



Delft University of Technology

Document Version

Final published version

Citation (APA)

Westerbeek, S. H. J. (2026). *Nonlinear Stability of Boundary Layers over Swept Wings with Sharp and Smooth Surface Features*. [Dissertation (TU Delft), Delft University of Technology]. <https://doi.org/10.4233/uuid:1684bdc4-4e42-4fa8-8931-9b51a78bfdba>

Important note

To cite this publication, please use the final published version (if applicable).
Please check the document version above.

Copyright

In case the licence states "Dutch Copyright Act (Article 25fa)", this publication was made available Green Open Access via the TU Delft Institutional Repository pursuant to Dutch Copyright Act (Article 25fa, the Taverne amendment). This provision does not affect copyright ownership.
Unless copyright is transferred by contract or statute, it remains with the copyright holder.

Sharing and reuse


Other than for strictly personal use, it is not permitted to download, forward or distribute the text or part of it, without the consent of the author(s) and/or copyright holder(s), unless the work is under an open content license such as Creative Commons.

Takedown policy

Please contact us and provide details if you believe this document breaches copyrights.
We will remove access to the work immediately and investigate your claim.

This work is downloaded from Delft University of Technology.

Nonlinear Stability of Boundary Layers over Swept Wings with Sharp and Smooth Surface Features



Sven Westerbeek

NONLINEAR STABILITY OF BOUNDARY LAYERS OVER SWEPT WINGS WITH SHARP AND SMOOTH SURFACE FEATURES

Proefschrift

ter verkrijging van de graad van doctor
aan de Technische Universiteit Delft,
op gezag van de Rector Magnificus,
Prof.dr.ir. H. Bijl,
voorzitter van het College voor Promoties,
in het openbaar te verdedigen op maandag 18 mei 2026 om 17:30 uur

door

Sven Hendrikus Johannes WESTERBEEK

Dit proefschrift is goedgekeurd door de (co)promotoren.

Samenstelling promotiecommissie:

Rector Magnificus,	chairperson
Prof. dr.	M. Kotsonis, promotor
Prof. dr.	H.M. Schuttelaars, promotor
Dr.	S.J. Hulshoff, copromotor

Onafhankelijke leden:

Prof. dr. ir. C. Vuik	Technische Universiteit Delft
Dr. A. Hanifi	KTH Royal institute of Technology, Zweden
Dr.-ing. M. Kloker	University of Stuttgart, Duitsland
Dr. G. Rigas	Imperial College London, Engeland
Prof. dr. -ing. habil. S. Hickel	Technische Universiteit Delft, <i>reservelid</i>



Keywords: Nonlinear Stability Analysis, Harmonic Navier-Stokes Equations, Non-linear Parabolized Stability Equations, Crossflow Instabilities

Printed by: Proefschriftmaken.nl

Front & Back: Schematic visualisation of a smooth hump on a swept wing

Financially supported by: European Research Council (ERC) Starting Grant #803082 "GloWing"

Copyright © 2025 by S. Westerbeek

ISBN 978-94-6384-948-7

An electronic version of this dissertation is available at
<http://repository.tudelft.nl/>.

*Human nature is both the source of our problems and the key to our solutions.
We must learn to live in balance with nature, or we will destroy ourselves.*

Jane Goodall

CONTENTS

Summary	ix
Samenvatting	xiii
Preface	xvii
1 Introduction	1
1.1 Motivation	2
1.2 Boundary Layer Flows	3
1.3 Swept Wings	3
1.4 Laminar Flow Control	5
1.5 Flow Stability Analysis	7
1.6 Objectives and Thesis Outline	8
1.7 Outline	9
2 Background of Flow Stability Analysis	11
2.1 Introduction	12
2.2 The Concept of Flow Stability	12
2.3 Laminar-Turbulent Transition	13
2.4 Swept Wings and Crossflow Instabilities	16
2.5 Flow Stability Analysis Frameworks	17
2.5.1 Domain Description	18
2.5.2 Reference and Non-Dimensional Quantities	18
2.5.3 Navier-Stokes Equations	19
2.5.4 Perturbation Analysis	19
2.5.5 Orr-Sommerfeld Equation	19
2.5.6 Parabolized Stability Equations	20
2.5.7 Harmonic Navier-Stokes Equations	21
2.5.8 Harmonic Balancing	23
2.5.9 Summary of Flow Stability Analysis Equations	24
3 Development of a Harmonic Navier-Stokes Solver	27
3.1 Introduction	28
3.2 Description of Example Problem Domains	28
3.3 Implementation	29
3.3.1 Discretisation	30
3.3.2 Memory Management and Construction of the Solution Matrix	31
3.3.3 Outflow Treatment	32
3.3.4 Embedded Boundary Method	33
3.3.5 Iterative Convergence of the Nonlinear Terms	34

3.4	Verification	37
3.5	Validation	40
3.5.1	Nonlinear Development of Two-dimensional TS Waves	40
3.5.2	Nonlinear Development of Stationary Crossflow Instabilities	42
3.5.3	Linear Development of Stationary Crossflow Instability over a Hump 44	
3.5.4	Linear Development of Stationary Crossflow Instability over a Forward- Facing Step.	47
3.6	Concluding Remarks	49
4	Receptivity of Crossflow Instability to Discrete Roughness Amplitude and Lo- cation (Experiments by G. Zoppini)	51
4.1	Introduction	52
4.1.1	Domain and Flow Description	52
4.1.2	Methodology.	54
4.1.3	Results	57
4.1.4	Conclusion.	59
5	Applications of Flow Stability Frameworks in Flows over Forward-Facing Steps (DNS by J. Casacuberta)	61
5.1	Interaction of Crossflow Instabilities with Forward-Facing Steps (DNS by J. Casacuberta)	62
5.1.1	Domain and Flow Description	63
5.1.2	Methodology.	64
5.1.3	Results	67
5.1.4	Conclusion.	71
5.2	Streaky Perturbations in Laminar Forward-Facing-Step Flow	72
5.2.1	Domain Flow and Description	73
5.2.2	Discretization	74
5.2.3	Results	75
5.2.4	Conclusion.	78
6	Interactions Between Stationary CFI and a Smooth Surface Hump	81
6.1	Introduction	82
6.1.1	Interaction Between Boundary Layer Instabilities and Surface Im- perfections.	83
6.1.2	Smooth surface humps as a laminar flow control method	85
6.1.3	Aim and Overview of this Chapter	86
6.2	Methodology	86
6.2.1	Flow Conditions and Domain Geometry	87
6.2.2	Base Flow Calculation	89
6.2.3	Flow Stability Simulation.	90
6.2.4	Energy-Balance Equations of the Stationary Spanwise-Harmonic Perturbations	92
6.2.5	Perturbation Advection Equations	94

6.3	Topology of the Base Flow.	94
6.3.1	Effect of the Hump on the Base Flow Boundary Layer	94
6.3.2	Crossflow Reversal	96
6.4	Evolution of the perturbation field	97
6.4.1	Effect of the Hump on Linear Perturbation Development	97
6.4.2	Local and Non-Local Linear Stability Analysis	100
6.4.3	Nonlinear effects in CFI-hump interactions	104
6.4.4	Energy Transfer Mechanisms in the Interaction Between Hump and CFI.	107
6.5	Concluding Discussion	111
7	Conclusion	115
7.1	Research aims.	115
7.2	Impact	116
7.3	Future Work.	117
7.3.1	Laminar Flow Control	117
7.3.2	Stability Frameworks.	118
	Epilogue	121
	Acknowledgements	123
A	Generalized Curvilinear Coordinate Transformation	135
A.1	Direct Derivative Formulation	135
A.2	Jacobian Formulation.	136
A.3	NPSE Matrix Contents	139
A.4	Left-Hand Side Matrix.	142
A.5	Performance	143
	Curriculum Vitæ	145
	List of Publications	147

SUMMARY

The global demand for air travel has grown strongly due to an increase in the global population and increased global welfare. The only exception to this staggering growth was caused by the COVID-19 pandemic between 2020 and 2024, limiting global mobility. Air travel comes with a large range of benefits such as speed, range, and convenience over other forms of transport, but comes at a great cost to the environment. Hence, a global push for a greener aviation industry is made via initiatives like the European Green Deal, which aims to significantly reduce greenhouse gas emissions. From an aerodynamic point of view, this largely comes down to improving the lift-to-drag ratio of aircraft. The latter, aircraft drag, is significantly increased by the transition to turbulence of the boundary layer. Two important ways to reduce the total drag, therefore, involve the delay of the onset of boundary layer turbulence and the reduction of turbulent boundary layer drag after transition. The former lies at the core of this thesis and requires in-depth knowledge of the boundary layer transition process.

Boundary layer transition is a complex process that is commonly studied via experiments, numerical stability analysis, or Direct Numerical Simulations (DNS). Although experiments and DNS can provide important insights into the flow, they are time-consuming and expensive. As an alternative, flow stability analysis, be it linear or nonlinear, can be used to study transition at a fraction of the cost. In flow stability analysis, the Navier-Stokes equations are subjected to a perturbation analysis where the flow is decomposed into a steady solution, the base flow, and perturbations. The perturbations are assumed to be periodic in both time and the spanwise direction. A limited set of perturbations contain a large part of the perturbation kinetic energy and dominate the transition process. Hence, a spectral cut-off is introduced above which perturbations are not accounted for. These assumptions reduce the problem to at most two dimensions with a limited spectral scope. The result is higher computational efficiency at the cost of applicability. Namely, flow stability analysis frameworks cannot account for variations of the wing in the spanwise direction, including three-dimensional surface features such as rivets. Moreover, the turbulent boundary layer state cannot be resolved as a result of the spectral cut-off in the description of the perturbation field.

Classical stability analysis makes use of the Orr-Sommerfeld (OS) equation, which are used to define a purely local eigenvalue problem. Although this method is still widely used today, it suffers strongly from its assumptions. The boundary layer is assumed parallel, i.e. the wall-normal velocity is zero and streamwise gradients are neglected. The history of the perturbation is also not accounted for. Furthermore, the OS equation does not account for nonlinearity, hence perturbation amplitudes of nonlinearly interacting wave packets cannot be tracked. However, the linear amplitude development of perturbations can be reconstructed from the calculated growth rates. As a consequence, the

OS method is strictly linear and thus does not allow for mode interactions fundamental to the analysis of perturbations at finite amplitudes. Although the methods based on the OS equation provide accurate perturbation shapes and growth rates locally when the boundary layer flow is nearly parallel and perturbation amplitudes are small, their accuracy deteriorates rapidly in other cases. Moreover, even when used for nearly parallel boundary layers, its errors in the predicted amplitude can accumulate rapidly in the streamwise direction.

An improvement over methods based on the OS equation can be found in the Parabolized Stability Equations (PSE), which maintain the first-order streamwise derivatives of the Navier-Stokes equations during their derivation. This allows the PSE to keep track of perturbation shapes and amplitudes, and as a result, account for mode interactions if nonlinear terms are maintained in the derivation. The PSE, however, still neglect higher-order streamwise derivatives. This results from parabolising the equations to enable a marching solution technique. While flow stability algorithms based on the PSE can be very efficient and maintain accuracy even for large-amplitude perturbations, their application is limited to what is commonly referred to as *slowly-varying* flows. It can be difficult to determine a priori whether PSE can be applied to a specific flow scenario. The algorithm can provide results when the neglected higher-order streamwise derivatives are important. Consequently, their application to cases with (locally) strongly non-parallel flows comes with uncertainties.

The last stability analysis framework considered in this thesis avoids these uncertainties at the cost of computational speed. The Harmonic Navier-Stokes (HNS) equations fully maintain all streamwise derivatives. Their application is therefore not limited to slowly-varying flows. However, they cannot be solved via a marching solution technique. Instead, the solution of the entire two-dimensional flow field is required. Although this is considerably more computationally expensive than solving the PSE, it is still orders of magnitude more computationally efficient than performing full DNS, while offering the same accuracy in predicting the development of primary instabilities in spanwise-invariant flows.

The result of all of the above flow stability analyses is the amplitude evolution of a pre-determined set of perturbations. However, the transition to turbulence itself is not simulated due to the use of a spectral cut-off to the benefit of computational efficiency. The method is hence unable to capture turbulent breakdown. The location where the flow transitions to turbulence is then generally inferred from a calibration with experiments. From this calibration, the total growth of primary instabilities, such as the Crossflow Instabilities (CFI), necessary to trigger the transition to turbulence is found. Once calibrated, studies of the effect of surface features on the transition location can quickly be made. Here, this approach is chosen to study the effect of a smooth surface hump on the stationary crossflow instabilities responsible for the transition of swept-wing boundary layers.

To arrive at an HNS framework, the generalised incompressible Navier-Stokes (NS)

equations are expanded in perturbed form. Spanwise and temporal periodicity of the perturbations is assumed, allowing for a Fourier ansatz in the spanwise direction and time. The resulting equations are discretised via spectral collocation in the wall-normal direction and finite-difference methods in the streamwise direction. The equations are then solved using a direct sparse-matrix solver. The nonlinear mode interaction terms, constructed via harmonic mode balancing, are converged iteratively. The solution implementation employs a generalised domain transformation to account for geometrically smooth surface features, such as humps. No-slip conditions are imposed in the interior of the domain to account for the presence of sharp surface features such as forward- or backwards-facing steps.

In this thesis, the Delft Harmonic Navier-Stokes Solver (DeHNSSo) is first developed. It is capable of accounting for both smooth and sharp surface deformations in the stability analysis of spanwise-invariant boundary layers. Second, this framework is used to assess the effects of surface deformations such as humps and steps on incoming perturbations in swept-wing boundary layers.

The developed framework is subjected to a verification and validation analysis through several canonical cases of boundary layer instability growth, including linear and non-linear growth of Tollmien-Schlichting waves in a Blasius boundary layer and the development of stationary crossflow instabilities in a swept flat-plate boundary layer. The latter problem is also treated in the presence of a geometrical smooth hump and a sharp forward-facing step at the wall. In all cases, a close match is observed with DNS.

Apart from these validation studies, HNS simulation results, such as perturbation amplitudes, growth rates, and shape functions, are compared to benchmark results from flow stability analysis methods such as Parabolized Stability Equations (PSE) in slowly-varying flows, and Adaptive Harmonic Linearised Navier-Stokes (AHLNS) otherwise. Good agreement is observed in all cases. Furthermore, studies of convergence, memory usage, and computational cost are performed. The computational cost is, in general, considerably lower than high-fidelity DNS at comparable grid resolutions.

DeHNSSo is then used to evaluate the effect of a smooth surface hump on the development of crossflow instabilities in a swept-wing boundary layer. The hump is shallow, and the base flow is shown to feature a region of crossflow reversal but crucially lacks a separation region. The hump has a negligible effect on the boundary layer away from it and boundary layer specifications quickly recover to the reference no-hump case, in less than one width of the hump itself downstream.

Linear flow stability simulations, i.e. infinitesimal perturbation amplitudes, show regions of inhibited and accelerated growth, with total perturbation kinetic energy being recovered shortly downstream of the hump. Then, reduced growth is observed over a large streamwise extent. Despite similarities in the base flow from that point onward, a distorted perturbation profile was shown to be responsible for the inhibited growth. The amplitude difference with the reference case is maintained until the outflow.

Increasing the inflow amplitude beyond the linear range results in a qualitatively similar response, however, the local destabilization at the leeward side of the hump is amplified resulting in locally elevated perturbation kinetic energy. For sufficiently large inflow amplitudes, this level surpasses even the reference case without a hump. Higher harmonics respond similarly, and linear mechanisms (e.g. production and lift-up) are found to be responsible for it. However, the relative increase is independent of inflow amplitude and reached values several times larger than observed in the reference case. For large inflow amplitudes, this causes quasi-saturation of amplitudes slightly downstream of the hump, much earlier than the reference no-hump case.

An energy budget analysis proves that primary mode growth is inhibited through sub-optimal growth caused by the reduced effectiveness of the lift-up mechanism. This is preceded by a perturbation shape deformation that is found to be governed by (spanwise) transport terms in the crossflow reversal region. As the inflow amplitude increased, a new unstable mode growing on a second branch of inflection points near the wall is found to provide locally positive energy transfer to the fundamental mode. The wall-normal and spanwise shears in the total flow increase significantly as a result of the hump, possibly destabilising secondary instabilities.

This finding strongly motivates the hypothesis that the stabilization of incoming stationary CFI via smooth surface protuberances is most effective at low amplitudes when flow mechanisms can be approximated by linear methods. The main reason for this is the inability of a second unstable mode pre-conditioned by the incoming CFI to reach amplitudes large enough to saturate and kickstart the transition process early.

The current research paves the way to further optimising the hump shape, size, and configuration, including explorations of asymmetric shapes and configurations with multiple humps. Efficient computational tools such as DeHNSSo, could be used to perform this optimisation for any wing shape, allowing for widespread use in real-world applications. Both new and existing aircraft could be outfitted with surface humps to delay the onset of boundary layer transition.

Regarding the developed stability frameworks, improvements could be made to address challenges in computational efficiency and robustness, particularly in handling large amplitudes. Coupling DeHNSSo with an NPSE (Nonlinear Parabolized Stability Equations) code in series could lead to significant improvements in computational efficiency by leveraging the advantages of both methods, provided limitations in inflow treatment and nonlinear continuation algorithms are managed effectively.

SAMENVATTING

De wereldwijde vraag naar luchtvaart is alsmāar toenemend door een toename in de wereldbevolking en globale welvaart. The enige uitzondering op deze ongekeende groei werd veroorzaakt door de Coronapandemie die tijdelijk deze groei stillegde tussen 2020 en 2024. Luchtvaart gaat gepaard met voordelen in snelheid, bereik en gemak ten opzichte van andere vormen van transport, maar tevens met grote kosten voor het milieu. Daarom is er een globale inzet om de luchtvaartindustrie te vergroenen.

Deze vergroening van de luchtvaartindustrie van een aerodynamisch opzich komt neer op het verlagen van de weerstand. De weerstand van vliegtuigen wordt significant beïnvloed door de transitie van de grenslaag van een laminaire naar een turbulente staat. Twee belangrijke manieren om de weerstand te verlagen hebben daarom te maken met het verlaten van de start van turbulentie en het verminderen van de weerstand die gegenereerd wordt door de turbulente grenslaag na transitie. Die eerste is de kern van dit proefschrift en vereist kennis van het transitieproces. Het bestuderen van grenslaagtransitie wordt vaak gedaan via experimenten, numerieke stabiliteitsanalyses, of Directe Numerieke Simulaties (DNS). Experimenten en DNS kunnen een erg accuraat en gedetailleerde beschrijving van de stroming geven, maar zijn vaak tijdrovend en duur. Stabiliteitsanalyse, lineair of niet-lineair, is een wiskunde methode die helpt in het voorspellen van transitie onderheven aan aannames over de stroming. Deze aannames zorgen voor een toename van de efficiëntie ten koste van de algemene toepasbaarheid van de methode. Daarbij kan deze methode door een afkapping van de spectrale inhoud van de simulatie de turbulente stroming niet representeren. De kosten in de algemene toepasbaarheid zijn onder andere dat de methode geen rekening kan houden met variaties in de richting van de spanwijdte of het representeren van 3-dimensionale objecten op de vleugel zoals bouten en moeren.

Klassieke stabiliteitsanalyses maakten gebruik van de Orr-Sommerfeld vergelijking, wat een puur lokaal eigenwaardeprobleem is. Hoewel deze methode hedendaags nog veel gebruikt wordt, lijkt het erg van de aanname van lokaliteit van het stabiliteitsprobleem. De grenslaag wordt aangenomen parallel te zijn. Met andere woorden, de verticale snelheid wordt nul verondersteld en er wordt geen rekening gehouden met de afgeleiden in de stromingsrichting. Daarnaast is de method, mede door deze aannames, een lineaire methode en kan het geen rekening houden met de interacties van verschillende verstoringsmodes. Hoewel de OS-vergelijking accuraat de vorm en het groeigetel van de verstoring kan voorspellen wanneer de grenslaag inderdaad bijna parallel is, wordt dit al snel minder bij niet-parallele grenslagen. Daarnaast accumuleren de fouten zelfs in bijna-parallele grenslagen snel bij herhaaldelijk gebruik in de richting van de stroming.

Een verbetering ten opzichte van de puur lokale aanpak neemt de vorm aan van de

geparaboliseerde stabiliteitsvergelijkingen (PSE), die de eerste orde afgeleiden in de stromingsrichting van de Navier-Stokes vergelijkingen niet verwaarloosd. Dit zorgt ervoor dat de PSE het verloop van de vorm en amplitude van de verstoring kan bijhouden en, als gevolg daarvan, rekening kan houden met de interacties van verstoringsmodes zolang gekozen is de niet-lineaire termen te behouden. Echter zijn de hogere orde afgeleiden wel verwaarloosd wat ervoor zorgt dat de methode bijna parabolisch is en de oplossing sequentieel gevonden kan worden in de stromingsrichting. Hoewel stromingsstabiliteits algoritmes gebaseerd op de PSE erg snel zijn, zorgen de aannames ervoor dat deze enkel toegepast kunnen worden op zogenaamde *langzaam veranderende stromingen*. Een exacte definitie van een langzaam veranderende stroming is niet beschikbaar. Bovendien is het algoritme robuster dan dat het accuraat is, wat leidt tot oplossingen waarbij de verwaarloosde hogere orde afgeleiden wel belangrijk waren geweest. Als een gevolg hiervan gaat de toepassing bij grenslagen met (lokaal) sterk veranderende stromingen gepaard met veel onzekerheid.

De laatste stabiliteitsanalysemethode die in dit proefschrift behandeld wordt vermijdt deze onzekerheid volledig ten koste van de benodigde rekenkracht. De Harmonische Navier-Stokes (HNS) vergelijkingen behouden alle afgeleiden in de stromingsrichting. De toepassing is daardoor niet gelimiteerd tot langzaam veranderende stromingen, maar kan daarentegen niet sequentieel opgelost worden. In plaats daarvan dient het gehele tweedimensionale stromingsveld in een keer opgelost te worden. Dit is vele malen meer intensief dan PSE, maar nog steeds veel goedkoper dan DNS terwijl het voor de primaire verstoringen dezelfde precisie op kan leveren in het antwoord.

In dit proefschrift is het bovenstaande gruikt om bij te dragen aan het vergroenen van de luchtvaartsector op twee manieren. Allereerst is er een niet-lineair harmonisch Navier-Stokes framework genaamd DeHNSSo opgezet dat met gladde en niet-gladde oppervlakteverstoringen rekening kan houden in de stabiliteitsanalyse van grenslagen. Daarnaast is dit kader gebruikt om de effecten, zowel positief als negatief, oppervlakteverstoringen als drempels en randen in grenslagen over gepijlde vleugels.

Het resultaat van een stabiliteitsanalyse van een stroming is de evolutie van de amplitude van een van te voren bepaalde set van verstoringen. Echter, de transitie naar een turbulente stroming wordt buiten beschouwing gelaten wat ten goede komt van de snelheid van deze simulaties. De transitielocatie wordt dan over het algemeen afgeleid van een calibratie met experimenten. Van deze calibratie wordt de totale groei, gecondenseerd tot een N-factor, die nodig was om de transitie te triggeren gevonden. Eenmaal gecalibreerd, kan de geometrie eenvoudig aangepast worden en kan dus het effect van oppervlakteverstoringen bestudeerd worden. Deze methode is daarom gekozen om het effect van een gladde drempel op de dwarsstrooinstabiliteiten in de grenslaag van gepijlde vleugels te bestuderen.

Een niet-lineair harmonisch Navier-Stokes kader is ontwikkel in dit proefschrift om de evolutie van instabiliteiten in laminaire grenslagen zonder variaties in de richting van de spanwijdte te simuleren. Dit kader verwaarloost geen afgeleiden in de stromingsrich-

ting en kan dus gebruikt worden in stromingen over vleugeloppervlakken met gladde en scherpe verstoringen.

Het ontwikkelde kader is onderhevig aan een verificatie- en validatieanalyse door middel van enkele canonische scenarios van grenslag instabiliteitsgroei, inclusief lineaire en niet-lineaire groei van Tollmien-Schlichting golven in een Blasius grenslaag en de evolutie van stationaire dwarsstroom instabiliteiten in een grenslag over gepijlde vleugels. Wat betreft die laatste, deze is daarnaast behandeld onder de aanwezigheid van een gladde drempel en een scherpe rand op het oppervlak. De resultaten van HNS zoals de verstoringamplitudes, groeisnelheden en vormfuncties vergeleken met benchmark resultaten van andere stromingsstabiliteitsanalyse methoden zoals de Geparaboliseerde StabiliteitsVergelijkingen (PSE), Adaptieve Harmonische Gelineariseerde Navier-Stokes Vergelijkingen (AHLNS) of Directe Numerieke Simulaties (DNS). Een goede overeenkomst is geobserveerd in alle gevallen. Daarnaast is het kader onderhevig aan een convergentieanalyse en een eenvoudige prestatiebenchmark (geheugengebruik en vereiste rekenkracht). De vereiste rekenkracht is over het algemeen significant lager dan hoogwaardige DNS bij gelijke fijnmazigheid.

De DeHNSSo code is vervolgens gebruikt om het effect van een gladde drempel op de ontwikkeling van dwarsstroominstabiliteiten in de grenslag over een gepijlde vleugel. De drempel is oppervlakkig en de basisstroming bevat daardoor een regio waarin de dwarsstroom van richting verandert, maar bevat geen separatieregio waar in de stroming lokaal stroomopwaarts (ten opzichte van de externe stroomrichting) gaat. Het effect van de drempel op de basisstroming is verwaarloosbaar ver weg van de drempel en de grenslaagwaarden zoals het snelheidsprofiel en de hoogte keren snel terug naar de referentiegrenslaag zonder drempel. Dit gebeurt binnen een afstand van een keer de breedte van de drempel stroomafwaarts.

De lineaire simulaties van de stromingsstabiliteit laat regio's zien van tegengehouden en versnelde groei waarbij de totale kinetische verstoringsenergie overeenkomt met de referentie kort stroomafwaarts van de drempel. Vervolgens is een lange regio te zien waarbij de groei lager is dan de referentiewaarde. Dit gebeurt ondanks een vergelijkbare basisstroming. Een verandering van het verstoringprofiel is verantwoordelijk voor deze lagere groei. Het verschil in amplitude dat vervolgens ontstaat houdt aan tot en met de uitstroom van het domein.

Bij toenemende instroomamplitude wordt een kwalitatief gelijke respons geobserveerd. Echter is een sterkere destabilisatie aan de weerszijde van de drempel te zien die resulteert in een lokaal toegenomen kinetische energie van de verstoringen vergeleken met de referentie zonder drempel. Hogere harmonische trillingen reageerden in een vergelijkbare manier en lineaire mechanismen (productie en lift-up) blijken hier verantwoordelijk voor. De relatieve toename van amplitude is daardoor tot op zekere hoogte onafhankelijk van de amplitude zelf wat leidt tot amplitudes vele malen hogere dan de referentie. Voor hoge instroomamplitudes zorgt dit voor quasi-verzadiging van de verstoringen vlak na de drempel. Dit is veel eerder, ofwel dichter bij de voorrand van de

vleugel, dan te zien in de referentie.

Een analyse van de energiebudgetten bewijst dat de groei van de primaire verstoring tegengehouden wordt doormiddel van een suboptimaal groeiproces veroorzaakt door een verminderde effectiviteit van het lift-up mechanisme. Dit fenomeen wordt veroorzaakt door een verandering van het verstoringprofiel door transport termen in de richting van de spanwijdte terplekke van de ommekeer van de dwarsstroomrichting. Bij hogere instroomamplitudes zorgt een nieuwe instabiele mode die groeit op een tweede tak inflectiepunten voor lokaal positieve overdracht van energie naar de primaire verstoring. De afschuivingen in de totale stroming loodrecht op de wand en in de richting van de spanwijdte nemen significant toe ten gevolge van de drempel die mogelijk secundaire instabiliteiten kunnen destabiliseren.

Deze bevindingen motiveren de hypothese dat de stabilisering van inkomende dwarsstroominstabiliteiten doormiddel van gladde verstoringen van het vleugel oppervlak het meest effectief zijn bij lage verstoringamplitudes en lineaire methodes de evolutie ervan goed kunnen benaderen. De belangrijkste reden hiervoor is het onvermogen van een tweede instabiele mode, die geconditioneerd wordt door de inkomende dwarsstroominstabiliteit, om amplitudes te halen die kunnen leiden tot het vroeg verzadigen in amplitude.

Het huidige onderzoek laat de potentie zien om de drempelvorm, grootte en configuratie te optimaliseren. Hierbij dienen asymetrische vormen en configuraties van meerdere drempels te worden inbegrepen. De rekenkundige programma's kunnen worden verbeterd om deze optimalisatie efficiënter en minder handmatig uit te voeren voor elke vleugelvorm. Hiermee kan de weg vrij worden gemaakt voor toepassingen op echte vleugels. Zowel nieuwe als bestaande vliegtuigen kunnen worden verbeterd met een gladde opperlakte verstoring en zo grenslaagtransitie op de vleugel uit te stellen.

Met betrekking tot de rekenprogramma's zijn er met name mogelijkheden tot het verbeteren van de efficiëntie (doormiddel van iteratieve oplosmethoden) en robuustheid, met name bij het doorrekenen van verstoring met grote amplitudes. Een koppeling van DeHNSSo met een NPSE kader in serie waarbij de sterktes van beide methodes worden gebruikt kan een significante verbetering opleveren in termen van efficiëntie. Hierbij dient wel met zorg gekeken te worden naar de behandeling van de in- en uitstroomkoppeling en de niet-lineaire continuering in algoritmes.

PREFACE

The thesis before you is the result of five years of research at the Flow Control research group at TU Delft. A deeply motivated, fun and above all intelligent group of people that all contribute to this unique learning environment. Most unique perhaps, although I cannot possibly compare it to any other research group, is the three-pillared approach that I see here over and over: experiments, simulations and theory. It is so powerful to look at flow phenomena from all three angles and this aids massively in the development of new insight as well as intuition.

A considerable part of my life on this planet is spent on this booklet and the articles that support it. Although its contents are, hopefully, interesting, the skills that I have had to develop to arrive at these insights are ultimately what I take along with me. In addition, and perhaps more importantly, it is the confirmation that this is what I want to do professionally: researching and applying numerical methods to investigate flow phenomena. It was with utmost joy and passion that I researched the topic of boundary layer stability and control, numerical discretisation and solution techniques, and the development of code. I hope this document conveys this as well as the scientific content.

*Sven Hendrikus Johannes WESTERBEEK
Delft, 2025*

1

INTRODUCTION

*If learning the truth is the scientist's goal, then he must make himself the enemy of all
that he reads*

ibn al-Haytham

The expected growth of the aviation sector, in conjunction with the need to decrease its emissions, calls for a drastic improvement in aircraft aerodynamic efficiency, i.e., drag reduction. An important contribution to the aircraft's drag is determined by the velocity shear near the wall. This contribution is vastly different between laminar and turbulent states of the boundary layer; the latter produces far more friction drag. Laminar flow control, the concept of maintaining a laminar boundary layer, is one of the most promising technologies for achieving drag reduction. This chapter will first cover some of the history of wing design leading to the modern swept wings seen on most commercial aircraft. Then, the concepts of boundary layers and transition to turbulence are introduced. Finally, the methods by which the process of transition is analysed and controlled are covered.




Figure 1.1: Crossflow boundary layer transition process from an initially laminar state via eigenmode growth of Crossflow Instabilities (CFI) to turbulence.

1.1. MOTIVATION

The number of European flight passengers surpassed pre-pandemic numbers in 2024 according to the latest ACI Airport Traffic Report ¹. Between the financial crisis of 2008 and 2018, global carbon emissions from aviation have increased by 4-5% per year ². However, to be able to meet the Paris Agreement target, emissions should decrease to 50% of 2005 levels by 2050. The Paris Agreement was signed in 2016 and aims to keep the increase in the global average temperature to well below 2°C above pre-industrial levels and limit the temperature increase to 1.5°C above pre-industrial levels. To reach these goals, the aviation sector should change drastically. This feat can only be achieved if the aircraft's performance is optimised along every possible vector. The use of some of the newest technologies might only be possible if optimisation along other vectors is achieved first. For example, alternative power sources such as batteries or hydrogen come with significant downsides (e.g., low energy density and specific energy) and engineering challenges. An increase in aircraft efficiency lowers the total energy demand, possibly enabling the use of green(er) energy sources.

The aerodynamic performance of an aircraft can be quantified as the lift-to-drag ratio (L/D). The lift must equal the aircraft's weight. Assuming the weight stays the same, and thus the lift does not change, performance improvements are achieved by reducing the drag. Any object moving with a relative velocity to the (viscous) fluid it is submerged in experiences drag. The total drag of an aircraft is made up for approximately 50% by friction drag (Beguet et al., 2017; Schrauf, 2005). The frictional drag, in turn, depends largely on the state of the boundary layer: laminar or turbulent. In fact, Beck et al., 2018 states that a turbulent boundary layer results in viscous drag five to ten times larger than its laminar counterpart. However, the boundary layer state depends on a wide range of factors; the study of these factors and the methods by which their effect can be predicted lies at the core of this thesis.

¹ACI EUROPE publishes a monthly Airport Traffic Report, compiling data provided from more than 450 airport members of ACI EUROPE, representing over 95% of European airport traffic, see <https://www.aci-europe.org/44-industry-data/40-airport-traffic.html>

²Data from the World Economic Forum (WEF), see <https://www.weforum.org/agenda/2022/12/aviation-net-zero-emissions/>

1.2. BOUNDARY LAYER FLOWS

Prandtl, 1905 first described the existence of a shear layer near a submerged surface moving with a relative speed to the fluid and provided a mathematical model for its behaviour. This shear layer, called the Boundary Layer (BL), is defined as a relatively thin region of fluid flow dominated by viscous forces close to the surface of an object moving through a fluid. In this region, the velocity is zero at the surface, i.e. the no-slip condition holds, and rapidly approaches the outer flow speed away from the surface. This region is where the aeroplane's friction drag originates. The friction is a function of the wall shear stress, τ_w , which critically depends on the velocity gradient at the wall. The velocity gradient is strongly affected by the boundary layer state, which can be laminar or turbulent.

A laminar flow is an organised, layered flow characterised by very little exchange of fluid particles perpendicular to the flow direction (Schlichting and Gersten, 2016), whereas a turbulent flow can be described as irregular and consisting of numerous space and time scales (Bailly and Comte-Bellot, 2015). Both states are illustrated in Fig. 1.1. In turbulent flows, the exchange of fluid particles in the direction perpendicular to the mean flow direction is much greater. The greater transport of fluid momentum perpendicular to the (mean) flow direction in turbulent boundary layers is the principal mechanism responsible for their higher skin-friction drag compared to laminar boundary layers under otherwise similar flow conditions. As high-momentum particles are moved near the wall, the mean flow profile is altered. This increased flow gradient at the wall results in a stronger shear according to the equation for the wall shear:

$$\tau_w = \mu \frac{\partial u}{\partial y}, \quad (1.1)$$


where τ_w is the wall shear stress, μ is the dynamic viscosity, u is the flow velocity and y is the wall-normal coordinate. To minimise the friction drag, the boundary layer over the wing surface should be laminar over as long a distance as possible, as the friction coefficient for otherwise equal flows is between 3-10 times lower for laminar as opposed to turbulent boundary layers, depending on the Reynolds number (Schlichting & Gersten, 2016). However, a laminar boundary layer is not guaranteed to stay laminar; the flow might develop into a turbulent state through a process called transition. The understanding, modelling and manipulation of this transition process is an important aspect of wing design.

1.3. SWEEP WINGS

As aircraft were equipped with ever more powerful engines, the speed aircraft could achieve approached the speed of sound. This happened around the 1930s with the introduction of the jet engine. Flying near the speed of sound results in unwanted compressibility effects, most notably the occurrence of shock waves. The appearance of shocks strongly increases wave drag and possibly causes flow separation. Reducing the incoming flow velocity normal to the leading edge helps reduce these compressibility effects

³Image (a) source: https://en.wikipedia.org/wiki/Swept_wing, retrieved on 19-04-2025.

⁴Image (b) source: <https://www.thisdayinaviation.com/28-october-1957>, retrieved on 19-04-2025.



(a) FJ-1 (left) and FJ-2 (right), 1952.³



(b) Introduction of the Boeing 707, 1957.⁴

Figure 1.2: Some of the first swept-wing aeroplanes: (a) the FJ-2 fighter jet and (b) the Boeing 707.

(Vos & Farokhi, 2015). This is achieved by applying a (usually backwards) sweep angle (Λ) to the wing such that the velocity normal to the leading edge is reduced geometrically by a factor of $\cos(\Lambda)$. This sweeping of the wing is easily observed in Figure 1.2a where a newer model of the FJ fighter jet (right) shows a swept-wing design. Since the introduction of the Boeing 707 in 1957, large-scale high-subsonic passenger aircraft have also featured swept-back wings, with the corresponding crossflow boundary layers. This makes the application of drag reduction methods to such flows an interesting and relevant topic to study.

Sweeping the wing backwards introduces a spanwise flow component in the presence of a favourable-to-adverse pressure gradient along the direction of the chord. This affects the external streamline and boundary layer profile as shown in Fig. 1.3, and makes the flow highly three-dimensional. The pressure gradient curves the inviscid streamlines, those at the upper edge of the boundary layer, inboard near the leading edge and outboard downstream (Haynes and Reed, 2000). Normal to the surface, the pressure is nearly constant throughout the boundary layer, as proved by Prandtl, 1905. Near the wall, lower velocities as a result of viscous forces disturb the balance between the inertia and pressure forces present outside of the boundary layer. The velocity of particles in the boundary layer adjusts to the new force balance, giving rise to an additional flow component. If the inviscid streamline direction is defined as the tangent of flow at the boundary layer edge, the crossflow component normal to this direction can be seen to rise within the boundary layer as a result. This three-dimensional boundary layer profile is depicted in Fig. 1.3.

In the coordinate system aligned with the local inviscid streamline, the crossflow component can be seen to reduce to zero at the wall as well as fade in the freestream by the definition of the coordinate system. As a result, this crossflow component must feature an inflection point, which is additionally shown in Fig. 1.3. This inflection point is responsible for the rise of Crossflow Instabilities (CFI) in the boundary layer (W. S. Saric et al., 2003).

Although the transition to turbulence in swept-wing boundary layers can be triggered by various types of instabilities, CFI normally dominate in cruise conditions. They manifest as co-rotating vortices, depicted by the ribbons in Fig. 1.1, roughly aligned with the inviscid streamline. Depending on the conditions (mostly the freestream turbu-




Figure 1.3: Three-dimensional streamline over a swept wing (a) and the boundary layer profile featuring an inflection point in the crossflow velocity component shown in blue. Figure adjusted from Yiming et al., 2019.

lence level), they can be stationary or travelling (Bippes, 1999). Travelling CFI are excited mainly by the vortical content of the freestream turbulence. These travelling Crossflow Instabilities are more unstable than their stationary counterpart (Joslin & Streett, 1994). However, because aeroplanes experience very low freestream turbulence in flight, the travelling CFI are excited at much lower amplitudes. This makes stationary CFI generally the dominant instability type in swept-wing boundary layers (Wassermann & Kloker, 2003). Their development is the focus of this doctoral thesis.

1.4. LAMINAR FLOW CONTROL

Laminar Flow Control (LFC) is an overarching term that dates back to the 1930s and 40s with the introduction of modern stability theory for laminar boundary layers (Fischer et al., 1983). It is used to describe techniques that delay the onset of turbulence in the boundary layer. The potential of LFC to reduce aircraft drag is immense; Schrauf, 2005 states that successful application of LFC to wings, horizontal tail planes, fins, and nacelles could lead to a reduction of the total drag of aircraft by up to 16%.

LFC can be subdivided into passive and active flow control devices, where the distinction is made between the need for an external energy source (active) or not (passive). Although active methods such as jets (Raju et al., 2009), plasma actuators (Dörr and Kloker, 2015; Duchmann et al., 2014), and boundary layer suction (G. A. Reynolds and Saric, 1986) have been shown to delay transition effectively, their real-world application is often hindered by mass and monetary costs as well as practical limitations (Denning et al., 1997; Fischer et al., 1983; Jabbar et al., 2010; Joslin, 1998). Instead, passive flow control methods are often leveraged to reduce the drag of aircraft without the large costs associated with active methods. Examples of such methods include the shape of the wing itself. Such Natural Laminar Flow (NLF) wings are designed so that the transition location is as far downstream as possible. W. S. Saric et al., 2008 and more recently Zoppini et al., 2022 showed that another passive device, Discrete Roughness Elements (DRE), is very effective in achieving transition delay in controlled environments. This concerns the placement of small cylindrical obstructions, usually in one or more rows, on the wing surface to induce vortical flow structures within the boundary layer. These structures, i.e. other (sub-critical) modes, then influence the growth of perturbations

that dominate the transition scenario via nonlinear interactions. The forced structures then decay, leaving nothing but the hampered growth of the lower wavenumber instabilities (W. S. Saric et al., 2008).

Independent of the chosen method, LFC methods are often hindered by surface features that can arise in the manufacturing of lifting surfaces. Some examples are steps and gaps that form at the junction of wing panels (Worner et al., 2003), as well as bumps and waviness resulting from plate rolling techniques and buckling under loads (Holmes et al., 1984). Improved manufacturing techniques have significantly diminished the presence of bumps and waviness, rendering them less prominent in contemporary wings. This advancement in manufacturing may be responsible for the attention shifting away from smooth surface irregularities in contemporary research, resulting instead in a focus on sharp surface geometries such as Forward-Facing Steps (FFS), Backward-Facing Steps (BFS), and gaps as seen in e.g. Beguet et al., 2017; Forte et al., 2015; Perraud and Seraudie, 2000; Rius-Vidales and Kotsonis, 2020; Wang and Gaster, 2005 and Casacuberta et al., 2022.

The goal of studies into these features has generally been to provide manufacturing tolerances such that the negative effects on the transition are minimised. However, the flow around surface features is rarely studied in detail due to the associated difficulty and costs. For example, in experimental studies, the flow field around such surface features has proven difficult to visualise and study directly, leading to discrepancies in the interaction mechanisms as described by Tufts et al., 2017 and Eppink, 2020 (see Rius-Vidales and Kotsonis, 2021). Instead, the focus is often placed on the resulting transition location (see e.g. Crouch et al., 2006). In numerical studies, often-used stability frameworks such as the Orr-Sommerfeld equation and Parabolised Stability Equations suffer from assumptions on the importance of streamwise gradients in such flows (Franco et al., 2018). Moreover, they are often derived from linearised equations, neglecting the interaction between perturbations and the distortion of the base flow itself; an assumption that is difficult to justify a priori. The main reason for this lack of detail is that surface features of even several microns in size can affect the transition process significantly, see e.g. Radeztsky Jr et al., 1999. Consequently, it is difficult to obtain detailed experimental measurements, while performing Computational Fluid Dynamics (CFD) is costly due to the fine mesh required to resolve the near-wall flow structures. This is even more true for studies examining all the parameters that might influence the flow dynamics.

The focus of previous studies has therefore been to discover empirical relations between e.g. the feature's size or location and the effect on the growth of instabilities, see e.g. the works of Holmes et al., 1985, Zuniga et al., 1994, and Drake et al., 2010 on critical Reynolds numbers Re_k and C. Edelmann and Rist, 2014, Crouch et al., 2006, and Tufts et al., 2017 working on N -factor corrections. The N -factor is a growth measure introduced in a simpler form independently by Smith and Gamberoni, 1956 and later by Van Ingen, 1956. It describes the growth factor expressed as e^N experienced by a perturbation from the point at which it first became unstable. Experiments are generally used to define a critical N -factor, i.e. the N -factor at which the transition occurred. The aforementioned empirical relations then aim to find an additional amplification factor, ΔN , to express the effect on the transition process, see e.g. C. A. Edelmann and Rist, 2015. A positive ΔN implies stabilisation because more growth is required for the flow to transi-

tion to turbulence. The opposite is true for a negative ΔN , which implies destabilisation. However, recent studies have found that this model is too simplistic. Eppink, 2020 and Rius-Vidales and Kotsonis, 2020 point to factors such as perturbation amplitude that are ignored in ΔN methods. These factors can be accounted for in higher-order stability methods such as the Harmonic Navier-Stokes equations.

1.5. FLOW STABILITY ANALYSIS

To overcome the high costs associated with solving the governing equations of fluid motion, a compartmentalised analysis technique is often leveraged. Several (systems of) equations fall under this umbrella; Some noteworthy examples are the Orr-Sommerfeld (OS) equation (Orr, 1907 and Sommerfeld, 1909), the Parabolised Stability Equations (PSE) (Bertolotti et al., 1992 and Herbert, 1993), and the Harmonic Navier-Stokes (HNS) equations (Streett, 1998). All of these follow from a perturbation analysis of the Navier-Stokes (NS) equations, where the total flow is described as the sum of a stationary, spanwise invariant base flow and perturbations. Then, the perturbations are described by a multiplication of a perturbation shape and a wave function. The latter describes the oscillatory behaviour and, in most cases, the growth of the perturbations in the streamwise direction. Assumptions regarding the spatial and temporal evolution of the base flow and the perturbations are then made to simplify the equations to varying degrees, requiring different numerical solution techniques to solve them.

