha\frac{\partial\hat{v}}{\partial x} + i\hat{v}\frac{\partial\alpha}{\partial x} + \frac{\partial^2\hat{v}}{\partial y^2} - \left(\alpha^2+\beta^2\right)\hat{v}\right] = 0, \tag{3.13c}$$

$$-i\omega\hat{w} + Ui\alpha\hat{w} + U\frac{\partial\hat{w}}{\partial x} + u'\frac{\partial W}{\partial x} + V\frac{\partial\hat{w}}{\partial y} + \hat{v}\frac{\partial W}{\partial y} + i\beta W\hat{w} + \frac{i\beta}{\rho}\hat{p} - v\left[i2\alpha\frac{\partial\hat{w}}{\partial x} + i\hat{w}\frac{\partial\alpha}{\partial x} + \frac{\partial^2\hat{w}}{\partial y^2} - \left(\alpha^2+\beta^2\right)\hat{w}\right] = 0 \tag{3.13d}$$

As mentioned before, the shape function q and wavenumber α are both allowed to vary in x. To limit the degrees of freedom in the system, an additional equation is thus required. The most desirable way to achieve this reduction in degrees of freedom is to make the wave function deal with most of the streamwise disturbances because the shape function needs to remain slowly varying compared to the wave function. This constraint on the shape function can be achieved through the normalisation condition given by,

$$\int \hat{q}^*\frac{\partial\hat{q}}{\partial x}dy = 0 \tag{3.14}$$

where the asterisk of the shape function indicates the complex conjugate. The normalization condition restricts q to mere shape changes and thus reduces the degrees of freedom as desired mathematically (Haynes & Reed, 2000).

With parallel flow no longer assumed, the flow field used becomes two-dimensional. This step improves accuracy compared to ILST, but also drastically increases computational cost. Though by disregarding the non-linearities in the flow, total cost still remains at only a fraction of CFD methods such as Large Eddy Simulation (LES) and Direct Numerical Simulation (DNS). Aerodynamicists use this technique to estimate disturbance growth in the linear stages of the transition process. Up to the point where the first non-linearities appear, this technique remains an accurate representation of the real world and can be trusted to depict correct disturbance growth. However, care has to be taken because the laminar to turbulent transition process is dominated by non-linearities. In numerical work, PSE can be used to estimate disturbance growth in the linear region of transition, when initial disturbances are still small. Either to help generate boundary conditions or to estimate the start of the non-linear disturbance region. Back-up with higher fidelity methods is desired to ensure the validity of the PSE application, and the boundaries within which it is used. In experimental work, where the measurements are taken discretely in space, this technique can also be used to estimate the upstream disturbance amplitudes within the linear region.

3.4. Nonlinear Parabolized Stability Equations

Nonlinear Parabolized Stability Equations are related to the previously discussed LPSE, but include nonlinearities in the flow field as well. Not disregarding the nonlinear terms allows for a better approximation of the disturbance growth beyond what LPSE is capable of. Although not able to simulate transition itself, nonlinear disturbance growth up to the transition point does get predicted accurately. Derivation of the NPSE starts again with the perturbation equations from section 3.2, given by Equation 3.2. Not neglecting the nonlinear terms, all multiplications of different perturbations are moved to the right side. As with LPSE, the derivation is based on a Fourier-ansatz. This time a slightly different form, written as,

$$q'(x,y,z,t) = \sum_{m=-M}^{M}\sum_{n=-N}^{N}\hat{q}_{m,n}(x,y)e^{i\left(\int_{x_0}^{x_e}\alpha(x)_{m,n}dx+\beta_m z-\omega_n t\right)}. \tag{3.15}$$

Although similar to the LPSE version, NPSE uses a double summation in the equation to account for a range of β wavenumbers and ω frequencies. Where ILST and PSE only introduced one disturbance

mode, multiple modes have to be introduced to obtain non-linear interaction. Because of the need for a complex conjugate, the summation inside the Fourier-ansatz includes the negative side of the summation as well, running from negative infinity to positive infinity for both inputs (Haynes & Reed, 2000). Practice will mandate the infinity values to be capped at a fixed but large number. The same assumptions discussed with LPSE hold for these equations, assuming the shape function q to vary slowly, while the exponential term allows rapid changes when required. Expanding the shape function and taking the derivatives of the Fourier-ansatz, similar to the derivation for LPSE, allows the introduction of these derivatives into the nonlinear set of equations. This yields the set of equations referred to as NPSE,

$$i\alpha\hat{u} + \frac{\partial\hat{u}}{\partial x} + \frac{\partial\hat{v}}{\partial y} + i\beta\hat{w} = 0, \tag{3.16a}$$

$$\begin{aligned} -i\omega\hat{u} + Ui\alpha\hat{u} + U\frac{\partial\hat{u}}{\partial x} + \hat{u}\frac{\partial U}{\partial x} + V\frac{d\hat{u}}{dy} + \hat{v}\frac{\partial U}{\partial y} + i\beta W\hat{u} + \frac{i\alpha}{\rho}\hat{p} + \frac{1}{\rho}\frac{\partial\hat{p}}{\partial x} \\ - v\left[i2\alpha\frac{\partial\hat{a}}{\partial x} + i\hat{u}\frac{\partial\alpha}{\partial x} + \frac{\partial^2\hat{u}}{\partial y^2} - \left(\alpha^2+\beta^2\right)\hat{u}\right] = -i\alpha\hat{u}^2 - \hat{u}\frac{\partial\hat{u}}{\partial x} - \hat{v}\frac{\partial\hat{u}}{\partial y} - i\beta\hat{w}\hat{u}, \end{aligned} \tag{3.16b}$$

$$\begin{aligned} -i\omega\hat{v} + Ui\alpha\hat{v} + U\frac{\partial\hat{v}}{\partial x} + u'\frac{\partial V}{\partial x} + V\frac{\partial\hat{v}}{\partial y} + \hat{v}\frac{\partial V}{\partial y} + i\beta W\hat{v} + \frac{1}{\rho}\frac{\partial\hat{p}}{\partial y} \\ - v\left[i2\alpha\frac{\partial\hat{v}}{\partial x} + i\hat{v}\frac{\partial\alpha}{\partial x} + \frac{\partial^2\hat{v}}{\partial y^2} - \left(\alpha^2+\beta^2\right)\hat{v}\right] = -i\alpha\hat{u}\hat{v} - \hat{u}\frac{\partial\hat{v}}{\partial x} - \hat{v}\frac{\partial\hat{v}}{\partial y} - i\beta\hat{w}\hat{v}, \end{aligned} \tag{3.16c}$$

$$\begin{aligned} -i\omega\hat{w} + Ui\alpha\hat{w} + U\frac{\partial\hat{w}}{\partial x} + u'\frac{\partial W}{\partial x} + \hat{v}\frac{\partial W}{\partial y} + V\frac{\partial\hat{w}}{\partial y} + i\beta W\hat{w} + \frac{i\beta}{\rho}\hat{p} \\ - v\left[i2\alpha\frac{\partial\hat{w}}{\partial x} + i\hat{w}\frac{\partial\alpha}{\partial x} + \frac{\partial^2}{\partial y^2}\hat{w} - \left(\alpha^2+\beta^2\right)\hat{w}\right] = -i\alpha\hat{u}\hat{w} - \hat{u}\frac{\partial\hat{w}}{\partial x} - \hat{v}\frac{\partial\hat{w}}{\partial y} - i\beta\hat{w}\hat{w}^2, \end{aligned} \tag{3.16d}$$

As with PSE, too many degrees of freedom are present in these equations, and one thus needs to be limited. By again using the normalisation equation used in Equation 3.14, the shape function growth can be limited to purely shape changes again (Westerbeek, 2020).

In work such as Casacuberta, Hickel, et al. (2022) and Casacuberta, Groot, et al. (2022), NPSE has been shown to accurately match DNS until close to the transition region. The accuracy of the disturbance growth prediction, compared with the computational cost that is orders of magnitude below DNS, indicates this is a valuable method in numerical research. Not only as a cheaper alternative to DNS but also as a standalone tool in numerical work. Prediction of disturbance amplitudes up and downstream of a discrete measurement point can give additional insight needed to draw conclusions. In numerical research where experimental results are used as an input, this is certainly beneficial, as being able to estimate upstream conditions helps set up boundary conditions. If necessary, this two-dimensional technique can also be expanded to generate a three-dimensional flow field, based on the Fourier-ansatz expansion of the results.

4

Methodology

To allow accurate replication of the experimental data set, the numerical setup needs to reflect the wind tunnel conditions as much as possible. In three steps, this work will attempt to numerically match the experimental data available. By first numerically reconstructing the baseflow, which is the steady-state solution to the Navier-Stokes equations, a solid foundation is built. The addition of stationary crossflow disturbances results in a steady DNS flow field, which can be directly validated with the available experimental hotwire data. Once the steady flow field is represented accurately, secondary disturbances are introduced to trigger laminar-to-turbulent transition. The detailed numerical setup to perform these steps, and how this is implemented in DNS, is discussed in this chapter.

4.1. Computational Domain

The numerical domain needs to model the flow over the 45°swept wing used during the experiment of Rius-Vidales and Kotsonis (2022). The resources to model the full wind tunnel are simply not available, and with computational cost in mind, the domain size is kept to a strict minimum. For this purpose, a box-shaped domain on the pressure side of the wing is used, aligned perpendicular to the airfoil leading edge. The numerical domain shape and size are visualised in Figure 4.1. All sides of the domain are flat, implying the airfoil surface is modelled as a flat plate. Curvature is emulated through prescribed pressure gradients instead. Spanning from 5% chord to 45% chord, the domain captures the majority of crossflow wave growth and the expected transition region. The domain width is limited to $8mm$, the exact spanwise fundamental wavelength to ensure the correct harmonics are introduced. The domain height needs to be kept to a minimum, without influencing the boundary layer development. This results in a domain with the dimensions of $0 \leq x/\delta_0 \leq 514.32$ in chordwise direction, $0 \leq y/\delta_0 \leq 25.93$ in wall-normal direction, and $-4.76 \leq z/\delta_0 \leq 4.76$ in spanwise direction. The coordinate system used is aligned with the airfoil leading edge, where $\mathbf{x} = [x, y, z]^T$ corresponds to the chordwise, wall-normal and spanwise directions, respectively.

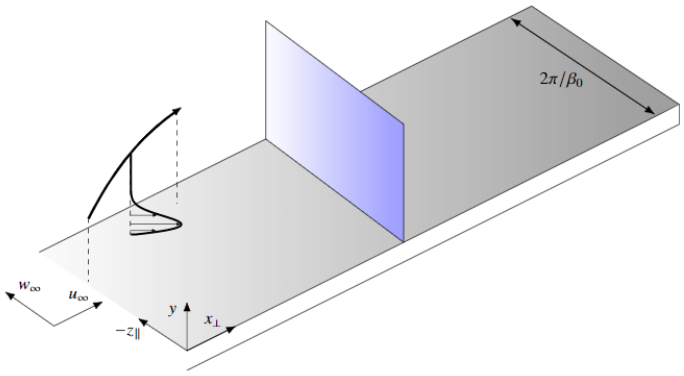

Figure 4.1: Visualisation of the numerical domain. Reprinted from Casacuberta, Hickel, et al. (2022).

4.1.1. Boundary Conditions

The simple geometry of the numerical domain requires a number of boundary conditions before it resembles the experimental flow field over the airfoil surface. To facilitate suitable boundary conditions on all sides, three layers of ghost cells are added around the domain. Emulating the curvature of the airfoil on the flat plate is achieved through an external pressure gradient prescribed at the top boundary of the domain. The pressure data used is acquired from the experiment performed by Rius-Vidales and Kotsonis (2022). Imposing the pressure distribution on the top boundary instead of the flat plate itself is allowed based on the Prandtl boundary layer equations for incompressible flow. Outlined in section 1.1, these equations showed that $\partial p/\partial y = 0$ in this case. The top boundary of the domain additionally damps any velocity and pressure fluctuations at that boundary through a decay function $u'(y) \propto e^{-\alpha y}$ applied to the ghost cells (Hickel & Adams, 2008). Prescribing the pressure gradient on the top boundary removes the need for complex boundary conditions at the flat-plate wall, where only the no-slip boundary condition $q = 0$ is applied, with $q = [u, v, w]^T$. The spanwise boundary conditions are periodic, to adhere to the infinite wing assumption. Periodicity also introduces harmonics of the spanwise fundamental wavenumber into the domain. The outflow boundary must meet the constant total pressure condition, as well as $\partial p/\partial x = 0$ and $\partial^2 q/\partial x^2 = 0$. The inflow boundary is where the crossflow disturbances will be prescribed, which is the topic of the following sections on baseflow, primary and secondary instabilities. These will include the free-stream velocity components and several disturbance modes. To prescribe the experimental velocity distribution in the numerical domain, a fourth-order polylogarithmic function is used which is curve-fitted to the experimental measurement data. Details on this function are given in section 4.2.

4.1.2. Mesh

To study the mechanisms behind laminar-to-turbulent transition, adequate grid points throughout the boundary layer are necessary to resolve all relevant flow structures. To obtain this, a y^+ value of below one is desired at the wall, retaining the same wall-normal grid spacing across the boundary layer. The resulting grid consists of 340,733,952 rectangular cells, 4108 x 576 x 144 in chordwise, wall-normal, and spanwise directions, respectively. To capture all structures in the boundary layer, half of the total grid points in wall-normal direction are equally spaced within the initial boundary layer height. The remaining points in this direction are hyperbolically spaced between the boundary layer and the top domain boundary. The result is a y^+ equal to 0.58, adequately below the required value of 1. An example of this grid spacing is shown in Figure 4.2. Grid points in chordwise and spanwise directions are all equally spaced, with a grid spacing in wall units of $x^+ = 5.28$ and $z^+ = 2.79$. The large refinement in z-direction is required to capture the spanwise features that drive the transition mechanisms in the experimental case. Similar refinement in x- and z-direction is also desired to keep the cell aspect ratios low. Since the flow moves at an angle of only slightly larger than 45 °to the domain, flow velocity is expected to be relatively similar in both directions throughout the domain. Very different refinements in both directions could lead to numerical errors introduced into the flow field. Coarser grids with fewer cells are used for testing the DNS setup. By first converging the solution to grids with respectively $1/16^{th}$ and $1/4^{th}$ the amount of cells, also allows initialisation and faster convergence of the fully refined domain. More refined grids are obtained by splitting the cells in each direction in two.

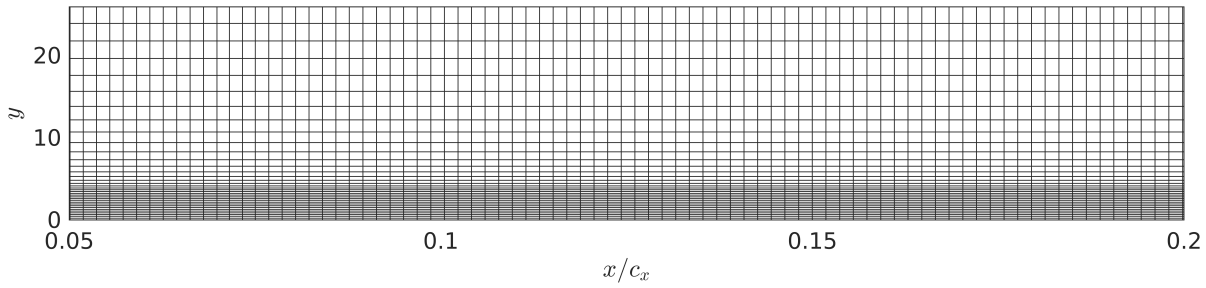

Figure 4.2: 2D x-y slice of the numerical mesh, coarsened for clarity by plotting every 16th grid point in x- and y-direction. Cropped in x-direction for comprehensibility, as the x-spacing remains identical.

4.2. Baseflow

Developing baseflow boundary conditions is the first step towards replicating the experimental flow field. An accurate match is essential at this point, to create a stable platform to add disturbances and ensure interactions between the introduced disturbances and the baseflow happen correctly. To achieve this match, experimental measurements are processed to serve as numerical input. The exact free-stream velocity magnitude in the experimental test section is first determined. Usually, this is calculated by subtracting the static pressure at the end of the wind tunnel contraction region from the total pressure between the turbulence grids in the settling chamber. From this pressure differential, the reference dynamic pressure (p_q) in the test section is determined based on an empty test-section calibration curve, resulting in the test section velocity $U_\infty = \sqrt{2p_q/\rho}$ (Rius Vidales, 2022). However, the blockage ratio of the wind tunnel model is 10%. With blockage ratios of 5% already known to result in meaningful blockage effects, this ratio of 10% is expected to noticeably impact the velocity measurements (Takeda & Kato, 1992). Therefore, a pitot tube in the test section is used to account for this effect. Deemed far enough upstream of the model to not be impacted by its presence, it gives a more representative reading of the test section velocity magnitude. With this information known, the local external velocity distribution over the airfoil can be determined using the available experimental pressure tap data.

4.2.1. External Velocity Magnitude

Two rows of pressure taps are available along the chord, as indicated in Figure 2.2. These two rows are processed individually at first, to avoid loss of information due to averaging. Additionally, processing both rows separately allows the evaluation of the infinite wing assumption. After processing, an external velocity profile is available for both rows, which can then be averaged before implementation in the numerical setup. To obtain both velocity profiles, the first step is to determine the pressure coefficient through

$$c_p = \frac{p_{s_{wing}} - p_{s_{pitot}}}{p_q}, \tag{4.1}$$

where the static (p_s) and dynamic (p_q) pressure of the pitot tube in the test section is used as input. The test section pitot data is deemed the closest representation of the actual test section conditions, because the calibration curve shows slight deviation from the expected results, most likely due to the blockage effects not being taken into account. Combined with the static pressure measurement on the wing surface, the pressure coefficient curves can now be obtained. Under the assumption of both incompressible and inviscid flow, the Bernoulli equation can then be combined with the pressure coefficient to result in an equation for the external flow velocity magnitude U_e,

$$U_e = \sqrt{1 - c_p} \cdot U_\infty, \tag{4.2}$$

where U_∞ is the free-stream velocity magnitude determined earlier. The incompressible flow assumption is an accurate depiction of the flow state here, as the velocity magnitude is approximately 24m/s or Mach 0.07, well below the limit for incompressible flows of Mach 0.3. The inviscid flow assumption allows more room for errors, as it does not take into account any viscous effects inside the boundary layer, which might affect the wall-normal pressure gradient. However, from the Prandtl boundary layer equations, it can be considered a valid assumption in a flow with these characteristics and the impact is deemed to be small enough to neglect at this stage. An external velocity magnitude curve (U_e) can therefore be established for each of the two pressure tap rows. The curves for both rows should match very closely based on the infinite swept wing assumption, and an average of both will be used in the remainder of this work. Keep in mind that this U_e is the external velocity magnitude, and will need to be decomposed into leading-edge normal and tangential components before it can be used as an input to DNS.

4.2.2. External Velocity Decomposition

Pitot tubes and pressure taps only allow the extraction of velocity magnitudes, which is inadequate as input for three-dimensional numerical simulations. Decomposition into leading-edge normal and tangential components is required and achieved based on the assumption of spanwise invariant flow, and the assumption that the leading-edge-normal velocity at the attachment line location is zero. Spanwise invariant flow implies negligible external pressure gradients and flow acceleration in the z-direction,

resulting in a constant free-stream spanwise velocity over the entire wing (Bippes, 1999; Wassermann & Kloker, 2002). This suggests the spanwise velocity only needs to be determined at one point, to know its magnitude at all points along the chord. Assuming the leading-edge-normal velocity at the attachment line location is zero means that only the spanwise velocity component is present at this point, so $w_e = U_e$. When this velocity magnitude at the attachment line can be extracted, it would therefore allow for velocity decomposition along the entire chord. With the spanwise velocity known, therefore, the leading-edge normal components can simply be extracted as $u_e = \sqrt{U_e^2 - w_e^2}$.

Without a pressure tap at the exact attachment line location, w_e can not be directly extracted from the measurement data. Instead, data from the pressure taps on the pressure and suction side of the leading edge have to be used. Interpolation of the experimentally measured U_e around the leading edge is not straightforward, as swept wings have an attachment line rather than a stagnation point. This implies U_e does not go to zero at any point around the leading edge. Therefore, w_e is determined in two steps. First, the attachment line location is determined. Through spline interpolation of all available U_e measurement points, the location of minimum U_e is found, as indicated in Figure 4.3. Located at $x/c_x = 0.0037$, this point is noted as the location of the attachment line. However, this type of interpolation is not deemed accurate enough to extract the velocity magnitude, since this point of minimum U_e lies below the lowest measurement point. Instead, a second interpolation procedure based on the attachment line location is performed. To extract w_e through a more accurate interpolation, the U_e distribution is adapted. U_e on the suction side of the wing are multiplied by -1, to generate a curve that crosses the x-axis. Because U_e is not supposed to reach zero at any point around the wing, this leads to a discontinuity around the attachment line location, as shown by the green dash-dotted line in Figure 4.4a. When decomposing U_e into u_e and w_e, two continuous curves should result. A constant-value w_e-component and a continuous u_e-component crossing the x-axis at the attachment line location ($x/c_x = 0.0037$). Since both components are unknown at this stage, the initial estimate for w_e obtained from Figure 4.3 is used. u_e resulting from decomposition should show a continuous curve, crossing the x-axis at the attachment line location. If this is not the case, the initial w_e estimate is deemed inaccurate and a new w_e estimate is obtained by subtracting or adding 0.01m/s to perform the decomposition again. This iterative procedure is repeated until the leading-edge-normal velocity curve (black line in Figure 4.4b) crosses zero at the attachment line location as a continuous curve.

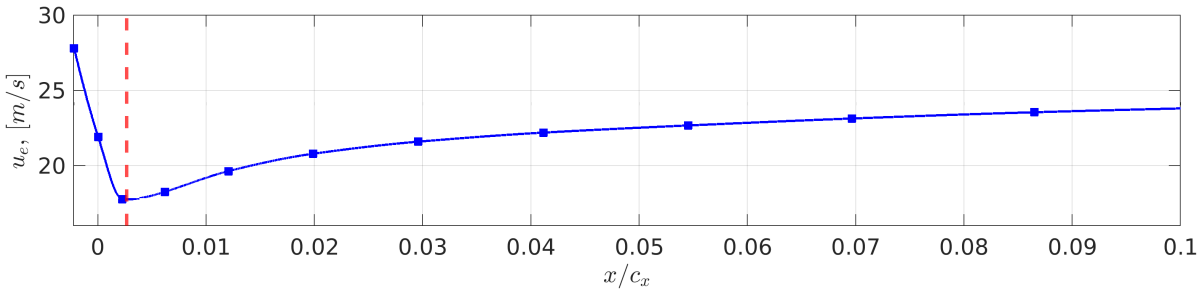

Figure 4.3: Spline interpolation of the discrete U_e-values to find the attachment line location, indicated by the red dashed line.

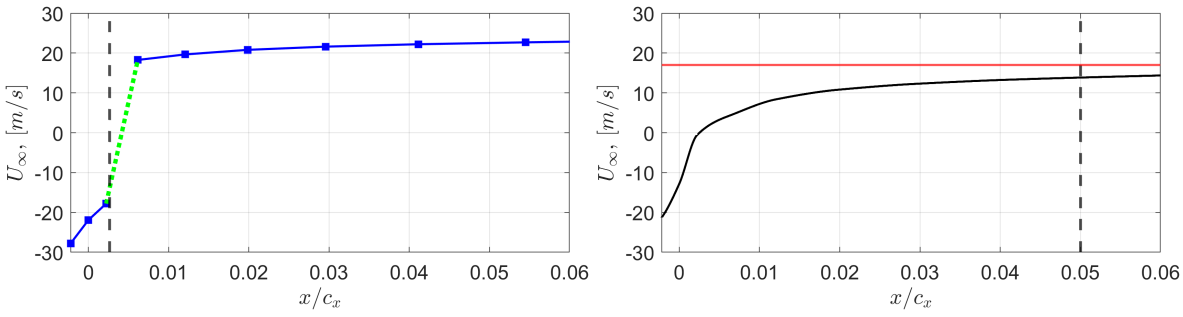

Figure 4.4: (a) Streamwise velocity magnitude U_e (-.), indicating the discontinuïty around the leading edge. Red dashed vertical line indicates the attachment line location. (b) U_e decomposed into spanwise velocity w_e (–) and chordwise velocity u_e (–) around the wing model. Black dashed vertical line indicates the numerical inflow boundary.

With w_e known, the inflow angle and u_e along the chord are known as well. To improve the match with the experiment further, the last step makes use of the HWA data planes. Far from the wall, free-stream velocity data can be extracted from these HWA planes and compared with the velocity curves obtained analytically earlier in this section. When the velocity distributions of both data sets do not line up, the numerical U_e is scaled until the data sets are aligned as well as possible. The inflow angle and relative contributions of u_e and w_e to U_e are known from previous calculations and, therefore, can remain constant throughout this correction process.

4.2.3. Numerical Implementation

The Reynolds number, the free-stream velocity magnitude, and the free-stream velocity components can directly be used in a numerical context. The free-stream velocity magnitude and velocity components are now known, and the Reynolds number was measured during the experiment by Rius-Vidales and Kotsonis (2022). However, there is one step left before DNS implementation. Experimental measurement locations are expressed in a percentage of the chord, but the flow travels a longer distance over the airfoil surface. On top of that, numerical data is non-dimensionalized with the initial boundary layer thickness. To match the experimental data with the numerical domain, the coordinates will thus have to be transformed and non-dimensionalized. Instead of using a percentage of the chord, the points are transformed to the airfoil surface and expressed as a distance along the surface to this point. This distance can directly be used in the flat plate numerical domain and allows the correct implementation of the data into the DNS solver.

The velocity function derived in the previous section, converted to the numerical coordinate system, is used as input for the DNS solver. Represented as a fourth-order polylogarithmic function, it is prescribed in DNS as

$$u_e/u_{\infty,0} = 0.0057\ln^4(x+c) + 0.0661\ln^3(x+c) + 0.2575\ln^2(x+c) + 0.6508\ln(x+c) + 1.9545, \quad (4.3)$$

where $c = 0.0509$ and represents the shift in starting point of the boundary layer from the leading edge to the stagnation point. This velocity distribution is also used as a basis to calculate the pressure distribution prescribed at the top boundary of the domain. With the numerical baseflow details known, the inflow boundary layer thickness is calculated at $\delta_0 = 8.41 \times 10^{-4}$m, and $Re = 750.7138$

4.3. Primary Crossflow Instability

Primary crossflow instability is generated by adding steady disturbances to the converged baseflow. Steady implies no secondary instabilities are allowed, and thus no breakdown happens at this stage yet. These primary CFI are prescribed at the inflow boundary of the numerical domain. The primary CFI can be composed of both travelling and stationary vortices. Due to the strong forcing of the DREs in the experimental setup, however, the stationary CFI is found dominant in the region of interest. Therefore, a strong stationary CFI has to be prescribed numerically as well.

4.3.1. Inflow Boundary Condition

The inflow boundary condition required to initiate the primary CFI consists of a set of eigenvalues, normalised eigenfunctions, and an adequate amplitude for these eigenfunctions. ILST is used to generate the set of eigenvalues and eigenfunctions. As described in section 3.2, the technique solves the eigenvalue problem to find α and the eigenfunction vector. In eqs. (3.4) to (3.7), the baseflow velocities and pressure are known from the baseflow DNS computation, and ω is set to zero since the primary crossflow is time-invariant. The spanwise fundamental wavenumber can be derived from the respective wavelength, known to be 8mm, resulting in a value of $785.40m^{-1}$ for β. Non-dimensionalized for numerical use, β becomes 0.6604. The result of the eigenvalue problems is a long list of eigenvalue and eigenfunction sets. These are either part of the discrete or continuous eigenvalue spectrum. They can be distinguished based on their characteristics. Eigenfunctions of the discrete spectrum show disturbances inside the boundary layer, but get damped out exponentially in the free-stream flow. The continuous spectrum eigenfunctions show oscillating behaviour in the free-stream flow and are damped out rapidly inside the boundary layer. Both spectra, therefore, affect different areas of the flow field. The eigenvalue required for primary crossflow is part of the discrete eigenvalue spectrum, as it needs to impose a disturbance directly on the boundary layer. Choosing the correct eigenvalue and eigenfunction combination within this discrete spectrum is done through filtering. This filtering procedure is based

on the desired exponential decay of the disturbance in the free-stream flow. Using the real component of the eigenvalue α, the spanwise wavenumber β, and the eigenfunction, the following expression is evaluated,

$$\phi = e^{-i\sqrt{\alpha_r^2+\beta^2}y}, \tag{4.4}$$

where ϕ is the eigenfunction, α the real component of the eigenvalue, β the spanwise wavenumber, and y the wall-normal coordinate. Through exponential fitting, the combined set of eigenvalue and eigenfunction that fits this expression the best is extracted. This set creates the boundary condition required to introduce the primary disturbance mode. The resulting set of eigenfunctions (u,v,w,p) and constants (α,β,ω) in y_t-direction can be expanded in z-direction to obtain a two-dimensional vector field prescribed at the domain inflow. This expansion is performed mathematically using,

$$q' = A \cdot \hat{q} \cdot \mathrm{e}^{\mathrm{i}(\alpha x+\beta z-\omega t)}, \tag{4.5}$$

where q' is the disturbance velocity vector prescribed as the inflow boundary condition in DNS, A is the prescribed disturbance amplitude, $\hat{q}$ is the eigenfunction vector, and the exponential terms include the constants α, β, and ω determined before. The ω-term is zero and can be neglected at this stage since the primary CFI is time-invariant. Even though $x = 0$ at the domain inflow boundary, the α component can not be neglected. Since the numerical domain makes use of ghost cells around the domain to prescribe boundary conditions (see Hickel and Adams (2008)), the primary CFI travel through three layers of grid cells before entering the simulation domain. Therefore, the primary CFI travel a finite distance in x-direction where α alters the disturbance shape before it reaches the inflow at $x = 0$. The simplified Fourier-ansatz that results only has one unknown left, which is the amplitude assigned to the disturbances. Based on the experiment, CFI amplitudes can be extracted from the HWA data available at eight chordwise locations. In the numerical domain, the inflow boundary where the amplitude must be prescribed is far upstream of these planes, complicating the process. A numerical optimisation will thus be required to find the appropriate CFI amplitude value at the inflow plane location. DNS is far too costly to be feasible, so stability tools are utilised in the form of an NPSE-based optimisation process. Attempting to find an initial CFI amplitude that results in matching stationary crossflow vortex amplitudes and growth rates throughout the full domain.

