A Direct-Adjoint Framework for Stability and Sensitivity Analyses of Turbomachinery Aeroacoustics



Anton Glazkov
Christ Church
University of Oxford

A thesis submitted for the degree of
Doctor of Philosophy

Trinity Term
2021

Acknowledgments

I count myself extremely fortunate to have met and had the support and guidance of so many wonderful people during my time at Oxford, without whom this experience would not have been nearly as rich and as rewarding as it has been.

My deepest gratitude is, of course, to my supervisors Li He and Peter Schmid, who have supported and encouraged me continuously ever since our first meeting in June 2016, where the foundations and ideas for this incredibly exciting project were laid. Li, thank you for the invaluable advice and insights into engineering, turbomachinery and CFD that have been indispensable throughout this project, and, most of all, for your trust and the freedom that you gave me in shaping my DPhil research in line with my interests and strengths. Peter, I am indebted to you for introducing me to computational modelling when I was an undergraduate at Imperial, for answering my many questions, and for the privilege of learning so many enlightening ideas from you, especially during our weekly calls our and meetings in London. These have shaped me both as a researcher and as a person, and I am, and will forever be, grateful for everything.

Perhaps the biggest highlights of this project have been my four research visits to the University of Cadiz and my interactions with Miguel Fosas de Pando, and all the folks at UCA, especially Alejandro Quirós Rodríguez. Miguel, working with you has been an absolute privilege. Thank you for always being so open and welcoming, for letting me ask you so many questions, for organising the trips to Cadiz and the access to the supercomputing facilities at CESGA. The most fruitful and engaging moments of my DPhil were spent working with you in Spain! Alejandro, it has been a great pleasure working with you, and I am so happy and excited for you joining Taraneh's group at Sorbonne for your PhD: you will have a great time there!

I am also grateful to my examiners, Richard Sandberg and Marko Bacic, for the wonderful discussion that we had during my viva, and for their incredibly insightful questions and comments.

Financial support for this research came from may sources, and I would like to acknowledge the Alberto del Vicario Scholarship from Christ Church and the funding provided by the Department of Engineering Science at Oxford. Travel to Cadiz was supported by MICINN through the Programa Estatal de I+D+i Orientada a los Retos de la Sociedad Grant No. DPI2016-75777-R, AEI/ERDF,EU. Access to the supercomputer resources was enabled by funding from CESGA/RES and the UK Turbulence Consortium, without whose support this work would not have been possible.

To my Osney labmates James, Hassan, Dan, Marios and Min, thank you for all great chats that we have had over the years at Osney Lab, and to Peter's Imperial group: Souvik, Thanos, Max and Calum, thank you for always being so open and indulging me in my obsession with mathematics, adjoints and block-circulant matrices!

I would also like to thank many of my friends, both from my college family at Christ Church, OURC and from Imperial. Smitha, thank you for always being there to listen and chat with me about everything ranging from mathematics to strange esoterica that I happened to discover earlier that day! Paul, although I never beat you at a single game of squash, or tennis, I'm

grateful that you let me keep trying, and were always up for meeting up to hang out in the GCR! Thank you also to Luc, Sibyl, Megan, Valentina, Ines, Layla, Tim, Matt, Tomasz, Sergio, Gabe and an countless others for making my time at Christ Church and Oxford so special!

None of this, however, would have been possible without you Alina. Although this process has been long, frustrating and, at times, uncertain, your love and support has been a constant during these years that has kept me pushing and striving to get to this point. Thank you for trusting and believing in me; I can feel it every day, no matter where we are in the world.

Finally, I do not know what I would do without my family, and especially my parents and my sister Tanya, who have been a constant pillar of support, not just during this DPhil, but throughout my whole life. I struggle to imagine where I would be if it were not for you.

Abstract

In order to meet new stringent emissions and noise reduction targets, the turbomachinery community is increasingly concerned with addressing inefficiencies in aeroengines, and of particular interest is the issue of sound generated through aeroacoustic processes within the components of the turbomachine. At the same time, precipitous increases in computing power over the last thirty years are driving larger and more detailed simulations. Increasing data output from these requires the development of new techniques that are able to identify the underlying processes and facilitate future efforts for optimisation.

This thesis seeks to address this by creating a direct-adjoint framework that can be applied in two sub-categories of flow problems commonly found in turbomachinery aeroacoustics, and lays the groundwork for future efforts in the efficient analysis and optimisation of these flows. For this, the in-house, high-order, time-domain pseudo-wave-based aeroacoustics code, previously used for isolated airfoil self-noise studies, has been extensively modified and adapted for application in turbomachinery geometries.

Using this code as a starting point, self-excited instabilities and the resulting acoustic fields are simulated within a representative compressor blade row, at off-design conditions. Prior experiments and numerical calculations of the flow at these conditions have confirmed the presence of an unsteady reattaching laminar separation bubble on the suction side of the blade, shedding vortices that are subsequently scattered by the blunt trailing edge. Nonlinear flow analyses are carried out first within a single blade passage, followed by a linear impulse response analysis. A detailed linear direct-adjoint mean flow global stability analysis is performed, and the direct spectrum obtained and correlated with the nonlinear pressure signals, showing excellent agreement. The direct global modes are categorised and corresponding adjoint modes leveraged to compute structural sensitivities with wavemaker regions. The sensitivity analysis confirms the presence of aeroacoustic feedback loops between the pressure and suction sides of the blade, and is consistent with earlier observations in similar studies of isolated airfoils. The theory of block-circulant matrices is then detailed and exploited to achieve computationally efficient decompositions of the governing system matrices, and the previous global stability analysis reproduced with ten passages per period window. An increase in linear growth rates is observed, new low-frequency structures identified and analysed in the context of new feedback loop pathways which propagate disturbances throughout the domain. The effects of aerodynamic mistuning are then briefly addressed.

Rotor-stator interaction is then considered in the context of an adjoint sensitivity problem. For this, a new high-order, time-domain adjoint sliding-plane interface treatment, which is capable of modelling and propagating self-excited instabilities and tonal blade row interaction noise, is developed and extensively validated. This work concludes with the application of this theory to a simplified rotor-stator sensitivity analysis within the context of a time-domain nonlinear-adjoint looping algorithm. Sensitivities identified here are briefly discussed.

Applicability of these methods to different codes and numerical treatments is highlighted throughout this thesis.

The work presented in this thesis is my own unless explicitly stated

Contents

Contents v

List of Figures ix

List of Tables xix

1 Introduction 1
- 1.1 Overview of Turbomachinery Aeroacoustics . . . 1
- 1.2 Theoretical Aeroacoustics and Analytic Techniques . . . 7
- 1.3 Computational Methods . . . 9
- 1.4 Techniques for Analysis . . . 13
- 1.5 Thesis Objectives, Contributions and Outline . . . 21
  - 1.5.1 Thesis Structure . . . 23
  - 1.5.2 Use of Open-Source Software . . . 24
  - 1.5.3 Comments on Figures . . . 25

2 Numerical Methods 27
- 2.1 Introduction . . . 27
- 2.2 Governing Equations . . . 29
  - 2.2.1 The Compressible Navier-Stokes Equations . . . 29
  - 2.2.2 Non-Dimensionalisation . . . 31
- 2.3 Modelling Domains . . . 32
  - 2.3.1 Curvilinear Coordinates and Transformations . . . 33
  - 2.3.2 Discretised Fields and Parallelisation . . . 36
- 2.4 Pseudo-Wave Form . . . 37
- 2.5 Boundary Conditions . . . 39
  - 2.5.1 Inflow Boundary Conditions . . . 40
  - 2.5.2 Outflow Boundary Conditions . . . 41
  - 2.5.3 Solid Wall Boundaries . . . 41
  - 2.5.4 Sponge Layers . . . 42
  - 2.5.5 Periodic Boundaries . . . 43
  - 2.5.6 Imposing Flow Disturbances . . . 43
- 2.6 Spatial Differentiation . . . 44
  - 2.6.1 Finite and Compact Difference Schemes . . . 44

2.6.2 Global and Local Compact Differentiation 47
2.6.3 Block Structure of Nonlinear Code 50
2.7 Efficient Linearisation and Direct and Adjoint Dynamics 50
2.8 Krylov Subspace Methods 53
2.9 Time Integration 54
2.9.1 Runge-Kutta 4th Order 54
2.9.2 Exponential Time Integration 56
2.10 Validation Cases 57
2.10.1 Summary of Validation Cases 57
2.10.2 Co-rotating Vortices 58
2.10.3 Taylor-Green Vortices 60
2.10.4 Linearisation, Direct and Adjoint 63
2.11 Conclusion 65

3 Global Stability Analysis of An Idealised Compressor Blade Row 66
3.1 Introduction 66
3.2 Numerical Methodology 72
3.2.1 Geometry 72
3.2.2 Numerical Simulations 75
3.2.3 Choice of Base Flow 78
3.2.4 Linearisation, Linear and Adjoint Dynamics 79
3.2.4.1 Linear Operators and Governing Equations 79
3.2.4.2 Choice of Inner Product 80
3.2.5 Integration in Time 81
3.2.6 Calculating Global Modes 81
3.2.7 Using the Global Modes 84
3.2.8 Details of Computer Wall Times 86
3.3 Single Passage Dynamics 88
3.3.1 Nonlinear Flow Field 88
3.3.2 Properties of the Mean Flow 93
3.3.3 Impulse Response Analysis 96
3.3.3.1 Initial Transients 99
3.3.3.2 Establishing the Feedback Mechanism 101
3.3.3.3 Asymptotic Behaviour 102
3.3.3.4 Comparison to Prior Work 104
3.3.4 Operator Spectrum 106
3.3.5 Structural, Sensitivity and Receptivity Analysis 108
3.3.5.1 Dominant Modes: Bubble-TE Coupling 109
3.3.5.2 High Frequency Modes: Bubble Dynamics 116
3.3.5.3 Confined Modes 117
3.3.6 Phase and Group Velocity Analysis 121
3.4 Multiple Passages 124
3.4.1 The N-Periodic Geometry 125
3.4.2 N-Periodic Modelling 125
3.4.3 Impulse Response Analysis 131
3.4.4 Global Stability 137
3.4.4.1 Direct Spectra 137

3.4.4.2 Most Unstable Modes . . . 139
3.4.4.3 New Unstable Structures . . . 143
3.4.4.4 Low-Frequency Dynamics . . . 143
3.5 Conclusion . . . 146

4 Rotor-Stator Modelling **151**
4.1 Introduction . . . 151
4.2 Numerical Method . . . 155
4.2.1 Direct Numerical Simulations . . . 155
4.2.1.1 Differentiation in Space . . . 157
4.2.1.2 Time Integration . . . 158
4.2.2 Lagrangian Framework . . . 158
4.2.3 The Sliding Plane . . . 162
4.2.3.1 Interpolation Procedure . . . 163
4.2.3.2 Effects of Interpolation on Modified Wavenumbers . 165
4.2.3.3 Communications for the Sliding Grids . . . 167
4.2.4 Linear and Adjoint Dynamics . . . 170
4.2.4.1 Linearisation, Direct and Adjoint . . . 170
4.3 Validations . . . 172
4.3.1 Stationary Interface . . . 173
4.3.1.1 Acoustic Pulse . . . 175
4.3.1.2 Isentropic Vortex . . . 175
4.3.1.3 Viscous Dissipation . . . 176
4.3.1.4 Heat Dissipation . . . 177
4.3.2 Interface Order of Accuracy . . . 179
4.3.3 Nonlinear Adjoint . . . 181
4.3.3.1 Adjoint Operator . . . 181
4.3.3.2 Checkpointing . . . 182
4.3.4 Parallel efficiency . . . 186
4.4 Applications . . . 188
4.4.1 Gradient Calculations . . . 189
4.4.1.1 Verifying the Adjoint Gradients . . . 190
4.4.1.2 Calculating Gradients for Large Parameter Spaces . 191
4.4.2 Sensitivity Analysis . . . 194
4.5 Conclusion . . . 194

5 Rotor-Stator Interactions **196**
5.1 Introduction . . . 196
5.2 Numerical Setup . . . 196
5.2.1 Domain and Flow Conditions . . . 196
5.2.2 Sensitivity Analysis . . . 197
5.3 Nonlinear-Adjoint Dynamics . . . 200
5.3.1 Overview of the Nonlinear Field . . . 200
5.3.2 Initial Gradients . . . 201
5.3.3 The Adjoint Sensitivities . . . 202
5.3.3.1 Wake-Rotor . . . 202

5.3.3.2 Convective Features and the Laminar Separation Bubble . . . 205
5.4 Notes on the Implementation . . . 208
5.5 Conclusion . . . 209

6 Conclusion **211**
6.1 Summary of Thesis . . . 211
6.2 Suggestions for Future Work . . . 215
6.2.1 Analysis Enabled by the Framework . . . 216
6.2.2 Solver Performance . . . 218
6.2.3 Developing Additional Tools . . . 220

Bibliography **223**

A Mode Clusters for Phase and Group Velocity Analysis **239**
A.1 Introduction . . . 239
A.2 The Plots . . . 239

B Normalisation Regions for Global Modes **242**
B.1 Near-Wake Axial Velocity Normalisation Regions . . . 242

C Sensitivities of Eigenvalues to Mistuning **244**
C.1 Introduction . . . 244
C.2 Formulation . . . 246
C.3 Mean Flow Modifications . . . 247
C.4 Acoustic Spectra . . . 250
C.5 Global Modes and Sensitivities . . . 255
C.6 Perspectives for Mistuning Analysis in Quasi-N-Periodic Systems . . . 262

D Numerical Stencils **264**
D.1 Compact difference schemes . . . 265

List of Figures

2.1 The mapping between the physical $\mathcal{P}$ and computational, $\mathcal{C}$, domains showing the contravariant, $\boldsymbol{\xi}_k$, and covariant, $\boldsymbol{\xi}^k$, bases defined on each space. The author of this thesis is grateful for the kind permission of M. Fosas de Pando for the reproduction of this figure from [Fosas De Pando, 2012]. 36

2.2 The dispersive and dissipative characteristics of the CULD5 and CC3 schemes, showing artificial dissipation at short wavelengths for the upwinding scheme, with no dissipation present in the CC3 scheme. The properties of the boundary closures (denoted by LB*, RB* and B*) do not affect the solution significantly. These are instead tuned to maintain the order of accuracy of the global scheme, and to ensure a numerically stable global solution [Adams and Shariff, 1996]. 46

2.3 Configuration of the processors in the computational space, and the ghost cells around a given processor in the *local compact* difference schemes. The entire pale red region is stored on the processor (i, j). Once the data is assembled in the ghost cells, with communications between the neighbouring processors, the derivatives are applied to the full region. Then the red region is extracted to form the distributed global vector. These schemes allow for high levels of parallel efficiency and significantly simplify the arrangement and storage of the relevant data in memory. . . 49

2.4 The nonlinear right-hand side function $\mathbf{F}$, of Equation (2.42), is decomposed into a calculation of the inviscid, $\mathbf{f}_2$, and the dissipation, $\mathbf{f}_1$, terms. Operations within each block are decoupled in space by considering the derivatives as independent input variables. 50

2.5 The inputs and outputs of the sub-functions (a) $\mathbf{f}_1$ and (b) $\mathbf{f}_2$. The vectors $\mathbf{q}_{i,k}$ are the vectors formed by taking the three derivatives using the central and the two upwinding directions. Vectors of the form $[\mathbf{a}; \mathbf{b}]$ denote a vertical stack of all the elements of $\mathbf{a}$ followed by all the elements of $\mathbf{b}$. 51

2.6 The linearisation procedure used to obtain the linear operator, with $A_{\star,\star,\star}$ matrices corresponding to the linearisation of code sub-blocks, D_k referring to the matrix-free differentiation operation and the C operator denoting the communications for the global/local compact schemes. . . . 52

2.7 The assembly of the adjoint operator simply reverses the order in which the forward linearisation is performed, substituting the corresponding adjoint code, $\star^H$, for each linear operation. Note that one must take into account the grid metric when computing the adjoint. 53
2.8 Temporal stability region of the time integration scheme superimposed with the spectra of the differentiation schemes (modified from [Fosas De Pando, 2012]). The classical Runge-Kutta 4 stability region is also included for comparison. 56
2.9 The grids, (a), used in the validation, (b) comparative snapshots of the analytic and numerical solutions at $t = 600$, with the pressure trace along $x = 0, y > 0$ shown in (c). 61
2.10 Relative pressure error along $y = 0, x > 0$ with varying n_{ghost}, compared to a reference solution with $n_{\text{ghost}} = 64$. Note that the domain size here has been reduced to $n_i = 612$, with the total length of the domain $L_x = L_y = 477.95$. 62
2.11 The kinetic energy E_k and kinetic energy dissipation rate, $\epsilon(\zeta)$, computed from the enstrophy, ζ, in (a). Turbulent decay of the vortex structures at $t = 5, 10, 15$, showing initial roll-ups, the turbulent solution and decay, respectively, in (b). 63
2.12 Quasi-linearisation error for the linear cascade problem of Section 3.2.4. . 64

3.1 The single-passage structured grid showing every tenth grid line in each dimension so that the grid metric may be visualised. Inflow boundary conditions are applied on the left and outflow on the right (blue). Nonslip, adiabatic conditions (red) are imposed on the blade surface, and periodic boundaries are imposed on the upper and lower surface of the upstream and downstream meshes (green). 74
3.2 The surface pressure coefficients $C_p = 2(p-1)/(\gamma M^2)$ at selected Mach and Reynolds numbers. Experimental data available for the pressure coefficient is plotted using solid dots. 76
3.3 The skin-friction coefficients $C_f = 2(\mathbf{t} \cdot \bar{\boldsymbol{\tau}} \cdot \mathbf{n})/\operatorname{Re}$ at selected Mach and Reynolds numbers. Here $\mathbf{t}$ and $\mathbf{n}$ are the tangent and normal vectors to the surface, and $\bar{\boldsymbol{\tau}}$ is the averaged stress tensor. Solid lines indicate the pressure surface, with dashed lines representing the suction side of the blade. Experimental data available for the pressure coefficient is plotted using solid dots. 77
3.4 A snapshot of the (a) nonlinear pressure and (b) dilatation field, with vorticity contours, at $t = 60$. We draw attention to the vortex shedding on the suction side separation bubble and at the resulting interaction at the trailing edge. The scattering at the trailing edge produces an acoustic source that propagates acoustics disturbances into the upstream. 89
3.5 The Fast Fourier Transform (FFT) of the acoustic field upstream of the blade leading edge at $\mathbf{x} = (-0.228, 0.138)$. 91
3.6 The snapshots of the pressure fields corresponding to the main peaks identified in the Fourier transform of the nonlinear pressure field. 92

3.7 The mean flow around the airfoil showing (a) the speed field, $\|\mathbf{u}\|$ superimposed with an arc-length parameterisation of the suction surface, (b) the velocity profiles normal to the blade surface identifying the separation and reattachment points of the laminar separation bubble on the suction surface. Panels (c) and (d) visualise the local velocity field at the separation bubble and the trailing edge. 94
3.8 The mean flow showing the active region of the semi-norm in color. The region is taken throughout the whole height of the mesh and extends from $x = -0.5$ to $x = 1.5$. The inset axis shows the horizontal component of velocity, u, for the initial linear disturbance use for the transient response analysis. 98
3.9 Temporal evolution of the Chu semi-norm for a divergence-free disturbance defined in Equation (3.28). Contributions of the state vector components are also shown. The unstable asymptotic growth is indicated with a dashed line set to a growth rate of $\mathrm{Re}\{\lambda\} = 0.342$. 98
3.10 Initial transient responses in the vicinity of the leading edge. Note the formation of the Tollmien-Schlichting waves packets behind the split vortices propagating parallel to both airfoil surfaces. Observe also the decay of the wave packet on the suction side side and the growth of the waves on the pressure side. 100
3.11 Tilted structures reminiscent of the Orr mechanism highlighted on the suction surface of the airfoil. 100
3.12 Receptivity to acoustic forcing of the upstream boundary layers on the pressure and suction surfaces, leading to the formation of TS waves. . . . 101
3.13 The amplification of disturbances within the pressure and suction side boundary layers for the time interval $t = 0 \rightarrow 10$, as illustrated by the entropy signal at the locations shown in the uppermost panel. Locations on the suction surface are indicated by the 'inverted' triangles and the entropy traces shown in the second graph, while the locations on the pressure surface are shown with 'upright' triangles and the results shown in the lower panel. 103
3.14 Asymptotic evolution of the pressure component of the semi-norm, showing the signal as well as the amplitude envelope function. The average peak spacing is $\Delta t = 1.42$. 104
3.15 The 'beating' phenomenon observed primarily on the suction side as the system enters the final phase of exponential growth. Note the stronger acoustic shedding at $t = 6$ and the amplified hydrodynamic features on the suction surface at $t = 6.5$. 104
3.16 The direct global modes of the linear operator $\mathbf{A}$ with respect to the mean flow, with the inset region below showing a magnified portion of the spectrum for $D_f = \{\lambda : \ \mathrm{Im}\{\lambda\} \in [20, 40]\}$. The eigenvalues are coloured by the L_2 relative error of the eigenvalue/eigenvector pair with respect to the time-stepping operator (see discussion above). The spectrum consists of two unstable eigenvalues with the rest falling in the stable half-plane. The fourteen most unstable modes in D_f are labelled $M_0 \rightarrow M_{13}$. 107

3.17 The real part of the pressure and axial velocity for the modes M_0, M_1, M_6 and M_{12} normalised with respect to the near-wake maximum axial velocity. 109
3.18 Surface pressure and velocity dynamics for the modes M_0, M_1, M_6 and M_{12}. 110
3.19 The direct pressure component of M_2, M_3, the downstream modes. The red circles indicate a modification of the vortex shedding at the trailing edge. 112
3.20 The adjoint horizontal velocity field of M_0, M_1, M_6 and M_{12} in (a), with the surface traces in (b). 113
3.21 The wavemakers for M_0, M_1, M_6 and M_{12}. 114
3.22 Time series of the earlier impulse response analysis projected onto selected direct modes by means of the adjoint operator. 115
3.23 The high-frequency direct modes associated with shear layer instabilities at the separation bubble. 117
3.24 The surface traces of the high-frequency modes for the (a) direct modes, and (b) the wavemakers. 118
3.25 The modes M_{11} and M_{13} show confined features between the pressure and suction surfaces. The direct fields are shown in (a), with adjoint surface traces in (b) and wavemakers in (c). 119
3.26 To the left, the modes selected for the phase and group velocity analysis are shown in colour while the rejected modes are represented as grey points. To the right, the interpolation contours along which the wavenumber is computed. The middle vertical contour is placed at $x_c = -1$ in the upstream and $x_c = 2$ in the downstream, with the neighbouring contours placed either side at $x_\pm = x_c \pm 0.15$, and extend to the full height of the passage. The horizontal contours are similarly equispaced in y, with y_c, the central contour, placed at the midpoint of the central vertical contour. The total length of these is fixed at 0.6 and the midpoint adjusted so that the entire contour fits into the single-passage flow domain. 122
3.27 The dispersion relation for each type of contour. Modes are filtered according to the standard deviation of the instantaneous wavenumber along the contour so that only the monochromatic modes with non-zero k are registered, and the average k computed. Colouring of these modes is by relative standard deviation (RSD) of k. 123
3.28 Showing the clustering of the modes obtained with DBSCAN, the corresponding modes and detailed $k - \omega$ plots of the two main clusters identified in this analysis. 124
3.29 The physical 10-passage geometry considered in this section. The axial lengths of the upstream and downstream sections of this geometry have been chosen so that wavelengths of up to $5h_p$ could be represented within the simulation. The figure has been rotated to save space. 125

3.30 Evaluation of the matrix-vector product $\hat{\mathbf{A}}_j\hat{\mathbf{v}}_j$ using a three-passage linearisation. The three-stage direct approach making use of sequential matrix-vector evaluation in (a) can be simplified in (b) by combining the three stages into one matrix-vector evaluation. Here the required matrix-vector products are confined to the respective passages and may be combined using MPI calls between the corresponding portions of the field. Note also that $\mathbf{A}_1$ and $\mathbf{A}_{N-1}$ are mostly sparse, so additional performance gains may be made by factoring this in, however this is outside the scope of this study. 129
3.31 The coupling fields $\mathbf{A}_1\hat{\mathbf{v}}_j$ and $\mathbf{A}_{N-1}\hat{\mathbf{v}}_j$ are shown for the computational, three-passage mesh. The assembly procedure taking these fields into the middle passage is also shown. 130
3.32 The evolution of the transient response to the divergence-free perturbation. See the supplementary material. 133
3.33 Evolution of the 10-passage norm, showing the initial transient dynamics for $t = 0 \to 5$ and the asymptotic regime for $t = 10 \to \infty$. 136
3.34 The 10-periodic direct modes for the 5 non-symmetric roots of unity. Note the absence of any new unstable structures for ρ_5, which means that the 2-periodic system has the same unstable modes as the single-passage case. 138
3.35 Parameterisation of the unit circle $z = e^{i\theta}$ to the curve eig $\left(\hat{A}(z)\right)$. With our choice of $\mathbf{T}_0$, $\mathbf{T}_1$, $\mathbf{T}_{N-1}$ the eigenvalues traverse separate curves, shown in blue. Each panel shows the mapping of the roots of unity for a given N, and shows how the system with one 'element' ($N = 1$) is stable, but with the addition of each subsequent element ($N > 1$), the most unstable mode alternates between the leftmost and rightmost curves. 139
3.36 The most unstable pressure modes in the (a) single, (b) 5-periodic and (c) 10-periodic cases, showing an adjustment of the upstream wavefronts to a more oblique configuration alongside dominant suction-side dynamics on the surface of the airfoil. 140
3.37 The surface direct and adjoint pressure distributions for the modes (ρ_0, M_0), (ρ_1, M_0) and (ρ_2, M_0) in (a), and the corresponding wavemakers in (b). 142
3.38 Pressure fields for (ρ_2, M_1), (ρ_3, M_0) and (ρ_4, M_1) in (a), with the pressure amplitudes within the passage shown in (b). In (b) note the high pressure amplitudes within the passage in (ρ_3, M_0), and appearance of pressure structures in the downstream and upstream of (ρ_2, M_1) and (ρ_4, M_1), respectively. 144
3.39 Selected low-frequency acoustic modes of the ten-periodic linear cascade. 145
3.40 Plots of the speed fields $|\mathbf{u}|$ for the low-frequency modes. Arrows indicate high-amplitude disturbances in the near wake, while for (ρ_1, M_{371}), the arrows show low-frequency flapping on the suction surface. 146
3.41 The local and global feedback mechanisms identified for N-periodic geometries in this paper. 150

4.1 For any given processor the local data is embedded within a layer of ghost points obtained from neighbouring processors. The compact differentiation stencil is then applied sequentially: first in the vertical and then in the horizontal direction. 157
4.2 The nonlinear-adjoint procedure showing the timestepping for the nonlinear and tangent linear model (TLM) (in blue and orange), and the adjoint iterative procedure that makes use of the checkpointing from the revolve library. The nonlinear solution is run first, with solutions saved at optimal points along the trajectory. The adjoint, which requires the nonlinear field at every timestep, first recovers the missing nonlinear data, then linearises the equations about the checkpoint state and finally advances the adjoint solution backwards in time. 161
4.3 The arrangement of the two grids at the sliding interface showing (a) the ghost layer from Grid 2 overlapping into Grid 1 and (b) the arrangement of ghost and physical points in the overlap region. We assume that the distribution of all points in the overlap region is the same along each vertical axis to simplify interpolation. 163
4.4 For each pair of consecutive donor points defining an interval $\{y\colon y \in [d_0, d_1)\}$ in the physical space, the inevitable map M^{-1} maps this interval to the interval $\{\eta\colon \eta \in [-0.5, 0.5)\}$. The image η_n of a neighbouring point n_0 – at which the value of the interpolate is sought – is computed using Equation (4.17) and the Newton-Raphson procedure by setting $f(\eta_n) = n_0$ and finding the root η_n. 164
4.5 At each step of the time integration procedure the coordinate of the interpolation point $y^{\dagger}_{\text{interp}} = y_0 - \mathbf{v}_{\text{sp}}t$ is computed, from which η_{interp} is calculated using the Newton-Raphson procedure. Once all η_{interp} are calculated, the interpolation matrix terms from Equation (4.18) are assembled in the form of a sparse linear operator, $\mathcal{L}_{d\to i}$, that takes a vector of donor values and outputs the values at $y^{\dagger}_{\text{interp}}$. The interpolation procedure just involves passing the function values and derivatives, with respect to the local coordinates system η (denoted by $\mathbf{D}_k$), as a vector to $\mathcal{L}_{d\to i}$. 165
4.6 The modified wavenumbers for the fourth-order Hermite and linear schemes are shown, at $\eta = 0$, alongside the modified wavenumbers for the interior spatial schemes CULD5 and CC3. Results for the Hermite scheme are shown with $\tilde{\alpha}_{deriv}\Delta\eta$ obtained with numerical differentiation using the CC3 scheme and for exact differentiation (labelled "Exact Hermite"). 166
4.7 Procedure for data management and communications at the sliding interface. Note that this occurs twice per time step, first for the donor points on the left, and then for the donor points on the right. The quantities $\tilde{f}_{\text{interp}}$ denote the interpolated value at the subset of interpolation points that overlap with the boundaries of the current processor, $\tilde{y}_{\text{interp}}$. 168
4.8 The distribution of processors either side of the sliding plane. For example, processor 4 sends and receives chunks of data from processors 1 and 2, and performs reconstruction back into a contiguous array when recovering the local data. 169

4.9 The linearisation procedure used to obtain the linear operator, with $\mathbf{A}_{\star,\star,\star}$ matrices corresponding to the linearisation of code sub-blocks, $\mathbf{D}_k$ referring to the matrix-free differentiation operation and the $\mathbf{C}$ operator denoting the sliding plane communications. 171
4.10 The assembly of the adjoint operator simply reverses the order in which the forward linearisation is performed, substituting the corresponding adjoint code, $\star^H$, for each linear operation. Note that one must take into account the grid metric when computing the adjoint. 171
4.11 Cartesian validation domain with the sliding plane interface conditions located between the red and black meshes. Note that every fourth grid line is shown, for clarity. 173
4.12 The coinciding grid, (a), and the grid displaced vertically by $dy/2$, (b), used in the validation of the communications and interpolation, respectively. 174
4.13 Pressure pulse positioned at $(x_0, y_0) = (0, 20)$ at $t = 0$, is evolved to $t = 15$, with the pressure field for the initial condition (a) and the final pressure field (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grids are shown in the respective figures. 176
4.14 An isentropic vortex positioned at the sliding plane, $(x, y) = (20, 20)$ at $t = 0$, is evolved to $t = 80$, with the pressure field for the initial condition (a) and the final pressure field (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grid are shown in the respective figures. . . 177
4.15 A dissipating vortex positioned at the sliding plane, $(x, y) = (20, 20)$ at $t = 0$, is evolved to $t = 80$, with the pressure field for the initial condition (a) and the final pressure field (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grid are shown in the respective figures. . . 178
4.16 Hot gas 'bump' positioned at $(x, y) = (20, 20)$ at $t = 0$, is evolved to $t = 80$, with the entropy field for the initial condition (a) and the final entropy field at $t = 80$ showing the heat diffusion process (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grid are shown in the respective figures. 179
4.17 Order of accuracy of the sliding plane at successively finer grids (solid blue) with the matching fit (dashed red), showing an order of accuracy consistent with that of the sliding plane interpolation scheme error characteristics, in (a). The timestepper direct-adjoint error with sliding interface in (b). 180
4.18 (a) The Cartesian domain used for the validation of the TLM and non-linear adjoint models. The mesh is uniform, with dimensions 256×128, the inflow on the left boundary of the black grid, outflow on the right of the red grid and periodic boundary conditions on the top and bottom of both the red and black grids. The grid is plotted for every 5th grid line so that the structure of the grid can be visualised. (b) The inlet jet velocity profile $u(y)$ at $x = -62$ (the inlet of the computational domain). 183
4.19 The unstable 1-periodic jet flow through the sliding plane illustrated at $t_t = 100$, with virtual periodic images above and below for a more complete illustration of the flow solution. 184

4.20 Snapshots of the interface region are shown here with the periodic images removed on the left but retained on the right for improved visualisation. The right mesh moves towards the top of the page at $\mathbf{v}_{\text{sp}} = 1.0$. As is clear to see, the wake passes cleanly through the interface. 184
4.21 (a) The initial conditions, $\mathrm{d}J/\mathrm{d}\mathbf{v}\ (\mathbf{v}_N)$, for the nonlinear adjoint time stepping. (b) The sensitivity of the cost function (solution of the nonlinear adjoint system) to initial conditions. The sliding plane is located at x=62. 186
4.22 As with the nonlinear solution, the snapshots of the periodic interface are shown here. Note that the time integration occurs backwards in time and the adjoint wake moves from right to left. 187
4.23 The scaling efficiency of a Cartesian domain case compared for both the single domain and multiple domains with a sliding plane, showing identical scaling performance. 188
4.24 Validation of the adjoint routine for computing the gradient. A single isentropic vortex is introduced above the jet shear layer and left to evolve for $t = 20$ units. Initial gradients of $\mathcal{J}$ are set as the initial condition for the adjoint variable, with the adjoint equation integrated to obtain sensitivities. Multiplication by the gradients of the nonlinear initial condition with respect to vortex height gives the gradient of $\mathcal{J}$. This is compared to a finite-difference approximation of the gradient, with the results showing relative errors on the order of 10^{-3}. 192
4.25 The initial field of three unsteady parallel jets, centred at $y = 7.5$, $y = 15.0$ and $y = 22.5$, with the sliding plane at $x = 62$, and $v_{\text{sp}} = 1$, (a). Forty identical vortices are added to this field in between these jets (see details in text), (b). The perturbation induced by these vortices to the pressure field, at $t = 120.0$ is shown in (c), with the adjoint pressure component shown at $t = 0.0$ shown in (d). The gradients, $\boldsymbol{\nabla}_{\mathbf{p}}\mathcal{J} = \partial\boldsymbol{\eta}/\partial\mathbf{p}\ \boldsymbol{\lambda}_0^{\intercal}$, are then calculated for $p =$ position vectors of vortices, and are shown in (e). 193

5.1 Geometry of the rotor-stator demonstration case. Sponge layers are placed at the inlet and outlet to suppress reflections at the boundaries. . 197
5.2 The initial randomised perturbation, at $t = 0.0$, illustrated for the entropy component of the state vector (although all field components are perturbed). The same field is illustrated at $t = 2.0$, showing the final position of the disturbance in the rotor before the initial cost function gradient is calculated, for the initial conditions of the adjoint solver. . . . 198
5.3 Evolution of the infinity norm of the adjoint field $\mathbf{w}$. 200
5.4 Initial gradients (adjoint variables) used as initial conditions for the adjoint calculation. 202
5.5 The nonlinear and adjoint field time series of the first wake-rotor interaction, showing the pressure pulse from the rotor leading edge and the corresponding adjoint. See the supplementary material. 203
5.6 Vortex sensitivity before impact of the vortex with the rotor leading edge at $t = 1.61$. 204

5.7 Adjoint horizontal velocity field associated with wake features that intersect the leading edge at $t = 1.70$, shown in Figure 5.5. The colourbar limits have been fixed here so that the growth of this sensitivity as it reaches the stator at earlier times may be compared visually. 205
5.8 Convective adjoint structures emanate from the leading edge of the blade located at $y \approx 1.12$ at $t = 1.0$, followed by a pressure pulse at $t = 0.7$ that is not associated with a wake interaction event. These features are labelled. See the supplementary material. 207

A.1 Cluster 1 240
A.2 Cluster 2 241

B.1 The normalisation regions for the single-passage/tuned case in (a), and the mistuned cases considered in Appendix C in (b). 243

C.1 Details of the perturbations to the mistuned geometries compared to the tuned 2-passage cascade. The variation in the y-coordinate of the leading edge in both cases is 5×10^{-3} in both the stagger angle and pitch perturbations. 247
C.2 Modifications to the mean flow as a result of perturbations to the stagger angle (in (a)-(d)) and the pitch (in (e)-(h)). The fields are scaled with the perturbation magnitude $\epsilon = 5 \times 10^{-3}$. Here the perturbed field is interpolated onto the tuned grid to make the comparison independent of changes to the grid metric. 249
C.3 The Fast Fourier Transforms (FFTs) for the pressure field taken upstream of the unperturbed blade. Peaks are labelled in each case to aid the discussion, with one peak being labelled with '*?*' as the features for that peak only appear in the tuned assembly. 250
C.4 The FFT pressure snapshots for the main peaks in the FFTs shown in Figure C.3. Arrows between plots indicate related peaks. The red cross shows the location of the pressure probe used to generate the FFT, and the green cross at the leading edge shows the location at which the phase is adjusted for each plot. 251
C.5 Selected FFT pressure snapshots for selected remaining peaks, as shown in Figure C.3. Note that fields [4], [11] and [12] are new peaks for the pitch-wise-mistuned case (see Figure C.3b), and hence do not have complementary tuned and stagger-angle mistuned fields shown alongside. 252
C.6 Global modes for the tuned and mistuned cases. The modes are labelled according to the corresponding frequency peak in the FFT plot in Figure C.3. Red triangles denote the locations of the peaks in the FFT frequencies in the tuned case. The modes labelled '*?*' correspond to the peak labelled '[*?*]' in Figure C.3. 256
C.7 Pressure snapshots of the leading direct modes for the tuned, mistuned stagger and mistuned pitch cases. 257
C.8 Surface pressure for the tuned and pitch-mistuned direct modes. The modes are labelled between the surface- and pressure-side plots. 259

C.9 Surface pressure and the wavemaker illustrated for each blade of the tuned and mis-staggered cases for the two leading modes. 260

List of Tables

2.1 The errors introduced in the modified wavenumbers as a function of the number of grid points used to represent the wave, n_{ppw}, for the CULD5 and CC3 schemes (replicated from [Fosas De Pando, 2012]). 46

3.1 Single-passage mesh parameters. Grid A refers to the prime grid used for all calculations in this paper, while Grid B and Grid C are used for the purposes of establishing the mesh-independence of the mean flow. Grid B is a local refinement of the blade-surface boundary layer, while grid Grid C is designed as a continuously differentiable subdivision of each grid-point interval in Grid A. 73
3.2 Computer resource utilisation for the different cases shown in this paper. 88
3.3 Eigenvalues of the selected modes. Modes attributable to peaks identified in the FFT are compared to the FFT results in Figure 3.5, and relative errors given. 108

4.1 Cartesian test case parameters 174
4.2 Parameters for TLM and Nonlinear Adjoint Case 183

D.1 Coefficients for the CULD (compact upwind with low dissipation) scheme. 265
D.2 Coefficients for the CULD (compact upwind with low dissipation) scheme (Continued). 265
D.3 Coefficients for the compact 3rd-order centred scheme. 265
D.4 Coefficients for the compact 3rd-order centred scheme (Continued). . . . 266
D.5 Butcher tableau for the Runge–Kutta scheme `RK4(3)5[2R+]C`. 266
D.6 Runge-Kutta coefficients for Table D.5. 266

Chapter 1

Introduction

1.1 Overview of Turbomachinery Aeroacoustics

Increasingly stringent air and noise pollution regulations coming into force at the present moment are presenting novel challenges to the aeronautics industry in the first half of the 21^{st} Century, as it faces the formidable task of reducing carbon dioxide by 75% per passenger kilometre, nitrogen oxide emissions by 90% and perceived noise by 65%, relative to typical new aircraft produced in 2000, by 2050 [European Commission, 2011]. Environmental noise arising from aviation is also a significant and ongoing concern. Increased air traffic and population centres encroaching on airport landing paths are predicted to expose more people than ever to disruptive levels of noise, with its associated detrimental effects [European Commission, 2011]. Mitigation of aircraft noise is of course a multi-faceted problem, with intelligent routing based on given weather conditions, adapted operation schedules at airports and noise reduction techniques for ground operations all playing an important role. Despite these mitigation strategies, the most important noise source remains the aircraft itself, and considerable research effort is presently being directed towards increasing the efficiency of various aspects of the airframe and the associated propulsion systems.

Since the advent of the jet age in the early 1950s, a large contributor to the overall noise signature of aircraft has been the associated gas-turbine power plants that were

quickly adopted for applications in civil aviation, enabling these first-generation jet airliners to fly significantly higher and faster than before. Early designs of these engines, such as the Rolls-Royce Avon, were in the axial turbojet configuration[1] and consist of a single engine core with the air moving axially along the main axis of the engine at the compression, combustion and exhaust stages, with the expulsion of high-temperature gasses at high speeds from the exhaust nozzle providing the means of propulsion. The high specific impulse and high exhaust velocities of these engines make them particularly well suited to supersonic flight but, by the 1970s, concerns for fuel efficiency in subsonic flight, and the jet noise associated with the high exhaust velocities[2], precipitated their effective replacement with turbofans in (subsonic) civil aviation. Greater efficiency has been achieved here by bifurcating the incoming air stream into the engine core and the surrounding bypass duct, which is fronted by a large fan, with ratio of the mass flow rates of air through the bypass duct compared to the engine core termed the *bypass ratio*. The reasons for this can be seen by examining the expression for the propulsive efficiency, which is given by $\eta_p = 2/(1 + u_{\text{exhaust}}/u_{\text{aircraft}}) < 1$, where u_{exhaust} is the velocity of the exhaust gasses from the engine, and u_{aircraft} is the relative velocity of the air surrounding the aircraft. With $\dot{m}$ denoting the mass-flow rate, the thrust can be approximated by $T \approx \dot{m}(u_{\text{exhaust}} - u_{\text{aircraft}}) > 0$, and so it follows that for a set thrust level, if one is to increase propulsive efficiency, then as u_{exhaust} tends to u_{aircraft}, $\dot{m}$ has to increase to compensate. This is achieved for subsonic aircraft by increasing the bypass ratio so that the primary contribution to the thrust comes from a large mass flow rate of air in the bypass duct moving at a low velocity, rather than through the high-velocity exhaust from the core. Modern aero-engines now achieve bypass ratios above $10:1$, with new designs, such as the Rolls Royce UltraFan, extending this to $15:1$ – three

[1] Another engine design is the variant with a centrifugal compressor, and was the basis of some of the first successful designs in the 1940s.

[2] Lighthill's eighth power law states that the far-field acoustic power, W, from turbulent motion scales as the eighth power of the characteristic flow velocity, U, as follows: $W = K\,\rho_0 L^2 U^8 / a_0^5$ [Lighthill, 1952].

times that of engines like the Trent 700 that was put into service in the early 1990s [Rolls Royce, 2021]. This has two major implications for the overall noise signature of the engine.

First, since jet noise arises as a result of multi-scale shear-layer instabilities that occur during the mixing of the hot, high-speed exhaust with the cold ambient air [Tam et al., 2007], the introduction of cold air though the bypass duct increases the mass flow rate through the engine, therefore allowing for the reduction of high exhaust velocities while increasing thrust (see e.g. Wilson and Korakianitis [2014]). On the other hand, increased fan diameters and the reduction in nacelle length[3] are expected to expose the fan noise sources further and lead to increases in the far field noise originating from the fan [Moreau, 2019; Hughes, 2013]. Progress in reducing jet noise, and these concurrent increases in fan noise, means that now fan noise comprises the dominant acoustic source in modern turbofans [Peake and Parry, 2012]. Fan noise, however, is not a single distinct source and may be further decomposed into tonal and broadband components.

Tonal noise refers to noise appearing at distinct frequencies and is characterised by a series of large, sharp peaks at a series of discrete frequencies in the spectrum. Periodic interactions between the fan wake and the outlet guide vanes (OGVs) are the principle source of tonal noise here and are known as the Tyler-Sofrin modes [Tyler and Sofrin, 1962]. Considering V OGVs and B fan blades leads to acoustic modes with orders $m = kV + nB$, $k, n \in \mathbb{Z}$, and the control of these modes involves the selection of blade counts V and B so that the mode is cut-off in the inlet, that is the mode does not propagate out of the engine nacelle. Further tonal noise sources include fan self-noise, which, for a fan rotating at angular frequency Ω results in modes of order $m = nB\Omega$, which is only cut-on in the near-casing position of the inlet duct when the blade tips are supersonic, rotating pressure fields of the rotor

[3]Nacelles are presently employed for locating acoustic liners with the aim of absorbing or redirecting noise in the upstream and ground-oriented directions. As fan size increases, nacelles become larger and are shortened to reduce weight and drag, therefore leaving less space for these liners [Lele and Nichols, 2014].

as perceived by a stationary observer, and the periodic interactions between the fan and macroscopic flow distortions arising due to the nacelle geometry and leading to modes with orders $m = 1$ (see e.g. [Peake and Parry, 2012]).

Flow turbulence also introduces a *broadband* stochastic component to the acoustics. Detailed aeroacoustic experiments by [Ganz et al., 1998] on a model fan rig identified rotor-stator interaction, rotor self-noise and the interaction of the rotor tips with turbulent boundary layers at the casing as the dominant sources of broadband noise. Here rotor-stator interaction noise, where the turbulent wakes from the rotor impinge upon the leading edges of the stators and cause unsteady loading, was determined to be the dominant noise source, followed by rotor self-noise arising from the scattering of turbulent structures at the trailing edge.

In many ways the flows and aeroacoustic processes in axial compressors share the characteristics and features of fan noise, although in aero-engine axial compressor applications, tonal noise dominates due to the blade row interactions within the multi-stage geometries. A characteristic common to both compressors, fans and turbines, is that these flows introduce adverse pressure gradients. This, coupled with the requirement to operate at off-design conditions, can lead to localised separation regions, followed by a downstream reattachment that has been suggested to be the cause of an additional dipole acoustic source concentrated on the blade suction surface [Jones et al., 2010]. A number of recent studies have considered this problem using flows over isolated NACA-0012 (e.g. [Desquesnes et al., 2007; Jones and Sandberg, 2011]) and controlled-diffusion airfoils at off-design conditions (e.g. [Sanjose et al., 2019]). This body of work has shown that both the separation bubble dynamics, the Kelvin-Helmholtz instabilities and the scattering of the Kelvin-Helmholtz vortices by the blunt trailing edges are sources of acoustic radiation which is characterised by broadband spectra with sharp peaks [Sanjose et al., 2019]. Although it should be noted that the case considered in [Sanjose et al., 2019] is performed at a relatively low Reynolds number of $150,000$, and separation region

is very close to the trailing edge, meaning that turbulent breakdown of coherent structures is minimised before scattering by the trailing edge.

Recent literature also reported rotating instabilities [Mailach et al., 2001] and nonharmonic acoustic resonances having been observed in a range of laboratory and industrial settings [Hellmich and Seume, 2008], which, in some cases have been implicated in fatigue failures of the blades (see e.g. [Rizk and Seymour, 1964]). In their literature review, Hellmich and Seume characterise these as sharp peaks in the acoustic spectrum occurring at frequencies that are generally not integer multiples of the shaft rotation frequency, seemingly triggered by self-excited flow instabilities and vortex shedding from the vanes and struts in the machine [Hellmich and Seume, 2008]. Initial work on this was performed by Parker [Parker, 1967, 1984; Parker and Stoneman, 1985, 1989] on flat-plate cascades within low-speed wind tunnels, and in the process identified trapped *Parker modes* that confine the acoustic energy between the blade pitches. Transmission of acoustic waves through blade rows (see e.g. [Koch, 1971; Glegg, 1999]) has also been considered, although studies looking at the interactions between the hydrodynamic generation process and the acoustic propagation through the blade row remain scarce.

A by-product of modern regulatory requirements, engine design trends, and the reduction of the dominant noise sources, is the increasing importance of other previously insignificant aeroacoustic processes, such as turbine and combustion noise [Nesbitt, 2010; Moreau, 2019]. Much like the fan and compressor, turbine noise consists of both tonal and broadband components[4], higher loading of the turbine is expected to increase the importance of both the acoustic sources due to the stronger wakes generated by the blades, with the reduced solidity of the turbine minimising attenuation, and enabling the propagation of the turbine noise into the environment. Coupling this with the lower shaft rotation speeds in current high-bypass designs, the generated tones are reduced to more uncomfortable frequencies.

[4]Which includes scattering within the turbulent shear layers of the exhaust jet as well as internally generated noise.

Recent research interest has also been directed towards combustion noise, which consists of two principal components (see for example the reviews by Dowling and Mahmoudi [2015] and Morgans and Duran [2016]). The first of these, termed *direct noise*, is related to the volume expansion associated with the unsteady heat release rates within the combustion process, and leads to a monopole-type acoustic source. The second, called *indirect noise*, arises when the inhomogeneous flow leaving the combustion chamber is accelerated as it passes through the nozzle guide vanes (NGVs) and turbine. Changes in the volume of this non-homogeneous fluid as it is accelerated creates a secondary acoustic source characterised by a dipole structure [Marble and Candel, 1977], with the component associated with convecting temperature 'hot-spots' termed *entropy noise*[5]. This is thought to dominate the indirect combustion noise component [Morgans and Duran, 2016], though the effects of non-uniform gas composition [Ihme, 2017; Magri et al., 2016] and vorticity [Cumpsty, 1979] have also been considered. Combustion noise has seen recently renewed interest for two reasons. The first, much like in the discussion above, is its increasing relative importance as a component of the overall aero-engine noise and the need to address all noise components concurrently in order to meet the more stringent emissions regulations [Morgans and Duran, 2016]. Furthermore, addressing the emission of atmospheric pollutants, especially nitrogen oxides, necessitates the introduction of lean premixed combustion that is known to be susceptible to thermoacoustic instabilities (see references in Morgans and Duran [2016]). Coupling the flame dynamics in the combustion chamber to the re-transmitted entropy noise from the duct and nozzle has been shown to introduce additional low-frequency modes, arising from the coupling of slow convective dynamics of the entropy waves and the fast acoustic propagation timescales, with implications for combustor stability [Hubbard and Dowling, 1998; Goh and Morgans, 2013; Motheau et al., 2014; Duran et al., 2014].

[5]Entropy noise formation in ducts and NGVs was side project of this work, and was the subject of our presentation at the 15$^{\text{th}}$ International Symposium on Unsteady Aerodynamics Aeroacoustics and Aeroelasticity of Turbomachines. The contribution of this project is not included but mentioned here for completeness.

Challenges related to moderate and high-frequency temperature measurements in experimental settings remain, however, and as a result experimental studies, such as the Entropy Wave Generator (EWG) experiment [Bake et al., 2009] at the DLR, have been restricted to simplified geometries.

1.2 Theoretical Aeroacoustics and Analytic Techniques

The beginnings of aeroacoustics as a scientific discipline may be traced to the works of Lighthill in the 1950s, who, in his two-part publication [Lighthill, 1952, 1954], derived an exact[6] inhomogeneous wave equation for quiescent fluids, with the source term consisting of derivatives of the Lighthill tensor. In the context of high-speed turbulent jets, this was simplified to considering quadrupole sources, which were subsequently solved using Green's functions, and the theory extended to account for the presence of solid boundaries [Curle, 1955]. Further extensions of Lighthill's analogy by Ffowcs Williams and Hawkings for sound generated by surfaces in motion [Ffowcs Williams and Hawkings, 1969a] have seen wide adoption for predicting noise in a range of contexts, including in rotating turbomachinery [Ffowcs Williams and Hawkings, 1969b], with the key to its success being its wide applicability in both numerical and computational contexts.

Many other successful tools for analytically predicting aeroacoustically generated noise have been developed since. For example, Amiet's theory [Amiet, 1975, 1976] for modelling leading and trailing edge noise over isolated airfoils assumes a 'frozen' turbulent forcing being convected past the respective ends of the airfoil. The magnitude of this forcing is assumed to be small in comparison with the mean flow, the interaction process is inviscid, so that the linearised Euler equations can be used, and a flat-plate geometry is assumed, although this has subsequently been extended

[6] That is, it is derived without simplifications directly from the compressible mass and momentum equations of gas flow.

to account for vortex shedding and airfoil thickness effects [Roger and Moreau, 2010]. The wave equation is then solved for the surface pressure and, assuming a dipole source structure, this is applied to Lighthill's analogy to propagate the acoustics in the far-field.

Similarly, for problems modelling the propagation of acoustic and entropy disturbances through blade rows, actuator disk theory has been developed for the modelling of indirect combustion noise [Cumpsty and Marble, 1977]. These work by assuming that the wavelength of the incoming disturbance is large relative to the chord of the blades making up the cascade, and so the blade row may be treated as 'compact' relative to the incoming wave, while, from the perspective of the wave, the blade rows are essentially jump conditions that should be satisfied. These approaches have seen significant success in early work, but again fail to account for local instability processes and, when the wavelengths of the incoming disturbance are comparable to or smaller than the size of the blade, this theory fails to take into account the effects of the blade geometries.

To summarise, the modelling of aeroacoustic processes in turbomachinery has been successfully achieved in a considerable body of work that has applied these aforementioned techniques to a large array of problems encountered in the field. These are useful because they distil fundamental physics into simple models that are fast and computationally efficient to apply at the design stage, and in control theory applications where fast turn-around times are required. Limitations arise here, however, because of the simplifications introduced in the modelling phase, with simplified geometries and flows substituted for real-world flow physics. In particular, prediction for realistic geometries remains challenging, and accounting for local and more complex effects, such as self-excited instabilities within unstable boundary layers, is rarely modelled, and treating the acoustic generation and propagation problems independently misses fundamental interactions between the hydrodynamics and the acoustics.

1.3 Computational Methods

Over the last thirty years, exponential increases in computing power have enabled researchers to tackle a much greater breadth and depth of fluid dynamics problems that were previously considered intractable. This is especially true for the field of aeroacoustics, which, in recent decades, has evolved from a sub-discipline of fluid dynamics, largely dominated by problems in simple geometries and solved with analytic methods, to a subject in its own right with simulations directly resolving the intricate noise-generating flow processes in complex geometries. New methods and specific computational requirements for *computational aeroacoustics* (CAA) have also been developed and contrast greatly with those for aerodynamic simulations of conventional *computational fluid dynamics* (CFD) codes.

Indeed, the successful prediction of radiated acoustic fields requires the numerical method to resolve the small-scale *generation* phenomena, such as shear layer instabilities within boundary layers, while at the same time *propagating* the acoustics fields, that are several orders of magnitude smaller in amplitude than the macroscopic pressure gradients. This introduces the requirements of high-order methods and especially careful numerical treatments at the boundaries for minimal dissipation and reduced numerical errors, which may otherwise contaminate the solution [Lele, 1997; Lele and Nichols, 2014].

This requirement of modelling phenomena at multiple scales presents considerable challenges to the computer simulation, and has been addressed through two broad strategies. The first and perhaps still the most common category of methods is the hybrid approach. Here the unsteady hydrodynamic problem is decoupled from the propagation, with the former being solved numerically, while the far-field noise is calculated using the FWH formulation. This makes these methods attractive because the complexity of the flow is confined to a relatively compact region of the domain, without the need to use extra grid points for simulating the acoustic propagation. The choice of solver used for modelling the unsteady hydrodynamics

depends largely on the computational cost that can be afforded for resolving the turbulent fluctuations for a given application, and the intended use case. While wall-resolved Large-Eddy Simulations (LES) presents the least amount of turbulence modelling and has considerably reduced grid-size requirements compared to DNS [Choi and Moin, 2012], at application Reynolds numbers this too becomes restrictive [Lele and Nichols, 2014]. For this reason, *hybrid* RANS-LES schemes have been adopted to relax this restriction by modelling near-wall areas, where grid sizes are too large to resolve the turbulent eddies, using a RANS model, and then switching to an LES solution when the grid is sufficiently fine. This is much more attractive in industrial applications, since the coarser grids allow for lower computational costs and enable multiple iterations during the design process. As a result, this remains a dynamic area of research (see the recent review in [Heinz, 2020] for an overview). Although beneficial for approximating the acoustic radiation in a design scenario, by modelling the processes in the boundary layer that are critical to the receptivity and transition processes, these hybrid approaches are unable to offer insights into the fundamental instability mechanisms. If insights into the dynamics of these systems are sought, rather than a prediction, lowering the Reynolds number to allow for wall-resolved LES or even DNS simulations is desirable, thus opening up opportunities for the investigation of fundamental processes.

By contrast, direct modelling of the full compressible Navier-Stokes equations, also known as Direct Noise Computation (DNC) or Comprehensive CAA [Lele and Nichols, 2014], computes and resolves both the unsteady flow and the acoustic propagation of the resulting sound waves using one, unified computation. Prediction of the radiated sound involves accurately resolving the flow at the scale of the turbulent eddies, and introduces strict grid-size requirements around obstacles and within detached shear layers. With Reynolds numbers in realistic turbomachinery applications easily reaching $\mathrm{Re} \approx 10^6$, and the associated grid size requirements scaling as $N \approx \mathcal{O}\left(\mathrm{Re}^{\frac{37}{14}}\right)$ – compared to wall-resolved LES which scales as $N \approx \mathcal{O}\left(\mathrm{Re}^{\frac{13}{7}}\right)$ –

[Choi and Moin, 2012] many successful applications of DNC have been performed at much reduced Reynolds numbers and for significantly simplified geometries. The benefit to the researcher, however, is that this approach provides detailed insights not only of the generation mechanisms, but also the increased understanding of the interactions between the acoustics and the hydrodynamic flow instabilities generating the former, which cannot always be achieved with hybrid methods.

Unlike other aerodynamics applications, we have seen in the preceding review that turbomachinery requires unique geometries, is characterised by flows through large-scale periodic blade row structures, and some form of coupling between the blade rows when they are assembled in compressor or turbine stages. The difficulty here is that stage designs usually have a large number of rotor and stator blades, with the numbers in each stage generally selected to be large co-prime pairs so as to avoid destructive tonal noise harmonics. Because of this, capturing the full temporal evolution of the flow for these configurations requires the solution of the full *three-dimensional* annulus, which, in general, is a formidable computational task, and so alternative methods have been developed for different situations to address this.

The mixing-plane [Denton, 1992] is perhaps the most commonly used method for the *aerodynamic* design of turbomachinery stages. This steady flow method performs an azimuthal averaging of the flow between the rotor and stator, and thus only requires the flow to be solved in one rotor and one stator passage despite the potentially disparate pitchwise periodicities that are encountered due to the differing blade counts in the full annulus. While this makes this method fast for design purposes, important unsteady information is lost due to the averaging, making the mixing plane unsuitable for aeroacoustic applications.

Harmonic balance [He and Ning, 1998; Hall et al., 2002] and chorochronic [Erdos et al., 1977; He, 1990, 1992; Gerolymos et al., 2002] methods were developed to address these deficiencies and enabled unsteady nonlinear time-periodic flow problems

to be considered. The former assumes that the unsteady flow is defined by a finite set of *predetermined frequencies*, usually the blade passing frequency in acoustic applications, that are superposed and the Fourier coefficients solved for, making it a spectral method. Periodicity is accounted for by applying appropriate phase shifts at the boundary conditions, thus allowing the periodic domain to be modelled in a single passage. Chorochronic methods instead exploit the fact that in a rotor-stator interaction setting, the *macroscopic* flow observed in any two passages in either blade row are temporally and spatially phase shifted relative to each other [Gerolymos et al., 2002], and so pairs of passages experience the same flow at different times. This periodicity is accounted for by phase-lagging the solution imposed at the boundaries of the single passage. This can be implemented by either saving a well-resolved time series [Erdos et al., 1977], which is then applied as the boundary condition at some later time, or by decomposing the solution at the boundaries using either Fourier series or POD [Mouret et al., 2016] to compress the data.

While it is well-known that accurate aeroacoustic predictions of turbomachinery blade rows depend on having the correct blade-passing frequencies of the interactions between the rotors and stators, if the goal is to gain a better understanding of transient phenomena and the effects of turbulence, seeking a time-domain solution is preferred. This because no assumptions about the characteristics of the flow spectrum are required. To reduce the cost of computing full-annulus solutions, the pitch[7] can be altered so that the interaction is periodic over smaller sectors of the annulus. For this the rotor-stator interface, which communicates between the rotor and stator domains in their respective frames of reference, must propagate flow disturbances with minimal dissipation so that the artificial effects of the domain interface do not contaminate the hydrodynamic and acoustic fields. Such methods are common and applications include helicopter rotor interaction [Steijl and Barakos, 2008], and turbomachinery [Bhaskaran and Lele, 2010; Johnstone et al., 2015], to

[7]Equivalently, the number of blades in each blade row.

name a few. It should be noted here, however, that sliding plane approaches are commonly criticised for generally being dissipative, introducing issues of grid metric continuity thus complicating grid generation, requiring large simulation domains for capturing realistic rotor-stator blade counts, and introducing parallelisation issues. Recently a new sliding plane method based on *Characteristic Interface Conditions* has been proposed [Johnstone et al., 2015], which allows for grid discontinuities and involves a single overlapping node, thereby minimising issues of grid generation and interpolation, and has been integrated successfully in a new DNS code HiPSTAR [Sandberg et al., 2015].

Despite these issues, sliding meshes based on overlapping grids will be used in this thesis to simplify the numerical formulation. As will be made clearer in the next section, the goal of this thesis is to introduce a direct-adjoint framework for the analysis of turbomachinery aeroacoustics in the context of self-excited instabilities, and so the numerical schemes used must be able to capture this unsteadiness without additional modelling in the time domain. Furthermore, for direct and adjoint approaches it becomes favourable from an implementation perspective because information exchanges between the rotor and stator domains become simple function compositions of a communications procedure and a linear interpolation. Both of these can be represented using matrices, thereby avoiding the need to perform linearisations of the sliding plane itself.

1.4 Techniques for Analysis

Ultimately, the goal of computational analysis is to gain a mechanistic perspective on the problem in question, so that a fundamental understanding of the processes and flow interactions is constructed, and to enable efficient strategies for flow control and optimisation to be deployed. With increased simulation fidelity and the continued increases of computing power predicted in the coming years, simulation datasets are expected to increase in both size and complexity which would render conventional

tools for analysis ineffective (see e.g. review of Lele and Nichols [2014]). Future work in this field should therefore focus on maximising the knowledge gain per terabyte of simulation data, instead of purely pursuing and relying on resolution and fidelity trends. In this way, low-dimensional behaviours are identified, models developed and effective control strategies deployed.

To achieve this, modal [Taira et al., 2017] and non-modal stability theory [Schmid, 2007] present a promising framework for performing this analysis, with these techniques facilitating a granular-level description of the underlying flow processes and physics in a wide range of experimental and numerical contexts.

Of these, perhaps the most familiar in engineering circles is the modal approach. Mathematical decompositions are applied to extract low-dimensional bases for capturing the structural and temporal evolution of the flow, and may be data-driven or deterministic, depending on the nature of the data and whether or not the underlying linear operators can be accessed with relative ease. Proper Orthogonal Decomposition (POD) [Lumley, 1967], spectral POD (SPOD) and Dynamic Mode Decomposition (DMD) [Schmid, 2010; Tu et al., 2014] are four examples of data-driven modal methods that have seen especially wide adoption in recent years. These methods are particularly attractive in experimental and computational settings since they require only a relatively small subset of flow field solution vectors[8] which can be extracted directly from nonlinear fields obtained using numerical and observed experimental data. A key property of POD and its derivative methods is that these yield orthogonal bases, ranked by energy content, and so are especially well-suited to reduced-order modelling [Taira et al., 2017], but amalgamate the dynamics with different frequencies into single modes. DMD addresses this by isolating modal structures associated with distinct frequencies, though with the DMD spectrum in its raw form it may be difficult to determine the optimal subset of modes that captures the dominant dynamics. Solutions such as Sparsity-Promoting DMD (DMDSP)

[8]These are referred to as *snapshots* from now on.

[Jovanović et al., 2014] attempt to remedy this by optimising the residual between the matrix of nonlinear flow snapshots and the DMD representation as a function of DMD amplitudes for a given sparsity[9], and therefore plotting the norm of the residual against the sparsity structure provides a possible ranking of the modes. A comprehensive review of these methods, including more detailed discussion of the strengths and weaknesses, is available in [Taira et al., 2017], with recent applications summarised in [Taira et al., 2020].

In spite of the numerous successful applications of these data-driven methods, a critical weakness is that these approaches do not provide any indication of the structural sensitivities for these modes. For experimental settings this indeed may be the only choice, but if a nonlinear numerical model and its linearisation are available, then the theory and methods of linear operators can be exploited, greatly improving the understanding of fundamental flow mechanisms.

At this point it is worth devoting some of this discussion to the meaning of the linearisation and the linear stability analysis of flows that are past the critical Reynolds number. In the earliest applications of global stability analysis, such as for Poiseuille flow and flows around cylinders, the conventional approach has been to take the steady base flow (i.e. the fixed point in dynamical systems theory) around which the linearisation is performed, and which naturally occurs provided that the Reynolds number is low enough to ensure a steady state. Evolution of linear perturbations, $\mathbf{v}'$, on top of this base flow, $\mathbf{v}_0$ are then readily given by the familiar Taylor expansion

$$\frac{\mathrm{d}}{\mathrm{d}t}(\mathbf{v}_0 + \epsilon\mathbf{v}') = \mathbf{v_0} + \epsilon\frac{\mathrm{d}\mathbf{F}}{\mathrm{d}\mathbf{v}}\bigg|_{\mathbf{v}_0}\mathbf{v}' + \mathcal{O}\big(\epsilon^2\big), \tag{1.1}$$

where $\mathbf{F}$ is the nonlinear right-hand side describing the temporal evolution of the nonlinear flow. The result is that the linear equations take the form

$$\frac{\mathrm{d}\mathbf{v}'}{\mathrm{d}t} = \frac{\mathrm{d}\mathbf{F}}{\mathrm{d}\mathbf{v}}\bigg|_{\mathbf{v}_0}\mathbf{v}' = \mathbf{A}\mathbf{v}', \tag{1.2}$$

[9]This can also be performed in reverse with a sparsity structure obtained for a given residual norm.

with **A** denoting the linear operator of the linearised Navier-Stokes equations. Eigenvalues of **A** determine the characteristics of the fixed point, and when the real parts of these transition to the unstable half-plane, typically as a result of an increased Reynolds number, the flow changes from being stable to an unstable state in which disturbances grow. Finding these steady states at higher Reynolds numbers becomes extremely challenging, however because the flow is no longer naturally steady, and an adaptation of this strategy is required.

It has been shown in the past that, for unstable flows, classical stability theory with steady base flows only provides an accurate prediction of the Strouhal number, St = St(Re), in the vicinity of the Hopf bifurcation [Barkley, 2006]. This is somewhat expected because it is known that nonlinearities play significant roles in selecting the dynamics of unsteady flows. It has been shown that introducing the mean flow as the linearisation state factors in these nonlinearities and results in accurate predictions of the dominant flow features. Subsequent application of mean-flow stability analyses has resulted in accurate low-dimensional descriptions of both the structures and frequencies in many settings, such as cylinder flows [Barkley, 2006], jets [Gudmundsson and Colonius, 2011], fuel injectors [Juniper, 2012; Tammisola and Juniper, 2016] and airfoils [Fosas de Pando et al., 2014], to name a few examples. Recent work by Beneddine [Beneddine et al., 2016] has provided the theoretical conditions for the validity of the mean-flow stability analysis in turbulent flows.

In many fluid problems, however, **A** is non-normal [Schmid, 2007] [10], and so special care should be taken when interpreting the results of the modal analysis. This is because such systems are characterised by large transient energy growth phenomena at small time scales [Trefethen and Bau III, 1997], and blind reliance on the operator spectrum would not characterise the full dynamical evolution of the system because these transient effects could result in the selection of alternative structures. Therefore, a non-modal approach is required. A well-known example of

[10] That is, the operator has non-orthogonal eigenvalues.

this is plane Poiseuille flow, where the spectrum is stable until $\mathrm{Re} = 5772$, but small disturbances nonetheless exhibit large growth for $87 \leq \mathrm{Re} \leq 5772$ before eventually decaying [Reddy et al., 1993; Schmid, 2007]. Such phenomena can be studied with ϵ-pseudospectra [Trefethen and Bau III, 1997], which define the sensitivity of the eigenvalues to small operator perturbations, but an alternative approach is to consider this from the perspective of sensitivities to initial conditions or the response of a forced system. This amounts to analysing the gain and resolvent functions

$$G(t) = \max_{\mathbf{v}_{\text{in}}} \frac{\|\exp(\mathbf{A}t)\mathbf{v}_{\text{in}}\|}{\|\mathbf{v}_{\text{in}}\|}, \quad R(\omega) = \max_{\hat{\mathbf{f}}} \frac{\|\hat{\mathbf{v}}\|}{\|\hat{\mathbf{f}}\|} = \|\mathbf{C}(i\omega\mathbf{I} - \mathbf{A})^{-1}\mathbf{B}\|, \tag{1.3}$$

where $\mathbf{v}_{\text{in}}$ is the initial condition, $\hat{\mathbf{f}}$ is the spatial distribution of the forcing function, and $\mathbf{B}$ and $\mathbf{C}$ denoting the input and output masks, respectively. The optimal forcing and response modes can be obtained using a singular-value decomposition (SVD) [Schmid and Henningson, 2001], but direct solutions of this problem are generally formidable due to the considerable degrees of freedom typically encountered for the system matrix. This has led to new techniques such as randomised SVD and randomised resolvent analysis to address these issues, which can then be applied in a computationally tractable manner for realistic, turbulent flows [Ribeiro et al., 2020].

From a control-theory perspective[11], it is desirable to extend this by identifying regions where these modal structures are most sensitive, so that control strategies can be localised efficiently. One way of achieving this is to use the adjoint of the direct linear operator $\mathbf{A}$, which, in essence, captures the gradient information of the problem [Gunzburger, 2002], that can be used subsequently in an appropriate gradient-based optimisation algorithm.

The adjoint approach was first introduced to the fluids community by Jameson who proposed a continuous derivation of the operator, which is referred to here as the *continuous adjoint*, by first deriving an analytic form of the linearisation, defining the

[11] A classical way of deriving the adjoint is through the unconstrained minimisation problem defined by a Lagrangian. The Lagrange multipliers are then the adjoint variables.

adjoint and then discretising it on the numerical grid with bespoke numerical schemes and boundary closures to ensure numerical stability [Jameson, 1988]. In this way the equations are obtained analytically, resulting in faster execution on a computer, but at the expense of a slower and more error-prone implementation, and a lack of consistency with the discretised linear operator, which can lead to convergence issues in optimisation. The alternative, suggested by Giles [Giles and Pierce, 2000], is to take a linear algebra approach by first discretising the linear operator followed by a simple[12] transpose-conjugation to arrive at the adjoint operator. In performing the discretisation first, the linear operator captures the characteristics of the numerical schemes and the choice of boundary closures, so that, when the adjoint is derived, an entirely self-consistent system is obtained without the need for additional labour-intensive adaptation of the adjoint equations. It is for this reason that the *discrete adjoint* is the preferred approach in this thesis. On the other hand the downside is that the code must be linearised numerically, which is generally an expensive operation that can potentially limit this approach, as will be shown later.

Structural sensitivity analysis and optimisation studies have seen extensive application of the adjoint (see e.g. [Hill, 1992; Giannetti and Luchini, 2007]). In the latter case, Giannetti and Luchini [2007] used *wavemaker* regions by defining them to be the spatial intersection of the direct and the adjoint mode, and applying them to the problem of cylinder wakes. Further contributions have also used this analysis on isolated tonal noise problems [Fosas de Pando et al., 2017], with the aid of which feedback mechanisms on both the pressure and suction sides of the airfoil were identified as instrumental to the aeroacoustic feedback loops and the selection of the dominant frequency of the airfoil self-noise problem.

The principal application of the adjoint does however remain optimisation, because exact gradients can be obtained for many control variables with only one

[12] Although conceptually simple, this requires some careful consideration as, in general, the linear operator is not available in closed form due to the considerable size of storing the full system matrix – which is generally dense.

solution of the adjoint differential equation required [Gunzburger, 2002]. These methods have seen extensive and successful use in aerodynamic shape optimisation, for example for isolated airfoils [Jameson, 1988], wings [Burgreen and Baysal, 1994; Reuther and Jameson, 1995], full aircraft configurations [Reuther and Jameson, 1995; Reuther et al., 1996] and, more recently, has been applied to turbomachinery [Wang and He, 2010; Wang et al., 2010]. Until very recently, however, most publications from the turbomachinery community making use of the adjoint method have considered primarily aerodynamic optimisation problems using steady flows. Effects of unsteadiness have been introduced in the past using the adjoint in combination with frequency domain harmonic balance approaches. These include Thomas et al. [2005] who computed the aerodynamic sensitivities for an isolated, harmonically pitching airfoil, Duta et al. [2002] investigated a minimisation problem for a wake-induced forced response problem of a turbine blade, while He and Wang [2011] performed a concurrent aerodynamic-aeroelastic adjoint optimisation of a single rotor blade row, and were able to demonstrate a 40% increase in aerodynamic damping and a 25% decrease in forced response, with an effectively negligible aerodynamic penalty.

Recent work has sought to remove the constraints of this frequency-domain analysis by reformulating the rotor-stator interaction problem in the time-domain so that non-periodic phenomena could be analysed [Ntanakas et al., 2018]. In this work the unsteady time-domain adjoint method was formulated for a three-dimensional rotor-stator URANS-based shape optimisation problem, and introduced an adjoint sliding plane formulation and a 'temporal coarsening' technique to reduce the computational costs of the adjoint solve. Despite being a time-domain formulation, and aside from being formulated in the context of a URANS solver, this approach is highly dissipative at the sliding plane and does not propagate either the nonlinear nor the adjoint field cleanly though the interface. This point is especially critical to address for aeroacoustic applications, and will be a core consideration of this thesis.

Extensions of these adjoint techniques to time-dependent, time-domain sensitivity

problems presents some additional complexities, and in particular the problem of storing the nonlinear trajectory. In steady and unsteady situations, the forward linear or nonlinear problem is solved first to determine the initial conditions for the adjoint, following which the adjoint equations are solved backwards in time to recover the gradient information [Gunzburger, 2002]. For steady problems the linear operator is autonomous (time-independent), but for unsteady problems there is an implicit dependence on the underlying trajectory, because the linear operator has to be re-linearised at each time step of the adjoint. Since the equations are solved backwards in time, this dependence requires the storage of the entire trajectory for the full time-horizon of the simulation with grave implications for system storage, especially in the case of large fields arising in compressible flow solvers. Checkpointing was introduced to remedy this by regularly saving solution snapshots and then recovering missing data when it is required. The `revolve` library [Griewank and Walther, 2000] is perhaps the most popular checkpointing library in applications. Multi-stage extensions that allow the use of both disk and random-access memory (RAM), and other libraries such as `a-revolve` and algorithms that dynamically allocate checkpoints have also been proposed [Wang et al., 2009].

The discussion of the various numerical techniques is not complete here without at least a brief mention of the high-performance numerical libraries that have enabled the considerable progress in recent computer implementations. Libraries such as BLAS and LAPACK [Anderson et al., 1999], provide respectively the low- and high-level implementations for matrix-vector algebra, and are the underpinnings of many more advanced libraries, such as ScaLAPACK [Blackford et al., 1997], which extends eigenvalue and SVD calculations of dense matrices represented on parallel architectures. For the work in this thesis these libraries are used in the more high-level context of PETSc [Abhyankar et al., 2018; Balay et al., 2019, 2020, 1997], which provides a tightly integrated collection of tools for the solution of large-scale parallel partial-differential equations, and extensions such as SLEPc [Hernandez et al., 2005],

which extend this to provide high-performance scalable eigenvalue solvers. PETSc also contains useful integration with the TAO optimisation framework, and, of particular interest to unsteady adjoint problems, the TSAdjoint [Zhang et al., 2019] module that enables optimal multi-level checkpointing using `revolve` and other methods.

1.5 Thesis Objectives, Contributions and Outline

The principle objective of this thesis is to develop an efficient framework for the application of time-domain direct and adjoint techniques to the study of turbomachinery aeroacoustics, in the various contexts that have been covered in the earlier sections of this introduction. Given the currently intractable costs of full three-dimensional time-domain DNC calculations, and the limitations to the available computing resources during this project, two classes of representative problems, that capture the essential complexities of turbomachinery aeroacoustics simulations, are considered.

The first case considers the flow through a linear cascade of controlled diffusion (CD) airfoils, which is modelled at a Mach number of $M = 0.3$ and Reynolds number $\mathrm{Re} = 1 \times 10^5$, and is adapted from earlier aeroelasticity experiments representative of a compressor cascade [Yang and He, 2004]. These experiments and subsequent scale-resolving numerical simulations [He and Yi, 2017; He, 2021] have shown the presence of an oscillating laminar separation bubble on the suction surface of the airfoil that sheds vortices that subsequently interact with the blunt trailing edge of the airfoil. The stability, sensitivity and receptivity properties of a single periodic passage, which we will refer to from now on as the one-periodic problem, will be obtained using a combination of a transient response analysis and a time-domain direct-adjoint mean-flow global stability analysis. We will then analyse the direct and adjoint global modes, compute the wavemakers, and, in the process, identify the feedback loops governing the dynamics of the system. A wavenumber analysis of the modes will also be performed to identify the large-scale structures and wave

packets that are formed under superposition of the global modes.

Given that many turbomachinery flows occur within domains defined by large-scale periodic structures supporting macroscopic behaviour that encompasses multiple blade passages, a major disadvantage of the one-periodic problem is that structures with azimuthal wavelengths greater than the blade pitch cannot be captured without additional modelling. Attempting to simulate domains with N blades per period window directly, which is henceforth referred to here as the N-periodic problem, may also be ill-advised due to the saturation in the parallel scaling and the large sizes of the simulation. This presents issues of storage, additional overheads for memory management and the limitations inherent with instantaneous access to the required computational resources. Instead, by assuming that the mean flow, the blade and the mesh geometries are identical in each passage, the linear operator governing the evolution of the linear initial value problem (IVP) attains block-circulant form. The operator is block-diagonalised, thereby breaking up the large N-periodic system into a series of smaller IVPs for a smaller, modified operator. This is inspired by similar earlier work on aeroelasticity problems [Pierre and Murthy, 1992] and extends a recent adaptation of this theory in the context of an incompressible flow solver [Schmid et al., 2017] to a high-fidelity compressible fluids code capable of representing aeroacoustics processes.

Modifications are introduced in the linear code segments, and a direct-adjoint global stability analysis is performed for $N = 10$, with the modes categorised by periodicity. The solution of the N-periodic IVP for an impulse response study further elucidates the presence of new feedback loops within the system, and helps explain the genesis of large-scale synchronisation waves that dominate the dynamics.

The effects of symmetry-breaking perturbations to an $N = 2$ system are subsequently considered for a simplified two-passage geometry in two aerodynamically mistuned configurations, something that, to the best of the author's knowledge, has not previously been considered in aeroacoustics global stability contexts.

The second class of problems considered in this thesis concerns the use of the adjoint for sensitivity and optimisation problems in rotor-stator interaction contexts. While previous work has successfully applied adjoint-based methodologies to steady-state and low-order approaches, these approaches are ultimately unable to model the generation of boundary layer instabilities and their propagation through the rotor-stator interface. The contribution presented in this thesis addresses this by introducing a high-order nonlinear-adjoint methodology that is capable of fully resolving these unsteady flow features in both the forward and adjoint simulations. Solution checkpointing is introduced with the aid of open source libraries, and extensive validations are performed. Care is taken throughout to ensure the universality of this methodology to other codes and schemes, with the development designed to be as independent of the numerical schemes as possible, so that it may be implemented easily in a wide range of contexts. The theory is then applied to a sample rotor-stator interaction problem, with two stator and three rotor blades defining the geometry, following which a sensitivity analysis is carried out.

1.5.1 Thesis Structure

Based on this, the structure of the thesis can now be summarised as follows.

Chapter 2 introduces the fifth-order, pseudo-wave, compressible Navier-Stokes code that is the foundation of the numerical work in this thesis. This was inherited by the current author from [Fosas De Pando, 2012] at the start of this project, and was used previously for the purposes of a global stability analysis of the tonal noise problem on isolated airfoils [Fosas de Pando et al., 2014, 2017]. The details of the implementation are summarised, with the modifications, made for adapting this code to the cascade geometry, provided. Procedures for the automatic linearisation of the nonlinear code are then described alongside structural details of the direct-adjoint components of the code, their application and interface with third-party libraries for the efficient calculation of global modes. This section of the thesis concludes with a

few examples validating and demonstrating the operation of the code.

Chapter 3 presents a blade row global stability analysis in a paper format. This begins with a nonlinear flow analysis, followed by the application of non-modal impulse response and global stability analyses, first to a one-periodic and then to a ten-periodic geometry, which are compared throughout this paper. At the time of writing this paper is in preparation. Appendix C briefly extends this work by considering the sensitivity of the self-excited instabilities and flow structures of tuned system to symmetry-breaking aerodynamic mistuning.

Chapter 4 concerns the development and validation of the nonlinear-adjoint procedure, which is demonstrated in a series of simplified Cartesian geometries. This again is presented in the form of a scientific publication in-preparation. To conclude this work, Chapter 5 applies this theory to the rotor-stator interaction problem, and observations about the sensitivity of the system to initial condition are presented.

Finally, Chapter 6 concludes this thesis with a discussion of the main findings and perspectives on future work.

This thesis is presented in an integrated format, with a mixture of chapters and papers forming the main document. Each paper in Chapter 3 and Chapter 4 is self-contained, is prefaced with a short introduction, and the structures of the papers adapted so that the layouts of the thesis and papers are consistent alongside correct labelling of equations, figures and tables. To facilitate the flow of the thesis, some references and explanations have also been adapted so that the papers fit better within the overall structure.

1.5.2 Use of Open-Source Software

Open source software enables much of the work in this thesis. Tools such as PETSc [Abhyankar et al., 2018; Balay et al., 2019, 2020, 1997; Zhang et al., 2019], SLEPc [Hernandez et al., 2005] and Eigen [Guennebaud et al., 2010] provide high-

performance libraries for the solution of large-scale parallel scientific problems defined by partial differential equations and sparse eigenvalue problems, respectively, with Eigen being used in this thesis for sparse matrix algebra. Parallelisation of the code is achieved through the Message-Passing Interface (MPI) standard [MPI Forum], and high-performance, parallel IO is implemented using the Hierarchical Data Format (HDF5) [The HDF Group, 2021].

All post-processing, with the exception of Figure 2.11b, is performed using Python 3.8 with NumPy [Harris et al., 2020] and SciPy [Virtanen et al., 2020] used for the bulk of the post-processing. Visualisations and plots are rendered using Matplotlib [Hunter, 2007]. Additional modules, such as `h5py` [Collette, 2013] and `mpi4py` [Dalcín et al., 2005, 2008, 2011], are used for data extraction from files and parallel processing of datasets, respectively.

1.5.3 Comments on Figures

Figures play a major part in this project, and are designed to enhance the reader's understanding of the text. For this reason specific care is taken in the selection of the colourmaps used for each figure. Figures depicting nonlinear flow fields are rendered with the perceptually uniform sequential colourmap `viridis`, which renders figures in a purple-yellow scheme. As is conventional in the literature, dilatation fields are show in grey-scale.

Linear calculations naturally generate diverging data that is centred at zero, and so these are shown in this thesis using a red-blue colourmap `RdBu_r`, showing positive values in red and negative values in blue. For the analysis of wavemakers and errors, absolute values of these quantities are shown using the sequential maps `Blues` and `BuPu`, which render the fields in gradients of blue and purple respectively (see also [Moreland, 2009] for discussions on the design and selection of perceptually uniform colourmaps).

For the comparison of linear modes in Chapter 3 and Appendix C, direct modes

are shown with $p \in [-2 \times 10^{-2}, 2 \times 10^{-2}]$ and $u \in [-5 \times 10^{-1}, 5 \times 10^{-1}]$, although colourbars are omitted because the mode plots are used to compare the structural properties of the modes rather than their quantitative values. Where comparisons are made, supplementary plots with flow values at the airfoil surfaces are provided to aid the discussion. Where a logarithmic colourmap is used, such as for wavemakers, colourbars are shown.

Finally, many of the discussion points in this thesis describe dynamic, time-dependent features of the flow fields that become increasingly difficult to demonstrate in the static format of a printed page. For this reason, certain plots throughout the report have been adapted into animations which have been included as supplementary material available with this thesis.

It is hoped that these considerations alongside this explanatory note will enable greater understanding and accessibility of this thesis for the reader.

Chapter 2

Numerical Methods

2.1 Introduction

Although the aim of this work is to develop and present an efficient and scalable computational framework for the analysis of various aeroacoustics problems commonly found in turbomachinery, we must first set out the foundations of the numerical methods used.

At its most elemental, the study of aeroacoustics is the study of the interactions between the hydrodynamic instabilities and the acoustic field. It is therefore clear that the numerical simulations, by means of which this analysis is to be carried out, must be capable of resolving both the *generation and propagation* of the instability waves as well as the acoustic field. These are coupled through feedback mechanisms, and it is the analysis of these that is the particular focus of this thesis. We have seen in the review above that this requirement to model the acoustic and hydrodynamic features simultaneously places specific and strict demands on the numerical models, thus making traditional, dissipative, low-order CFD methods largely unsuitable, and therefore requiring the development of bespoke CAA codes.

In combination with the nonlinear solution, the linearised dynamics are also desired so that the rich theory of linear operators can be leveraged in the analysis. In particular, efficient access both the direct and adjoint dynamics should be established.

The basis of the numerical code used here is the Direct Noise Computation (DNC) implementation `aaX`, that was used previously to model high-speed compressible flows at moderate Reynolds numbers, and successfully applied to a mean flow global stability and sensitivity analysis of an airfoil tonal noise problem [Fosas De Pando, 2012; Fosas de Pando et al., 2014, 2017]. This code was inherited by us at the start of this project, and forms the proven numerical foundations for our subsequent work in the later chapters.

Nonlinear components of the solver model the time-accurate evolution of the compressible Navier-Stokes equations, with various flow configurations for open and closed flows being defined through the exact form of the computational domains, and the corresponding boundary conditions. Furthermore, these computational domains are discretised into structured grids, with high-order compact finite difference schemes applied to minimise numerical dissipation, and the solutions integrated in time through the method of lines.

In addition to the nonlinear predictive capabilities, this code features completely numerically consistent and efficient linearisation facilities that will be leveraged for modelling the *direct* and *adjoint* dynamics. The linearisation is performed through finite difference approximations of the first derivative for the program sub-blocks defining the spatially decoupled portions of the nonlinear program. This results in sparse system matrices that are easily combined through simple arithmetic operations to yield the direct operator. Sparsity is further exploited through sequential transpose conjugation of the expression for the direct operator to generate the adjoint, without the need for transposition of dense matrices. All operations are defined through matrix-vector multiplications, and therefore storage of the dense system matrices is not required.

Our code is implemented with a specific focus on maintainability, modularity and parallelism. The main source code is written in C++ with object-oriented programming paradigms used throughout to maximise code reuse, and with For-

tran components integrated through compiled libraries. Specialist libraries, such as PETSc [Abhyankar et al., 2018] and SLEPc [Hernandez et al., 2005], interface with this code through wrappers around the linear portions of the code, and provide performant and validated routines for the necessary linear algebra algorithms. Parallelisation of the code is handled through the Message Passing Interface (MPI) standard [MPI Forum], and parallel I/O implemented with HDF5 [The HDF Group, 2021], to ensure optimal scaling performance on modern supercomputers.

Although this code is the basis of our analysis here, its modular structure alongside our implementation-agnostic numerical methods in later chapters means that the interested reader can easily substitute their own codes and schemes into this framework, provided that certain code structures, such as availability of a linearisation component, are maintained.