Crucially, one such assumption concerns streamwise derivatives of the perturbation (shape) and/or base flow. In the case of the OS equation, all streamwise dependencies are ignored. This includes both the change of the perturbation shape as well as the growth of the boundary layer. This goes so far as to assume that the vertical velocity component of the base flow is strictly zero, consistent with an order of magnitude analysis of the equations (Herbert, 1993). In other words, the boundary layer is assumed not to grow locally. The resulting problem can be written as an eigenvalue problem. This explains why the OS equation is still commonly used to generate initial conditions for more involved flow stability analysis frameworks.

One such framework is the PSE. Its main defining assumption is that the flow varies slowly and hence the shape function of the perturbation changes gradually. Consequently, second-order derivatives of this shape function can be neglected. The resulting system of equations is essentially parabolic in the streamwise direction (Bertolotti et al., 1992; Herbert, 1993). Some ellipticity remains from the pressure derivative term, as explained by Li and Malik, 1996, resulting in a minimum step size or a stabilisation procedure as shown by Andersson et al., 1998 to ensure numerical stability. The (mostly) parabolic nature of the PSE allows for the discretised system to be solved in a marching fashion. For cases with sufficiently small streamwise variation, e.g. the surface geometry, the PSE can yield highly accurate results for a fraction of the cost of a Direct Numerical Simulation (DNS).

When surface features are present in the flow, the validity of these assumptions is not guaranteed. In other words, the contribution of spatial derivatives of the shape function might be of the same order as the remaining terms. Recent works by Cooke et al., 2019; Dobrinsky, 2003; Franco et al., 2020, and Casacuberta et al., 2021 on both sharp and smooth surface features proved that the PSE equations cannot reliably predict the de-

velopment of incoming perturbations. Moreover, a method to estimate the importance of (higher-order) streamwise derivatives a priori does not yet exist, and it is unknown where the limits of the slowly varying flow assumption of the PSE lie.

Consequently, a framework based on the HNS equations is required to provide trustworthy insight into the interaction of perturbations with surface features. The harmonic Navier-Stokes equations fully maintain higher-order streamwise derivatives. As a result, they are solved for the entire domain at once, rather than in a marching fashion. This brings forth a greater computational cost than OS and PSE. The HNS framework lies at the core of this thesis and is discussed in Chapt. 3. A more in-depth overview of the different stability methods will be given in Sect. 2.5.

1.6. OBJECTIVES AND THESIS OUTLINE

This thesis focuses on numerically predicting the growth of boundary layer instability, accounting for nonlinear interactions and surface features. The aim is to develop a robust yet efficient framework to aid investigations of the effects of surface features in boundary layer transition, be it from the perspective of understanding the flow mechanics, finding manufacturing tolerances or for developing novel active or passive LFC devices. The developed framework is then used to perform simulations of transitional flows over smooth and sharp surface modifications.

The first part of this work is dedicated to developing a novel stability framework. This framework is the Delft Harmonic Navier-Stokes framework (DeHNSSo) and is based on the discretised HNS equations. Special attention is put toward the computational efficiency of the method in terms of both the use of Random Access Memory (RAM) and computational speed. The Research Aims (RAs) are formulated as:

- RA 1:** *Develop and implement a nonlinear stability framework based on the Harmonic Navier-Stokes equations capable of accounting for smooth and sharp surface features.* (Chapter 3)

- RA 2:** *Investigate and implement methods to accelerate computations and reduce the amount of system RAM required for nonlinear simulations.* (Chapter 3).

The second and final part of this work involves the use of the developed stability framework, as well as auxiliary developed frameworks, to gain further insight into the interaction of incoming perturbations with various smooth and sharp surface features. Additionally, the slowly varying flow assumption present in PSE is carefully considered by applying PSE to cases with varying nonparallelism in the flow. The aims relating to this part are:

 - RA 3:** *Identify the effects and influence of the slowly varying flow assumption of the Parabolized Stability Equations in flows over smooth and sharp surface features.* (Chapters 4 and 5)

 - RA 4:** *Investigate the effect of smooth surface humps as a method for laminar flow control in Crossflow Instability-dominated swept-wing boundary layers.* (Chapter 6)

1.7. OUTLINE

This dissertation will proceed with a detailed background on the numerical prediction of transition in Ch. 2. The gap in the current literature is then established. In Ch. 3, the derivation and implementation of a new flow stability analysis framework, called DeHNSSo, will be explained. This chapter will also feature a validation of the framework. Then, DeHNSSo, as well as existing classical stability methods, are used to study the effect of DREs, forward-facing steps and humps on the transition process of swept-wing boundary layers. First, NPSE is used to find a relation between the resulting CFI and DRE heights in Ch. 4. Then, in Ch. 5, NPSE and HNS frameworks are used to study the interaction between incoming CFI and forward-facing steps. Lastly, DeHNSSo is used to study the interaction between incoming CFI and smooth surface humps at amplitudes ranging from fully linear to strongly nonlinear in Ch. 6. Concluding remarks on both the developed frameworks and the results of the flow analyses are presented in the final chapter, Ch. 7. Recommendations for future work are also considered in Ch. 7.

2

BACKGROUND OF FLOW STABILITY ANALYSIS

In mathematics, you don't understand things. You just get used to them.

Johann von Neumann

How much easier it is to be critical than to be correct.

Benjamin Disraeli

Over the wall of an object moving relative to a viscous flow that surrounds it, boundary layers develop. This flow region close to the wall is initially laminar, i.e. smooth. Through the inception and growth of disturbances, such flows can transition to turbulence. The study of transition is complex and predicting the state of a boundary layer can be (computationally) expensive, especially when imperfections in the surface (e.g. steps, humps, and gaps) are incorporated into the study. Computationally efficient frameworks exist such as those based on the Orr-Sommerfeld (OS) or the Parabolized Stability Equations (PSE), but cannot always be relied upon when these features are included due to the rise of non-negligible streamwise gradients in the flow. Hence, there exists a need for a computationally efficient framework to study the effects of surface imperfections. This chapter introduces various commonly used flow stability prediction frameworks. The assumptions necessary for their derivation are analysed and summarised.

Parts of this chapter have been published in Westerbeek, S., Hulshoff, S., Schuttelaars, H., & Kotsonis, M. (2024). DeHNSSo: The Delft Harmonic Navier-Stokes Solver for nonlinear stability problems with complex geometric features. *Computer Physics Communications*, 302, 109250.

2.1. INTRODUCTION

In the previous chapter, the need for computationally efficient flow stability analysis frameworks was explained. These frameworks serve a critical role during the design of wings and in particular studies of laminar flow control. Several frameworks were named that use equations that follow from a perturbation analysis of the governing equations, the Orr-Sommerfeld, Parabolized Stability Equations, and Harmonic Navier-Stokes equations. This thesis focuses on these (most common) frameworks for flow stability analysis. Other methods, such as weakly non-linear analysis (see e.g. Wu, 2019), Floquet theory as shown in e.g. Herbert, 1988, triple deck theory (A. H. Nayfeh, 1991), and the extension to secondary instability analysis (Ambrosino et al., 2025; Malik et al., 1994), are considered out of its scope.

This chapter will provide a deeper dive into the frameworks that are considered state-of-the-art in flow stability analysis. First, the description of the flow is formalised in section 2.4. Then, the OS, PSE, and HNS equations are derived from the governing equations in section 2.5. The differences in the assumptions between each framework are explained and will be summarized at the end of the chapter in Tab. 2.1 for future reference.

2.2. THE CONCEPT OF FLOW STABILITY

In the context of (boundary layer) flows, A. V. Boiko et al., 2011 states that a flow is stable if disturbances do not result in a conversion of the flow to another (not necessarily turbulent) state. A stable system returns to the unperturbed state after the application of a small disturbance. This very general definition is necessary due to the large variety of instability types and mechanisms. For example, flows can be convectively unstable or absolutely unstable as indicated in Fig. 2.1. The difference between the two appears to be a mere nuance in the figure but has drastic consequences for the numerical treatment of the problem. In short, an absolutely unstable system sees the instability grow in time even at a constant location, whereas a convectively unstable system's perturbation growth is associated with movement in the spatial dimension. In this thesis, the focus will lie on convectively unstable systems (Fig. 2.1b) due to their dominance in the transition scenario of swept-wing boundary layers.

In fluid mechanics, the unperturbed state is often referred to as the *base flow*: the laminar solution to the problem void of perturbations. It is, by definition, a solution to the Navier-Stokes equations. In practice, however, since this solution itself might be computationally expensive to find, a close approximation is often used via e.g. a Reynolds-Averaged Navier-Stokes (RANS) simulation or using a solution to the steady Prandtl boundary layer equations. The solutions of the RANS or boundary layer equations are not strictly solutions to the Navier-Stokes equations, violating the definition of the base flow. During the derivation of flow stability analysis frameworks, the base flow is assumed a solution to the steady unperturbed problem. The base flow is subtracted from the system of equations, resulting in the perturbation equations. When the base flow is not strictly a solution to the full Navier-Stokes equations, this introduces an error. Such errors might equally arise from the use of experimental measurements as base flows as shown in Bottaro et al., 2003. However, the errors of using e.g. a mean flow from

RANS are assumed negligible if the approximation of the base flow via such methods is sufficiently close to the exact solution (Sipp & Lebedev, 2007). Otherwise, direct numerical simulations of the unperturbed Navier-Stokes equations might be required to serve as a base flow. Although the stability of many flows can be considered, e.g. jets, (atmospheric) shear layers, boundary layers, pipe flows, etc, in this thesis, the term *base flow* is used to refer to the laminar unperturbed boundary layer over an aerodynamic surface.

Instability waves are introduced in the boundary layer through a process called receptivity. This process defines the inception of perturbations, rather than their evolution (W. S. Saric et al., 2002). Receptivity is the process by which the disturbances in the freestream and fluctuations in the flow establish the initial conditions of perturbation amplitude, frequency, and phase based on the ambient disturbance environment and the projection mechanisms of disturbances into eigenmodes (W. S. Saric et al., 2002; Schlichting and Gersten, 2016). Surface imperfections, vortical structures in the incoming flow, or sound waves are examples of environmental disturbances that can lead to the inception of instability waves. The amplitude of the perturbations is strongly influenced by the aforementioned freestream vorticity strength and surface imperfection size (Morkovin, 1994; W. S. Saric et al., 2002). Moreover, the amplitude also affects the path to transition as will be covered in Sect. 2.3.

Perturbations can grow or decay in space or time. It is stressed that some perturbations may decay while others may grow at the same location and time. Hence, to call a flow stable or unstable can be misleading. This would strictly mean that the flow is stable to all perturbations. Hence, here the stability of a specific perturbation or set of perturbations within a flow is considered instead. A perturbation is considered stable when it decays and unstable when it grows.

2.3. LAMINAR-TURBULENT TRANSITION

Instabilities as a result of certain perturbations can eventually lead to a large distortion of the base flow and result in laminar-turbulent transition. This process can be scrutinised via flow stability analysis. O. Reynolds, 1883 first experimentally explored the factors that influence the nature (straight or sinuous) of the flow and described the scaling laws that govern the tendency of the flow to either state. This was done by releasing a streak of




Figure 2.1: Disturbance propagation of absolutely unstable systems (a) and convectively unstable systems (b). Illustration taken from A. V. Boiko et al., 2011.




Figure 2.2: Depiction of (a) Reynolds' experimental setup and (b) a schematic representation of fully laminar flow, then the rise of eddies and lastly the development of a turbulent flow in a pipe as the flow velocity was increased.¹

coloured dye into a flow through a glass tube as shown in Fig. 2.2¹. O. Reynolds, 1883 states: "As the velocity was increased, at some point in the tube, always at a considerable distance from the trumpet or intake, the colour band would all at once mix up with the surrounding water, and fill the rest of the tube with a mass of coloured water." The colour residing within a smooth band indicates a laminar flow state, while the mixing of the coloured band with the surrounding flow indicates a turbulent flow state. The velocity was proven to affect the location where the transition to turbulence occurs, and the onset of turbulence was found to happen rapidly. Moreover, once the flow became turbulent, it quickly spread through the width of the tube. Reynolds showed the dependency of the transition location on the (kinematic) viscosity ν , distance L , and velocity u for equivalent geometries (in this case, the pipe diameter). This gave rise to the well-known Reynolds number:

$$Re = \frac{uL}{\nu}. \quad (2.1)$$

Transition to turbulence is a complex process that is not unique, i.e., not all flows transition the same way. Morkovin, 1994 and later Reshotko and Tumin, 2006 categorised the paths to transition into five main branches denoted by the letters A-E and provided the sketch shown in Fig. 2.3. The figure shows that the dominant path depends on the disturbance level. Examples of this are freestream turbulence and acoustic noise as found by Morkovin, 1969. Later surface imperfections, even those as small as several microns, were found to play an important role (Deyhle and Bippes, 1996). Receptivity, a mechanism fed by the forcing of environmental disturbances, precedes all transition paths. Depending on the initial condition, breakdown to turbulence can occur via any of the paths indicated in Fig. 2.3. This work focuses mainly on the eigenmode and transient growth processes found in paths A, B, and C. Hence, they will be given most attention.

¹Image (a) from Reynolds (1883), retrieved from Wikimedia Commons on June 13, 2025: https://commons.wikimedia.org/wiki/File:Reynolds_fluid_turbulence_experiment_1883.jpg. Originally published in Papers on Mechanical and Physical Subjects, vol. 2 (1901), now reissued by Forgotten Books and Image (b) taken from CFD Direct (n.d.). Reynolds' experiment. Retrieved June 13, 2025, from <https://doc.cfd.direct/notes/cfd-general-principles/reynolds-experiment>.




Figure 2.3: Paths to transition as shown in Reshotko and Tumin, 2006.

The route to turbulence described by path A is described by initially weak disturbances that grow exponentially, over a distance or in time. The exponential growth of these primary instabilities is known as eigenmode growth. Due to the small amplitude of the mode initially, compared to the base flow velocities, this growth is accurately predicted by linear methods. As the amplitude of the perturbation increases, the effect of the disturbance on the mean flow starts to be noticeable and nonlinear interactions become non-negligible (W. S. Saric et al., 2002). In other words, the perturbations are able to distort the flow that they travel through. This deformed flow is in turn destabilised by other so-called secondary instabilities (Kohama, 1987; H. L. Reed and Saric, 1989 Herbert, 1988). Secondary instabilities feed on the newly developed shears in the flow, starting the cascade of perturbation energy to ever-smaller scales. While this is a critical stage in the transition process, this thesis will focus solely on the development of the primary instabilities to maintain a workable scope and isolate and characterise the development of stationary CFI.

The path to transition was long thought to be dominated solely by path A, or eigenmode growth. Any transition scenario that did not fit this model was considered bypass transition (see Morkovin, 1969, Fig. 7.) It was only 30-some years ago that the concept of transient growth was introduced. It was found that a (local) dominance of transient growth over eigenmode growth can occur as in paths B, C, and D of Fig. 2.3. Transient growth, also referred to as algebraic growth, results from the non-orthogonal nature of the functions describing the stability of flows (Reshotko, 2001). The effect of this non-




Figure 2.4: Schematic representation of temporal transient growth recreated from Schmid et al., 2002 showing the growth of the resulting vector despite the decline of both its contributing vectors.

orthogonality is that the solutions can partially cancel out, i.e. the total perturbation kinetic energy might not resemble the scalar sum of the kinetic energy from the underlying perturbations. As a consequence, even if both modes decay, if one of the modes decays more rapidly than the other, the vectorial sum of the two modes can grow. A geometric visualisation of this process in two dimensions is shown in Schmid et al., 2002. This rapid growth can either pre-condition the flow for eigenmode growth (path B), achieve large enough amplitudes to distort the mean flow and give rise to secondary instabilities (path C), or grow so strongly that a breakdown to turbulence via a bypass mechanism occurs (path D).

2.4. SWEEP WINGS AND CROSSFLOW INSTABILITIES

The paths to transition shown in Fig. 2.3 are not limited to one specific type of flow and instead describe trends that are observed in many different flows. In the receptivity phase, various instabilities can be generated. However, for specific flows, say the flow over a swept wing on a commercial aircraft, the transition scenario is commonly dominated by a specific instability type. The sweep angle and airfoil shape define the base flow. At high Reynold's numbers, sweep angles of over 15° result in crossflow instabilities dominating the transition scenario (Schrauf, 2005). Whether the nature of the CFI is travelling or stationary depends mainly on the freestream turbulence intensity and surface roughness (Joslin & Streett, 1994).

Around the turn of the last century, Arnal et al., 1997; H. L. Reed and Saric, 1989 and Bippes, 1997 among others performed crossflow-related research, which led to a much better understanding of this phenomenon. Crossflow Instability is an inviscid instability, i.e. one that arises due to the presence of an inflection point in the boundary layer profile (W. Saric & Reed, 2003). This inflection point originates from the crossflow component in the boundary layer previously shown in Fig. 1.3. Crossflow Instability appears in either a stationary or travelling form, although the two can and often do co-exist within, e.g. a transitioning boundary layer. Generally, however, transition to turbulence is triggered




Figure 2.5: Infrared thermography of a flow (coming from the left) over a swept wing. Lighter areas correspond to laminar flow, while darker areas (on the right) correspond to turbulent flow with a jagged front around $x/c_x = 0.4$ from Zoppini, 2023

by one of the two, depending on various factors. Of those factors, freestream turbulence intensity is the most important (Bippes, 1999). For the low atmospheric turbulence at cruise flight altitude, stationary Crossflow Instabilities are most prominent in swept-wing boundary layers. Hence, stationary Crossflow Instabilities will feature prominently in this thesis. The travelling CFI is excluded from the current scope. By extension, the steady-unsteady interactions of CFI are excluded in this work.

The crossflow instability manifests itself as co-rotating vortices that can strongly modify the flow by sweeping low-momentum fluid upward and high-momentum fluid particles downward (Serpieri and Kotsonis, 2016). As the name suggests, stationary Crossflow Instabilities appear frozen in time. In contrast, for a travelling CFI, the up- and downwelling regions travel in the spanwise direction (z) over time. Mathematically, this can be described as a phase change due to the periodicity in the z -direction. In experiments, travelling Crossflow Instabilities are quickly recognised due to a blurry uniform transition front approximately parallel to the leading edge due to their spanwise motion over the wing surface. Stationary CFI instead lead to a more jagged and well-defined transition front (Serpieri, 2018; Zoppini, 2023) providing a clear distinguishing feature to determine the main cause of transition in experiments. Fig. 2.5 shows such a jagged transition front from experiments by Zoppini, 2023 in the Low-Turbulence Tunnel (LTT) that approximate cruise conditions due to the low turbulence levels in the flow.

2.5. FLOW STABILITY ANALYSIS FRAMEWORKS

The development of boundary layer perturbations can be numerically predicted via various stability analysis frameworks that follow from the Navier-Stokes equations. The assumptions made during their derivation drastically impact their application, robustness, and numerical solution method. In this section, the most relevant frameworks for the current thesis, namely those based on Local Linear Stability Theory (Orr-Sommerfeld), the Parabolized Stability Equations and the Harmonic Navier-Stokes equations are described, after an introduction to the problem definition below.




Figure 2.6: Numerical domain of interest for the Orr-Sommerfeld (red), Parabolized Stability Equations (blue, arrows indicate marching direction), and Harmonic Navier-Stokes equations (grey). Global x , y , and z coordinates are shown in addition to the curvilinear ξ , η , and z coordinates that adhere to the smooth wall deformation. The dashed lines represent the boundary layer thickness. Note that all the mentioned frameworks leverage the OS framework to provide inflow conditions at $x = x_0$.

2.5.1. DOMAIN DESCRIPTION

The aforementioned frameworks are derived from the Navier-Stokes equations with different assumptions. While not all methods directly utilise a two-dimensional domain, e.g. the Orr-Sommerfeld framework, a generalised domain description is necessary. Even in the case of the Orr-Sommerfeld equations, they are often evaluated at various stream-wise locations in this two-dimensional domain to predict instability growth. A generalised schematic of the four-sided problem domain is presented in Fig. 2.6 showing the different numerical solution strategies for the OS, PSE, and HNS frameworks. All cases in this work can be described using this generalised domain. It features an inflow plane, a (possibly locally curved) bottom wall, a top boundary, and an outflow plane. The semi-infinite physical domain has to be truncated to allow for a numerical treatment of the problem. This results in an artificial inflow, outflow, and top boundary. The choice of flow stability analysis framework affects the boundary conditions, especially regarding the outflow plane. When flow stability frameworks are used that can carry information upstream, the interaction of instabilities with the outflow boundary can rapidly deteriorate results throughout the domain. The outflow boundary implementation in the Harmonic Navier-Stokes framework developed in this research is considered in detail in Sect. 3.3.3. The effect of approximations at the top boundary is mitigated by ensuring a sufficiently tall domain for all methods.

2.5.2. REFERENCE AND NON-DIMENSIONAL QUANTITIES

All quantities throughout this thesis will be shown in non-dimensional form by default. Dimensional quantities are denoted with an overbar. Global characteristic scales are used to nondimensionalise all equations. Length quantities are normalized by the reference length $\bar{l}_{\text{ref}} = \sqrt{\bar{x}_0 \bar{\nu}_{\text{ref}} / \bar{U}_{\text{ref}}}$, defined as the Blasius length scale at the inflow (i.e., $\bar{x}_0$) of the computational domain as shown in Fig. 2.6. The reference kinematic viscosity is defined by $\bar{\nu}_{\text{ref}}$, and velocity quantities are normalized by the reference velocity $\bar{U}_{\text{ref}}$, taken as the external (i.e., inviscid freestream) base flow streamwise (i.e., along x) velocity at the inflow. A global Reynolds number can then be defined as $Re = \bar{l}_{\text{ref}} \bar{U}_{\text{ref}} / \bar{\nu}_{\text{ref}}$.

2.5.3. NAVIER-STOKES EQUATIONS

The motion of fluid particles in low-speed flow ($M < 0.3$, with M the Mach number) can be accurately described by the incompressible Navier-Stokes equations. In primitive variables (i.e., in velocity-pressure formulation), this leads to the following equations for momentum and continuity:

$$\frac{\partial \mathbf{q}}{\partial t} + (\mathbf{q} \cdot \nabla) \mathbf{q} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{q}, \quad (2.2a)$$

$$\nabla \cdot \mathbf{q} = 0, \quad (2.2b)$$

where $\mathbf{q} = [u, v, w]^T$ is the vector of instantaneous velocities, p is the instantaneous static pressure and ∇ is the three-dimensional vector differential operator:

$$\nabla = \begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \\ \frac{\partial}{\partial z} \end{bmatrix}. \quad (2.3)$$

2.5.4. PERTURBATION ANALYSIS

To assess the behaviour of perturbations within a pre-existing base flow, a perturbation analysis of the Navier-Stokes equations (Eqs. 2.2) is performed. This involves the subdivision of the total flow into the steady base flow and perturbations such that $\mathbf{q}(x, y, z, t) = \mathbf{Q}(x, y, z) + \mathbf{q}'(x, y, z, t)$ and $p(x, y, z, t) = P(x, y, z) + p'(x, y, z, t)$, where $\mathbf{Q}$ and P are the velocity and pressure vectors of the base flow and $\mathbf{q}'$ and p' the perturbations of velocity and pressure. If this decomposition is substituted for $\mathbf{q}$ and p in Eqs. 2.2(a-b) and the base flow terms are subtracted, the perturbation equations are retrieved:

$$\frac{\partial \mathbf{q}'}{\partial t} + (\mathbf{Q} \cdot \nabla) \mathbf{q}' + (\mathbf{q}' \cdot \nabla) \mathbf{Q} + (\mathbf{q}' \cdot \nabla) \mathbf{q}' = -\nabla p' + \frac{1}{Re} \nabla^2 \mathbf{q}', \quad (2.4a)$$

$$\nabla \cdot \mathbf{q}' = 0, \quad (2.4b)$$

which are the basis for all stability analysis frameworks. Assumptions on the solution shape and spatial variance and the importance of nonlinear terms lead to different stability analysis frameworks.

2.5.5. ORR-SOMMERFELD EQUATION

The Orr-Sommerfeld equation due to and despite its simplicity is still widely used today to analyse the stability of flows, particularly in the aerospace industry. A unique quality of the OS equation, as opposed to the other flow stability analysis equations considered in this thesis, is its ability to provide a local solution, which includes a perturbation shape, from nothing more than a local base flow profile and the perturbation characteristics. This comes at the cost of accuracy, especially for long domains where small errors accumulate, mainly due to the assumption of parallelism. Consequently, the contemporary use of OS equation revolves around generating initial conditions for more advanced stability analysis frameworks or even DNS. In some cases, however, it can provide insight

into the stability of boundary layer profiles distorted by surface features and help explain the perturbation development over these features.

The derivation of the Orr-Sommerfeld equation starts from a perturbation ansatz. Here, only the ansatz and outcome of the derivation are presented. For a complete derivation, the reader is referred to Groot, 2018. The Orr-Sommerfeld equation assumes small amplitudes so that nonlinear terms can be neglected. Moreover, this means that the perturbation ansatz needs to consider only a single perturbation, rather than an ensemble. The ansatz for a single harmonic reads

$$\phi'(\xi, \eta, z, t) = \hat{\phi}(\eta)e^{i(\alpha\xi + \beta z - \omega t)} + \text{c.c.}, \quad (2.5)$$

where $\phi = [u, v, w, p]$ for both the total perturbation indicated by the prime as well as the shape function indicated by the hat. The total perturbation is a function of the generalised coordinates ξ and η (see Fig. 2.6), the spanwise direction z and time t . The exponential part of the ansatz is also referred to as the *wave function* containing α , the complex-valued streamwise wavenumber, β , the real-valued spanwise wavenumber, and ω , the real angular frequency. Crucially, the shape function $\hat{\phi}$ is only a function of the wall-normal direction. The same assumption holds for the base flow. The ansatz is substituted into Eqs. 2.4(a-b).

The equations can then be reduced to a single equation with unknowns α , and $\hat{v}$, namely:

$$\left[\left(i\omega - i\alpha U - i\beta W \right) \left(\frac{\partial^2}{\partial y^2} - \alpha^2 - \beta^2 \right) + i\alpha \frac{d^2 U}{dy^2} + i\beta \frac{d^2 W}{dy^2} + \frac{1}{Re} \left(\frac{\partial^2}{\partial y^2} - \alpha^2 - \beta^2 \right)^2 \right] \hat{v} = 0, \quad (2.6)$$

with $\hat{v} = 0$ at the wall and in the freestream. In this work, the Orr-Sommerfeld equation is often mentioned as a method to provide initial conditions for the PSE and HNS frameworks. However, the Orr-Sommerfeld equation provides only the shape of the wall-normal perturbation velocity and the streamwise wavenumber as can be seen in Eq. 2.6. The initial condition of the aforementioned requires shape functions for u , v , w , and p . Hence, an expanded form of the Orr-Sommerfeld is used instead which solves the problem for all perturbation velocities and pressure. The derivation of this set of equation follows from the substitution of ansatz 2.5 into the Navier-Stokes equations without further reducing the problem. The downside of this formulation is that it comes with a more complicated solution algorithm and a larger system of equations, see Groot, 2018, though that is not of much concern when it is merely used to generate initial conditions. In terms of assumptions, the two equations are equivalent and hence the term OS equation is used to refer to both.

2.5.6. PARABOLIZED STABILITY EQUATIONS

The PSE are derived from the perturbation equations by assuming a perturbation ansatz where the base flow, shape function and streamwise wavenumber can vary in the streamwise direction:

$$\mathbf{q}'(\xi, \eta, z, t) = \sum_{m=-M}^M \sum_{n=-N}^N \hat{\mathbf{q}}_{m,n}(\xi, \eta) e^{i(\int_{\xi_s}^{\xi_e} \alpha_{m,n} d\xi^* + \beta_n z - \omega_m t)}, \quad \beta, \omega \in \mathcal{R}; \quad \hat{\mathbf{q}}, \alpha \in \mathcal{C}, \quad (2.7)$$

where now the wave specifications α , β and ω and the shape function $\hat{\mathbf{q}}$ contain subscripts m and n . These indicate the wave characteristics of mode (m, n) with respect to the primary mode $(1, 1)$, specifically $\beta_n = n\beta_1$ and $\omega_m = m\omega_1$, truncated at $n = N$ and $m = M$. The imaginary unit is denoted i . The stability domain ranges from ξ_s to ξ_e . The $*$ annotation, in $d\xi^*$, indicates the integration variable. For the shape function and the streamwise wave number, no direct relation to the primary mode exists and their solution needs to be found through solving the stability equations. A more detailed derivation in Cartesian coordinates can be found in Bertolotti et al., 1992; Herbert, 1997. One particularity of the PSE that needs to be mentioned is the additional degree of freedom that the ansatz has for representing growth in the streamwise direction as can be seen in Eq. 2.7. Namely, there is a closure problem due to the dependency on the streamwise coordinate present in both the shape function and the wave function. Without further restraints, the shape function could capture (part of) the growth of the perturbation. In some cases, this growth could be rapid, possibly breaking the assumption that $\partial \hat{\mathbf{q}} / \partial \xi$ is small, i.e. the slowly-varying flow assumption. Hence, during the solution process, all the amplitude growth is redirected to the complex component of α and only shape changes are represented in the shape function.

Here, however, the derivation is performed in generalised coordinates to be able to account for (local) curvature. The parabolization assumption is made in the ξ direction, and the grid for this case is generated ensuring orthogonality at the wall. The Nonlinear Parabolized Stability Equations (NPSE) can be written in matrix form as:

$$\sum_{m=-M}^M \sum_{n=-N}^N \mathbb{L} \hat{\mathbf{q}}_{m,n} + \sum_{m=-M}^M \sum_{n=-N}^N \mathbb{M} \frac{\partial \hat{\mathbf{q}}_{m,n}}{\partial \xi} + \sum_{m=-M}^M \sum_{n=-N}^N \mathbb{N} \frac{\partial \alpha_{m,n}}{\partial \xi} \hat{\mathbf{q}}_{m,n} = \sum_{m=-M}^M \sum_{n=-N}^N \mathbf{r}_{m,n}, \quad (2.8)$$

where the right-hand side forcing, $\mathbf{r}$, comprises the nonlinear terms that are neglected in the Linear Parabolized Stability Equations (LPSE). Moreover, the summation over the modes m and n drops out and only a single mode is considered. The contents of the matrices $\mathbb{L}$, $\mathbb{M}$, and $\mathbb{N}$ can be found in appendix A.3.

2.5.7. HARMONIC NAVIER-STOKES EQUATIONS

As for the PSE and OS equation, a spatial coordinate transformation is adopted to bring the perturbation equations from the physical solution domain (x, y, z) to the generalised curvilinear numerical domain (ξ, η, z) . It must be stressed that the transformation is only applied in the streamwise and wall-normal directions (i.e., x and y , and not z). A derivation and explanation of the transformation is presented in Appendix A. To facilitate the harmonic expansion of the equations, a vector of perturbation state variables (i.e., velocity and pressure) is introduced: $\boldsymbol{\phi}' = [\mathbf{q}', p']^T$, where $\mathbf{q}'$ and p' are defined in the transformed numerical domain: $\mathbf{q}' = \mathbf{q}'(\xi, \eta, z, t)$ and $p' = p'(\xi, \eta, z, t)$. The perturbations are discretised using a finite sum of Fourier modes in the temporal and spanwise

dimensions. The evolution of the perturbations is then again described by the multiplication of an amplitude function $\hat{\phi}(\xi, \eta)$ and an exponential wave function in z and t . The employed solution form is reminiscent of those for the PSE (Bertolotti et al., 1992) and AHLNS (Franco et al., 2018), albeit here the wave function does not make use of a streamwise wavenumber. Hence, the streamwise-periodic nature of the perturbations is captured fully by the shape function. This comes at a computational cost, as finer discretization in the streamwise direction is necessary, although this allows the method to accurately predict rapid streamwise variations. No further assumptions on the importance of higher-order streamwise derivatives are made (neither with regard to the base flow) to maintain the generality of the perturbation field. The perturbations are described in the form:

$$\phi'(\xi, \eta, z, t) = \sum_{m=-\infty}^{\infty} \sum_{n=-\infty}^{\infty} \hat{\phi}_{m,n}(\xi, \eta) e^{i(\beta_n z - \omega_m t)}, \quad (2.9)$$

where $\beta, \omega \in \mathcal{R}$ and $\hat{\phi} \in \mathcal{C}$. The spanwise modal wavenumber is β , ω is the modal angular frequency, and t is the time. The common notation (m, n) is used to describe the relation of the wave specifications of higher harmonic modes, i.e., $\omega_m = m\omega_1$ and $\beta_n = n\beta_1$. Here ω_1 and β_1 indicate the fundamental frequency and spanwise wavenumber, respectively. In addition, the $(0, 0)$ mode has a spanwise wavenumber and angular frequency of zero. This mode specifically is referred to as the Mean Flow Distortion (MFD) and can be interpreted as the (nonlinear) modification of the base flow by the perturbations.

The ansatz of Eq. 2.9 is truncated to $-N \leq n \leq N$ and $-M \leq m \leq M$ to find a numerical solution. This leads to:

$$\phi'(\xi, \eta, z, t) = \sum_{m=-M}^M \sum_{n=-N}^N \hat{\phi}_{m,n}(\xi, \eta) e^{i(\beta_n z - \omega_m t)}, \quad (2.10)$$

where the truncation limits of m and n are denoted by M and N , respectively. These must be chosen such that most of the perturbation energy is captured by the solution. Consequently, higher harmonics can be assumed to have a marginal effect on the instability development and can be neglected. The Fourier ansatz employed here is nearly identical to NPSE formulations in e.g., Haynes and Reed, 2000; Herbert, 1993 aside from the absence of the streamwise wavenumber α . Inserting the perturbation ansatz of Eq. 2.10 into Eq. 2.4 leads to the Harmonic Navier-Stokes equations in generalised coordinates:

$$\sum_{m=-M}^M \sum_{n=-N}^N \mathcal{L}_{m,n} \hat{\phi}_{m,n} + \mathcal{M}_{m,n} \frac{\partial \hat{\phi}_{m,n}}{\partial \xi} + \mathcal{N}_{m,n} \frac{\partial^2 \hat{\phi}_{m,n}}{\partial \xi^2} = \sum_{m=-M}^M \sum_{n=-N}^N \mathbf{r}_{m,n}, \quad (2.11)$$

where the mode-specific operators $\mathcal{L}_{m,n}$, $\mathcal{M}_{m,n}$, and $\mathcal{N}_{m,n}$ comprise the respective products with $\hat{\phi}_{m,n}$, $\frac{\partial \hat{\phi}_{m,n}}{\partial \xi}$, and $\frac{\partial^2 \hat{\phi}_{m,n}}{\partial \xi^2}$. Within $\mathcal{L}_{m,n}$ one can also find the derivatives with respect to η which are dealt with differently than the derivatives with respect to ξ due to the use of spectral differentiation in the η direction. The right-hand side contains the double sum of nonlinear interactions over all considered modes. The contents and treatment of $\mathbf{r}_{m,n}$, i.e., the forcing of mode (m, n) , are considered in more detail in Sect. 2.5.8.

As Herbert, 1993 notes, the physical field must be real and thus it must hold that




Figure 2.7: Schematic of the Fourier mode space, adapted from Herbert, 1993. An example of a nonlinear interaction is shown between two modes via the arrows (from black source modes to white recipient mode). The dashed outline indicates the truncated ensemble of modes (M, N) and the dark-shaded area shows modes that can be found via symmetries. The light-shaded area shows modes that can be found via symmetries in 2D flows.

$$\phi'_{-n,-m} = \phi'_{m,n}{}^{\dagger}, \quad (2.12)$$

where $\dagger$ denotes the complex conjugate. This simple operation nearly halves the required number of calculations as shown in Fig. 2.7. The light-shaded area indicates the modes that can be generated from symmetry when base flow is purely two-dimensional (i.e. $W=0$) and $\omega_1 = 0$ or $\beta_1 = 0$. Then, an additional axis of symmetry is introduced in z such that:

$$\phi'_{m,-n} = [u'_{m,n}, v'_{m,n}, -w'_{m,n}, p'_{m,n}]^T. \quad (2.13)$$

2.5.8. HARMONIC BALANCING

The HNS and PSE (2.8 and 2.11) are both coupled systems, which complicates the development of an efficient solution algorithm. This can be simplified by employing harmonic balancing, which is based on the idea that the nonlinear interaction of two modes (i.e. their product), (ω_j, β_k) and (ω_l, β_p) , results in the forcing of a third mode (ω_g, β_h) as illustrated by an example in Fig. 2.7. This idea originates from the NPSE methodology as shown in Haynes and Reed, 2000; Herbert, 1993, and Westerbeek, 2020 and is easily visualised by observing the multiplication of the respective wave functions. As an example, consider the term $-u' \frac{\partial u'}{\partial \xi}$, where u' contains the double truncated sum shown in Eq. 2.10. Isolating the interaction of any two modes (j, k) and (l, p) , $-u' \frac{\partial u'}{\partial \xi}$ can be rewritten as:

$$\sum_{j=-M}^M \sum_{k=-N}^N \sum_{l=-M}^M \sum_{p=-N}^N -\hat{u}_{(j,k)} e^{i(k\beta_1 z - j\omega_1 t)} \frac{\partial \hat{u}_{(l,p)}}{\partial \xi} e^{i(p\beta_1 z - l\omega_1 t)}, \quad (2.14a)$$

$$= \sum_{j=-M}^M \sum_{k=-N}^N \sum_{l=-M}^M \sum_{p=-N}^N -\hat{u}_{(j,k)} \frac{\partial \hat{u}_{(l,p)}}{\partial \xi} e^{i((k+p)\beta_1 z - (j+l)\omega_1 t)}, \quad (2.14b)$$

$$= \sum_{g=-M}^M \sum_{j=\max(g-M, -M)}^{\min(M+g, M)} \sum_{h=-N}^N \sum_{l=\max(h-N, -N)}^{\min(N+h, N)} -\hat{u}_{(j,k)} \frac{\partial \hat{u}_{(g-j, h-k)}}{\partial \xi} e^{i(h\beta_1 z - g\omega_1 t)}, \quad (2.14c)$$

where in the last step interactions that exceed the spectral domain are excluded from the calculation. Thus, the mode (g, h) will be forced by every interaction of active modes (j, k) and (l, p) for which it holds that $j + l = g$ and $k + p = h$. In a given simulation, it is thus known a priori which interaction affects which mode. Consequently, only the relevant interactions are summed to find the source term that is presented on the right-hand side of the equations for this particular mode. This results in Eqs. 2.8 and 2.11 becoming at most a system of $(M + 1) \times (2N + 1)$ inhomogeneous linear equations accounting for symmetries. For PSE, this reads:

$$\mathbb{L} \hat{\mathbf{q}}_{m,n} + \mathbb{M} \frac{\partial \hat{\mathbf{q}}_{m,n}}{\partial \xi} + \mathbb{N} \frac{\partial \alpha_{m,n}}{\partial \xi} \hat{\mathbf{q}}_{m,n} = \mathbf{r}_{m,n}, \quad (2.15)$$

while for HNS the system can be written as:

$$\mathcal{L}_{m,n} \hat{\boldsymbol{\phi}}_{m,n} + \mathcal{M}_{m,n} \frac{\partial \hat{\boldsymbol{\phi}}_{m,n}}{\partial \xi} + \mathcal{N}_{m,n} \frac{\partial^2 \hat{\boldsymbol{\phi}}_{m,n}}{\partial \xi^2} = \mathbf{r}_{m,n}, \quad (2.16)$$

coupled via the nonlinear forcing terms. The matrices $\mathcal{L}$, $\mathcal{M}$, and $\mathcal{N}$ contain the information of base flow and mode characteristics. These matrices will be further scrutinised in Chapt. 3.

Table 2.1: Overview of flow stability analysis frameworks.

Framework	Ansatz, $\boldsymbol{\phi}'$	Assumptions	Solution Technique
OS	$\hat{\boldsymbol{\phi}}_{m,n}(\eta) e^{i(\alpha_{m,n}\xi + \beta_n z - \omega_m t)}$	$\frac{\partial}{\partial \xi} = 0, V = 0$	Local eigenproblem
PSE	$\hat{\boldsymbol{\phi}}_{m,n}(\xi, \eta) e^{i\left(\int_{\xi_s}^{\xi} \alpha(\xi)_{m,n} d\xi + \beta_n z - \omega_m t\right)}$	$\frac{\partial^2 \hat{\boldsymbol{\phi}}}{\partial \xi^2} = 0$	Streamwise marching
HNS	$\hat{\boldsymbol{\phi}}_{m,n}(\xi, \eta) e^{i(\beta_n z - \omega_m t)}$	-	Direct solver

2.5.9. SUMMARY OF FLOW STABILITY ANALYSIS EQUATIONS

The flow stability analysis equations considered in this thesis crucially differ in terms of the solution ansatz and solution technique, which results in benefits and limitations that are summarised in table 2.1 for future reference. The table does not explicitly state the common assumptions of incompressible flow and a steady, spanwise invariant base

flow, and periodicity in the spanwise direction and time of the perturbations, but it is good to keep these assumptions in mind.

In practice, the use-cases of the three methods partially overlap, but can be simplified as follows: The Orr-Sommerfeld is used to find the initial shape of a mode with certain wave specifications (β and ω) or in some cases as an a posteriori analysis tool for specific (deformed) base flow states. Although the OS framework can be used to find the linear development of a mode, it will not be used for that purpose here. The PSE framework is used to perform linear and nonlinear flow stability analysis over flat surfaces due to its high computational efficiency and accuracy. The PSE can be used to perform flow stability analysis over (locally) curved surfaces, but the HNS framework is used for such cases. The HNS framework does not suffer from any assumptions on the streamwise gradients of the flow and hence prevents an invalid application in that regard. Flow stability analysis of flows over surfaces with smooth and sharp features generally also includes flat surface sections. The HNS framework is then applied in the entire domain, as a combination of the PSE and HNS frameworks falls out of the implementation scope of this thesis.

3

DEVELOPMENT OF A HARMONIC NAVIER-STOKES SOLVER

Life need not be easy, provided only that it is not empty.

Lise Meitner

*Keep your mind open to change all the time. Welcome it. Court it.
It is only by examining and re-examining your opinions and ideas that you can progress.*

Dale Carnegie

A harmonic Navier-Stokes solver was developed that can account for the nonlinear interactions of perturbations and account for both smooth and sharp surface features. The equations are rewritten by factoring out mode-specific terms in conjunction with an explicit formulation to strongly reduce the computer's Random Access Memory (RAM) required to find the solution. Additionally, the memory requirements scale favourably with the number of modes considered. The outflow boundary is preceded by a buffer region that damps incoming perturbations and prevents upstream travelling reflections. The implementation is validated by means of four canonical boundary layer transition scenarios and verified through a grid convergence study. Simulation results closely match state-of-the-art solvers (NPSE, AHLNS, DNS). Computation times are an order of magnitude larger than NPSE, but several orders of magnitude lower than DNS.

Parts of this chapter have been published in Westerbeek, S., Hulshoff, S., Schuttelaars, H., & Kotsonis, M. (2024). DeHNSSo: The Delft Harmonic Navier-Stokes Solver for nonlinear stability problems with complex geometric features. *Computer Physics Communications*, 302, 109250.

3.1. INTRODUCTION

THE proposed HNS framework that lies at the core of this thesis, named the Delft Harmonic Navier-Stokes Solver (DeHNSSo), is envisioned as a generalised and flexible simulation tool for a wide range of transitional flows. Wall-bounded laminar boundary layers form the key type of flows typically explored in such analyses and will hence be used to validate and verify the implementation.

In this chapter, the details of the implementation of DeHNSSo are discussed in meticulous detail such that a reader should be able to reproduce the code. However, the latest version of the code implemented in Matlab is available on GitHub¹. This chapter will treat the generalised domain description, implementation covering everything from discretisation to outflow treatment and the treatment of sharp wall features, and finally a verification and validation.

3.2. DESCRIPTION OF EXAMPLE PROBLEM DOMAINS

In this section, the physical and numerical domains are outlined wherein the HNS framework is employed. This includes the particularities of geometrical wall features. Due to their dominance in relevant aerodynamic applications such as unswept or swept wings, two respective transition problems are used as target cases, namely Tollmien-Schlichting waves developing in boundary layers with mild or no sweep (typically under zero or adverse pressure gradients) and stationary crossflow instabilities developing in swept boundary layers (typically under favorable pressure gradients).

A schematic of the four-sided numerical domain is presented in Fig. 3.1. It features an inflow plane, a (possibly curved) bottom wall, a top boundary, and an outflow plane. The semi-infinite physical domain has to be truncated to allow for a numerical assessment of the problem. This results in an imposed outflow and top boundary. To mitigate the unphysical effect of the domain truncation, an outflow buffer is implemented that prevents upstream reflections. This buffer is indicated by the shaded region in Fig. 3.1 and its implementation is considered in detail in Sect. 3.3.3. The effect of approximations at the top boundary is mitigated by ensuring a sufficiently tall domain.