4.3.2. Amplitude Determination

Trial-and-error using DNS, to find an upstream boundary condition that results in the required downstream flow field, is much too costly for the current computational capabilities. Instead, different techniques are used to estimate the inflow disturbance amplitude, based on downstream experimental data. In chapter 3, a number of stability tools were outlined, capable of doing just this. Predicting the growth of disturbances in different parts of the transition region. Because stability tools only work marching downstream, an iterative process is worked out to find the amplitude at the inflow plane required to get the correct flow field downstream. NPSE is the stability tool of choice for this process. Because of the extensive list of assumptions (i.e. incompressible flow, parallel flow, small disturbances, linear growth), ILST is too simplified to accurately reproduce the experimental data required, and LPSE also only works in the linear regime of the flow field. Knowing strong crossflow growth is present in the experiment, it is assumed nonlinear disturbances might already be present at the first available HWA plane location. NPSE is thus determined to be most suitable for accurate optimisation of the inflow boundary disturbance amplitude. Based on work performed by Westerbeek (2020), an NPSE simulation is set up on a grid with 500 equidistantly spaces streamwise points and 80 wall-normal Chebyshev collocation points, using 11 spanwise Fourier modes throughout the simulation. The streamwise derivatives are discretized based on a first-order backward Euler discretisation scheme, similar to the setup used by Casacuberta, Hickel, et al. (2022).

Up to this point, all information to run NPSE is known, except for the disturbance amplitude. This is convenient since an optimisation loop can be constructed based on only one input variable. The Matlab $fminsearch$ function is used for this purpose, based on the Nelder-Mead simplex algorithm (Lagarias et al., 1998). This direct search method does not make use of derivatives, which means it is ideal for nonlinear optimisation procedures and arrives at the optimal value rather fast. With the amplitude as input, it minimises a predefined error. The correct definition of the error magnitude is thus crucial for the functioning of the optimisation procedure. Using the disturbance amplitude or the distorted baseflow profiles as shown by Figure 4.5 are two valid options, but to combine the largest amount of information

into this single error metric, the use of the disturbance profiles is preferred. These are obtained by taking the spanwise root-mean-square of the time-averaged disturbance, as given by,

$$\langle \hat{Q}(y) \rangle_z = \sqrt{\frac{1}{n} \sum_{j=1}^{n} (\bar{Q}(y, z_j) - \bar{Q}_{DB}(y))^2}, \tag{4.6}$$

where $\langle \hat{Q}(y) \rangle_z$ is the disturbance profile, $\bar{Q}(y, z)$ the time-averaged velocity plane, and $\bar{Q}_{DB}(y)$ is the disturbed baseflow profile (Rius-Vidales & Kotsonis, 2022). As the HWA probe only measures a specific sum of velocity components $Q = \sqrt{(ucos\Lambda + wsin\Lambda)^2 + v^2}$, the numerical velocity components have to be processed in the same way. As stated in chapter 3, the two-dimensional NPSE results can be expanded in three dimensions using the Fourier-ansatz given by Equation 3.15, which allows for easy extraction of the NPSE data at the location of the HWA planes. Summing the numerical velocity components to obtain a numerical value for Q, and then calculating the disturbance profile with Equation 4.6, allows for a direct comparison. The error is then calculated by integrating the area between the overlapping sections of the curves as shown in Figure 4.5. The hotwire curves are cropped near the wall to ensure the data used for error estimation is not affected by measurement inaccuracies near the wall. By subtracting the overlapping sections of the curves and integrating the difference over the y-axis, a single value that describes the error is obtained. Combining the error values for each of the eight planes results in the quantification of the overall error, as given by Equation 4.7. This ϵ is the value $fminsearch$ minimises by varying the inflow amplitude. As initial amplitude, $A_0 = 3.5x10^{-3}u_\infty$ is used, based on DNS research work on very similar flow fields performed by Casacuberta, Hickel, et al. (2022).

$$\epsilon = \int_x \int_y \frac{\left(\langle \hat{Q}(x,y) \rangle_{z,HWA} - \langle \hat{Q}(x,y) \rangle_{z,NPSE}\right)}{\overline{Q_e}(x)} dydx \tag{4.7}$$

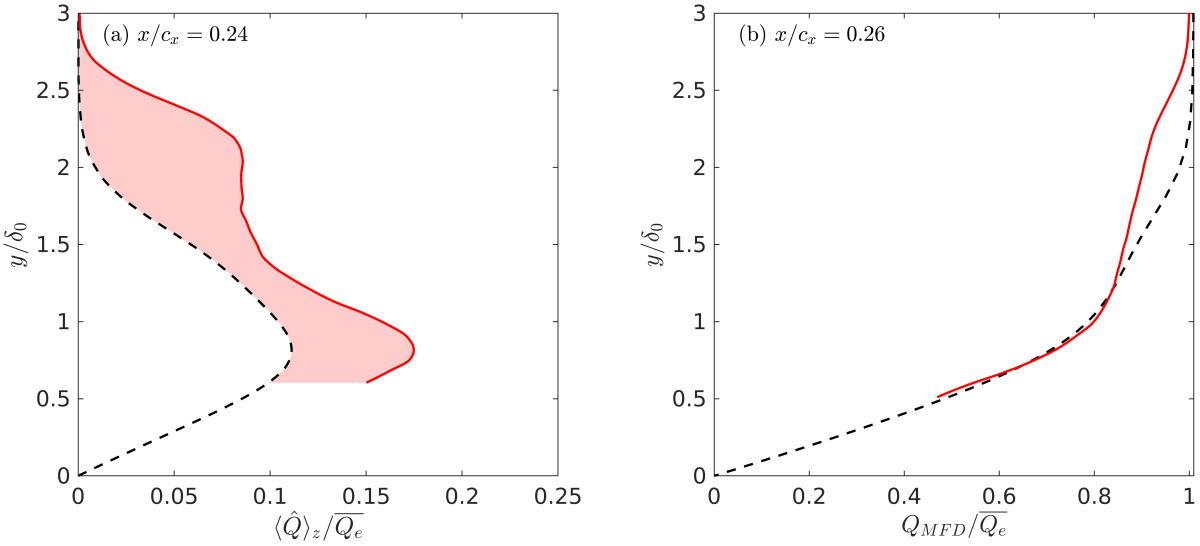

Figure 4.5: (a) Disturbed baseflow example profile indicating error between HWA (–) and NPSE (--) data. (b) MFD example profile comparing HWA (–) and NPSE (--) data.

Matching the stationary crossflow growth with the experimental data is the second crucial step in this research, and the process described so far can allow for errors to creep into the solution. Keeping in mind several checks to perform during this optimisation procedure helps prevent errors. Before running the optimisation code, the NPSE code used is validated through previous work. The NPSE solver used during this project was developed in-house by Westerbeek (2020), and used throughout projects in the past. Casacuberta, Hickel, et al. (2022) achieved a close match between NPSE and DNS with a very similar crossflow setup, using the same in-house solver. With the DNS data of Casacuberta, Hickel, et al. (2022) available, the correct implementation of the NPSE solver in this new research

is validated by matching the DNS data again. Having ensured the NPSE solver works correctly, the optimisation tool can be run, but care has to be taken to avoid local minima in the error. Running several lower fidelity optimisations from very different initial amplitudes helps confirm the absolute minimum is being reached. This should result in the conclusion that the initial amplitude set at $A_0 = 3.5 \times 10^{-3} u_\infty$ does not affect the result, but merely speeds up the optimisation process when it is close to the final result. To also avoid the disproportional influence of a single hotwire plane on the overall results, multiple optimisations with different sets of hotwire planes are run to ensure the result remains the same. Testing different methods to determine the error is a secondary check to ensure the different error determination methods available do not result in significantly different results, but merely in better accuracy. Although the disturbance profiles are objectively the most accurate way to compare, it is thus good to also keep track of the disturbance amplitude and the perturbed baseflow shape. A last check could well be to directly compare the hotwire plane contours with the NPSE data, as two-dimensional NPSE data can be made three-dimensional using the Fourier-ansatz approach.

4.4. Free-Stream Turbulence Generation

Implementation of free-stream turbulence into the flow simulation is fundamentally based on the same principles used in the primary crossflow setup, with an added degree of randomisation. Instead of using a single mode to initiate the disturbance, free-stream turbulence utilises a sum of randomised modes to replicate the characteristics of isotropic turbulence. The approach is adapted from works such as Schlatter (2000) and Brandt et al. (2004), who both used the same principles applied to different types of flow fields. Through the selection of pseudo-random sets of wavenumbers, complementary sets of eigenvalues and eigenfunctions will be extracted that serve as inflow conditions for the free-stream turbulence simulation. Instead of the discrete eigenvalue spectrum, the continuous spectrum is now used, with eigenfunctions that oscillate in the free-stream flow and are damped inside the boundary layer. The weight of the eigenfunctions is determined through wind tunnel data, including turbulence intensity and the turbulence integral length scale. Exactly how many inflow modes are prescribed to form the free-stream turbulence depends on the desired randomness, though literature points towards a number between 400 and 1200 modes in total (Balzer & Fasel, 2016; Brandt et al., 2004; Durovic, 2022). Tests using different numbers of points have been performed on a small DNS domain, concluding the lower limit of 400 free-stream turbulence modes is sufficient to mimic isotropic turbulence in this setup. The first step is selecting 400 sets of wavenumbers to use in the ILST equations. The resulting eigenvalues and eigenfunctions are then filtered to obtain one eigenvalue and eigenfunction for each wavenumber set. These eigenfunctions are then given an amplitude before being used as input in DNS.

4.4.1. Wavenumber Selection

Free-stream turbulence inflow modes can be generated using wavenumbers. Each set consists of three wavenumbers, describing the wave characteristics in streamwise direction (α), wall-normal direction (γ), and spanwise direction (β). These wavenumbers can then be used either as input for the ILST equations to obtain eigenvalues and eigenfunction or for filtering out the correct eigenfunction from the many results provided by ILST. Based on Taylor's Frozen Turbulence Hypothesis, $\omega = \alpha U$, the ω input required for ILST is also provided. β is directly used as input to ILST, and γ dictates the oscillation pattern of the eigenfunction during the filter process. As a first step, 400 pseudo-random sets of wavenumbers will have to be selected, before isotropic turbulence can be replicated. Although the principle behind the wavenumber selection is randomness, a range from which to choose each wavenumber randomly is required. The limits to these ranges are chosen based on both limits in flow physics and grid resolution throughout the domain. Each wavenumber can be converted to a corresponding wavelength $\lambda = 2\pi/\kappa$, where κ is the wavenumber vector. Keeping in mind the Nyquist theorem, waves need to be sampled at a certain rate to be detectable. At least two points need to be located on each wave before they can be characterised. In this work, more margin is left, and at least four grid points need to be within each wavelength. This provides the upper limit of the wavenumber ranges, while the lower limit needs to make sure the wavelengths do not exceed twice the domain size in the respective direction. Depending on the wavenumber, additional specific restrictions are also in place. The spanwise wavenumber is most restricted, and can only be a multiple of the fundamental spanwise wavenumber, fixed to $8mm$ by the experiment. Therefore, the upper limit of the spanwise

wavenumber range is $2\pi\delta_0/0.008m$, where the lower limit is based on the grid refinement and set to four times the spanwise cell size equalling $2\pi\delta_0/(1.11 \times 10^{-4})$ or 47.5497. Limits of the streamwise wavenumber α are entirely based on the grid used. To ensure the grid picks up the oscillations generated, the wavelength must be larger than four grid cells, or $0.5008 \times \delta_0$m. But not larger than twice the domain length, or $1028.5 \times \delta_0$m. The range of α thus lies between 0.0061 and 12.5463. In the wall-normal direction, the only physical limit requires γ to be larger than zero. Using the same grid criteria as with α, γ is limited to a wavelength between four grid cells and twice the domain height. This results in γ wavenumbers between 1 and 11.1150.

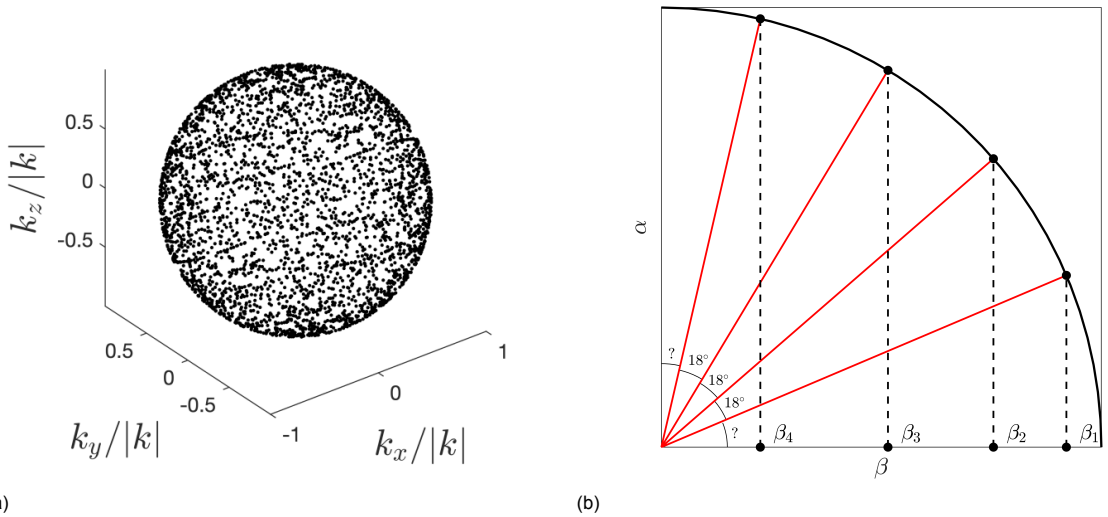

(a) (b)

Figure 4.6: (a) Distribution of points on a sphere with fixed wavenumber magnitude (Durovic, 2022). (b) β distribution on each sphere section, where the question marks indicate the randomly chosen angles between $0°$ and $36°$.

With the allowed ranges of wavenumbers determined, 400 sets need to be randomly chosen within these ranges. A reliable and repeatable method to select these wavenumbers is using concentric spheres, as also used and described by Schlatter (2000), Brandt et al. (2004) and Balzer and Fasel (2016). This approach starts by generating a three-dimensional axis system with each axis representing one of the wavenumbers, within which a number of concentric spheres are placed on the origin, as shown by Figure 4.6a. The radius of the spheres is the total wavenumber magnitude $\kappa = \sqrt{\alpha^2 + \gamma^2 + \beta^2}$. The smallest and largest spheres are limited by the minima and maxima of the allowable ranges of each wavenumber type. To ensure the three wavenumbers can be varied freely for each wavenumber magnitude, the limiting spheres do not have the smallest and largest possible diameters. The smallest sphere is limited by $\kappa = \sqrt{(2min(\alpha))^2 + (2min(\gamma))^2 + (2min(\beta))^2}$ and the largest sphere by $\kappa = \sqrt{min(\alpha)^2 + max(\gamma)^2 + min(\beta)^2}$. The minimum sphere radius used is therefore 2.3968, and the maximum is 11.1346. Equally spaced between the smallest and largest spheres, a predetermined number of concentric spheres are added. Brandt et al. (2004) proposes 20 spheres, and this is the same amount of spheres used in this research. More can be used, as Balzer and Fasel (2016) proposes the use of 40 spheres, but 20 is deemed enough for sufficiently isotropic free-stream turbulence in this flow field. On each of these spheres, an equal number of points is selected. To get to the 400 modes required here, 20 points per sphere are therefore selected. These points need to be randomly chosen while being equally spaced. A commonly used technique is to use polyhedrons with an equal amount of corners as points are required. Placing this polyhedron inside each sphere at a random angle, and selecting the intersection points of both, indicates the points needed. Changing the angle of the polyhedron for each sphere adds the degree of randomness required, as suggested by Schlatter (2000). Because the spanwise wavenumber is limited to discrete multiples of the fundamental value, this technique is not appropriate here. Instead, the choice is made to select the wavenumber sets in two steps. First, the spanwise wavenumbers for each sphere are selected, using a circle with the same radius as the sphere. On this circle, a number of points need to be distributed, equal to the rounded square root of the total points on the sphere. With all wavenumbers being positive, only a quarter of the sphere is used, and therefore the β points are also selected on a quarter of a circle (see Figure 4.6b).

In this case, within each quarter circle, four points have to be distributed equally but randomly. This is done using congruent angles, splitting the 90°quarter-circle into five 18°parts. This ensures each point is spaced at an 18°angle apart, as shown in Figure 4.6b. Enough room is left to randomly rotate these points by an angle between 0°and 18°to ensure randomness between each of the spheres. By projecting these points onto the x-axis and rounding these projected values to the nearest allowed value, the spanwise wavenumbers are chosen. This β distribution system is represented visually by Figure 4.6b. These wavenumbers indicate the different allowed spanwise values which, dependent on the total amount of points required, can return multiple times throughout the selection process.

The amount of times each spanwise wavenumber returns is based on its magnitude. Each spanwise wavenumber essentially generates a new circle when imagining a plane perpendicular to the β-axis at the selected β value, intersecting with the sphere. The larger the value for β, the smaller the circle generated by the intersection will be. Figure 4.7a gives a visual representation of these 2D circles generated from the intersection with the 3D sphere. Four of these circles are generated in this work, for four positive β values. The 20 points that need to be distributed on each sphere have to be distributed on these four circles. This is done based on each circle's circumference, where the circumference of all circles is summed and each circle gets assigned a fraction of the total points based on its circumference compared to the total circumference sum. Once the amount of points for each circle is defined, the points are distributed on the circles in the exact same way as done previously for β. Only now, the angle is different based on the number of points, but the principle remains the same. Because both α and γ are selected simultaneously, no projection is necessary, and the wavenumbers can be directly extracted from both axes. This principle is visualised in Figure 4.7b. Repeating this process for each sphere results in 400 sets of three wavenumbers, which can be used to generate 400 inflow eigenmodes.

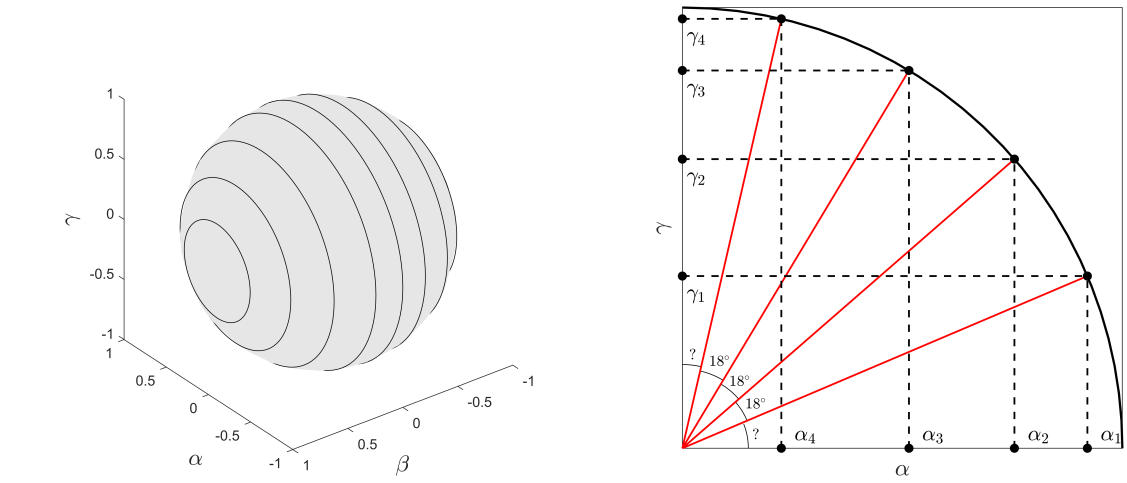

(a) (b)

Figure 4.7: (a) 2D slices of the 3D wavenumber spheres at constant, previously determined β-values. (b) α and γ distribution system on each 2D slice, where the question marks indicate randomly chosen angles between 0° and 36°.

4.4.2. Eigenfunction Extraction

Just as with the primary crossflow mode, ILST is used to generate an eigenvalue and eigenfunction spectrum based on each of the wavenumber sets. Where the discrete part of the spectrum was used for the primary crossflow disturbance, the continuous spectrum is used to generate free-stream turbulence. This theoretically infinite spectrum of eigenvalues features eigenfunctions exhibiting oscillating behaviour outside the boundary layer. Inside the boundary layer, these eigenfunctions are damped, while outside the boundary layer, they converge towards a constant oscillation when $y \to \infty$. The damped behaviour inside the boundary layer is the result of shear sheltering, referring to the inability of the continuous Orr-Sommerfeld modes to penetrate into the boundary layer (Jacobs & Durbin, 1998). When a large number of these oscillatory modes are summed, the chaotic-looking velocity profile that results resembles isotropic free-stream turbulence. An example of such an eigenfunction of the continuous spectrum showing oscillatory motion is presented by Figure 4.8.

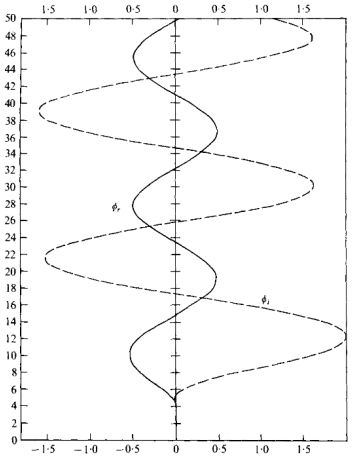

Figure 4.8: Real (ϕ_r) and imaginary (ϕ_i) parts of an eigenfunction in the continuous spectrum (Grosch & Salwen, 1978).

This is only one of a large spectrum ILST outputs for a single wavenumber set. To generate free-stream turbulence, a specific eigenfunction and eigenvalue from this continuous spectrum have to be selected for each mode. This is done using the wall-normal wavenumber γ, which dictates the wall-normal oscillation pattern desired. If the continuous spectrum would truly be infinite, an exact eigenfunction and eigenvalue pair could be picked. However, because of computational limits the continuous spectrum that results from ILST is not infinite in size, but sampled based on a Chebyshev polynomial. Therefore curve-fitting to find the closest matching result is used. From the wall-normal wavenumber γ, combined with the eigenvalue α, an oscillatory curve can be generated using,

$$\tilde{v} = e^{\gamma y} + e^{-\gamma y} + e^{\alpha y}; \quad y \to \infty \tag{4.8}$$

This curve represents the desired eigenfunction behaviour, for large y-coordinates. Through curve-fitting the continuous mode eigenfunctions generated with ILST to this curve, the closest match is selected. This eigenfunction and its accompanying eigenvalue are then chosen as inputs to the DNS solver. The process is repeated for all 400 disturbance modes. With the eigenfunctions and eigenvalues for each of the 400 disturbance modes known, the only missing information before all inflow boundary conditions can be generated are the weights. The eigenfunctions all have a unit amplitude and therefore require adequate weights to determine their contribution to the free-stream turbulence. The weight of each eigenfunction is assigned as a wave amplitude, calculated in the next section.

4.4.3. Amplitude Determination

With the eigenvalues and eigenfunctions for each mode known, the only unknown left is the amplitude assigned to each eigenfunction. Determining this amplitude requires two flow characteristics to be calculated, the turbulence integral length scale and the turbulence intensity. The latter is a known characteristic of the wind tunnel, but can also be extracted from pitot and HWA data, given the accuracy required for this research. Using a HWA probe placed outside of the boundary layer, an instantaneous velocity profile is extracted. Subtracting the mean velocity from this profile and taking the root-mean-squared value, results in the turbulence intensity. Understand that this turbulence intensity is calculated based on velocity magnitude Q measured by the HWA probe and not via the standard way with individual velocity components. A correction will need to be applied to take this into account, though this only happens right before implementation into DNS. For now, this value is being used. The turbulence integral length scale is the only unknown left. Representing the largest turbulent scales in the flow field, this value can be extracted using the auto-correlation function on the free-stream instantaneous velocity data. Since measurements on a single point are used, the turbulence integral time scale is first determined. The auto-correlation function assesses velocity fluctuations over a period of time and looks for self-similarities in the signal, to determine how long it takes for a signal to repeat itself. Several ways of assessing the auto-correlation function exist, but most commonly the point where the function reaches zero is taken (Trush et al., 2020). With the turbulence integral time scale known, the length scale is obtained by multiplying the time scale by the free-stream velocity magnitude. In mathematical form, this auto-correlation function is represented as,

$$L_{11} = \int_0^{\infty} \frac{\overline{u(t)u(t+\tau)}}{\overline{u^2}} d\tau. \tag{4.9}$$

With both the turbulence integral length scale ($L11$) and the turbulence intensity known, the energy in the flow field associated with each wavenumber magnitude can be determined through the following formula,

$$E(\kappa) = TuL_{11} \frac{1.196\,(\kappa L_{11})^4}{\left[0.558 + (\kappa L_{11})^2\right]^{\frac{17}{6}}}, \tag{4.10}$$

where Tu is the turbulence intensity value as a fraction of one, κ is the wave number magnitude, and L_{11} is the turbulence integral length scale. The result is an energy over wavenumber curve, of which an example is shown in Figure 4.9a. With the energy assigned to each wavenumber known, the energy per sphere is known as well. Through a simple equation, this energy of each sphere can be distributed over the number of points on each sphere, as given by,

$$A(\kappa) = \sqrt{\frac{2E(\kappa)\Delta\kappa}{N_p}}, \tag{4.11}$$

where $\Delta\kappa$ is the spacing between two spheres, and N_p is the number of points per sphere. With all variables required for DNS implementation known, only one thing is left and that is applying the turbulence intensity correction mentioned earlier. For this, the free-stream turbulence inflow plane is first constructed using Equation 4.12, summing all 400 modes.

$$q' = \sum_{n=1}^{400} A_n \cdot q_n' \cdot e^{-i(\alpha_n x + \beta_n z - \omega_n t)}. \tag{4.12}$$

This results in a plane similar to the example given by Figure 4.9b, where isotropic velocity fluctuations impose turbulence on the free-stream flow but get damped out rapidly inside the boundary layer. The turbulence intensity imposed on the free-stream flow by this plane can now be calculated as measured by the hotwire. Both the plane and the hotwire data should result in the same value. If this is not the case, the difference needs to be determined, and every amplitude multiplied by this correction factor before input in DNS. This ensures the turbulence intensity imposed on the DNS is exactly the same as measured during the experiment.

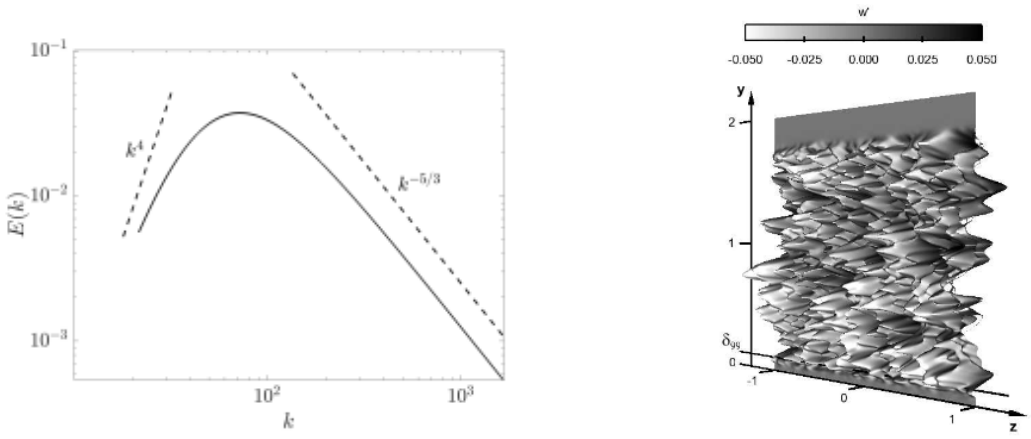

(a) (b)

Figure 4.9: (a) Energy distribution over the wavenumber space (Durovic, 2022). (b) Exemplary inflow boundary conditions for free-stream turbulence implementation (Balzer & Fasel, 2016).