The following chapter is a brief summary of the comprehensive discussion found in [Fosas De Pando, 2012] on the implementation of the nonlinear code in our solver, which we include here for clarity and completeness. This is then followed by a description of the linearisation component, and the integration of the direct and adjoint dynamics with high-performance matrix algebra libraries, that forms the foundation for future chapters. We conclude here with a section containing two validation cases that demonstrate the validity and accuracy of this code.

2.2 Governing Equations

2.2.1 The Compressible Navier-Stokes Equations

In this work we use the compressible Navier-Stokes equations to model aeroacoustic processes in subsonic gas flow at moderate Reynolds numbers. Derived from the principles of the conservation of mass, momentum and energy for a fluid particle, these are written below in terms of the density ρ, velocity $\mathbf{u}$ and internal energy e in Equations (2.1)–(2.3). These equations are closed by an expression for the

viscous stress tensor $\boldsymbol{\tau}$ in Equation (2.4), and Fourier's law of heat conduction in Equation (2.5) providing an expression for the thermal flux $\mathbf{q}$.

$$\frac{\partial \rho}{\partial t} + \boldsymbol{\nabla} \cdot (\rho \mathbf{u}) = 0, \tag{2.1}$$

$$\rho\left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \boldsymbol{\nabla}\mathbf{u}\right) = -\boldsymbol{\nabla} p + \boldsymbol{\nabla} \cdot \tau + \mathbf{f}_v, \tag{2.2}$$

$$\rho\left(\frac{\partial e}{\partial t} + \mathbf{u} \cdot \boldsymbol{\nabla} e\right) = -p\boldsymbol{\nabla} \cdot \mathbf{u} + \boldsymbol{\tau} \colon \boldsymbol{\nabla}\mathbf{u} - \boldsymbol{\nabla} \cdot \mathbf{q} + Q_r, \tag{2.3}$$

$$\boldsymbol{\tau} = \mu\big(\boldsymbol{\nabla}\mathbf{u} + \boldsymbol{\nabla}\mathbf{u}^T\big) + \left(\mu_v + \frac{2}{3}\mu\right)(\boldsymbol{\nabla} \cdot \mathbf{u})\mathbf{I}, \tag{2.4}$$

$$\mathbf{q} = -k\boldsymbol{\nabla} T, \tag{2.5}$$

$$p = \rho R T \qquad \text{and} \qquad e = \frac{R}{\gamma - 1} T, \tag{2.6}$$

with volume forcing, $\mathbf{f}_v$; heat sources, Q_r; viscosity, μ; volume viscosity, μ_v; thermal conductivity, k; the universal gas constant, R; and the heat capacity ratio, γ.

Modelling a *calorically perfect gas* with Equation (2.6), we close the dimensional system. This is a valid assumption for gases moving with subsonic or low supersonic Mach numbers and in situations where the temperature variations are sufficiently small. We will assume that air, with a homogeneous composition and a constant ratio of specific heats $\gamma = 1.4$, is the working fluid. For the purposes of this study, this is sufficient, however these assumptions may need to be revisited in the case of the combustor or turbine aeroacoustics as the convected combustion products and large thermal gradients have non-trivial effects on the acoustic response of the turbine. For example, Magri et al. [2016] showed that indirect noise arising from the acceleration of a heterogeneous gas depends depends strongly on the mixture composition, although further experimental work is required to establish the significance of these findings for realistic aero engines.

To apply the pseudo-wave scheme of Sesterhenn [2000], we recast these equations in terms of pressure; entropy, s; and velocity fields. Using the combined first and

second laws of thermodynamics $\mathrm{d}e = T\mathrm{d}s - p\mathrm{d}(1/\rho)$, and the definition of the speed of sound $a^2 = (\partial p/\partial \rho)_s$, the energy and mass-conservation equations transform into

$$\frac{\partial p}{\partial t} + \mathbf{u}\cdot\boldsymbol{\nabla} p = -\rho a^2 \boldsymbol{\nabla}\cdot\mathbf{u} + (\gamma-1)\rho T\left(\frac{\partial s}{\partial t} + \mathbf{u}\cdot\boldsymbol{\nabla} s\right), \tag{2.7}$$

$$\rho T\left(\frac{\partial s}{\partial t} + \mathbf{u}\cdot\boldsymbol{\nabla} s\right) = \tau\colon \boldsymbol{\nabla}\mathbf{u} - \boldsymbol{\nabla}\cdot\mathbf{q} + Q_r. \tag{2.8}$$

2.2.2 Non-Dimensionalisation

Following standard procedure, we cast Equation (2.7), Equation (2.8), and Equation (2.2) in terms of non-dimensional variables, which are denoted by means of a superscript $'$, and defined as follows

$$\mathbf{x}' = \frac{\mathbf{x}}{x_0},\ t' = \frac{u_0 t}{x_0},\ p' = \frac{p}{p_0},\ s' = \frac{s}{R},\ \rho' = \frac{\rho}{\rho_0},\ \mu' = \frac{\mu}{\mu_0},\ k' = \frac{k}{k_0}. \tag{2.9}$$

The non-dimensionalisation is performed by selecting certain characteristic scales of the physical problem, defined by the subscript $\star_0$, which are chosen based upon the problem being solved. For example we can choose these dimensions to be the chord length of the airfoil alongside the flow parameters at fixed point within the flow domain for the values of p_0, u_0, ρ_0, μ_0 and k_0.

Applying the transformations from 2.9, the non-dimensional equations become

$$\frac{1}{\rho}\left(\frac{\partial p}{\partial t} + \mathbf{u}\cdot\boldsymbol{\nabla} p\right) = -\gamma\boldsymbol{\nabla}\cdot\mathbf{u}\left(\frac{\partial s}{\partial t} + \mathbf{u}\cdot\boldsymbol{\nabla} s\right), \tag{2.10}$$

$$\rho\left(\frac{\partial u}{\partial t} + \mathbf{u}\cdot\boldsymbol{\nabla}\mathbf{u}\right) = -\frac{1}{\gamma M^2}\boldsymbol{\nabla} p + \frac{1}{\mathrm{Re}}\boldsymbol{\nabla}\cdot\tau + \mathbf{f}_v, \tag{2.11}$$

$$p\left(\frac{s}{t} + \mathbf{u}\cdot s\right) = \frac{\gamma M^2}{\mathrm{Re}}\tau\colon \boldsymbol{\nabla}\mathbf{u} - \frac{\gamma}{(\gamma-1)\,\mathrm{Pr}\,\mathrm{Re}}\boldsymbol{\nabla}\cdot\mathbf{q} + Q_r, \tag{2.12}$$

$$\tau = \mu\left(\boldsymbol{\nabla}\mathbf{u} + \boldsymbol{\nabla}\mathbf{u}^T\right) + \left(\mu_v - \frac{2}{3}\mu\right)(\boldsymbol{\nabla}\cdot\mathbf{u})\mathbf{I}, \tag{2.13}$$

$$\mathbf{q} = -k\boldsymbol{\nabla} T, \tag{2.14}$$

$$p = \rho T, \tag{2.15}$$

$$(\gamma - 1)s = \log \frac{p}{\rho^\gamma}, \tag{2.16}$$

$$a^2 = \frac{p}{\rho}, \tag{2.17}$$

$$\mathrm{Re} = \frac{\rho_0 u_0 x_0}{\mu_0}, \tag{2.18}$$

$$\mathrm{Pr} = \frac{\gamma R}{\gamma - 1}\frac{\mu_0}{k_0}, \tag{2.19}$$

$$M = \frac{u_0}{a_0}. \tag{2.20}$$

Here Re is the Reynolds number, Pr is the Prandtl number and M is the Mach number for this flow. For convenience the primes denoting the non-dimensional quantities have been dropped here and through to the end of this thesis.

2.3 Modelling Domains

The numerical solver solves for the flow variables $\mathbf{v} = (p, s, \mathbf{u})$, using Equations (2.10)–(2.20), at all points within some *physical domain*, $\mathcal{P}$, which contains the solid surfaces that interact with the flow and acoustic fields.

Here we are primarily concerned with external flows, where $\mathcal{P}$ is considered to have a 'boundary' at infinity alongside a physical boundary at the blade surface. Since modelling an infinite, or a very large domain is computationally impractical, we truncate $\mathcal{P}$ so that the spatial extent of the domain is of a similar size as that of the solid object[1], thereby minimising the cost of solving the propagation problem in the far-field regions. With this smaller domain, the computational cost is necessarily reduced at the expense of adding additional complications at the now *fictitious* boundaries delineating the inflow and outflow regions of the truncated region. These boundaries must therefore be able to handle large unsteady disturbances without

[1]For instance the horizontal length of the domain could be a small multiple of the airfoil chord.

introducing artificial feedback into the interior domain from the numerical treatment at the boundary.

In all but the simplest Cartesian geometries, it is no longer practical or desirable to design bespoke numerical schemes for each new geometry that one considers. Doing so is not only time consuming but also results in complicated expressions for the derivatives and boundary conditions, many of which rely on normal derivatives at the solid interfaces. This significantly limits the flexibility and practicality for the numerical code while at the same time increasing the likelihood of errors and inconsistencies in the implementation.

The numerical schemes modelling the evolution of Equations (2.10)–(2.20) make use of *structured* Cartesian grids. These subdivide $\mathcal{P}$ into several regions into each of which a simple Cartesian mesh is embedded. Boundary conditions at solid walls that contain wall-normal or wall-tangential components (e.g. non-slip and no-flux conditions) are simplified by stipulating that the mesh is locally orthogonal at the boundary surfaces.

2.3.1 Curvilinear Coordinates and Transformations

From a computational perspective it is convenient to define a bijective mapping between $\mathcal{P}$ and a standardised computational domain $\mathcal{C}$, in which a standard set of operations is performed for all geometries, and adding small portions of the code that modify its behaviour to account for differing boundary conditions, for example. To that end, we define a complementary computational domain $\mathcal{C}$, with uniform curvilinear coordinates $\boldsymbol{\xi}$, $\xi^k \in [0, 1]$, in which the calculations are performed. The utility of representing $\mathcal{P}$ in terms of this new domain $\mathcal{C}$ is that, in $\mathcal{P}$, the forms of the governing equations, boundary conditions and differentiation schemes may take on complex forms that would depend on the chosen geometry. Instead, by delegating the change of coordinates to the bijective map between $\mathcal{P}$ and $\mathcal{C}$, these forms are simplified, and allow for changes to the geometry of $\mathcal{P}$ without requiring manual

adjustment of the calculation procedures that are performed in $\mathcal{C}$.

The bijective mapping, that we define as $\mathbf{x}(\boldsymbol{\xi})$, with $\mathbf{x} \in \mathcal{P}$ and $\boldsymbol{\xi} \in \mathcal{C}$ maps these domains so that if

$$X = \{\mathbf{x} \colon \mathbf{x} \in \mathcal{P}\}, \qquad \Xi = \{\boldsymbol{\xi} \colon \boldsymbol{\xi} \in \mathcal{C}\}, \tag{2.21}$$

$$\partial X = \{\mathbf{x} \colon \mathbf{x} \in \partial\mathcal{P}\}, \qquad \partial\Xi = \{\boldsymbol{\xi} \colon \boldsymbol{\xi} \in \partial\mathcal{C}\}, \tag{2.22}$$

then $\mathbf{x} \colon \Xi \to X$ and $\mathbf{x} \colon \partial\Xi \to \partial X$. From this, the fields in $\mathcal{P}$ are correspondingly mapped to $\mathcal{C}$ by $f(\mathbf{x}(\boldsymbol{\xi})) = f(\boldsymbol{\xi})$.

With this formulation, grid Jacobians naturally arise in the calculations involving derivatives of the field quantities, and are

$$j_{ik} = \frac{\partial x_i}{\partial \xi^k}, \qquad k^{ki} = \frac{\partial \xi^k}{\partial x_i}, \tag{2.23}$$

with the Cartesian basis transforming to

$$\boldsymbol{\xi}_k = j_{ik}\mathbf{x}_i, \qquad \boldsymbol{\xi}^k = k_{ki}\mathbf{x}_i, \tag{2.24}$$

and the metric tensors

$$g_{kl} = \boldsymbol{\xi}_k \cdot \boldsymbol{\xi}_l, \qquad g^{kl} = \boldsymbol{\xi}^k \cdot \boldsymbol{\xi}^l, \tag{2.25}$$

$$h_k = \sqrt{g_{kk}}, \qquad h^k = \sqrt{g^{kk}}. \tag{2.26}$$

Transformations of the vector quantities can then be written in terms of these grid metric tensors as[2]

$$a^k = k^{ki}a_i, \qquad a_k = g^{kl}a_l, \tag{2.27}$$

$$a^k = g^{kl}a_l, \qquad a_k = g_{kl}a^l, \tag{2.28}$$

and the gradient in $\mathcal{C}$ is defined as

$$\boldsymbol{\nabla} = \xi^k \frac{\partial}{\partial \xi^k}, \tag{2.29}$$

where the Einstein summation convention for repeated indices is implied.

[2]Any vector $\mathbf{a}$ can be represented in either a covariant, $a^k\boldsymbol{\xi}_k$, or contravariant, $a_k\boldsymbol{\xi}^k$, basis.

In more complex configurations it may be desirable to subdivide the domain into sub-domains (or blocks) instead of considering $\mathcal{C}$ as a single block. These may include regions with similar boundary conditions, changes between different grid topologies or simply multiple blocks that are arranged to construct a more complex configuration. In such cases, the mapping is designed so that each region is mapped to its own sub-region of the unit cube $\mathcal{C}$. A good example of this can be found in later chapters where the blade row is subdivided into single upstream and downstream regions with periodic boundary conditions, and individual regions for each blade passage with the solid wall boundaries applied to the upper and lower surfaces.

Although this contingency does not arise in the work here, in the case of an O-grid alongside an H-grid, which is commonly used to represent turbomachinery blades with blunt trailing edges, it may also be advantageous to map the physical domain to a computational domain consisting of a union of smaller unit cube domains, as this may simplify the accounting of the indices and relevant dimensions in the implementation.

A key limitation of this approach for arbitrary geometries, however, is the need to devise a generalised grid connectivity program[3] so that data can be exchanged correctly between the constituent sub-blocks in such a way that is independent of sub-block topology. Various strategies exist, for example the popular overset (or chimera) grid method (e.g. [Steger et al., 1982; Delfs, 2001; Wang et al., 2014a]), but in this thesis we will restrict ourselves to cases where the grids at the interfaces coincide and have a globally continuous grid metric in the entire domain. This undoubtedly introduces additional constraints on the grid generation process, and does present difficulties when generating grids for turbine blades where the flow turning is high, however for compressor-type blading considered in this thesis, this was found to retain sufficient flexibility to enable the generation of high-quality grids.

[3]This is generally true for schemes utilising structured grids, since in an unstructured approach there is considerable freedom in choosing grid points.

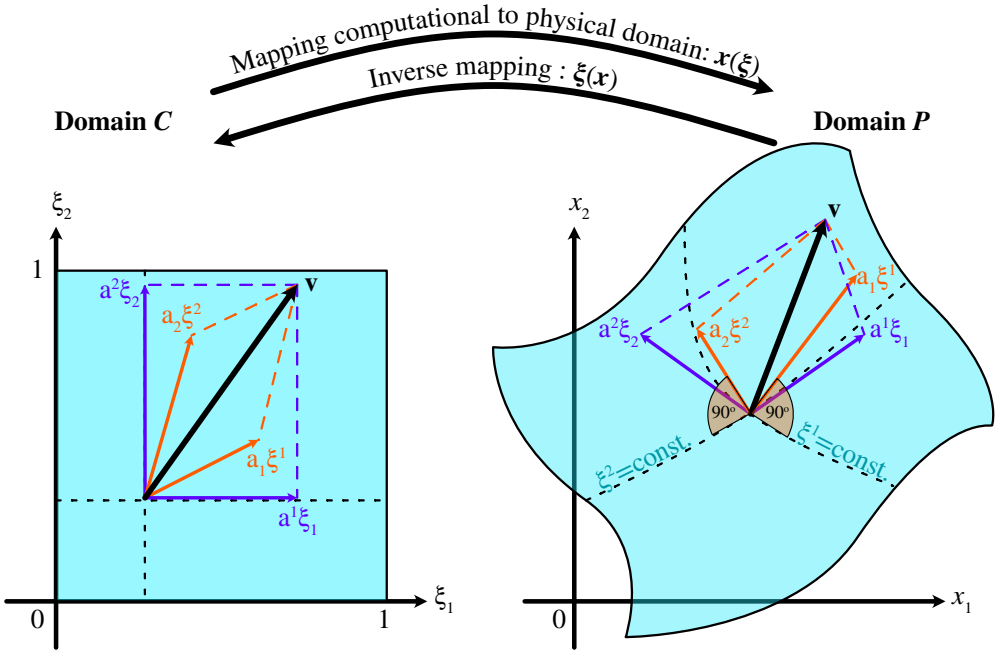


Figure 2.1: The mapping between the physical $\mathcal{P}$ and computational, $\mathcal{C}$, domains showing the contravariant, $\boldsymbol{\xi}_k$, and covariant, $\boldsymbol{\xi}^k$, bases defined on each space. The author of this thesis is grateful for the kind permission of M. Fosas de Pando for the reproduction of this figure from [Fosas De Pando, 2012].

2.3.2 Discretised Fields and Parallelisation

In this implementation the governing equations are discretised on a structured grid, with finite difference schemes being used to approximate the derivative terms.

The simplest choice in this case is a uniform, equispaced partition of $\mathcal{C}$ along each of the basis vectors. The quantity fields are represented here as

$$a_{pi_0i_1\ldots i_n} = a_p\left(\frac{i_0-1}{N_0-1}, \frac{i_1-1}{N_1-1}, \ldots, \frac{i_n-1}{N_n-1}\right), \qquad (2.30)$$
$$p \in \{1,\ldots,P\}, \quad i_a \in \{1,\ldots,N_a\},$$

for an n-dimensional field $\mathbf{a}$.

We define the state as a column vector of $P \times N_0 \times \cdots \times N_n$ elements so that they are grouped by variable and are contiguous in memory. For an arbitrary three-dimensional field, this is equivalent to

$$a_{P\times N_0\times N_1\times(n_2-1)+P\times N_0\times(n_1-1)+P\times(n_0-1)+p} = a_{pn_0n_1n_2}. \qquad (2.31)$$

Modern parallel computer architectures are exploited for their efficiency gains by breaking up these fields onto disjoint blocks, with each block allocated to its own unique processor. Our approach for this is to partition the field with a Cartesian topology in MPI, and, unless otherwise stated, we will assume that in all cases all the data required for all computations is stored locally on the processor once the field is assembled.

Communications between the blocks, when they are required, occur only between neighbouring blocks in the topology, which in turn minimises the communication overheads. This will be elaborated on in Section 2.6.

2.4 Pseudo-Wave Form

The central concept behind the pseudo-wave formulation of [Sesterhenn, 2000] is the idea that the inviscid transport terms of the governing equations may be rewritten in terms of wave-like structures that propagate along the structured grid lines, in a manner similar to the familiar concept of the method of characteristics from hyperbolic PDE theory. Unlike characteristic waves, however, this decomposition does not preserve any invariant quantity, and is therefore a purely numerical concept that is leveraged to achieve the desired numerical stability characteristics of the solution.

Based on the sign in the advection speed of the pseudo-wave terms, upwind-biased differentiation schemes can be aligned with the wave propagation direction along the grid lines. This addresses the known high-frequency numerical stability issues encountered when high-order schemes with low dissipation are applied. Diffusive terms on the other hand are not as sensitive to the numerical stability of the scheme provided that the correct time step is employed, and so these terms are differentiated independently with centred schemes.

An advantageous consequence of this characteristics-like formulation is that the *inviscid* boundary conditions follow directly from the formulation. This will

be the subject of Section 2.5.1 later on, but for now we will point out that for the inflow/outflow boundaries of $\mathcal{C}$, for example, the temporal evolution of the flow quantities there is determined by the desired behaviour in the far-field, and a complex argument may be required to justify certain boundary conditions for the flow information coming into $\mathcal{C}$ from the far-field. Since the grid is orthogonal at the boundary, by construction, however, the information from the far-field is determined precisely by the incoming pseudo-wave, the form of which is determined by the governing equations and the desired boundary condition.

In what follows, a brief outline of the pseudo-wave form of the governing equations is presented. Further discussion and additional details may be found in [Fosas De Pando, 2012].

The pseudo-waves, consisting of the acoustic pseudo-waves $\Theta_k^{\pm}$, cross-transport terms Θ^{kl} and the entropy wave Θ_s^k are

$$\Theta_k^{\pm} = \left(u^{(k)} \pm \frac{a}{M} h^{(k)}\right)\left(\frac{1}{\gamma\rho a}\frac{\partial p}{\partial \xi^{(k)}} \pm M \frac{\partial u_i}{\partial \xi^{(k)}}\frac{k^{(k)i}}{h^{(k)}}\right), \tag{2.32}$$

$$\Theta^{kl} = u^{(l)}\frac{\partial u_i}{\partial \xi^{(l)}}\left(k^{ki} - k^{(l)i}\frac{g^{k(l)}}{g^{(ll)}}\right), \tag{2.33}$$

$$\Theta_s^k = u^{(k)}\frac{\partial s}{\partial \xi^{(k)}}. \tag{2.34}$$

With this, it can be shown that the governing equations Equations (2.10)–(2.20) become [Sesterhenn, 2000]

$$\frac{\partial p}{\partial t} + \gamma\rho a \sum_l \frac{\Theta_l^+ + \Theta_l^-}{2} = (\gamma - 1) p \,\mathrm{RHS}_s, \tag{2.35}$$

$$\frac{\partial u^k}{\partial t} + \sum_l \frac{\Theta_{(l)}^+ - \Theta_{(l)}^-}{2M}\frac{g^{(l)k}}{h^{(l)}} + \sum_l \Theta^{kl} = \mathrm{RHS}^k, \tag{2.36}$$

$$\frac{\partial s}{\partial t} + \sum_l \Theta_s^l = \mathrm{RHS}_s, \tag{2.37}$$

where we have defined the right-hand side dissipation terms as

$$\mathrm{RHS}^k = \frac{1}{\rho}\left(\frac{1}{\mathrm{Re}} k^{ki}\frac{\partial \tau_{ij}}{\partial \xi^l} k^{lj} + k^{ki}\left(f_v\right)_i\right), \tag{2.38}$$

$$\mathrm{RHS}_s = \frac{1}{\rho T}\left(\frac{\gamma M^2}{\mathrm{Re}}\tau_{ji}\frac{\partial u_i}{\partial \xi^l}k^{lj} - \frac{\gamma}{(\gamma-1)\,\mathrm{Pr}\,\mathrm{Re}}\frac{\partial q_i}{\partial \xi^l}k^{li} + Q_v\right). \tag{2.39}$$

In a similar way, the constitutive laws become

$$q_i = -\chi \frac{\partial T}{\partial \xi^l} k^{li}, \tag{2.40}$$

$$\tau_{ij} = \mu \left(\frac{\partial u_i}{\partial \xi^l} k^{lj} + \frac{\partial u_j}{\partial \xi^l} k^{li} \right) + \left(\mu_v - \frac{2}{3}\mu \right) \frac{\partial u_m}{\partial \xi^l} k^{lm} \delta_{ij}. \tag{2.41}$$

As we have shown above, the metric terms arise from the bijective mapping between the domains $\mathcal{P}$ and $\mathcal{C}$, and in turn these depend on the precise point distributions of the mesh. It is clear, then, that the pseudo-waves are not unique and depend on the particular choice of grid, and that converting the expressions in Equations (2.32)–(2.41) between different meshes with the same boundaries is a non-trivial procedure. This point in particular informed our decision of selecting communications with primitive variables between grid sub-blocks and at the sliding plane described in Chapter 4.

Henceforth the vector $\mathbf{v} = [p \quad s \quad u_0 \quad u_1 \quad (u_2)]^\mathsf{T}$ will denote the nonlinear state vector, and in so doing, we can now rewrite the nonlinear evolution equations as

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \mathbf{F}\left(\mathbf{v}, \frac{\mathrm{d}\mathbf{v}}{\mathrm{d}\boldsymbol{\xi}}, \frac{\mathrm{d}\mathbf{v}}{\mathrm{d}\boldsymbol{\xi}}^+, \frac{\mathrm{d}\mathbf{v}}{\mathrm{d}\boldsymbol{\xi}}^-, \frac{\mathrm{d}^2\mathbf{v}}{\mathrm{d}\boldsymbol{\xi}^2} \right). \tag{2.42}$$

The dependence on the derivatives of the state vector of the right-hand side function $\mathbf{F}$ is shown explicitly, with the superscripts $\star^+$ and $\star^-$ denoting the positively and negatively biased upwinding schemes, while the derivatives without subscripts denote the use of a centred scheme. We do so in order to decouple the spatial dependence in $\mathbf{F}$, the reasons for which which will be explored further in Section 2.7.

2.5 Boundary Conditions

In the context of Navier-Stokes solvers, the evolution of the flow field at subsequent times is determined by the initial conditions and the conditions imposed at the boundaries of $\mathcal{C}$. The latter define the interactions at the solid surfaces and at the far-field, fictitious boundaries that arise as a result of the truncation of the physical domain.

Special care must be taken at the far-field to ensure that these boundaries do not introduce spurious reflections that contaminate the interior solution. This is addressed here through a combination of the well-known Navier-Stokes Characteristics Boundary Conditions (NSCBCs) [Poinsot and Lele, 1992] and viscous closures suggested in [Lodato et al., 2008], the forms of which follow directly from the pseudo-wave formulation. Sponge layers augment this partially non-reflective treatment by providing additional damping to outgoing flow disturbances in damping regions surrounding the inflow and outflow boundaries of $\mathcal{C}$.

The nature of the problems considered here, as well as the lack of a shock-capturing capability in this implementation[4], means that we consider just the subsonic (i.e. $|u|M < a$) forms of the boundary conditions.

2.5.1 Inflow Boundary Conditions

Inflow conditions are defined by regions of the boundary satisfying $u^n > 0$ for $\xi^n = 0$ and $u^n < 0$ for $\xi^n = 1$, with the local pseudo-wave advection speed determining the waves in Equations (2.32)–(2.34) that need to be imposed at that particular boundary. Applying the theory of [Poinsot and Lele, 1992] and [Lodato et al., 2008] leads to the incoming waves taking the forms

$$\Theta^n_\pm = \pm\frac{\eta^n_\pm}{L}\big(1-M^2\big)\frac{a}{h^n}(u^n + u^n_{\text{ref}}) + \text{RHS}^\pm_n, \tag{2.43}$$

$$\Theta^{tn} = \pm\frac{\eta^{(t)}}{L}\frac{a}{M}\Big(u^{(t)} - u^{(t)}_{\text{ref}}\Big) + \text{RHS}^{tn}, \tag{2.44}$$

$$\Theta^n_s = \frac{\eta_s}{L}\frac{a}{M}(s - s_{\text{ref}}) + \text{RHS}^n_s\,. \tag{2.45}$$

L is the non-dimensional length of the domain, with relaxation parameters $\eta^n_\pm > 0$, $\eta^t > 0$ and $\eta_s > 0$ determining the characteristics of the filtering performed by

[4] Although it has not been a focus of this project, it is possible to integrate shock capturing in this framework. Adding this functionality to the nonlinear code is established procedure in the literature, however extra care is needed for adapting the linearisation components. For typical shock-capturing schemes, this leads to nonlinear differentiation schemes, which, in principle, can be linearised by using similar methods as detailed by Fosas de Pando et al. [2012]. Linearisations of solutions with shocks have been performed previously, such as by Crouch et al. [2007], who used local solution smoothing in the vicinity of the shock in order to control numerical oscillations arising from the sharp gradients.

these boundary conditions, that remove the outgoing disturbance while enforcing the reference flow values, $\star_{\text{ref}}$.

Additional constraints in the form of Equation (2.46), are applied at the inlet to ensure numerical stability.

$$k^{ni}\frac{\partial \tau_{ij}}{\partial \xi^n}k^{nj} = 0. \tag{2.46}$$

2.5.2 Outflow Boundary Conditions

Similarly, at the outflow boundary, the incoming acoustic characteristic is set to satisfy

$$\Theta^k_{\pm} = \frac{\sigma}{L}\frac{1-M^2}{\gamma M}\frac{p-1}{\rho} + (1-\beta)\,\text{RHS}^{\pm}_k + \beta\,\text{RHS}^{\pm}_{k,ref}, \tag{2.47}$$

with the suggested values for the new relaxation coefficients $\sigma = 0.528$ and $\beta = M$ [Lodato et al., 2008].

As with the inflow, viscous conditions are also imposed for numerical stability, with

$$\frac{\partial q_n}{\partial \xi^n} = 0, \quad \frac{\partial \tau_{tn}}{\partial \xi^n} = 0. \tag{2.48}$$

2.5.3 Solid Wall Boundaries

At the solid wall, the non-slip and wall-normal velocity boundary condition are imposed, with $\mathbf{u} = 0$ and $\partial(\mathbf{u}\cdot\boldsymbol{\xi}^n)/\partial t = 0$, respectively. For the wall normal velocity condition, we should once again solve for the incoming waves, since these are featured in both Equation (2.35) and Equation (2.36). If this process is ignored and the wall-normal derivative condition is imposed directly through $\partial(\mathbf{u}\cdot\boldsymbol{\xi}^n)/\partial t = 0$, then the pressure equation becomes ill-posed. Rearranging Equation (2.36) once again to solve for the incoming characteristics gives [Fosas De Pando, 2012]

$$\Theta^{\pm}_n = \Theta^{\mp}_n \mp \frac{2M}{h^n}\left(\sum_{l\neq n}\frac{\Theta^+_{(l)} - \Theta^-_{(l)}}{2M}\frac{g^{(l)n}}{h^{(l)}} + \Theta^n - \text{RHS}^n\right). \tag{2.49}$$

For the wall-tangent condition, it is sufficient to set $\partial u^t/\partial t = 0$ directly, where u^t is the projection of the velocity in the wall-tangent direction.

Temperature evolves according to

$$\frac{\partial T}{\partial t} = \frac{\gamma - 1}{\gamma} T \left(\frac{1}{p} \frac{\partial p}{\partial t} + \frac{\partial s}{\partial t} \right), \tag{2.50}$$

through which the term RHS_s can be set using the expressions for $\partial p/\partial t$, $\partial s/\partial t$ derived earlier. Hence, at the boundary, the wall temperature, T_w, satisfies

$$\mathrm{RHS}_s = \frac{1}{\gamma - 1} \frac{1}{T} \frac{\partial T_w}{\partial t} + \frac{1}{a} \sum_l \frac{\Theta_l^+ + \Theta_l^-}{2} + \frac{1}{\gamma} \sum_l \Theta_s^l. \tag{2.51}$$

To model heat transfer between the solid and the gas, the wall-normal heat flux is defined at the boundary. To do this, the heat flux $\mathbf{q}$ can be modified by subtracting the wall-normal component of $\mathbf{q}$: $\boldsymbol{\xi}_{(n)} \cdot \mathbf{q} \xi^{(n)}$, and adding back in the desired wall normal heat flux $\mathbf{q}_w = q_{(n)} \xi^n$. Hence defining

$$\mathbf{q}_{\text{new}} \coloneqq \mathbf{q} + \left(q_{(n)} - \boldsymbol{\xi}_{(n)} \cdot \mathbf{q} \right) \xi^{(n)}, \tag{2.52}$$

would automatically satisfy the required boundary conditions.

2.5.4 Sponge Layers

Although the NSCBCs are nominally non-reflecting, these act more as low-pass filters [Selle et al., 2004], and may have reflection coefficients close to unity for strong disturbances with larger wavelengths. Situations where this boundary treatment results in reflections includes unsteady wakes moving through the outlet boundary [Fosas De Pando, 2012], as well as in the case of low-frequency acoustic waves with wavelengths on the order of the domain size.

A strategy to address this is to augment the boundaries with a damping zone, known as a sponge layer [Bodony, 2006]. Our code implements the damping zone in a spatial distribution defined by non-zero positive values of $\sigma(\mathbf{x})$, in which the damping is proportional to the difference between solution vector $\mathbf{v}$ and a reference solution $\mathbf{v}_{\text{ref}}$. This term, $-\sigma(\mathbf{x})(\mathbf{v} - \mathbf{v}_{\text{ref}})$, is then applied directly to the non-linear right-hand side function of Equation (2.42).

2.5.5 Periodic Boundaries

Periodic boundaries are treated somewhat differently to those in the sections above. Though it is tempting to transfer the pseudo-waves between the points at the periodic interface, it can be shown that doing so results in dispersive errors as a consequence of the properties of the boundary spatial scheme. Similarly, approaches with periodic stencils would require the 'global' differentiation approach, as data would be needed along the entire grid line linking the two periodic boundaries.

A simpler implementation, at least in the context of this code, is to exploit the 'local' compact schemes in Section 2.6.1. This approach imposes the periodic condition by treating processors at the boundary as neighbours, so that the ghost points located across the boundary receive the correct information from the interior of their respective periodic 'neighbour'. In this way, communications between processor blocks on the interior and at the boundary are consistent, and hence the performance penalty is minimised.

2.5.6 Imposing Flow Disturbances

It is worth mentioning here that flow disturbances, such as vortices, entropy waves, or inlet turbulence may sometimes need to be specified within the computational domain. In the interior of the domain, this forcing can be imposed either as a volume source term in the governing equations or by directly applying a forcing term to the pseudo-waves.

At the inlet, however, the boundary conditions can be modified too by solving for the time-dependent incoming pseudo-waves directly, much in the same manner as in Section 2.5.1. Doing this requires a compromise between the non-reflective treatment and specifying the incoming disturbance, and so a variety of techniques to address this can be applied, such as introducing a relaxation term to control outgoing waves while still resolving the generated disturbance, for example.

Although not used for any cases in this thesis, this flexibility allows for consider-

able freedom in the form of the flow disturbances that can be introduced into the simulations.

2.6 Spatial Differentiation

2.6.1 Finite and Compact Difference Schemes

Obtaining numerical solutions of Equation (2.42) on discretised domains introduces the need to approximate the derivative terms $\partial f/\partial \xi^n$ of functions f defined on the nodes of the structured grid, with the accuracy of this discretisation determined by the asymptotic characteristics of the truncation error, and the effect of the scheme on the dispersion relation of a model ODE (see e.g. the discussion in the introduction of Adams and Shariff [1996]). A popular procedure for derivative approximations on structured grids is the finite difference method, and two main categories of these exist: explicit schemes and implicit, both of which we will mention here.

A general finite difference approximation can be written as

$$\sum_{k\in\mathbb{Z}} a_{i,i+k} f'_{i,i+k} = \sum_{k\in\mathbb{Z}} b_{i,i+k} f_{i,i+k}, \tag{2.53}$$

or equivalently

$$\mathbf{A}\mathbf{f}' = \mathbf{B}\mathbf{f}. \tag{2.54}$$

Schemes with $a_{i,i+k} = 0$, $\forall k \neq 0$ are explicit, with the implicit schemes having non-zero entries outside of the main diagonal of the matrix $\mathbf{A}$, thus requiring a solution of the matrix equation in Equation (2.54).

Though the explicit schemes of any order n are easy to derive through simple Taylor expansions of a polynomial fitted to the data points, high-order schemes require stencils of sizes at least $n + 1$ as well as the corresponding stable boundary closures with 'desirable' numerical characteristics. Moreover, it is well known that these schemes may require much finer meshes to represent the dispersion relation than the corresponding waveform itself, thereby increasing the computational cost.

In contrast, implicit schemes can have much smaller stencils, and for this reason are known as *compact* schemes. For a given order of approximation, the extra coefficients that appear in the stencils of both **A** and **B** can then be used to optimise the resolution characteristics of the scheme. Stencil weights can also be made asymmetric to create *upwind schemes*, with controlled numerical dissipation at the high wavenumbers added to mitigate stability issues in advection problems.

Used here are the fifth-order Compact Upwind Low-Dispersion schemes (CULD5) of [Adams and Shariff, 1996] for the derivatives present in the pseudo-waves, while a centred compact non-dissipative third-order scheme (CC3) of [Lele, 1992] with spectral-like resolution is adopted for the diffusive terms, as the numerical stability of the diffusion terms is determined solely by the timestep restriction[5]. By applying the Fourier component $f(x,t) = \hat{f}(t)\exp(i\alpha x)$ to a model advection or diffusion equations, we determine the effects of the spatial scheme of the resulting solution through the modified wavenumber $\tilde{\alpha}$. For this scheme, the dissipation and dispersion characteristics of the interior and boundary schemes are shown in Figure 2.2.

The effect of the schemes on the modified wavenumber has implications for grid generation too. For a harmonic disturbance with spatial wavenumber α, and a grid point spacing of Δx, the number of grid points per wavelength is defined as

$$n_{\text{ppw}} = \left\lfloor \frac{2\pi}{\alpha \Delta x} \right\rfloor. \tag{2.55}$$

Rearranging Equation (2.55) for the wavenumber and comparing it to the complex modified wavenumber shown in Figure 2.2, we can determine the error introduced in the representation of the original wave by the differentiation scheme. These are summarised in Table 2.1.

[5]This is because the diffusive terms are dissipative, and so any high-wavenumber numerical noise that is introduced is dissipated by the flow physics without amplification. The time-step restriction is discussed further in Section 2.9.1.

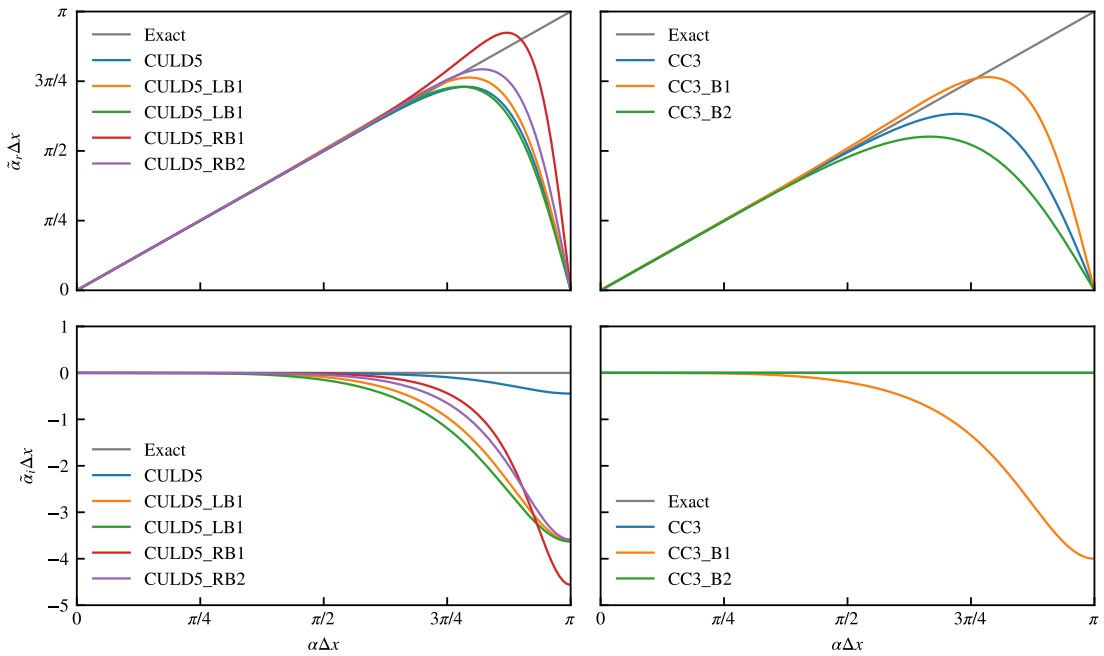


Figure 2.2: The dispersive and dissipative characteristics of the CULD5 and CC3 schemes, showing artificial dissipation at short wavelengths for the upwinding scheme, with no dissipation present in the CC3 scheme. The properties of the boundary closures (denoted by LB*, RB* and B*) do not affect the solution significantly. These are instead tuned to maintain the order of accuracy of the global scheme, and to ensure a numerically stable global solution [Adams and Shariff, 1996].

n_{ppw}	CULD5		CC3	
	$\lvert\tilde{\alpha}_r - \alpha\rvert/\alpha$	$\tilde{\alpha}_i$	$\lvert\tilde{\alpha}_r - \alpha\rvert/\alpha$	$\tilde{\alpha}_i$
2	1.00×10^{-0}	-4.46×10^{-1}	1.00×10^{-0}	0
4	3.20×10^{-4}	-5.64×10^{-3}	9.70×10^{-3}	0
8	1.56×10^{-5}	-7.00×10^{-5}	1.20×10^{-4}	0
12	1.45×10^{-6}	-5.89×10^{-6}	1.01×10^{-5}	0
16	2.64×10^{-7}	-1.03×10^{-6}	1.78×10^{-6}	0
24	2.34×10^{-8}	-8.96×10^{-8}	1.55×10^{-7}	0

Table 2.1: The errors introduced in the modified wavenumbers as a function of the number of grid points used to represent the wave, n_{ppw}, for the CULD5 and CC3 schemes (replicated from [Fosas De Pando, 2012]).

2.6.2 Global and Local Compact Differentiation

The equations of the compact difference schemes, which we introduced in the previous section, impart a global spatial coupling of all the points in the field due to the 'inversion' of the coefficient matrix that is required to obtain the derivative terms, which we denote from now on by the operator

$$\mathbf{D}_k = \mathbf{A}_k^{-1}\mathbf{B}_k, \tag{2.56}$$

with the subscript k representing the differentiation scheme being used.

The challenge that this presents in the context of a parallel architecture is that the data is distributed over many processes. At the same time, the solution to the implicit linear equation requires access to the entire data field, and this apparent dichotomy must be reconciled so that the parallel efficiency of the full scheme is preserved. Two approaches exist to address this.

The first of these is what we will define as the *global* approach. This method computes the full compact difference scheme as defined in Equation (2.54), and is summarised in Algorithm 1. We observe that in line 7 and line 9 of this algorithm

Algorithm 1 Global Compact Schemes

1: Let $p_{\text{dim}} = [p_3, p_2, p_1]$ be the number of processors in each dimension
2: Let d be the direction in which the derivative is taken
3: **for** axis in $[z, y, x]$ **do**
4: $d =$ axis
5: Let p_{ind} be the processor index in direction d
6: Color data on each processor along d into $p_{\text{dim}}[d]$ equal, disjoint partitions Π_i
7: `MPI_SendRecv` on each processor p_{ind} assemble $\Pi_{p_{\text{ind}}}$ into a single vector
8: Perform differentiation
9: `MPI_SendRecv` to reassemble the local data back onto each processor

that significant amounts of data must be transmitted between the processors for each direction in which the derivative is taken. Though for small problems this may not be a problem, in cases with larger state vectors increasing the number of parallel units participating in the computation quickly saturates the parallel communications

buses that connect the independent nodes in a distributed memory system. As such, the gain in parallel efficiency here is minimal for the types of problems that we wish to consider.

An alternative variant arises from the observation that the derivative is an essentially *local* property, by which we mean that, despite the global coupling in Equation (2.54), perturbations to an arbitrary single point $\mathbf{x}_p$ of $\mathbf{v}$ would predominantly result in changes of $\mathbf{v}'$ for points $\mathbf{x} \in B_r(\mathbf{x}_p)$, $r > 0$, where $B_r(\mathbf{x}_p)$ is a ball centred on $\mathbf{x}_p$ of radius r with periodicity if it intersects a periodic boundary. For points $x \in \mathcal{C} \setminus B_r(\mathbf{x}_p)$, the perturbed and unperturbed $\mathbf{f}'$ coincide up to machine precision.

For a given processor we construct an augmented local sub-field that consists of a halo of ghost points surrounding the local data, as shown in Figure 2.3. In this thesis the thickness of this ghost layer is 16 points thick as this was found to be a good compromise between the accuracy of the derivative and the computational efficiency. We will make further comments on this choice in the validation section, Section 2.10.2.

Differentiation is then performed using Algorithm 2, where the memory-intensive assembly and distribution of the data from the global schemes has been replaced by asynchronous communications that populate only the ghost cells prior to differentiation. Once both sets[6] of ghost cells in a particular direction are populated, the calculations of the derivatives are commenced and the flow of control passes to the next axis. The observant reader may notice that the differentiation is performed

Algorithm 2 Local Compact Schemes

1: Let d be the direction in which the derivative is to be taken
2: **for** axis in $[z, y, x]$ **do**
3: $d =$ axis
4: Asynchronous `MPI_SendRecv` between neighbours to ghost cells on the 'left'
5: Asynchronous `MPI_SendRecv` between neighbours to ghost cells on the 'right'
6: Perform differentiation in direction d

[6]Unless the ghost cells are at a non-periodic boundary, they contain two regions of ghost cells along each dimension.

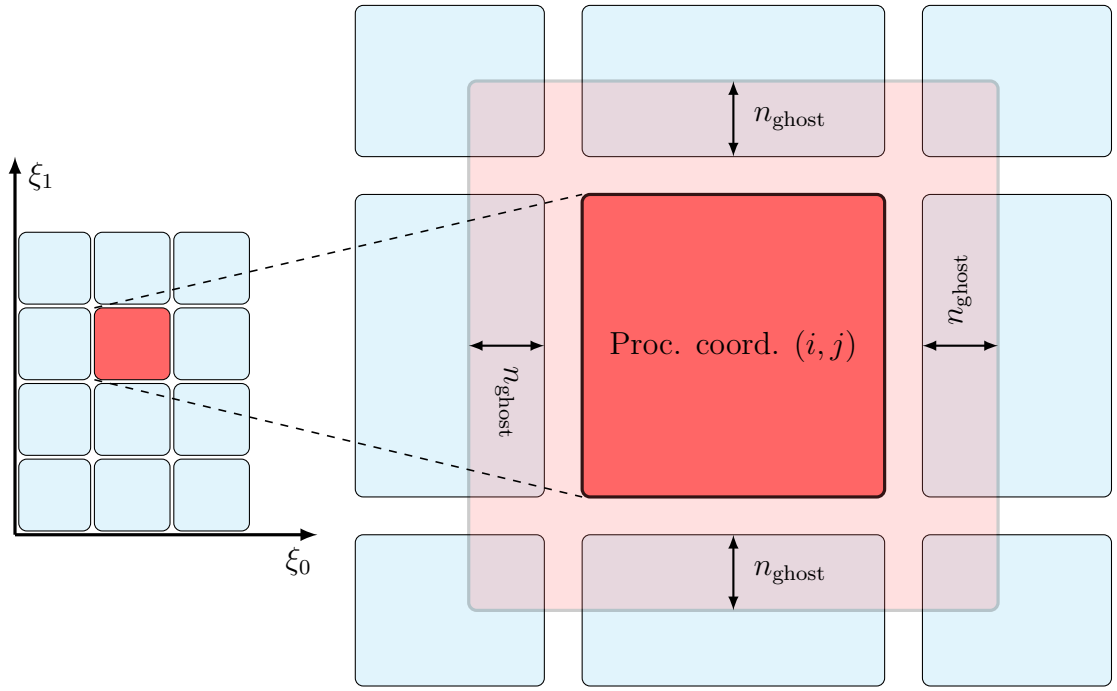


Figure 2.3: Configuration of the processors in the computational space, and the ghost cells around a given processor in the *local compact* difference schemes. The entire pale red region is stored on the processor (i, j). Once the data is assembled in the ghost cells, with communications between the neighbouring processors, the derivatives are applied to the full region. Then the red region is extracted to form the distributed global vector. These schemes allow for high levels of parallel efficiency and significantly simplify the arrangement and storage of the relevant data in memory.

in the order of the z, y and x axes. This is done so that the sliding plane interpolation (see Chapter 4) may commence as soon as the required data is available, without waiting for other communications to complete. Notice also, that unlike the global scheme, this does not require the data to be reassembled, since all the data already exists on the processor where it is required. This significantly minimises the communications overhead and improves parallel efficiency, since, assuming a fixed number of grid points per processor, the amount of data transferred to the ghost cells remains constant, regardless of the global domain size.

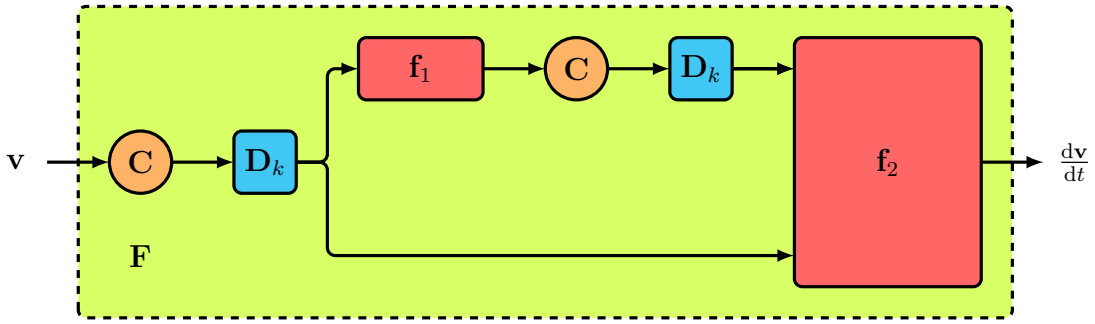


Figure 2.4: The nonlinear right-hand side function $\mathbf{F}$, of Equation (2.42), is decomposed into a calculation of the inviscid, $\mathbf{f}_2$, and the dissipation, $\mathbf{f}_1$, terms. Operations within each block are decoupled in space by considering the derivatives as independent input variables.

2.6.3 Block Structure of Nonlinear Code

The right-hand side of Equations (2.35)–(2.37) is implemented as a single vector-valued function interface taking as inputs the time signal, `t`; solution vector, `v`; and the spatial derivatives, $d\mathbf{v}$, in the form

$$\mathbf{f} = \mathbf{f}(t, \mathbf{v}, d\mathbf{v}). \tag{2.57}$$

We then proceed by decomposing the function $\mathbf{F}$ into two further sub-functions, $\mathbf{f}_1$ and $\mathbf{f}_2$, that compute the inviscid and viscous parts of Equations (2.35)–(2.37) respectively. This is shown diagrammatically in Figure 2.4, with the inputs and outputs of the sub-functions shown in Figure 2.5. The matrix $\mathbf{C}$ is the linear operation responsible for the communications which populate the ghost points prior to the differentiation, represented here by $\mathbf{D}_k$.

2.7 Efficient Linearisation and Direct and Adjoint Dynamics

With the spatial coupling of the equations, the naive approach to linearising the governing equations, in Equation (2.42), would apply individual perturbations sequentially to every degree of freedom. Of course this incurs significant computational

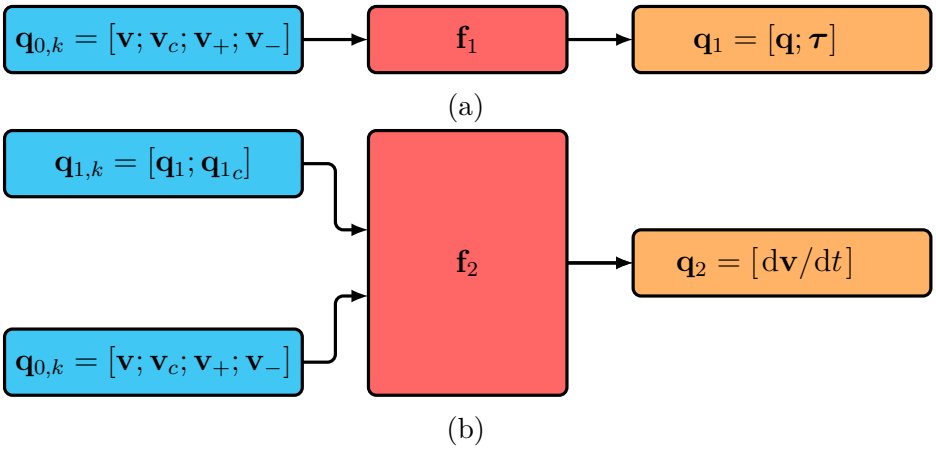


Figure 2.5: The inputs and outputs of the sub-functions (a) $\mathbf{f}_1$ and (b) $\mathbf{f}_2$. The vectors $\mathbf{q}_{i,k}$ are the vectors formed by taking the three derivatives using the central and the two upwinding directions. Vectors of the form $[\mathbf{a}; \mathbf{b}]$ denote a vertical stack of all the elements of $\mathbf{a}$ followed by all the elements of $\mathbf{b}$.

costs on the order of $4(N+1)$ nonlinear function evaluations for a system with N degrees of freedom.

An alternative method, derived in [Fosas de Pando et al., 2012], observes that the spatial coupling in the governing equations is introduced through the spatial differentiation operators, $\mathbf{D}_k$. The solution, then, is to break down the Navier-Stokes equations into sub-functions that depend on a set of variables $\mathbf{q}$ and their derivatives $\mathbf{q}_+$, $\mathbf{q}_-$, and $\mathbf{q}_c$, where $\star_\pm$ and $\star_c$ denote the positive and negative upwinding and the central derivatives respectively. Now, the value of each these sub-functions is independent at each spatial point, and hence the linearisation of that particular function may be performed simultaneously at all points within the domain.

For each sub-function $\mathbf{f}_i = \mathbf{f}_i(\mathbf{x}, \mathbf{q}_k)$, where $k \in \{0, +, -, c\}$ represents the corresponding first derivative of $\mathbf{q}$, we compute the finite-difference approximation

$$\mathbf{A}_{ijk} = \frac{\mathbf{f}_i(\mathbf{q}_{j,k} + \epsilon \mathbf{1}_j) - \mathbf{f}_i(\mathbf{q}_{j,k})}{\epsilon}, \tag{2.58}$$

with ϵ selected to minimise the cancellation error. Alternative methods such as complex-step derivatives, where the perturbation is applied in the from $\mathbf{q}_{\mathrm{bf}} + \epsilon i \tilde{\mathbf{q}}$ and has an $\mathcal{O}(\epsilon^2)$ cancellation error, or the even introducing a field of hyper-dual numbers [Fike and Alonso, 2011] where there is no error in the first and second derivative

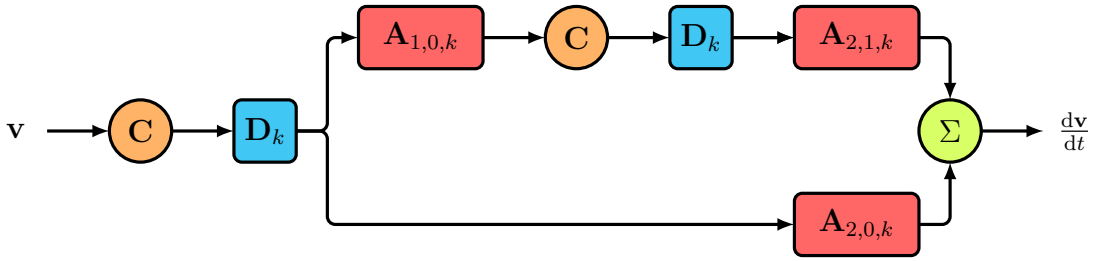


Figure 2.6: The linearisation procedure used to obtain the linear operator, with $A_{\star,\star,\star}$ matrices corresponding to the linearisation of code sub-blocks, D_k referring to the matrix-free differentiation operation and the C operator denoting the communications for the global/local compact schemes.

calculations, are superior compared to the finite difference method. However, the adaptation of existing codes to use these methods often requires significant modifications of the existing codebase. It was decided, therefore, that we will proceed here with the finite difference method in the knowledge that we can make use of the existing code, and, critically, that we are mindful of, and are able to control, the linearisation error so that it is of the same order as the square root of the machine round-off error, $\sqrt{\epsilon_{\text{machine}}} \approx 4 \times 10^{-8}$.

The assembly of the linear operator is shown in Figure 2.6, with the algebraic form given by Equation (2.59),

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \left(\mathbf{A}_{2,0,k} + \mathbf{A}_{2,1,k}\mathbf{D}_k\mathbf{C}\mathbf{A}_{1,0,k}\right)\mathbf{D}_k\mathbf{C}\mathbf{v}, \tag{2.59}$$

and with $\mathbf{C}$ representing the communications in the local compact schemes.

The adjoint operator $\mathbf{A}^H$ follows trivially and the adjoint right-hand side can be written as follows

$$\frac{\mathrm{d}\mathbf{w}}{\mathrm{d}t} = \mathbf{C}^H\mathbf{D}_k^H\left(\mathbf{A}_{2,0,k}^H + \mathbf{A}_{1,0,k}^H\mathbf{C}^H\mathbf{D}_k^H\mathbf{A}_{2,1,k}^H\right)\mathbf{w}. \tag{2.60}$$

This procedure is shown in Figure 2.7.

Having shown above how the forward and adjoint operators are constructed, we should emphasise that the linear operator is never assembled directly, and, as with the nonlinear code, we perform all calculations with the linear operators in a matrix-free way. This is done by storing the $\mathbf{A}_{i,j,k}$, which are diagonal, as sparse

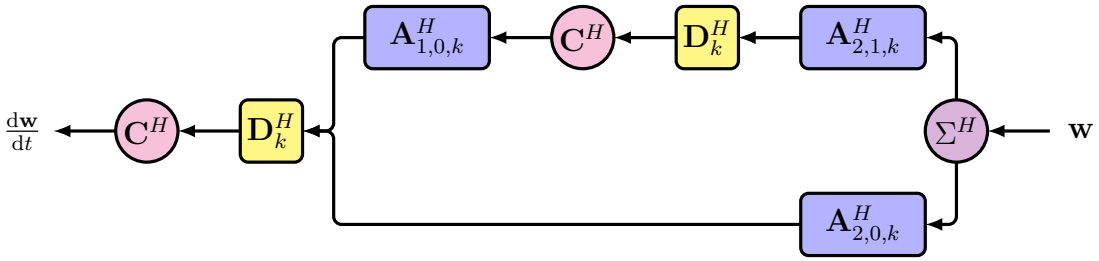


Figure 2.7: The assembly of the adjoint operator simply reverses the order in which the forward linearisation is performed, substituting the corresponding adjoint code, $\star^H$, for each linear operation. Note that one must take into account the grid metric when computing the adjoint.

matrices and simply performing point-wise vector multiplications. The derivatives, $\mathbf{D}_k$, by virtue of being linear, are calculated using the same matrix-free schemes used in the nonlinear code. Finally, the communication matrices, $\mathbf{C}$, are also linear and reuse the existing communications used in the nonlinear code.

The adjoint operator takes the complex conjugate entries of $\mathbf{A}_{i,j,k}$ and transpose conjugates the spatial derivative operators, $\mathbf{D}_k$, which amount to no more than a reordering of the addressing indices.

2.8 Krylov Subspace Methods

For large-scale problems one of the key issues is the dimensionality of the linear system. One way around this is to introduce the *Krylov subspace of size m of the linear operator $\mathbf{A}$ and vector $\mathbf{v}$*, defined as

$$\mathcal{K}_m(\mathbf{A}, \mathbf{v}) = \operatorname{span}\left\{\mathbf{v}, \mathbf{A}\mathbf{v}, \mathbf{A}^2\mathbf{v}, \ldots, \mathbf{A}^{m-1}\mathbf{v}\right\}. \tag{2.61}$$

The orthogonal basis of this subspace is constructed using the Arnoldi algorithm with modified Gram-Schmidt reorthogonalisation. This is preferred to the direct application of Equation (2.61) as not only is the resulting basis orthogonal, but it also avoids the errors that would be encountered when evaluating the matrix-vector products $\mathbf{A}^i\mathbf{v}$ with finite-precision arithmetic. The result of this is that we can write

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m\mathbf{H}_m + h_{m+1,m}\mathbf{v}_{m+1}\mathbf{e}_m^H, \tag{2.62}$$

with $\mathbf{V}_m$ denoting the m-dimensional orthogonal basis of the Krylov subspace $\mathcal{K}_m(\mathbf{A}, \mathbf{v})$ and $\mathbf{H}_m \in \mathbb{R}^{m \times m}$ is upper Hessenberg. Error estimates may be determined as functions of $\mathbf{v}_{m+1}$. Function evaluations of the linear operator $\mathbf{A}$ are subsequently replaced by analogous function evaluations of the projected and reduced operator $\mathbf{H}$.

2.9 Time Integration

Having transformed the nonlinear equations in Equation (2.42), and by extension the corresponding linearised form in Equations (2.59)–(2.60), to a set of ODEs, these now need to be integrated in time to solve for the state vector $\mathbf{v}$. This is achieved using Runge-Kutta methods, predominantly for the nonlinear equations, and the exponential time stepping method in the case of linear ODEs[7]. These are introduced briefly in the following sections.

2.9.1 Runge-Kutta 4th Order

For the nonlinear equations the popular Runge-Kutta family of time integrators is used. The nonlinear solver and the checkpointing algorithms both use a low-storage variant of the Runge-Kutta 4 algorithm `RK4(3)5[2R+]C` [Kennedy et al., 2000], with Butcher Tableau and its values provided in Tables D.5–D.6. This scheme is characterised as low-storage since at any given stage only two state vectors are required to store the intermediate values. In cases where there are many degrees of freedom, the savings to system memory are no longer inconsequential.

Another attractive feature of this scheme is its large stability region. Stability regions of time integration schemes are determined by the curve in λ-space where

[7]Note that *adjoint* Runge-Kutta methods are used for the nonlinear-linear adjoint looping in Chapter 4.

the linear equation

$$\frac{\mathrm{d}v}{\mathrm{d}t} = \lambda v, \tag{2.63}$$

integrated using that particular time integration scheme with a time step Δt, is bounded for all time. This occurs when the recursion

$$v_{k+1} = (1 + p(\lambda \Delta t))v_k, \tag{2.64}$$

is bounded, and is equivalent to finding the region where

$$|1 + p(\lambda \Delta t)| \leq 1. \tag{2.65}$$

Here p is the characteristic polynomial for the time stepping scheme.

It is well known (e.g. [Ferziger and Perić, 2002]) that for the temporal stability of an advection-diffusion-type equation the spectrum of the spatially discredited equation must fall within the temporal stability region of the time integration scheme. Using this, the restrictions on the time step Δt for a given minimum Δx in the mesh, and maximum wave propagation speed U, may be written as

$$\frac{U\Delta t}{\Delta x} \leq \frac{3.38}{2.29}, \quad \text{and} \quad \frac{\Delta t}{\mathrm{Re}\,\Delta x^2} \leq \frac{4.82}{3.96}, \tag{2.66}$$

which determines the time step restrictions of the advection and viscous terms respectively. The former inequality in Equation (2.66) is more restrictive than the viscous bound, and is the global timestep selected for the time integration. Indeed in the cases presented later on in this thesis, the convective timestep is $\mathcal{O}(10^{-5})$, while the diffusive condition is $\mathcal{O}(10^{-4})$. Furthermore, the point at which the viscous time step becomes more restrictive occurs approximately when $\mathrm{Re} \approx (\Delta x)^{-1}$, and this is not the case here. The selection of the explicit time integration scheme is therefore justified in this case, although once Reynolds numbers increase and/or when one desires greater resolution in the thin viscous boundary layer, the stiffness of the resulting system may make implicit integration more desirable.

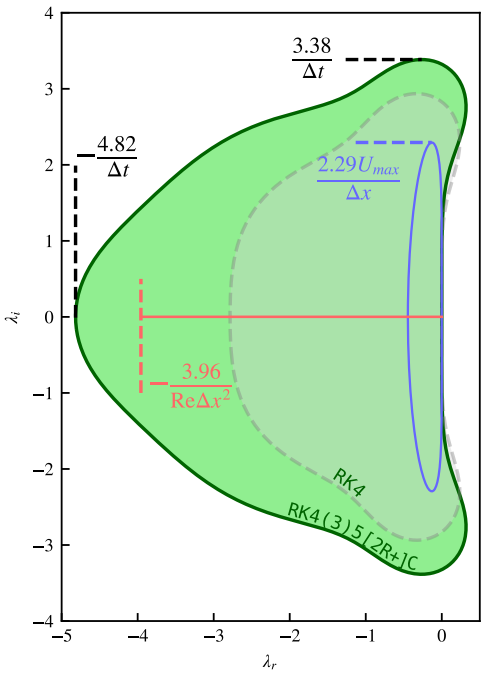


Figure 2.8: Temporal stability region of the time integration scheme superimposed with the spectra of the differentiation schemes (modified from [Fosas De Pando, 2012]). The classical Runge-Kutta 4 stability region is also included for comparison.

2.9.2 Exponential Time Integration

For linear autonomous[8] initial value problems (IVPs) of the type

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \mathbf{A}\mathbf{v} + \mathbf{f}, \quad \text{with} \quad \mathbf{v}(0) = \mathbf{v}_0, \tag{2.67}$$

the exact solution is known exactly, and is, for $t \geq 0$

$$\mathbf{v}(t) = e^{t\mathbf{A}}\mathbf{v}_0 + \int_0^t e^{(t-s)\mathbf{A}}\mathbf{f}(s)\,\mathrm{d}s. \tag{2.68}$$

The task of solving the IVPs in Equation (2.67) therefore condenses to the approximation of the matrix exponential term $e^{t\mathbf{A}}$ in Equation (2.68), where the matrix exponential is defined here as the usual infinite series

$$e^{t\mathbf{A}}\mathbf{v} = \sum_{n=0}^{\infty} \frac{1}{n!}(t\mathbf{A})^n\mathbf{v}, \tag{2.69}$$

and $\mathbf{A}$ is a bounded linear operator. For linear equations, methods such as the Runge-Kutta scheme in Section 2.9.1 essentially approximate the matrix exponential

[8]Where the governing linear matrix $\mathbf{A}$ does not explicitly depend on time.

with a truncation of the infinite series in Equation (2.69), though forming this truncated series is considered undesirable [Moler and Van Loan, 2003].

An alternative way is to use a degree $(m-1)$ polynomial p_{m-1} to approximate the matrix exponential [Schulze et al., 2009],

$$e^{t\mathbf{A}}\mathbf{v} \approx p_{m-1}(t\mathbf{A})\mathbf{v}, \tag{2.70}$$

which is an element of the m-dimensional Krylov subspace defined in Equation (2.61), and, by construction, there exists an optimal element of this subspace that minimises the error in the approximation in Equation (2.70). The Arnoldi algorithm is used to construct an orthonormal basis of $\mathcal{K}_m(\mathbf{A}, \mathbf{v})$ alongside the upper Hessenberg matrix $\mathbf{H}_m \in \mathbb{R}^{m\times m}$. With some linear algebra [Schulze et al., 2009], the original problem of approximating $e^{t\mathbf{A}}\mathbf{v}$, with $\mathbf{A} \in \mathbb{R}^{n\times n}$, is exchanged for an evaluation of $e^{t\mathbf{H}_m}\mathbf{v}$ with $m \ll n$. This matrix exponential is then easily computed using Padé methods [Schulze et al., 2009; Moler and Van Loan, 2003].

The time integration is cast as a time-stepping problem with discrete time intervals satisfying $t_{n+1} = t_n + \tau$. The Krylov subspace dimension m is fixed for the duration of the linear simulation, and the error estimate for the approximation in Equation (2.70), derived from the Arnoldi procedure, is used to set τ so that this error is below some pre-determined threshold. Our applications set $m = 16 - 32$, allowing for an equivalent CFL number of approximately 10, although there are some variations depending on the particular case being considered.

2.10 Validation Cases

2.10.1 Summary of Validation Cases

We now conclude this chapter by providing two test cases that validate the operation of the nonlinear portion of our code, and evaluate its suitability for aeroacoustics applications.

The first of these consists of two-dimensional co-rotating equal-strength vortices that radiate acoustic waves into the far-field, which evaluates the inviscid portion of the code, which is reproduced here courtesy of the author of [Fosas De Pando, 2012]. The second is a three-dimensional turbulent energy dissipation problem validating the viscous dissipation terms. Though the three-dimensional problem will not be considered in this thesis, we nonetheless perform this validation to confirm that the code is indeed capable of performing three-dimensional DNS/DNC. This section then concludes with a check of the generation and evaluation of the direct and adjoint linear operators obtained from the linearisation component of our code.

In all cases the validations are performed in a single-block Cartesian geometry due to the readily available analytic two-dimensional solutions, simpler grid geometries with faster turn-around time that are unencumbered by the stiffness issues resulting from viscous boundary layers, and ultimately do not require additional geometries to be developed in order to model more complex test cases.

2.10.2 Co-rotating Vortices

The first case presented here is rerun from earlier validations performed by Fosas De Pando [2012] in the preceding project, in which the core functionality of this code was developed.

Pairs of co-rotating vortices with equal strength are well-known acoustic sources, having been extensively studied in free shear flows [Mitchell et al., 1995], with the case particularly well-suited for aeroacoustic validation problems because it captures a quadrupole acoustic source that is significantly weaker than monopoles and dipoles in subsonic flow settings [Kim, 2013]. Accurate prediction of the far-field noise in this case also relies on capturing both the hydrodynamics of the two vortices and the resulting propagation of the acoustic field.

Using the case from [Kim, 2013] and the far-field model of [Mitchell et al., 1995], and setting $M = 1$, $\gamma = 1.4$, $\mathrm{Re} = \infty$, $\mathrm{Pr} = 1$, circulation $\Gamma_0 = \pi^2/10$ and rotation

frequency $\Omega_0 = \pi\Gamma M/4$, the initial conditions, as from [Kim, 2013], read as follows,

$$p = \left[1 - \frac{\gamma - 1}{2}\left(u^2 + v^2\right)\right]^{\gamma/(\gamma-1)}, \tag{2.71}$$

$$(u, v) = \sum_{m=0}^{1} \frac{(1+\delta)\Gamma_0}{2\pi(r_m^2 + r_c^2)}(y, -x_m), \tag{2.72}$$

with

$$r_m^2 = x_m^2 + y^2, \quad x_m = x + (-1)^m, \quad r_c = 10^{-2r_m^2}\sigma. \tag{2.73}$$

In the above, x_m is the position of the two vortices, r_m is the radial distance from each vortex, and r_c is a correction constant that removes the point-wise singularity arising within the vortex core. Further, the factor $\delta = 0.059395$ is chosen to correct the rotation speed to match the ideal potential flow solution, with $\delta = r_c = 0$ and point vortices.

In polar coordinates (r, θ), the analytic solution for the latter case derived in [Mitchell et al., 1995] describes the induced pressure as

$$p_{\text{analytic}}(r, \theta, t) = -\frac{\Gamma_0^4}{64\pi^3 L^4}\left[J_2(2\Omega_0 r)\sin(2\theta + 2\Omega_0 t)\right. \tag{2.74}$$

$$\left. + Y_2(2\Omega_0 r)\cos(2\theta + 2\Omega_0 t)\right], \tag{2.75}$$

with $J_2(z)$ and $Y_2(z)$ denoting second-order Bessel functions of the first and second kinds, respectively.

The grids used for this case are

$$x(\xi, \eta) = x_{\text{rect.}} + A_x \sin(\alpha\xi)\sin(\beta\eta), \tag{2.76}$$

$$y(\xi, \eta) = y_{\text{rect.}} + A_y \sin(\alpha\xi)\sin(\beta\eta), \tag{2.77}$$

with a rectilinear grid $\mathbf{x}_{\text{rect.}}$, which creates a smooth grid with increased grid density near the centre and an inflated grid in the far-field. This intermediate grid is then augmented with a sinusoidal perturbation to break grid orthogonality, with $A_x = A_y = 0.375$ and $\alpha = \beta = 16\pi$. Overall dimensions of the physical computational domain are $L_x = L_y = 1392.34$, with $n_x = n_y = 720$ points in each direction.

A sponge layer with strength $\sigma = 0.5$ further reduces the reflections from the boundaries.

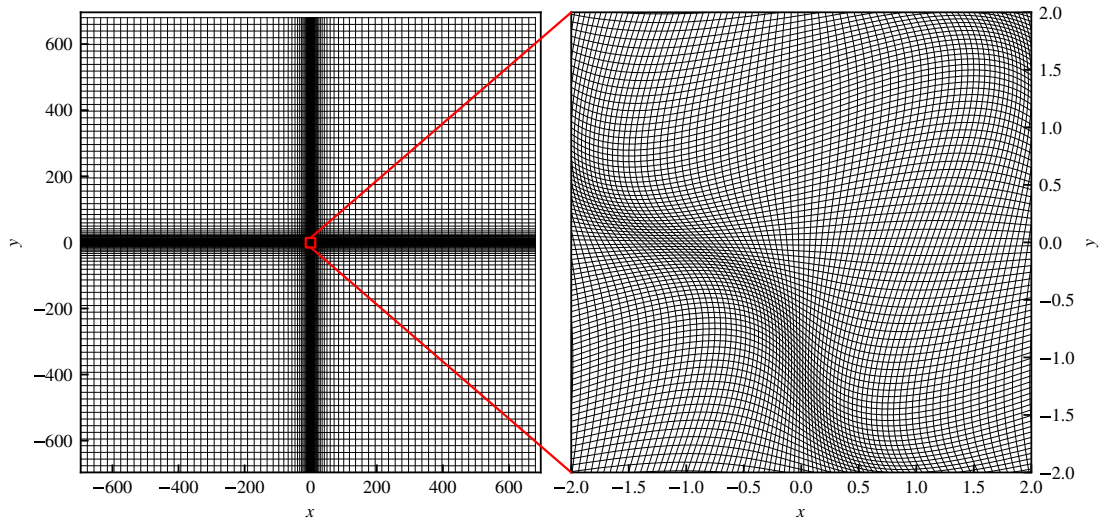
Comparative snapshots of the analytic and numerical solutions at $t = 600$ are shown in Figure 2.9b, with pressure traces given in Figure 2.9, and show excellent agreement in the far-field.

Making use of this opportunity, we can also examine the effects of varying the number of ghost points in the locally compact schemes. To arrive at stable solutions, and minimise computational overheads we place the restriction $8 \leq n_{\text{ghost}} \leq 32$. For the purposes of this exercise, we select $n_{\text{ghost}} \in \{8, 16, 32\}$, and $A_x = A_y = 0$, so that the grid is uniform, and reproduce the pressure traces from Figure 2.9 for each of these with the results in Figure 2.10. Although the relative error is minimised for $n_{\text{ghost}} = 32$, throughout this thesis we use $n_{\text{ghost}} = 16$ as this gives a reasonably low error for a significantly lower computational cost[9].

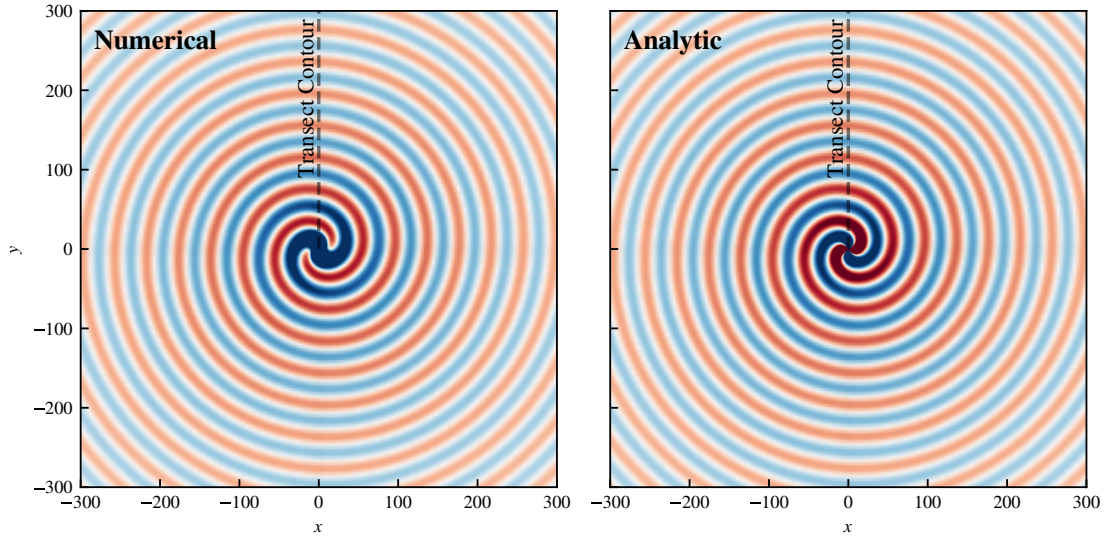
2.10.3 Taylor-Green Vortices

To validate the viscous terms we break with the two-dimensional nature of this thesis, and evaluate a three-dimensional DNS Taylor-Green vortex dissipation problem. This serves two purposes. The first of these is that this case is extremely sensitive to implementation errors, and with a comparison to analytic models of turbulent energy decay, one readily confirms that the viscous terms are indeed implemented correctly. The second more practical reason is that it serves to show how this numerical code is readily adapted to consider the effects of turbulence on the results in the subsequent chapters, despite us not performing these here due to limitations in our computer resources.

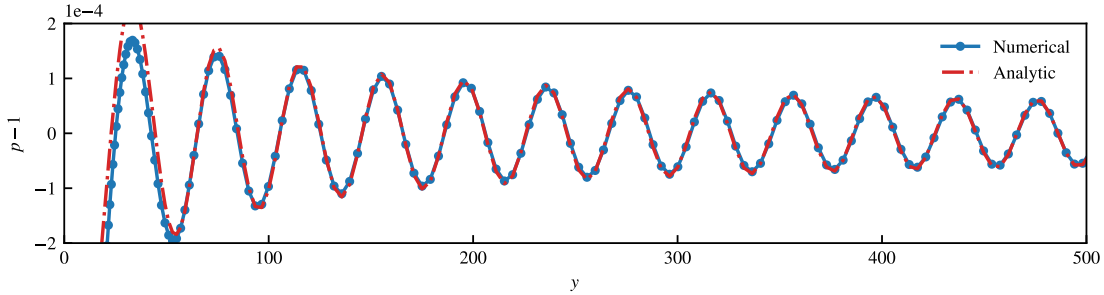
[9]With $n_x = 612$, $n_{\text{ghost}} = 32$ and a distribution over 36 processors, more than one-third of the points in the domain are located in the ghost-cell regions. With $n_{\text{ghost}} = 16$ this falls to 19%.



(a)



(b)



(c)

Figure 2.9: The grids, (a), used in the validation, (b) comparative snapshots of the analytic and numerical solutions at $t = 600$, with the pressure trace along $x = 0, y > 0$ shown in (c).

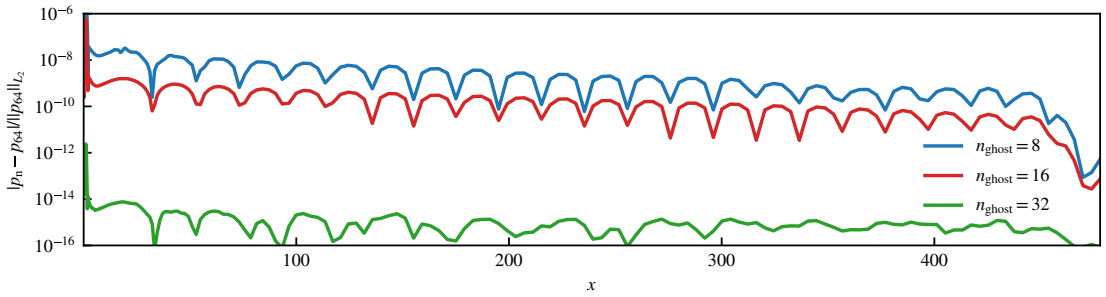


Figure 2.10: Relative pressure error along $y = 0, x > 0$ with varying n_{ghost}, compared to a reference solution with $n_{\text{ghost}} = 64$. Note that the domain size here has been reduced to $n_i = 612$, with the total length of the domain $L_x = L_y = 477.95$.

Taylor-Green vortices are given by the initial conditions

$$p = 1 + \frac{\gamma M^2}{16}\cos(2x) + \cos(2y)\cos(2x) + 2, \quad s = 0, \tag{2.78}$$

$$u = \sin(x)\cos(y)\cos(z), \quad v = -\cos(x)\sin(y)\cos(z), \quad w = 0, \tag{2.79}$$

within a rectilinear cubic domain with $x, y, z \in [-\pi, \pi]$, and grid dimensions $n_{x,y,z} = 128$, and $n_{x,y,z} = 256$ for the two cases considered here. The flow conditions are set to Re = 1600, M = 0.1, Pr = 1.

Kinetic energy, E_k, and kinetic energy dissipation rates, $\epsilon(\zeta)$, are calculated as follows

$$E_k = \frac{1}{2V}\sum_{\mathcal{P}}\left[u^2 + v^2 + w^2\right](\Delta x)^3, \tag{2.80}$$

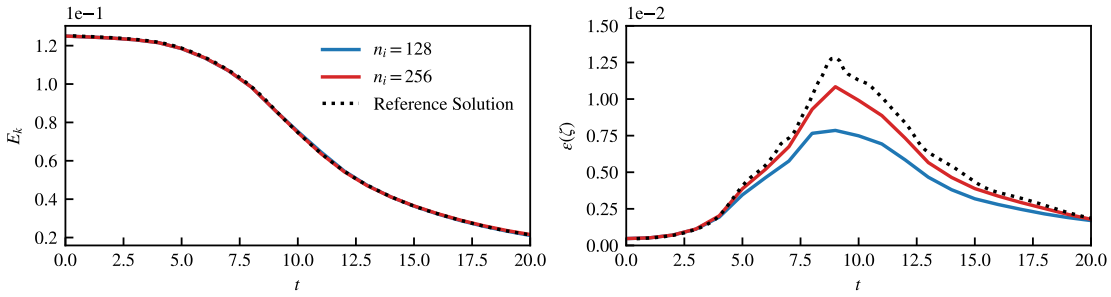
$$\epsilon(\zeta) = 2\mu\zeta, \tag{2.81}$$

$$\zeta = \frac{1}{V}\sum_{\mathcal{P}}\left[\omega_x^2 + \omega_y^2 + \omega_z^2\right](\Delta x)^3, \tag{2.82}$$

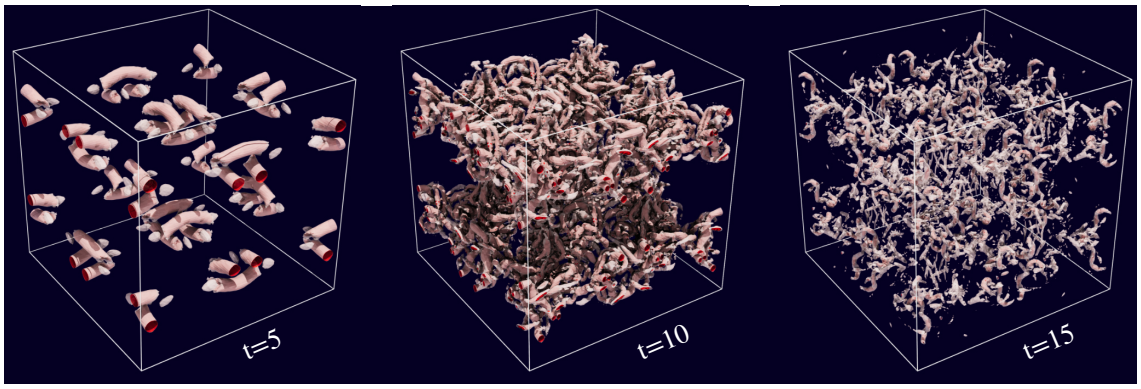
with ζ denoting the integrated enstrophy[10] over the physical domain, ω_i the vorticity and $V = 8\pi^3$ the total volume of the domain. The evolution of both E_k and ϵ over twenty time units is compared against the reference solution calculated in [DeBonis, 2013] in Figure 2.11.

In both cases the kinetic energy of the system matches very well at all times, while the integrated enstrophy shows a deviation consistent with an under-resolved solution, and subsequent improvements as the grid size is increased.

[10]The flow here is largely incompressible, and so enstrophy is defined using the vorticity.



(a)



(b)

Figure 2.11: The kinetic energy E_k and kinetic energy dissipation rate, $\epsilon(\zeta)$, computed from the enstrophy, ζ, in (a). Turbulent decay of the vortex structures at $t = 5, 10, 15$, showing initial roll-ups, the turbulent solution and decay, respectively, in (b).

2.10.4 Linearisation, Direct and Adjoint

As the core concept of this thesis, we conclude this chapter by making a brief mention of the methodologies used to verify the linearisation alongside the direct and adjoint operators.

Tempting as it is to present a simple series of examples alongside their linearised counterparts, we feel that this would somewhat obscure the main message which is that *the linearisation must be checked for each individual case that is considered.* As a consequence of the decision to maintain a quasi-linearisation approach, instead of using complex-step linearisation or hyper-dual number fields, we are forced to optimise the linearisation ϵ for each $\mathbf{A}_{i,j,k}$ block that we compute, and the optimal epsilon varies on a case-by-case basis. After the linear operator is computed, it is

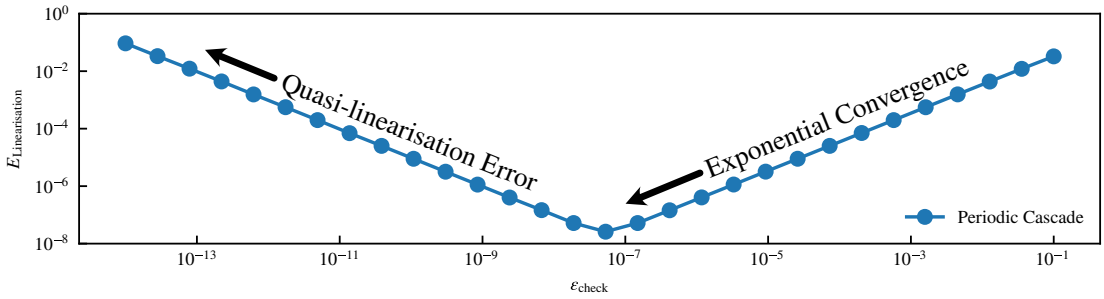


Figure 2.12: Quasi-linearisation error for the linear cascade problem of Section 3.2.4.

verified directly by comparing the relative residual

$$E_{\text{Linearisation}} = \|\mathbf{A}\mathbf{v}' - \frac{\mathbf{F}(\mathbf{v} + \epsilon_{\text{check}}\mathbf{v}') - \mathbf{F}(\mathbf{v})}{\epsilon_{\text{check}}}\|_2, \tag{2.83}$$

with the linear operator $\mathbf{A}$ being the fully assembled right-hand-side function that is used in the direct time-stepper. Equation (2.83) is then calculated for a discrete set of between 8 and 30 points for $\epsilon_{\text{check}} \in [1 \times 10^{-1}, 1 \times 10^{-12}]$, for a random vector $\mathbf{v}'$ with the minimum expected at $\epsilon_{\text{check}} \approx 4 \times 10^{-8}$ with $E \approx \mathcal{O}(10^{-8})$.

Performing these checks ensures that both the linearisation procedure and the matrix-vector product for the direct calculation are implemented correctly. As an example of this, we show this final stage check for the linear cascade case, that we will consider later on in Chapter 3, in Figure 2.12.

Deriving the adjoint equation through Equation (2.60) is a matrix operation however, and so is independent of the quality of the linearisation used to derive $\mathbf{A}$. This means that the duality relation $\langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle_{\mathbf{M}} = \langle \mathbf{A}^H \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}}$ should hold for all test vectors $\mathbf{v}$, $\mathbf{w}$ and all valid inner products $\langle \cdot, \cdot \rangle_{\mathbf{M}}$ up to machine precision. The free choice of test vectors allows for the substitution of random fields here which perturb all degrees of freedom to ensure that all degrees of freedom are implemented correctly. Note that the inclusion of the derivative terms in both Equation (2.59) and Equation (2.60) do not invalidate this analysis, despite the test vectors being non-differentiable, since we are more concerned with the duality relation holding rather than accurately computing the derivative.

2.11 Conclusion

This chapter summarises the fifth-order DNC solver that is the basis for this project, and which was inherited by us for the purposes of the analysis and development that is to follow this section of the thesis.