Capitalising on the ability of the HNS approach to incorporate surface wall features, DeHNSSo will be subjected to two types of wall geometry features commonly found on lifting surfaces: smooth humps and sharp forward-facing steps. These two geometric features are of scientific interest as both humps and steps naturally occur on wing surfaces due to manufacturing limitations (Casacuberta et al., 2022; Eppink, 2020; Perraud et al., 2004; Rius-Vidales & Kotsonis, 2021). Their presence has been noted to significantly modify the spatial development of disturbances responsible for boundary layer transition (Franco et al., 2020). Nevertheless, the details of the interactions between incoming boundary layer disturbances and surface features remain largely unknown. DeHNSSo enables research to consider the vast parameter space determining the effect of these features on the transition process as shown in Chaps. 5 and 6.

Sharp and smooth geometric features are dealt with in different ways within DeHNSSo. Namely, smooth features are accounted for via a domain coordinate transformation while sharp features are implemented using an Embedded Boundary Method (EBM), although

¹see <https://github.com/SvenWesterbeek/DeHNSSo>

both features could be accounted for using either method, there are advantages to retaining two techniques. Specifically, accounting for smooth features using an embedded boundary technique means a loss of accuracy at the wall or a large increase in the number of grid points required. On the other hand, accounting for sharp features using curvilinear coordinate transformations introduces singularities in the transformations. In principle, the latter could be avoided by rounding sharp corners as done in Hildebrand et al. (2020), but this introduces potentially large model errors. The two methods (i.e., coordinate transformations and embedded boundary method) are fundamentally different and thus introduced separately.

A smooth hump at the wall is used as a nominal example of strong local wall curvature and can be treated using a generalised domain coordinate transformation relating the Cartesian system (x, y, z) to a generalised curvilinear coordinate system (ξ, η, z) . In the current implementation, the ξ -axis is fitted to the curved wall as indicated in Fig. 3.1 (a). This figure is the same figure as Fig. 2.6, but is reposted here for ease of reference. The hump center location is indicated by x_m . The Virtual Leading Edge (VLE) can be seen in Fig. 3.1 upstream of the stability domain's inflow location present at x_0 with the flow coming from the left.

Sharp steps are treated using an embedded boundary technique as shown in Fig. 3.1 (b), where the no-slip condition is imposed on interior grid nodes that match the step wall coordinates. The value for x_m now corresponds to the step location. The ξ -axis is thus not conformed to the step and follows the wing surface as if no step is present. The numerical grid is arranged such that the wall and the step corner can be described exactly by the grid points. Note that the coordinate transformation and embedded boundary techniques can be combined for geometries more complex than those considered here.


3.3. IMPLEMENTATION

The solution framework as described in Sect. 2.5.7 is designed with several requirements in mind; The intention is to offer a flexible and accessible method for those working in the field of stability analysis, ideally able to provide accurate results on desktop machine architectures within a reasonable time. This is greatly enabled by the efficiency and speed of the harmonic expansion, which effectively removes the need for numerical discretisation in two out of four problem dimensions (i.e., z and t), as described in Sect. 3.3.1. This reduction of the HNS system to an essentially two-dimensional numerical problem substantially reduces computational costs.

The HNS system is solved through the inversion of the equation

$$\mathbb{M}_{m,n} \hat{\phi}_{m,n} = \mathbf{r}_{m,n}, \quad (3.1)$$

sequentially for each mode (m, n) comprising the discretised form of Eq. 2.16. The state vector $\hat{\phi}_{m,n}$ contains all perturbation velocities and pressures in the domain organised as $[\hat{u}_{1,1}, \hat{u}_{1,2} \dots \hat{p}_{\Gamma_\xi, \Gamma_\eta}]^T$ for a specific mode. If the number of elements in the ξ -direction is described by Γ_ξ and likewise for η , the length of $\hat{\phi}_{m,n}$ is $4 \times \Gamma_\xi \times \Gamma_\eta$. The matrix $\mathbb{M}$ is a banded sparse square matrix of order $4 \times \Gamma_\xi \times \Gamma_\eta$ with a bandwidth, m_b , equal to $20 \times \Gamma_\eta$. The factor 20 originates from the fourth-order finite difference discretisation scheme



3.3.1. DISCRETISATION

The discretisation of Eq. 2.16 is performed differently for the streamwise (i.e., along ξ) and wall-normal (i.e., along η) directions. The wall-normal derivatives are discretised using a spectral collocation method with Chebyshev polynomial bases. Spectral collocation methods are characterised by rapid convergence as a function of the number of degrees of freedom, however, come at the cost of fuller matrices. The current implementation uses the differentiation matrix suite by Weideman and Reddy (2000). Based on the theoretical collocation point distribution, the suite provides first and second-order differentiation matrices, $\mathcal{D}_1$ and $\mathcal{D}_2$ respectively. For increased resolution in the near-wall region, the distribution of collocation points is modified and clustered near the wall, following an algebraic mapping from the computational domain $[0, 1]$ to the physical domain $[0, H]$ as proposed by Malik (1990). The mapping uses a median collocation point to define the height η_i under which half of the collocation points reside. This distribution is stretched when a step is considered to contain a near-constant $\Delta\eta$ up to a point slightly over the step height. From that point upward, the grid becomes more coarse with




Figure 3.2: Nonzero entries represented as a blue dot in the sparse banded matrix $\mathbb{M}$ with $\Gamma_\xi = \Gamma_\eta = 7$. For the first index, the shape functions $\hat{u}$, $\hat{v}$, $\hat{w}$, and $\hat{p}$ are shown that will be multiplied by the entries of that column. Likewise, M_x , M_y , M_z and the C indicate the x , y , z momentum and continuity equations respectively solved per row.

the distance from the wall.

The first and second-order streamwise derivatives in the HNS equations are discretised using a fourth-order central finite-difference scheme. This scheme becomes progressively downwind at the inflow and upwind at the outflow, as shown in Sect. 3.3.3. Similarly, the scheme becomes progressively upwind in the vicinity of sharp features for $\eta \leq h_{step}$ if present. The scheme is unaltered, i.e. remains centered, for $\eta > h_{step}$. More details on the treatment of these features are provided in Sect. 3.3.4.

3.3.2. MEMORY MANAGEMENT AND CONSTRUCTION OF THE SOLUTION MATRIX

Matrix $\mathbb{M}_{m,n}$ from Eq. 3.1 is a highly sparse banded matrix. Direct sparse matrix solvers are ideal tools for this type of problem. However, while these methods are fast and robust, memory can become a bottleneck for large problems due to a phenomenon known as fill-in when factorising a sparse matrix (Xia et al., 2023). Moreover, considering a given nonlinear stability problem simulation, Eq. 3.1 needs to be solved successively for each

of the modes in the entire mode ensemble as well as repeatedly to achieve convergence of the nonlinear interaction terms (contained in $\mathbf{r}_{m,n}$). As such, accounting for symmetries, the total memory allocation necessary can reach $\left((M+1) \times (2N+1) - N\right) \times S$ in 3D flows ($W \neq 0$) and $(M+1) \times (N+1) \times S$ in 2D ($W = 0$) flows if $\omega_1 = 0$ or $\beta_1 = 0$, where M is the maximum number of frequency (ω) modes, N is the maximum number of spanwise wavenumber (β) modes and S is the memory size of matrix $\mathbb{M}_{m,n}$. A strategy for addressing this can be devised by first rearranging Eq. 2.16 to:

$$[A]\mathbf{q}_{m,n} + [B]\omega_m\mathbf{q}_{m,n} + [C]\beta_n\mathbf{q} + [D]\beta_n^2\mathbf{q}_{m,n} + [E]\frac{\partial\mathbf{q}_{m,n}}{\partial\xi} + [F]\frac{\partial^2\mathbf{q}_{m,n}}{\partial\xi^2} = \mathbf{r}_{m,n}, \quad (3.2)$$

where matrices A-F contain base flow values and coefficients associated with wall-normal derivative operators only and are mode-independent (see A.3). Matrices A-F form the building blocks for the left-hand side matrix $\mathbb{M}_{m,n}$. Using Eq. 3.2, the left-hand side of Eq. 3.1 can be created at every iteration using

$$\mathbb{M}_{m,n} = \mathbb{M}_1 + \mathbb{M}_2\omega_m + \mathbb{M}_3\beta_n + \mathbb{M}_4\beta_n^2, \quad (3.3)$$

where the streamwise derivatives are approximated via finite-difference operators contained in $\mathbb{M}_1$. This allows common and mode-invariant matrices $\mathbb{M}_1$ to $\mathbb{M}_4$ to be recycled during the simulation. In fact, sparse matrices $\mathbb{M}_1$ to $\mathbb{M}_4$ need only to be constructed once before the beginning of the simulation and can be accessed in parallel by separate computational cores, each evaluating different modes. In summary, with the described procedure, the total memory used to store the left-hand side matrices for a given simulation is significantly reduced to S , (i.e., the size of $\mathbb{M}_1$) plus the overhead associated with storing the three additional highly-sparse matrices ($\mathbb{M}_2 - \mathbb{M}_4$). In addition, due to the avoidance of constructing the lengthy matrix $\mathbb{M}_{m,n}$ repeatedly, gains in computational time are achieved. The generalised construction of the matrices $\mathbb{M}_1$ to $\mathbb{M}_4$ can be found in A.4. The memory required to solve Eq. 3.1 scales with the numerical grid only. Specifically, it was found that the memory requirements saving $\mathbb{M}$ and solving the problem both scale linearly with Γ_ξ and quadratically with Γ_η .

3.3.3. OUTFLOW TREATMENT

A considerable challenge in incompressible NS solution frameworks is the treatment of spurious reflections from numerical domain boundaries. This is a direct result of the introduction of an outflow boundary when modelling the semi-infinite domain as a finite domain (Dobrinisky, 2003). Streett and Macaraeg (1989) note that the ellipticity in the incompressible Navier-Stokes equations allows for the upstream travel of information which arises primarily from the viscous terms and the pressure field term. Joslin (1992) notes that the interaction of local velocity perturbations with the pressure boundary causes a pressure pulse that travels upstream at infinite speed for incompressible simulations. Consequently, neither Dirichlet nor Neumann boundary conditions can be imposed directly at the outflow.

To avoid this unphysical behaviour in the present implementation, streamwise and spanwise perturbation velocities are damped within a buffer region located upstream




Figure 3.3: Streamwise dependency of the buffer function, $g(x)$, starting at $x_b = 85\%$ of the domain for a normalised domain length $x/L = [0, 1]$. The buffer has no effect for a value of 1 and maximum effect at the end of the domain with a value of $\approx 6 \times 10^{-6}$ for $\kappa = 6$.

of the outflow (see Fig. 3.1). This buffer significantly reduces the strength of reflections originating from the outflow. Damping is performed by multiplying all terms by a buffer function $G(x)$ and then adding its reciprocal $(1 - G(x))$ to the diagonal terms of matrix $\mathbb{M}$ corresponding to the streamwise and spanwise perturbation components. This adapts the equations to ensure the streamwise and spanwise perturbation components become smoothly negligible. Mass continuity ensures that the wall-normal velocity and pressure perturbation components are attenuated in a similar fashion while not over-constraining the equations. The attenuation function is adapted from Joslin (1992) as:

$$g(x) = \frac{1}{2} \left(1 + \tanh \left(\kappa \left[1 - 2 \frac{x - x_b}{x_e - x_b} \right] \right) \right), \quad (3.4)$$

where x_b and x_e refer to the starting and ending location of the buffer respectively and κ is a stretching factor, equal to 6 for the cases presented here. This function is scaled so that no discontinuity exists in the buffer function $G(x)$ at $x = x_b$. The buffer formulation then reads:

$$G(x) = \begin{cases} 1 & x < x_b \\ 1 - \frac{g(x)}{g(x_b)} & x \geq x_b \end{cases} \quad (3.5a)$$

where $g(x)$ is used to distinguish between the attenuation function shape and the buffer formulation $G(x)$. For cases presented here, this method has been effective in preventing outflow boundary interactions in the domain of interest.

Fig. 3.3 shows the streamwise dependency of the attenuation function. In addition to directly attenuating the disturbance velocities, the buffer is also applied to the nonlinear terms. This is performed via a direct multiplication with the local buffer function value. This ensures that the right-hand side of Eq. 3.1 is near-zero at the outflow.

3.3.4. EMBEDDED BOUNDARY METHOD

In the present implementation, sharp geometric wall features (e.g., a step) are introduced in a way similar to the immersed boundary method commonly used in DNS applications (see e.g., Meyer et al. (2010) and Örley et al. (2015)). As such, the grid domain




Figure 3.4: Three finite-difference stencils color-coded for several streamwise stations i around the step with dots, circles, and crosses representing internal, surface, and sub-surface grid points respectively. The red line represents the solid wall.

is not conformed to the step, but the interior nodes are adjusted to the step wall to ensure that a collocation point coincides exactly with the step corner. The left-hand side matrix $\mathbb{M}$ can now be manipulated such that the proper boundary conditions are enforced. More specifically, no-slip and impermeability conditions are imposed at points that coincide with the step wall while perturbation pressures are solved implicitly. Grid points that exist underneath the step surface do not enter the finite difference stencil as the scheme is adjusted to become progressively upwind in the proximity of the step. This shift toward a more upwind stencil is only performed for grid points below the step corner. Above the step, the stencil is unaffected, and a fourth-order central finite-difference stencil is maintained as is visualised in Fig. 3.4. The wall-normal derivative operators are unaffected by the presence of a step.

3.3.5. ITERATIVE CONVERGENCE OF THE NONLINEAR TERMS

One of the distinguishing features of the presented HNS framework is its ability to treat nonlinear stability problems efficiently. In classical NPSE approaches, the nonlinear terms can be converged at every streamwise station separately (see Bertolotti et al., 1992; Haynes and Reed, 2000; Herbert, 1993). This is made possible by the parabolization of the equations allowing them to be solved in a marching fashion. In HNS, the entire flow field must be considered per iteration. The nonlinear terms thus have to be converged from an initial guess, which is typically the solution to the linear problem. Based on this approach, Rigas et al. (2021) solved the coupled nonlinear system using an iterative Newton algorithm. This requires the nonlinear terms to be maintained on the left-hand side of Eq. 2.11 and the full Jacobian to be calculated for the nonlinear system. Although proven effective for the studied cases, its high computational cost is noteworthy. Specifically, due to the need to solve for the entire modal ensemble simultaneously, the problem scales rapidly with the number of modes considered.

An iterative approach is adopted in this work, which is independent of the total number of modes considered. The nonlinear terms are maintained as a lagged source term on the right-hand side (see Eq. 2.16) and can be iteratively constructed for each mode. This is analogous to the treatment of nonlinear terms in the NPSE, albeit considering the entire flow field. Generally, lagged approaches are less robust than, e.g., an iterative Newton algorithm. To overcome this shortcoming, an amplitude ramping procedure is employed.

AMPLITUDE RAMPING

A key challenge in the iterative construction of the nonlinear terms in Eq. 2.16 is the appearance of nonphysical and diverging solutions. Specifically, a non-linear simulation needs to be initialised with at least one perturbation mode (e.g. the fundamental mode), typically extracted from a linear eigenvalue solution. However, for particular types of instabilities, linear stability analysis can lead to unrealistically large amplitudes in the domain. In crossflow-dominated flows especially, disturbance amplitudes several times greater than the base flow freestream velocity value can be reached purely based on linear growth. Fig. 3.5 provides an example thereof for a stationary CFI mode in a swept-wing boundary layer. While the nonlinearly developing fundamental mode reaches amplitude saturation (i.e., $u'_{max} \approx 0.3$, see e.g. Haynes and Reed, 2000 and Casacuberta et al., 2020), the linear solution (initialised with the same inflow amplitude) would approximate or surpass a value of 1 within the considered domain. Initialising a nonlinear simulation with this linear solution would force the second, third, etc. harmonics to reach amplitudes of $O(u'_{max})^2$, $O(u'_{max})^3$, etc. This causes a divergence of the nonlinear terms and an eventual blow-up of the simulation. Herein lies the challenge for nonlinear convergence of HNS, as the calculation of nonlinear terms that enforce amplitude saturation relies on the presence of mode harmonics, most notably the mean-flow distortion (White & Saric, 2005).

In the present implementation, this problem is overcome by reducing the inflow amplitude for the linear simulation used to provide the initial guess for the HNS solution. This amplitude is found a posteriori such that the maximum amplitude does not exceed a pre-determined threshold (AMAX) within the considered domain. Since the simulation was linear, it can be scaled to the desired inflow amplitude directly. However, the proper amplitude threshold is difficult to predict as it depends on the significance of nonlinear dynamics. In practice, an effective threshold is easily found by trial and error, and setting it too low can at most result in slightly longer computation times. In the present work, the threshold is set to 20% of the reference velocity for all cases. The nonlinear continuation approach is then started; The inflow amplitude is slowly increased until the desired inflow amplitude is reached. The amplitude ramping is done using two modes and the MFD such that nonlinear divergence is prevented. For the test cases presented here, this was sufficient. The expansion of the spectral ensemble by the generation of new harmonics is performed under two conditions: If there is only one active mode (e.g. the fundamental mode), the second harmonic and the mean flow distortion are introduced. If two modes are already active, higher harmonics (e.g. $n > 2$) are only introduced in sequence when the goal amplitude is reached and each new harmonic is converged before a higher harmonic is introduced.

The convergence criterion for a mode (m, n) can be calculated as:




Figure 3.5: Initial (light) to final (dark) stationary crossflow instability amplitude development (based on maximum streamwise perturbation velocity) of the fundamental mode in a swept-wing boundary layer with $N = 5$ during amplitude ramping. Higher harmonics are not shown. Red and blue lines represent linear amplitude developments corresponding to the highest and lowest inflow amplitude respectively. The dashed line indicates the approximate amplitude limit over which non-physical solutions would occur in the current domain.

$$\sigma_{m,n}^k = \frac{1}{\Gamma_x} \sum_{i=1}^{\Gamma_x} \left(\frac{|A_{i,\max}^k - A_{i,\max}^{k-1}|}{A_{i,\max}^k} \right)_{m,n} \quad (3.6)$$

where the amplitude measure, A , is the peak amplitude of the streamwise perturbation velocity u' at a streamwise index i and k is the iteration counter. The convergence criterion σ is evaluated for all harmonics in z and t separately and none should exceed the criterion before proceeding. A relative amplitude convergence criterion of 10^{-6} (in units of U_{ref}) is used for the results presented here. An exception is made for the higher harmonics whose maximum amplitude can be of the order of 10^{-10} such that this criterion reaches the order of machine accuracy (10^{-16}). An absolute limit of 10^{-14} is therefore applied to the numerator of Eq. 3.6, below which the mode is always assumed to have converged.

CONTINUATION PROCEDURE

Considering the previously described amplitude ramping approach, in the first iteration the fundamental mode amplitude is computed linearly. This result is used to cal-

culate a nonlinear forcing term for the MFD and the second harmonic. This allows for the calculation of a nonlinear forcing term of the fundamental mode. During this continuation loop, new modes are introduced if their forcing exceeds the pre-determined case-dependent threshold.

The iterative scheme can be shortly summarised as

$$\mathbb{M}_{(m,n)}(A_{(m,n)}^{k=0}) = 0, \quad \mathbf{f}_{(m,n)}^{k=1} = \mathbb{NLT}(\hat{\phi}_{(m,n)}^{k=0}), \quad (3.7a)$$

$$\mathbb{M}_{(m,n)}(A_{(m,n)}^{k+1}) = \mathbf{f}_{(m,n)}^k, \quad \mathbf{f}_{(m,n)}^{k+1} = \mathbb{NLT}(\hat{\phi}_{(m,n)}^k), \quad (3.7b)$$

where $\mathbb{M}$ is the linear HNS operator, written here as a function of the initial amplitude given to the fundamental mode(s). The operator $\mathbb{NLT}$ comprises the nonlinear terms that define $\mathbf{f}^{k+1}$ based on an intermediate solution, $\hat{\phi}^k$, which should be converged before proceeding to greater inflow amplitudes. For sufficiently small initial amplitude increment, γ , $\mathbf{f}^{k-1} \approx \mathbf{f}^k$ and the problem does not diverge. However, a small γ is associated with large computation times. For simulations presented here, $\gamma = 1.1$.

The forcing term shown in Eq. 3.7 (b) is linearly extrapolated to aid in the numerical stability and speed of the iteration procedure from $k = 3$ onward. This extrapolation can be described as:

$$\mathbf{f}_{(m,n,i)}^k = \mathbf{f}_{(m,n,i)}^{k-1} + \frac{\mathbf{f}_{(m,n,i)}^{k-1} - \mathbf{f}_{(m,n,i)}^{k-2}}{A^{k-1} - A^{k-2}} (A^k - A^{k-1}), \quad (3.8)$$

for all active modes in both the current and previous iterations. In the strongly nonlinear saturated regimes within the computational domain, $\mathbf{f}^k \approx \mathbf{f}^{k-1}$ since $\mathbf{f}^{k-2} \approx \mathbf{f}^{k-1}$. Consequently, the effect is largest for lower amplitudes where the behaviour is linear in amplitude. The procedure for solving the HNS iteratively is presented as pseudo-code in algorithm 1.

3.4. VERIFICATION

Before the developed HNS framework is compared to benchmark methods such as NPSE, AHLNS, and DNS, it is desirable to establish the dependence of its output on numerical and discretisation parameters. In this section, the HNS implementation is verified by assessing the observed convergence rate and comparing it to the expected convergence rate anticipated from the discretisation schemes used in the current implementation.

As a verification test case, the nonlinear development of TS waves in the Blasius boundary layer ($M = 5$) was used, essentially replicating the conditions simulated by Herbert in his seminal PSE report (Herbert, 1993). The case will be described in detail in Sect. 3.5.1. For details on nondimensionalisation, the reader is referred to Sect. 2.5.2. In short, the flow is described by a constant external velocity $\bar{U} = 10$ m/s. The Blasius length at the inflow ($\bar{\delta}_0 = 6.075 \times 10^{-4}$ m) serves as the reference length. A global Reynolds number $Re = \frac{\bar{\delta}_0 \bar{U}_e}{\bar{\nu}} = 400$ is then defined using a kinematic viscosity of $\bar{\nu} = 1.5188 \times 10^{-5}$ m²/s. The primary mode with fundamental angular frequency is $\omega_1 = 0.0344$ is introduced at the domain inflow as the solution to the local eigenvalue problem on which an initial amplitude of $A_0 = 0.0025\sqrt{2}$ is imposed. All higher harmonics are set to 0 at the inflow and are generated nonlinearly in the interior domain. The natural forcing of the higher

Algorithm 1 HNS Solver Pseudo Code

```

1: procedure FIND A LINEAR INFLOW AMPLITUDE
2:   Inflow Eigenvalue Problem:
3:     Impose inflow condition (LST, LOAD, or ZERO)
4:     Present the LST solution on the right-hand side for the first stage
5:   Initial Linear Estimate:
6:     Solve the linear problem
7:     Find  $A_{\max}$ 
8:     Reduce the inflow amplitude  $A_0$  to ensure  $A_{\max} = \text{AMAX}$ 

9: procedure CONVERGE THE NONLINEAR PROBLEM
10: while  $A_0 \leq A_{\text{goal}}$ 
11:   while  $\epsilon > \text{nonlinear convergence threshold}$ 
12:     Calculate Nonlinear Terms:
13:       Compute nonlinear terms based on the latest iteration
14:     Solve Nonlinear Problem:
15:       Solve HNS
16:     Introduce Harmonics:
17:     if Number of active modes is less than 2 or goal initial amplitude is reached
18:       Introduce new modes
19:     end if
20:     Perform Convergence Check:
21:       Calculate the change in normalised amplitude evolution,  $\epsilon$ 
22:   end while
23:   Increase inflow amplitude as  $A_0^{k+1} = \gamma * A_0^k$ 
24:   Extrapolate forcing term
25: end while

26:   CLOSE;

```

harmonics ensures a constant phase relation to the fundamental mode which is introduced with an arbitrary phase. Aside from this phase shift, the downstream results are confirmed to not be affected by the phase of the inflow condition.

The results are subjected to variations in the streamwise discretisation (for constant $\Gamma_\eta = 100$) and the number of Chebyshev collocation points in the wall-normal direction (for constant $\Gamma_\xi = 2674$). The streamwise discretisation is equidistant while the wall-normal discretisation is clustered near the wall as described in Sect. 3.3.1. Other important parameters can affect the result through their influence on the grid, such as the median collocation point, y_i , and domain size parameters, H , L , and x_b . Preliminary convergence tests were performed (not shown here for brevity) for these parameters. Consequently, they are not varied in the convergence study and their value will be constant for all simulations: $H = 99$, $y_i/H = 0.05$, $L = 2.625 \times 10^3$, and $x_b/L = 0.85$. The convergence threshold, $\sigma_{m,n}^k$, was lowered significantly here to 10^{-8} to ensure its effect

on the outcomes is orders of magnitude smaller than the discretisation error.

Two evaluation metrics are used to evaluate the grid convergence. Firstly, an integral measure of perturbation kinetic energy E over the total physically relevant domain (x_0 x_b), between the inflow and buffer starting location respectively, is used. This can be calculated as:

$$E_t = \int_{x_0}^{x_b} \int_0^H \frac{1}{2} (|u'|^2 + |v'|^2 + |w'|^2) dy dx, \quad (3.9)$$

This integral perturbation kinetic energy measure is calculated as the modal total perturbation energy, i.e. the sum of all modes, and shown in Fig. 3.6. In addition, a local measure of perturbation kinetic energy is evaluated locally at $x = 1700$:

$$E_l = \int_0^H \frac{1}{2} (|u'|^2 + |v'|^2 + |w'|^2) dy|_{x=1700}. \quad (3.10)$$




Figure 3.6: Perturbation kinetic energy difference (ϵ) with respect to the most refined simulation as a function of Γ_x for constant $\Gamma_\eta = 100$ (a) and Γ_η for constant $\Gamma_\xi = 2674$ (b) based on E_t (black) and E_l at $x = 1700$ (red). Markers represent simulation results. The expected convergence is indicated by the dashed line.

The results are expected to converge with the order of $(\Delta\xi)^4$ as indicated by the dashed lines in Fig 3.6 (a). The absolute difference ϵ is introduced to describe the difference in integrated nondimensional perturbation kinetic energy E between the current case and the finest case. The result of HNS on the finest grid closely matches PSE results, as will

be shown in Sect. 3.5, thus allowing for its use as ground truth for the purpose of this verification. The convergence behaviour is as expected. Concerning the wall-normal discretisation, the shape functions for this case feature no singularities in the domain. Boyd (2001) states that the Chebyshev series of functions without singularities can be expected to have supergeometric convergence. For such functions, the exponential index of convergence should increase with the number of Chebyshev polynomials included as is seen in Fig. 3.6 (b). Similar convergence trends are observed for all considered cases in the present work and results are available upon request.

3

3.5. VALIDATION

The HNS framework described above was used to analyze four classical and practically relevant boundary layer stability cases. The first case considers the nonlinear growth of Tollmien-Schlichting waves in a Blasius boundary layer, outlined in Sect. 3.5.1. In the second example (Sect. 3.5.2), the nonlinear development of stationary Crossflow Instabilities (CFI) in a swept-wing boundary layer is considered. This case requires the amplitude ramping procedure described in Sect. 3.3.5. Lastly, two modifications of the latter swept-wing scenario were considered to demonstrate the implementations of smooth and sharp wall features as explained in Sect. 3.2. Additional non-homogeneous wall boundary conditions (here blowing and suction) are used to trigger the CFI mode in the final case. For each case, benchmark solutions from well-established stability frameworks and/or DNS are used for validation. The four presented reference cases collectively call all terms of the Harmonic Navier-Stokes equations (i.e. frequency and span-wise wavenumber terms) and make use of all features of the implementation (i.e., generalised coordinates, embedded boundary method, amplitude ramping, outflow buffer and non-homogeneous boundary conditions). This ensures that the current implementation is thoroughly tested. The individual cases are explained in more detail in their respective sections (i.e. Sects. 3.5.1 - 3.5.4.)

It is stressed here that the results presented in this section were calculated on grids with refinement that far exceeded the minimum requirements for sufficient accuracy. It was ensured that the chosen grid did not affect the results. As was shown in Fig. 3.6, results converge rapidly, and coarser grids still provide accurate solutions ($\epsilon \approx 10^{-5}$). A.5 additionally shows the computational requirements needed to solve the HNS on grids with medium refinement. The use of a coarser grid leads to computation times orders of magnitude lower than the refined simulations with a negligible loss of accuracy of the result.

3.5.1. NONLINEAR DEVELOPMENT OF TWO-DIMENSIONAL TS WAVES

The first case considering the nonlinear growth of Tollmien-Schlichting waves in a Blasius boundary layer has been studied in detail before in e.g. Bertolotti et al. (1992) and Chang et al. (1993) and Herbert (1997). In short, the case is defined by a constant external velocity $\bar{U}_{ref} = \bar{U}_e(x) = 10$ m/s (i.e., zero pressure gradient) and a flat plate geometry. The global Reynolds number $Re = \frac{\delta_0 \bar{U}_e}{\bar{\nu}} = 400$ is defined by the Blasius length scale at the inflow, $\delta_0 = 6.075 \times 10^{-4}$ m, and the kinematic viscosity $\bar{\nu} = 1.5188$ m²/s. The primary TS wave mode (1,0) is enforced at the inflow as a solution to the local eigenvalue stability




Figure 3.7: Amplitude evolution of TS waves in a Blasius boundary layer at $Re_0 = 400$ and $\omega_1 = 0.0344$ showing modes (1,0) (o), (2,0) (+), (3,0) (*), (4,0) (x) and (5,0) (□) and mean flow distortion (0,0) (no marker) as calculated by NPSE (dashed black) and the current HNS (red solid lines). The linear amplitude evolution of mode (1,0) is indicated by the blue dashed line.

problem. The primary mode is characterised by $\omega_1 = 0.0344 = \frac{2\pi\bar{f}\bar{\delta}_0}{U_0}$ where $\bar{f} = 90.6$ Hz.

The basic flow state used for this problem is the steady solution to the incompressible boundary layer equations initiated from the Falkner Skan Cooke (FSC) solution at the inflow. The equations were subject to Neumann boundary conditions for the velocity at the freestream and Dirichlet boundary conditions for pressure. No slip and impermeability conditions are imposed at the wall. For comparison and validation of the HNS results, the amplitude evolution of the first five harmonics and the mean-flow distortion are calculated using an in-house cross-validated NPSE solver used in Westerbeek (2020) and Zoppini et al. (2022), and Casacuberta et al. (2022). The primary (1,0) mode was introduced at the inflow with a peak streamwise perturbation velocity of $A_0 = 0.0025\sqrt{2}$, corresponding to the target test case in Bertolotti et al. (1992). All other modes are set to zero at the inflow and are then generated in the domain through nonlinear interactions.

For HNS simulations, a rectangular Cartesian grid was discretised with 2674 equidistant grid points along x , resulting in approximately 65 grid points per wavelength for the fundamental mode and 13 grid points per wavelength for the highest harmonic. In the y direction, the grid features 100 Chebyshev collocation points. The median collocation point, $y_i = 4.95$, is slightly smaller than the average boundary layer thickness in the domain based on δ_{99} . For the reference NPSE run, the domain was discretised in 200

streamwise locations with the wall-normal discretisation identical to that of HNS. The streamwise extent of the domain for both NPSE and HNS was from $x = 400$ to $x = 2487$. The HNS domain was artificially extended to account for the addition of the buffer region, ranging from $x = 2487$ to the outflow at $x = 2854$. For the present case setup and considered domain extent, the primary TS wave mode (1,0) essentially convects through both branches I and II of the linear neutral stability curve. As such, a natural decay of the primary mode can be expected downstream of $x \approx 2000$. Consequently, nonlinear terms did not need to be damped upstream of the buffer region as the amplitudes decayed significantly between the peak value around $x = 2000$ and the start of the buffer region. Furthermore, the linear evolution of the primary mode from $A_0 = 0.0025\sqrt{2}$ does not reach unphysical amplitudes, making the suppression of nonlinear terms by means of amplitude ramping unnecessary for this case. The NPSE methodology was significantly faster for this problem; On a desktop machine (see system A in A.5) it took 1.5 hours (= 12 CPU core hours) to solve the nonlinear problem with NPSE while DeHNSSo required nearly 12 hours (= 96 CPU core hours).

Fig. 3.7 shows a good agreement between NPSE and HNS. This is expected given that the difference between the two sets of equations lies in the second-order streamwise derivative terms which are very small in this case. Only minor differences are observed in the results present in the higher harmonics near the inflow, possibly related to the sequential introduction of harmonics in the NPSE procedure. However, this does not affect the downstream result as the initial amplitudes are small, resulting in a largely linearly governed perturbation development at the inflow for the primary mode (1,0). The higher harmonics rapidly converge shortly downstream of the inflow (at around $x = 800$) where a good match with HNS is seen. A strong decline of the higher harmonics' amplitudes can be seen around $x = 700$. This behaviour is not directly visible in earlier works due to the linear scaling present in the y axis, see Bertolotti et al. (1992) and Herbert (1993), or it is less pronounced as shown in Chang et al. (1993). In the latter, it is likely a result of phase differences between the various modes. The evolution of the mean-flow distortion matches closely between the two simulations as can be seen from the kink in the amplitude curve at $x \approx 1700$. This kink is the result of tracing the maximum amplitude of the mode, which for the MFD switches abruptly between two local maxima at $x = 1700$ (Bertolotti et al., 1992; Chang et al., 1993; Herbert, 1993).

3.5.2. NONLINEAR DEVELOPMENT OF STATIONARY CROSSFLOW INSTABILITIES

Next, the nonlinear growth of stationary CFI is simulated for the flow case examined in Casacuberta et al. (2021), specifically their case A. The results of steady-state DNS void of perturbations are used as the base flow to the stability problem. The case considers a flat plate geometry with an externally imposed favorable pressure gradient, obtained from wind tunnel experiments by Rius-Vidales and Kotsonis (2021). The resulting external inviscid velocity depends on the streamwise coordinate and can be approximated by the following polynomial fit:

$$U_e(x) = 1.8574 + 0.5303 \ln(x) + 0.1752 \ln^2(x) + 0.0377 \ln^3(x) + 0.0023 \ln^4(x). \quad (3.11)$$




Figure 3.8: Amplitude development (a) of the first (o), second(+), third(*), fourth(x) and fifth ($\square$) crossflow instability modes and the mean-flow distortion (no marker) as calculated by HNS (black dashed), DNS (magenta dotted) and NPSE (red solid lines). The linear solution is indicated by the dashed blue line. The streamwise perturbation shape is shown for $x =$ (b) 651, (c) 829, (d) 1009, (e) 1189, and (f) 1359.

normal to the leading edge, while the external spanwise velocity $W_e = -1.24$ is constant over the domain. The simulation domain starts at 5% of the wing chord length (corresponding to $x = 220\delta_0$), where $\delta_0 = 2.14 \times 10^{-4}$ m. This results in a global Reynolds number $Re = \delta_0 \bar{U}_0 / \bar{\nu} = 220$ for $\bar{\nu} = 1.47 \times 10^{-5}$ m²/s and $\bar{U}_0 = 15.1$ m/s.

The domain ranges from $x = 220$ to $x = 1560$ and has a domain height $H = 89$. It is discretised by 1272 equidistant grid points in x , resulting in approximately 35 grid points per wavelength for the fundamental mode and 8 grid points per wavelength for the highest harmonic. The grid is defined on 100 collocation points in y clustered near the wall with the median collocation point present at $y_i = 4.5$. The outflow buffer starts 15% upstream of the outflow boundary from $x = 1360$. The nonlinear forcing terms are instead suppressed from 20% upstream of the outflow.

The primary stationary CFI mode (0,1) is characterised by a spanwise wavenumber $\beta_1 = 0.18$ corresponding to a spanwise wavelength of 7.5 mm and $\omega_1 = 0$. The enforced perturbation shape at the inflow, $\hat{\phi}$, is found by solving the local linear stability eigenvalue problem. A peak streamwise perturbation amplitude of $A_0 = 3.5 \times 10^{-3}$ is then

imposed on the solution. All higher harmonic modes are set to zero at the inflow and rise naturally from nonlinear interactions downstream.

A linear calculation of the development of the primary mode (0,1) results in large amplitudes for this case as explained in Sect. 3.3.5. The nonlinear iteration procedure diverges rapidly if this initial linear guess is used directly. Hence, amplitude ramping was required for this case. The initial amplitude was reduced to $A_0 = 2.6 \times 10^{-4}$, a factor 13.3 lower than the goal amplitude. The inflow amplitude was then gradually increased by $\gamma = 10\%$ per iteration. Consequently, 27 iterations were required to converge the nonlinear terms.

The results are compared to two benchmark simulation methods: DNS and NPSE. Both methods were found to be in good agreement in earlier work (Casacuberta et al. (2022)). DNS were performed using the conservative finite volume solver INCA (see Hickel and Adams (2008) and Hickel et al. (2014)). Herein, the domain was discretised in 6760 cells in the streamwise direction. In the wall-normal direction, the domain is discretised in 576 cells. Lastly, 72 cells were present in the spanwise direction. INCA features a fifth-order upwind scheme for the convective terms while viscous terms are discretised using a second-order central difference scheme. It took a total of 8.8 days to find the solution using 2000 CPUs in parallel for a total of approximately 500.000 CPU core hours. The NPSE simulations were performed using the same basic state calculated in INCA. The flow was interpolated on the numerical domain discretised in 80 Chebyshev polynomials in the wall-normal direction and 500 streamwise stations. NPSE found the solution in approximately 3 hours (24 CPU core hours), whereas DeHNSSo took nearly 4 days (768 CPU core hours) on a desktop machine (see system A in A.5).

Using both the HNS and NPSE, five stationary CFI modes ($M = 0, N = 5$) and the mean flow distortion mode (0,0) were computed and are shown alongside DNS results in Fig. 3.8. Furthermore, the absolute perturbation shapes shown in Fig. 3.8 (b-f) demonstrate a close match between all methods throughout the domain. The secondary lobe present in the shape functions displayed in 3.8 (f) has a nonlinear origin and is captured by all three methods that accounted for nonlinear interactions.

3.5.3. LINEAR DEVELOPMENT OF STATIONARY CROSSFLOW INSTABILITY OVER A HUMP

Smooth surface protuberances (i.e., humps) are unavoidable on wing surfaces (Worner et al., 2003). The presence of such humps is known to lead to a local acceleration of the base flow (A. Boiko et al., 1990), which in turn affects the stability of the boundary layer. Moreover, classical stability methods have repeatedly been shown to be unreliable in such cases, see Cooke et al. (2019) and Franco et al. (2018), and Westerbeek et al. (2023). Here, the stability of a three-dimensional boundary layer featuring a hump is considered to validate the implementation of the generalised transformation approach in the present HNS framework. The examined case was previously studied in Westerbeek et al. (2023), there denoted case B1, to which the reader is referred for a more detailed (nonlinear) analysis.

The flow setup is defined by the reference Blasius length at the inflow, $\bar{\delta}_0 = 2.14 \times 10^{-4}$ m, and reference velocity $\bar{U}_0 = 15.1$ m/s. This results in $Re = 220$. The domain ranges from $x = 220$ to $x = 2165$ and has a height $H = 89$. The external flow is identical to the




Figure 3.9: Linear CFI amplitude development over the full domain (a), around the smooth hump (b), and growth rate (c) as calculated by the current HNS framework (solid, black), and the AHLNS (green, dashed). The flat-plate reference (i.e., no hump present) CFI amplitude development is shown in red.

previous validation case of Sect. 3.5.2. The external velocity normal to the leading edge thus follows Eq. 3.11 and the spanwise velocity is constant over the domain, namely $W_e = -1.24$. The hump is described as a smooth surface protuberance with a width b of $39.9 \delta_0$ and a height h of $2.5 \delta_0$. The wall is described by

$$y_{\text{wall}}(x) = h \exp \left[- \left(\frac{x_m - x}{b} \right)^2 \right], \quad (3.12)$$

with the hump center located at $x_m = 859$. The buffer starts from $x = 1873$.

Due to the presence of the hump, the boundary layer equations cannot reliably predict the evolution of the base flow. As such, the base flow solution was found using the finite-element solver COMSOL (*COMSOL Multiphysics® v. 5.6 user guide*, 2022). The setup featured second-order elements for velocities and first-order elements for pressure. The grid was clustered near the wall using 115 elements in the wall-normal direction and additionally refined around the hump in the streamwise direction, in which 945 elements are present. Two elements are present in the spanwise direction in which peri-




Figure 3.10: Grid lines up to $y = 30$ indicated by the black lines. One in every fifteen ξ and one in every three η lines are shown.

odic boundary conditions were imposed. A top pressure condition was imposed as was done for the swept flat plate case in Sect. 3.5.2 (see Casacuberta et al. (2021)). At the top boundary, velocities were subjected to a homogeneous boundary condition for the second-order derivatives. At the inflow, a Falkner Skan Cooke boundary layer profile was used, corresponding to the local pressure gradient. No-slip conditions were imposed on the (curved) wall and a static pressure equal to the local top boundary pressure was imposed at the outflow. The static pressure at the top boundary follows directly from the external velocity presented earlier in Eq. 3.11. As the purpose of this case was merely to validate the implementation of the generalised transformation approach, the hump was made shallow to avoid flow separation. The base flow solution was interpolated onto the numerical grid via *COMSOL for Matlab* which uses COMSOL's built-in interpolation algorithm. This interpolation has the accuracy of the used elements (i.e., second-order for velocities and first-order for pressures).

The HNS domain was discretised with 2000 grid points in the streamwise direction. The wall-normal direction was represented by 100 collocation points with the median collocation point at a constant distance of $4.45\bar{\delta}_0$ from the wall. This discretisation results in approximately 48 grid points per wavelength. The grid was conformed to the hump as shown in Sect. 3.2, while the η coordinate remains parallel to the global y -axis for HNS computations. The grid can be seen in Fig. 3.10.

A limited selection of modelling tools that can solve this problem nonlinearly are available. Hence, the specific validation is performed linearly. Even in a linear formulation, the PSE framework could suffer from limitations in the vicinity of the hump, resulting from the slowly-varying flow assumption. To avoid such modelling errors, the Adaptive Harmonic Linearised Navier-Stokes (AHLNS) framework was used as a reference. For linear computations, this tool is equivalent to the current HNS. The adaptivity of the AHLNS comes from the (re-)introduction of the streamwise wavenumber in the perturbation ansatz, to reduce the grid requirements when modal perturbation behaviour is found. In contrast to PSE, however, no assumptions on the importance of higher-order




Figure 3.11: Grid lines indicated in black. One in every forty ξ and one in every eight η lines are shown. The step surface is indicated in red.

streamwise gradients are made. For details on AHLNS, the reader is referred to Franco et al. (2020), Franco et al. (2018).

The AHLNS calculations were provided by J.A. Franco and S. Hein of the German Aerospace Center (DLR) using the same base flow. Fig. 3.9 (a) shows the primary mode's linear amplitude evolution. A zoomed-in view of the amplitude development around the hump can be seen in Fig. 3.9 (b). A slight stabilization can be seen downstream of the hump. Nevertheless, the total effect on the amplitude is minimal. The growth rate as seen in Fig. 3.9 (c) shows the effect of the hump on the disturbance evolution and the good agreement between the two simulation frameworks. The small oscillations in all growth rate curves are an artifact of the calculation of growth rates from amplitude data.

3.5.4. LINEAR DEVELOPMENT OF STATIONARY CROSSFLOW INSTABILITY OVER A FORWARD-FACING STEP

The difficulty of maintaining a laminar boundary layer in operational environments is partially caused by manufacturing limitations resulting in (critical) steps on the wing surface. The experiments of Perraud and Seraudie (2000) and later Wang and Gaster (2005) showed that the forward-facing step is less critical, i.e. less likely to advance transition than the backward-facing step for the same step height. When a step cannot be prevented at a junction, a forward-facing step becomes the preferred solution as a result. The final validation case will therefore study the application of HNS to solve the stability problem of stationary CFI, interacting with a forward-facing step.

The flow around a step is complex and solving solely for either the base flow or its stability can be challenging. Specifically, separation bubbles can form both upstream and downstream of the step as shown by Wilhelm et al. (2003) in two-dimensional flows. For the purpose of the present validation, the base flow, as well as the developed flow used here for comparison purposes, are produced using the conservative finite-volume solver INCA, described in Sect. 3.5.2. The case is based on the experiments of Rius-Vidales and Kotsonis (2021) who performed a detailed experiment of the topological changes of the

transitional flow field with a step compared to a reference case (i.e., no step).

One of the desired features of DeHNSSo is the ability to simulate non-homogeneous boundary conditions, which can be efficiently used in studies of receptivity or flow control. To demonstrate this functionality, a region of wall blowing and suction is employed slightly upstream of the step to force the dominant CFI mode. This steady blowing and suction has a sinusoidal spanwise dependence with a wavelength equal to that of the primary mode. The suction and blowing strength depends on the streamwise coordinate described by:

$$v_{\text{wall}}(x) = \begin{cases} A_{\text{bs}} \left(4 \frac{(x-x_1)-(x_2-x)}{(x_2-x_1)^2} \right)^3 & \text{if } x > x_1 \vee x < x_2 \\ 0 & \text{if } x \leq x_1 \vee x \geq x_2, \end{cases} \quad (3.13)$$

where the blowing and suction is active from $x_1 = 709$ to $x_2 = 738$ with a strength of $A_{\text{bs}} = 5 \times 10^{-5}$ defined as the imposed wall-normal peak velocity.

