

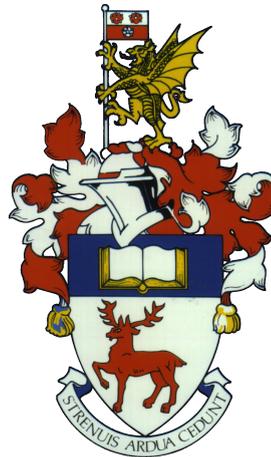
UNIVERSITY OF SOUTHAMPTON

FACULTY OF ENGINEERING AND THE ENVIRONMENT

**Development and Application of an Adjoint-based Optimal Flow Control
Framework for Compressible Direct Numerical Simulations**

by

Javier Otero



Thesis for the degree of Doctor of Philosophy

May 2017

UNIVERSITY OF SOUTHAMPTON

ABSTRACT

FACULTY OF ENGINEERING AND THE ENVIRONMENT

Doctor of Philosophy

DEVELOPMENT AND APPLICATION OF AN ADJOINT-BASED OPTIMAL FLOW
CONTROL FRAMEWORK FOR COMPRESSIBLE DIRECT NUMERICAL
SIMULATIONS

by [Javier Otero](#)

The advances in the understanding of compressible turbulent flows challenge researchers to tackle more ambitious problems that surpass human intuition. Hence, it becomes necessary to rely on methods that give an insight into the underlying physical mechanisms that govern these turbulent flows to perform non-trivial tasks such as flow control. This thesis describes the development and application of an adjoint-based optimal flow control framework for compressible flows, which has been appended to an existing in-house DNS code (HiPSTAR). With minor coding effort, this framework is extended to permit the computation of both (stable and unstable) exact steady and periodic flow solutions in compressible flows over complex geometries. In this work we present for the first time a family of exact periodic solutions in compressible flows ($Re = 2000$), which we study in detail, paying special attention to the flow-acoustic interaction characteristic from cavity flows. Furthermore, a family of equilibrium flow solutions associated to these periodic orbits is also introduced, where it is shown that both families meet at the quasi-incompressible Mach number range. A stability analysis of both periodic and equilibrium solutions shows that the bifurcation is of subcritical Hopf type. Moreover, this flow-acoustic interaction is optimally controlled in a cavity flow of a higher Reynolds number ($Re = 5000$) by using the adjoint-based optimal flow control framework. The target of the flow actuation consists in reducing noise levels at the sensor location, by either reducing the overall sound radiation or altering the sound directivity. Lastly, we also carry out adjoint-based optimal flow control on a three-dimensional backward-facing step flow ($Re = 3000$). The investigation is oriented to cabin noise reduction, where we optimise an upstream flow actuation to reduce the flow impinging scales which present the largest impact on cabin noise. In addition, the complexity present in the flow studied pushes the adjoint-based approach at its limits. These intrinsic limitations of the method are discussed, where we also suggest and demonstrate a successful alternative which permits the use of adjoint-based optimisation methods in separated flows.

Declaration of Authorship

I, [Javier Otero](#), declare that the thesis entitled *Development and Application of an Adjoint-based Optimal Flow Control Framework for Compressible Direct Numerical Simulations* and the work presented in the thesis are both my own, and have been generated by me as the result of my own original research. I confirm that:

- this work was done wholly or mainly while in candidature for a research degree at this University;
- where any part of this thesis has previously been submitted for a degree or any other qualification at this University or any other institution, this has been clearly stated;
- where I have consulted the published work of others, this is always clearly attributed;
- where I have quoted from the work of others, the source is always given. With the exception of such quotations, this thesis is entirely my own work;
- I have acknowledged all main sources of help;
- where the thesis is based on work done by myself jointly with others, I have made clear exactly what was done by others and what I have contributed myself;
- parts of this work have been published as:
 - Javier Otero, Richard D. Sandberg, and Ati S. Sharma. *Direct Numerical Simulations for Adjoint-based Optimal Flow and Noise Control of a Backward-Facing Step*, 22nd AIAA/CEAS Aeroacoustics Conference, Aeroacoustics Conferences, (AIAA 2016-2889)

Signed:

Date:

Acknowledgements

I would like to thank my supervisors Dr Ati Sharma and Prof Richard Sandberg for their valuable input through these years, and also the long lasting meetings were I was letting out all my frustration.

Also, special thanks to my officemates and rest of the research group members, who were always willing to give a hand when needed. This thesis wouldn't be possible without all the fruitful chats during the numerous coffee/water breaks. Keep it like that guys, you make this university such a spectacular place to work.

I am very grateful to my parents, who were always remarkably supportive. Although it might have been a bit difficult to understand to them, they were always curious to know what I was doing at work. As well, thank you to the rest of my family. We all form an amazing team taking care of each other.

I am extremely thankful to my girlfriend, Claudia, for listening to all the troubles that I had throughout this PhD. Thank you for looking always on the bright side of things and being so positive.

A very special thanks goes to my WR52 housemates Patrick, Andrea, Fran, Jayanth, Markus and Vincenzo (and the rat). We had such a great time in that house, in particular, making BBQs as many times as possible regardless of the weather conditions.

I am also grateful to the guys from the 'chino' group. I really miss those very very very long conversations after lunch in the 'soup place'.

Finally, many thanks to everyone who joined me for a run or bike ride under the never rainy, never cold and never windy Southampton weather. Especially, to Jacques and Jake, with whom I shared most of my extreme bike rides, for waiting for me every time I had a puncture.

Contents

Declaration of Authorship	v
Acknowledgements	vii
List of Figures	xiii
List of Tables	xxi
Nomenclature	xxiii
1 Introduction	1
1.1 Motivation	9
1.2 Contributions	11
1.3 Outline of the Report	12
2 Background Theory	13
2.1 Flow Control and Optimisation	13
2.1.1 Adjoint Method	14
2.1.1.1 Adjoint-based Gradient	15
2.1.1.2 Limitations of Adjoint-based Optimisation: Model Predictive Control	16
2.2 Flow-field Decomposition and Reduced Order Modelling	18
2.2.1 Spectral Decomposition	18
2.2.2 Proper Orthogonal Decomposition	19
2.2.3 Dynamic Mode Decomposition	21
2.3 Stability Analysis	23
2.4 Summary	27
3 Governing Equations	29
3.1 Navier-Stokes Equations	29
3.2 Adjoint Equations	30
3.2.1 Adjoint Euler Equations	31
3.2.2 Adjoint Navier-Stokes Equations	32
3.3 Adjoint Equations in Characteristic Form	33
3.3.1 Characteristic Equations – ξ	33
3.3.2 Characteristic Equations – η	36
3.4 Summary	37

4	Computational Setup	39
4.1	Direct Numerical Simulations	39
4.1.1	Numerical Method	39
4.1.2	Grid	40
4.1.3	Initial and Boundary Conditions	40
4.1.3.1	Isothermal no-slip wall	41
4.1.3.2	Free-stream Boundary Condition	41
4.1.3.3	Outflow Boundary Condition	41
4.1.3.4	Laminar Inflow Boundary Condition	42
4.2	Adjoint Simulations	42
4.2.1	Storage Saving Scheme	42
4.2.2	Control Update	43
4.2.3	Boundary Conditions	44
4.2.3.1	Terminal Condition	44
4.2.3.2	Isothermal No-Slip Wall	45
4.2.3.3	Non-Reflecting Characteristic Conditions	45
4.2.3.4	Soft Adjoint Inflow/Outflow Characteristic Conditions	45
4.2.4	Characteristic Interface Conditions	46
4.3	Summary	47
5	Verification	49
5.1	Green-Lagrange Identity	49
5.1.1	Multi-block Periodic Case	50
5.1.2	Laminar Free-stream Case	51
5.2	Finite Difference Gradient	52
5.2.1	Free-stream Noise Cancelling Case	52
5.2.2	Laminar Boundary Layer Noise Cancelling Case	54
5.3	Summary	55
6	Periodic Orbits & Equilibria in a 2D Cavity Flow	57
6.1	Introduction	57
6.2	Numerical Method	59
6.2.1	Steady Solutions	60
6.2.2	Periodic Solutions	61
6.3	Periodic Solution at $M=0.5$	62
6.4	Family of Periodic Solutions across Mach Number	70
6.5	Other Periodic Solutions and Stability	83
6.6	Family of Equilibrium Solutions across Mach Number	87
6.7	Summary and Conclusions	91
7	Stability Analysis of 2D Cavity Flow Equilibria	93
7.1	Introduction	93
7.2	Global Stability Analysis at $M=0.5$	95
7.3	Stability Evolution across Mach Number	102
7.4	Summary and Conclusions	108
8	Noise Radiation and Directivity Control of a 2D Cavity Flow	111
8.1	Introduction	111

8.2	Computational Setup	113
8.3	Optimal Control Results	115
8.4	Summary and Conclusions	126
9	Control of a 3D Backward-Facing Step Flow	127
9.1	Introduction	127
9.2	Computational Setup	128
9.3	Optimal Flow Control	132
9.4	Bypassing Adjoint-based Limitations: Gradient Normalisation	136
9.5	Summary and Conclusions	138
10	Summary and Conclusions	141
A	Derivation of the Adjoint Euler Equations	145
B	Adjoint Boundary Terms	147
B.1	Boundary Conditions	149
C	Exact Solutions - Simulation Parameters	151
D	Grid Independence of 2D Cavity Flows	153
E	Limitations of Adjoint-based Optimisation in Separated Flows	155
E.1	Introduction	155
E.2	Simulation Details	157
E.3	Sensing Region Influence on Adjoint Divergence	158
E.4	Adjoint-Field Decomposition	160
E.5	Sensitivity Analysis on the Large Scales	163
E.6	Summary and Conclusions	164
	References	167

List of Figures

2.1	Illustration of the receding horizon technique	17
2.2	Sketch of one of the flow control setups of the 2D cavity flow studied in chapter 8. Figures (a) and (b) represent the forward and adjoint simulations respectively. The flow actuation is depicted by the green box and the blue dashed circle represents the sensing region.	17
2.3	Illustration of a wavenumber-frequency spectrum for surface pressure of a subsonic separated flow. Φ_c and Φ_o represent the turbulent and acoustic components of the spectrum, respectively (figure reproduced from Bremner and Wilby, 2002)	19
5.1	Left and right figures show contours of ρ' and p^* respectively, both after 2 forcing periods. Contour ranges from $4.5 \cdot 10^{-4}$ to $-4.5 \cdot 10^{-4}$. The dashed lines represent the location of the block characteristic interfaces.	50
5.2	Spatial integration as a function of time for ρ' and p^* (a) , and $(\rho u)'$ and m_u^* (b) .	51
5.3	Spatial integration as a function of time for ρ' and p^* (a) , and $(\rho u)'$ and m_u^* (b) .	51
5.4	Illustration of the case setup. The red dashed circle represents the location of the noise source and controller, whereas the blue dashed circle represents the location of the sensing region. The contours show pressure fluctuations which range from $-0.5 \cdot 10^{-4}$ to $0.5 \cdot 10^{-4}$	53
5.5	Time dependent gradient of the cost functional obtained with a finite difference approach ($\nabla_\phi \mathcal{J}_{fd}$) and adjoint method ($\nabla_\phi \mathcal{J}_{adj}$).	54
5.6	Time dependent gradient of the cost functional obtained with a finite difference approach ($\nabla_\phi \mathcal{J}_{fd}$) and adjoint method ($\nabla_\phi \mathcal{J}_{adj}$).	55
6.1	Illustration of the time horizon T optimisation. After L-BFGS returns a new iterate, the algorithm updates the old value of T (indicated by the grid line at $t_0 + T$), by the new value (highlighted with the vertical red dashed line).	62
6.2	4D representation of the periodic orbit at $M = 0.5$. The trajectory is projected over the norm of dilatation ($\ \nabla \cdot \vec{u}\ $), vorticity ($\ \nabla \times \vec{u}\ $), dissipation rate ($\ \varepsilon\ $) and kinetic energy ($\ e_{kin}\ $). The red marks represent key physical phenomena occurring along the periodic orbit. For the sake of clarity, the symbols are labelled with letters from a to e in (d) . The grey line shows the natural flow evolution from the vicinity of the initial condition towards the periodic orbit.	63
6.3	Zoomed-in area at the orbit cross-point from figure 6.2b. Black and red lines represent the trajectories of the periodic orbit and fully developed flow, respectively.	64
6.4	Illustration of the trailing edge - shear layer dipole interaction.	65
6.5	Illustration of the leading edge - shear layer dipole interaction.	65

6.6	Contours of instantaneous density at the time steps represented by the red marks in figure 6.2. Subfigures (a) to (e) show the density field at the instants represented in figure 6.2d with the same letter. This is, (a) corresponds to the mark labelled with a , (b) with the mark labelled with b and so on. Density contours range from 0.96 to 1.04, where blue and red represent low and high density respectively.	67
6.7	Contours of instantaneous dilatation at the time steps represented by the red marks in figure 6.2. Contour plots are overlapped with the instantaneous streamlines. Subfigures (a) to (e) show the dilatation field at the instants represented in figure 6.2d with the same letter. This is, (a) corresponds to the mark labelled with a , (b) with the mark labelled with b and so on. Dilatation contours range from -0.01 to 0.01, where black and white represent flow expansion and compression respectively. This contour range has been chosen specifically to highlight the sound generation mechanisms.	68
6.8	Contours of instantaneous vorticity at the time steps represented by the red marks in figure 6.2. Contour plots are overlapped with the instantaneous streamlines. Subfigures (a) to (e) show the vorticity field at the instants represented in figure 6.2d with the same letter. This is, (a) corresponds to the mark labelled with a , (b) with the mark labelled with b and so on. Vorticity contours range from -15 to 15, where blue and red represent negative and positive vorticity values respectively. This contour range has been chosen specifically to highlight the shear layer.	69
6.9	Evolution of the period (T), and time-averaged norms of dilatation ($\overline{\ \nabla \cdot \vec{u}\ }$), kinetic energy ($\overline{\ e_{kin}\ }$), vorticity ($\overline{\ \nabla \times \vec{u}\ }$) and dissipation rate ($\overline{\ \varepsilon\ }$) as a function of Mach number. Red solid line in (a) shows the periods predicted by Rossiter's formula.	71
6.10	Evolution of Re_{Θ} as a function Mach number. The Re_{Θ} values are computed respect to the averaged flow quantities through one complete period at each period T	73
6.11	4D representation of the family of solutions across Mach number, ranging from 0.25 to 0.65 with increments of 0.05. As a reference, the periodic orbit that corresponds to $M = 0.5$ is coloured in red. The trajectories are projected over the norm of dilatation ($\ \nabla \cdot \vec{u}\ $), vorticity ($\ \nabla \times \vec{u}\ $), dissipation rate ($\ \varepsilon\ $) and kinetic energy ($\ e_{kin}\ $). The periodic orbits are labelled in (c) . The equilibrium solutions (see section 6.6) from Mach numbers 0.25 to 0.65 (ordered with increasing kinetic energy) are represented as red dots. Note how both families of solutions collapse in the quasi-incompressible regime at $M = 0.30$, which can be seen as a bifurcation point.	74
6.12	Instantaneous contours of vorticity at Mach numbers 0.30 (a) and 0.35 (b)	75
6.13	4D representation of the family of solutions across Mach number, ranging from 0.65 to 0.75. As a reference, the periodic orbit that corresponds to $M = 0.73$ is coloured in red. Similarly to figure 6.2, the trajectories are projected over the norm of dilatation ($\ \nabla \cdot \vec{u}\ $), vorticity ($\ \nabla \times \vec{u}\ $), dissipation rate ($\ \varepsilon\ $) and kinetic energy ($\ e_{kin}\ $). The periodic orbits are labelled in (a)	77
6.14	Snapshots of the dilatation field at the maximum value of the norm of dilatation for Mach numbers 0.65 and 0.74. Similarly to figure 6.7, dilatation contours range from -0.01 to 0.01, where black and white represent flow expansion and compression respectively.	78

6.15	4D representation of the family of solutions across Mach number, ranging from 0.75 to 0.8. As a reference, the periodic orbit that corresponds to $M = 0.78$ is coloured in red. Similarly to figure 6.2, the trajectories are projected over the norm of dilatation ($\ \nabla \cdot \vec{u}\ $), vorticity ($\ \nabla \times \vec{u}\ $), dissipation rate ($\ \varepsilon\ $) and kinetic energy ($\ e_{kin}\ $). The periodic orbits are labelled in (d)	79
6.16	Instantaneous contours of the dilatation field at Mach numbers 0.76 (a) , 0.78 (b) and 0.8 (c) . These instants correspond to the maximum norm of the vorticity field along the periodic orbit.	80
6.17	OSPL from 35 equally distributed monitor points along an arc of radius 5, each of them separated by 5 degrees. The centre of the arc is located at coordinates 1.5 and 1 in the streamwise and vertical directions, respectively.	82
6.18	Comparison of the continued and original orbits at Mach numbers 0.65 and 0.8. The continued orbits are represented with solid lines. The trajectories found from the developed flows are shown as dashed lines. The solutions corresponding to Mach numbers 0.65 and 0.8 are coloured with red and black, respectively.	84
6.19	Snapshot of velocity magnitude at the point of maximum dilatation for the continued (a) and original (b) orbits at $M = 0.8$	84
6.20	Evolution of both periodic orbits at $M = 0.8$ represented earlier in this section, following a large amplitude body force perturbation at the shear layer (black lines). Both solutions return back to their respective limit cycles (red lines) showing an stable character.	85
6.21	(a) Frequency spectrum of the perturbed flow trajectories at $M = 0.8$ and the exact (continued) solution also at $M = 0.8$. (b) Final stage of p-M080 decaying to M080.	86
6.22	Evolution of the norms of dilatation ($\ \nabla \cdot \vec{u}\ $), kinetic energy ($\ e_{kin}\ $), vorticity ($\ \nabla \times \vec{u}\ $), dissipation rate ($\ \varepsilon\ $) and Reynolds number based on the boundary layer's momentum thickness (Re_{Θ}) as a function of Mach number.	88
6.23	Contour plots of dilatation on the steady solutions at Mach numbers 0.25 (a) and 0.5 (b)	89
6.24	Bifurcation of the periodic and steady families of solutions across Mach number. The red and black marks represent the data points corresponding to the steady and periodic solutions, respectively. The black solid and red dashed lines show the expected trend of both families of solutions in between the points.	90
6.25	Transition from the steady solution to a fully developed flow at $M = 0.5$ (grey line). The steady or equilibrium solution is represented with a red dot, whereas the periodic solution at this Mach number is illustrated with a black line.	91
7.1	Steady-state cost function evolution as a function of non-dimensional time units. The red line represents the best fit of the linear flow behaviour during the transition.	96
7.2	Eigenvalue spectrum for the steady solution at $M = 0.50$. The red circular markers show the eigenvalues calculated through the adjoint simulations. The blue squared symbols represent the approximate eigenvalues obtained using the non-linear forward simulations. The smaller squared symbols show purely numerical modes (noise) product of the insufficiently large forward snapshot sequence, where their colour intensity gives a measure of each mode's amplitude.	97