Here we have given a brief description of our high-order, Navier-Stokes solver, with an overview of the governing equations, the decomposition into pseudo-wave form and the corresponding boundary conditions. Differentiation in space and integration in time are also addressed with the relevant features of the schemes summarised. Linearisation capabilities of this code are carefully described, with specific attention given to the efficient extraction of the linear operators.

We concluded this chapter with a brief summary of the tests used to carry out the validation of this code, and to demonstrate its suitability to the simulation and prediction of aeroacoustic processes.

Chapter 3

Global Stability Analysis of An Idealised Compressor Blade Row

Abstract

A direct-adjoint mean flow global stability investigation of self-excited instabilities in an idealised, two-dimensional compressor blade row at off-design conditions is carried out. The work is split into two parts. First, nonlinear, impulse responses and the global modes are computed for a single passage system, with good agreement between the linear and nonlinear structures. Structural sensitivities and feedback loops are identified with the aid of wavemakers, and show that dominant structures arise due to feedback mechanisms linking the pressure and suction sides of the airfoil. The second part extends this to multiple blade passages per period window by exploiting the theory of block-circulant matrices and Bloch-wave theory. Analysis from the single passage is repeated on a system with ten blade passages, with more unstable ten- and five-periodic synchronisation structures arising as a result of a more optimal phase relationship that is supported by the larger system. Low-frequency structures are briefly analysed, and comments made regarding the extension of this analysis to aerodynamically mistuned systems[1].

3.1 Introduction

Increasingly stringent emissions regulations are presenting novel challenges to the turbomachinery community. With the need to produce more efficient, lighter and quieter aeroengines, this requires additional understanding of flow instabilities and unsteady flow phenomena present in the various components of the turbomachine [Peake and Parry, 2012]. Despite this, early-stage design in industry still relies mostly

[1] In the thesis this is included in Appendix C.

on steady-state RANS models over a single passage that capture the aerodynamics, but not the instabilities. Access to computationally efficient models, as well as an in-depth understanding of the instabilities and their feedback mechanisms, is therefore required for optimising these systems in the future.

Modal and non-modal analyses provide a powerful framework for the analysis of these systems [Schmid, 2007; Taira et al., 2017], and have, in the past, been applied to a wide variety of flow problems with some considerable success.

Of these, perhaps the most used in industrial settings are the data-driven modal-decomposition approaches of the proper orthogonal decomposition (POD) [Lumley, 1967], its spectral counterpart (SPOD) and the dynamic mode decomposition (DMD) [Schmid, 2010]. These have been applied extensively in turbomachinery contexts, with success of these methods largely attributable to the fact that they may be applied both to numerical and experimental data directly, using snapshots of the nonlinear flow solution or observations of the flow. Many comprehensive reviews demonstrating the theory, strengths and weaknesses [Taira et al., 2017] and the application [Taira et al., 2020] of these methods exist. Although these methods offer a convenient, and frequently the only, option for a modal decomposition of the flow, perhaps their greatest weakness is that they fail to provide any direct indication of the structural sensitivities of the various modal structures that are obtained. For numerical studies, however, the availability of a linearised component (i.e. the direct linear operator) can be used to address this by making use of the tools of numerical analysis, and enables global stability analyses to be carried out [Schmid and Henningson, 2001].

The global modes approach is most familiar in the context of steady base flows undergoing transition to instability, with early forays into stability theory starting with the study of simple flows at low Reynolds numbers [Schmid and Henningson, 2001]. By contrast, the vast majority of flows in turbomachinery are highly unsteady, with flow conditions taking the flow well beyond criticality, and with this the

argument for the application of classical stability theory around steady base flows becomes increasingly tenuous. Recent developments have instead considered time-averaged mean flows as the linearisation state. Unlike steady base flows these always exist, but are no longer the steady states of the Navier-Stokes equations. It has been shown in Beneddine et al. [2016] that this analysis remains valid however because important nonlinearities of the flow field are accounted for in the mean flow. For this reason, many studies that have attempted mean-flow global stability analyses of fuel injectors [Juniper, 2012; Tammisola and Juniper, 2016], jets [Gudmundsson and Colonius, 2011] and airfoils [Fosas de Pando et al., 2014, 2017], have shown a very good agreement between both the spatial structures and their corresponding temporal frequencies when compared to the spectrum of the nonlinear flow field. Furthermore, the adjoint of the direct operator[2] can provide additional structural sensitivity information for the given global modes [Hill, 1992; Giannetti and Luchini, 2007; Chomaz, 2005; Schmid and Brandt, 2014], and has seen great utility in many turbomachinery optimisation applications [Wang and He, 2010; Wang et al., 2010], and sensitivity analyses [Fosas de Pando et al., 2017]. Two versions of the adjoint approach exist: the continuous [Jameson, 1988] and the discrete [Giles and Pierce, 2000], with the latter preferred here for the consistency with the discretised direct operator. At this point it is also worth highlighting that for many flow problems the Navier-Stokes operator is highly non-normal, the result of which is that it has the potential to support large scale transient dynamics that differ considerably from the asymptotic behaviour obtained from the global modes [Trefethen and Bau III, 1997]. For this reason, performing an impulse response analysis prior to calculating the spectrum and the corresponding eigenvectors is a prudent step to evaluate the effects that non-normality has on the flow.

With turbomachinery blade rows consisting of multiple periodic structural elements interacting with the flow, it is natural to consider extensions of the global

[2]Which is equivalently the Hermitian conjugate of the direct linear system.

stability problem to N passages, so that the large-scale behaviour, that previously could not be modelled with the restriction to a single-passage periodic domains, is captured. Unfortunately global stability analysis becomes intractable, even in situations when N is relatively small (e.g. $N < 5$), and so the problem must be simplified in some way to account for this periodicity without increasing the instantaneous demand for computer resources. This difficulty has been addressed previously using the harmonic balance method [He and Ning, 1998; Hall et al., 2002; Ekici and Hall, 2007; Gopinath et al., 2007], phase-lagged [Erdos et al., 1977; He, 1990, 1992] and chorochronic approaches [Gerolymos et al., 2002] to account for the periodicity in blade row simulations, although note that none of these can account for self-excited instabilities as the formulations of these approaches are subject to prescribed phase-shift periodicities of *known* frequencies. More recently, Schmid et al. [2017] proposed a fully algebraic linear approach, where the direct system matrix is rearranged into block-circulant form by reconfiguration of the state vector. Assumptions made regarding the numerical nearest-neighbour coupling[3] of the blade passages can be exploited to increase the sparsity of the linear operator without eliminating any physical dynamics. This approach is reminiscent of Bloch-wave theory [Bloch, 1929], and similar methods have been used in the past in the field of structural mechanics to study vibrations in bladed disks [Olson et al., 2014]. The diversity of potential applications was demonstrated in [Schmid et al., 2017], with dynamics in periodic viscous wakes, annular combustors, burners and blade rows provided as examples. In the latter instance a fifty-four-blade linear cascade was studied in the context of incompressible flow, with a continuously-parameterised spectrum obtained alongside optimal forcing and response structures. In this paper we seek to extend this to a fully compressible flow problem with the aim of identifying the macro-dynamics originating from the feedback mechanisms linking the micro-scale self-excited hydrodynamics processes and the propagating acoustics field.

[3]In other words, the code enforces the coupling, but, by construction, communications in the code occur only between neighbouring passages.

Perhaps the closest analogue of this in turbomachinery applications are compressor non-engine order vibration phenomena [Hellmich and Seume, 2008] and acoustic resonance structures [Parker, 1967, 1984; Parker and Stoneman, 1985, 1989], which have been suggested to occur as a consequence of vortex shedding from separation bubbles forming on the suction surfaces of the compressor blades when operating at off-design conditions.

The demonstration problem selected in this analysis is inspired by earlier flow receptivity studies previously conducted on isolated airfoils, which we use here and apply to simplified cascade geometries to demonstrate this methodology. Briefly, earlier work by Paterson et al. [1973] identified the presence of discrete tones on isolated airfoils and suggested empirical laws for the evolution of the main peak for the tonal noise structure. Theoretical and experimental work by Tam, Fink, Nash and others [Tam, 1974; Fink, 1975; Nash et al., 1999; Pröbsting et al., 2014] later proposed that the presence of these discrete tones can be attributed to localised acoustic feedback mechanisms linking the acoustic source at the trailing edge to the upstream boundary layers of the airfoil, with subsequent numerical work considering this problem on NACA0012 [Desquesnes et al., 2007; Jones and Sandberg, 2011] and controlled-diffusion geometries [Sanjose et al., 2019; Wu et al., 2021]. In his earlier work, Fosas de Pando et al. [2014] analysed this tonal noise problem with the NACA0012 airfoil and performed a global stability analysis that correlated these tonal noise peaks to global mode structures. Furthermore, later calculations by Fosas de Pando et al. [2017] extended this by obtaining the adjoint global modes, which, with the aid of wavemaker[4] regions, showed that both the pressure and suction sides of the airfoil participate in setting the feedback mechanisms. This latter analysis serves as the inspiration for the work carried out here.

Here the geometry and flow conditions are chosen so that the flow exhibits many of the qualitative flow features that are present in more realistic settings. These

[4]As will be explained later, wavemakers are the element-wise products of the global adjoint and direct modes, which identify the structural sensitivities of the modes.

include oscillating laminar separation bubbles that form under adverse pressure gradients on the airfoil suction side, unsteady boundary layers on the pressure surface, trailing edge shedding processes, inhomogeneous mean-flow density gradients, and acoustic scattering from the leading and trailing edges. Therefore, the goal here is not only to demonstrate the stability analysis framework, but also to inform the reader of potentially interesting phenomena that may soon be within reach of this methodology in future large-scale LES and DNS simulations.

This paper is organised as follows. First, Section 3.2 introduces the representative compressor cascade flow case on which this global stability analysis is to be performed, with further details of the compressible, viscous Navier-Stokes solver, its linear counterpart, and the methodology for obtaining the global modes concluding this section. The single-passage case is then considered in Section 3.3, so that the *local* feedback mechanisms over one blade can be identified. Here the nonlinear flow field is analysed using a global Fourier transform in Section 3.3.1, which identifies the main features of the acoustic spectrum for this problem, and the mean flows obtained and described in Section 3.3.2. A linearisation around this mean flow is performed, and an impulse response analysis carried out in Section 3.3.3, to identify transient growth effects that occur due to operator non-normality. Following this, the global modes are calculated in Section 3.3.4, the spectrum analysed and decomposed into groups of modes that capture distinct families of behaviour for the dominant modes in Section 3.3.5.1, high-frequency separation bubble dynamics in Section 3.3.5.2 and for confined mode structures in Section 3.3.5.3. The consideration of the single passage concludes with a phase and group velocity analysis in Section 3.3.6, which seeks to determine relationships between the different modes in the global spectrum. The second portion of this paper extends this analysis to a ten-periodic blade row using the algebraic method introduced by Schmid et al. [2017], in Section 3.4. Details of the geometry and implementation are provided in Section 3.4.1 and Section 3.4.2, respectively. This is followed once again by an im-

pulse response analysis in Section 3.4.3, and global stability analysis in Section 3.4.4, where we observe the new structures that arise in the larger domain. To complete the discussion, Section 3.5 summarises the findings and provides possible extensions and future perspectives of this work.

3.2 Numerical Methodology

3.2.1 Geometry

The single-passage geometry consists of an idealised controlled-diffusion compressor blade developed by [Sanger and Shreeve, 1986; Sasaki and Breugelmans, 1998], and later used in a low-speed aero-elasticity experiment [Yang and He, 2004]. All simulations presented here are performed at a chord-based Reynolds number of $\mathrm{Re} = 100,000$ and an exit Mach number of $M = 0.3$. The inlet flow angle is set to $\theta_{\mathrm{in}} = 37.5^{\circ}$ to match the inlet angle of the blades, and the non-dimensional pitch length for the periodic boundary of the linear cascade is set to $h_p = 0.6$. These flow conditions are adapted from the experiment and previous numerical studies [He and Yi, 2017; He, 2021] by decreasing the Reynolds number from $\mathrm{Re} = 195,000$, and increasing from $M = 0.05$, so that a coarser mesh, and consequently faster time stepping, could be used to increase calculation speed and enable us to model some of the larger, more demanding cases in Section 3.4. Furthermore, as the purpose of this study is to develop and demonstrate the numerical framework, the changes introduced to the Reynolds and Mach numbers are applied for numerical convenience to expedite the numerical calculations rather than to make definitive statements regarding the flow properties of the airfoil at the experimental conditions. Despite this, the new conditions are chosen so that the main mean flow features, in particular the laminar separation bubble on the suction side and its location, match qualitatively with the experiment so that the dynamics that the system exhibits are *qualitatively similar*. This will be discussed further in Section 3.2.2.

Parameter	Grid A	Grid B	Grid C
n_{upstream}	540	576	1080
$n_{\text{downstream}}$	540	680	1080
n_{airfoil}	680	840	1360
n_{vertical}	256	256	512
L_{upstream}	3.075	3.233	3.075
$L_{\text{downstream}}$	4.024	4.250	4.024
y^+	1.4	0.8	0.7

Table 3.1: Single-passage mesh parameters. Grid A refers to the prime grid used for all calculations in this paper, while Grid B and Grid C are used for the purposes of establishing the mesh-independence of the mean flow. Grid B is a local refinement of the blade-surface boundary layer, while grid Grid C is designed as a continuously differentiable subdivision of each grid-point interval in Grid A.

The flow field is discretised on a structured H-grid mesh, with the mesh parameters given in Table 3.1, and the computational domain shown in Figure 3.1. To remain consistent with the analysis performed later in Section 3.4, the axial lengths of the upstream and downstream sections of this grid are fixed at $L_{\text{upstream}} = 3.075$, $L_{\text{downstream}} = 4.024$. Vertical refinement of the mesh in the downstream section is also gradually smoothed over $L_{\text{smoothing}} = 1.5$ so that unsteady wake structures are represented with sufficient resolution everywhere in the downstream region. Periodic conditions are applied along the top and bottom edges of the upstream and downstream meshes, with a steady inflow and outflow boundary conditions specified at the left and right sides of the mesh, respectively, in order to match the required flow angle and outlet Mach number. The minimum grid spacing is $\Delta s_0 = 4.78 \times 10^{-4}$, $\Delta s_1 = 1.89 \times 10^{-4}$ at the airfoil surface, with the maximum spacing found in the upstream region with $\Delta s_0 = 8.30 \times 10^{-3}$, $\Delta s_1 = 6.39 \times 10^{-3}$, where s_i is the arc-length along direction $i \in 0, 1$[5]. The maximum stretching ratio, $|\Delta s_i - \Delta s_{i-1}|/|\Delta s_{i-1}|$, between consecutive nodes is set to 5%.

At these flow conditions the main unsteady feature is a reattaching laminar separation bubble which forms on the suction side and is centred at $0.55 \times C_{\text{ax}}$, where C_{ax} is the axial chord, before reattaching, as seen in Figures 3.2–3.3, which

[5]In other words, $i = 0$ is the horizontal direction, and $i = 1$ the vertical, for a rectilinear Cartesian mesh.

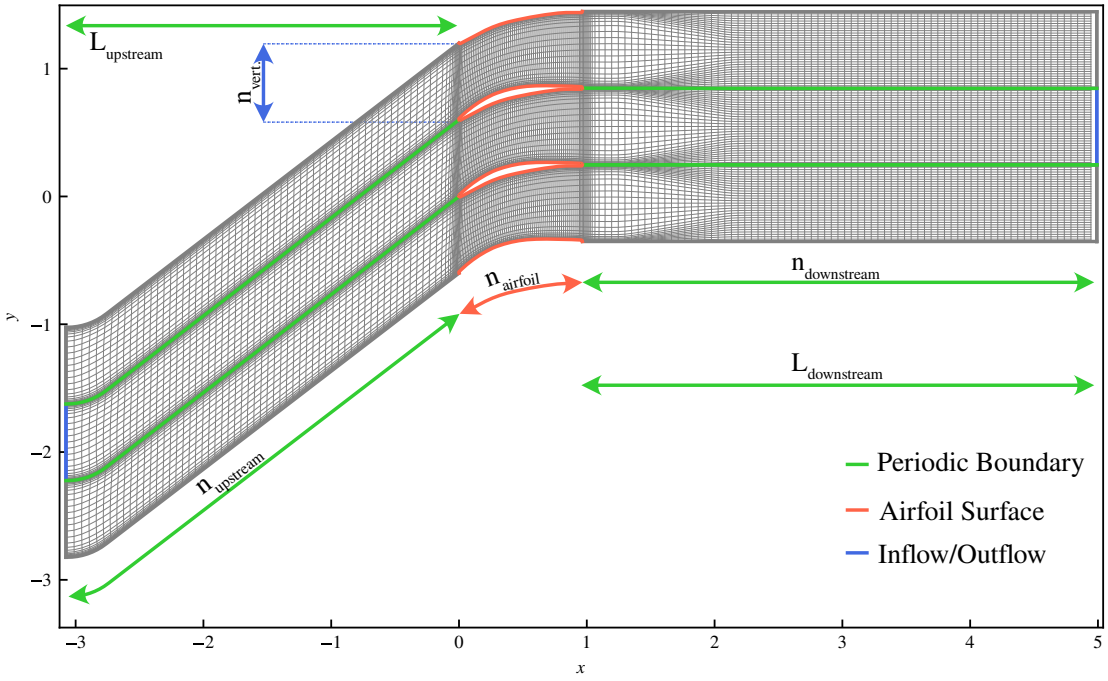


Figure 3.1: The single-passage structured grid showing every tenth grid line in each dimension so that the grid metric may be visualised. Inflow boundary conditions are applied on the left and outflow on the right (blue). Nonslip, adiabatic conditions (red) are imposed on the blade surface, and periodic boundaries are imposed on the upper and lower surface of the upstream and downstream meshes (green).

also characterise the sensitivity of the flow to Mach and Reynolds numbers, as well as demonstrating grid independence of the solution for $M = 0.3$, $\mathrm{Re} = 1 \times 10^5$.

In the results presented there, lowering the Reynolds number increases the size of the separation bubble and moves the reattachment point downstream towards the trailing edge, while increasing the Mach number moves this point back upstream. Furthermore, the separation point remains invariant to changes in both parameters, as seen in Figure 3.3. As a result, by lowering the Reynolds number and increasing the Mach number, the two-dimensional mean flow can be made to resemble the form of the experimental data obtained for the flow conditions at $M = 0.05$ and $\mathrm{Re} = 1.95 \times 10^5$. In so doing, this two-dimensional simulation has qualitatively the same features, with the laminar separation bubble at a similar location and with similar separation and reattachment locations. Although these changes will alter the shedding frequencies and sizes of the resulting vortices, the interactions between

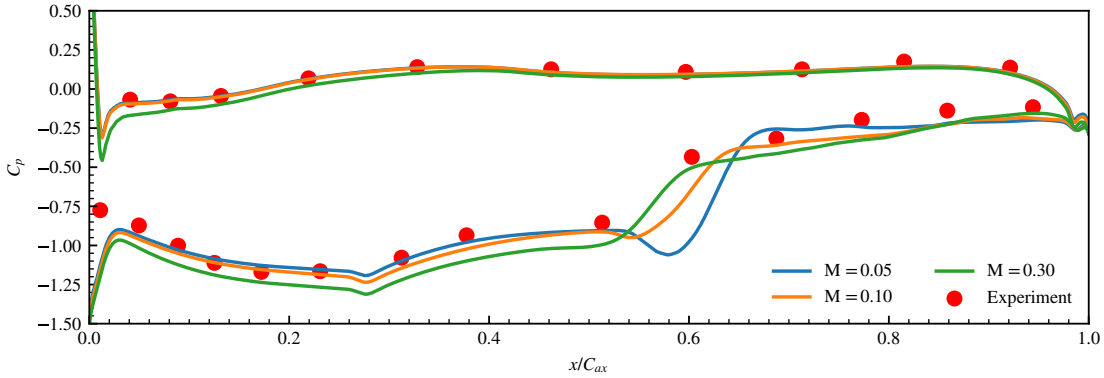
the boundary layer, separation bubble and trailing-edge processes will remain, and can be studied using the machinery developed here.

Numerical and experimental nonlinear unsteady simulations of this blade have also shown strong vortex shedding from both this detached shear layer as well as at the blunt trailing edge. At higher Reynolds numbers, both the laminar and turbulent solutions of this cascade [He and Yi, 2017] also show two-dimensional structures within the pressure side boundary layer which is subject to a positive pressure gradient. These two-dimensional features subsequently turn into Görtler vortices, or streamwise streaks, on the final quarter of the pressure surface, before finally interacting with the trailing edge. Note that these features are absent in the laminar solution of [He and Yi, 2017], although present in the scale-resolving turbulent flow simulation of [He, 2021]. These flow conditions provide rich dynamics under which multiple complex flow features interact and are coupled by the resulting acoustic field arising from this unsteadiness.

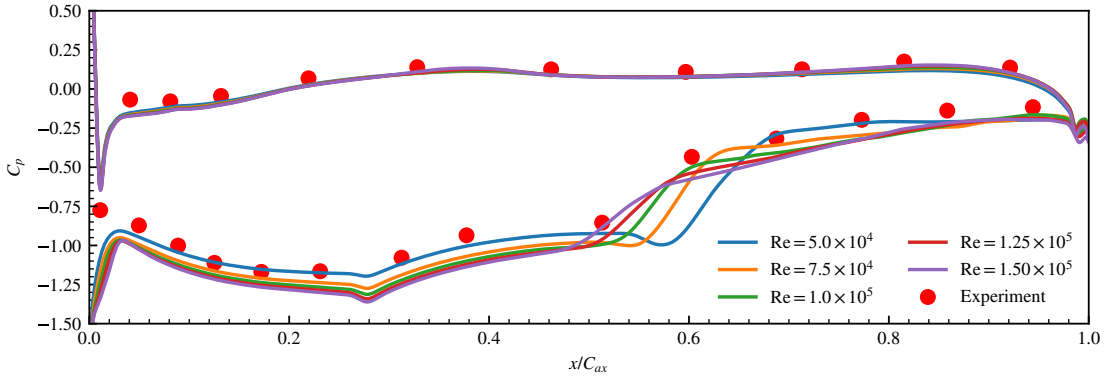
3.2.2 Numerical Simulations

A two-dimensional compressible, high-order Navier-Stokes code based on the pseudo-wave method [Sesterhenn, 2000; Fosas De Pando, 2012] is used for the numerical work in this paper. Advection terms are computed using fifth-order compact upwinding schemes [Adams and Shariff, 1996], with the dissipation terms differentiated using a compact third-order scheme [Lele, 1992]. A low-storage Runge-Kutta 4 [Kennedy et al., 2000] is then used to march the solution forwards in time using the method of lines. Parallelisation is achieved using the Message Passing Interface (MPI) [MPI Forum], so that these codes may be deployed on modern high-performance parallel computing hardware with maximum scaling efficiency.

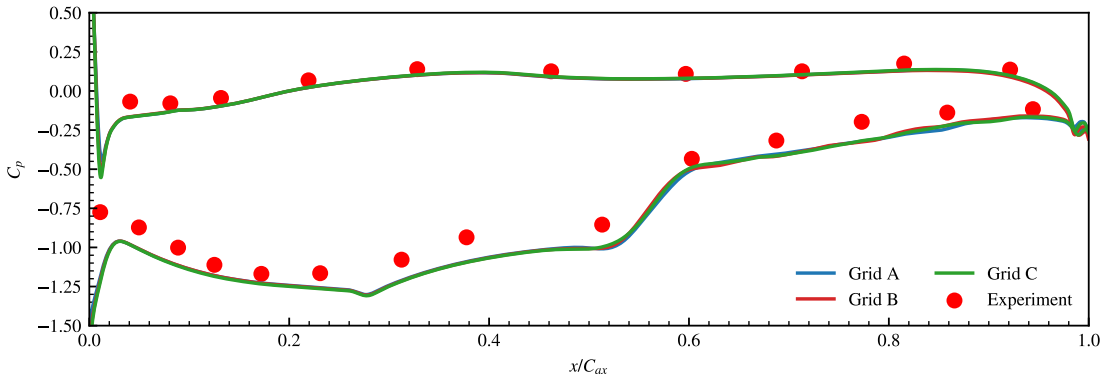
Direct and adjoint dynamics are obtained from an efficient linearisation [Fosas de Pando et al., 2012] of the nonlinear code. Matrix-vector multiplications and timestepping functionalities of this code interface with high-performance linear al-



(a)

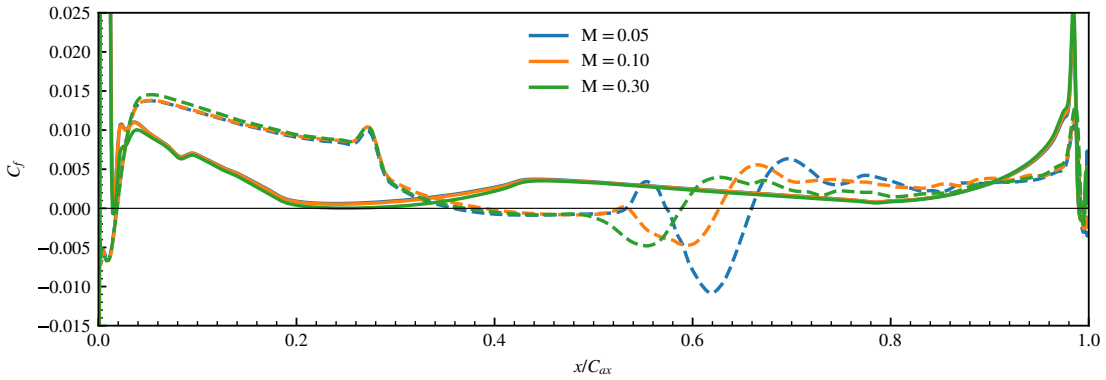


(b)

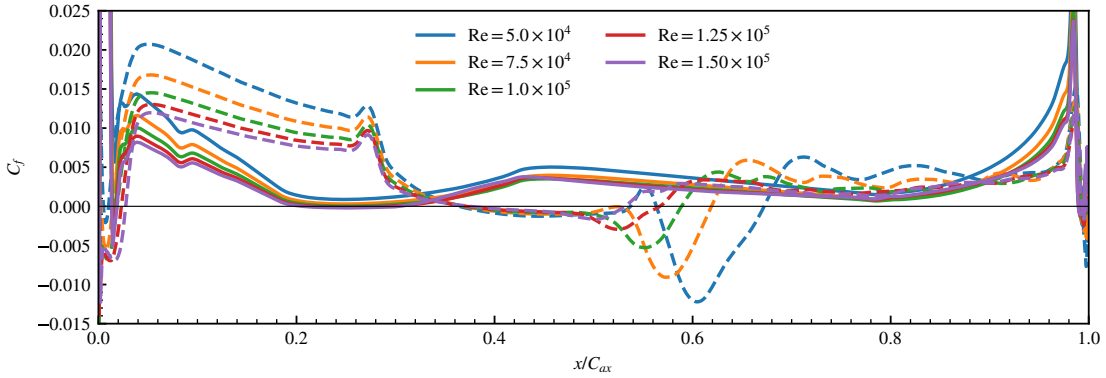


(c)

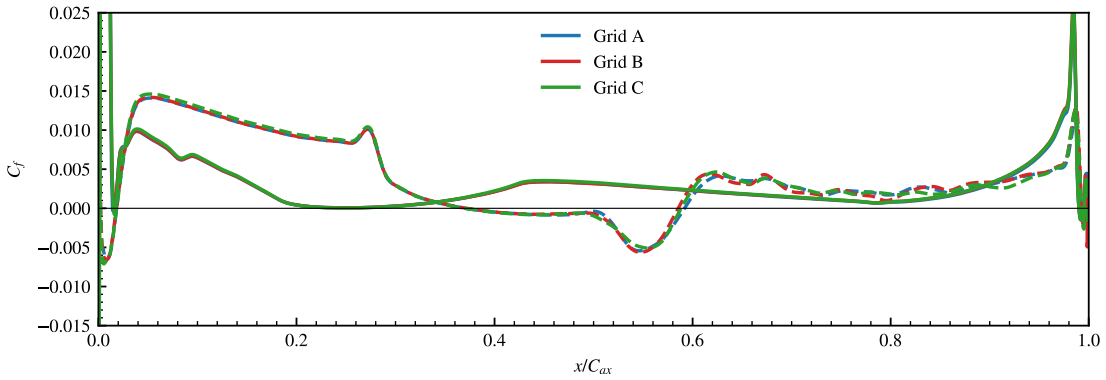
Figure 3.2: The surface pressure coefficients $C_p = 2(p-1)/(\gamma M^2)$ at selected Mach and Reynolds numbers. Experimental data available for the pressure coefficient is plotted using solid dots.



(a)



(b)



(c)

Figure 3.3: The skin-friction coefficients $C_f = 2(\mathbf{t} \cdot \bar{\boldsymbol{\tau}} \cdot \mathbf{n})/\,\mathrm{Re}$ at selected Mach and Reynolds numbers. Here $\mathbf{t}$ and $\mathbf{n}$ are the tangent and normal vectors to the surface, and $\bar{\boldsymbol{\tau}}$ is the averaged stress tensor. Solid lines indicate the pressure surface, with dashed lines representing the suction side of the blade. Experimental data available for the pressure coefficient is plotted using solid dots.

gebra libraries and solvers integrated into PETSc [Abhyankar et al., 2018; Balay et al., 2020, 2019] and SLEPc [Hernandez et al., 2005] through wrappers providing interfaces to the linearised right-hand side functions.

3.2.3 Choice of Base Flow

Integral to any global stability analysis is the choice of base flow that is made as the linearisation point around which the linear dynamics are considered. Snapshots of the flow shown in laminar and scale-resolving simulations in [He and Yi, 2017; He, 2021], respectively, and later on in Section 3.3.1 show that the flow field is dominated by instabilities, with no natural steady flow around which a linearisation may be taken. Based on the evidence of prior attempts to obtain steady flow states for isolated airfoil flows (see e.g. [Fosas De Pando, 2012]) by reducing the Reynolds number and using stabilisation techniques such as selective frequency damping methods [Åkervik et al., 2006], such approaches are likely to lead to large separation regions that do not adequately reflect the realistic flow dynamics in the vicinity of the blade surface. Classical stability analyses performed on steady base flows, derived through these methods, are similarly incapable of accurately modelling the unsteady spatial structures and characteristic frequencies due to the removal of intrinsic nonlinearities that govern the unsteady characteristics of the flow (see discussion in e.g. [Jones and Sandberg, 2011]). For this reason, and following from the discussion in the introduction, we select the time-averaged mean flow as the linearisation state.

This mean flow is computed from the nonlinear flow field once it has reached a quasi-periodic limit cycle on a single periodic passage, for $t \geq 20$. This flow is then averaged over an additional time interval of size $t_T = 30$, so that the flow in the vicinity the blade surface is converged. Despite this long averaging time, rare vortex merging phenomena occurring in the downstream region result in a far-field mean flow that is not be fully converged until much later. This is not considered to be a significant problem, however, since we focus our analyses primarily on the

hydrodynamics and acoustics in the vicinity of the airfoil and in the near wake, with the dominant dynamics expected to occur in these regions.

3.2.4 Linearisation, Linear and Adjoint Dynamics

3.2.4.1 Linear Operators and Governing Equations

The nonlinear Navier-Stokes equations, partially discretised in space, are written as

$$\frac{\mathrm{d}\mathbf{q}}{\mathrm{d}t} = \mathbf{F}(\mathbf{q}, t), \tag{3.1}$$

with $\mathbf{F}$ representing the nonlinear Navier-Stokes operator.

In the discretised framework, the derivatives within $\mathbf{F}$ introduce spatial coupling between the individual grid points, which make the naive approach of obtaining the linearisation of $\mathbf{F}$, by sequential perturbations around the mean flow of each degree of freedom, prohibitively expensive. Instead, by treating the derivatives as inputs [Fosas De Pando, 2012; Fosas de Pando et al., 2012][6], it is possible to break this down into a sequence of block operations that significantly speed up the process, and enable computationally inexpensive access to both the direct, $\mathbf{A}$, and the adjoint, $\mathbf{A}^\star$, operators. Then the linear equations are

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \mathbf{A}\mathbf{v} + \mathbf{f}, \tag{3.2}$$

$$\frac{\mathrm{d}\mathbf{w}}{\mathrm{d}t} = \mathbf{A}^\star\mathbf{w}, \tag{3.3}$$

with a forcing term $\mathbf{f}$, which can be used to model the action of a time-dependent physical source on the linear system, such as a harmonic acoustic source or an entropy wave, for example. From this point on, Equation (3.2) is referred to as the direct equation and Equation (3.3) as the adjoint equation.

The analysis of the stability properties of the operators $\mathbf{A}$ and $\mathbf{A}^\star$ forms the core component of this paper.

[6] See also Section 2.7 and Chapter 4 in the thesis.

3.2.4.2 Choice of Inner Product

The small-perturbation energy norm for compressible flows [Chu, 1965] is used throughout this study to quantify the magnitude of the linear disturbance. Written in the variables $(p, s, \mathbf{u})$, the energy $\mathcal{E}$ within a fluid enclosed by a fixed volume $\mathcal{V}$ is[7]

$$\mathcal{E} = \frac{1}{2}\int_{\mathcal{V}} \left(\frac{1}{\gamma^2 M^2}\frac{p^2}{\bar{p}} + \frac{\gamma-1}{\gamma^2 M^2}\bar{p}s^2 + \bar{\rho}\|u\|^2 \right) \mathrm{d}\mathcal{V}, \tag{3.4}$$

with $\bar{p}$ and $\bar{\rho}$ denoting, respectively, the mean flow[8] pressure and density.

On a discretised domain, the inner product is

$$\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = \mathbf{w}^H \mathbf{M} \mathbf{v}, \tag{3.5}$$

where the positive definite matrix $\mathbf{M} = \mathrm{diag}\,(\mathbf{m}_0, \mathbf{m}_1, \ldots, \mathbf{m}_{n-1})$ is block diagonal with

$$\mathbf{m}_i = \begin{pmatrix} \frac{1}{\gamma^2 M^2}\frac{1}{\bar{p}} & 0 & 0 & 0 \\ 0 & \frac{\gamma-1}{\gamma^2 M^2}\bar{p} & 0 & 0 \\ 0 & 0 & \bar{\rho} & 0 \\ 0 & 0 & 0 & \bar{\rho} \end{pmatrix}_i |\mathbf{x}(\xi)_i| \Delta\xi\Delta\eta, \tag{3.6}$$

for $i \in \{0, 1, \ldots, n-1\}$, n denoting the number of grid points in the field, and the determinant $|\mathbf{x}(\xi)|$ describing the scaling between the Cartesian physical and curvilinear computational spaces (in the thesis see Chapter 2).

With this choice of norm, the fundamental relation

$$\langle \mathbf{A}^\star \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = \langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle_{\mathbf{M}}, \tag{3.7}$$

must be satisfied to machine precision – typically $\mathcal{O}(10^{-16})$ on modern computers. From this we can also write the adjoint operator as $\mathbf{A}^\star = \mathbf{M}^{-1}\mathbf{A}^{\mathbf{H}}\mathbf{M}$, so that we

[7] Chu's energy norm generalises the disturbance kinetic energy norm to account for potential and thermal energy contributions. Therefore, the three terms in Equation (3.4) account for the effects of compression, heat exchange and kinetic energy, respectively, although it should be pointed out that the exact forms of these differ when compared to [Chu, 1965] because of the non-dimensionalisation used in this thesis.

[8] The mean flow can be substituted for any field around which the linearisation is taken, such as a steady base flow or a flow snapshot.

can perform all the calculations with $\mathbf{A}^H$ before scaling this accordingly to obtain the desired properties of $\mathbf{A}^\star$[9].

For visualisations the adjoint and wavemaker fields are divided by $\Delta\xi\Delta\eta$ so that the physical amplitudes displayed are independent of the number of grid points, since the discretised form of the inner-product integral in Equation (3.5) and Equation (3.6) is weighted to account for the integration over the unit square or cube.

3.2.5 Integration in Time

A general initial value problem

$$\frac{\mathrm{d}\mathbf{q}}{\mathrm{d}t} = \mathbf{F}(\mathbf{q}, t), \tag{3.8}$$

$$\mathbf{q}(0) = \mathbf{q}_0, \tag{3.9}$$

for some initial field $\mathbf{q}_0$, is integrated here in one of two ways. If $\mathbf{F}$ is nonlinear then a low-storage Runge-Kutta scheme [Kennedy et al., 2000] is used to march the solution forward in time at a CFL number[10] of 0.8. Alternatively if $\mathbf{F}$ is linear, then the time integration approximates the exact exponential solution to the linear IVP by using an exponential Krylov timestepping scheme [Schulze et al., 2009]. This scheme is used with Krylov subspace dimension of between 16 and 32 vectors, with a CFL number of 16.

3.2.6 Calculating Global Modes

The global modes are the eigenvalue-eigenvector pairs $(\lambda, \mathbf{v})$, or equivalently $(\mu, \mathbf{w})$ for the adjoint operator, of the familiar eigenvalue problem

$$\mathbf{A}\mathbf{v} = \lambda\mathbf{v}. \tag{3.10}$$

[9]It should be noted that this scaling back is mandatory on non-uniform grids, since the inner product accounts for the grid metric.

[10]The CFL number is defined globally, and is based on the minimum grid spacing and maximum wave speed.

In typical global stability applications the main purpose of the analysis is to identify features in the most unstable portion of the spectrum alongside the corresponding least stable eigenvalue, which determines the asymptotic stability properties of the system. For this we again leverage the properties of Krylov subspaces. The kernel of

$$(\lambda \mathbf{I} - \mathbf{A})\,\mathbf{v} = \mathbf{0}, \tag{3.11}$$

may be approximated by using the Rayleigh-Ritz procedure on the much smaller upper-Hessenberg matrix, $\mathbf{H}_m$, as

$$\left(\tilde{\lambda}\mathbf{I} - \mathbf{H}_m\right)\boldsymbol{\xi} = \mathbf{0}, \tag{3.12}$$

where the $\tilde{\lambda}$ are referred to as the Ritz values and the $\boldsymbol{\xi}$ as the Ritz vectors. Here the matrix $\mathbf{H}_m$ is a projection of $\mathbf{A}$ onto the Krylov subspace $\mathcal{K}_m(\mathbf{A}, \mathbf{v})$, defined as

$$\mathcal{K}_m(\mathbf{A}, \mathbf{v}) = \text{span}\left\{\mathbf{v}, \mathbf{A}\mathbf{v}, \mathbf{A}^2\mathbf{v}, \ldots, \mathbf{A}^{m-1}\mathbf{v}\right\}, \tag{3.13}$$

which can be constructed by using the Arnoldi algorithm. In our implementation, the Krylov subspace is constructed by selecting a random vector $\mathbf{v}$ and iteratively applying the linear operator to it.

Equation (3.12) much smaller problem would typically be solved using the QZ algorithm [Moler and Stewart, 1973], but direct application of this method tends to converge modes that are on the periphery of the spectrum. Instead we seek the modes of $e^{\Delta t \mathbf{A}}$, the time advancement operator, using what is commonly known as the time-stepper method [Edwards et al., 1994]. Based on the trailing edge vortex shedding period and the spectral analysis of the nonlinear flow field, we set $\Delta t = 0.016$.

The iterative Krylov-Schur technique [Stewart, 2002] is used to converge the eigenvalues with 2048 column vectors, with a single iteration of this method used to converge the spectrum. This choice is made necessary by slow-converging and erroneous modes near $\text{Im}\{\lambda\} = 0$, which occur as a by-product of the time-stepping

method. For a given error threshold, modes are returned from most to least stable until the first mode that no longer satisfies the error condition. After one iteration, the subspace is deflated, and the iteration repeats on the non-converged modes. Given that the erroneous eigenvalues near the origin are determined by the algorithm to have relatively large growth rates, it means that these block the return of the eigenvalues that have smaller growth rates but that may be well-converged otherwise. By increasing the subspace dimension and decreasing the error tolerance, we ensure that we are able to converge and return the maximum number of modes, avoiding the blocking phenomena. The result of this, however, is that it is no longer possible to perform multiple iterations of the algorithm since it now uses poorly converged modes when deflating the subspace. This process may be augmented with shift-and-invert or harmonic-extraction [Morgan and Zeng, 2006; Hernandez et al., 2005] strategies that would target the most unstable modes, but this is not performed here as it was found that picking a large subspace was sufficient to obtain the modes required within reasonable error tolerances.

Due to the non-uniformity of the grid, error estimates based on the L_2-norm overemphasise regions with high mesh-point densities, and, as a result, the error estimate returned by the Krylov-Schur algorithm will be higher. It was found that this issue can be mitigated to some degree by performing a Cholesky decomposition $\mathbf{M} = \mathbf{F}^H\mathbf{F}$, and redefining the linear operator so that the eigenvalues are computed for the operator $\mathbf{FAF}^{-1}$, so that the L_2 error norm becomes an error norm induced by $\mathbf{M}$. However, in this paper we found that the error estimates computed for the time-advancement operator, $e^{\Delta t\mathbf{A}}$, provided a good proxy for the weighted norm induced by $\mathbf{M}$, and, to reduce additional computation and post-processing time, all errors stated here are for the time-advancement operator.

3.2.7 Using the Global Modes

Unless otherwise stated, the following analysis will normalise the direct modes, $\mathbf{v}$, by maximum near-wake velocity, as shown in Figure B.1a, and the adjoint modes so that $\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = 1$. In this way the energy content of the mode can be compared directly alongside the sensitivity.

Let $(\lambda, \mathbf{v})$ *not* be an eigenvalue-eigenvector pair of the direct operator $\mathbf{A}$ on a finite-dimensional complex inner product space. This holds if and only if $\mathbf{A} - \lambda\mathbf{I}$ is invertible and if and only if there exists a matrix $\mathbf{S}$ defined on the same inner product space such that

$$\mathbf{S}(\mathbf{A} - \lambda\mathbf{I}) = (\mathbf{A} - \lambda\mathbf{I})\mathbf{S} = I, \tag{3.14}$$

taking the adjoint of which gives

$$\mathbf{S}^\star(\mathbf{A}^\star - \bar{\lambda}\mathbf{I}) = (\mathbf{A}^\star - \bar{\lambda}\mathbf{I})\mathbf{S}^\star = I, \tag{3.15}$$

in other words, if and only if $\mathbf{A}^\star - \bar{\lambda}\mathbf{I}$ is invertible. Taking the negative of this chain, this proves that the eigenvalues μ of $\mathbf{A}^\star$ are the complex conjugates of the eigenvalues $\mathbf{A}$, $\mu = \bar{\lambda}$. For real operators $\mathbf{A}$, as is the case here, the spectrum of eigenvalues are complex conjugate pairs. Despite this relationship, the eigenvectors of $\mathbf{A}^\star$ do not follow trivially from the eigenvectors of $\mathbf{A}$ and must be computed explicitly. Moreover, as we have alluded to in Section 3.3.3, the linearised Navier-Stokes operator $\mathbf{A}$ is non-normal (i.e. $\mathbf{A}\mathbf{A}^H \neq \mathbf{A}^H\mathbf{A}$), meaning that it is expected that the spatial support of the direct and adjoint modes will not, in general, coincide [Schmid and Brandt, 2014].

So far the notion of the adjoint has been introduced in this paper in an abstract way through operator analysis, however this concept, and the relationship between the adjoint and direct eigenvectors, can be made concrete. Despite the fact that the direct eigenvectors may not form a spanning set of the vector space of n-dimensional solution vectors, in principle it is possible to extend this subset so that it is a basis for the resulting inner product space and the additional vectors that complete the

basis are orthogonal to all the global modes (in other words, they do not project onto the global modes). Then, for any given perturbation $\mathbf{v}$ we can represent it as a linear combination of the global modes and the extension set

$$\mathbf{v} = \sum_{i=0}^{n_0-1} \alpha_i \mathbf{v}_i + \sum_{j=0}^{n_1-1} \beta_j \mathbf{u}_j, \tag{3.16}$$

and taking the inner product with respect to the adjoint mode $\mathbf{w}_k$, it follows that

$$\langle \mathbf{w}_k, \mathbf{v} \rangle = \sum_{i=0}^{n_0-1} \alpha_i \langle \mathbf{w}_k, \mathbf{v}_i \rangle + \sum_{j=0}^{n_1-1} \beta_j \langle \mathbf{w}_k, \mathbf{u}_j \rangle. \tag{3.17}$$

Using Equation (3.7) we note that, for eigenmodes $(\lambda_i, \mathbf{v}_i)$ and $\left(\bar{\lambda}_j, \mathbf{w}_j\right)$,

$$\langle \mathbf{A}^\star \mathbf{w}_j, \mathbf{v}_i \rangle = \langle \mathbf{w}_j, \mathbf{A}\mathbf{v}_i \rangle \quad \Longrightarrow \quad (\bar{\lambda}_j - \lambda_i)\langle \mathbf{w}_j, \mathbf{v}_i \rangle = 0, \tag{3.18}$$

and if $i \neq j$, then $\bar{\lambda}_j - \lambda_i \neq 0$, and thus $\langle \mathbf{w}_j, \mathbf{v}_i \rangle = 0$.

We also note that if $\mathbf{v}_j \in \{\mathbf{v}_0, \ldots, \mathbf{v}_{n_0-1}\}$, then it follows that $\langle \mathbf{w}_k, \mathbf{v}_j \rangle = \delta_{kj}$ for all j and k. Equation (3.17) therefore leads to $\alpha_i = \langle \mathbf{w}_i, \mathbf{v} \rangle$, meaning that taking the inner product of $\mathbf{v}$ and the adjoint mode gives the projection of the vector $\mathbf{v}$ onto the corresponding direct mode.

We can exploit this projection property of the adjoint mode and introduce the concept of the wavemaker, which we use here in preference to the calculation of the ϵ-pseudospectrum [Trefethen, 1999] so as to reduce the expensive calculations needed for characterising the structural sensitivities of the modes. To do so, let us suppose that the operator is perturbed by a small amount [Schmid and Brandt, 2014], with commensurate changes to the resulting eigenvectors and eigenvalues, so that $\mathbf{A}' = \mathbf{A} + \delta\mathbf{A}$, results in $\mathbf{v}' = \mathbf{v} + \delta\mathbf{v}$ and $\lambda' = \lambda + \delta\lambda$ that satisfy

$$\mathbf{A}'\mathbf{v}' = \lambda'\mathbf{v}', \tag{3.19}$$

which expanded to first order becomes

$$(\mathbf{A} - \lambda\mathbf{I})\delta\mathbf{v} = (\delta\mathbf{A} + \delta\lambda\mathbf{I})\mathbf{v}. \tag{3.20}$$

Taking the inner product of both sides with respect to a trial vector $\mathbf{s}$ leads to

$$\langle \mathbf{s}, (\mathbf{A} - \lambda \mathbf{I})\delta\mathbf{v} \rangle_{\mathbf{M}} = \langle \mathbf{s}, \delta\mathbf{A}\mathbf{v} \rangle_{\mathbf{M}} + \delta\lambda \langle \mathbf{s}, \mathbf{v} \rangle_{\mathbf{M}}. \tag{3.21}$$

Let us suppose that we can eliminate the dependence of the left-hand side. In this case

$$\langle \mathbf{s}, (\mathbf{A} - \lambda \mathbf{I})\delta\mathbf{v} \rangle_{\mathbf{M}} = 0, \tag{3.22}$$

which is easily rearranged to

$$\langle \delta\mathbf{v}, \left(\mathbf{A}^{\star} - \bar{\lambda}\mathbf{I}\right)\mathbf{s} \rangle_{\mathbf{M}} = 0, \tag{3.23}$$

which is satisfied if we pick $\mathbf{s} = \mathbf{w}$. As such the sensitivity of the eigenvalue can be defined as

$$\delta\lambda = \frac{\langle \mathbf{w}, \delta\mathbf{A}\mathbf{v} \rangle_{\mathbf{M}}}{\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}}}, \tag{3.24}$$

where, if we consider the matrix $\mathbf{A}$ to be a delta function which identifies the local response from localised forcing, we can build up a spatial map, called the wavemaker, defined as simply the pointwise multiplication of the direct and adjoint modes: $\delta\lambda = \mathbf{w} \odot \mathbf{v}$, where we once again use the convention in this paper that $\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = 1$. Another interpretation of this is that the wavemaker defines the local projection of the spatial vector components of the mode onto the same mode. As a consequence, the wavemaker defines a spatial map indicating the locations in the flow field where the mean flow is most receptive to forcing by the mode, and represents the feedback mechanisms that drive and sustain that mode.

3.2.8 Details of Computer Wall Times

All numerical calculations within this paper require the use of high-performance computing (HPC) facilities, and therefore it bears a brief mention of how this is achieved here.

Calculations in Section 3.3 use the *FinisTerrae II* supercomputer provided by CESGA, using 2.5GHz Intel Xeon E5-2680v3 (Haswell) processors, while the results

in Section 3.4 are obtained using the Edinburgh Parallel Computing Centre (EPCC) *ARCHER* system, which used 2.7GHz Intel Xeon E5-2697 (Ivy Bridge) processors, although some additional results were also obtained using the implementation phase of the successor system: *ARCHER2*.

The single-passage mesh described above is the basis for all calculations in this paper. For the fast nonlinear cases, the simulations are run on 80 processors, with 20 in the horizontal/axial and 4 in the vertical/pitchwise directions, and amounts to a load of $22,528$ degrees of freedom per processor. For the larger eigenvalue calculations, this is increased to 164 processors in a 41×4 distribution, which was found to give sufficiently fast turn-around times.

For the multi-passage problem in Section 3.4, the methodology consists of performing similar calculations on three identical passages that are stacked pitchwise. Processor loading and distributions are identical in this case, with 492 ($= 3 \times 164$) processors used. It appears that the single-passage adjoint eigenvalue calculation was bottle-necked by disk write speeds for the simulation that returned > 2000 eigenvalues, and hence terminated after 48 hours. A later test simulation produced a nominal runtime of $34:00$, where the number of returned vectors was limited to 65. This seems to have been a problem with the HDF5 library on FinisTerrae II, since a re-installation of the code on ARCHER results in 'multiple-passages' figures that suggest a scaling between the direct and adjoint closer to 1.3. This is larger than the maximum direct-adjoint scaling factor of 1.1 observed by Fosas de Pando et al. [2012], and suggests that there may be additional performance losses incurred in either IO performance or through the amended N-periodic structuring of the right-hand side function (see Section 3.4.2 for additional details), the latter of which is more likely here.

A summary of computational times and costs is given in Table 3.2.

	Case	Grid Size	Procs.	Wall Time (h)	CPUh
Single Passage	Nonlinear	450,560	80	07 : 56	635
	Impulse Response			09 : 18	744
	Direct Eigenvalues		164	17 : 00	2,788
	Adjoint Eigenvalues			48 : 00	4,873
Multiple Passages	Impulse Response	1,351,680	492	∼08 : 30 (avg)	∼48,734 (total)
	Direct Eigenvalues			∼ 20 : 48 (avg)	∼ 102,336 (total)
	Adjoint Eigenvalues			∼ 27 : 40 (avg)	∼ 136,120 (total)

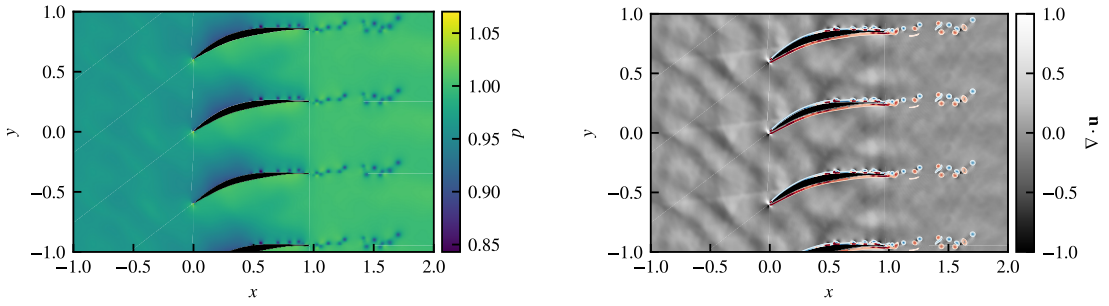
Table 3.2: Computer resource utilisation for the different cases shown in this paper.

3.3 Single Passage Dynamics

3.3.1 Nonlinear Flow Field

We begin by performing a brief analysis of the nonlinear flow field, computed directly using the nonlinear Navier-Stokes solver from Section 3.2.2. We follow the methodology of Section 3.2.3 and calculate an initial starting flow until $T = 20$, so that the initial transients arising from the initial conditions are removed. This is then integrated for 40 additional time units with snapshots saved every $\Delta t = 0.02$, giving 2000 samples. Positive frequencies[11] up to $f_{\max} = 24.96$, equivalent to $\mathrm{Im}\{\lambda\}_{\max} \approx 156.8$, with a spectral resolution/bin size of $\Delta f = 0.024$ are resolved with this sampling rate when this field is processed with a Fast Fourier Transform (FFT). Our choice for the range of frequencies considered here was made by estimating the frequency of the vortex shedding at the separation bubble and at the trailing edge. A snapshot of the nonlinear pressure and dilatation fields is shown in Figure 3.4. These plots demonstrate that the unsteadiness at the separation bubble and the trailing edge results in strong acoustic radiation that is directed primarily upstream through the

[11] All quantities here are given in non-dimensional form in accordance with the non-dimensionalisation given in Chapter 2.



(a) (b)

Figure 3.4: A snapshot of the (a) nonlinear pressure and (b) dilatation field, with vorticity contours, at $t = 60$. We draw attention to the vortex shedding on the suction side separation bubble and at the resulting interaction at the trailing edge. The scattering at the trailing edge produces an acoustic source that propagates acoustics disturbances into the upstream.

blade passage and into the inlet portion of the domain. The consequence of this is that there exists a potential pathway through which the acoustic field, which is generated by the sources at the separation bubble and the trailing edge, can couple and interact with highly sensitive regions where this flow first becomes convectively unstable. This mechanism will be analysed in depth later on when we consider the adjoint and wavemaker fields.

At this point we should note that, in three-dimensional simulations of this flow, the two-dimensional coherent structures, identified in Figure 3.4, will likely break down before the trailing edge. The resulting acoustic field would therefore have a broadband spectrum, possibly augmented by the presence of a number of peaks corresponding to the tonal acoustic modes. Despite this, prior simulations, especially in He and Yi [2017]; He [2021], indicate that the qualitative structure of the mean flow prior to transition is captured remarkably well by the two-dimensional flow fields shown here, with the laminar boundary layers and the reattaching separation bubble exhibiting qualitatively similar dynamics. As a result, the conclusions made in this paper regarding the receptivity and sensitivities of the laminar boundary layers are likely to be transferable to the three dimensional settings, at least to some degree. In situations where inlet turbulence or periodic wakes impinge upon the blade surface,

the two-dimensional behaviour observed here is likely to be masked further by more complex transition phenomena and interaction effects, such as bypass transition, however the presence of this two-dimensional analysis would allow for a comparison point against which these three-dimensional effects can be evaluated.

The high-resolution dataset described above is used to determine the features of the acoustic spectrum, with the pressure measurements taken at $\mathbf{x} = (-0.228, 0.138)$, which is sited upstream of the blade leading edge and at the midpoint between the periodic boundaries, and is also situated within the main 'lobe' of the upstream-propagating acoustic field (see also Figure 3.6, where this point is shown with a red cross). This spectrum is shown in Figure 3.5. The observed spectrum consists of a series of eleven discernible peaks that dominate the pressure field across a range of frequencies from $\mathrm{Im}\{\lambda\} = 20.73$ to 71.47, suggesting a certain low-dimensional aspect to this problem. Above this, the acoustic spectrum falls away rapidly, and this is to be expected as dominant unsteady features involve the trailing edge and separation bubble dynamics, which we observe are predominantly intermediate-wavelength phenomena, between $\mathrm{Im}\{\lambda\} \in [20, 60]$. Low frequency dynamics, $\mathrm{Im}\{\lambda\} < 20$, are also largely absent from the spectrum. While very low-frequency phenomena, such as the low-frequency breathing motion of the separation bubble, typically with $\mathrm{Im}\{\lambda\} < 1$ (see for example [Fosas de Pando et al., 2014]), are expected to feature in the spectrum, these are observed to be acoustically silent at the measurement location. Furthermore, in the context of the spectrum shown in Figure 3.5, this low-frequency region is only resolved by seven 'bins', and so much greater refinement of the frequency space is necessary to capture these modes. That is to say that low-frequency hydrodynamic phenomena may be present, but these are not active acoustic sources that can be captured through the acoustic field, and are likely to require much finer sampling to resolve fully.

Further potential for low-frequency system responses ($\mathrm{Im}\{\lambda\} < 20$) exist through resonance pathways. Acoustic standing waves with wavelengths on the order of

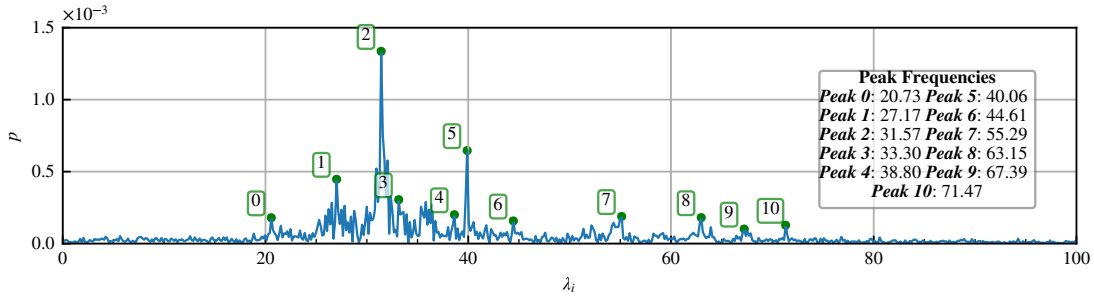


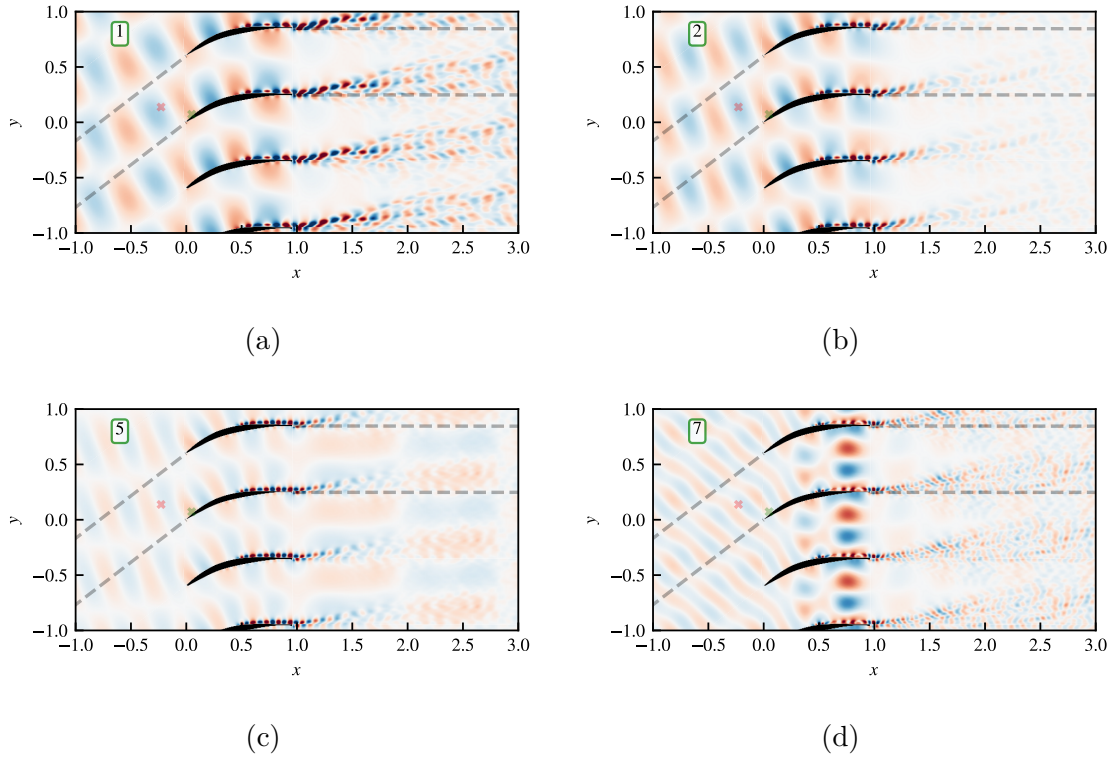
Figure 3.5: The Fast Fourier Transform (FFT) of the acoustic field upstream of the blade leading edge at $\mathbf{x} = (-0.228, 0.138)$.

L_{upstream} or $L_{\text{downstream}}$,would be associated with plane waves with frequencies on the order of $\text{Im}\{\lambda\} = 2\pi n(Mu + a)/2L$, which leads to the fundamental mode being $\text{Im}\{\lambda\} \approx 1$. These are effectively suppressed by the sponge layers at both ends of the domain, that create essentially anechoic conditions. For modes propagating vertically, the maximum wavelength that can be sustained by the periodic system is therefore $L = 0.6$, for which the associated frequency is $\text{Im}\{\lambda\} = Mv + a/L \approx 10.5$. Doubling this gives $\text{Im}\{\lambda\} \approx 21$, but verification with the field for peak [0][12] shows that this is not a resonant structure.

The net result of this is that the peaks observed in Figure 3.5 can be associated with amplified, self-excited instabilities rather than artefacts of the computational domain. Additionally, due to the laminar nature of this two-dimensional solution, the flow field lacks the transition to turbulence that characterises the three-dimensional flow, and therefore the broadband component of the acoustics is missing here. However, in this paper we are concerned primarily with establishing the mechanics of the sensitivity analysis, and therefore we leave the extension to fully physical, three-dimensional case for later work.

Extending this calculation of the FFT to every point within the domain, we isolate the fields for the observed peaks, which are plotted in Figure 3.6. Here, peaks [1] and [2] are associated with Kelvin-Helmholtz (KH) instabilities originating from the laminar separation bubble (to be characterised in Section 3.3.2) on the

[12]Here [⋆] corresponds to the labelled peak in the spectrum shown in Figure 3.5.



(a) (b)

(c) (d)

Figure 3.6: The snapshots of the pressure fields corresponding to the main peaks identified in the Fourier transform of the nonlinear pressure field.

suction surface of the airfoil, with the scattering of these coherent structures by the blunt trailing edge resulting in upstream-propagating acoustic waves. These two modes are also characterised by in- or out-of-phase dynamics, between the leading edge and the acoustic source at the trailing edge, for modes [1] and [2], respectively. Similarly, peak [5] is a higher-frequency shedding mode of the separation bubble, although the participation of pressure surface, and the associated interactions at the trailing edge, now result in a downstream pressure field. Finally, the field for peak [7] illustrates the formation of a resonant structure forming within the blade passage, but rather than being excited by the previous KH modes, the pressure field shows high amplitude fluctuations within the separated shear layer itself.

From these snapshots, it is clear that the flow is characterised by the interactions of the hydrodynamic and acoustic fields arising as a result of shear-layer instabilities interacting with the dynamics at the trailing edge. We therefore seek to establish

the stability and sensitivity properties of these modes in single and multiple passages in the remainder of this paper.

3.3.2 Properties of the Mean Flow

Having confirmed the separation bubble and trailing edge dynamics as the key contributing sources to the acoustic spectrum, we briefly characterise these regions in the mean flow prior to our modal and non-modal analyses in the subsequent sections of this paper.

Inspecting Figure 3.7b, where we have parameterised the airfoil surface as a function of arc-length, s, inflection points in the mean suction-side boundary layer velocity profiles appear at $s \approx 0.35$ due to the adverse pressure gradient experienced by the flow at this point. Inflectional profiles are known to be a necessary condition for instability in the inviscid and incompressible limit, as shown by Rayleigh [1879].

Briefly, the inflection point results are derived by assuming perturbation of the type $v(x_i) = \hat{v}(y)\exp(i(\alpha x - \alpha c t))$, $x_i \in x, y$ applied to the linearised Euler equations for a mean parallel flow, $U(x_i) = U(y)\delta_{1,i}$, which leads to the Rayleigh equation (see [Schmid and Henningson, 2001] for a complete derivation):

$$(U - c)(\mathcal{D}^2 - \alpha^2)\hat{v} - U''\hat{v} = 0, \tag{3.25}$$

where $\mathcal{D}$ denotes derivatives in the y-direction. If there are disturbances with $c > 0$, then multiplying the Rayleigh equation by the complex conjugate velocity $\bar{\hat{v}}$ and integrating by parts in the vertical direction gives

$$\int_{y_0}^{y_1} |\mathcal{D}\hat{v}|^2 + \alpha^2|\hat{v}|^2 \,\mathrm{d}y + \int_{y_0}^{y_1} \frac{U''}{U - c}|\hat{v}|^2 \,\mathrm{d}y = 0. \tag{3.26}$$

The first integral is real, but the second integrand is not necessarily. Thus taking the imaginary part of the above expression,

$$\mathrm{Im}\left\{\int_{y_0}^{y_1} \frac{U''}{U - c}|\hat{v}|^2 \,\mathrm{d}y\right\} = \int_{y_0}^{y_1} \frac{U''c}{|U - c|^2}|\hat{v}|^2 \,\mathrm{d}y = 0, \tag{3.27}$$

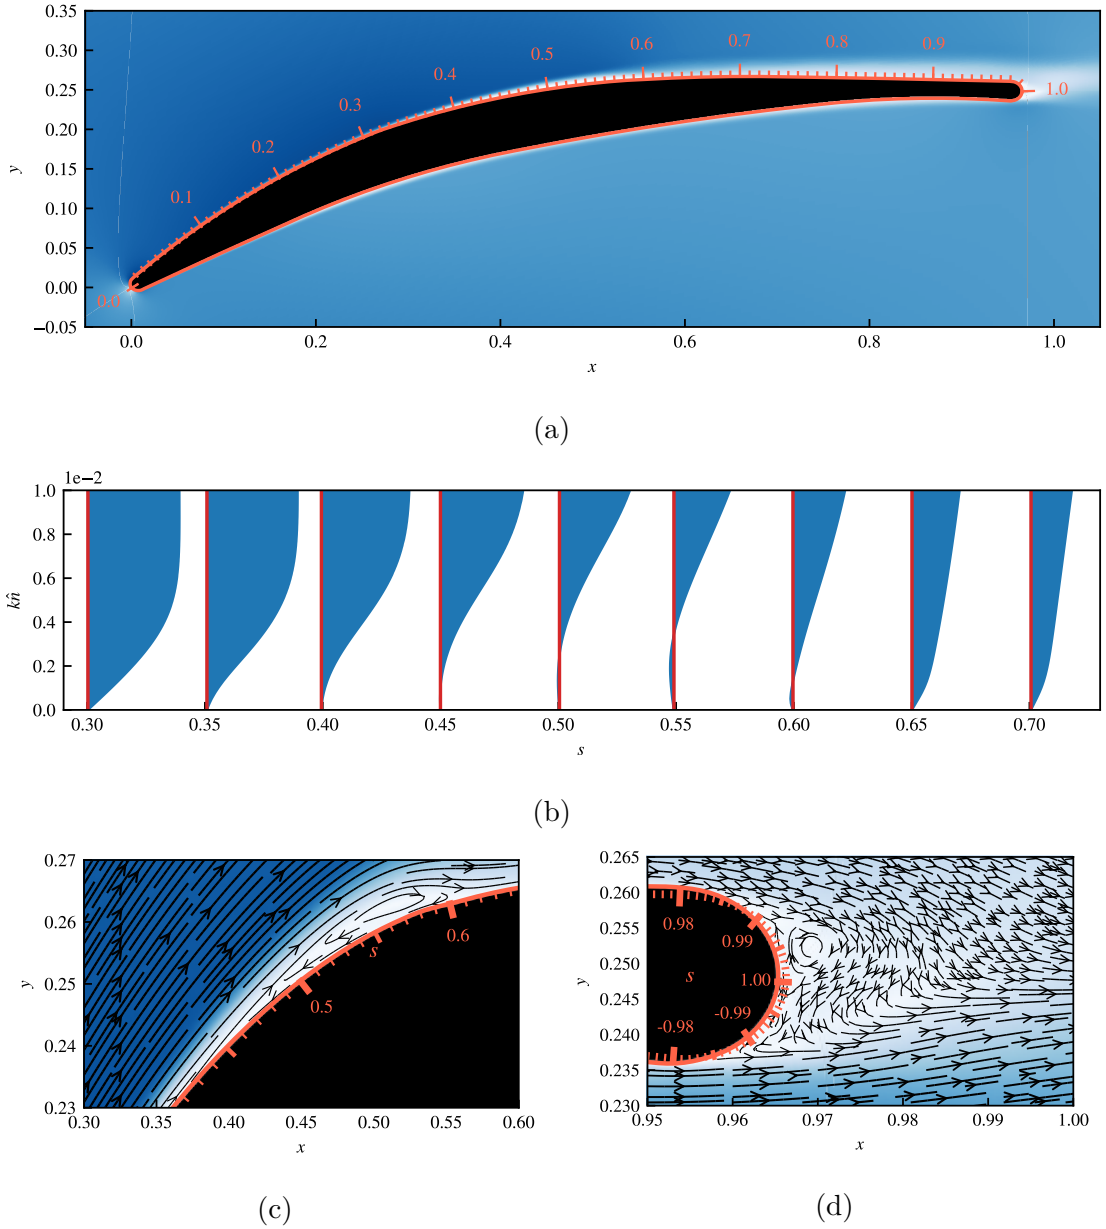


Figure 3.7: The mean flow around the airfoil showing (a) the speed field, $\|\mathbf{u}\|$ superimposed with an arc-length parameterisation of the suction surface, (b) the velocity profiles normal to the blade surface identifying the separation and reattachment points of the laminar separation bubble on the suction surface. Panels (c) and (d) visualise the local velocity field at the separation bubble and the trailing edge.

and since $c > 0$, it follows that U'' should change sign to make the second integral vanish. In other words, it follows that the mean flow profile has an inflection point. Stronger necessary conditions are derived by Fjørtoft [1950], which state that the inflection point has to be a mean vorticity maximum.

The incompressibility condition can also be relaxed and a *generalised inflection point criterion* derived (see [Schmid and Henningson, 2001] for full details). This suggests that the flow can support convective instabilities in the boundary layer even upstream of the separation point.

Further along at $s = 0.4$ the flow separates, as seen in both in Figure 3.2 and Figure 3.3, the reversed flow in the velocity profiles in Figure 3.7b, and the streamlines of the separation bubble in Figure 3.7c, before reattaching by $s = 0.65$. Our observations of the bubble are in agreement with previous experimental [Yang and He, 2004] and three-dimensional laminar and turbulent numerical calculations [He and Yi, 2017; He, 2021].

Downstream of the reattachment point, the vortex shedding from the separation bubble leads to a thickening of the mean flow boundary layer. Although the velocity profiles here may suggest a propensity for the formation of convective instabilities in the flow, we should remember that the nonlinear flow here is already highly unsteady, and we expect that this region of the mean flow serves more to support and amplify the instabilities that have grown downstream at the separation bubble, rather than modelling the growth of new instability mechanisms.

Fields along the pressure surface are similarly subject to adverse pressure gradients, however the combination of the concave geometry and the flow conditions results in a flow that, although close to separating, remains attached at all points. This has two implications. First, the flow here is again expected to be highly sensitive to flow disturbances, and, as will be seen later on in Section 3.3.4, will participate in the global instability dynamics through receptivity mechanisms. Secondly, small changes in the flow conditions, such as alterations to the stagger angle, lead to sepa-

ration on the pressure side, resulting in additional unsteadiness at the trailing edge as well as altered receptivity characteristics that dramatically change the acoustic signature of the airfoil.

Finally, the mean velocity at the trailing edge, shown in Figure 3.7d consists of a low-energy recirculating region immediately downstream of the trailing edge, with a high-speed 'jet' entering from the pressure side. The flow then entrains the lower-velocity separated boundary layer from the suction surface downwards, with the resulting change in momentum deflecting the jet upwards, around one trailing-edge diameter downstream of the trailing edge. This change in momentum may explain why the PS-SS vortex interaction at the trailing edge results in vortices being deflected upwards even though in the far-field the flow remains axial. We should note here that while the mean flow does provide us with insights into the macroscopic features of the flow, in areas where the flow field is highly unsteady the computed field may never be attained physically at any point in the simulation. Furthermore, even though we have shown that the two-dimensional meanflow captures the key properties of the experimental mean flows, how these features carry on into the three-dimensional case should be the subject of a careful and detailed future study.

3.3.3 Impulse Response Analysis

We now turn our attention to the impulse response analysis of the linearised direct operator $\mathbf{A}$, to analyse the transient growth phenomenon. It is an established fact that the linearised Navier-Stokes operator is non-normal[13], and the corresponding non-orthogonality of the eigenvectors results in systems that may significantly amplify non-modal perturbations, even in the presence of stable spectra [Schmid, 2007]. A transient response analysis is therefore a prudent step, since not only does it identify non-modal amplification mechanisms that are absent from the modal perspective, but also identifies whether the global stability analysis of this system is an

[13]In other words the operator does not commute with its Hermitian adjoint: $\mathbf{A}\mathbf{A}^\star \neq \mathbf{A}^\star\mathbf{A}$.

accurate description for the evolution of infinitesimal flow perturbations.

We consider the initial value problem (IVP) given by Equation (3.2), with $\mathbf{f} = \mathbf{0}$, and initial conditions $\mathbf{v}(t = 0) = \mathbf{v}_0$. The initial perturbation, with norm $||\mathbf{v}_0||_\mathbf{M} = 1$, takes the form of a divergence-free disturbance to the two velocity field components. This is done so as to reduce the initial pressure transients, which would otherwise contaminate the linear field as they reflect within the blade passage and propagate through the periodic boundary in the upstream and downstream sections of the computational domain. The exact form of the perturbation is

$$p_0 = 0, \quad s_0 = 0, \quad u_0 = \frac{-2(y - y_0 - h_p\alpha)}{\sigma^2}e^{-r}, \quad v_0 = \frac{2(x - x_0)}{\sigma^2}e^{-r}, \tag{3.28}$$

with $\alpha \in \{-2, -1, 0, 1, 2\}$ to ensure that the initial condition is periodic, $\sigma = 5\times10^{-3}$ and

$$r = \frac{(x - x_0)^2 + (y - y_0 - h_p\alpha)^2}{\sigma^2}, \tag{3.29}$$

and is centred directly upstream of the leading edge at

$$(x_0, y_0) = (-0.02, -0.02\tan(37.5\pi/180)). \tag{3.30}$$

This way the disturbance will be convected downstream, therefore interacting with both the pressure and suction surfaces during the initial transient response. The mean flow, initial condition and the location of the semi-norm used to measure the growth of this disturbance are shown in Figure 3.8.

Equation (3.2) is integrated until $t = 20$ using an exponential Krylov time stepping algorithm, with snapshots saved at intervals of $\Delta t = 0.02$. Evaluating these snapshots and processing them using this semi-norm gives the temporal evolution shown in Figure 3.9.

From this, two main stages of the evolution may be identified. The first of these is the initial transient response, which we can define here to be between $t = 0$ and $t = 5$, which consists of transient growth, of one order of magnitude between $t = 0.4$ until $t = 0.75$, followed by an overall decay in the perturbation energy. Furthermore,

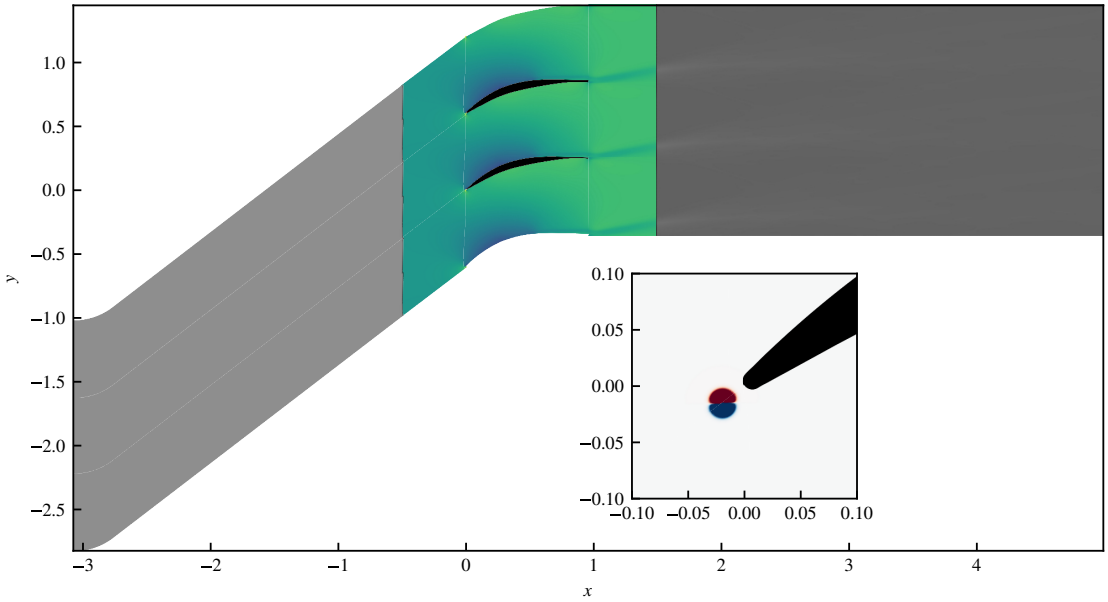


Figure 3.8: The mean flow showing the active region of the semi-norm in color. The region is taken throughout the whole height of the mesh and extends from $x = -0.5$ to $x = 1.5$. The inset axis shows the horizontal component of velocity, u, for the initial linear disturbance use for the transient response analysis.

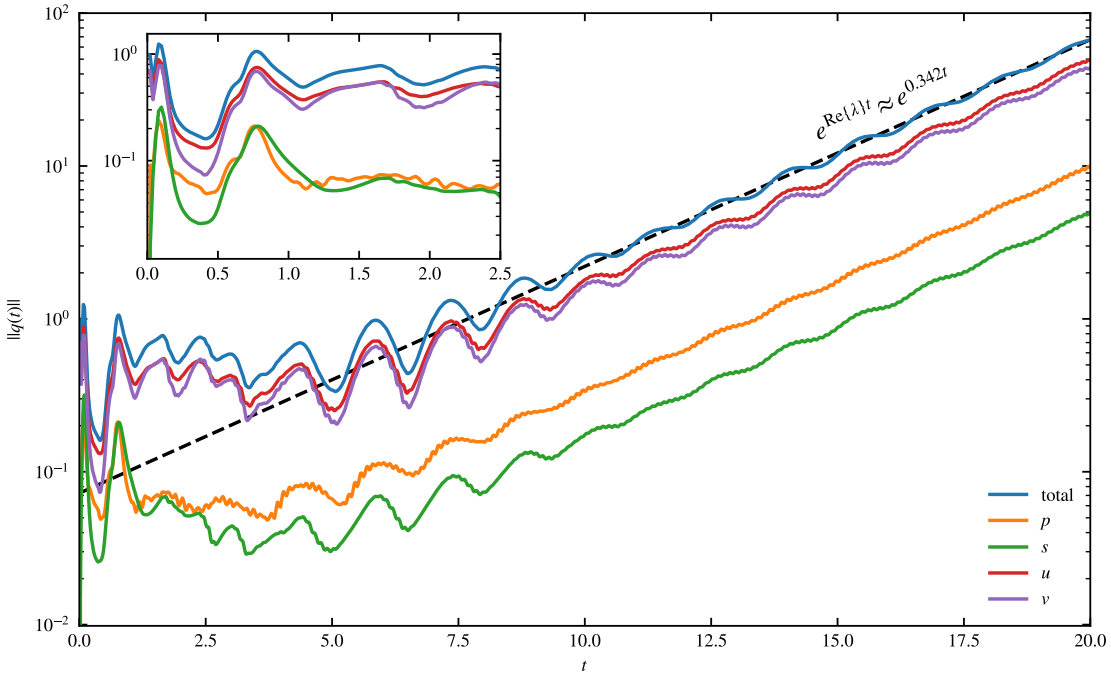


Figure 3.9: Temporal evolution of the Chu semi-norm for a divergence-free disturbance defined in Equation (3.28). Contributions of the state vector components are also shown. The unstable asymptotic growth is indicated with a dashed line set to a growth rate of $\mathrm{Re}\{\lambda\} = 0.342$.

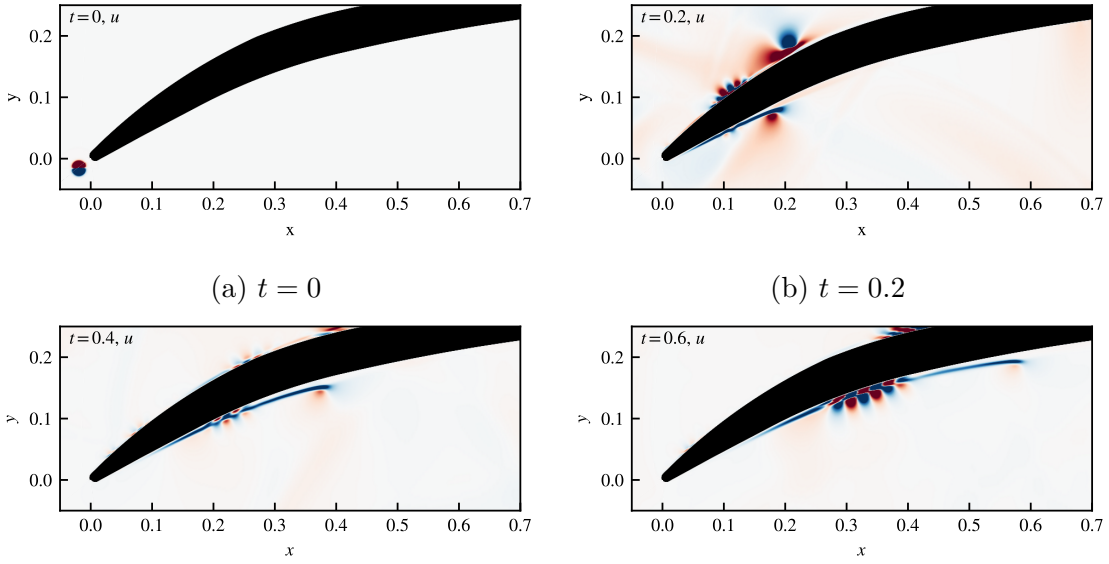
for the first portion of this time window (until approximately $t = 2.5$) the evolution of the pressure and entropy components of the norm are closely correlated. This suggests that the system response is associated with the growth of boundary layer instabilities, without corresponding noise generation. Between $t = 2.5$ and $t = 5$ the pressure component increases, indicative of new acoustic sources, but despite this increase in the pressure, the two velocity components and the entropy largely remain decaying. After $t = 5$ the system enters the phase of exponential growth punctuated by beating, where the system acts as an amplifier, and all components of the norm increase exponentially with the same growth rate of 0.342.

This breakdown into these two regimes not only highlights the non-normality of the operator, which was touched on above, but also the presence of feedback mechanisms that drive and sustain the growth. In what follows, we will consider the two phases of evolution from a mechanistic perspective, in the process we will attempt to identify the flow structures corresponding to features in Figure 3.9.

3.3.3.1 Initial Transients

Large scale transient effects first appear when the perturbation interacts with the leading edge and the resulting energy is scattered as pressure waves from this region into the far-field, corresponding to the initial spike in the semi-norm of Figure 3.9 at $t \approx 0.1$. In Figure 3.10, the remaining convective component of the disturbance subsequently splits into two vortices that are advected parallel to both boundary layers of the airfoil. Behind these, wave packets of Tollmien-Schlichting (TS) type are formed through an energy transfer from the vortices to the boundary layers. Note that between $t = 0.1$ and approximately $t = 0.4$ the system remains dissipative, and while the pressure side TS wave packet can be seen to grow, the suction side packet is highly damped as a result of the favourable pressure gradients at that point.

The dynamics of the norm change, once the wave packet reaches the separa-



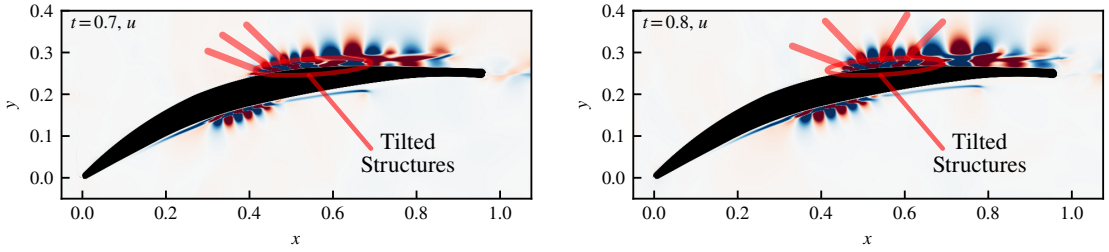
(a) $t = 0$ (b) $t = 0.2$

(c) $t = 0.4$ (d) $t = 0.6$

Figure 3.10: Initial transient responses in the vicinity of the leading edge. Note the formation of the Tollmien-Schlichting waves packets behind the split vortices propagating parallel to both airfoil surfaces. Observe also the decay of the wave packet on the suction side side and the growth of the waves on the pressure side.

tion point of the separation bubble, and enters the detached shear layer. Kelvin-Helmholtz instabilities amplify the incoming disturbance here, with the amplified features sustained until they are convected past the trailing edge, at which point they are scattered as pressure waves into the free-stream.

Meanwhile within the separation bubble, tilted structures reminiscent to those of the Orr mechanism are observed (see Figure 3.11).



(a) $t = 0.7$ (b) $t = 0.8$

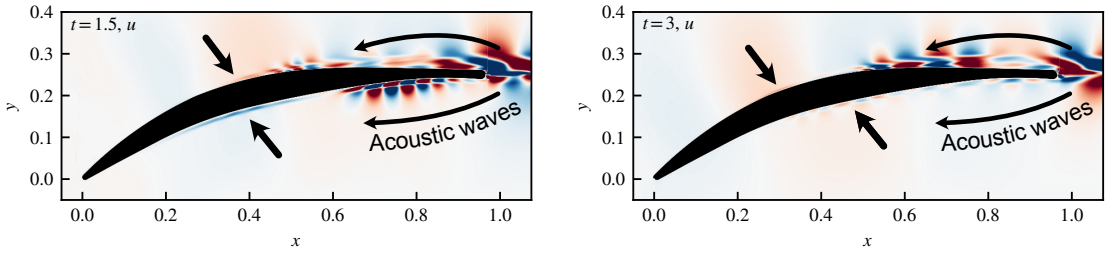
Figure 3.11: Tilted structures reminiscent of the Orr mechanism highlighted on the suction surface of the airfoil.

Briefly, in a shear flow such as the one we observe here, quantities that are

convected by the local mean flow (such as entropy or vorticity) are distorted as the faster fluid transports the portion of the flow further away from the slower fluid, and transforms them into fine-scale elongated tilted structures. By constructing a corresponding perturbation structure at earlier times that is tilted in the opposite direction to the shear, the stress terms will gradually reorient the structure so that it becomes coherent or vertical, before stretching it once more under the action of the stresses. As is summarised in [Roy and Govindarajan, 2010], this amounts to a transient growth phenomenon whereby the stresses extract energy from the mean flow so that the perturbation grows before tilting the perturbation in the direction of the flow and returning the energy back to the mean flow.