The simulation domain is defined by the intervals $x \in [685 \ 993]$ and $y \in [0 \ 89]$ with the Blasius length $\bar{\delta}_0 = 2.14 \times 10^{-4}$ m. The reference velocity is 15.1 m/s. This leads to a global $Re = 220$. The external velocity normal to the leading edge follows Eq. 3.11 with a constant spanwise velocity of $W_e = -1.24$. The buffer was initiated at 85% of the domain ($x = 947$).

The interaction of an incoming stationary CFI with a sharp forward-facing step results in strong local gradients that demand a finer grid than the previously described three cases. For the present case, a total of 4000 grid points were used in the streamwise direction, refined around the forward-facing step using a Gaussian function resulting in a minimum discretisation step size $\Delta x_{\text{min}} = 2.1 \times 10^{-2}$ at the step and a $\Delta x_{\text{max}} = 8.5 \times 10^{-2}$ at $20 \bar{\delta}_0$ away from the step. For reference, the streamwise wavelength upstream of the step is $40 \bar{\delta}_0$. Therefore, there exist 469 grid points per wavelength in the coarse region. This is increased to approximately 2000 at the step. In the wall-normal direction, 350 collocation points were used. The collocation points were clustered around the step height throughout the domain with a $\Delta y_{\text{min}} = 3.37 \times 10^{-3}$ increasing to $\Delta y_{\text{min}} = 0.9$ at $y = 70$. This resulted in a grid of 1.4 million points. The grid is visualised in Fig. 3.11. It took DeHNSSo nearly 6 hours on 48 CPU cores (288 CPU core hours) to find the linear solution to the problem. INCA took 8.8 days on 2000 CPU cores (422400 CPU core hours) to find the solution. Moreover, INCA solved the Navier-Stokes equations nonlinearly resolving an additional five higher harmonics, although they were found to be negligible due to their low amplitude.

Due to the large number of grid points required to converge this case, the choice was made to consider the linear development of the stationary crossflow instability governed by $\beta_1 = 0.18$ only. No amplitude ramping was thus necessary.

A comparison with DNS in terms of amplitude and shape functions can be seen in Fig. 3.12. It is noted that although the DNS did not neglect nonlinear terms, it was ensured that linear dynamics dominate the behaviour of the perturbations by imposing a sufficiently small blowing and suction amplitude. Good agreement with DNS is seen in which a near-wall structure is accurately captured in the downstream vicinity of the step. This near-wall structure is important for both the local and downstream development of the incoming perturbation as was also found in similar cases experimentally in




Figure 3.12: Linear interaction of a stationary CFI initiated by blowing and suction around $x = 709$ and a forward-facing step present at $x = 738$, indicated by the dashed vertical line. Results are shown from DNS (magenta, dotted), HNS (black, solid), and AHLNS (green, dashed). Maximum perturbation amplitude development is shown in (a) and absolute streamwise perturbation shape is shown for $x =$ (b) 849, (c) 858, (d) 860, (e) 864, and (f) 867.

Eppink (2020) and Rius-Vidales and Kotsonis (2021) as well as numerically in Casacuberta et al. (2022). Fig. 3.13 shows the streamwise perturbation velocity contours of the primary mode around the step. Here, too, the near-wall structure is clearly observed slightly downstream of the step corner.

3.6. CONCLUDING REMARKS

The present chapter outlines the development and implementation of an efficient non-linear Harmonic Navier-Stokes solver for the simulation of boundary layer and flow stability problems in complex geometrical domains. Smooth wall protuberances as well as local grid refinements are accounted for using a generalized domain transformation. Sharp geometries are accounted for through the use of embedded boundaries. The developed HNS solver proved accurate, fast, and robust in predicting the (non)linear stability of TS waves and stationary CFI developing on flat plate geometries or interacting




Figure 3.13: Real streamwise perturbation velocity around the step with the dashed line indicating the δ_{99} of the base flow and the gray box represents the step.

with a localized hump or forward-facing step.

The developed HNS is envisioned to be a valuable tool for parametric studies of stability calculations in flows around steep humps, steps, and cases involving receptivity to inflow or wall forcing and non-modal growth. For these cases, PSE frameworks could fail to produce reliable and accurate results due to inherent assumptions of a slowly varying modal perturbation evolution. While DNS could be used to perform similar work, the high computational cost makes it impractical for many applications. The current implementation allows for rapid nonlinear stability analysis in flows featuring inhomogeneous boundary conditions (e.g. blowing and suction), humps, and steps, enabling parametric studies of detailed wing geometries. The results can provide valuable insight into the interaction of incoming perturbations with surface features to further assist engineers and scientists in the pursuit of laminar wall-bounded flows.

4

RECEPTIVITY OF CROSSFLOW INSTABILITY TO DISCRETE ROUGHNESS AMPLITUDE AND LOCATION (EXPERIMENTS BY G. ZOPPINI)

*The more clearly we can focus our attention on the wonders and realities of the universe
about us, the less taste we shall have for destruction.*

Rachel Carson

The NPSE are used to find the equivalent initial amplitude of crossflow instabilities in lab experiments where Discrete Roughness Elements (DREs) are used to force stationary CFI on a swept wing. An amplitude sweep of these CFI is performed with NPSE. The initial amplitude of the CFI, at the DRE-location, is found via a match of the downstream primary perturbation amplitude and shape function. This process is repeated for several DRE heights. Increased DRE heights are found to correspond to a greater equivalent initial CFI amplitude. However, the discovered relation between forcing configuration and the equivalent instability amplitude is not linear. It is hypothesized that transient perturbation mechanisms are responsible for the observed initial amplitude trend. For the highest DRE amplitudes the match with experimental results deteriorates, which is likely the result of neglecting the unsteady perturbations and wing curvature in the numerical framework.

Parts of this chapter are published in
Zoppini, G., **Westerbeek, S.**, Ragni, D., Kotsonis, M. (2022) Receptivity of crossflow instability to discrete roughness amplitude and location. *Journal of Fluid Mechanics*, 939, A33.

4.1. INTRODUCTION

Swept wings play a pivotal role in modern aviation. Understanding the intricate dynamics of Crossflow Instabilities on these wings is crucial for enhancing their aerodynamic performance. Commonly, wind tunnel tests are an important part of the wing design process and research. The use of Discrete Roughness Element (DRE) is ubiquitous in such tests as a method to incite perturbations in the boundary layer with the desired characteristics (e.g. spanwise wavelength and amplitude), see e.g. the works of Reibert et al., 1996, Reibert et al., 1996 and Serpieri and Kotsonis, 2016. The relation between the shape and size of the elements and the initial instability amplitude is subject to debate. However, the small scales associated with the flow features surrounding DREs complicate a fully experimental approach to determine the relation.

This chapter explores the use of NPSE to estimate the corresponding onset amplitude of CFI at the location of the discrete roughness elements on swept wings. By integrating NPSE with experimental and theoretical methodologies, a link and comprehensive understanding of the intricate interplay between roughness elements and crossflow instabilities is demonstrated.

The numerical part of the investigation builds upon experimental work that explores the effect of discrete roughness elements on the development and breakdown of stationary crossflow instabilities. Zoppini et al., 2022 focused on assessing the receptivity of swept wings to various roughness element heights and chordwise locations, shedding light on the nuanced relationship between primary instability amplitudes as well as the effect of roughness size and chordwise location. Experimental measurements employing infrared thermography and planar Particle Image Velocimetry (PIV) were used in conjunction with simulations based on NPSE. As part of this thesis, the setup, role, and findings relating to the use of NPSE for the estimation of initial perturbation amplitudes are considered. The methodology and its limitations are covered here in more detail with respect to the original work.

4.1.1. DOMAIN AND FLOW DESCRIPTION

The experimental measurements presented in Zoppini et al., 2022 were conducted in the low-speed Low-Turbulence wind Tunnel (LTT) at TU Delft. The problem details are summarized here completeness. The reader is referred to Zoppini et al., 2022 for a complete overview and conclusions relating to flow physics.

To simplify the comparison with experimental results, dimensional units are used in the original work (Zoppini et al., 2022), although streamwise coordinates are normalized by the chord length c . The chord, c , is 1.27 meters, measured in the freestream direction. Measured normal to the leading edge, $c_\perp$, the chord is 0.9 m long. Here, the focus lies on numerical results and nondimensional units are preferred for comparison. However, the figures from the original paper are not altered, so in this section both the dimensional (indicated with an overbar) and nondimensional units are presented. The following reference units are used for the nondimensionalisation: The boundary layer thickness at the inflow, $\bar{\delta}_0 = 2.18 \times 10^{-4}$ m, the freestream velocity, $\bar{U}_{\text{ref}} = 18.4$ m/s, and the kinematic viscosity, 1.45×10^{-5} m²/s. This results in a global Reynolds number based on the boundary layer thickness at the inflow of $Re = 264.7$.

The considered flow develops over the in-house designed M3J 45°swept-wing model




Figure 4.1: M3J airfoil shape (no marker), pressure coefficient (red), and its streamwise derivative (blue) from Zoppini et al., 2022.

with a constant chord. This wing model, detailed in Serpieri and Kotsonis, 2015, features a favourable pressure gradient up to 65% of the chord. Test conditions were performed with the wing at an angle of attack of -3.36° at a chord-based Reynolds number of 2.17×10^6 , based on the velocity in the freestream direction and the chord length of 1.27 m. The wing is symmetrical, hence the negative angle of attack is equivalent to performing measurements on the pressure side with a positive angle of attack. However, this setup allowed for easier (optical) access in the experiments. Fig. 4.1 shows the M3J airfoil shape and the pressure coefficient for the current angle of attack measured at the surface.

At these wind tunnel conditions, the turbulence intensity (Tu) was sufficiently low ($Tu/U_\infty \approx 0.025\%$) to allow stationary crossflow waves to dominate the stability and transition scenario. This implies the formation of a spanwise invariant boundary layer with its transition scenario primarily governed by the stationary CFI (Serpieri, 2018). Hence, NPSE simulations were performed, focusing solely on primary CFI characterized by $\omega = 0$. The spanwise periodicity of the DRE array was chosen to be 8 mm as the stationary crossflow instability corresponding to this spanwise wavelength was found to reach the largest amplitudes within the PIV domain.

The height of the DREs is described by k . The roughness-based Reynolds number was calculated with the roughness height in the unperturbed boundary layer as $Re_k = k|u(k)|/\nu$ where the streamwise velocity $u(k)$ is the velocity at a distance k from the wall (see Gregory et al., 1956; Reibert et al., 1996 and Kurz and Kloker, 2016). Different values of Re_k are considered by varying the DRE height. This is described in Tab. 4.1.

Measurements using PIV were taken in zy planes aligned with the leading edge. This means that only the spanwise and wall-normal velocity components were measured. Hence, as opposed to the fully numerical studies in this thesis, amplitudes are calculated based on the spanwise velocity component w in this chapter. A spanwise spatial Fourier analysis was used to discern between various modal contents of the signal. The tracing of

Table 4.1: Geometric parameters of DRE arrays from Zoppini et al., 2022.

Case	λ (mm)	D (mm)	k (mm)	k/D	Re_k range
k1	8	1.772 ± 0.017	0.1147 ± 0.0023	0.083	16-31
k2	8	1.732 ± 0.008	0.2179 ± 0.0031	0.126	63-101
k3	8	1.767 ± 0.004	0.3292 ± 0.0009	0.186	139-185
k4	8	1.721 ± 0.009	0.4374 ± 0.0035	0.261	235-302

these signals for the various streamwise locations provides an estimate of the amplitude development of the individual modes. The various modal contents of the velocity fields are denoted R_1 when only the mode (0, 1) is included and, for example, $R_{1,2,3}$ when the modes (0, 2) and (0, 3) are included as well.

4

4.1.2. METHODOLOGY

First the solution to the base flow problem was found, which was used in the flow stability problem. The receptivity process was not resolved numerically and its effects were instead captured via the matching of the downstream perturbation development. A CFI mode was directly imposed at the inflow of NPSE simulations as opposed to imposing the perturbation field observed from experiments. In this case, the base flow is unaffected by the choice of DRE configuration, i.e. changes in Re_k , and the DRE wake are not accounted for in the base flow.

BASE FLOW

Stability solutions require a known reference base flow (as explained previously in Sect. 2.5.4). In this case, the steady incompressible solution of the spanwise invariant Boundary Layer equations was used. Specifically, an in-house solver for the boundary layer equations (see Serpieri, 2018) employing the experimentally acquired free-stream flow characteristics and pressure distribution was used. Airfoil curvature was neglected during this procedure. The inflow profile was the solution to the local Falkner-Skan-Cooke equation with the imposed local pressure gradient Hartree parameter. No-slip boundary conditions were imposed at the wing surface. At the top boundary, the streamwise and spanwise velocities were prescribed corresponding to the experimental pressure measurements. These approximate the values presented in Fig. 4.1. The external spanwise velocity, $\bar{W}_e$ was kept constant at 18.14 m/s, i.e. $W_e = 0.986$. In reality, the external spanwise velocity can vary due to wind-tunnel blockage effects. However, the variations of W_e were small and were therefore ignored in the numerical model. The V velocity was left free to allow for boundary layer growth.

The boundary layer equations are discretized using a second-order finite-difference approach in the streamwise direction and spectral differentiation in the wall-normal direction. In the streamwise direction, the domain is discretised by 750 equispaced points. The wall-normal direction is discretised using 100 Chebyshev collocation points. The system of equations is solved in a marching fashion in the streamwise direction. Details on the assumptions and solving procedure can be found in Serpieri, 2018.




Figure 4.2: Sketch of real wing (M3J) showing the location of the DRE array and the numerical domain. The flow is coming from the left at an angle of attack α . The start of the numerical domain depends on the location of DRE at a distance x_0 from the leading edge. The boundary layer originates from upstream the numerical domain as indicated by the dashed line cutting through the numerical domain. Wing curvature is not maintained in the numerical domain as indicated by the straight wall.

FLOW STABILITY SIMULATION

Equations 2.8 were solved for the base flow calculated as described in the previous section (Sect. 4.1.2) on a rectangular domain. The domain ranges from the (variable) DRE location ($x/c = 0.02$ to $x/c = 0.3$, i.e. $x = 82.5$ to $x = 1237$) depending on the case to the final PIV measurement plane at $x/c = 0.35$, $x = 1443$. The domain height is 1 cm or $H = 45.8$. The numerical domain was oriented at a 90-degree angle to the leading edge, which constitutes the x -direction. The numerical domain and its relation to the real wing used in the experimental setup is shown in Fig. 4.2

At the inflow plane, the solution to the local stability eigenvalue problem (see Eq. 2.6) with $\beta = 0.1714$ and $\omega = 0$ was imposed. The eigensolution was scaled to achieve a specific peak spanwise perturbation component for comparison with experimental results. Brackets or a subscript containing the respective w -field (defining the spectral cut-off) are added to the symbol for the amplitude, A , to indicate this.

To ensure that the spectral cut-off does not affect the results, seven harmonics ($N = 7, M = 0$) and the mean flow distortion mode $(0, 0)$ were accounted for in the nonlinear simulations. Tests containing more higher harmonics showed that a negligible part of the perturbation kinetic energy is contained in these higher modes ($N > 7$) and do not significantly affect the amplitude evolution of the first seven harmonics in the problem.

MODELLING PROCEDURE

The flow behaviour and perturbation development in the near-wake of the DRE array cannot be accurately measured in the experimental setup. Instead, the downstream perturbation development was tracked using PIV at several streamwise locations (see Zoppini et al., 2022). For some forcing conditions, the perturbation amplitude in the experimental measurement domain reaches levels where nonlinear interactions become significant. Consequently, linear methods (LPSE, OS, or adjoint methods) would pro-

vide inaccurate estimates of the amplitude upstream of the measurement domain. This necessitates the use of nonlinear stability methods to find an *equivalent* amplitude at the DRE location. The wing (surface) geometry downstream of the DRE array is smooth, lacking any surface features that could affect the perturbation development. Therefore, the NPSE framework is well-suited to provide stability downstream of the DRE wake. The input amplitude must be found through trial and error as the nonlinear nature of the downstream evolution complicates upstream marching. The following approach is thus followed instead:

4

1. Solve the base flow starting from experimental measurements on a simplified rectangular domain.
2. Solve the linear perturbation development of mode (0, 1) from the DRE location to the end of the experimental measurement window.
3. Fit the linear amplitude development to the most upstream station of experimental amplitude data.
4. The amplitude from the previous step forms the upper bound of primary mode inflow amplitude, $A_{0,\max}$.
5. Solve the nonlinear PSE equations with up to mode (0, 5) for 20 inflow amplitudes between $A_{0,\max}/10$ to $A_{0,\max}$.
6. Determine the closest matching downstream perturbation evolution by minimizing the squared differences between the numerical and perturbation shapes (including amplitude) of the mode (0, 1) across the entire PIV domain ($0.25 < x/c < 0.36$). Every measurement plane is given an equal weighting in the least mean squares minimization process.
7. If necessary, run further amplitude sweeps until a close match is found between experimental and numerical results.
8. Find the initial amplitude of the closest match between numerical and experimental results at the DRE location.

The numerical component of this research involves generating a data set for the stability of the base flow described in Sec. 4.1.2. The fundamental mode (0, 1) with a spanwise wavenumber $\beta = 0.17$ based on ($\lambda_z = 36.7$) is imposed with varying amplitudes such that the downstream results can be matched with experimental results. In the experimental investigation, the DRE height k was varied that affects the CFI amplitude shortly downstream in an otherwise similar flow. A single numerical amplitude sweep could thus be used to find the initial amplitude estimation per DRE location for several DRE heights.




Figure 4.3: Integral perturbation amplitude along x from (a–c) $\bar{w}_{R1}$ fields and (d–f) $\bar{w}_{R2}$ fields. Columns refer to a fixed k , matching NPSE solution (– lines), NPSE solution sensitivity (shaded regions), NPSE initial amplitude estimations (full markers). Data relative to x/c locations downstream of the transition point are excluded from the plot. Figure taken from Zoppini et al., 2022.

4.1.3. RESULTS

A nonlinear stability development dataset was successfully created using NPSE for the current study. Following the approach of 4.1.2, one of the results in the dataset was selected per case, i.e., per DRE height (k) and location, that matched the nonlinear perturbation development in the experimental measurement domain.

Fig. 4.3 shows the experimental and numerical amplitude development for all cases. Since the dataset contained results at intervals of the inflow amplitude, an uncertainty range in the results exists which is indicated by the shaded region. Close matches were found in nearly all cases. High-amplitude cases in which nonlinear saturation occurred suffered from the lack of curvature effects in the numerical approach. These are known to significantly stabilise the perturbation development, see Haynes and Reed, 2000. The saturation level of the instabilities could thus be affected by curvature. For example, the experimental results of Figs. 4.3 a-c show a saturation of the amplitude in the experimental results near the domain outflow, while the numerical perturbation development shows continued growth. However, the saturation amplitude was also found to be dependent on the forcing condition. Namely, higher DRE heights lead to lower saturation amplitudes. The receptivity problem was not solved numerically but instead modelled by adjusting the amplitude of the solution to the local eigenvalue problem. As a result, NPSE cannot capture the effect of the forcing condition on the downstream perturbation development. The effect of the DREs on the base flow and non-modal effects were also not accounted for. The result of these processes are instead visible in the equiva-




Figure 4.4: Particle image velocimetry and NPSE mode (0,1) (a) N-factor and (d) amplitude curves for $Re_k = 24$ (o) and $Re_k = 18$ (*) with x_{DRE}/c (-.-), $A_{w,0}$ (x). (b-e) Shape functions of the mode (0,1) and (c-f) (0,2) mode at $x/c = 0.30$ for the NPSE (full line), LPSE (dashed), and Fourier shape functions from experimental data (symbols). Experimental δ_{99} (horizontal dashed line), NPSE matching sensitivity is shown via shadowed areas. Figure taken from Zoppini et al., 2022.

lent amplitude calculation. Moreover, unsteady perturbations were not incorporated in the current analysis where their effect might be significant directly or indirectly (through nonlinear effects) on the transition scenario and perturbation development. These likely play a more important role as the Re_k increases (Breuer and Kuraishi, 1994; Lucas, 2014 and Loiseau et al., 2014). Hence, it is hypothesised that neglecting unsteady perturbations deteriorates the match for the higher Re_k cases, especially for the perturbation field of the second harmonic $\bar{w}_{R2}$ (see Fig. 4.3 e and f). Further research will be necessary to find the source of the mismatch present in some cases.

The results are further scrutinised in Fig. 4.4, which shows the amplitude and N-factor development as well as the evolution of the shape functions of LPSE, NPSE, and experimental measurements for cases $Re_k = 24$ and $Re_k = 18$ placed further upstream and downstream, respectively, as indicated by the red crosses. At the location of the red crosses, the equivalent amplitudes are extracted, allowing observation of the perturbation development downstream. A good match is observed between experimental and numerical results. The linear PSE accurately predicts the amplitude development for the downstream forcing case ($Re_k = 18$). However, the nonlinear terms were found to dominate the perturbation behaviour for the cases shown at $x/c = 0.3$ as seen by the mismatch between experimental and LPSE perturbation shapes in Figs. 4.4 b, c, e, and f. In Fig. 4.4 c, a slight mismatch in the shape function is seen where the perturbation velocity does not change monotonously. However, from the uncertainty region, it can be seen that the exact perturbation shape is highly sensitive to the inflow amplitude. The mismatch in the shape function is therefore assumed acceptable and does not appear to affect the amplitude development as shown in Fig. 4.4 a and d.

The amplitudes found using the equivalent amplitude methodology point toward a nonlinear relationship between forcing configuration and the equivalent instability am-

plitude. The trend points towards non-modal mechanisms governing the total perturbation energy in the near-wake of the DREs. Inspired by these results, follow-up work by Zoppini et al., 2023 discovered transient decay in the vicinity of the DRE. Transient decay follows the same flow mechanisms as transient growth, but results instead in a decay of perturbation kinetic energy. Further details of the physical implications of these results are out of the scope of this work. For a more complete overview of the flow physics, the reader is referred to Zoppini et al., 2022 and Zoppini et al., 2023.

4.1.4. CONCLUSION

The nonlinear development of primary CFI was modelled through the Nonlinear Parabolized Stability Equations (NPSE). The solution successfully matches the experimental observations of instabilities as they develop downstream from an introduction at the virtual DRE location. The perturbation response was often nonlinearly dominated in the boundary layer where experimental measurements occurred. This necessitated the use of NPSE for the upstream extrapolation of the perturbation development.

Differences between numerical and experimental results were found in some cases that featured large DRE forcing conditions. Not incorporating the base flow deformation by the DREs, the inability to solve for non-modal growth, incorporating only steady disturbances in the analysis (i.e., $\omega = 0$), and neglecting the airfoil curvature are hypothesized to be responsible for the observed differences. However, the trends appear unaffected by this.

The upstream initialisation of the numerical solution provides a nonlinear estimate of the equivalent initial amplitude at the DRE array location. The equivalent amplitude methodology proved powerful in discovering trends that suggest non-modal growth in the near-element flow, despite not being able to solve the perturbation behaviour in the flow region that contained this non-modal growth. This was later confirmed in Zoppini et al., 2023.

5

APPLICATIONS OF FLOW STABILITY FRAMEWORKS IN FLOWS OVER FORWARD-FACING STEPS (DNS BY J. CASACUBERTA)

If you know you are on the right track, if you have this inner knowledge, then nobody can turn you off-no matter what they say.

Barbara McClintock

The application of PSE and HNS frameworks are used to assess the stability of flows forward-facing steps. The PSE are used to confirm that the slowly-varying flow assumptions are only violated in a very short streamwise extent around the step. In this region, the PSE fails to capture near-wall structures found in the DNS results. The HNS was then used to overcome these limitations. The HNS framework accurately captures all flow features found in the DNS results. Additionally, a study is done on the effect of wing sweep on the development of the discovered near wall vortical structures, leveraging the computational efficiency of HNS. It is found that the most unstable primary crossflow instability in the clean case does not necessarily result in the strongest near-wall disturbance growth.

Parts of this chapter have been published in

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.

Casacuberta, J., **Westerbeek, S.**, Franco, J. A., Groot, K. J., Hickel, S., Hein, S., & Kotsonis, M. (2025). Streaky perturbations in swept-wing flow over forward-facing step. *Physical Review Fluids*, 10(2), 023902.

5.1. INTERACTION OF CROSSFLOW INSTABILITIES WITH FORWARD-FACING STEPS (DNS BY J. CASACUBERTA)

The LFC technology was proved by Henke, 1999 among others to be an effective method to delay the laminar–turbulent transition in aerospace applications, yet challenges persist in maintaining laminar flow over large swept surfaces, such as subsonic transport aircraft wings. These challenges originate primarily due to surface roughness. A particular emphasis is put on Forward-Facing Steps (FFS). When steps at plate junctions cannot be avoided due to manufacturing limitations, forward-facing steps are preferred. This configuration proved to be less detrimental (i.e., advance transition less) than its backwards-facing counterpart (Crouch et al., 2006; Perraud & Seraudie, 2000; Wang & Gaster, 2005). Forward-facing steps are therefore ubiquitous on lifting surfaces of aircraft.

Contemporary research on the effect of steps focuses on the flow mechanics as opposed to studying the effect of a step via the transition location as was often done in early works on the topic such as Crouch et al., 2006 and Tufts et al., 2017. Flows over steps feature complex fluid dynamics that complicate the analysis both experimentally and numerically. For example, Wilhelm et al., 2003 found that flows over FFS can feature upstream and downstream separation bubbles. Eppink, 2020 points to these separated flow regions as important destabilising factors. Tufts et al., 2017 proposed that in swept wings localised recirculation regions arise both upstream and downstream of the step. The presence of spanwise pressure gradients leads to helical vortices that travel along the span of the wing. The amplification of stationary CFI modes during transition was linked to a constructive interaction between the CF vortices and the downstream helical vortex. Specifically, beyond the step edge, the interaction was found to be stronger when the height of the FFS surpassed the core height of the CF vortices. This led to an understanding that the step height should not exceed to vortex core height to avoid destabilising incoming CFI. Later studies also found destabilisation, i.e. promoted growth, of incoming CFI through an interaction with an FFS. However, their findings did not support the importance of the unperturbed vortex core height (Eppink, 2020; Eppink & Casper, 2019; Rius-Vidales & Kotsonis, 2021). Rius-Vidales and Kotsonis, 2021 additionally find destabilisation of higher harmonics in the flow and argue that this is likely of a nonlinear nature, forced by the destabilisation of the fundamental mode.

Stüer et al., 1999 experimentally and Lanzerstorfer and Kuhlmann, 2012b numerically found the rise of streaky structures in the flow downstream of a forward-facing step. These flow structures appear to be governed by the lift-up effect (Landahl, 1975) and exist in addition to the incoming instability. The flow behaviour hence appears non-modal and rapidly varying in the streamwise direction; two flow features that commonly limit the application of NPSE to such flow scenarios. To investigate the importance of these flow features, the NPSE framework will be used to predict the nonlinear perturbation development of stationary CFI upstream and downstream of the step. The ability of NPSE to provide accurate results downstream of the step corner indicates that second-order streamwise derivatives of the shape profile can be ignored and, hence, that the perturbation varies in the streamwise direction within the limits of the slowly-varying flow assumption.




Figure 5.1: Sketch of the computational domains (I) upstream of the step and (II) downstream of the step located at x_{st} . The hatched region indicates part of the domain not considered in PSE (not to scale).

In this section, the effect of a forward-facing two-dimensional step on the development of stationary crossflow instabilities is investigated by means of DNS and NPSE. The focus lies on the application and limitations of NPSE to the problem and the implications thereof. The reader is referred to Casacuberta et al., 2022 for a thorough look into the physics of the flow problem.

5.1.1. DOMAIN AND FLOW DESCRIPTION

The swept-wing flow is modelled as a flat-plate flow, i.e. airfoil curvature is neglected in this work, resulting in a rectangular numerical domain as is shown in Fig. 5.1. A Cartesian xyz coordinate system is introduced with the x -direction orthogonal to the wing leading edge. The y -axis points up from the surface and the z -axis runs parallel to the leading edge (pointing inboard such that W_e is negative). For the purpose of NPSE, the domain is split into a part upstream and a part downstream of the step as the flow directly surrounding the step violates the assumptions of NPSE. At the step itself, the flow stability problem is not solved. For the downstream part, initial conditions are instead taken from DNS results at a position slightly downstream of the step corner.

The characteristic length and velocity scales used for non-dimensionalisation are the inflow boundary layer thickness ($\delta_0 = 7.71 \times 10^{-4}$ m) as indicated in Fig. 5.1 and chordwise inflow free-stream velocity ($U_\infty = 15.10$ m/s), respectively, resulting in a global Reynolds number $Re = \frac{U_\infty \delta_0}{\nu} = 791.37$. The domain height is 26 times the inflow boundary layer thickness for all cases. This ensures that the effect of the top boundary on the flow is negligible.

The study explores the flow dynamics around forward-facing steps of various heights. DNS are conducted for different step configurations, including a smooth flat-plate reference case. The external pressure distribution is used to find the laminar solution as explained later in Sect. 5.1.2 is obtained from independent experiments conducted by Rius-Vidales and Kotsonis, 2021 on a 45-degree swept wing (M3J). The numerical domain's inlet is positioned at 5% of the wing chord. The steps are located at $x/\delta_0 = 177.62$, corresponding to 20% of the chord of the wing model. An additional streamwise coordinate $x_{st} = x - 177.62$ is introduced to refer to locations with respect to the step. The spanwise wavelength of the fundamental stationary crossflow mode in DNS is chosen such that $\beta_1 = 0.646$ based on a critical mode analysis for maximum amplification based on PSE at the outflow of the computational domain.

The step width h is varied per case for a total of 4 cases including the no-step reference case where $h = 0$. These are used to study the effect of step height on the perturbation behaviour as well as the streamwise extent of the invalidity of the NPSE assumptions. The step heights per case are $0.59\delta_0$ (Case 1), $0.76\delta_0$ (Case 2), and $0.97\delta_0$ (Case 3) with $\delta_0 = \delta_{99}(x = x_0)$. These step heights correspond to between 30-50 % of the unperturbed boundary layer thickness ($\delta_{99}(x = x_s)$) at the step location.

5.1.2. METHODOLOGY

The incompressible Nonlinear Parabolized Stability Equations (NPSE) are solved on a grid consisting of 500 equispaced streamwise stations and 80 Chebyshev collocation points in the wall-normal direction. Streamwise derivatives are discretised through a first-order backward-Euler scheme, incorporating 10 stationary spanwise Fourier modes and the mean-flow distortion ($N = 10$) in the simulations. The number of harmonics is set to approximate the spanwise discretisation accuracy of the DNS ($N_z = 72$) which was able to capture 6 modes and the mean flow distortion. The initial condition for the fundamental crossflow mode is derived from Linear Stability Theory (LST) applied to the local base-flow velocity profile from Direct Numerical Simulation (DNS).

High-order harmonics are automatically initiated by the nonlinear forcing terms, progressively introduced in the chordwise-marching scheme once their strength exceeds a threshold of 10^{-8} in non-dimensional units. The strength is determined based on the order of magnitude of the associated nonlinear forcing term. Upon introducing a new harmonic component, its amplitude is assumed to be zero upstream of its point of introduction, resulting in immediate downstream growth. As a characteristic of the Parabolized Stability Equations (PSE) approximations, the complex streamwise wavenumber of each mode is iteratively updated at each streamwise station until it reaches a threshold of $10^{-6}/Re$, ensuring gradual streamwise changes in the perturbation shape function.

BASE FLOW SOLUTION

DNS of the unperturbed laminar base flow solution to the incompressible Navier-Stokes equations were performed using the conservative finite-volume code INCA. For details on its implementation, the reader is referred to Hickel and Adams, 2008 and Hickel et al., 2014, although case-specific settings and boundary conditions will be discussed later in this section for completeness. The results serve as the base flow for the stability calculations. Additionally, it serves as the starting point for the reference transitional flow simulation.

The computational setup included a spanwise invariant three-dimensional laminar boundary layer. The inflow boundary condition is the solution to the Falkner–Skan–Cooke equations. The flow develops in the x -direction and encounters a forward-facing step of height h as described in Sect. 5.1.1.

At the top boundary, a Dirichlet condition for the static pressure applied using a logarithmic fit of surface pressure measurements from the no-step case in the experiments of Rius-Vidales and Kotsonis, 2021. In nondimensional units, this results in a freestream velocity

$$U_e(x) = 0.0023\ln^4(x) + 0.0377\ln^3(x) + 0.1752^2(x) + 0.5303\ln(x) + 1.8574, \quad (5.1)$$




Figure 5.2: (a) Base-flow chordwise velocity (color map) with lines of constant stream function (solid black) (b,c) Wall shear $(\partial U/\partial y)|_w$: smooth case (thick solid black), step case 1 (dotted orange), step case 2 (dash-dotted blue), and step case 3 (thin solid red) from Casacuberta et al., 2022.

where x is the streamwise coordinate normalized by δ_0 originating from the attachment line. At the outflow, a Neumann-type boundary condition is applied to the velocities. Similarly, this condition is applied in the spanwise direction to model the infinite swept-wing assumption. Lastly, no-slip and impermeability conditions are applied to the wall boundary.

The base flow features two recirculation regions in correspondence with results by Wilhelm et al., 2003 among others. The two recirculation regions are shown by the streamlines in Fig. 5.2a as well as the negative wall-normal derivative of the streamwise velocity as measured at the wall in Fig. 5.2b. Naturally, the size of the recirculation region scales with the step height. For the tallest step, the recirculation region extends approximately 3 characteristic lengths upstream.

FLOW STABILITY SETUP

The stability problem is set up in a similar fashion as the described setup in the previous chapter (Ch. 4). In short, the stability calculations are performed on a subdomain of the DNS numerical domain. This subdomain covers part or all of the flow up or downstream of the step. Aside from the inflow boundary, there are no differences between these two

scenarios. Namely, both feature no-slip and impermeability conditions that translate as a Dirichlet condition for velocities at the wall. Moreover, perturbations are expected to fade into the freestream, which is modelled again by a Dirichlet condition for velocities at $y = H$. The mean flow distortion is exempt from this condition to allow for boundary layer growth due to nonlinear interactions. No boundary conditions for pressure exist. Pressure gradients are solved implicitly. The solution is marched in the positive x direction.

For simulations upstream of the step, the initial condition is again the solution to the local Orr-Sommerfeld eigenvalue problem. The result is normalised and multiplied by the desired initial amplitude $A = 3.5 \times 10^{-3}$ equal to the imposed CFI inflow amplitude in the DNS. Downstream of the step ($x_s - x_s + 7\delta_0$), the local eigenvalue problem does not guarantee an accurate solution. More specifically, in the vicinity of the step, the streamwise velocity and pressure gradients in the base flow, as well as the wall-normal velocity, can become significant. Moreover, the incoming perturbation profile might be strongly distorted due to the presence of the step or nonlinear mechanisms. The local eigenvalue problem cannot account for this. Hence, the PSE simulation downstream of the step is initiated from DNS results for all harmonics (up to $n = 5$). This initialisation requires the description of the local perturbation solution from DNS to be Fourier transformed and provided in the form of the PSE ansatz of Eq. 2.7 per harmonic. The relative phase difference of the harmonics must be preserved in this process, and a streamwise wavenumber must be calculated from the results. The latter is performed following Eissler and Bestek, 1996, Ma and Zhong, 2005 and Hu et al., 2020:

$$\alpha_r = \frac{\partial |\Psi|}{\partial x}, \quad (5.2a)$$

$$\alpha_i = \frac{\partial (\ln(A/A_0))}{\partial x}, \quad (5.2b)$$

with Ψ the spatial component of the phase function present in the wave function. However, for the calculation of this spatial component, a slightly different form is used, namely:

$$\mathbf{q}'(x, y, z, t) = \hat{\mathbf{q}}(x, y, z, t) e^{i[\Psi(x, z) - \omega t]}. \quad (5.3)$$

In some cases, the local perturbation behaviour violates the assumptions of PSE (see Tab. 2.1). In a physical sense, this implies that the perturbation evolution cannot be accurately described by a single streamwise wavenumber ($\in \mathbf{C}$) such that rapid variations of the shape function should occur. Initializing the PSE problem with a perturbation profile that would develop in such a manner is therefore expected to either fail to converge or provide an inaccurate solution.

In the reference transitional DNS, the inflow boundary layer profile is amended with a solution to the Orr-Sommerfeld eigenvalue problem for the Falkner-Skan-Cooke boundary layer profile. As with stability calculations, the result is multiplied by a factor ensuring that the peak streamwise perturbation yields $A_0 = 3.5 \times 10^{-3}$ in units of the freestream velocity U_{ref} . Therefore, the two methods feature identical inflow conditions before the interpolation on the respective numerical grids.




Figure 5.3: Chordwise evolution of the amplitude of the crossflow perturbation in the smooth reference case from LPSE (blue symbols), NPSE (white symbols), and DNS (solid lines): Fourier modes $j = 1-5$ (thick-to-thin) and $j = 0$ (dotted). The vertical red line indicates the virtual step location from Casacuberta et al., 2022

VALIDATION

Figure 5.3 illustrates the comparison between DNS and NPSE amplitudes for the no-step case. This comparison is performed using the amplitude based on the peak streamwise perturbation profile, denoted A^u . Considering only the fundamental mode and disabling nonlinear interactions, the procedure yields solutions to the linear Parabolized Stability Equations (LPSE). Introducing the fundamental mode and its complex conjugate, while allowing for nonlinear interactions, leads to the nonlinear perturbation development. In this scenario, the solutions remain independent of the initial-mode amplitude, arbitrarily matched to the NPSE and DNS amplitudes at the inflow. The match is strong throughout the entire flow domain including the long streamwise extent of saturated primary modes. The perturbation behaviour is initially linear as indicated by the close match between the results up to $x/\delta_0 = 200$, shortly downstream of the virtual step location.

Both methods are subject to their own truncation, discretisation, and modelling errors. Nevertheless, a close match is achieved for the nonlinear perturbation development throughout the domain. Hence, both methodologies are considered trustworthy for the current investigation.

5.1.3. RESULTS

The limit to the PSE assumptions is difficult to calculate a priori. Although a limit to its applicability exists in the current geometry from the rapid variation of the wall and base flow in the vicinity of the step, the exact point at which the shape function variation becomes significant is unknown. Shape function variations are a second-order effect from the wall geometry, with the first-order being the base flow changes. Hence, a generalisation cannot be made from the current result and it is not the goal of this work. Instead, the importance of specific terms in the perturbation equations of Eqs. 2.4 are evaluated via the use of PSE in the proximity of the step.




Figure 5.4: Amplitude of the fundamental chordwise-velocity perturbation component from DNS (lines) and PSE (symbols) upstream of the step: smooth case (thin solid black line and circles), step case 1 (dash-dotted orange line and triangles), step case 2 (dash-dotted blue line and diamonds), and step case 3 (thick solid red line and squares). Vertical lines indicate the beginning of disagreement between DNS and PSE (1% relative error). Figure taken from Casacuberta et al., 2022.

5

UPSTREAM PERTURBATION EVOLUTION SIMULATED BY PSE

Figure 5.4 shows the amplitude development of the fundamental crossflow instability upstream of steps for all cases examined in this study and the reference no-step case. The amplitudes match closely to the DNS results far upstream of the step ($x_{st}/\delta_0 < -5$). The CFI are destabilized by the presence of the step upstream of the step location. The presence of the step is felt by the perturbations upstream via changes in the base flow. A clear example of this is the flow recirculation region. For the largest step case, the recirculation region commences from $x_{st} = -2.8$, see Fig. 5.2. In the middle and small steps, the recirculation region is smaller, starting at $x_{st} = -1.4$ and 0.7 respectively.

The error threshold of 1% in amplitude is crossed upstream of the presence of a circulation region for all cases. However, it is noted that the base flow profile thins upstream of the recirculation region. With increased step height, the terms neglected in the PSE equations appear to become significant further upstream. The direct effect of the recirculation region is not visible in the amplitude development. The amplitude development is based on the peak amplitude, which resides at some distance away from the wall. Any near-wall effects, e.g. the recirculation bubble upstream of the step, see Fig. 5.2, thus do not strongly affect the perturbation structure where the peak velocity rises. This is also shown in Fig. 5.5, which compares PSE and DNS results upstream of the step. Around the wall-normal location where the peak velocity resides, the results from both methods are comparable. The development of the primary structure does not appear strongly altered by the recirculation region. Hence, the PSE does not directly fail to converge. Instead, the relative error increases slowly in the streamwise direction as the errors caused by the PSE assumptions compound in the streamwise direction.

The current results are in line with findings of Hetsch and Rist, 2009 and most recently de Rosa et al., 2024 which show that the PSE methodology can be used to perform stability analysis over small recirculation regions. They also show that the PSE framework strongly outperforms local stability frameworks, proving the importance of first-order streamwise derivatives in the perturbation equations and the relatively small importance of higher-order streamwise derivatives for small recirculation regions. Hjort,




Figure 5.5: Profiles of the normalised amplitude function $|\hat{u}|(0,1)$ from DNS (solid line) and PSE (symbols) upstream of the step at x_{st}/δ_0 values: (a) -4.5, (b) -2.5, and (c) -0.5. Figure taken from Casacuberta et al., 2022.

2004 additionally states that the PSE is able to find flow stability solutions provided that the recirculation is *small*. However, it is difficult to state whether a recirculation is sufficiently small. The PSE framework should hence be used as an investigative tool in such cases only. Fully elliptical frameworks such as HNS are preferred in studies concerning perturbation development.

In Fig. 5.5 disturbance profiles for the streamwise perturbation velocity of the mode (0, 1) are presented for three locations upstream of the largest step. The match between PSE and DNS is strong. However, the rise of a new flow structure near the wall is seen in the results of DNS that is not captured by the PSE framework. Its onset is fast, occurring from $x_{st} = -2.5$ to -0.5 just upstream of the step. The incoming CFI profile does not appear strongly altered by its presence or the recirculation region that exists in the base flow. Its development is therefore accurately predicted by PSE as it lifts and passes over the recirculation region. Closer to the step, the PSE framework fails to provide a solution as the shape function changes more rapidly.

The rapid base flow changes at the step corner as seen in Fig. 5.2 warrant rapid perturbation changes similar to the rise of the near-wall structure upstream of the step. The streamlines near the wall are initially aligned with the x -direction before turning almost 90 degrees to the y -direction at the step within 4 characteristic length scales. Base flow changes of such a scale are likely to result in rapid changes in the streamwise and wall-normal perturbation profiles. The former decreases while the latter increases. The rapid near-wall decay of the streamwise perturbation velocity is seen in Fig. 12 from Casacuberta et al., 2022. The non-modal nature of the perturbation development is also shown by the growth rates of the perturbation vector and a perturbation profile projected onto the base flow streamlines in Fig. 16 from Casacuberta et al., 2022. Since the PSE ansatz features a single wavenumber for u' , v' , w' , and p' , these streamwise variations much return in the shape functions for the respective perturbation velocities. This results in a non-negligible second-order streamwise derivative and the assumptions of PSE are violated.

DOWNSTREAM PERTURBATION EVOLUTION SIMULATED BY PSE

The PSE framework was not able to provide flow stability solutions over the step from one δ_0 upstream to $6\delta_0$ downstream. Instead, the domain was split into a part upstream of the step and downstream of the step. Perturbation amplitude development was matched up to approximately 3 characteristic lengths upstream of the step. After this point, amplitude differences exceed 1% despite the growth being similar.

Shortly upstream of the step, a near-wall structure identified in the DNS results was not captured by PSE. The convective nature of the problem demands a separate simulation downstream of the step corner, starting from the DNS results. This prevents the errors made upstream from contaminating the comparison of the downstream simulations. Regardless, the current PSE code was unable to march over the step, thus demanding this procedure. The downstream simulation is hence initiated from DNS results as was explained in Sect. 5.1.2. The DNS solution is described at the inflow location using the PSE ansatz following a Fast-Fourier Transform (FFT) and an a posteriori calculation of the streamwise wave number. This concerns a nonlinear procedure, i.e. multiple modes can be provided at the inflow and the phase-difference between harmonics is maintained. The perturbation behaviour downstream of the step can immediately be governed by finite-amplitude effects. Therefore, the immediate presence of higher harmonics at amplitudes as provided per DNS is necessary.

The results of the downstream simulation using linear and nonlinear PSE, in addition to the reference results per DNS, for the flow case containing the largest step are presented in Fig. 5.6. It is noted that the initiation of PSE closer to the step than shown here prevented convergence of the PSE immediately after the marching procedure had started. Therefore, it can already be concluded that the perturbation behaviour immediately downstream of the step features rapid shape function variation when only a single streamwise wavenumber is present to describe the streamwise perturbation behaviour. This holds for both the linear and nonlinear PSE simulations.