List of Figures

7.3	Forward ((a), (c), (e), (g)) and adjoint ((b), (d), (f), (h)) modes associated with the leading eigenvalue for the equilibrium solution at $M = 0.5$. From top to bottom, the figures show the components of Q and Q^* . Contour plots are overlapped with the streamlines from the steady solution.	98
7.4	(a) Time signal of the streamwise momentum at $\vec{x} = (2.25, 1)$. (b) Estimated flow-acoustic feedback loop's frequency evolution across time. Red dashed and blue dashed horizontal lines indicate the frequencies associated the unstable mode and periodic limit cycle, respectively. (c) Frequency spectrum of the time signal show in (a)	100
7.5	Structural sensitivity maps of the streamwise (a) and vertical (b) momentum components, calculated as $ \Psi_{m_u^*}^* \cdot \Psi_{\rho u} $ and $ \Psi_{m_v^*}^* \cdot \Psi_{\rho v} $, respectively. Contour levels are normalised with $6 \cdot 10^{-6}$. Figures are overlapped with the streamlines from the steady solution.	102
7.6	Eigenspectrum of the family of equilibrium solutions across Mach number. The red dashed lines show the rise of the unstable branch at Mach numbers 0.3, 0.35 and 0.4. The blue dash-dotted lines represent the evolution of a decaying branch as the Mach number is increased. The thin black dashed line highlights the branch appearing due to noise in the stable steady states.	103
7.7	Illustration of the subcritical Hopf bifurcation of the periodic and steady families of solutions across Mach number. The solid lines represent stable flow states, whereas the dashed lines show unstable flow configurations. The red and black coloured lines indicate the periodic and steady families of solutions, respectively. Also with the same colour coding, the circle symbols show the actual flow solutions computed in this investigation. The hysteresis loop is depicted with the thick grey solid lines, where the arrows show the loop direction. For this illustration, the Mach number range where bistability occurs is represented with the red coloured horizontal axis.	104
7.8	Streamwise momentum component of the adjoint mode corresponding to the leading eigenvalue of the decaying branch (highlighted with blue dash-dotted lines in figure 7.6) at $M = 0.3$ (a) and $M = 0.35$ (b)	105
7.9	Evolution of the streamwise momentum receptivity of some of the eigenvalues from the unstable branch across Mach number. The first mode is shown in figures (a) , (c) and (e) at Mach numbers 0.30, 0.35 and 0.4, respectively. The fourth mode at $M = 0.3$ and the second mode at $M = 0.35$ are also shown in (b) and (d)	106
7.10	Streamwise (a) and vertical (b) adjoint momentum component of the unstable mode at $M = 0.65$	107
8.1	(a) Illustration of the 2D cavity flow. Contours show instantaneous z -vorticity. Red dashed lines represent the sound radiation from the trailing edge. (b) Phase portrait of the trajectory followed by the non-actuated cavity flow-field.	112
8.2	Root mean square of the magnitude of the adjoint momentum \vec{m}^* when using the downstream sensor. Red and blue contours indicate high and low values, respectively. The green dashed rectangle indicates the actuating sub-volume. The red-dashed circle represents the downstream sensing region.	114
8.3	Downstream sensor case. Evolution of the cost function for the five receding horizons.	116

8.4	Root mean square pressure (p_{rms}) contours for uncontrolled (a) and controlled cases. Ordered alphabetically, (b) , (c) and (d) represent the p_{rms} contours for the downstream, upstream and arc sensor cases, respectively. Locations of upstream and downstream regions are indicated with blue-dashed and red-dashed circles.	117
8.5	Power spectrum of pressure at the centre of the downstream (a) and upstream (b) monitor points. The black solid and blue dotted lines represent the non-actuated and arc cases. The red dashed lines show the results from the downstream case in (a) and upstream case in (b) . The green dash dotted vertical line indicates the sampling frequency of the flow sensors used for the adjoint-based optimisation.	118
8.6	Overall sound pressure level plots, with 0° and 180° as the upstream and downstream directions. The red solid line shows the non-actuated case. Dotted and dashed lines represent the downstream and upstream sensor cases, respectively. The dash-dotted line indicates the result of the arc sensor case. The OSPL levels are normalised with the maximum value of the non-actuated case.	120
8.7	Streamwise (left) and vertical (right) momentum components of the leading dynamic modes for each case. From top to bottom are the non-actuated, downstream, upstream and arc cases. Contour levels are normalised with 0.007.	122
8.8	Dominant actuation modes for the downstream case. Ordered from mode D1 to D4, (a) to (d) and (e) to (h) show the streamwise and vertical forcing components. Contour levels are normalised with 0.3741.	125
8.9	Dominant actuation modes for the arc case. Ordered from mode A1 to A4, (a) to (d) and (e) to (h) show the streamwise and vertical forcing components. Contour levels are normalised with 0.3741.	125
9.1	Illustration of the flow topology over a backwawrd-facing step. The turbulent structures (represented as isosurfaces of Q-criterion coloured with streamwise velocity) impinge on the target region (dark grey surface) causing the pressure fluctuations that would generate flow-induced vibrations on a structure at that same location.	129
9.2	Root mean square of the magnitude of the adjoint momentum \vec{m}^* . Red and blue contours indicate high and low values, respectively. The green rectangle indicates the actuating sub-volume. The sensing region is represented with a red line at the flow impingement area.	130
9.3	Illustration of the adjoint topology over a backwawrd-facing step with the sensing region located in the flow impingement area. The adjoint structures (represented as isosurfaces of ρ^* coloured with p^*) travel upstream from the sensing region (dark green surface) towards the flow actuation.	131
9.4	(a) Fourier transformed pressure signal at three different downstream locations of the step for the uncontrolled case. (b) Streamwise, vertical and spanwise components of the body-forcing vector $\vec{f}(\vec{x}_0, t)$ at an arbitrary location \vec{x}_0 within the actuation sub-domain. Note the pronounced increase in magnitude of the actuation parameters caused by the divergence of the adjoint field with backwards time marching.	132
9.5	Wavenumber spectrum of the pressure fluctuations gathered at the sensing region for the non-actuated (a) and (b) , and actuated scenarios (c) and (d) . Contour levels range from -200 to -40 dB in (a) and (c) , and from -100 to -40 dB in (b) and (d) . In both cases the total number of contours is 22.	133

List of Figures

9.6	Instantaneous snapshot of the density-field in the lower wall of the step for the non-actuated (a) and actuated (b) scenarios. Both density-fields were captured at $t = T_0 + 3Ta$. Sensing region is delimited by the two vertical green dashed lines.	134
9.7	Wave number frequency spectrum at $St = 2.5$. Figures (a) and (b) correspond to the non-actuated and actuated cases, respectively. Contour levels range from -122 to 102. The acoustic circle (black dotted circle) represents the propagating speed of sound in the $x - z$ plane.	135
9.8	Instantaneous contours of the flow actuation at $t = T_0$, showing the stream-wise (a) , vertical (b) and spanwise (c) forcing components. Frontal view in the upstream direction. Contour levels are normalised by 0.0027.	136
9.9	(a) Illustration of the sensitivity re-scaling method based on the norm of the adjoint field. Original and re-scaled norms are represented as black and green dashed lines, respectively. The red dash-dotted line shows the logarithmic line fit of the saturated norm of the adjoint field. (b) Streamwise, vertical and spanwise components of the preliminary body-forcing vector $\vec{f}(\vec{x}_0, t)$ when using the gradient normalisation method. Time signals correspond to an arbitrary location \vec{x}_0 within the actuation sub-domain.	137
D.1	Grid independence at $Re = 5000$ shown with four different Favre-averaged quantities at four different streamwise locations along the cavity. Subfigures (a) and (b) show profiles of the streamwise and vertical velocity components, respectively. On the other hand, subfigures (c) and (d) represent the σ_{11} and σ_{22} components of the Reynolds stress tensor σ . The magnitudes of σ_{11} and σ_{22} have been scaled up (x200) in their respective plots for clarity purposes.	154
E.1	Illustration of a forward (a) and adjoint (b) simulations in a flow control case. The green and orange boundaries represent the flow actuation location and the target region respectively.	156
E.2	Instantaneous contours of z -vorticity. Dashed lines represent the computational block interfaces.	158
E.3	Mean velocity field for the thin (a) medium (b) and thick (c) cases. The sensing region is represented with the orange line. Contours of velocity magnitude range from 0 to 1.1.	159
E.4	Density fluctuations at the centre of the sensing zone. The fluctuations are computed as $\rho' = (\rho - \rho_0)$. The simulation time is relative to the start of the data capturing shown.	160
E.5	Adjoint field norm for the three boundary layer cases. Note that the adjoint simulations march backwards in time.	160
E.6	DMD spectrum for the thin BL (a) and medium BL (b) cases. Each eigenvalue is coloured by its energy content by projecting the dynamic modes onto a POD basis, which also gives a measure of coherence (Schmid, 2010). The levels are independent for each case.	161
E.7	DMD spectrum for the thin BL (a) and medium BL (b) cases. Labels (1), (2) and (3) indicate the three leading unstable modes.	162
E.8	Dynamic mode shapes of the mode (1) (see figure E.7b) for m_u^* (a) and m_v^* (b) . The high level of activity in the near flow reattachment zone shows the origin of the divergence in the adjoint simulation.	162
E.9	Illustration of the propagation of the flow sensitivities for the adjoint simulations of the medium BL (a) and thin BL (b) cases.	163

- E.10 Comparison of the temporal evolution of the adjoint field norm for the medium BL case using different number of POD modes for the flow-field reconstruction. 164

List of Tables

7.1	Comparison across Mach number of the non-dimensional frequency associated with the periodic solutions and the first eigenvalue of the unstable branch of the corresponding steady solution.	108
8.1	Cost function values for the three different sensors with and without flow actuation.	116
8.2	Growth rate μ_r and frequency (Strouhal number) μ_i of the dominant dynamic modes for each case. Last column indicates the Rossiter mode number associated to these leading dynamic modes.	121
8.3	Dominant actuation modes for the downstream and arc cases.	124
C.1	Parameters and stability of the 2D cavity equilibrium solutions. D and Θ are the cavity depth and momentum thickness at the cavity's leading edge, respectively. The columns μ_r and μ_i contain the growth rate and frequency of the leading (first) eigenvalue of the unstable branch.	151
C.2	Simulation parameters of the 2D cavity periodic solutions. D and Θ are the cavity depth and momentum thickness at the cavity's leading edge, respectively. St shows the Strouhal number of each periodic trajectory.	152

Nomenclature

Latin Symbols

BT	Boundary terms
c	Speed of sound
$c.c.$	Complex conjugate
c_p	Specific heat
E	Total energy
e_{int}	Internal energy
e_{kin}	kinetic energy
\vec{f}	External volume force
J	Determinant of the coordinate transformation
\mathcal{J}	Cost function
k	Spatial frequency or wavenumber
\mathcal{L}	Amplitude of the characteristic waves
L	Characteristic length
\mathcal{L}	Lagrangian functional
M	Mach number
\vec{m}^*	Adjoint momentum vector
p	Pressure
Pr	Prandtl number
Q	State variables
Q^*	Adjoint state variables
\vec{q}	Heat flux vector
R	Ideal gas constant
Re	Reynolds number
R_{su}	Sutherland's constant
St	Strouhal number
t	Time
T	Time horizon
T_e	Temperature

Nomenclature

u	Streamwise velocity component
\vec{u}	Velocity vector
U	Contravariant velocity
v	Vertical velocity component
w	Spanwise velocity component
\vec{x}	Spatial vector
x	Streamwise direction
y	Vertical direction
z	Spanwise or lateral direction

Greek Symbols

α	Example variable
γ	Heat capacity ratio
Γ	Boundary of the domain Ω
ε	Viscous dissipation rate
η	Vertical direction in generalised coordinates
Θ	Momentum thickness
κ	Thermal conductivity
λ	Propagating velocity of the characteristic waves
λ_n	Ritz value
μ	Dynamic viscosity
μ_n	System's eigenvalue
ξ	Streamwise direction in generalised coordinates
ρ	Density
τ	Shear stress tensor
ϕ	Control parameters
ψ	Adjoint boundary variables
Ψ_n	Mode shape or spatial mode function
Ψ_n^*	Adjoint mode shape or spatial mode function
ω	Temporal frequency
Ω	Computational domain

Abbreviations

BFS	Backward-Facing Step
BL	Boundary Layer
CFD	Computational Fluid Dynamics
CIC	Characteristic Interface Conditions
DMD	Dynamic Mode Decomposition

DNS	Direct Numerical Simulation
FFT	Fast Fourier Transform
FSI	Fluid-Structure Interaction
LES	Large Eddy Simulation
LHS	Left Hand Side
LODI	Local One Dimensional Inviscid
MPI	Message Passing Interface
OMP	Open Multi-Processing
OSPL	Overall Sound Pressure Level
PDE	Partial Differential Equation
PIV	Particle Image Velocimetry
POD	Proper Orthogonal Decomposition
RANS	Reynolds Averaged Navier-Stokes
RHS	Right Hand Side
ROM	Reduced Order Model
SVD	Singular Value Decomposition
WFS	Wavenumber-Frequency Spectrum

Subscripts & Superscripts

$\tilde{\alpha}$	Approximated variable α - it must not be confused with the arbitrary state \tilde{Q}
$\bar{\alpha}$	Time-averaged variable
α^*	Adjoint variable
α_η	Metric term $\frac{\partial \alpha}{\partial \eta}$
α_ξ	Metric term $\frac{\partial \alpha}{\partial \xi}$
α_w	Variable at wall location
α_∞	Free stream variable

Symbols & Operators

$tr(\alpha)$	Trace of the matrix α
α^T	Transpose of the matrix α
δ	Kronecker delta
Δ	Delta operator
\mathcal{N}	Navier-Stokes operator
\cdot	Inner product
∇	Nabla operator
\otimes	Outer product

Chapter 1

Introduction

One of the first experiments in flow control for scientific purposes was carried out by Prandtl in 1904 ([H. Schlichting and K. Gersten, 2000](#)), who studied the flow separation around a cylinder by exerting active flow control on the boundary layer. Since then, the research community has carried out numerous flow control experiments, which helped to unveil some of the underlying physical mechanisms that govern the fluid motion. Progress demanded more complex flow investigations which were not feasible when Prandtl did his first flow control experiment. Given the enormous advances in experimental apparatus and computational resources since then, it became possible to perform some of these more ambitious studies. On the other hand, especially on the computational fluid dynamics (CFD) side, these available resources are yet not enough to tackle most of the flow engineering applications through direct numerical simulation (DNS) of the flow governing equations. This comes as a consequence of the ‘cost’ of simulating a fluid flow increases with the difference in size between the reference length of the problem and the scale of the smallest eddy in the flow ([Kolmogorov, 1991](#)). Hence, the unavailability of the necessary computational capacity to deal with such applications forced the development of simpler flow models, such as Reynolds-averaged Navier-Stokes (RANS) equations or large eddy simulations (LES), which resulted into affordable simulations by means of degrading the accuracy of the flow solution (e.g. [Pope, 2000](#)). In particular, these simpler methods usually struggle to correctly predict flows with chaotic behaviour, such as turbulent flows. Moreover, the level of chaos in fluid systems is enhanced in the case of separated flows, where the detached flow area consists in a highly turbulent (chaotic) region, which even lacks a dominant flow convection direction. As discussed in the present thesis, the prediction of the turbulent flow behaviour in this kind of scenario plays an essential role in the applicability of methods oriented to flow control, such as the adjoint method. In addition, another very extended flow approximation, mainly for low-speed flows, is the assumption of flow incompressibility. Compared to a low Mach number (compressible) simulation, this approach permits a much faster computation of the flow field, and also a considerably easier implementation of the boundary conditions, which make this method very attractive for many applications. On the other hand, incompressible simulations

Chapter 1 Introduction

are unable to compute the flow acoustics¹ and other compressible flow events, which for some cases, these events are essential for the correct prediction of the flow behaviour. In any case, if the compressible effects do not play a significant role in the flow dynamics, it is still possible to approximate the far field acoustics with, for example, integral methods (Lighthill, 1952, 1954) or linearised Euler equations, among others.

Cavity Flows

In this thesis, we step away from any of the assumptions stated above, restricting our simulations to compressible DNS. This results in a more computationally intensive method in comparison with other approaches, but it prevents dragging potential modelling errors along the present investigation. In particular, we focus our efforts on flows over a rectangular cavity and a backward-facing step (BFS), which are both considered canonical cases due to the large number of publications available in the literature about these two scenarios. Especially, the study of cavity flows has been tackled by numerous researchers for several decades. Perhaps the most interesting features of this flow are the constant flow separation and reattachment, which make this special case very interesting from a fundamental research perspective in fluid dynamics. Furthermore, by just increasing the Reynolds number, the complexity of the flow dynamics inside the cavity can be raised almost as desired. The first investigations carried out in cavity flows were almost exclusively motivated by military applications, such as aircraft bomb bays (e.g. Bilanin and Covert, 1973) or aircraft landing gear wells (e.g. Tam and Block, 1978), among others. As well, the extremely limited computational resources available to this researchers forced them to rely exclusively on experimental techniques to gain an insight into compressible cavity flows (Block, 1976; Heller et al., 1971; Sarohia, 1977). In fact, in one of the first military motivated cavity experimental studies, Rossiter (1964) reported for the first time the presence of compressible self-sustained flow oscillations in cavity flows, which are commonly known as Rossiter modes. These self-sustained periodic oscillations consist of Kelvin-Helmholtz instabilities which grow along the separated shear layer and eventually impinge on the trailing edge of the cavity. This flow impingement originates an acoustic wave which also propagates upstream, and interacts with the incoming shear layer and boundary layer. This flow-acoustic interaction triggers the appearance of new instabilities, which sets off again this feedback loop. Due to the great demand of computational resources required to solve compressible DNS, the contributions from the computational fluid dynamics side were initially limited to incompressible flows (e.g. Pereira and Sousa, 1994; Chang et al., 2006). Despite the aforementioned acoustic-convective feedback loop is not active in their incompressible setup, Kelvin-Helmholtz instabilities are still present in the separated shear layer, provided that the Reynolds number and the cavity's aspect ratio L/D are high enough. Moreover, despite that such a mechanism is exclusive to compressible flows, de Vicente et al. (2014) suggested that a similar pressure feedback process occurs instantaneously for incompressible flows. In addition to computational studies, incompressible cavity flows have also been investigated through experimental setups. For example, Basley et al.

¹Note that the acoustics (or sound waves) are density fluctuations that linearly propagate through the flow-field.