4.5. DNS Implementation

Throughout this work, the INCA solver is used to perform DNS. Developed in-house at the Delft University of Technology, this is an extensive multi-physics CFD solver capable of performing high-fidelity complex LES and DNS simulations (Hickel et al., 2014). This system provides the flexibility required for the use and implementation of the numerical techniques proposed in this work. Through an explicit third-order Runge-Kutta method, the Navier-Stokes equations are marched forward in time to solve the flow field. As a general convergence criteria, the L_2-norm of the temporal derivatives, $\epsilon_{L_2} = 10^{-8}$, is chosen. For an efficient way of reproducing the experimental flow field, the methodology introduced in this chapter is implemented in steps. Following the order of this chapter, baseflow is the first flow field that needs to be matched with experimental data. Due to the spanwise invariance of the baseflow, the flow in z-direction does not need to be resolved. This allows the removal of all but two spanwise grid cells, saving considerable resources at this stage of the simulation. Baseflow implementation into DNS requires several parameters to be passed to the solver, including the Reynolds number, and the free-stream velocity components at the inflow plane. The external velocity distribution in the streamwise direction, calculated in subsection 4.2.2, is prescribed through a logarithmic polynomial. The coefficients, as well as the offset added to take into account the attachment line location, are presented by Equation 4.3. Once the solution based on these parameters converges, a three-dimensional baseflow solution remains. The only experimental data that can be used for comparison is the measured pressure distribution of the wing. Boundary layer theory is used as a second method to validate the DNS baseflow results. When all three data types agree well with each other, the steady flow field simulation can be initiated.

To simulate this steady flow field by introducing primary CFI, a flow field resolved in all three dimensions is required again, achieved by multiplying the total number of grid cells by 72. After this mesh adaptation, only adding the primary crossflow disturbance mode to the baseflow setup is required to trigger primary crossflow instabilities. This consists of a set of velocity and pressure profiles, an eigenvalue, a spanwise wavenumber and a temporal frequency, along with an amplitude assigned to the normalised eigenfunctions. From these inputs, the inflow disturbance plane is obtained through

$$q' = A \cdot \hat{q} \cdot e^{-i(\alpha x + \beta z)}, \tag{4.13}$$

where $q = [u, v, w]^T$, A the disturbance amplitude, α the eigenvalue, and β the spanwise fundamental wavenumber. Adding this disturbance plane to the existing inflow boundary condition risks introducing convergence issues. To reach full convergence and enforce a stationary crossflow field, Selective Frequency damping (SFD) is therefore applied from $x/\delta_0 = 230$ until the outflow boundary of the domain. Downstream of $x/\delta_0 = 230$, SFD is applied across the full width and height of the domain. In this region, the convergence criterion is calculated differently, taking the L_2-norm of the difference between the instantaneous flow solution and the low-pass filtered solution $\bar{q}$. Referred to as ϵ_{SFD}, this new convergence target is set to $\epsilon_{SFD} = \|q - \bar{q}\|_{L_2} = 10^{-10}$. The goal of this method is to obtain a steady-state solution q_s, equal to the instantaneous solution q. Because q_s is not a known state, the low-pass filtered solution is instead used in the evolution equation of q, which reads as

$$\dot{q} = f(\mathbf{q}) - \chi(\mathbf{q} - \bar{\mathbf{q}}), \qquad \chi \in \mathbb{R}, \tag{4.14}$$

where $\dot{q}$ is time derivative of the instantaneous solution q, and f the non-linear Navier-Stokes operator (Casacuberta et al., 2018). χ is the feedback control coefficient, set to suppress instabilities with unstable frequencies. Because the integral definition of a low-pass time filter would be too resource intensive, since the complete time history of the signal would need to be stored, an equivalent differential form is being used,

$$\dot{\bar{q}} = \frac{q - \bar{q}}{\Delta}, \qquad \Delta \in \mathbb{R}_{>0}. \tag{4.15}$$

Δ is referred to as the time constant, replacing the cut-off frequency of the low-pass filter as $\Delta = 1/\omega_c$ (Åkervik et al., 2006). This differential form can be time-marched less computationally costly, and results in an easier way of implementing SFD. The efficiency of SFD, however, is strongly dependent on the values of χ and Δ. These parameters in the INCA solver have been extensively studied by Casacuberta et al. (2018) among others and, therefore, will not be altered during this research.

Adding free-stream turbulence onto the converged steady flow field results in the final unsteady flow field. By summing the 400 disturbance modes calculated in section 4.4, processing them using

Equation 4.12 and adding them to the steady flow field, all required boundary conditions are present. Since free-stream turbulence is inherently unstable, the SFD zone used in the steady case will have to be reduced in size. Otherwise, it will rapidly filter out the free-stream turbulence fluctuations. SFD is still applied in the unsteady simulation, but only from $x/c_x = 0.489$ until the outflow boundary, to ensure a stable outflow boundary condition. Through the process of receptivity, these disturbances will introduce secondary instabilities into the boundary layer. These secondary instabilities grow rapidly and ultimately cause transition to turbulence. Because less free-stream shear stress is expected to be present in the numerical simulation compared to the wind tunnel, it is expected that the free-stream turbulence will dissipate quicker in DNS. Considering the most essential part of the receptivity process occurs far upstream of the transition region, where dissipation is in any case still low, reduces concerns of dissipation playing a crucial role (Brandt et al., 2004). When transition occurs in the flow field, full convergence will no longer be possible in DNS. Turbulent boundary layers are unsteady and thus ever-changing. Instead, this simulation is run until the boundary layer characteristics and the transition front sufficiently stabilise. After which the necessary flow characteristics can be extracted to assess the match with experimental data. The DNS setup of a second numerical simulation without any unsteady forcing uses the same steady flow field as the free-stream turbulence simulation. Instead of introducing a large number of disturbance modes at the inflow, this time the flow is left to develop by itself, and only the SFD zone imposed on the transition region gets removed. The removal of this SFD zone ensures the flow is allowed to transition from laminar to turbulent when unsteadiness develops along the domain. This is the ideal baseline simulation to assess the impact of the FST on both CFI development and transition throughout the numerical domain.

To assess the transition location and mechanisms of both simulations, several statistical metrics must be extracted from the flow field. When the unsteady simulation has run sufficiently long, and reached a repeating pattern, the flow field is averaged to extract the shear stress in the flow. Based on the friction coefficient c_f derived from these shear stresses, the transition location is extracted as the point of maximum $\partial c_f/\partial x$. This metric is comparable to the IR thermography used in the experimental data by Rius-Vidales and Kotsonis (2022), as it produces the most visible separation between the laminar and turbulent flow states. Probe data is also sampled for each iteration, which provides point-sampled instantaneous velocity data of the flow field. Used to assess the frequency content of the flow field, this probe data can be directly compared to the instantaneous HWA data sampled during the experiments. Sampling of the entire flow field every 4 iterations is also done to analyse the intensity and distribution of the velocity fluctuation across the crossflow vortex. By taking the standard deviation (σ_Q) of the instantaneous velocity fluctuations at each point in a z-y plane, this distribution can be assessed spatially through,

$$\sigma_Q = \sqrt{\frac{1}{n-1}\sum_{i=1}^{N} |Q - \overline{Q}|^2}. \tag{4.16}$$

As a metric to keep track of the overall instability growth a_Q, this σ_Q can be integrated over the entire z-y plane, given as,

$$a_Q = \frac{1}{\overline{Q_e}} \iint_S \sigma(y,z) dz dy, \tag{4.17}$$

where $\overline{Q_e}$ is the free-stream component of the HWA-measured time-averaged Q-velocity magnitude. By applying a frequency filter to the instantaneous velocity signal used to calculate σ_Q, the instability growth of specific frequencies can be extracted as well. To investigate which frequencies are dominating the transition region inside the boundary layer, the signal frequency content is extracted through a Power Spectral Density (PSD) analysis. The numerical velocity signal used spans a sampling time of 0.025 seconds, much lower than the two seconds samples extracted experimentally. However, the sampling rate is significantly higher at $1MHz$ compared to $51.2kHz$, related to the small time step size required to resolve the turbulent structures. This implies the Nyquist frequency is well above the frequency range of interest between $1kHz$ and $10kHz$, determined from the experimental data presented in chapter 2.

A large amount of computational resources is needed to solve the DNS flow field and sample the results. Several different computational clusters are used during this research, depending on the needs during each simulation stage. Baseflow DNS, together with any optimisations through Matlab are run in HPC12. This smaller cluster is part of the aerodynamics department at TU Delft and allows anything

from single-core computations up to simulations with several hundred CPU cores. Although insufficient for full unsteady DNS, this computer allows fast turn-around times for these smaller simulations. For the steady primary crossflow and the full unsteady solution, the Delft High-Performance Computing Centre (DHPC) is used (Delft High Performance Computing Centre (DHPC), 2022). With a speed of 2 petaFLOPS, this new computing system is ranked among the top 500 fastest supercomputers in the world, allowing several thousands of CPU cores to be used for these computationally expensive unsteady DNS runs. For the final stages of the unsteady simulations, and to obtain flow statistics related to these simulations, this work used the Dutch national e-infrastructure with the support of the SURF Cooperative using grant no. EINF-5929. The 14 petaFLOPS Snellius system managed by SURF (www.surf.nl) is able to provide the large number of processors required for a large-scale DNS project as is performed here. The EINF grant SURF provided the researcher with, included 1,000,000 CPU hours on Snellius to conclude the final stages of the project. Something the author of this document is very thankful for.

5

Results

Following the methodology results in a three-dimensional numerical flow field, post-processed equivalently to the experimental data to allow evaluation of the match between both data types. Analysis of the baseflow, steady and unsteady flow field enables an extensive discussion on the value of this work for future numerical and experimental research on incompressible flows on swept wings. To put the unsteady results into perspective, the numerical baseflow and steady crossflow results are discussed first. This serves as a starting point for unsteady flow simulations and provides an idea of the method's accuracy for each step of the process, indicating that numerical replication of crossflow instabilities is feasible for each step of the process individually. When discussing the unsteady flow field towards the end of this chapter, the results focus on the three main pillars of transition location, transition type, and frequency content of the crossflow vortices. Showcasing the capabilities of this FST-based method and its value for future research. The latter is backed up by the last section of this chapter, showing a comparison between the unsteady numerical flow fields obtained with and without the addition of free-stream turbulence. This provides a clear overview of the benefits and downsides of the methodology related to incompressible CFI research on swept wings.

5.1. Baseflow Simulation

Based on the test section pitot tube, the wind tunnel model is found to encounter a free-stream velocity magnitude equalling 24.47 m/s. Based on the identification of the attachment line described in subsection 4.2.2, the velocity magnitude is decomposed into a chordwise velocity component, $U_\infty = 17.58$ m/s, and a spanwise velocity component, $W_\infty = 17.02$ m/s, at the attachment line location. These components make a 45.94°inflow angle, relative to the model's leading edge. By use of the pressure distribution over the airfoil, the free-stream velocity magnitude (U_e) along the chord (x/c_x) is calculated. Divided into two rows in X-direction, the surface pressure ports allowed the generation of two velocity profiles, as given by Figure 5.1. The difference between both profiles is considered minor, leading to the author's decision to accept the spanwise invariant flow assumption. This acceptance simultaneously confirms the validity of the infinite swept-wing approach, therefore allowing for both profiles to be averaged to a single U_e distribution (see red dotted line in Figure 5.1). As the W_∞ is assumed constant across the chord, following the infinite swept wing assumption, the spanwise velocity profile (w_e) is now determined as well. The averaged u_e profile that results from decomposition (see dotted line in Figure 5.1) is then used as an input for the CFD solver. The raw u_e profile is too variable so a fourth-order polynomial is used to approximate this u_e distribution instead (see green line in Figure 5.1). Before implementation, the polynomial is cropped to the numerical domain length, from $x/c_x = 0.05$ to the outflow boundary at $x/c_x = 0.47$. Both numerical boundaries are indicated by black dashed lines on Figure 5.1. At the numerical inflow plane, the free-stream velocities are determined to be $u_\infty = 12.81$ m/s along the chord (x-direction) and $w_\infty = w_e = 17.02$ m/s along the span (z-direction). All numerical input values are normalised by the chordwise inflow velocity u_∞, resulting in non-dimensional velocities at the inflow plane of 1 and 1.3286, respectively. Based on this velocity data extracted from the experimental flow field, the velocity distribution can be decomposed and implemented into DNS, resulting in a baseflow tailored to match the experiment.

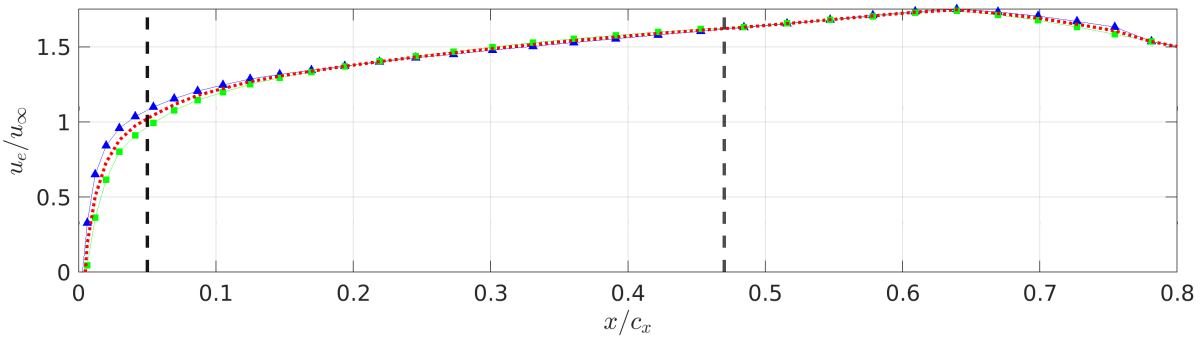

Figure 5.1: Distribution of experimental u/u_∞ along x/c_x for both upper (△) and lower (□) pressure port rows, averaged (··) for numerical processing. Black dashed line indicates the numerical in- and outflow boundaries.

The resulting numerical baseflow is simulated in INCA as a three-dimensional flow field with two cells in z-direction, because of the spanwise invariance of the parameters. This implies the spanwise dimension of the domain is allowed to remain unresolved at this stage. However, because of the already three-dimensional geometry, the baseflow can directly serve as input for the steady crossflow simulation. The chordwise velocity distribution u/u_∞ of the baseflow can be represented as a z-y_t slice (see Figure 5.2) and shows a growing boundary layer in x-direction. The overall increase of free-stream velocity u_e imposed on the domain to emulate the flow acceleration along the curved airfoil surface in the wind tunnel experiment is also apparent in Figure 5.2, as the boundary layer is only a fraction of the total domain size.

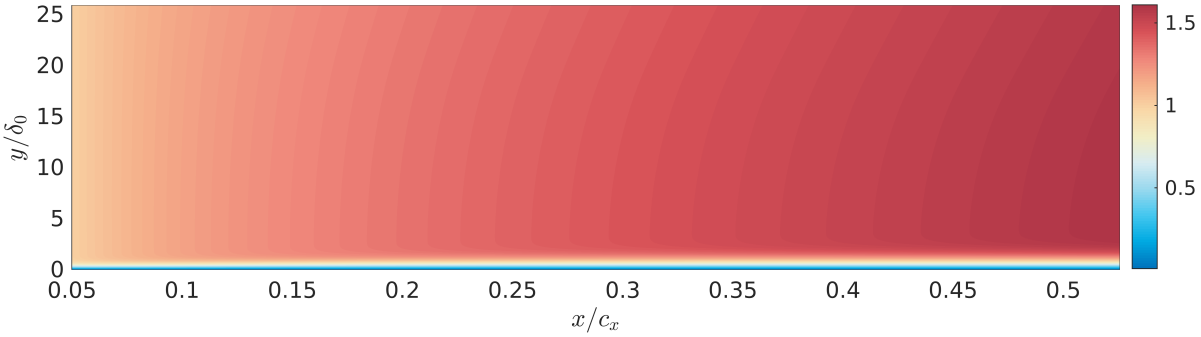

Figure 5.2: Numerical baseflow z-y slice indicating the two-dimensional u/u_∞ distribution, simulated using INCA and converged to $\epsilon_{L_2} = 1 \times 10^{-8}$.

Since the HWA data of the experiment only provides details on the instantaneous flow field and lacks isolated baseflow information, there is no possibility for a comprehensive comparison with the experiment. Neither visually, nor based on HWA data points. Therefore, the baseflow simulation is compared back to the experimental data through the airfoil pressure ports. The velocity distribution obtained from the experimental pressure data through Equation 4.2 is based on incompressible and inviscid flow assumptions. Nevertheless, the velocities considered here are small, ensuring the validity of these assumptions at this stage of the research. Comparison between the numerical and experimental velocity distributions shows minimal differences between both data sets (see Figure 5.3). The slight differences that can still be distinguished between the different data types are thought to be largely attributed to the averaging of the pressure port data and expected variation in the experiment due to unsteadiness. Considering the amount of processing required to arrive at this numerical flow field, the differences between both curves are labelled negligible, confirming the experimental data is correctly applied in the numerical setup.

Additionally, a theoretical boundary layer (BL) solver based on the Falkner-Skan Cooke (FSC) equations is used to verify if the numerical baseflow is generated correctly (Falkner & Skan, 1931). Rather than a comparison with experimental data, this BL solver is used for verification purposes only. A direct comparison between the numerical external velocity and the experimental velocity based on the pressure ports already indicated the differences in u_e remain negligible (see Figure 5.3), and already partially confirms the validity of the baseflow used during the remainder of this work. To fully verify, the

results are compared more extensively to the BL solver. This includes the wall-normal velocity v_e/v_∞, pressure p/p_∞, and the boundary layer thicknesses δ^*/δ_0 and δ_{99}/δ_0. All of the aforementioned comparisons are visualised in Figure 5.4, showing both solvers are in good agreement with each other, increasing confidence in the numerical baseflow solution. Although small differences between the BL solver and the DNS results are still noticeable, these are expected based on the assumptions the BL solver is based on. This is mainly thought to be related, but is not limited to, the assumption that $\partial p/\partial x = 0$ in the BL solver solution.

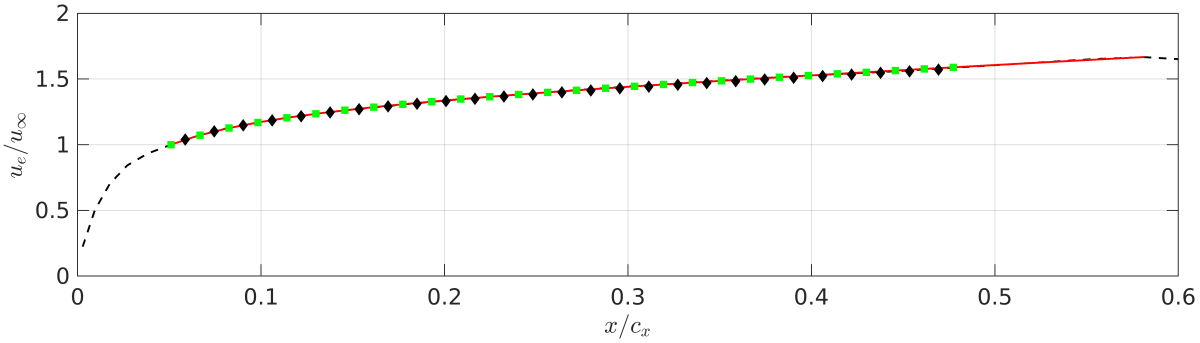

Figure 5.3: Chordwise distribution of u/u_∞ for the interpolated experimental data (--), the polynomial fit to the experimental data (-), the BL solver result (□), and the DNS results from the INCA simulation (◇).

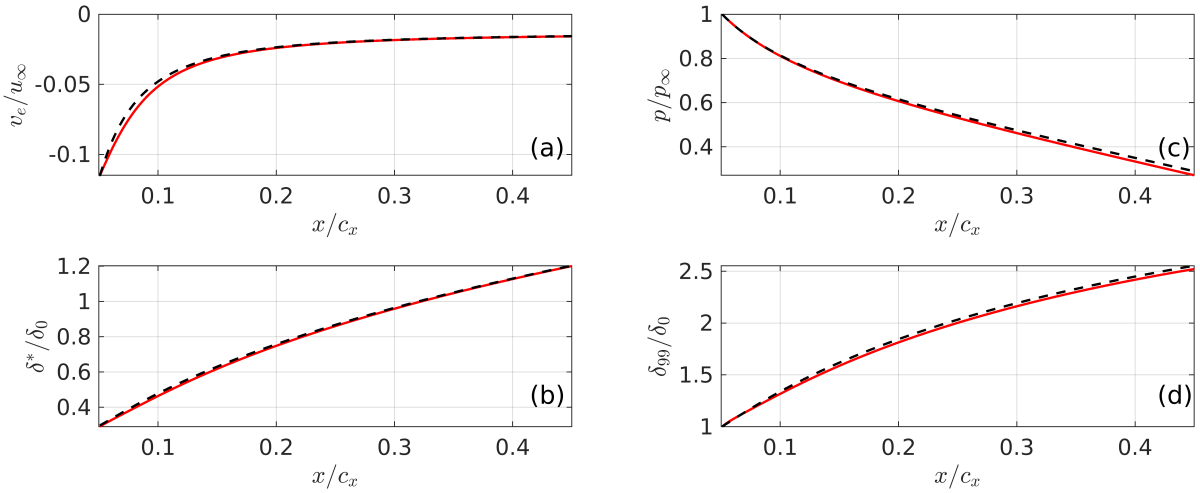

Figure 5.4: DNS baseflow comparison with the FSC boundary layer equations for wall-normal velocity v_e/u_∞ (a), pressure p/p_∞ (b), and boundary layer thicknesses δ^*/δ_0 (c) and δ_{99}/δ_0 (d).

5.2. Primary Crossflow Instability Development

Adding primary CFI to the numerical baseflow results in the steady flow field, which is characterised by a wave-like pattern generated by the crossflow vortices inside the boundary layer. At this stage of the numerical process, the HWA data comes into play as a reference for setting up the steady flow field simulation. Although the HWA data is instantaneous, time-averaging allows the recreation of the steady flow field from unsteady measurement data. The time-averaged HWA data can therefore be directly compared to its numerical equivalent, assessing the accuracy of the steady simulation. Until close to the transition region, primary CFI dominates the boundary layer. Their increase in amplitude when propagating downstream leads to the formation of secondary crossflow instabilities, which tend to grow exponentially and lead to rapid breakdown to turbulence. As outlined in section 4.3, a trial-and-error-based optimisation is used at this stage of the process. A single disturbance mode is generated based on the fundamental spanwise wavelength of $8mm$ and imposed as an inflow condition on the baseflow. The addition of this single disturbance results in the numerical steady flow field. Through ILST, the shape of this disturbance and the accompanying eigenvalue and wavenumbers that form the inflow

boundary condition are determined. The shape of the disturbance consists of a velocity vector along the y-direction for the u, v, and w components, along with the pressure. These vectors, which can be made two-dimensional to cover the inflow plane using Equation 4.5, are given by Figure 5.5 for the readers' information. Next to the velocity and pressure vectors, the non-dimensional spanwise fundamental wavenumber $\beta = 0.6604$ is part of the disturbance mode, next to the temporal angular frequency $\omega = 0$. These wavenumbers are necessary to recreate the two-dimensional field from the previously determined vectors. $\omega = 0$ because the flow at this stage is steady, and no temporal fluctuations are desired. The complex eigenvalue $\alpha = 0.7669 - 0.0141i$ is also part of the primary CFI disturbance mode, influencing the disturbance shape throughout the ghost-cells in front of the domain inflow, as discussed on subsection 4.3.1. These inputs are part of the discrete eigenvalue spectrum of the Orr-Sommerfeld equation, with the main characteristic that the velocity is exponentially damped outside the boundary layer. This agrees with the desire to only introduce this steady disturbance inside the boundary layer. At this stage, these inputs are still normalised by the chordwise velocity magnitude at the inflow plane. With the inflow boundary condition shapes known, only the amplitude for these normalised functions now needs to be determined, which is discussed in subsection 5.2.1.

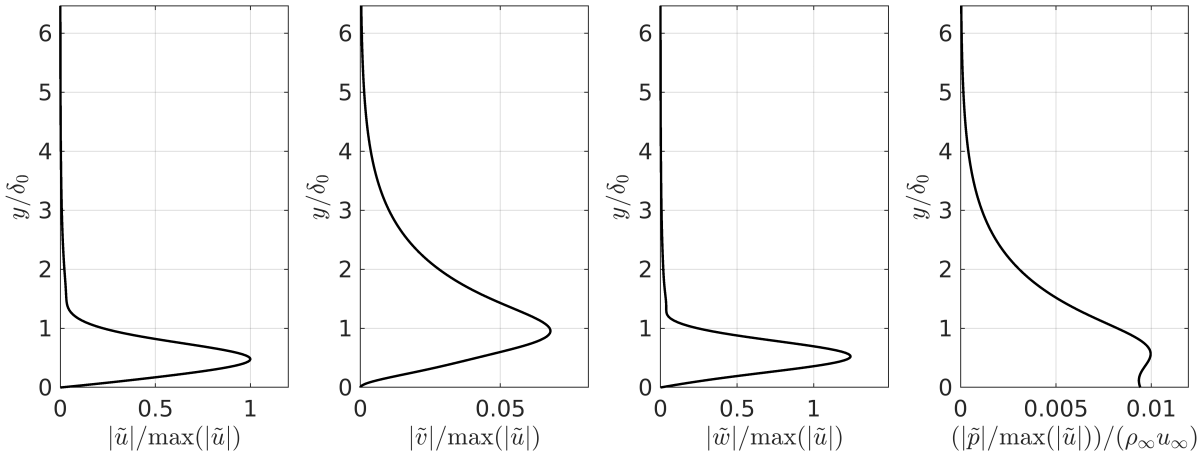

Figure 5.5: Normalised $\tilde{u}$, $\tilde{v}$, $\tilde{w}$ and $\tilde{p}$ vectors in y-direction representing the primary CFI at the inflow boundary.

5.2.1. Disturbance Amplitude Determination

The velocity and pressure values at the inflow are normalised, implying the need for an adequate amplitude to scale the disturbances to the levels measured by the hotwire during the experiment. Based on the methodology outlined in subsection 5.2.1, an optimisation script is run to determine the amplitude required to obtain the closest match with experimental data possible. To reduce cost, this optimisation procedure uses NPSE instead of DNS. From similar research performed by Casacuberta et al. (2021), an initial amplitude $A_0/u_\infty = 3.5 \times 10^{-3}$ is chosen as initial estimation. To show a baseline comparison with experiments, before optimisation, Figure 5.6 compares the NPSE results with $A_0/u_\infty = 3.5 \times 10^{-3}$ for maximum disturbance amplitude, mean flow distortion (MFD), and disturbance profile. The differences between the numerical and experimental data at this stage are clear, as the experimental profiles are visibly more saturated than the numerical results with $A_0/u_\infty = 3.5 \times 10^{-3}$. A much closer match is required to reproduce the conditions of the experiments, pointing out the need for an optimisation procedure.

The iterative optimiser walks through a number of different amplitudes to localise the point of minimum error between the disturbance profiles. Based on the integral of the difference between numerical and experimental disturbance profiles, the error metric converges gradually to a minimum. The steps taken by the optimiser while iterating towards the optimal amplitude are shown in Figure 5.7a. The resulting amplitude for the primary CFI inflow boundary condition equals $A_0/u_\infty = 1.9401 \times 10^{-2}$. The new maximum amplitude curve in Figure 5.7b creates a much-improved match with experimental data. Experimental data inevitably contains outliers, because of the lack of time dependence between the different hotwire measurement points. Due to the time-intensive process of acquiring hotwire data, the different planes used to match NPSE were measured on different days. Different data points on the experimental curve are therefore independent of each other. Details on this variation in wind tunnel

parameters can be found in section 2.2, of which the impact can not be underestimated. However, by taking all eight planes into account, the optimisation is based on all the experimental data points available. The optimisation criteria based on the disturbance profile, calculated by Equation 4.7 and chosen in subsection 5.2.1, lead the author to believe the results in Figure 5.7b are the closest fit to the experimental data possible with the tools available. Potential limitations in the current numerical setup and the inflow boundary conditions used not known to the author at this point, could in the future lead to further increased accuracy. However, given the variability of the experimental data for the steady flow field, this work considers the matching procedure a success, limiting the error in disturbance amplitude to well below 5%.

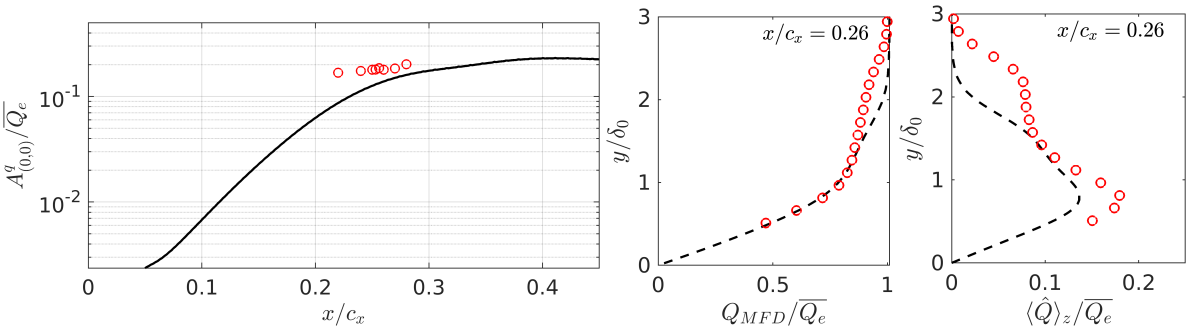

(a) (b) (c)

Figure 5.6: Comparison between NPSE (--) and experimental data (∘) with $A_0/u_\infty = 3.5 \times 10^{-3}$ for (a) maximum disturbance amplitude, (b) MFD and (c) disturbance profile

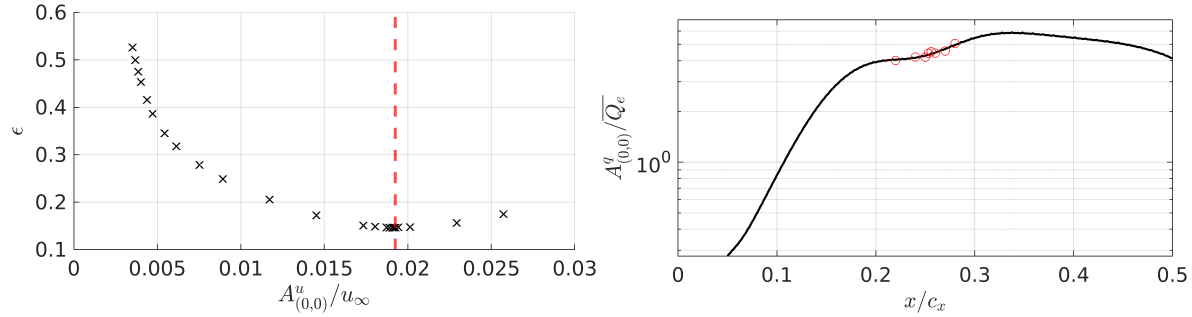

(a) (b)

Figure 5.7: (a) Convergence plot of $A^q_{(0,0)}$ based on the error criterion evaluation (ϵ). (b) $A^q_{(0,0)}$ comparison between NPSE with $A_0/u_\infty = 1.9401 \times 10^{-2}$ and experimental data.