3.3.3.2 Establishing the Feedback Mechanism

Returning back to the pressure waves scattered at the trailing edge, we note that the acoustic waves are propagated predominantly upstream, as seen in the earlier nonlinear solution. This exposes the boundary layers to periodic forcing arising from the unsteady interactions between the unstable shear layer, and in the process visibly generating TS waves on the pressure surface in Figure 3.12. These waves propagate downstream before arriving at the trailing edge to be scattered into acoustic waves.



(a) $t = 1.5$ (b) $t = 3.0$

Figure 3.12: Receptivity to acoustic forcing of the upstream boundary layers on the pressure and suction surfaces, leading to the formation of TS waves.

A similar receptivity mechanism is also observed upstream of the separation bubble in Figure 3.12. The process here is different to the one on the pressure side, however, since after the waves are generated, they are amplified by the separation

bubble before propagating downstream, in a manner similar to that observed with the Orr mechanism in the section above. These amplified structures are then convected downstream, and sustained by the detached shear layer before once again scattering at the blunt trailing edge. It appears that this extraction of energy from the mean flow to the perturbation through this feedback mechanism is the driver of the exponential growth seen in the latter part of Figure 3.9.

Amplification of structures on the pressure and suction surfaces is shown in Figure 3.13, using the entropy field to isolate the convective features. At the suction surface the reattachment of the flow between the $s = 0.60$ and $s = 0.65$ demonstrates amplification by a factor of $k = 5.21$, compared to the amplification of $k = 5.07$ between $s = 0.37$ and $s = 0.60$. As a result, the processes within the separation bubble act as efficient amplifiers, with an overall growth factor of $k = 26.45$ between the origin point of the instability and the eventual structure that is convected past the trailing edge, although nonlinear processes likely regulate this growth rate in the fully nonlinear flow. Following our observation in Figure 3.12, waves are similarly seen on the pressure side, although maximum amplification of the flow now occurs in the vicinity of the mid-chord, with the probe at $s = -0.56$ recording a growth rate of $k = 2.55$ against the signal at $s = -0.37$. These waves then decay, and so the relative growth rate between $s = -0.37$ and $s = -0.76$ is only $k = 1.33$. In the future, a further resolvent analysis could be performed here to determine the transfer functions governing the growth rates of harmonic disturbances introduced within the boundary layers.

3.3.3.3 Asymptotic Behaviour

From approximately $t = 5$ until the end of the simulation, the system enters a period of unbounded exponential growth, punctuated by a beating phenomenon occurring at time scales on the order of the convection time for structures on both blade surfaces, which is also seen with the formation of wave packets in Figure 3.13.

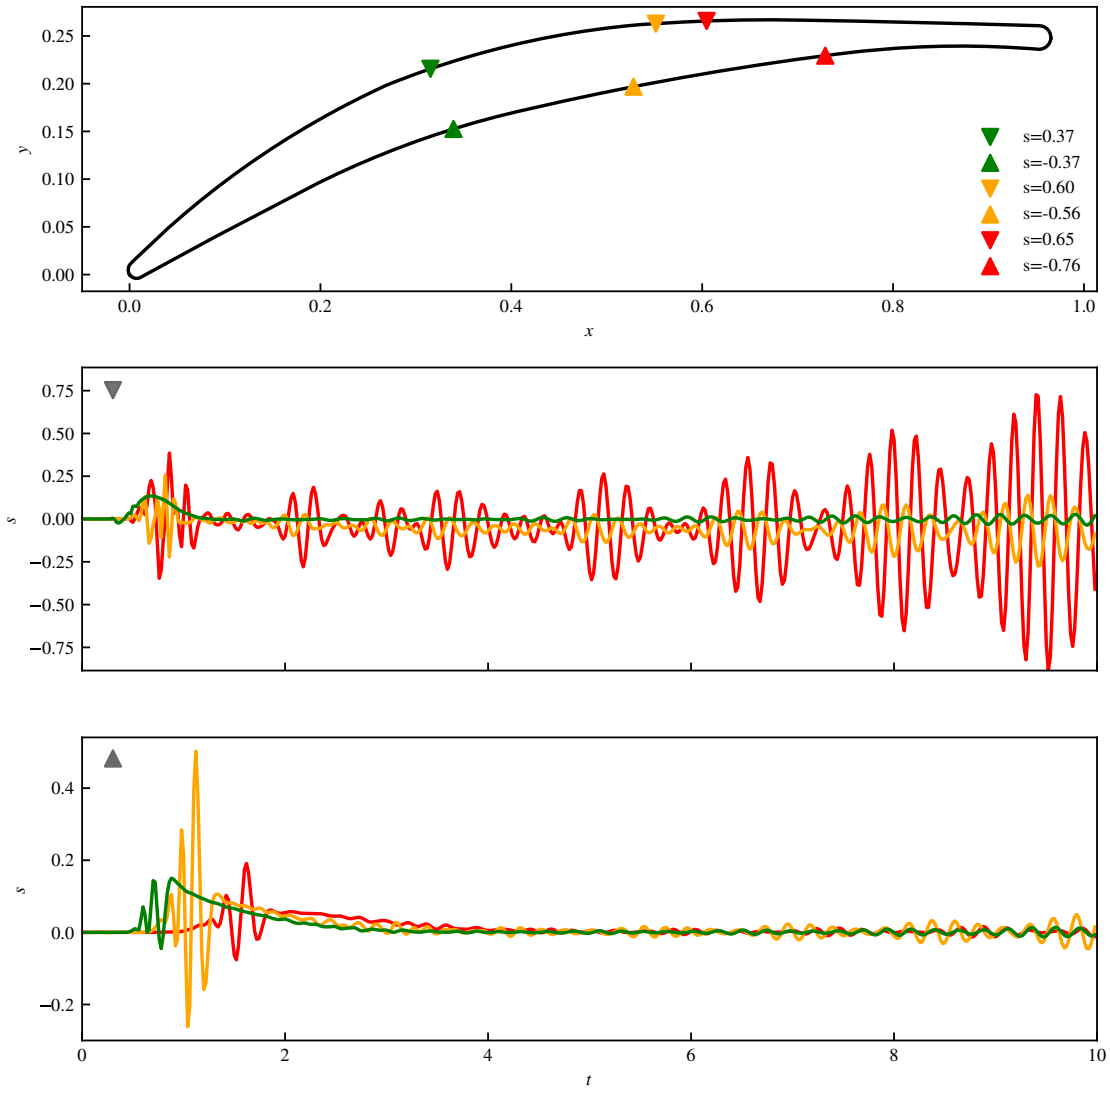


Figure 3.13: The amplification of disturbances within the pressure and suction side boundary layers for the time interval $t = 0 \rightarrow 10$, as illustrated by the entropy signal at the locations shown in the uppermost panel. Locations on the suction surface are indicated by the 'inverted' triangles and the entropy traces shown in the second graph, while the locations on the pressure surface are shown with 'upright' triangles and the results shown in the lower panel.

Having identified the asymptotic growth rate in Figure 3.9 as $\mathrm{Re}\{\lambda\} \approx 0.342$, we form $\|\mathbf{q}(t)\| \exp\{-\mathrm{Re}\{\lambda\}t\}$ to obtain Figure 3.14. From this the beating is clearly seen as a modulation of the pressure signal, with a time period of $\Delta t = 1.42$, with an expected frequency gap of $\Delta\lambda_i = 2\pi/\Delta t \approx 4.43$ between two unstable modes.

Axial velocity fields, given in Figure 3.15 for the respective 'falling' ($t = 6.0$) and 'rising' ($t = 6.5$) parts of the norm. These show that the flow alternates between

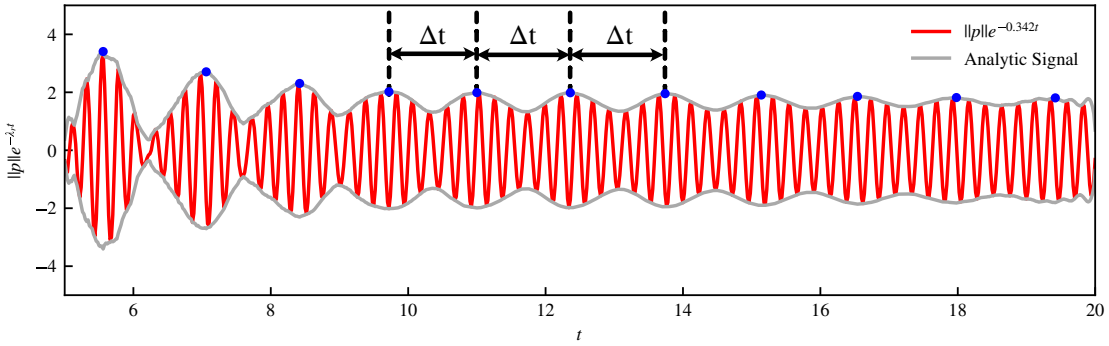
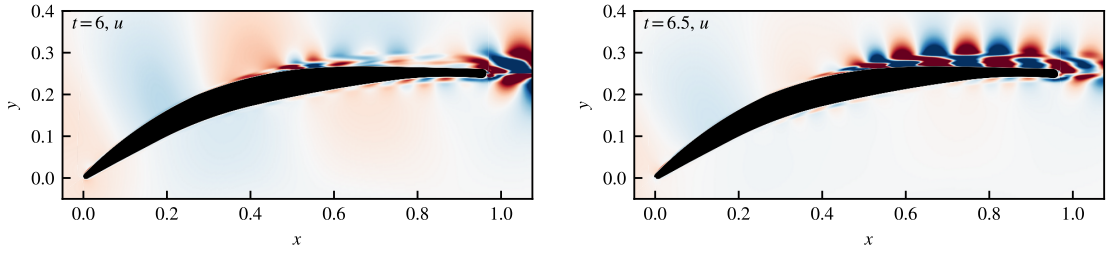


Figure 3.14: Asymptotic evolution of the pressure component of the semi-norm, showing the signal as well as the amplitude envelope function. The average peak spacing is $\Delta t = 1.42$.



(a) $t = 6$ (b) $t = 6.5$

Figure 3.15: The 'beating' phenomenon observed primarily on the suction side as the system enters the final phase of exponential growth. Note the stronger acoustic shedding at $t = 6$ and the amplified hydrodynamic features on the suction surface at $t = 6.5$.

states of highly amplified hydrodynamic features on the suction side, with low levels of acoustic radiation, and states of high acoustic excitation when these features reach the trailing edge and are scattered[14].

3.3.3.4 Comparison to Prior Work

Similar[15] impulse-response observations have also been made for isolated CD airfoils [Wu et al., 2021], where the 5-degree-angle-of-attack case produced a spanwise-homogeneous mode that exhibited a rapid decay in magnitude followed by a sustained

[14]The times of the snapshots are not the local maximum and minimum of the norm oscillation, as seen in Figure 3.14, because we want to visualise the modulation of the convective structures on the suction surface, and there is a time lag between the formation of the structure and the response of the acoustic source registered in the norm.

[15]These impulse response analyses are conducted using the forced Navier-Stokes approach rather than using a consistent linear-operator analysis, as is done here.

growth alongside a beating. This was attributed to the presence of aeroacoustic feedback mechanisms. The upstream-biased directivity of the acoustic source at the trailing edge was also noted to be similar to the observations in this section, although there also appears to be an absence of sustained hydrodynamic instabilities on the pressure surface, unlike what is observed here. An increase of the angle of attack to 8-degrees was shown to stabilise the impulse response. This is attributed to the removal of the laminar separation bubble located near the trailing edge on the suction side, and instead increasing the size of the reattaching separation bubble near the suction-side leading edge, which, in nonlinear simulations, subsequently transitions to turbulence along a large portion of the suction surface.

Impulse stability analyses of this kind are extremely rare for cascade configurations, however, with one of the few being an analysis of a single-passage high-pressure turbine (HPT) cascade [Sandberg and Wheeler, 2019]. This analysis specifically considered trailing edge blowing as a means to control the trailing edge source, by reducing the maximum velocity in the re-circulation region downstream of the trailing edge. Much like this study, the case with no trailing edge blowing exhibits a decay followed by an exponential growth of the disturbance, punctuated by beating. The mode that is spanwise-homogeneous also exhibits the highest growth rates. As the trailing edge blowing is increased, the system first transitions to a damped oscillatory mode and subsequently to a fully damped response, where the disturbance decays exponentially. The conclusions drawn by the authors included a hypothesis that the suppression of the acoustic source is associated with the removal of the recirculation region at the trailing edge, however, this disagrees with the CD-airfoil calculations [Wu et al., 2021], where the presence of the laminar separation bubble was observed to be the primary factor in determining the aeroacoustic feedback loops. This case is also interesting from the point of view of the present study because the HPT cascade contains shocks that prevent the propagation of the acoustic waves upstream. Instead, the interaction between the blades is facilitated by the propagation

of waves from the trailing edge source to the suction-side boundary boundary layer of the blade below. Therefore, unlike our study, the interaction here is exclusively unidirectional, although the limitation to a single passage restricts the observation of potential macroscopic behaviours that might exist here. By comparison, it will be shown in Section 3.4, which relaxes the restriction of this analysis to a single passage, that the subsonic linear CD cascade presented here exhibits feedback mechanisms that allow for disturbances to propagate both upwards and downwards along the axis of the cascade.

3.3.4 Operator Spectrum

The spectrum of the direct operator $\mathbf{A}$, calculated using the methods described in Section 3.2.6, is shown in Figure 3.16. As expected from our earlier analysis of the impulse response in Section 3.3.3, this spectrum is globally unstable with respect to the mean flow. It consists of two eigenvalues in the unstable half-plane, corresponding to Kelvin-Helmholtz modes on the suction surface, as well as a collection of marginally stable modes with non-zero frequencies. For some of these marginally stable modes, the mode is present in association with branch-like structures that are more heavily damped than the 'root' mode located on $\mathrm{Re}\{\lambda\} = 0$, with the angle of the branch (i.e. frequency shift of the mode with respect to the damping) corresponding to the physical location within the domain where the dynamics of the mode are represented, suggesting, and we will confirm this later, that this phenomenon arises as a result of weak resonance structures between the blade and the domain boundaries[16]. Note also that we have removed the modes with zero frequencies here since these are poorly converged as a result of the time-stepping method used here to compute the spectrum. For the purpose of this work, however, this is not a significant concern, since our goal here is to identify the key mechanisms responsible for the dominant

[16]In other words, if the acoustic component of the mode is located predominantly in the downstream, the branch is shallower with respect to the real axis as a result of the lower velocity of the mean flow in that region.

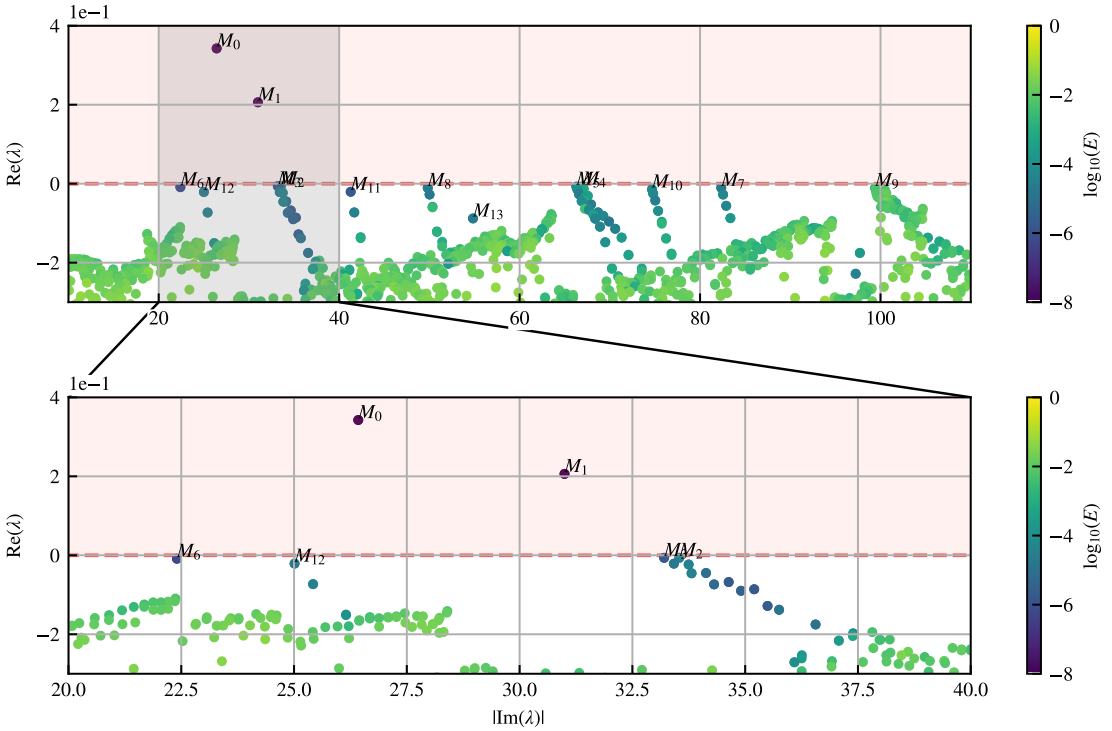


Figure 3.16: The direct global modes of the linear operator $\mathbf{A}$ with respect to the mean flow, with the inset region below showing a magnified portion of the spectrum for $D_f = \{\lambda : \ \mathrm{Im}\{\lambda\} \in [20, 40]\}$. The eigenvalues are coloured by the L_2 relative error of the eigenvalue/eigenvector pair with respect to the time-stepping operator (see discussion above). The spectrum consists of two unstable eigenvalues with the rest falling in the stable half-plane. The fourteen most unstable modes in D_f are labelled $M_0 \rightarrow M_{13}$.

features present in Figure 3.5, which are seen predominantly for $\mathrm{Im}\{\lambda\} \in [20, 90]$.

Temporal frequencies of the modes, and the relative errors between the modes and the frequencies identified in the FFT (see Figure 3.5) are calculated. The dominant peak in the FFT corresponds to the mode M_1, with a relative frequency error of 1.82%, while the modes M_0, M_2 and M_3 all have errors below 3%. Note, however, that the most unstable mode, M_0, does not correspond to the dominant peak in the FFT, suggesting that important nonlinear effects, that are not captured by the modal analysis, play a role in selecting the frequency.

Mode	Eigenvalue	FFT frequency	$\Delta \operatorname{Im}\{\lambda\}$	FFT peak index
M_0	$0.3426 + 26.4230i$	27.17	2.75%	[1]
M_1	$0.2061 + 30.9954i$	31.57	1.82%	[2]
M_2	$-0.0063 + 33.5399i$	33.30	0.72%	[3]
M_3	$-0.0064 + 33.2027i$	33.30	0.29%	[3]
M_4	$-0.0069 + 66.9689i$			
M_5	$-0.0075 + 66.2931i$			
M_6	$-0.0087 + 22.4041i$	20.73	8.08%	[0]
M_7	$-0.0104 + 82.3487i$			
M_8	$-0.0105 + 49.8186i$			
M_9	$-0.0108 + 99.4057i$			
M_{10}	$-0.0143 + 74.6719i$			
M_{11}	$-0.0205 + 41.2794i$			
M_{12}	$-0.0211 + 25.0113i$			
M_{13}	$-0.0880 + 54.8335i$			

Table 3.3: Eigenvalues of the selected modes. Modes attributable to peaks identified in the FFT are compared to the FFT results in Figure 3.5, and relative errors given.

3.3.5 Structural, Sensitivity and Receptivity Analysis

Direct Modes Having obtained the operator spectrum in the previous section, we are now in a position to analyse the structures within the eigenvectors. To do so, we restrict ourselves to modes in $|\operatorname{Im}\{\lambda\}| \in [20, 90]$, and, using the representative modes M_i, we are able to identify three main categories of modes. These are modes describing the coupling effects arising from the scattering of unsteady features from the suction side at the trailing edge; unamplified high-frequency modes from the unsteady shear layer of the separation bubble; and finally trapped modes demonstrating the possibility of self-excited resonance.

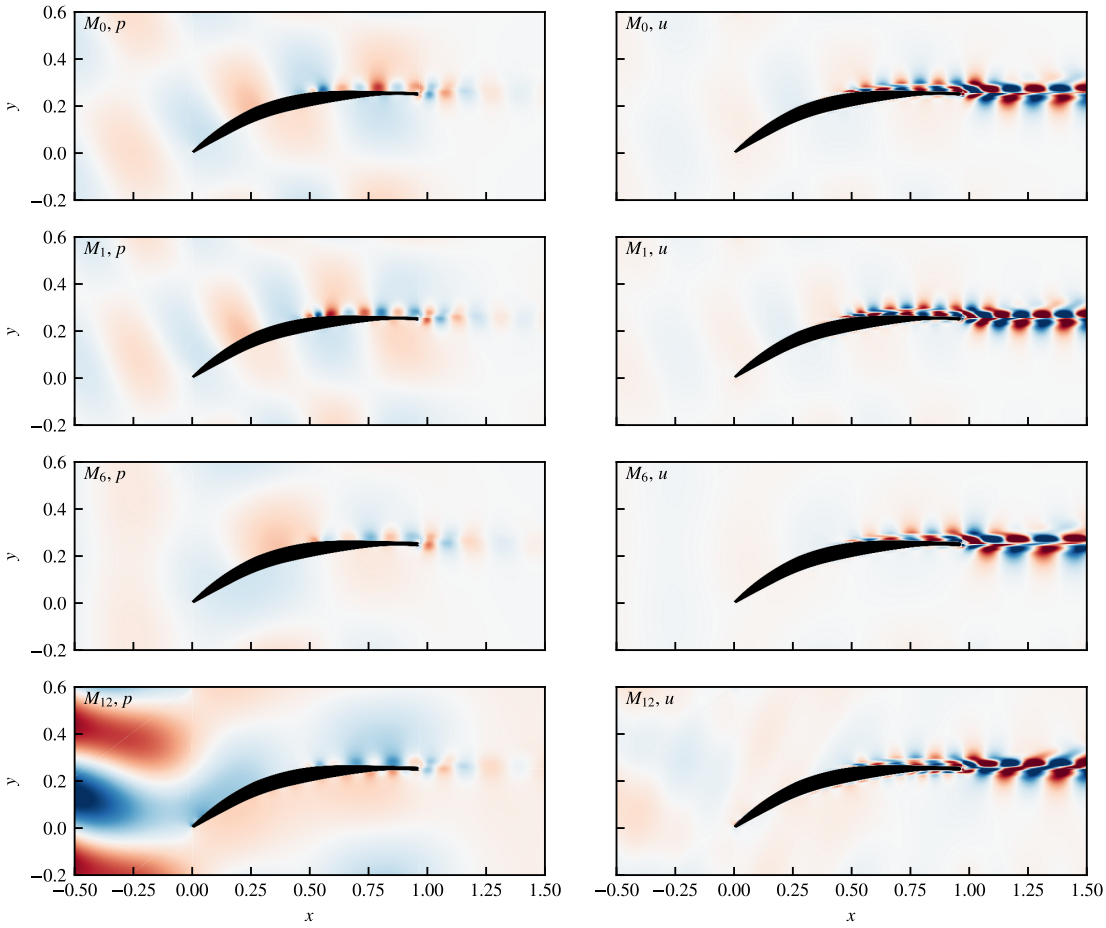


Figure 3.17: The real part of the pressure and axial velocity for the modes M_0, M_1, M_6 and M_{12} normalised with respect to the near-wake maximum axial velocity.

3.3.5.1 Dominant Modes: Bubble-TE Coupling

Modes M_0, M_1, M_6 and M_{12} represent the most dominant structures of this system, and correspond to the largest peaks in the FFT of the upstream pressure field, as established in Table 3.3.

Figure 3.17 shows the normalised pressure and axial velocity for these modes. It is clear from this figure that the modes all share a dipole-like structure at the trailing edge which propagates acoustic waves upstream through the passage and into the far-field, with minimal acoustic propagation downstream of the trailing edge. Instantaneous surface pressure distributions of these fields suggest that the modes being selected are those that are able to support small-fraction rational axial wavelengths, ranging from in-phase dynamics on each surface in M_{12} to M_0

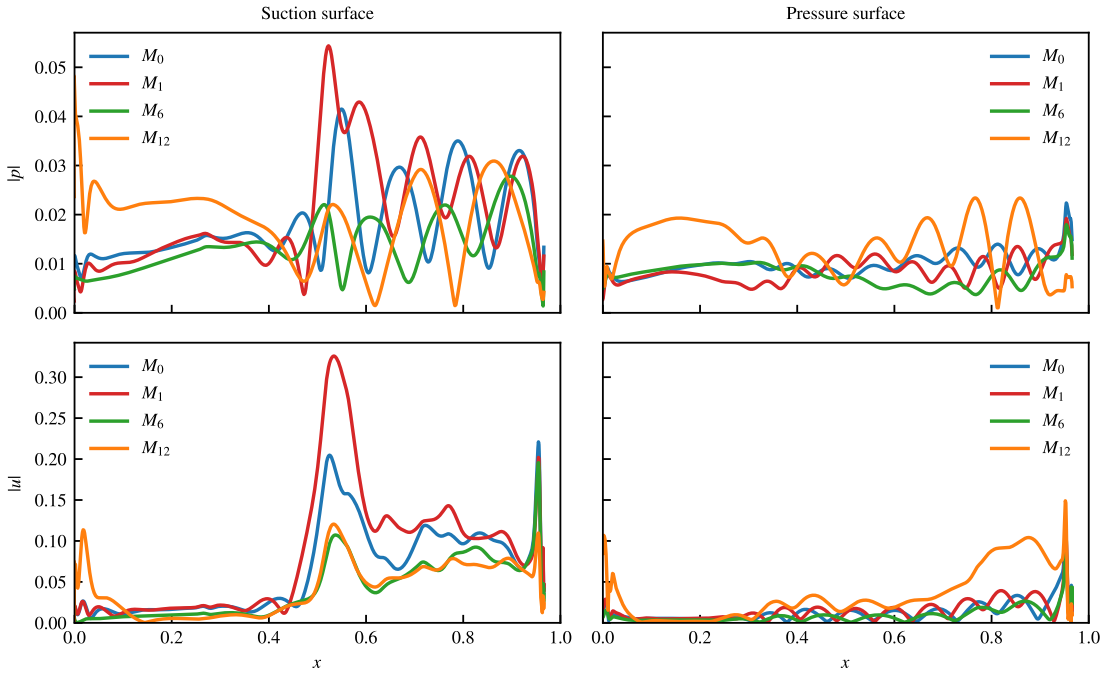


Figure 3.18: Surface pressure and velocity dynamics for the modes M_0, M_1, M_6 and M_{12}.

supporting two wavelengths on its suction side and 3/2 wavelengths on the pressure side, for example.

The wake regions for all these cases are correspondingly similar, at least up to the spatial wavenumbers of the structures. What is clear, however, is the participation of the suction side instability in these dynamics, with the appearance of the convective structures, suggestive once again of an upstream receptivity mechanism. This coupling appears to be strongest for M_0 and M_1, while M_6 and M_{12} have lower pressure and velocity amplitudes on the suction surface, suggesting that the dynamics of these modes are related more to trailing edge processes than to the amplification at the separation bubble. This is borne out in Figure 3.18, where the pressure traces confirm that the pressure distributions on the suction surface for the latter modes are amplified further downstream compared to the more unstable structures, which experience the main amplification at the separation bubble.

Interestingly, the amplification of M_1 in both the p and u fields is 1.31 and 1.59 times larger, respectively, than for M_0, and occurs marginally earlier on the suction

surface for the latter mode, but decays to a level smaller than M_0 on both surfaces by the time it reaches the trailing edge. Also, the near-wall flow on both the pressure and suction sides consists of a superposition between the vortex dynamics and the upstream-propagating acoustics field leading to nodes and anti-nodes in the pressure field. The higher frequency (both spatial and temporal) of M_1 may result in a node that is located further upstream on the laminar separation bubble when compared to M_0. This may be the mechanism through which M_1 is selected in the nonlinear field as the dominant acoustic source, although a careful nonlinear analysis is required to determine this, and is outside the scope of this study.

Pressure-side traces show that pressure signals are amplified in $x > 0.5$ for all modes except M_6, with M_{12} also experiencing amplification in the region $x \in [0.1, 0.35]$, which is unlike the other modes of this family. Spatial growth rates for M_0 and M_1 differ most following the onset of the instability for $x \in [0.2, 0.4]$, where the pressure amplitudes of M_1 are below those of M_0. However, the M_1 mode is strongly amplified from $x = 0.35$, attaining $|p| = 0.013$ although this growth is not reflected in the amplitudes of the axial velocity u, for which all modes (except M_{12}) have comparable amplitudes of $|u| \approx 0.026$.

Downstream Modes This Kelvin-Helmholtz instability is also responsible for the modes M_2 and M_3 that result in downstream propagating waves, as shown in Figure 3.19. The differences between these modes and the ones in Figure 3.17 appear to be the modification of the trailing edge shedding, with a high-pressure zone originating from the suction surface that is unmatched by a similar strength feature originating from the pressure side.

Analysis of blade-surface pressure and velocity traces (not shown here) show peak pressure amplitudes on the suction surface of $||p|| = 0.055$, which is equivalent to the maximum amplitude attained by M_1, and similarly $||p|| = 0.024$ on the pressure surface: double that of M_1. However, pressure-side velocity amplitudes for $M_{2/3}$ only have a maximum of $||u|| = 0.022$ compared to M_1, where the amplitude is

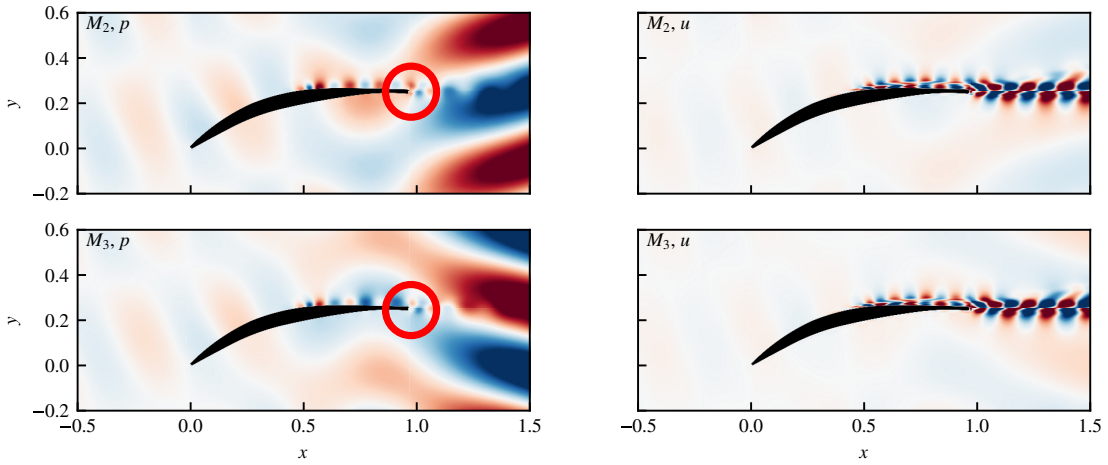
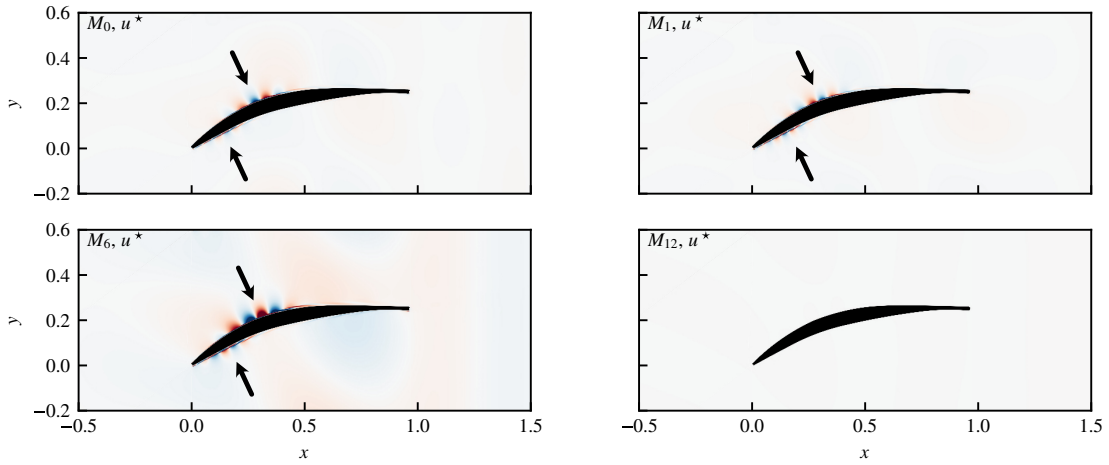


Figure 3.19: The direct pressure component of M_2, M_3, the downstream modes. The red circles indicate a modification of the vortex shedding at the trailing edge.

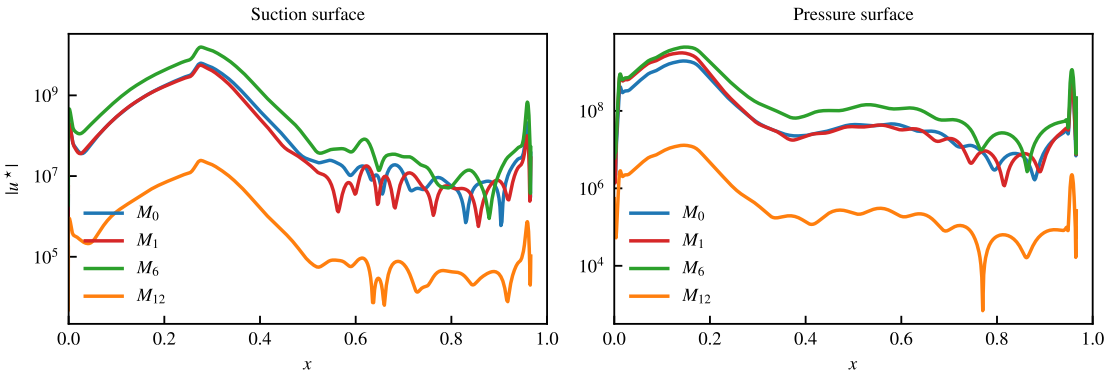
$||u|| = 0.032$. This indicates that, despite the additional forcing of the boundary layer on the pressure surface, the response of this boundary layer is diminished, which leads to this modification of the shedding at the trailing edge.

Adjoints and Wavemakers Adjoint structures corresponding to the dominant modes are shown for the streamwise velocity component in Figure 3.20a, where they are normalised so that $\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = 1$. These show that M_0, M_1, M_6 and M_{12} have compact supports for sensitivities in regions upstream of the separation bubble as well as on the pressure surface in the vicinity of the leading edge, with peak sensitivities attained at $x = 0.277$ and $x = 0.122$ on the suction and pressure sides respectively, based on the adjoint traces at the blade surface in Figure 3.20b. M_6 is found to be the most sensitive of the modes, with M_{12} the least, which may help to explain why M_{12} does not appear as a peak in Figure 3.5. M_0 and M_1 have similar sensitivities, except for a (maximum) 13% relative increase (attained at $x = 0.30$) in the sensitivity of M_0 relative to M_1 downstream of the separation point on the suction side.

Prior to advancing this discussion, we first present the wavemaker regions that locate the origin of the instability characterising the mode. Figure 3.21 shows that the



(a)



(b)

Figure 3.20: The adjoint horizontal velocity field of M_0, M_1, M_6 and M_{12} in (a), with the surface traces in (b).

wavemamakers for M_0, M_1, M_6 and M_{12} share similar supports along both surfaces of the airfoil. The overarching pattern here being regions of sensitivity upstream of the laminar bubble on the suction surface and, perhaps more surprisingly, a smaller, though similarly significant, region of sensitivity close to the leading edge on the pressure side. Numerical simulations performed earlier for the self-noise problem on isolated NACA-0012 airfoils [Fosas de Pando et al., 2017] have shown similar receptivity properties on both surfaces of the isolated airfoil, with the conclusion being made that, since the two surfaces are not *convectively linked*,[17] this constitutes

[17]In other words there is no pathway of purely convective structures forming on one of the surfaces to interact at later times with the other surface without being scattered into acoustic waves.

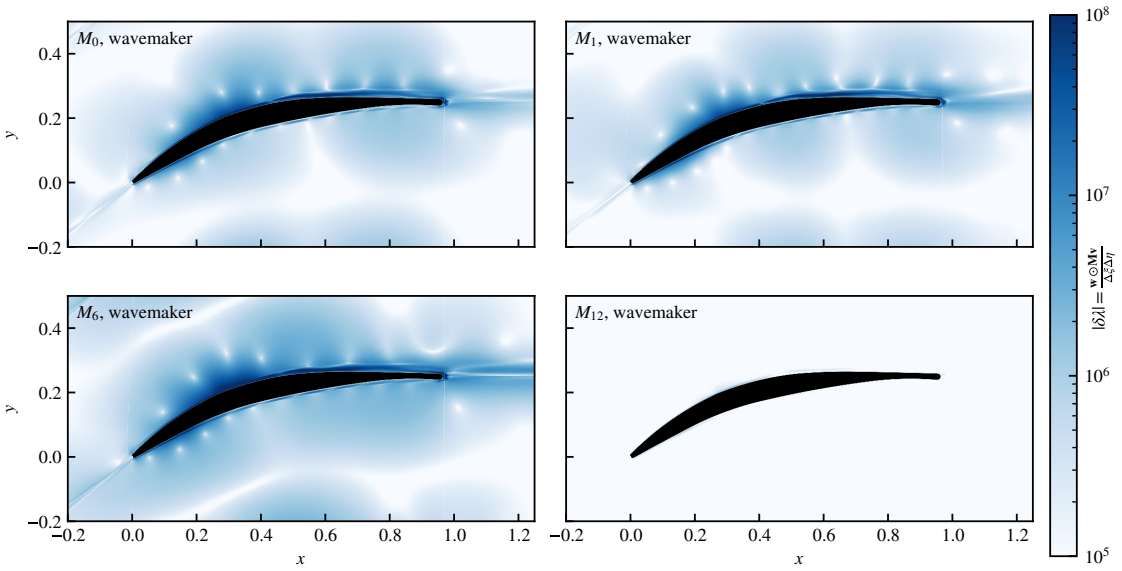


Figure 3.21: The wavemakers for M_0, M_1, M_6 and M_{12}.

as direct evidence of the receptivity of the associated boundary layers to upstream-propagating acoustic waves. Indeed we have seen in the impulse response, and from the modes in Figure 3.17, that these two regions are responsible for the generation of the Tollmien-Schlichting waves that are amplified and convected down to the trailing edge.

Application of control strategies for suppressing these dynamics would be most effective in these wavemaker regions, and would result in the greatest change to the eigenvalue.

Unlike the earlier numerical experiments on isolated airfoils, however, the pathways for the interaction are more complex in the cascade configuration. The scattered acoustic waves are now not only interacting with the blade through an upstream propagation, but also through interactions between the suction side of the dipole and the pressure side boundary layer, and vice-versa. What is more, the acoustic waves can also be scattered by the leading edge and redirected onto the pressure or suction surfaces. This introduces additional coupling effects, although we will postpone this analysis until Section 3.4.3, where the extra degrees of freedom from the N-periodic blade row will make these processes clearer.

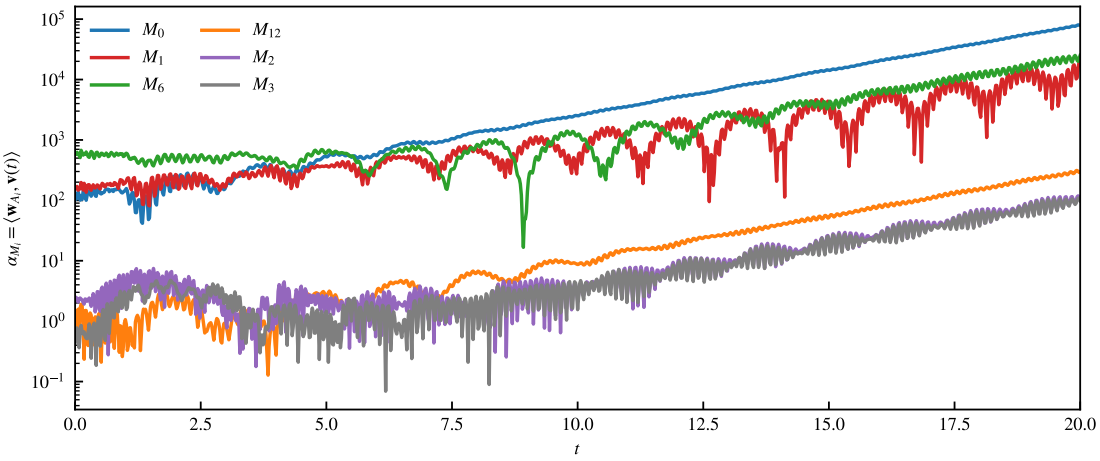


Figure 3.22: Time series of the earlier impulse response analysis projected onto selected direct modes by means of the adjoint operator.

Relationship to the Impulse Response Now that the modal dynamics have been established, we are now in a position to refer back to the impulse response analysis of Section 3.3.3 and contextualise our observations there in terms of the modal decomposition. We can do this by referring back to the discussion in Section 3.2.7, which stated that the projection of a direct vector $\mathbf{v}$ onto the direct mode $\mathbf{v}_{M_i}$ is precisely $\alpha_{M_i} = \langle \mathbf{w}_{A_i}, \mathbf{v} \rangle_{\mathbf{M}}$. Applying this projection to the impulse response time series with each of the modes discussed here, we obtain Figure 3.22. Of note here is the dominance of the M_6 mode for $t \lesssim 5$, when the larger structures, such as the TS wave packets that project more favourably onto the lower-wavenumber mode. This persists until the dynamics and feedback mechanisms for the M_0 mode are established, resulting in the growth of this mode, which the impulse response analysis above showed occurs at $t > 5$. It should be noted here, however, that for the modes $M_i, \quad i \notin [0, 1]$, that are known to have $\mathrm{Re}\{\lambda\} < 0$, the projection of the time series onto these modes does not give decaying growth rates, which is due to small numerical errors that violate the bi-orthogonality relation.

The superposition of two sinusoidal waves can be written as

$$\cos\left(\mathrm{Im}\{\lambda_0\}t\right) + \cos\left(\mathrm{Im}\{\lambda_1\}t\right) \tag{3.31}$$

$$= 2\cos\left(\frac{\mathrm{Im}\{\lambda_0\} + \mathrm{Im}\{\lambda_1\}}{2}t\right)\cos\left(\frac{\mathrm{Im}\{\lambda_0\} - \mathrm{Im}\{\lambda_1\}}{2}t\right), \tag{3.32}$$

with the envelope function taking the form

$$f_{\text{envelope}} = 2\cos\left(\frac{\Delta\,\mathrm{Im}\{\lambda\}}{2}t\right). \tag{3.33}$$

Though the pairwise frequency difference between many of the modes in Figure 3.16 is similar, only M_0 and M_1 are non-decaying and thus are the only ones able to sustain the beating for $t > 15$. From Table 3.3, $\Delta\,\mathrm{Im}\{\lambda\} \approx 4.43$, which matches the difference in frequencies of M_0 and M_1: $\Delta\,\mathrm{Im}\{\lambda\} \approx 4.57$.

3.3.5.2 High Frequency Modes: Bubble Dynamics

In addition to the dynamics examined above, the direct model similarly identifies high-frequency ($\mathrm{Im}\{\lambda\} > 50$)) structures associated with the shearing on the surface of the laminar bubble. These modes are M_4, M_5, M_7, M_8 and M_9, and are shown in Figure 3.23.

The hydrodynamics of these modes are confined primarily to the free shear layer of the separation bubble, with a peak for all modes at $x = 0.51$, and the structures downstream have negative growth rates, resulting in minimal interactions with the trailing edge, as evidenced by the u plots in Figure 3.23. Furthermore, the amplified structure on the separation bubble is approximately three wavelengths in size before decaying in the streamwise direction.

Despite this, the acoustic source nevertheless produces waves in the downstream region, for M_4 and M_5, and in the upstream region for M_7, M_8 and M_9. The relationship between these two groups appears to be the inclusion of pressure-side structures in the latter group, with amplified wave packets persisting between $x = 0.25$ and $x = 0.6$, as seen in Figure 3.24. The adjoint field (not shown) is sensitive once again at the separation point at $x = 0.30$, although now the direct mode is

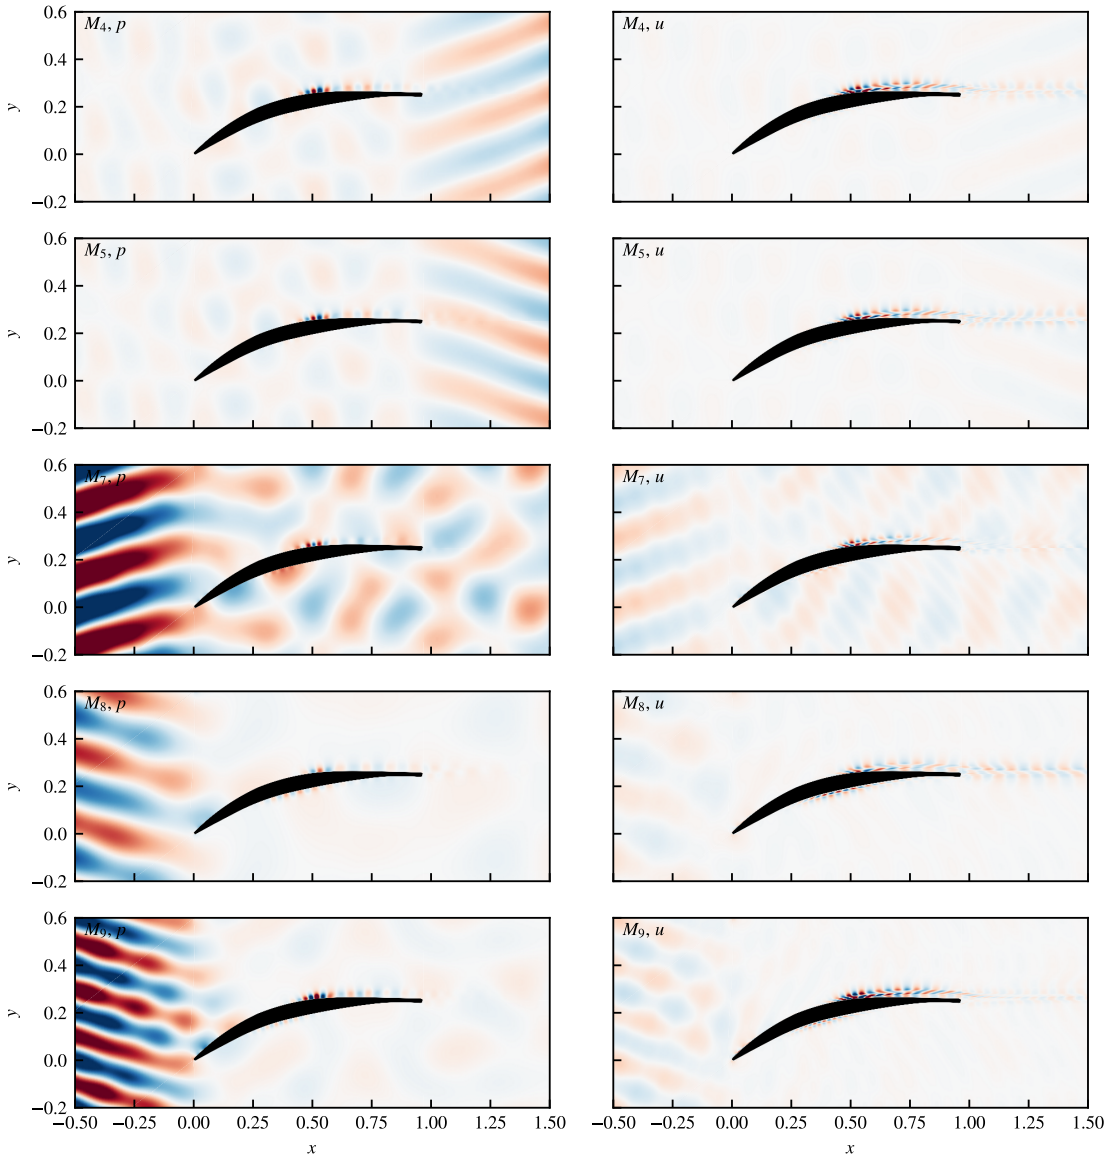
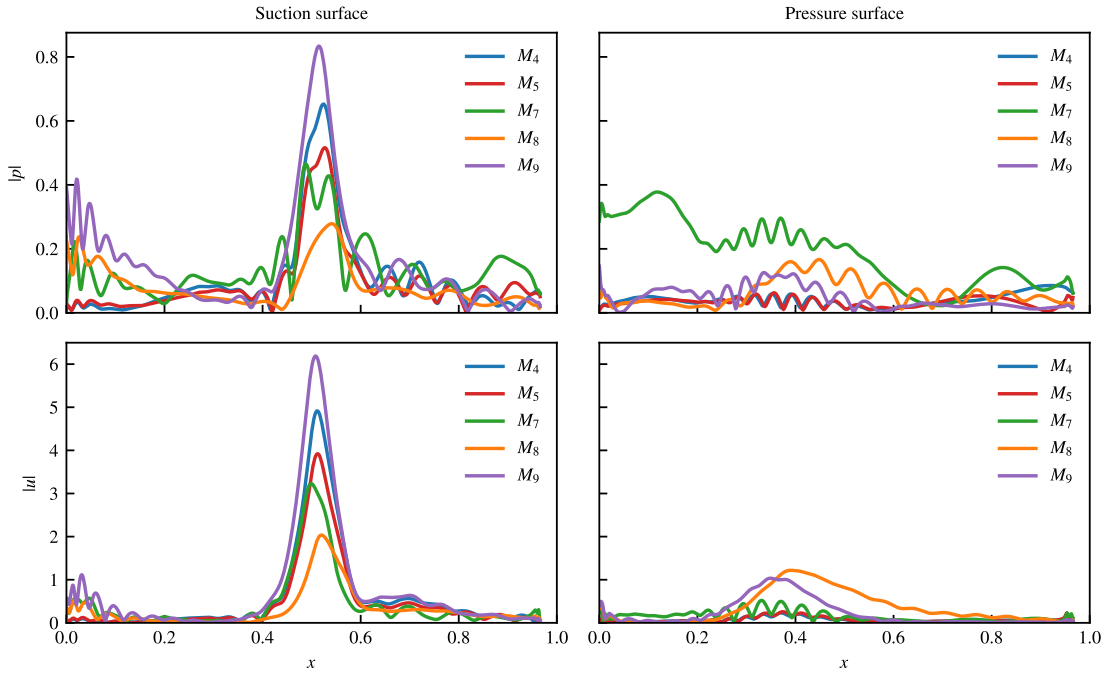


Figure 3.23: The high-frequency direct modes associated with shear layer instabilities at the separation bubble.

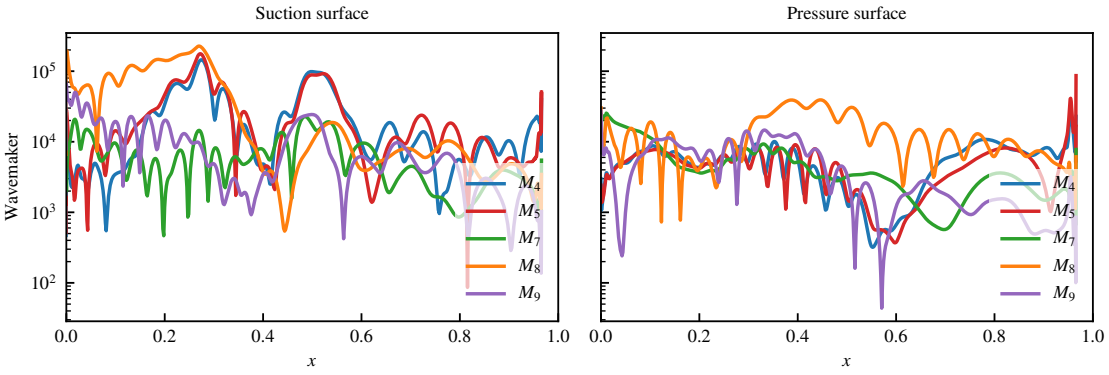
concentrated at $x = 0.51$, and hence the wavemaker in Figure 3.24b is now three orders of magnitude smaller than the wavemakers calculated for the dominant modes, suggesting that these modes are robust instabilities that are intrinsic to the flow.

3.3.5.3 Confined Modes

The final category of modes that have acoustic structures that are suggestive of acoustic resonances within the blade passage are shown in Figure 3.25, alongside



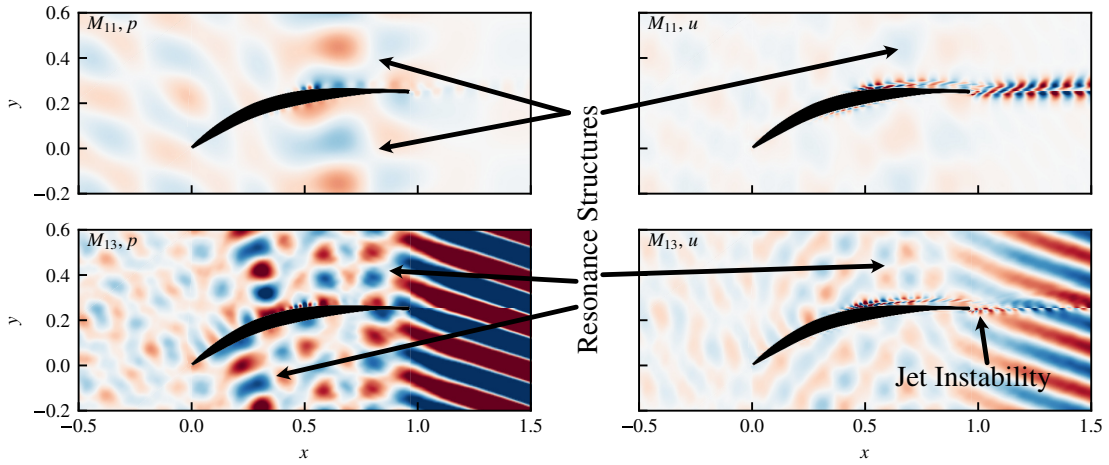
(a)



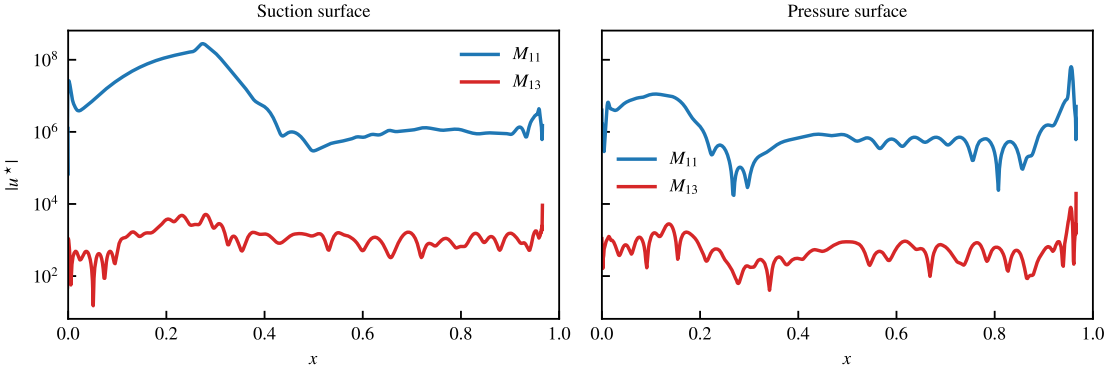
(b)

Figure 3.24: The surface traces of the high-frequency modes for the (a) direct modes, and (b) the wavemakers.

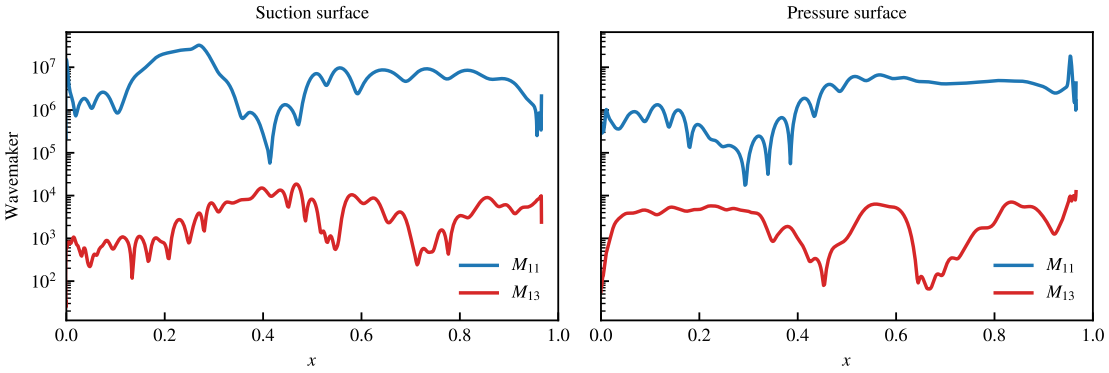
the adjoint and wavemaker surface pressure traces. Though these modes are not trapped in the sense that they still 'leak' acoustic energy into the upstream far-field, their structures contain regions of intensified pressure that are nonetheless confined between the two airfoil surfaces. Similar effects have been observed in both experimental [Parker, 1967] and in industrial settings (see review in [Hellmich and Seume, 2008]), especially in multi-stage compressors, and are highly sensitive to the particular flow conditions.



(a)



(b)



(c)

Figure 3.25: The modes M_{11} and M_{13} show confined features between the pressure and suction surfaces. The direct fields are shown in (a), with adjoint surface traces in (b) and wavemakers in (c).

Although we will not be performing a parametric flow analysis for this case, we nevertheless note that the form of the resonance in M_{11} and M_{13} are confined to a pitchwise orientation near the trailing edge, and between the separation bubble, and near the leading edge on the pressure side in the case of the former. On the other hand, from Figure 3.25, M_{13} does not appear to be exciting any instabilities on the pressure side boundary layer and is otherwise insensitive, especially when compared to M_{11}, which is confirmed by the wavemakers in Figure 3.25c. This suggests that the mechanism behind the structures described *in this particular case* is a combination of directivities of the acoustic source aligning with acoustic rays that are perpendicular, or at least close to perpendicular, to both airfoil surfaces, and this combination traps a portion of the acoustic energy inside the blade passage

As a final observation in this section, note that unlike virtually all other modes considered previously which exhibited either localised structures or had a coupling between the unsteadiness at the blade surfaces and the full wakes, M_{13} exhibits localised effects in the near wake, and is confined to instability waves along the 'jet' that was identified at the trailing edge in Section 3.3.2.

Relationship to the Nonlinear Dynamics The relationship between the dominant linear direct modes and the FFT analysis of the nonlinear field is clear to see, even through a cursory examination of the results above. For instance, the fields of M_0, M_1, M_2, M_3 and M_6 match within a maximum frequency error of 3% (see Table 3.3), and the overall structures of the modes are captured well. This suggests that, despite the action of nonlinearities in the flow, the linear receptivity mechanisms that we have identified during the course of our analysis remain mostly intact and sustain the structures even through to the nonlinear regime, and even past the point of nonlinear saturation for the unstable modes. On the other hand, while having accurately predicted the set of dominant modes, identifying the most dominant tonal noise peak with frequency $\lambda_i = 31.57$ in Figure 3.5 has proven to be a greater challenge with the field corresponding to M_1 rather than M_0, which is

otherwise suggested by a more positive real-part of the selected eigenvalue. With the tools presented in this paper, identifying the causes of this change is difficult, and a more detailed analysis is required.

The linear and nonlinear models similarly agree on the compactness of the tonal noise phenomena with the leading modes confined to the interval $|\operatorname{Im}\{\lambda\}| \in [20, 80]$, which is precisely the region seen in the nonlinear analysis in the frequency domain. Despite this, it would be unwise to generalise this finding to arbitrary configurations, since the appearance of other small-scale features, like separation bubbles near the leading edge, will create additional high-frequency sources that must be captured with much finer meshes in order to resolve the corresponding instabilities.

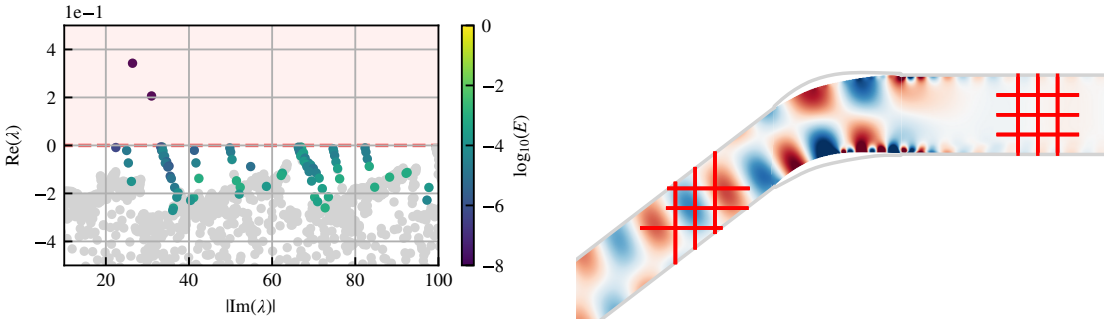
3.3.6 Phase and Group Velocity Analysis

One final question that we wish to address in this section is how the modes in the Figure 3.16 are related to each other, and whether or not they superpose to form larger structures that dominate the global dynamics.

To do this, direct modes are first filtered by taking into account only those where the convergence error is $E < 1 \times 10^{-3}$, which leaves 156 modes with non-zero frequencies, shown in Figure 3.26a.

The instantaneous wavenumber, k, of each mode is computed using a Pisarenko Harmonic Decompositon (PHD), which is a super-resolution method for frequency calculation of signals contaminated with noise. These calculations are performed by constructing a series of perpendicular contours in the regions upstream and downstream of the blade, as shown in Figure 3.26b, onto which the acoustic field for the calculation is interpolated. To automate the process, the exact form of each mode is not assumed a priori, and so two additional parallel contours are used at each point in each direction to minimise the possibility that the contour falls into a node of the acoustic field which would not correctly identify erroneous frequencies.

The contours are subdivided into 60 equispaced arc-length segments with a



(a) (b)

Figure 3.26: To the left, the modes selected for the phase and group velocity analysis are shown in colour while the rejected modes are represented as grey points. To the right, the interpolation contours along which the wavenumber is computed. The middle vertical contour is placed at $x_c = -1$ in the upstream and $x_c = 2$ in the downstream, with the neighbouring contours placed either side at $x_\pm = x_c \pm 0.15$, and extend to the full height of the passage. The horizontal contours are similarly equispaced in y, with y_c, the central contour, placed at the midpoint of the central vertical contour. The total length of these is fixed at 0.6 and the midpoint adjusted so that the entire contour fits into the single-passage flow domain.

two-dimensional piecewise cubic spline using the Clough-Tocher scheme (see `scipy interpolate` documentation) interpolating the grid data onto this new contour. As a result, the Nyquist frequency is $k_n = 50$. The PHD is applied to this equi-spaced contour with two complex exponentials. From our earlier observations in Section 3.3.5.1 and in manual inspections of these modes, we observed that the pressure field in the regions surrounding the contours tends to be relatively regular and exhibits a dominant wavenumber in its structure. To that end we make the assumption that the instantaneous frequency should be constant along each of the contours, and so for each mode and contour we average to obtain the corresponding averaged wavenumber, along with the relative standard deviation/coefficient of variance (RSD/CV), c_v of the raw data. The resulting dispersion relations are plotted below for each type of contour where $c_v < 0.1$ in Figure 3.27. For each of these we observe two dominant features. First that groups of modes occur in tight clusters with similar phase velocities, $v_p = \omega/k$, and secondly that for almost all the modes there is a constant group velocity, $v_g = \mathrm{d}\omega/\mathrm{d}k$, confirming that these modes form a

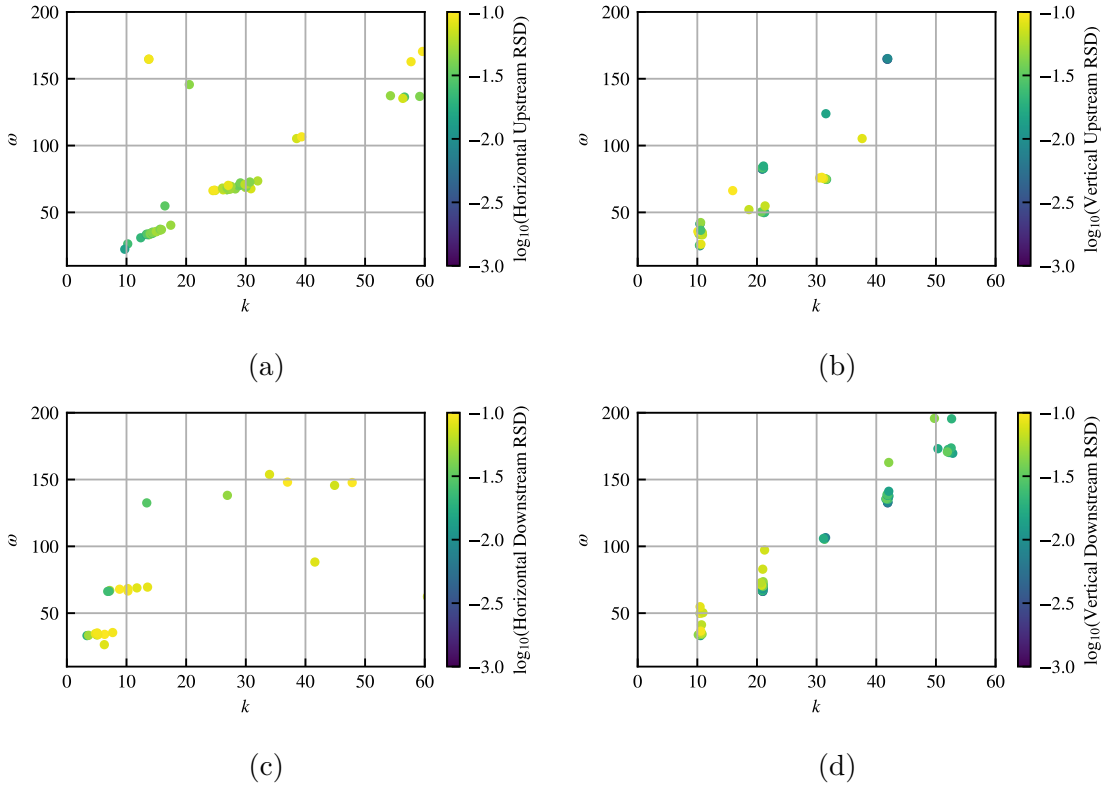


Figure 3.27: The dispersion relation for each type of contour. Modes are filtered according to the standard deviation of the instantaneous wavenumber along the contour so that only the monochromatic modes with non-zero k are registered, and the average k computed. Colouring of these modes is by relative standard deviation (RSD) of k.

large structure under superposition that propagates together.

Identification of modes with consistent phase velocities is performed using the density-based clustering algorithm `DBSCAN` [Hahsler et al., 2019], which is able to identify clusters of modes based on the density of the data as well as being robust to outlying points. Applying this to the contours in Figure 3.27a and Figure 3.27d, where the clustering of the modes is most visible, we obtain the following clusters, in Figure 3.28, from which it appears that M_1, M_2 and M_3 (along with their axial resonant structures) propagate as one, with the other modes superposing to form a large wavepacket as a result of the constant group velocity. It is important to remember, however, that the disparate growth rates of the modes within this structure mean that most of the modes will decay quickly, leaving only the superposition of M_0 and M_1, at least in the linear regime.

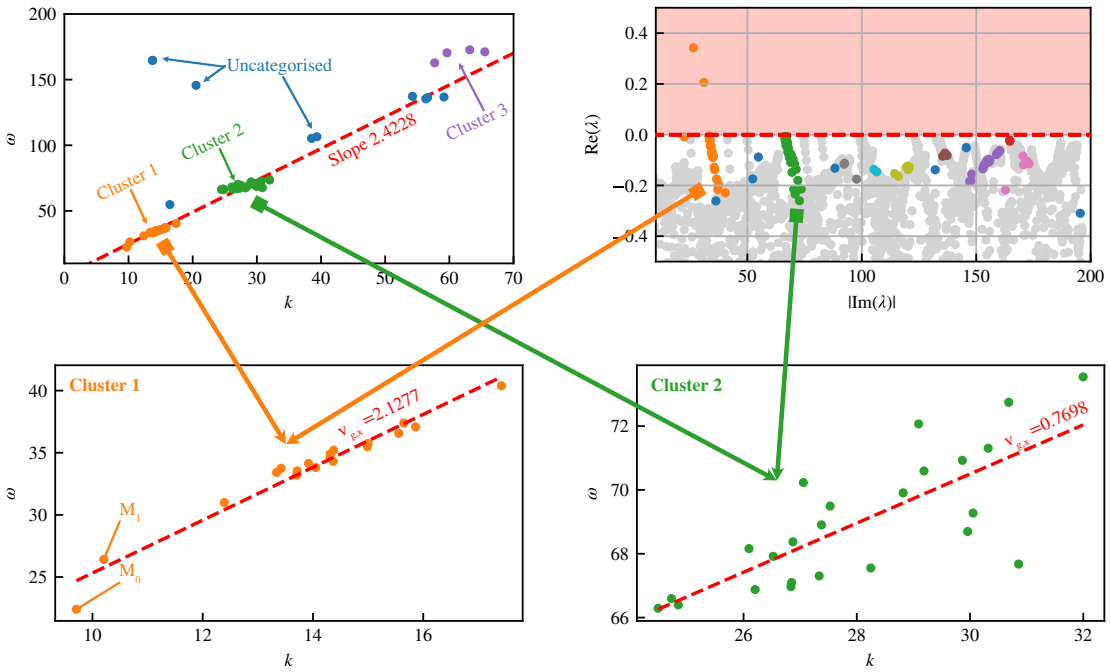


Figure 3.28: Showing the clustering of the modes obtained with DBSCAN, the corresponding modes and detailed $k-\omega$ plots of the two main clusters identified in this analysis.

On the other hand, in the nonlinear case we have seen that the growth rates of the modes are not necessarily a good way of predicting which modes are selected by the nonlinear processes. Instead, those with the sensitive wavemakers are the ones that appear able to sustain the dynamics into the nonlinear regime.

3.4 Multiple Passages

So far in this work we have focused exclusively on single-passage dynamics, which has provided us some valuable insights into local stability characteristics of the flow over the single blade. Realistic annular blade rows, however, typically contain on the order of 10 to 50 blade passages and are capable of supporting large-wavelength disturbances, synchronisation and standing waves that are impossible to resolve in a simulation of a single periodic passage. In this section, we consider a 10-passage configuration and extend the global stability analysis above to this case. We achieve this by using a N-periodic matrix decomposition method, that allows large system

matrices to be decomposed into smaller independent problems that are simpler to solve. With this direct modelling of entire blade rows, we will then investigate the effects of aerodynamic mistuning on individual cascade elements using this framework.

3.4.1 The N-Periodic Geometry

For the purpose of this study, we will restrict ourselves to a 10-passage linear cascade, and make use of the model compressor passage geometry from Section 3.3 that is tessellated in the vertical direction as shown in Figure 3.29. This is to minimise the computational costs. We again apply periodic boundary conditions in the upstream and downstream regions of the mesh corresponding to passages P_0 and P_9.

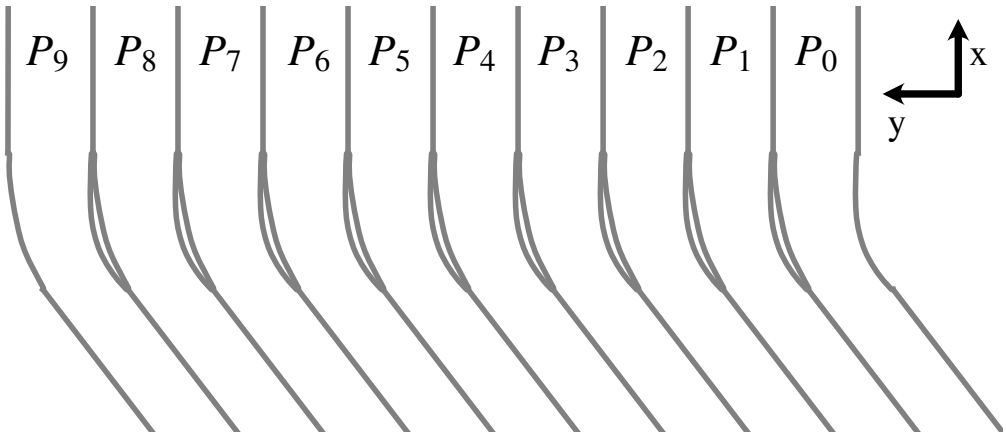


Figure 3.29: The physical 10-passage geometry considered in this section. The axial lengths of the upstream and downstream sections of this geometry have been chosen so that wavelengths of up to $5h_p$ could be represented within the simulation. The figure has been rotated to save space.

3.4.2 N-Periodic Modelling

While in theory it is possible to model entire N-periodic configurations directly by simply modelling N blade passages in a single mesh, this becomes impractical as the number of passages is increased. This arises independently of the parallel scaling performance of the code, since the main constraints come from having to store

large snapshots of data for the Krylov-Schur algorithm in Random-Access Memory (RAM), and finding computational resources on which instantaneous access to $N \times$ (number of processors per passage) processors is available for extended periods of time. More pertinent to the analysis itself, however, is that larger time intervals between the snapshots may be required to converge the low-frequency modes with the exponential time stepping method, since with larger blade counts structures with very low azimuthal wavenumbers and low temporal frequencies are expected, resulting in convergence issues without prior adjustment of the time step. Although this may seem like a hopeless endeavour, by making a couple of sufficiently general assumptions these issues are easily mitigated.

Linearising the N-periodic system we once again arrive at the IVP

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \mathbf{A}\mathbf{v}, \quad \mathbf{v}(0) = \mathbf{v}_0, \tag{3.34}$$

where $\mathbf{A}$ and $\mathbf{v}$ now represent the $nN \times nN$ and nN size matrix and vector, respectively, with n being the total degrees of freedom in each passage and N the number of passages. We will make two assumptions about this system:

- that the blade surface geometries, and by extension the meshes, for each passage are identical, and;
- that the mean flow within each passage is the same.

Grouping the entries of $\mathbf{A}$ and $\mathbf{v}$ by passages, so that $\mathbf{v} = (\mathbf{v}_0, \mathbf{v}_1, \ldots, \mathbf{v}_{N-1})^\intercal$, with $\mathbf{v}_j$ denoting the portion of the state vector in the j^{th} passage, under these assumptions, transforms the state matrix into block-circulant form

$$\mathbf{A} = \begin{pmatrix} \mathbf{A}_0 & \mathbf{A}_1 & \cdots & \mathbf{A}_{N-1} \\ \mathbf{A}_{N-1} & \mathbf{A}_0 & \cdots & \mathbf{A}_{N-2} \\ \vdots & \ddots & \ddots & \vdots \\ \mathbf{A}_1 & \cdots & \mathbf{A}_{N-1} & \mathbf{A}_0 \end{pmatrix}. \tag{3.35}$$

We can interpret this form of $\mathbf{A}$ as local dynamics of a given passage in $\mathbf{v}_j$, governed by $\mathbf{A}_0$, and influence of all the other passages $\mathbf{v}_{i \neq j}$ effected through the off-diagonal

matrices $\mathbf{A}_{i\neq 0}$ on the dynamics of $\mathbf{v}_j$, so that

$$\frac{\mathrm{d}\mathbf{v}_j}{\mathrm{d}t} = \mathbf{A}_0\mathbf{v}_j + \sum_{i\neq j} \mathbf{A}_{(i-j) \mod N}\mathbf{v}_i. \tag{3.36}$$

In other words, the number of off-diagonal non-zero blocks is determined by the spatial coupling between the j^{th} passage and its neighbours. Generally this is imposed by the spatial differentiation scheme as in the case of finite-difference and finite-volume schemes, or through the boundary nodes of neighbouring elements in a finite-element formulation, for example.

Many applications, including this present one, exhibit either nearest-neighbour coupling or a weak global coupling. This means that the only passages that contribute to the solution are the dynamics in the passage above and below[18]. Hence, Equation (3.35) becomes a *block-tri-diagonal* system with only $\mathbf{A}_{N-1}$, $\mathbf{A}_0$ and $\mathbf{A}_1 \neq 0$. Although smaller, the system is not yet decoupled, however, and further manipulation is required to make progress.

Fortunately, among the many properties of circulant matrices is the diagonalisation transformation that takes a block-circulant matrix into block diagonal form (see e.g. [Schmid et al., 2017]). Picking the matrices so that $J \in \mathbb{C}^{N\times N}$ and $P \in \mathbb{C}^{nN\times nN}$,

$$\mathbf{J}_{j+1,k+1} = \rho_j^k \Big/ \sqrt{N}\,, \quad \rho = e^{2\pi i j/N} \tag{3.37}$$

$$\mathbf{P} = \mathbf{J} \otimes \mathbf{I}_{n\times n}, \tag{3.38}$$

and for any matrix function, f, of the system matrix is diagonalised as

$$\mathbf{P}^H f(\mathbf{A})\mathbf{P} = f\left(\hat{\mathbf{A}}\right), \tag{3.39}$$

[18]In the case of weak global coupling, assuming nearest-neighbour coupling would introduce a small error. Indeed, in many numerical codes spatial coupling is introduced through the derivative, which is known to be an essentially local property. Therefore, if the passage height between blades is sufficiently large, this error will typically be below that of the machine precision. Note that this an assumption about the structure of the code and not the physical system.

where

$$f\left(\hat{\mathbf{A}}\right) = \begin{pmatrix} f\left(\hat{\mathbf{A}}_0\right) & & & \\ & f\left(\hat{\mathbf{A}}_1\right) & & \\ & & \ddots & \\ & & & f\left(\hat{\mathbf{A}}_{N-1}\right) \end{pmatrix}, \tag{3.40}$$

and moreover,

$$\hat{\mathbf{A}}_j = \mathbf{A}_0 + \rho_j \mathbf{A}_1 + 1/\rho_j \, \mathbf{A}_{N-1}. \tag{3.41}$$

In a matrix-free setting we are only concerned with matrix-vector products such as $\hat{\mathbf{A}}_j \hat{\mathbf{v}}_j$, and so from Equation (3.41), it is clear that this amounts to taking the three matrix-vector products $\mathbf{A}_0 \hat{\mathbf{v}}_j$, $\mathbf{A}_1 \hat{\mathbf{v}}_j$ and $\mathbf{A}_{N-1} \hat{\mathbf{v}}_j$.

While performing each of these calculations individually is possible, to reduce the computational overhead, we instead linearise around a three-passage geometry with identical mean flows in each passage, giving a $3n \times 3n$ block-circulant matrix

$${}^3\mathbf{A} = \begin{pmatrix} \mathbf{A}_0 & \mathbf{A}_{N-1} & \mathbf{A}_1 \\ \mathbf{A}_1 & \mathbf{A}_0 & \mathbf{A}_{N-1} \\ \mathbf{A}_{N-1} & \mathbf{A}_1 & \mathbf{A}_0 \end{pmatrix}. \tag{3.42}$$

Multiplying by the augmented vector ${}^3\hat{\mathbf{v}}_j = (\mathbf{0}, \hat{\mathbf{v}}_j, \mathbf{0})$ gives

$${}^3\mathbf{A}\left({}^3\hat{\mathbf{v}}_j\right) = \begin{pmatrix} \mathbf{A}_{N-1}\hat{\mathbf{v}}_j \\ \mathbf{A}_0\hat{\mathbf{v}}_j \\ \mathbf{A}_1\hat{\mathbf{v}}_j \end{pmatrix}, \tag{3.43}$$

and it is now easy to combine into the form of Equation (3.41). Assuming that all passages have the same processor distributions, the local solution of the top and bottom passage is simply transferred to the corresponding processor on the middle passage, pre-multiplied by ρ_j or ρ_j^{-1} and added to the local solution of the corresponding destination processor (i.e. the processor where the final solution is assembled). This is easily achieved in MPI by using `MPI_Cart_create` and a

sequence of `MPI_SendRecv` calls, and has the advantage of using pre-existing linear codes with minimal programming intervention required by the end user. A diagram of this procedure is shown in Figure 3.30. Figure 3.31 shows how these coupling

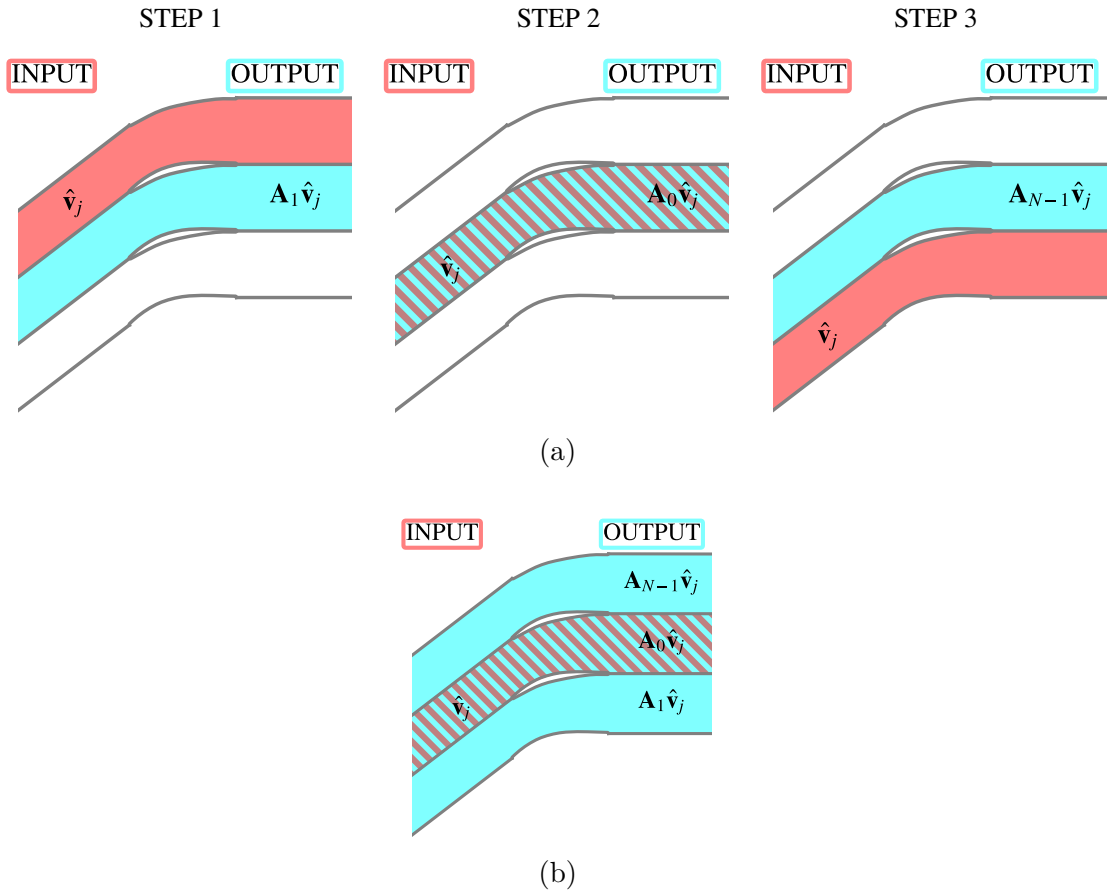


Figure 3.30: Evaluation of the matrix-vector product $\hat{\mathbf{A}}_j\hat{\mathbf{v}}_j$ using a three-passage linearisation. The three-stage direct approach making use of sequential matrix-vector evaluation in (a) can be simplified in (b) by combining the three stages into one matrix-vector evaluation. Here the required matrix-vector products are confined to the respective passages and may be combined using MPI calls between the corresponding portions of the field. Note also that $\mathbf{A}_1$ and $\mathbf{A}_{N-1}$ are mostly sparse, so additional performance gains may be made by factoring this in, however this is outside the scope of this study.

fields, $\mathbf{A}_1\hat{\mathbf{v}}_j$ and $\mathbf{A}_{N-1}\hat{\mathbf{v}}_j$, manifest themselves in the periodic blade row that is considered here. A linearisation of the three-passage system is performed and a random field for $\hat{\mathbf{v}}_j$ is introduced into the middle passage. In the figure, the solution in the middle passage is removed to visualise coupling fields, and to demonstrate how these are reintroduced into the central passage. This case also serves as a validation

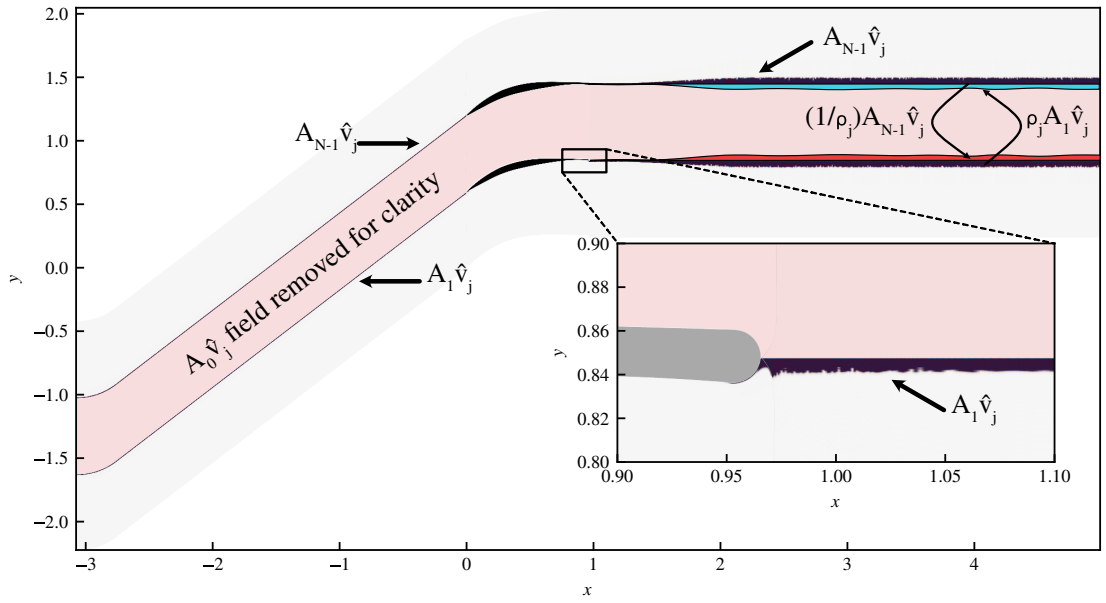


Figure 3.31: The coupling fields $\mathbf{A}_1\hat{\mathbf{v}}_j$ and $\mathbf{A}_{N-1}\hat{\mathbf{v}}_j$ are shown for the computational, three-passage mesh. The assembly procedure taking these fields into the middle passage is also shown.

of the adjoint, with a corresponding adjoint matrix-vector product computed and compared using the familiar duality relation, giving a relative error of 4.571×10^{-14}.

In essence, the large original problem in Equation (3.34) is recast as N independent 3-passage simulations, since now we only have to solve problems of the kind

$$\frac{\mathrm{d}({}^3\mathbf{v}_j)}{\mathrm{d}t} = {}^{3\to1}\mathbf{C}\big({}^3\mathbf{A}\big)\big({}^3\mathbf{v}_j\big), \tag{3.44}$$

where the matrix ${}^{3\to1}\mathbf{C}$ represents the communications in Figure 3.30b, the assembly of the sum Equation (3.41) and zeroing-out the top and bottom passages to return the resulting evolution vector to the standard triplet form, in other words: $\mathrm{d}({}^3\mathbf{v}_j)/\mathrm{d}t = ({}^3\,\mathrm{d}\mathbf{v}_j/\mathrm{d}t)$.