(2011) recovered two-dimensional global structures from an incompressible three-dimensional cavity flow using time-resolved particle image velocimetry (PIV). One of the first relevant studies addressing the flow over an open cavity using compressible DNS was carried out by Rowley et al. (2002). In their study, they performed a large parametric study (changing Re_D , M , L/D , etc.) and documented in thorough detail the dynamics in each case. For low aspect ratio cavities (such as the one investigated herein) they found that the system's dynamics were governed by a shear layer (Rossiter) mode, exhibiting the same feedback mechanism as described above. On the other hand, for higher aspect ratios, the cavity flow was found to abandon the shear layer mode and undergo a transition towards a wake mode type of motion. For further reading about open cavity flows, the reader is referred to a summary of the recent computations and experimental techniques in Lawson and Barakos (2011).

Invariant Solutions

The existence of the above mentioned self-sustained oscillations in compressible cavity flows make this particular setup very attractive for seeking compressible exact solutions of the Navier-Stokes equations. This kind of solutions are also referred to as 'invariant solutions' and 'recurrent' or 'exact coherent' flows. The objective behind the calculation of such exact flow solutions consists in reaching the understanding of complex flow phenomena by breaking down the state-space into simpler flow solutions without the use of any mathematical modelling (i.e., Fourier analysis, POD, DMD, and so on). Thus, highly chaotic systems (turbulent flows) can be thought of a continuous transition from the proximity of one solution to others. Consequently, these chaotic systems may be approximated as an expansion over periodic orbits (e.g. Cvitanović et al., 2016). Hence, exploiting this possibility in the future, we may be able to carry out control of chaotic flows by evaluating the sensitivities of the periodic orbits to the actuation. This would remove the need of computing the sensitivities of turbulent flows with the adjoint method, which for this type of flows, it presents major limitations (e.g. Blonigan, 2016). At this point is also worth highlighting that stable exact solutions can act as a flow attractor, where the system undergoes a continuous decay towards it.

Due to the high complexity involved in solving compressible flows alongside with the much higher computational demands, the research field associated with invariant flow solutions has focused almost exclusively in incompressible flows. Previous investigations searched for steady states in incompressible flows using methods such as continuation (Keller, 1977), selective frequency damping (Åkervik et al., 2006), and even an adjoint-based approach (Farazmand, 2016). In particular, Åkervik et al. (2006) used his method in open cavity flows with a very large aspect ratio L/D . The numerical computation of exact steady solutions was only expanded to compressible flows by Yamouni et al. (2013), where they used the iterative Newton method for its calculation. Back to incompressible flows, this type of work was first carried out computationally with the aim of seeking non-steady states by Nagata (1990), and since then, the search of such flow solutions remained exclusive of incompressible flows. For further details of similar

investigations, the reader is referred to the reviews from [Kawahara et al. \(2012\)](#) and [Cvitanović \(2013\)](#). More recently, [Farazmand \(2016\)](#) extended the available methods to obtain steady and travelling wave solutions using an adjoint-based approach. With his framework, he guarantees convergence to an exact solution from any initial condition. Note that this is a remarkable advantage over more classic methods (e.g. Newton-Raphson iterations), which require of a starting point within the proximity of the sought flow solution to yield satisfactorily converged results. Although it was not his main objective, the reader is also encouraged to go through this article for an interesting introduction/review of invariant flow solutions.

Stability of Cavity Flows

The stability of open cavity flows has been widely studied in the past recent years ([Theofilis, 2011](#)). These investigations were mainly motivated to reach the understanding of the physical mechanisms which drive this very particular flow. This more fundamental perspective of the flow mechanisms is key for the successful design of more complex setups such as active or passive control of this flow ([Liu, 2016](#)). Despite the system's dynamics are strongly governed by compressible events (as described above), some work has been done assessing the stability of incompressible cavities. For example, [Sipp and Lebedev \(2007\)](#) found that the eigenvalue frequencies arising from their linear stability analysis do not match those of the non-linear periodic cavity flow. We later show that this phenomenon also translates to compressible cavity flows. One of the most relevant studies carried out so far investigating the linear stability of compressible cavity flows was done by [Brès and Colonius \(2008\)](#), where they used the mean flow as the base state. In their study, they evaluated the stability of various 2D and 3D cavity flows at several Mach number, where it was found that the origin of the spanwise instabilities was unrelated to the flow compressibility. Hence, the fact that these three-dimensional instabilities are of purely convective nature motivated numerous studies to further investigate this phenomenon from an incompressible flow perspective, such as [de Vicente et al. \(2014\)](#), [Meseguer-Garrido et al. \(2014\)](#) or [Citro et al. \(2015\)](#). Of especial interest, the neutral stability curves gathered in [Meseguer-Garrido et al. \(2014\)](#) show that the flow becomes spanwise unstable before (in a Reynolds number sense) the two-dimensional instabilities commence appearing. Thus, if the flow is 2D unstable, this implies that 3D instabilities are likely to be present too. This coexistence was found to exert a frequency modulation on the instabilities respect to their isolated behaviour ([Brès and Colonius, 2008](#)). Moreover, apart from the computational studies stated above, the spanwise stability of cavity flows has also been looked into experimentally ([Douay et al., 2016](#)). The first stability analysis which used an exact solution as basic flow was carried out by [Åkervik et al. \(2007\)](#). Despite that their two-dimensional cavity flow was modelled as incompressible with a relatively low Reynolds number ($Re = 350$), this equilibrium solution was unstable due to the large cavity's aspect ratio ($L/D \approx 25$), which caused the separated flow to undergo transition towards a wake mode type ([Rowley et al., 2002](#)) periodic cycle. Further, a similar study was carried out by [Giannetti and Luchini \(2007\)](#), this time investigating the instability in a 2D cylinder's wake, and also using several incompressible steady solutions as the base

flow for the stability problem. The use of steady flow solutions to carry out stability analysis was extended to compressible flows by [Yamouni et al. \(2013\)](#), where they studied the evolution of the unstable eigenvalues of a 2D cavity flow across Mach number. With the aspect ratio $L/D = 1$ and the Reynolds number $Re_D = 7500$ used in their study, the flow was already unstable in the incompressible regime. This prevented them from isolating the origin of the compressible instabilities as a function of the Mach number. Note that the appearance of additional unstable eigenvalues with rising Mach number appeared as a surprise to them since compressible effects have stabilising effects in other flows such as shear layer flows (e.g. [Miles, 1958](#); [Pavithran and Redekopp, 1989](#); [Meliga et al., 2010](#)). In any case, similarly to the three-dimensional scenario described above, the constant presence of the unstable modes from purely convective origin is likely to cause a modulation over the onset and evolution of compressible events as a function of Mach number.

The Adjoint Method

A great part of the investigation presented in this thesis relies on the use of the adjoint method in fluid flows. If the reader is not familiar with it, at this point it should only be regarded as a very efficient approach to compute the system's sensitivities, where its application extends well beyond the field of aerodynamics. In most of these applications, the adjoint method is used to efficiently solve an optimisation problem just as shown in the following chapter. For example, a classic problem in seismology consists in minimising the differences between the observed and modelled fields in order to recover the origin of the seismic event ([Fichtner et al., 2006a,b](#)). Also in the field of geophysics, [Bunge et al. \(1998\)](#) used the adjoint method on their Earth's mantle convection model to recover the past conditions of the mantle structure. Further, also following a similar approach, [Talagrand and Courtier \(1987\)](#) were able to recover the initial conditions of a meteorological event using the adjoint method on their models, by minimising the difference between the outcome of the model and the observed data. In oceanography, the use of the adjoint allows the recovery of unobservable sea conditions (such as the sea below large extensions of ice) or the efficient computation of the ocean sensitivities to changing conditions ([Heimbach et al., 2005](#)). For example, such calculations permit the impact estimation in the oceans produced by the rapid loss of ice in Antarctica due to climate change ([Dinniman et al., 2016](#)).

On the other hand, the application of the adjoint method for optimisation purposes in fluid dynamics is slightly different to the cases mentioned above. Instead of recovering an initial or previous state, the optimisation is carried out instead to optimise a flow actuation, which can be either active or passive. [Bewley et al. \(2001\)](#) carried out one of the first relevant investigations which used adjoint simulations for optimal flow control purposes. Their objective consisted in the relaminarisation of an incompressible turbulent channel flow by exerting control at the boundaries with zero-net-flux actuators. Further, one of the most extended applications of the adjoint method is related to the control of far-field noise. Shear layer noise ([Wei and Freund,](#)

2006; Spagnoli and Airiau, 2008; Vishnampet et al., 2015), jet noise (Cerviño et al., 2002; Freund, 2011; Kim et al., 2014; Schulze and Sesterhenn, 2011), and even trailing edge noise (Schulze and Sesterhenn, 2013) amongst others, have been successfully optimally controlled using the adjoint method in recent investigations. As well, other applications are, for example, the control of the flow separation in a cylinder at low Reynolds numbers (Boujo and Gallaire, 2014), or aerodynamic shape optimisation (Jameson, 2003; Schneider, 2006). The primary advantage of the adjoint method resides in the computation of the gradients of the control parameters to the cost function, where the overall computational demands are independent of the number of control parameters. This allows researchers to implement and optimise flow actuations with numerous control parameters at no extra cost, such as numerous time-dependent source terms in the governing equations (Wei and Freund, 2006), or even some porous media distribution over an aerofoil's trailing edge (Schulze and Sesterhenn, 2013) and also at the nozzle exit of a supersonic jet (Schulze and Sesterhenn, 2011). Further applications of the adjoint method apart from optimal flow control, such as stability analysis (Luchini and Bottaro, 2014) or the search of invariant solutions of the Navier-Stokes equations (Farazmand, 2016) are discussed later in this report.

Control of Cavity Flows

The application of the adjoint method in cavity flows seeking an optimal flow actuation is still remains unexplored by the literature. Instead, the studies which use the adjoint method in cavity flows for flow control purposes limit its application to assessing the stability the flow. Thus, based on such stability analysis, they are able to compute receptivity and structural sensitivity maps (Luchini and Bottaro, 2014; Schmid and Brandt, 2014), which show where to most efficiently place the flow actuation. One of the most recent contributions to passive control in cavity flows which uses this adjoint stability-based approach is the work of Liu et al. (2016). Their study regarded a three-dimensional spanwise wall-bounded cavity at $Re \approx 1000$. After carrying out a full stability analysis of the flow, they performed two different passive flow control strategies in the regions of the flow with high values of the structural sensitivity, attempting to control the spanwise shear layer instability. The two passive actuations consisted in adding wall roughness and placing a rod in the area near the leading. From these two approaches, the rod actuation was the most successful. Moreover, away from the adjoint method, the most popular method to carry out flow control in open cavity flows is through closed-loop control. Some of the most relevant published articles using this method are Rowley et al. (2005, 2006); Henningson and Åkervik (2008); Barbagallo et al. (2009); S. J. Illingworth and Rowley (2012); Nagarajan et al. (2013). For further details and other control methodologies in open cavity flows, the reader is referred to the review articles from Rowley and Williams (2006) and Cattafesta et al. (2008).

Backward-facing Step Flow

In addition to the open cavity, this thesis also focuses on the three-dimensional flow over an open backward-facing step. The available literature studying BFS flows is considerable from both experimental and numerical perspectives. The majority of these investigations assume the BFS flow as a sudden expansion inside a channel flow, where the inclusion of an upper wall is likely to modify the flow characteristics in comparison to an open BFS flow. The first study tackling a three-dimensional DNS of a BFS was accomplished by [Le et al. \(1997\)](#), which was validated with the experimental results from [Jović and Driver \(1994\)](#). The main characteristics of this especial flow are the constant flow separation leading to a shear layer which separates the detached flow area from the free stream. Similarly to the flow over an open cavity, the complexity of the flow dynamics within the region with separated flow can be raised almost as desired by just increasing the Reynolds number. On the other hand, note that the length of the recirculation region only keeps a relation with the Reynolds number up to $Re \approx 6600$ ([Armaly et al., 1983](#)). In addition, the investigation by [Adams and Johnston \(1988\)](#) measured the reattachment length for several laminar and turbulent configurations of the incoming boundary layer, where they found that this reattachment length was approximately 30% shorter for the laminar scenarios. A study carried out by [Jović \(1998\)](#) showed that the near-wall boundary layer can take up to 15 step heights in the streamwise direction to recover for $Re > 25500$ (Re based on the step height), where this length experienced a decrease for lower Re . Further, when tracking turbulent quantities such as turbulent kinetic energy or eddy viscosity, they observed that the outer layer required more than 100 step heights to recover equilibrium values. [Kaikitis et al. \(1991\)](#) and [Hasan \(1992\)](#) noticed the existence of two characteristic frequencies present in the BFS flow, through both numerical and experimental investigations, respectively. These two frequencies were referred to by [Hasan \(1992\)](#) as ‘modes of instability’. The lowest frequency is associated with the shear layer oscillation, and it is known as ‘shear layer flapping’ or ‘shear layer mode’. On the other hand, the higher frequency or ‘step mode’, was found to be related to the free-shear layer instability. These two characteristic frequencies were also observed in other studies such as [Eaton and Johnston \(1980\)](#), [Silveira et al. \(1993\)](#) or [Lee and Sung \(2000\)](#).

When a particular flow field is studied from a three-dimensional point of view, the spanwise extent of the domain (both computationally and experimentally) plays an essential role in achieving a correct three-dimensional flow behaviour. In the potential case where the spanwise extent of the domain is insufficient, the flow might experience three-dimensional events induced by the limits of the domain. Similarly, in computational setups, the spanwise boundary conditions could also have an effect on the required spanwise extent. In particular, [Barkley et al. \(2002\)](#) observed that the larger spanwise events were always shorter than the flow reattachment length. For instance, [Le et al. \(1997\)](#) used a domain with of 4 step heights for his simulation at $Re = 5100$, which appears to not satisfy the above criterion. Other studies from [Kaikitis et al. \(1991\)](#) and [Kopera \(2011\)](#) used $L_z = 2\pi$ step heights as their domain width, whereas [Schäfer et al. \(2009\)](#) and [von Terzi \(2004\)](#) set their spanwise extension as π and 3.2 step heights, respectively. At this point, it is worth mentioning that [von Terzi \(2004\)](#) used a symmetry boundary

condition in the spanwise direction for this simulation at $Re = 3000$, which should be equivalent to a total spanwise extent of 6.4 step heights. Moreover, the three-dimensional features of the BFS flow have also been explored experimentally by [Furuichi and Kumada \(2002\)](#), where they focus on the spanwise dynamics at the flow reattachment location.

The stability analysis of the BFS has also caught the attention of several researchers. In fact, assuming incompressible flow, [Barkley et al. \(2002\)](#) found the critical Reynolds number at which three-dimensional instabilities commence to appear at $Re = 748$, where the instability had a wavelength of 6.9 step heights. The reader should bear in mind that the averaged reattachment length reduces with Re (up to approximately $Re = 6600$ where it becomes unrelated). Thus, as stated above, assuming that the wavelength of this spanwise instability stays below the reattachment length, this particular wavelength is expected to decrease with Reynolds number. Additionally, [Blackburn et al. \(2008\)](#) analysed the stability of an incompressible three-dimensional BFS flow, where the Reynolds numbers investigated ranged up to 500. In addition, they also used the adjoint method to compute receptivity maps, which would, later on, be used to study the optimal transient growth of the instabilities. Additional investigations assessing the stability of BFS flows can be found in [Kaiktis et al. \(1991\)](#), [Kaiktis et al. \(1996\)](#), [Beaudoin et al. \(2004\)](#), [Wee et al. \(2004\)](#) or [Lanzerstorfer and Kuhlmann \(2012\)](#).

Control of Backward-facing Step Flows

Over the past decades, the control of BFS flows has been assessed in several experimental studies. The most popular flow actuation is by far the implementation of time-periodic suction and blowing actuators in the vicinity of the top corner of the step, where the objective was normally the reduction of the flow separation area ([Chun and Sung, 1996](#); [Yoshioka et al., 2001](#); [Gautier and Aider, 2013](#)). In particular, [Uruba et al. \(2007\)](#) identified the suction as the leading contributor of the control effectiveness for this type of actuation setups. Also, [Hasan \(1992\)](#) attributed the enhanced shear layer instability growth as the primary cause of a shortened flow separation region. This higher instability growth also lead to a higher turbulent intensity at the impingement location. Apart from a time-periodic suction and blowing actuation, other relevant flow control setup consisted in a flapping foil within the separated flow ([Lai et al., 2002](#)), plasma actuation ([Ruisi et al., 2016](#)), and also the use of a permeable membrane at the flow impingement location ([Heenan and Morrison, 1998](#)). Furthermore, on the computational side, this problem has also been tackled with similar setups as stated above, where regardless of the flow actuation, these studies mostly attempted to develop a closed-loop control. One of the earliest computational studies was carried out by [Choi et al. \(1999\)](#), where they actuated an incompressible two-dimensional flow through mass injection on the vertical wall of the step. [Kang and Choi \(2002\)](#) raised the complexity of this flow control setup when controlling a three-dimensional BFS flow using LES. The actuation also consisted in blowing and suction at the trailing edge of the step seeking a reduction of the flow separation, where the sensing region was located further downstream. Using the same actuation mechanism, [Mehrez et al. \(2011\)](#) attempted to enhance

the heat and mass transfer in the BFS flow. Additionally, other investigations replaced the synthetic jet by an obstacle in the proximity of the trailing edge, also with the aim of increasing the heat and mass transfer. The obstacles used in these investigations were, for example, a rotating cylinder (Selimefendigil and Öztop, 2014), a fixed cylinder (Kumar and Dhiman, 2012) or a squared obstacle (Selimefendigil and Oztop, 2014). Similarly, Neumann and Wengle (2003) (computationally) and Miao et al. (1991) (experimentally) also placed an obstacle in the upstream flow to enhance mixing, but in both cases, their ultimate goal was the reduction of the flow separation. Lastly, one of the most relevant investigation in active control of BFS flows is the work by Barbagallo et al. (2012). In particular, they developed a reduced order model (ROM) closed-loop control strategy to mitigate the unsteadiness of the flow. Their Reynolds number was $Re = 600$ and the flow actuation consisted in two body forces perturbing the vertical velocity field of the upstream flow. Moreover, they carried out a thorough estimation study prior to the actuation. Despite the sensors located further upstream do not capture the actuation signal fully amplified by the flow, they yet provide a more favourable estimation compared to the sensors located further downstream, mainly due to their considerably lower time delay.