Next to disturbance amplitude, more detailed techniques to compare NPSE with experimental data use MFD and disturbance profiles. Eight HWA planes are available to generate profiles for comparison. Four are used here as a basis for analysis for both comparison types (see Figure 5.8a and Figure 5.8b). To avoid misinterpretation of the difference between numerical and experimental curves, these are normalised by the local free-stream velocity, to reduce the effect of individual HWA planes being gathered under slightly different conditions. Because a difference in free-stream velocity could have an influence on the near-wall velocity gradient, small errors are still to be expected. The MFD comparison in Figure 5.8a confirms the accurate match of the steady crossflow field, despite the small errors near the wall starting to creep in near $x/c_x = 0.28$. Unable to pinpoint the exact reason for these discrepancies, the optimiser is still confident this disturbance amplitude represents the best numerical setup, based on the available reference data. However, MFD profiles are still limited in the information they provide. In an attempt to include more information on the experimental flow field, disturbance profiles are compared as well. Their complex profiles are harder to match, as they do not contain spatially averaged data like the MFD. Using the root-mean-square of the time-averaged velocity data subtracted by the MFD, they contain more information about the flow than the MFD which is averaged in z-direction, or the amplitude profile which only takes the maximum disturbance amplitude into account. The overall dis-

turbance profile shape matches closely between both data sets, albeit with noticeable differences near the upper region of the boundary layer. Although the lower of the two peaks of each curve almost overlap, persisting differences remain around the upper peak. With inter-plane variations between profiles near this upper peak too significant, the choice was made to find the best overall match possible overall HWA planes. The variation between planes is thought to be a feature of variation in experimental data and the varying wind tunnel parameters for individual HWA planes. Individual planes could have been matched closer, but the challenge in this research concerns the matching of the full flow field. Given that previous work from the likes of Nishino and Shariff (2010) and Tempelmann et al. (2011) reported significant differences from 60% to orders of magnitude, the matching procedure outlined here can be considered satisfactory.

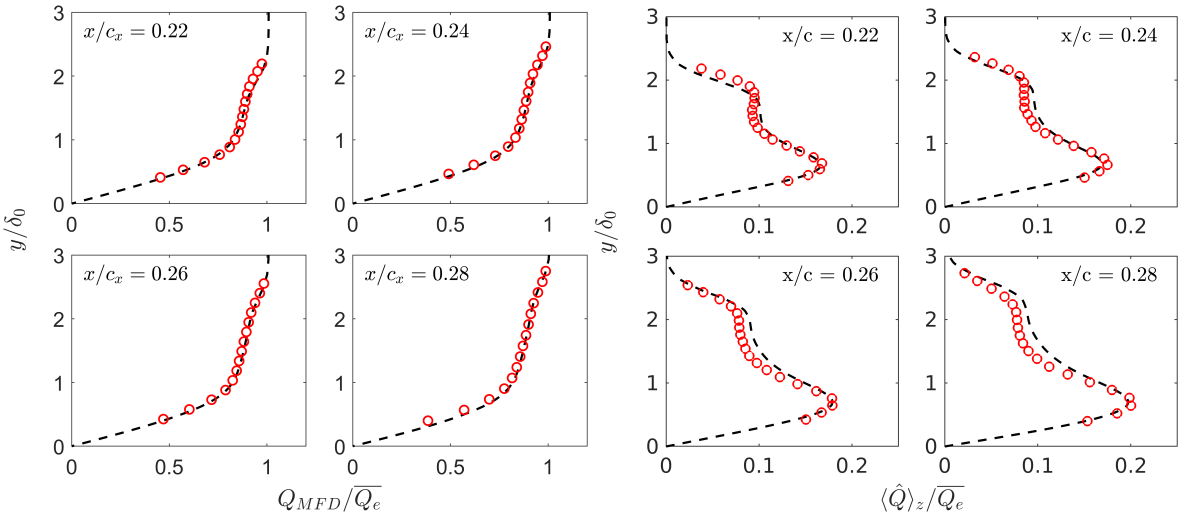

(a) MFD profiles.

(b) Disturbance profiles.

Figure 5.8: NPSE (--) MFD and disturbance profile comparison with time-averaged HWA planes (◦) for x/c_x = 0.22, 0.24, 0.26 and 0.28.

The resulting optimal amplitude turned out to be over five times larger than the initial amplitude fed to the optimisation tool. To ensure the global minimum error is found and local minima are disregarded, several tests are run with varying initial amplitudes. Next to the initial guess of $A_0/u_\infty = 3.5 \times 10^{-3}$, therefore, $A_0/u_\infty = 0.026$ and $A_0/u_\infty = 0.05$ are run as well. Keeping in mind the historically widespread research into this type of flow field, the researcher is confident the optimal amplitude is within these bounds. No irregularities are found within any of the convergence curves, which all appear to converge smoothly to a single minimum value. Direct comparison of all the three initial amplitude convergence curves (see Figure 5.9) confirms the effectiveness of the optimiser. The convergence curves largely overlap and the same optimal disturbance amplitude is found each time, increasing confidence in the NPSE optimisation results.

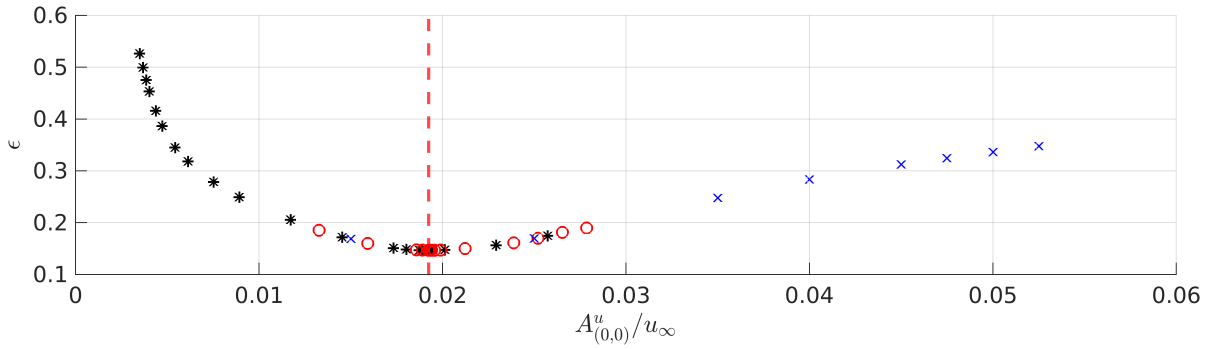

Figure 5.9: ϵ convergence plots for different initial amplitudes: (*) $A_0/u_\infty = 3.5\times10^{-3}$, (◦) $A_0/u_\infty = 0.026$ and (x) $A_0/u_\infty = 0.05$.

Linear PSE is run for the same inflow amplitude of $A_0/u_\infty = 1.9401 \times 10^{-2}$. Although linear PSE

would make for a significantly faster optimisation tool, the problem with the equations is the inability to assess its limitations efficiently. Results will always be returned, but it is hard to assess until what point in the streamwise direction the results remain accurate. In this specific case, PSE was only accurate up until 15% chord, as derived from the direct comparison with NPSE in Figure 5.10. Downstream of $x/c_x = 0.15$, PSE and NPSE solutions start to differ noticeably. Because this research focuses on the transition region, further downstream between $x/c_x = 0.25$ and $x/c_x = 0.35$ chord, PSE would not result in a suitable inflow amplitude for the steady DNS. Therefore, the decision to use NPSE for the optimisation tool was correct, and the researcher still deems it the most suitable tool for the job.

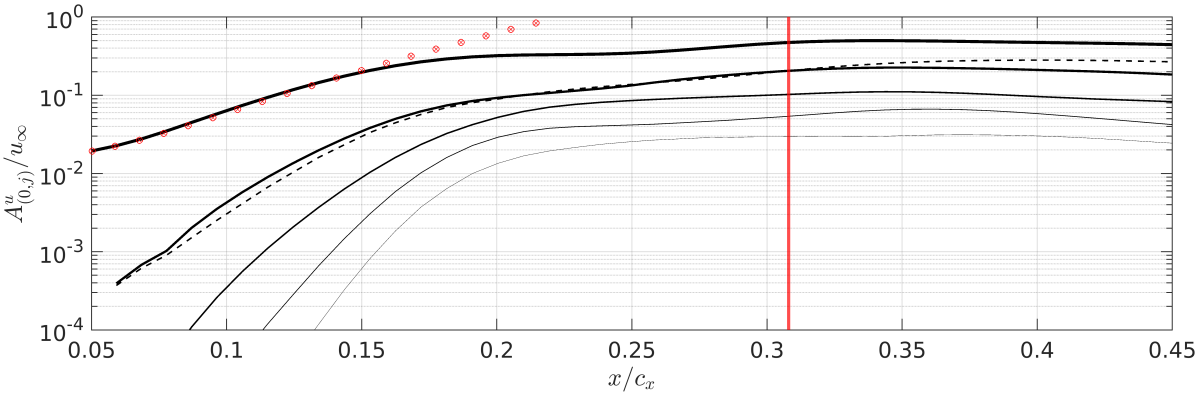

Figure 5.10: Chordwise evolution of $A^u_{(0,j)}/u_\infty$ with $A_0/u_\infty = 1.9401 \times 10^{-2}$, for linear PSE (◦) and NPSE Fourier modes j = 1-5 (-, thick-to-thin) and j = 0 (--). Vertical red line indicating experimental laminar-to-turbulent transition location.

5.2.2. DNS Results

The inflow disturbance amplitude resulting from the NPSE optimisation is used as an input for DNS. Previous work by the likes of Casacuberta, Hickel, et al. (2022) showed DNS results to be really close to its NPSE precursor. The DNS flow field that results indicates growing stationary crossflow vortices, spaced 8mm apart as visualised by Figure 5.11. At $x/c_x \approx 0.18$, the vortices reach a first saturation amplitude, after which a second amplitude increase is discovered. This double saturation amplitude growth is present in both NPSE and DNS. The double saturation curve could have a range of explanations and is thought to be related to the pressure distribution imposed onto the flat plate geometry. After the second saturation in amplitude growth at $x/c_x \approx 0.30$, the amplitude noticeably starts to decline. Most likely, this behaviour is related to the proximity of the outflow condition and the change in pressure gradient moving more aft on the airfoil geometry. At the outflow, a constant pressure boundary condition is applied, influencing the region of flow in its vicinity. There is no reason for concern, as the influence remains local and the region of interest lies far upstream of the outflow, between 20% and 35% chord. So although the amplitudes downstream of $x/c_x = 0.32$ are considered irrelevant for comparison with the experimental results, since transition should already have occurred, it is still useful to ensure the different numerical tools used throughout to generate the steady flow field agree in this downstream region. Nevertheless, the main region of interest still lies upstream of the expected transition front. Therefore it is positive that the three-dimensional boundary layer shows the presence of strong stationary crossflow vortices, which is in line with the experimental results and can be linked to the DREs used in the experimental model.

The CFI amplitude growth in DNS, and the growth of its individual spanwise Fourier modes $A^u_{(0,j)}/u_\infty$, can be directly compared with the NPSE data to ensure the NPSE-based optimiser correctly represented the DNS simulations. By splitting up the total disturbance amplitude for DNS and NPSE into modes $j = 1$-5 and the mean flow distortion ($j = 0$), these modes can be graphically compared (see Figure 5.12). The match between NPSE and DNS is very close until well past the transition region. It is only downstream of $x/c_x = 0.4$ where the solutions start to diverge very slowly. As the region of interest lies far upstream of this point, it is not expected to have any influence on the results. This effect is thought to be related to the outflow boundary condition in DNS, set to a constant static pressure with a second-order Neumann-type boundary condition imposed on the velocity (Casacuberta, Hickel, et al., 2022). This is an artificially generated boundary condition not directly based on the experimental data, therefore, not likely to be a fully accurate depiction of reality. NPSE uses spatial marching in the

x-direction, not solving the entire flow field at once. For this reason, the outflow boundary condition does not affect the NPSE results as it marches towards it. DNS on the other hand, does get affected by the outflow boundary conditions while solving for the complete flow field, making it a probable explanation for the differences observed near the outflow. Among other possible causes for discrepancies is the use of a flat plate domain, imposing the airfoil curvature of the experiment through boundary conditions instead of simulating it. The use of a flat plate to emulate curved surfaces is generally accepted in numerical studies when the curvature remains low (Casacuberta et al., 2021). Small errors based on this assumption may have been left undiscovered as this is the first direct numerical replication of this type of experiment. Although considered unlikely to introduce significant errors, this is an assumption to keep in mind for future work, where further investigation into curvature effects could be relevant. Disregarding these small differences near the outflow, DNS and NPSE agree very closely in modal disturbance amplitude. The same also holds for the MFD and disturbance profiles as well, where the difference is barely noticeable. Both the MFD and disturbance profiles at the locations of the HWA planes, are compared with NPSE in Figure 5.14 and Figure 5.15, showing remarkably good agreement. The resemblance of these curves for NPSE, DNS and HWA is further discussed at once in the next section.

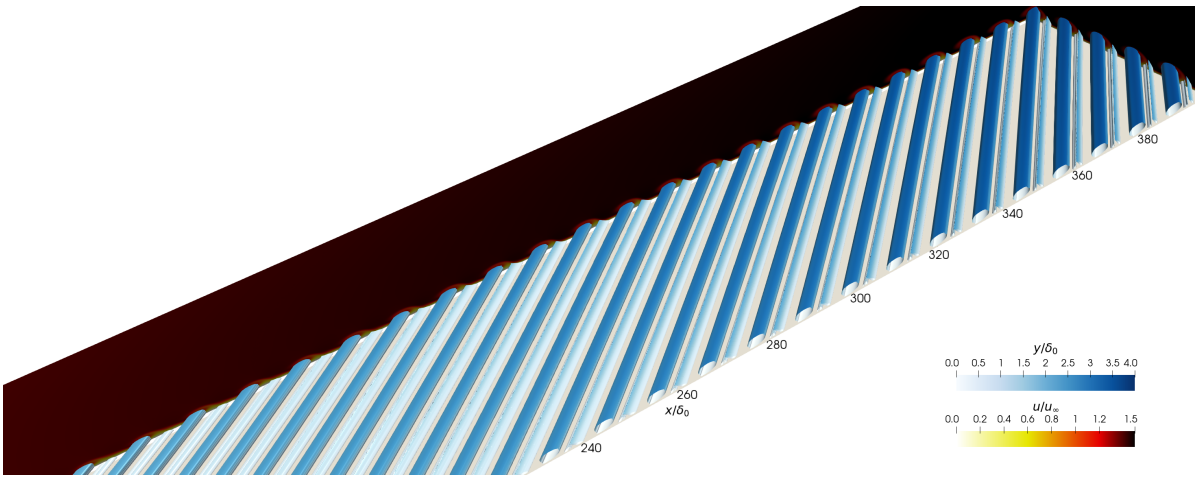

Figure 5.11: Steady DNS instantaneous Q-criterion isosurface coloured by wall-distance (y/δ_0), along with y-z and y-x u/u_∞ slices. Domain duplicated four times in z-direction based on periodic boundary conditions.

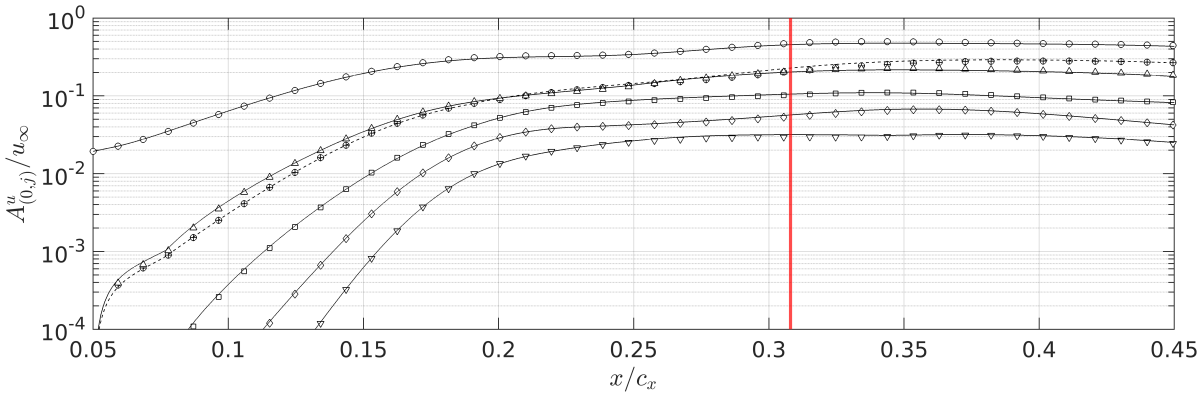

Figure 5.12: Chordwise evolution of the disturbance amplitude of inflow amplitude of $A_0/u_\infty = 1.9401 \times 10^{-2}$, for NPSE (white symbols) and DNS Fourier modes j = 1-5 (-, thick-to-thin) and j = 0 (--). Experimental laminar-to-turbulent transition location indicated by vertical red line.

5.2.3. Steady Flow Field Review

The DNS simulation and NPSE results match closely up until $x/c_x = 0.45$. Due to the fact that NPSE results are optimised to be as close as possible to the experimental data, the DNS simulation should therefore also resemble the experimental data very well. Briefly comparing the maximum disturbance

amplitude, this also appears to be the case, even though the experimental data points are limited compared to the high-resolution numerical data set. Figure 5.13 provides this three-way comparison for the disturbance amplitude ($A^q_{(0,0)}$) between NPSE, DNS and experimental data in graphical form, normalised by the local free-stream velocity $\overline{Q_e}$.

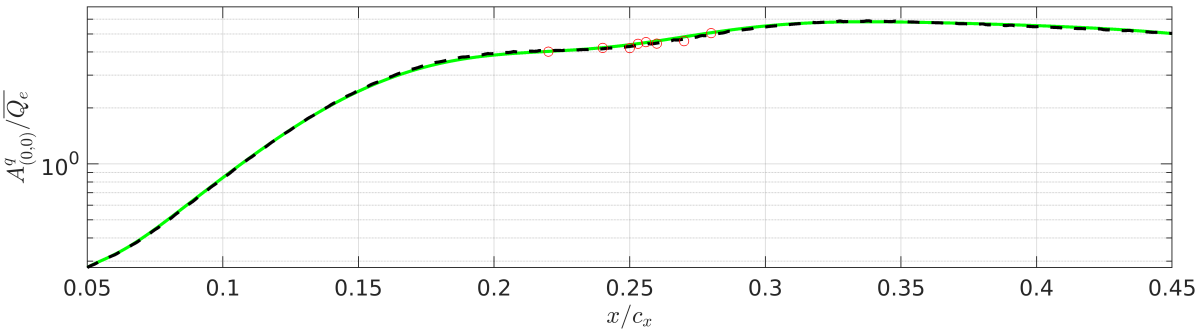

Figure 5.13: Chordwise evolution of $A^q_{(0,0)}/\overline{Q_e}$ for NPSE (--), DNS (–) and experimental data (◦).

MFD profiles give a much better indication of the method's accuracy at the HWA plane locations than merely the maximum amplitude. All eight MFD profiles are shown for comparison in Figure 5.14, and although there is variability between individual planes, the overall match with experimental data is very close. The invariability between NPSE, DNS and experimental data is in this instance also noticeable, with all three data sets producing very similar results. Even though the HWA data close to the wall is not available due to the limitations of the measurement techniques, the characteristic shape of the velocity profile that differentiates it from the baseflow boundary layer profiles is clearly visible and incorporated within the area of comparison.

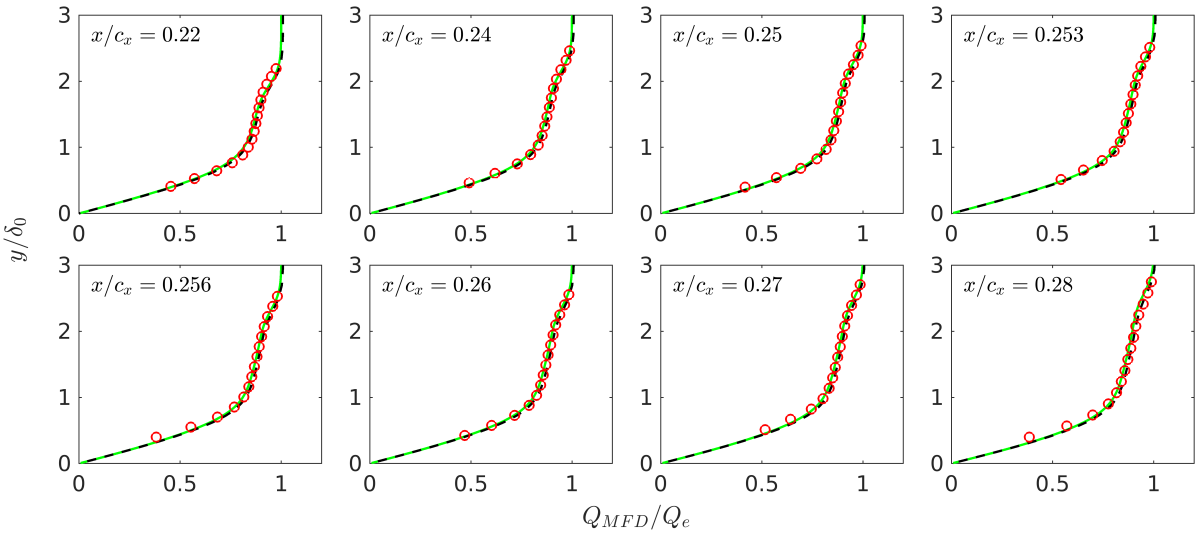

Figure 5.14: Mean Flow Distortion (MFD) profiles for all available HWA plane locations, comparing NPSE (--), DNS (–) and experimental data (◦).

The disturbance profiles show a similar trend when comparing all three data sets, with the NPSE and DNS results lying very close to each other. The direct comparisons for all available HWA planes are given by Figure 5.13. The difference between the almost identical NPSE and DNS profiles when compared to the experimental profiles is still noticeable, especially around the upper sections of the boundary layer. Depending on the chordwise location, the upper peak of the different profiles is misaligned differently. The error does remain within a couple of percent difference between the local peak amplitudes, and the three different curves all remain close to each other when considering all eight planes available. In view of the overall match across all planes, the numerical simulation is considered an acceptable replication of the experiment at this stage. This numerical twin of the experiment is

by no means perfect but remains a significant improvement compared to previous work by the likes of Nishino and Shariff (2010) and Tempelmann et al. (2011). The individual differences in the results when analysing different HWA planes can potentially be linked to a multitude of factors, but can with certainty partially be attributed to the variation in experimental measurements. As outlined in section 2.2, the time-intensive process of acquiring HWA data led to the planes being measured on different days and under slightly different environmental conditions, leading to a variation in data which is impossible to simulate in one numerical equivalent. Additionally, the wind tunnel may contain flow features that the numerical simulation does not take into account.

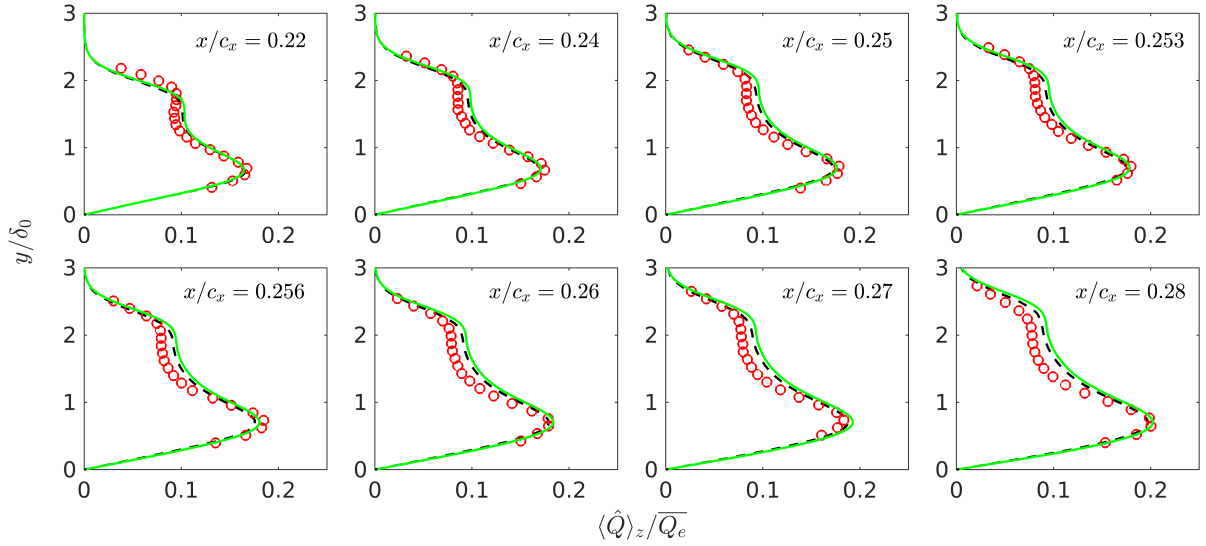

Figure 5.15: Disturbance profiles for all available hotwire plane locations, comparing NPSE (--), DNS (–) and experimental data (◦).

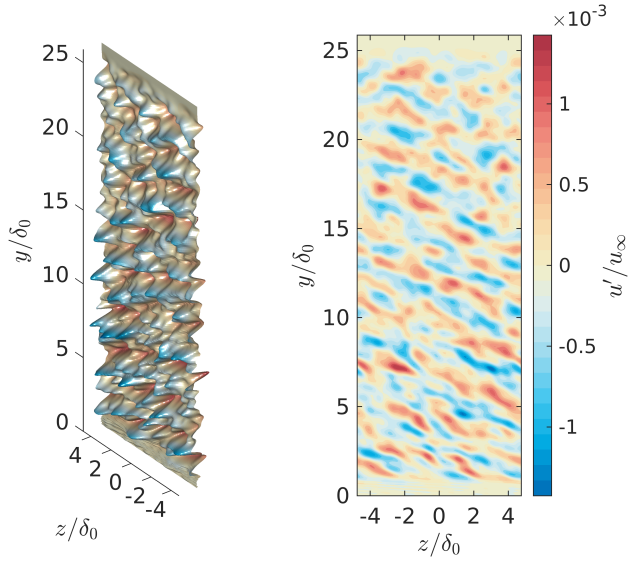

Figure 5.16: u'/u_∞ at the domain inflow boundary, resembling free-stream turbulence as a 3D (left) and 2D (right) z-y plane at $x/c_x = 0.05$.

5.3. Secondary Crossflow Instability Development and Transition

Secondary crossflow instabilities develop inside boundary layers modulated by primary CFI (Malik et al., 1994; Malik et al., 1999). These secondary CFI are thought to be the leading cause of flow breakdown into turbulence. The development of secondary CFI is accelerated by the introduction of free-stream turbulence (FST) through receptivity, a process this work attempts to replicate numerically. The

implementation of FST in the numerical domain is based on a combination of recent methodology improvements described by Schlatter (2000), Brandt et al. (2004) and Balzer and Fasel (2016). These improvements in numerical FST generation allow for a structured approach going forward, increasing confidence and repeatability for similar work in the future. FST is reconstructed by creating and summing a large number of pseudo-random disturbance modes. Using the approach discussed in section 4.4, a large number of disturbance modes are generated and summed. The two-dimensional disturbance field that results can be visualised as a z-y plane (see Figure 5.16), showing the FST inflow boundary condition in both two and three dimensions, for u'/u_∞. Designed to be added on top of the steady flow field at the inflow boundary, this plane exhibits random disturbances across the majority of its area, resembling the FST levels in the wind tunnel environment. Two exemptions include the areas near the wall and the top boundaries. The former is related to the shear sheltering, described in subsection 4.4.2. The zone near the top boundary is artificially damped through exponential decay to avoid unwanted interaction between the FST and the domain top boundary causing numerical stability problems.