Without optimisation the computational cost is nominally three times larger for the full periodic system than the full computation, however these N simulations may be run independently, with lower demands on the computational resources at any given time and with a lower RAM requirement. Further improvements to the

overall cost can also be made by carefully examining the exact form of the coupling matrices $\mathbf{A}_1$ and $\mathbf{A}_{N-1}$, since these are expected to have substantial sparsity, and also carefully examining the code structure that leads to the coupling matrices. As an example of the latter, where ghost regions are used, the coupling matrices are essentially a proxy for the communications between the ghost cells and the interior of the neighbouring domain in the triplet. Instead of communicating with the neighbouring mesh, however, the procedure can be modified for a single periodic passage so that the ghost cell portion of the data is intercepted, multiplied by the appropriate root of unity and then applied to the destination processors across the periodic boundary. Doing this would bring the total cost from $3N$ to just N as no additional mesh segments have to be considered. On the other hand, the cost of significant additional development time is a consideration that must be taken into account, especially since many existing code structures would have to be rewritten.

In the following we use the same mean flow and passage geometries that have been computed for the single passage to keep the analysis consistent with the work performed earlier in Section 3.3. The only alteration between our two cases has been the reduction in the size of the sponge layer, in the linear analysis, from $\Delta x_{\text{sponge}} = 2.0$ to $\Delta x_{\text{sponge}} = 1.0$, so that the larger wavelengths can be represented with reduced dissipation within the domain[19]. We compute all the results in this section using the N-periodic decomposition method detailed above.

3.4.3 Impulse Response Analysis

Following the methodology set out in Section 3.3, we begin with an impulse response analysis. Since the local dynamics near the blade surface have already been established, we turn our attention instead to the large-scale structures present in

[19]This has the effect of marginally increasing $\text{Re}\{\lambda\}$ for M_6, making it unstable. We should note too that, in the context of a mean flow stability analysis, this small change in growth rates is largely unimportant because dynamics of the problem are defined predominantly by the frequencies and structures of the global modes rather than the growth rates. Mode frequencies are unaffected by the changes in the parameters of the sponge layer.

this 10-periodic configuration. For ease of comparison, the initial condition is kept identical to the single-passage case but is applied only between the third and fourth passages (P_2 and P_3 in Figure 3.29) in the configuration, as shown in the upper left panel of Figure 3.32. Of course there is a greater flexibility in this particular case regarding the choice for the exact form of the initial perturbation, but it has been kept consistent with Section 3.3.3 here so that the mechanism by which the instability develops within the perturbed passage, and how it communicates and interacts with the neighbouring passages, is made clearer and easier to interpret. The other panels in Figure 3.32 show the disturbance evolving over the same $t = 20$ time scale used in the single-passage, for an improved comparative analysis.

Initially the perturbation evolves in much the same way as in Section 3.3.3, with scattering from the perturbed blade as the disturbance is advected into the path of the airfoil leading edge. This is followed by the now familiar growth of Tollmien-Schlichting wave packets and subsequent amplification by the Orr mechanism at the separation bubble. The difference here, at least for $t \leq 1.0$, is that while the convective disturbance is left at the surface of the perturbed airfoil, the pressure waves scattered from the leading edges can now propagate and interact with neighbouring blades, at which point secondary reflections and further scattering of the pressure waves from the other blades can be observed at $t = 0.5$. Furthermore, through the acoustic receptivity mechanisms identified in the previous section, these secondary scattered waves propagate along the airfoil surfaces and introduce additional convective instabilities on the affected airfoils, which can be seen at the trailing edges at $t = 1.0$. These are nevertheless dominated by the blade to which the perturbation is applied, since most of the energy is retained in the convective structures rather than the scattered acoustic waves of the initial pulse.

When the disturbance reaches the trailing edge the acoustic source becomes active, which once again occurs at approximately $t = 2$, based on the increase in the pressure component of the norm in Figure 3.33 above the level of the entropy

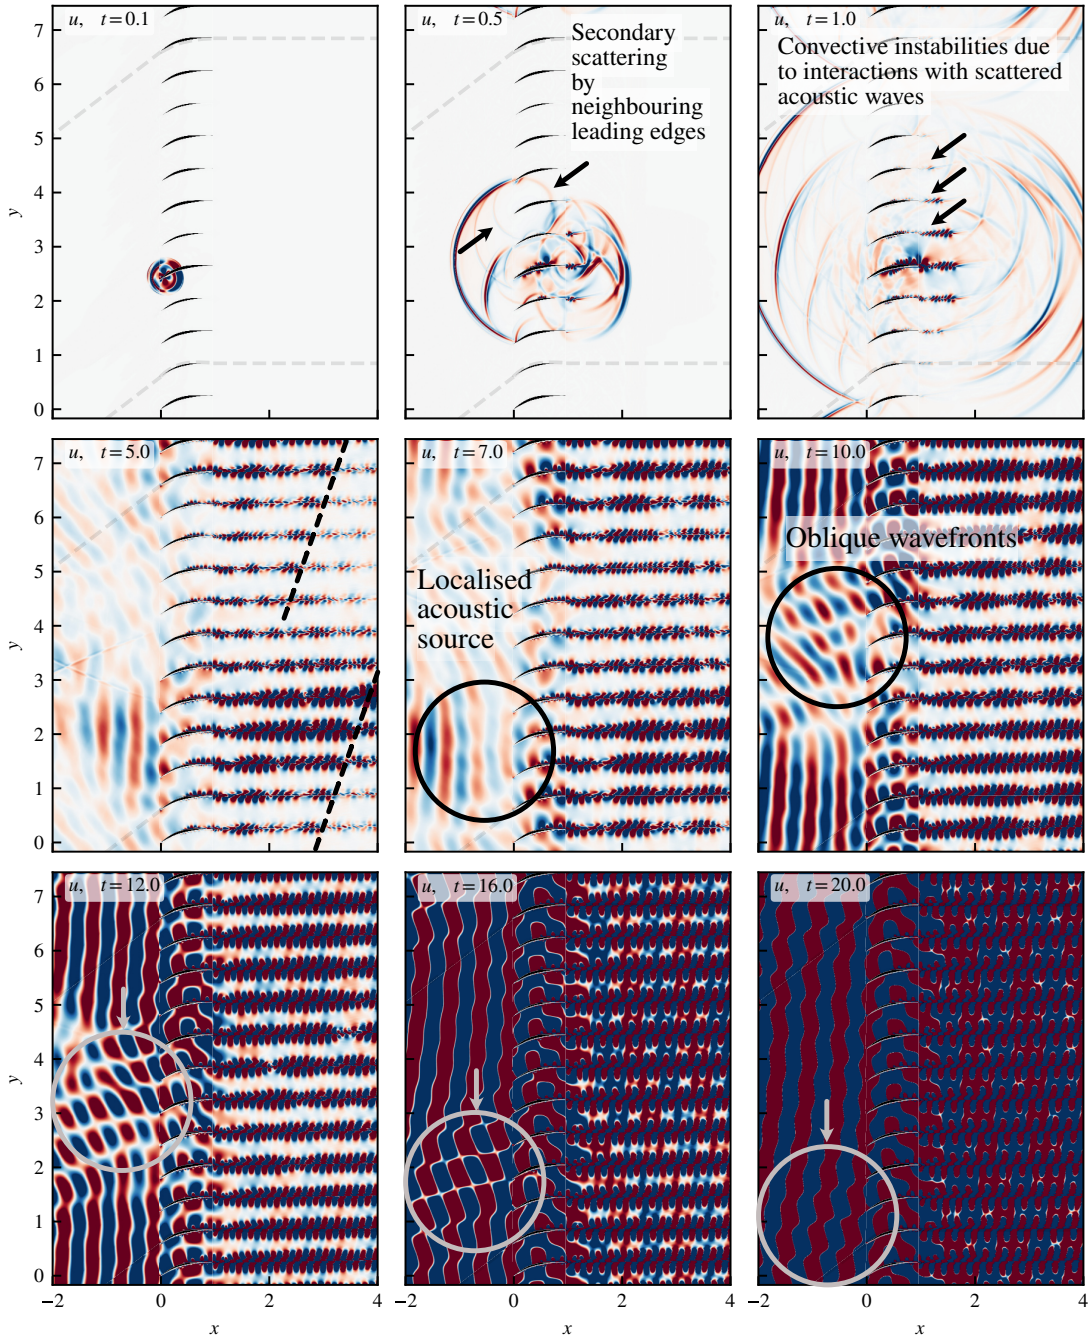


Figure 3.32: The evolution of the transient response to the divergence-free perturbation. See the supplementary material.

component. As before the acoustic radiation exhibits a directivity pattern that is focused primarily in the upstream direction, resulting in the interaction of the acoustic waves with the regions of sensitivity identified in the wavemaker analysis in Section 3.3.4. This time, however, these interactions occur on the perturbed blade as well as on the pressure surface of the blade above and on the suction surface of the blade below. Indeed, aside from the small amounts of scattering of the acoustic waves by the leading and trailing edges of the unperturbed blades, the primary forcing that spreads the initial disturbance to the neighbouring passages occurs through this mechanism. This is confirmed by localised growth of the disturbance despite the acoustic waves propagating through the whole domain at time scales much shorter[20] than the observed evolution of the perturbation.

As we noted in Section 3.3.5.1, both the pressure and suction surfaces have regions of high receptivity to acoustic waves with frequencies associated with the Kelvin-Helmholtz instability (i.e. $\mathrm{Im}\{\lambda\} \approx 31$), with the suction side being the more sensitive of the two regions. Although we expect modifications to the M_0 mode of the single passage, due to the removal of the restrictive single-passage periodicity, we hypothesise, and will later confirm, that the receptivity mechanism for this mode will be similar to those identified in Figure 3.21, since the mean flow is identical in both cases. This explains why the disturbance propagates downwards through to the passage below, while the growth in the passages above remains constrained. This can be seen by drawing a line between comparable features in the downstream wakes, and is shown for $t = 5.0$.

At intermediate time scales, between approximately $t = 5.0$ and $t = 10.0$, the main receptivity mechanism begins to operate and propagate the disturbance throughout the domain. This can be seen in Figure 3.32 by the successive growth in the lengths of the wakes from the perturbed blade and downwards (shown with a line for the $t = 5.0$ panel), although at this point the asymptotic ($t \to \infty$) phase

[20]The acoustic waves propagate at speeds of $u \pm Ma \in [0.7, 1.3]$, while the convective speed is $Mc \approx 0.3$.

relationship in the pressure dynamics between the blades is not yet fully established. Localised wave packets of intensified acoustic radiation are observed periodically emanating from the blades followed by periods of quiescence (see the region $x < 0$, $y \in [0, 3]$ at $t = 5.0$ and $t = 7.0$ of Figure 3.32), which is consistent once again with the beating mechanism observed previously and occurring, according to Figure 3.33, at similar time scales to those observed in Section 3.3.3, suggesting that this growth occurs through local feedback mechanisms on the surfaces of the blades observed in the single-passage setting.

Here we also observe the establishment of large-scale structures. Throughout Figure 3.32 a slowly moving region of approximately five to six passages propagates downwards, characterised by wavefronts angled at approximately 45 degrees from the axis of the blade row. Although a definitive description awaits further analysis, the fact that it appears above the perturbed blade suggests that the mechanism at play is related to scattered acoustic waves, as these are the only structures capable of bypassing the asymptotically dominant receptivity mechanism, which, as we have mentioned extensively, is mostly a localised interaction.

We can explain the persistence of this structure by noting that it, like the dominant mechanism, propagates downwards one passage at a time, potentially through similar receptivity mechanisms as discussed above, given the similarity of the acoustic wavelengths between M_0 from Section 3.3.5.1 and this structure. As such the time scales for the feedback are similar, and so the 'group velocity' of both structures through the cascade are comparable.

The de-phased wavepacket is eventually dominated by the structure with a constant acoustic source phase relationship, as is seen developing on the blade surfaces at $t = 20.0$. This seems reasonable because each blade has identical receptivity characteristics, by construction. Structures with constant phase relationships between the passages are therefore expected to dominate, with frequencies and growth rates dictated by the geometry of the problem.

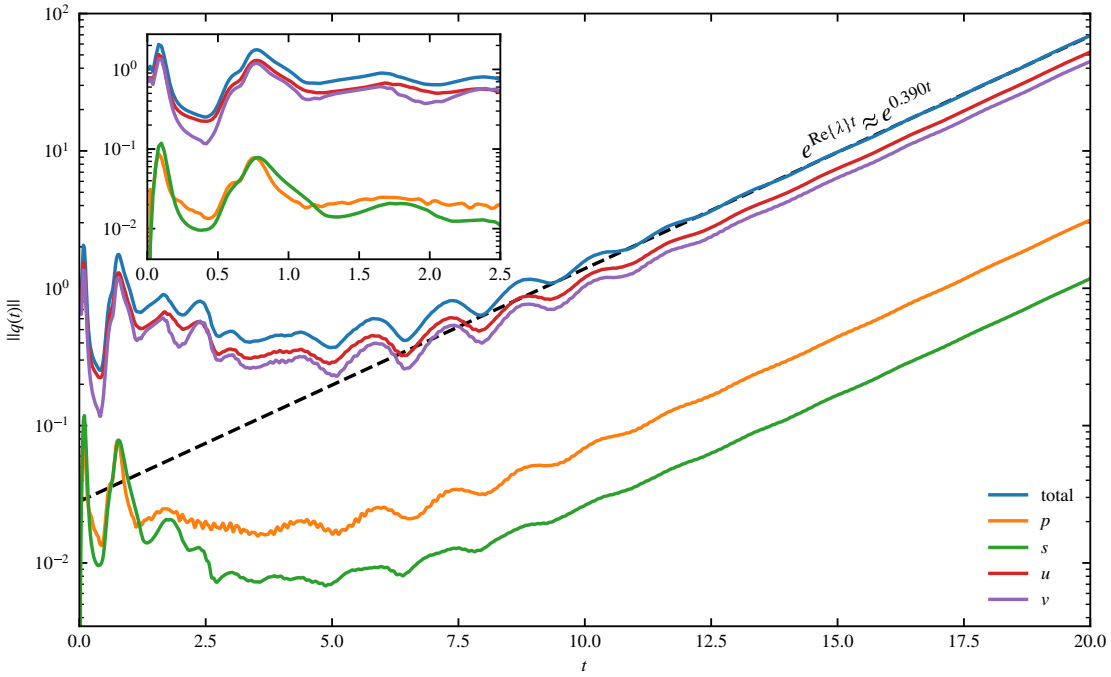


Figure 3.33: Evolution of the 10-passage norm, showing the initial transient dynamics for $t = 0 \to 5$ and the asymptotic regime for $t = 10 \to \infty$.

For asymptotically large times ($t \to \infty$), the most unstable eigenvalue once again dominates. Based on our earlier investigations we expect this system to be unstable, with *at least* the same growth rate as the single passage, following our observations of the spectral structure of the N-periodic system in Section 3.4.2. In actual fact, an exponential fit of the norm growth rate between $t = 10 \to 20$ shows a growth rate of $\mathrm{Re}\{\lambda\} = 0.390$, which is higher than that observed in the single passage. The corresponding snapshot at $t = 20$, in Figure 3.32, confirms that this is a full ten-periodic structure and that the mode is, at least qualitatively, similar to that seen in the single passage with a near-constant phase-shift between each consecutive passage. This justifies this analysis of the N-periodic configuration, since it clearly demonstrates that this system contains dynamics that are substantially different from, and impossible to model in, the single-passage. Modelling larger sectors of the bladed disk are naturally expected to produce additional, larger-period modes and behaviours, but this paper concentrates on a ten-periodic sector in order to reduce the associated (and sizable) computational costs.

3.4.4 Global Stability

Global spectra are obtained directly from the definition above, using the triplet system in Equation (3.43), by applying the SLEPc methodology in Section 3.2.6. We observe [Schmid et al., 2017] that, for an eigenvalue-eigenvector pair $(\lambda_j, \mathbf{v}_j)$ of the triplet system for a root of unity ρ_j, the vector $\mathbf{v} = \left(\mathbf{v}_j, \rho_j \mathbf{v}_j, \dots \rho_j^{N-1}\mathbf{v}_j\right)^\intercal$ is an eigenvector of the full N-periodic system matrix $\mathbf{A}$, with similar results for the adjoint operator following directly.

Note that, since $\mathbf{P}^H\mathbf{A}\mathbf{P} = \hat{\mathbf{A}}$, taking the Hermitian conjugate of both sides gives $\mathbf{P}^H\mathbf{A}^H\mathbf{P} = \hat{\mathbf{A}}^H$, showing that, for one, the same process can be used to compute the adjoint alongside the direct eigenvectors, and that, since $\mathbf{w} = \left(\mathbf{w}_j, \rho_j \mathbf{w}_j, \dots, \rho_j^{N-1}\mathbf{w}_j\right)^\intercal$, in the inner product $\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}}$ the conjugation of $\mathbf{w}$, would result in the multiplication $\bar{\rho_j}\rho_j = 1$ in each passage. It follows therefore that the *wavemaker is invariant with respect to the passage*, and so the wavemakers in the following will be presented within a single passage domain to save space.

3.4.4.1 Direct Spectra

The direct spectrum is calculated here for the roots of unity $\{\rho_j\}_{j=1}^{9}$, and focusing primarily on the modes with frequencies $|\mathrm{Im}\{\lambda\}| < 60$, $\mathrm{Im}\{\lambda\} \neq 0$. This is because we have identified the most unstable eigenvalue to be in this region in Section 3.4.4, and because we expect large-scale, low frequency structures (with wavelengths on the order of $10h_p$), such as the one observed in the upstream region of the impulse response, to manifest with low temporal frequencies (e.g. $|\mathrm{Im}\{\lambda\}| \lesssim 20.0$).

Instead of considering the modes for each ρ_i sequentially, we are instead interested how a given mode in the one-periodic case (i.e. the fields corresponding to modes in ρ_0) changes based on the periodicity of the problem as we move ρ_j along the boundary of the unit circle. In other words, the N-periodic analysis is a map of the unit circle to the closed curve $\mathrm{eig}\left(\hat{A}\left(e^{i\theta}\right)\right)$, onto which the roots of unity $\rho_j = e^{2\pi i \theta_j / N}$ are mapped.

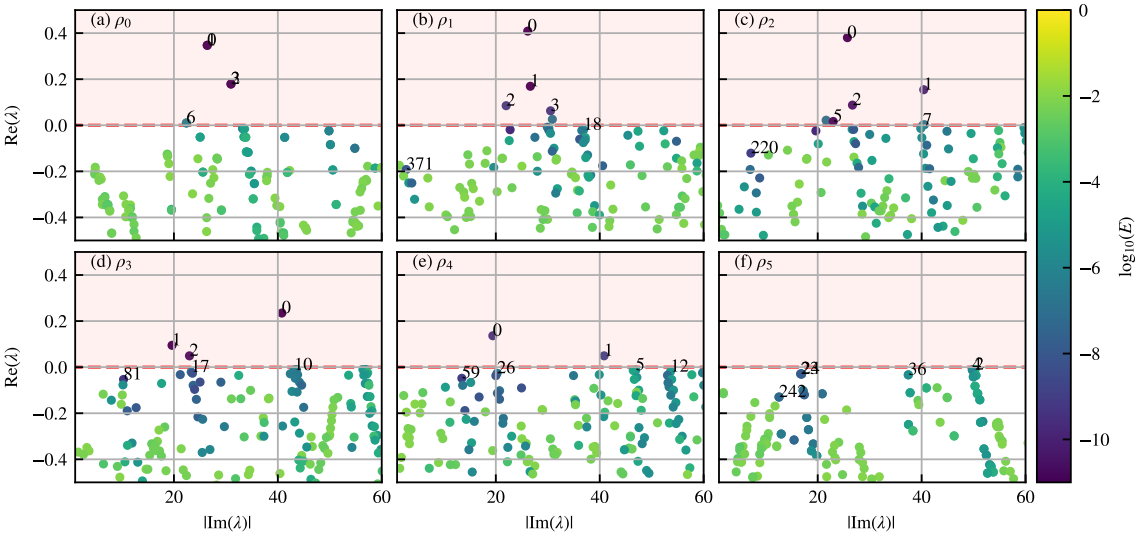


Figure 3.34: The 10-periodic direct modes for the 5 non-symmetric roots of unity. Note the absence of any new unstable structures for ρ_5, which means that the 2-periodic system has the same unstable modes as the single-passage case.

We can illustrate this mechanism on a smaller toy problem by plotting out the symbol curve of $\hat{\mathbf{T}}(z) = \mathbf{T}_0 + z\mathbf{T}_1 + \mathbf{T}_{N-1}/z$, for points on the unit circle $z = e^{i\theta}$, representing all the possible roots of unity, for all possible $N \in \mathbb{N}$. With an arbitrary choice of 3×3 matrices[21], $\hat{\mathbf{T}}(z)$ is shown for the three eigenvalues in each case in Figure 3.35. Each of the eigenvalues follows its own path through the complex λ-space. This example clearly shows that while the 'single-element' problem is stable, the addition of other 'elements' can lead to the system becoming unstable with the dominant structure in the dynamics alternating between different mode families (e.g. compare the most unstable eigenvalue in $N = 2$ against $N = 3$ in Figure 3.35).

For the airfoil case, the problem appears to be effectively symmetric about ρ_5, which may be due to the flow being essentially uniform and parallel to the mesh boundaries at which the periodic boundary condition is applied, although the causes of this warrant future investigation. As a result, we will look only at modes for $\rho_0 \to \rho_5$ in the preliminary analysis presented here.

[21]The toy-problem matrices are $\mathbf{T}_0 = -\begin{pmatrix} 3/2 - 2i & 0 & 0 \\ 0 & 1/2 + 4i & 0 \\ 0 & 0 & 3 \end{pmatrix}$, $\mathbf{T}_1 = -\begin{pmatrix} -1 & 1 & -1 \\ 0 & -1 & 1 \\ -1 & 0 & 1 \end{pmatrix}$ and $\mathbf{T}_{N-1} = -\begin{pmatrix} 0 & -1 & -1 \\ 1-i & 0 & 1 \\ -1 & 1-i & 0 \end{pmatrix}$

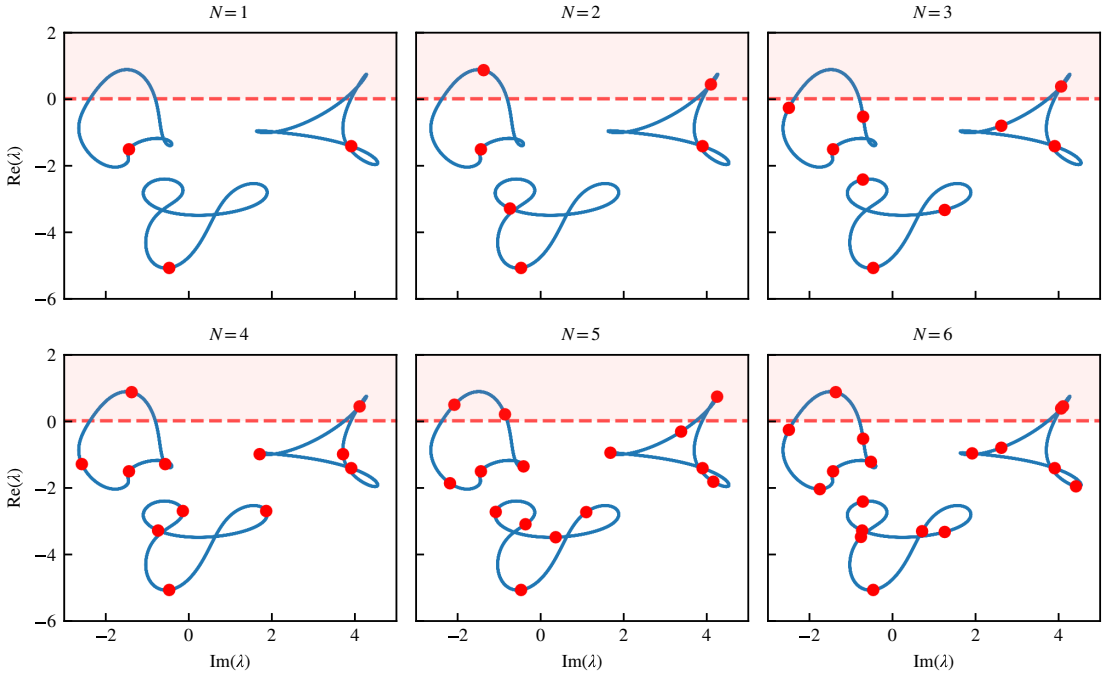


Figure 3.35: Parameterisation of the unit circle $z = e^{i\theta}$ to the curve eig $\left(\hat{A}(z)\right)$. With our choice of $\mathbf{T}_0$, $\mathbf{T}_1$, $\mathbf{T}_{N-1}$ the eigenvalues traverse separate curves, shown in blue. Each panel shows the mapping of the roots of unity for a given N, and shows how the system with one 'element' ($N = 1$) is stable, but with the addition of each subsequent element ($N > 1$), the most unstable mode alternates between the leftmost and rightmost curves.

Observe also that there are no additional unstable two-periodic modes (ρ_5) in Figure 3.34. Since the spectrum for the two periodic problem is the union of the ρ_0 and ρ_5 spectra, it follows that the dominant structures observed for this case would give the same response as a single-passage. This may lead to false conclusions regarding the dynamics of multi-bladed systems when simulations that are performed on a two-passage mesh do not show de-synchronisation of the passages.

3.4.4.2 Most Unstable Modes

Let us first consider the most unstable mode (ρ_0, M_0). From Figure 3.34, this mode becomes more unstable in the five- (ρ_2) and ten-periodic (ρ_1) cases compared to the one-periodic passage, with the growth rate increasing from $\mathrm{Re}\{\lambda\} = 0.347$, for

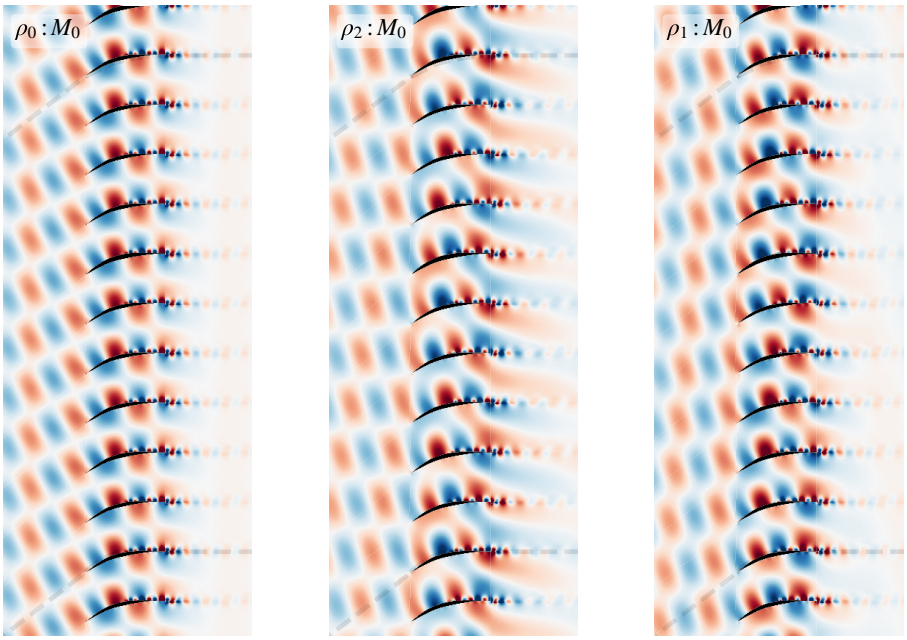


Figure 3.36: The most unstable pressure modes in the (a) single, (b) 5-periodic and (c) 10-periodic cases, showing an adjustment of the upstream wavefronts to a more oblique configuration alongside dominant suction-side dynamics on the surface of the airfoil.

(ρ_0, M_0), to $\mathrm{Re}\{\lambda\} = 0.380$, for (ρ_2, M_0), and finally to $\mathrm{Re}\{\lambda\} = 0.408$, for (ρ_1, M_0)[22].

The visualisation of these modes in Figure 3.36 shows that, as the periodicity is increased, the wavefronts of the upstream acoustic field deviate away from the vertical plane waves seen for ρ_0, as expected, due to the changing phase relationship defined through the N-periodic decomposition. Observe also that the form of the most unstable mode, (ρ_1, M_0) in Figure 3.36, matches the asymptotic behaviour of Figure 3.32. The only difference is that the asymptotic growth rate of 0.390, determined in Figure 3.33, is lower than the eigenvalue growth rate $\mathrm{Re}\{\lambda\} = 0.408$. This is attributable to the length of the impulse response simulation and the slow decay of the low-frequency dynamics associated with the de-phased wave packet, which even at $T = 20$ is still not saturated by the most unstable mode. Based on these growth rates, it is reasonable to assume that the real part of the mapping of

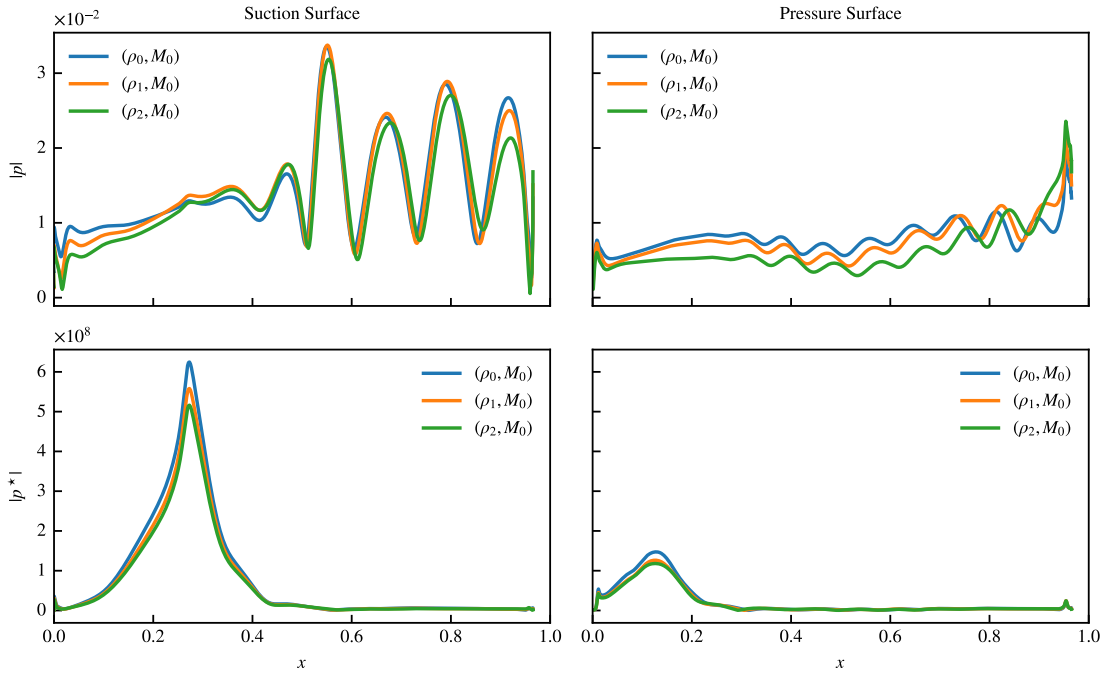
[22]The modes are $(\rho_0, M_0) = 0.409 - 26.143i$, $(\rho_1, M_0) = 0.409 - 26.143i$ and $(\rho_2, M_0) = 0.380 - 25.714i$.

the unit circle to the symbol curve grows for small arguments θ before gradually decreasing as the argument increases. However, with only the one-, two-, five- and ten-periodic solutions resolved with this analysis, it it not possible to obtain the exact mapping of the unit circle to the symbol curve, in a manner similar to that shown in the demonstration example of Figure 3.35. By choosing a highly composite number[23], further resolution can be achieved, thereby allowing for a more detailed interpretation of the effects of having more blades within the blade row, with the downside being the increased cost, albeit one that increases linearly with N.

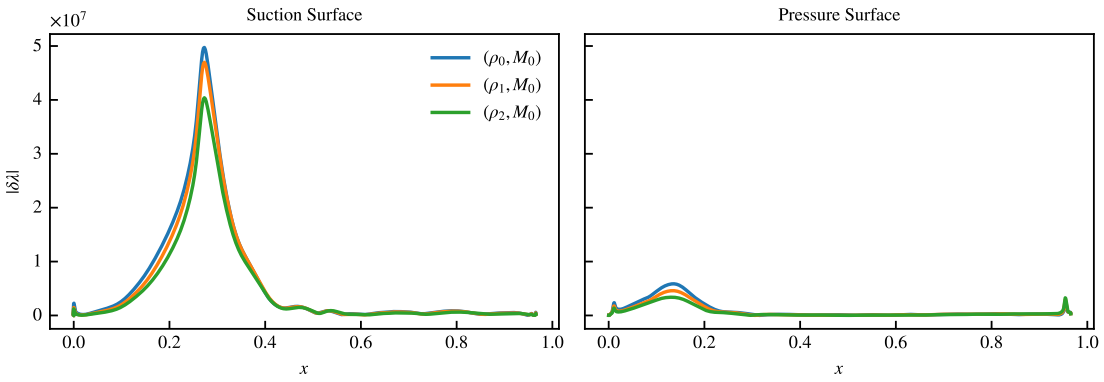
Continuing with the analysis, the pressure distributions for the three scaled modes are shown in Figure 3.37. Peak pressure amplitudes are consistently higher for the ten-periodic case when compared to the five-periodic configuration on the suction surface, with higher pressures seen for the five-periodic case only on the pressure surface for $x > 0.9$. On the suction side pressures for (ρ_0, M_0) are higher than the high-periodicity modes for $x < 0.220$, with the latter dominant for $0.220 \leq x \leq 0.507$, and crucially, (ρ_1, M_0) attains higher amplitudes ($\sim 5\%$ higher) at the point of peak sensitivity, at $x = 0.273$, as evidenced by the adjoint fields. Downstream of $x = 0.507$, the peak pressures of (ρ_0, M_0) increase and are greater than those of (ρ_2, M_0), and comparable to those of (ρ_1, M_0) ($|p| \approx 3.356 \times 10^{-2}$). Note also that peaks in $x > 0.507$, corresponding to ρ_1 and ρ_2, show small displacements in x due to the changes in the standing wave pattern resulting from the small variations in frequency between the modes.

Adjoint fields demonstrate reduced sensitivity for the N-periodic modes on both the suction and pressure surfaces, but the deficit is reduced when the wavemaker is calculated, as shown in Figure 3.37b. The wavemakers show that the (ρ_0, M_0) mode has the greatest structural sensitivity, with sensitivity decreasing while traversing the unit circle (i.e. as $\arg(\rho)$ increases). It is important to note, however, that

[23] A highly composite number is one that has a large number of divisors. This means that many periodicities are captured by a periodic analysis. For example, $N = 10$ has divisors $1, 2, 5, 10$, while $N = 12$ has $1, 2, 3, 4, 6, 12$.



(a)



(b)

Figure 3.37: The surface direct and adjoint pressure distributions for the modes (ρ_0, M_0), (ρ_1, M_0) and (ρ_2, M_0) in (a), and the corresponding wavemakers in (b).

a lower value of the wavemaker field is not necessarily indicative of less dominant structure, only one that is less sensitive. Highly amplified robust structures may therefore result in lower wavemaker fields, which could be the process that leads to smaller wavemaker peaks for (ρ_1, M_0) and (ρ_2, M_0).

The effect of changing the argument of ρ is to alter the phase relationship between the pressure and suction surfaces. By changing ρ, the phase relationship could become more optimal between the features on the suction and pressure surfaces,

although this requires further study, and greater resolution in θ to verify. Furthermore, we have seen in the impulse response analyses, carried out above, that the system is defined by convective and acoustic timescales, and hence determining the dominant modes may involve considering the phase relationships at the point of receptivity *and* the form of the resulting waves as they interact at the trailing edge. Therefore, an acoustic ray-tracing argument should be adopted to further explore the relationship between periodicity and mode growth rates.

3.4.4.3 New Unstable Structures

The higher periodicity also leads to new, unstable resonance structures that trap acoustics waves within the blade passages, reminiscent of Parker modes. This mode family is seen for $(\rho_2, M_1) = 0.154 - 40.402i$, $(\rho_3, M_0) = 0.235 - 40.769i$ and $(\rho_4, M_1) = 0.0493 - 40.838i$. While a trapped acoustic mode is seen within the blade passages for all these modes, note how the change in periodicity in (ρ_3, M_0) concentrates the mode within the passage by removing propagating waves in the upstream and downstream regions, and that this corresponds to the most unstable mode of this group. Recall from Figure 3.34 that there are no unstable modes for ρ_0 in this frequency range, and hence this example demonstrates the appearance of new families of modes that are captured in a multi-passage setting. The modes discussed here, and the pressure amplitudes highlighting the trapped acoustic waves, are shown in Figure 3.38. Further analysis of these structures is left for future work.

3.4.4.4 Low-Frequency Dynamics

Low-frequency dynamics are also captured here, with a spectrum that is significantly denser in $|\mathrm{Im}\{\lambda\}| < 20$ for all ρ_i where $i \neq 0$, in comparison to ρ_0. Despite the challenges of converging these modes, we have been able to select the sufficiently well converged modes (ρ_1, M_{371}), (ρ_2, M_{220}), (ρ_3, M_{81}), (ρ_4, M_{59}) and (ρ_5, M_{242}), so as to

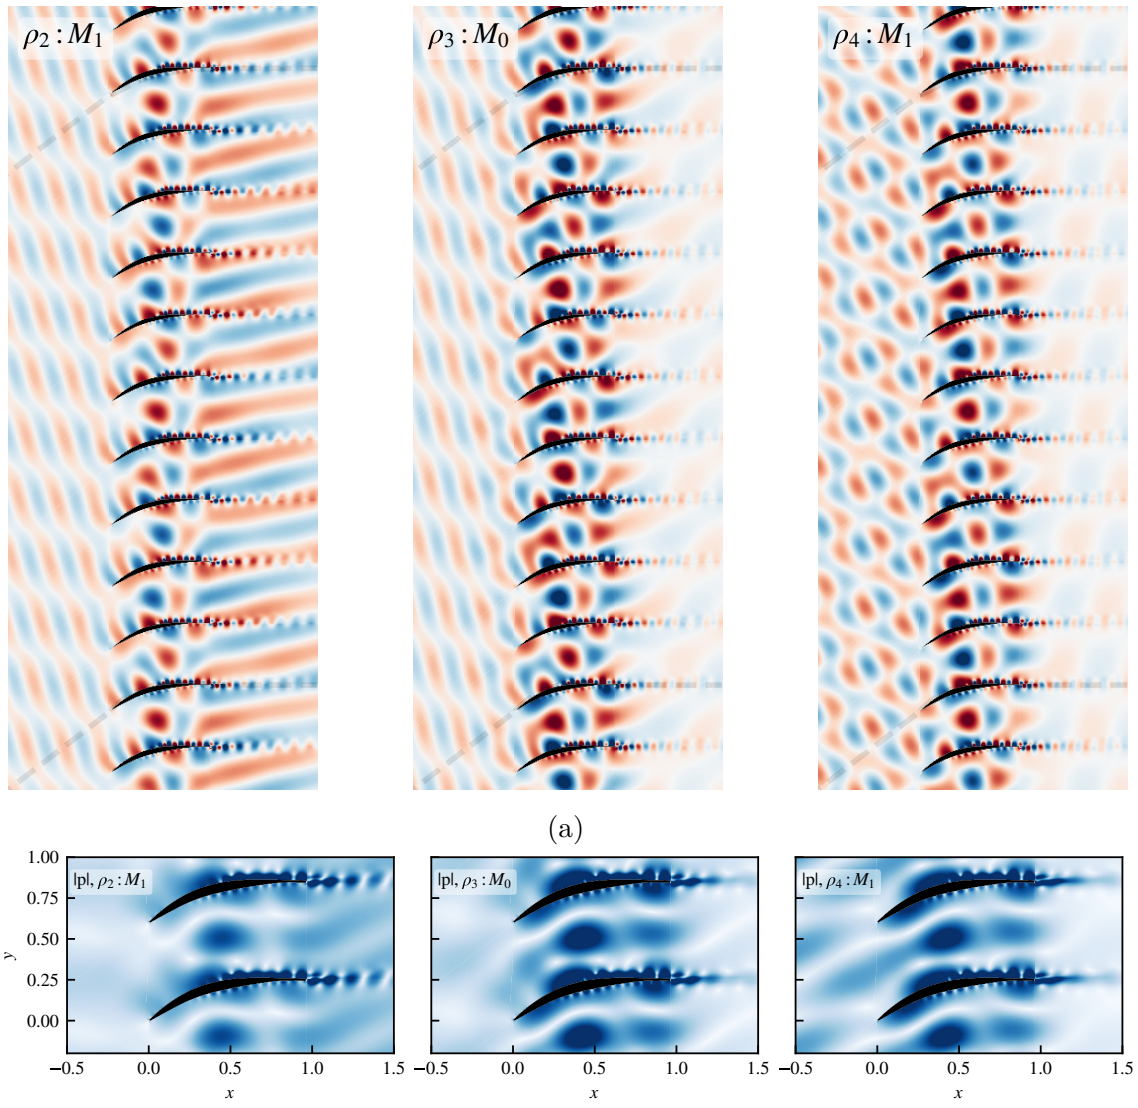


Figure 3.38: Pressure fields for (ρ_2, M_1), (ρ_3, M_0) and (ρ_4, M_1) in (a), with the pressure amplitudes within the passage shown in (b). In (b) note the high pressure amplitudes within the passage in (ρ_3, M_0), and appearance of pressure structures in the downstream and upstream of (ρ_2, M_1) and (ρ_4, M_1), respectively.

visualise the low-frequency dynamics here. These modes are shown in Figure 3.39, and all cases shown here, the pressure dynamics of the modes show large-scale plane-wave structures that span multiple passages, both upstream and downstream of the blade row. Furthermore, note that, unlike the modes seen previously within this paper, the acoustic mode structure within the passages takes the form of plane-waves, with wavelengths greater than 1.

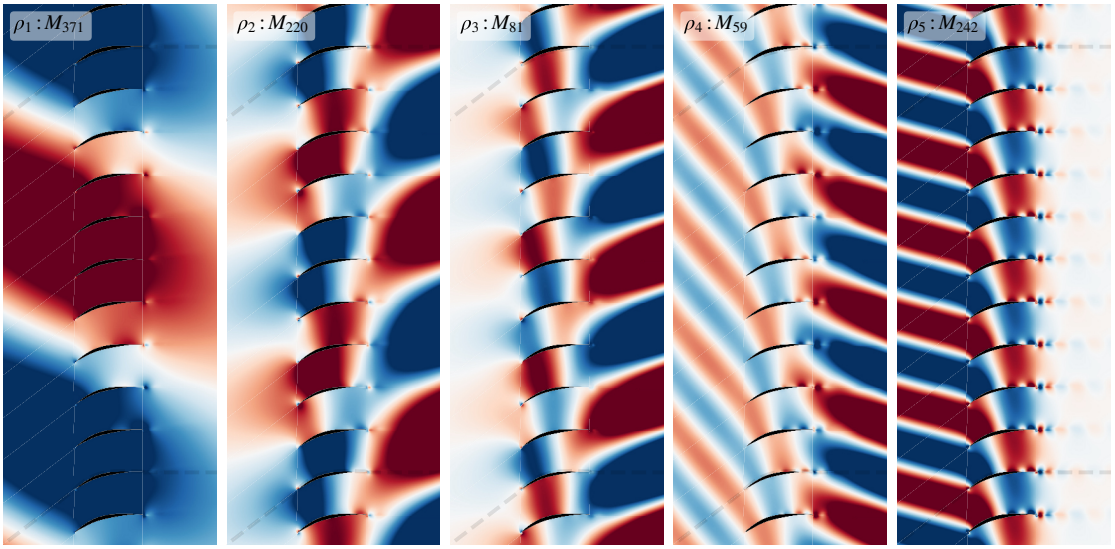


Figure 3.39: Selected low-frequency acoustic modes of the ten-periodic linear cascade.

To determine the driver of these acoustic dynamics, plots of the speed, $|\mathbf{u}|$, in Figure 3.40, show the hydrodynamics localised at the blade surface. These show that the unsteady hydrodynamic activity is confined primarily to the suction-side boundary layer and within the wake downstream of the trailing edge. For the lowest-frequency mode considered here, (ρ_1, M_{371}), these dynamics take the form of acoustic plane waves occupying the upstream portion of the domain, driven by a low-frequency ($\mathrm{Im}\{\lambda\} = 2.76$) 'flapping' motion of the suction-side laminar separation bubble, while the other modes, with higher frequencies, localise the fluctuations closer to the trailing-edge and the near-wake regions.

Note, however, that the selected modes in Figure 3.34, especially (ρ_1, M_{371}), are located in regions of the spectrum that are surrounded by other modes with poorer convergence characteristics, although ones that are suggestive of higher growth rates. As such, while we have been able to observe new low-frequency structures in these modes, further analysis, including a careful receptivity study, are required for a more detailed understanding of the hydrodynamic features driving these modes. This can be achieved by using a combination of shift-invert and harmonic extraction techniques to target the low-frequency branches of the spectrum.

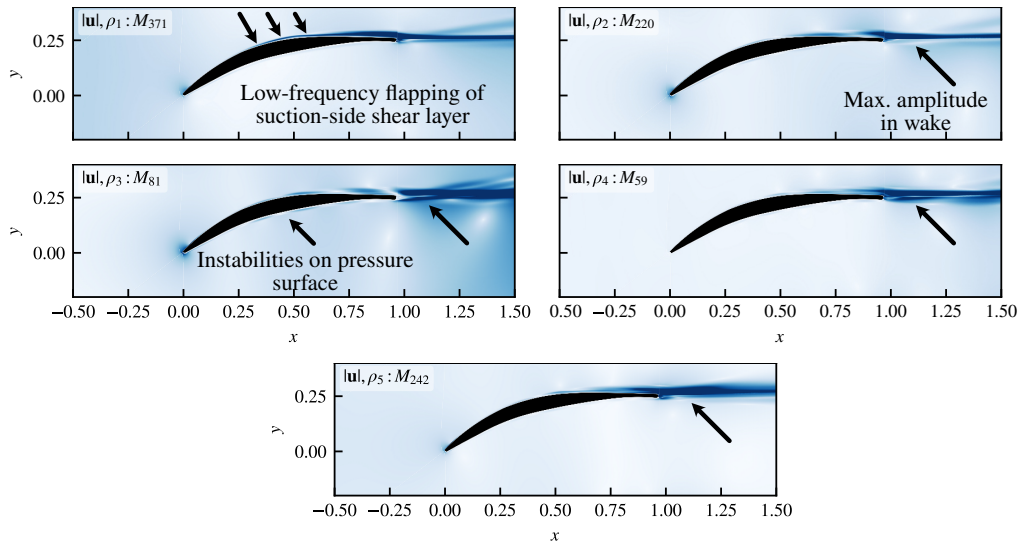


Figure 3.40: Plots of the speed fields $|\mathbf{u}|$ for the low-frequency modes. Arrows indicate high-amplitude disturbances in the near wake, while for (ρ_1, M_{371}), the arrows show low-frequency flapping on the suction surface.

3.5 Conclusion

This paper considered the modal and non-modal aeroacoustic stability properties of a representative linear compressor cascade at $\mathrm{Re} = 100,000$ and $M = 0.3$. This used a controlled-diffusion geometry that was adapted from earlier aeroelasticity experiments, and features an unstable suction-side laminar separation bubble that sheds vortices downstream to the trailing edge.

In the first part of this paper, a mean-flow global stability analysis was performed on a single-passage periodic domain. A transient response analysis highlighted regions of instability on the suction side, as well as a convectively unstable region on the pressure surface. The linear system was seen to be asymptotically unstable, and the mechanism for this instability was shown to manifest itself through a feedback process that links the upstream-propagating acoustic waves, originating at the trailing edge, to the sensitive boundary layers at the separation point of the laminar separation bubble.

A global stability analysis was then carried out, following which it was determined that the spectrum of the linear operator is unstable and is characterised by two

unstable and a sequence of twelve marginally stable modes, with the mode shapes and frequencies found to match very well with the pressure spectrum obtained from the nonlinear solution using the Fast-Fourier Transform (FFT). From this it was also shown that the main peak in the pressure spectrum corresponds to the second-most unstable mode.

Through an analysis of the eigenvectors for these fourteen modes, three main groups of modes were identified. These were the Kelvin-Helmholtz modes, that drive the dominant dynamics of the system; high-frequency modes related to separation bubble dynamics; and two degenerate modes, which were found in association with acoustic resonances within the blade passage.

To augment this description, adjoint modes were calculated. These confirmed that for the most unstable modes, the sensitivities were located upstream of the separation bubble, and also on the pressure surface for the most unstable modes. Applying the wavemaker concept, the structural sensitivities showed that the dominant dynamics of this system are determined by feedback processes involving the separation bubble and, more interestingly, the convectively unstable pressure side region near the leading edge. This region is not linked convectively to the suction side separation bubble and so the feedback mechanism occurs through acoustic means.

In many ways the single-passage results here can be seen to mirror earlier studies with isolated airfoils. For example, much like the numerical studies [Desquesnes et al., 2007; Jones and Sandberg, 2011] and the global stability analysis conducted by Fosas de Pando et al. [2014, 2017] for the NACA-0012 airfoil, the linear cascade is shown to have similar receptivity and sensitivity characteristics on the pressure and suction surfaces of the blade, with aeroacoustic feedback mechanisms shown to exploit these sensitivities. The generated acoustics field also shares considerable similarities and characteristics to the isolated CD airfoil studies at similar flow conditions [Sanjose et al., 2019; Wu et al., 2021], especially with respect to the characteristics of the acoustic source and the upstream directivity of the propagating acoustic waves. Of

course the exact properties of the flow change due to the presence of the cascade, with separation bubble located at approximately the mid-chord, instead of near the leading or trailing edges, the global instability to impulse perturbations remains. Unlike earlier work with isolated airfoils, however, this case also captures new features, such as acoustic resonances in the blade passage and pressure side instabilities. By making use of the direct-adjoint methodology provided to us by this framework, we can also be more definitive about the nature and origin of the processes that sustain the aeroacoustic feedback loop, compared to purely numerical studies that have to infer the feedback mechanisms from nonlinear time-series data. Furthermore, such global stability analyses have rarely been applied to cascade configurations in the past, with one of the few works by Sandberg and Wheeler [2019] applying this analysis to high-pressure turbine flows, and so this work provides new insights into the cascade problem.

Using a cyclic decomposition of the system matrix, this paper demonstrates how a tuned N-periodic blade row problem may be decomposed into smaller sub-problems, that enable computationally tractable stability calculations for cascades of arbitrary sizes, by making use of the properties of block-circulant matrices. This extended earlier work by Schmid et al. [2017], that described this method in the context of an incompressible flow and much-simplified cascade geometries, to a full aeroacoustic simulation.

Through the use of this method, the single-passage analysis was replicated for a ten-periodic linear blade row. The impulse-response analysis showed that disturbances propagate through this cascade by exploiting suction-side sensitivities that are exposed to the dipole source at the trailing edge of the blade above. Scattering of an initial pressure pulse by the leading edges of the cascade blades also resulted in the formation of a persistent transient region occupying five passages. This region is characterised by weak acoustic radiation in the upstream and a perturbed wave vector. The persistence of this feature was explained by the same suction-side

propagation arguments. These observations are summarised in Figure 3.41, panel 3.

This analysis of the N-periodic system is, to the best of our knowledge, the first time that a detailed global stability analysis has been conducted for a cascade configuration with $N > 1$. By extending this calculation to ten passages, it was shown that the local feedback mechanisms that exist in the single passage case can result in the formation of large scale synchronisation structures that span the entire blade row. These are shown to arise due to the interaction of acoustic sources with neighbouring blades, which is not something that can be resolved by single-passage impulse-response analyses, such as those presented by Sandberg and Wheeler [2019]. Due to the absence of shocks in this model compressor blade row, interactions here are also more complex than in the high-pressure turbine simulations by Sandberg and Wheeler [2019], as here acoustic perturbations can travel further upstream and therefore interact with blades both above and below, as shown in Figure 3.41.

The direct and adjoint spectra for the ten-periodic problem were also obtained successfully. From this, it was shown that the ten-periodic domain results in more unstable dynamics that supersede the most unstable mode for the single-passage. New low-frequency modes were also obtained and analysed using adjoint quantities and the wavemaker concept. Of more practical interest, however, is the observation that the two-periodic system does not admit any additional unstable modes. This demonstrates that computational flow solutions seeking to investigate flow instabilities in one or two passages may not capture the dynamics of the full N-periodic system, even if the one- and two-passage solutions have the same dominant dynamics.

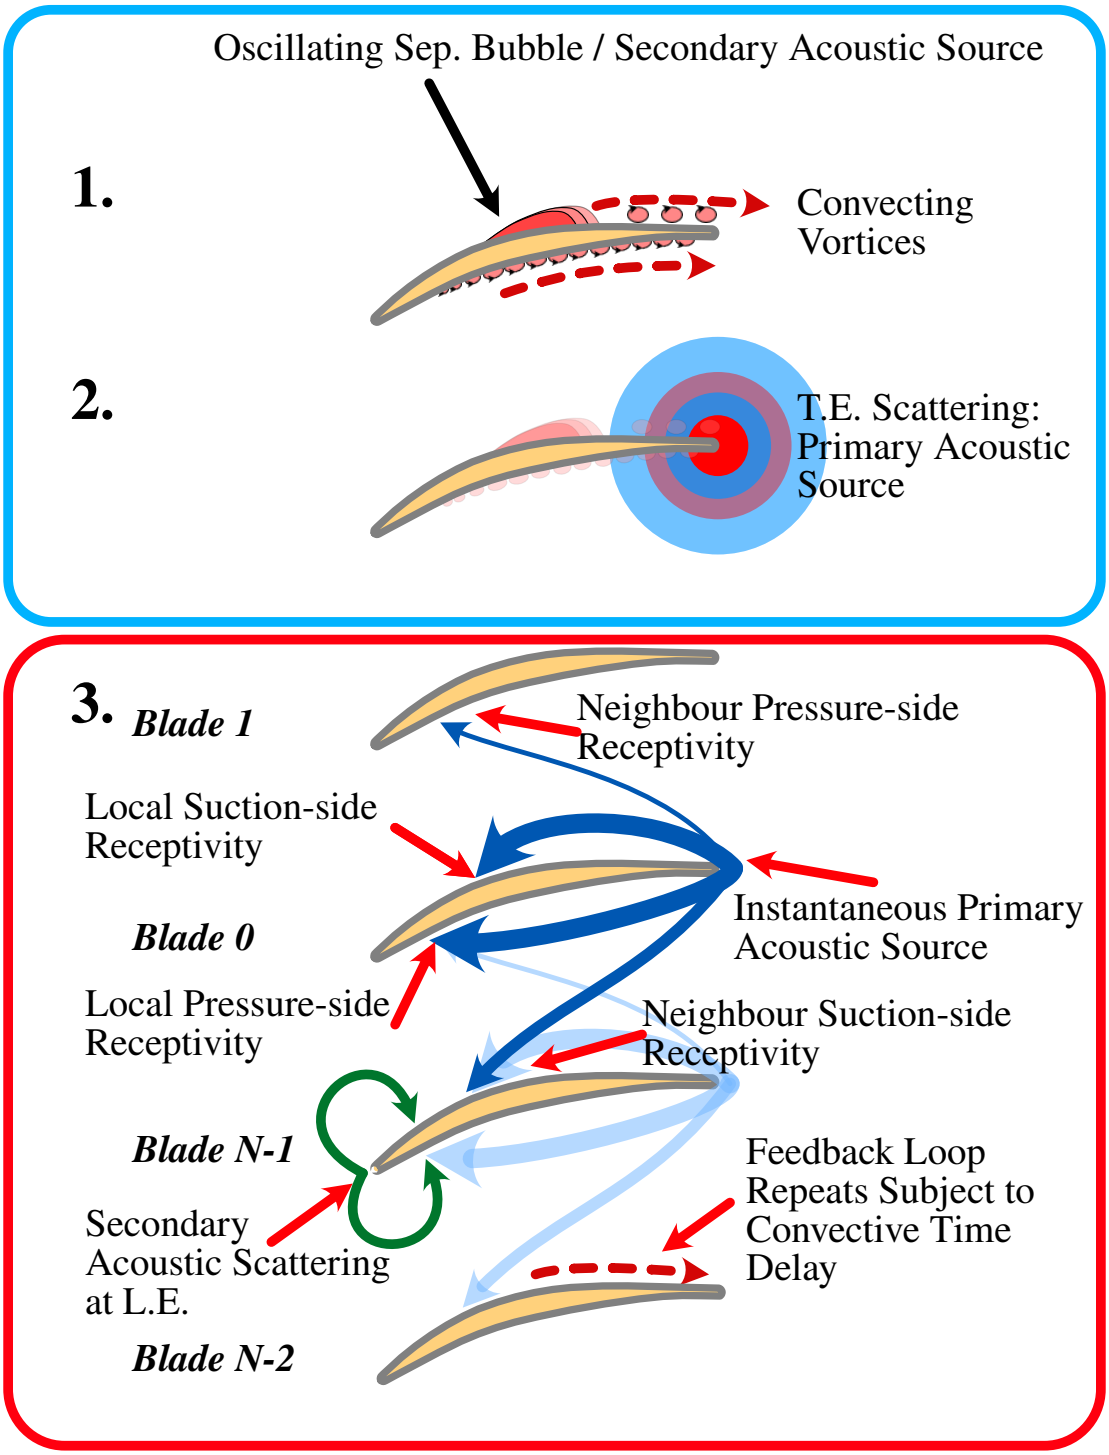


Figure 3.41: The local and global feedback mechanisms identified for N-periodic geometries in this paper.

Chapter 4

Sliding Plane Formalism for High-Order Aeroacoustic and Sensitivity Calculations in Multi-Element Domains

Abstract

In many engineering fields the analysis and optimisation of flow interactions in systems defined by both static and rotating periodic arrangements of immersed bodies is of academic and practical interest. Cases such as these may exhibit complex aeroacoustic feedback mechanisms through self-excited instabilities and through the receptivity of the flow to the resulting acoustic field. At best this may manifest itself as minor performance losses and at worst result in the complete structural failure of the affected machinery. Though the application of nonlinear computational aeroacoustics to such cases is becoming more prevalent, a more complete understanding may be gained from the direct and adjoint dynamics. This remains uncommon, however, due to the complexity of the implementation, computational costs and the challenges of resolving the instability at the interface. This paper introduces a new methodology for the nonlinear, direct and adjoint analysis of unsteady flows and aeroacoustics for multi-element systems in relative motion, using a lightweight, high-order sliding-plane interface, and within a unified numerical framework.[1]

4.1 Introduction

Over the past 50 years the traditional approach for aeroacoustic problems has been to decouple the noise generation from its propagation into the far field. This is typically

[1]Paper in preparation as of May 2021.

achieved using either an analytic representation for the noise source [Lighthill, 1954; Ffowcs Williams and Hawkings, 1969a,b] or through a hybrid approach, which uses a small-scale computational fluid dynamics (CFD) calculation to capture the main instability mechanism, after which the acoustic waves are propagated to the far field using an acoustic propagation model, for example [Wang et al., 2006; Manoha et al., 2002].

Despite significant strides made using this approach, a rich subset of problems remains that requires both capturing the generation of the acoustics from the unsteady flow field, and simultaneously coupling this with the propagation. This is especially relevant in complex domains where the flow field is dominated by instabilities and multi-scale feedback mechanisms, where the stability properties and interactions within the flow may be critical in capturing a particular flow state.

The rarity of unified codes that resolve both generation and propagation of acoustic waves is due to the need to capture microscopic perturbations in the acoustics field and macroscopic variations in the flow field. Many commercial codes that make use of low-order schemes are ill-suited to this task [Lele, 1997; Colonius and Lele, 2004; Tam, 2004], since, with the high numerical dissipation, the acoustics field is quickly dissipated. For this reason, bespoke high order Computational Aeroacoustics (CAA) codes must be used. Such implementations are specifically designed to propagate disturbances at a wide range of frequencies and amplitudes with minimal dissipation.

One group of problems that is of current interest for researchers is the interaction of rotating periodic systems, examples of which include radial impellers in motor vehicle turbochargers [Broatch et al., 2015], rotor-body interaction in helicopters [Steijl and Barakos, 2008] and rotor-tower interactions in wind turbines [Zahle et al., 2009], or perhaps the most prominent example, rotor-stator interactions at various points in turbomachinery [Peake and Parry, 2012]. All of these require the numerical code to handle the relative motion between the solid surfaces, and a vast array of

methodologies have been deployed to model these in computational codes. Perhaps the most common and simplest method is the sliding plane or sliding mesh algorithm. This approach models the governing equations in rotating reference frames, within which the mesh is stationary, and provides a methodology for linking the flow solutions between multiple domains in relative motion. In aeroacoustic applications, however, a lot of care must be taken to ensure the correct propagation through the interface without adverse effects to the flow or acoustic fields arising due to dissipation, dispersion or distortion. Significant progress for nonlinear simulations have been made in this regard, but a more fundamental understanding may be gained by analysing the linearised dynamics of such systems, with gradient-based optimisation techniques also becoming accessible [Schmid and Brandt, 2014]. These simulations remain rare, however, due in part to the unique computational and implementation challenges.

Adjoint-based techniques are one such group of methods that offer a powerful pathway for analysing sensitivities and computing gradients of the flow with respect to large parameter spaces, for both analysis and optimisation. Two main approaches exist, each with their own advantages and disadvantages [Gunzburger, 2002].

The first is the *continuous* adjoint pioneered in the field of aeronautics by Jameson [1988], which derives the continuous adjoint PDEs analytically from the governing equations before discretising them. This is typically the slowest to implement due to the manual effort and care required to derive the adjoint equations and the compatible numerical scheme, but it makes up for this by being fast to execute, as the Jacobian calculation is a simple function evaluation. An example of this, and one of the first applications of this approach to multi-blade-row turbomachinery, was the work in [Wang and He, 2010]. In this publication, a continuous adjoint was derived for a steady 3D Reynolds Averaged Navier Stokes (RANS) solver and made use of the mixing-plane treatment [Denton, 1992]. This method was then successfully applied to the problem of constrained aerodynamic shape optimisation

of a multi-stage geometry [Wang et al., 2010].

Alternatively, the *discrete* adjoint discretises the forward problem along with the time integration scheme and computes the adjoint of this operator by taking the complex conjugate transpose of the resulting matrix [Giles and Pierce, 2000]. The adjoint here is fully consistent with the forward problem and is fast to implement as it only requires the forward linear operator to be available. While the cost of the matrix-vector product is similar to a function evaluation in the continuous adjoint, the main drawback consists of obtaining the discrete linear operator in the first place. This can be a costly operation, especially if it is repeated at every time step as in the nonlinear-adjoint looping approach [Gunzburger, 2002]. Some recent work making use of the unsteady discrete adjoint method with sliding interfaces includes [Nielsen and Diskin, 2013; Ntanakas and Meyer, 2015; Ntanakas et al., 2018]. These focused on steady and unsteady RANS solutions, and while these approaches have been successful for general aerodynamic shape optimisation, the numerical methods and the treatment of the sliding interface are unsuitable for high-order aeroacoustics simulations. This is due, in part, to the dissipative treatment at the sliding plane which does not propagate the nonlinear nor the adjoint cleanly through the interface between domains.

The nonlinear-adjoint equations are non-autonomous and solved backwards in time, so the nonlinear forward trajectory must be stored for the linearisation process. A brute-force approach should be avoided due to the considerable storage requirements and the time-cost of disk storage access, and so a checkpointing procedure should be used to optimise performance (e.g. see Gunzburger [2002] for the theory of time-dependent adjoint problems, and Griewank and Walther [2000] for details of the `revolve` checkpointing library). Using both the system memory and hard drives while saving only a small subset of the required trajectory checkpoints, this procedure mitigates the issue of storage at the expense of having to recompute the missing data.

Due to the complexity of the governing equations and boundary closures used in our code, in this paper we will proceed with the discrete adjoint for ease of implementation. Despite this, the methodology described here is highly modular and a continuous approach is easily integrated into this framework.

The central aim of this study is to present a new time-domain, high-order and lightweight sliding plane implementation for computational aeroacoustics applications and optimisation, with numerically consistent forward and adjoint capabilities. Our specific focus is on multi-blade-row geometries in turbomachinery applications. In such applications flow solvers are typically low-order, and, as we have seen, very few adjoint treatments exist for dealing with multiple blade rows, especially at the high-order accuracy required for CAA. In this paper we describe the numerical background of our solver, the sliding plane algorithm and our treatment of unsteady flows in Section 4.2, and perform extensive validations of the nonlinear, direct and adjoint computations using the sliding plane in Section 4.3. We conclude the paper in Section 4.4 by illustrating two representative problems that may be studied using this methodology[2]. Although we will focus this paper primarily towards turbomachinery applications, throughout we emphasise the modularity and adaptability of this approach to different codes and flow solvers.

4.2 Numerical Method

4.2.1 Direct Numerical Simulations

The DNS solver used in this paper is based on the pseudo-wave method of Sesterhenn [Sesterhenn, 2000; Fosas De Pando, 2012], and implements the non-dimensional

[2]In the thesis, the rotor-stator interaction section is in Chapter 5.

compressible gas flow equations for a calorically perfect gas,

$$\frac{1}{p}\left(\frac{\partial p}{\partial t}+\mathbf{u}\cdot\nabla p\right)=-\gamma\nabla\cdot\mathbf{u}+(\gamma-1)\left(\frac{\partial s}{\partial t}+\mathbf{u}\cdot\nabla s\right), \tag{4.1}$$

$$\rho\left(\frac{\partial \mathbf{u}}{\partial t}+\mathbf{u}\cdot\nabla\mathbf{u}\right)=-\frac{1}{\gamma M^2}\nabla p+\frac{1}{\mathrm{Re}}\nabla\cdot\boldsymbol{\tau}+\mathbf{f}_v, \tag{4.2}$$

$$p\left(\frac{\partial s}{\partial t}+\mathbf{u}\cdot\nabla s\right)=\frac{\gamma M^2}{\mathrm{Re}}\boldsymbol{\tau}\colon\nabla\mathbf{u}-\frac{\gamma}{(\gamma-1)\mathrm{RePr}}\nabla\cdot\mathbf{q}+Q_r, \tag{4.3}$$

$$\boldsymbol{\tau}=\mu\left(\nabla\mathbf{u}+\nabla\mathbf{u}^{\intercal}\right)+\left(\mu_v-\frac{2}{3}\mu\right)(\nabla\cdot\mathbf{u})\,I, \tag{4.4}$$

$$\mathbf{q}=-k\nabla T, \tag{4.5}$$

with the relations

$$p=\rho T,\quad (\gamma-1)s=\log\frac{p}{\rho^\gamma},\quad a^2=\frac{p}{\rho}, \tag{4.6}$$

where the non-dimensional variables, which we define in Equations (4.7)–(4.8), are defined by normalising the dimensional variables ($\star'$) by a characteristic scale, denoted by a subscript $\star_0$ at a fixed point in the flow. These scales are usually motivated by the initial or boundary conditions, typical length scales or physical properties of the fluid. For our purposes, we will define the $\star_0$ quantities at the outlet of the computational domain, unless otherwise stated.

$$\mathbf{x}=\frac{\mathbf{x}'}{x_0},\quad t=\frac{u_0t'}{x_0},\quad \mathbf{u}=\frac{\mathbf{u}'}{u_0},\quad p=\frac{p'}{p_0}, \tag{4.7}$$

$$s=\frac{s'}{s_0},\quad \rho=\frac{\rho'}{\rho_0},\quad \mu=\frac{\mu'}{\mu_0},\quad k=\frac{k'}{k_0}. \tag{4.8}$$

Equations (4.1)–(4.3) are then discretised in space and decomposed into pseudo-waves on a curvilinear domain[3]. Further details of this as well as complete functional forms of these terms may be found in [Fosas De Pando, 2012].

Defining the state vector $\mathbf{v}=[p\quad s\quad \mathbf{u}]$, we represent these partially discretised nonlinear governing equations as

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t}=\mathbf{F}\left(\mathbf{v},\frac{\partial\mathbf{v}}{\partial\boldsymbol{\xi}},\frac{\partial^2\mathbf{v}}{\partial\boldsymbol{\xi}^2}\right), \tag{4.9}$$

which has an implied dependence on the grid metric that accounts for transformations from the Cartesian to the curvilinear computational space.

[3]The Cartesian variables $\mathbf{x}=(x,y)$ transform to the curvilinear (or computational domain) variables $\boldsymbol{\xi}=(\xi,\eta)$ under the action of a bijective map M.

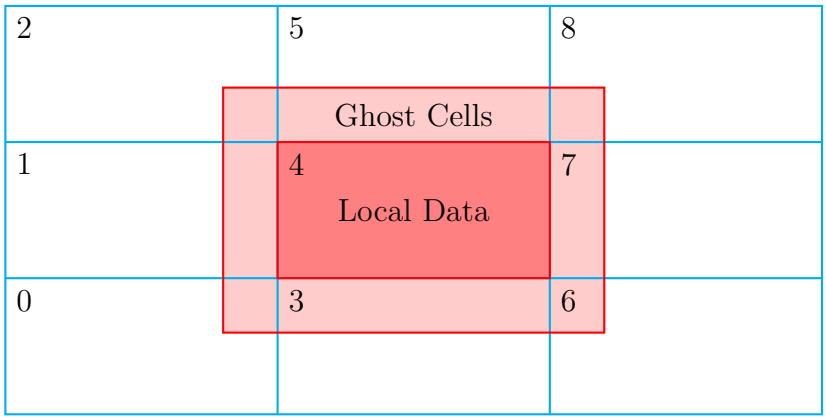


Figure 4.1: For any given processor the local data is embedded within a layer of ghost points obtained from neighbouring processors. The compact differentiation stencil is then applied sequentially: first in the vertical and then in the horizontal direction.

4.2.1.1 Differentiation in Space

The fifth-order compact, low-dispersion upwinding (CULD5) [Adams and Shariff, 1996] is used for the differentiation of the inviscid terms within the computational domain. The upwinding direction is chosen based on the direction of the pseudo-waves so as to retain high-order differentiation accuracy while simultaneously introducing damping at the high-wavenumbers in order to control numerical noise occurring as a result of the discretisation. Dissipation terms are differentiated using a third-order compact scheme [Lele, 1992], and stencil information for both schemes is provided in Tables D.1–D.4.

In general, compact schemes have a global support over the entire computational domain, and hence, a parallel code that partitions the global domain into regions of local data would require a significant communications overhead in order to calculate the derivatives. This communication process quickly saturates with increasing problem size, and the speed gains of a parallel approach rapidly disappear. To mitigate this the two compact schemes described above are cast into a *locally compact variant*, whereby the stencil makes use of the local data surrounded by ghost points containing data from the neighbouring processor and is shown in Figure 4.1. It was found that between 16 and 32 ghost points are required around each processor in order

for the differentiation error from the truncated stencil to become negligible, when compared to the global scheme. This approach has the additional benefit that each processor only needs to fetch the missing ghost data from its immediate neighbours, thus considerably reducing the volume of data exchanged.

4.2.1.2 Time Integration

Both the nonlinear and linearised equations [4] use the method of lines to propagate their respective state vector in time for the Initial Value Problem (IVP) in Equations (4.10)–(4.11),

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \mathbf{F}(t, \mathbf{v}), \tag{4.10}$$

$$\mathbf{v}\Big|_{t=0} = \mathbf{v}_0. \tag{4.11}$$

The nonlinear solver and the checkpointing algorithm both use a low-storage variant of the Runge-Kutta 4 algorithm `RK4(3)5[2R+]C` [Kennedy et al., 2000], with the Butcher tableau given in Table D.5 and coefficients in Table D.6. Simulations of the linearised Navier-Stokes make use of exponential timestepping with Krylov subspaces [Schulze et al., 2009; Moler and Van Loan, 2003], which offers superior speed and accuracy when compared to the Runge-Kutta methods for computing the linearised dynamics.

4.2.2 Lagrangian Framework

Of particular interest for us are sensitivity and optimisation problems, many of which arise naturally in many theoretical and engineering applications. Typically one seeks an optimal solution that either maximises or minimises a cost function, often subject to some additional side constraints that define the problem, such as insisting that the solution obeys the governing equations, boundary and terminal conditions. Gradients are then obtained for the cost function, and, in the case of

[4] Which may take the form of a direct, adjoint or checkpointed solution.

an optimisation problem, used in a gradient descent algorithm of choice. A popular methodology for solving these constrained optimisation problems is the Lagrangian framework. This casts the constrained optimisation problem to an unconstrained setting that is, at least conceptually, easier to solve.

For now let us consider a simple case where we seek the gradients of a cost function, $\mathcal{J}$, with respect to parameters $\mathbf{p}$ parameterising some initial condition $\boldsymbol{\eta} = \boldsymbol{\eta}(\mathbf{p})$, and naturally subject to the governing equations in the IVP in Equations (4.10)–(4.11) over some period $t \in [t_0, t_1]$. In time-discretised form, we represent the nonlinear evolution in Equation (4.10) by

$$\mathbf{v}_{i+1} = \mathcal{N}(\mathbf{v}_i), \quad \text{with} \quad \mathbf{v}_0 = \boldsymbol{\eta}, \tag{4.12}$$

where $\mathcal{N}$ denotes the nonlinear propagation operator that takes the solution at time index i to the next step at time index $i+1$. The *Lagrangian*, $\mathcal{L}$ can then be written as [Zhang et al., 2019]

$$\mathcal{L}(\boldsymbol{\eta}) = \mathcal{J}(\mathbf{v}_N) - \boldsymbol{\lambda}_0^\intercal(\mathbf{v}_0 - \boldsymbol{\eta}) - \sum_{n=0}^{N-1} \boldsymbol{\lambda}_{\boldsymbol{n+1}}^\intercal(\mathbf{v}_{n+1} - \mathcal{N}(\mathbf{v}_n)), \tag{4.13}$$

with the Lagrange multipliers $\boldsymbol{\lambda}_n$, $n = 0, \ldots, N$. Here the first term enforces the cost function; the second, the initial condition; and the last constraining the timestepping procedure for all time. Using standard techniques [Gunzburger, 2002], a local minimum is sought for $\mathcal{L}$ as a function of the initial condition $\boldsymbol{\eta}$. Considering first order variations

$$\frac{\mathrm{d}\mathcal{L}}{\mathrm{d}\boldsymbol{\eta}} = \boldsymbol{\lambda}_0^\intercal - \left(\frac{\mathrm{d}\mathcal{J}}{\mathrm{d}\mathbf{v}}(\mathbf{v_n}) - \boldsymbol{\lambda}_N^\intercal\right)\frac{\partial \mathbf{v}_N}{\partial \boldsymbol{\eta}} - \sum_{n=0}^{N-1}\left(\boldsymbol{\lambda}_n^\intercal - \boldsymbol{\lambda}_{n+1}^\intercal \frac{\mathrm{d}\mathcal{N}}{\mathrm{d}\mathbf{v}}(\mathbf{v}_n)\right)\frac{\partial \mathbf{v}_n}{\partial \boldsymbol{\eta}}, \tag{4.14}$$

Then, since we have a free choice of $\boldsymbol{\lambda}_i$, we can set the terms multiplying $\frac{\partial \mathbf{v}_N}{\partial \boldsymbol{\eta}}$, which we would like to avoid computing since it is computationally intensive, to zero so that

$$\boldsymbol{\lambda}_n = \left(\frac{\mathrm{d}\mathcal{N}}{\mathrm{d}\mathbf{v}}(\mathbf{v}_n)\right)^\intercal \boldsymbol{\lambda}_{n+1}, \quad \text{with the initial condition,} \quad \boldsymbol{\lambda}_N = \left(\frac{\mathrm{d}\mathcal{J}}{\mathrm{d}\mathbf{v}}(\mathbf{v}_N)\right)^\intercal. \tag{4.15}$$

Note that for the formulation of the Lagrangian in Equation (4.13) above that $\mathcal{L}(\mathbf{v}, \boldsymbol{\eta}, \boldsymbol{\lambda}) = \mathcal{J}(\mathbf{v})$, and so the gradient of the cost function with respect to the

parameters is simply $\boldsymbol{\nabla}_{\mathbf{p}}\mathcal{J} = \partial\boldsymbol{\eta}/\partial\mathbf{p}\,\boldsymbol{\lambda}_0^\intercal$. Equation (4.15) is solved backwards in time from time index $n = N$ to $n = 0$.

Note that, critically, the linearisation of the nonlinear time advancement operator $\mathrm{d}\mathcal{N}/\mathrm{d}\mathbf{v}$ must be evaluated around $\mathbf{v}_n$. In the case of an autonomous system, such as one where the linearisation is performed around a steady mean flow, this is inconsequential, since the linear operator $\mathrm{d}\mathcal{N}/\mathrm{d}\mathbf{v}$ is autonomous. In cases where such a physically meaningful mean flow does not exist, however, there is an additional dependence on time through $\mathbf{v}_n$. Here we are required to linearise around the nonlinear snapshot at each time step, and so access to the entire *trajectory* of the nonlinear state vector $\{\mathbf{v}_n\}_0^N$ becomes necessary, with grave implications for storage if done directly on any sizable problem.

This issue of storage may be addressed with an appropriately chosen checkpointing scheme, whereby an optimal subset[5] of nonlinear flow fields is stored on either Random-Access Memory (RAM) or the hard drive (HDD) from which intermediate solutions may be recovered. The `revolve` library [Griewank and Walther, 2000], and its corresponding PETSc implementation [Zhang et al., 2019] are used here to interface with our nonlinear, linearisation, direct, and adjoint codes. Multilevel checkpointing, which uses both RAM and the HDD for the storage of snapshots, is available here and is used to address the issue of limited storage capacity.

Figure 4.2 shows the checkpointing algorithm utilised in our code. First the nonlinear solution is stepped forwards in time, with the PETSc `TSTrajectory` object being used to save the trajectory, and provides the option of running a concurrent forward sensitivity propagation problem [Zhang et al., 2019]. Then, once the nonlinear trajectory is saved, the `TSAdjoint` object handles both the adjoint timestepping and the recovery of the checkpointed base flow, with the timestepper computing the Jacobian, $\mathbf{A}$, using Equation (4.25).

[5]See Figure 4.2 for an explanation of optimality.

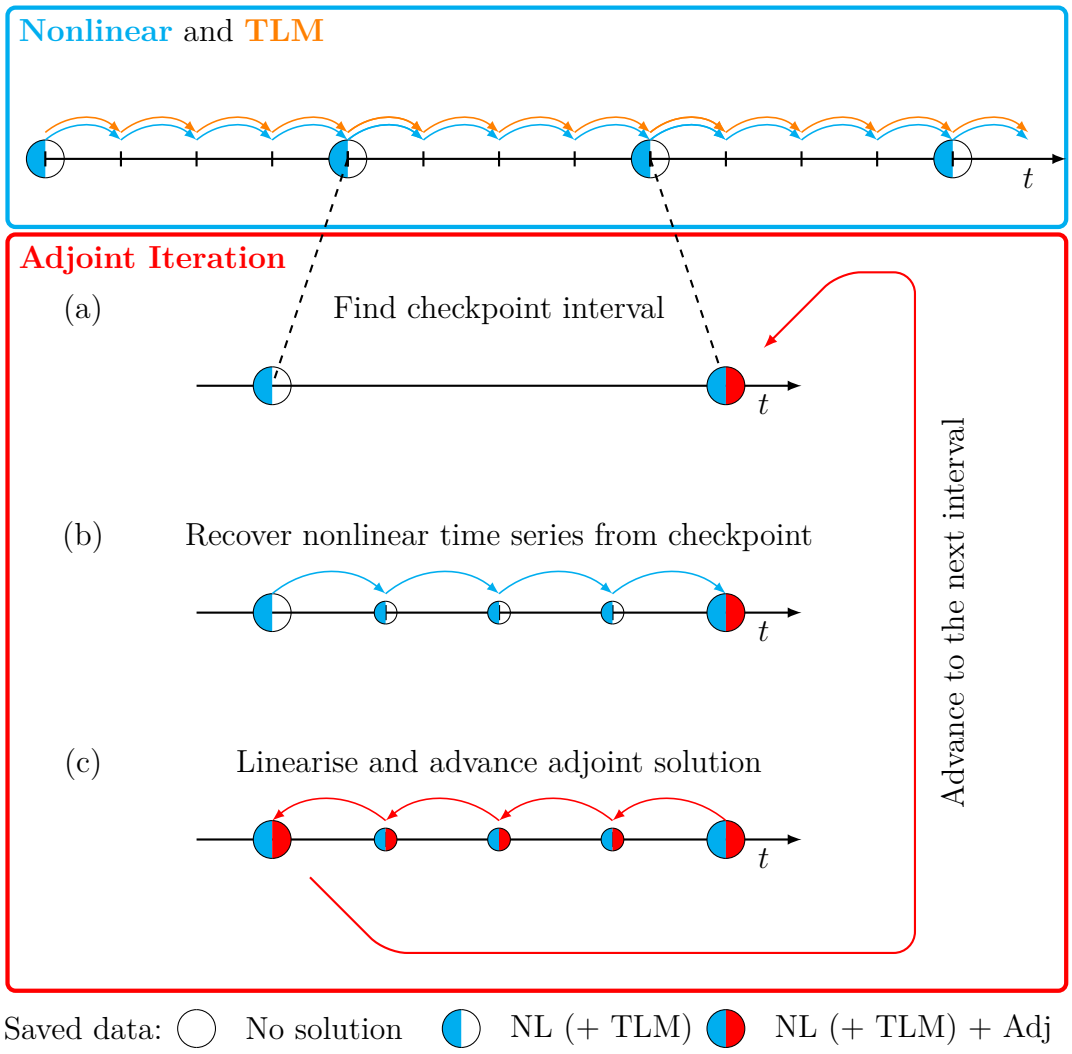


Figure 4.2: The nonlinear-adjoint procedure showing the timestepping for the nonlinear and tangent linear model (TLM) (in blue and orange), and the adjoint iterative procedure that makes use of the checkpointing from the revolve library. The nonlinear solution is run first, with solutions saved at optimal points along the trajectory. The adjoint, which requires the nonlinear field at every timestep, first recovers the missing nonlinear data, then linearises the equations about the checkpoint state and finally advances the adjoint solution backwards in time. Optimality in this context refers to minimising the total number of forward calculations required for the combined forward and adjoint loop. Consider the case where there is a finite number of checkpoints that can be stored, and which are freed once they are no longer required. In this case, both the checkpoints saved on the forward run (large circles) and the intermediate values (small circles), computed during procedure (b), must fit within the allocated storage. As the adjoint simulation progresses checkpoints are freed and more space is available for intermediate checkpoints, and so the intervals between the forward checkpoints can be made larger. It can be shown (see e.g. Griewank and Walther [2000]) that, for explicit time integration, this amounts to selecting checkpoints in a binomial fashion.

4.2.3 The Sliding Plane

In many practical applications, solid bodies immersed in fluid move relative to each other either as a result of fluid-structure interactions or due to external forcing. Modifications must therefore be made to our solver in order to model such cases, which includes handling the communications between the two domains and changes to the linearisation, direct, and adjoint solvers.

The simplest case of constant-speed relative 'sliding' motion between two computational domains is one that we will implement here. In general, the coordinate transformations into the stationary reference frame may be complicated, for instance in the case of two co-axially rotating 3D meshes capturing two bladed disks inside a jet engine, which would require us to modify the governing equations to account for the rotating reference frame. In our 2D case, however, this manifests itself as a simple Galilean coordinate transformation

$$\mathbf{x}^{\dagger} = \mathbf{x} - \mathbf{v}_{\text{sp}} t, \quad \text{where} \quad \mathbf{v}_{\text{sp}} = \text{const. vector}, \tag{4.16}$$

for which the invariance of the Navier-Stokes equations Equations (4.1)–(4.3) is well-known and easily verified. In the above, $\mathbf{v}_{\text{sp}}$ represents the sliding plane velocity, which can be positive or negative.

The geometry we consider consists of two adjacent grids as shown in Figure 4.3. The overlap between the grids occurs for the ghost points only, and not for the internal domain. In this region each vertical row of ghost points coincides with the vertical row in the neighbouring internal domain for all time as the mesh translates. Although this can be generalised to remove this restriction, we found such a setup desirable as it reduces the support of the interpolation operator to just two (instead of four) neighbouring points. In addition, if a constant vertical distribution of points is taken in both overlapping regions, a single interpolation operator is required for the entire interface.

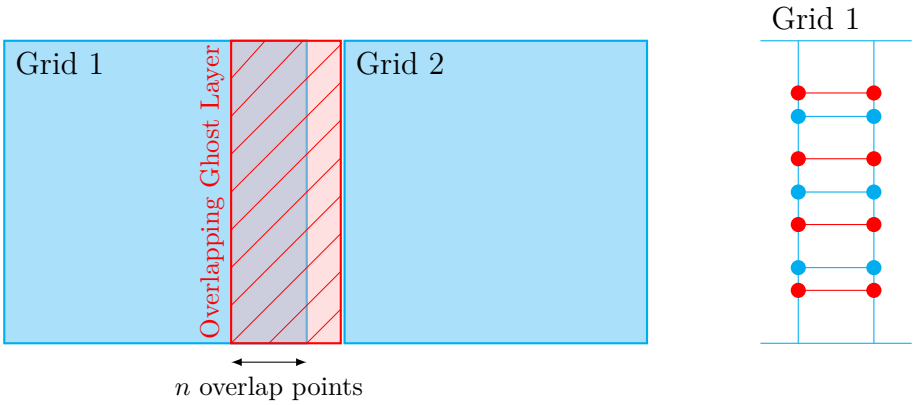


Figure 4.3: The arrangement of the two grids at the sliding interface showing (a) the ghost layer from Grid 2 overlapping into Grid 1 and (b) the arrangement of ghost and physical points in the overlap region. We assume that the distribution of all points in the overlap region is the same along each vertical axis to simplify interpolation.

4.2.3.1 Interpolation Procedure

For the interpolation scheme, we use the 4th-order Hermite interpolation scheme of Delfs [2001], that was used previously by Bhaskaran and Lele [2010] for an LES calculation of a high-pressure turbine. This is chosen because it has been shown by, e.g., Bhaskaran and Lele [2010] that the high-order interpolation is necessary in order to maintain the high-order accuracy of the code. Furthermore, this method is favoured in preference to spectral interpolation methods due to its simplicity; compact support that requires information only at the two neighbouring nodes when performed in one dimension, for sparse linear operators; removes the need to integrate external libraries into our code; and also avoids the use non-uniform Fourier transforms, which are necessary because an interpolation between the computational ($\boldsymbol{\xi}$) and physical ($\mathbf{x}$) coordinates is necessary to determine the bijective linear map, M, that maps between these two spaces.