1.1 Motivation

The investigation presented in the current report explores two different compressible separated flow configurations over complex geometries. In particular, we focus our efforts on flows over a rectangular cavity and a BFS, where we apply a newly developed adjoint-based framework for optimal flow control purposes. As seen above, both flow geometries are considered as canonical cases due to the large number of publications available in the literature about these two scenarios. On the other hand, the majority of the computational studies are carried out under the assumption of flow incompressibility, and here we show that these neglected compressible events can have a dominant repercussion on the flows studied. For instance, it is widely believed in a general sense that compressibility has a stabilising effect on shear layer instabilities (e.g. Miles, 1958; Pavithran and Redekopp, 1989; Meliga et al., 2010), and here we show how this does not hold for compressible cavity flows. In fact, we prove that these compressible effects in open cavity flows fuel further the growth of stronger shear layer instabilities, due to the strong acoustic-convective feedback mechanism² characteristic from this type of flows, known also as Rossiter modes (Rossiter, 1964). In particular, this acoustic-convective feedback dominates the sound radiation in cavity flows, where its control provoked a great interest amongst the fluid dynamics researchers over the past recent years. More precisely, the flow actuations available in the literature are usually oriented to modify the shear layer impingement over the cavity's trailing edge, to either change the leading sound radiation direction or to completely diminish the sound radiation in all directions. To this end, in the present investigation, we develop and use an adjoint-based optimal flow control framework to actuate and modify the sound radiation

²Despite that such a mechanism is exclusive to compressible flows, de Vicente et al. (2014) suggest that a similar pressure feedback process occurs instantaneously for incompressible flows.

Chapter 1 Introduction

of a two-dimensional cavity flow. Similarly, this framework is also applied to the flow over a three-dimensional BFS, which represents a simplified fluid-structure interaction (FSI) cabin noise reduction problem. The most active sector in this field is the automotive industry, where their main interest (passenger comfort) resides in reducing the flow-induced vibrations on windows and car panels which radiate sound into the cabin. Especially, the front window is the area of the car where the largest research effort is expended (e.g. [Bremner, 2012](#); [Van Herpe and Duarte, 2012](#)). The A-pillar of a car (see [Haruna et al. \(1990\)](#) for an illustration) and also the wing mirror are the responsible for a flow separation and impingement over the front side window. The pressure fluctuations associated with both detached flow and impingement induce relevant vibrations on this window, which radiate noise into the passengers' cabin. Hence, the three-dimensional BFS approximates such flow events through a much simpler geometry, where the target of our flow actuation is the pressure fluctuations in a section of the lower wall of the BFS. Note that these pressure fluctuations occur due to both convective and acoustic events. Also, despite the energy of the acoustic waves is lower than the convective ones, they present a considerable contribution to cabin noise since their characteristic wavelength matches the leading structural modes of the window ([Bremner and Wilby, 2002](#)). Thus, the use of a compressible flow solver is justified for both cavity and BFS problems.

Before delving into the flow control exercises of the above-described separated flows, this report also analyses the two-dimensional open cavity flow from a much more fundamental perspective. In particular, we apply our optimisation framework to compute families of periodic and steady exact compressible flow solutions across Mach number for this geometry. The motivation behind the calculation of such exact flow solutions is to understand the complex flow phenomena by breaking down the flow-field into simpler flow solutions without the use of any mathematical modelling. Following this approach, if such solutions are unstable, highly chaotic systems can be thought of a continuous transition from the proximity of one solution to others ([Cvitanović et al., 2016](#)). Contrarily, stable exact solutions can act as a flow attractor, where the system continuously decays towards it. In the context of this thesis, one potential application of finding these periodic flow solutions uses the possibility of approximating chaotic systems as an expansion over periodic orbits. Thus, in the future, we may be able to carry out control of chaotic flows by evaluating the sensitivities of the periodic orbits to the actuation. Therefore, this would remove the need of calculating the sensitivities of chaotic systems with the adjoint method, which as discussed later on, it presents some major limitations for such scenarios. For the present case, the leading interest on computing these periodic solutions in a cavity flow resides in analysing the evolution of the acoustic-convective feedback mechanism across Mach number. Given the non-dimensional character of our framework, the reader must bear in mind that a change in Mach number, apart from changing the flow compressibility, it also modifies the propagating speed of sound relative to the flow velocity. In addition, to show and prove the destabilising effect of a higher Mach number in cavity flows, we also carry out a linear stability analysis of the family of steady flow solutions. To this purpose, since the development of a linearised Navier-Stokes solver is out of the scope of this report, the eigenvalue decomposition of

each steady solution can be calculated using the adjoint equations (Luchini and Bottaro, 2014) with just minor modifications to the adjoint-based optimal flow control framework.

1.2 Contributions

Chapter 6 reports for the first time a family of exact periodic solutions in compressible flows with the added difficulty of a complex geometry, which is the leading contribution from this chapter. In addition, we also present a novel framework which permits the computation of such solutions in any other complex geometry. Furthermore, a family of equilibrium flow solutions associated to the periodic orbits is also introduced, where it is shown that both families meet at the quasi-incompressible regime, forming a Hopf bifurcation of subcritical type. This proves the destabilising effect of compressibility in cavity flows, which is analysed in detail, covering an existing gap in the literature. Moreover, the flow-acoustic interaction present in cavity flows is generally overlooked by literature, where here we provide a thorough description of its evolution across Mach number. For the aspect ratio used, these results extend the data available in the literature of open cavity flows. Chapter 7 describes the linear stability analysis of the equilibrium flow solutions computed in chapter 6. The leading contribution from this chapter is the insight over the compressible flow mechanisms which trigger an instability in these steady solutions. In addition, this analysis sheds light on the origin and nature of the bifurcation amongst the periodic and equilibrium families of solutions. Furthermore, this chapter aims to establish a connection with previous studies regarding the stability of 2D and 3D cavity flows, isolating the effect of purely compressible two-dimensional flow phenomena across Mach number. Moreover, these results also extend the available data about stability analysis of cavity flows in terms of both aspect ratio and Reynolds number. Chapter 8 presents a new application of adjoint-based optimisation. In particular, we make use of this method to perform noise radiation and directivity control of a two-dimensional open cavity flow, which is the main contribution of this chapter. In addition, we present the resulting optimal actuation, showing the ideal forcing patterns for such control in cavity flows. These results are valuable information which can be used in future investigations to model other flow actuation strategies. Lastly, in chapter 9 we show another new application of adjoint-based optimisation, this time on a three-dimensional BFS flow. The investigation is oriented to cabin-noise reduction, where we optimise an upstream flow actuation to reduce the flow impinging scales which present the largest impact on cabin noise. In addition, the complexity present in the flow studied pushes the adjoint-based approach at its limits, which results in an added complexity compared to other applications of the method. These intrinsic limitations of the method are discussed, where we also suggest and demonstrate a successful alternative which permits the use of adjoint-based optimisation methods in separated flows.

1.3 Outline of the Report

In chapter 2 the theory that supports the current investigation is discussed. The flow and adjoint governing equations are presented in chapter 3. The numerical setup that describes the implementation of the DNS code with its new adjoint framework is detailed in chapter 4. The verification of the adjoint-based optimisation framework is carried out in chapter 5 using simple test cases. The results presented in this report begin in chapter 6. There, we show and analyse the families of periodic and steady exact flow solution for a two-dimensional cavity flow. The stability of the steady solutions is further investigated in chapter 7 through an adjoint-based linear stability analysis. Chapter 8 contains the adjoint-based optimal flow control of a two-dimensional cavity flow, seeking an overall sound and directivity control. Finally, the adjoint-based optimal flow control approach is pushed to its limit in chapter 9, when attempting the optimal control of a three-dimensional BFS flow, as a simplified cabin noise reduction problem. For chapters 6 to 9, their contributions are summarised in their respective initial paragraph. Similarly, the conclusions and future work suggestions are gathered in the summary section at the end of each of these chapters. Lastly, chapter 10 recapitulates the thesis from a more overall perspective, summarising the leading contributions and conclusions from this work. In addition, the appendices complement the report, showing additional derivation of the adjoint equations (appendix A) and their boundary terms (appendix B) and also presenting the relevant simulation parameters relative to chapters 6 and 7 (appendix C). Of particular interest, a preliminary study exploring the limitations of adjoint-based optimisation on separated flows is reported in appendix E.

Chapter 2

Background Theory

The current chapter introduces the essential theory this research project is based on. Firstly, we present the basics of flow control and optimisation, which contains an introduction and derivation of the adjoint method. This is followed by an insight on the flow decomposition techniques used throughout this report. Finally, we describe the fundamentals of linear stability theory, with special emphasis on the valuable use of the adjoint method.

2.1 Flow Control and Optimisation

To evaluate the performance of a flow control strategy a generic quadratic cost functional is defined as

$$\mathcal{J}(Q, \phi) = \frac{1}{2} \int_0^T \int_{\Omega} [W(\vec{x}) Q^T \mathbf{M}_1 Q + \phi^T \mathbf{M}_2 \phi] d\vec{x} dt, \quad (2.1)$$

where Q represents the state variables and ϕ is the vector of control parameters. \mathbf{M}_1 sets the relation of the state variables in the cost function and \mathbf{M}_2 accounts for the cost in the actuation parameters. The function $W(\vec{x})$ weights the cost in the domain Ω . The optimal control will be the values of the actuation parameters that result in the absolute minimum of the cost functional for the time horizon T .

To drive the cost function towards its minimum, the gradient

$$\frac{D\mathcal{J}}{D\phi} = \frac{\partial \mathcal{J}}{\partial Q} \frac{dQ}{d\phi} + \frac{\partial \mathcal{J}}{\partial \phi} \quad (2.2)$$

must be computed. The challenging part of obtaining the gradient resides in the computation of the *sensitivities* of the system $dQ/d\phi$. The most intuitive method to calculate this term is perturbing individually each of the control parameters and then using a finite difference scheme to approximate the derivatives. This process would be repeated until the minimum value of the cost functional is reached, where $D\mathcal{J}/D\phi = 0$. With this method, the amount of required fluid simulations increases proportionally with the number of control parameters, which makes

this approach unattainable when dealing with practical CFD simulations¹. Fortunately for the present investigation, there are alternatives which provide optimal control with a much lower computational cost.

2.1.1 Adjoint Method

The Lagrangian functional \mathcal{L} is defined as

$$\mathcal{L}(Q, Q^*, \phi) = \mathcal{J}(Q, \phi) - (Q^*)^T \cdot \mathcal{N}(Q, \phi), \quad (2.3)$$

where \mathcal{N} is the operator representing the set of governing equations² and Q^* are the *co-state* or *adjoint* variables. The governing equations, the adjoint equations and the optimality condition can be derived from the Lagrangian functional (M. D. Gunzburger, 2003) by setting its first variation with respect to Q^* , Q and ϕ equal to zero, respectively

$$\frac{\delta \mathcal{L}}{\delta Q^*} = 0 \quad \Rightarrow \quad \text{State equations} \quad (2.4)$$

$$\frac{\delta \mathcal{L}}{\delta Q} = 0 \quad \Rightarrow \quad \text{Adjoint equations} \quad (2.5)$$

$$\frac{\delta \mathcal{L}}{\delta \phi} = 0 \quad \Rightarrow \quad \text{Optimality condition.} \quad (2.6)$$

Bear in mind that if the operator \mathcal{N} or the cost function \mathcal{J} were non-linear, the differentiation defined in (2.5) would not remove the forward state variables Q from the adjoint equations. Hence, if that is the case, solving the adjoint equations for a time-dependent problem would require the state variables to be available at every timestep³.

From the three systems of equations above, the optimal control can be achieved by solving iteratively the state equations, followed by the solution of the adjoint equations, to finally use these two results to compute the updated control parameters with the optimality condition. Unfortunately, this method is equivalent to the steepest descent algorithm with a fixed step size (Gunzburger and Lee, 1994), which can converge very slowly. Due to the great cost of performing DNS, we instead use the L-BFGS algorithm (Morales and Nocedal, 2011), which is more efficient (Badreddine et al., 2014). Note that the tolerances of the optimisation algorithm might require modification⁴ to satisfy the Wolfe conditions (Nocedal and Wright, 2006). As shown later in chapter 3, the reader should bear in mind that the state equations march forward in time starting from an initial condition at $t = 0$, whereas the adjoint equations are posed backwards in time starting from a terminal condition at $t = T$.

¹An intuitive example is shown in Schulze et al. (2011).

²A forward flow simulation will satisfy $\mathcal{N}(Q) = 0$.

³See subsection 4.2.1 for further details.

⁴The author recommends a low tolerance in the descent condition. High tolerances are likely to not satisfy this condition, resulting from the line search algorithm changing direction, converging to the initial control state.

2.1.1.1 Adjoint-based Gradient

The sensitivity equations can be derived by differentiating the state equations with respect to ϕ , resulting in

$$\frac{\partial \mathcal{N}}{\partial Q} \frac{dQ}{d\phi} + \frac{\partial \mathcal{N}}{\partial \phi} = 0. \quad (2.7)$$

Rearranging (2.7) and substituting it into (2.2) leads to

$$\frac{D\mathcal{J}}{D\phi} = \frac{\partial \mathcal{J}}{\partial \phi} - \frac{\partial \mathcal{J}}{\partial Q} \left(\frac{\partial \mathcal{N}}{\partial Q} \right)^{-1} \left(\frac{\partial \mathcal{N}}{\partial \phi} \right), \quad (2.8)$$

where the solution to the adjoint system (2.5) is

$$Q^* = \left[\frac{\partial \mathcal{J}}{\partial Q} \left(\frac{\partial \mathcal{N}}{\partial Q} \right)^{-1} \right]^T. \quad (2.9)$$

The adjoint equations can now be expressed as

$$\left(\frac{\partial \mathcal{N}}{\partial Q} \right)^T Q^* = \left(\frac{\partial \mathcal{J}}{\partial Q} \right)^T. \quad (2.10)$$

Its solution can be used to compute the gradient of the functional as

$$\frac{D\mathcal{J}}{D\phi} = \frac{\partial \mathcal{J}}{\partial \phi} - (Q^*)^T \frac{\partial \mathcal{N}}{\partial \phi}. \quad (2.11)$$

Note that for (2.11), the cost of computing the gradient is independent of the number of control parameters and is only necessary to solve the adjoint system once to obtain the gradients for every control parameter. An alternative introduction to the adjoint method can be found, for example in the supplemental material of [Luchini and Bottaro \(2014\)](#).

There are two different approaches to derive the adjoint of a system, where both of them have their advantages and disadvantages with respect to each other. These two methods are the continuous or *differentiate then discretise*, and the discrete or *discretise then differentiate*. The continuous approach consists of deriving the adjoint equations analytically following (2.5), where they are then discretised with a numerical scheme to be solved numerically. On the other hand, the discrete approach starts from the already discretised forward system which is then differentiated to obtain the discrete adjoint operator. Note that with this method, the numerics used to solve the forward system of equations are kept during the entire derivation. This includes the numerical scheme used, boundary conditions, and so on. If any of these numerical methods are changed the whole derivation procedure must be repeated. To make this process simpler, automatic differentiation tools ([Mader and Martins, 2008](#)) or matrix-free methods ([Fosas de Pando et al., 2012](#)) can be used. The main characteristic of the discrete approach is that the sensitivities and the gradients obtained are consistent with the discretised problem. This means that, for example, in a very under-resolved simulation, the discrete adjoint will provide the ‘right’ gradients to minimise the cost function even though the physics of the problem are corrupted due to poor

grid resolution. On the contrary, the continuous adjoint would fail to give the ‘correct’ gradients to minimise the given cost function. Theoretically, both continuous and discrete approaches would converge to the same result when the discretisation errors vanish ($\Delta\vec{x} \rightarrow 0$) and the equations are physically correct. [Vishnampet et al. \(2015\)](#) claimed a higher ‘accuracy’ in the sense above of their discrete adjoint code against their continuous one when comparing the gradients from both adjoint solvers with a finite difference approach. In their study, they considered a shear-layer in an open flow, where the boundary conditions were of characteristic-type. In the author’s opinion, it is unsurprising that they claimed an accuracy down to numerical precision for the discrete adjoint, as both forward and discrete adjoint systems are the same. Moreover, a common practice when dealing with characteristic boundary conditions in continuous adjoint codes is to neglect the boundary terms that arise from the derivation of the continuous adjoint equations (see chapter 4). Note that this might lead to inconsistent gradients if these boundaries are not placed far enough from the domain of interest. Further discussion about the differences of both approaches can be found, for example, in [M. D. Gunzburger \(2003\)](#) and [Giles and Pierce \(2000\)](#).

[Bewley et al. \(2001\)](#) carried out one of the first relevant investigations which used adjoint simulations for optimal flow control purposes. Their objective consisted in the relaminarisation of an incompressible turbulent channel flow by exerting control at the boundaries with zero-net-flux actuators. The most popular application of the adjoint method amongst the flow control community is probably on far-field noise reduction. Noise optimal control in shear layer flows ([Wei and Freund, 2006](#); [Spagnoli and Airiau, 2008](#); [Vishnampet et al., 2015](#)) and jet flows ([Cerviño et al., 2002](#); [Kim et al., 2014](#)) have been widely studied in the past recent years. In relation to jet noise reduction, [Schulze et al. \(2011\)](#) presented a convenient cost function which only targets specific ranges of the frequency domain. Other successful applications of adjoint-based control are aerodynamic shape optimisation ([Jameson, 2003](#); [Schneider, 2006](#)) or optimisation of porous media for trailing edge noise reduction ([Schulze and Sesterhenn, 2013](#)). Further applications of the adjoint method apart from optimal flow control, such as stability analysis ([Luchini and Bottaro, 2014](#)) or the search of invariant solutions of the Navier-Stokes equations ([Farazmand, 2016](#)) are discussed later in this report. A summary on the ‘essentials’ for model-based control in fluid mechanics for both iterative adjoint-based and also direct Riccati-based feedback can be found in [Kim and Bewley \(2007\)](#). The derivation of the adjoint equations used in the present investigation can be found in chapter 3 and appendix A.

2.1.1.2 Limitations of Adjoint-based Optimisation: Model Predictive Control

In a separated flow, the trajectories of two neighbouring fluid particles diverge exponentially with time due to the chaotic behaviour of turbulence. Consequently, the initial condition is ‘forgotten’ by the system and the system’s sensitivities cannot be computed because the initial and final conditions are unrelated. This means that if the prediction time horizon is longer than the time it takes two neighbouring trajectories to diverge, the computed sensitivities will be

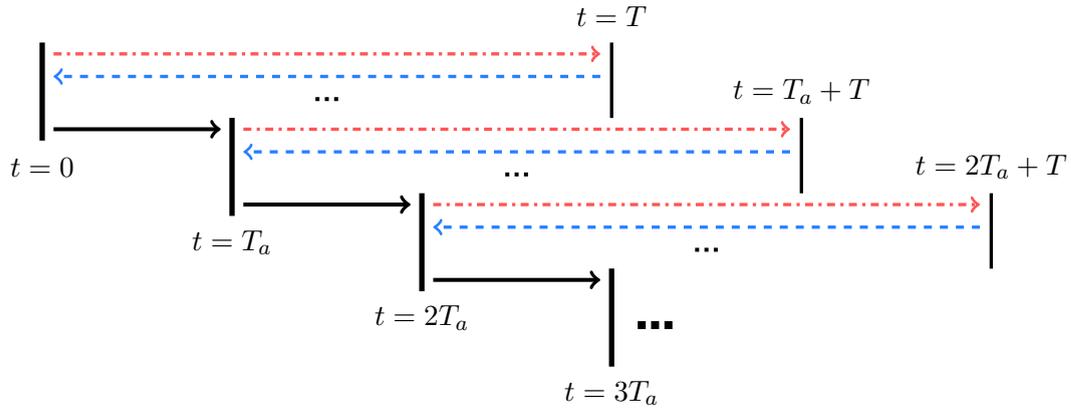


Figure 2.1: Illustration of the receding horizon technique, where T is the horizon length. Forward (---) and adjoint (---) simulations are run iteratively in the same horizon until the convergence of the control parameters is achieved. After convergence, the optimisation horizon is shifted forward in time T_a and the entire procedure is repeated until required.