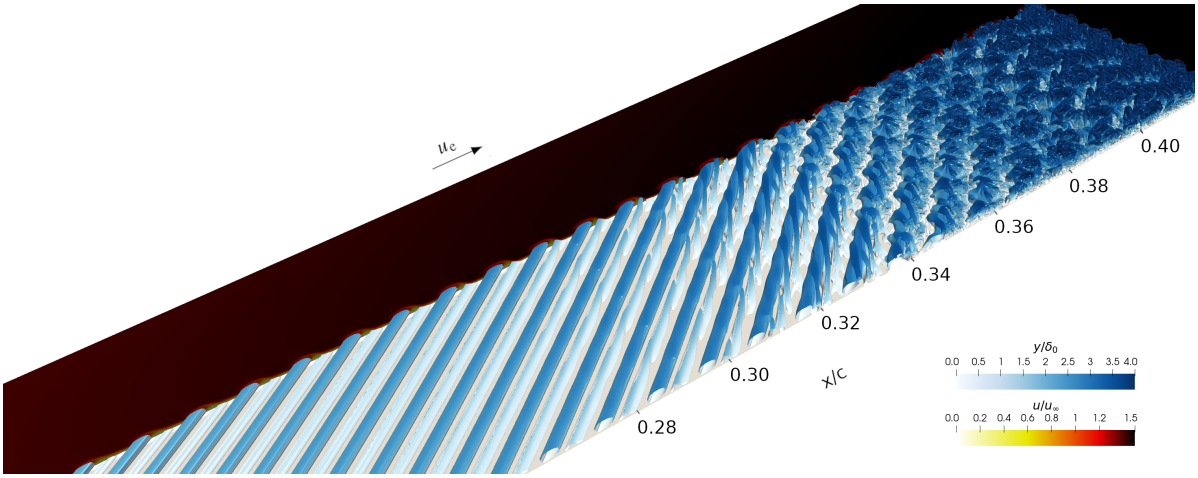

Figure 5.17: Instantaneous snapshot of the unsteady DNS Q-criterion isosurface coloured by wall-distance (y/δ_0), along with z-y and x-y u/u_∞ slices. The periodic domain is duplicated four times in z-direction.

The many additional inflow modes inevitably increase the complexity of the flow field. This complexity of FST implies it is more computationally expensive compared to more discrete forcing methods (e.g. Blowing & Suction), using fewer disturbance modes. Nevertheless, it is able to obtain a representative flow field and laminar-to-turbulent transition process without the need for iteration, leading to a reduction in total computational cost across the project. This method also repeatedly produces an unsteady flow field that captures crossflow-driven transition and allows the investigation of the transition location, type, and frequency content. Along with the potential for further extensive analysis. Figure 5.17 shows a section of the Q-criterion isosurface ($Q = 10^{-4}$), within which the majority of the laminar-to-turbulent transition process can be identified. Although this type of numerical flow field can potentially be achieved with previously established methods, this particular result is obtained with just a single simulation. All while maintaining excellent correlation with experimental data, as this section will point out. There is no need for iterating the boundary conditions and input parameters in a trial-and-error-based approach. With the wind tunnel parameters known, the correct settings can be analytically determined and implemented. The Q-criterion isosurface of the unsteady simulation, presented by Figure 5.17, shows a rapid transition from a laminar to a turbulent state. Visually indicating a type-I-driven process, this already indicates qualitatively the type of transition is similar to the experimental equivalent (Rius-Vidales & Kotsonis, 2022). In the following sections, the transition location, transition type, and frequency content will be discussed further. This along with a discussion on the FST implementation and its impact on the entire numerical domain, as well as a comparison with numerical simulations without FST implemented.

5.3.1. Laminar-Turbulent Transition Location

During the experiment, the transition location on the wind tunnel model is extracted through infrared (IR) thermography (Rius-Vidales & Kotsonis, 2022). Although not a possibility in the numerical simulation, the skin friction forces at the airfoil surface can be used as an alternative measure to IR thermography. This can be explained through the Reynolds analogy, directly linking the skin friction coefficient to the amount of heat transfer that occurs at the wall (Anderson, 2016). Therefore, higher skin friction generates more heat transfer resulting in a hotter airfoil surface, indicating a similar temperature distribution pattern to IR thermography at the airfoil surface (see Figure 5.19). Although differences in absolute magnitude are to be expected, for which an error margin needs to be kept in mind when directly comparing the spanwise-averaged transition location, the skin friction is considered a good qualitative measurement to compare with the IR thermography-based transition location. In DNS, the first cell layer at the wall is used to extract this skin friction across the two-dimensional boundary. Due to the fact that turbulence is inherently unsteady, the shear from the numerical simulation as a function of x/c_x is both time-averaged and spanwise-averaged over 0.4 flow-through cycles. This sampling length equals 6.5 periods of 1kHz or roughly 0.025 seconds per data point. This averaging window is deemed sufficient to compare with the experimental data, knowing the bulk of the dominant frequencies are located between 2kHz and 7kHz. The skin friction coefficient values that result from the statistical sampling are represented by the curve in Figure 5.18, showing the spanwise-averaged C_f along x/c_x. The two-dimensional shear across the wall suggests again that the one-dimensional representation conceals information, given that the shear along the span varies for the same x/c_x positions (see Figure 5.19).

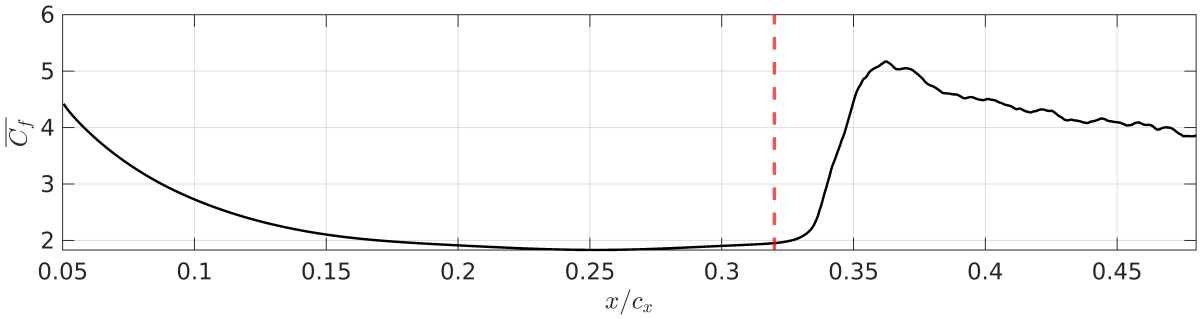

Figure 5.18: Time- and spanwise-averaged C_f as a function of x/c_x along the wall boundary in the DNS domain with FST, indicating the experimental laminar-to-turbulent transition location (--) (Rius-Vidales & Kotsonis, 2022).

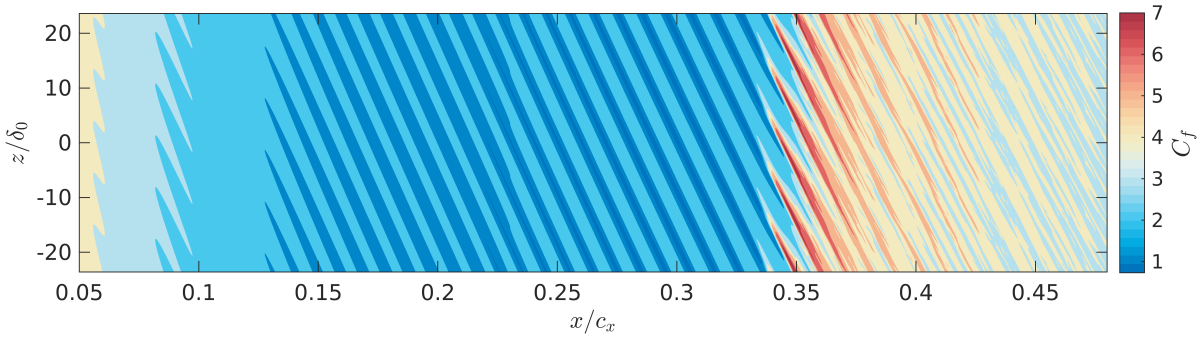

Figure 5.19: Time-averaged C_f in the x-z plane along the wall boundary in the DNS domain with FST.

Figure 5.18 highlights that the friction coefficient in the turbulent region is threefold the magnitude of the laminar equivalent just before transition. As the skin friction does not instantly switch from this low-shear laminar state to a high-shear turbulent state, there is a transition region with high gradients in between. The onset of transition, where the rapid increase of skin friction starts, is found at $x/c_x = 0.33$. A maximum skin friction coefficient is reached downstream of $x/c_x = 0.35$. To determine the exact transition location, the location with the highest C_f-gradient in the transition region is chosen. This point ($x/c_x \approx 0.34$) is deemed most representative for comparison with qualitative methods such as the IR

thermography used here since it represents the centre of the boundary between two clearly-defined C_f regions, namely the laminar and turbulent boundary layer states. The results of the IR thermography in Rius-Vidales and Kotsonis (2022) indicate transition from laminar to turbulent flow at $x/c_x \approx 0.32$. The red dashed line indicates this location. Both transition locations are within $x/c_x = 0.02$ of each other, with the numerical transition occurring downstream of the experimental transition. Although the numerical transition front location is considered an accurate representation of the experiment by the author, possible reasons for this discrepancy will be brought up during the remainder of this work, based on a more detailed analysis of the CFI development and transition. To put the accuracy of the numerical simulation into perspective, the time-averaged C_f at the wall for both the numerical and experimental flows is also compared qualitatively. Figure 5.20 shows the IR thermography of the wind tunnel model on the left and the time-averaged numerical C_f at the wall on the right. These two-dimensional planes confirm the average transition location differs by $x/c_x \approx 0.02$, and both show a type of saw-tooth pattern. However, this saw-tooth pattern appears to generate a significantly higher transition location variability in the experimental IR thermography. Variations of $x/c_x = 0.05$ are no exception, whereas the numerical field variations remain around $x/c_x = 0.01$. This variation shows experimental transition occurs as early as $x/c_x = 0.29$, but also partially at $x/c_x = 0.34$, matching the numerical transition location. This qualitative analysis indicates this numerical simulation is representative of the experimental flow field, but most probably lacks some of the CFI characteristics inherent to the wind tunnel or free-flight conditions which are causing larger transition location variability.

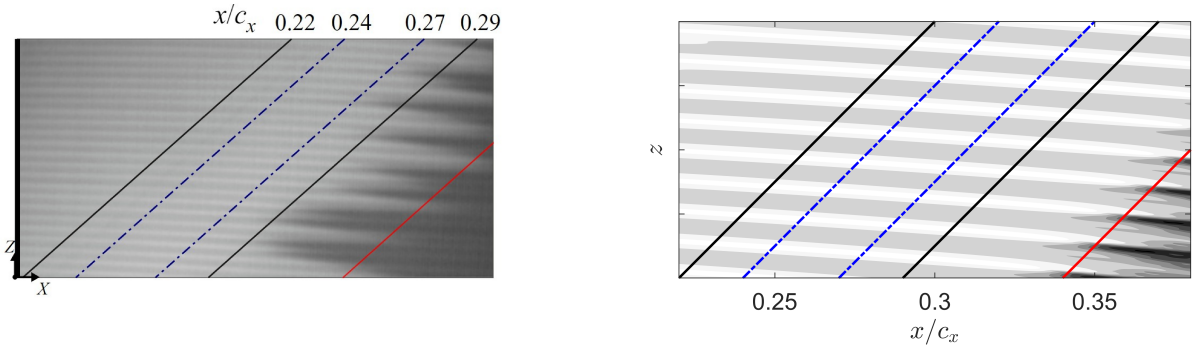

Figure 5.20: Experimental IR thermography of the wind tunnel model (left) and numerical wall shear (right), indicating the numerical laminar-to-turbulent transition location (--) on both plots. IR thermography image reworked from Rius-Vidales and Kotsonis (2022)

5.3.2. Secondary Crossflow Instability Type

As previously discussed in subsection 1.2.1, laminar-turbulent transition is triggered by secondary CFI growth. Depending on the location within the boundary layer this secondary CFI manifests most dominantly, this is referred to as Type-I, Type-II or Type-III secondary CFI (White & Saric, 2005). Rius-Vidales and Kotsonis (2022) shows that in agreement with Serpieri and Kotsonis (2016), the secondary CFI mode of Type-I drives the transition behaviour in the experiment considered in this work, showing both a dominant Type-I secondary CFI, next to a noticeably present Type-II secondary CFI. It is a well-documented transition process and is a crucial part of assessing the success of replication. A detailed analysis of this secondary CFI is, therefore, desired to explain the small difference in transition location between numerical and experimental. The DNS results with FST applied appear to show qualitative confirmation of both Type-I and Type-II secondary instability behaviour (see Figure 5.21 and Figure 5.22). From the isosurfaces, Type-I secondary CFI is the most clearly recognisable, showing vortex elongation in the upwelling region of the vortex creating diagonal streaks in both Figure 5.21 and Figure 5.22. Downstream, the rapid increase in secondary CFI in this region appears to break down the elongated structures rapidly from around $x/c_x = 0.34$. During this process, the vortex elongation in the Type-I region appears to influence the Type-II region more than seen in similar works using the blowing and suction methodology instead of FST (e.g. Casacuberta, Groot, et al. (2022)). The wall-distance colouring of the Q-criterion isosurface in Figure 5.21 accentuates this potential involvement of the Type-II secondary CFI by colouring this region of instability darker blue, where close to the turbulent region variations in colour start to appear. Both secondary instabilities of Type-I and Type-II have been related to the development of the Kevin-Helmholtz type instability (Bonfigli & Kloker,

2007). The Type-III secondary CFI is attributed to the interaction between stationary and travelling crossflow vortices and is significantly harder to detect on the isosurfaces presented here (Rius-Vidales & Kotsonis, 2022). Although further analysis is required to confirm if the Type-III instability is involved in the transition process discussed here, the lack of Type-III instabilities in the numerical flow field is to be expected due to the deliberate forcing of stationary disturbance modes at the numerical inflow boundary. This, as well as the dominance of the Type-I and Type-II secondary CFI, can be confirmed more definitively by analysing the velocity fluctuations in all three regions in the next few paragraphs. Looking further downstream, from $x/c_x = 0.36$ onward, the breakdown of the flow structures does not lead to the clearly distinguishable hairpin vortex structures described by Casacuberta, Groot, et al. (2022). Instead, the turbulent boundary layer Q-criterion isosurface in Figure 5.21 shows much less coherence. The second Q-criterion isosurface coloured by chordwise velocity (Figure 5.22) is also included to allow for a different viewpoint on the transition process. Although the same conclusions can be drawn from this second isosurface, the velocity differences around the different regions of instability become more recognisable, as well as the chaotic-looking turbulent structures downstream. The blending of the individual crossfloww vortices into a single turbulent motion is also shown more clearly through the diminishing velocity differences when moving downstream. As additional information, the same isosurface coloured by v/u_∞ and w/u_∞ can be found in Appendix B as well.

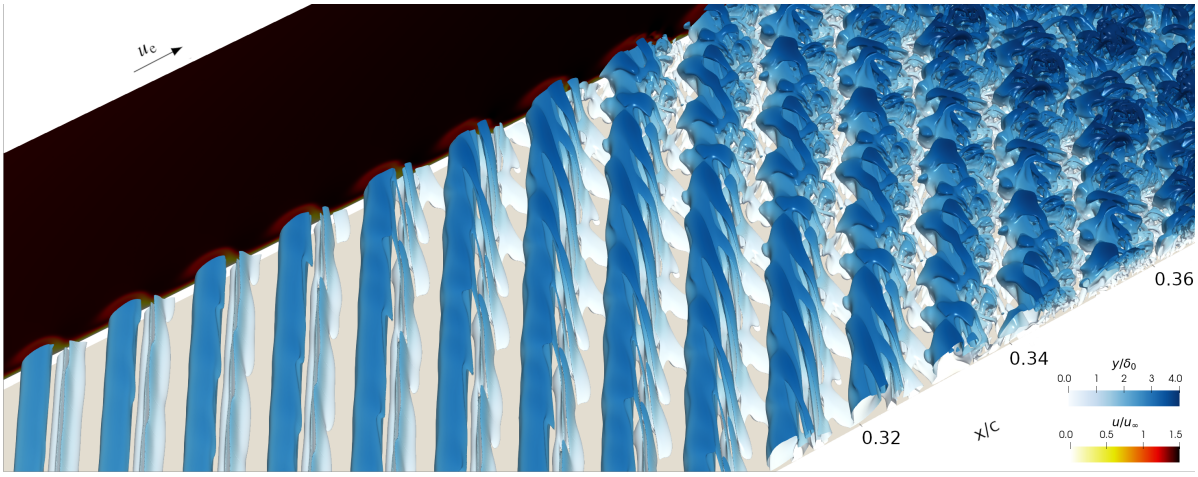

Figure 5.21: Detailed snapshot of the unsteady DNS instantaneous Q-criterion isosurface coloured by y/δ_0, along with y-z and y-x u/u_∞ slices. Domain duplicated four times in z-direction based on periodic boundary conditions.

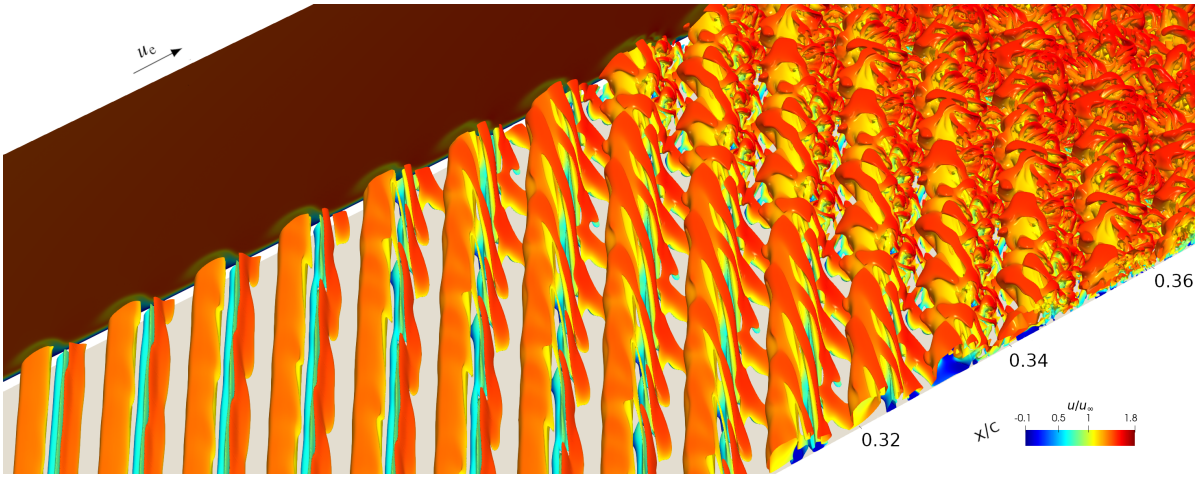

Figure 5.22: Detailed snapshot of the unsteady DNS instantaneous Q-criterion isosurface coloured by u/u_∞, along with y-z and y-x u/u_∞ slices. Domain duplicated four times in z-direction based on periodic boundary conditions.

When comparing back to the transition behaviour seen by Casacuberta, Groot, et al. (2022), the

Type-II region in both Figure 5.21 and Figure 5.22 shows noticeably more modulation with FST. This increased modulation in both Type-I and Type-II regions might, in part, be caused by FST forcing a broadband set of frequencies. Whereas the Blowing & Suction used in Casacuberta, Groot, et al. (2022) and described in Appendix A forces discrete frequencies. This difference is then thought to also be the cause of the less coherent turbulent structures downstream of transition in the FST-based simulation. Qualitative analysis of the transition type already established some similarities between numerical and experimental data through the isosurfaces. As no definitive conclusions can be drawn from just isosurfaces of instantaneous flow variables, a quantitative analysis based on the standard deviation of the temporal velocity fluctuations is performed as well (see Equation 4.16). By use of the numerical instantaneous velocity data (Q) sampled at a frequency of $1MHz$, the velocity fluctuations (Q') measured by the HWA are numerically re-created. The numerical measurement points used throughout this analysis form two-dimensional z-y arrays at a range of x-locations. The point spacing within the arrays replicates the HWA plane resolution in y- and z-direction. This approach makes it possible to extract the numerical equivalents of the HWA planes. Through the standard deviation of the instantaneous velocity measurements at each point, a spatial representation of the velocity fluctuation magnitude and distribution is obtained in the form of z-y planes (see Figure 5.23).

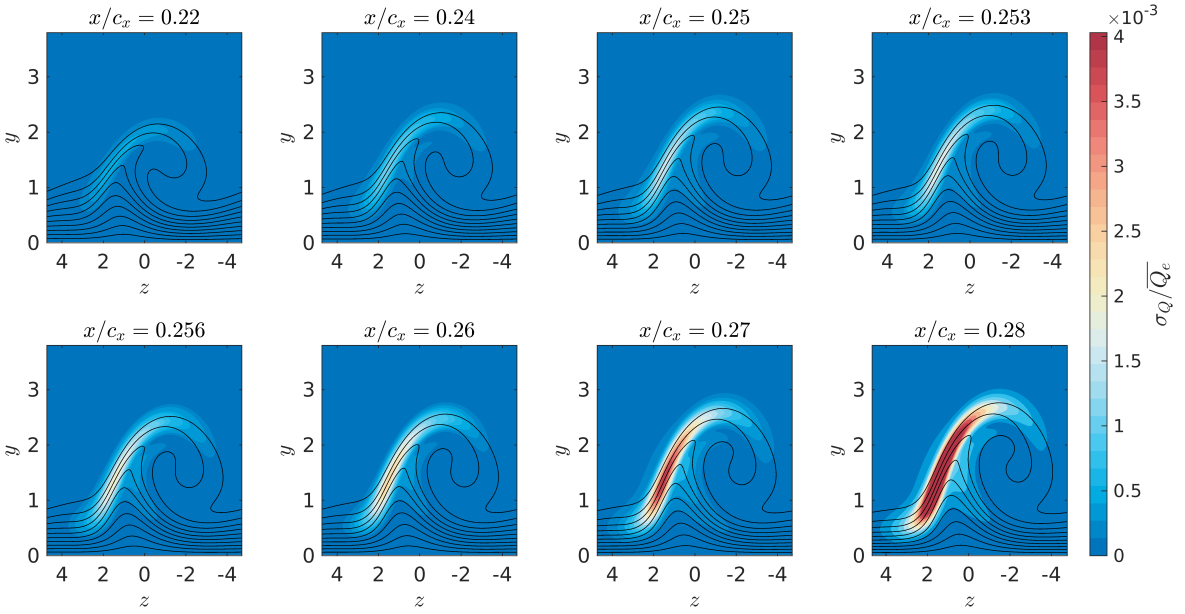

Figure 5.23: Normalised standard deviation of the temporal velocity fluctuations ($\sigma_Q/\overline{Q_e}$, coloured contour plots) and matching time-averaged velocity contour lines (Q/Q_e, solid lines).

Based on the full numerical data set (i.e. without frequency filtering), the eight planes in Figure 5.23 indicate the bulk of fluctuations is present in the upwelling region of the crossflow vortex, in the area associated with Type-I secondary CFI. The evolution of these velocity fluctuations across the different planes points towards a rapidly increasing growth rate in the downstream direction. The dominant Type-I fluctuations also gradually start to expand towards the wall, and upwards towards the cusp of the crossflow vortex, influencing the Type-II region. This confirms the qualitative analysis based on the isosurfaces performed before that the transition behaviour is dominated by Type-I secondary CFI. It does raise questions about the Type-II modulations observed in the isosurfaces of Figure 5.21 and Figure 5.22, and their origin. While the isosurfaces hinted towards a separate Type-II secondary CFI, it now becomes clear that Type-II itself is weak compared to the dominant Type-I secondary CFI in the numerical flow field. Therefore, the modulations visible in the isosurface figures are thought to be related more to the expansion of the Type-I towards the cusp of the vortex than a separate Type-II, as visible in the downstream planes of Figure 5.23. With the numerical transition occurring further downstream than the experimental data, additional planes downstream of the available HWA data are extracted as well (see Figure 5.25). These additional z-y planes show rapid growth in unsteadiness close to the transition location at $x/c_x \approx 0.34$. Ultimately, the largest unsteadiness leading to transition appears to start near the wall and spreads upward towards the cusp of the vortex. However, the bulk

of the unsteadiness remains centred around the Type-I region, in line with the experimental results.

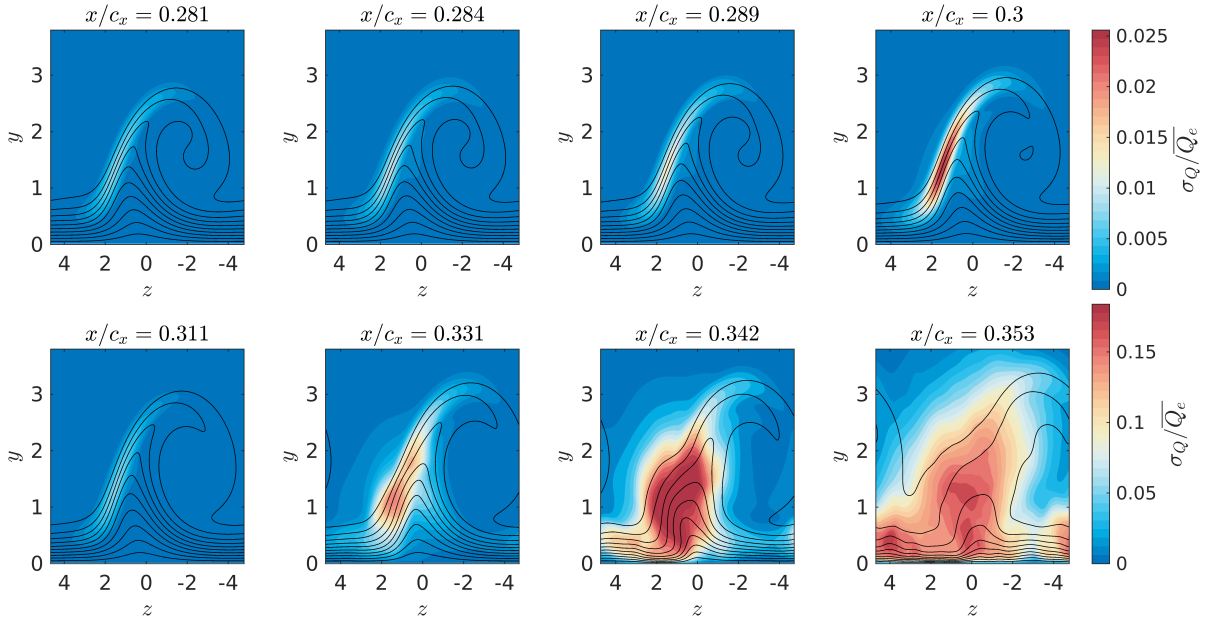

Figure 5.24: Normalised standard deviation of the temporal velocity fluctuations ($\sigma_Q/\overline{Q_e}$, coloured contour plots) and matching time-averaged velocity contour lines (Q/Q_e, solid lines).

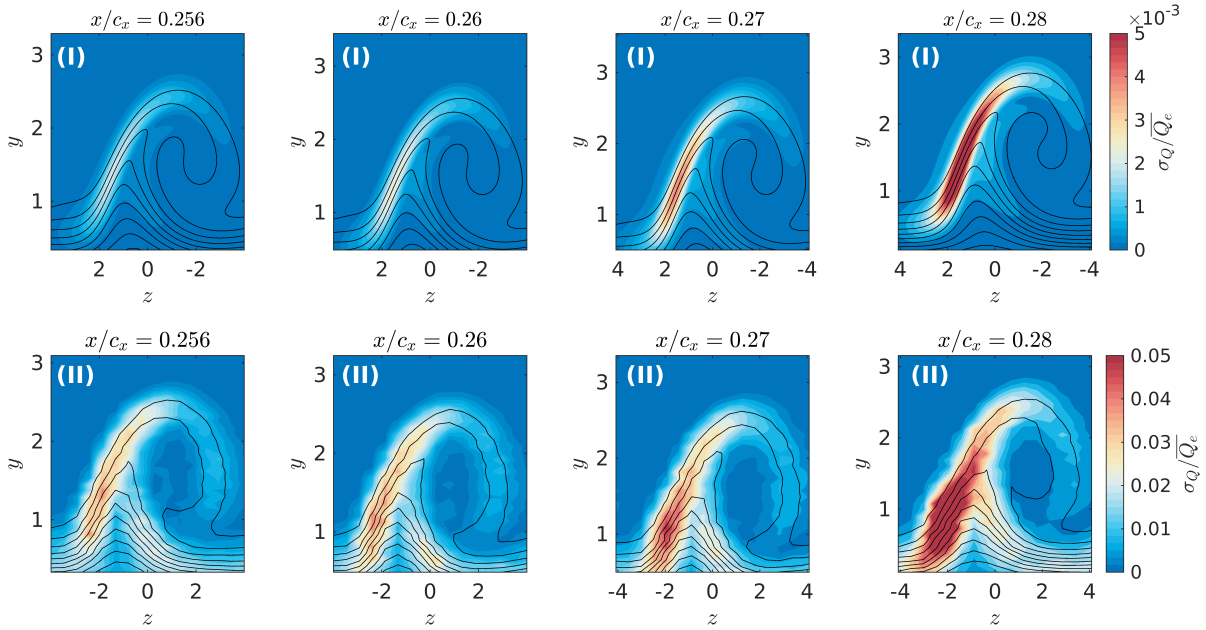

Figure 5.25: Normalised standard deviation of the temporal velocity fluctuations for both DNS (I) and HWA (II) data sets ($\sigma_Q/\overline{Q_e}$, coloured contour plots) and matching time-averaged velocity contour lines (Q/Q_e, solid lines).