In the perturbation development as calculated per DNS, a near-wall structure is observed developing shortly downstream of the step corner. Its amplitude surpasses that of the incoming CFI for some streamwise extent. The coexistence of these two structures within the same spanwise wavelength results in flow behaviour not accurately described by a single wavenumber for all perturbation velocity components (as is assumed in PSE and OS frameworks). The analysis of the flow physics is out of the scope of this thesis. For that, the reader is instead referred to Casacuberta et al., 2022 regarding its discovery as well as a later work describing its origin and development in Casacuberta et al., 2025. However, in the case of smooth surface geometries, a similar double-peaked perturbation profile will be observed in Chapt. 6 and the analysis of that flow is in the current thesis' scope.

Through PSE simulations, the behaviour of the fundamental mode shortly downstream of the step, between $x = [185 - 260]$, was shown to be governed by linear mechanisms as linear PSE was able to accurately predict the perturbation development shortly after the near-wall structure had faded. The perturbation development of higher harmonics at approximately the same location is accurately captured by NPSE. Namely, a rapid decay of between 50-80% of the amplitude of higher harmonics is observed. This decay even continues as the primary mode continues to grow between $x = [230 - 300]$.




Figure 5.6: Chordwise evolution of the peak streamwise perturbation velocity amplitude $u'_{\max}$ for modes 1–5 (thick-to-thin) and the mean flow distortion (dotted), from DNS (solid lines), LPSE (blue symbols), NPSE (white symbols): (a) in the near-step regime, (b) downstream of the step, for step case 3. The vertical red line indicates the step location. Figure taken from Casacuberta et al., 2021.

Given the limitations of the PSE framework, the shape functions should slowly vary in this region of the flow.

5.1.4. CONCLUSION

The PSE framework was used in conjunction with DNS to predict the nonlinear evolution of CFI upstream and downstream of steps of varying heights in an otherwise similar Crossflow Boundary Layer. The results are used to analyse the perturbation development in the proximity of the step and further downstream. The limitations of the PSE framework are used to investigate the nature of the perturbation dynamics. Lastly, the effect of the step height on these limitations is considered.

Upstream of the step, the PSE can provide accurate results up to $x_{st} = -1$ for the incoming perturbation profile. The step height affected this distance only marginally. With increasing step height, the 1% amplitude error threshold was crossed further upstream. Moreover, a near-wall flow structure that exists in the same Fourier space as the incoming perturbation is not captured by the PSE framework. The inability of PSE to capture this structure is attributed to the structure's rapid onset, demanding non-negligible second-order derivatives of the shape functions. Furthermore, this structure possibly possesses a distinct streamwise wavelength. This renders the single streamwise wavenumber of the PSE assumption inadequate.

The incoming perturbation profile lifts up from the wall as it approaches the step, seemingly following the base flow streamlines to a large degree. This process is accurately captured by PSE with errors in amplitude compared to DNS being small, namely a maximum amplitude difference of 1% between the PSE and DNS results exists up to the aforementioned location of $x_{st} = -1$. The conclusion can be made that the incoming perturbation is not significantly affected by the near-wall structure that develops closer to the wall for the studied cases.

Further downstream, the discrepancy between PSE and DNS is more profound. How-

ever, the PSE can still provide valuable information about the importance of some terms in the perturbation equations, most notably $\partial^2 \mathbf{q}' / \partial x^2$, and non-modal flow behaviour. When either of these, or both, significantly affect the perturbation development, the PSE framework will struggle to accurately predict the perturbation development. In the direct vicinity of the step, the PSE framework failed to provide any results for all studied cases due to the perturbation development not conforming to the PSE ansatz. From $x_{st} = 8$ onward, the PSE framework was able to find solutions to the flow stability problem. This is approximately 6 characteristic lengths, δ_0 , downstream of the recirculation bubble that exists downstream of the step corner. These results accurately matched DNS results for all harmonics despite the strong variations in amplitude in this region.

It can be concluded that higher-order streamwise gradients cannot be ignored in the direct vicinity of the step. This region extends further than the recirculation bubble in the base flow downstream of the step apex. As a result, the PSE should not be relied upon for the prediction of the perturbation development over the step in general. This is in line with the findings of Cooke et al., 2019, and Franco et al., 2020. However, shortly downstream of the step, the PSE appears well-suited to solve the flow stability problem if provided with accurate initial conditions derived from other methods.

5

5.2. STREAKY PERTURBATIONS IN LAMINAR FORWARD-FACING-STEP FLOW

In the previous section (Sect. 5.1), NPSE was used to confirm the nature of perturbation behaviour upstream and downstream of a forward-facing step in a crossflow boundary layer. Additionally, the limitations of NPSE were tested and it was shown that the NPSE can not reliably provide stability results in proximity of the step. This highlighted the need for a stability framework that can overcome these limitations while maintaining as much of the computational efficiency of the NPSE framework as possible. Chapt. 3 introduced DeHNSSo, a harmonic Navier-Stokes framework capable of finding flow stability solutions in flows featuring complex 2-dimensional surface geometries.

The flow over a forward-facing step is complex and was proven to contain regions of strongly non-modal perturbation behaviour by Casacuberta et al., 2022. Moreover, a near-wall flow structure was discovered that required further research. The current section will cover the use of harmonic Navier-Stokes frameworks in flow scenarios featuring a forward-facing step. More specifically, it concerns the contribution of DeHNSSo to the discovery of streaky perturbations in laminar forward-facing step flow (see Casacuberta et al., 2025).

The main questions arising from the aforementioned discovery of a near-wall structure, or streak, in the flow over a forward-facing step relates to its origin, type, and dependency on the incoming perturbation. More specifically, recent investigations into the interaction mechanism of crossflow instabilities with a forward-facing step have led to controversy on this topic. While the streak is ascribed to the interaction of the incoming CFI and the recirculation region present downstream of the step corner by Tufts et al., 2017, a different explanation is provided by Rius-Vidales and Kotsonis, 2020 and Eppink, 2020. They argue that the incoming CFI is destabilized by additional inflection points in the base flow profile. Casacuberta et al., 2022 indicate the structure to be at

most *conditioned* by the incoming CFI.

To aid in the investigation, harmonic Navier-Stokes frameworks, HNS and the Adaptive Harmonic Linearized Navier-Stokes (AHLNS), are validated with DNS solutions and used to perform parametric studies of the effect of the sweep angle and the spanwise wavelength of the perturbation. The efficiency of the HNS frameworks enables this approach, and the current investigation provides a valuable test case for DeHNSSo.

The AHLNS code is made by the German Aerospace Center (DLR) that combines the benefits of the PSE and HNS (Franco et al., 2018). Namely, the periodic nature of perturbations in the streamwise direction is leveraged in a similar fashion as PSE. This results in a much slower variation of the shape function than the classical Harmonic Navier Stokes description, which contains no streamwise wavenumber in the wave function by default. The main benefit of this approach is that the streamwise grid spacing can be coarser resulting in lower computation times and costs. However, this comes at the cost of an additional convergence loop to redirect growth and streamwise oscillations to the streamwise wavenumber.

5.2.1. DOMAIN FLOW AND DESCRIPTION

The flow geometry closely resembles the flow over a forward-facing step shown previously in Sect. 5.1. For the sake of brevity, the set-up and flow geometry is therefore not described again in this section. Only the additional information relating to the change of sweep angle and spanwise wavelength are described here.

The current investigation employs the same wall geometry shown in Sect. 5.1 with a variation in the sweep angle and fundamental spanwise wavelength using the same wall geometry. The choice to vary these two parameters follows from the hypotheses that the streak is conditioned by the incoming crossflow instability, which affects the amplitude, phase, and spanwise wavenumber, and that the most unstable primary CFI for the considered domain might not result in the most unstable streak. The response of the streak to these parameters can help identify its origin and aid the understanding of its effect on the downstream perturbation development. Moreover, the streak could be responsible for the supercritical behaviour, i.e. transition advancement, found in some experiments such as those by Rius-Vidales and Kotsonis, 2021. This fact is crucial in determining manufacturing tolerances that aim to minimise the effect of such surface features on the transition scenario.

The inflow conditions differ from the previous study on a similar geometry, see Sect. 5.1.1. Here, rather than introducing the solution of the local eigenvalue problem at the inflow, a region of blowing and suction is used to introduce the fundamental mode. This change was made to ensure that the receptivity process governs the rise of the CFI mode, ensuring the results are comparable between sweep angles. The forcing follows the same procedure described previously in the chapter on the development of DeHNSSo, see section 3.5.4. In short, the receptivity process is fully simulated by imposing inhomogeneous Dirichlet boundary conditions for the wall-normal velocity at the wall for the fundamental mode. This represents a spanwise periodic blowing and suction with a spanwise wavelength equal to that of λ_1 . The streamwise dependency of the blowing and suction strength was previously shown in equation 3.13 and is case-independent. The implementation is part of the boundary conditions and is thus accounted for di-

rectly when the linear problem is solved. The current application of this boundary condition is as part of the linear problem only. Hence, amplitude ramping is not necessary. At the inflow, perturbations are forced to zero. Higher harmonics and the mean flow distortion are quenched everywhere in the domain, resulting in a linear approximation. Although DeHNSSo is capable of nonlinear simulations for this case, the highly refined grid around the step would result in large computation times. Regardless, the main flow dynamics, including the origin of the near-wall structure, are hypothesised to be linear based on the DNS results. The current implementation can confirm its linear nature without the need for an energy budget analysis if the near-wall structure is captured using a spatio-temporally monochromatic perturbation ansatz.

The sweep angle is varied in the current study from $\Lambda = 51^\circ$ (case A) down to the unswept case $\Lambda = 0^\circ$ (case F) as shown in Tab. 5.1. Tab. 5.1 shows the relation between W_∞ and U_∞ at the domain inflow. This study of the sweep angle effect serves two purposes: 1. to find an optimal sweep angle for the excited near-wall streaks, and 2. to discover whether a spanwise flow component is a prerequisite for said near-wall streak. For the unswept case, the spanwise flow component $W_\infty = 0$. The perturbation that forms as a result is not a crossflow instability and shows spatial similarity to a Klebanoff mode (see Klebanoff, 1971 and Kendall, 1985) manifesting as spanwise-alternating regions of velocity excess and deficit. However, Klebanoff modes feature a temporal component which the current streak does not. The near-wall steady streak develops during the interaction with an FFS akin to that seen in the interaction of CFI with forward-facing steps.

The sweep of the fundamental wavelength is performed for case A only ($\Lambda = 45^\circ$). A range of spanwise wavelengths is not given a priori. Rather, it is adjusted based on the results until a maximum amplitude response is found for the near-wall structure.

5.2.2. DISCRETIZATION

In DeHNSSo, the same discretization was used for all cases which was previously shown in Fig. 3.11 and proven to ensure that the grid size does not affect the results. This results in 4000 streamwise stations that are locally clustered around the step. Similarly, 350 collocation points are locally clustered around the step. A nearly-constant grid spacing is maintained from the wing surface (i.e., $\eta = 0$) to the step corner. This is necessary to capture the perturbation behaviour just upstream of the step (see Fig. 5.8a and f). From this point onward, the distance between collocation points smoothly increases away from the wall.

For all the cases presented in this work, a mapping of the wall-normal coordinate is used that follows Malik, 1990 to which the reader is referred for a more detailed description. In short, the boundary layer coordinates between $y = 0$ and $y = y_{\max}$ are mapped

Table 5.1: Ratio between the freestream spanwise and chordwise velocities and corresponding sweep angle per case.

Case	A	B	C	D	E	F
$\frac{W_\infty}{U_\infty}$	-1.241	-0.993	-0.745	-0.496	-0.248	0
Λ	51.136°	44.791°	36.670°	26.398°	13.938°	0°

onto a computational domain $0 \leq \zeta \leq 1$ according to

$$y = \frac{a\zeta}{b - \zeta}, \quad (5.4)$$

where b depends on the scaling parameter a as $b = 1 + a/y_{\max}$. The scaling parameter a is defined as $a = y_{\max}y_i/(y_{\max} - 2y_i)$ with y_i the median collocation point. In other words, half of the collocation points exist under the input parameter y_i . This mapping suits boundary layer stability calculations well due to the ability to cluster collocation points where large shears exist. However, the distribution can only be tuned by a single parameter: y_i . This is insufficient for cases that feature a step. Namely, using this mapping the distance between collocation points increases between $y = 0$ and $y = y_i$, while the accuracy at the step height $y = h$ should be similar, if not greater than at wall $y = 0$. Hence, simulations featuring a step use a more complicated point distribution. By combining several mappings with different y_i , a more constant distance between the collocation points $\Delta y(y)$ until the step height can be achieved before it increases into the freestream. As a result, fewer collocation points are necessary to resolve the flow physics around the step corner. Moreover, fewer points are lost under the step surface as a result of the embedded boundary method (see Sect. 3.3.4).

5

5.2.3. RESULTS

All perturbations were introduced via spanwise-periodic wall blowing and suction. The receptivity process is not examined in detail in Casacuberta et al., 2025 due to the focus on the streaks. However, the ability of the DeHNSSo implementation to solve the receptivity problem and match DNS is significant, and proves the robustness of the framework to the perturbation profile and initialisation techniques. This additionally serves as a validation for the inhomogeneous boundary condition implementation allowing for e.g. flow control studies.

In order to effectively compare the downstream perturbation evolution between DeHNSSo and DNS methodologies, a close match in the incoming perturbation is of critical importance. Failure to accurately resolve the receptivity problem could introduce errors in the analysis. Hence, Fig. 5.7a is shown to compare the amplitude evolution of both approaches for case A through the evolution of the peak streamwise perturbation amplitude. The wall boundary condition mimics spanwise periodic blowing and suction which rapidly displaces the momentum of particles in the wall-normal direction. This results in a strong perturbation amplitude near the wall around the location of the actuation ($x = 730$). The evolution closely matches the results of DNS, showing a strong initial peak that slowly decays from $x = 730$ onward. This decay is partially explained by the redistribution of perturbation momentum seen in Fig. 5.7b, by which the perturbation amplitude is spread out more in the wall-normal direction as it convects downstream before being distorted by the presence of the step. Given the current results, a comparison of the perturbation behaviour at the step is not expected to be influenced by differences in the upstream results. Although not shown here, a similar match is achieved for all other flow cases (A-G) in the receptivity region.

The streamwise perturbation velocity amplitude function is shown in Fig. 5.8 as per DNS, AHLNS, and HNS for cases A and G. A close match is achieved between the three




Figure 5.7: Case A streamwise perturbation amplitude (a) with DNS results in black and HNS results in dashed red. Step location is presented by the vertical dashed line. Absolute streamwise perturbation xy contour of $|u'|$ from HNS results (b).

methods. In both cases, a near-wall structure arises shortly downstream of the step. Crucially, this structure exists very close to the wall and has a sharp peak in the streamwise perturbation velocity. A highly refined grid is thus necessary, regardless of the sweep around $\eta = h$ to capture it.

For the top row, i.e. where $W_\infty/U_\infty = 1.24$, the near-wall structure does not appear to merge with the incoming perturbation structure. Instead, it rapidly decays within 10 characteristic lengths. In the unswept case, the near-wall structure rises in a similar fashion as in the maximum sweep case, but the downstream development is different. Namely, the near-wall structure appears to be maintained for a longer streamwise extent and merges with the incoming perturbation profile. The HNS methodology is able to capture both distinctly different perturbation developments and closely matches the DNS results. For a further analysis of the flow physics, the reader is referred to Casacuberta et al., 2025.

To maintain a manageable scope, the spanwise wavelength was chosen based on the most unstable mode ($\bar{\lambda}_z = 7.5\text{mm} = 9.73\delta_0$, $\beta_1 = 0.646$) in the clean reference case, i.e. without a step. However, the near-wall structure does not need to exhibit a similar most unstable behaviour at this wavelength. Leveraging the efficiency of HNS frameworks (DeHNSSo and AHLNS), the spanwise wavelength was varied between $\bar{\lambda}_z = 4 - 7.5\text{mm} = 5.19 - 9.73\delta_0$. This range was determined based on the response to slightly greater and smaller spanwise wavelengths which showed a destabilisation of the near-wall structure for smaller wavelengths. Then, the spanwise wavelength was reduced at an interval of




Figure 5.8: Profiles along y of the amplitude function $|\tilde{u}|(0,1)$ at the step from DNS (solid black), AHLNS (dash-dotted green), and linear HNS (dashed red) at $x/\delta_0 = -0.5$ (a,f), 0.5 (b,g), 2.5 (c,h), 5 (d,i), 10 (e,j). The top row of figures presents the case with $\max(|W_\infty/U_\infty|)$ and the bottom row illustrates case $|W_\infty/U_\infty| = 0$. Figure taken from Casacuberta et al., 2025

0.5 mm to find an approximation of the maximum response.

The choice of fundamental wavenumber influences the development of the incoming mode. Different modes will experience different growth paths up to and around the step. Consequently, a comparison of the amplitudes directly would not yield great insight. Instead, the amplitudes were normalised by the amplitude where the near-wall structure streamwise velocity peak overtook the incoming structure's peak streamwise velocity, resulting in a ΔN shown in Fig. 5.9. Since the amplitude development is the result of linear HNS simulations, this does not affect the interpretation of the results. It is noted that only AHLNS results are shown in Casacuberta et al., 2025, while the current results are made with DeHNSSo.

Fig. 5.9 indicates that the largest amplitude response is found for a smaller spanwise wavelength than the mode governing the transition scenario in the clean reference case without a step. Namely, the near-wall structure achieves the largest amplitude relative to the incoming perturbation for a wavelength around $\bar{\lambda}_z = 6$ mm for flow case A (see Tab. 5.1.) This is slightly higher than the 5 mm (20%) reported in Casacuberta et al., 2025. This could be a result of the differences found between DNS, DeHNSSo and ALHNS as also visible in Fig. 3.12. The differences are small; However, they are significant enough in this flow region to affect the optimisation problem, leading to a different result.




Figure 5.9: Development of instability amplitude normalized by the amplitude at peak streamwise velocity change-over, for $\bar{\lambda}_z = 4.5$ mm (light red) to $\bar{\lambda}_z = 7.5$ mm (black) in steps of $\Delta\bar{\lambda}_z = 0.5$ mm around the step (x_s).

5

The current results point to a risk in considering only the most unstable mode of the reference case in boundary layers over FFS. The most unstable incoming CFI may not result in the strongest near-wall streak response downstream of the step corner. A considerable implication is that this allows for optimisation in the step location and geometry. For some combination of geometry and step location, the most unstable near-wall streak might align with the most unstable incoming CFI. A large incoming CFI conditioning a rapidly growing near-wall streak could anticipate transition and should be avoided. This demands a considerable expansion of the parameters considered, such as the incoming CFI wavelength and step location. However, linear simulations likely suffice to discover such interactions, making such endeavours feasible.

Despite the importance of nonlinear interactions in such flows, Casacuberta et al., 2025 finds that the rise of near-wall streaks is dominated by linear mechanisms. Hence, nonlinear simulations are hypothesized to be required only for the discovered critical case in linear studies to find the limits in incoming CFI amplitude.

5.2.4. CONCLUSION

The HNS framework, here represented both by the AHLNS code from DLR (Franco et al., 2020; Franco et al., 2018) and the current implementation, DeHNSSo, was used to simulate the formation and development of Crossflow Instabilities and resulting steady streaks in swept-wing boundary layers featuring a step at various sweep angles. The HNS framework was shown to be well-suited to tackle the receptivity problem accompanied by spanwise periodic blowing and suction as well as the intricate flow dynamics of the interaction of the resulting boundary layer instabilities with forward-facing steps. The two implementations, AHLNS and DeHNSSo, closely match with DNS results of perturbation development at various sweep angles. This holds both for the receptivity of perturbations following blowing and suction actuation as well as the interaction of the resulting perturbations with a forward-facing step. All three methods predict the rise of a near-wall structure, or streak, downstream of the step corner that exists at the same spanwise wavelength as the incoming mode. Given the assumption of linearity made for these simulations in DeHNSSo and AHLNS, this structure was proven to be purely of

linear origin.

The HNS methodologies accurately capture the development of the streak for all considered sweep angles. The perturbation development downstream of the step corner varies based on the sweep angle. The streak rapidly decays for the larger sweep angles, while persisting further downstream for the unswept case.

The efficiency of the stability frameworks was then leveraged to perform a spanwise wavelength sweep to find the most unstable response of the near-wall structure. It was found that a maximum streak response was influenced by the spanwise wavelength of the incoming mode. Compared to the amplitude of the incoming mode, a fundamental wavelength of 6 mm resulted in the strongest streak response. This is a smaller than the most unstable modal response in the reference case without a hump (7.5 mm), but slightly higher than reported in Casacuberta et al., 2025 based on AHLNS (5 mm). In the paper, a small mismatch between DNS, DeHNSSo and AHLNS was observed in the amplitude development of the streak. It is likely that the same numerical details, such as the interpolation of the base flow and the definition of the step geometry, are responsible for this.

6

INTERACTIONS BETWEEN STATIONARY CFI AND A SMOOTH SURFACE HUMPS

Failure is not the opposite of success; it's part of success.

Arianna Huffington

DeHNSSo is leveraged to perform a parametric sweep of the amplitude of stationary CFI over a smooth surface hump in a swept boundary layer. A near-wall structure is seen to arise on the leeward side of the hump. It is found that increasing the initial CFI amplitude results in an amplitude response of the near-wall perturbation structure that is qualitatively similar, but exhibits increasing local destabilization. The near-wall structure shows larger peak velocities than the incoming CFI.

An energy budget analysis reveals that the growth of the near-wall structure is dominated by linear mechanisms after the initial formation. Moreover, the growth of the fundamental incoming CFI is inhibited through the reduced effectiveness of the wall-normal momentum redistribution mechanism downstream of the hump. This is preceded by a spatial perturbation shape deformation, governed by (spanwise) transport terms.

The results suggest that stabilization of incoming stationary CFI via smooth surface humps is most effective at low incoming perturbation amplitudes. At higher perturbation amplitudes, newly-formed near-wall structures, pre-conditioned by the incoming CFI, overtake the incoming CFI, and possibly hastening the transition process.

Parts of this chapter are under review for publication in Journal of Fluid Mechanics.

Westerbeek, S., Casacuberta, J., & Kotsonis, M. (2025). Linear and nonlinear interactions between stationary crossflow instabilities and a smooth surface hump. Journal of Fluid Mechanics, under review.

6.1. INTRODUCTION

In this chapter, the efficiency and robustness of the DeHNSSo are leveraged to study the effect of incoming crossflow instability amplitude on the interaction with a smooth hump. A collaboration that led to an article by Rius-Vidales et al., 2025 showed significant transition delay of low-amplitude crossflow instabilities due to the presence of a smooth surface hump. However, higher amplitude crossflow instabilities, forced via larger discrete roughness elements on the wing, caused rapid transition advancement up to the hump location. A numerical amplitude sweep with DeHNSSo is analysed based on (perturbation) velocity fields, Reynolds-Orr energy budgets, and a comparison with local linear stability results.

In critical engineering applications such as subsonic transport aircraft, skin friction drag strongly depends on the state of the boundary layer (i.e. laminar or turbulent) developing on main wings and stabilisers. The considerable contribution of skin friction drag to the total aerodynamic drag can be reduced by either manipulating the turbulent boundary layer resulting from laminar-turbulent transition or by delaying the onset of transition, i.e. through Laminar Flow Control (LFC) (Arnal, 1996). The use of LFC showed promise in both laboratory studies and flight tests (Joslin, 1998) and largely motivates the present work. However, the effectiveness of LFC is known to be hindered by structural and geometrical imperfections present on *realistic* wings, such as steps, gaps, or grooves that affect the boundary layer transition process (Crouch et al., 2006; Wang & Gaster, 2005). Such imperfections arise due to manufacturing tolerances as well as operational wear and tear (Worner et al., 2003). Various studies have found that steps (Casacuberta et al., 2022; Rius-Vidales & Kotsonis, 2020), gaps (Forte et al., 2015), waviness (Holmes et al., 1985; Westerbeek & Kotsonis, 2022), and humps (Franco et al., 2020) alter the stability of the boundary layer and are a source of laminar flow deterioration (i.e. the transition front is moved upstream, thus reducing the extent of laminar flow). The flow behaviour and eventual transition scenario vary with the type of instability interacting with the aforementioned surface imperfections.

Swept wings are a key focus area for LFC, inasmuch as they are widely used in subsonic transport aircraft. The laminar-turbulent transition of subsonic boundary layers over swept wings is driven by the growth of four distinct instability types; namely Görtler vortices, attachment line instabilities, Tollmien-Schlichting (TS) waves, and Crossflow Instability (CFI). The Görtler instability arises due to concave surface curvature and is commonly controlled through opportune wing profile design (H. Reed & Saric, 2002). The leading-edge attachment line has been observed to be unstable to viscous TS-like waves as found in Blasius flow (Theofilis, 1998). A reduced leading-edge radius can mitigate the attachment line instability Lin and Malik, 1997. Consequently, the most critical instability kinds in swept wing flow are Tollmien-Schlichting waves and Crossflow Instability (CFI) (W. S. Saric et al., 2003). The latter may manifest in either stationary or travelling form, see Deyhle and Bippes, 1996 for a discussion on pertinent receptivity mechanisms. Swept wing boundary layers typically experience significant growth of stationary CFI near the leading edge. These instabilities are generally responsible for transition in environments of low freestream turbulence (Deyhle & Bippes, 1996; Downs & White, 2013) and are strongly destabilised under favourable pressure gradients (Bippes, 1999). If transition is not anticipated due to CFI growth and the favourable pressure gradient is

not maintained, the growth of TS waves can lead to transition further downstream in the mid-chord region (H. L. Reed & Saric, 1989).

6.1.1. INTERACTION BETWEEN BOUNDARY LAYER INSTABILITIES AND SURFACE IMPERFECTIONS

Two-dimensional, i.e. invariant in one spatial direction, surface imperfections are commonly found to promote transition governed by TS instability (see for instance Masad and Iyer, 1994; A. Nayfeh, 1992; Perraud and Seraudie, 2000 and Worner et al., 2003). Gao et al., 2011 used the Parabolised Stability Equations (PSE) method to study the interaction between incoming TS waves and a smooth hump and found significant instability growth downstream of the hump. The main hump shape prevented the local formation of recirculating flow potentially giving rise to a global instability mechanism (Xu et al., 2017). However, for additional variations of the hump shape, the effects of the induced adverse pressure gradient and flow separation contribute to the destabilising influence of the hump. Park and Park, 2013 carried out a similar analysis accounting for nonlinear perturbation effects. They showed that the hump promotes instability growth even for perturbations with a small initial amplitude. Furthermore, it was noted that the amplitude of the incoming perturbation conditions affects the secondary instability growth more prominently than the shape of the hump. Franco et al., 2020 used the Adaptive Harmonic Linearized Navier-Stokes (AHLNS) method to investigate the effect of two-dimensional humps that are either smooth or rectangular (i.e. a combination of forward- and backward-facing step). They reported an increase in the perturbations amplitude downstream of the hump that is proportional to the height of the hump. Moreover, they found that the shape of the hump, i.e. smooth or rectangular, becomes increasingly important when the hump is of greater height.

It has been widely reported that the shape of the surface influences the mechanisms of perturbation interaction. For instance, Forward-Facing Steps (FFS) are less critical than Backward-Facing Steps (BFS) in promoting transition (Crouch et al., 2006; Perraud & Seraudie, 2000; Wang & Gaster, 2005). At the same time, under certain conditions, FFS have the potential to stabilise incoming TS instability, as demonstrated in the Direct Numerical Simulations (DNS) by Worner et al., 2003. To further investigate the mechanisms of TS wave stabilisation by a FFS, Barahona et al., 2022 combined experimental and numerical techniques to assess the interaction of a monochromatic (i.e. single frequency) TS wave with FFS. They hypothesize that the Orr mechanism (Åkervik et al., 2008) governs the perturbation evolution locally at the FFS.

Wu and Hogg, 2006 argued that incoming TS waves interacting with a localised surface feature are scattered due to the rapid spatial distortion of the base flow around the feature. The effect of the surface feature on the instability amplitude is quantified through a so-called transmission coefficient. Xu et al., 2016 further investigated the behaviour of incoming TS waves interacting with rapid localised base-flow distortions ($< \mathcal{O}(\lambda_x)$, i.e. smaller than the streamwise wavelength of the perturbation). Here the earlier-hypothesised linear dependence of amplification on the surface feature height could not be explained by small humps and indentations both destabilising incoming TS waves. Moreover, the statements in Wu and Hogg, 2006 describing surface protuberances locally increasing the disturbance amplitude are nuanced by the finding that the

effect occurs over a longer streamwise extent than the protuberance itself, imply that some effects are not localised.

Similar to TS waves, the interaction of CFI with forward- and backward-facing steps has been studied in detail. Both experimental and numerical studies found that the incoming crossflow instabilities experience a wall-normal diversion as they approach the step, essentially being transported in the direction of the local mean flow streamlines. Meanwhile, the wall-normal perturbation velocity component shows a distinct secondary peak near the wall (Casacuberta et al., 2021; Eppink, 2020; Eppink et al., 2018; Tufts et al., 2017) and shortly downstream of the step corner near-wall vortical structures arise (Casacuberta et al., 2021; Eppink, 2020; Eppink et al., 2018; Rius-Vidales & Kotsonis, 2021). Casacuberta et al., 2022 demonstrated that numerical solutions of the Nonlinear Parabolised Stability Equations (NPSE) initiated from DNS data in the immediate proximity of the step were unsuccessful. However, when the initialisation point was displaced marginally downstream, close alignment with DNS outcomes was achieved. This observation was instrumental in establishing the non-modal characteristics inherent in the instability development in the immediate proximity of the step. Furthermore, a detailed examination of the near-wall vortical structures, concluded that these may be influenced by the original incoming crossflow instability, yet exhibit autonomous behaviour distinct from it (Casacuberta et al., 2025).

6

In contrast to literature describing sharp surface modifications such as steps, existing literature pertaining to the interaction between CFI and smooth surface modifications, such as humps, is limited. Cooke et al., 2019 numerically investigated the interaction of both stationary and travelling CFI with localised bumps (i.e. a rapid succession of a smoothed forward- and backward-facing step). They employed both Linear Harmonic Navier-Stokes (LHNS) and Linear Parabolised Stability Equation (LPSE) frameworks but found the LPSE to be incapable of providing accurate results. The presence of surface features generally proved destabilising for the incoming instabilities. However, in two scenarios where the bump height was low, travelling crossflow instabilities were found to undergo stabilisation. Westerbeek et al., 2023 performed linear and nonlinear simulations of the interaction between stationary CFI and large aspect-ratio (i.e. width/height > 50) smooth humps of varying height and width that featured no recirculation regions. They observed consistent stabilisation of the fundamental mode of crossflow instability, defined as the most unstable mode in the baseline case (i.e. no hump present). Higher harmonics were instead consistently destabilised by the hump. Both behaviours were found to be governed by linear production mechanisms.

More recently, motivated by the numerical work of Westerbeek et al., 2023, Rius-Vidales et al., 2025 performed an exploratory proof-of-concept wind tunnel experiment, where a single smooth hump was installed on a 45° swept wing and tested at chord Reynolds number of about 2 million. The outcomes pointed to two distinct scenarios, notably influenced by the amplitude of incoming CFI. Specifically, compared to reference no-hump conditions, the hump was found to stabilise relatively low-amplitude incoming stationary CFI at a linear growth stage, and produced a significant extension of laminar flow of approximately 14% of chord length. In contrast, when the incoming CFI was at relatively high amplitude and undergoing non-linear growth, the hump was found to cause an abrupt transition advancement, essentially tripping the flow.

The majority of aforementioned studies treated the effect of single smooth surface modifications on boundary layer stability. However, plate rolling techniques and wing bending can lead to the formation of multiple and successive smooth surface modifications, which can be generalised as surface waviness in wings and profoundly affect the transition scenario (A. H. Nayfeh et al., 1988). The effect of waviness on CFI has received some, but limited, attention. Masad, 1996 found that the presence of surface undulations resulted in a destabilising effect on incoming crossflow instabilities. Moreover, this effect was positively correlated to the undulation amplitude. Thomas et al., 2016 instead found a scenario that led to the stabilisation of the dominant CFI mode, but cited a sensitivity to undulation wavelength and phase with respect to the incoming mode. Lastly, Westerbeek and Kotsonis, 2022 performed a numerical investigation into the effect of surface undulation shape on the nonlinear amplitude development of CFI but found the effect of the undulation shape to be minor for the studied cases. The presence of undulations yielded a destabilising effect on the perturbations. This is in contrast to the positive effects of single smooth surface protuberances reported by Westerbeek et al., 2023. The negative effects of surface waviness were likely caused by the proximity of shape repetition, since Westerbeek et al., 2023 found stabilisation to occur significantly downstream of the single surface modification.

6.1.2. SMOOTH SURFACE HUMPS AS A LAMINAR FLOW CONTROL METHOD

The recent observations by Westerbeek et al., 2023 and Rius-Vidales et al., 2025, suggest the potential of smooth surface modifications, such as humps, as a practically feasible passive LFC method, in the specific use-case of swept wings. Preliminary results identified the robustness of the stabilisation to variations of chord Reynolds number, perturbation forcing conditions and pressure gradient (i.e. angle of attack). Under the effect of such a hump, the incoming stationary crossflow instability is shown to undergo a series of stabilising and destabilising events that take place mostly downstream of the hump apex. The stabilisation is often found to be stronger than the preceding destabilisation, resulting in a downstream amplitude that is lower than in the clean reference case. Additionally, in cases where the transition location is advanced, higher harmonics appear to play an important role (Rius-Vidales et al., 2025). Nonlinear simulations of CFI interacting with a smooth hump in Westerbeek et al., 2023 showed a consistent destabilisation of higher harmonics. Despite these preliminary observations, the details of the underlying mechanism of interaction between the incoming CFI and the hump are still largely unknown. Specifically, the importance of incoming instability amplitude needs to be systematically tested to gain more insight into the (nonlinear) flow dynamics.

Navigating the large parameter space relating to the surface modification (size, shape, location, etc.), boundary layer characteristics and incoming instabilities (wavenumber, frequency, amplitude, etc.) requires a numerically efficient yet accurate approach. The effect of surface imperfections is often measured by considering the direct effect on transition location. In a wealth of past works, universal roughness-based Reynolds number thresholds were sought, above which surface features promote transition in e.g. Holmes et al., 1985; Tufts et al., 2017; Wang and Gaster, 2005. Alternatively, reduced-order models resort to finding a correction to the integral amplification factor (i.e. ΔN -factor) based on surface features (see Crouch and Kosorygin, 2020; Crouch et al., 2006). A linear nu-

merical study by C. Edelmann and Rist, 2014 found that streamwise position, step location, step height, and Mach number affect the N-factor. However, as the ΔN -factor method is linear per definition, the effect of perturbation amplitude is generally ignored. Indeed, both Eppink, 2020 and Rius-Vidales and Kotsonis, 2020 noted that single-parameter models neglect key aspects of the underlying physics such as the incoming perturbation amplitude. These can be conditionally extracted in experiments or high-fidelity simulations. However, detailed parametric studies of the interaction between surface modifications and incoming instabilities are hindered in experiments by measurement limitations (noise, field of view, reach, and intrusion) near the surface and numerically by the large computational costs involved with running transitional DNS. For those reasons, few studies are concerned with the effect of the incoming perturbation amplitude on the interaction with surface modifications at the scale of the boundary layer instabilities. Weakly nonlinear stability theory, as shown in e.g. Wu, 2019, could be applied to perform stability calculations more efficiently. However, weakly nonlinear theories require a hierarchical ordering of the terms based on amplitude (Bertolotti et al., 1992). Since this hierarchy is not generally known a priori, at least for the presently investigated flows, its application may not be valid for all inflow amplitudes. To prevent the weakly nonlinear assumption from affecting the analysis, a fully nonlinear framework is used here instead. In summary, the limitations of both traditional numerical and experimental approaches in tackling the interaction of boundary layer instabilities with surface modifications motivated the development of an efficient Harmonic Navier-Stokes (HNS) framework as used in the present work (Westerbeek et al., 2024).

6.1.3. AIM AND OVERVIEW OF THIS CHAPTER

The effect of humps, on laminar-turbulent transition appears to be influenced largely by the amplitude of incoming disturbances, as demonstrated by the experiments of Rius-Vidales et al., 2025. The present work aims to improve understanding of the stabilising and/or destabilising effect of smooth surface humps on incoming CFI, further enabling the design of such geometries as a realisable and practical passive LFC method. To elucidate these effects, a parametric sweep of the amplitude of the incoming CFI is deployed based on Harmonic Navier-Stokes simulations. A stability and energy budget examination is conducted to assess how certain hump-induced base flow features, e.g. changes of the base flow profiles, inflection points, crossflow reversal, or local pressure gradients, qualitatively and quantitatively influence the development of the incoming CFI.

The work is structured as follows. The methodology is outlined in section 6.2. Herein, the domain and flow cases are described, followed by a presentation of the techniques used for extracting the base flow, perturbation development, and energy budget. The topology of the base flow is scrutinised in section 6.3, while in section 6.4 the non-linear interaction of incoming CFI with varying amplitude and the hump is analysed. Local linear stability frameworks are also leveraged for this analysis.

6.2. METHODOLOGY

In this chapter, the geometry of the hump and the flow conditions are guided by the experiments of Rius-Vidales et al., 2025, which provided evidence of transition delay of

Table 6.1: Reference quantities used in the present work

Symbol	Description	Reference quantity
$\bar{l}_{\text{ref}}$	reference length	$2.1798 \cdot 10^{-4} \text{ m}$
$\bar{\nu}_{\text{ref}}$	kinematic viscosity	$1.4550 \cdot 10^{-5} \text{ m}^2/\text{s}$
$\bar{U}_{\text{ref}}$	reference velocity	14.3 m/s
Re	Reynolds number	214

about 14% chord length by a surface hump in CFI-dominated flow, when the incoming CFI was intercepting the hump at relatively low amplitude. The observed transition delay was found to be robust to changes in the Reynolds number. However, when the incoming CFI was at a higher amplitude, the hump was found to promote early transition (i.e. the transition front is shifted upstream when compared to reference no-hump conditions). In these experiments, the amplitude of incoming CFI was controlled via the height of micrometric Discrete Roughness Elements (DRE) located near the wing leading edge, see for instance Zoppini et al., 2022.

In the present work, mechanisms of unsteady instability and laminar breakdown are not accounted for in the numerical simulations. The goal of this chapter is particularly placed on the characterisation of the effects of the initial CFI amplitude on the mechanisms of stationary perturbation growth and decay due to the hump. The truncation of the flow to the dominant stationary spectral components of the instability system greatly simplifies the computational complexity and resource demand. By the assumption that transition is driven by the amplification of classic modal primary and secondary CFI instabilities (i.e. path A in the Reshotko/Morkovin map, Reshotko and Tumin, 2006), the growth behaviour and amplitude of the primary instability are used as proxies for laminar-turbulent transition.

6.2.1. FLOW CONDITIONS AND DOMAIN GEOMETRY

All quantities in this chapter will be shown in non-dimensional form unless specified. Dimensional quantities are denoted with an overbar. Global characteristic scales are used to non-dimensionalise the equations due to the non-local nature of the Harmonic Navier-Stokes (HNS) framework used in this chapter. The reference length is defined as the Blasius length scale $\bar{l}_{\text{ref}} = \sqrt{\bar{x}_0 \bar{U}_{\text{ref}} / \bar{\nu}_{\text{ref}}}$ at the inflow of the numerical domain, i.e. at $\bar{x} = \bar{x}_0$ (see figure 6.1). The reference kinematic viscosity is denoted as $\bar{\nu}_{\text{ref}}$ and all velocity quantities are normalised by the reference velocity $\bar{U}_{\text{ref}}$, taken as the external (i.e. inviscid freestream) chordwise (i.e. pointing in the direction orthogonal to the leading edge, x) velocity at the simulation domain inflow. The Reynolds number is defined as $Re = \bar{l}_{\text{ref}} \bar{U}_{\text{ref}} / \bar{\nu}_{\text{ref}}$. The reference values used in the present work are summarised in Table 6.1.

The experiments of Rius-Vidales et al., 2025 were performed on the M3J swept-wing model featuring a symmetric aerofoil and a 45° sweep angle. The maximum profile thickness is placed at 65% chord to maintain a favourable pressure gradient over a large chordwise extent. Measurements were performed on the pressure side of the wing and at an angle of attack, α , of 3° . This ensures favourable conditions for the growth of crossflow instabilities while suppressing other types of instabilities such as Tollmien-

Schlichting waves and Görtler vortices. For a detailed description of the wing design, the reader is referred to Serpieri and Kotsonis, 2016. Stationary CFI was introduced via DREs placed at a distance of 2.5% of chord from the leading edge. The resulting perturbation was shown to trigger the most unstable mode for the considered domain. This mode is characterised by the angular frequency $\omega_1 = 0$ (i.e. a stationary mode) and the spanwise wavenumber $\beta_1 = 2\pi/\lambda_z = 0.1831$, where λ_z denotes the spanwise wavelength. The subscript 1 refers to the fundamental component of the truncated Fourier series. This notation is employed due to the use of a nonlinear stability framework, as outlined in section 6.2.3. The DRE height in the experiments was adjusted to introduce CFI at varying amplitudes (Zoppini et al., 2022) which is modelled in the present simulations by imposing the inflow amplitude and shape of the fundamental mode (0, 1) directly. The treatment of the inflow boundary is covered in more detail in section 6.2.3.

In the wind tunnel measurements of Rius-Vidales et al., 2025, the local airfoil curvature at the location of the hump was orders of magnitude smaller than the curvature of the hump and hence it is neglected for the purpose of this investigation. The flow problem correspondingly considers a rectangular domain containing an inflow boundary, a top boundary, a wall featuring the smooth hump, and an outflow boundary (see Fig. 6.1). A global Cartesian coordinate system is considered, denoted by (x, y, z) , with x indicating the direction orthogonal to the leading edge, and z is the direction parallel to the leading edge. The direction y completes the coordinate system. The numerical coordinates read (ξ, η, z) with the ξ -axis pointing orthogonal to the leading edge, but following the surface hump, in a body-fitted sense. The experimental hump shape is approximated by

$$y(x) = h \exp\left(0.5 \left(\frac{x - 618.96}{37.7}\right)^2\right) \quad (6.1)$$

where all parameters are non-dimensionalised with the reference length ($\bar{l}_{ref}$). The resulting hump has a width $b = 229 \equiv 0.056\bar{c}$ and a height $h = 4.58$, thus $b/h = 50$. The hump is symmetrical around its center, which is located at $x_m (\equiv 619\bar{l}_{ref} = 0.15\bar{c})$, relative to the AL. The η -axis is offset by the surface hump, though it remains parallel to the y -axis (see Fig. 6.1). The perturbation equations are numerically solved in the aforementioned domain in the (ξ, η, z) coordinates via a coordinate transformation from the Cartesian coordinate system, (x, y, z) . The u -, v -, and w -velocity components are defined in the direction of x , y , and z , respectively. The perturbation kinetic energy is defined as $E_k = \sqrt{u'^2 + v'^2 + w'^2}$, where the prime indicates perturbation quantity.

The inflow boundary of the numerical domain, $x = x_0 = 214$, corresponds virtually to 5% of the chord of the reference wing and is located at a distance of $214\bar{l}_{ref} (= 0.05\bar{c})$, where $\bar{c} = 0.9$ m) reference lengths downstream of the effective attachment line (AL). The chordwise length of the domain is $L = 1156 (\equiv 0.28\bar{c})$ and the height of the numerical domain is $H = 91$, equivalent to $27\delta_{99}$ at the inflow. The choice of domain height was based on a compromise between accuracy and numerical efficiency. The chosen value was taken through a sensitivity study and was based on solution invariance with domain height, while maintaining a sufficiently fine grid spacing.




Figure 6.1: Sketch of the M3J airfoil shape in the physical domain (shown here inverted) and numerical domain neglecting airfoil curvature. The flow is coming from the left at an angle of attack α and sweep angle Λ . The global coordinate system (x, y, z) and the numerical curvilinear coordinate system (ξ, η, ζ) are shown. The outflow buffer is denoted by the shaded region with attenuation strength going from 0 (white) at the dotted line to 1 at the outflow (dark gray). The dashed line conceptually represents the boundary layer thickness at 0.99 of U_e (δ_{99}).

6.2.2. BASE FLOW CALCULATION

The numerical framework for stability analysis is based on the perturbation equations, which require a pre-calculated solution of the unperturbed base flow. To this end, computations of the incompressible base flow are carried out independently by use of the commercial finite-element solver COMSOL. The setup features a square domain discretised in second-order elements for velocities and first-order elements for pressure. In the wall-normal direction, 96 elements were present, which are clustered near the wall with an exponential function. In the chordwise direction, 1320 elements were present with additional refinement around the hump. By virtue of spanwise invariance in the base flow, only two elements were considered in the z -direction.

BASE FLOW BOUNDARY CONDITIONS

The boundary conditions used in the computations of the base flow are described next. For both hump and reference cases, the static pressure at the top boundary of the base flow domain is prescribed and set to be identical to the pressure measured in the experiments of Rius-Vidales et al., 2025 on the wing surface at reference no-hump conditions. The experimentally-measured surface static pressure results in an external (i.e. inviscid) velocity distribution as described in Casacuberta et al., 2022, equation 2.14. Additionally, second-order wall-normal velocity derivatives are suppressed at the top boundary. This accommodates the presence of a streamwise pressure gradient which results in non-zero first-order wall-normal derivatives for velocities at the top boundary. Lastly, a constant external spanwise velocity, $\bar{W}_e = -1.35\bar{U}_{\text{ref}}$, is prescribed at the top boundary, to model the swept flow conditions.