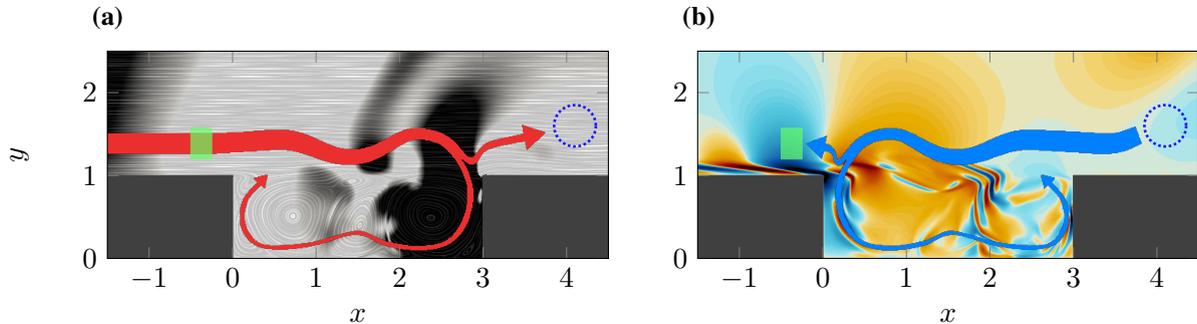


Figure 2.2: Sketch of one of the flow control setups of the 2D cavity flow studied in chapter 8. Figures (a) and (b) represent the forward and adjoint simulations respectively. The flow actuation is depicted by the green box and the blue dashed circle represents the sensing region.

incorrect. [Lea et al. \(2000\)](#) used the Lorenz attractor as an example of chaotic systems, showing how the adjoint analysis failed to give the correct time-averaged sensitivities. They suggested that for relatively long time horizons with respect to the predictability of the system's time scales, the adjoint quantities diverge exponentially due to cumulative error growth. Hence, the use of techniques such as receding horizon (figure 2.1 and, for example, [Lee, 2011](#)) are necessary.

In the above-mentioned investigation by [Bewley et al. \(2001\)](#), a larger time horizon was shown to be beneficial when controlling time-averaged flow quantities. On the other hand, they stated that as the prediction horizon is increased, the required effort to optimise the dynamical system for a given control scheme increases exponentially. An intrinsic limitation in the choice of this time horizon T resides in the travelling time of the control flow scales from the actuation location to the target region, which depends exclusively on the flow characteristics between these two locations (figure 2.2a). Note that the same limitation occurs in adjoint simulations in the opposite direction (figure 2.2b). These limitations are further explored in a preliminary study

on a two-dimensional BFS, gathered in appendix E. In addition, if the reader is not familiar with the use of the adjoint method in chaotic systems, it is strongly recommended to read appendix E before delving into chapters 8 and 9.

2.2 Flow-field Decomposition and Reduced Order Modelling

In complex systems, such as any time evolving fluid flow with non-linear structures present in it, it is essential to post-process the raw data to unveil the underlying physical mechanisms that govern the system's dynamics. In the current section, several methods to decompose and analyse the raw data from the direct and adjoint simulations (applied later in this report) are presented. Additionally, regarding the flow actuations carried out in following chapters, these tools will enable a further understanding of their effect on the rest of the flow-field.

2.2.1 Spectral Decomposition

Fourier transforms allow to map the physical domain onto a wave domain, providing an additional perspective of the flow-field by observing the magnitude, frequency and phase of the fluctuations. These transformations can be applied for both time and spatial domains, where the space-equivalent of frequency ω is the wavenumber k

$$\alpha(\omega) = \frac{1}{2\pi} \int_{-\infty}^{\infty} \alpha(t) e^{-i\omega t} dt. \quad (2.12)$$

This transformation represents all the fluctuations of the variable $\alpha(t)$ as a sum of sine and cosine waves with different amplitude and phase. If this is now extended to a two-dimensional set of data, for example, time series of the same quantity α gathered over an streamwise array of monitor points, $\alpha(x, t)$ will be transformed to

$$\alpha(k_x, \omega) = \frac{1}{(2\pi)^2} \iint_{-\infty}^{\infty} \alpha(x, t) e^{-i\omega t} e^{-ik_x x} dt dk_x. \quad (2.13)$$

The dimensions of frequency and wavenumber are $1/\text{length}$ and $1/\text{time}$ respectively. Hence, the wavenumber-frequency spectrum will show the energy content for every wave propagation velocity as a function of their frequency and wavenumber. This analysis can be further extended to a second spatial dimension as

$$\alpha(k_x, k_z, \omega) = \frac{1}{(2\pi)^3} \iiint_{-\infty}^{\infty} \alpha(x, z, t) e^{-i\omega t} e^{-ik_x x} e^{-ik_z z} dt dk_x dk_z, \quad (2.14)$$

where now for each frequency ω_i , the directivity of the flow structures can be observed in the $k_x(\omega_i) - k_z(\omega_i)$ plane.

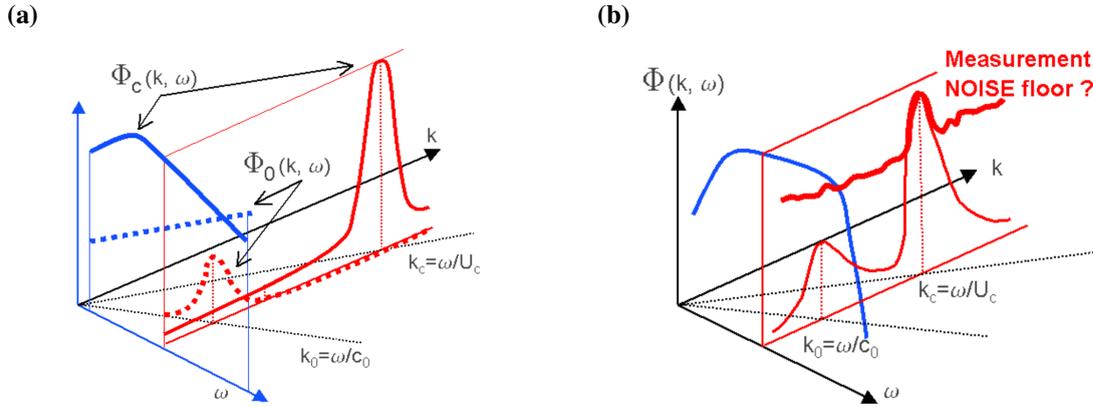


Figure 2.3: Illustration of a wavenumber-frequency spectrum for surface pressure of a subsonic separated flow. Φ_c and Φ_o represent the turbulent and acoustic components of the spectrum, respectively (figure reproduced from [Bremner and Wilby, 2002](#))

Wavenumber-Frequency Spectrum

When dealing with compressible flows, [Bremner and Wilby \(2002\)](#) suggested that for low Mach numbers, using the difference of the turbulent and acoustic structures in the wavenumber-frequency space, it is possible to decompose the pressure-field and analyse both wave-types separately (figure 2.3a). They described the acoustic field as the result of the propagation of turbulent disturbances far from the original eddy, where the amplitude of pressure fluctuations has decreased to the acoustic range. It is characteristic for acoustic waves to have a low wavenumber, and they will be laying on the acoustic cone present in the wavenumber-frequency spectrum (WFS), limited by both positive and negative propagating speed of sound. Turbulent fluctuations will have a higher wavenumber and they will be situated on both sides of the acoustic cone, limited by the maximum upstream and downstream flow velocity. Note that for supersonic flows this decomposition is not feasible since the turbulent and acoustic structures will overlap in the WFS.

This type of flow-field decomposition is almost exclusively used in studies dealing with aero-vibro-acoustics and cabin noise (e.g. [Van Herpe et al., 2011](#); [Bremner, 2012](#)), as it clearly shows each of the loading mechanisms on the structure which cause noise radiation. The highest energy point on the convective side of the WFS indicates the main flow convective velocity. Both turbulent and acoustic diffusion are also present in the WFS. Acoustic diffusion is located inside the acoustic cone, whereas the turbulent diffusion sits in between positive and negative flow convective velocities, what may lead to obscuring the acoustic content in the WFS (figure 2.3b).

2.2.2 Proper Orthogonal Decomposition

[Sirovich \(1987\)](#) originally suggested the use of proper orthogonal decomposition (POD) as a method to identify coherent structures from PIV snapshots. Since then, the use of POD as a tool

Chapter 2 Background Theory

to break down the flow and analyse the coherent structures has rapidly gained in popularity. This method can be seen as an energetic filter that has been shown already to be very effective for turbulent flow scales separation (Holmes et al., 1996; Druault et al., 2012). Additionally, POD can be used as an order reduction method (Smith, 2003) by disregarding the low-energy modes. On the other hand, Rowley (2005) highlighted that POD provides an optimal decomposition of the flow energy-wise, which does not guarantee that these modes give the best representation of the flow dynamics.

The main idea behind POD is to approximate a function $\alpha(\vec{x}, t)$ for a given domain as an infinite summation

$$\tilde{\alpha}(\vec{x}, t) = \sum_{n=1}^{N_{mod}} a_n(t) \Psi_n(\vec{x}), \quad (2.15)$$

where $\Psi_n(\vec{x})$ is a function for the mode n with a time-invariant orthonormal basis. To compute the modes, the flow-field data from each snapshot is stored in columns of a matrix \mathbf{X} with dimensions $m \times n$, where m represents the dimensions of the captured flow-field and n is the total number of snapshots. The formulation of this decomposition allows to select a subset of the entire flow domain disregarding the rest of the flow-field. Hence, the events that occur in the chosen sub-domain are enhanced. A singular value decomposition (SVD) on \mathbf{X} results in

$$\mathbf{X} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H, \quad (2.16)$$

where the proper orthogonal modes $\Psi_n(\vec{x})$ can be found as the left singular vectors (columns of \mathbf{U}) and the temporal coefficients $a_n(t)$ are the right singular vectors (rows of \mathbf{V}^H). The diagonal elements of $\mathbf{\Sigma}$ are sorted in decreasing order and represent the energy of each proper orthogonal mode. The singular value decomposition sorts the modes (energy-wise) in a least squares sense, i.e. the first two functions give the best-two term approximation, the three first give the first three-term approximation and so on. Given the chosen basis, the functions Ψ_n are not time dependent and the coefficient function $a_n(t)$ can be determined as

$$a_n(t) = \int_{\Omega} \alpha(\vec{x}, t) \Psi_n(\vec{x}) dx. \quad (2.17)$$

Meyer et al. (2007) used POD to study the structures involved in a jet in cross-flow, Kostas et al. (2005) identified coherent structures from velocity and vorticity field on a backward-facing step flow and Bergmann et al. (2005) developed a POD-based reduced order model to control the wake of a rotating cylinder. In the field of vibro-acoustics, Druault et al. (2013) attempted to discern between turbulent and acoustic pressure fluctuations on a car side window flow. As mentioned by Bremner and Wilby (2002); Bremner and Zhu (2003) the acoustic waves are several orders of magnitude below the energy content of the turbulent structures, which cause the ‘acoustic’ POD modes to be sometimes obscured by low energy random events in the flow-field.

2.2.3 Dynamic Mode Decomposition

Another way that allows the extraction of additional dynamic information from flow field raw data is the dynamic mode decomposition (DMD) introduced by Schmid (2010). This method uses uniformly sampled snapshots (captured with constant Δt) to decompose the flow variables into flow structures or modes that have a major contribution to the flow dynamics, i.e. a growth rate and frequency can be assigned to each of these modes. In a similar fashion to POD, it is possible to choose a sub-domain of interest to carry out DMD. Given a snapshot of the flow field at a given time step i , the DMD method assumes that the snapshot from the time step $i + 1$ can be linked with the previous capture by a linear transformation or mapping \mathbf{A} , such as

$$\mathbf{X}^N = \mathbf{A}\mathbf{X}^{N-1}. \quad (2.18)$$

If dealing with a highly non-linear system, sufficiently small time steps are required to assume linear behaviour in between each frame. After the linear mapping is assumed, the set of time-frames is expressed as Krylov sequence

$$\mathbf{X}^N = \{x, \mathbf{A}x, \mathbf{A}^2x, \dots, \mathbf{A}^{N-1}x\}. \quad (2.19)$$

where \mathbf{X} is the snapshot matrix and N is the total number of time steps. The next assumption of this approach, is to consider that after N snapshots, where all the relevant flow features have been captured, the following time steps can be formulated as a linear combination of the former ones

$$\mathbf{x}_N = \mathbf{X}^{N-1}\mathbf{a} + \mathbf{r}, \quad (2.20)$$

where \mathbf{r} is a residual vector. With \mathbf{e}_{N-1}^T as the $(N - 1)$ order unit vector, the series can be continued as

$$\mathbf{A}\mathbf{X}^{N-1} = \mathbf{X}^N = \mathbf{X}^{N-1}\mathbf{S} + \mathbf{r}\mathbf{e}_{N-1}^T, \quad (2.21)$$

where \mathbf{X}^N is the shifted snapshot matrix. The matrix \mathbf{S} is of companion type, and its eigenvalues approximate some of the eigenvalues from \mathbf{A} . According to Schmid (2010), the computation of \mathbf{S} with the above method would result in an ill-conditioned problem. For the sake of robustness, he suggested the projection of the operator \mathbf{A} onto a POD basis as

$$\mathbf{X}^{N-1} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^H, \quad (2.22)$$

which after substituting and rearranging in (2.20) leads to

$$\mathbf{U}^H\mathbf{A}\mathbf{U} = \mathbf{U}^H\mathbf{X}^N\mathbf{V}\mathbf{\Sigma}^{-1} \equiv \tilde{\mathbf{S}}. \quad (2.23)$$

The eigenvalues of $\tilde{\mathbf{S}}$ (λ_n) give a robust approximation of the ones from \mathbf{A} . These eigenvalues are also known as Ritz values, where the unstable ones present the largest magnitude, having $|\lambda_n| > 1$ and $|\lambda_n| < 1$ as unstable and stable eigenvalues, respectively. Lastly, the eigenvalues which sit at the unit circle $|\lambda_n| = 1$ correspond to complex eigenvalues with purely oscillatory

Chapter 2 Background Theory

behaviour. Hence, to revert this transformation and recover the ‘standard’ eigenvector form we must apply the following relation

$$\mu_n = \log(\lambda_n) / \Delta t, \quad (2.24)$$

where the real and imaginary parts of μ_n indicate the growth rate $\mu_r = \Re\{\log(\lambda_n)\} / \Delta t$ and frequency $\mu_i = \Im\{\log(\lambda_n)\} / (2\pi\Delta t)$ of the corresponding dynamic mode. The dynamic modes can be finally calculated as

$$\Psi_n = \mathbf{U}\mathbf{w}, \quad (2.25)$$

where \mathbf{w} are the eigenvectors of $\tilde{\mathbf{S}}$. In the same publication, [Schmid \(2010\)](#) proved DMD to be an effective tool to identify the spatial structures and their associated frequencies using PIV data from a flow over a square cavity, a flow over a flexible membrane and a jet flow between two cylinders. Seeking a measure of coherence, [Schmid \(2010\)](#) projected the dynamic modes onto a POD basis to rank them by energy content. [Sampath and Chakravarthy \(2014\)](#) used POD and DMD to study the flow over a backward-facing step, where they presented an alternative measure of coherence between POD and dynamic modes. By a Fourier transform of the time coefficients of the most energetic POD modes they were able to extract the dominant frequencies for each POD mode. This frequencies enabled them to link each POD mode to a dynamic mode with the same frequency.

[Tu et al. \(2014\)](#) extended the use of DMD to a wider range of datasets. This included non-sequential data samples and they also presented some sampling techniques to reduce the noise present on the dynamic modes. Additionally, they reformulated the way the dynamic modes are computed as

$$\Psi_n = \frac{1}{\lambda_n} \mathbf{X}^N \mathbf{V} \Sigma^{-1} \mathbf{w}, \quad (2.26)$$

which (according to them) increases the connection between the DMD and Koopman modes ([Mezić, 2005](#); [Rowley et al., 2009](#)). For this reason, they referred to their method as ‘Exact DMD’.

[Mezić \(2013\)](#) stated that the major advantage of Koopman modal analysis and DMD is that they can deal with snapshot sequences rather than relying on the linearisation around a state of interest, e.g., global mode analysis ([Huerre and Monkewitz, 1990](#); [Bagheri, 2010](#)), where, as detailed in the following section, the flow is decomposed into modes that oscillate at a single frequency with a particular growth rate. [Song et al. \(2013\)](#) tested the performance of the linear global mode techniques and DMD methods when capturing linear and non-linear sound generation mechanisms, suggesting that the DMD method captures the dynamics of the near and far field acoustics with a small number of modes.

DMD with Control

Proctor et al. (2014) extended the application of DMD to account for the effect of control in actuated systems. This enables the possibility to generate an input-output reduced order model of the system, providing an insight into the interaction of the control with the system dynamics. Note that the previous definition of DMD cannot be used as is not able to include the effects of the external forcing over the dynamics represented by each of the modes. The linear mapping (2.18) is now transformed into

$$\mathbf{X}^N = \mathbf{A}\mathbf{X}^{N-1} + \mathbf{B}\mathbf{Y}^{N-1}, \quad (2.27)$$

where \mathbf{Y}^{N-1} is the input matrix. In their report, they present two approaches to computing the dynamic modes depending on whether the matrix \mathbf{B} is known or not. From a fluid dynamics perspective, the matrix \mathbf{B} is rarely known as the system is highly non-linear. Hence, this report is only concerned about the method for an unknown matrix \mathbf{B} .