Because the exact HWA planes can be extracted from the numerical flow field, a direct comparison to confirm this analysis is also made possible. By processing the available Q' data for both data sets in the same manner, a side-by-side comparison is generated in Figure 5.25. The top row in Figure 5.25 indicates the numerical instability evolution for four planes at different x/c_x, while the bottom row represents the equivalent processed HWA planes. Both the top and bottom rows show a similar story, with clear Type-I dominant secondary CFI. The numerical secondary CFI remains centred closer around this Type-I location and is largely limited to the upwelling region of the crossflow vortex, although downstream these secondary CFI start to show signs of spreading towards the Type-II region as well.

The HWA data, on the other hand, shows much more diffuse secondary CFI. While also centred around the Type-I region, the temporal velocity fluctuations appear to spread out farther into both Type-II and Type-III regions. Nevertheless, the largest magnitudes of $\sigma_Q/\overline{Q_e}$ are achieved at roughly the same locations as the numerical equivalent, albeit a little closer to the wall within the upwelling region of the vortex. These findings indicate that the numerical simulation follows a similar transition pattern as the one observed in the experimental study performed by Rius-Vidales and Kotsonis (2022), where the dominant temporal velocity fluctuations occur in the Type-I region.

Compared to the HWA data which only provides velocity magnitude ($Q = \sqrt{(ucos\Lambda + wsin\Lambda)^2 + v^2}$), the numerical simulation provides the individual u-, v- and w-components. This benefit of numerical work can provide a wealth of additional information on the secondary CFI components in different directions such as the standard deviation of the velocity fluctuations of the individual components σ_u, σ_v, and σ_w (see Figure 5.26). σ_Q is added to the right of Figure 5.26 to compare with the velocity magnitude measured by HWA. The most noticeable is the dominance σ_u and σ_w, compared to the v-component. The ratio between σ_u and σ_w remains roughly equal to the ratio of their mean velocities, which results in equal orders of magnitude in Figure 5.26 as σ_u and σ_w are normalised by $\overline{U_e}$ and $\overline{W_e}$, respectively. At the most upstream planes in the data set, σ_v remain small, several orders of magnitude lower than σ_u and σ_w. Because of the normalisation with V_e, which is several orders of magnitude smaller than U_e and W_e as well, this difference remains somewhat masked in the figures used. From Figure 5.27, it appears the increase in the magnitude of σ_v relative to the other components is larger when moving downstream. Although only becoming manifesting larger values further downstream than the σ_u and σ_w components, this indicates σ_v is nevertheless a contributor in the transition process. The different frequency components of these temporal velocity fluctuations, which have currently only been examined in terms of absolute magnitude, are the topic of the last CFI analysis between HWA and numerical data. This is a necessary analysis in order to confirm the Type-I dominant secondary CFI displays the same characteristics in both the numerical simulation and the experiment. Even though the secondary CFI type appears similar between both data sets, the frequency analysis might lead to an explanation for the slight difference in laminar-turbulent transition location.

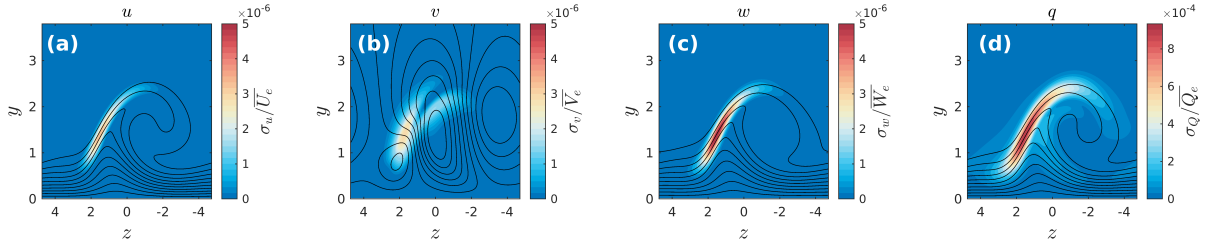

Figure 5.26: Normalised standard deviation of the temporal velocity fluctuations in DNS at $x/c_x = 0.253$, for u-, v- and w-components (a-c) and the velocity magnitude Q (d). Time-averaged velocity contours are included as solid lines.

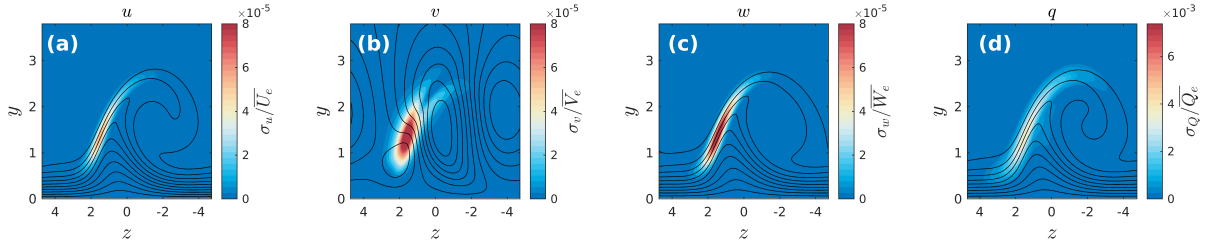

Figure 5.27: Normalised standard deviation of the temporal velocity fluctuations in DNS at $x/c_x = 0.28$, for u-, v- and w-components (a-c) and the velocity magnitude Q (d). Time-averaged velocity contours are included as solid lines.

5.3.3. Frequency Content

To attempt to find an explanation for the slight delay in transition to turbulence, the similarity between numerical and experimental data in terms of frequency content is evaluated through a Power Spectral Density (PSD) analysis. This is a method to expose any discrepancies in the numerical flow field the standard deviation approach (Equation 4.16) was not able to detect. Each type of secondary CFI has a

characteristic frequency range, allowing an assessment of the flow field transition mechanics compared to the experimental data, as well as historical work on this topic. The Type-I-driven laminar-turbulent transition has been observed during the experiment of Rius-Vidales and Kotsonis (2022). Both the numerical isosurfaces, as well as the temporal velocity fluctuations, have highlighted the numerical flow field produces a similar-looking transition behaviour. The unfiltered standard deviation data σ_q, confirms Type-I is the dominant secondary CFI in the numerical flow field, and shows its spatial dependency. Analysing the DNS measurement points located in the Type-I, Type-II and Type-III regions now allows for a qualitative comparison of the frequency content between the experimental and numerical flow fields, which is a key step in the identification of the Type-I secondary CFI dominance. At a single x/c_x-location, the PSD analysis related to each secondary CFI type is extracted for direct comparison (see Figure 5.28). A PSD plot for a point in the free-stream flow is added for comparison as well. All four measurement points are indicated on the average velocity contour on the left of Figure 5.28. The plane depicted here is just downstream of the $x/c_x = 0.28$ HWA plane. This is the most downstream HWA plane, which is deemed most relevant for the transition process as it is closest to the transition location. The conclusions drawn from Figure 5.28 are two-fold. The broadband lower frequency component of the fluctuations, between $1kHz$ and $3kHz$, is clearly present in both data sets. This agrees well with the observations reported in Rius-Vidales and Kotsonis (2022), where a low-frequency peak centred around $2kHz$ was also distinguishable. Within this band, the Type-I CFI is found most dominant, confirming the transition type discussed before. Type-II frequency amplitudes are lower magnitude than Type-I, but still much more pronounced than Type-III instabilities. From this single plane can already be observed that the numerical simulation is lacking the high-frequency components of q'. (Rius-Vidales & Kotsonis, 2022) points out a high-frequency component centred around $5kHz$ not distinguishable in the numerical results. While already present upstream of $x/c_x = 0.253$ in experimental data, the high-frequency component does not start to show up until very close to the transition region, downstream of the last available HWA plane at $x/c_x = 0.28$ (Figure 5.28). The apparent sudden damping of these lacking high-frequency components in the numerical results is thought to be related to discrepancies in the numerical setup, either through the inflow boundary conditions or the dissipation of high frequencies throughout the domain. If these are valid suspicions, and if the numerical setup is to blame for these discrepancies, will be discussed in the next section of this chapter.

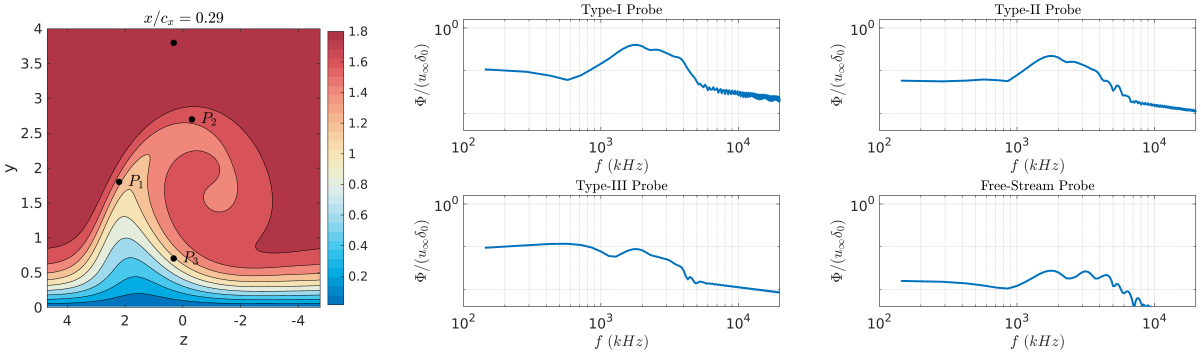

Figure 5.28: Power Spectral Density (PSD) analysis of the unfiltered temporal velocity data at $x/c_x = 0.29$ using probes in the Type-I (P1), Type-II (P2), Type-III (P3) and free-stream flow regions.

Although it makes sense that the amplification of higher frequencies happens further downstream, given transition occurs further downstream as well, frequencies above $4kHz$ only start to become significant very close to transition. When performing a spectral analysis for the Type-I secondary CFI point at a range of x/c_x planes (Figure 5.29), this trend in frequency growth can be assessed more clearly. The dominant broadband frequency of 1 to 3 kHz is highlighted even further, monotonically growing in magnitude when moving downstream. Although the higher frequencies also gain in magnitude, the amplitude of all frequencies above $3kHz$ rises without a specific band visible. The most explicit feature remains the prominent low-frequency band, which stays dominant until transition occurs. Whereas no explicit signs of the experimental high-frequency band around $5kHz$ reported in Rius-Vidales and Kotsonis (2022) appear across all x-locations, it is relevant to point out that broadband high-frequency content is still rapidly gaining amplitude. A higher-frequency component starts to develop around $4kHz$ very close to transition, pointing towards its relevance for the transition behaviour.

Nevertheless, considering the entire region upstream of transition, the high-frequency behaviour seen during the experiment is not reflected in the numerical data.

The Type-II spectral analysis, performed in an identical way to the Type-I analysis, confirms both areas of the flow are strongly linked in their transition behaviour (Figure 5.30). Although Type-I remains dominant by several orders of magnitude, Type-II indicates the same dominant frequency band of $1kHz$ to $3kHz$ and a growth scenario centred around the $2kHz$ most amplified frequency. The larger amplitude jumps between the most downstream planes do appear to indicate a significant increase in growth rate close to transition. The higher frequencies also catch up with the lower frequencies, indicating a noteworthy difference in growth rates. The Type-III frequency spectrum, on the other hand, remains much flatter for longer. Plotted in an identical way to the Type-I and Type-II instabilities, Figure 5.31 again indicates maximum spectral amplitude around $2kHz$, although slightly biased towards lower frequencies. A secondary peak around $4kHz$ starts to form slowly, but the amplitude remains inferior to the lower frequency band. Next to the missing high-frequency band centred around $5kHz$, another observation for all three instability types is the apparent lack of variation in high-frequency content at all, upstream of $x/c_x = 0.33$. At least up until very close to the transition region, where high-frequency oscillations are amplified rapidly. This behaviour is somewhat different from the experimental data, where high-frequency content is already detectable in the HWA measurements at $x/c_x = 0.22$ (Rius-Vidales & Kotsonis, 2022). Even though FST at the inflow boundary produces frequencies between $26Hz$ and $26kHz$, it appears from this spectral analysis that the receptivity process favours the lower frequency bands. Higher frequencies do show rapid increases in growth rates very close to transition, however, indicating they are still introduced into the flow. Their initial amplitude upstream could be a potential reason for this behaviour, assuming their initial amplitude is significantly lower than lower frequencies.

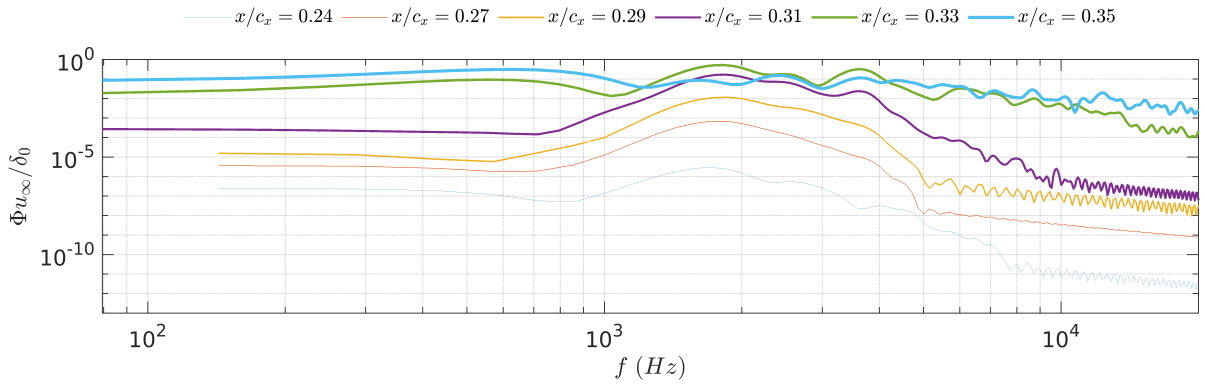

Figure 5.29: Spectral analysis of Type-I probes at different chordwise locations in the outer upwelling region P_1, ranging from $x/c = 0.24$ (thinnest line) to $x/c = 0.35$ (thickest line).

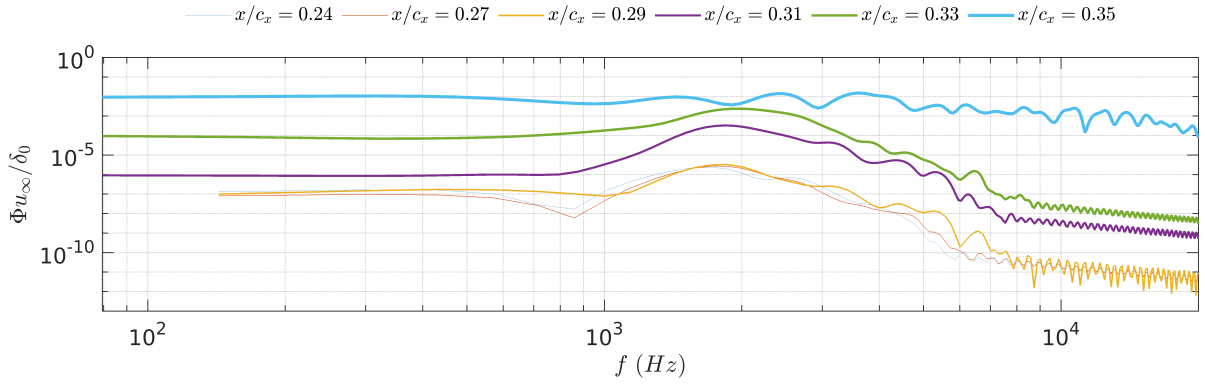

Figure 5.30: Spectral analysis of Type-II probes at different chordwise locations in the upper region of the crossflow vortex P_2, ranging from $x/c = 0.24$ (thinnest line) to $x/c = 0.35$ (thickest line).

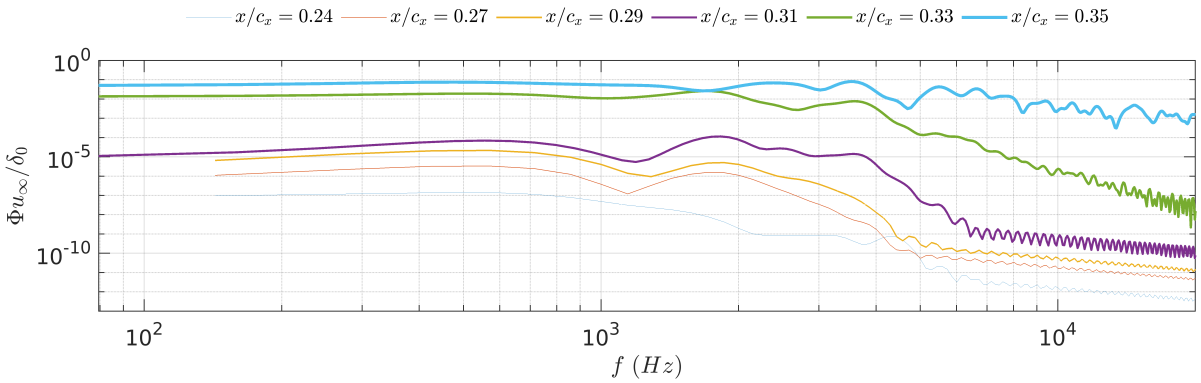

Figure 5.31: Spectral analysis of Type-III probes at different chordwise locations in the inner upwelling region P_3, ranging from $x/c = 0.24$ (thinnest line) to $x/c = 0.35$ (thickest line).

This late exponential growth of high-frequency instabilities becomes even more visible when dividing the signal per kHz, and analysing the overall picture for each x/c_x plane used. By calculating the PSD for the temporal velocity fluctuations at each measurement point and dividing the PSD amplitudes into $1kHz$ bins, centred around each integer value $(+/-500Hz)$, a simplified analysis can be performed. Integrating the amplitudes for each frequency bin across the entire plane at each x/c_x location, results in a single amplitude value for each kHz bin at each plane location. Figure 5.32 shows the growth of this integral amplitude for the frequency bins from $1kHz$ to $7kHz$. From this comparison, $2kHz$ is again found to be dominant across all planes, where the high frequencies start to catch up moving downstream. When converting these amplitudes to non-dimensional values by dividing each curve by their amplitude at $x/c_x = 0.24$, shows the exponential growth of the high frequencies even more clearly. Although the growth of these higher frequencies seems much more gradual in the experimental data, it is positive that these frequencies eventually appear in the numerical flow field as well, albeit very late. Certainly when considering they are deemed responsible for the Type-I transition process in crossflow-driven transition on swept wing flows. This difference in frequency content compared to experimental HWA data does therefore not disprove these results. The neutral points for both high and low frequencies as presented by Casacuberta, Groot, et al. (2022) could still be respected as well, since also here, the amplification of the higher frequencies appears to start further downstream compared to the low frequencies. This thought process banks on the initial amplitudes of the high frequencies being significantly lower than the lower frequencies. When no longer integrating across the entire plane, but looking specifically at the maximum amplitudes in the Type-I region, this hypothesis appears to get confirmed. Figure 5.34 indicates the higher frequency components start at amplitudes several orders of magnitude below the lower frequencies. Past $x/c_x = 0.30$, the higher frequencies then start to increase in growth rate, almost catching up with the lower frequencies before transition occurs at $x/c_x = 0.34$. This leads the author to believe the Type-I transition behaviour is more driven by lower frequencies than higher frequencies. As lower frequencies tend to produce a less unstable crossflow vortex, this would also explain the slight delay in transition when analysing the location observed in subsection 5.3.1.

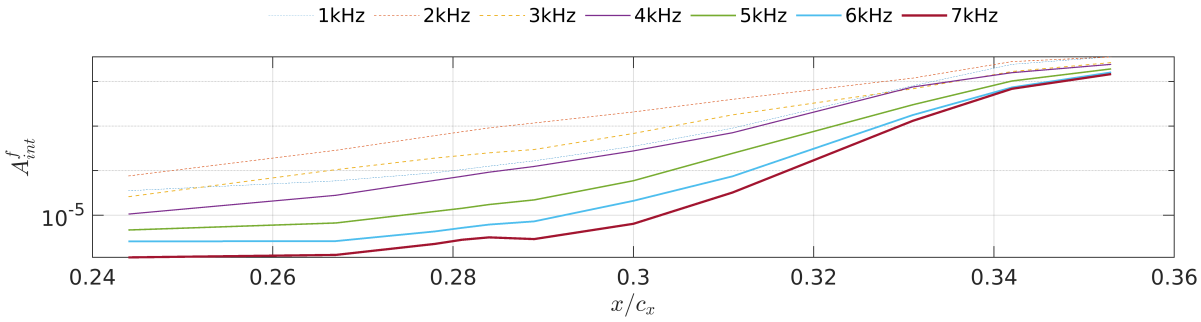

Figure 5.32: Secondary CFI PSD amplitude development as a function of x/c_x for frequencies from $1kHz$ to $7kHz$, in bands of $1kHz$. Values are based on the integration of the amplitudes of each measurement point at a specific x/c_x location.

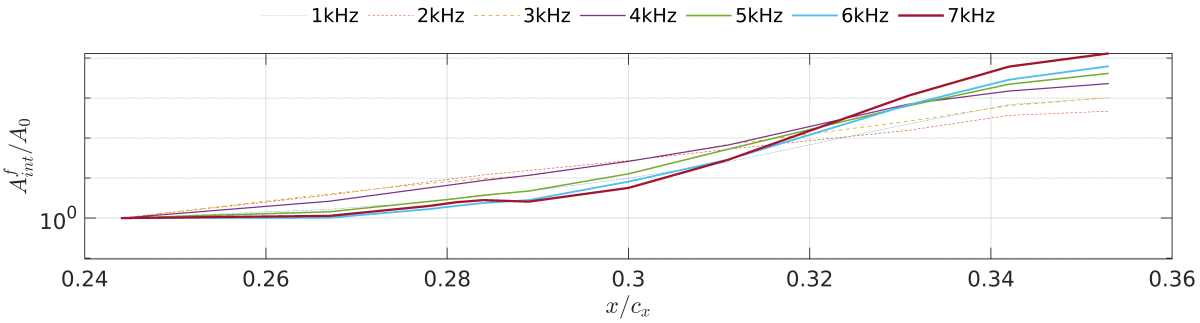

Figure 5.33: Secondary CFI PSD amplitude development as a function of x/c_x for frequencies from $1kHz$ to $7kHz$, in bands of $1kHz$. Values are based on the integration of the amplitudes of each measurement point at a specific x/c_x location. Each curve is normalised by their respective amplitudes at $x/c = 0.22$.

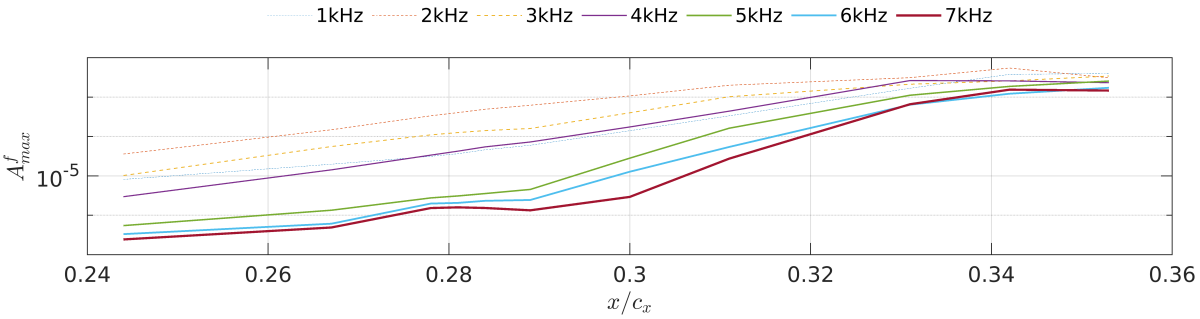

Figure 5.34: Secondary CFI PSD amplitude development as a function of x/c_x for frequencies from $1kHz$ to $7kHz$, in bands of $1kHz$. Values are based on the maximum PSD amplitude found at a specific x/c_x location.

To confirm this lower-frequency dominance and the differences in high-frequency amplitudes between DNS and experiments, a spatial analysis of the secondary CFI frequency bands was performed as well. The author returns to the standard deviation analysis ($\sigma_Q/\overline{Q_e}$) like was performed in subsection 5.3.2. By now filtering the raw Q' signals, and using this to calculate $\sigma_Q/\overline{Q_e}$ across each HWA plane, the spread of frequency amplitude inside the crossflow vortices can be analysed. Following the same steps outlined in Rius-Vidales and Kotsonis (2022), two frequency bands are filtered. A low-frequency band (B_L, $450Hz \leq f \leq 3kHz$) and a high-frequency band (B_H, $3.5kHz \leq f \leq 9kHz$). Spatially represented by Figure 5.35 and Figure 5.36, respectively, they both compare four z-y planes upstream of transition. For B_L, the magnitudes of σ remain very similar between both DNS and HWA data. Like in the unfiltered data, the HWA data is again more diffused than the numerical equivalent, but the similarities are apparent. The rapid growth in magnitude when moving downstream is also similar between DNS and HWA. The differences in the magnitude of σ are more pronounced for B_H in Figure 5.36, where the maximum HWA $\sigma_Q/\overline{Q_e}$ magnitudes are higher in every plane. More relevant perhaps, is the noticeably faster growth in the magnitude of the DNS fluctuations, even though the initial magnitude of σ_Q in DNS at $x/c_x = 0.284$ is significantly lower than its HWA equivalent. This again backs up the earlier findings that high-frequency content only becomes significant much further downstream in DNS, but then starts to grow much faster than the HWA data indicates.

The numerical signal length of only 0.025s does not offer an explanation for the discrepancies in the high-frequency band. With a sampling rate of over $250kHz$, the Nyquist frequency is as high as $125kHz$, more than sufficient for capturing $5kHz$ velocity fluctuations. The shorter signal only influences the lowest possible frequency captured, although this is well below the range of interest, at $180Hz$. The free-stream turbulence modes prescribed at the inflow boundary also introduce a large range of frequencies, from $26Hz$ to $26kHz$, excluding insufficient frequency range from the potential issues. The amplitudes assigned to the disturbance modes introducing these frequencies do potentially cause problems, though it is an unlikely scenario. The distribution of free-stream turbulence inflow modes is based on energy. Distributed around a central frequency based on the turbulence integral length scale, it provides an optimal spread, even though the frequency bands furthest away from this central frequency are less represented in the flow field. This under-representation of the higher frequency

bands could point towards a discrepancy in dissipation rates, where higher frequencies are dissipated faster. However, one thing becomes increasingly apparent, which is the lack of high-frequency content far upstream of the numerical transition region. All other factors considered when comparing DNS and experimental data have a good agreement, only a tangible explanation for the lack of high-frequency components is still missing to complete this study.

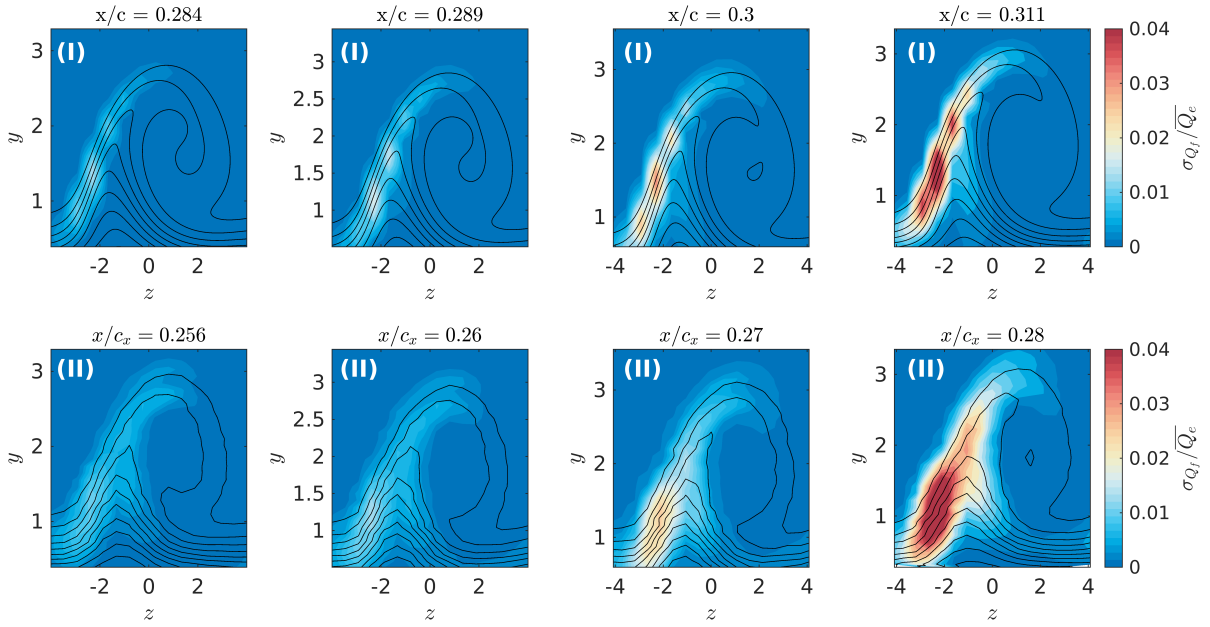

Figure 5.35: Bandpass filtered (B_L,$450Hz \leq f \leq 3000Hz$) normalised standard deviation of the raw temporal velocity fluctuations for both DNS (I) and HWA (II) data (coloured contour plots) and matching time-averaged velocity field $Q/\overline{Q_e}$ (solid lines).