The solution to the Falkner-Skan-Cooke equation corresponding to the local pressure gradient is imposed at the inflow, in conjunction with a first-order Neumann con-

dition for the pressure. At the outflow, static pressure is prescribed equal to the static pressure at the top boundary at the corresponding x location. At the wall, the no-slip and no-penetration conditions are enforced. Periodic boundary conditions are imposed in the transverse boundaries.

6.2.3. FLOW STABILITY SIMULATION

The in-house solver for nonlinear HNS, DeHNSSo, is employed to compute numerical solutions of the perturbation equations. The reader is referred to Westerbeek et al., 2024 for details on the HNS method and the numerical framework. The unsteady incompressible Navier-Stokes Equations are considered, which in nondimensional form read as

$$\frac{\partial \mathbf{q}}{\partial t} + (\mathbf{q} \cdot \nabla) \mathbf{q} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{q}, \quad (6.2a)$$

$$\nabla \cdot \mathbf{q} = 0. \quad (6.2b)$$

The p is the static pressure and the velocity vector $\mathbf{q} = [u \ v \ w]^T$ expresses the *total* flow, i.e. the superposition of base flow, $\mathbf{Q}$, and perturbation fields, $\mathbf{q}'$, such that

$$\mathbf{q}(\xi, \eta, z, t) = \mathbf{Q}(\xi, \eta) + \mathbf{q}'(\xi, \eta, z, t), \quad (6.3)$$

where $\mathbf{q}'$ assumes a harmonic expansion along the z -axis and in time t so that the perturbation field can be expressed as a truncated Fourier series, i.e.

$$\mathbf{q}'(\xi, \eta, z, t) = \sum_{n=0}^N \sum_{m=0}^M \hat{\mathbf{q}}_{(n,m)}(\xi, \eta) e^{i(\beta_n z - \omega_m t)} + \text{c.c.}, \quad (6.4)$$

and likewise for the pressure, p , where $\omega \in \mathbb{R}$, $\beta \in \mathbb{R}$, and $i^2 = -1$. The moduli of $\hat{\mathbf{q}}_{(n,m)} \in \mathbb{C}$ will be referred to as the amplitude functions, the fundamental spanwise mode is denoted by β_1 , and the c.c. expresses the complex conjugate. The non-oscillatory perturbation, or so-called Mean Flow Distortion (MFD), corresponds to $(m, n) = (0, 0)$. In the present nonlinear simulations, higher-order harmonics and the mean-flow distortion are allowed to rise by modal interactions up to $N = 5$ downstream of the inflow, and only stationary modes are considered (i.e. $M = 0$). The $N > 5$ harmonic modes are disregarded from the numerical simulations, inasmuch as an order of magnitude analysis suggests that they do not contribute to more than 0.42% to the total perturbation kinetic energy for any of the considered cases.

A posteriori calculation of growth-rate metrics can be instructive for the analysis and interpretation of results. Here, growth rate is computed as

$$\alpha_i = -\frac{1}{|f'_{(m,n)}|} \frac{d|f'_{(m,n)}|}{dx}, \quad (6.5)$$

with $|f'_{(m,n)}|$ measured at a chosen wall-normal location. Due to the ability of the HNS equations to describe the development of both modal and non-modal perturbations in fully non-parallel flows, the choice of perturbation quantity f' has an effect on the value of the resulting growth rate. In this chapter, the maximum amplitude of chordwise perturbation velocity $|u'|$ and the wall-normal integral of perturbation kinetic energy are used as perturbation quantities f' .




Figure 6.2: Spanwise perturbation velocity profile for mode (0, 1) ($\beta_1 = 0.1831$) at the most upstream measurement plane in the experiments of Rius-Vidales et al., 2025, corresponding to $x = 527$ in the present work. HNS results (black lines) and experimental measurements (red markers) for the low-amplitude DRE case (solid line and small markers) and high-amplitude DRE case (dashed line and large markers).

PERTURBATION BOUNDARY CONDITIONS

The boundary conditions for the HNS simulations are discussed here briefly for completeness, albeit a more in-depth overview of the implementation can be found in Westerbeek et al., 2024. A non-homogeneous Dirichlet condition is used at the inflow, namely, the solution to the local Orr-Sommerfeld (OS) eigenvalue problem considering $\beta = \beta_1$. The perturbation profile is normalised and the desired inflow amplitude is assigned to it. All additional modes of the truncated Fourier series (equation 6.4) are initiated with a zero amplitude and arise naturally further downstream as a result of nonlinear interactions. The spanwise wavelength and the range of investigated inflow amplitudes of the inflow mode are chosen to respectively match the spanwise spacing of DREs in the experiments of Rius-Vidales et al., 2025 ($\tilde{\lambda}_{DRE} = 7.5\text{mm}$) and the amplitude range dictated by the DRE arrangement. Figure 6.2 shows a comparison between the experimental spanwise perturbation profiles measured in Rius-Vidales et al., 2025 and the current simulation results. The comparison is made at the most upstream experimental measurement plane located at $x/c = 0.1$. The experimental low-amplitude and high-amplitude DRE cases and the closest available HNS result are shown. In the present work, the amplitude of the low-amplitude DRE case is most closely approximated by an HNS inflow amplitude of $A_0 = 4.5 \times 10^{-3}$ and the high-amplitude DRE case by $A_0 = 1.2 \times 10^{-2}$. Thus, the range of forcing conditions responsible for the experimentally observed upstream movement of the transition front is contained in the current amplitude sweep that ranges from $A_0 = 1.1 \times 10^{-3}$ up to $A_0 = 1.4 \times 10^{-2}$.

At the solid wall, ($\eta = 0$), all velocity-perturbation components are subject to homogeneous Dirichlet boundary conditions, enforcing no-slip and no-penetration. The perturbation pressure is solved implicitly and does not require a specific boundary condition. The outflow boundary is commonly associated with spurious reflections if left untreated as it interacts with incoming disturbances. This effect is particularly notorious in incompressible time-asymptotic formulations, see Dobrinsky, 2003, Kloker et al., 1993,

and Streett and Macaraeg, 1989 among others. To address this, an outflow buffer region in which downstream-travelling instabilities are artificially damped is implemented (see Fig. 6.1). The buffer region starts at 85% of the total domain length L and extends to the outflow boundary. Herein, perturbations are gradually damped following a hyperbolic tangent distribution in x , see Westerbeek et al., 2024 for details.

Homogeneous Dirichlet boundary conditions are imposed on the top boundary for all perturbation velocities, to enforce the decay of perturbation modes in the inviscid freestream. An exception is made for the wall-normal velocity of the non-oscillatory MFD mode, which is left free to accommodate the spatial growth of the mean flow. A similar approach has been considered in NPSE computations (Herbert, 1997). Additionally, the top boundary is placed sufficiently far from the wall to minimise potential artefacts stemming from the adaptation of the mean flow distortion in the vertical direction.

DISCRETISATION

The system of spanwise-harmonic perturbation modes equation ?? is discretised in the wall-normal direction using a spectral collocation method with a Chebyshev polynomial basis. The differentiation matrix suite by Weideman and Reddy, 2000 is used to provide the first and second-order differentiation matrices. Derivatives in the streamwise direction are approximated using a fourth-order central finite difference scheme for both the first- and second-order streamwise derivatives. For all investigated cases in this chapter, the domain is discretised using 80 Chebyshev collocation points in η that were clustered near the wall. In the chordwise direction, 1500 equispaced grid points are used. The results were subjected to a grid sensitivity study where both spatial axes were independently refined to ensure results were not significantly affected by the discretisation.

6.2.4. ENERGY-BALANCE EQUATIONS OF THE STATIONARY SPANWISE-HARMONIC PERTURBATIONS

The mechanisms of interaction between incoming CFI and the hump will be scrutinised by the use of energy-balance equations of the stationary spanwise-harmonic perturbation modes. These read as

$$0 = \mathcal{P}_{\beta_n} + \mathcal{T}_{\beta_n} + \mathcal{D}_{\beta_n} + \mathcal{W}_{\beta_n} + \mathcal{N}_{\beta_n}, \quad \text{with } n = 0, 1, \dots, N, \quad (6.6)$$

where n expresses the Fourier space of perturbations with spanwise wavenumber β_n and N is the number of spanwise modes considered.

Each term in equation 6.6 characterises a distinct mechanism of kinetic-energy transfer particular to perturbations of spanwise wavenumber β_n . Namely, $\mathcal{P}_{\beta_n} = \mathcal{P}_{\beta_n}(x, y)$ expresses the exchange of kinetic energy between the unperturbed base flow and perturbations β_n , $\mathcal{T}_{\beta_n} = \mathcal{T}_{\beta_n}(x, y)$ is the advection of kinetic energy of perturbations β_n by the base flow, $\mathcal{D}_{\beta_n} = \mathcal{D}_{\beta_n}(x, y)$ characterises the viscous effects, $\mathcal{W}_{\beta_n} = \mathcal{W}_{\beta_n}(x, y)$ and $\mathcal{N}_{\beta_n} = \mathcal{N}_{\beta_n}(x, y)$ respectively express the work of perturbation pressure and nonlinear interactions between harmonic modes and the mean flow distortion. By the stationary nature of perturbations, equation 6.6 expresses a balance; i.e. the corresponding physical mechanism remains in temporal equilibrium in the space of perturbations β_n . Inspection of the spatial behaviour of the terms of equation 6.6 provides insight into the evolution of corresponding physical mechanisms. Furthermore, their respective sign

informs whether they act towards destabilising, i.e. > 0 , or stabilising, i.e. < 0 , perturbations β_n locally in space.

To aid in the analysis and interpretation of linear perturbation effects, the production term of equation 6.6 considering the fundamental perturbation, β_1 , i.e.

$$\mathcal{P}_{\beta_1} = -\frac{2\pi}{\beta_1} \int_S \Lambda_{\beta_1} \, dx dy, \quad (6.7)$$

is decomposed as

$$\begin{aligned} \Lambda_{\beta_1}(x, y) = & \left(\hat{u}_{(0,1)} \hat{u}_{(0,1)}^\dagger + \text{c.c.} \right) \frac{\partial u_B}{\partial x} + \left(\hat{u}_{(0,1)} \hat{v}_{(0,1)}^\dagger + \text{c.c.} \right) \frac{\partial u_B}{\partial y} \\ & + \left(\hat{v}_{(0,1)} \hat{u}_{(0,1)}^\dagger + \text{c.c.} \right) \frac{\partial v_B}{\partial x} + \left(\hat{v}_{(0,1)} \hat{v}_{(0,1)}^\dagger + \text{c.c.} \right) \frac{\partial v_B}{\partial y} \\ & + \left(\hat{w}_{(0,1)} \hat{u}_{(0,1)}^\dagger + \text{c.c.} \right) \frac{\partial w_B}{\partial x} + \left(\hat{w}_{(0,1)} \hat{v}_{(0,1)}^\dagger + \text{c.c.} \right) \frac{\partial w_B}{\partial y}, \end{aligned} \quad (6.8)$$

where the dagger represents a complex conjugate and u_B , v_B , and w_B refer to the base flow velocities in the x , y , and z -directions respectively. The integration space S denotes a surface in the x - y plane that encompasses the hump in x and vertically extends from the wall to the free-stream. Following the approach of Albensoeder et al., 2001, production is expressed into four contributions,

$$\mathcal{P}_{\beta_1} = I_1^{\beta_1} + I_2^{\beta_1} + I_3^{\beta_1} + I_4^{\beta_1}, \quad (6.9)$$

each of which expresses a distinct mechanism relating the exchange of kinetic energy between the unperturbed base flow and fundamental perturbations (Lanzerstorfer & Kuhlmann, 2012a; Loiseau et al., 2016; Picella et al., 2018). Equation 6.9 is obtained by introducing the perturbation decomposition

$$\begin{aligned} \mathbf{q}'_{(0,1)} &= \mathbf{q}'_{t,(0,1)} + \mathbf{q}'_{n,(0,1)} = [q'^1_{t,(0,1)} \, q'^2_{t,(0,1)} \, q'^3_{t,(0,1)}]^T + [q'^1_{n,(0,1)} \, q'^2_{n,(0,1)} \, q'^3_{n,(0,1)}]^T \\ &= [\hat{q}^1_{t,(0,1)} e^{i\beta z} \hat{q}^2_{t,(0,1)} e^{i\beta z} \hat{q}^3_{t,(0,1)} e^{i\beta z}]^T + [\hat{q}^1_{n,(0,1)} e^{i\beta z} \hat{q}^2_{n,(0,1)} e^{i\beta z} \hat{q}^3_{n,(0,1)} e^{i\beta z}]^T \end{aligned} \quad (6.10)$$

into the original form of production, where the superscripts 1, 2, and 3 here denote vector components in the x -, y -, and z -directions, respectively. In equation 6.10, the vector field $\mathbf{q}'_t$ characterises the perturbation that is locally tangent to the base flow streamlines; i.e. pointing in the local direction of the (base) flow. The complementary vector field $\mathbf{q}'_n$ characterises the perturbation that is locally normal to the base flow streamlines; i.e. which acts in the cross-stream direction. It must be noted that both $\mathbf{q}'_t$ and $\mathbf{q}'_n$ are stationary spanwise-periodic perturbation fields. The reader is referred to Casacuberta et al., 2022 and Casacuberta et al., 2024 for explicit analytic expressions of $\mathbf{q}'_t$ and $\mathbf{q}'_n$, as well as details on their mathematical formulation.

The term $I_2^{\beta_1}$ of equation 6.9 will be shown to add the main contribution to production $\mathcal{P}_{\beta_1}$. It reads as

$$I_2^{\beta_1} = -\frac{2\pi}{\beta_1} \int_S \Lambda_2^{\beta_1} \, dx dy, \quad (6.11)$$

with

$$\begin{aligned}\Lambda_2^{\beta_1}(x, y) = & \left(\hat{q}_t^1 \hat{q}_n^{1\dagger} + \text{c.c.} \right) \frac{\partial u_B}{\partial x} + \left(\hat{q}_t^1 \hat{q}_n^{2\dagger} + \text{c.c.} \right) \frac{\partial u_B}{\partial y} \\ & + \left(\hat{q}_t^2 \hat{q}_n^{1\dagger} + \text{c.c.} \right) \frac{\partial v_B}{\partial x} + \left(\hat{q}_t^2 \hat{q}_n^{2\dagger} + \text{c.c.} \right) \frac{\partial v_B}{\partial y} \\ & + \left(\hat{q}_t^3 \hat{q}_n^{1\dagger} + \text{c.c.} \right) \frac{\partial w_B}{\partial x} + \left(\hat{q}_t^3 \hat{q}_n^{2\dagger} + \text{c.c.} \right) \frac{\partial w_B}{\partial y},\end{aligned}\quad (6.12)$$

where the subindex (0, 1) is removed for conciseness. The $I_2^{\beta_1}$ term describes the exchange of kinetic energy between the unperturbed base flow and streamwise-tangent perturbations, $\mathbf{q}'_t$, by the action of cross-stream perturbations, $\mathbf{q}'_n$. A mechanism also seen in the well-known lift-up effect (Ellingsen & Palm, 1975; Landahl, 1975). A treatment analogous to equations 6.11 and 6.12 may be performed for the remaining terms of decomposition 6.9, namely $I_1^{\beta_1}$, $I_3^{\beta_1}$, and $I_4^{\beta_1}$, see Casacuberta et al., 2025.

6.2.5. PERTURBATION ADVECTION EQUATIONS

A framework based on the so-called perturbation advection equations complements the analysis of the energy-balance equations presented above. As will be shown in the discussion below, the presence of the smooth surface hump significantly distorts the organisation of both the base flow and the stationary perturbation field. Towards isolating this effect, a framework is introduced to simulate base-flow advection of incoming perturbations. Isolating the perturbation convection term from equation ??a, i.e. neglecting production, pressure effects, nonlinear interactions and viscous effects for a stationary flow, yields:

$$(\mathbf{Q} \cdot \nabla) \mathbf{q}' = 0, \quad (6.13)$$

where the perturbation ansatz is identical to that of the HNS framework (see equation 6.4). Physically, the solution of equation 6.13 governs the advection of an incoming perturbation profile according to the three-dimensional streamlines of the base flow, akin to a scalar transport process. The perturbation advection framework is numerically implemented as an adaptation of the HNS framework and employs the same discretisation and boundary conditions as a result.

6.3. TOPOLOGY OF THE BASE FLOW

With the hump present, the organisation of the unperturbed steady base flow is altered significantly. This section describes the evolution of the base flow in the vicinity of the hump and compares it with the reference case (i.e. no hump present). The focus lies on the development of base flow quantities in the vicinity of the hump that are known to be important to the flow stability.

6.3.1. EFFECT OF THE HUMP ON THE BASE FLOW BOUNDARY LAYER

The spatial evolution of the chordwise pressure gradient, $\partial P / \partial x$, and key boundary layer parameters are illustrated in Fig. 6.3.




Figure 6.3: Evolution of the base flow in the vicinity of the hump. (a) Chordwise pressure gradient. Black dashed line indicates boundary layer thickness, δ_{99} . The hump is shown in orange. (b) Boundary-layer parameters based on the chordwise velocity (U), namely δ_{99} (black), δ_{disp} (purple), $10 \times \delta_{mom}$ (blue), and chordwise skin friction (red, right axis). Dashed lines correspond to the hump case and solid lines correspond to the reference case. The orange region denotes the chordwise extent of the hump and vertical dashed line indicates its center. Dash-dotted lines segregate the main region of influence of the hump.

The boundary layer thicknesses are measured from the local (deformed) wall in the direction normal to the reference clean wing surface and presented against the η coordinate. The hump strongly affects all boundary layer quantities at the hump location and in its proximity. Distortions can be seen in the chordwise region from $x = 450$ (all spatial coordinates are non-dimensionalised with l_{ref} , see table 6.1) onward while the boundary layer recovers to the reference conditions around $x \approx 820$. The hump incurs a notable distortion as well on the pressure gradient field, initially decelerating the flow and thickening the boundary layer. As the flow traverses the hump, it experiences successive acceleration and deceleration events.

The skin friction coefficient remains positive throughout the streamwise extent of the hump, thus signifying the absence of laminar separation and reverse flow in the chordwise direction. Previous studies on the effect of surface features on incoming CFI linked the formation of a local region of reverse flow with the rapid spatial growth of CFI (Hosseini & Fasel, 2016; Westerbeek et al., 2023; Xu et al., 2017). The absence of reverse flow around the hump in the present work ensures that the alteration of the behavior of CFI at the hump is not explicitly driven by such effects.

The effect of the hump on the base flow velocity profiles is quantitatively illustrated in Fig. 6.4. The hump alters the evolution of the wall-normal velocity of the base flow, effectively creating a transversal motion of successive upwash and downwash events in x . This effect extends further than the viscous region (i.e. $\eta > \delta_{99}$). The chordwise velocity experiences a rapid deceleration in x as a result of the hump-induced pressure gradient.




Figure 6.4: (a-e) Base flow components of chordwise velocity (U) (red dotted and black solid lines), and wall-normal velocity ($V \times 10$) (red dash-dotted and black dashed lines) in the hump (red) and reference (black) cases. (f) Spatial evolution of the crossflow velocity in the hump case. Black dashed line indicates the δ_{99} . Grey dashed lines indicate loci of inflection points. The hump is indicated in orange.

6

This is particularly accentuated in the near-wall region (i.e. $\eta < 4$). However, in contrast to the wall-normal velocity, this effect largely diminishes outside the boundary layer.

6.3.2. CROSSFLOW REVERSAL

Crossflow reversal, i.e. the inversion of the direction of the crossflow velocity component, is captured in the downstream vicinity of the hump apex in Fig. 6.4f. The crossflow component is defined as the velocity component orthogonal to the external inviscid streamline. Granted its inflectional shape, it is widely acknowledged to be the driving source of the crossflow instability (W. S. Saric et al., 2003). The local and rapid chordwise change of pressure gradients by the hump leads to crossflow reversal. Within the scope of incoming CFI interacting with forward-facing step flow, Tufts et al., 2017 ascribe crossflow reversal to the effect of local adverse pressure gradients. Moreover, Eppink, 2020 states that the modification of the crossflow velocity in the vicinity of a forward-facing step influences the rotation direction of the crossflow instability at the step. Despite the shallow hump shape, crossflow reversal arises prominently in the presently investigated flow (Fig. 6.4f).

In contrast to geometry-induced modification of pressure gradients, Wassermann and Kloker, 2005 investigated the effect of external (i.e. inviscid) pressure gradient change-over on the development of stationary CFI. They report that the pressure gradient change-over induces reversal of the crossflow component in the boundary layer. The results of a linear local stability analysis demonstrated that the Adverse Pressure Gradient (APG) and subsequent crossflow reversal act towards stabilising the incoming CFI in a modal sense.

The authors concluded that “the reversed basic crossflow close to the wall leads to a redistribution of spanwise shear in the mean flow distorted by the steady crossflow vortices coming from upstream.” This may be partially reconciled with the observed stabilisation of CFI by the hump in the experiments of Rius-Vidales et al., 2025, but raises questions on how the stabilisation effect is dependent on the amplitude of incoming instabilities, as quantitatively characterised in this present work. Furthermore, Wassermann and Kloker, 2005 studied the response of the boundary layer to the change-over from Favorable Pressure Gradient (FPG) to APG without reverting to FPG, which is the flow environment in the hump case.

In the present scenario, the crossflow component reverts direction almost immediately downstream of the hump apex. The region of crossflow reversal extends from $x = 640$ to $x = 725$. The peak of reverse crossflow velocity reaches a maximum in x of $\approx -0.05\bar{U}_{\text{ref}}$. However, just upstream of the hump apex, the crossflow component is significantly strengthened, i.e. in the classic direction, in comparison to reference conditions, due to the enhanced favourable pressure gradient by the hump. Crossflow reversal is also localised in the wall-normal direction, i.e. it does not extend far into the upper boundary layer. In a fashion similar to the work of Hosseinverdi and Fasel, 2016, crossflow reversal here extends from the wall to approximately $\eta = 2$. This can be attributed to the lower local spanwise momentum in this region of the boundary layer. Sufficiently downstream of the hump, the base flow relaxes to reference conditions; from $x = 730$, the crossflow component returns to the classic positive direction.

6.4. EVOLUTION OF THE PERTURBATION FIELD

To analyse the response of the perturbation field to the presence of the hump, linear and non-linear formulations of the HNS framework are used. In both cases, the perturbation evolution is initiated by the amplification in x of the fundamental CFI mode. As described above in section 6.2, a stationary (i.e. zero temporal frequency) crossflow eigenmode with fundamental spanwise wavelength, $\beta_1 = 0.1831$, is prescribed at the inflow. This is computed as the solution to the local Orr-Sommerfeld eigenvalue problem. Linear HNS simulations account solely for the growth of the fundamental crossflow mode. In contrast, the effect of finite-amplitude crossflow perturbations on the interaction mechanism with the hump is assessed using the fully nonlinear HNS framework.

6.4.1. EFFECT OF THE HUMP ON LINEAR PERTURBATION DEVELOPMENT

The spatial amplification of the fundamental CFI is quantified in Fig. 6.5. Instability growth is largely unaffected by the presence of the hump up to its leading edge (i.e. $x < 500$). In contrast, the impact of the hump on the CFI development manifests both locally in the vicinity of the hump as well as further downstream of it.

The perturbation amplification factor, N , is depicted in Fig. 6.5a. The N -factor is calculated locally as:

$$N(x) = \ln \left(\frac{A(x)}{A_0} \right), \quad (6.14)$$

with the amplitude measure A based on either the maximum of the chordwise velocity perturbation or the wall-normal integral of perturbation kinetic energy (section 6.2).




Figure 6.5: Spatial evolution of the fundamental CFI mode ($\beta_1 = 0.1831$) in the hump (red) and reference (black) cases quantified as: (a) amplification factor, N , based on the peak of chordwise velocity perturbation (solid lines) and total perturbation kinetic energy E_k (dash-dotted lines), and (b) spatial growth rate. Growth metrics are based on the peak of chordwise velocity perturbation (solid and dashed lines) and on total perturbation kinetic energy (dash-dotted lines). The orange-shaded region indicates the hump's spatial extent.

6

The reference amplitude $A_0 = A(x_0)$ is taken at the domain inflow ($x_0 = 214l_{\text{ref}}$). Growth rates based on these two quantities (equation 6.5) are additionally illustrated in Fig. 6.5c. The evolution of growth rates is qualitatively similar. Specifically, the effect of the hump on the incoming instability can be decomposed into three distinct regions. The first region spans approximately the upstream half of the hump (i.e. $500 \leq x \leq 580$). In this range, the CFI is destabilised, when compared to reference conditions. This trend correlates with the spatial deceleration of chordwise base flow velocity and concurrent increase in the crossflow component, as evident in Fig. 6.4. Conversely, the second region (i.e. $580 \leq x \leq 800$), coincides partly with the crossflow reversal seen in the base flow. In this region, the CFI stabilises in the vicinity of the hump's apex and shortly downstream of it. In contrast, upstream of the hump's trailing edge ($x \approx 720$) the growth rates become again lower than reference values.

Finally, a third region of evident growth rate modification due to the hump is identified (i.e. $800 \leq x \leq 1100$). Within this region, the hump appears to have a significant and persistent effect on the perturbation growth rate as far downstream as $x \approx 1100$, in a range of approximately two hump lengths downstream of the hump's apex. This is in stark contrast to the effect of the hump on the base flow. The latter is effectively recovered to the reference case from $x \approx 820$ onward, as shown in figures 6.3 and 6.4. In combination, the incoming CFI experiences successive events of stabilisation and destabilisation as it traverses over the hump. However, in an integral manner, the instability exits the hump's region of influence and recovers to the reference growth rate at $x \approx 1100$ with almost 60% of the amplitude of the reference case ($\Delta N \approx -0.5$ based on $|u'|$), fur-




Figure 6.6: (a-f) Chordwise (solid line) and wall-normal (dashed line) perturbation velocity profiles over the hump (red) and reference case (black), at selected chordwise stations. For ease of visualisation, $|\hat{v}|$ is scaled by a factor 10. Horizontal dashed blue lines indicate the location of inflection points in the crossflow profile of the base flow in the hump case. (g) Contour map of the amplitude function of the chordwise perturbation velocity from $|\hat{u}| = 0$ (white) to $|\hat{u}| = 0.5$ in 10 levels. Red solid and black dashed lines indicate the location of local maxima of $|\hat{u}|$ in the hump and reference case, respectively.

ther reconciling with the experimental observations of transition delay by Rius-Vidales et al., 2025.

To gain insight into the effect of the hump on the perturbation evolution locally at the hump, as well as downstream of it, the amplitude shape profile (see equation 6.4) is inspected (Fig. 6.6). For the purpose of this analysis, only the chordwise ($\hat{u}$) and wall-normal ($\hat{v}$) amplitude functions are considered. The spanwise velocity component $\hat{w}$ is not shown here as it largely resembles $\hat{u}$. As the incoming CFI ascends the hump, the wall-normal perturbation velocity is amplified, which is then largely compensated during the descent downstream of the hump apex. In contrast, the chordwise component ($\hat{u}$) demonstrates a more complex response. The shape profile $\hat{u}$ shows a near-wall deficit compared to the reference case upstream of the hump apex. This is accompanied by an overall broadening of the shape profile and displacement of its local maximum further away from the wall. In this regard, the perturbation shape follows the behaviour of the base flow closely (Fig. 6.4). This is further reconciled by a consistent and evident correlation of the wall-normal location of maximum $|\hat{u}|$ with the location of inflection points in the crossflow velocity profiles of the base flow (identified by blue dashed lines in Fig. 6.6 and visualised in Fig. 6.4f).

The resultant modifications in perturbation shape are also linked to changes in amplitude. In the vicinity of the hump apex and immediately downstream of it, the perturbation experiences a rapid spatial amplification accompanied by a displacement of the peak perturbation velocity towards the wall. Within a short chordwise range ($718 < x < 739$), the $\hat{u}$ profile features two local maxima. For comparison, the base flow crossflow

component experiences reversal and as a consequence features two inflection points in the chordwise region $719 < x < 740$.

The concurrent presence of two maxima is responsible for a pronounced discontinuity in the growth rate derived from the maximum of chordwise perturbation velocity at $x \approx 740$ as shown in Fig. 6.5c. At this location, the maximum perturbation velocity switches between the two local peaks, thus the evolution of the growth rate in x is not continuous. Naturally, this does not manifest in the energy-based growth rate due to its integral evaluation. A double-peak topology of the u -perturbation shape has been reported in previous studies involving CFI interacting with a hump (Westerbeek et al., 2023), as well as a forward-facing step (Casacuberta et al., 2021; Casacuberta et al., 2022). Casacuberta et al., 2021 argue that the near-wall peak is the manifestation of an additional perturbation structure (namely, perturbation streaks) that forms locally at the upper step corner and co-exist with the CFI developing above and farther from the wall.

6.4.2. LOCAL AND NON-LOCAL LINEAR STABILITY ANALYSIS

Due to the convective nature of the CFI, the perturbation development at any chordwise location is expected to be strongly dependent on its spatial history. The global nature of the HNS formulation, which is based on the full NS equations, can capture this dependency. In contrast, the unique quality of locality of the Orr-Sommerfeld (OS) operator renders it insensitive to the history of the perturbation and the chordwise variation of the base flow. Thus, these effects can be segregated by comparing the solutions of local stability analysis for the fundamental mode with the corresponding solution of the linear HNS equations. Towards this goal, the alteration of the perturbation shape at the hump may be ascribed to the changes in the local eigensolution and the rise of new unstable modes as a consequence of the altered base flow.

The Orr-Sommerfeld analysis is carried out in the range $x \in [214, 1200]$ at 1500 equally distributed stations for the hump and reference cases. Selected eigensolutions of the local OS eigenspectrum are portrayed in Fig. 6.7. Based on the spatial stability formulation of the problem, results are interpreted in terms of the complex eigenvalue $\alpha = \alpha_r + i\alpha_i$, for which the real component corresponds to the chordwise wavenumber of the instability mode (not shown here) while the imaginary component expresses the chordwise growth rate. The spatial growth rate computed from the peak chordwise perturbation velocity obtained from HNS (equation 6.5) is shown for reference. It must be stressed that the HNS solution considers a series of terms in the Navier-Stokes equations that are neglected in the OS formulation, such as the chordwise flow derivatives. As such, even for the reference case, slight deviations in the growth rate predicted by the two approaches are measured. This is particularly evident at upstream locations in the domain ($x < 500$), where the boundary layer is strongly accelerating in x and non-parallel effects are non-negligible in the HNS solution (Fig. 6.7).

The growth rates obtained as solutions to the OS and HNS equations are significantly affected by the presence of the hump. The fundamental local CFI mode (cyan markers in Fig. 6.7a) starts to deviate from the reference case at approximately $x = 500$. Within the range occupied by the hump ($500 < x < 720$) the differences between OS and HNS solutions are also evident in the shape of corresponding eigensolutions. The shape function of the OS eigenmode develops a double-peak structure rapidly in x (see figure 6.7d-f)




Figure 6.7: (a) Orr-Sommerfeld spatial growth rate α_i for $\beta_1 = 0.1831$ in the hump case corresponding to the original CFI mode tracked from the inflow (cyan markers) and a novel near-wall mode in the CF reversal region (purple markers). Reference case is shown in black markers. Corresponding HNS solution in the hump case is shown for comparison (red dashed line). Orange shaded region denotes the spatial extent of the hump. Vertical dashed lines indicate the region of crossflow reversal in the base flow. (b-h) Local eigenfunction shape of chordwise velocity at selected chordwise stations in the hump case (cyan and purple solid lines), corresponding to the original and new near-wall CFI modes, respectively) and reference case (black solid line). Corresponding HNS solution in the hump case is shown for comparison (red dashed line). Horizontal dashed lines indicate the location of inflection points in the base flow crossflow profile.

and returns to the single-peaked structure immediately downstream of the hump. Particularly evident in a short region downstream of the hump's apex ($621 < x < 740$), is the secondary peak in the near-wall region ($\eta \approx 2$). This range shows a strong correlation to the region of crossflow reversal (indicated by the vertical dashed lines in Fig. 6.7) and is accompanied by rapid spatial stabilisation in x , evident by the positive OS growth rate. In contrast, the HNS solution experiences a largely milder spatial evolution, with a double-peak structure only briefly manifesting at the downstream end of the hump ($x \approx 723$). Similarly, the growth rate calculated by HNS experiences a shift to less unstable values, albeit remaining negative. The inability of the OS solution to reconcile with the corresponding HNS growth rate in the range $500 < x < 1100$ is strongly suggestive of growth mechanisms active in this region which cannot be described on the basis of local purely parallel eigensolutions. This will be discussed in sections 6.4.3 and 6.4.4.

Interestingly, the modified boundary layer due to the hump supports a new instability eigenmode. This appears in the OS eigenspectrum downstream of the hump apex




Figure 6.8: Spanwise harmonic reconstruction of (a-g) local OS and (h-n) HNS solutions of the incoming CFI mode normalised by $u'_{\max}$. The chordwise perturbation velocity is illustrated, considering contour levels from -1 (blue) to 1 (red) at selected chordwise stations

6

at $x = 647$ and is located in the unstable half-plane in the short chordwise extent between $x \approx 660$ and $x \approx 710$ (see purple markers in Fig. 6.7a). The OS operator admits this mode as a valid, albeit only briefly unstable, solution to the local boundary layer stability problem. This mode is not directly evident in the HNS solution, but similarities in the shape profile are observed. The presence of an OS eigensolution, additional to the incoming CFI and qualitatively resembling the behaviour of the one identified here, was described by Krochak et al., 2022 in backward-facing step flow and associated to a new inflection point appearing in addition to the classical inflection of crossflow velocity profiles shortly downstream of the step apex.

To further elucidate the character of the modified CFI in relation to the novel near-wall mode induced by the hump, their mode shapes are examined in the $\eta - z$ plane, normalised by their respective maximum. Specifically, the chordwise velocity component of the local stability solution is expanded in the spanwise direction through the application of the OS harmonic ansatz and displayed in Fig. 6.8a-g for the incoming CFI mode (cyan lines in Fig. 6.7b-h). For reference, the HNS solution is shown in Fig. 6.8h-n. The local CFI mode predicted by the OS solution experiences significant topological changes due to the hump. At $x = 647$, a localised deformation in the profile occurs as the structure tilts to the positive z -axis near the wall. This is strongly reminiscent of the observations of Wassermann and Kloker, 2005 who studied the effect of a Favourable to Adverse Pressure Gradient (FPG to APG) change-over on the evolution of stationary CFI. The evolution of the mode topology downstream of the hump's apex is noteworthy. The mode is continuously *stretched* in the z -direction and evolves from a classical crossflow instability shape to a distinct out-of-phase double-peak structure (Fig. 6.8d). Concurrently, in the region where the double peak forms and near-wall phase inversion is mostly active ($650 < x < 720$), the mode is predicted to be linearly stable by the OS




Figure 6.9: Scalar advection of CFI-like perturbations in (a-e) the hump and (g-k) the reference case, seeded at $x = 600$, i.e. the virtual location of the hump apex. Chordwise velocity-perturbation contours normalised relative to $u'_{\max}$ ranging from -1 (blue) to 1 (red) at selected chordwise stations. The hump is indicated in orange.

operator, i.e. $\alpha_i > 0$ (figure 6.7). In the same region, the HNS solution is found to be stabilised, in comparison with the reference case.

The origin of deformation of the CFI into a double-peak structure is further explored here through a simple heuristic experiment, where a pre-existing perturbation system is exposed to pure scalar transport by the base flow. This is facilitated directly within the HNS framework, by disabling all terms in the HNS equations except the advection term, as outlined in section 6.2.5. The effect of scalar transport on a pre-computed perturbation obtained as a solution to the Orr-Sommerfeld eigenvalue problem and introduced at the x -location of the hump apex is shown in Fig. 6.9 for both hump and reference cases. The perturbation system that advects over the hump is found to be prominently tilted in the $y-z$ plane, as opposed to the reference case, due to the crossflow reversal. Although many other important flow mechanisms are neglected, the resulting shape closely resembles the double-peaked shape observed in the full HNS simulation and the OS analysis above (Fig. 6.8). This further confirms a key element in the overall mechanism of stability manipulation by the hump, namely the spanwise transport varying with the distance from the wall and the resulting distortion of incoming CFI.

Finally, the origin and influence of the novel near-wall unstable eigenmode shown in Fig. 6.10 are briefly discussed. This perturbation resembles a classic crossflow instability in a structural sense. However, in contrast to the incoming CFI shown in Fig. 6.8, the mode shape in the $\eta-z$ plane of the new near-wall mode is tilted anti-clockwise. This can also be traced back to the topology of the local base flow, as this new eigenmode destabilises purely within the region of crossflow reversal (Fig. 6.4). This agrees




Figure 6.10: Spanwise harmonic reconstruction of the local OS (secondary) unstable near-wall mode normalised by $u'_{\max}$. The chordwise perturbation velocity is illustrated, considering contour levels from -1 (blue) to 1 (red) at selected chordwise stations.

with the results of Eppink, 2020 and Krochak et al., 2022 for CFI-dominated flows over backward-facing steps, where the presence of step-induced inflectional profiles embedded in a region of crossflow reversal was reported to give rise to new stationary instability modes near the wall. The newly identified OS eigenmode in the present work rapidly stabilises once it exits the crossflow reversal region. As such, in conditions of pure linear instability growth, the near-wall mode is unlikely to reach a significant amplitude to be observable as a separate mode co-existing with the original incoming CFI. Therefore, its role in linear or low instability amplitude scenarios is likely to be minor. However, as will be shown in the following sections, the role of such near-wall modes can be elevated in importance in environments of high-amplitude and non-linear instability development.

6

6.4.3. NONLINEAR EFFECTS IN CFI-HUMP INTERACTIONS

The finite amplitude of incoming CFI interacting with a hump was found to significantly influence the transition location in the experiments of Rius-Vidales et al., 2025, thus suggesting a key role played by the non-linear perturbation effects. In the present section, these effects are investigated numerically by varying the inflow amplitude of the pre-imposed fundamental CFI mode (section 6.2.3) in a fully non-linear HNS framework. All other parameters, such as base flow, geometry, and wave specifications, are kept identical to the linear analysis described above in sections 6.4.1 and 6.4.2. The evolution of the fundamental CFI (i.e. $\beta = \beta_1$), its four high-order harmonics (i.e. $\beta = \beta_n, n = 2-5$), and the mean flow distortion (MFD, i.e. $\beta = 0$) are considered here. The initial amplitude of the fundamental perturbation is varied from $A_0 = 1.1 \times 10^{-3}$ to 1.4×10^{-2} , while all other harmonics and the MFD are initiated as null at the domain inflow. The MFD and higher harmonics arise naturally in x through nonlinear interactions.

The chordwise evolution of the amplitude of the fundamental mode $\beta = \beta_1$ and the harmonic $\beta = \beta_2$ are illustrated in figure 6.11. For the lowest initial amplitude considered in this chapter, the perturbation evolution is linearly dominated in the entire domain. Consequently, nonlinear mechanisms increasingly dominate the behaviour as the inflow amplitude of the fundamental perturbation is increased. Significant qualitative differences can also be observed between the amplitude evolution of the $\beta = \beta_1$ and $\beta = \beta_2$ modes as a function of inflow amplitude. However, higher harmonics follow




Figure 6.11: Amplitude of (a) fundamental mode $\beta = \beta_1$ and (b) first harmonic mode $\beta = \beta_2$ in the hump (red) and reference cases (black) for inflow amplitudes $A_0 = [0.12, 0.20, 0.32, 0.53, 0.86, 1.40] \times 10^{-2}$ marked respectively by the circle, plus, triangle, asterisk, diamond, and square markers at the inflow. Increasing opacity of lines indicates an increase in amplitude. The hump and its chordwise extent are indicated in orange. Purely linear amplitude evolution, obtained by linearised HNS, is indicated by dotted blue line.

a trend closely resembling the $\beta = \beta_2$ mode. Hence, the behaviour of modes with $\beta > \beta_2$ is here disregarded in the discussion of results. For low initial amplitudes, the fundamental mode is generally stabilised and exits the simulation domain at a significantly lower amplitude than in the corresponding reference case (Fig. 6.11a). This is also observed in the results of the purely linear simulation. The amplitude reduction occurs effectively at all chordwise positions downstream of the hump's apex. In contrast, as the inflow amplitude increases, an increasingly strong destabilisation is observed at the downstream end of the hump (i.e. $x \approx 720$). For initial mode amplitudes greater than $A_0 = 5.3 \times 10^{-3}$, the CFI downstream of the hump attains amplitudes higher than the associated reference case despite then decaying to a lower amplitude further downstream.

Linear growth of CFI is followed by subsequent stages of non-linear growth behaviour that appears similar to saturation, i.e. quasi-saturation, as reported in numerical (Casacuberta et al., 2022; Wassermann & Kloker, 2002) and experimental (Arnal, 1994; Reibert et al., 1996) work as well as non-linear stability analysis (Haynes & Reed, 2000). This originates from nonlinear energy exchange between harmonic modes and particularly between the fundamental mode and MFD. Specifically, the growth of mean flow distortion inhibits further growth of the fundamental mode, thus leading to a new (local) equilibrium among the spectral components of the perturbation system. This nonlinear saturation of fundamental CFI modes is generally a good predictor of flow breakdown to turbulence (Reibert et al., 1996). For the two highest inflow amplitudes shown here, namely $A_0 = 0.86 \times 10^{-2}$ and $A_0 = 1.4 \times 10^{-2}$, the fundamental CFI amplitude appears to saturate at chordwise locations $x = 952$ and $x = 890$, respectively¹. For comparison, preliminary simulations on longer domains reveal that the reference clean cases reach quasi-saturation only downstream of the presently simulated domain outflow.

The intermediate outcomes of the analysis so far provide a first insight into the experimental results of Rius-Vidales et al., 2025. On the one hand, at low inflow forcing amplitudes, the modification of incoming CFI by the hump is dominated linearly, i.e. the amplitude trend largely resembles that of the solution of linear HNS. A sequence of stabilising and destabilising events occurs, which, in a chordwise integral manner, lead to an overall decrease in mode amplitude. This corresponds to the low amplitude DRE case in the experiments of Rius-Vidales et al., 2025, where a transition delay of 14% chord length was achieved in the presence of the hump. On the other hand, as the inflow forcing amplitude increases, the localised destabilisation event at the downstream end of the hump rapidly increases the stationary instability amplitude above reference levels, thus creating enhanced spanwise and wall-normal flow shears, known to be fertile ground for the development of unsteady secondary instabilities (Malik et al., 1999; Wassermann & Kloker, 2002). Notwithstanding the fact that secondary instability evolution and breakdown are not modelled in the current HNS framework, the upstream shift of stationary instability growth is likely to be the cause of transition anticipation observed in the high-amplitude DRE case of Rius-Vidales et al., 2025. Time-resolved simulations or experiments are necessary to confirm this hypothesis and further explore the nature of unsteady breakdown.

¹A universal criterion for identifying CFI amplitude saturation is not generally available. In the present work, the onset of quasi-saturation is defined as the chordwise location at which the chordwise growth rate of the fundamental mode becomes zero and does not become negative again

For the harmonic mode $\beta = \beta_2$ (and by association higher harmonics not shown here for brevity), destabilisation shortly downstream of the hump's apex is seen for all choices of inflow amplitude assigned to the fundamental mode (figure 6.11b). This effect is qualitatively similar among all tested cases and is not spatially correlated to the stabilisation or destabilisation events governing the fundamental mode. The combination of these observations suggests a largely linear response of higher harmonics to the hump. A similar response has been observed in the interaction between harmonic CFI perturbations and a forward-facing step, which was found to be dominated by a secondary inflection point past the step (Casacuberta et al., 2022). Notwithstanding, in the downstream vicinity of the hump ($650 < x < 850$), higher harmonics are sufficiently destabilised to reach comparable amplitude to the fundamental mode.

6.4.4. ENERGY TRANSFER MECHANISMS IN THE INTERACTION BETWEEN HUMP AND CFI

The analysis above has suggested that the mechanisms of interaction between CFI and the hump are largely linearly dominated, inasmuch as each harmonic perturbation mode is primarily affected by hump-induced base flow modifications and only mildly influenced by mode-to-mode interactions. This is particularly evident for low inflow amplitudes of the fundamental mode. However, as the incoming CFI amplitude is increased, nonlinear effects appear to play an increasingly important role downstream of the hump, through fundamental mode saturation and strong growth of higher harmonics. To elucidate the role played by linear and non-linear perturbation mechanisms, an energy-budget analysis of the stationary spanwise perturbation modes is carried out, see section 6.2.4 for details. For ease of discussion, three cases are chosen, hereafter referred to as the *low*- ($A_{(0,1)} = 1.1 \times 10^{-3}$), *middle*- ($A_{(0,1)} = 3.7 \times 10^{-3}$) and *high*- ($A_{(0,1)} = 1.4 \times 10^{-2}$) amplitude cases. The spatial evolution of the energy budget terms (equation 6.6) applied to the fundamental perturbation, namely production ($\mathcal{P}_{\beta_1}$), advection ($\mathcal{T}_{\beta_1}$), viscous effects ($\mathcal{D}_{\beta_1}$), and the work of nonlinear interactions ($\mathcal{N}_{\beta_1}$), is depicted in Fig. 6.12 for both the flow over the hump and the reference case without hump. The work of perturbation pressure ($\mathcal{W}_{\beta_1}$) is significantly weaker than the other terms and is therefore neglected from the remainder of the analysis. The discussion of this section is exclusive to the fundamental perturbation, i.e. mode $\beta = \beta_1$. For the sake of illustration, figure 6.13 reproduces figure 6.12 considering mode (0, 2).