The system from (2.27) can be rearranged in the form of

$$\mathbf{X}^N = \begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix} \begin{bmatrix} \mathbf{X}^{N-1} \\ \mathbf{Y}^{N-1} \end{bmatrix} = \mathbf{G}\mathbf{\Omega}, \quad (2.28)$$

to then carry out a SVD on $\mathbf{\Omega}$, which gives

$$\begin{bmatrix} \mathbf{A} & \mathbf{B} \end{bmatrix} \approx \begin{bmatrix} \mathbf{X}^N \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}_1^H \\ \mathbf{Y}^N \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}_2^H \end{bmatrix}. \quad (2.29)$$

Additionally, a second SVD on \mathbf{X}^N is required to find a reduced model of the output space

$$\mathbf{X}^N = \hat{\mathbf{U}} \hat{\mathbf{\Sigma}} \hat{\mathbf{V}}^H. \quad (2.30)$$

Hence, (2.29) can be reorganised as

$$\tilde{\mathbf{A}} = \hat{\mathbf{U}}^H \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}_1^H \hat{\mathbf{U}} \quad (2.31)$$

$$\tilde{\mathbf{B}} = \hat{\mathbf{U}}^H \mathbf{Y}^N \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}_2^H \hat{\mathbf{U}}. \quad (2.32)$$

In a similar fashion as for the exact DMD, the dynamic modes of \mathbf{A} can be computed as

$$\Psi_n = \mathbf{X}^N \mathbf{V} \mathbf{\Sigma}^{-1} \mathbf{U}_1^H \hat{\mathbf{U}} \mathbf{w}. \quad (2.33)$$

2.3 Stability Analysis

The stability analysis in fluid flows has provoked a great interest in researchers during the past decades, where it can still be regarded as a ‘hot topic’ amongst investigators in fluid dynamics.

Chapter 2 Background Theory

In this section, we show how the adjoint method introduced earlier in section 2.1 can shed light on the stability of a fluid system. This time, the sensitivity information provided by the adjoint method is not used to compute very efficiently gradients to update some flow actuation. Instead, the adjoint method is applied yielding sensitivity information respect to the stability of the system. This is achieved using the same adjoint solver with just very minor modifications to the numerical framework used for flow control purposes (see chapters 3 and 4). This results in a very attractive situation where we can easily extract a new feature from our numerical framework without any extra coding effort.

One of the simplest ways to define the stability of a system under equilibrium (from a finite-time perspective) is by tracking the amplification (or decay) of an infinitesimal energy perturbation of the form

$$\mathbf{X}'_0 = \mathbf{X}_0 + \epsilon \hat{\mathbf{X}} \quad (2.34)$$

through a given time horizon (e.g. Schmid and Henningson, 2000). Respectively, \mathbf{X}'_0 and $\epsilon \hat{\mathbf{X}}$ represent the perturbed state and the energy perturbation of infinitesimal amplitude, where \mathbf{X}_0 stands for the unperturbed state in equilibrium, also referred to as basic flow or base flow. Hence, the most straightforward of the stability analysis can be carried out just by perturbing the base flow and measuring the drift in the state respect to the non-perturbed state after some time. If the altered system would diverge from the original state with time-marching, the system would be regarded as unstable. Contrarily, a system where the disturbances introduced decay, rendering the unperturbed system, is considered stable. Note that from this simple exercise, one can only establish whether the fluid system in question is stable or not under the particular perturbation introduced. Unfortunately, the answer to other relevant questions remain unknown, such as *is our system stable under other different perturbations?*, and if unstable, *how should the system be perturbed to maximise the instability growth?*, *where in our system does the instability grow?* or *where in our system is the instability originated?* Fortunately for the present investigation, there are other methods which provide a deeper insight on the stability properties of fluid flows.

In fluid dynamics, one of the most popular methods to investigate the stability of fluid flow configurations is global linear stability (e.g. Theofilis, 2011). As the method's name suggests, this approach assumes that our system is governed by a linear operator, which is expressed as indicated by (2.18), or in a general form, as shown by (2.27). Similarly to DMD, this technique seeks a flow decomposition of the form

$$\tilde{\mathbf{X}}(t) = \mathbf{X}_0 + \sum_{n=1}^{N_{mod}} \Psi_n e^{\mu_n t} + c.c., \quad (2.35)$$

where \mathbf{X}_0 represents the base flow and Ψ_n and μ_n are the system's eigenvectors and eigenvalues, respectively. These eigenvalues μ_n are simply the roots of the characteristic polynomial of \mathbf{A} , where each eigenvalue has its associated eigenvector Ψ_n . The election of the base flow \mathbf{X}_0 is a critical step to obtain meaningful results from this type of stability analysis. The most common approach (which we use later in appendix E on a two-dimensional BFS flow) is to choose the

time-averaged flow field as the base flow. Bear in mind that the use of the time-averaged flow field as the base flow is only meaningful if the frequencies of the resulting eigenvalues are related to those of the non-linear flow, which does not necessarily always hold (Sipp and Lebedev, 2007). Alternatively, it is also possible to set the base flow as an exact steady flow solution (see chapter 7). For this particular case, the information yielded by the eigendecomposition is always useful, since it directly provides an insight into the stability of such particular flow solution.

One way of carrying out such a decomposition would be coupling an eigenvalue solver (see for example PARPACK by Lehoucq and G. (2001) or SLEPc by Hernandez et al. (2005)) with the numerical framework which solves the governing equations. At this point, the reader might realise that the equations which govern the fluid dynamics are non-linear (see chapter 3). In any case, one could readily obtain the forward spectrum by coupling any of the linear eigenvalue solvers stated above to the forward non-linear code. As well, this method could also be used on the adjoint equations since they are already linear. Thus, we are able to compute the eigenspectrum of a fluid system since the eigenvalues obtained through both adjoint or forward eigendecompositions are the same (see for example the supplemental appendix of Luchini and Bottaro, 2014). Alternatively to coupling an eigenvalue solver to our framework, the adjoint global eigenvalues of the linear adjoint process can also be extracted applying DMD (Schmid, 2010), as introduced earlier in subsection 2.2.3. Differently from above, where the eigenvalue solver requires flow solves on demand until convergence is met, the DMD user must ensure beforehand that the snapshot sequence is large enough. Also bear in mind that when computing the adjoint snapshot sequence, the base flow is not the initial condition of the adjoint simulation, as indicated previously in (2.35) for a forward flow decomposition. Instead, the base flow is accounted for in the adjoint simulation as the forward state Q , which is present in the adjoint equations due to the non-linearity of the Navier-Stokes equations (see subsection 2.1.1 and chapter 3). Hence, it is a common practice to use a random noise field as the initial (terminal) condition of the adjoint simulation.

Global Modes and Structural Sensitivity

Generally, the characteristic eigenvalues of a fluid-dynamics problem sit in the complex plane, where the real and imaginary parts indicate the growth rate and frequency, respectively. Similarly, the right eigenvectors Ψ (or forward global modes) associated with the eigenvalues of the system are also complex and satisfy

$$\mathbf{A}\Psi = \mu\Psi. \quad (2.36)$$

Note that the governing equations are formulated in the real number space \mathbb{R} . This forces every complex eigenpair (μ_n, Ψ_n) to have its corresponding complex conjugate (*c.c.*), which cancels out the imaginary part of Ψ_n in equation (2.35). These forward modes highlight the flow events which evolve in time with the same growth rate and frequencies as indicated by their respective eigenvalues. Also, the forward modes provide an insight into the system's response to a perturbation. Hence, stable ($\Re\{\mu_n\} < 0$) and unstable ($\Re\{\mu_n\} > 0$) modes show the regions of the

Chapter 2 Background Theory

flow-field where this perturbation is damped and amplified, respectively. On the other hand, the left eigenvectors Ψ^* or adjoint global modes shed light on the system's sensitivities and fulfil

$$\Psi^* \mathbf{A} = \mu \Psi^*, \quad (2.37)$$

where, as mentioned above, the eigenvalues match the ones from equation (2.36). Note that this is a great advantage for the present investigation since we can obtain the system's eigenvalues through the adjoint equations, without the need of developing a linearised Navier-Stokes solver. These adjoint modes are strongly linked to the forward modes, where particularly, as remarked by Schmid and Brandt (2014), they reveal receptivity information of the associated forward mode to external perturbations. Further, the adjoint modes can be thought of spatial maps which reveal the locations of the flow-field where an external perturbation would set off efficiently the respective forward mode, being the local magnitude of the adjoint mode a measure of efficiency. For a more detailed review of the use of the adjoint equations for stability analysis see Luchini and Bottaro (2014). Unfortunately, at this point, the forward modes remain yet unknown, but they can be approximated, for example, using DMD. When studying the stability of a given base flow \mathbf{X}_0 , the flow behaviour can be assumed linear on the initial stages of time marching, even when solving the full non-linear equations. Hence, one could use DMD on this limited set of data to approximate the forward modes of the system⁵. Finally, the correctness of these approximated forward modes can be measured comparing the resulting eigenvalues with the ones obtained through the adjoint simulations.

Intuitively, the region of the flow-field where the instability is more efficiently generated would be the one with both high receptivity (high values of the adjoint mode) and response (high values of the forward mode). Hence, if both forward and adjoint modes are available, we can extract such information combining them as their point-wise product. This results in a spatial sensitivity map of the eigenvalue, highlighting the locations where an external (structural) perturbation would cause the largest drift of the eigenvalue. To clarify, Luchini and Bottaro (2014) defined a structural perturbation as a variation in the coefficients of \mathbf{A} rather than forcing the right-hand side of the governing equations. Hence, the locus of the flow-field where this structural sensitivity map reaches maximum values represents the origin of the instability, which is commonly referred to as the wave maker. For a mathematical derivation and definition, the reader is referred to Giannetti and Luchini (2007). Furthermore, linking stability analysis back with flow control, the structural sensitivity map or wave maker is a powerful tool to choose the right location of a flow actuation. An actuator located where the wave maker reaches high values would give it a large control authority since it can efficiently cause a drift in the corresponding leading eigenvalue.

⁵Further extending the dataset would result in non-linear flow events being accounted for by the DMD, which might contaminate the desired linear modes.

2.4 Summary

The adjoint method was presented in section 2.1 as the chosen procedure to perform optimal flow control in the following chapters of the current report. The advantages and the mathematical derivation of the adjoint problem were demonstrated. The different approaches and more relevant applications of the method were also presented. Section 2.2 described the flow decomposition techniques used throughout this report. Section 2.3 described the fundamentals of linear stability theory, with special emphasis on the valuable use of the adjoint method.

Chapter 3

Governing Equations

This chapter presents the equations that govern the simulations carried out through this research project. In particular, both ‘forward’ and adjoint governing equations are solved directly, without any modelling assumptions. This results in a more computationally intensive method in comparison with other approaches, but it prevents dragging the potential modelling errors along the present investigation. Additionally, as part of the development of the adjoint framework, the adjoint equations are also expressed in characteristic form, to, later on, derive the appropriate adjoint boundary conditions for the current formulation of these equations.

3.1 Navier-Stokes Equations

The full compressible Navier-Stokes equations govern the fluid flow in our simulations. This set of equations are derived by applying the mass conservation, Newton’s second law (momentum conservation) and the conservation of energy, resulting into a non-linear set of five partial differential equations (PDEs), known as continuity (3.1), momentum (3.2) and total energy (3.3) equations, which are defined for the entire flow domain Ω . In particular, the equations are used in their conservative form, where the variables are density ρ , the three-dimensional momentum vector $\rho\vec{u}$ and the conserved energy ρE

$$\frac{\partial \rho}{\partial t} + \nabla \cdot (\rho\vec{u}) = 0 \quad (3.1)$$

$$\frac{\partial \rho\vec{u}}{\partial t} + \nabla \cdot (\rho\vec{u} \otimes \vec{u}) = -\nabla p + \nabla \cdot \tau + \rho\vec{f} \quad (3.2)$$

$$\frac{\partial \rho E}{\partial t} + \nabla \cdot (\rho E\vec{u}) = -\nabla \cdot (p\vec{u}) + \nabla \cdot (\tau \cdot \vec{u}) + \nabla \cdot \vec{q} + \rho\vec{u} \cdot \vec{f}, \quad (3.3)$$

where \vec{u} is the velocity vector, \vec{f} represents an external volume force, and E is the total energy defined as the sum of internal and kinetic energy $E = e_{int} + \frac{1}{2}\vec{u} \cdot \vec{u}$. Note that the non-linearities of the system reside on the second term of the left hand side (LHS) of equations (3.2) and (3.3), where only the first one is the responsible of the generation of spatial structures that evolve in

Chapter 3 Governing Equations

time that we know as turbulence. Since the flow is considered compressible, the equation of state (3.4) gives closure to the system linking pressure p with density ρ , temperature T_e and internal energy e_{int}

$$p = \frac{\rho T_e}{\gamma M_\infty^2} = (\gamma - 1)\rho e_{int}. \quad (3.4)$$

The shear stress tensor τ is symmetric and is defined as

$$\tau = \frac{\mu}{Re_\infty} \left(\nabla \vec{u} + (\nabla \vec{u})^T \right) - \left(\frac{2}{3} \frac{\mu}{Re_\infty} \nabla \cdot \vec{u} \right) \delta, \quad (3.5)$$

with μ as dynamic viscosity and δ as the identity matrix. Note the contribution of the viscous stress tensor in the momentum (3.2) and energy equations (3.3). The term \vec{q} in (3.3) is known as heat flux vector, which is defined as

$$\vec{q} = \frac{\mu}{Re_\infty (\gamma - 1) M_\infty^2 Pr_\infty} \nabla T_e \quad (3.6)$$

where γ is the isentropic coefficient of the fluid which is assumed to be 1.4.

A common practice is to non-dimensionalise the above set of governing equations, giving rise to a more general solution and also, with regards to the numerical computation, reducing truncation error. This non-dimensionalisation is achieved by introducing the dimensionless Reynolds (Re), Mach (M) and Prandtl (Pr) numbers; which are defined as

$$Re = \frac{\rho_\infty u_\infty L}{\mu_\infty}, \quad M = \frac{u_\infty}{\sqrt{\gamma R T_e_\infty}} \quad \text{and} \quad Pr = \frac{\mu_\infty c_p}{\kappa_\infty}. \quad (3.7)$$

The sub-script ∞ indicates free stream variables, and L is the characteristic length of the problem, R is the ideal gas constant ($287 [m^2/Ks^2]$), c_p is the specific heat and κ is the thermal conductivity. The Prandtl number is assumed to be constant at 0.72 in all the simulations presented in this report. The molecular viscosity μ is computed using Sutherland's law (White, 1991)

$$\mu = T_e^{3/2} \left(\frac{1 + R_{Su}}{T + R_{Su}} \right), \quad (3.8)$$

setting the ratio of the Sutherland constant over free-stream temperature to 0.36867.

3.2 Adjoint Equations

The present research relies on the adjoint method to carry out optimal control, stability analysis and study the flow sensitivities. As introduced in chapter 2, there are two main approaches to tackle the adjoint problem, the *discrete* and the *continuous* approaches. One of the fundamental sources of disagreement between both methods resides in the resulting gradient of the cost functional as the spatial resolution decreases. DNS are carried out in the current research, which implies that no modelling assumptions are considered, and also the temporal and spatial

discretisations are fine enough to capture all the events contained in the flow field. Given these conditions, the discrepancies between the continuous and discrete approaches should be minimal. For convenience and flexibility in the implementation, the continuous approach is the one chosen in the present investigation.

The adjoint formulation used assumes the energy equation to be rearranged in terms of pressure. This is done by deriving the internal energy equation and then substituting in the equation of state. The internal energy equation is obtained by subtracting the kinetic energy equation from the total energy equation. The kinetic energy is calculated as

$$\frac{\partial \rho e_{kin}}{\partial t} = \vec{u} \cdot \text{momentum} - \frac{1}{2} \vec{u}^2 \cdot \text{continuity}, \quad (3.9)$$

where the derivatives can be summed up using the product rule. Hence, subtracting equation (3.9) from the total energy equation (3.3) and using the equation of state (3.4), gives

$$\frac{\partial p}{\partial t} + \nabla \cdot (p\vec{u}) + (\gamma - 1) [p(\nabla \cdot \vec{u}) + \nabla \cdot \vec{q} + tr(\tau \cdot \nabla \vec{u})] = 0. \quad (3.10)$$

According to the definition given in equation (2.5), the adjoint equations follow from the linearisation of the Navier-Stokes equations (3.1), (3.2) and (3.10) as

$$\lim_{\epsilon \rightarrow 0} \left(\frac{\mathcal{L}(Q + \epsilon \tilde{Q}, Q^*, \phi) - \mathcal{L}(Q, Q^*, \phi)}{\epsilon} \right) = 0, \quad (3.11)$$

where \tilde{Q} is an arbitrary state. As the equations considered here are PDEs, integration by parts is required to remove, if possible, the derivatives from the arbitrary state \tilde{Q} . This last step introduces additional terms which are only defined at the boundary of the domain Γ , known as *boundary terms*. Note that equation (3.11) satisfies the Green-Lagrange identity

$$\langle Q^*, \mathcal{N}'(Q, \phi) Q' \rangle = \langle \mathcal{N}^*(Q, \phi) Q^*, Q' \rangle + BT, \quad (3.12)$$

where BT represents the boundary terms and $Q' = \epsilon \tilde{Q}$. Generally, these boundary terms should not be neglected since they contribute to determine the ‘correct’ boundary conditions of the adjoint problem (see chapter 4).

3.2.1 Adjoint Euler Equations

The state variables chosen for the derivation of the adjoint governing equations are $Q = [p, \rho u, \rho v, \rho w, \rho]$, where their corresponding co-state or adjoint variables are $Q^* = [\rho^*, m_u^*, m_v^*, m_w^*, p^*]$. For the sake of clarity, the viscous terms of the governing equations are grouped as the RHS Sv terms

Chapter 3 Governing Equations

on each equation. Hence, the adjoint Euler equations can be written as

$$-\frac{\partial \rho^*}{\partial t} - \vec{u} \cdot \nabla \rho^* + (\gamma - 1) \rho^* \nabla \cdot \vec{u} - \nabla \cdot \vec{m}^* = Sv_{\rho^*} \quad (3.13)$$

$$-\frac{\partial \vec{m}^*}{\partial t} - \gamma \frac{p}{\rho} \nabla \rho^* - (\gamma - 1) \frac{\rho^*}{\rho} \nabla p - \vec{u} \left(\nabla \vec{m}^* + (\nabla \vec{m}^*)^T \right) - \nabla p^* = Sv_{\vec{m}^*} \quad (3.14)$$

$$-\frac{\partial p^*}{\partial t} + \gamma \frac{p \vec{u}}{\rho} \cdot \nabla \rho^* + (\gamma - 1) \frac{\rho^*}{\rho} \vec{u} \cdot \nabla p + \vec{u} \cdot (\vec{u} \cdot \nabla) \vec{m}^* = Sv_{p^*}. \quad (3.15)$$

Note that the derivation of the above adjoint equations relates the adjoint density with pressure and the adjoint pressure with density, as proposed by [Cerviño et al. \(2002\)](#). Hence, both forward and adjoint equations present a “logical zero-Mach-number limit”, where in case of constant ρ and ρ^* the equations (3.10) and (3.15) cancel out. For a detailed derivation of the above equations¹ the reader should refer to appendix A.