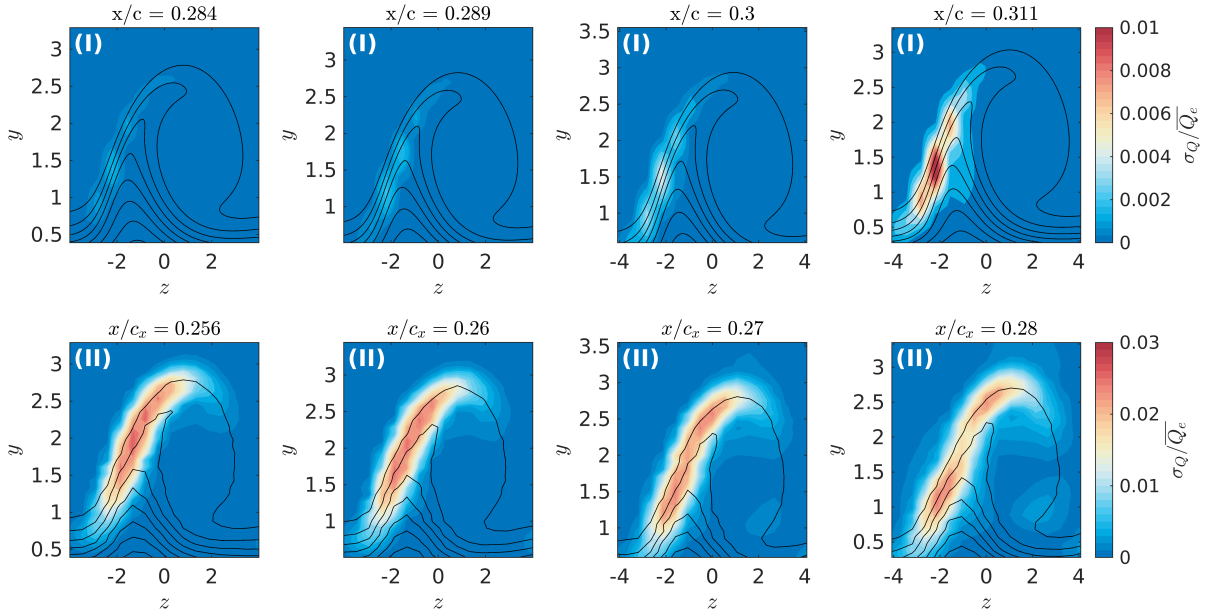

Figure 5.36: Bandpass filtered (B_H,$3500Hz \leq f \leq 9000Hz$) normalised standard deviation of the raw temporal velocity fluctuations for both DNS (I) and HWA (II) data (coloured contour plots) and matching time-averaged velocity field $Q/\overline{Q_e}$ (solid lines)..

5.3.4. Receptivity: Free-Stream Turbulence

With free-stream turbulence being used as the source of unsteadiness in this simulation, it is most likely to harbour the explanation as to why the high-frequency component of the Type-I secondary CFI is underrepresented. Free-stream turbulence is continuously generated and prescribed at the inflow

boundary of the domain. Its exact magnitude is based on separate wind tunnel tests measuring the turbulence intensity at a range of different velocities with the model in the test section (Serpieri, 2018). Due to the large relative spread of turbulence intensity values for the different velocities (see Figure 5.37a), the decision was made to curve-fit the measurement points to get the best estimate of Tu. The result is $Tu = 0.022\%$, which the numerical methodology introduces at the inflow plane using frequencies between $46Hz$ and $26kHz$. Due to the methodology used to assign which frequencies get prescribed, the spread is somewhat biased towards the lower side of the frequency spectrum, as visualised in Figure 5.37b. After numerical implementation at the inflow, the energy contained within this FST spectrum starts to dissipate while propagating downstream. Unlike in the wind tunnel, there are no external factors influencing and therefore dissipation of the FST in this case always starts immediately after initiation at the inflow. This implies the turbulence intensity value in the free-stream flow field decreases monotonically and reduces an order of magnitude in intensity in a matter of 10% of the chord. The extracted turbulence intensity values at several different x-locations along the chord, for two different wall-normal distances in the free-stream flow, prove this high dissipation rate (see Figure 5.38a). It is expected that the turbulence intensity value in the experiment remains relatively constant across the test section, which is now proven not to be the case in the numerical domain.

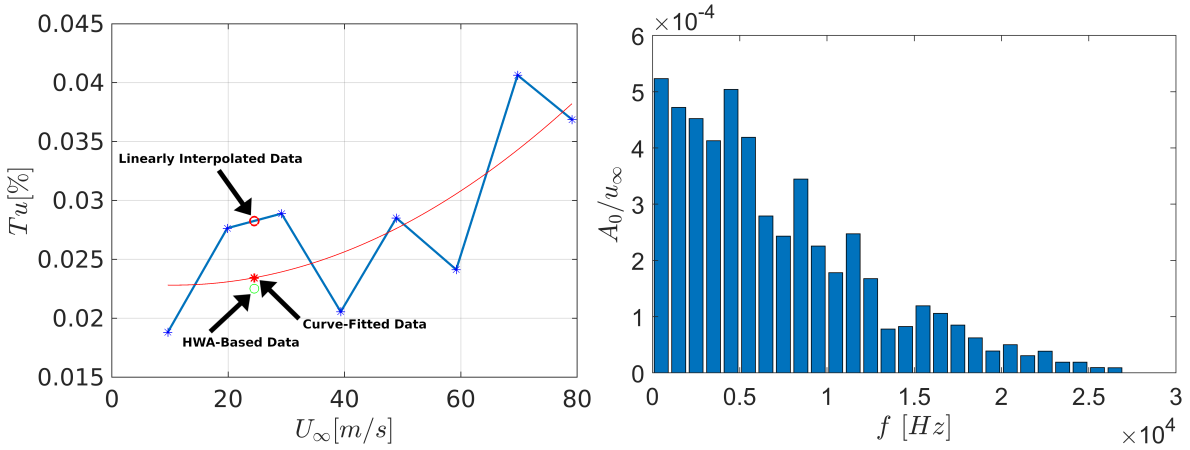

(a) (b)

Figure 5.37: a) Wind tunnel Tu-levels without a model in the test section at a range of free-stream velocities (*), curve-fitted (–) and interpolated for the required velocity (*). HWA-measured Tu extracted from the sampled planes indicated for the readers' information (◦) Turbulence intensity measurements performed by Serpieri (2018). b) Distribution of the relative amplitudes assigned to each frequency at the inflow of the numerical domain as part of the free-stream turbulence generation.

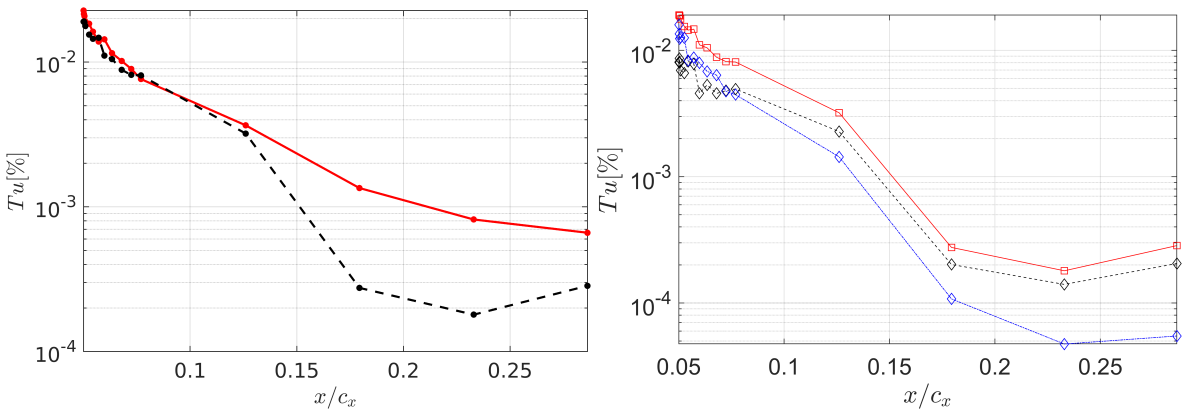

(a) (b)

Figure 5.38: a) DNS Tu-levels as a function of x/c_x for $y/\delta_0 = 10$ (–) and $y/\delta_0 = 20$ (--). b) Tu-levels in the DNS free-stream flow as a function of x/c_x for the low-frequency band ($f \leq 3000Hz$, black line), the high-frequency band ($3000Hz \leq f$, blue line) and the total free-stream flow (red line).

Now that a monotonically dissipating trend in the total FST propagation has been identified, a closer look also reveals differences between the dissipation of high and low frequencies. By dividing the total FST content into a low ($f \leq 3000Hz$) and high ($3000Hz \leq f$) frequency segment, the overall dissipation trend discussed previously can be reconstructed for these separate segments (see Figure 5.38b). Although the turbulence intensity levels for both segments very close to the inflow are similar, this analysis indicates these values start to diverge rather quickly further downstream. With the high-frequency content dissipation faster, the difference in Tu with the low-frequency content is already an order of magnitude after only 5% of the chord. How large this impacts the frequency spectrum at positions downstream of the inflow boundary is also clear when comparing overlapped PSD plots for different x/c_x positions (see Figure 5.39). While near the inflow boundary, the PSD spectrum is still roughly flat, the higher frequency sections start to drop more quickly in magnitude than the lowest frequencies. This difference in dissipation rate is therefore another major factor that could play a role in the suppression of high-frequency content inside the Type-I secondary CFI region.

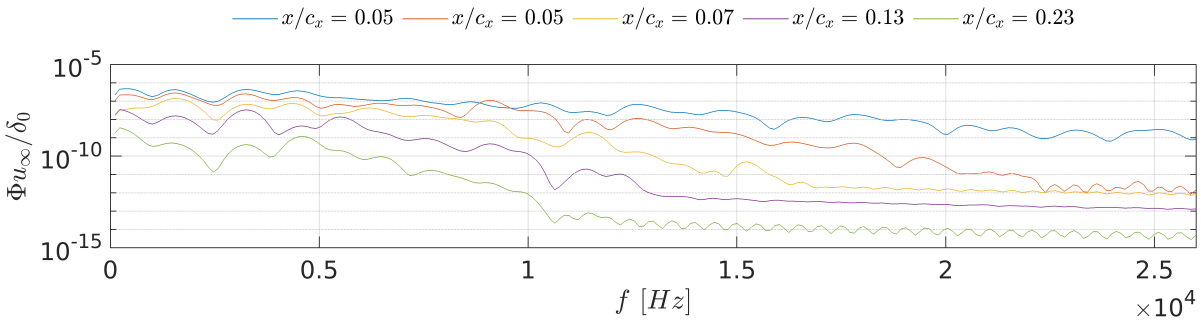

Figure 5.39: Spectral analysis of probes at different chordwise locations in the free-stream flow, ranging from $x/c_x = 0.05$ to $x/c_x = 0.23$, indicating the increased dissipation rate of the higher frequencies within the spectrum.

The HWA measurements at the upper boundary of each extracted plane provide data close to free-stream conditions, although potentially contaminated by the boundary layer instabilities, these appear to confirm the relatively constant turbulence intensity across the measurement domain assumed earlier. Using the HWA planes is merely an estimation of the turbulence intensity values, however, and most likely an exaggeration due to influences of the boundary layer. Even if the exact turbulent energy dissipation rates across the experimental test section would be known, it would require significant research and numerical implementation to impose the exact turbulence intensity values on the DNS domain. On top of this inherent dissipation problem, grid resolution in the free-stream flow is coarser than inside the boundary layer. To avoid additional dissipation increases related to grid resolution increases in free-stream grid refinement are to be considered in future research. Being too costly for this research, a higher grid resolution is not considered here. However, it is expected to improve future research outcomes. Keep in mind that increasing resolution is not going to inverse turbulence dissipation, but the process will most likely be delayed. Especially higher frequencies could benefit from this improvement, potentially increasing the high-frequency content in the boundary layer as well.

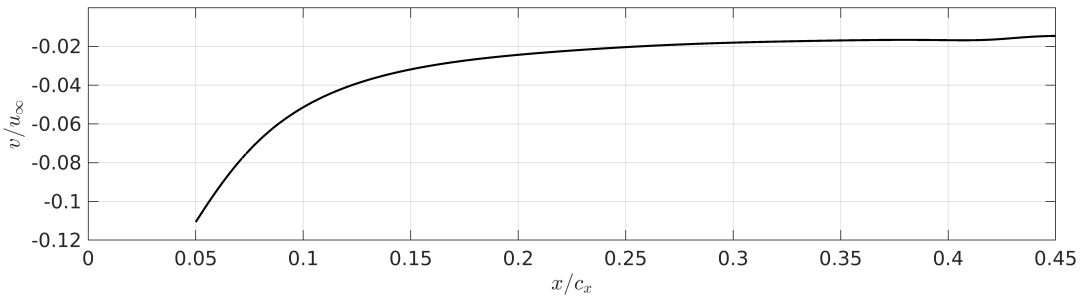

Figure 5.40: Wall-normal velocity v/u_∞ along the domain top boundary ($y/\delta_0 = 25$), where negative values indicate flow moving towards the wing.

With the wall-normal velocity component in the free-stream flow pointing towards the wall, there is an additional influx of unperturbed flow through the top domain boundary. In this particular simulation,

free-stream turbulence is only being prescribed at the inflow boundary, and therefore the influx from the top boundary does not contain any disturbances. Although the effect is minimal in the upstream sections of the flow, its influence slowly gets closer to the boundary layer. Figure 5.40 indicates the velocity through the top boundary, giving the reader an idea of the influx of flow not containing free-stream turbulence. Based on this, it could be argued that future research should include free-stream turbulence introduction through the top boundary if the domain increases in streamwise length.

5.4. Laminar-Turbulent Transition Without FST

Multiple other techniques to trigger unsteady disturbances in the flow field exist, such as the Blowing & Suction technique commonly used when simulating incompressible flow turbulent transition over a swept wing. Previously used in work as Casacuberta et al. (2021) and Casacuberta, Hickel, et al. (2022), this Blowing & Suction technique is discussed more in detail in Appendix A. Proven to work well in the past, this is a suitable technique to trigger unsteadiness inside the boundary layer. However, it is more likely to require an iterative approach to arrive at the desired numerical solution than the FST methodology presented in this work. To emphasise the effectiveness of unsteady forcing, and more specifically the FST methodology used here, a simulation without any unsteady forcing is run as well to briefly compare the impact on the transition front location and corresponding isosurfaces. This implies continuing to run the steady flow field simulation used in the FST methodology, but instead of adding unsteadiness, simply removing the artificial stabilisation (SFD). When left to develop without stabilisation, transition is expected to occur if the domain is large enough, regardless if unsteady forcing is used or not. This approach is unlikely to result in a flow field that resembles the experimental measurements since unsteadiness is inherently present near the numerical inflow boundary in every possible flight or wind tunnel scenario. Nevertheless, it is an excellent confirmation numerical free-stream turbulence improves the simulation capabilities of crossflow-driven transition over a swept wing. As the domain remains the same as the previously discussed free-stream turbulence simulation, the transition location of both cases can be assessed in the same way (see Figure 5.41). Through an analysis of the skin friction coefficient C_f in Figure 5.41, the large difference with the experimental transition location becomes clear. With transition in the numerical case without unsteady forcing occurring at roughly $x/c_x = 0.42$, this results in a difference of about $x/c_x = 0.1$ with the experiment.

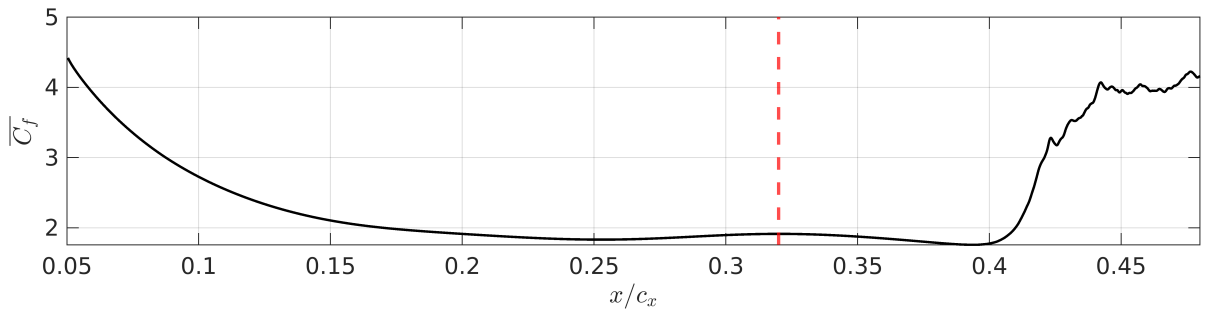

Figure 5.41: Time- and spanwise-averaged C_f as a function of x/c_x along the wall boundary in the DNS domain without FST, indicating the experimental laminar-to-turbulent transition location (--) (Rius-Vidales & Kotsonis, 2022).

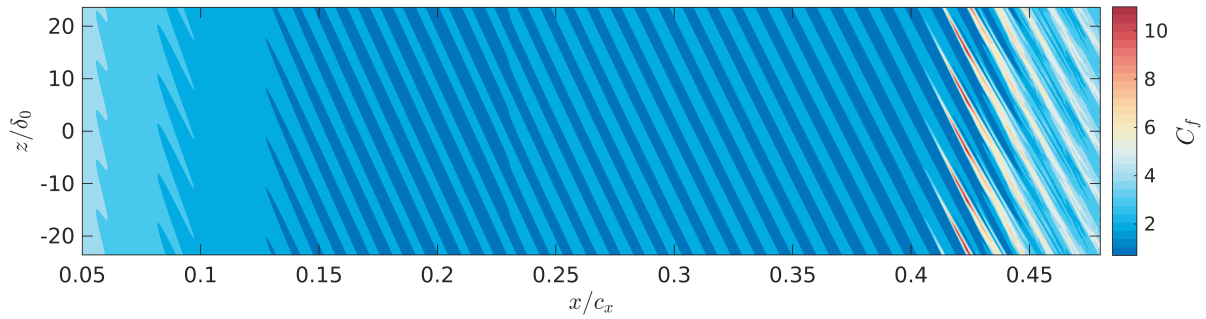

Figure 5.42: Time-averaged C_f in the x-z plane along the wall boundary in the DNS domain without FST.

The two-dimensional plane of C_f across the wing (see Figure 5.42) backs up this finding that without unsteady forcing through FST, the transition front moves about 8% of the chord downstream, while this difference is 10% chord compared to the experimental data. This behaviour is the logical result of not introducing unsteadiness at the inflow boundary. Although the secondary CFI naturally starts to develop moving downstream, these take longer to develop than disturbances already introduced with a measurable amplitude at the inflow boundary. Looking at the Q-criterion isosurface of this simulation in Figure 5.43, it shows very coherent structures leading to a seemingly slower transition process compared to FST. Also starting with modulations in the upwelling Type-I region of the crossflow vortex, these disturbances slowly start to spread around the Type-I region, eventually leading to breakdown. The breakdown appears to start closer to the wall than with FST, affecting the upper parts of the boundary layer less in the early stages of transition. Instead of Type-II modulation, Type-III seems the second most dominant instability in this simulation, although conclusions can not be drawn from this isosurface alone. The flow field after transition also shows a much less chaotic pattern than the FST simulation, and the formation of hairpin vortices is clearly visible in a pattern similar to that seen with Blowing & Suction in Casacuberta, Groot, et al. (2022).

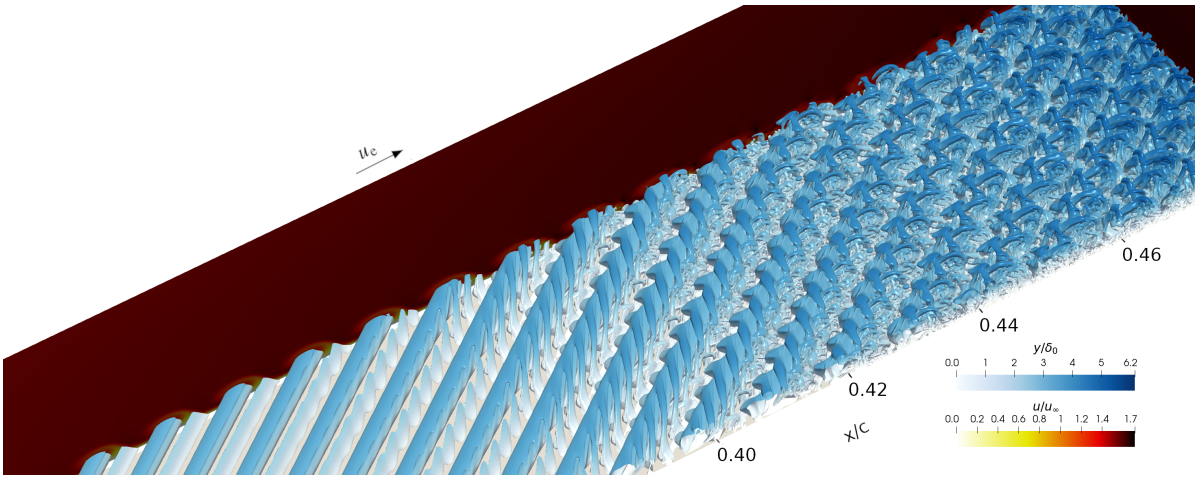

Figure 5.43: Detailed snapshot of the unsteady DNS instantaneous Q-criterion isosurface without FST, coloured by v/u_∞, along with y-z and y-x u/u_∞ slices. Domain duplicated four times in z-direction based on periodic boundary conditions.

These clear differences with the FST simulation in transition type and behaviour amplify the need for a suitable numerical framework to simulate crossflow development and transition to turbulence on swept wings. The numerical use of free-stream turbulence for this purpose has proven to be a practical tool to represent the experimental flow field. Therefore, it can already be concluded that there is a clear need for artificial unsteadiness introduction in this type of simulation and that FST is an excellent candidate to set up DNS in future research.

6

Conclusion & Recommendations

This work's main aim lay in determining how to generate a framework to numerically simulate the development of crossflow instabilities (CFI) and transition over a swept wing. The structure ultimately provided, based around a global free-stream turbulence (FST) method, offers flexibility towards future research because of its lacking need for a priori transition information. This chapter concludes the Master's Thesis report by bundling the main findings of this work. Not only focusing on the agreement between the experimental and numerical data sets but also on the possible reasons for discrepancies between both. Assessing how the numerical tools played a role in the outcome, and what experimental data proved to be most useful, leads to recommendations for future work. These recommendations will be made on two main fronts: the reference experimental data available, along with the possibilities for future improvement on the numerical methodology.

6.1. Conclusion

The generation of a framework for the simulation of CFI development and transition to turbulence over a swept wing was based on the replication of experimental data obtained by Rius-Vidales and Kotsonis (2022). Throughout the numerical reconstruction of the unsteady wind tunnel flow field, three major components have been identified. The first of which was the baseflow, where baseflow refers to the steady-state solution of the Navier-Stokes equations. To subsequently obtain the steady flow field, primary crossflow instabilities (CFI) were then introduced in the baseflow, followed by the addition of FST to trigger secondary CFI growth in the last step towards transition to turbulence. Resembling three different stages within the transition process, the three steps were implemented chronologically. With limited experimental data on the baseflow available, the spanwise invariant simulation was based on the airfoil static pressure data and wind tunnel parameters such as inflow velocity magnitude and total pressure. A model for incompressible swept-wing flow was laid out to generate this type of baseflow repeatedly. Both experimental pressure data and the FSC boundary layer equations used for reference indicate variations of less than 2% compared to the computed DNS free-stream velocity magnitude distribution, U_e. Even though in-depth confirmation of the numerical accuracy remains challenging due to the lack of experimental data available on the baseflow, overall, a solid basis for the addition of primary CFI was obtained. These primary CFI were generated by a linear local stability method (ILST) and assigned an amplitude through NPSE-based optimisation before being added to the numerical domain at the inflow boundary. Converged in time, the resulting flow field shows crossflow vortices creating wave-like structures inside the boundary layer and is referred to as the steady flow field. Directly comparing numerical and experimental primary CFI growth rates, MFD and disturbance profiles showed good agreement in both shape and amplitude.

The numerical match with experimental data at this stage was by no means perfect, however, but remained a significant improvement compared to previous attempts by Nishino and Shariff (2010) and Tempelmann et al. (2011). By building a numerical framework for the generation of inflow boundary conditions with stability tools (ILST & NPSE), a fairly low-cost method to tailor the numerical flow field to the experimental data has been developed. Proven to work for flow over the pressure side of an airfoil, one of the main reasons for discrepancies in the results can be traced back to the experimental

data. Sampling HWA data takes significant time and resources, leading to each of the HWA planes being sampled on different days. The Low Turbulence wind tunnel at TU Delft is operated based on the Reynolds number alone, therefore, the ambient parameters varied over the different experimental test days. Among others, this led to noticeable differences in free-stream velocity and pressure. The variation in free-stream parameters the wind tunnel operated on while measuring each HWA plane were discussed in section 2.2. Full correction for these differences is difficult, but their effects have been mitigated as much as possible by normalising the data before comparison. However, small errors varying per plane are therefore still expected, and only the average fit to all planes is taken into account across this research. Neglecting the geometric curvature of the airfoil can be considered another source of error. The effect of the curvature is reproduced through the boundary conditions with success, but exact replication remains an assumption based on small surface curvature. This effect is nevertheless considered minor compared to not simulating the flow around the leading edge. Receptivity to disturbances is generally strongest near the front of the airfoil, which therefore plays a large role in transition behaviour. Although mostly relevant for secondary instabilities, primary crossflow instabilities are also first found near the leading edge. Without the most upstream 5% of the airfoil chord being simulated, this work still managed to produce sufficiently accurate inflow boundary conditions to reproduce the steady experimental flow field.

Other minor effects that could play a role in the result accuracy, such as wind tunnel wall influences and the infinite span wing assumption, are excluded from this work but are thought to have limited impact on this experimental setup (Rius-Vidales & Kotsonis, 2022). Based on the results presented in this section, this method has proven to produce a steady numerical flow field very close to the experimental equivalent. This method remains relatively low-cost and repeatable because of the use of stability tools and the limited input necessary. Most NPSE input parameters are not case-specific, and therefore the main challenge for future projects is related to obtaining an accurate baseflow as input for the NPSE optimisation. If no DNS baseflow is available, the required input can also be generated using numerical boundary layer solvers instead. When the baseflow is available and the experimental data is then formatted correctly, the optimisation procedure is able to calculate the remaining DNS inputs required. The experimental data used is preferably obtained through HWA or PIV techniques, which allows for the straightforward generation of disturbance profiles. Other types of experimental data can also be used, considering they can provide the necessary data to generate disturbance profiles. If that is not possible, the method will have to be adapted for a different error metric. This flexibility implies this methodology does not only work in a narrow range of applications but can be tailored to a number of different inputs, depending on the experimental data and numerical tools available. The steady flow field discussion can be closed here for now, as a well-matching numerical result is available as a basis to start simulating unsteadiness.

By then numerically generating FST based on the wind tunnel measurements, secondary CFI growth was triggered within the steady flow field. Without the addition of FST, transition was shown to occur 10% of the chord further downstream than the experimental data suggested. Through the addition of FST, this difference was reduced to a mere 2% of the chord downstream of the experimentally determined transition front. Considering the spanwise differences in laminar-turbulent transition location of up to $x/c_x = 0.05$ throughout the experiment, visible in the IR thermography images from Rius-Vidales and Kotsonis (2022), the numerical transition-front location is still deemed an acceptable compared to experiments. The Type-I-dominant transition behaviour present in the simulation with FST also matches the experimental observations, even though the noticeable Type-II secondary CFI response observed during the experiment remains largely subdued in DNS. Potential causes for this discrepancy can be found in the secondary CFI frequency content. The lower frequency content ($450Hz \leq f \leq 3000Hz$) in DNS is the most prominent in the Power Spectral Density (PSD) analysis. This band is also distinguishable in the experimental PSD analysis, based on the HWA data from Rius-Vidales and Kotsonis (2022). The most dominant frequencies measured during the experiment in the Type-I region ($4000Hz \leq f \leq 6000Hz$), are noticeably less strong in the DNS. These high-frequency secondary CFI are much less prominent in the earlier stage of transition in the numerical simulation. However, the rapid growth of the high-frequency CFI very close to transition still indicates the same Type-I-dominated transition mechanisms discovered in experiments. This delayed growth of higher frequencies is most likely the cause of the small difference in transition-front location due to the inherent high-frequency nature of the Type-I instability (Koch et al., 2000).

Despite the resemblance in transition-front location between the numerical and experimental data,

the delayed amplitude growth of high-frequency secondary CFI thus indicated the necessity for further research. The assessment of the similarity between both data sets hinged on two main pillars: the NPSE optimisation leading to a representative steady flow field in DNS, and the FST methodology triggering secondary CFI. While the former was proven to lead to an acceptable match with experiments, FST required additional analysis to identify its current benefits and shortcomings. The capability of FST to reproduce the unsteady flow field without the need for iteration on the boundary conditions is its major advantage. Previous techniques such as Blowing & Suction, required a number of DNS iterations to obtain representative inputs, while FST makes use of actual wind tunnel parameters, reducing the need for iterative DNS simulations. In spite of it showing good agreement in transition location and secondary CFI types, discrepancies in terms of frequency content were still noticeable with the experimental equivalent. Considering the size of the numerical domain is significantly smaller than the wind tunnel test section, an aspect of the unsteady experimental flow field is most likely under-represented at the inflow boundary. The current methodology to generate FST focuses on the random distribution of wavenumbers, which are all assigned certain energy levels based on the turbulence integral length scale and the desired turbulence intensity. Despite high frequencies being introduced at the DNS inflow boundary, the current FST methodology indicates a distribution bias towards the lower frequencies within the used range. Because of the turbulence integral length scale used, the energy distribution within this frequency range is also biased towards the lower frequencies. Adding the problem that higher frequencies dissipate rapidly within the free-stream flow, this is thought to be the main contributor to the delay of high-frequency growth within the Type-I and Type-II secondary CFI. Regardless of the remaining discrepancies, this framework has shown potential benefits for future research. Next to the optimisation technique to simulate the steady flow field reliably, the promising nature of FST generation in DNS has also been presented. Despite the need for more research on this front, the current framework that has been developed for numerical simulation of crossflow development and transition on swept wings showed room for both flexibility and repeatability. Therefore, the author is confident the framework will prove valuable for future numerical and experimental research on incompressible crossflow over swept wings. Not only as a whole, but also smaller sections to generate an accurate baseflow, or create a trustworthy steady flow field for transition studies.