For each consecutive pair of donor points, we construct a local coordinate system η within the computational space that is centred at the midpoint of these two donor points, as shown in Figure 4.4. Since the map between the computational space and the physical space is bijective, we can define a invertible map M that maps the

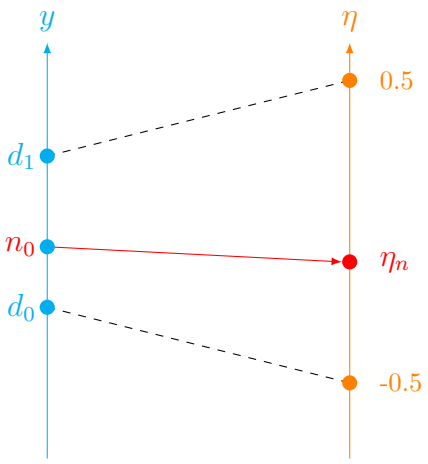


Figure 4.4: For each pair of consecutive donor points defining an interval $\{y\colon y \in [d_0, d_1)\}$ in the physical space, the inevitable map M^{-1} maps this interval to the interval $\{\eta\colon \eta \in [-0.5, 0.5)\}$. The image η_n of a neighbouring point n_0 – at which the value of the interpolate is sought – is computed using Equation (4.17) and the Newton-Raphson procedure by setting $f(\eta_n) = n_0$ and finding the root η_n.

physical coordinate to the computational space y, $(\eta) = M^{-1}(y)$. In this space, the interpolation formula is given by Equation (4.17).

$$f(\eta) = \sum_{k=0}^{1} \left[C_k^0(\eta) f_{j+k} + C_k^\eta(\eta) \left[\frac{\partial f}{\partial \eta}\right]_{j+k} \right] + \mathcal{O}(\Delta^4), \tag{4.17}$$

where

$$C_k^0 = \alpha_k^0 \beta_k^0, \quad C_k^\eta = \alpha_k^\eta \beta_k^\eta, \tag{4.18}$$

and

$$\alpha_k^0 = \frac{1}{2} - (-1)^k \eta, \quad \beta_k^0 = 1 - 2m\eta(-1)^k \left[\frac{1}{2} + (-1)^k \eta\right] \tag{4.19}$$

$$\alpha_k^\eta = m(-1)^k \left(\frac{1}{4} - \eta^2\right), \quad \beta_k^\eta = \frac{1}{2} - (-1)^k \eta. \tag{4.20}$$

Since the grid metric is assumed to be differentiable by construction, we can calculate the coordinate of the interpolation points setting $y_{\text{interp}} = f(\eta_{\text{interp}})$ in the interpolation formula Equation (4.17), and solving the resulting system by means of the Newton-Raphson procedure. Note also that this formula admits a parameter $m \in \{0, 1\}$ that switches between a simple linear interpolation, $m = 0$, and the

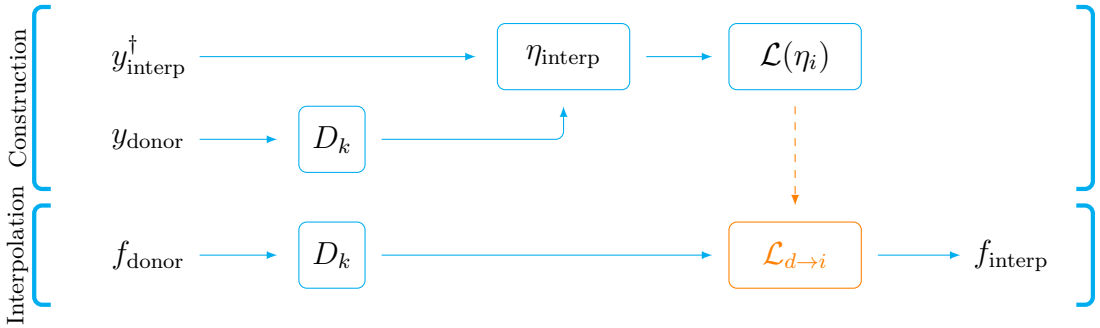


Figure 4.5: At each step of the time integration procedure the coordinate of the interpolation point $y^{\dagger}_{\text{interp}} = y_0 - \mathbf{v}_{\text{sp}}t$ is computed, from which η_{interp} is calculated using the Newton-Raphson procedure. Once all η_{interp} are calculated, the interpolation matrix terms from Equation (4.18) are assembled in the form of a sparse linear operator, $\mathcal{L}_{d\to i}$, that takes a vector of donor values and outputs the values at $y^{\dagger}_{\text{interp}}$. The interpolation procedure just involves passing the function values and derivatives, with respect to the local coordinates system η (denoted by $\mathbf{D}_k$), as a vector to $\mathcal{L}_{d\to i}$.

4th-order Hermite scheme, with $m = 1$. The DNS code then implements this as illustrated in Figure 4.5. By construction this procedure is linear, and so does not require any additional modifications to implement it for the evaluation of the direct or adjoint Jacobian.

4.2.3.2 Effects of Interpolation on Modified Wavenumbers

It is worth noting here that, much like the well-known effects of finite difference schemes on the modified wavenumber of a harmonic signal[6], the interpolation too introduces additional modifications to the wavenumber. In addition, the interaction between the finite difference scheme, used to compute $\mathrm{d}f/\mathrm{d}\eta$, and the interpolation scheme should also be quantified.

To do so, consider the Fourier component $f(\eta) = \exp\left(i\alpha\Delta\eta(\eta + 0.5)\right)$, where α is the spatial wavenumber, $\Delta\eta$ represents the regular spacing between consecutive points in the computational space, and the coordinate transformation $(\eta + 0.5)$ is

[6]In the thesis, see Figure 2.2 and the surrounding discussion.

added so that $f(-0.5) = 1$. Upon applying this to Equation (4.17), one arrives at

$$\tilde{f}(\eta) = C_0^0(\eta) + i\tilde{\alpha}_{deriv}\Delta\eta C_0^\eta(\eta) + C_1^0(\eta)e^{i\alpha\Delta\eta} + i\tilde{\alpha}_{deriv}\Delta\eta C_1^\eta(\eta)e^{i\tilde{\alpha}_{deriv}\Delta\eta}, \tag{4.21}$$

where $\tilde{\alpha}_{deriv}\Delta\eta$ is the modified wavenumber for the spatial differentiation scheme. The properties modified interpolation wavenumbers are shown in Figure 4.6 for $\eta = 0$, as this is the point of maximum interpolation error, and results are also shown for $\tilde{\alpha}_{deriv}\Delta\eta = \tilde{\alpha}_{CC3}\Delta\eta$ and $\tilde{\alpha}_{deriv}\Delta\eta = \alpha\Delta\eta$, to establish the effect of the spatial scheme.

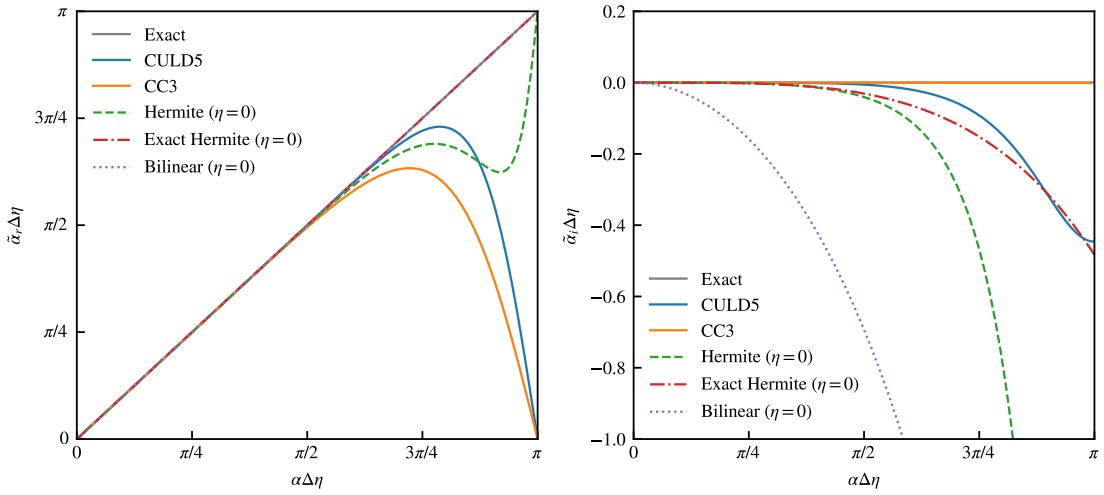


Figure 4.6: The modified wavenumbers for the fourth-order Hermite and linear schemes are shown, at $\eta = 0$, alongside the modified wavenumbers for the interior spatial schemes CULD5 and CC3. Results for the Hermite scheme are shown with $\tilde{\alpha}_{deriv}\Delta\eta$ obtained with numerical differentiation using the CC3 scheme and for exact differentiation (labelled "Exact Hermite").

The results in Figure 4.6 show that numerical differentiation has a considerable effect on the modified wavenumber of the fourth-order Hermite scheme[7] when compared to the exact differentiation. These primarily manifest as higher dispersion at wavenumbers $\alpha\Delta\eta > \pi/2$ and dissipation at wavenumbers $\alpha\Delta\eta > \pi/4$, which is equivalent to waves resolved by fewer than four and eight points per wavelength,

[7]Note also that the bilinear scheme does not make use of the derivatives, and so it is unaffected by the choice of spatial schemes.

respectively. In practice, girds are designed to resolve waves at much higher resolutions, with between 16 and 24 points per wavelength considered acceptable for the resolution properties of the CULD5 and CC3 schemes (see Table 2.1 in the thesis). It should also be noted that the expansion of the grid in the vertical direction, to become essentially equispaced in the vicinity of the sliding plane, has the effect of lowering the wavenumber of the incoming disturbances[8]. This means that the differences between the properties of the CULD5 scheme and the combined fourth-order Hermite and CC3 scheme at the sliding plane are likely to be less significant than indicated by the results in Figure 4.6.

For three-dimensional calculations these properties do pose an additional challenge because turbulent structures would largely manifest at high-wavenumbers that would be damped if they are not sufficiently resolved. Using Figure 4.6 it is clear to see that, with $\alpha\Delta\eta < \pi/2$, the interpolation scheme is less dispersive than the spatial schemes, and its dissipation properties are largely invariant to the choice of the differentiation scheme. Despite this, the interpolation is more dissipative than the interior spatial schemes for $\alpha\Delta\eta = \pi/2$, although note that, the grid resolution is normally selected to have at least $\alpha\Delta\eta \leq \pi/4$ to resolve the unsteady flow features. At this point the interpolation is both non-dispersive and only weakly dissipative. Further improvements to this framework can be made through the selection of spectral interpolation schemes, but, as mentioned before, these are likely to introduce dense interpolation matrices that may be difficult to manipulate for the adjoint mode. For this reason the fourth-order Hermite scheme is preferred.

4.2.3.3 Communications for the Sliding Grids

Communications at the sliding plane are implemented using the Message Passing Interface (MPI), and are designed to minimise the amount of data that needs to be exchanged, with the procedure summarised in Figure 4.7.

[8]This can be thought of in much the same way as the prolongation step in a multigrid method.

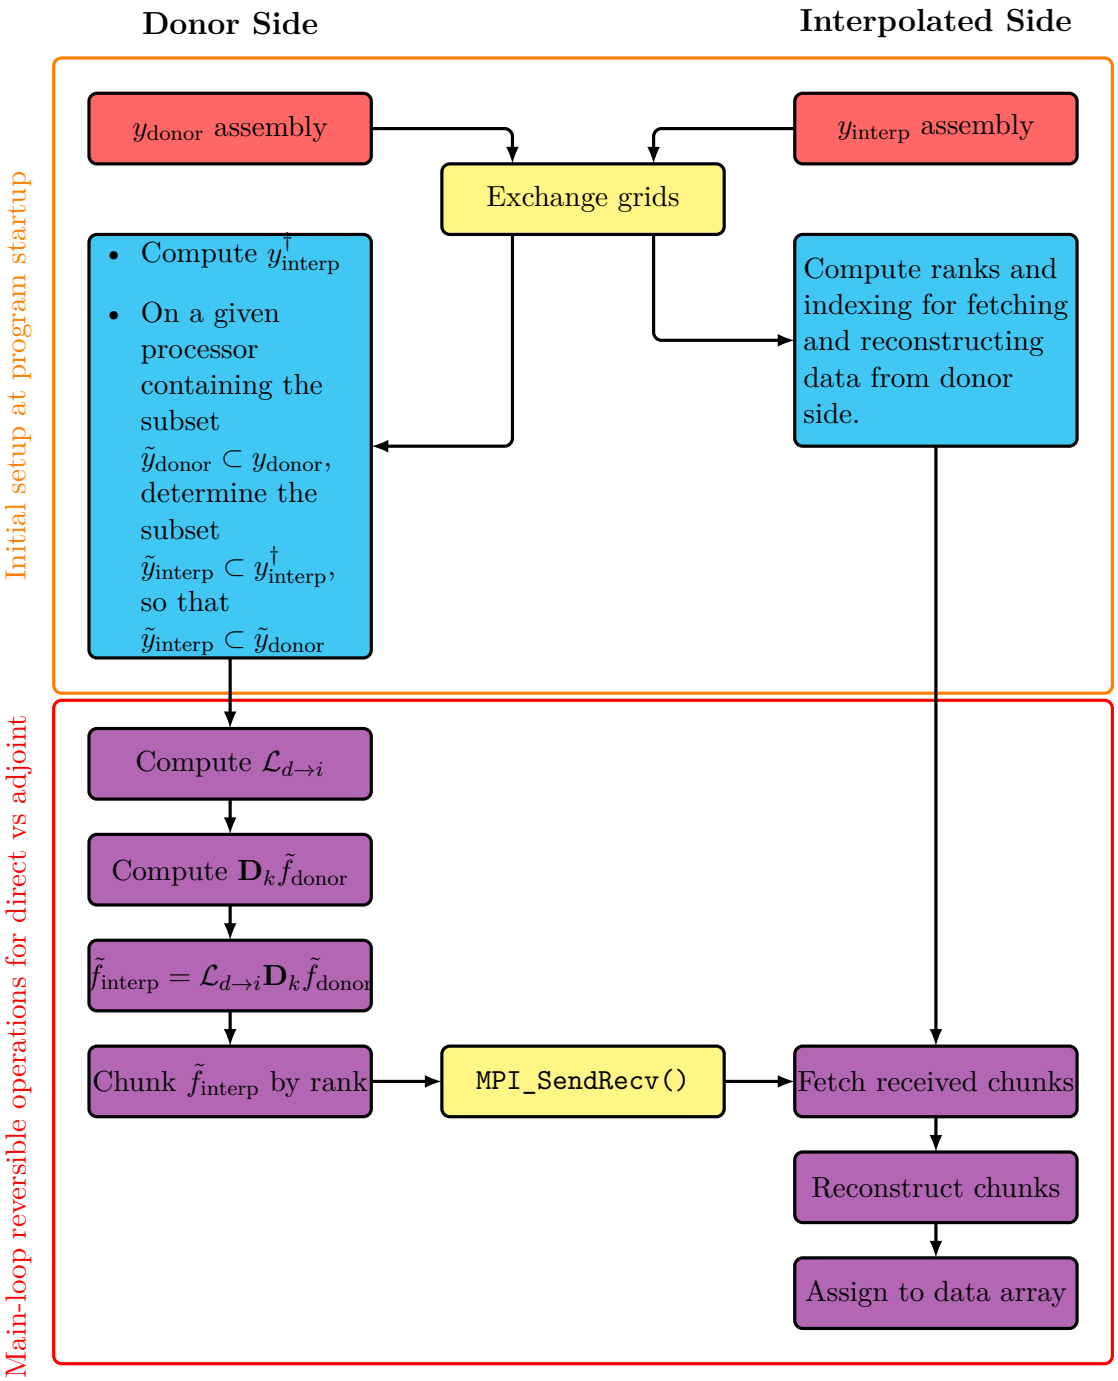


Figure 4.7: Procedure for data management and communications at the sliding interface. Note that this occurs twice per time step, first for the donor points on the left, and then for the donor points on the right. The quantities $\tilde{f}_{\text{interp}}$ denote the interpolated value at the subset of interpolation points that overlap with the boundaries of the current processor, $\tilde{y}_{\text{interp}}$.

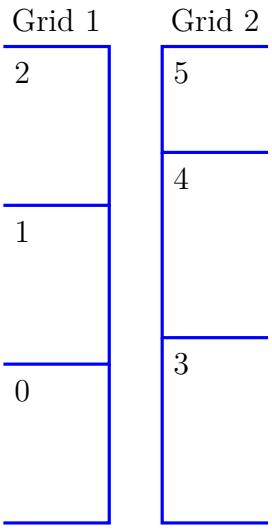


Figure 4.8: The distribution of processors either side of the sliding plane. For example, processor 4 sends and receives chunks of data from processors 1 and 2, and performs reconstruction back into a contiguous array when recovering the local data.

Since the interpolation requires the metric and the derivative of the state variables at the donor points, we perform all the calculations on the 'donor side' of the grid. The coordinates of the interpolation points are passed over to the donor side at the start of the computation with the coordinate transformation performed at each timestep on locally stored data. Then, a simple interval search determines the pair of donor points used in the construction of the local coordinate system. The sparse interpolation operator is created and the data interpolated as per the procedure in Section 4.2.3.1.

In practice, the data is distributed over a number of different processors arranged vertically at the sliding plane, as shown in Figure 4.8. As a result, each interpolation point is associated with an origin and destination processors so that they may be 'chunked' and reconstructed in the correct order once the data is sent and received at the donor and interpolation end, respectively. Doing so allows for asynchronous communications with each segment of data being sent directly to the processors that require it.

4.2.4 Linear and Adjoint Dynamics

4.2.4.1 Linearisation, Direct and Adjoint

Let $\mathbf{q}_d$ and $\mathbf{q}_i$ be the data and $\mathbf{x}_d$ and $\mathbf{x}_i$ be the coordinates of the donor and interpolation points, respectively. Then, let $\mathcal{L}_{d\to i} = \mathcal{L}_{d\to i}(\mathbf{x}_d, \mathbf{x}_i)$ be the linear operator that interpolates the solution from the donor to the interpolation points.

The entire communications procedure is then

$$\mathbf{q}_i = \mathbf{M}\mathcal{L}_{d\to i} D_k \mathbf{q}_d \tag{4.22}$$

Now, if we define a vector $\mathbf{x} = (\mathbf{x}_d; \mathbf{x}_i)$, we can represent the communications matrix as a block matrix of the form

$$\mathbf{M} = \begin{bmatrix} \mathbf{I}_{n_i} & \mathbf{0}_{n_i} \\ \mathbf{I}_{n_i} & \mathbf{0}_{n_i} \end{bmatrix}. \tag{4.23}$$

We note however that although we implement this using the equivalent matrix-free operation $\mathbf{q_i} = \mathbf{q_d^i}$, followed by $\mathbf{q_d^i} = \mathbf{0}$, Equations (4.22)–(4.23) are nonetheless a helpful aid for visualising the communications.

Due to the spatial coupling of the equations, the naive approach to linearising the governing equations Equations (4.1)–(4.3) would require the individual perturbation of every degree of freedom sequentially. Of course this incurs significant computational costs on the order of $4(N+1)$ nonlinear function evaluations for a system with N degrees of freedom.

An alternative method, derived by Fosas de Pando et al. [2012], uses the observation that the spatial coupling in the governing equations is introduced through the spatial differentiation operators. The solution is to break down the Navier-Stokes equations into sub-functions that depend on a set of variables $\mathbf{q}$ and their derivatives $\mathbf{q}_+$, $\mathbf{q}_-$, and $\mathbf{q}_c$, where $\star_\pm$ and $\star_c$ denote the positive and negative upwinding and the central derivatives respectively. Now the value of each of these sub-functions is independent at each spatial point, and hence the linearisation of that particular function may be performed simultaneously at all points within the domain.

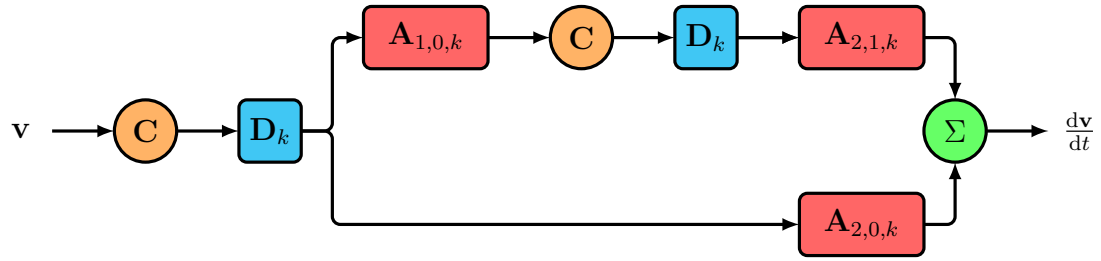


Figure 4.9: The linearisation procedure used to obtain the linear operator, with $\mathbf{A}_{\star,\star,\star}$ matrices corresponding to the linearisation of code sub-blocks, $\mathbf{D}_k$ referring to the matrix-free differentiation operation and the $\mathbf{C}$ operator denoting the sliding plane communications.

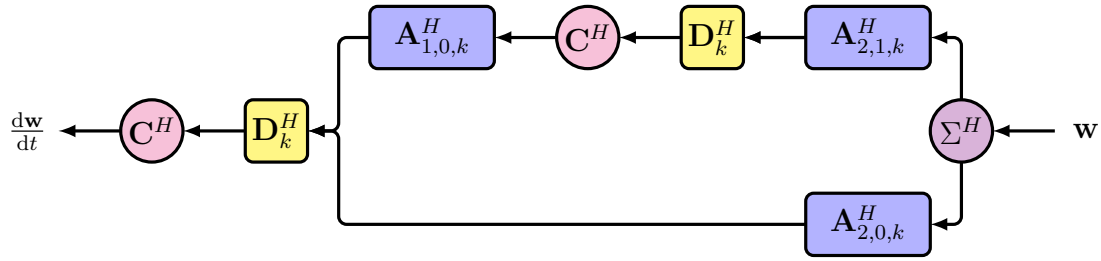


Figure 4.10: The assembly of the adjoint operator simply reverses the order in which the forward linearisation is performed, substituting the corresponding adjoint code, $\star^H$, for each linear operation. Note that one must take into account the grid metric when computing the adjoint.

For each sub-function $f = f(\mathbf{x}, \mathbf{q}_k)$, where $k \in \{0, +, -, c\}$ denotes the corresponding first derivative of $\mathbf{q}$, we compute the finite-difference derivative

$$\mathbf{A}_{ijk} = \frac{f(\mathbf{q}_{\mathrm{bf}} + \epsilon\mathbf{1}) - f(\mathbf{q}_{\mathrm{bf}})}{\epsilon}, \tag{4.24}$$

with ϵ picked to minimise the cancellation error. We note also that the sliding plane communications operate independently of the domain and therefore the linearised sub-functions may be computed on all domains simultaneously.

Following the linearisation, the direct right-hand side is reassembled as in Figure 4.9,

with the algebraic form given by Equation (4.25),

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \left(\mathbf{A}_{2,0,k} + \mathbf{A}_{2,1,k}\mathbf{D}_k\mathbf{C}\mathbf{A}_{1,0,k}\right)\mathbf{D}_k\mathbf{C}\mathbf{v}, \tag{4.25}$$

with the operator $\mathbf{C}$ denoting the sliding plane communications.

To this end, the adjoint operator $\mathbf{A}^H$ follows trivially and the adjoint right-hand side can be written as follows

$$\frac{\mathrm{d}\mathbf{w}}{\mathrm{d}t} = \mathbf{C}^H\mathbf{D}_k^H\left(\mathbf{A}_{2,0,k}^H + \mathbf{A}_{1,0,k}^H\mathbf{C}^H\mathbf{D}_k^H\mathbf{A}_{2,1,k}^H\right)\mathbf{w}, \tag{4.26}$$

with $\mathbf{w} = \begin{bmatrix} p^\dagger & s^\dagger & \mathbf{u}^\dagger \end{bmatrix}$ as the adjoint variable, and the adjoint communications and interpolation

$$\mathbf{C}^H = \mathcal{L}_{d\to i}^H\mathbf{M}^H, \tag{4.27}$$

where now we see that

$$\mathbf{M}^H = \begin{bmatrix} \mathbf{I}_{n_i} & \mathbf{I}_{n_i} \\ \mathbf{0}_{n_i} & \mathbf{0}_{n_i} \end{bmatrix}, \tag{4.28}$$

which is equivalent to $\mathbf{w}_d = \mathbf{w}_d + \mathbf{w}_i$ and $\mathbf{w}_i = \mathbf{0}$. Moreover, the sparse matrix $\mathcal{L}_{d\to i}^H$ is easily obtained using the built-in transposition functions within many popular sparse matrix libraries (e.g. Eigen [Guennebaud et al., 2010]). This procedure is shown in Figure 4.10.

Note that the MPI communications at the sliding interfaces are simply reversed in the adjoint, so that if a particular processor rank is used as an origin from which data is received in the direct case, then the same rank becomes a destination in the adjoint. The data must also be chunked so that a chunk of the same shape is sent in the adjoint as is received in the direct.

4.3 Validations

The numerical code with the sliding interfaces has been validated against exact and approximate solutions to the compressible Navier-Stokes equations. Here we present a selection of test cases that demonstrate the accuracy and validity of the sliding plane implementation.

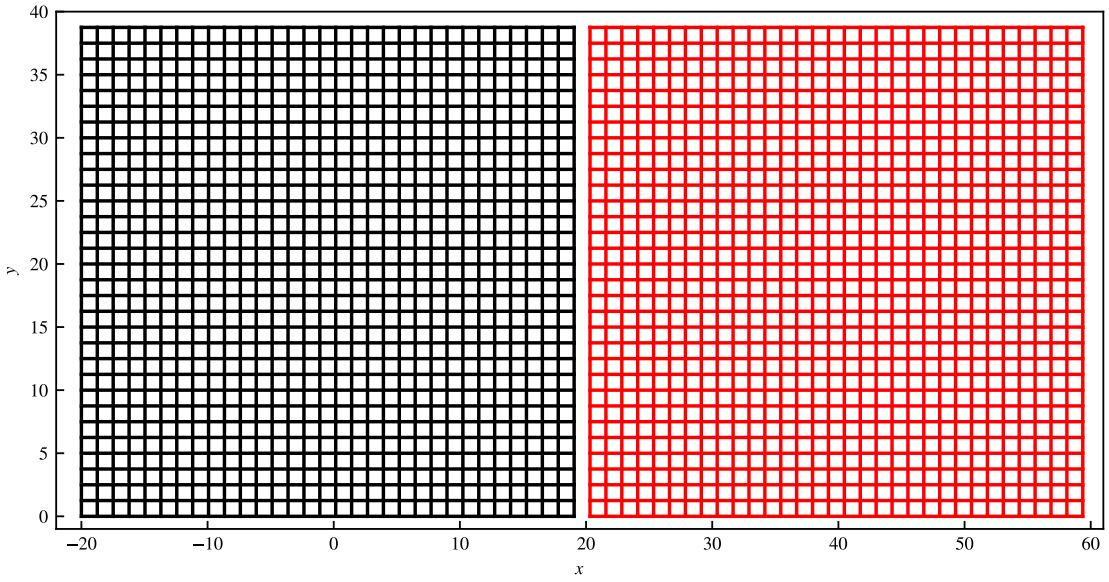


Figure 4.11: Cartesian validation domain with the sliding plane interface conditions located between the red and black meshes. Note that every fourth grid line is shown, for clarity.

4.3.1 Stationary Interface

For illustrative purposes, the initial validation of the sliding plane is performed on a simple Cartesian domain, as shown in Figure 4.11. The domain consists of two uniform mesh blocks arranged adjacent to each other along the x-axis, with the interface conditions implemented at the edges closest to the respective neighbouring domain.

Since the validity of the numerical scheme is known [Fosas De Pando, 2012][9], the effect of the interface is evaluated by comparing to solution within a single mesh, which is generated by combining the two separate meshes. Disturbances of various types are then superposed upon a uniform base flow and the solution is left to evolve over a set time period. After this the relative L_2 error

$$E_{L_2} = \frac{\|(\mathbf{v} - \mathbf{v}_{sb})\|_{L_2}}{\|\mathbf{v}_{sb}\|_{L_2}}, \tag{4.29}$$

where $\mathbf{v}_{sb}$ is the single block solution with no sliding plane, is computed and the local error is plotted over the entire domain.

[9] In the thesis, also see Chapter 2.



(a) (b)

Figure 4.12: The coinciding grid, (a), and the grid displaced vertically by $dy/2$, (b), used in the validation of the communications and interpolation, respectively.

The tests run in three stages. First, the two blocks are aligned so that the horizontal grid lines coincide across the boundary. This evaluates the validity of the communications routines that arrange and transfer data across the interface. The second part offsets the right grid vertically by $dy' = dy/2$, so that now the interface experiences maximum interpolation error. This is shown diagrammatically in Figure 4.12. Finally, we set $\mathbf{v}_{sp} \neq 0$ for the red domain in Figure 4.11, and evaluate the errors for a case where there is relative motion between the two sides.

Numerical details of the test cases are given in Table 4.1.

Test case	Mesh size	Ma	Re	Pr	u_0	v_{sp}	$t = t_{\text{end}}$
Acoustic pulse	128×128	0.4	∞	1	0.05	0.5	15.0
Isentropic vortex	128×128	0.4	∞	1	0.05	0.5	80.0
Heat dissipation	128×128	0.4	2.0	0.72	0.05	0.5	80.0
Viscous dissipation	128×128	0.4	2.0	1×10^{10}	0.05	0.5	80.0

Table 4.1: Cartesian test case parameters

4.3.1.1 Acoustic Pulse

The acoustic pulse disturbance, as shown in panel (a) of Figure 4.13, is defined by (see e.g. [Tam and Webb, 1993])

$$p = 1 + \frac{\epsilon}{2a}\int_0^\infty \exp\left\{-\frac{\xi^2}{4a}\right\}\cos\left(\frac{\xi(t+2.5)}{M}\right)J_0(\xi r)\xi\, d\xi, \tag{4.30}$$

$$s = 0, \tag{4.31}$$

$$u = u_0 + \frac{\epsilon x}{2ar\gamma M}\int_0^\infty \exp\left\{-\frac{\xi^2}{4a}\right\}\sin\left(\frac{\xi(t+2.5)}{M}\right)J_1(\xi r)\xi\, d\xi, \tag{4.32}$$

$$v = \frac{\epsilon y}{2ar\gamma M}\int_0^\infty \exp\left\{-\frac{\xi^2}{4a}\right\}\sin\left(\frac{\xi(t+2.5)}{M}\right)J_1(\xi r)\xi\, d\xi - v_{\text{sp}} \tag{4.33}$$

where J_0 and J_1 are the Bessel functions of the first kind of order 0 and 1 respectively, $\mathbf{v}_{\text{sp}}$ is the *vertical* speed of the sliding interface, and the user-defined parameters are $\epsilon = 0.01$, $u_0 = 0.05$, $a = \log 2$, and $r = \sqrt{x^2 + y^2}$.

The pulse is then placed at $t = 0$ in the middle of the left domain, $(x_0, y_0) = (0, 20)$, and the solution evolved until $t = 15$. At this point the acoustic pulse propagates through the sliding interface as in Figure 4.13, and is compared to the single domain solution. The error for the displaced and moving interface grids is then calculated and shown in the bottom panels of Figure 4.13, with the relative error 1.024×10^{-8} and 6.376×10^{-6}, respectively.

4.3.1.2 Isentropic Vortex

An isentropic vortex is given by (see e.g. Spiegel et al. [2015] and references therein)

$$p = \left(1 - \frac{(\gamma-1)S^2M^2{e_f}^2}{8\pi^2}\right)^{\frac{\gamma}{\gamma-1}}, \quad s = 0, \tag{4.34}$$

$$u = u_0 - \frac{Sye_f}{2\pi R}, \quad v = \frac{Sxe_f}{2\pi R} - v_{\text{sp}}, \tag{4.35}$$

with e_f defined as

$$e_f = e^{-\frac{-\left(x^2+y^2\right)}{2R^2}}, \tag{4.36}$$

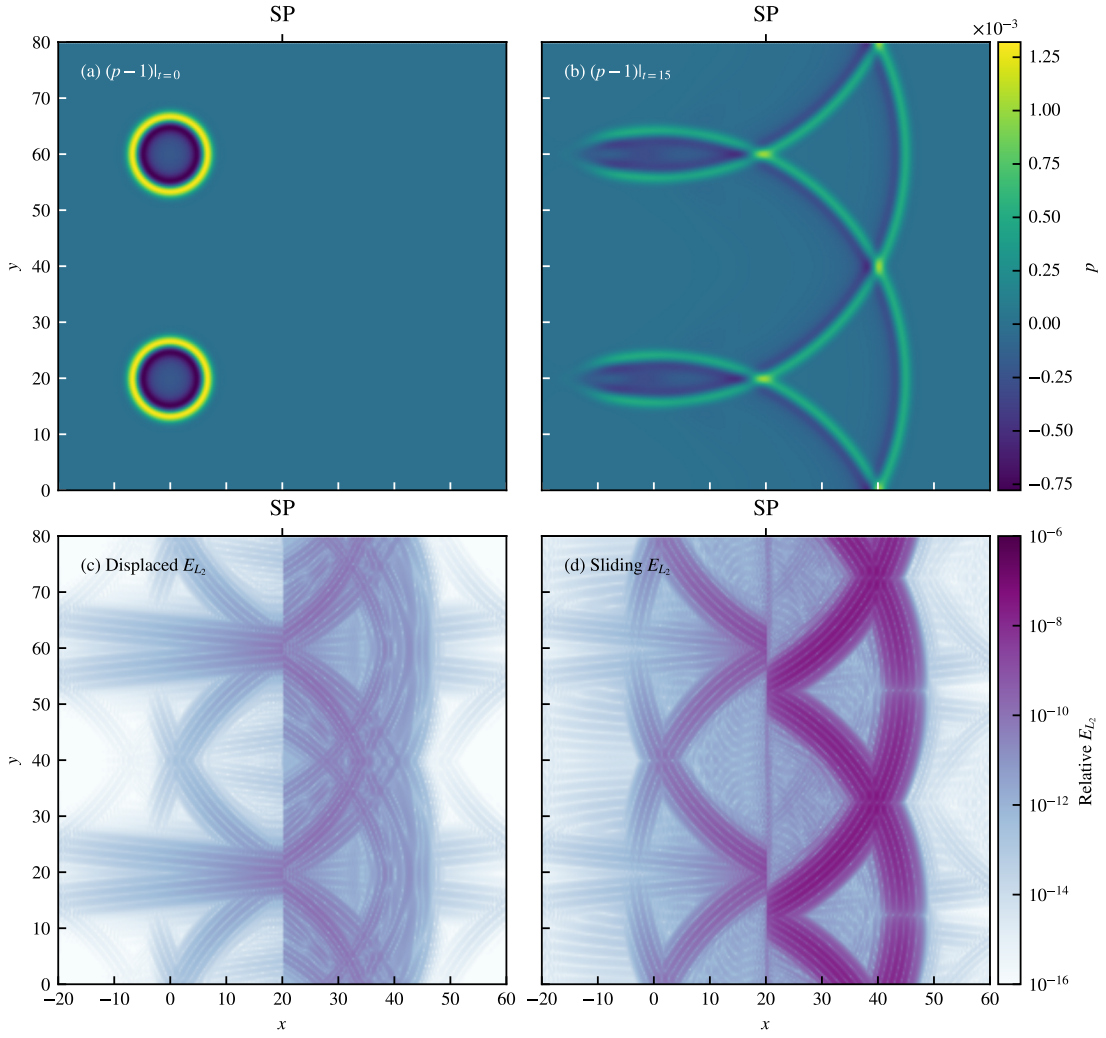


Figure 4.13: Pressure pulse positioned at $(x_0, y_0) = (0, 20)$ at $t = 0$, is evolved to $t = 15$, with the pressure field for the initial condition (a) and the final pressure field (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grids are shown in the respective figures.

and free parameters $S = 3.5$, $R = 3.0$ and $u = 0.05$.

This structure is placed at the sliding plane and left to evolve for $T = 80$ time units, and the results shown in Figure 4.14. Relative errors for the stationary displaced and sliding solutions are 7.873×10^{-8} and 7.758×10^{-8}, respectively.

4.3.1.3 Viscous Dissipation

To evaluate viscous dissipation across the interface, an identical vortex to that defined in Equations (4.34)–(4.36) is generated at the sliding plane. The Reynolds number

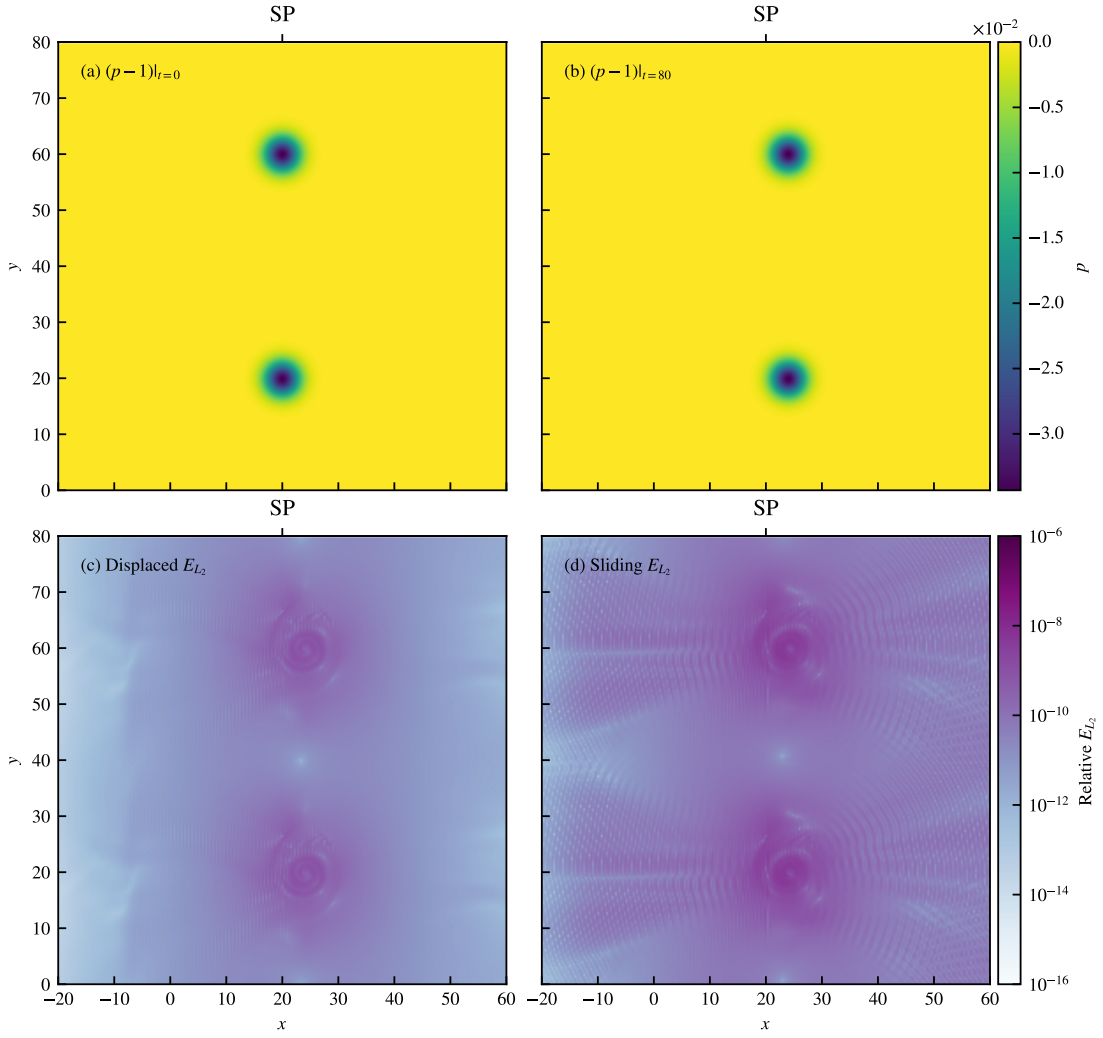


Figure 4.14: An isentropic vortex positioned at the sliding plane, $(x, y) = (20, 20)$ at $t = 0$, is evolved to $t = 80$, with the pressure field for the initial condition (a) and the final pressure field (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grid are shown in the respective figures.

is then decreased to Re = 2 so that viscous forces dominate the flow dynamics. The vortex is again left to dissipate for 80 time units, and the errors taken as before. These are illustrated in Figure 4.15, with the relative errors 5.346×10^{-10} and 1.855×10^{-6}, respectively.

4.3.1.4 Heat Dissipation

The final test case that we present here is the evaluation of the error in the heat dissipation term of the governing equations. For this case, a low Reynolds number

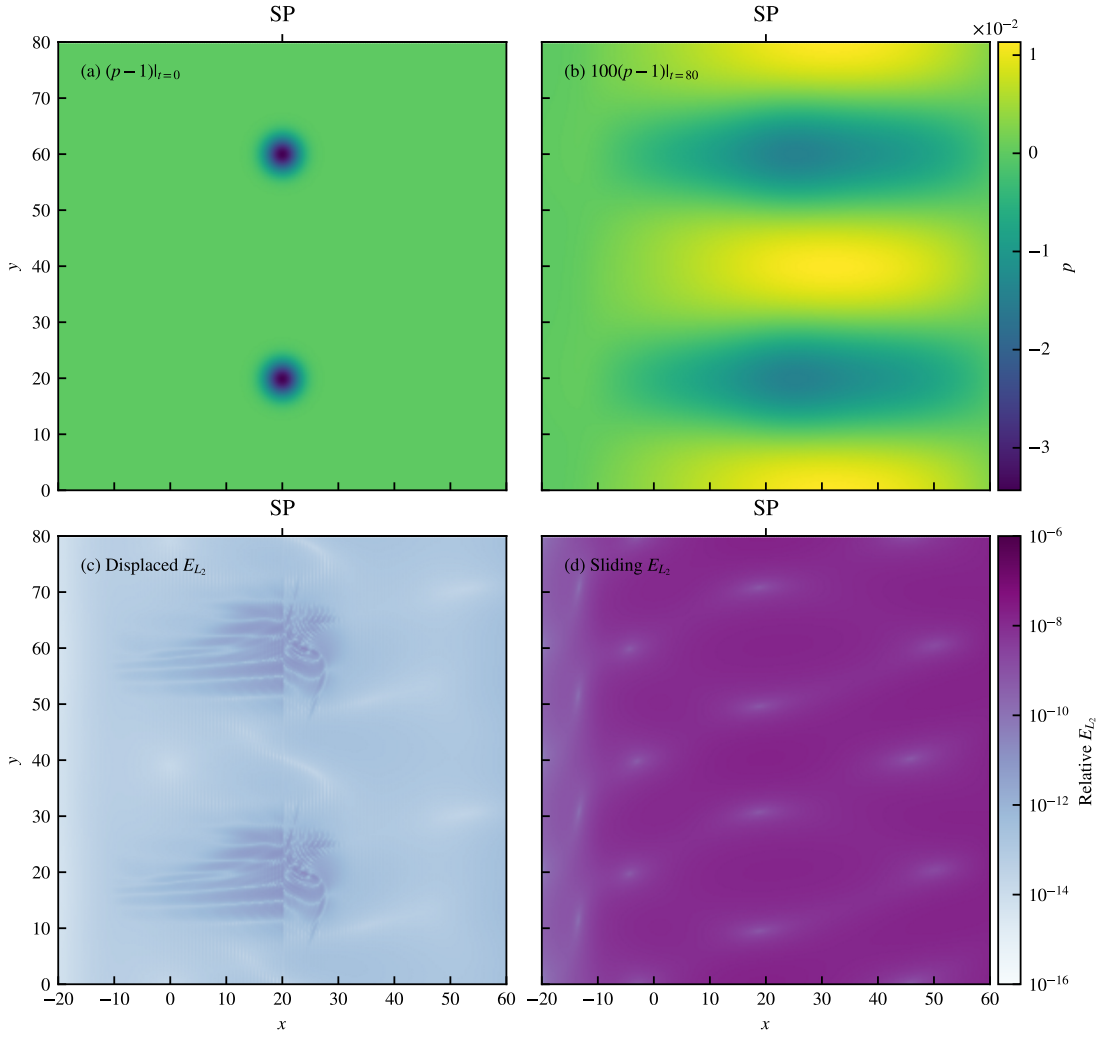


Figure 4.15: A dissipating vortex positioned at the sliding plane, $(x, y) = (20, 20)$ at $t = 0$, is evolved to $t = 80$, with the pressure field for the initial condition (a) and the final pressure field (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grid are shown in the respective figures.

and a high Prandtl number are set so that the heat dissipation term dominates the viscous effects. The same base flow is used as before and a hot spot defined by

$$p = 1, \quad s = \frac{e_f}{2\pi R}, \quad u = u_0, \quad v = -v_{\text{sp}}, \tag{4.37}$$

is placed at $(x_0, y_0) = (15, 20)$. This asymmetric placement was chosen so that the front from the dissipating hot spot has a perpendicular component across the sliding plane, rather than tangent to it had it been placed directly at the sliding plane at $(20, 20)$. As before, results are shown in Figure 4.16, with the relative errors

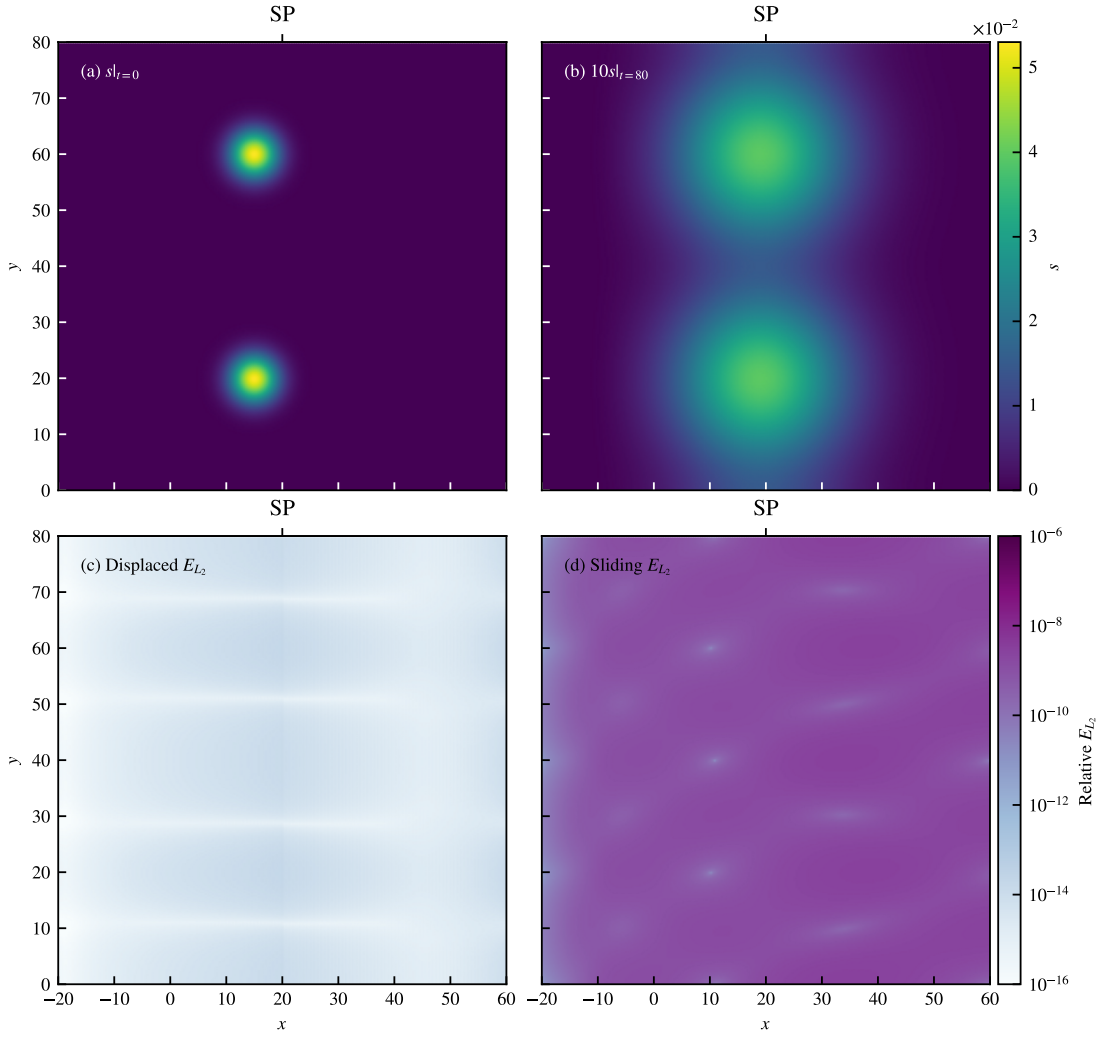


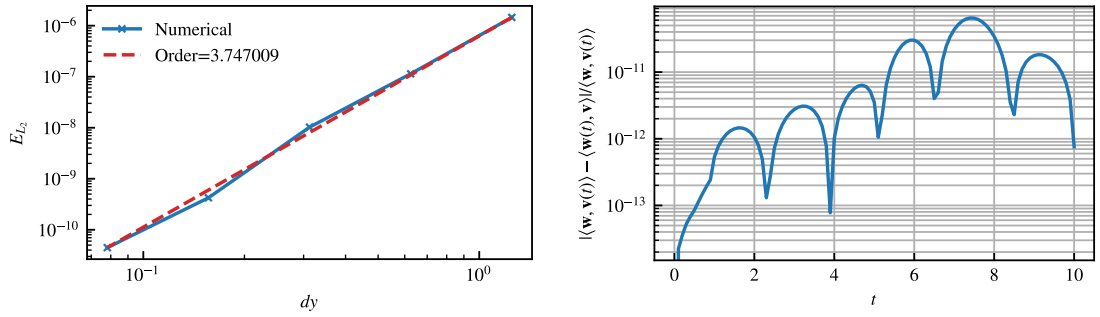
Figure 4.16: Hot gas 'bump' positioned at $(x, y) = (20, 20)$ at $t = 0$, is evolved to $t = 80$, with the entropy field for the initial condition (a) and the final entropy field at $t = 80$ showing the heat diffusion process (b). The relative L_2 errors for a stationary displaced (c) and sliding (d) grid are shown in the respective figures.

1.406×10^{-12} and 2.889×10^{-7}, for the stationary and sliding meshes, respectively.

4.3.2 Interface Order of Accuracy

The error derived from the sliding interface is computed by comparing the acoustic pulse case in Section 4.3.1.1 at varying grid resolutions for the displaced grid case as illustrated in Figure 4.12b. This case specifically checks the effect of the interpolation operator on the overall order of the scheme.

It is clear to see from this that the additional error from the sliding plane and



(a) (b)

Figure 4.17: Order of accuracy of the sliding plane at successively finer grids (solid blue) with the matching fit (dashed red), showing an order of accuracy consistent with that of the sliding plane interpolation scheme error characteristics, in (a). The timestepper direct-adjoint error with sliding interface in (b).

the interpolation scheme remains almost 4th-order accurate. Although this is lower than the overall order of the internal numerical scheme, it is nonetheless high enough for our intended CAA applications.

It should of course be noted that this is valid for when the grid on either side of the interface is approximately equal in size. An interpolation from a very coarse grid to a very fine one would necessarily result in large interpolation errors. To mitigate this problem, we found that expanding regions in refined grid, such as areas around the wake and boundary layers, so that vertical grid sizes are comparable on either side of the sliding plane, yielded very satisfactory results. It is not necessary, and indeed may not be feasible, to prescribe an identical point distribution for both grids; moreover the interpolation and communications procedure was tested with a large variety of grids and was confirmed to be robust and efficient for all tested cases.

4.3.3 Nonlinear Adjoint

4.3.3.1 Adjoint Operator

Validation of the adjoint occurs in two stages. For this we check the fundamental relation for the domain with the sliding plane

$$\langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle = \left\langle \mathbf{A}^H \mathbf{w}, \mathbf{v} \right\rangle \tag{4.38}$$

for which we obtain

$$\frac{\left| \langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle - \left\langle \mathbf{A}^H \mathbf{w}, \mathbf{v} \right\rangle \right|}{|\langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle|} = 1.839 \times 10^{-16}, \tag{4.39}$$

when checked with random input fields $\mathbf{w}$ and $\mathbf{v}$.

Secondly, for the same input $\mathbf{w}$ and $\mathbf{v}$, we check that the timestepping algorithm also satisfies the relation

$$\langle \mathbf{w}, \mathbf{v}(t) \rangle = \langle \mathbf{w}(t), \mathbf{v} \rangle \,, \quad \forall t \in [t_0, t_1], \tag{4.40}$$

where $\mathbf{v}(t)$ and $\mathbf{w}(t)$ represent the forward and adjoint at time t, respectively. The base flow is taken to be homogeneous, with flow conditions replicating the acoustic pulse case in Table 4.1. This choice is made as it avoids the need to checkpoint the nonlinear solution for all time, although it should be noted that this does not affect the result because the error is determined by the manipulation of the linear field. The results are shown in Figure 4.17b, where it can be seen that the error is bounded above by 10^{-10}, with the 'beating' phenomenon arising due to imprecise alignment of the grid in the direct and adjoint, as a consequence of a negative sign introduced in the sliding plane velocity in the adjoint. This is not an issue, at least in subsequent work presented here, because the checkpointing algorithms handle the adjoint by sending a decreasing 'forward' time signal, rather than using an increasing 'adjoint pesudo-time', and so, unlike this validation, the sliding grid calculation is identical for both the forward and backward problems.

4.3.3.2 Checkpointing

For many CAA applications, the *clean* propagation of unsteady wakes through the sliding interface is critical in order to avoid spurious results and the introduction of nonphysical feedback from the interactions between the interior and interface schemes. We note, moreover, that operator validations of the type described above in Equations (4.39)–(4.40), while essential, *do not provide any guarantee of how the interface treatment interacts with the direct and adjoint solutions.* For instance, using this scheme as described above but with a reduced overlap region at the sliding plane of, say, four points would appear to give correct nonlinear and direct fields as well as satisfying the adjoint relations, but the adjoint field would be completely contaminated by the sliding plane due to the different dispersion characteristics of the boundary closure in the spatial discretisation scheme, and the additional sensitivity that this would impose.

Unstable shear layers are particularly well-suited to validating both the TLM and the nonlinear-adjoint model, since they can be constructed to exhibit a global instability that produces wakes that are qualitatively similar to unsteady vortex shedding seen in many fluid-structure interaction problems. These flows may be created without introducing additional solid bodies into the flow, thus simplifying the mesh, and consequently the timestep requirements of the simulation.

To confirm this, we use the case of a perfectly expanded, isothermal jet injected into a slower-moving fluid, that will undergo a transition to instability upstream of the the sliding plane. Pressure is therefore assumed to be uniform everywhere and the velocity profile within the jet is given by Equation (4.42) [Reichert and Biringen, 1997], and illustrated in Figure 4.18b

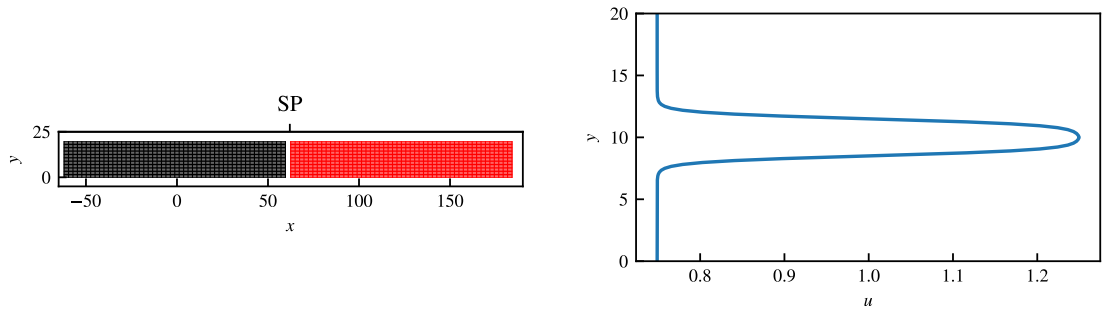
$$u(y) = \mathcal{K}\tanh(-2\,|y-10|+3)+1+\alpha\delta\psi, \tag{4.41}$$

$$\mathcal{K} = \frac{U_0 - U_\infty}{U_0 + U_\infty}, \tag{4.42}$$

with symmetry about $y = 10$, where U_0 and U_∞ are the flow velocities of the fluid

Parameter	Value	Notes
Grid dimensions	256×128 (per side)	See Figure 4.18a
Physical dimensions	124×20 (per side)	See Figure 4.18a
Ma	0.8	
γ	1.4	
Re	10,000	
Pr	0.72	
U_0	1.67	
U_∞	1.0	

Table 4.2: Parameters for TLM and Nonlinear Adjoint Case



(a) (b)

Figure 4.18: (a) The Cartesian domain used for the validation of the TLM and nonlinear adjoint models. The mesh is uniform, with dimensions 256×128, the inflow on the left boundary of the black grid, outflow on the right of the red grid and periodic boundary conditions on the top and bottom of both the red and black grids. The grid is plotted for every 5th grid line so that the structure of the grid can be visualised. (b) The inlet jet velocity profile $u(y)$ at $x = -62$ (the inlet of the computational domain).

inside the jet and in the far field respectively, and $u = u(x, y)$ is the x-component of the velocity field. The vertical component of velocity is set to $-v_{sp} = -1$, which is the local mesh velocity in the global reference frame for the moving mesh. To accelerate the transition to instability, we also seed the velocity field with a uniformly distributed random variable $\delta\psi$ in the region enclosed by the lines $|x - 10| - 0.5 = 0$ and $x = -40$, with an amplitude $\alpha = 0.05$.

As before a uniform Cartesian grid is used, and is illustrated in Figure 4.18a, with the key parameters of the simulation summarised in Table 4.2.

We obtain the well-developed jet flow by integrating the nonlinear equations

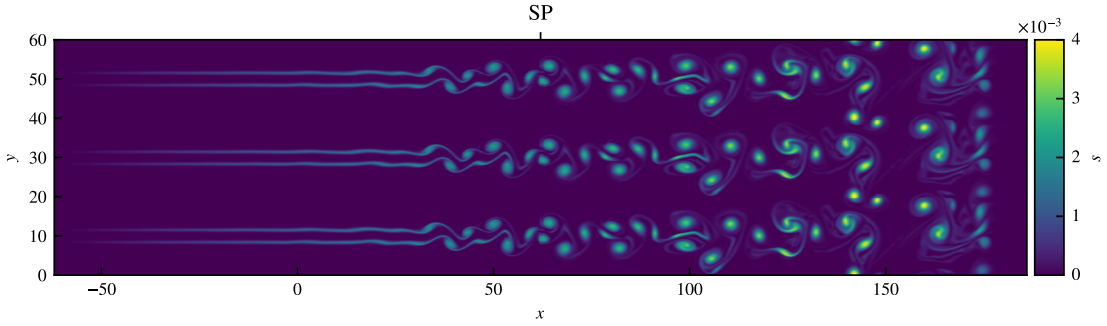


Figure 4.19: The unstable 1-periodic jet flow through the sliding plane illustrated at $t_t = 100$, with virtual periodic images above and below for a more complete illustration of the flow solution.

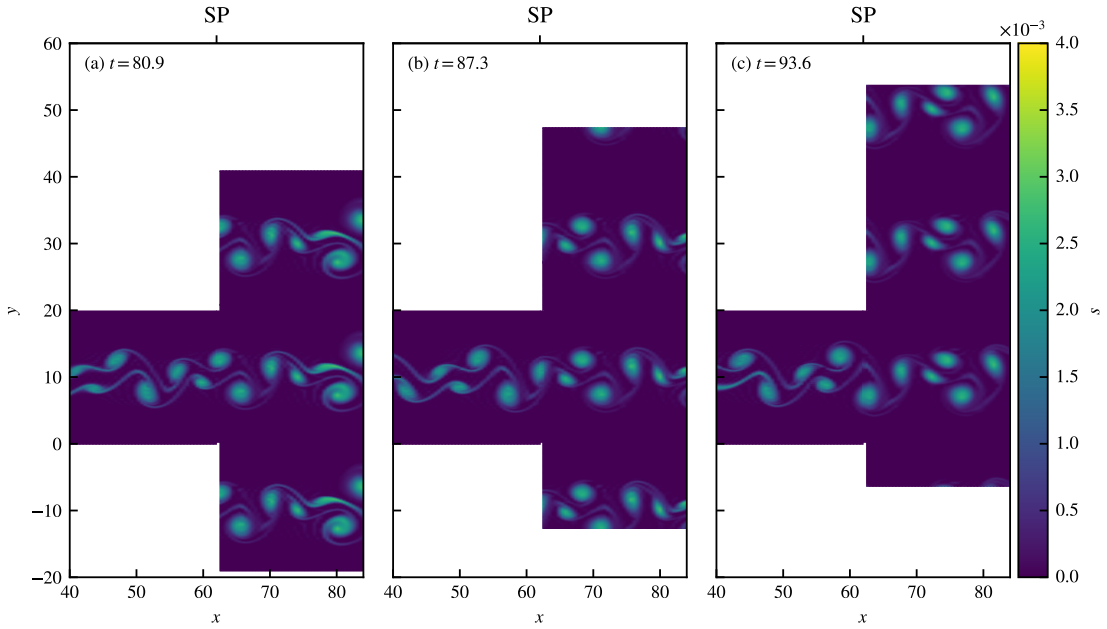


Figure 4.20: Snapshots of the interface region are shown here with the periodic images removed on the left but retained on the right for improved visualisation. The right mesh moves towards the top of the page at $\mathbf{v}_{\text{sp}} = 1.0$. As is clear to see, the wake passes cleanly through the interface.

until $t = 2000$, with the last snapshot being used as the new initial condition for the nonlinear-adjoint solution process. As in the process shown in Figure 4.2, the nonlinear solver is run using this initial condition until $t_t = 100$. The limit cycle is shown using the entropy field at the end of the simulation in Figure 4.19, showing the instability fully developed within the whole domain.

Typically the adjoint method is reserved for computing gradients of systems

governed by a large numbers of parameters, and for a set of N parameters the break-even point of the adjoint method against the cost of a sequential parameter-by-parameter finite difference approach occurs at around $N = 16 - 32$ in this code. This is strongly determined by the cost of re-linearising along the trajectory against the cost of a forward nonlinear function evaluation.

We begin by defining the cost function $\mathcal{J} = \mathcal{J}(\mathbf{v}, \boldsymbol{\eta})$ where $\boldsymbol{\eta}$ is the parameterised initial condition. A common norm used to study small fluctuations in compressible flows is the small-perturbation energy norm [Chu, 1965], where an energy of a volume of fluid $\mathcal{V}$ is

$$\mathcal{E} = \frac{1}{2}\int_{\mathcal{V}} \left(\frac{1}{\gamma^2 M^2}\frac{p^2}{\bar{p}} + \frac{\gamma - 1}{\gamma^2 M^2}\bar{p}s^2 + \rho^2 \|\mathbf{u}\|^2 \right) \mathrm{d}\mathcal{V}, \tag{4.43}$$

with the quantities $\bar{p}$ and $\bar{\rho}$ representing the pressure and density of the *time dependent flow* on top of which the perturbation $(p, s, \mathbf{u})$ is imposed. Naturally this is translated into the discrete form of the inner product

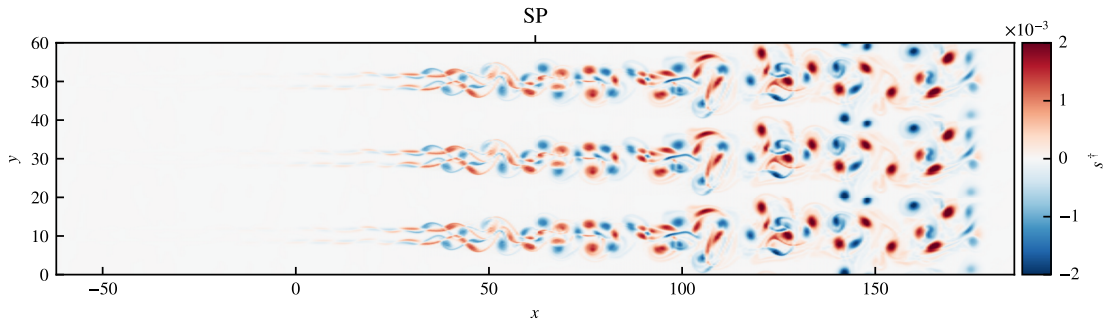
$$\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = \mathbf{w}^H \mathbf{M} \mathbf{v}, \quad \text{with} \quad \mathbf{M} = \mathrm{diag}\,\{\mathbf{m}_i\}, \tag{4.44}$$

and

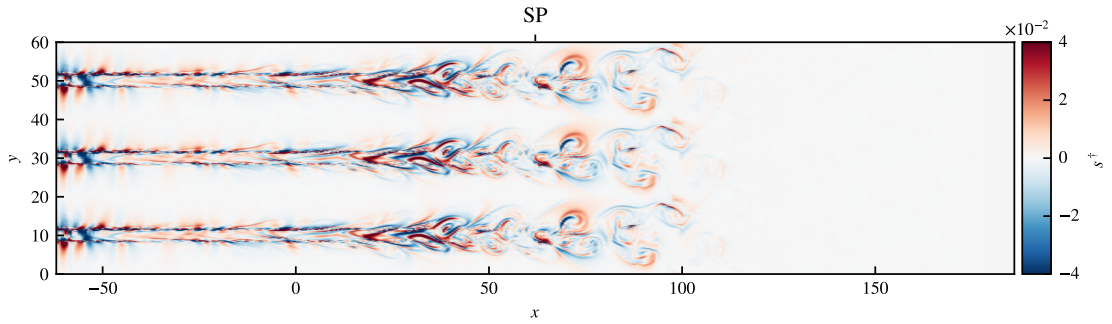
$$\mathbf{m}_i = \begin{pmatrix} \frac{1}{\gamma^2 M^2}\frac{1}{\bar{p}} & 0 & 0 & 0 \\ 0 & \frac{\gamma-1}{\gamma^2 M^2}\bar{p} & 0 & 0 \\ 0 & 0 & \bar{\rho} & 0 \\ 0 & 0 & 0 & \bar{\rho} \end{pmatrix}_i |\mathbf{x}(\xi)|_i \Delta\xi_i \Delta\eta_i \mu_i, \tag{4.45}$$

where we note the dependence of the discrete norm on the grid metric, which accounts for varying grid density, and denote the region of the domain that is used in the inner product by a masking term μ that takes the value 1 or 0 depending on whether the point is or is not taken in the inner product.

The gradient of this cost function with respect to the initial conditions is computed to show the propagation of the adjoint field across the sliding interface. Perturbation quantities are defined as $\tilde{\mathbf{v}} = \mathbf{v} - \bar{\mathbf{v}}$, where $\mathbf{v}$ is a nonlinear solution with



(a)



(b)

Figure 4.21: (a) The initial conditions, $\mathrm{d}J/\mathrm{d}\mathbf{v}\ (\mathbf{v}_N)$, for the nonlinear adjoint time stepping. (b) The sensitivity of the cost function (solution of the nonlinear adjoint system) to initial conditions. The sliding plane is located at x=62.

$\alpha = 0$, and $\bar{\mathbf{v}}$ is computed with $\alpha = 1 \times 10^{-3}$. Initial conditions for the adjoint $\mathrm{d}J/\mathrm{d}\mathbf{v}\ (\mathbf{v}_N)$ are shown in Figure 4.21a, with the final solution shown in Figure 4.21b. It is clear from this that the adjoint field is not affected by the sliding plane treatment, meaning that the gradients calculated by this code are independent of the sliding plane. This is significant as aeroacoustics demands such performance for both the large and small-scale structures of the flow.

4.3.4 Parallel efficiency

Deployment of this code to parallel clusters is of paramount importance for almost all practical applications, and limited computational resources require problems of this kind to scale efficiently when the number of processors and grid dimensions both increase.

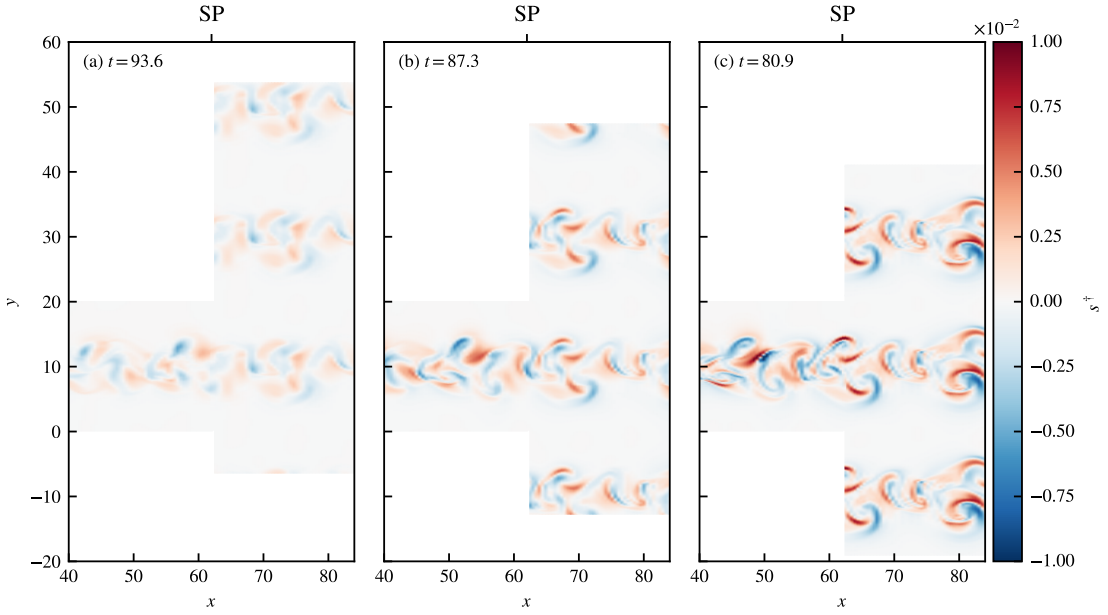


Figure 4.22: As with the nonlinear solution, the snapshots of the periodic interface are shown here. Note that the time integration occurs backwards in time and the adjoint wake moves from right to left.

To that end, we verify the parallel scaling of this code by using the test case from Section 4.3.1.1 with a mesh dimension of 1024×1024 per side, which is typical of the mesh dimensions encountered in many of our application cases. It is also advantageous to evaluate scaling losses using a Cartesian geometry as there is considerable freedom in assigning optimal processor distributions that one may not have when considering a general rotor-stator problem as in Section 4.4, as normally the grids for multi-passage meshes are optimised for a particular number of processors so that each has similar numbers of degrees of freedom.

In a fixed-size problem this code is expected to exhibit two sources of scaling losses. First is the increased communications overheads for block-to-block data transfer used for derivative calculations, as described in Section 4.2.1.1, and second is the procedure governing the communications and interpolation at the sliding interface between neighbouring grids, as shown in Section 4.2.3.3. A single-grid computation is then compared for a similar size grid partitioned in the middle with a sliding plane. The code is run for 200 time units and the total wall time

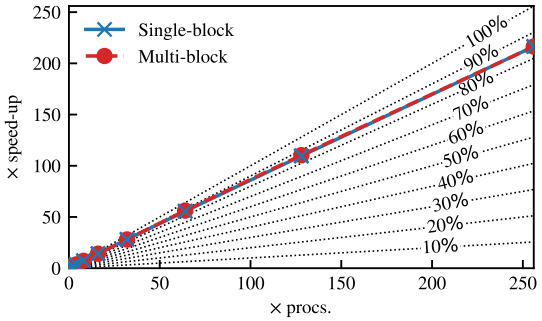


Figure 4.23: The scaling efficiency of a Cartesian domain case compared for both the single domain and multiple domains with a sliding plane, showing identical scaling performance.

for the evaluation of the right-hand side function is obtained and the timings of both the single and multi-block cases are obtained. This process is repeated for $N_{\text{procs}} = \{2^n\}_{i=1}^{9}$ and the results scaled with respect to the timings for $N_{\text{procs}} = 2$.

The results are shown in Figure 4.23, showing between an $85 - 90\%$ scaling efficiency for all tested cases. Typically we choose the CPU count to be as low as possible while still maintaining reasonable execution time, although the scalability offered by the code is advantageous for large-scale computations. The execution time is also not significantly affected by the sliding plane overheads as the vertical mesh dimension is relatively small. This is typical of many practical applications since the mesh has many more degrees of freedom in the interior domain than it does in the vertical direction at the sliding plane, and so the vertical dimensions of such meshes tend to be quite small with small numbers of processors at the boundary.

4.4 Applications

While we have shown the applicability, robustness and efficiency of our code with various simple examples in Section 4.3, they are hardly representative of real-world research problems. As a more concrete example, we present here two examples of problems that may be studied using this framework.

4.4.1 Gradient Calculations

At the start of this paper we motivated our use of the adjoint by demonstrating how the adjoint field could be used to derive the gradients of the cost function. Access to this gradient information immediately opens up many avenues for optimisation. Though we will not demonstrate optimisation in this paper, we nonetheless summarise the conjugate gradient method so that the reader is acquainted with the practical considerations of integrating the nonlinear adjoint gradient calculation with the optimisation.

Nonlinear Conjugate Gradient (CG) methods [Hestenes et al., 1952; Fletcher, 1964] are a family of particularly well-suited algorithms for large-scale unconstrained minimisation problems since they are matrix-free, low-storage and, at any point in the optimisation of $\mathbf{v}_k$, they require only access to the cost function value $\mathcal{J}(\mathbf{v}_k)$ and its gradient $\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_k)$. An outline of the CG algorithm is given in Algorithm 3. The

Algorithm 3 Nonlinear Conjugate Gradient Method

1: $k = 0$
2: Initial state $\mathbf{v}_0$
3: Nonlinear solve to get $\mathcal{J}(\mathbf{v}_0)$
4: Adjoint solve to get $\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_0)$
5: $\mathbf{p}_0 = -\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_0)$,
6: **while** $\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_k) \neq 0$ **do**
7: Compute step-length α_k (line search for min. of $\mathcal{J}$ along $\mathbf{p}_k$)
8: $\mathbf{v}_{k+1} = \mathbf{v}_k + \alpha_k \mathbf{p}_k$
9: Adjoint solve to get $\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_{k+1})$
10: Compute β^*_{k+1} (the conjugate gradient parameter)
11: $\mathbf{p}_{k+1} = -\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_{k+1}) + \beta^*_{k+1}\mathbf{p}_k$
12: $k = k + 1$

step length α_k is chosen to satisfy the *strong Wolfe conditions*, which ensure that $\mathbf{p}_k$ is always a direction of descent. Multiple choices also exist for β_k[10] too, with the original Fletcher-Reeves [Fletcher, 1964] form

$$\beta_k = \frac{\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_k)^\intercal \boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_k)}{\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_{k-1})^\intercal \boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_{k-1})}, \tag{4.46}$$

[10] Although appropriate modifications to the line search may be necessary [Nocedal and Wright, 2006] to ensure that a descent direction is chosen for $\mathbf{p}_k$.

and the Polak-Ribiére-Polyak form [Polak and Ribiere, 1969; Polyak, 1969],

$$\beta_k = \frac{\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_k)^\intercal(\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_k) - \boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_{k-1}))^2}{\|\boldsymbol{\nabla}\mathcal{J}(\mathbf{v}_{k-1})\|}, \tag{4.47}$$

or others (e.g. see review by Hager and Zhang [2006]). Occasional restarts of the method (i.e. setting $\beta_l = 0$ for some step l) may also be used to improve convergence by resetting the algorithm, albeit with potential adverse effects on numerical performance, since historical data is erased [Nocedal and Wright, 2006].

4.4.1.1 Verifying the Adjoint Gradients

The gradients obtained using the adjoint are verified by computing the gradient of the cost function (as defined above), $\boldsymbol{\nabla}\mathcal{J}_p$, with respect to a parameter p, and then repeating this gradient calculation by using a finite-difference approximation.

To do this, a small isentropic vortex is placed at $(x, y) = (0, 13)$, above the upper shear layer of the jet defined using the same flow conditions as in the case above. The vortex perturbs the shear layer and the flow is allowed to evolve until $t = 20.0$. Initial conditions for the adjoint variable are calculated, and the adjoint equation is solved backwards in time until $t = 0$. The parameter $p = y_{\text{vort}}$ here is the vertical position of the vortex, and so the gradient of the initial condition $\boldsymbol{\eta}$ with respect to p is calculated. This is multiplied by the adjoint variable to form $\boldsymbol{\nabla}_p\mathcal{J} = \partial\boldsymbol{\eta}/\partial p\ \boldsymbol{\lambda}_0^\intercal$, which is then compared to the finite-difference approximation. The results of this are shown in Figure 4.24. Note how $\mathrm{d}\boldsymbol{\eta}/\mathrm{d}p$ and the adjoint solution at $t = 0$, $\boldsymbol{\lambda}_0$, have overlapping supports. Furthermore, the relative gradient error here is found to be on the order of 0.1%, although as the flow is integrated over longer time horizons, the gradient error does increase.

This may be attributed to two main factors. The first is the inherent error of the method, since it attempts to approximate a chaotic nonlinear trajectory with locally linearised segments, leading to the well-known issues of trajectory divergence for problems with long time horizons. Various techniques have been proposed to address this in chaotic dynamical systems, such as the Least Squares Shadowing

(LSS) method [Wang et al., 2014b], although we leave this for future work. Second, and potentially equally significant in this case, is the error in the linearisation itself. In this implementation, the linearisation is computed through simple finite-difference approximations of the nonlinear right-hand side function, with errors on the order of $E \approx \mathcal{O}(10^{-8})$. As the equations are advanced further in time these errors naturally accumulate, though this may be mitigated by choosing error-free linearisation techniques, such as complex-step derivatives or the method that extends this concept using hyper-dual numbers [Fike and Alonso, 2011]. Though an attractive proposition, this may require significant modifications to the code, and is left as future work here.

4.4.1.2 Calculating Gradients for Large Parameter Spaces

The attractive property of the adjoint method over the finite difference approach is the simultaneous optimisation over high-dimensional parameter spaces with only a single evaluation of the adjoint equation required. We can demonstrate this here by extending the previous example to have three parallel jets, centred at $y = 7.5$, $y = 15.0$ and $y = 22.5$, as shown in Figure 4.25, panel (a). This case is calculated on a 124×30 domain using two, 512×256 point, uniform, rectilinear sub-domains; with the sliding plane located again at $x = 62$ and a sliding plane velocity of $v_{\text{sp}} = 1.0$. The starting flow for the gradient calculation is obtained by integrating for $t = 2040$ to produce the flow in Figure 4.25, panel (a). Around these jets forty isentropic vortices are placed, parameterised by the two-dimensional position vectors (i.e. eighty control parameters), that are located at $\mathbf{p} = (x_p, y_p) = \left(\{-30 + 5i\}_{i=0}^{9}, \left\{15 + (-1)^k 3.75 + (-1)^k 7.5i\right\}_{i=0}^{1}\right)$, with $k \in 0, 1$, as shown in Figure 4.25, panel (b). The perturbed and target flows, $\mathbf{v}$ and $\mathbf{v}_{\text{targ}}$, are integrated for $t = 120.0$, giving the perturbation time to cross into the second flow domain, as evidenced by the perturbation field in Figure 4.25, panel (c). Taking this as the initial gradient, the adjoint equation is integrated backwards in time to

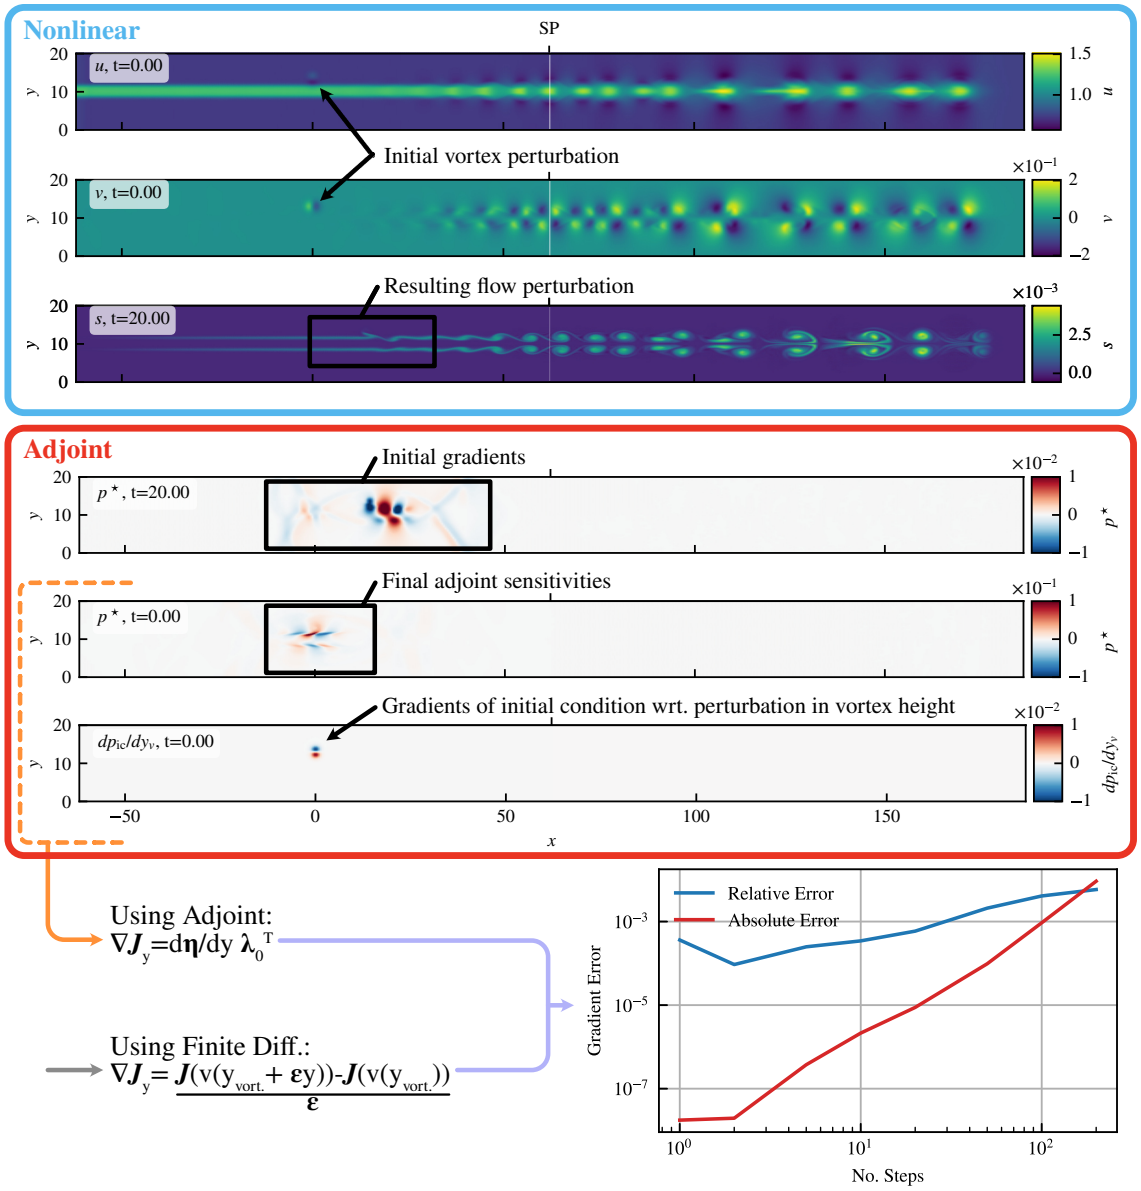


Figure 4.24: Validation of the adjoint routine for computing the gradient. A single isentropic vortex is introduced above the jet shear layer and left to evolve for $t = 20$ units. Initial gradients of $\mathcal{J}$ are set as the initial condition for the adjoint variable, with the adjoint equation integrated to obtain sensitivities. Multiplication by the gradients of the nonlinear initial condition with respect to vortex height gives the gradient of $\mathcal{J}$. This is compared to a finite-difference approximation of the gradient, with the results showing relative errors on the order of 10^{-3}.

$t = 0.0$, with the final adjoint field shown for the pressure component in Figure 4.25, panel (d). This adjoint field is concentrated near $x = -50$, with tilted structures seen within the shear layers for $x \in [-62, 50]$, and crucially, these overlap with the locations of the initial vortex positions. Computing $\mathrm{d}\boldsymbol{\eta}/\mathrm{d}\mathbf{p}$, which is the gradient

of the perturbed initial condition with respect to the position/parameter vector $\mathbf{p}$, and multiplying by the adjoint variable gives the gradients as shown in Figure 4.25, panel (e). Applying this adjoint method, we obtain the gradients of the position vectors in Figure 4.25, panel (e).

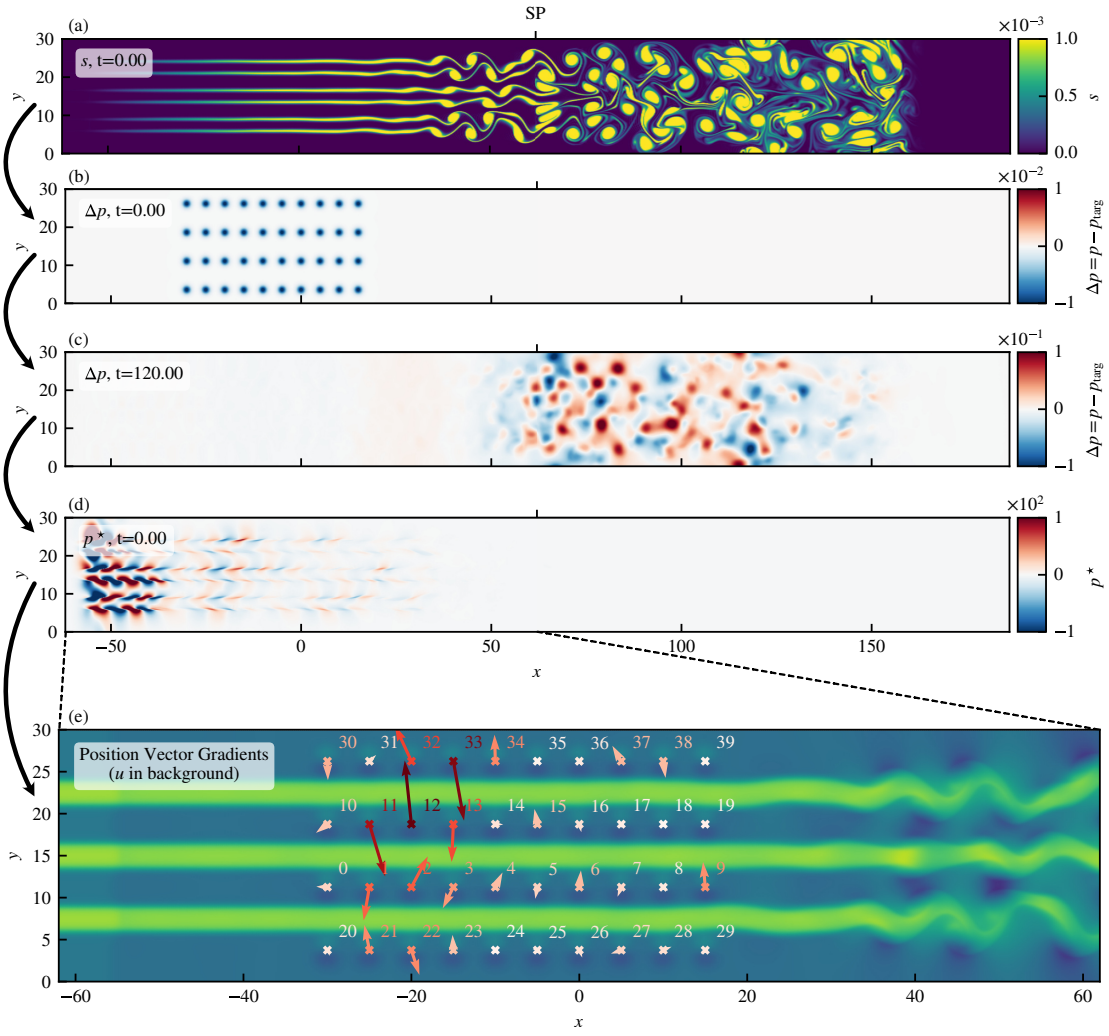


Figure 4.25: The initial field of three unsteady parallel jets, centred at $y = 7.5$, $y = 15.0$ and $y = 22.5$, with the sliding plane at $x = 62$, and $v_{\text{sp}} = 1$, (a). Forty identical vortices are added to this field in between these jets (see details in text), (b). The perturbation induced by these vortices to the pressure field, at $t = 120.0$ is shown in (c), with the adjoint pressure component shown at $t = 0.0$ shown in (d). The gradients, $\boldsymbol{\nabla}_{\mathbf{p}}\mathcal{J} = \partial\boldsymbol{\eta}/\partial\mathbf{p}\ \boldsymbol{\lambda}_0^{\intercal}$, are then calculated for $p =$ position vectors of vortices, and are shown in (e).