3.2.2 Adjoint Navier-Stokes Equations

In previous studies based on Lagrangian optimisation, the derivation of the adjoint viscous terms has been usually carried under certain assumptions (e.g. [Moret-Gabarro, 2009](#); [Spagnoli and Airiau, 2008](#)). No assumptions were made during the derivation of the viscous terms in the present investigation. Hence, the full compressible adjoint Navier-Stokes equations are solved. The viscous source terms of equations (3.13), (3.14) and (3.15) are

$$Sv_{\rho^*} = \bar{q}^* + \mu^* \chi \quad (3.16)$$

$$Sv_{\vec{m}^*} = \frac{1}{\rho} \nabla \cdot \sigma^* \quad (3.17)$$

$$Sv_{p^*} = -\frac{\vec{u}}{\rho} \cdot \nabla \cdot \sigma^* - \frac{p}{\rho} (\bar{q}^* + \mu^* \chi), \quad (3.18)$$

where \bar{q}^* and σ^* are

$$\bar{q}^* = \frac{\gamma}{Pr_\infty Re_\infty \rho} \nabla \cdot (\mu \nabla \rho^*), \quad \sigma^* = \tau^* - 2(\gamma - 1) \rho^* \tau. \quad (3.19)$$

As the adjoint simulation propagates information in a backward sense, energy must be transmitted from small to bigger structures, hence σ^* must consist of a combination of dissipation and production terms. The dissipation is accounted for with

$$\tau^* = \frac{\mu}{Re_\infty} \left(\nabla \vec{m}^* + (\nabla \vec{m}^*)^T \right) - \left(\frac{2}{3} \frac{\mu}{Re_\infty} \nabla \cdot \vec{m}^* \right) \delta, \quad (3.20)$$

¹Note that appendix A shows the derivation of the adjoint Euler equations. The adjoint Navier-Stokes equations can be derived appending the viscous terms to the forward Euler equations and following the same steps as indicated in appendix A.

which could be seen as an adjoint momentum shear stress tensor, whereas the production term depends on the forward shear stress tensor τ . Note that if the viscous dissipation term is neglected in (3.10) this production term would vanish. For simplicity, the adjoint viscosity (μ^*) has been introduced, which leads to the following algebraic equation which closes the full adjoint Navier-Stokes equations

$$\mu^* = tr \left((\nabla \vec{m}^* - (\gamma - 1) \rho^* \nabla \vec{u}) \cdot \left(\frac{\tau}{\mu} \right) \right) + \frac{1}{Pr M_\infty^2 Re_\infty} \left(\nabla \rho^* \cdot (\nabla T_e)^T \right). \quad (3.21)$$

Note that the need to introduce (3.21) vanishes if constant viscosity is assumed. The adjoint viscosity term is weighted in the adjoint continuity and pressure equations by χ , which follows directly from assuming the variations in viscosity according to the Sutherland's law, giving

$$\chi = \frac{T_e^{3/2}}{p} \left(\frac{1 + R_{Su}}{T_e + R_{Su}} \right) \left(\frac{T_e}{(T_e + R_{Su})} - \frac{3}{2} \right). \quad (3.22)$$

3.3 Adjoint Equations in Characteristic Form

To achieve an equivalent boundary treatment as done in the forward case, the adjoint equations must be transformed into characteristic form. With this method, the derivation of conditions like non-reflecting boundary conditions (see section 4.1.3) or even characteristic interface conditions (see section 4.2.4) is feasible. For the adjoint governing equations, the method of characteristics applies in the same manner as for the forward equations (Thompson, 1987). Note that equations (3.13), (3.14) and (3.15) lack of a conservative or flux form. Hence, the procedure gets easier since the conversion to primitive variables is already done (Thompson, 1990). As the code used in this project only allows periodic boundary conditions in the spanwise direction, the characteristic equations must only be derived in the streamwise (ξ) and vertical (η) directions. The equations can be rewritten in generalised Cartesian coordinates by multiplying them by the determinant of the coordinate transformation² (J), and using the relations

$$J \frac{\partial}{\partial x} = \frac{\partial y}{\partial \eta} \frac{\partial}{\partial \xi} - \frac{\partial y}{\partial \xi} \frac{\partial}{\partial \eta} \quad (3.23)$$

$$J \frac{\partial}{\partial y} = \frac{\partial x}{\partial \xi} \frac{\partial}{\partial \eta} - \frac{\partial x}{\partial \eta} \frac{\partial}{\partial \xi}. \quad (3.24)$$

3.3.1 Characteristic Equations – ξ

The idea behind the method is to transform the derivatives that are normal to the computational boundary into combinations of the amplitudes of the characteristic waves \mathcal{L} , where each one of them has its own characteristic velocity λ . The \mathcal{L} values, which are initially computed from the data within the domain, are corrected afterwards by imposing the appropriate boundary or

²For more information regarding the transformation into generalised coordinates, the reader is referred to Anderson (1995).

Chapter 3 Governing Equations

interface conditions. To achieve this transformation, the governing equations have first to be written in the form

$$-\frac{\partial Q^*}{\partial t} + \frac{1}{J} A_\xi \frac{\partial Q^*}{\partial \xi} + C_\xi = 0, \quad (3.25)$$

where C_ξ contains the non-normal derivatives to the boundary, the source terms and the viscous terms. If the full viscous equations are considered, the system of equations becomes non-hyperbolic, which makes necessary the use of certain assumptions. The Local One-Dimensional Inviscid (LODI) Relations (Poinsot and Lele, 1992) assume that the system is one-dimensional and inviscid at the boundaries, which implies that the contribution of C_ξ can be neglected. Following this strategy, the boundary conditions are ‘designed’ for the LODI problem, but the solution is advanced in time without any assumptions. The matrix A_ξ contains all the coefficients multiplying the normal derivatives to the computational boundary. Hence, after carrying out the coordinate transformation to equations (3.13), (3.14) and (3.15), the matrix A_ξ can be written as

$$A_\xi = \begin{pmatrix} vx_\eta - uy_\eta & -y_\eta & x_\eta & 0 & 0 \\ -c^2 y_\eta & vx_\eta - 2uy_\eta & -vy_\eta & -wy_\eta & -y_\eta \\ c^2 x_\eta & ux_\eta & 2vx_\eta - uy_\eta & wx_\eta & x_\eta \\ 0 & 0 & 0 & vx_\eta - uy_\eta & 0 \\ c^2 (uy_\eta - vx_\eta) & u^2 y_\eta - uvx_\eta & uvy_\eta - v^2 x_\eta & uwy_\eta - vwx_\eta & 0 \end{pmatrix}. \quad (3.26)$$

The speed of sound c and the contravariant velocity U^3 can be expressed as

$$c^2 = \frac{\gamma p}{\rho} \quad (3.27)$$

$$U = uy_\eta - vx_\eta. \quad (3.28)$$

The characteristic propagation velocities can be obtained by calculating the eigenvalues of A_ξ , giving

$$\lambda = \begin{pmatrix} -U \\ -U \\ -U \\ -U - c\sqrt{x_\eta^2 + y_\eta^2} \\ -U + c\sqrt{x_\eta^2 + y_\eta^2} \end{pmatrix}. \quad (3.29)$$

At this point it can be seen already that the propagation of the adjoint characteristics occurs (as expected) in the opposite direction as for the forward system (e.g. Kim and Lee, 2000). The amplitude of the characteristics is defined as

$$\mathcal{L} \equiv \frac{1}{J} \lambda \left[S_\xi^{-1} \frac{\partial Q^*}{\partial \xi} \right]^T, \quad (3.30)$$

³In this section the coordinate transformation metrics $\frac{\partial y}{\partial \eta}$, $\frac{\partial x}{\partial \eta}$, $\frac{\partial y}{\partial \xi}$ and $\frac{\partial x}{\partial \xi}$ are abbreviated as y_η , x_η , y_ξ and x_ξ respectively.

where the rows of the matrix S_ξ^{-1} are left eigenvectors of A_ξ . Note that equation (3.30) is divided by the determinant of the coordinate transformation J . The characteristic terms must be computed in ‘physical space’ rather than ‘computational space’⁴ due to possible metric-terms discontinuities when dealing with multi-block characteristic interfaces (see section 4.2.4). The matrix S_ξ^{-1} can be written as

$$S_\xi^{-1} = \begin{pmatrix} 0 & \frac{y_\eta U}{x_\eta^2 + y_\eta^2} & \frac{-x_\eta U}{x_\eta^2 + y_\eta^2} & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & \frac{x_\eta}{x_\eta^2 + y_\eta^2} & \frac{y_\eta}{x_\eta^2 + y_\eta^2} & 0 & 0 \\ \frac{-c}{2\sqrt{x_\eta^2 + y_\eta^2}} & \frac{-cy_\eta + u\sqrt{x_\eta^2 + y_\eta^2}}{2c(x_\eta^2 + y_\eta^2)} & \frac{cx_\eta - v\sqrt{x_\eta^2 + y_\eta^2}}{2c(x_\eta^2 + y_\eta^2)} & \frac{-w}{2c\sqrt{x_\eta^2 + y_\eta^2}} & \frac{-1}{2c\sqrt{x_\eta^2 + y_\eta^2}} \\ \frac{c}{2\sqrt{x_\eta^2 + y_\eta^2}} & \frac{-cy_\eta + u\sqrt{x_\eta^2 + y_\eta^2}}{2c(x_\eta^2 + y_\eta^2)} & \frac{cx_\eta + v\sqrt{x_\eta^2 + y_\eta^2}}{2c(x_\eta^2 + y_\eta^2)} & \frac{w}{2c\sqrt{x_\eta^2 + y_\eta^2}} & \frac{1}{2c\sqrt{x_\eta^2 + y_\eta^2}} \end{pmatrix}. \quad (3.31)$$

Bear in mind that this transformation into characteristic form introduces non-linearities in the metric terms. In any case, these equations can still be regarded as linear since our simulations only use fixed grids, where these metric parameters remain constant throughout the simulation at every grid point.

At this stage, boundary conditions can be applied by modifying \mathcal{L} , giving corrected \mathcal{L}' values. Hence, equation (3.25) must be rewritten to get the corrected time derivative as

$$-\frac{\partial Q^*}{\partial t} + S_\xi \mathcal{L}' + C_\xi = 0, \quad (3.32)$$

with S_ξ being a matrix whose columns are the right eigenvalues of A_ξ . The $S_\xi \mathcal{L}'$ term can be written as

$$S_\xi \mathcal{L}' = \begin{pmatrix} -\frac{1}{c^2} & -\frac{w}{c^2} & -\frac{ux_\eta + vy_\eta}{c^2} & -\frac{\sqrt{x_\eta^2 + y_\eta^2}}{c} & \frac{\sqrt{x_\eta^2 + y_\eta^2}}{c} \\ 0 & 0 & x_\eta & -y_\eta & -y_\eta \\ 0 & 0 & y_\eta & x_\eta & x_\eta \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & U & U \end{pmatrix} \cdot \begin{pmatrix} \mathcal{L}'_1 \\ \mathcal{L}'_2 \\ \mathcal{L}'_3 \\ \mathcal{L}'_4 \\ \mathcal{L}'_5 \end{pmatrix}. \quad (3.33)$$

Note that for the sake of computational stability, the matrices S_ξ and S_ξ^{-1} have been rearranged to avoid potential singularities caused by the metric terms.

⁴Due to the generalised coordinate form of the equations, it is more convenient to advance the time integration if the grid coordinates are mapped into a uniform-spaced domain $\left(J \frac{\partial Q^*}{\partial t}\right)$, which is referred to as ‘computational space’. On the other hand, if the derivatives are computed in the stretched or non-uniform grid $\left(\frac{\partial Q^*}{\partial t}\right)$, it is said that they are in ‘physical space’.

3.3.2 Characteristic Equations – η

Following the same procedure as for the ξ boundary, the governing equations are written in the form

$$-\frac{\partial Q^*}{\partial t} + \frac{1}{J} A_\eta \frac{\partial Q^*}{\partial \eta} + C_\eta = 0, \quad (3.34)$$

where the matrix A_η is

$$A_\eta = \begin{pmatrix} -vx_\xi + uy_\xi & y_\xi & -x_\xi & 0 & 0 \\ c^2 y_\xi & -vx_\xi + 2uy_\xi & vy_\xi & wy_\xi & y_\xi \\ -c^2 x_\xi & -ux_\xi & -2vx_\xi + uy_\xi & -wx_\xi & -x_\xi \\ 0 & 0 & 0 & -vx_\xi + uy_\xi & 0 \\ c^2(vx_\xi - uy_\xi) & uvx_\xi - u^2 y_\xi & v^2 x_\xi - uv y_\xi & vwx_\xi - uwy_\xi & 0 \end{pmatrix}, \quad (3.35)$$

and for the vertical direction, the contravariant velocity is

$$V = vx_\xi - uy_\xi. \quad (3.36)$$

The characteristic propagation velocities, this time in the η direction, are obtained by calculating the eigenvalues of the matrix A_η , giving

$$\lambda = \begin{pmatrix} -V \\ -V \\ -V \\ -V - c\sqrt{x_\xi^2 + y_\xi^2} \\ -V + c\sqrt{x_\xi^2 + y_\xi^2} \end{pmatrix}. \quad (3.37)$$

Similarly to the ξ boundary, the amplitudes of the characteristic waves \mathcal{L} have first to be computed with the current data within the computational domain, before they are corrected with the given boundary conditions

$$\mathcal{L} \equiv \frac{1}{J} \lambda \left[S_\eta^{-1} \frac{\partial Q^*}{\partial \eta} \right]^T, \quad (3.38)$$

where S_η^{-1} can be written as

$$S_\eta^{-1} = \begin{pmatrix} 0 & \frac{-y_\xi V}{x_\xi^2 + y_\xi^2} & \frac{x_\xi V}{x_\xi^2 + y_\xi^2} & 0 & 1 \\ 0 & 0 & 0 & 1 & 0 \\ 0 & \frac{x_\xi}{x_\xi^2 + y_\xi^2} & \frac{y_\xi}{x_\xi^2 + y_\xi^2} & 0 & 0 \\ \frac{-c}{2\sqrt{x_\xi^2 + y_\xi^2}} & \frac{cy_\xi - u\sqrt{x_\xi^2 + y_\xi^2}}{2c(x_\xi^2 + y_\xi^2)} & -\frac{cx_\xi + v\sqrt{x_\xi^2 + y_\xi^2}}{2c(x_\xi^2 + y_\xi^2)} & \frac{-w}{2c\sqrt{x_\xi^2 + y_\xi^2}} & \frac{-1}{2c\sqrt{x_\xi^2 + y_\xi^2}} \\ \frac{c}{2\sqrt{x_\xi^2 + y_\xi^2}} & \frac{cy_\xi + u\sqrt{x_\xi^2 + y_\xi^2}}{2c(x_\xi^2 + y_\xi^2)} & \frac{-cx_\xi + v\sqrt{x_\xi^2 + y_\xi^2}}{2c(x_\xi^2 + y_\xi^2)} & \frac{w}{2c\sqrt{x_\xi^2 + y_\xi^2}} & \frac{1}{2c\sqrt{x_\xi^2 + y_\xi^2}} \end{pmatrix}. \quad (3.39)$$

To represent the derivatives normal to the computational boundary in terms of the corrected characteristic amplitudes, equation (3.34) is rearranged as

$$-\frac{\partial Q^*}{\partial t} + S_\eta \mathcal{L}' + C_\eta = 0, \quad (3.40)$$

where

$$S_\eta \mathcal{L}' = \begin{pmatrix} -\frac{1}{c^2} & -\frac{w}{c^2} & -\frac{ux_\xi + vy_\xi}{c^2} & -\frac{\sqrt{x_\xi^2 + y_\xi^2}}{c} & \frac{\sqrt{x_\xi^2 + y_\xi^2}}{c} \\ 0 & 0 & x_\xi & y_\xi & y_\xi \\ 0 & 0 & y_\xi & -x_\xi & -x_\xi \\ 0 & 1 & 0 & 0 & 0 \\ 1 & 0 & 0 & V & V \end{pmatrix} \cdot \begin{pmatrix} \mathcal{L}'_1 \\ \mathcal{L}'_2 \\ \mathcal{L}'_3 \\ \mathcal{L}'_4 \\ \mathcal{L}'_5 \end{pmatrix}. \quad (3.41)$$

3.4 Summary

The governing equations for the forward and adjoint simulations were introduced in the current chapter. The reasons behind the election of a continuous adjoint approach were mentioned in section 3.2. To be able to perform adjoint simulations of open flows, the derivation of the adjoint equations in characteristic form was shown in section 3.3. Additionally, this will also be useful to derive the characteristic interface conditions (see section 4.2.4).

Chapter 4

Computational Setup

This chapter describes the computational setup used in the simulations carried out in the present research. Special emphasis is put on the numerics related to the newly developed adjoint framework, which has been implemented alongside an existing in-house DNS code. This results in a unique framework capable of dealing with high fidelity optimal flow control simulations, and also adjoint-based stability analysis for compressible flows in complex geometries.

4.1 Direct Numerical Simulations

The fact that a general analytical solution for the set of governing equations ((3.1) to (3.3)) is not yet known, forces us to discretise them in space and time and use numerical methods that approximate the exact space-time solution. Since we intend to perform a DNS, a fine space-time discretisation of the governing PDEs is required to resolve all the spatiotemporal events within the flow starting from a given initial condition. Thus, high-accuracy numerical schemes are required due to the broad range of wavelengths present in the flow.

4.1.1 Numerical Method

The software used to carry out the numerical simulations is a parallel in-house FORTRAN code called HiPSTAR (High-Performance Solver for Turbulence and Aeroacoustics Research, [Sandberg, 2015](#)), which solves the governing equations by using generalised Cartesian coordinates. This code assumes periodicity in the spanwise direction, hence this direction is discretised using Fourier transformations (FFTW library, [Frigo and Johnson, 2005](#)). The streamwise and vertical¹ directions are discretised with an explicit fourth-order accurate standard central finite difference scheme. At the boundaries, one-sided explicit schemes are used ([Carpenter et al., 1999](#)), being as accurate as the interior scheme. The solution advances in time with an explicit ultra-low storage

¹The vertical direction is also referred to as wall-normal direction.

Chapter 4 Computational Setup

Runge-Kutta method, which achieves fourth order accuracy with a five step method (Kennedy et al., 2000). To enhance the stability of the numerical scheme, a skew-symmetric splitting of the non-linear terms present in the forward equations is applied (Kennedy and Gruber, 2008). A sixth-order accurate explicit filter (Bogey et al., 2009) with a weighting of 0.2 is applied after every Runge-Kutta cycle to remove the spurious numerical high-wavenumber oscillations. The parallelisation in the streamwise and vertical directions is achieved with MPI communication between the neighbouring processes. To allow the entire spanwise domain to be available for the FFT, OMP parallelisation in this direction is used.

4.1.2 Grid

The grid generation is one of the key points to achieve correct results in the simulation with an optimum performance. To make sure that there is a sufficient number of points across any structure and the numerical scheme is able to resolve it, a correct estimation of the wavelength of the scales present in the flow has to be done before the grid is generated. The size of the scales is strongly linked to the Reynolds number of the case (Pope, 2000), and the geometry in particular, which will determine the areas that are more sensitive to the grid size. For this reason, the grid generation is case-dependent.