6.2. Recommendations

Based on the numerical match achieved, the framework for numerical simulation developed during this research shows benefits for future research into crossflow, both with and without complementary experimental data. Despite current shortcomings in terms of high-frequency oscillations in the laminar-turbulent transition region, future work based on the findings of this research can fine-tune the functioning of this framework. This in turn can lead to achieving replication of more complex crossflow-driven transition scenarios, at the same time enlarging the experimental data set noticeably.

6.2.1. Experimental Data

Acquiring experimental data of CFI-dominated boundary layers is a time- and cost-intensive process, often limited in the type and amount of measurements that can be taken at once. Therefore, not all measurement data desired to set up a numerical equivalent can always be provided. In this particular work, a wealth of experimental data is made available, which is found still to be rather limiting when it comes to setting up a simulation. To allow for baseflow simulation within this framework, surface pressure ports are a necessity. Ideally, free-stream velocity components are available (i.e. through multi-axis HWA measurements), though this work has proven decomposition is possible for incompressible flow without major assumptions is possible. For steady and unsteady flow field generation, there is a need for some sort of three-dimensional measurement data. Either through multiple two-dimensional HWA planes, or PIV. Alternatives can replace HWA in the future, but this type of measurement is sufficient for this framework. To increase confidence in the NPSE-based optimiser, individual velocity components are still preferred over the velocity magnitude Q ($= \sqrt{(u\cos\Lambda + w\sin\Lambda)^2 + v^2}$), to ensure all components are adequately represented in the steady flow field. The main recommendation on the HWA or PIV data for future work is to include several free-stream measurements, sufficiently far from the boundary layer. Given the FST methodology works with free-stream wind tunnel conditions, it is beneficial to have direct wind tunnel measurements available, rather than convoluted ways to obtain statistical wind tunnel data. Despite being proven sufficient to produce an accurate numerical simulation within this

framework, all measurement data used is considered strictly required. The additional measurements mentioned in this section are therefore only desired to facilitate the process and reduce the number of assumptions made.

6.2.2. Numerical Crossflow Replication

The baseflow and primary CFI components of the framework introduced here produce an acceptable representation of the experimental flow field. The secondary CFI component based on FST requires further research, in spite of its promising match with the experimental transition-front location and CFI types. Therefore, the main recommendations for future work centre around potentially improving the secondary CFI frequency content. Since the free-stream frequency content can be directly linked to both the FST technique and the overall simulation setup, improvements to the methodology will be required. A broader range of frequencies should be introduced, which is currently limited by the free-stream grid resolution. By increasing this resolution, the range of wavenumbers used in the methodology becomes larger, resulting in a wider frequency spread at the inflow. Despite the inherent bias of this method to lower frequencies, it is still believed this will shift the frequency range to higher frequencies overall. To allow for these higher frequencies, and potentially lower dissipation rates, a higher free-stream grid resolution will also have to be investigated. Higher resolution will inevitably result in an increase in computational resources required but is deemed necessary to resolve these frequencies across the domain. This increased resolution is thought to also lower dissipation rates, potentially supporting the highest frequencies within the FST until further downstream. If this does not achieve the desired results, a different method should be researched to keep the FST levels high throughout the domain. If this is successful, one might think about moving the inflow boundary of the DNS domain further upstream on the wing if computational resources allow. Since DNS currently only simulates a small portion of the wind tunnel test section, a large number of assumptions were made about the non-simulated flow around the domain. A larger simulation domain thus implies fewer assumptions in this regard, as well as a long domain for sustained FST to influence the boundary layer and instability growth. Additionally, an equivalent numerical simulation with this Blowing & Suction should also be run to directly compare the FST methodology with existing methods for unsteady forcing. This is to provide a more comprehensive study of the benefits and downsides of FST compared to established methods. These recommended studies will further provide proof of the value of this work. The ultimate test left will then be to apply the framework to a different low-speed experiment over a swept-wing model.

Bibliography

Åkervik, E., Brandt, L., Henningson, D. S., Hœpffner, J., Marxen, O., & Schlatter, P. (2006). Steady solutions of the Navier-Stokes equations by selective frequency damping. *Physics of Fluids*, *18*(6). https://doi.org/10.1063/1.2211705/920038

Anderson, J. (2016). *Fundamentals of Aerodynamics* (6th Edition). McGraw Hill.

Arnal, D., Coustols, E., & Juillen, J. C. (1984). Etude expérimentale et théorique de la transition sur une aile en flèche infinie. *La Recherche aérospatiale*, (4), 275–290.

Arnal, D., Gasparian, G., & Salinas, H. (1998). Recent advances in theoretical methods for laminar-turbulent transition prediction. *36th AIAA Aerospace Sciences Meeting and Exhibit*. https://doi.org/10.2514/6.1998-223

Balzer, W., & Fasel, H. F. (2016). Numerical Investigation of the Role of Free-stream Turbulence on Boundary-Layer Separation. *Journal of Fluid Mechanics*, *801*, 289–321. https://doi.org/http://dx.doi.org/10.1017/jfm.2016.424

Bippes, H. (1999). Basic experiments on transition in three-dimensional boundary layers dominated by crossflow instability. *Progress in Aerospace Sciences*, *35*(4), 363–412. https://doi.org/10.1016/S0376-0421(99)00002-0

Blasius, H. (1908). Grenzschichten in Flüssigkeiten mit kleiner Reibung. *Zeitschrift für Angewandte Mathematik und Physik*, *56*, 1–37.

Bonfigli, G., & Kloker, M. (2007). Secondary instability of crossflow vortices: Validation of the stability theory by direct numerical simulation. *Journal of Fluid Mechanics*, *583*, 229–272. https://doi.org/10.1017/S0022112007006179

Borodulin, V. I., Ivanov, A. V., Kachanov, Y. S., Mischenko, D. A., Örlü, R., Hanifi, A., & Hein, S. (2019). Experimental and theoretical study of swept-wing boundary-layer instabilities. Unsteady cross-flow instability. *Physics of Fluids*, *31*(6). https://doi.org/10.1063/1.5094609

Borodulin, V. I., Ivanov, A. V., Mischenko, D. A., Örlü, R., Hanifi, A., & Hein, S. (2019). Experimental and theoretical study of swept-wing boundary-layer instabilities. Three-dimensional Tollmien-Schlichting instability. *Physics of Fluids*, *31*.

Brandt, L., Schlatter, P., & Henningson, D. S. (2004). Transition in boundary layers subject to free-stream turbulence. *Journal of Fluid Mechanics*, *517*, 167–198. https://doi.org/10.1017/S0022112004000941

Casacuberta, J., Hickel, S., & Kotsonis, M. (2021). Mechanisms of interaction between stationary cross-flow instabilities and forward-facing steps. *AIAA Scitech 2021 Forum*. https://doi.org/10.2514/6.2021-0854

Casacuberta, J., Hickel, S., Westerbeek, S., & Kotsonis, M. (2022). Direct numerical simulation of interaction between a stationary crossflow instability and forward-facing steps. *Journal of Fluid Mechanics*, *943*, A46. https://doi.org/10.1017/JFM.2022.456

Casacuberta, J., Groot, K. J., Hickel, S., & Kotsonis, M. (2022). Secondary instabilities in swept-wing boundary layers: Direct Numerical Simulations and BiGlobal stability analysis. *AIAA Scitech 2022 Forum*. https://doi.org/10.2514/6.2022-2330

Casacuberta, J., Groot, K. J., Tol, H. J., & Hickel, S. (2018). Effectivity and efficiency of selective frequency damping for the computation of unstable steady-state solutions. *Journal of Computational Physics*, *375*, 481–497. https://doi.org/10.1016/J.JCP.2018.08.056

Çengel, Y. A., & Cimbala, J. M. (2006). *Fluid Mechanics: Fundamentals and Applications*. McGraw-Hill.

Dagenhart, J. R., & Saric, W. S. (1999). *Crossflow Stability and Transition Experiments in Swept-Wing Flow* (tech. rep.). Langley Research Center. Hampton, Virginia. http://www.sti.nasa.gov

Delft High Performance Computing Centre (DHPC). (2022). DelftBlue Supercomputer (Phase 1). https://www.tudelft.nl/dhpc/ark:/44463/DelftBluePhase1

Deyhle, H., & Bippes, H. (1996). Disturbance growth in an unstable three-dimensional boundary layer and its dependence on environmental conditions. *Journal of Fluid Mechanics*, *316*, 73–113. https://doi.org/10.1017/S0022112096000456

Deyhle, H., Höhler, G., & Bippes, H. (1993). Experimental investigation of instability wave propagation in a three-dimensional boundary-layer flow. *AIAA Journal*, *31*(4), 637–645. https://doi.org/10.2514/3.11597

Donlan, C. J., & Weil, J. (1952). *Characteristics of Swept Wings at High Speeds* (tech. rep.). National Advisory Committee for Aeronautics.

Durovic, K. (2022). *Direct Numerical Simulation of Boundary-Layer Transition with Free-stream Turbulence* (Doctoral dissertation). KTH. Sweden, Stockholm. https://www.diva-portal.org/smash/get/diva2:1631337/FULLTEXT01.pdf

Eckert, M. (2010). The troublesome birth of hydrodynamic stability theory: Sommerfeld and the turbulence problem. *Eur. Phys. J. H*, *35*, 29–51. https://doi.org/10.1140/epjh/e2010-00003-3

Falkner, V., & Skan, S. (1931). LXXXV. Solutions of the boundary-layer equations. *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science*, *12*(80), 865–896. https://doi.org/10.1080/14786443109461870

Fischer, T., & Dallmann, U. (1991). Primary and Secondary Stability Analysis of a Three-Dimensional Boundary-Layer Flow. *7. DGLR-Fachsymposium 'Strömungen mit Ablösung'*, 91–95.

Friederich, T. A. (2013). Control of the secondary crossflow instability : direct numerical simulation of localized suction in three-dimensional boundary layers. https://doi.org/10.18419/OPUS-3920

Gray, W. E. (1952a). The effect of wing sweep on laminar flow. *RAE TM 255*.

Gray, W. E. (1952b). The nature of the boundary layer flow at the nose of a swept wing. *RAE TM 256*.

Grosch, C. E., & Salwen, H. (1978). The continuous spectrum of the Orr-Sommerfeld equation. Part 1. The spectrum and the eigenfunctions. *J. FZuid Mech*, *87*(1), 33–54. https://doi.org/10.1017/S0022112078002918

Hader, C., & Fasel, H. F. (2022). Direct Numerical Simulations (DNS) of Natural Transition in High-Speed Boundary Layers Using a Broadband Random Forcing Approach. *IUTAM Bookseries*, *38*, 565–574. https://doi.org/10.1007/978-3-030-67902-6{_}49

Hall, P., Malik, M., & Poll, D. (1984). On the stability of an infinite swept attachment line boundary layer. *Proc. R. Soc. Lond. A*, *395*(1809), 229–245. https://doi.org/10.1098/rspa.1984.0099

Haynes, T. S., & Reed, H. L. (2000). Simulation of swept-wing vortices using nonlinear parabolized stability equations. *J. Fluid Mech.*, *405*, 325–349. https://doi.org/10.1017/S0022112099007260

Hickel, S., & Adams, N. A. (2008). Implicit LES applied to zero-pressure-gradient and adverse-pressure-gradient boundary-layer turbulence. *International Journal of Heat and Fluid Flow*, *29*(3), 626–639. https://doi.org/10.1016/J.IJHEATFLUIDFLOW.2008.03.008

Hickel, S., Egerer, C. P., & Larsson, J. (2014). Subgrid-scale modeling for implicit large eddy simulation of compressible flows and shock-turbulence interaction. *Physics of Fluids*, *26*(10). https://doi.org/10.1063/1.4898641

Högberg, M., & Henningson, D. (1998). Secondary instability of cross-flow vortices in Falkner–Skan–Cooke boundary layers. *Journal of Fluid Mechanics*, *368*, 339–357. https://doi.org/10.1017/S0022112098001931

Jacobs, R. G., & Durbin, P. A. (1998). Shear sheltering and the continuous spectrum of the Orr-Sommerfeld equation. *Physics of Fluids*, *10*. https://doi.org/10.1063/1.869716

Joseph, D. D. (1976). *Stability of Fluid Motions I* (1st ed.). Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-80991-0

Joslin, R. D., & Streett, C. L. (1994). The role of stationary cross-flow vortices in boundary-layer transition on swept wings. *Physics of Fluids*, *6*, 3442–3453. https://doi.org/10.1063/1.868401

Koch, W., Bertolotti, F. P., Stolte, A., & Hein, S. (2000). Nonlinear equilibrium solutions in a three-dimensional boundary layer and their secondary instability. *Journal of Fluid Mechanics*, *406*, 131–174. https://doi.org/10.1017/S0022112099007387

Kohama, Y., Saric, W., & Hoos, J. A. (1991). A high-frequency, secondary instability of crossflow vortices that leads to transition.

Lagarias, J. C., Reeds, J. A., Wright, M. H., & Wright, P. E. (1998). Convergence properties of the Nelder-Mead simplex method in low dimensions. *SIAM Journal on Optimization*, *9*(1), 112–147. https://doi.org/10.1137/S1052623496303470

Lighthill, M. J. (1958). On displacement thickness. *Journal of Fluid Mechanics*, *4*(4), 383–392. https://doi.org/10.1017/S0022112058000525

Malik, M. R., Li, F., & Chang, C. L. (1994). Crossflow disturbances in three-dimensional boundary layers: nonlinear development, wave interaction and secondary instability. *Journal of Fluid Mechanics*, *268*, 1–36. https://doi.org/10.1017/S0022112094001242

Malik, M. R., Li, F., Choudhari, M. M., & Chang, C. L. (1999). Secondary instability of crossflow vortices and swept-wing boundary-layer transition. *Journal of Fluid Mechanics*, *399*, 85–115. https://doi.org/10.1017/S0022112099006291

Meersman, J. A., Hader, C., & Fasel, H. F. (2018). Hypersonic boundary-layer transition: Comparison of the fundamental resonance breakdown for a flared and straight cone at mach 6. *2018 Fluid Dynamics Conference*. https://doi.org/10.2514/6.2018-3851

Ng, L. L., & Crouch, J. D. (1999). Roughness-induced receptivity to crossflow vortices on a swept wing. *Physics of Fluids*, *11*(2), 432. https://doi.org/10.1063/1.869860

Nishino, T., & Shariff, K. (2010). Direct Numerical Simulation of a Swept-Wing Boundary Layer with an Array of Discrete Roughness Elements. *Seventh IUTAM Symposium on Laminar-Turbulent Transition*, *18*, 289–294. https://doi.org/10.1007/978-90-481-3723-7{_}46

Poll, D. I. (1985). Some observations of the transition process on the windward face of a long yawed cylinder. *Journal of Fluid Mechanics*, *150*, 329–356. https://doi.org/10.1017/S0022112085000155

Pope, S. B. (2000). *Turbulent Flows* (1st ed.). Cambridge University Press. https://doi.org/10.1017/CBO9780511840531

Prandtl, L. (1904). Über Flüssigkeitsbewegung bei sehr kleiner Reibung. *Verhandlinger 3. Int. Math. Kongr. Heidelberg*, 484–491.

Radeztsky, R. H. J. (1994). *Growth and development of roughness-induced stationary crossflow vortices* (Doctoral dissertation). Arizona State University. Tempe. https://www.proquest.com/openview/787312d95428e342a94f3e658942759d/1?pq-origsite=gscholar&cbl=18750&diss=y

Reed, H. L., Saric, W. S., & Arnal, D. (1996). Linear Stability Theory Applied to Boundary Layers. *Annu. Rev. Fluid. Mech.*, *28*, 389–428. https://doi.org/10.1146/annurev.fl.28.010196.002133

Reibert, M., Saric, W., Carrillo, R. J., & Chapman, K. (1996). Experiments in nonlinear saturation of stationary crossflow vortices in a swept-wing boundary layer. *34th Aerospace Sciences Meeting and Exhibit*. https://doi.org/10.2514/6.1996-184

Rius Vidales, A. F. (2022). *Influence of a forward-facing step on crossflow instability and transition: An experimental study in a swept wing boundary-layer.* (Doctoral dissertation). Delft University of Technology. Delft.

Rius-Vidales, A. F., & Kotsonis, M. (2021). Impact of a forward-facing step on the development of crossflow instability. *Journal of Fluid Mechanics*, *924*. https://doi.org/10.1017/JFM.2021.497

Rius-Vidales, A. F., & Kotsonis, M. (2022). Unsteady interaction of crossflow instability with a forward-facing step. *Journal of Fluid Mechanics*, *939*. https://doi.org/10.1017/JFM.2022.151

Rizzetta, D. P., Visbal, M. R., Reed, H. L., & Saric, W. S. (2010). Direct Numerical Simulation of Discrete Roughness on a Swept-Wing Leading Edge. *AIAA Journal*, *48*(11), 2660–2673. https://doi.org/10.2514/1.J050548

Saric, W. S., Carpenter, A. L., & Reed, H. L. (2011). Passive control of transition in three-dimensional boundary layers, with emphasis on discrete roughness elements. *Trans. R. Soc. A*, *369*, 1352–1364. https://doi.org/10.1098/rsta.2010.0368

Saric, W. S., Carrillo, R. B., & Reibert, M. S. (1998). Leading-Edge Roughness as a Transition Control Mechanism. *36th AIAA Aerospace Sciences Meeting and Exhibit*. https://doi.org/10.2514/6.1998-781

Saric, W. S., Reed, H. L., & White, E. B. (2003). Stability and Transition of Three-Dimensional Boundary Layers. *Annu. Rev. Fluid Mech*, *35*, 413–440. https://doi.org/10.1146/annurev.fluid.35.101101.161045

Saric, W. S., West, D. E., Tufts, M. W., & Reed, H. L. (2015). Flight test experiments on discrete roughness element technology for laminar flow control. *53rd AIAA Aerospace Sciences Meeting*. https://doi.org/10.2514/6.2015-0539

Saric, W. S., West, D. E., Tufts, M. W., & Reed, H. L. (2019). Experiments on discrete roughness element technology for swept-wing laminar flow control. *AIAA Journal*, *57*(2), 641–654. https://doi.org/10.2514/1.J056897

Schlatter, P. (2000). Direct numerical simulation of laminar-turbulent transition in boundary layer subject to free-stream turbulence.

Schmid, P. J., & Henningson, D. S. (2001). Stability and Transition in Shear Flows. *142*. https://doi.org/10.1007/978-1-4613-0185-1

Serpieri, J. (2018). Cross-Flow Instability: Flow diagnostics and control of swept wing boundary layers. https://doi.org/10.4233/UUID:3DAC1E78-FCC3-437F-9579-048B74439F55

Serpieri, J., & Kotsonis, M. (2015). Design of a swept wing wind tunnel model for study of cross-flow instability. *33rd AIAA Applied Aerodynamics Conference*. https://doi.org/10.2514/6.2015-2576

Serpieri, J., & Kotsonis, M. (2016). Three-dimensional organisation of primary and secondary crossflow instability. *Journal of Fluid Mechanics*, *799*, 200–245. https://doi.org/10.1017/JFM.2016.379

Sforza, P. (2014). Drag Estimation. *Commercial Airplane Design Principles*, 349–404. https://doi.org/10.1016/B978-0-12-419953-8.00009-7

Szodruch, J. (1991). Viscous Drag Reduction On Transport Aircraft. *29th Aerospace Sciences Meeting*. https://doi.org/10.2514/6.1991-685

Takeda, K., & Kato, M. (1992). Wind tunnel blockage effects on drag coefficient and wind-induced vibration. *Journal of Wind Engineering and Industrial Aerodynamics*, *42*(1-3), 897–908. https://doi.org/10.1016/0167-6105(92)90096-S

Tempelmann, D., Schrader, L.-U., Hanifi, A., Brandt, L., & Henningson, D. S. (2011). Numerical study of boundary-layer receptivity on a swept wing. *6th AIAA Theoretical Fluid Mechanics Conference*. https://doi.org/10.2514/6.2011-3294

Trush, A., Pospíšil, S., & Kozmar, H. (2020). Comparison of turbulence integral length scale determination methods. *WIT Transactions on Engineering Sciences*, *128*, 113–123. https://doi.org/10.2495/AFM200111

Ustinov, M. V. (2013). Laminar-Turbulent Transition in Boundary Layers (Review). Part. 1: Main Types of Laminar-Turbulent Transition in a Swept-Wing Boundary Layer. *TsAGI Science Journal*, *44*(1), 1–63. https://doi.org/10.1615/TSAGISCIJ.2013007932

Van Ingen, J. L. (2008). The eN method for transition prediction. Historical review of work at TU Delft. *38th AIAA Fluid Dynamics Conference and Exhibit*. https://doi.org/10.2514/6.2008-3830

Wassermann, P., & Kloker, M. (2002). Mechanisms and passive control of crossflow-vortex-induced transition in a three-dimensional boundary layer. *Journal of Fluid Mechanics*, *456*, 49–84. https://doi.org/10.1017/S0022112001007418

Westerbeek, S. (2020). *Development of a Nonlinear Parabolized Stability Equation (NPSE) Analysis Tool for Spanwise Invariant Boundary Layers* (Doctoral dissertation). Delft University of Technology. Delft.

White, E. B., & Saric, W. S. (2005). Secondary instability of crossflow vortices. *Journal of Fluid Mechanics*, *525*, 275–308. https://doi.org/10.1017/S002211200400268X

A

Blowing & Suction

The most often used technique to introduce unsteadiness into the numerical flow field, other than free-stream turbulence, is blowing & suction. This is a forcing method, that prescribes disturbances of a set frequency, wavelength and amplitude at a fixed location on the wall. One can visually think of this method as a vibrating membrane at a set location on the wall, which vibrates at a certain frequency and amplitude with a preset wavelength. Since it has been the go-to method at TU Delft by work such as Casacuberta et al. (2021) and Casacuberta, Hickel, et al. (2022), this work also aims to compare the blowing & suction technique to the newer free-stream turbulence methodology. An assessment can so be made to which technique is more applicable to similar numerical work. Blowing & suction is commonly applied in combination with the numerical solver used at TU Delft, and therefore requires minimal implementation. What is left is to determine the inputs necessary for the method to work. This includes the surface section of the wall where the forcing is applied, and the different frequencies, wavelengths and amplitudes forced.

The surface section used is often modelled as a strip in spanwise direction across the entire domain. Forcing takes place near the inflow to allow the disturbances to develop naturally downstream. In this work, settings are used from a very similar numerical flow field, set up by Casacuberta, Groot, et al. (2022). Therefore the disturbance strip runs from $x/\delta_0 = 20$ to $x/\delta_0 = 32$, and across the full spanwise length of the domain. The disturbance is prescribed within this strip through the wall-normal velocity, calculated as

$$v\,(x, 0, z, t) = A_0 f_f\,(x) \cos\left(\beta z + \omega t\right), \tag{A.1}$$

where A_0 is the amplitude of the disturbances, β is a spanwise wavelength, and ω the desired angular frequency for the disturbance. The f-term ensures the disturbance strip blends smoothly into the rest of the surface, without any irregularities (Casacuberta, Groot, et al., 2022). f is set using purely the x-coordinates of the disturbance strip, and calculated as,

$$f_f\,(x) = \left(\frac{4\,(x - x_{\text{start}})\,(x_{\text{end}} - x)}{(x_{\text{end}} - x_{\text{start}})^2} \right)^3 . \tag{A.2}$$

Multiple spanwise wavenumbers and angular frequencies can be forced with this method, in which case the different v-profiles calculated with Equation A.1 are summed before being added to the numerical domain. For this simulation, the setup parameters of fellow researchers at TU Delft are replicated, where only the fundamental spanwise wavenumber 0.6604 is used to force multiple angular frequencies ω. These frequencies range from $3kHz$ to $14kHz$, in increments of $1kHz$. Simply adding this v-component to the domain at the predefined points results in the unsteady numerical simulation. Although an easy and repeatable method to implement, it is time and cost-intensive to set up independently. If wavenumbers and frequencies to implement need to be chosen without prior research to base settings on, a trial-and-error approach would likely be needed. Inflicting a considerable amount of computational cost, the free-stream turbulence method would therefore be a more efficient method if its working is proven in this research. It is arguably also a less natural way of introducing disturbances into the boundary layer. Instead of receptivity to the free-stream disturbances, blowing & suction enforces

disturbances from the wall. This could potentially trigger different transition mechanisms than experimental flows, and also the discrete frequencies being forced do not help with matching experimental data.

B

Additional DNS Unsteady Flow Field Isosurfaces

B.1. Q-Criterion Isosurface: Coloured by Wall-Normal Velocity

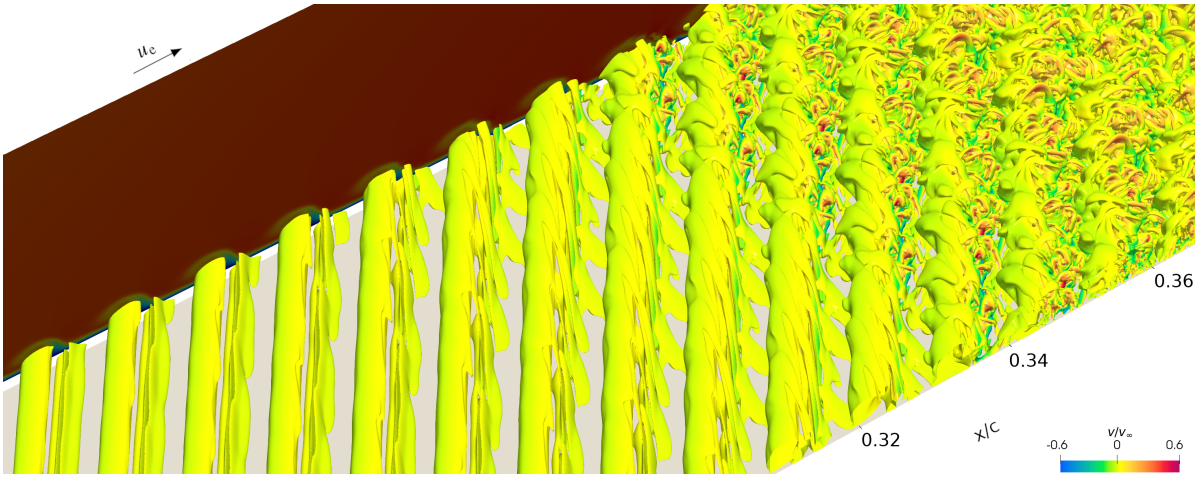

Figure B.1: Detailed snapshot of the unsteady DNS instantaneous Q-criterion isosurface coloured by v/u_∞, along with y-z and y-x u/u_∞ slices. Domain duplicated four times in z-direction based on periodic boundary conditions.

B.2. Q-Criterion Isosurface: Coloured by Spanwise Velocity

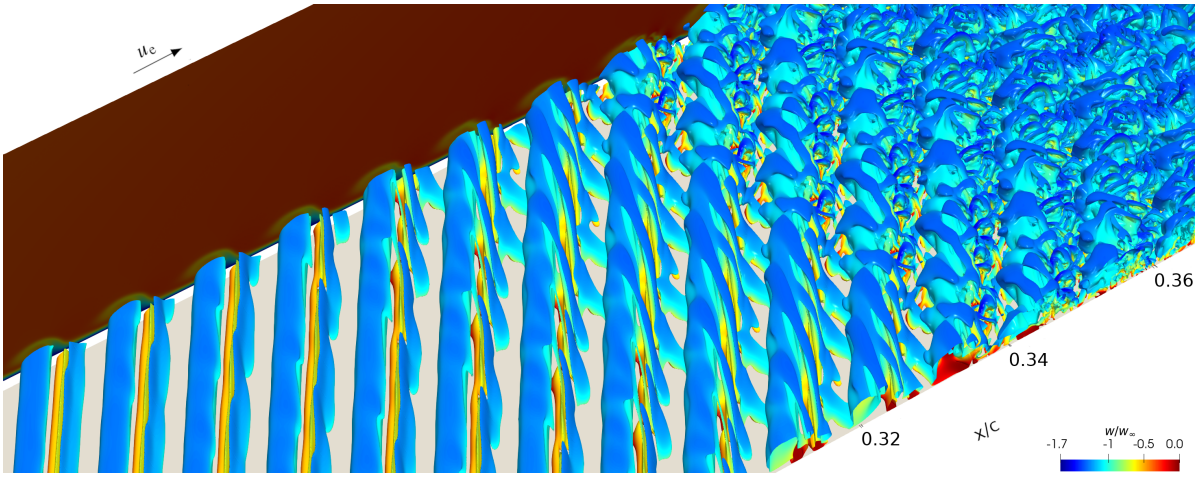

Figure B.2: Detailed snapshot of the unsteady DNS instantaneous Q-criterion isosurface coloured by w/u_∞, along with y-z and y-x u/u_∞ slices. Domain duplicated four times in z-direction based on periodic boundary conditions.