4.4.2 Sensitivity Analysis

Although optimisation may be the ultimate goal in many practical applications, a sensitivity analysis can be similarly instructive in problems that are intractable in optimisation settings due to their size, or as a result of computational constraints.

For an example of this, we refer the reader to Chapter 5 of this thesis, where we will consider a simplified rotor-stator interaction problem.

4.5 Conclusion

Throughout this paper we have described and developed a novel lightweight high-order time-domain sliding plane algorithm for studying flow instabilities and aeroacoustics in flows with moving interfaces.

In the first sections we described our code and governing equations, showcased our lightweight linearisation procedure and extended the formulation to meshes in relative motion. This was then augmented by our discussion of the sliding plane communications procedure and how this is easily integrated into our forward and adjoint code. We then derived the discrete nonlinear-adjoint problem, using the standard Lagrangian approach. This methodology was thoroughly validated for both direct and adjoint problems, and the sliding plane was shown to have a minimal impact on the parallel scaling.

This paper concluded with three demonstration cases that showcased the adjoint method. First the calculation of the gradients for an unstable jet problem, where the gradient derived using the adjoint, was validated against the finite difference approach. This showed good agreement to within 10^{-3} of relative error[11]. However, the errors from the linearisation scheme amplified this error, and hence more accurate gradients can be calculated by using a more accurate linearisation.

[11] A further grid dependence study is required to obtain the characteristics of this error as the grid is refined.

Subsequently, this case was extended to three unstable jets to demonstrate the advantages of this approach compared to a finite difference method. The highly sensitive and chaotic flow generated by these jets was perturbed by forty isentropic vortices positioned to disrupt the shear layer. The nonlinear-adjoint calculation with the sliding plane was then used to calculate the gradients for the initial vortex position vectors that maximise the Kelvin-Helmholtz instability.

The final case demonstrated a sensitivity analysis applied to a model rotor-stator interaction problem (see Chapter 5 in the thesis version of this paper).

The chaotic nature of this dynamical system is also known to lead to a trajectory divergence between the nonlinear solution and the tangent model, resulting in a degradation in the quality of the gradient. Attempts can be made to mitigate this by using, for example, Least-Squares Shadowing (LSS) approaches (see, for example Wang et al. [2014b]), however this is generally computationally expensive, and outside the scope of this paper.

Further work should also address the need for time-dependent control and time-dependent cost functions, while integrating this technology into the TAO optimisation framework.

While we have focused on linear sliding motion between neighbouring meshes in a DNS setting here, the approach can be generalised to other codes and setups, including LES and RANS-based methodologies, and co-axial rotation of the domains.

Chapter 5

Rotor-Stator Interactions

5.1 Introduction

This chapter is the extension of the paper presented in Chapter 4 and demonstrates an application of the adjoint sliding plane code to an unsteady, instability-resolving rotor-stator interaction problem for an idealised compressor geometry. Although the paper in Chapter 4 will eventually contain elements of this chapter, we have decided to present it independently here so that relevant connections may be established with various portions of this thesis in a way that is independent to the self-contained nature of the paper.

5.2 Numerical Setup

5.2.1 Domain and Flow Conditions

The numerical domain for this, preliminary, rotor-stator sensitivity analysis consists of a periodic mesh with two stators placed upstream of three rotor blades. While we reuse the surface geometry and flow conditions of the rotor mesh from our earlier work in Chapter 3, the stator blades are 1 : 10 aspect-ratio ellipses[1] placed in a

[1]Thin ellipses are chosen for the stator in order to simplify the mesh geometry. Unlike a realistic compressor stator, the ellipses chosen here do not introduce any complex blade surfaces, which avoids any additional time-step restrictions that might occur with a more complex mesh. As the goal here is to demonstrate the functionality of the nonlinear-adjoint looping, it was decided to retain a simpler case, and to leave the more realistic geometry as future work.

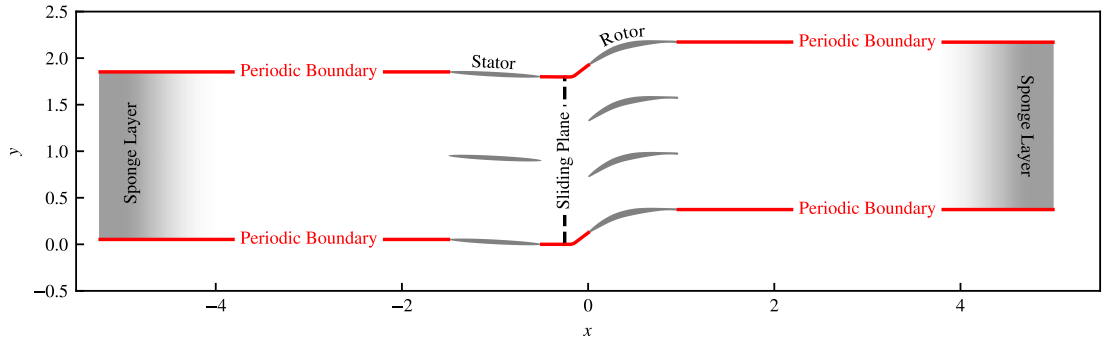


Figure 5.1: Geometry of the rotor-stator demonstration case. Sponge layers are placed at the inlet and outlet to suppress reflections at the boundaries.

uniform axial flow and staggered so that they are at a 3° angle of attack relative to the incoming flow, and resulting in unsteady von Kármán wakes from the stator trailing edges being generated. We assume that the flow turning introduced by these stator blades is insignificant, and that the resulting flow remains predominantly axial before interacting with the rotor blades. The gap between the rotor and stator blade rows is set to $C_{ax}/2$, where C_{ax} is the axial chord of the rotor blade. This geometry is detailed in Figure 5.1.

Flow in the upstream region of the blade row cases in Chapter 3 is used as the basis for setting the initial conditions, with $\mathrm{Re} = 100,000$ and $M = 0.3$. The axial projection of the velocity is used as the axial velocity in the stator mesh, while the vertical component is extracted for use as the mesh velocity, $v_{sp} = -0.791$. The case is then run for 20 time units so that the initial transients are removed from the flow field, before attempting the sensitivity analysis.

5.2.2 Sensitivity Analysis

As with the final sections of the previous chapter, our purpose here is to demonstrate a sensitivity analysis of the rotor-stator problem. To that end we again consider sensitivities to the initial conditions, using the same cost function as before.

To excite all possible frequencies, the endpoint of the converged nonlinear simu-

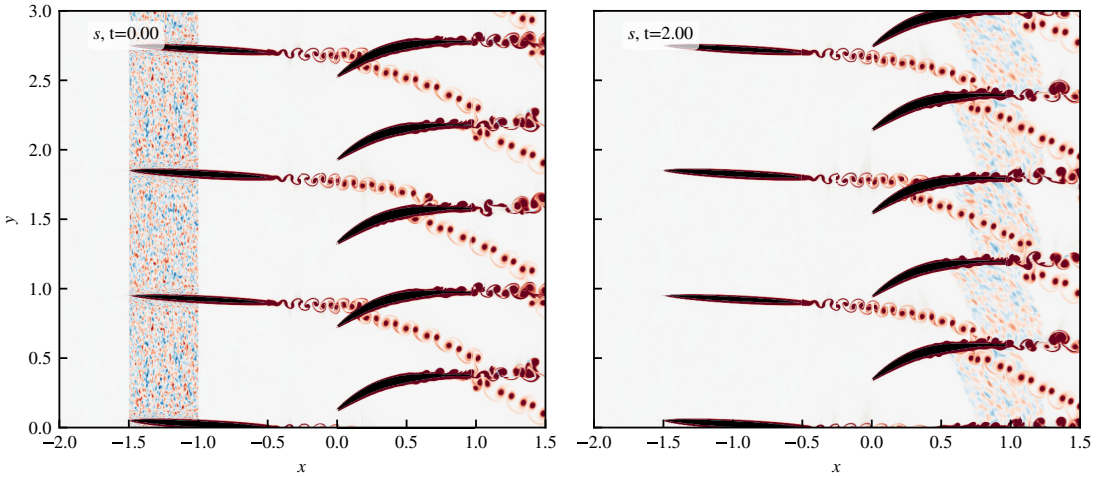


Figure 5.2: The initial randomised perturbation, at $t = 0.0$, illustrated for the entropy component of the state vector (although all field components are perturbed). The same field is illustrated at $t = 2.0$, showing the final position of the disturbance in the rotor before the initial cost function gradient is calculated, for the initial conditions of the adjoint solver.

lation is augmented by a random perturbation. The perturbation is placed in both stator passages for $-1.5 \leq x \leq 1.0$. A Gaussian filter[2] is then applied twice to this field so that the perturbation field is sufficiently smooth, to mitigate undesirable sensitivities from the numerical scheme, and to render the nonlinear flow fields differentiable at all times. In the vicinity of the blade surfaces the perturbation is set to zero[3] so that the non-slip, no-flux boundary condition holds for all time. This initial condition field is shown in the left panel of Figure 5.2.

As in Chapter 4, the cost function is

$$\mathcal{J} = \frac{\langle \mathbf{v} - \mathbf{v}_{\text{targ}}, \mathbf{v} - \mathbf{v}_{\text{targ}} \rangle_{\mathbf{M}}}{2}, \tag{5.1}$$

and the target field, $\mathbf{v}_{\text{targ}}$ is obtained simply by running the unperturbed simulation until $t_1 = 2$. This is performed prior to running the nonlinear-adjoint looping, and

[2]The choice of a Gaussian filter is made here because of a need to remove the very-high-wavenumber detail introduced by the random field. The Gaussian acts as a local averaging of the field values in the vicinity of the chosen point, it blurs the discontinuities in the flow fields without adding extra structure, and is easy to implement using SciPy. Furthermore, the amount of blurring can be controlled by changing the standard deviation of the filter.

[3]This layer consists of the first 30 grid points normal to the blade surface, and has an average thickness of $\Delta y = 7.50 \times 10^{-3}$.

can in fact be performed concurrently with the computation of the starting field, as described above.

Constraints in computation time meant that both the forward and adjoint problems are run with $t_1 = 2$ and a CFL number of 1[4] to reduce computation time. Selecting $t_1 = 2$ limits the range of phenomena that can be investigated, because because a short time scale does not allow the slower feedback mechanisms to become established, however this is still sufficient to make some preliminary observations, and to demonstrate the mechanics of this analysis at a reasonable computational cost.

In the following, the time series for the evolution of the nonlinear and adjoint fields is presented for certain specific features in the respective fields. For clarity, both fields are shown side-by-side at the same time index. This choice is made because the time propagation is reversed in the adjoint, meaning that features in the adjoint correspond to interactions at later times in the nonlinear field. Finding the point at which the adjoint disturbance appears and correlating that with the nonlinear field at that time index suggests the interaction phenomenon that is responsible for the increase in sensitivity.

We should also note that the adjoint field grows exponentially almost everywhere in reverse time, as is expected, and to visualise the adjoint fields better this exponential growth rate is estimated and removed from the adjoint field so that a visualisation of the unsteadiness is possible. An infinity-norm of the adjoint time series is shown in Figure 5.3, the growth of which is approximated using the exponential $10^{-5.208t}$ (note time decreases from $t = 2$ to $t = 0$ in the adjoint). This does mean, however, that certain features are non-resolvable with this colour scale. Where it is appropriate this is adjusted and clearly stated so that observations may

[4]The choice is made to increase the CFL number because the initial conditions considered here are modified snapshots of a previous nonlinear simulation. Therefore, the solution does not experience large velocity transients that might affect its temporal stability. As such, the time step can be increased because the maximum value of U, as described in Equation (2.66) is known more accurately.

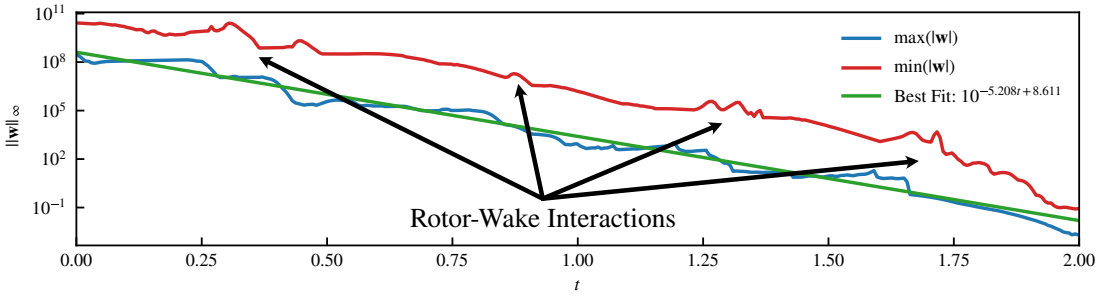


Figure 5.3: Evolution of the infinity norm of the adjoint field $\mathbf{w}$.

be made for all features of interest.

The randomised disturbance applied to all four fields of the state vector breaks down into a 'fast' acoustic field and the slower vorticity and entropy fields with disparate interaction time scales. Although the goal of this initial disturbance is to excite all possible instabilities in the flow, rather than demonstrating a physical phenomenon, we nevertheless keep track of the disturbance so that it can be correlated to the adjoint.

5.3 Nonlinear-Adjoint Dynamics

5.3.1 Overview of the Nonlinear Field

In the configuration and geometry considered here, the stators produce weak wakes from the trailing edges, which interact with the two unsteady laminar separation bubbles on the pressure and suction sides of the ellipse, and subsequently cross the sliding plane and impinge on the rotor. The wake itself is distorted through shearing within the passage, with occasional collisions and vortex breakdown near the rotor leading edge, yielding acoustic pulses. Subsequent interactions occur between the vortices in the wake and the sensitive boundary layers of the rotor, identified earlier in Chapter 3. This interaction is significant because we have shown that the wavemaker region is responsible for both setting the receptivity dynamics of the dominant modes and governing the acoustic coupling between the individual

rotor blades. Disruption of these feedback mechanisms by the wake is therefore expected to yield high sensitivities due to the global nature of the acoustic field. Finally, small modifications to the vortex structures near the trailing edge affect the scattering process that occurs here, so this is also a potential pathway towards high sensitivities.

5.3.2 Initial Gradients

The initial gradient of the cost function is computed by taking the difference between the final time steps of the target, $\mathbf{v}_{\text{targ}}$, and perturbed fields $\mathbf{v}$, with the resulting field shown in Figure 5.4.

First, note that the gradients in the wake of the middle stator (i.e. with leading edge positioned at $y \approx 1.0$) are significantly less impacted than the stator at the periodic boundary (with leading edge at ~ 1.8). It is possible that this occurs due to a receptivity mechanism between the pressure pulse from the rotor leading edge and the separation bubbles on that stator blade.

Effects of the disturbance are also seen on the pressure and suction sides of the rotor, with the gradients attaining similar magnitudes on both surfaces. In some cases, such as on the second rotor blade from the bottom of Figure 5.4, the interaction is not limited to points where the boundary layers interact with the wakes. In this case, convective instabilities are observed upstream of the interaction point suggesting that the feedback mechanism identified earlier has indeed been modified. Evidence of this is observed by pressure gradients in the acoustic field in the middle of the rotor passages, as well as in the rotor wake.

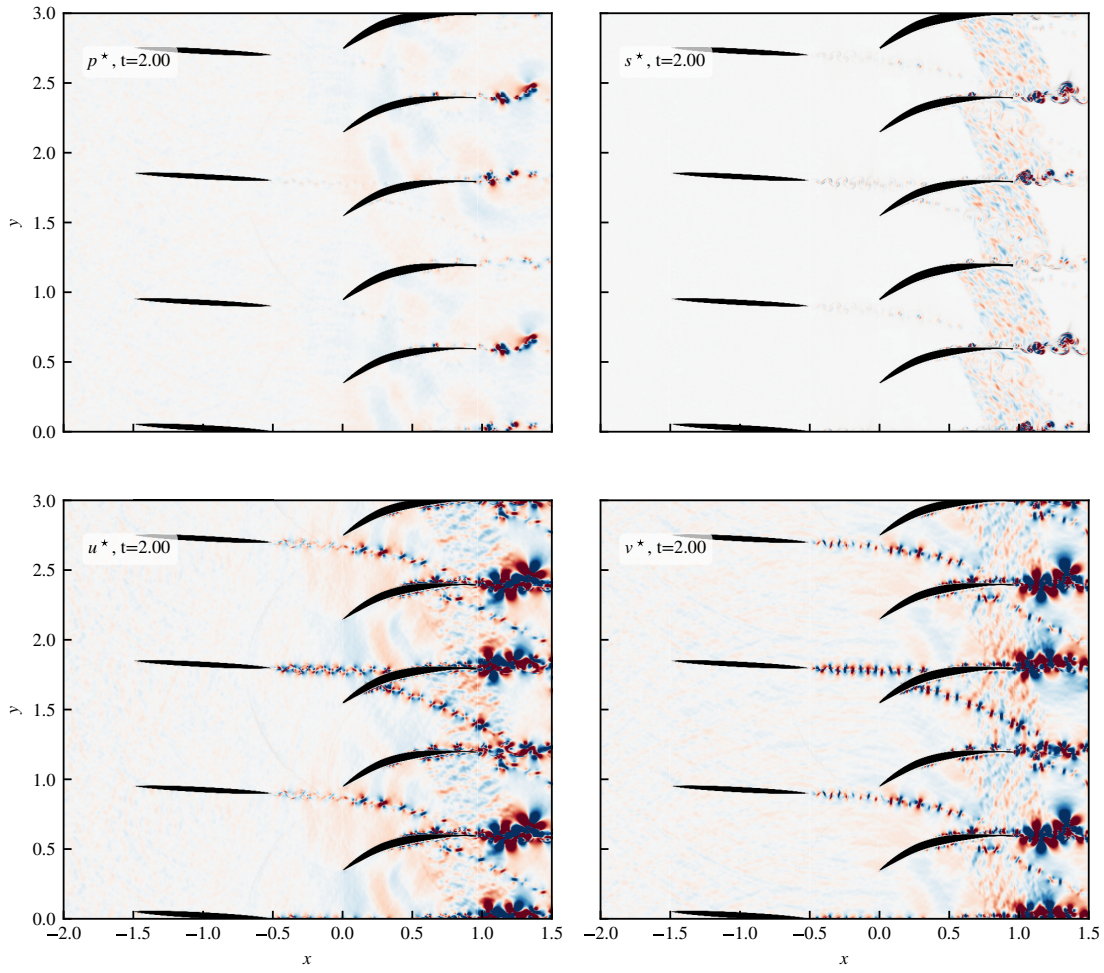


Figure 5.4: Initial gradients (adjoint variables) used as initial conditions for the adjoint calculation.

5.3.3 The Adjoint Sensitivities

5.3.3.1 Wake-Rotor

Figure 5.5 shows a time series of the first[5] jump in the magnitude of the adjoint (see Figure 5.3). The nonlinear field shows that at this point the wake of the stator intersects the leading edge of the rotor, thus forming the acoustic pulse in the nonlinear pressure field. At the time of contact, adjoint convective structures on the pressure and suction sides arrive at the leading edge and produce adjoint pressure waves that propagate into the rest of the domain, which follows from the

[5]Here we refer to the reversed time scale of the adjoint simulation.

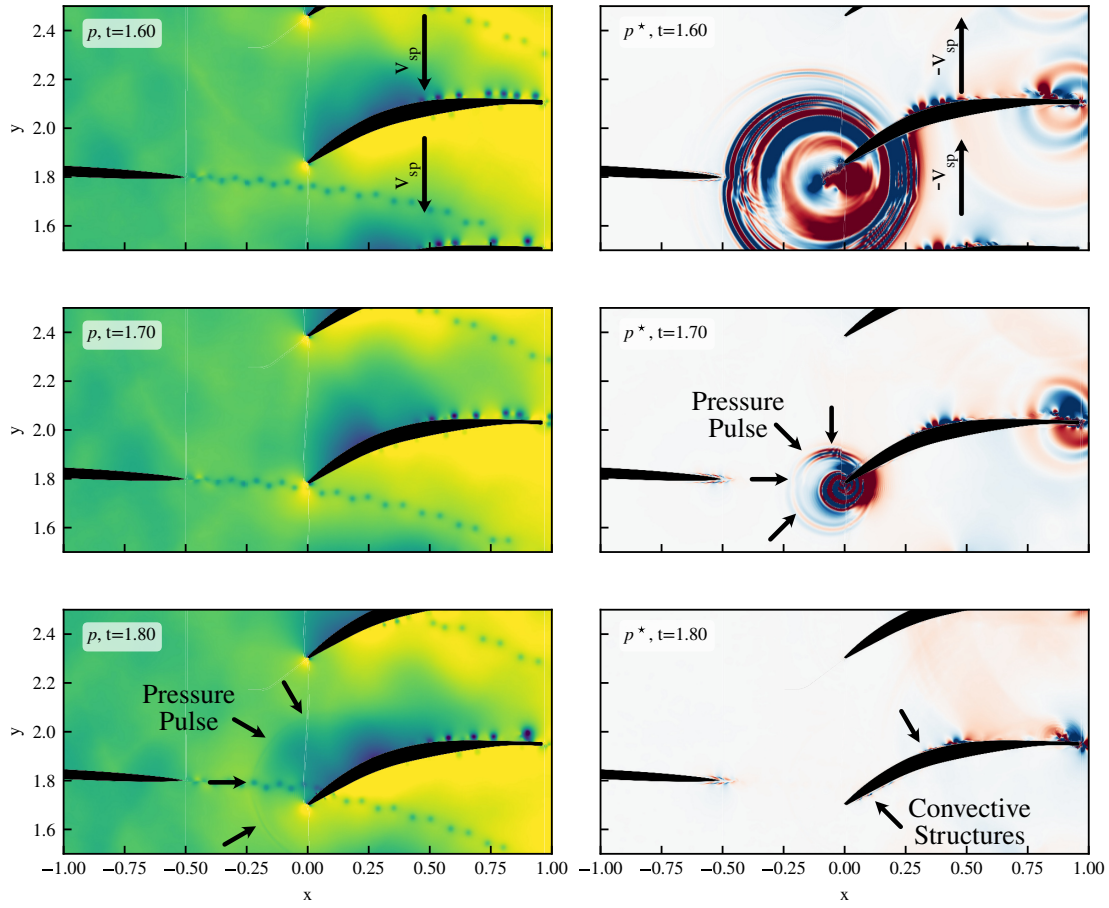


Figure 5.5: The nonlinear and adjoint field time series of the first wake-rotor interaction, showing the pressure pulse from the rotor leading edge and the corresponding adjoint. See the supplementary material.

observation that the wave equation is invariant to the adjoint[6]. Continuing with the adjoint simulation, we note the convective structures propagating upstream from the rotor and intersecting with the portion of the stator wake that in the forward problem intersects with the leading edge of the rotor. In the context of a gradient optimisation problem, it appears that the adjoint here is generating gradients that would deflect the wake as much as possible onto the unstable boundary layers on both the pressure and suction sides. The pressure component of the adjoint is then necessary to maximise the disturbance at the point of interaction. It is possible

[6]It is easy to show from the linearity properties of the adjoint and the derivative operator that the second derivative does not change sign.

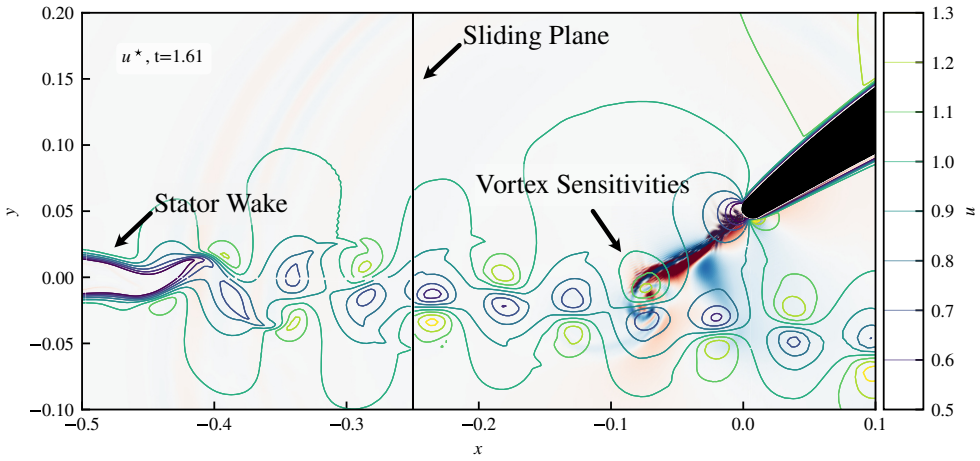


Figure 5.6: Vortex sensitivity before impact of the vortex with the rotor leading edge at $t = 1.61$.

that this occurs due to the disparate receptivity characteristics of the pressure- and suction-side boundary layers, therefore small changes in propagation speed, size or strength of a flow feature potentially affect which side of the stagnation point the disturbance ends up on, and consequently the value of the cost function at the end time varies dramatically. This selection process is best seen by superimposing the entropy contours of the wake at $t = 1.615$, just before the wake vortex impinges upon the leading edge of the rotor blade, and is demonstrated in Figure 5.6.

As instabilities are by definition points of high sensitivity, small modifications to the wake-generation processes at certain, *strategic* times will result in potentially large changes to the wake. Propagating the solution above in adjoint time and plotting the results in Figure 5.7, we see the packet of sensitivity from the rotor-wake interaction impinge on the stator blade and spread along the boundary layer in tilted structures resembling the Orr mechanism. As with the global stability analysis in Chapter 3, these concentrations are located upstream of the laminar separation bubbles on both surfaces, although with a slightly higher adjoint on the pressure-side boundary layer as compared to the suction side. Further upstream, the adjoint decays within the boundary layers, suggesting that processes near the trailing edge dominate.

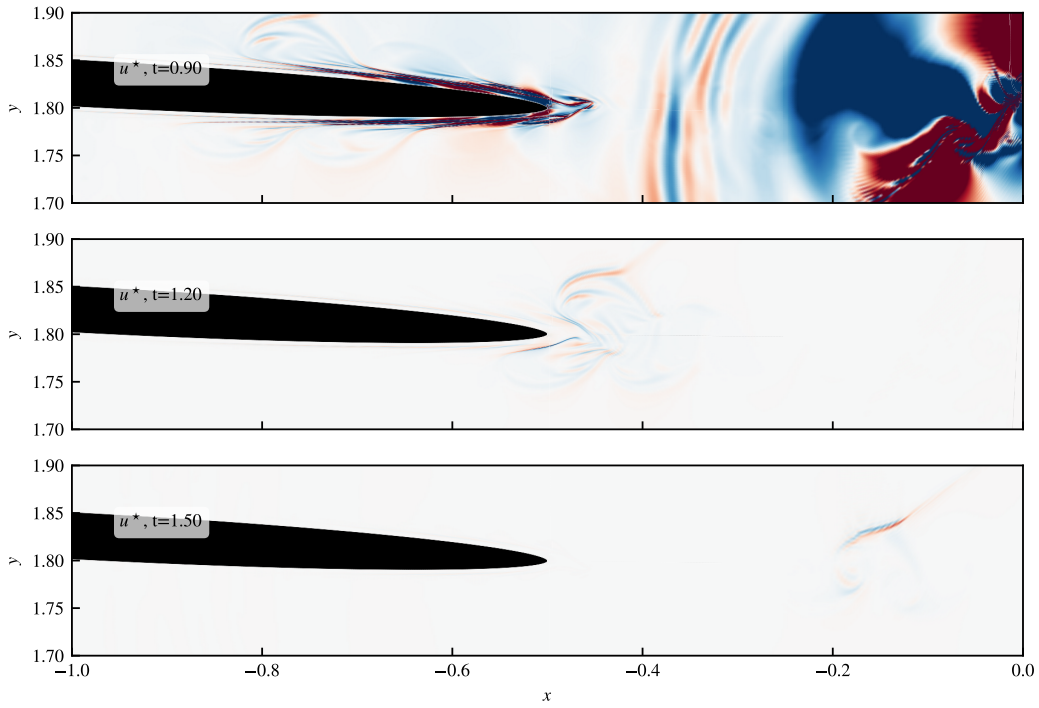


Figure 5.7: Adjoint horizontal velocity field associated with wake features that intersect the leading edge at $t = 1.70$, shown in Figure 5.5. The colourbar limits have been fixed here so that the growth of this sensitivity as it reaches the stator at earlier times may be compared visually.

5.3.3.2 Convective Features and the Laminar Separation Bubble

In our analysis of the receptivity characteristics of the isolated blade row, we multiplied the mode pair **w** and **v** pointwise in order to compute the wavemaker region for a mode at a given frequency. This provided a simple argument for how the acoustic waves influence the hydrodynamic features of the pressure and suction sides in order to set and sustain global instabilities. For the dominant mode M_0, for example, the acoustic and velocity fields overlap with regions where the corresponding adjoint is large to show that, upstream of the separation bubble on the suction side and downstream of the leading edge of the pressure side, the flow is sensitive to disturbances at that frequency. Due to the absence of convective features originating from the inlet, the adjoint velocity field does not have a corresponding 'direct' feature to generate a wavemaker region in that part of the domain.

In the rotor-stator interaction problem, however, the situation changes funda-

mentally. Now the wake from the stator blades creates a new pathway for convective disturbances to affect the laminar separation bubble.

A good example of this can be seen by following the evolution of the rotor blade with the leading edge located at $y \approx 1.12$ at $t = 1.0$, as in Figure 5.8. At $t = 1.65$[7], the suction surface of this blade, and the laminar separation bubble, interact with the wake from the stator (with leading edge at $y_0 \approx 1$). At $t = 1.34$, the leading edge of the rotor intersects the wake, causing pressure impulses, from the leading edge, similar to those identified in Section 5.3.3.1. This is followed by convective structures emanating from the leading edge of the rotor at $t = 1.0$ and $t = 0.9$, as highlighted in Figure 5.8. This continues until $t = 0.70$ at which point the structures from the separation bubble 'arrive' at the leading edge, leading to another pressure pulse at the leading edge, but this time it is not associated with a wake intersection effect. A visual comparison between the advecting adjoint waves in the suction-side boundary layer, and to the time between the adjoint pressure pulses, suggests that this is indeed the mechanism involved.

By contrast, pressure-side features propagate over much longer time scales and the length of this demonstrative simulation is insufficient to capture multiple pressure side interactions across multiple blades so that conclusions may be drawn. Despite this, the nonlinear snapshots in Figure 5.8 do show that the pressure-side boundary layer is unstable when forced by the wake, and due to the relationship of these features with the TE acoustic source and the suction-side dynamics, as seen earlier in Chapter 3, it is expected that the adjoint field would reflect this through sensitivities that alternate between the pressure and suction sides.

In summary, the main differences between the rotor-stator and rotor-alone cases, at least when compared to the adjoint mean flow global modes analysis presented in Chapter 3, are the strong nonlinear interactions at the leading edge and separation bubble in the rotor-stator setting, that results in the coupling of the two blade rows

[7]This is not shown here, but see the animation in the supplementary material.

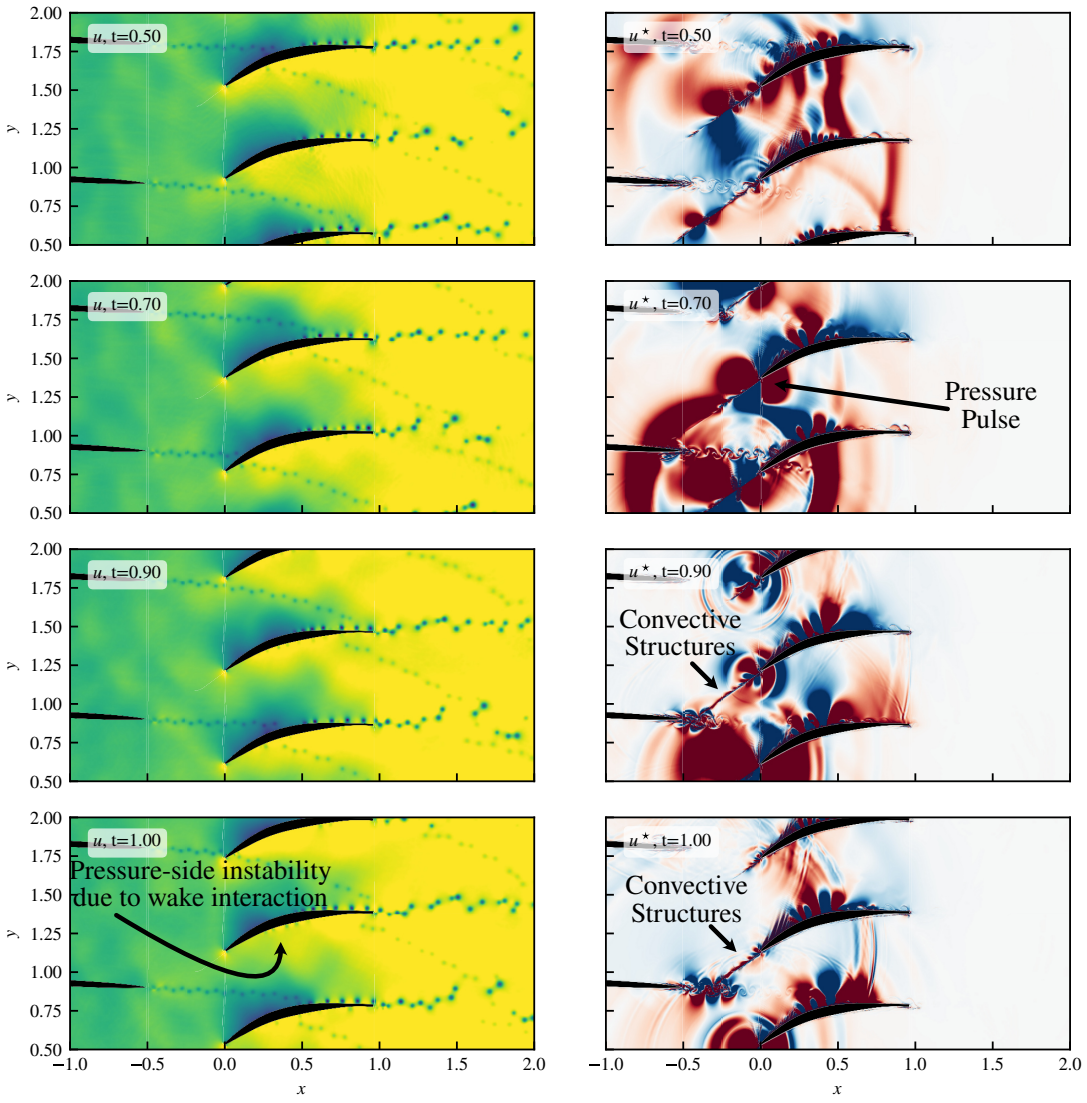


Figure 5.8: Convective adjoint structures emanate from the leading edge of the blade located at $y \approx 1.12$ at $t = 1.0$, followed by a pressure pulse at $t = 0.7$ that is not associated with a wake interaction event. These features are labelled. See the supplementary material.

through acoustic and convective adjoint structures. In particular, the adjoint convective structures, or 'wakes', are seen emanating from the leading edges of the rotor blades to target features within the nonlinear wakes from the stator blades. These are subsequently concentrated at the trailing edge of the stator blades. Although upstream propagating 'wakes' in the $u^\star$ field are also observed in the rotor-alone case, these are not targeted by convective features in of the direct modes, and are hence absent from the wakemaker description. Periodic disturbance of the rotor

pressure-side boundary layers is also observed, with high adjoint sensitivities close to the leading edges. However, these features are not static and propagate all the way to the leading edge, at which point these contribute to the pressure pulse process, as seen in Figure 5.5. Observe also that in Chapter 3, the adjoint fields for the rotor-alone case do not indicate the presence of pressure spikes near the leading edge, and instead are concentrated at the laminar separation bubble on the rotor suction surface, and in the vicinity of the leading edge. Despite these differences, further analysis is needed to quantify the exact variations between the rotor-alone and the rotor-stator interaction problem, by comparing the nonlinear-adjoint solution for an isolated rotor blade.

5.4 Notes on the Implementation

This demonstration case was calculated using the EPCC ARCHER supercomputer, using a total of 492 processors. Of these, 240 processors were assigned to the stator mesh, and 252 for the rotor domain, ensuring a balanced parallel loading. This case is run with 5388 hard-disk checkpoints in a single-level checkpointing scheme, as issues with the PETSc library installation prevented the use of the multi-level checkpointing functionalities, and hence an additional performance penalty was incurred through increased disk-access overheads. For this case, the nonlinear solution took $1:06$ hours, with the total wall-time of the checkpointed solution taking $45:23$ hours, which amounts to the adjoint checkpointing taking $40.26\times$ longer than the nonlinear solution. This is larger than the ~ 32 scaling seen for the test cases in Chapter 4, and can be explained by the lack of multi-level checkpointing for this solution. In total, this solution consumed $22,238$CPUh, not counting the allocation used for testing.

5.5 Conclusion

In this very brief demonstration of the nonlinear adjoint method we have shown the power of this unsteady time-domain nonlinear-adjoint analysis, in combination with an instability-resolving sliding-plane treatment, for the analysis of aeroacoustics problems in turbomachinery applications.

Propagation of the adjoint sensitivities through the sliding mesh was observed with minimal dissipation and distortion, and nonlinear and adjoint structures, arising from the interactions between the stator and rotor, were captured. These features included the adjoint targeting vortices, in the stator wake, that impact the leading edge of the rotor. This situation is thought to be particularly sensitive because allowing more of the vortex to propagate onto the suction side of the rotor would result in greater amplification by the laminar separation bubble situated there. Furthermore, the formation of elongated streamwise streaks at the stator controls the structure of the wake that results in later interactions. Finally, adjoint convective features resulting from the interaction of the wake with the laminar separation bubble were also observed. It is clear, then, that a more thorough and considered analysis is required to unravel the details of these processes. This opens up a clear pathway for future work.

To extend the utility of this method, efforts should focus on improving the scaling ratio between the direct and adjoint solutions, and specifically on the efficiency of the linearisation and the integration of the multi-level checkpointing scheme. At present, one linearisation of the flow field is approximately 17 times more expensive than a nonlinear function evaluation [Fosas de Pando et al., 2012], as the nonlinear right-hand side, $\mathbf{F}$, is evaluated multiple times during the linearisation in Equation (4.25), with an additional evaluation required to reset the linearisation object in the code. By deriving analytic expressions for the discrete adjoint, the performance advantages of the continuous approach could be leveraged in combination with the consistency properties of the discrete form (see Section 4.1). If the nonlinear-to-adjoint cost ratio

can be reduced, then this opens up possibilities for modelling much larger domains and for much longer run-times, which is necessary for the analysis of system responses evolving over long timescales; for example, the convective features on the pressure surfaces discussed above.

Chapter 6

Conclusion

6.1 Summary of Thesis

In this thesis we have sought to address the need for new analysis tools for the study of turbomachinery aeroacoustics, by developing a framework leveraging the power of time-domain direct and adjoint techniques, and applying these to a subset of representative problems to demonstrate the operation and power of this analysis.

This thesis began, in Chapter 2, with summaries of the fifth-order DNC solver, the linearisation component and the direct-adjoint functionality that drive the analysis, with specific emphasis placed on ensuring the efficiency of the linear calculations. Validations were then performed for a two-dimensional, co-rotating vortex pair and a three-dimensional Taylor-Green vortex turbulent dissipation problem, both of which showed good agreement with the literature, thereby validating the implementation of the inviscid transport and viscous dissipation terms, respectively. Comments were also made about the accuracy of the local compact schemes, the linearisation quality, and the validation processes that are used prior to performing any linear analysis in this thesis.

The first application of this analysis framework, in Chapter 3, considered the modal and non-modal global linear stability properties of an idealised compressor blade row at $\mathrm{Re} = 100,000$ and $M = 0.3$, with the blade geometries defined by controlled-diffusion airfoils. In this setting, the flow is characterised by an unstable

suction-side laminar separation bubble, which sheds vortices into the downstream, which later interact with the blunt trailing edges of the blades. In order to separate the local interactions from the global coupling within the blade row, the investigation performed in this chapter was sub-divided into an analysis of the stability properties of a single, periodic passage, followed by an extension of this to N identical passages forming the full blade row.

For the single-blade passage, an initial transient response analysis was performed. This highlighted unstable and amplifying regions in association with the boundary layer and the laminar separation bubble on the suction-side of the blade, as well as convective instabilities propagating downstream along both surfaces. Furthermore, the linear system was also found to be asymptotically unstable, with aeroacoustic feedback processes, linking the acoustic sources at the trailing edge and separation bubble to the unstable shear layers on the two airfoil surfaces, found to underpin the growth of the instability. By combining the acoustic and convective timescales in the context of this feedback mechanism, the asymptotic growth of this disturbance is punctuated by a beating, resulting from the time delay between the receptivity within the boundary layer, the subsequent convection and growth, and the eventual activity of the acoustic source when the disturbance passes over the trailing edge.

By carrying out a global stability analysis, the spectrum of the linear operator was found to be unstable with respect to the mean flow, with two unstable modes and twelve marginally stable modes. Of these, the dominant modes were found to match well with peaks in the Fast Fourier Transform of the nonlinear pressure field, with relative errors between frequencies on the order of 1%, and comparable flow structures seen in the modes and the snapshots of the FFT. Through a further analysis of the eigenvectors for these fourteen modes, three main groups were identified. These are the large-scale Kelvin-Helmholtz modes, driving the dominant dynamics of the system; the high-frequency instabilities associated with the detached shear layer of the suction-side laminar separation bubble; and two degenerate modes, that

arise due to acoustic resonances within the blade passage. This description was then augmented through the introduction of the adjoint modes, after which the wavemakers were calculated. For the most unstable modes, these showed that the peak sensitivities of these modes are located upstream of the separation bubble on the suction surface, and on the pressure surface downstream of the leading edge. Since these two regions are not linked through convective instabilities, this confirms that the feedback process is the result of an aeroacoustic receptivity mechanism that links the acoustic source to the boundary layers at these locations.

In the final part of the chapter, the efficient extension of the single-passage analysis to a multi-passage blade row was achieved by leveraging the properties of block-circulant matrices. By constructing a linear, N-periodic system of identical blades, the flow problem within a ten-periodic array was simplified to ten computationally tractable three-passage calculations, following which the analysis of the single-passage system was replicated for the larger system, extending earlier demonstration work performed in the context of a low-Reynolds number, incompressible flow solver. An impulse response analysis was conducted, and showed that macroscopic disturbances propagate through this cascade by exploiting suction-side receptivity processes on the blade surfaces. Moreover, the scattering of initial acoustic waves by the leading edges of the cascade led to the formation of a persistent transient feature that was observed through the duration of the simulation, and was explained through similar receptivity arguments. A subsequent direct-adjoint global mode analysis demonstrated an increased growth rate for the dominant aeroacoustic modes, with the ten-periodic structure now dominating the linear dynamics. New unstable modes were also observed, and low-frequency eigenvectors, supported by the ten-periodic system, analysed. From a practical perspective, however, it was shown that the two-periodic system does not introduce any additional unstable structures, from which it follows that computational investigations over one- and two-passage domains, which are common approaches taken in the literature, may

not capture the dominant dynamics present in the N-periodic system.

To conclude the discussion, Appendix C presented an exploratory analysis of a two-passage aerodynamic mistuning problem, with one blade perturbed by altering either the pitch height or the stagger angle. Although rather limited in scope, this brief investigation showed that aerodynamic mistuning can both amplify and dampen the aeroacoustic response of the cascade, introduce new peaks in the acoustic spectrum and alter the frequencies of existing peaks. Further comments were also made about how the N-periodic formulation in Chapter 3 can be exploited to obtain eigenvalue sensitivity estimates for small changes to the tuned geometry, although this remains an active area of research at this time.

Future extensions of the work in Chapter 3 could focus on increasing the resolution of the discretisation of the unit circle by taking N to be a highly composite number, so that a high-resolution representation of the symbol curve is obtained. This would help identify the pathways taken by the eigenvalue in systems supporting many different periodicities. Studies of the optimal perturbations and responses should also be conducted through the use of singular-value decompositions.

The second part of this thesis considered the extension of the adjoint framework to an unsteady, time-domain rotor-stator interaction problem. In Chapter 4, the numerical mechanics of a novel adjoint sliding plane treatment were developed, with specific attention devoted to the implementation of a high-order sliding plane, alongside a corresponding, numerically consistent adjoint treatment that is insensitive to the moving interface, the interpolation and the communications processes. Careful verification and validation was performed throughout to ensure that the requisite adjoint duality relations are maintained, both for the domain-decomposed operator and the time-stepping problem. For the non-autonomous problem, checkpointing was introduced by wrapping the relevant portions of the code into the PETSc framework, which handles the management of checkpoints and the looping of the nonlinear and adjoint solutions. Calculations of the gradients, with respect to the parame-

terised initial conditions, were validated against finite-difference calculations, and a discussion was provided of how the adjoint method could be used for optimisation problems conducted over large parameter spaces. As an example of this, the chapter concluded with gradients calculated for an initial perturbation problem applied to three unstable parallel jets.

Concluding the thesis, Chapter 5 applied the unsteady nonlinear-adjoint framework, of Chapter 4, to a sensitivity analysis of a rotor-stator interaction problem. Simulations were performed for two elliptical stator blades and three rotor blades, adapted from the compressor blade geometry in Chapter 3. Although further analysis is required, including in more realistic flow geometries, preliminary results indicated the suitability of the instability-resolving sliding-plane treatment for rotor-stator interaction problems, with minimal dissipation and distortion of the unsteady flow features in both the nonlinear and adjoint flow fields. Interactions between the rotors and stators were reflected in the adjoint field, with adjoint features targeting the regions of the stator wake that impinge upon the leading edges of the rotor blades. This is believed to arise due to the differing sensitivity characteristics of the pressure and suction surfaces, with the adjoint targeting the vortex so that it is displaced onto the more sensitive suction side. Furthermore, the formation of elongated streamwise streaks at the trailing edges of the stator blades control the wake structures, that subsequently interact with the rotors and the laminar separation bubbles on the rotor suction surface. For a clearer understanding of these phenomena, a more thorough analysis is necessary to dissect the flow field, and to unravel the details of these processes, which opens a clear pathway for future work.

6.2 Suggestions for Future Work

Work performed in this thesis has shown that direct-adjoint analyses present a rich and viable framework for the study of turbomachinery aeroacoustics and other highly sensitive and unsteady flows. To that end, there are three possible future directions

for the extension of this framework. The following considers a critical overview of the applications enabled by the tools developed, potential improvements to the performance of the numerical methods, and the development of additional tools that may be integrated with the foundations that are laid out in the preceding chapters.

6.2.1 Analysis Enabled by the Framework

Throughout this thesis, extensive references have been made to the extensions of the existing analysis, details of which may be found throughout the summary conclusions for each chapter, in the work above. In this section, we will instead focus briefly on the larger perspective, and summarise the alternative analyses that are now possible with the application of this framework.

Having established the direct-adjoint framework developed in this code, and demonstrated its applications in an idealised context, it is natural to consider extensions of this analysis to three-dimensional contexts. Although it may still be computationally intractable to consider fully three-dimensional blade geometries, a span-wise extrusion of the two-dimensional geometry by 64 (or 128) points, amounting to approximately 6.5% (or 13%) of the airfoil chord, should provide some indication of the change of the sensitivities derived in Chapter 3, and the robustness of the observed structures. Estimated costs for this kind of analysis are nominally $64 - 128\times$ higher than the comparable two-dimensional simulation, with the single-passage nonlinear analysis expected to consume $\approx 81,280$CPUh for the nonlinear solution, with the direct eigenvalue problem expected to consume $356,864$CPUh, based on the cost data given in Chapter 3, although not accounting for parallel scaling losses expected to be in the region of $10 - 20\%$, assuming a perfectly balanced loading on all processors. Adding more processors, lowering the Reynolds number to $\mathrm{Re} \leq 50,000$, reducing the resolution of the grid, and using smaller Krylov subspaces in the calculation of the direct and adjoint eigenvalues, is expected to reduce these costs to some degree, and address the limitations of performing this

analysis in three-dimensional settings. Although the reduced Reynolds number is likely to make the flow more two-dimensional and less representative of the fully three-dimensional airfoil at $\mathrm{Re} \geq 100,000$, increasing the degrees of freedom in the spanwise direction would nevertheless allow the two-dimensional structures to break down, and for structures with higher spanwise wavenumbers to develop. By doing this at a lower Reynolds number, the outputs of the framework, when applied in a three-dimensional setting, can be evaluated at a significantly reduced computational cost. This is something that is a limiting factor particularly for the modal analysis. When this is done, further consideration can be given to the problem at higher Reynolds numbers, once baseline observations have been established.

It was shown in Chapter 4 how the nonlinear-adjoint algorithms can be used to generate gradients of cost functions with respect to various parameters that define the problem. From this, the gradient calculation can be integrated directly into an optimisation framework, making use of a gradient descent algorithm. Recent distributions of PETSc include the TAO optimisation framework, which enables tight integration between the PETSc wrappers of the code and high-performance compiled optimisation libraries, and hence transferring the existing code into this context is a relatively simple matter. By considering optimisation problems, with appropriately sensitive cost functions, one could apply this methodology for optimising the initial conditions, to identify the optimal transient response over a finite time horizon, so that the regions of sensitivity and the mechanisms, identified in Chapter 5, are pinpointed with greater precision.

The code developed here can also be recycled and applied to an optimal control problem, with relatively little additional effort. Control vectors could include shape parameters governing the blade geometry, or the injection of mass flow at the blade surface. Insights provided by this analysis could be used to understand the flow feedback mechanisms in the unsteady flow cases, and how these interactions can be disrupted in an optimal way, according to some predefined measure.

To conclude, it should also be mentioned that the scalability of the parallel code does admit the possibility of extending this rotor-stator analysis to multiple stages, although with a commensurate increase in the costs of the communication and IO losses for saving the larger datasets.

6.2.2 Solver Performance

Code and solver performance still presents perhaps the biggest challenge for the adoption of this methodology in wider contexts. Before we continue with this discussion, we should clarify that we will not refer to the performance of the underlying Navier-Stokes code because it is taken as given that an improved implementation, that is able to model turbulence and resolve ever-finer details of the flow field, will inevitably lead to results that are more accurate and physically relevant in comparison to our use of the two-dimensional implementation in this thesis. Instead, we are concerned with the mechanics and limitations of the algorithms that handle the analysis, and which have been the subject and core focus of this thesis.

Perhaps the biggest limitation in this body of work concerns the performance of the linearisation. In situations where the linear system matrix is autonomous, such as in Chapter 3, this is of relatively little concern because the linearisation only needs to be performed once, so the computational cost of this is absorbed into the pre-processing step. On the other hand, Chapter 4 considered a non-autonomous scenario, requiring the linearisation of the nonlinear flow at each timestep of the solution. From Section 2.7, the pictorial representation in Figure 4.9 and the definition of the sparse sub-matrices $\mathbf{A}_{i,j,k}$ in Equation (2.58), the nonlinear right-hand side, $\mathbf{F}(\mathbf{v}, \mathbf{v}_k)$ is effectively evaluated multiple times when computing the linearisation. Furthermore, the right-hand side function $\mathbf{F}(\mathbf{v}_{\mathrm{bf}}, \mathbf{v}_{\mathrm{bf},k})$ is evaluated at each step, and the states of the internal system objects saved, so that the base flow state for all objects can be restored after each perturbation is applied. Performance is lost here not only because of the additional function evaluations, but also due to memory

bandwidth issues encountered with the reset procedure. Current configurations of the code make this part difficult to optimise, and, even if some optimisation is possible, it is unlikely to yield substantial improvements in performance that would bring the computational costs of the tangent-linear model and adjoint simulations to parity, or at least reduce it to the same order of magnitude, as with the nonlinear problem.

Continuous adjoint approaches are of course a possible solution to this problem, since these would replace the direct and adjoint matrices and codes with analytically-derived expressions, though, given the extensive discussion in Chapter 1, we consider this approach as undesirable. Instead, a possible solution is to derive analytic expressions for the $\mathbf{A}_{i,j,k}$ matrices directly from the discretised code, and in so doing, considerable performance gains could be realised by the removal of the linearisation and memory management issues. By focusing on the $\mathbf{A}_{i,j,k}$'s instead of deriving the full linearisation of $\mathbf{F}$, all that is required is an expansion of the existing expressions with the aid of a chain rule. The derivatives and inter-block communications, such as periodicity, are external to the sub-functions and so are accounted for automatically in the terms denoted by $\mathbf{C}$ and $\mathbf{D}_k$.

On the other hand, derivations performed manually are prone to errors, especially when considering the more complex terms arising in the boundary conditions. These would naturally require lengthy code rewrites when modifying the grid topology, for example. Code reuse is also important to consider here, since each derivative term of the input vector requires its own expression for the linearised equations with respect to that variable, the codebase rapidly grows in size and becomes virtually unmaintainable.

Despite this, if these problems can be solved and the user desires performant nonlinear-adjoint looping, replacing the linearisation could be considered a worthwhile exercise, as computational costs will be reduced by at least one order of magnitude – significantly expanding the range of potential applications.

Trajectory divergence is a well-known issue for tangent-linear approximations of chaotic systems, and leads to the degradation of the gradient approximation, and an exponential growth in sensitivities. Various strategies have been proposed to address this, such as Least Squares Shadowing (LSS), which constructs complementary tangent and adjoint trajectories that are bounded in time and do not diverge exponentially in time. This assumes that the system is *ergodic*, in other words that the long-time averaged behaviour of the system is invariant with almost all initial conditions. The evolution of the tangent trajectory is then defined through the Karush-Kuhn-Tucker (KKT) equations, which are solved, and sensitivities for long-time averaged cost functions obtained. The costs of solving this system can be significant, however, with costs multiple orders of magnitude higher for the LSS approach when compared to the conventional CFD analysis. We defer further commentary and instead refer the reader to [Blonigan et al., 2018] for an application of LSS to the example of a chaotic flow around a two-dimensional airfoil.

Finally, improvements to the speed of the adjoint components of this code would enable gradients, derived from the nonlinear-adjoint procedures to be utilised in gradient optimisation frameworks. Tight integration between PETSc and TAO could be exploited to facilitate these efforts, especially for parametric sensitivity analyses, as mentioned above.

6.2.3 Developing Additional Tools

By basing this framework on the theory of linear systems, a rich collection of techniques and methods is readily available to extend the analyses carried out in this thesis. While the global modes in Chapter 3 give an indication of the dynamics at the final time, the non-normal linear operator produces rich dynamics over a finite time horizon that may be substantially different to those observed in the direct modes. Optimal forcing and response structures, and frequency response analyses of the harmonically-driven system are of particular scientific interest here,

and especially so in the N-periodic system in Section 3.4, which may exhibit both localised responses and dynamics that encompass multiple blade passages, which are yet to be categorised.

Obtaining the optimal forcing and response structures is achieved using the SVD, with SLEPc providing matrix-free iterative methods similar to the eigenvalue solutions described earlier. Special care must be taken here, however, since each iteration must run until the end of the desired time horizon, and this may be prohibitively expensive for very large systems and long time horizons. Possible solutions include reducing the number of basis vectors for the Krylov subspace so that the iterations proceed faster, and reducing the time horizons over which the optimal response is sought.

Harmonically driven systems take the form

$$\frac{\mathrm{d}\mathbf{v}}{\mathrm{d}t} = \mathbf{A}\mathbf{v} + \mathbf{f}, \tag{6.1}$$

with $\mathbf{f} = \hat{\mathbf{f}} \exp\{i\omega t\}$. The resolvent operator is then defined as

$$R(\omega) = \max_{\hat{\mathbf{f}}} \frac{\|\hat{\mathbf{v}}\|}{\|\hat{\mathbf{f}}\|} = \|(i\omega\mathbf{I} - \mathbf{A})^{-1}\|. \tag{6.2}$$

Since the matrix $(i\omega\mathbf{I} - \mathbf{A})$ is dense, direct inversions of this operator are impossible and so alternative techniques, such as randomised resolvent analysis should be used in this case, especially since the resolvent must be evaluated with a sufficiently dense subset of $\omega \in \mathbb{R}$ over some interval of interest.

In many situations, such as in the case of oscillating separation bubbles, periodic wake impingement or in a general unsteady rotor-stator interaction problem, all of which have been seen here, the evolution of the system can be considered to have a quasi-periodic time dependence. Studies of these systems using Floquet analyses, with the aid of the sliding plane formalism applied to rotor-stator interaction problems as explored in Chapter 5, aim to identify periodic instabilities by computing the associated Floquet multipliers and modes of the monodromy operator. The quasi-periodicity of the flow potentially poses additional challenges as the macroscopic

flow structures may not be fully periodic, and so some form of additional filtering may be required to ensure periodicity.

Finally, having demonstrated the utility of the N-periodic decomposition in Section 3.4, it is natural to wonder whether it is possible to extend this to a rotor-stator problem of the kind seen in Chapter 5. As a result of non-matching blade counts between the rotor and stator, the flow within each passage is no longer the same, it is unsteady and it is non-periodic because of superimposed chaotic unsteadiness that may not return identical flow states at time intervals equal to the rotation period at the sliding plane. In the latter case this lack of periodicity could be addressed by filtering out all harmonics not equal to multiples of the blade-passing frequency, thereby defining a quasi-periodic state that is representative of a blade row interaction base flow. Following these modifications the block-tri-diagonal matrix is reconstructed but it now takes a twisted-Toeplitz form for both the rotor and the stator, with these matrices having many interesting properties, such as defining a collection of pseudomodes that can have both global and local support (as previously shown by Fosas de Pando and Schmid [2019]). As the coefficients of the matrices are periodic in time, the flow problem is solved using the methods of Floquet analysis, although the efficient implementation of this is a subject for future research.

Bibliography

S. Abhyankar, J. Brown, E. M. Constantinescu, D. Ghosh, B. F. Smith, and H. Zhang. Petsc/ts: A modern scalable ode/dae solver library. *arXiv preprint arXiv:1806.01437*, 2018.

N. A. Adams and K. Shariff. A high-resolution hybrid compact-ENO scheme for shock-turbulence interaction problems. *Journal of Computational Physics*, 127: 27–51, Aug. 1996. ISSN 00219991. doi: 10.1006/jcph.1996.0156.

R. Amiet. Acoustic radiation from an airfoil in a turbulent stream. *Journal of Sound and Vibration*, 41, Aug. 1975. ISSN 0022460X. doi: 10.1016/S0022-460X(75)80105-2.

R. Amiet. Noise due to turbulent flow past a trailing edge. *Journal of Sound and Vibration*, 47, Aug. 1976. ISSN 0022460X. doi: 10.1016/0022-460X(76)90948-2.

E. Anderson, Z. Bai, C. Bischof, S. Blackford, J. Demmel, J. Dongarra, J. D. Croz, A. Greenbaum, S. Hammarling, A. McKenney, and D. Sorensen. *LAPACK Users' Guide*, volume 1. Society for Industrial and Applied Mathematics, 3 edition, 1999.

P. W. Anderson. Absence of diffusion in certain random lattices. *Physical Review*, 109, Mar. 1958. ISSN 0031-899X. doi: 10.1103/PhysRev.109.1492.

E. Åkervik, L. Brandt, D. S. Henningson, J. Hœpffner, O. Marxen, and P. Schlatter. Steady solutions of the navier-stokes equations by selective frequency damping. *Physics of Fluids*, 18, June 2006. ISSN 10706631. doi: 10.1063/1.2211705.

F. Bake, C. Richter, B. Mühlbauer, N. Kings, I. Röhle, F. Thiele, and B. Noll. The entropy wave generator (EWG): A reference case on entropy noise. *Journal of Sound and Vibration*, 326, Oct. 2009. ISSN 0022460X. doi: 10.1016/j.jsv.2009.05.018.

S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, pages 163–202. Birkhäuser Press, 1997.

S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc Web page. `https://www.mcs.anl.gov/petsc`, 2019. URL `https://www.mcs.anl.gov/petsc`.

S. Balay, S. Abhyankar, M. F. Adams, J. Brown, P. Brune, K. Buschelman, L. Dalcin, A. Dener, V. Eijkhout, W. D. Gropp, D. Karpeyev, D. Kaushik, M. G. Knepley, D. A. May, L. C. McInnes, R. T. Mills, T. Munson, K. Rupp, P. Sanan, B. F. Smith, S. Zampini, H. Zhang, and H. Zhang. PETSc users manual. Technical Report ANL-95/11 - Revision 3.13, Argonne National Laboratory, 2020. URL `https://www.mcs.anl.gov/petsc`.

D. Barkley. Linear analysis of the cylinder wake mean flow. *Europhysics Letters (EPL)*, 75, Sept. 2006. ISSN 0295-5075. doi: 10.1209/epl/i2006-10168-7.

S. Beneddine, D. Sipp, A. Arnault, J. Dandois, and L. Lesshafft. Conditions for validity of mean flow stability analysis. *Journal of Fluid Mechanics*, 798, July 2016. ISSN 0022-1120. doi: 10.1017/jfm.2016.331.

R. Bhaskaran and S. K. Lele. Large eddy simulation of free-stream turbulence effects on heat transfer to a high-pressure turbine cascade. *Journal of Turbulence*, 11:N6, Jan. 2010. ISSN 1468-5248. doi: 10.1080/14685241003705756. URL `http://www.tandfonline.com/doi/abs/10.1080/14685241003705756`.

L. S. Blackford, J. Choi, A. Cleary, E. D'Azevedo, J. Demmel, I. Dhillon, J. Dongarra, S. Hammarling, G. Henry, A. Petitet, K. Stanley, D. Walker, and R. C. Whaley. *ScaLAPACK Users' Guide*, volume 1. Society for Industrial and Applied Mathematics, 1 edition, 1997.

F. Bloch. über die quantenmechanik der elektronen in kristallgittern. *Zeitschrift für Physik*, 52, July 1929. ISSN 1434-6001. doi: 10.1007/BF01339455.

P. J. Blonigan, Q. Wang, E. J. Nielsen, and B. Diskin. Least-squares shadowing sensitivity analysis of chaotic flow around a two-dimensional airfoil. *AIAA Journal*, 56, Feb. 2018. ISSN 0001-1452. doi: 10.2514/1.J055389.

D. J. Bodony. Analysis of sponge zones for computational fluid mechanics. *Journal of Computational Physics*, 212, Mar. 2006. ISSN 00219991. doi: 10.1016/j.jcp.2005.07.014.

A. Broatch, J. Galindo, R. Navarro, J. García-Tíscar, A. Daglish, and R. K. Sharma. Simulations and measurements of automotive turbocharger compressor whoosh noise. *Engineering Applications of Computational Fluid Mechanics*, 9:12–20, Jan. 2015. ISSN 1994-2060. doi: 10.1080/19942060.2015.1004788. URL `http://www.tandfonline.com/doi/full/10.1080/19942060.2015.1004788`.

G. Burgreen and O. Baysal. Three-dimensional aerodynamic shape optimization of wings using sensitivity analysis. American Institute of Aeronautics and Astronautics, Jan. 1994. doi: 10.2514/6.1994-94.

M. P. Castanier and C. Pierre. Using intentional mistuning in the design of turbomachinery rotors. *AIAA Journal*, 40, Oct. 2002. ISSN 0001-1452. doi: 10.2514/2.1542.

M. P. Castanier and C. Pierre. Modeling and analysis of mistuned bladed disk vibration: Current status and emerging directions. *Journal of Propulsion and Power*, 22, Mar. 2006. ISSN 0748-4658. doi: 10.2514/1.16345.

H. Choi and P. Moin. Grid-point requirements for large eddy simulation: Chapman's estimates revisited. *Physics of Fluids*, 24, Jan. 2012. ISSN 1070-6631. doi: 10.1063/1.3676783.

J. Chomaz. Global instabilities in spatially developing flows: Non-normality and nonlinearity. *Annual Review of Fluid Mechanics*, 37:357–392, Jan. 2005. ISSN 0066-4189. doi: 10.1146/annurev.fluid.37.061903.175810. URL `http://www.annualreviews.org/doi/10.1146/annurev.fluid.37.061903.175810`.

B. T. Chu. On the energy transfer to small disturbances in fluid flow (Part I). *Acta Mechanica*, 1:215–234, Sept. 1965. ISSN 00015970. doi: 10.1007/BF01387235. URL `https://link.springer.com/article/10.1007/BF01387235`.

A. Collette. *Python and HDF5*. O'Reilly Media, Inc., 2013. ISBN 9781449367831.

T. Colonius and S. K. Lele. Computational aeroacoustics: Progress on nonlinear problems of sound generation. *Progress in Aerospace Sciences*, 40:345–416, Aug. 2004. ISSN 03760421. doi: 10.1016/j.paerosci.2004.09.001.

E. F. Crawley and K. C. Hall. Optimization and mechanisms of mistuning in cascades. *Journal of Engineering for Gas Turbines and Power*, 107, Apr. 1985. ISSN 0742-4795. doi: 10.1115/1.3239742.

J. Crouch, A. Garbaruk, and D. Magidov. Predicting the onset of flow unsteadiness based on global instability. *Journal of Computational Physics*, 224, May 2007. ISSN 00219991. doi: 10.1016/j.jcp.2006.10.035.

N. Cumpsty. Jet engine combustion noise: Pressure, entropy and vorticity perturbations produced by unsteady combustion or heat addition. *Journal of Sound and Vibration*, 66, Oct. 1979. ISSN 0022460X. doi: 10.1016/0022-460X(79)90697-7.

N. A. Cumpsty and F. E. Marble. The interaction of entropy fluctuations with turbine blade rows; a mechanism of turbojet engine noise. *Proceedings of the Royal Society of London. A. Mathematical and Physical Sciences*, 357, Nov. 1977. ISSN 0080-4630. doi: 10.1098/rspa.1977.0171.

N. Curle. The influence of solid boundaries upon aerodynamic sound. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 231, Sept. 1955. ISSN 0080-4630. doi: 10.1098/rspa.1955.0191.

L. Dalcín, R. Paz, and M. Storti. Mpi for python. *Journal of Parallel and Distributed Computing*, 65, Sept. 2005. ISSN 07437315. doi: 10.1016/j.jpdc.2005.03.010.

L. Dalcín, R. Paz, M. Storti, and J. D'Elía. Mpi for python: Performance improvements and mpi-2 extensions. *Journal of Parallel and Distributed Computing*, 68, May 2008. ISSN 07437315. doi: 10.1016/j.jpdc.2007.09.005.

L. D. Dalcín, R. R. Paz, P. A. Kler, and A. Cosimo. Parallel distributed computing using python. *Advances in Water Resources*, 34, Sept. 2011. ISSN 03091708. doi: 10.1016/j.advwatres.2011.04.013.

J. DeBonis. Solutions of the Taylor-Green vortex problem using high-resolution explicit finite difference methods. American Institute of Aeronautics and Astronautics, Jan. 2013. ISBN 978-1-62410-181-6. doi: 10.2514/6.2013-382.

J. W. Delfs. An overlapped grid technique for high resolution CAA schemes for complex geometries. 2001. doi: 10.2514/6.2001-2199.

J. D. Denton. The calculation of three-dimensional viscous flow through multistage turbomachines. *Journal of Turbomachinery*, 114:18–26, Jan. 1992. ISSN 15288900. doi: 10.1115/1.2927983.

G. Desquesnes, M. Terracol, and P. Sagaut. Numerical investigation of the tone noise mechanism over laminar airfoils. *Journal of Fluid Mechanics*, 591, Nov. 2007. ISSN 0022-1120. doi: 10.1017/S0022112007007896.

A. P. Dowling and Y. Mahmoudi. Combustion noise. *Proceedings of the Combustion Institute*, 35, 2015. ISSN 15407489. doi: 10.1016/j.proci.2014.08.016.

I. Duran, S. Moreau, F. Nicoud, T. Livebardon, E. Bouty, and T. Poinsot. Combustion noise in modern aero-engines. *AerospaceLab*, pages 1–11, June 2014.

M. C. Duta, M. B. Giles, and M. S. Campobasso. The harmonic adjoint approach to unsteady turbomachinery design. *International Journal for Numerical Methods in Fluids*, 40, Sept. 2002. ISSN 0271-2091. doi: 10.1002/fld.289.

W. Edwards, L. Tuckerman, R. Friesner, and D. Sorensen. Krylov methods for the incompressible navier-stokes equations. *Journal of Computational Physics*, 110, Jan. 1994. ISSN 00219991. doi: 10.1006/jcph.1994.1007.

K. Ekici and K. C. Hall. Nonlinear analysis of unsteady flows in multistage turbomachines using harmonic balance. *AIAA Journal*, 45, May 2007. ISSN 0001-1452. doi: 10.2514/1.22888.

K. Ekici, R. E. Kielb, and K. C. Hall. Aerodynamic asymmetry analysis of unsteady flows in turbomachinery. *Journal of Turbomachinery*, 132, Jan. 2010. ISSN 0889-504X. doi: 10.1115/1.3103922.

J. I. Erdos, E. Alzner, and W. McNally. Numerical solution of periodic transonic flow through a fan stage. *AIAA Journal*, 15, Nov. 1977. ISSN 0001-1452. doi: 10.2514/3.60823.

European Commission. Flightpath 2050 Europe's vision for aviation. Technical report, 2011.

J. H. Ferziger and M. Perić. *Computational Methods for Fluid Dynamics*. Springer Berlin Heidelberg, 3 edition, 2002. ISBN 978-3-540-42074-3. doi: 10.1007/978-3-642-56026-2.

J. E. Ffowcs Williams and D. L. Hawkings. Sound generation by turbulence and surfaces in arbitrary motion. *Philosophical Transactions of the Royal Society of London. Series A, Mathematical and Physical Sciences*, 264, May 1969a. ISSN 0080-4614. doi: 10.1098/rsta.1969.0031.

J. E. Ffowcs Williams and D. L. Hawkings. Theory relating to the noise of rotating machinery. *Journal of Sound and Vibration*, 10, July 1969b. ISSN 0022460X. doi: 10.1016/0022-460X(69)90125-4.

J. Fike and J. Alonso. The development of hyper-dual numbers for exact second-derivative calculations. American Institute of Aeronautics and Astronautics, Jan. 2011. ISBN 978-1-60086-950-1. doi: 10.2514/6.2011-886.

M. R. Fink. Prediction of airfoil tone frequencies. *Journal of Aircraft*, 12, Feb. 1975. ISSN 0021-8669. doi: 10.2514/3.44421.

R. Fjørtoft. *Application of integral theorems in deriving criteria of stability for laminar flows and for the baroclinic circular vortex*. Grøndahl & søns boktr., I kommisjon hos Cammermeyers boghandel, 1950.

R. Fletcher. Function minimization by conjugate gradients. *The Computer Journal*, 7, Feb. 1964. ISSN 0010-4620. doi: 10.1093/comjnl/7.2.149.

M. Fosas De Pando. *Tonal noise generation in flows around aerofoils: a global stability analysis*. PhD thesis, École Polytechnique Laboratoire d'Hydrodynamique, Apr. 2012. URL `https://tel.archives-ouvertes.fr/tel-00816987`.

M. Fosas de Pando and P. J. Schmid. Linear stability of flow through a compressor stage. In *Bull. Amer. Phys. Soc.*, page 205, 2019.

M. Fosas de Pando, D. Sipp, and P. J. Schmid. Efficient evaluation of the direct and adjoint linearized dynamics from compressible flow solvers. *Journal of Computational Physics*, 231, Oct. 2012. ISSN 00219991. doi: 10.1016/j.jcp.2012.06.038.

M. Fosas de Pando, P. J. Schmid, and D. Sipp. A global analysis of tonal noise in flows around aerofoils. *Journal of Fluid Mechanics*, 754, Sept. 2014. ISSN 0022-1120. doi: 10.1017/jfm.2014.356.

M. Fosas de Pando, P. J. Schmid, and D. Sipp. On the receptivity of aerofoil tonal noise: an adjoint analysis. *Journal of Fluid Mechanics*, 812, Feb. 2017. ISSN 0022-1120. doi: 10.1017/jfm.2016.736.

U. W. Ganz, P. D. Joppa, T. J. Patten, and D. F. Scharpf. Boeing 18-inch fan rig broadband noise test. Technical Report NASA/CR-1998-208704, Langley Research Center, National Aeronautics and Space Administration, 1998.

G. A. Gerolymos, G. J. Michon, and J. Neubauer. Analysis and application of chorochronic periodicity in turbomachinery rotor/stator interaction computations. *Journal of Propulsion and Power*, 18, Nov. 2002. ISSN 0748-4658. doi: 10.2514/2.6065.

F. Giannetti and P. Luchini. Structural sensitivity of the first instability of the cylinder wake. *Journal of Fluid Mechanics*, 581, June 2007. ISSN 0022-1120. doi: 10.1017/S0022112007005654.

M. B. Giles and N. A. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65, 2000. ISSN 13866184. doi: 10.1023/A:1011430410075.

S. Glegg. The response of a swept blade row to a three-dimensional gust. *Journal of Sound and Vibration*, 227(1):29–64, 1999. ISSN 0022-460X. doi: https://doi.org/10.1006/jsvi.1999.2327. URL `https://www.sciencedirect.com/science/article/pii/S0022460X99923271`.

C. S. Goh and A. S. Morgans. The influence of entropy waves on the thermoacoustic stability of a model combustor. *Combustion Science and Technology*, 185, Feb. 2013. ISSN 0010-2202. doi: 10.1080/00102202.2012.715828.

A. Gopinath, E. van der Weide, J. Alonso, A. Jameson, K. Ekici, and K. Hall. Three-dimensional unsteady multi-stage turbomachinery simulations using the harmonic balance technique. American Institute of Aeronautics and Astronautics, Jan. 2007. ISBN 978-1-62410-012-3. doi: 10.2514/6.2007-892.

A. Griewank and A. Walther. Algorithm 799: revolve. *ACM Transactions on Mathematical Software*, 26, Mar. 2000. ISSN 0098-3500. doi: 10.1145/347837.347846.

K. Gudmundsson and T. Colonius. Instability wave models for the near-field fluctuations of turbulent jets. *Journal of Fluid Mechanics*, 689, Dec. 2011. ISSN 0022-1120. doi: 10.1017/jfm.2011.401.

G. Guennebaud, B. Jacob, et al. Eigen v3. http://eigen.tuxfamily.org, 2010.

M. D. Gunzburger. *Perspectives in Flow Control and Optimization*. Society for Industrial and Applied Mathematics, Jan. 2002. ISBN 978-0-89871-527-9. doi: 10.1137/1.9780898718720.

W. W. Hager and H. Zhang. A survey of nonlinear conjugate gradient methods. *Pacific journal of Optimization*, 2:35–58, 2006.

M. Hahsler, M. Piekenbrock, and D. Doran. dbscan: Fast density-based clustering with R. *Journal of Statistical Software*, 91, 2019. ISSN 1548-7660. doi: 10.18637/jss.v091.i01.

K. C. Hall, J. P. Thomas, and W. S. Clark. Computation of unsteady nonlinear flows in cascades using a harmonic balance technique. *AIAA Journal*, 40, May 2002. ISSN 0001-1452. doi: 10.2514/2.1754.

C. R. Harris, K. J. Millman, S. J. van der Walt, R. Gommers, P. Virtanen, D. Cournapeau, E. Wieser, J. Taylor, S. Berg, N. J. Smith, R. Kern, M. Picus, S. Hoyer, M. H. van Kerkwijk, M. Brett, A. Haldane, J. F. del Río, M. Wiebe, P. Peterson, P. Gérard-Marchant, K. Sheppard, T. Reddy, W. Weckesser, H. Abbasi, C. Gohlke, and T. E. Oliphant. Array programming with NumPy. *Nature*, 585(7825):357–362, Sept. 2020. doi: 10.1038/s41586-020-2649-2. URL `https://doi.org/10.1038/s41586-020-2649-2`.

L. He. An euler solution for unsteady flows around oscillating blades. *Journal of Turbomachinery*, 112, Oct. 1990. ISSN 0889-504X. doi: 10.1115/1.2927714.

L. He. Method of simulating unsteady turbomachinery flows with multiple perturbations. *AIAA Journal*, 30, Nov. 1992. ISSN 0001-1452. doi: 10.2514/3.11291.

L. He. Averaging for high fidelity modeling—toward large eddy simulations in multipassage multi-row configurations. *Journal of Turbomachinery*, 143, Feb. 2021. ISSN 0889-504X. doi: 10.1115/1.4049616.

L. He and W. Ning. Efficient approach for analysis of unsteady viscous flows in turbomachines. *AIAA Journal*, 36, Nov. 1998. ISSN 0001-1452. doi: 10.2514/2.328.

L. He and D. X. Wang. Concurrent blade aerodynamic-aero-elastic design optimization using adjoint method. *Journal of Turbomachinery*, 133, Jan. 2011. ISSN 0889-504X. doi: 10.1115/1.4000544.

L. He and J. Yi. Two-scale methodology for URANS/Large Eddy Simulation solutions of unsteady turbomachinery flows. *Journal of Turbomachinery*, 139, Oct. 2017. ISSN 0889-504X. doi: 10.1115/1.4036765.

S. Heinz. A review of hybrid RANS-LES methods for turbulent flows: Concepts and applications. *Progress in Aerospace Sciences*, 114, Apr. 2020. ISSN 03760421. doi: 10.1016/j.paerosci.2019.100597.

B. Hellmich and J. R. Seume. Causes of acoustic resonance in a high-speed axial compressor. *Journal of Turbomachinery*, 130, July 2008. ISSN 0889-504X. doi: 10.1115/1.2775487.

V. Hernandez, J. E. Roman, and V. Vidal. Slepc. *ACM Transactions on Mathematical Software*, 31, Sept. 2005. ISSN 0098-3500. doi: 10.1145/1089014.1089019.

M. R. Hestenes, E. Stiefel, et al. *Methods of conjugate gradients for solving linear systems*, volume 49. NBS Washington, DC, 1952.

D. C. Hill. A theoretical approach for analyzing the restabilization of wakes. Technical Report NASA Technical Memorandum 103858, Ames Research Center, Apr. 1992.

C. Hodges. Confinement of vibration by structural irregularity. *Journal of Sound and Vibration*, 82, June 1982. ISSN 0022460X. doi: 10.1016/S0022-460X(82)80022-9.

D. Hoyniak and S. Fleeter. Aerodynamic detuning analysis of an unstalled supersonic turbofan cascade. *Journal of Engineering for Gas Turbines and Power*, 108, Jan. 1986. ISSN 0742-4795. doi: 10.1115/1.3239886.

S. Hubbard and A. Dowling. Acoustic instabilities in premix burners. American Institute of Aeronautics and Astronautics, June 1998. doi: 10.2514/6.1998-2272.

C. E. Hughes. NASA collaborative research on the ultra high bypass engine cycle and potential benefits for noise, performance, and emissions. Technical report, NASA, Glenn Research Center, 2013.

J. D. Hunter. Matplotlib: A 2d graphics environment. *Computing in Science & Engineering*, 9(3):90–95, 2007. doi: 10.1109/MCSE.2007.55.

M. Ihme. Combustion and engine-core noise. *Annual Review of Fluid Mechanics*, 49, Jan. 2017. ISSN 0066-4189. doi: 10.1146/annurev-fluid-122414-034542.

A. Jameson. Aerodynamic design via control theory. *Journal of Scientific Computing*, 3, Sept. 1988. ISSN 0885-7474. doi: 10.1007/BF01061285.

R. Johnstone, L. Chen, and R. D. Sandberg. A sliding characteristic interface condition for direct numerical simulations. *Computers and Fluids*, 107, Jan. 2015. ISSN 00457930. doi: 10.1016/j.compfluid.2014.10.006.

L. E. Jones and R. D. Sandberg. Numerical analysis of tonal airfoil self-noise and acoustic feedback-loops. *Journal of Sound and Vibration*, 330, Dec. 2011. ISSN 0022460X. doi: 10.1016/j.jsv.2011.07.009.

L. E. Jones, N. D. Sandham, and R. D. Sandberg. Acoustic source identification for transitional airfoil flows using cross correlations. *AIAA Journal*, 48, Oct. 2010. ISSN 0001-1452. doi: 10.2514/1.J050345.

M. R. Jovanović, P. J. Schmid, and J. W. Nichols. Sparsity-promoting dynamic mode decomposition. *Physics of Fluids*, 26, Feb. 2014. ISSN 1070-6631. doi: 10.1063/1.4863670.

M. P. Juniper. Absolute and convective instability in gas turbine fuel injectors. American Society of Mechanical Engineers, June 2012. ISBN 978-0-7918-4468-7. doi: 10.1115/GT2012-68253.

K. R. V. Kaza and R. E. Kielb. Flutter and response of a mistuned cascade in incompressible flow. *AIAA Journal*, 20, Aug. 1982. ISSN 0001-1452. doi: 10.2514/3.51172.

C. A. Kennedy, M. H. Carpenter, and R. M. Lewis. Low-storage, explicit Runge-Kutta schemes for the compressible navier-stokes equations. *Applied Numerical Mathematics*, 35:177–219, Nov. 2000. ISSN 01689274. doi: 10.1016/S0168-9274(99)00141-5.

J. W. Kim. Quasi-disjoint pentadiagonal matrix systems for the parallelization of compact finite-difference schemes and filters. *Journal of Computational Physics*, 241, May 2013. ISSN 00219991. doi: 10.1016/j.jcp.2013.01.046.

W. Koch. On the transmission of sound waves through a blade row. *Journal of Sound and Vibration*, 18, Sept. 1971. ISSN 0022460X. doi: 10.1016/0022-460X(71)90635-3.

S. K. Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103:16–42, Nov. 1992. ISSN 10902716. doi: 10.1016/0021-9991(92)90324-R.

S. K. Lele. Computational aeroacoustics: A review. American Institute of Aeronautics and Astronautics Inc, AIAA, 1997. doi: 10.2514/6.1997-18.

S. K. Lele and J. W. Nichols. A second golden age of aeroacoustics? *Philosophical Transactions of the Royal Society A: Mathematical, Physical and Engineering Sciences*, 372, Aug. 2014. ISSN 1364503X. doi: 10.1098/rsta.2013.0321.

M. J. Lighthill. On sound generated aerodynamically I. general theory. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 211, Mar. 1952. ISSN 0080-4630. doi: 10.1098/rspa.1952.0060.

M. J. Lighthill. On sound generated aerodynamically II. turbulence as a source of sound. *Proceedings of the Royal Society of London. Series A. Mathematical and Physical Sciences*, 222, Feb. 1954. ISSN 0080-4630. doi: 10.1098/rspa.1954.0049.

G. Lodato, P. Domingo, and L. Vervisch. Three-dimensional boundary conditions for direct and large-eddy simulation of compressible viscous flows. *Journal of Computational Physics*, 227, May 2008. ISSN 00219991. doi: 10.1016/j.jcp.2008.01.038.

Y. Lu, J. Green, S. C. Stapelfeldt, and M. Vahdati. Effect of geometric variability on running shape and performance of a transonic fan. *Journal of Turbomachinery*, 141, Oct. 2019. ISSN 0889-504X. doi: 10.1115/1.4044676.

J. L. Lumley. The structure of inhomogeneous turbulent flows. Proceedings of the USSR Academy of Sciences, 1967.

L. Magri, J. O'Brien, and M. Ihme. Compositional inhomogeneities as a source of indirect combustion noise. *Journal of Fluid Mechanics*, 799, July 2016. ISSN 0022-1120. doi: 10.1017/jfm.2016.397.

R. Mailach, I. Lehmann, and K. Vogeler. Rotating instabilities in an axial compressor originating from the fluctuating blade tip vortex. *Journal of Turbomachinery*, 123, July 2001. ISSN 0889-504X. doi: 10.1115/1.1370160.

E. Manoha, C. Herrero, S. B. Khelil, P. Guillen, P. Sagaut, and I. Mary. Numerical prediction of airfoil aerodynamic noise. American Institute of Aeronautics and Astronautics, June 2002. ISBN 978-1-62410-119-9. doi: 10.2514/6.2002-2573.

F. Marble and S. Candel. Acoustic disturbance from gas non-uniformities convected through a nozzle. *Journal of Sound and Vibration*, 55, Nov. 1977. ISSN 0022460X. doi: 10.1016/0022-460X(77)90596-X.

C. Martel, R. Corral, and J. M. Llorens. Stability increase of aerodynamically unstable rotors using intentional mistuning. *Journal of Turbomachinery*, 130, Jan. 2008. ISSN 0889-504X. doi: 10.1115/1.2720503.

B. E. Mitchell, S. K. Lele, and P. Moin. Direct computation of the sound from a compressible co-rotating vortex pair. *Journal of Fluid Mechanics*, 285, Feb. 1995. ISSN 0022-1120. doi: 10.1017/S0022112095000504.

C. Moler and C. Van Loan. Nineteen dubious ways to compute the exponential of a matrix, twenty-five years later. Technical report, 2003. URL `http://www.siam.org/journals/sirev/45-1/41801.html`.

C. B. Moler and G. W. Stewart. An algorithm for generalized matrix eigenvalue problems. *SIAM Journal on Numerical Analysis*, 10, Apr. 1973. ISSN 0036-1429. doi: 10.1137/0710024.

S. Moreau. Turbomachinery noise predictions: Present and future. *Acoustics*, 1, Jan. 2019. ISSN 2624-599X. doi: 10.3390/acoustics1010008.

K. Moreland. Diverging color maps for scientific visualization. Springer-Verlag, 2009. doi: 10.1007/978-3-642-10520-3_9.

R. B. Morgan and M. Zeng. A harmonic restarted Arnoldi algorithm for calculating eigenvalues and determining multiplicity. *Linear Algebra and its Applications*, 415, May 2006. ISSN 00243795. doi: 10.1016/j.laa.2005.07.024.

A. S. Morgans and I. Duran. Entropy noise: A review of theory, progress and challenges. *International Journal of Spray and Combustion Dynamics*, 8, Dec. 2016. ISSN 1756-8277. doi: 10.1177/1756827716651791.

E. Motheau, F. Nicoud, and T. Poinsot. Mixed acoustic–entropy combustion instabilities in gas turbines. *Journal of Fluid Mechanics*, 749, June 2014. ISSN 0022-1120. doi: 10.1017/jfm.2014.245.

G. Mouret, N. Gourdain, and L. Castillon. Adaptation of phase-lagged boundary conditions to Large Eddy Simulation in turbomachinery configurations. *Journal of Turbomachinery*, 138, Apr. 2016. ISSN 0889-504X. doi: 10.1115/1.4032044.

MPI Forum. MPI Forum. https://www.mpi-forum.org/, 2021. Accessed: 12/06/2021.