Figure 6.12: Energy budget distribution (equation 6.6) for the fundamental perturbation, $\beta = \beta_1$ in the (a-i) hump and (j-r) reference case. *Low*- (first column), *medium*- (second column), and *high*- (third column) amplitude cases. (a-c, j-l) Production ($\mathcal{P}_{\beta_1}$), (d-f, m-o) advection and viscous effects ($\mathcal{T}_{\beta_1} + \mathcal{D}_{\beta_1}$), and (g-i, p-r) work of nonlinear interactions ($\mathcal{N}_{\beta_1}$). The black dashed line corresponds to δ_{99} . Grey dashed lines indicate loci of crossflow inflection points. (s-u) y-integrated energy budgets for hump (red) and reference case (black): $\mathcal{P}_{\beta_1}$ (solid lines), $\mathcal{D}_{\beta_1} + \mathcal{T}_{\beta_1}$ (dashed lines) and $\mathcal{N}_{\beta_1}$ (dotted lines).




Figure 6.13: Energy budgets (Eq. 6.6) for mode $(0, 2)$, $\beta = \beta_2$. Identical notation as Fig. 6.12

In reference conditions, i.e. without the hump present, a dominant mechanism of linear production drives the spatial amplification of the fundamental CFI until, at least, the virtual location of the hump in all cases considered. This is exemplified by $\mathcal{P}_{\beta_1}$ adding the largest positive (i.e. destabilising) contribution to the energy budget. Considering the *low*- and *medium*-amplitude cases, the amplification of the fundamental CFI is linearly dominated until $x \approx 1100$ and $x \approx 800$, respectively; in both cases, well downstream of the virtual x -location of the hump. This is substantiated by the agreement between the solutions of linear and non-linear HNS simulations in the corresponding x -ranges, as demonstrated in figure 6.11.




Figure 6.14: Spatial distribution of the term $I_2 = -(2\pi/\beta_1)\Lambda_2^{\beta_1}$ (equation 6.11) applied to the fundamental perturbation, $\beta = \beta_1$. (a,b) *low*-, (c,d) *medium*-, and (e,f) *high*-amplitude cases for reference (left column) and hump (right column) conditions. The hump is indicated in orange. Dashed black line indicates δ_{gg} . Crossflow inflection points are indicated by the white dashed line. Red isolines identify regions exceeding 0.5 and 1 times the maximum contour level.

When the hump is present in the *low*-amplitude case, the mechanism of production maintains a dominant destabilising contribution locally around the hump, i.e. $\mathcal{P}_{\beta_1}$ is the only positive term (Fig. 6.12a). The terms of advection and viscous effects yield a negative contribution (i.e. they act stabilising, $\mathcal{T}_{\beta_1} + \mathcal{D}_{\beta_1} < 0$) and thus mainly balance the effect of production. The nonlinear mechanism does not yield a significant contribution to the energy budget. When the amplitude of the incoming CFI is increased in the *medium*- and *high*-amplitude cases, the influence of the non-linear mechanism shortly downstream of the hump appears more evident (figure 6.12h,i), albeit mainly stabilising (i.e. $\mathcal{N}_{\beta_1} < 0$). As such, the local CFI destabilisation downstream of the hump for the higher-amplitude cases is ascribed to the enhancement of the mechanism of linear production by the hump (Fig. 6.12u). An opposite trend is observed in the *low*-amplitude cases. That is, the mechanism of production is weakened by the presence of the hump, when compared to the reference case, which is associated to the decay of CFI amplitude locally after the hump apex, as illustrated in figure 6.11

The analysis further leverages the decomposition of the production term $\mathcal{P}_{\beta_1}$ as outlined in equations (6.9), (6.11), and (6.12). For all the cases explored in the present work, the term $I_2^{\beta_1}$ is found to add the main contribution to production (equation 6.9). The spatial distribution of $I_2^{\beta_1}$ is depicted in Fig. 6.14 pertaining to the fundamental mode $\beta = \beta_1$ of the perturbation. In reference conditions, the *strength* of production increases with the amplitude of incoming CFI, however, no significant qualitative or topological

differences are observed in the spatial evolution of $I_2^{\beta_1}$. In contrast, in the presence of the hump, both the magnitude and spatial topology of $I_2^{\beta_1}$ are significantly altered by the amplitude of the incoming CFI. In the *low*- and *medium*-amplitude cases, the *strength* of $I_2^{\beta_1}$ is notably reduced, as compared to reference conditions. This is particularly evident in the vicinity of the core of the incoming crossflow vortex, which approximately follows the upper inflectional point in the crossflow component of the boundary layer (upper white dashed line in Fig. 6.14). This agrees with the trends discussed above regarding the decay of both production and CFI amplitude in the same spatial region. However, in the region of crossflow reversal, delimited by the lower inflectional point (lower white dashed line in Fig. 6.14) and further downstream of the hump ($710 < x < 850$, $y \approx 2$), production is seen to increase locally. This near-wall destabilisation is particularly evident in the *high*-amplitude case. The emergence of enhanced momentum redistribution and production in the near-wall region can be traced to the outcomes of the linear analysis presented in section 6.4.1. More specifically, the novel near-wall structure emerging in the region of crossflow reversal (i.e. lower peak in the double-peak structure in Fig. 6.7) carries a strong spatial phase modification, evident as a spanwise shift, compared to the upper body of the incoming CFI perturbation. Such strong modifications in phase have been shown to impart significant changes on the wall-normal momentum distribution (Casacuberta et al., 2024). In addition, the presence of the lower inflectional point in the base flow in combination with the formation of these near-wall structures provide the necessary conditions for rapid spatial perturbation growth near the wall, as predicted by the emergence of a new unstable OS eigenmode in this region (see section 6.4.2). In the case of high-amplitude incoming CFI, the enhanced near-wall production is able to overwhelm the original incoming perturbation in terms of amplitude, leading to strong stationary structures that can drive transition as discussed in the following section.

6.5. CONCLUDING DISCUSSION

The present work strives to analyse and identify the underlying mechanisms of interaction between stationary CFI and a smooth surface hump. Specifically, the effect of incoming CFI amplitude in the mechanisms of such interaction is investigated. As shown by linear and nonlinear HNS simulations (see Fig. 6.11) as well as the experiments of Rius-Vidales et al., 2025, whether the hump produces an overall stabilising or destabilising effect appears to be strongly influenced by the amplitude of the incoming CFI. The analysis of changes to both base flow and perturbation systems covered in the above sections provides an initial set of principles for these interactions. These are to be taken as merely the starting point towards a more holistic understanding, since a wealth of key parameters including Reynolds number, hump shape and CFI wavelength have been kept constant in this chapter.

The effect of the hump on the laminar base flow is found to manifest as strong modifications in the chordwise and wall-normal velocities, and by association, changes in the boundary layer crossflow component. Crossflow reversal is detected at the leeward side of the hump. Nevertheless, the base flow recovers to the reference case conditions shortly downstream of the hump, in terms of velocity profile and integral quantities. The base flow changes are found to impact the stability of the flow to incoming CFI. These

are traced and quantified in this chapter using local linear stability analysis as well as linear and non-linear HNS simulations.

Energy budgets of the harmonic spanwise-perturbation modes based on the Reynolds-Orr identity are used to evaluate the relative importance of kinetic energy terms, highlighting the dominance of linear production processes. More specifically, the effect of the I_2 term in the vicinity and shortly downstream of the hump was found to dominate the flow. The localised near-wall increase in the I_2 term for all amplitudes appears to be important for understanding the destabilisation of CFI in the vicinity of the hump. The spatial region where I_2 experiences a strong increase ($710 < x < 850$) corresponds largely to the region in which a secondary near-wall peak was identified in the linear perturbation profiles in Fig. 6.7. In the same region, local parallel stability analysis based on the Orr-Sommerfeld (OS) equation predicts a strong spanwise deformation and a double-peak shape for the primary unstable mode near the wall. Furthermore, the OS analysis also identifies a shortly-lived new near-wall mode, emerging purely within the region of crossflow reversal. Such near-wall structures are, in fact, present in all cases in this chapter, and are preconditioned in amplitude by the incoming CFI. Considering the above, a key delineation can be made between *beneficial* and *adverse* effects on the transition scenario. Specifically, for medium-to-high CFI amplitudes, near-wall production dominates local instability growth and leads to the development of highly amplified near-wall structures. Presumably, the strong shears resulting from these structures provide fertile ground for the development of unsteady secondary instabilities, in a manner similar to the development of *type-I* and *type-II* modes in classical saturated CFI (Malik et al., 1999). These in turn can lead to laminar breakdown and accelerated transition observed in the high-amplitude DRE experiments of Rius-Vidales et al., 2025. In contrast, for low to medium incoming CFI amplitudes, near-wall production is present but not strong enough to overtake the overall stabilising effect of the hump on the original incoming CFI. Instead, lasting stabilisation is observed downstream of the hump, which can in turn be linked to the low-amplitude DRE case in the experiments of Rius-Vidales et al., 2025, where transition delay was achieved.

While the above sheds some light onto *adverse* cases where transition is expected to advance, the origin and nature of the stabilising effect of the hump necessitate further discussion. Pertinent effects can again be traced back to the spanwise deformation of the perturbation system due to the hump. The chordwise region of about one hump length downstream of the hump ($800 < x < 1100$) appears to play a key role in the stabilisation effect on CFI. This is a region where the spanwise invariant base flow has largely recovered to the reference (i.e. no hump) case (Fig. 6.3). As a consequence, a purely modal local stability analysis (i.e. based on OS solutions) in this region indicates an almost identical growth rate as the reference case (Fig. 6.7a). In contrast, the predictions of linear HNS in the same chordwise range and on the same base flow show some differences from the modal solution. This can be further traced to the key component of growth in this region, namely linear production and, more specifically, the momentum redistribution term, I_2 (Fig. 6.14). An examination of the constituents of the I_2 term as outlined in equation 6.12 reveals that under minimal changes in the base flow, the key source of modifications of energy growth must lie in the actual shape of the perturbation. A key impact of the hump on incoming CFI is a strong shearing effect in the spanwise di-

rection. This process is energy-neutral and thus not found in an energy budget analysis. Instead, its effect is visible downstream as momentum redistribution is reduced. The origin of this effect can be largely modelled through a pure spanwise momentum transport of the perturbation shape in a purely advection-based simulation of an incoming CFI perturbation profile. As evident in figures 6.11 and 6.7f, the perturbation entering the recovered base flow region ($800 < x < 1100$) is highly distorted due to the aforementioned spanwise shearing effect. In essence, this provides fertile conditions for perturbation growth that is *less efficient* than a classical non-deformed CFI in this region. In other words, the hump effectively deforms the incoming CFI into a locally less unstable structure.

7

CONCLUSION

This above research is now distilled into conclusions regarding the developed frameworks and the discovered flow physics. Firstly, it is reviewed to what extent the research aims have been achieved. Secondly, the impact of the research findings and numerical tools is considered. Lastly, ideas and suggestions for future work are presented. These indicate the limitations of the current work and aid further development of the numerical tools and laminar flow control of crossflow instabilities.

7.1. RESEARCH AIMS

This thesis considered smooth and sharp surface features and several flow stability analysis frameworks to analyse their effects on the transition process. The higher goals here were to improve the performance of swept wings and the flow stability analysis frameworks used to perform the analysis of them. This higher goal was captured by four research aims where the first two relate to the development of numerical tools and the final two relate to the flow physics of boundary layer perturbations interacting with smooth and sharp surface features. Each of the research aims is discussed below.

[RA 1:] *Develop and implement a nonlinear stability framework based on the Harmonic Navier-Stokes equations capable of accounting for smooth and sharp surface features.* (Chapter 3) A nonlinear harmonic Navier-Stokes solver was developed that uses a generalized coordinate transform to treat smooth non-flat surfaces and an embedded boundary method to treat sharp surface features. The solver was subjected to a verification and validation process which demonstrated its correct implementation and accuracy (see Sects. 3.5 and Sects. 3.4).

[RA 2:] *Investigate and implement methods to accelerate computations and reduce the amount of system RAM required for nonlinear simulations.* (Chapter 3). The nonlinear HNS equations consider multiple modes that interact via the nonlinear terms. The left-hand side of the problem contains the spanwise wavenumber and angular frequency of the considered mode, resulting in a mode-dependent left-hand side. The equations were rewritten into a form that removes the mode specifications from the matrices present on

the left-hand side. Consequently, these matrices need only be saved once, and multiplication with the respective wave specifications can be performed quickly (i.e. in sparse formulation) to construct the left-hand side for the considered mode. This results in a significant reduction in memory requirements for solving the nonlinear problem and almost completely removes the scaling of these requirements with the number of modes. In the case of $N = 5$, $M = 0$, this results in 83% less memory required to save the left-hand side. Moreover, during the construction of the left-hand side, measures were taken to reduce the overhead costs (computational and memory) using sparse storage techniques.

[RA 3:] *Identify the effects and influence of the slowly varying flow assumption of the Parabolized Stability Equations in flows over smooth and sharp surface features.* (Chapters 4 and 5) The limitations of the slowly-varying flow assumption in PSE were investigated for smooth and sharp surface features. It was shown that even seemingly superficial surface humps can have a profound effect on the accuracy of results from the PSE framework. No general method for an a priori estimation of the validity of the PSE assumptions was derived. However, it was shown that the HNS framework overcomes all these limitations with a limited increase in computation requirements. This holds even for nonlinear simulations which can still be performed on a desktop computer in a matter of hours.

[RA 4:] *Investigate the effect of smooth surface humps as a method for laminar flow control in Crossflow Instability-dominated swept wing boundary layers.* (Chapter 6) Lastly, the effect of smooth surface features on the linear and nonlinear development of crossflow instabilities in swept-wing boundary layers was investigated. A robust stabilisation (often manifested as a growth reduction rather than decay) was found for incoming perturbations of small amplitude. This stabilisation was scrutinised in detail. An energy budget analysis showed that the stabilisation of the fundamental mode was of linear origin and caused by a reduction of linear production terms. Similarly, the destabilisation of higher harmonics was found to be the result of linear mechanisms as well. Topologically, the stabilisation of the fundamental mode was found to be the result of an additional structure rising near the wall in addition to the incoming CFI which resulted in a reduced effectiveness of trans-shear momentum redistribution as captured by the I2 term.. For higher harmonics, the near-wall structure could reach amplitudes far surpassing that of the incoming mode. Although a change of sign in the I2 term can occur downstream of the hump, in reality, its effect is nullified by the direct effect of nonlinear interactions of the near-wall structure itself. Hence, the limitation to the use of smooth humps as a method for laminar flow control was found to come from large amplitudes of the incoming modes accompanied by large higher harmonics which experience a strong linear destabilisation, possibly leading to an early onset of transition to turbulence. This means that effect laminar flow control can be achieved by preventing large amplitude CFI onset and/or placing the hump near the inception of CFI.

7.2. IMPACT

The impact of the work performed during this research is not just expected but can already be measured by the many publications supported by the developed frameworks. A considerable part of this thesis was also dedicated to the application of the developed tools, which have made an impact on the body of scientific literature and our under-

standing of several flows, e.g. post-DRE flow in Zoppini et al., 2022 and Sequeira, 2022, CFI-FFS interactions in Casacuberta et al., 2022 and Casacuberta et al., 2025, interactions of CFI and smooth surface features (Rius-Vidales et al., 2025; Westerbeek & Kotsonis, 2022; Westerbeek et al., 2026; Westerbeek et al., 2023) and active Flow Control (Mohammadikalakoo et al., 2024). The research described in the aforementioned articles directly influences the ability to design aerodynamic surfaces. Moreover, the research has partially led to a patent (WO 2024/151166 A1) on the use of smooth surface humps for laminar flow control on aerodynamic surfaces. As mentioned before, currently many methods exist for effective LFC. However, many of these either fail outside of lab conditions due to the presence of insect contamination, dust, or scaling issues and/or are too expensive to be economically viable either due to upfront costs or the cost of maintenance. The hump has the potential to overcome all these limitations and provide a practically feasible solution to passive laminar flow control. Moreover, it could be used on new aircraft as well as retrofitted onto existing aircraft. This is crucial given the long lifespan of aircraft, 13-20 years on average depending on the type (Kito, 2021), and the need to reduce emissions as soon as possible.

7.3. FUTURE WORK

The current research has open ends that leave opportunities for future research. The purpose of this section is first and foremost to give insight into the limitations of the current work and how these could be solved with some time and dedication. Furthermore, it can function as a starting point for further researchers who take on this task. This section is divided into a flow control part focusing on passive laminar flow control via smooth surface humps and a numerical part that tackles improvements to DeHNSSO and other codes that could accelerate the numerical treatment of boundary layer stability problems.

7.3.1. LAMINAR FLOW CONTROL

The discovery of effective passive flow control in swept-wing boundary layers through smooth surface protuberances was made in part by this research in close collaboration with colleagues assessing this problem from various angles (experimental and numerical), in particular the works of Casacuberta et al., 2024; Rius-Vidales et al., 2025 and Ambrosino et al., 2025. However, the device has not yet left the controlled environment of simulations and lab tests and even within these environments, the device could not be optimised within the allocated time. This leaves ample opportunity for future research which are considered next.

HUMP SYSTEM OPTIMISATION

The stabilisation of Crossflow Instabilities via humps depends on the shape and size of the hump. Currently, only shallow ($50 < h/b < 100$), symmetric humps were tested. These were placed downstream of the neutral point of the most unstable CFI mode in the reference clean case. In the future, the location and shape of the hump can be optimised and configurations that feature multiple humps successively positioned in the streamwise direction (possibly of varying size and shape) could be designed to improve the maximum achievable amplitude reduction. Given the stabilisation that was found to oc-

cur mostly downstream of the hump, the effect of placing successive humps in the flow needs to be better understood to prevent undermining the main stabilisation mechanism of the hump. Moreover, the size of the hump becomes a more critical parameter in a successive hump configuration. Namely, multiple humps would have to have a streamwise extent limited to where linear perturbation behaviour is expected to occur. This can significantly alter the optimal hump shape. It can reasonably be expected that a small hump (in streamwise extent) might not yield the optimal growth inhibition in itself but can allow for a second hump to be added such that the total configuration exceeds the stabilisation of any one hump.

Asymmetric shapes could aid in avoiding flow separation on the leeward side of the hump for particularly steep humps while possibly maintaining the stabilizing effects. Asymmetry could thus prove valuable for optimisation purposes. This holds especially for configurations with multiple humps as smaller hump widths can be introduced without causing flow separation known to destabilise the flow. Preliminary tests show promise in allowing for greater hump height and preventing flow separation while still stabilizing CFI although these have not been shown in this work due to their preliminary nature. Further investigation is necessary to make a definitive conclusion on the benefits of using asymmetric shapes.

The frameworks developed in this thesis can be used to optimise the hump shape, length, height, and overall configuration (e.g. multiple humps). As shown in this work, a location far downstream of the most downstream hump should be used to accurately measure the effect of the configuration to allow for (de)stabilisation downstream. It is recommended to do this analysis first using linear frameworks. It might be possible to maintain low perturbation amplitudes throughout the domain such that nonlinear effects do not become important within the design space. Linear optimisation is expected to be computationally cheap, although it is currently bottlenecked by the requirement to solve for a new base flow with each configuration. This step is currently more computationally expensive than solving the linear flow stability problem.

Part of the continuation of this research should focus on the scaling and production of the hump. The aerospace sector is a notoriously difficult sector to innovate in due to the high requirements for safety and the large costs associated with testing and manufacturing. If these difficulties are overcome, a significant reduction in fuel consumption and its associated emissions can be achieved. However, it is in the hands of policymakers to ensure this benefits the climate, i.e. by not allowing for more flights as a result of the reduction in emissions per flight.

Aerospace innovations have a track record of long lead times before applications due to the large scales, costly tests, and strict safety regulations. Nevertheless, there is ample opportunity for the hump to find its way onto real wings with the potential to drastically reduce fuel consumption of transonic aircraft.

7.3.2. STABILITY FRAMEWORKS

The stability frameworks developed in the current research have been able to support important research within the Flow Control research group at TU Delft for the past years. However, future research regarding e.g. shape optimisation, secondary instability analysis, or compressible flows invites future improvement and additions.

OPTIMISATION REQUIREMENTS

The future work suggested in Sect. 7.3.1 requires an overarching optimisation setup. To this end, the hump shape should be parametrized such that the design space can reasonably be expected to contain optimal stabilisation. This means that it should include the possibility of asymmetric hump shapes and repeated humps. The optimisation problem appears to fit well the use of adjoint optimisation solvers as shown in e.g. Rigas et al., 2021; Shi et al., 2021 and Pralits and Hanifi, 2003 among others. A dedicated 2-dimensional base flow solver that can efficiently find the laminar base flow for every configuration would greatly streamline the optimisation routine and reduce the manual work currently required to run base flows using separate solvers. Incorporating this into one framework that can find the optimal hump shape and configuration for a given set of flow conditions (Re , inflow amplitude, pressure gradient, sweep angle, etc.) would greatly support future experimental research as well as pave the way to actual implementation.

IMPROVEMENTS TO DeHNSSo

The current implementation can be improved in several ways. The main limitations currently relate to the computational efficiency of Matlab's backslash operator, i.e. the linear system solver algorithm, and the robustness when large amplitudes are considered. Tackling these two factors can severely improve DeHNSSo. Additionally, two additional features could be implemented that relate to the PSE solver, which are considered at the end of this subsection.

The backslash operator is currently employed to solve the problem. During this work, attempts were made to use iterative solvers such as the Generalised Minimum Residual (GMRES), biconjugate gradient (stabilised) (BiCGStab), and Quasi-Minimal Residual (QMR) algorithms. The memory use of all the aforementioned is several orders of magnitude smaller than the direct solver and scales much more favourably with problem size. Other HNS implementations appear to employ a GMRES solution technique, hence it should be possible as the physical system of equations is equivalent. However, it could be that specific numerical choices such as the outflow buffer complicate the application of this algorithm. As such, it has not proven successful in the current implementation. Consequently, BiCG(stab) and QMR techniques were tested as well, but the same numerical choices likely also prevented their fruitful application. No iterative solutions to the linear system could be found. The effective use of these iterative solvers depends on preconditioners, a close initial estimate, and appropriate settings. Failure to provide any of the aforementioned within some degree of accuracy can lead to failure to (effectively) converge to the solution.

The robustness of the nonlinear continuation algorithm could be significantly improved for cases involving (quasi-)saturation behaviour as commonly observed in transitional flows dominated by crossflow instabilities. The current implementation requires an ever-smaller inflow amplitude increase to converge the highly nonlinear equilibrium, which can significantly increase the solution time.

The outflow buffer as it is currently implemented has been effective in preventing unwanted interactions with the outflow boundary. However, the buffer takes up approximately 15% of the domain, effectively increasing the number of grid points for a given

region of interest. Preliminary tests with a parabolic outflow boundary by the current author in DeHNSSo, i.e. a transition to parabolic equations accompanied by a backwards-oriented numerical scheme, proved very effective, requiring only several additional stations ($<1\%$ of the domain) for similar benefits as the buffer region. By slowly introducing a parabolic scheme and neglecting higher-order streamwise derivatives, the behaviour of the flow could be made numerically parabolic. This disables any upstream influence the outflow boundary might have. If the geometry is flat near the outflow for a significant streamwise extent, this approach should be effective. Numerical problems can arise, especially if higher-order streamwise derivatives are approximated with a purely upwind scheme. These terms must therefore be small near the outflow and eliminated from the equations as part of the outflow treatment.

Lastly, the HNS implementation could be coupled with the NPSE code. The HNS and NPSE frameworks complement each other. Where HNS excels when accounting for surface features, body forces, and non-modal initial conditions, and overall robustness, NPSE excels in computational efficiency in slowly-varying flows, and maintains this when simulating strongly nonlinear behaviour of many harmonics. The PSE could be used in series with HNS. This would allow for the HNS to only solve for parts of the domain that feature flow behaviour that breaks the assumptions of NPSE. The benefits are twofold: The inability of the current implementation to solve for fully saturated primary instability behaviour efficiently can be overcome with NPSE and the overall solution time should be several times smaller due to the application of NPSE in a large segment of the domain. The domain will should divided where the flow behaviour is purely convective to avoid requiring implicit coupling between the domains. One limitation to this coupling should be noted, however, that originates from the inflow treatment and nonlinear continuation algorithm within DeHNSSo currently. The inflow presented to DeHNSSo should not feature large amplitudes or lead to amplitudes associated with primary instability saturation within the domain as the amplitude ramping algorithm requires starting from amplitudes so small that nonlinear interactions are negligible everywhere in the domain. In fully linear scenarios, however, this coupling should not be subject to any limitations.

EPILOGUE

I recall applying for a Master's thesis project in the research group of Marios Kotsonis; We had a coffee in the break room at the top floor of the low-speed lab, where it was located then. We discussed the usual topics that were important during an interview for such a commitment, but I recall vividly the atmosphere that Marios had created in the lab and the calmness he radiated already in that moment. I was a little nervous at first, but when I sat there talking to him and my future Master thesis supervisor, and later good friend, Jordi, I felt welcome, respected and excited. Those feelings were correct. I was welcome and respected; and doing research in this wonderful group of people has been incredibly exciting.

The past 5 years, or 6 if you count the master's thesis of which this research was an extension, have gone by incredibly fast. I started the PhD thesis, already dreading the moment I had to write the actual thesis. For it was not the academic world per se that drew me to commit to a PhD, but rather the chance of doing research on such an interesting topic. Having to write a thesis eventually was almost an afterthought. That is, until you have to start writing it or the articles that support it. This became abundantly clear when the time came to start publishing about my research, which required a translation from what seems so clear in my head onto paper: My Achilles' heel. The precision with which one is required to write in academic publications is a skill I have never mastered and have come to respect all the more in others. Especially because some articles are rather difficult to understand, proving the need for well-written articles. To perform research is not what aids to develop the state-of-the art, but publishing about the results does. One cannot push the boundaries of science alone. A well-written article aids the reader from the problem statement to the conclusion and onward with a better understanding of the topic. Although I do still believe that it is only by doing something yourself that you can truly understand it, but here too, a well-written article should help in that process.

Although I will never recommend any potential candidate to start a PhD and consider the thesis an afterthought, cause it really is not, I would do it all again if I had to. I firmly believe that the skills, both hard and soft, learned during a PhD are incredibly powerful in any career that might succeed it, whether it be industry to academia. And more than that, it simply proves that one is able to push through for a very long time, diving into some of the most difficult topics and exploring where no one has previously explored. All while making the scientific method second nature.

That being said, the hump was really just a fortunate accident and not a product of exploration guided by the scientific method. Initially, this research was intended to push the PSE framework to new, more complex, applications: steps. Looking at some of the first DNS results by Jordi, we noticed a smooth dividing streamline between the up- and downstream recirculation regions and the rest of the flow. The idea was thus formed to implement a curvilinear coordinate transformation to allow for stability cal-

culations over a step where the sharp step and the surrounding recirculation regions are excluded from the domain, leaving a smooth surface. Even before being able to do any such calculations, more refined DNS results indicated that such a dividing streamline did not indicate that no interaction with near-wall features in the perturbation field existed. Hence, the thesis objective was changed to look into the Harmonic Navier-Stokes equations capable of tackling the problem of forward-facing steps fully, i.e. accounting for the sharp surface geometry and rapid streamwise flow variations that result from it. However, all this work on coordinate transformations and curvilinear (orthogonal) grid generation should not be for nought. Hence, we looked into smooth surface features as well. A topic not often covered literature, and it was low-hanging fruit given our simulation capabilities. And so, after a detour via sinusoidal surface undulations, the hump was born: Not from a deep understanding of flow dynamics, or a hunch that it could stabilize crossflow instabilities, but because we had all the tools ready to do such a study with very little effort and it was vacant from the literature. And because linear simulations are simply cheaper than nonlinear ones, the first simulations were performed linearly and showed strong stabilisation of the incoming CFI. That just happened to be where the hump was most effective. It seems that research sometimes just requires a strike of luck. A lot of skill to make working code, software, and analysis tools, but a healthy dose of luck sure does help.

ACKNOWLEDGEMENTS

Truly, a doctoral thesis is not a work done solely by the main author. It is a work that combines the effort and ideas of many people both directly (through scientific collaborations) and indirectly (through inspiration, motivation and/or support). I have many people to thank, without whom I am certain this thesis would not have come to be. If you are reading this, you are mostly likely one of those people. So thank you! I want to thank all my friends, family and colleagues for their contributions to this work, but I do want to single out some in particular.

First and foremost, I have to thank my family. My partner Karlijn, my parents Peter and Jacqueline, my sister Elise and her partner Dennis, and of course my grandparents for being an inexhaustible source of support and inspiration. I cannot find the words to describe how important you are to me. I am overwhelmed looking back at the past 5 years (and honestly, all 29 years) and remembering everything you have done for me.

As for the friends I made along the way, I want to thank Jordi, Beto, Giulia, Thomas, Kaisheng, Theo, Marina, Giulio, Marina, Babak, and Parisa. You are all wonderful individuals with bright minds and kind hearts. In particular, Jordi, Giulia and Beto have played a very important role through their direct involvement in the articles that support this thesis. I have thoroughly enjoyed discovering flow phenomena and ways to describe or predict them with you and hope our paths cross professionally in the future once more. And Thomas in particular has been so kind to use his precious lunch break to work out with me in the gym. I have to say that that helped me through many days more than you can imagine.

I also want to mention some people who have honoured me with their long-term friendship. They will very likely not read this. Roeland, Laurens, Bart, Jeroen, Cyrus, Rahul, Tomas, Sid, Boaz, Emmy, Joey, Demi, Angelica, Rik, thank you all for being amazing friends that I am so lucky to have.

Lastly, I have to extend my heartfelt gratitude to my supervisors and co-promoters: Marios, Henk and Steve. Your guidance has been invaluable and has presented me with a solid foundation from which to carry out my own research. I have learned that I work best knowing that there is a backup plan. It gives me the confidence to try, fail, and learn without the fear of coming up empty-handed in the end. Your combined knowledge that I was able to tap into and the hours and hours of reading and correcting my (mostly terribly written) drafts, were vital to this thesis. Thank you.

BIBLIOGRAPHY

- Åkervik, E., Ehrenstein, U., Gallaire, F., & Henningson, D. S. (2008). Global two-dimensional stability measures of the flat plate boundary-layer flow. *European Journal of Mechanics-B/Fluids*, 27(5), 501–513.
- Albensoeder, S., Kuhlmann, H., & Rath, H. (2001). Three-dimensional centrifugal-flow instabilities in the lid-driven-cavity problem. *Physics of fluids*, 13(1), 121–135.
- Ambrosino, B., Mohammad, M., Kotsonis, M., Hanifi, A., & Hein, S. (2025). Stability analysis on the effect of a smooth surface hump on secondary instabilities of cross-flow vortices. *11th European Conference for AeroSpace Sciences (EUCASS) 2025*.
- Andersson, P., Henningson, D., & Hanifi, A. (1998). On a stabilization procedure for the parabolic stability equations. *Journal of Engineering Mathematics*, 33(3), 311–332.
- Arnal, D. (1994). Boundary layer transition: Predictions based on linear theory. *AGARD Report*, 793.
- Arnal, D. (1996). *Control of laminar-turbulent transition for skin friction drag reduction*. Springer.
- Arnal, D., Casalis, G., Cousteix, J., Reneaux, J., Arnal, D., Casalis, G., Cousteix, J., & Reneaux, J. (1997). Laminar-turbulent transition in subsonic boundary layers-research and applications in france. *28th Fluid Dynamics Conference*, 1905.
- Bailly, C., & Comte-Bellot, G. (2015). *Turbulence*. Springer.
- Barahona, M., Rius-Vidales, A. F., Tocci, F., Ziegler, P., Hein, S., & Kotsonis, M. (2022). Tollmien-schlichting waves over forward-facing steps: An experimental and numerical study. *12th International Symposium on Turbulence and Shear Flow Phenomena (TSFP12)*, 1–6.
- Beck, N., Landa, T., Seitz, A., Boermans, L., Liu, Y., & Radespiel, R. (2018). Drag reduction by laminar flow control. *Energies*, 11(1), 252.
- Beguet, S., Perraud, J., Forte, M., & Brazier, J.-P. (2017). Modeling of transverse gaps effects on boundary-layer transition. *Journal of Aircraft*, 54(2), 794–801.
- Bertolotti, F., Herbert, T., & Spalart, P. (1992). Linear and nonlinear stability of the blasius boundary layer. *Journal of Fluid Mechanics*, 242, 441–474.
- Bippes, H. (1997). Environmental conditions and transition prediction in 3-d boundary layers. *AIAA paper*, 1906.
- Bippes, H. (1999). Basic experiments on transition in three-dimensional boundary layers dominated by crossflow instability. *Progress in aerospace sciences*, 35(4), 363–412.
- Boiko, A. V., Dovgal, A. V., Grek, G. R., & Kozlov, V. V. (2011). *Physics of transitional shear flows: Instability and laminar-turbulent transition in incompressible near-wall shear layers* (Vol. 98). Springer Science & Business Media.

- Boiko, A., Dovgal', A., Kozlov, V., & Shcherbakov, V. (1990). Flow instability in the laminar boundary layer separation zone created by a small roughness element. *Fluid Dynamics*, 25(1), 12–17.
- Bottaro, A., Corbett, P., & Luchini, P. (2003). The effect of base flow variation on flow stability. *Journal of Fluid Mechanics*, 476, 293–302.
- Boyd, J. P. (2001). *Chebyshev and fourier spectral methods*. Courier Corporation.
- Breuer, K. S., & Kuraishi, T. (1994). Transient growth in two- and three-dimensional boundary layers. *Physics of Fluids*, 6(6), 1983–1993.
- Casacuberta, J., Groot, K., Ye, Q., & Hickel, S. (2020). Transitional flow dynamics behind a micro-ramp. *Flow, Turbulence and Combustion*, 104(2), 533–552.
- Casacuberta, J., Hickel, S., & Kotsonis, M. (2021). Mechanisms of interaction between stationary crossflow instabilities and forward-facing steps. *AIAA Scitech 2021 Forum*, 0854.
- Casacuberta, J., Hickel, S., & Kotsonis, M. (2024). Passive stabilization of crossflow instabilities by a reverse lift-up effect. *Physical Review Fluids*, 9(4), 043903.
- 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.
- Casacuberta, J., Westerbeek, S., Franco, J., Groot, K., Hickel, S., Hein, S., & Kotsonis, M. (2025). Streaky perturbations in swept-wing flow over forward-facing step. *Physical Review Fluids*, 10(2), 023902.
- Chang, C., Malik, M., Erlebacher, G., & Hussaini, M. (1993). Linear and nonlinear pse for compressible boundary layers.
- Comsol multiphysics® v. 5.6 user guide. (2022). Version Version 5.6. COMSOL AB.
- Cooke, E. E., Mughal, M. S., Sherwin, S., Ashworth, R., & Rolston, S. (2019). Destabilisation of stationary and travelling crossflow disturbances due to steps over a swept wing. *AIAA Aviation 2019 Forum*, 3533.
- Crouch, J., & Kosorygin, V. (2020). Surface step effects on boundary-layer transition dominated by tollmien–schlichting instability. *AIAA Journal*, 58(7), 2943–2950.
- Crouch, J., Kosorygin, V., & Ng, L. (2006). Modeling the effects of steps on boundary-layer transition. *IUTAM Symposium on Laminar-Turbulent Transition*, 37–44.
- Denning, R., Allen, J., & Armstrong, F. (1997). Future large aircraft design—the delta with suction. *The Aeronautical Journal*, 101(1005), 187–198.
- de Rosa, D., Hein, S., Franco, J. A., & Tocci, F. (2024). Stability analysis of laminar separation bubbles on airfoils. *AIAA SciTech 2024 Forum*, 1358.
- 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.
- (DHPC), D. H. P. C. C. (2022). DelftBlue Supercomputer (Phase 1).
- Dobrinisky, A. Y. (2003). *Adjoint analysis for receptivity prediction*. Rice University.
- Dörr, P., & Kloker, M. (2015). Stabilisation of a three-dimensional boundary layer by base-flow manipulation using plasma actuators. *Journal of Physics D: Applied Physics*, 48(28), 285205.
- Downs, R. S., & White, E. B. (2013). Free-stream turbulence and the development of cross-flow disturbances. *Journal of Fluid Mechanics*, 735, 347–380.

- Drake, A., Bender, A., Korntheuer, A., Westphal, R., Rohe, W., Dale, G., McKeon, B., & Geraschchenko, S. (2010). Step excrescence effects for manufacturing tolerances on laminar flow wings. *48th AIAA Aerospace Sciences Meeting Including the New Horizons Forum and Aerospace Exposition*, 375.
- Duchmann, A., Simon, B., Tropea, C., & Grundmann, S. (2014). Dielectric barrier discharge plasma actuators for in-flight transition delay. *AIAA Journal*, 52(2), 358–367.
- Edelmann, C., & Rist, U. (2014). Impact of forward-facing steps on laminar-turbulent transition in subsonic flows. *New Results in Numerical and Experimental Fluid Mechanics IX: Contributions to the 18th STAB/DGLR Symposium, Stuttgart, Germany, 2012*, 155–162.
- Edelmann, C. A., & Rist, U. (2015). Impact of forward-facing steps on laminar-turbulent transition in transonic flows. *Aiaa Journal*, 53(9), 2504–2511.
- Eissler, W., & Bestek, H. (1996). Spatial numerical simulations of linear and weakly non-linear wave instabilities in supersonic boundary layers. *Theoretical and Computational Fluid Dynamics*, 8(3), 219–235.
- Ellingsen, T., & Palm, E. (1975). Stability of linear flow. *The Physics of Fluids*, 18(4), 487–488.
- Eppink, J. L. (2020). Mechanisms of stationary cross-flow instability growth and breakdown induced by forward-facing steps. *Journal of Fluid Mechanics*, 897.
- Eppink, J. L., Wlezien, R. W., King, R. A., & Choudhari, M. (2018). Interaction of a backward-facing step and crossflow instabilities in boundary-layer transition. *AIAA Journal*, 56(2), 497–509.
- Eppink, J. L., & Casper, C. (2019). Effects of forward-facing step shape on stationary crossflow instability growth and breakdown. *AIAA Aviation 2019 Forum*, 3532.
- Fischer, M., WRIGHT, A., JR, & Wagner, R. (1983). A flight test of laminar flow control leading-edge systems. *Aircraft Design, Systems and Technology Meeting*, 2508.
- Forte, M., Perraud, J., Séraudie, A., Beguet, S., & Casalis, L. G. G. (2015). Experimental and numerical study of the effect of gaps on laminar turbulent transition of incompressible boundary layers. *Procedia IUTAM*, 14, 448–458.
- Franco, J. A., Hein, S., & Valero, E. (2020). On the influence of two-dimensional hump roughness on laminar–turbulent transition. *Physics of Fluids*, 32(3), 034102.
- Franco, J., Hein, S., & Valero, E. (2018). Effect of humps and indentations on boundary-layer transition of compressible flows using the ahlns methodology. *7th European Conference on Computational Fluid Dynamics (ECFD 7-ECCOMAS)*.
- Gao, B., Park, D.-H., & Park, S.-O. (2011). Stability analysis of a boundary layer over a hump using parabolized stability equations. *Fluid Dynamics Research*, 43(5), 055503.
- Gregory, N., Walker, W., & Johnson, D. (1956). *The effect on transition of isolated surface excrescences in the boundary layer*. Citeseer.
- Groot, K. (2018). *Biglobal stability of shear flows: Spanwise & streamwise analyses* (Doctoral dissertation). Delft University of Technology. <https://doi.org/10.4233/UUID:60EF07B2-00DB-418B-9495-5A9BAF6105DF>
- Haynes, T. S., & Reed, H. L. (2000). Simulation of swept-wing vortices using nonlinear parabolized stability equations. *Journal of Fluid Mechanics*, 405, 325–349.

- Henke, R. (1999). "a 320 hlf fin" flight tests completed. *Air & Space Europe*, 1(2), 76–79.
- Herbert, T. (1993). Parabolized stability equations, agard r-793.
- Herbert, T. (1988). Secondary instability of boundary layers. *Annual review of fluid mechanics*, 20, 487–526.
- Herbert, T. (1997). Parabolized stability equations. *Annual Review of Fluid Mechanics*, 29(1), 245–283.
- Hetsch, T., & Rist, U. (2009). Accuracy of local and nonlocal linear stability theory in swept separation bubbles. *AIAA journal*, 47(5), 1116–1122.
- Hickel, S., & Adams, N. (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.
- 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).
- Hildebrand, N., Choudhari, M. M., & Paredes, P. (2020). Predicting boundary-layer transition over backward-facing steps via linear stability analysis. *AIAA Journal*, 58(9), 3728–3734.
- Hjort, S. (2004). Incompressible boundary layer instability and transition.
- Holmes, B. J., Obara, C. J., Martin, G. L., & Domack, C. S. (1985). Manufacturing tolerances for natural laminar flow airframe surfaces. *SAE Transactions*, 522–531.
- Holmes, B. J., Obara, C. J., & Yip, L. P. (1984). Natural laminar flow experiments on modern airplane surfaces. *NASA Tech. Pap.*, 2256.
- Hosseinverdi, S., & Fasel, H. F. (2016). Direct numerical simulations of laminar-to-turbulent transition in laminar separation bubbles in three-dimensional boundary-layer. *46th AIAA Fluid Dynamics Conference*, 3793.
- Hu, W., Hickel, S., & Van Oudheusden, B. (2020). Influence of upstream disturbances on the primary and secondary instabilities in a supersonic separated flow over a backward-facing step. *Physics of Fluids*, 32(5).
- Jabbal, M., Liddle, S., & Crowther, W. (2010). Active flow control systems architectures for civil transport aircraft. *Journal of Aircraft*, 47(6), 1966–1981.
- Joslin, R. D. (1992). *Validation of three-dimensional incompressible spatial direct numerical simulation code: A comparison with linear stability and parabolic stability equation theories for boundary-layer transition on a flat plate* (Vol. 3205). National Aeronautics; Space Administration, Office of Management . . .
- Joslin, R. D. (1998). Aircraft laminar flow control. *Annual review of fluid mechanics*, 30(1), 1–29.
- Joslin, R., & Streett, C. (1994). The role of stationary cross-flow vortices in boundary-layer transition on swept wings. *Physics of Fluids*, 6(10), 3442–3453.
- Kendall, J. (1985). Experimental study of disturbances produced in a pre-transitional laminar boundary layer by weak freestream turbulence. *18th Fluid Dynamics and Plasmadynamics and Lasers Conference*, 1695.
- Kito, M. (2021). Impact of aircraft lifetime change on lifecycle co2 emissions and costs in japan. *Ecological Economics*, 188, 107104.
- Klebanoff, P. (1971). Effect of free-stream turbulence on a laminar boundary layer. *Bulletin of the American Physical Society*, 16(11), 1323–+.

- Kloker, M., Konzelmann, U., & Fasel, H. (1993). Outflow boundary conditions for spatial navier-stokes simulations of transition boundary layers. *AIAA journal*, 31(4), 620–628.
- Kohama, Y. (1987). Some expectation on the mechanism of cross-flow instability in a swept wing flow. *Acta mechanica*, 66, 21–38.
- Krochak, O., Casacuberta, J., Hickel, S., & Kotsonis, M. (2022). Mechanisms of interaction between a stationary crossflow instability and backward-facing steps. *12th International Symposium on Turbulence and Shear Flow Phenomena, TSFP 2022*.
- Kurz, H. B., & Kloker, M. J. (2016). Mechanisms of flow tripping by discrete roughness elements in a swept-wing boundary layer. *Journal of Fluid Mechanics*, 796, 158–194.
- Landahl, M. T. (1975). Wave breakdown and turbulence. *SIAM Journal on Applied Mathematics*, 28(4), 735–756.
- Lanzerstorfer, D., & Kuhlmann, H. C. (2012a). Global stability of the two-dimensional flow over a backward-facing step. *Journal of Fluid Mechanics*, 693, 1–27.
- Lanzerstorfer, D., & Kuhlmann, H. C. (2012b). Three-dimensional instability of the flow over a forward-facing step. *Journal of fluid mechanics*, 695, 390–404.
- Li, F., & Malik, M. R. (1996). On the nature of pse approximation. *Theoretical and Computational Fluid Dynamics*, 8(4), 253–273.
- Lin, R.-S., & Malik, M. R. (1997). On the stability of attachment-line boundary layers. part 2. the effect of leading-edge curvature. *Journal of Fluid Mechanics*, 333, 125–137.
- Loiseau, J.-C., Robinet, J.-C., & Leriche, E. (2016). Intermittency and transition to chaos in the cubical lid-driven cavity flow. *Fluid Dynamics Research*, 48(6), 061421.
- Loiseau, J.-C., Robinet, J.-C., Cherubini, S., & Leriche, E. (2014). Investigation of the roughness-induced transition: Global stability analyses and direct numerical simulations. *Journal of Fluid Mechanics*, 760, 175–211.
- Lucas, J. (2014). *Spatial optimal perturbations for transient growth analysis in three-dimensional boundary layers* (Doctoral dissertation). Université de Toulouse.
- Ma, Y., & Zhong, X. (2005). Receptivity of a supersonic boundary layer over a flat plate. part 3. effects of different types of free-stream disturbances. *Journal of Fluid Mechanics*, 532, 63–109.
- Malik, M. R., Li, F., Choudhari, M. M., & Chang, C. (1999). Secondary instability of cross-flow vortices and swept-wing boundary-layer transition. *Journal of Fluid Mechanics*, 399, 85–115.
- Malik, M. R. (1990). Numerical methods for hypersonic boundary layer stability. *Journal of computational physics*, 86(2), 376–413.
- 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.
- Masad, J. A. (1996). *Effect of surface waviness on transition in three-dimensional boundary-layer flow* (tech. rep.).
- Masad, J. A., & Iyer, V. (1994). Transition prediction and control in subsonic flow over a hump. *Physics of Fluids*, 6(1), 313–327.

- Meyer, M., Devesa, A., Hickel, S., Hu, X., & Adams, N. A. (2010). A conservative immersed interface method for large-eddy simulation of incompressible flows. *Journal of Computational Physics*, 229(18), 6300–6317.
- Mohammadikalakoo, B., Kotsonis, M., & Doan, N. A. K. (2024). Control of tollmien–schlichting waves using particle swarm optimization. *Physics of Fluids*, 36(12).
- Morkovin, M. V. (1969). On the many faces of transition. *Viscous Drag Reduction: Proceedings of the Symposium on Viscous Drag Reduction held at the LTV Research Center, Dallas, Texas, September 24 and 25, 1968*, 1–31.
- Morkovin, M. V. (1994). Transition in open flow systems—a reassessment. *Bull. Am. Phys. Soc.*, 39, 1882.
- Nayfeh, A. H., Ragab, S. A., & Al-Maaitah, A. A. (1988). Effect of bulges on the stability of boundary layers. *The Physics of fluids*, 31(4), 796–806.
- Nayfeh, A. (1992). *Influence of two-dimensional imperfections on laminar flow* (tech. rep.). SAE Technical Paper.
- Nayfeh, A. H. (1991). Triple-deck structure. *Computers & fluids*, 20(3), 269–292.
- Örley, F., Pasquariello, V., Hickel, S., & Adams, N. A. (2015). Cut-element based immersed boundary method for moving geometries in compressible liquid flows with cavitation. *Journal of Computational Physics*, 283, 1–22.
- Orr, W. (1907). The stability or instability of the steady motions of a perfect liquid and a viscous liquid. *Proc. Roy. Irish Acad. A*, 27(9), 1–26.
- Park, D., & Park, S. O. (2013). Linear and non-linear stability analysis of incompressible boundary layer over a two-dimensional hump. *Computers & Fluids*, 73, 80–96.
- Perraud, J., Arnal, D., Seraudie, A., & Tran, D. (2004). Laminar-turbulent transition on aerodynamic surfaces with imperfections. *ONERA: Tire a Part*, 13, 2004.
- Perraud, J., & Seraudie, A. (2000). Effects of steps and gaps on 2d and 3d transition. *European Congress on Comp. Methods in Applied Science and Eng., ECCOMAS*, 11–14.
- Picella, F., Loiseau, J.-C., Lusseyran, F., Robinet, J.-C., Cherubini, S., & Pastur, L. (2018). Successive bifurcations in a fully three-dimensional open cavity flow. *Journal of Fluid Mechanics*, 844, 855–877.
- Pralits, J. O., & Hanifi, A. (2003). Optimization of steady suction for disturbance control on infinite swept wings. *Physics of fluids*, 15(9), 2756–2772.
- Prandtl, L. (1905). Über flüssigkeitsbewegung bei sehr kleiner reibung. *Verhandl. 3rd Int. Math. Kongr. Heidelberg (1904), Leipzig*.
- Radeztsky Jr, R. H., Reibert, M. S., & Saric, W. S. (1999). Effect of isolated micron-sized roughness on transition in swept-wing flows. *AIAA journal*, 37(11), 1370–1377.
- Raju, R., Aram, E., Mittal, R., & Cattafesta, L. (2009). Simple models of zero-net mass-flux jets for flow control simulations. *International Journal of Flow Control*, 1(3), 179–197.
- Reed, H., & Saric, W. (2002). Supersonic laminar flow control on swept wings using distributed roughness. *40th AIAA Aerospace Sciences Meeting & Exhibit*, 147.
- Reed, H. L., & Saric, W. S. (1989). Stability of three-dimensional boundary layers. *Annual Review of Fluid Mechanics*, 21(1), 235–284.

- Reibert, M., Saric, W., Carrillo, R., Jr, & Chapman, K. (1996). Experiments in nonlinear saturation of stationary crossflow vortices in a swept-wing boundary layer. *34th Aerospace Sciences Meeting and Exhibit*, 184.
- Reshotko, E. (2001). Transient growth: A factor in bypass transition. *Physics of Fluids*, 13(5), 1067–1075.
- Reshotko, E., & Tumin, A. (2006). Application of transient growth theory to bypass transition. *IUTAM Symposium on One Hundred Years of Boundary Layer Research: Proceedings of the IUTAM Symposium held at DLR-Göttingen, Germany, August 12-14, 2004*, 83–93.
- Reynolds, G. A., & Saric, W. (1986). Experiments on the stability of the flat-plate boundary layer with suction. *AIAA journal*, 24(2), 202–207.
- Reynolds, O. (1883). Xxix. an experimental investigation of the circumstances which determine whether the motion of water shall be direct or sinuous, and of the law of resistance in parallel channels. *Philosophical Transactions of the Royal society of London*, (174), 935–982.
- Rigas, G., Sipp, D., & Colonius, T. (2021). Nonlinear input/output analysis: Application to boundary layer transition. *Journal of Fluid Mechanics*, 911.
- Rius-Vidales, A. F., & Kotsonis, M. (2020). Influence of a forward-facing step surface irregularity on swept wing transition. *AIAA Journal*, 58(12), 5243–5253.
- Rius-Vidales, A. F., & Kotsonis, M. (2021). Impact of a forward-facing step on the development of crossflow instability. *Journal of Fluid Mechanics*, 924, A34.
- Rius-Vidales, A. F., Morais, L., Westerbeek, S., Casacuberta, J., Soyler, M., & Kotsonis, M. (2025). Delay of swept-wing transition using a surface hump. *Journal of Fluid Mechanics*, 1014, A35.
- Saric, W., & Reed, H. (2003). Crossflow instabilities-theory & technology. *41st Aerospace Sciences Meeting and Exhibit*, 771.
- Saric, W. S., Carpenter, A. L., & Reed, H. L. (2008). Flight experiments on swept wing transition using distributed roughness. *ICAS Paper No*, 2(5).
- Saric, W. S., Reed, H. L., & Kerschen, E. J. (2002). Boundary-layer receptivity to freestream disturbances. *Annual review of fluid mechanics*, 34(1), 291–319.
- Saric, W. S., Reed, H. L., & White, E. B. (2003). Stability and transition of three-dimensional boundary layers. *Annual review of fluid mechanics*, 35(1), 413–440.
- Schlichting, H., & Gersten, K. (2016). *Boundary-layer theory*. springer.
- Schmid, P. J., Henningson, D. S., & Jankowski, D. (2002). Stability and transition in shear flows. applied mathematical sciences, vol. 142. *Appl. Mech. Rev.*, 55(3), B57–B59.
- Schrauf, G. (2005). Status and perspectives of laminar flow. *The aeronautical journal*, 109(1102), 639–644.
- Sequeira, A. (2022). Non-modal growth induced by discrete roughness elements in a swept wing boundary layer: Modelling and diagnostics, msc thesis.
- Serpieri, J. (2018). *Cross-flow instability: Flow diagnostics and control of swept wing boundary layers* (Doctoral dissertation). Delft University of Technology.
- Serpieri, J., & Kotsonis, M. (2015). Design of a swept wing wind tunnel model for study of cross-flow instability. *33rd AIAA Applied Aerodynamics Conference*, 2576.
- Serpieri, J., & Kotsonis, M. (2016). Three-dimensional organisation of primary and secondary crossflow instability. *Journal of Fluid Mechanics*, 799, 200–245.

- Shi, Y., Mader, C. A., & Martins, J. R. (2021). Natural laminar flow wing optimization using a discrete adjoint approach. *Structural and Multidisciplinary Optimization*, 64(2), 541–562.
- Sipp, D., & Lebedev, A. (2007). Global stability of base and mean flows: A general approach and its applications to cylinder and open cavity flows. *Journal of Fluid Mechanics*, 593, 333–358.
- Smith, A., & Gamberoni, N. (1956). Transition, pressure gradient, and stability theory. report no. es. 26388, douglas aircraft co. Inc., El Segundo, CA, 1, 24.
- Sommerfeld, A. (1909). Ein beitrag zur hydrodynamischen erklärang der turbulenten flüssigkeitsbewegung. *Proceedings of the 4th International Mathematical Congress, Rome 1908*, 3, 116–124.
- Streett, C. (1998). Direct harmonic linear navier-stokes methods for efficient simulation of wave packets. *36th AIAA Aerospace Sciences Meeting and Exhibit*, 784.
- Streett, C., & Macaraeg, M. (1989). Spectral multi-domain for large-scale fluid dynamic simulations. *Applied Numerical Mathematics*, 6(1-2), 123–139.
- Stüer, H., Gyr, A., & Kinzelbach, W. (1999). Laminar separation on a forward facing step. *European Journal of Mechanics-B/Fluids*, 18(4), 675–692.
- Theofilis, V. (1998). On linear and nonlinear instability of the incompressible swept attachment-line boundary layer. *Journal of Fluid Mechanics*, 355, 193–227.
- Thomas, S. M., Cand Mughal, Gipon, M., Ashworth, R., & Martinez-Cava, A. (2016). Stability of an infinite swept wing boundary layer with surface waviness. *AIAA Journal*, 54(10), 3024–3038.
- Tufts, M. W., Reed, H. L., Crawford, B. K., Duncan Jr, G. T., & Saric, W. S. (2017). Computational investigation of step excrescence sensitivity in a swept-wing boundary layer. *Journal of Aircraft*, 54(2), 602–626.
- Van Ingen, J. (1956). A suggested semi-empirical method for the calculation of the boundary layer transition region. *Technische Hogeschool Delft, Vliegtuigbouwkunde, Rapport VTH-74*.
- Vos, R., & Farokhi, S. (2015). *Introduction to transonic aerodynamics* (Vol. 110). Springer.
- Wang, Y., & Gaster, M. (2005). Effect of surface steps on boundary layer transition. *Experiments in Fluids*, 39, 679–686.
- 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.
- Wassermann, P., & Kloker, M. (2003). Transition mechanisms induced by travelling cross-flow vortices in a three-dimensional boundary layer. *Journal of Fluid Mechanics*, 483, 67–89.
- Wassermann, P., & Kloker, M. (2005). Transition mechanisms in a three-dimensional boundary-layer flow with pressure-gradient changeover. *Journal of Fluid Mechanics*, 530, 265–293.
- Weideman, J. A., & Reddy, S. C. (2000). A matlab differentiation matrix suite. *ACM Transactions on Mathematical Software (TOMS)*, 26(4), 465–519.
- Westerbeek, S. (2020). Development of a nonlinear parabolized stability equation (npse) analysis tool for spanwise invariant boundary layers, msc thesis.

- Westerbeek, S., Hulshoff, S., Schuttelaars, H., & Kotsonis, M. (2024). Dehnsso: The delft harmonic navier-stokes solver for nonlinear stability problems with complex geometric features. *Computer Physics Communications*, 302, 109250.
- Westerbeek, S., & Kotsonis, M. (2022). Effect of streamwise surface undulations on the nonlinear stability of crossflow instabilities. *12th International Symposium on Turbulence and Shear Flow Phenomena, TSFP 2022*.
- Westerbeek, S., Casacuberta, J., & Kotsonis, M. (2026). Linear and nonlinear interactions between stationary cross-flow instabilities and a smooth surface hump. *Journal of Fluid Mechanics*, 1026, A43.
- Westerbeek, S., Franco Sumariva, J. A., Michelis, T., Hein, S., & Kotsonis, M. (2023). Linear and nonlinear stability analysis of a three-dimensional boundary layer over a hump. *AIAA SciTech 2023 Forum*, 0678.
- White, E. B., & Saric, W. S. (2005). Secondary instability of crossflow vortices. *Journal of Fluid Mechanics*, 525, 275–308.
- Wilhelm, D., Härtel, C., & Kleiser, L. (2003). Computational analysis of the two-dimensional–three-dimensional transition in forward-facing step flow. *Journal of Fluid Mechanics*, 489, 1–27.
- Worner, A., Rist, U., & Wagner, S. (2003). Humps/steps influence on stability characteristics of two-dimensional laminar boundary layer. *AIAA journal*, 41(2), 192–197.
- Wu, X., & Hogg, L. W. (2006). Acoustic radiation of tollmien–schlichting waves as they undergo rapid distortion. *Journal of Fluid Mechanics*, 550, 307–347.
- Wu, X. (2019). Nonlinear theories for shear flow instabilities: Physical insights and practical implications. *Annual Review of Fluid Mechanics*, 51(1), 451–485.
- Xia, Y., Jiang, P., Agrawal, G., & Ramnath, R. (2023). End-to-end lu factorization of large matrices on gpus. *Proceedings of the 28th ACM SIGPLAN Annual Symposium on Principles and Practice of Parallel Programming*, 288–300.
- Xu, H., Lombard, J.-E. W., & Sherwin, S. J. (2017). Influence of localised smooth steps on the instability of a boundary layer. *Journal of Fluid Mechanics*, 817, 138–170.
- Xu, H., Sherwin, S. J., Hall, P., & Wu, X. (2016). The behaviour of tollmien–schlichting waves undergoing small-scale localised distortions. *Journal of Fluid Mechanics*, 792, 499–525.
- Yiming, D., Zhenghong, G., Chao, W., et al. (2019). Boundary-layer transition of advanced fighter wings at high-speed cruise conditions. *Chinese Journal of Aeronautics*, 32(4), 799–814.
- Zoppini, G., Michelis, T., Ragni, D., & Kotsonis, M. (2023). The near wake of discrete roughness elements on swept wings. *Journal of Fluid Mechanics*, 960, A11.
- Zoppini, G., Westerbeek, S., Ragni, D., & Kotsonis, M. (2022). Receptivity of crossflow instability to discrete roughness amplitude and location. *Journal of Fluid Mechanics*, 939, A33.
- Zoppini, G. (2023). *Receptivity of swept wing boundary layers to surface roughness: Diagnostics and extension to flow control* (Doctoral dissertation). Delft University of Technology. <https://doi.org/10.4233/uuid:edfbbf98-2530-463b-94d0-43dee5435786>
- Zuniga, F. A., Drake, A., Kennelly, R. A., Koga, D. J., & Von Westphal, R. (1994). *Transonic flight test of a laminar flow leading edge with surface excrescences*. NASA.

A

GENERALIZED CURVILINEAR COORDINATE TRANSFORMATION

A.1. DIRECT DERIVATIVE FORMULATION

Assuming two coordinate systems that are smooth, regular, and time-independent: a Cartesian x, y, z coordinate system and body-fitted numerical ξ, η, z where the third dimension, z , is common to both systems. The ξ and η serve as the numerical grid. Applying the chain rule, the differentials in the Cartesian grid can be written as

$$\partial x = \frac{\partial x}{\partial \xi} \partial \xi + \frac{\partial x}{\partial \eta} \partial \eta, \quad (\text{A.1a})$$

$$\partial y = \frac{\partial y}{\partial \xi} \partial \xi + \frac{\partial y}{\partial \eta} \partial \eta, \quad (\text{A.1b})$$

and the local partial derivatives as

$$\frac{\partial}{\partial x} = \left(\frac{\partial}{\partial \xi} \right) \frac{\partial \xi}{\partial x} + \left(\frac{\partial}{\partial \eta} \right) \frac{\partial \eta}{\partial x}, \quad (\text{A.2a})$$

$$\frac{\partial}{\partial y} = \left(\frac{\partial}{\partial \xi} \right) \frac{\partial \xi}{\partial y} + \left(\frac{\partial}{\partial \eta} \right) \frac{\partial \eta}{\partial y}, \quad (\text{A.2b})$$

Where the derivatives $\frac{dx}{d\xi}$, $\frac{dx}{d\eta}$, $\frac{dy}{d\xi}$, $\frac{dy}{d\eta}$ are known as they can be calculated numerically from the grid definitions. The second-order derivative needs also be defined and can be found by differentiating equation A.2b with respect to y once more. This yields

$$\frac{\partial^2}{\partial y^2} = \frac{\partial}{\partial y} \left[\left(\frac{\partial}{\partial \xi} \right) \frac{\partial \xi}{\partial y} + \left(\frac{\partial}{\partial \eta} \right) \frac{\partial \eta}{\partial y} \right], \quad (\text{A.3})$$

for y and

A

$$\frac{\partial^2}{\partial x^2} = \frac{\partial}{\partial x} \left[\left(\frac{\partial}{\partial \xi} \right) \frac{\partial \xi}{\partial x} + \left(\frac{\partial}{\partial \eta} \right) \frac{\partial \eta}{\partial x} \right], \quad (\text{A.4})$$

for x . This can be rewritten as

$$\frac{\partial^2}{\partial y^2} = \left(\frac{\partial^2}{\partial \xi \partial y} \right) \frac{\partial \xi}{\partial y} + \left(\frac{\partial}{\partial \xi} \right) \frac{\partial^2 \xi}{\partial y^2} + \left(\frac{\partial^2}{\partial \eta \partial y} \right) \frac{\partial \eta}{\partial y} + \left(\frac{\partial}{\partial \eta} \right) \frac{\partial^2 \eta}{\partial y^2}, \quad (\text{A.5})$$

and

$$\frac{\partial^2}{\partial x^2} = \left(\frac{\partial^2}{\partial \xi \partial x} \right) \frac{\partial \xi}{\partial x} + \left(\frac{\partial}{\partial \xi} \right) \frac{\partial^2 \xi}{\partial x^2} + \left(\frac{\partial^2}{\partial \eta \partial x} \right) \frac{\partial \eta}{\partial x} + \left(\frac{\partial}{\partial \eta} \right) \frac{\partial^2 \eta}{\partial x^2}, \quad (\text{A.6})$$

The $\frac{\partial}{\partial y}$ and $\frac{\partial}{\partial x}$ can be replaced by the definitions of equation A.2. A subscript is used to denote the derivative with a double subscript denoting the second-order derivative. This results in

$$\frac{\partial^2}{\partial y^2} = \frac{\partial}{\partial \xi} \left[\left(\frac{\partial}{\partial \xi} \right) \xi_y + \left(\frac{\partial}{\partial \eta} \right) \eta_y \right] \xi_y + \left(\frac{\partial}{\partial \xi} \right) \xi_{yy} + \left(\frac{\partial}{\partial \eta} \right) \left[\left(\frac{\partial}{\partial \xi} \right) \xi_y + \left(\frac{\partial}{\partial \eta} \right) \eta_y \right] \eta_y + \left(\frac{\partial}{\partial \eta} \right) \eta_{yy}, \quad (\text{A.7})$$

and

$$\frac{\partial^2}{\partial x^2} = \frac{\partial}{\partial \xi} \left[\left(\frac{\partial}{\partial \xi} \right) \xi_x + \left(\frac{\partial}{\partial \eta} \right) \eta_x \right] \xi_x + \left(\frac{\partial}{\partial \xi} \right) \xi_{xx} + \left(\frac{\partial}{\partial \eta} \right) \left[\left(\frac{\partial}{\partial \xi} \right) \xi_x + \left(\frac{\partial}{\partial \eta} \right) \eta_x \right] \eta_x + \left(\frac{\partial}{\partial \eta} \right) \eta_{xx}. \quad (\text{A.8})$$

Which can finally be simplified to the definitions of the second-order derivatives used in this document:

$$\frac{\partial^2}{\partial y^2} = \left(\frac{\partial^2}{\partial \xi^2} \right) \xi_y^2 + \left(\frac{\partial^2}{\partial \eta^2} \right) \eta_y^2 + 2\eta_y \xi_y \left(\frac{\partial^2}{\partial \xi \partial \eta} \right) + \left(\frac{\partial}{\partial \xi} \right) \xi_{yy} + \left(\frac{\partial}{\partial \eta} \right) \eta_{yy}, \quad (\text{A.9})$$

and

$$\frac{\partial^2}{\partial x^2} = \left(\frac{\partial^2}{\partial \xi^2} \right) \xi_x^2 + \left(\frac{\partial^2}{\partial \eta^2} \right) \eta_x^2 + 2\eta_x \xi_x \left(\frac{\partial^2}{\partial \xi \partial \eta} \right) + \left(\frac{\partial}{\partial \xi} \right) \xi_{xx} + \left(\frac{\partial}{\partial \eta} \right) \eta_{xx}. \quad (\text{A.10})$$

A.2. JACOBIAN FORMULATION

Equation A.1b are given in the direct form. In the case these forms are not available, the Jacobian formulation must be used. Starting from the matrix form, this is equal to

$$\begin{bmatrix} \frac{\partial}{\partial \xi} \\ \frac{\partial}{\partial \eta} \end{bmatrix} = \begin{bmatrix} \frac{\partial x}{\partial \xi} & \frac{\partial y}{\partial \xi} \\ \frac{\partial x}{\partial \eta} & \frac{\partial y}{\partial \eta} \end{bmatrix} \begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{bmatrix}, \quad (\text{A.11})$$

of which the inverse is

$$\begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{bmatrix} = \begin{bmatrix} \frac{\partial x}{\partial \xi} & \frac{\partial y}{\partial \xi} \\ \frac{\partial x}{\partial \eta} & \frac{\partial y}{\partial \eta} \end{bmatrix}^{-1} \begin{bmatrix} \frac{\partial}{\partial \xi} \\ \frac{\partial}{\partial \eta} \end{bmatrix}. \quad (\text{A.12})$$

Considering the partial derivatives in the body-fitted grid directly, we can write instead

$$\begin{bmatrix} \frac{\partial}{\partial x} \\ \frac{\partial}{\partial y} \end{bmatrix} = \begin{bmatrix} \frac{\partial \xi}{\partial x} & \frac{\partial \eta}{\partial x} \\ \frac{\partial \xi}{\partial y} & \frac{\partial \eta}{\partial y} \end{bmatrix} \begin{bmatrix} \frac{\partial}{\partial \xi} \\ \frac{\partial}{\partial \eta} \end{bmatrix}, \quad (\text{A.13})$$

Combining Eqs. A.12 and A.13, we find:

$$\begin{bmatrix} \frac{\partial \xi}{\partial x} & \frac{\partial \eta}{\partial x} \\ \frac{\partial \xi}{\partial y} & \frac{\partial \eta}{\partial y} \end{bmatrix} = \begin{bmatrix} \frac{\partial x}{\partial \xi} & \frac{\partial y}{\partial \xi} \\ \frac{\partial x}{\partial \eta} & \frac{\partial y}{\partial \eta} \end{bmatrix}^{-1}, \quad (\text{A.14})$$

and we introduce a shorthand notation for the transformation coefficients with a subscript, e.g. $\xi_x = \frac{\partial \xi}{\partial x}$ and calculate the inverse of the right-hand side to give:

$$\begin{bmatrix} \xi_x & \eta_x \\ \xi_y & \eta_y \end{bmatrix} = \frac{1}{x_\xi y_\eta - y_\xi x_\eta} \begin{bmatrix} y_\eta & -y_\xi \\ -x_\eta & x_\xi \end{bmatrix} \quad (\text{A.15})$$

leading to

$$\xi_x = \frac{1}{J} y_\eta, \quad \eta_x = -\frac{1}{J} y_\xi, \quad \xi_y = -\frac{1}{J} x_\eta, \quad \eta_y = \frac{1}{J} x_\xi. \quad (\text{A.16})$$

where $\frac{1}{x_\xi y_\eta - x_\eta y_\xi}$ is the Jacobian of the coordinate transformation. This formulation is necessary since only the coordinates ξ and η form an orthogonal set that allows for partial differentiation. In the equations, e.g. ξ_x is found which is hard to calculate directly. Consequently, the Jacobian is calculated first and the desired correction terms are found.

Then the following formulation can be employed:

$$\frac{\partial}{\partial x} = \frac{1}{J} \left[\left(\frac{\partial}{\partial \xi} \right) y_\eta - \left(\frac{\partial}{\partial \eta} \right) y_\xi \right] = \frac{\partial}{\partial \xi} \xi_x + \frac{\partial}{\partial \eta} \eta_x \quad (\text{A.17a})$$

$$\frac{\partial}{\partial y} = \frac{1}{J} \left[\left(\frac{\partial}{\partial \eta} \right) x_\xi - \left(\frac{\partial}{\partial \xi} \right) x_\eta \right] = \frac{\partial}{\partial \xi} \xi_y + \frac{\partial}{\partial \eta} \eta_y \quad (\text{A.17b})$$

to rewrite the boundary layer and stability equations. However, a second-order derivative can be found with respect to the wall-normal coordinate. In order to find ξ_{yy} and η_{yy} , the operator is applied to ξ and η respectively. This results in

$$\xi_{yy} = \frac{1}{J} \left[\left(\frac{1}{J} \left[\left(\frac{\partial \xi}{\partial \eta} \right) x_\xi - \left(\frac{\partial \xi}{\partial \xi} \right) x_\eta \right] \right)_\eta x_\xi - \left(\frac{1}{J} \left[\left(\frac{\partial \xi}{\partial \eta} \right) x_\xi - \left(\frac{\partial \xi}{\partial \xi} \right) x_\eta \right] \right)_\xi x_\eta \right], \quad (\text{A.18a})$$

$$\eta_{yy} = \frac{1}{J} \left[\left(\frac{1}{J} \left[\left(\frac{\partial \eta}{\partial \eta} \right) x_\xi - \left(\frac{\partial \eta}{\partial \xi} \right) x_\eta \right] \right)_\eta x_\xi - \left(\frac{1}{J} \left[\left(\frac{\partial \eta}{\partial \eta} \right) x_\xi - \left(\frac{\partial \eta}{\partial \xi} \right) x_\eta \right] \right)_\xi x_\eta \right], \quad (\text{A.18b})$$

and

$$\xi_{xx} = \frac{1}{J} \left[\left(\frac{1}{J} \left[\left(\frac{\partial \xi}{\partial \xi} \right) y_\eta - \left(\frac{\partial \xi}{\partial \eta} \right) y_\xi \right] \right)_\xi y_\eta - \left(\frac{1}{J} \left[\left(\frac{\partial \xi}{\partial \xi} \right) y_\eta - \left(\frac{\partial \xi}{\partial \eta} \right) y_\xi \right] \right)_\eta y_\xi \right], \quad (\text{A.19a})$$

$$\eta_{xx} = \frac{1}{J} \left[\left(\frac{1}{J} \left[\left(\frac{\partial \eta}{\partial \xi} \right) y_\eta - \left(\frac{\partial \eta}{\partial \eta} \right) y_\xi \right] \right)_\xi y_\eta - \left(\frac{1}{J} \left[\left(\frac{\partial \eta}{\partial \xi} \right) y_\eta - \left(\frac{\partial \eta}{\partial \eta} \right) y_\xi \right] \right)_\eta y_\xi \right], \quad (\text{A.19b})$$

A

This can be further simplified by seeing that $\frac{\partial \eta}{\partial \eta} = 1$ and likewise for ξ . Similarly $\frac{\partial \xi}{\partial \eta}$ and $\frac{\partial \eta}{\partial \xi}$ are 0. This simplifies the equations of A.18 and A.19 to

$$\xi_{yy} = \frac{1}{J} \left[\left(\frac{-x_\eta}{J} \right)_\eta x_\xi - \left(\frac{-x_\eta}{J} \right)_\xi x_\eta \right], \quad (\text{A.20a})$$

$$\eta_{yy} = \frac{1}{J} \left[\left(\frac{x_\xi}{J} \right)_\eta x_\xi - \left(\frac{x_\xi}{J} \right)_\xi x_\eta \right], \quad (\text{A.20b})$$

and

$$\xi_{xx} = \frac{1}{J} \left[\left(\frac{y_\eta}{J} \right)_\xi y_\eta - \left(\frac{y_\eta}{J} \right)_\eta y_\xi \right], \quad (\text{A.21a})$$

$$\eta_{xx} = \frac{1}{J} \left[\left(\frac{-y_\xi}{J} \right)_\xi y_\eta - \left(\frac{-y_\xi}{J} \right)_\eta y_\xi \right], \quad (\text{A.21b})$$

respectively. All derivative fields are now known and the derivatives can be calculated with respect to the numerical domain ξ, η to calculate ξ_{yy} and η_{yy} respectively.

A.3. NPSE MATRIX CONTENTS

$$\mathbb{L}\hat{\mathbf{q}}_{m,n} + \mathbb{M}\frac{\partial\hat{\mathbf{q}}_{m,n}}{\partial\xi} + \mathbb{N}\frac{\partial\alpha_{m,n}}{\partial\xi}\hat{\mathbf{q}}_{m,n} = \mathbf{r}_{m,n}, \quad (\text{A.22})$$

with

$$\mathbb{L} = \begin{bmatrix} L + \frac{\partial U}{\partial x} & \frac{\partial U}{\partial y} & 0 & i\alpha\xi_x + \eta_x D_1 \\ \frac{\partial V}{\partial x} & L + \frac{\partial V}{\partial y} & 0 & \eta_y D_1 + i\alpha\xi_y \\ \frac{\partial W}{\partial x} & \frac{\partial W}{\partial y} & L & i\beta \\ i\alpha\xi_x + \eta_x D_1 & i\alpha\xi_y + \eta_y D_1 & i\beta & 0 \end{bmatrix},$$

$$\mathbb{M} = \begin{bmatrix} M & 0 & 0 & \xi_x \\ 0 & M & 0 & \xi_y \\ 0 & 0 & M & 0 \\ \xi_x & \xi_y & 0 & 0 \end{bmatrix},$$

$$\mathbb{N} = \begin{bmatrix} -\frac{i}{Re}(\xi_x^2 + \xi_y^2) & 0 & 0 & 0 \\ 0 & -\frac{i}{Re}(\xi_x^2 + \xi_y^2) & 0 & 0 \\ 0 & 0 & -\frac{i}{Re}(\xi_x^2 + \xi_y^2) & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\mathbf{r} = \begin{bmatrix} -\hat{u}\frac{\partial\hat{u}}{\partial\xi}\xi_x - \hat{u}\frac{\partial\hat{u}}{\partial\eta}\eta_x - \hat{v}\frac{\partial\hat{u}}{\partial\xi}\xi_y - \hat{v}\frac{\partial\hat{u}}{\partial\eta}\eta_y - i\beta\hat{u}\hat{w} \\ -\hat{u}\frac{\partial\hat{v}}{\partial\xi}\xi_x - \hat{u}\frac{\partial\hat{v}}{\partial\eta}\eta_x - \hat{v}\frac{\partial\hat{v}}{\partial\xi}\xi_y - \hat{v}\frac{\partial\hat{v}}{\partial\eta}\eta_y - i\beta\hat{v}\hat{w} \\ -\hat{u}\frac{\partial\hat{w}}{\partial\xi}\xi_x - \hat{u}\frac{\partial\hat{w}}{\partial\eta}\eta_x - \hat{v}\frac{\partial\hat{w}}{\partial\xi}\xi_y - \hat{v}\frac{\partial\hat{w}}{\partial\eta}\eta_y - i\beta\hat{w}^2 \\ 0 \end{bmatrix},$$

where

$$\begin{aligned} L &= -i\omega + iU\alpha\xi_x + U\eta_x D_1 + i\alpha V\xi_y + V\eta_y D_1 \\ &+ i\beta W + \frac{1}{Re}(\alpha^2(\xi_x^2 + \xi_y^2) + \beta^2) - \frac{1}{Re}(\eta_x^2 + \eta_y^2)D_2 \\ &- \frac{1}{Re}2i\alpha(\eta_y\xi_y + \eta_x\xi_x)D_1 - \frac{1}{Re}i\alpha(\xi_{yy} + \xi_{xx}) \\ &- \frac{1}{Re}(\eta_{xx} + \eta_{yy})D_1, \end{aligned} \quad (\text{A.23a})$$

A

and

$$M = U\xi_x + V\xi_y - \frac{2i\alpha}{Re}(\xi_x^2 + \xi_y^2) - (\xi_{xx} + \xi_{yy})\frac{1}{Re} - \frac{2(\eta_y\xi_y + \eta_x\xi_x)}{Re}D_1, \quad (\text{A.24a})$$

where D_1 and D_2 are first-order and second-order wall-normal derivative operators respectively.

HNS MATRIX CONTENTS

$$[A]\hat{\mathbf{q}}_{m,n} + [B]\omega\hat{\mathbf{q}}_{m,n} + [C]\beta\hat{\mathbf{q}}_{m,n} + [D]\beta^2\hat{\mathbf{q}}_{m,n} + [E]\frac{\partial\hat{\mathbf{q}}_{m,n}}{\partial\xi} + [F]\frac{\partial^2\hat{\mathbf{q}}_{m,n}}{\partial\xi^2} = \mathbf{r}_{m,n},$$

with

$$A = \begin{bmatrix} A_1 + \frac{\partial U}{\partial x} & \frac{\partial U}{\partial y} & 0 & \eta_x D_1 \\ \frac{\partial V}{\partial x} & A_1 + \frac{\partial V}{\partial y} & 0 & \eta_y D_1 \\ \frac{\partial W}{\partial x} & \frac{\partial W}{\partial y} & A_1 & 0 \\ \eta_x D_1 & \eta_y D_1 & 0 & 0 \end{bmatrix},$$

$$B = \begin{bmatrix} -i & 0 & 0 & 0 \\ 0 & -i & 0 & 0 \\ 0 & 0 & -i & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

$$C = \begin{bmatrix} iW & 0 & 0 & 0 \\ 0 & iW & 0 & 0 \\ 0 & 0 & iW & i \\ 0 & 0 & i & 0 \end{bmatrix},$$

$$D = \begin{bmatrix} \frac{1}{Re} & 0 & 0 & 0 \\ 0 & \frac{1}{Re} & 0 & 0 \\ 0 & 0 & \frac{1}{Re} & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

$$E = \begin{bmatrix} E_1 & 0 & 0 & \xi_x \\ 0 & E_1 & 0 & \xi_y \\ 0 & 0 & E_1 & 0 \\ \xi_x & \xi_y & 0 & 0 \end{bmatrix},$$

$$F = \begin{bmatrix} -\frac{1}{Re}\xi_x^2 - \frac{1}{Re}\xi_y^2 & 0 & 0 & 0 \\ 0 & -\frac{1}{Re}\xi_x^2 - \frac{1}{Re}\xi_y^2 & 0 & 0 \\ 0 & 0 & -\frac{1}{Re}\xi_x^2 - \frac{1}{Re}\xi_y^2 & 0 \\ 0 & 0 & 0 & 0 \end{bmatrix},$$

$$\mathbf{r} = \begin{bmatrix} -\hat{u}\frac{\partial\hat{u}}{\partial\xi}\xi_x - \hat{u}\frac{\partial\hat{u}}{\partial\eta}\eta_x - \hat{v}\frac{\partial\hat{u}}{\partial\xi}\xi_y - \hat{v}\frac{\partial\hat{u}}{\partial\eta}\eta_y - i\beta\hat{u}\hat{w} \\ -\hat{u}\frac{\partial\hat{v}}{\partial\xi}\xi_x - \hat{u}\frac{\partial\hat{v}}{\partial\eta}\eta_x - \hat{v}\frac{\partial\hat{v}}{\partial\xi}\xi_y - \hat{v}\frac{\partial\hat{v}}{\partial\eta}\eta_y - i\beta\hat{v}\hat{w} \\ -\hat{u}\frac{\partial\hat{w}}{\partial\xi}\xi_x - \hat{u}\frac{\partial\hat{w}}{\partial\eta}\eta_x - \hat{v}\frac{\partial\hat{w}}{\partial\xi}\xi_y - \hat{v}\frac{\partial\hat{w}}{\partial\eta}\eta_y - i\beta\hat{w}^2 \\ 0 \end{bmatrix}, \text{ and}$$

$$A_1 = U\eta_x D_1 + V\eta_y D_1 - \frac{1}{Re} \left(\eta_x^2 D_2 + \eta_{xx} D_1 + \eta_y^2 D_2 + \eta_{yy} D_1 \right),$$

$$E_1 = U\xi_x + V\xi_y - \frac{1}{Re}\xi_{xx} - \frac{1}{Re}2\eta_x\xi_x\frac{\partial}{\partial\eta} - \frac{1}{Re}2\eta_y\xi_y\frac{\partial}{\partial\eta} - \frac{1}{Re}\xi_{yy},$$

where D_1 and D_2 are first-order and second-order wall-normal derivative operators, respectively.

A.4. LEFT-HAND SIDE MATRIX

$$\mathbb{M}_1 = \begin{bmatrix} I & 0 & & & \dots & & 0 \\ -\frac{2E}{6\Delta x} + \frac{F}{(\Delta x)^2} & A - \frac{3E}{6\Delta x} - \frac{2F}{6(\Delta x)^2} & \frac{6E}{6\Delta x} + \frac{F}{(\Delta x)^2} & -\frac{E}{6\Delta x} & \ddots & & \\ \frac{E}{12\Delta x} - \frac{F}{12(\Delta x)^2} & -\frac{8E}{12\Delta x} + \frac{16F}{12(\Delta x)^2} & A - \frac{30F}{12(\Delta x)^2} & \frac{8E}{12\Delta x} + \frac{16F}{12(\Delta x)^2} & -\frac{E}{12\Delta x} - \frac{F}{12(\Delta x)^2} & & \\ \ddots & & \ddots & \ddots & \ddots & & \\ & \frac{E}{12\Delta x} - \frac{F}{12(\Delta x)^2} & -\frac{8E}{12\Delta x} + \frac{16F}{12(\Delta x)^2} & A - \frac{30F}{12(\Delta x)^2} & \frac{8E}{12\Delta x} + \frac{16F}{12(\Delta x)^2} & -\frac{E}{12\Delta x} - \frac{F}{12(\Delta x)^2} & \\ & & \frac{E}{6\Delta x} & -\frac{6E}{6\Delta x} + \frac{F}{(\Delta x)^2} & A - \frac{3E}{6\Delta x} - \frac{2F}{(\Delta x)^2} & \frac{2E}{6\Delta x} + \frac{F}{(\Delta x)^2} & \\ & \dots & 0 & \frac{E}{2\Delta x} + \frac{F}{(\Delta x)^2} & -\frac{4E}{2\Delta x} - \frac{2F}{(\Delta x)^2} & A + \frac{3E}{2\Delta x} + \frac{F}{(\Delta x)^2} & \end{bmatrix}$$

$$\mathbb{M}_2 = \begin{bmatrix} 0 & \dots & 0 \\ & B & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & B \end{bmatrix}, \quad \mathbb{M}_3 = \begin{bmatrix} 0 & \dots & 0 \\ & C & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & C \end{bmatrix}, \quad \mathbb{M}_4 = \begin{bmatrix} 0 & \dots & 0 \\ & D & \vdots \\ \vdots & & \ddots & 0 \\ 0 & \dots & 0 & D \end{bmatrix}$$

A.5. PERFORMANCE

DeHNSSo is meant to run simply and fast on regular desktop machines. As such, performance is evaluated here on similar architectures. The performance of the HNS implementation is shown via Tab. A.1 containing the computational cost associated with solving the problems of Sect. 3.5. Both the requirements for the highly refined simulations necessary for validation and medium refinement are shown for the nonlinear simulations. The number of degrees of freedom per mode relates to the grid parameters as $\text{DOF/Mode} = \Gamma_\xi \times \Gamma_\eta \times 4$, where the constant is derived from the 4 state variables of velocity in x , y , and z and pressure.

A server workstation, here referred to as system A, features 2 Intel Xeon Gold 5222 CPUs operating at a clock speed of 3.8 GHz and 128 Gb of DDR4 Memory. However, note that for the cases in this work, only around 26 Gb was needed to solve the system of equations directly.

The interaction of a CFI with a step was assessed using the Delft High-Performance Computing Centre (DHPC) (DHPC), 2022 1.5 Tb memory node due to the high grid refinement required around the step. This node features 2 Intel XEON E5-6248R 24C for a total of 48 cores operating at a clock speed of 3.0GHz equipped with a total of 1.5 Tb of RAM. This system is referred to in Tab. A.1 as system B.

Lastly, the linear calculation of a crossflow instability over a hump was performed on a regular desktop computer. This system, referred to as system C, features an Intel(R) Xeon(R) W-2123 CPU @ 3.60GHz with 4 cores. The system is outfitted with a total of 32 Gb of DDR4 RAM.

Table A.1: Performance and system information for all considered cases. NL = Nonlinear, L = Linear, Ref. = Refined, Sys. = System, Mem = Memory, DOF = degrees of freedom, h = hours.

Case	# Modes	$\Gamma_\xi \times \Gamma_\eta$	DOF/Mode	Sys.	Cores	Sys. Mem.	Time
NL TS	5	800×40	128000	C	8	32 Gb	0.45 h
NL TS Ref.	5	2672×100	1155168	A	8	128 Gb	11.70 h
NL CFI	5	1200×40	192000	C	8	32 Gb	1.29 h
NL CFI Ref.	5	2544×100	1017600	A	8	128 Gb	92.97 h
L CFI Hump	1	2000×100	800000	C	4	32 Gb	0.59 h
L CFI Step	1	4000×350	5600000	B	48	1.5 Tb	5.86 h

CURRICULUM VITÆ

Sven Hendrikus Johannes WESTERBEEK

12-06-1996 Born in Nootdorp, The Netherlands.

EDUCATION

2014-2017 Bachelor of Science in Maritime Engineering
Delft University of Technology

2017-2020 Master of Science in Aerospace Engineering - Aerodynamics, cum laude
Delft University of Technology
Thesis: Development of a Nonlinear Parabolized Stability Equation (NPSE) Analysis Tool for Spanwise Invariant Boundary Layers
Supervisor: Prof. Dr. M. Kotsonis

2020-2026 PhD. Aerodynamics
Delft University of Technology
Thesis: Nonlinear Stability of Boundary Layers over Swept Wings with Sharp and Smooth Surface Features
(Co-)promotors: Prof. Dr. M. Kotsonis,
Prof. Dr. H.M. Schuttelaars,
Dr. S.J. Hulshoff

CAREER

2024-present Scientific software engineer and researcher
Deltares

LIST OF PUBLICATIONS

1. Zoppini, G, **Westerbeek, S**, Ragni, D, Kotsonis, M, "*Receptivity of crossflow instability to discrete roughness amplitude and location*", in *Journal of Fluid Mechanics*, volume 939, page A33, publisher Cambridge University Press.
2. Casacuberta, J, Hickel, S, **Westerbeek, S**, Kotsonis, M, "*Direct numerical simulation of interaction between a stationary crossflow instability and forward-facing steps*", in *Journal of Fluid Mechanics*, volume 943, page A46, publisher Cambridge University Press.
3. **Westerbeek, S**, Kotsonis, M, "*Effect of streamwise surface undulations on the nonlinear stability of crossflow instabilities*", in *12th International Symposium on Turbulence and Shear Flow Phenomena*, TSFP 2022.
4. **Westerbeek, S**, Franco Sumariva, JA, Michelis, T, Hein, S, Kotsonis, M, "*Linear and Nonlinear Stability Analysis of a Three-Dimensional Boundary Layer over a Hump*", in *AIAA SciTech 2023 Forum*, page 0678.
5. **Westerbeek, S.**, Hulshoff, S., Schuttelaars, H., & Kotsonis, M. (2024). "*DeHNSSo: The Delft Harmonic Navier-Stokes Solver for nonlinear stability problems with complex geometric features*", in *Computer Physics Communications*, 302, 109250.
6. Casacuberta, J., **Westerbeek, S.**, Franco, J. A., Groot, K. J., Hickel, S., Hein, S., & Kotsonis, M. (2025). "*Streaky perturbations in swept-wing flow over forward-facing step*", in *Physical Review Fluids*, 10(2), 023902.
7. Rius-Vidales, A. F., Morais, L., **Westerbeek, S.**, Casacuberta, J., Soyler, M., & Kotsonis, M. (2025). "*Delay of swept-wing transition using a surface hump*", *Journal of Fluid Mechanics*, 1014, A35..
8. **Westerbeek, S.**, Casacuberta, J., & Kotsonis, M. (2026). "*Linear and nonlinear interactions between stationary crossflow instabilities and a smooth surface hump*", *Journal of Fluid Mechanics* 1026, A43.