E. C. Nash, M. V. Lowson, and A. McAlpine. Boundary-layer instability noise on aerofoils. *Journal of Fluid Mechanics*, 382, Mar. 1999. ISSN 0022-1120. doi: 10.1017/S002211209800367X.

E. Nesbitt. Towards a quieter low pressure turbine: Design characteristics and prediction needs. *International Journal of Aeroacoustics*, 10, Feb. 2010. ISSN 1475-472X. doi: 10.1260/1475-472X.10.1.1.

E. J. Nielsen and B. Diskin. Discrete adjoint-based design for unsteady turbulent flows on dynamic overset unstructured grids. *AIAA Journal*, 51:1355–1373, June 2013. ISSN 0001-1452. doi: 10.2514/1.J051859. URL `http://arc.aiaa.org/doi/10.2514/1.J051859`.

J. Nocedal and S. Wright. *Numerical Optimization*. Springer New York, 2006. ISBN 978-0-387-30303-1. doi: 10.1007/978-0-387-40065-5.

G. Ntanakas and M. Meyer. The unsteady discrete adjoint method for turbomachinery applications. Sept. 2015. doi: 10.5281/ZENODO.823136.

G. Ntanakas, M. Meyer, and K. C. Giannakoglou. Employing the time-domain unsteady discrete adjoint method for shape optimization of three-dimensional multirow turbomachinery configurations. *Journal of Turbomachinery*, 140, Aug. 2018. ISSN 15288900. doi: 10.1115/1.4040564.

B. Olson, S. Shaw, C. Shi, C. Pierre, R. Parker, R. P. Circulant, B. J. Olson, S. W. Shaw, L. S. Randolph, and D. Head. Circulant matrices and their application to vibration analysis. *Applied Mechanics Reviews*, 66:10, 2014. doi: 10.1115/1.4027722ï. URL `https://hal.archives-ouvertes.fr/hal-01378829`.

R. Parker. Resonance effects in wake shedding from compressor blading. *Journal of Sound and Vibration*, 6, Nov. 1967. ISSN 0022460X. doi: 10.1016/0022-460X(67)90202-7.

R. Parker. Acoustic resonances and blade vibration in axial flow compressors. *Journal of Sound and Vibration*, 92, Feb. 1984. ISSN 0022460X. doi: 10.1016/0022-460X(84)90196-2.

R. Parker and S. Stoneman. An experimental investigation of the generation and consequences of acoustic waves in an axial flow compressor: Large axial spacings between blade rows. *Journal of Sound and Vibration*, 99, Mar. 1985. ISSN 0022460X. doi: 10.1016/0022-460X(85)90355-4.

R. Parker and S. A. T. Stoneman. The excitation and consequences of acoustic resonances in enclosed fluid flow around solid bodies. *Proceedings of the Institution of Mechanical Engineers, Part C: Mechanical Engineering Science*, 203, Jan. 1989. ISSN 0263-7154. doi: 10.1243/PIME_PROC_1989_203_081_02.

R. W. Paterson, P. G. Vogt, M. R. Fink, and C. Lee Munch. Vortex noise of isolated airfoils. *Journal of Aircraft*, 10, May 1973. ISSN 0021-8669. doi: 10.2514/3.60229.

N. Peake and A. B. Parry. Modern challenges facing turbomachinery aeroacoustics. *Annual Review of Fluid Mechanics*, 44, Jan. 2012. ISSN 0066-4189. doi: 10.1146/annurev-fluid-120710-101231.

C. Pierre and D. V. Murthy. Aeroelastic modal characteristics of mistuned blade assemblies - mode localization and loss of eigenstructure. *AIAA Journal*, 30, Oct. 1992. ISSN 0001-1452. doi: 10.2514/3.11251.

T. Poinsot and S. Lele. Boundary conditions for direct simulations of compressible viscous flows. *Journal of Computational Physics*, 101, July 1992. ISSN 00219991. doi: 10.1016/0021-9991(92)90046-2.

E. Polak and G. Ribiere. Note sur la convergence de méthodes de directions conjuguées. *ESAIM: Mathematical Modeling and Numerical Analysis-Modeling Mathématique et Analyze Numérique*, 3:35–43, 1969.

B. T. Polyak. The conjugate gradient method in extremal problems. *USSR Computational Mathematics and Mathematical Physics*, 9:94–112, 1969.

S. Pröbsting, J. Serpieri, and F. Scarano. Experimental investigation of aerofoil tonal noise generation. *Journal of Fluid Mechanics*, 747, May 2014. ISSN 0022-1120. doi: 10.1017/jfm.2014.156.

L. Rayleigh. On the stability, or instability, of certain fluid motions. *Proceedings of the London Mathematical Society*, s1-11, Nov. 1879. ISSN 00246115. doi: 10.1112/plms/s1-11.1.57.

S. C. Reddy, P. J. Schmid, and D. S. Henningson. Pseudospectra of the Orr–Sommerfeld operator. *SIAM Journal on Applied Mathematics*, 53, Feb. 1993. ISSN 0036-1399. doi: 10.1137/0153002.

R. Reichert and S. Biringen. Numerical simulation of compressible plane jets. American Institute of Aeronautics and Astronautics, June 1997. doi: 10.2514/6.1997-1924.

J. Reuther and A. Jameson. Aerodynamic shape optimization of wing and wing-body configurations using control theory. American Institute of Aeronautics and Astronautics, Jan. 1995. doi: 10.2514/6.1995-123.

J. Reuther, A. Jameson, J. Farmer, L. Martinelli, and D. Saunders. Aerodynamic shape optimization of complex aircraft configurations via an adjoint formulation. American Institute of Aeronautics and Astronautics, Jan. 1996. doi: 10.2514/6.1996-94.

J. H. M. Ribeiro, C. Yeh, and K. Taira. Randomized resolvent analysis. *Physical Review Fluids*, 5, Mar. 2020. ISSN 2469-990X. doi: 10.1103/PhysRevFluids.5.033902.

W. Rizk and D. F. Seymour. Investigations into the failure of gas circulators and circuit components at Hinkley Point nuclear power station. *Proceedings of the Institution of Mechanical Engineers*, 179, June 1964. ISSN 0020-3483. doi: 10.1243/PIME_PROC_1964_179_042_02.

M. Roger and S. Moreau. Extensions and limitations of analytical airfoil broadband noise models. *International Journal of Aeroacoustics*, 9, May 2010. ISSN 1475-472X. doi: 10.1260/1475-472X.9.3.273.

Rolls Royce. Rolls Royce Future Products, May 2021. URL `https://www.rolls-royce.com/products-and-services/civil-aerospace/future-products.aspx#/`.

A. Roy and R. Govindarajan. *An Introduction to Hydrodynamic Stability*, pages 131–147. Springer New York, 1 edition, 2010. ISBN 978-1-4419-6493-9. doi: 10.1007/978-1-4419-6494-6.

R. Sandberg and A. Wheeler. Effect of trailing-edge boundary conditions on acoustic feedback loops in high-pressure turbines. *Journal of Sound and Vibration*, 461, Nov. 2019. ISSN 0022460X. doi: 10.1016/j.jsv.2019.114917.

R. D. Sandberg, V. Michelassi, R. Pichler, L. Chen, and R. Johnstone. Compressible direct numerical simulation of low-pressure turbines—Part I: Methodology. *Journal of Turbomachinery*, 137, May 2015. ISSN 0889-504X. doi: 10.1115/1.4028731.

N. L. Sanger and R. P. Shreeve. Comparison of calculated and experimental cascade performance for controlled-diffusion compressor stator blading. *Journal of Turbomachinery*, 108, July 1986. ISSN 0889-504X. doi: 10.1115/1.3262023.

M. Sanjose, A. Towne, P. Jaiswal, S. Moreau, S. Lele, and A. Mann. Modal analysis of the laminar boundary layer instability and tonal noise of an airfoil at Reynolds number 150, 000. *International Journal of Aeroacoustics*, 18, Apr. 2019. ISSN 1475-472X. doi: 10.1177/1475472X18812798.

T. Sasaki and F. Breugelmans. Comparison of sweep and dihedral effects on compressor cascade performance. *Journal of Turbomachinery*, 120, July 1998. ISSN 0889-504X. doi: 10.1115/1.2841738.

P. J. Schmid. Nonmodal stability theory. *Annual Review of Fluid Mechanics*, 39:129–162, Jan. 2007. ISSN 0066-4189. doi: 10.1146/annurev.fluid.38.050304.092139.

P. J. Schmid. Dynamic mode decomposition of numerical and experimental data. *Journal of Fluid Mechanics*, 656, Aug. 2010. ISSN 0022-1120. doi: 10.1017/S0022112010001217.

P. J. Schmid and L. Brandt. Analysis of fluid systems: Stability, receptivity, sensitivity. *Applied Mechanics Reviews*, 66, Mar. 2014. ISSN 0003-6900. doi: 10.1115/1.4026375.

P. J. Schmid and D. S. Henningson. *Stability and Transition in Shear Flows*, volume 142. Springer New York, 2001. ISBN 978-1-4612-6564-1. doi: 10.1007/978-1-4613-0185-1.

P. J. Schmid, M. Fosas De Pando, and N. Peake. Stability analysis for n-periodic arrays of fluid systems. *Physical Review Fluids*, 2:113902, Nov. 2017. ISSN 2469990X. doi: 10.1103/PhysRevFluids.2.113902. URL `https://journals.aps.org/prfluids/abstract/10.1103/PhysRevFluids.2.113902`.

J. C. Schulze, P. J. Schmid, and J. L. Sesterhenn. Exponential time integration using Krylov subspaces. *International Journal for Numerical Methods in Fluids*, 60:591–609, June 2009. ISSN 02712091. doi: 10.1002/fld.1902. URL `http://doi.wiley.com/10.1002/fld.1902`.

L. Selle, F. Nicoud, and T. Poinsot. Actual impedance of nonreflecting boundary conditions: Implications for computation of resonators. *AIAA Journal*, 42, May 2004. ISSN 0001-1452. doi: 10.2514/1.1883.

J. Sesterhenn. A characteristic-type formulation of the Navier-Stokes equations for high order upwind schemes. *Computers and Fluids*, 30:37–67, Sept. 2000. ISSN 00457930. doi: 10.1016/s0045-7930(00)00002-5.

S. C. Spiegel, H. Huynh, and J. R. DeBonis. A survey of the isentropic euler vortex problem using high-order methods. In *22nd AIAA Computational Fluid Dynamics Conference*, page 2444, June 2015. doi: 10.2514/6.2015-2444.

A. Srinivasan. Vibrations of bladed-disk assemblies- a selected survey. *ASME, Transactions, Journal of Vibration, Acoustics, Stress and Reliability in Design*, 106:165–168, 1984.

J. Steger, F. Dougherty, and J. Benek. A chimera grid scheme. June 1982.

R. Steijl and G. Barakos. Sliding mesh algorithm for CFD analysis of helicopter rotor-fuselage aerodynamics. *International Journal for Numerical Methods in Fluids*, 58, Oct. 2008. ISSN 02712091. doi: 10.1002/fld.1757.

G. W. Stewart. A Krylov–Schur algorithm for large eigenproblems. *SIAM Journal on Matrix Analysis and Applications*, 23, Jan. 2002. ISSN 0895-4798. doi: 10.1137/S0895479800371529.

K. Taira, S. L. Brunton, S. T. M. Dawson, C. W. Rowley, T. Colonius, B. J. McKeon, O. T. Schmidt, S. Gordeyev, V. Theofilis, and L. S. Ukeiley. Modal analysis of fluid flows: An overview. *AIAA Journal*, 55, Dec. 2017. ISSN 0001-1452. doi: 10.2514/1.J056060.

K. Taira, M. S. Hemati, S. L. Brunton, Y. Sun, K. Duraisamy, S. Bagheri, S. T. M. Dawson, and C.-A. Yeh. Modal analysis of fluid flows: Applications and outlook. *AIAA Journal*, 58, Mar. 2020. ISSN 0001-1452. doi: 10.2514/1.J058462.

C. Tam, K. Viswanathan, K. K. Ahuja, and J. Panda. The sources of jet noise: Experimental evidence. American Institute of Aeronautics and Astronautics, May 2007. ISBN 978-1-62410-003-1. doi: 10.2514/6.2007-3641.

C. K. Tam. Computational aeroacoustics: An overview of computational challenges and applications. *International Journal of Computational Fluid Dynamics*, 18: 547–567, 8 2004. ISSN 10618562. doi: 10.1080/10618560410001673551.

C. K. Tam and J. C. Webb. Dispersion-relation-preserving finite difference schemes for computational acoustics. *Journal of Computational Physics*, 107(2):262–281, Aug. 1993. doi: 10.1006/jcph.1993.1142.

C. K. W. Tam. Discrete tones of isolated airfoils. *The Journal of the Acoustical Society of America*, 55, June 1974. ISSN 0001-4966. doi: 10.1121/1.1914682.

O. Tammisola and M. P. Juniper. Coherent structures in a swirl injector at $Re = 4800$ by nonlinear simulations and linear global modes. *Journal of Fluid Mechanics*, 792, Apr. 2016. ISSN 0022-1120. doi: 10.1017/jfm.2016.86.

The HDF Group. https://www.hdfgroup.org/, 2021. Accessed: 12/06/2021.

J. P. Thomas, K. C. Hall, and E. H. Dowell. Discrete adjoint approach for modeling unsteady aerodynamic design sensitivities. *AIAA Journal*, 43, Sept. 2005. ISSN 0001-1452. doi: 10.2514/1.731.

L. N. Trefethen. Computation of pseudospectra. *Acta Numerica*, 8, Jan. 1999. ISSN 0962-4929. doi: 10.1017/S0962492900002932.

L. N. Trefethen and D. Bau III. *Numerical Linear Algebra*. SIAM, 1 edition, 1997.

J. H. Tu, C. W. Rowley, D. M. Luchtenburg, S. L. Brunton, and J. N. Kutz. On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1, June 2014. ISSN 2158-2505. doi: 10.3934/jcd.2014.1.391.

J. M. Tyler and T. G. Sofrin. Axial flow compressor noise studies. Technical report, SAE Technical Paper, Jan. 1962.

P. Virtanen, R. Gommers, T. E. Oliphant, M. Haberland, T. Reddy, D. Cournapeau, E. Burovski, P. Peterson, W. Weckesser, J. Bright, S. J. van der Walt, M. Brett, J. Wilson, K. J. Millman, N. Mayorov, A. R. J. Nelson, E. Jones, R. Kern, E. Larson, C. J. Carey, İ. Polat, Y. Feng, E. W. Moore, J. VanderPlas, D. Laxalde, J. Perktold, R. Cimrman, I. Henriksen, E. A. Quintero, C. R. Harris, A. M. Archibald, A. H. Ribeiro, F. Pedregosa, P. van Mulbregt, and SciPy 1.0 Contributors. SciPy 1.0: Fundamental Algorithms for Scientific Computing in Python. *Nature Methods*, 17:261–272, 2020. doi: 10.1038/s41592-019-0686-2.

D. X. Wang and L. He. Adjoint aerodynamic design optimization for blades in multistage turbomachines-part I: Methodology and verification. *Journal of Turbomachinery*, 132, Apr. 2010. ISSN 0889504X. doi: 10.1115/1.3072498.

D. X. Wang, L. He, Y. S. Li, and R. G. Wells. Adjoint aerodynamic design optimization for blades in multistage turbomachines-part II: Validation and application. *Journal of Turbomachinery*, 132, Apr. 2010. ISSN 0889504X. doi: 10.1115/1.3103928.

G. Wang, F. Duchaine, D. Papadogiannis, I. Duran, S. Moreau, and L. Y. Gicquel. An overset grid method for large eddy simulation of turbomachinery stages. *Journal of Computational Physics*, 274, Oct. 2014a. ISSN 00219991. doi: 10.1016/j.jcp.2014.06.006.

M. Wang, J. B. Freund, and S. K. Lele. Computational prediction of flow-generated sound. *Annu. Rev. Fluid Mech*, 38:483–512, 2006. doi: 10.1146/annurev.fluid.

Q. Wang, P. Moin, and G. Iaccarino. Minimal repetition dynamic checkpointing algorithm for unsteady adjoint calculation. *SIAM Journal on Scientific Computing*, 31, Jan. 2009. ISSN 1064-8275. doi: 10.1137/080727890.

Q. Wang, R. Hu, and P. Blonigan. Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267, June 2014b. ISSN 00219991. doi: 10.1016/j.jcp.2014.03.002.

S.-T. Wei and C. Pierre. Localization phenomena in mistuned assemblies with cyclic symmetry part I: Free vibrations. *Journal of Vibration and Acoustics*, 110, Oct. 1988a. ISSN 1048-9002. doi: 10.1115/1.3269547.

S.-T. Wei and C. Pierre. Localization phenomena in mistuned assemblies with cyclic symmetry part II: Forced vibrations. *Journal of Vibration and Acoustics*, 110, Oct. 1988b. ISSN 1048-9002. doi: 10.1115/1.3269548.

D. G. Wilson and T. Korakianitis. *The Design of High-Efficiency Turbomachinery and Gas Turbines*. The MIT Press, 2 edition, Sept. 2014.

M. J. Wilson, M. Imregun, and A. I. Sayma. The effect of stagger variability in gas turbine fan assemblies. *Journal of Turbomachinery*, 129, Apr. 2007. ISSN 0889-504X. doi: 10.1115/1.2437776.

H. Wu, R. D. Sandberg, and S. Moreau. Stability characteristics of different aerofoil flows at $re_c = 150,000$ and the implications for aerofoil self-noise. *Journal of Sound and Vibration*, 506, Aug. 2021. ISSN 0022460X. doi: 10.1016/j.jsv.2021.116152.

H. Yang and L. He. Experimental study on linear compressor cascade with three-dimensional blade oscillation. *Journal of Propulsion and Power*, 20, Jan. 2004. ISSN 0748-4658. doi: 10.2514/1.1280.

F. Zahle, N. N. Sørensen, and J. Johansen. Wind turbine rotor-tower interaction using an incompressible overset grid method. *Wind Energy*, 12, Sept. 2009. ISSN 10954244. doi: 10.1002/we.327.

H. Zhang, E. M. Constantinescu, and B. F. Smith. PETSc TSAdjoint: a discrete adjoint ODE solver for first-order and second-order sensitivity analysis. *arXiv e-preprints*, 2019.

Appendix A

Mode Clusters for Phase and Group Velocity Analysis

A.1 Introduction

To save space and improve the flow of the thesis, additional plots of the mode clusters identified in the phase-speed analysis in the latter part of Section 3.3 have been moved into this appendix, so that there is sufficient space to render the mode plots in full size. We hope that the inclusion of these will help the reader understand the different families of dynamics discussed in that section of the our work.

A.2 The Plots

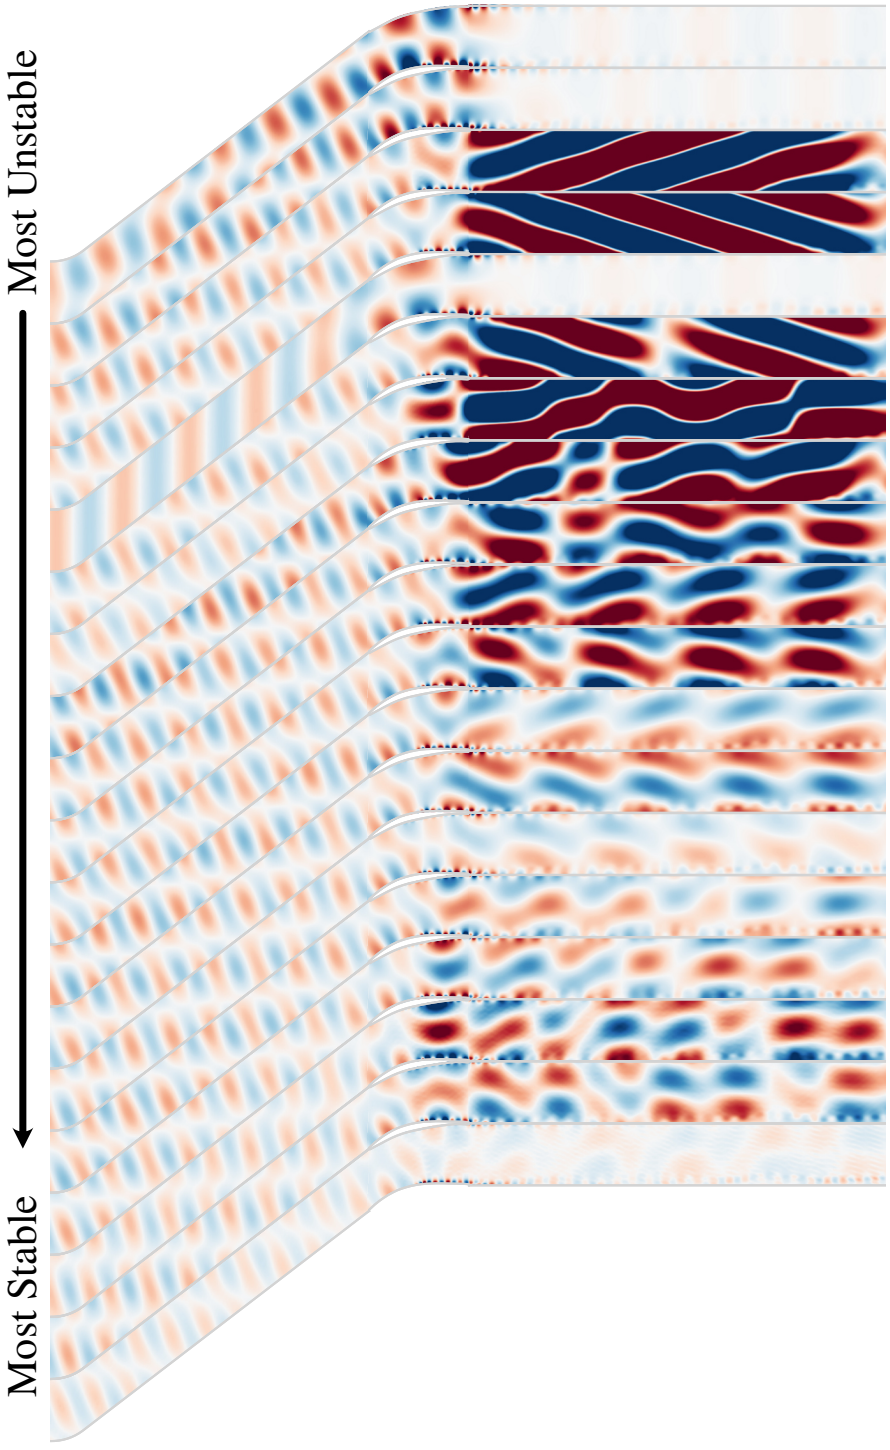


Figure A.1: Cluster 1

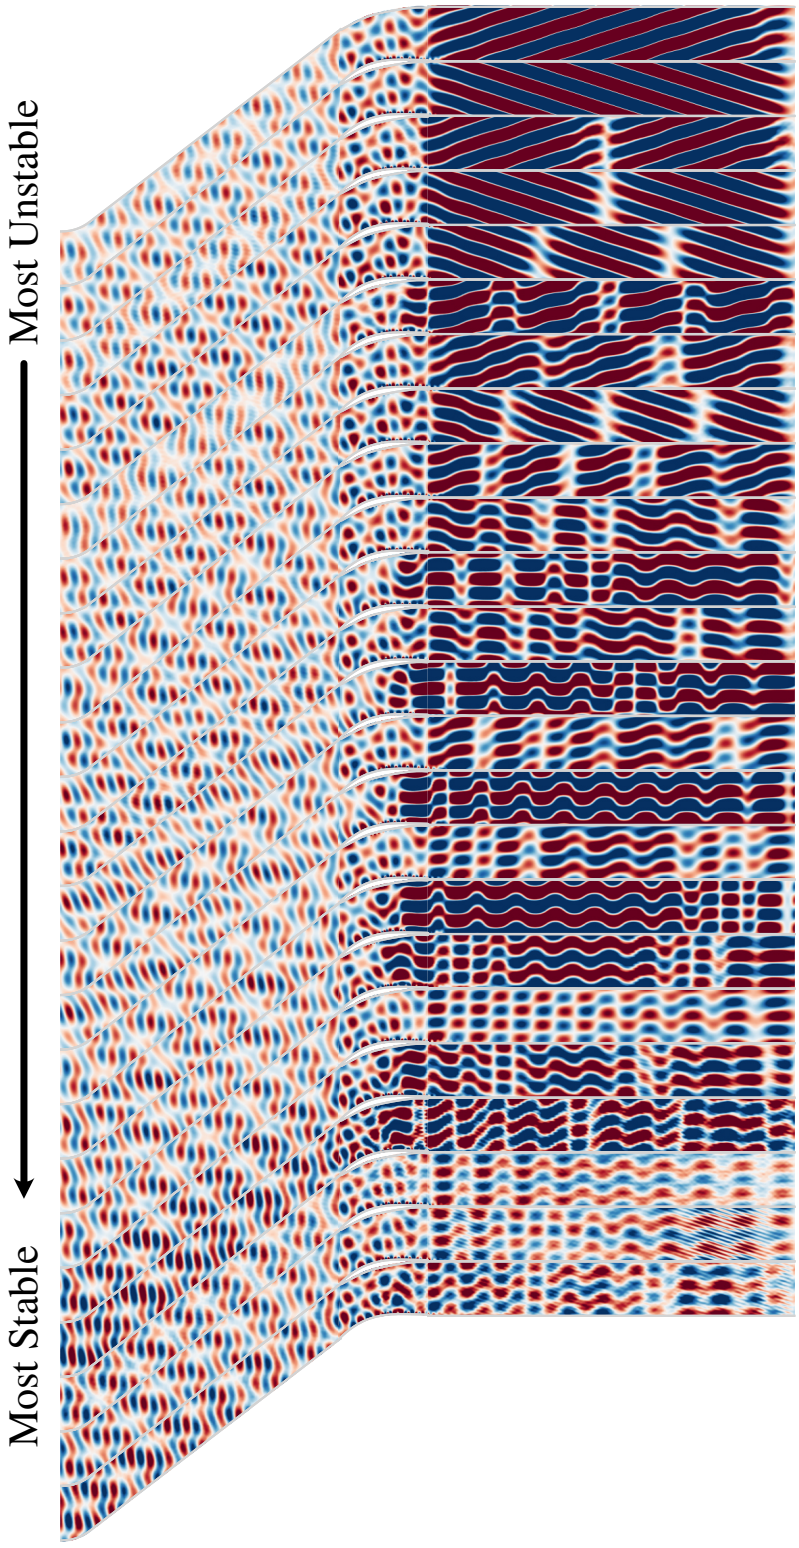


Figure A.2: Cluster 2

Appendix B

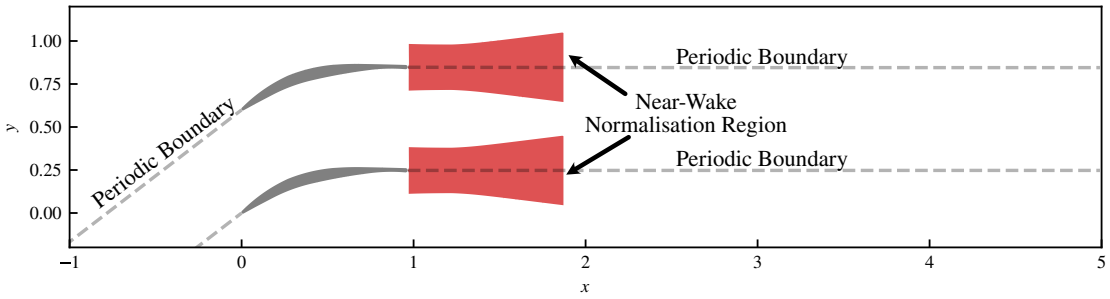
Normalisation Regions for Global Modes

B.1 Near-Wake Axial Velocity Normalisation Regions

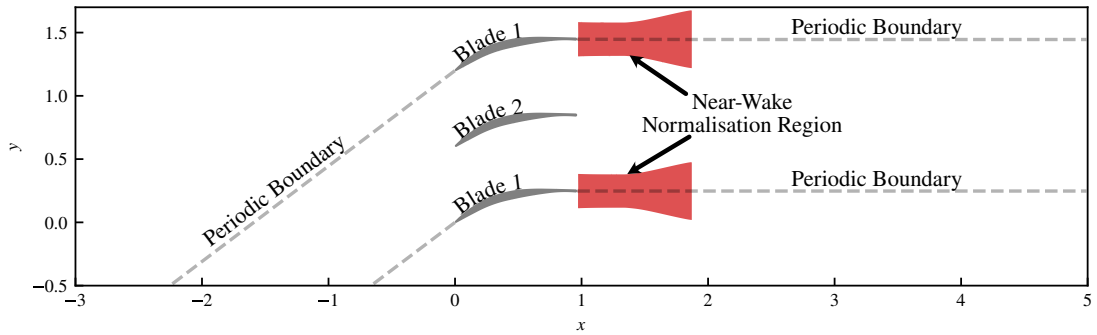
All the direct global modes in this thesis are normalised by the near-wake axial velocity component, u, of the state vector. This choice is made because, for most modes here, the dynamics of the mode are driven by a self-excited instability that generates a wake that is convected downstream of the trailing edge. By performing the normalisation on these wake structures, and considering modes describing similar dynamics, a comparative analysis of the mode structures can be performed.

Since the normalisation procedure has to account for multiple blade passages, the region in which the normalisation is performed is made compact, so that one blade wake is selected at a time. To simplify the procedure, the closed downstream region, bounded by the grid lines $n_y = \pm 100$, $n_x = 1081$ and $n_x = 1281$, is selected, thereby avoiding interpolation and complex array slicing. This gives a region with an initial height of $\Delta y \approx 0.77$ centred vertically on the trailing edge, and extending for $\Delta x \approx 1.35$ before expanding as a trapezoid until $x \approx 1.87$, with a final height of $\Delta y \approx 0.95$. This is shown in Figure B.1a.

For the mistuning cases in Appendix C, the normalisation region is located at the trailing edge of the unperturbed blade, 'blade 1', as the flow there is expected to



(a)



(b)

Figure B.1: The normalisation regions for the single-passage/tuned case in (a), and the mistuned cases considered in Appendix C in (b).

be affected less than on the surface of the perturbed blade, 'blade 2'. This is shown in Figure B.1b.

Appendix C

Sensitivities of Eigenvalues to Mistuning

C.1 Introduction

The work in the latter sections of Chapter 3 considered periodic cascades with blade geometries and mean flows that were identical in each blade passage. However, although blades disks and fans are normally designed with perfect cyclic symmetry, realistic scenarios in working machinery rarely satisfy this assumption. Instead these blades may be *mistuned* in some way, such as in their mechanical properties, stagger angle, pitch height, surface geometry or the effects of other components of the machine altering the mean flow in some portion of the cascade. This can be by design, as a result of wear and damage to the aerodynamic surfaces, or due to blade untwist altering the stagger angles of individual blades under the applied aerodynamic loads [Wilson et al., 2007; Lu et al., 2019].

Prior research into mistuning in turbomachinery contexts has focused predominately on the forced response and aeroelasticity problems for *structural* perturbations (see reviews in [Srinivasan, 1984], [Castanier and Pierre, 2006] for the details of earlier work on bladed disks), where the mechanical properties of the bladed disk and/or blisk are varied, and a number of general observations have been made. For instance, structural mistuning is generally considered stabilising to the flutter response but may lead to mode localisation, where small portions of the cascade have high vi-

brational amplitudes, leading to high-cycle fatigue for the structural components [Castanier and Pierre, 2006]. This phenomenon was first identified by Anderson[1] in the context of solid-state physics [Anderson, 1958] but was later adapted for vibration localisation phenomena by Hodges [1982], who showed that, for a coupled pendulum system, localisation depends on the ratio of disorder to coupling strength. Work by Wei and Pierre extended this to bladed disk assemblies in free [Wei and Pierre, 1988a] and forced [Wei and Pierre, 1988b] vibration scenarios, and in the process identified strong localisation and eigenvalue veering phenomena associated with small amounts of mistuning in systems where the oscillating sub-structures are weakly coupled.

More recently *aerodynamic* mistuning, where the aerodynamic properties of the geometry are altered, has been investigated. For example, Wilson et al. [2007] considered the effect of mis-staggering (i.e. mistuning by altering the stagger angle) of fan blade cascades at conditions where the mistuning resulted in the shock waves being alternately 'swallowed' by the blade passages. Due to varying aerodynamic loading, the fan blades untwist and result in the propagation of saw-tooth mis-staggering patterns throughout the annulus; a phenomenon termed *alternate passage divergence* [Lu et al., 2019]. Furthermore, Ekici et al. [2010] showed that alternating variations applied to blade pitch resulted in stabilisation of flutter, while the response to mis-staggering depended on the direction the mistuning was applied, with the effects attributed to local changes to the steady flow field. Potential benefits of mistuning for passive vibration and flutter control have also prompted a number of studies into applying intentional perturbations to control flutter and forced response of these structures (see for example [Kaza and Kielb, 1982; Crawley and Hall, 1985; Hoyniak and Fleeter, 1986; Castanier and Pierre, 2002; Martel et al., 2008].).

To the best of our knowledge, however, limited work has been performed to identify effects of aerodynamic mistuning on self-excited flow instabilities and the re-

[1]This is also known as Anderson localisation.

sulting acoustic responses. To that end, in this appendix we will perform exploratory calculations to evaluate the effects of mistuning the stagger angle and pitch of a two-passage geometry, in preparation for larger, future simulations.

As we have observed above for the tuned case, the costs of a direct calculation of the modes on the mistuned geometry are prohibitive. It should also be noted that a block-circulant-based decomposition, of the kind shown in Section 3.4.2, no longer leads to a block-diagonal $\hat{\mathbf{A}}$, and so a decoupling of the large problem into smaller sub-problems is no longer possible. Despite this, the N-periodic approach can still be leveraged for this analysis, under certain assumptions that we will discuss here.

This appendix presents some observations of the aerodynamically mistuned aeroacoustics problem from the global stability point of view, using a toy problem with a two-blade linear cascade. The goal here is to demonstrate the mechanics of the problem, and the effects that a fixed-amplitude mistuning of the stagger angle or blade pitch has on self-excited instabilities in the context of the feedback mechanisms identified earlier. We also make use of this opportunity to discus the theoretical basis of how the N-periodic analysis presented in Section 3.4 can be leveraged to address the problem of large cascades without having to resort to calculations of full cascade geometries, which will be a subject for future research.

C.2 Formulation

As mentioned in the introduction, a simplified two-periodic domain is considered here, at flow conditions matching those of the one-periodic mean flow calculation in Section 3.2.1, with the blade situated on the periodic boundary henceforth referred to as 'blade 1', while the middle blade is labelled 'blade 2' (see Figure B.1b for a diagram). For convenience and consistency, blade 2 is chosen as the location where the perturbation to the geometry is applied. As our goal here is to test the mechanics and obtain some preliminary observations of the aeroacoustic response to aerodynamic mistuning, we will consider a two-passage case to enable us to compute

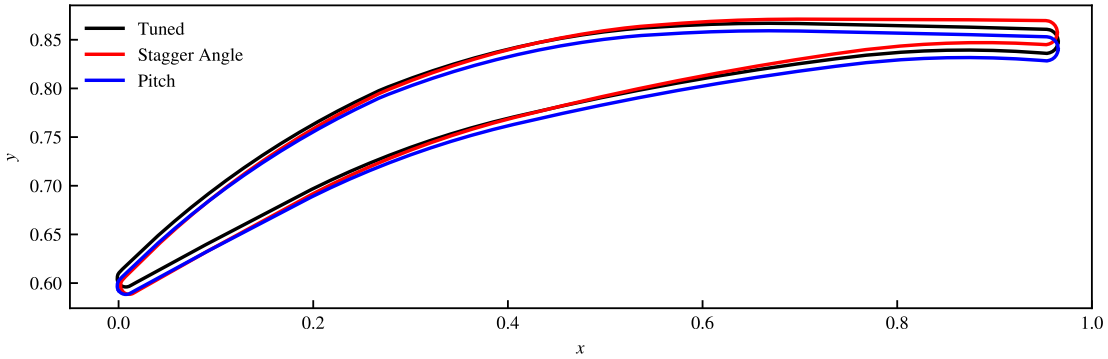


Figure C.1: Details of the perturbations to the mistuned geometries compared to the tuned 2-passage cascade. The variation in the y-coordinate of the leading edge in both cases is 5×10^{-3} in both the stagger angle and pitch perturbations.

the eigenvalues directly. The geometries considered here are shown in Figure C.1 and consist of a tuned case, a case with the stagger angle decreased by one degree, $\delta\gamma = -1^\circ$ about the mid-chord, and finally a case where the pitch is perturbed so that the leading edge of the blade moves downwards by approximately $\delta s = 5 \times 10^{-3}$. This change in pitch is equivalent to a change in the leading edge position in the mistuned stagger angle case, thus providing a degree of consistency between the two mistuning patterns. Since our goal here is not to do a parametric analysis, and as a result of limitations to our computer resources, we restrict ourselves to these three cases and leave the parameter sweep for future work.

C.3 Mean Flow Modifications

Once again mean flows are computed here for the three cases in line with the procedure described in Section 3.2.3. Unlike the tuned cases that we investigated in Chapter 3, to aid comparison both the tuned and mistuned mean flows have been computed for the two-passage case directly, instead of tessellating a single passage flow. These are compared in Figure C.2 by subtracting the mistuned field from the tuned one and scaling by the perturbation epsilon $\epsilon = 5 \times 10^{-3}$, and we interpolate the mistuned field onto the tuned mesh in order to allow comparison of the mean

flows.

As expected, alterations to the stagger angle alter both the pressure and velocity distributions within the blade passages, with an increase in the velocity observed in the lower passage due to the reduced 'cross-sectional area' of the that passage. Note that in this case it appears that the modifications to the near-wall velocity are substantially outside the linear regime with ϵ-normalised variations on the order of $\mathcal{O}(10^2)$. Pressure field variations are smaller and $\mathcal{O}(1)$. Entropy fields suggest that the suction side separation point has moved further downstream, with increased separation and shedding in the vicinity of the leading edge on the pressure side. Vortex shedding near the trailing edge appears suppressed on the pressure side, though, with increased shedding on the suction side from the separation bubble. Further to the large changes in the flow in the vicinity of the perturbed blade, it is interesting to observe that the changes to the flow around the unperturbed blade consist largely of adjustments to the surface pressure distribution and modifications of the azimuthal velocity component near the leading edge, with only very minor changes to the entropy and axial velocity components. Despite the small blade count in this cascade, this nevertheless presents us with a convenient opportunity to compare the local and global feedback mechanisms in this setting.

By virtue of not having to adjust the flow angle, the pitch perturbation localises the response in the mean flow to the blade surface more than the former case. Broadly speaking, the increase in passage height in the upper passage has reduced the flow velocity over the suction side while increasing it on the pressure side of the perturbed blade. Shedding strength from the laminar bubble is therefore decreased while increasing significantly on the pressure side. Once again we observe that except for small pressure field changes, the unperturbed blade remains largely unaffected by the modified geometry.

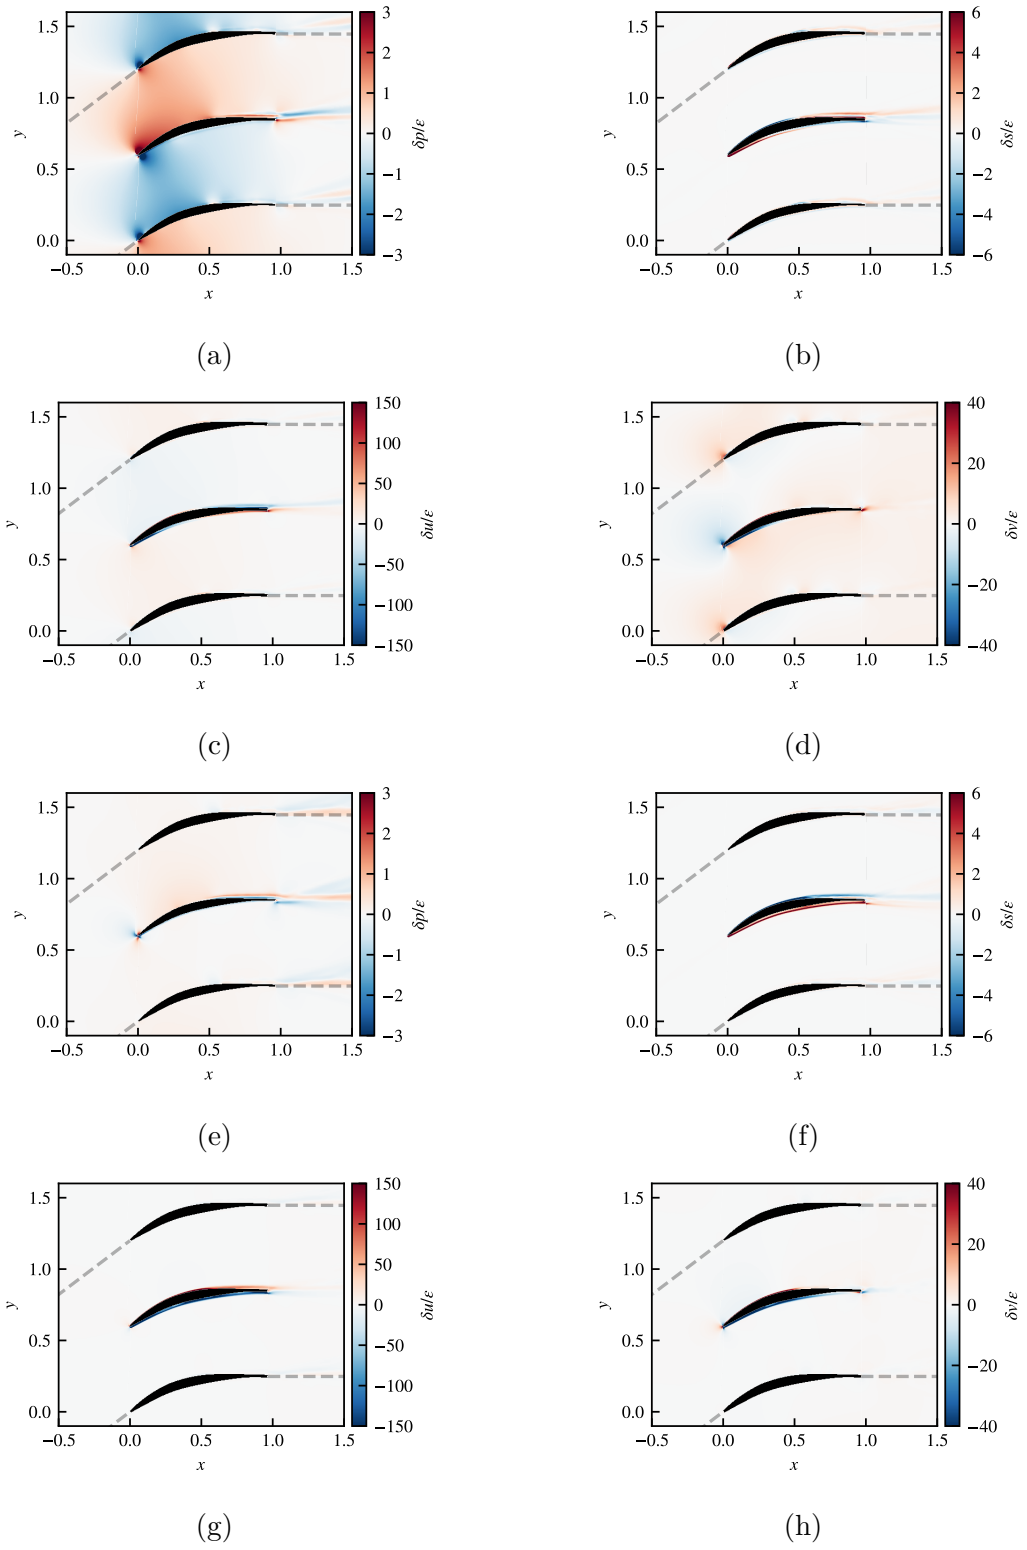
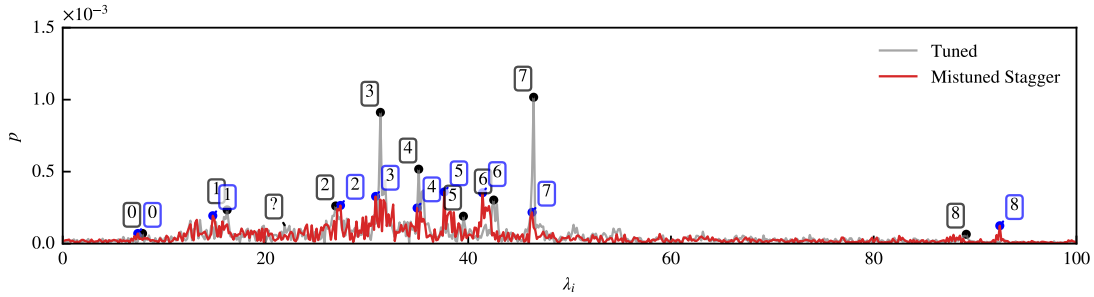
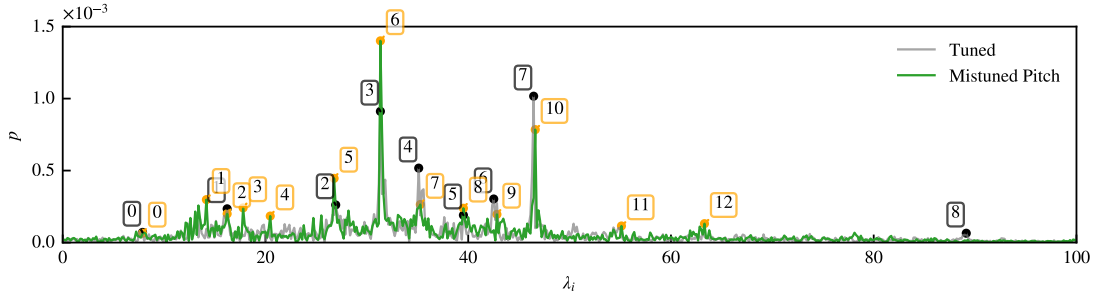


Figure C.2: Modifications to the mean flow as a result of perturbations to the stagger angle (in (a)-(d)) and the pitch (in (e)-(h)). The fields are scaled with the perturbation magnitude $\epsilon = 5 \times 10^{-3}$. Here the perturbed field is interpolated onto the tuned grid to make the comparison independent of changes to the grid metric.



(a)



(b)

Figure C.3: The Fast Fourier Transforms (FFTs) for the pressure field taken upstream of the unperturbed blade. Peaks are labelled in each case to aid the discussion, with one peak being labelled with '*?*' as the features for that peak only appear in the tuned assembly.

C.4 Acoustic Spectra

Following the procedure in Section 3.3.1, finely sampled datasets are obtained for the nonlinear fields for the three cases, with the FFT computed for the fields at all points within the domain. To guide us, we have illustrated the pressure spectra for non-zero frequencies in Figure C.3 for a point $\mathbf{x} = (-0.212, 0.151)$, which is shown with a red cross in the figure, upstream of the unperturbed blade, in order to estimate the response in the acoustic field. This position is chosen because we have identified in Section 3.3 and Section 3.4 that the dominant acoustic feedback mechanisms propagate the acoustic waves in a predominantly upstream direction. In this way we can identify the peaks that exhibit the largest changes, and compare a representative sample of the necessary spectral snapshots below (see Figure C.4 and Figure C.5 for the snapshots of the FFT at these peaks).

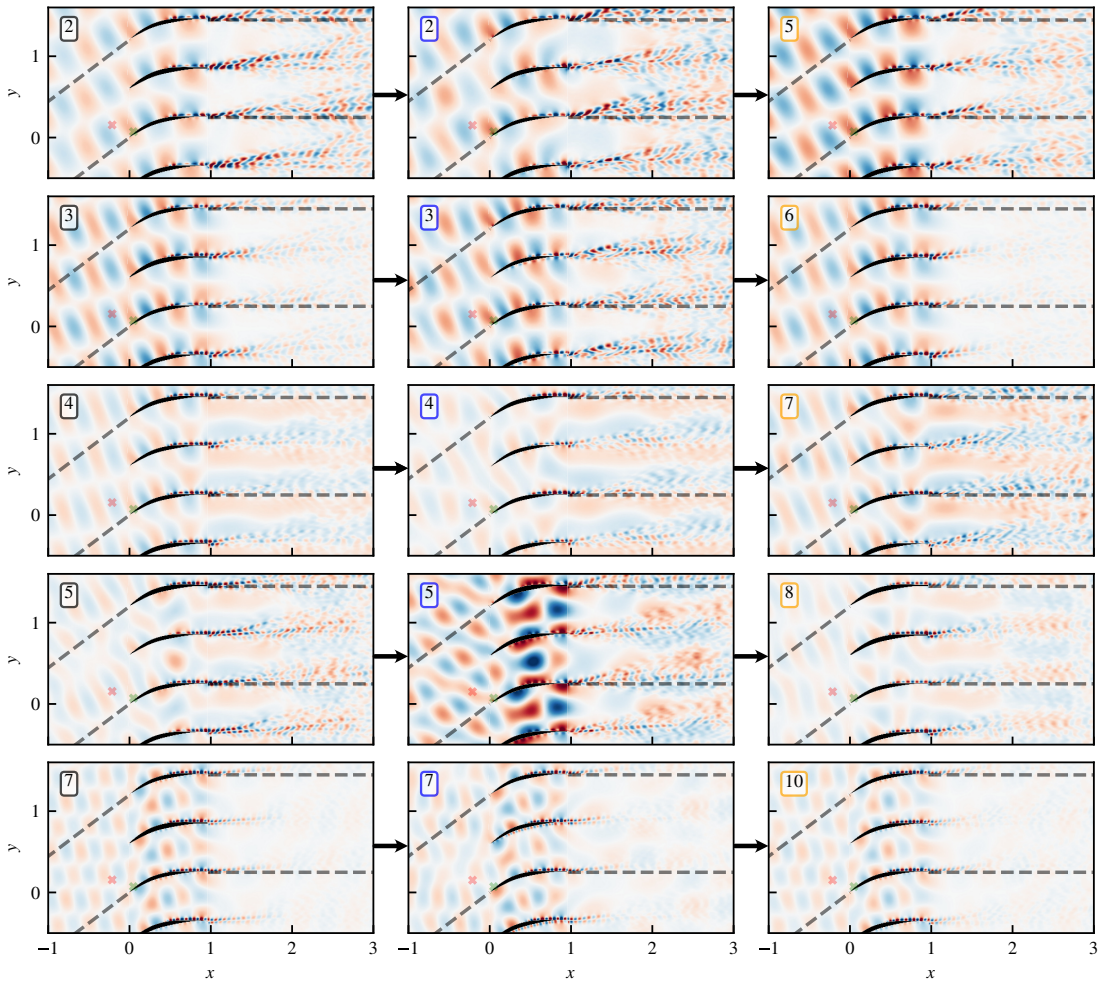


Figure C.4: The FFT pressure snapshots for the main peaks in the FFTs shown in Figure C.3. Arrows between plots indicate related peaks. The red cross shows the location of the pressure probe used to generate the FFT, and the green cross at the leading edge shows the location at which the phase is adjusted for each plot.

From Figure C.3a, it is clear to see that changes to the stagger angle see a dramatic reduction in the magnitude of the main tonal noise peaks ([3] → [3]) , ([4] → [4]) and ([7] → [7]), while variations in pitch strongly amplify ([3] → [6]) while reducing ([4] → [7]) and ([7] → [10]), although the reduction in the latter is not as dramatic as in the stagger angle case. Fields for selected modes that are mentioned in the discussion below are shown in Figure C.4 and Figure C.5.

The response of peak [7] is perhaps easiest to explain here since we have seen before in Section 3.3 that this feature is formed by a partially confined mode which creates a resonant standing-wave structure at the trailing edge with a wavelength

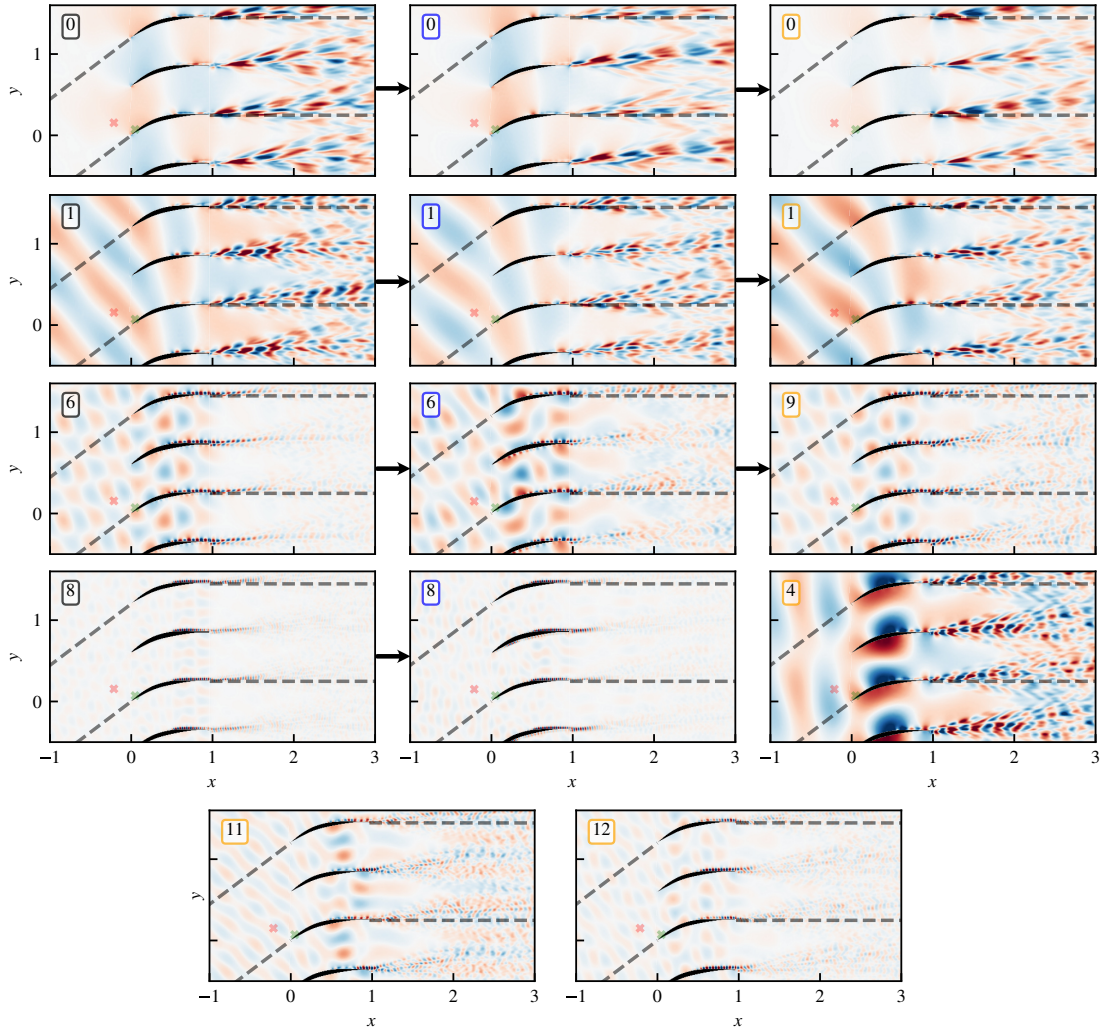


Figure C.5: Selected FFT pressure snapshots for selected remaining peaks, as shown in Figure C.3. Note that fields [4], [11] and [12] are new peaks for the pitch-wise-mistuned case (see Figure C.3b), and hence do not have complementary tuned and stagger-angle mistuned fields shown alongside.

matching the pitch of the blade. The orientation of the two blade surfaces in this region are responsible for reflecting the acoustic waves back towards the source confining the pressure waves at the trailing edge, and so while the pitch-wise mistuning disrupts the standing wave somewhat, the energy is still confined between the blades. By contrast, mistuning of the stagger angle changes the orientation of the blade surface and allows a portion of the acoustic energy to escape into the downstream, as is clearly demonstrated in Figure C.4.

While stagger angle modifications are responsible for suppressing some resonant

structures, others, such as [5] → [5], are seen to grow instead. Looking closely at this case in Figure C.4 demonstrates mode *localisation* effects in the upper passage of the mistuned snapshot. The change in blade surface orientation allows for standing-wave structures to become amplified between the perturbed suction side and the unperturbed blade, occurring between the trailing edges and between the separation bubble and the pressure surface. By contrast, the tuned cascade is dominated mostly by hydrodynamic features and direct propagation of acoustics from the trailing-edge scattering. While the propensity of mistuned structures to exhibit mode localisation and eigenvalue *veering* (see introduction and specifically [Hodges, 1982; Wei and Pierre, 1988a,b]) is well-established in the literature in the context of structural mechanics, to the best of our knowledge the application of this theory to aeroacoustics problems remains limited. This is likely due to the extreme computational cost of multi-element aeroacoustics simulations, although asymptotic perturbation techniques introduced in the structural vibration community can be leveraged, as we shall see in Appendix C.6.

The suppression of peak [4] on the other hand occurs through a significant reduction in suction-side shedding locally on the perturbed blade, indicative of a reduction in suction-side receptivity further upstream on that blade, which perhaps is not surprising given the local changes in the flow velocity in the vicinity of the blade in both cases. Reducing the vortex shedding here has the effect of reducing the corresponding acoustic source from the scattering at the trailing edge, and even destroying the in-phase relationship of the acoustic sources as in the case of the stagger-angle mistuning.

Dynamics in peak [3], however, provide a contrast to the other modes discussed above, since ([3] → [6]) exhibits amplification while ([3] → [3]) demonstrates decay. From plots of the relevant snapshots in Figure C.4, the main difference appears to be a de-phasing of the acoustic source and the presence of wave structures on the pressure side on the blade with the altered stagger angle. Energy transfer into

the acoustic field similarly shows the reduced efficiency of the acoustic source, as demonstrated by the more visible wave structures in the downstream. In contrast, [6] is visually very similar to [3], but the acoustic source is more efficient, with a greater transfer of energy from the hydrodynamics to the source. Although we will leave the detailed sensitivity analysis for the following section, our observations in Section 3.4.4 has shown that the system is very sensitive to small changes in the phase of the acoustic field in the wavemaker region, and the mechanism we suggest is responsible for the higher growth rates of the five- and ten-periodic structures associated with these frequencies. Given that the two-periodic blade row can only support in-phase dynamics of this particular acoustic source, it appears that altering the passage height produces the same phase-change effect which enables greater receptivity in the boundary layer. Similar conclusions are also applicable to ([2] $\rightarrow$ [5]).

To conclude, we briefly summarise the remaining peaks labelled in Figure C.3. Peaks [0] and [1] are two-periodic structures associated with low-frequency dynamics in the separation bubble and in the wake, with peak [1] producing more in the way of an upstream acoustic field. Much like peak [5], [6] represents a resonant structure between the pressure side of blade 2 and the suction side of blade 1 in the vicinity of the separation bubble, with the frequencies and peak amplitudes changing through much the same mechanisms as [5]. Peaks [8] and [8] are related to high-frequency shear structures on the suction side, with [8] being perturbed as a result of the mistuning and featuring amplified structures on the pressure side of the perturbed blade 2, as expected with local changes in the incident flow angle.

Unlike the case with mistuned stagger angle, the pitch-wise perturbation introduces a number of additional modes not present in the other two cases. Peaks [1] and [3] are structurally similar to [1], but represent stronger acoustic waves in the upstream. [4] is an entirely new structure with the geometry supporting wavelengths on the order of two times the chord length on the pressure surfaces, with the suction surface supporting structures with wavelengths on the order of the blade chord. Fi-

nally, [11] and [12] are acoustic resonance structures supported by locally intensified shedding from the separation bubble on blade 2.

All the modes mentioned here are shown in Figure C.5.

C.5 Global Modes and Sensitivities

Unlike our global mode analyses in Chapter 3, we will restrict our attention here to only the four most unstable modes so that a smaller Krylov subspace, with 1024 column vectors, can be used. This is critical since larger subspaces would lead to longer calculation times, and the two-passage calculations would otherwise result in significant usage of the limited computational resources. These modes are shown in Figure C.6. The slight variations in the two-passage mean flow, as compared to the mean flows computed earlier in Section 3.3, result in elevated growth rates, with reduced frequencies for the mis-staggered case and elevated frequencies for the pitch-wise mistuning, when compared to the tuned system. We have also computed modes for the two-passage tuned case once again, so that the comparison is more accurate. Ultimately our goal is to establish the changes to the structural sensitivity of these modes. Once again the frequencies of the global modes are generally slightly under-predicted, which is to be expected when comparing a linear to a nonlinear simulation, although a detailed analysis of this is left for future work. We will also limit the commentary on the changes to the growth rates of the modes, as despite predicting a trend for higher growth in the pitch-mistuned and decay in for the stagger-angle cases, we note that modes such as ([2] $\rightarrow$ [2]), ([4] $\rightarrow$ [4]) and ([4] $\rightarrow$ [7]) do not adequately reflect the changes in growth rates that are observed from the FFT of the nonlinear field. On the other hand, by visualising the FFT in Figure C.3 at a fixed observation point, we have already seen in the section above how changes to the geometry of the problem can alter the characteristics of the acoustic sources, and so if the linear mode shows a more efficient energy extraction from the mean flow in some portion of the domain – with a correspondingly higher

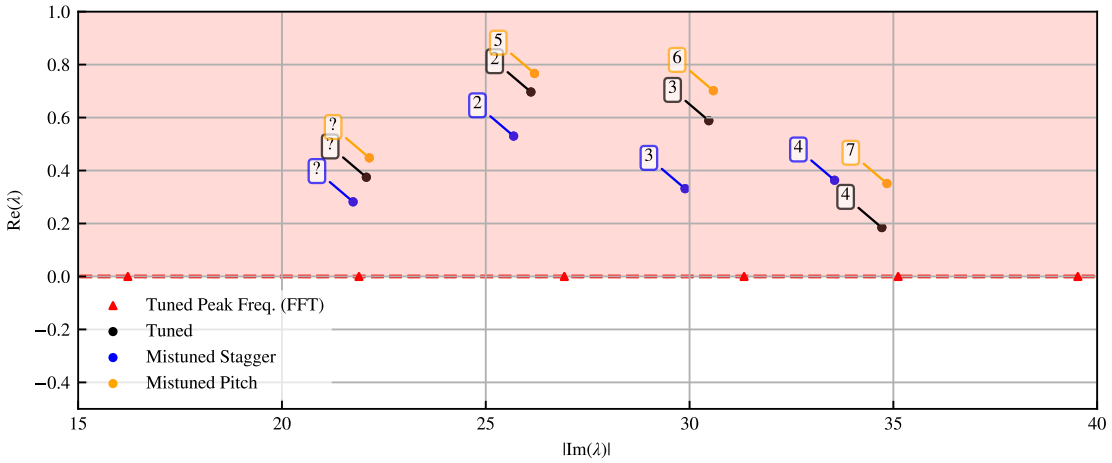


Figure C.6: Global modes for the tuned and mistuned cases. The modes are labelled according to the corresponding frequency peak in the FFT plot in Figure C.3. Red triangles denote the locations of the peaks in the FFT frequencies in the tuned case. The modes labelled '*?*' correspond to the peak labelled '[*?*]' in Figure C.3.

linear growth rate – there is no requirement for the acoustic energy to propagate in the same direction and attain similar intensities at that particular observation point[2]. We can see also from Section 3.3.1 and Section 3.3.4, and indeed from Figure C.6, that the most unstable mode is not necessarily the one responsible for the largest peak in the acoustic spectrum. To understand the reasons for this, a nonlinear stability analysis is required, but this is left as an exercise for the complete paper in the future.

Instead, we will concern ourselves with the structures of these modes as these match well with the FFT snapshots shown earlier. The direct modes are again normalised with respect to the maximum axial velocity in the near wake, using the normalisation region as shown in Figure B.1b. Pressure snapshots of the normalised modes are shown in Figure C.7. From a quick visual comparison, the first two columns of Figure C.7 show that the stagger angle mistuning leads to the largest structural changes to the acoustic field, with the one-periodic symmetry of the modes[3] broken by the change. In all cases except [4], pressure fluctuations on the

[2]In other words, the observation point of the FFT pressure probe may be located in a 'weaker' part of the structure.

[3]We established in Section 3.4 that the two-periodic case does not add any modes that are more unstable than the one-periodic case.

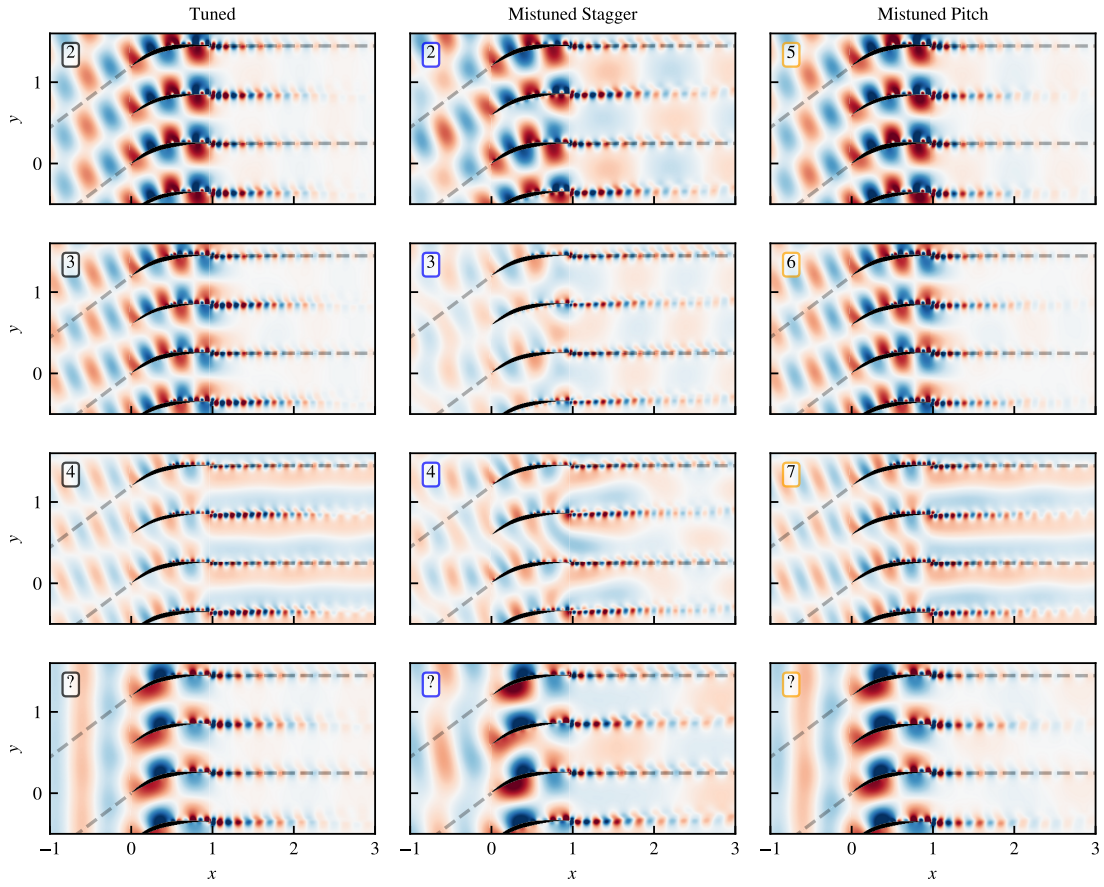


Figure C.7: Pressure snapshots of the leading direct modes for the tuned, mistuned stagger and mistuned pitch cases.

pressure side of blade 2 are diminished, and with the 'node' between the suction side of blade 1 and the pressure side of blade 2 disappearing. This change in pressure distribution affects the pattern of acoustic radiation from the blade row with modes such as modes [2] and [?], showing intensified acoustic fields upstream of the perturbed suction surface of blade 2, and with the unperturbed blade 1 producing reduced pressure fluctuations. On the other hand, modes [3] and [4] instead appear to have suppressed acoustic waves from blade 2 while making up for the deficit on blade 1. The mechanism by which this occurs requires further study since the process of the superposition of the waves at the trailing edges is more complex than in the tuned case, as for this case the acoustic sources from the scattering are in-phase and so the pressure disturbance is confined to upstream propagation along the blade

surface, but in the mistuned scenario the in-phase relationship breaks down leading to acoustic coupling in the blades.

The stagger angle mistuning in particular, and to a lesser extent the pitch-wise mistuning, introduce additional 'leakage' of large-wavelength downstream propagating waves from the blade row, with the two passages in anti-phase. This is seen as a chequerboard pressure pattern downstream of the blades. While the unsteady convective features at the trailing edge are now in anti-phase, it is interesting to see that near the separation point the dynamics are in-phase but towards to trailing edge this relationship is reversed.

Pitch-wise mistuning, by contrast retains the single-passage periodicity of the tuned system for these modes, based on the evidence of Figure C.7. Surface pressure plots comparing the tuned and pitch-wise-mistuned pressure distributions are shown in Figure C.8. These plots show that the pressure distribution around the blades for all modes is higher on the suction side than in the tuned case, leading to greater acoustic forcing on the unstable boundary layer. Another aspect worthy of further exploration is whether there exists a link between the mistuning of the blade pitch and the increased growth rates seen in the five- and ten-periodic modes. Recalling the discussion in Section 3.4.4, we hypothesised that the increased growth could be the result of a more optimal phase relationship in the feedback loop sustaining the mode, and so what needs to be investigated at a later time is whether a similar mechanism operates here.

Limitations to supercomputer time mean that only the tuned and stagger-angle mistuned cases are available with the corresponding adjoint modes. The normalisation $\langle \mathbf{w}, \mathbf{v} \rangle_{\mathbf{M}} = 1$ is applied, and the direct and wavemaker surface traces are depicted in Figure C.9. As expected, in the tuned cascade, pressure and sensitivity distributions of both modes largely overlap on both blades, with the blades exhibiting the same dynamics and the same receptivity in both cases, as was expected based on our reasoning in Section 3.4.4, which stated that the wavemakers are the same in all

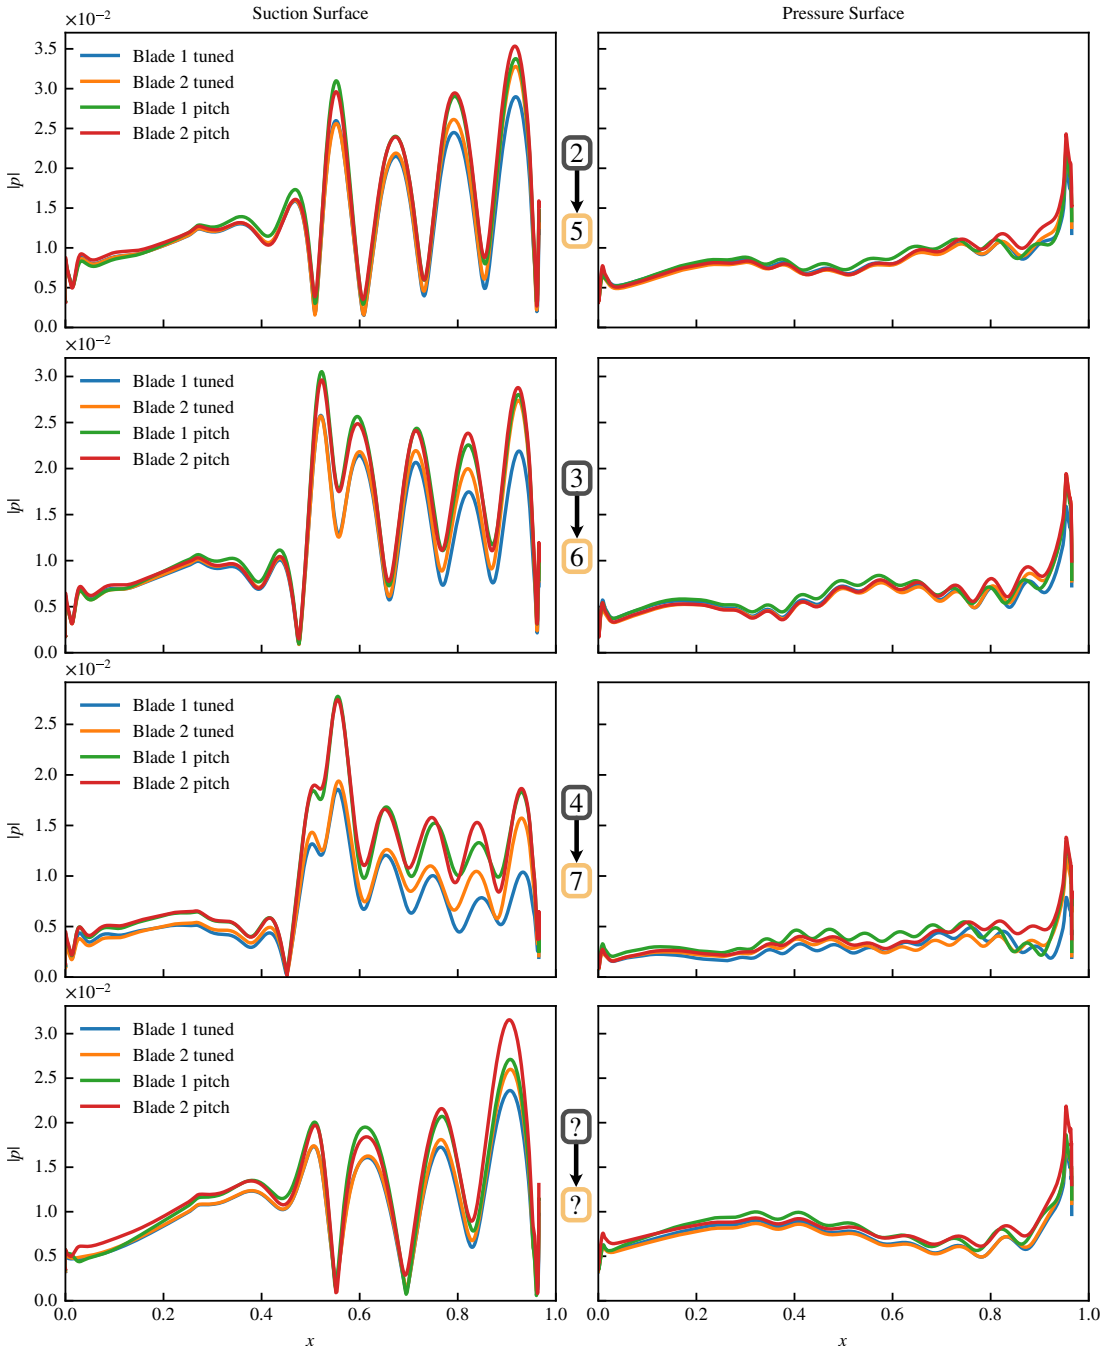
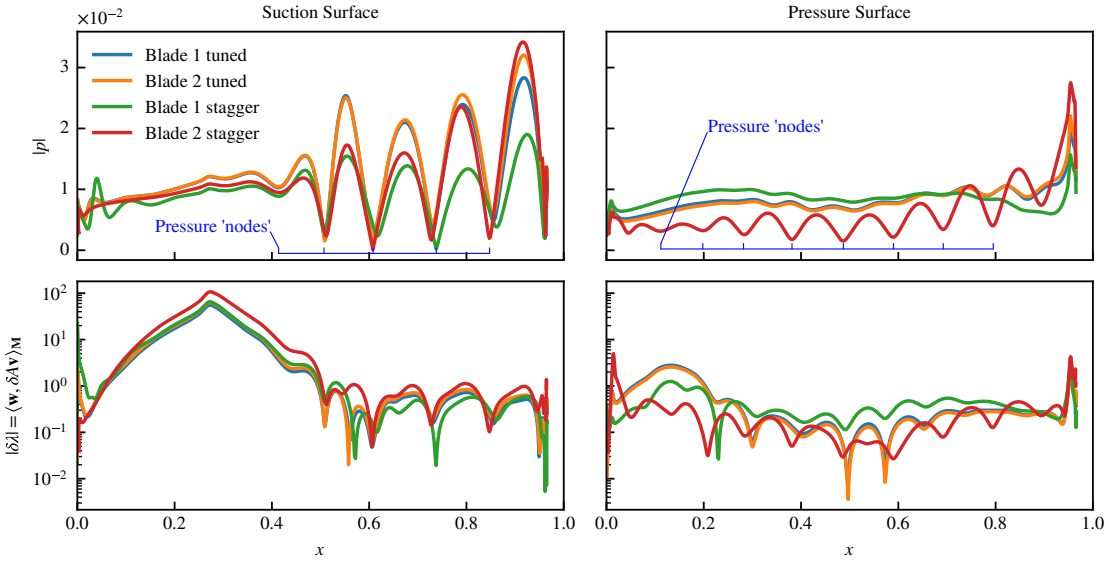
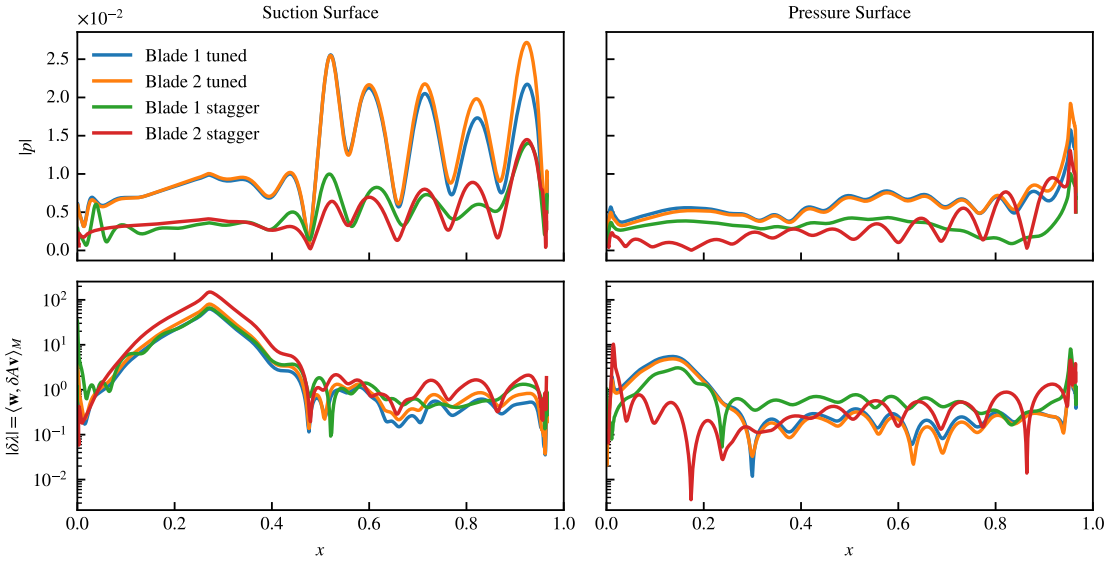


Figure C.8: Surface pressure for the tuned and pitch-mistuned direct modes. The modes are labelled between the surface- and pressure-side plots.



(a) Mode ([2] and [2]). Pressure 'nodes' are labelled here for the pressure field.



(b) Mode ([3] and [3])

Figure C.9: Surface pressure and the wavemaker illustrated for each blade of the tuned and mis-staggered cases for the two leading modes.

passages. To recap, these modes are characterised by the amplification of convective boundary layer instabilities by the suction-side bubble and an acoustic feedback mechanism that involves both the pressure and suction surfaces, as evidenced by the wavemakers.

Mistuning breaks this symmetry with both modes, [2] and [3], attaining lower amplitudes in the suction-side separation region, with the greatest changes seen in mode [3], where surface pressure amplitude upstream of the separation is up to 63.4% lower than the corresponding tuned mode. As the mode is normalised in each case by the maximum near-wake axial velocity, it follows then that mode [3] in particular is less efficient at scattering the convective disturbance on both blades at the respective trailing edges, because, for a given wake strength, the pressure amplitude is lower. Curiously, the wavemakers indicate that, despite the lower surface pressure, the suction surface of the perturbed blade, is significantly more sensitive than in the tuned scenario with the wavemakers in Figure C.9a and Figure C.9b indicating a larger peak receptivity on blade 2. Blade 1 on the other hand has a similar wavemaker to the tuned case albeit with a slight phase difference in the nodes within the superposition region between $x = 0.5$ and $x = 1.0$, which is possibly a consequence of the lower frequencies seen in Figure C.6.

Pressure-side dynamics tell a different story, however. In both cases the wavemaker for the perturbed blade 2 is approximately one order of magnitude lower in the interval $x = [0.1, 0.2]$ than the corresponding tuned case, and we have seen from our earlier analysis and visualisations in Figure 3.21 that receptivity in this region is key to sustaining the mode. Instead the perturbed blade sees higher sensitivities at the leading edge, where we have observed additional separation in the mean flow. This change in the wavemaker is not just a decrease in the sensitivity, however, since the direct modes for blade 2 is exposed to much lower pressure fluctuations for both modes, and evidence of additional vortex shedding (with more pressure nodes in the pressure trace) from the separation region at the leading edge is clear to see.

C.6 Perspectives for Mistuning Analysis in Quasi-N-Periodic Systems

An obvious question that arises is how one should address the large N-periodic arrays encountered in Section 3.4 now that the two key assumptions – the cyclic symmetry of the mesh and of the flow – are violated. This is a significant problem, even if the perturbations occur only in $\mathbf{A}_0$, as now the expression $\hat{\mathbf{A}} = \mathbf{P}^H\mathbf{A}\mathbf{P}$ is not diagonal. This is because the penultimate step in computing the diagonal matrix involves the sum $\sum_k \rho_j^k \mathbf{K}_j$ in the off-diagonal blocks, which in the tuned case has all the matrices $\mathbf{K}_i = \mathbf{K}_j$ for all $i, j \in [0, 1, \ldots, N-1]$, and since ρ_j is a root of unity, it follows that $\rho_j^N - 1 = 0$ and $1 + \rho_j + \rho_j^2 + \cdots + \rho_j^{N-1} = 0$. Now that the former no longer holds, off-diagonal blocks are to be expected, and though these may be small, they are precisely the terms that encode the information about mistuning that we seek, and so cannot be ignored. This presents us with a challenge since the perturbed system matrix once again exhibits global coupling and requires the solution of the full matrix system.

Our methodology for this analysis depends very much upon the magnitude of the perturbation to the surface geometry and critically *the degree to which this perturbation influences the mean flow.* Indeed, in the section above, we have seen that even small changes to the stagger angle, that is changes of 1° or equivalently $\mathcal{O}(10^{-3})$ of vertical displacement, result in significant changes to the pressure-side sensitivities. These nonlinear changes to the entire structure of the problem mean that the most realistic option becomes the direct employment of model reduction techniques such as POD or sparsity-promoting DMD (DMDSP) [Jovanović et al., 2014] applied to time-domain snapshots of the nonlinear flow field. In this instance the nonlinear solution is highly parallel and is constrained only by the IO time taken to save the required snapshots for the DMD, and so does not suffer the loss in performance seen in the calculation of the eigenvalues in large systems. N-periodic solutions may still be helpful in this case if the perturbations are not too large, as

the temporal frequencies of the modes are likely to remain similar. This makes the interpretation of the DMD modes significantly easier.

The disadvantage of the dimensionality reduction approach, however, is that we do not have access to the sensitivity information afforded to us by the adjoint modes, which is key to establishing the feedback mechanisms and causality characterising the system. On the other hand, if the perturbation is small and the resulting response of the mean flow is linear, then we can make use of the linear sensitivity derived in Equation (3.24) for the wavemaker. In that earlier section we considered $\delta\mathbf{A}$s that were defined as spatial delta functions, with which we introduced the wavemaker and by extension the structural sensitivity of the flow. This form of $\delta\mathbf{A}$ is not a general restriction. By calculating the mean flows for the tuned and mistuned cases and linearising, we can construct, respectively, $\mathbf{A}$ and $\tilde{\mathbf{A}} = \mathbf{A}+\delta\mathbf{A}$, with $\mathbf{A}$ denoting the unperturbed state. Further exploiting the linearity of the problem we obtain

$$\delta\mathbf{A}\mathbf{v} = (\mathbf{A} + \delta\mathbf{A})\mathbf{v} - \mathbf{A}\mathbf{v}, \tag{C.1}$$

and so the product $\delta\mathbf{A}\mathbf{v}$ can still be computed in a matrix-free way. In this way

$$\delta\lambda = \frac{\langle \mathbf{w}, \delta\mathbf{A}\mathbf{v}\rangle_{\mathbf{M}}}{\langle \mathbf{w}, \mathbf{v}\rangle_{\mathbf{M}}}, \tag{C.2}$$

precisely defines the change in the eigenvalue as a result of the mistuning.

Appendix D

Numerical Stencils

k	-2	-1	0	1	2
$a_{1,1+k}$			1	-3.0607612480196	-9.2825687875052
$a_{2,2+k}$		-1.3841732050877193	1	6.620636765789473684	-2.12293497649122807
$a_{j,j+k}$	0.027844800835977185	0.45586882948901069	1	0.477479190809459099	0.038825537571973596
$a_{n-1,n-1+k}$	15.74839510877193	27.04931558947368	1	-0.3211204877192982	
$a_{n,n+k}$	4.074311507053876	6.583618303429768	1		

Table D.1: Coefficients for the CULD (compact upwind with low dissipation) scheme.

D.1 Compact difference schemes

k	-2	-1	0	1	2
$b_{1,1+k}$	-2.0916904202872	12.7390135650198	-7.5911418720294	-3.324665234327	0.2684839616238
$b_{2,2+k}$	3.3623251557894737	-11.84525148578947	8.83692208	0.29908109	-0.65307684
$b_{j,j+k}$	-0.118727423955	-0.7482059425881687	0.00052666822731662	0.7184558284234524	0.1479508698923997
$b_{n-1,n-1+k}$	-1.8466026	-30.228379	21.1592312	12.27649478947368	-1.36074438947368
$b_{n,n+k}$	0.040891100302009	-0.7577318529877	-6.875427455144652	4.202556257560724	3.389711950269619

Table D.2: Coefficients for the CULD (compact upwind with low dissipation) scheme (Continued).

k	-1	0	1
$a_{1,1+k}$		2	4
$a_{2,2+k}$	1	4	1
$a_{i,i+k}$	1	3	1
$a_{n-1,n+k-1}$	1	4	1
$a_{n,n+k}$	4	2	

Table D.3: Coefficients for the compact 3rd-order centred scheme.

k	-2	-1	0	1	2
$b_{1,1+k}$			-5	4	1
$b_{2,2+k}$		-3	0	3	
$b_{i,i+k}$	-1/12	-7/3	0	7/3	1/12
$b_{n-1,n+k-1}$		-3	0	3	
$b_{n,n+k}$	-1	-4	5		

Table D.4: Coefficients for the compact 3rd-order centred scheme (Continued).

0					
c_2	a_{21}				
c_3	b_1	a_{32}			
c_4	0	b_2	a_{43}		
c_5	0	0	b_3	a_{54}	
	b_1	b_2	b_3	b_4	b_5

Table D.5: Butcher tableau for the Runge–Kutta scheme `RK4(3)5[2R+]C`.

a_{21}	970286171893/4311952581923
a_{32}	6584761158862/12103376702013
a_{43}	2251764453980/15575788980749
a_{54}	26877169314380/34165994151039
b_1	1153189308089/22510343858157
b_2	1772645290293/4653164025191
b_3	−1672844663538/4480602732383
b_4	2114624349019/3568978502595
b_5	5198255086312/14908931495163

Table D.6: Runge-Kutta coefficients for Table D.5.