HiPSTAR only allows structured (Cartesian) grids to discretise the flow field. To avoid unaffordable and extremely costly grids, clustering points in the streamwise and vertical directions is possible. Hence, the resolution of the domain can be kept high in the area we are interested in, while placing the boundaries of the domain far away without this being an issue cost-wise. Once the grid is read in, a coordinate transformation is carried out (see section 3.3) to an uniform-spaced grid to not degrade the performance of the numerical scheme. When dealing with complex geometries, there is no other option than using a multi-block setup to satisfy the grid generation requirements of the software used. Due to this complexity, is not always possible to make the grid stretching continuous (up to second order) to communicate the blocks through halo points, and compute the spatial derivatives at the interfaces with interior schemes. Hence, characteristic interface conditions (CIC - Kim and Lee, 2003) are available, where metric discontinuities are permitted allowing a greater flexibility to generate the computational grid.

4.1.3 Initial and Boundary Conditions

Once the case is discretised, the starting point of the simulation is the initial condition imposed, from where the governing equations advance the flow solution in time as far as required. The ideal initial condition would satisfy the set of governing equations, but such scenario is generally unknown. Hence, a common practice is to use a theoretical approximation to this initial condition which does not necessarily need to satisfy the governing equations. If this is the case, a buffer time must be allowed for the equations to drive this initial condition through the transient

towards a physical flow solution². Note that the initial condition must be close enough to a physical flow-field to prevent the simulation from crashing. For unstable and complex simulations (e.g. flow over a compressor blade), the election of the right initial condition is a critical point, where a poor initial condition might lead to an undesirable flow behaviour (in the compressor blade case, the blade might stall leading to a reverse flow).

To avoid the flow quantities to drift and become non-physical as the simulation marches forward in time, it is required to restrict the flow behaviour at the boundaries by setting appropriate boundary conditions. There are three main types of boundary conditions. Dirichlet conditions impose the actual value of a given quantity at the boundary, Neuman conditions fix the normal spatial derivative of the quantity and the characteristic type which restricts the time derivative.

4.1.3.1 Isothermal no-slip wall

The necessary conditions required to impose an equivalent behaviour as for an isothermal no-slip wall are

$$u|_w = v|_w = w|_w = 0, \quad T_e|_w = const, \quad (4.1)$$

where density is the only unconstrained variable at the wall.

4.1.3.2 Free-stream Boundary Condition

One of the principal challenges when dealing with compressible flows is the modelling of external or open boundary conditions. The flow state at those boundary locations is generally unknown and time varying, which complicates the use of Dirichlet or Neuman conditions there. When large non-linear flow structures cross these boundaries, they might produce spurious reflections which propagate back to the flow-field, acting as a noise source. For this reason, in an attempt to eradicate or at least mitigate these reflections, characteristic boundary conditions (CBC - [Thompson, 1987, 1990](#)) are used for this sort of boundaries. In particular, the free-stream boundaries are modelled as ‘perfectly non-reflecting’, which set to zero the amplitude of every characteristic wave entering the domain. Additionally, to further reduce the impact of these boundary conditions, they must be placed as far as possible from the region of interest, stretching the grid size towards the boundary to increase the numerical dissipation.

4.1.3.3 Outflow Boundary Condition

In contrast with the above-mentioned free-stream boundary condition, the outflow boundary is the one with a higher number of waves travelling outwards through it. Hence, the use of

²A common practice is to allow at least two domain ‘flow-through’ periods before regarding the flow solution as physical.

characteristic boundary condition on that boundary might not be enough to avoid all reflections and a buffer zone can be used in addition to the CBC (Sandberg and Sandham, 2006). These buffer zones use a ramping function that damps the strong fluctuations in the near boundary area, which means that the original flow field is progressively distorted towards the boundary. Additionally, in the case of a three-dimensional subsonic flow, the characteristic decomposition dictates that there are four waves leaving the domain and one travelling inwards. As indicated by Poinso and Lele (1992), if the same conditions are applied as for the free-stream, this could originate a pressure drift. To avoid this undesired behaviour, the pressure in the infinity is imposed and linked to the pressure at the outflow boundary in such a way that, pressure waves will enter the domain through the outflow to maintain the mean value of pressure close to the chosen value at the infinity (Kim and Lee, 2000). Note that this will not be the case for supersonic flows, where at the outflow, all the characteristic waves leave the domain.

4.1.3.4 Laminar Inflow Boundary Condition

A steady two-dimensional boundary layer is generated at the inflow boundary computing a Blasius solution for the laminar boundary layer (Blasius, 1908). This requires choosing beforehand a boundary layer origin and Reynolds number to determine the boundary layer profile. To avoid the above-mentioned reflections, this time for upstream-travelling waves, the inflow profile is integrated over time alongside with a CBC (Jones, 2008).

4.2 Adjoint Simulations

The adjoint framework has been developed and implemented within HiPSTAR, which permits sharing an important part of the computational setup, such as derivative routines or the parallelisation, both introduced in subsection 4.1.1. The adjoint equations are implemented in generalised Cartesian coordinates on the $\xi - \eta$ plane. Similarly to the forward equations, the spanwise direction is restricted to be periodic. Additionally, it is still necessary to address the storage problem associated with non-linear time-dependent adjoint simulations and as well the derivation of the adjoint boundary conditions for the present formulation.

4.2.1 Storage Saving Scheme

As indicated in chapter 3, the forward state variables are present in the adjoint equations due to the non-linear nature of the forward operator. This implies that the time-dependent forward state must be available during the adjoint computations. Hence, if the forward state were stored every sub-step of the Runge-Kutta scheme for every time-step, this would turn into a considerable storage problem if the case were large in size. Fortunately for the present investigation,

there are several techniques to lighten the storage requirements. The method known as ‘check-pointing’ consists in storing instantaneous snapshots every given number of time-steps. During the adjoint simulations, several forward simulations are re-started on demand from the captured snapshots to fill the gap between the snapshots, storing all the intermediate flow states in memory, which are then available to the adjoint solver. This process is repeated sequentially and has the shortcoming that the forward state has to be computed twice. Less computationally intensive techniques are based on the fact that DNS are over-resolved in time due to stability requirements of the time marching scheme, permitting a lower sampling frequency without losing relevant flow information. The ‘frozen coefficient’ technique (also known as zero order hold) updates the forward state every given number of time-steps, leaving them unchanged in between updates. In the present investigation, we use a slightly more advanced method, which linearly interpolates the forward state between the snapshots. Similarly to the numerical schemes that we use to compute derivatives, the order of the interpolation could be increased as required, which could potentially allow an even lower sampling frequency of the forward state.

4.2.2 Control Update

If performing flow control, it is required the use of an optimisation algorithm which updates the control parameters, using the given gradient of the functional computed during the adjoint simulation. Such kind of methods are known as gradient-based optimisation algorithms and can be written as

$$\phi^{k+1} = \phi^k - \alpha^k (\nabla \mathcal{J})^k, \quad (4.2)$$

where k represents the iteration number and α determines the size of the control update for the given $\nabla \mathcal{J}$ direction. The above method is known as steepest descent method and is the simplest algorithm from the gradient-based family. One of the main disadvantages of this method is that the convergence might be very slow in the presence of ‘narrow valleys’ in the cost functional. In such situations, the conjugate gradient algorithm offers a much better performance. This method combines the direction of maximum descent with a ‘momentum’ term that accounts for the update directions from previous steps. One of the most popular methods used in previous studies based on adjoint optimisation (e.g. [Bewley et al., 2001](#); [Schulze et al., 2011](#)) is the *Polak-Ribiere* variant of the conjugate gradient method ([Polak and Ribiere, 1969](#)).

An alternative to the above algorithms resides in the use of higher order methods based on Newton’s method, which use higher derivatives to acquire information about the curvature of the cost function, achieving a more direct path to the minimum. On a generalised form, these methods can be expressed as

$$\phi^{k+1} = \phi^k - \left[H \mathcal{J} \left(Q, \phi^k \right) \right]^{-1} \nabla \mathcal{J} \left(Q, \phi^k \right), \quad (4.3)$$

where $H \mathcal{J} \left(Q, \phi^k \right)$ represents the Hessian matrix of the cost function. For this approach, the increase in accuracy comes along with a considerable increase in computational cost, mainly

caused by the inversion of the Hessian matrix. Especially for large dimensional problems, it is convenient the use of algorithms based on Quasi-Newton methods, where the inverse of the Hessian matrix is estimated instead of computed. From this family of methods, the Broyden–Fletcher–Goldfarb–Shanno (BFGS) (Nocedal and Wright, 2006) algorithm is one of the most widely used. This particular method has also a limited memory version that permits box constraints (L-BFGS-B) (Morales and Nocedal, 2011), which allows handling optimisation problems with a large number of parameters. This algorithm has been proven more efficient than the conjugate gradient (Badreddine et al., 2014), and is the one used in the present investigation.

4.2.3 Boundary Conditions

The derivation of appropriate boundary conditions (temporal and spatial) for a continuous or ‘differentiate then discretise’ adjoint approach starts by accounting the forward boundary conditions (BC) in the Lagrange functional. This will introduce additional adjoint variables (ψ) which are only defined at the boundaries and they are uncoupled from the adjoint governing equations

$$\mathcal{L}(Q, Q^*, \phi) = \mathcal{J}(Q, \phi) - Q^{*\text{T}} \cdot \mathcal{N}(Q, \phi) - \psi^{\text{T}} \cdot BC(Q, \phi). \quad (4.4)$$

Hence, additional boundary terms are also introduced, closing the algebraic system of equations that the boundary terms form, whose solution is the adjoint boundary conditions. For the sake of clarity, the boundary terms and the detailed derivation of the boundary conditions shown below is included in appendix B.

4.2.3.1 Terminal Condition

As the adjoint equations are posed backwards in time, it is required an adjoint initial condition or ‘terminal condition’ from where to start the simulation. If no specific term is included in the Lagrangian functional to impose a given terminal condition, the adjoint variables are initialized as $\rho^*|_{t=T} = 0$, $\vec{m}^*|_{t=T} = 0$ and $p^*|_{t=T} = 0$, what is known as the ‘null condition’.

For some cases, it can be convenient the enforcement of the terminal condition, by including in the Lagrangian functional a term such as

$$\frac{\gamma_p}{2} \int_{\Omega} (\phi|_{t=T} - \Lambda)^2 d\Omega \quad (4.5)$$

where γ_p is a penalty term and Λ is a given function. According to M. D. Gunzburger (2003), this could improve the behaviour of the control states near the terminal time.

4.2.3.2 Isothermal No-Slip Wall

Substituting (4.1) terms into (4.4) and solving the system of equations of the full viscous boundary terms, leads to the adjoint isothermal no-slip wall condition

$$\vec{m}^*|_w = \rho^*|_w = 0. \quad (4.6)$$

4.2.3.3 Non-Reflecting Characteristic Conditions

Differently from wall boundaries, the inflow, outflow and free-stream boundary conditions can be derived under the assumption that they are located far away from the region of interest, which allows neglecting the boundary terms. Since the equations are available in characteristic form (see section 3.3), these boundaries can be modelled as perfectly non-reflecting by setting the amplitude \mathcal{L} of all incoming waves to zero. As for the forward simulations, this would result in an ill-posed problem where the quantities will drift as the simulation marches forward in time (Poinsot and Lele, 1992). Hence, similarly to the forward simulations, this condition can only be applied to the free-stream boundary.

4.2.3.4 Soft Adjoint Inflow/Outflow Characteristic Conditions

If the terminal condition of the adjoint simulation is the null condition, the far stream or infinity conditions could be modelled as if all the quantities far away from the computational domain are zero. Hence, LODI relations (Poinsot and Lele, 1992) can be derived at the inflow and outflow boundaries, leading to suitable boundary conditions that will prevent the aforementioned drift of the adjoint quantities. The boundary conditions for inflow and outflow as streamwise boundaries follow from equation (3.32), where after assuming the flow at the boundaries to be locally one-dimensional and inviscid the source term C_ξ can be neglected giving the LODI system

$$-\frac{\partial Q^*}{\partial t} + S_\xi \mathcal{L} = 0. \quad (4.7)$$

Chapter 4 Computational Setup

Solving the above system of equations for \mathcal{L} the adjoint LODI relations can be written as

$$\mathcal{L}_1 = \frac{\partial p^*}{\partial t} + U \left(\frac{y_\eta \frac{\partial m_u^*}{\partial t} - x_\eta \frac{\partial m_v^*}{\partial t}}{x_\eta^2 + y_\eta^2} \right) \quad (4.8)$$

$$\mathcal{L}_2 = \frac{\partial m_w^*}{\partial t} \quad (4.9)$$

$$\mathcal{L}_3 = \frac{x_\eta \frac{\partial m_u^*}{\partial t} + y_\eta \frac{\partial m_v^*}{\partial t}}{x_\eta^2 + y_\eta^2} \quad (4.10)$$

$$\begin{aligned} \mathcal{L}_4 = c^2 \left(- (x_\eta^2 + y_\eta^2) \left(c^2 \frac{\partial \rho^*}{\partial t} + \frac{\partial m_u^*}{\partial t} u + \frac{\partial m_v^*}{\partial t} v + \frac{\partial m_w^*}{\partial t} w \right) - \frac{\partial m_u^*}{\partial t} y_\eta \sqrt{c^2 (x_\eta^2 + y_\eta^2)} \right. \\ \left. + \frac{\partial m_v^*}{\partial t} x_\eta \sqrt{c^2 (x_\eta^2 + y_\eta^2)} - \frac{\partial p^*}{\partial t} (x_\eta^2 + y_\eta^2) \right) / \left(2 (c^2 (x_\eta^2 + y_\eta^2))^{3/2} \right) \end{aligned} \quad (4.11)$$

$$\begin{aligned} \mathcal{L}_5 = c^2 \left((x_\eta^2 + y_\eta^2) \left(c^2 \frac{\partial \rho^*}{\partial t} + \frac{\partial m_u^*}{\partial t} u + \frac{\partial m_v^*}{\partial t} v + \frac{\partial m_w^*}{\partial t} w \right) - \frac{\partial m_u^*}{\partial t} y_\eta \sqrt{c^2 (x_\eta^2 + y_\eta^2)} \right. \\ \left. + \frac{\partial m_v^*}{\partial t} x_\eta \sqrt{c^2 (x_\eta^2 + y_\eta^2)} + \frac{\partial p^*}{\partial t} (x_\eta^2 + y_\eta^2) \right) / \left(2 (c^2 (x_\eta^2 + y_\eta^2))^{3/2} \right), \end{aligned} \quad (4.12)$$

where each time derivative can be replaced by

$$\frac{\partial Q^*}{\partial t} = K [Q^* - Q_\infty^*]. \quad (4.13)$$

The constant K determines the strength of the incoming waves, and it can be expressed as

$$K = \sigma_{in/out} (1 - M^2) \frac{c}{L}. \quad (4.14)$$

The value of $\sigma_{in/out}$ is chosen to be 0.25, as recommended by [Poinsot and Lele \(1992\)](#) for the Navier-Stokes equations. These boundary conditions ensure a well-posed system, where for the case of an outflow (adjoint inflow) the values of \mathcal{L}_1 , \mathcal{L}_2 , \mathcal{L}_3 and \mathcal{L}_4 are determined by equations (4.8) to (4.11), and in case of an inflow (adjoint outflow) \mathcal{L}_5 is only modified by (4.12).

4.2.4 Characteristic Interface Conditions

The CIC framework has also been extended to the adjoint equations, which gives the code the special feature of using the adjoint framework for complex geometries without losing numerical accuracy. The idea behind this method is to treat each single block as an individual domain, where boundary conditions are imposed at every block boundary. This includes as well the block interfaces, where the condition at this particular block boundary should satisfy that the time derivatives for both neighbouring blocks are the same. Hence, the interface condition

follows from (3.32) and can be expressed as

$$\mathcal{L}_{b1} + S^{-1}C_{b1} = \mathcal{L}_{b2} + S^{-1}C_{b2}, \quad (4.15)$$

where the different subscripts indicate different computational blocks and the term $S^{-1}C$ is known as the modified source term.

Note that the possible differences in grid stretching approaching the interface from both blocks makes (4.15) to be valid only in physical space (see section 3.3). At this point, the amplitude of the incoming waves follows directly from (4.15), where the incoming wave in one block through the interface corresponds to the outgoing wave in the neighbouring block through the same interface.

4.3 Summary

The numerical treatment for the forward simulations, alongside with the boundary and initial conditions were presented in section 4.1. This is followed by the description of the numerics that are exclusive for the adjoint simulations (section 4.2). The most relevant approaches to deal with the forward flow-field storage problem and the gradient-based update of the control parameters were introduced. Subsection 4.2.3 showed the boundary conditions required for the adjoint simulations, where in subsection 4.2.3.4 the adjoint boundary conditions for both inflow and outflow were derived. These boundary conditions account for far field information to avoid the adjoint variables to drift. The chapter concluded with the derivation of the characteristic interface conditions that enable the code to carry out adjoint simulations in complex geometries without compromising the numerical accuracy.

Chapter 5

Verification

A strict verification of the newly developed adjoint framework is required before any further study is carried out. To evaluate each of the framework's features individually, the testing is broken down into several simple cases. Additionally, the verification is regarded from the perspective of two different methods to ensure a correct behaviour of the adjoint code. These two methods are based respectively on the Green-Lagrange identity (3.12) and a finite difference approximation to the gradient of the functional.

5.1 Green-Lagrange Identity

One possible form of verification for the adjoint equations is through the use of the already mentioned Green-Lagrange identity (3.12), as proposed by Spagnoli and Airiau (2008). Defining the forcing terms of the forward and adjoint systems as f and g respectively, and dropping the boundary terms as they are cancelled out by the boundary conditions, (3.12) can be rewritten as

$$\langle Q^*, f \rangle = \langle g, Q' \rangle. \quad (5.1)$$

Using the definition of the inner product $\langle \alpha_a, \alpha_b \rangle = \int_0^T \int_{\Omega} \alpha_a^T \alpha_b d\Omega dt$, and choosing density from the array of state variables Q , (5.1) can be rewritten as

$$\int_T \int_{\Omega} p^* f d\Omega dt = \int_T \int_{\Omega} g \rho' d\Omega dt. \quad (5.2)$$

Hence, if the forward and adjoint forcing is the same, the time-space integration of the perturbations of the given state variable and its adjoint must be the equal. Note that this will only hold for closed domains or if both forward and adjoint systems are in conservative form as the boundary terms have been neglected in the derivation of the CBC.

Chapter 5 Verification

As suggested by [Moret-Gabarro \(2009\)](#), a time-periodic Gaussian pulse in the centre of the domain is chosen as forcing for both forward and adjoint¹ simulations

$$f = g = 0.01 \sin\left(\frac{2\pi}{100}t\right) W_a(x, y), \quad (5.3)$$

where $W_a(x, y)$ is the spatial shape of the forcing, which follows a Gaussian distribution

$$W_a(x, y) = e^{-\frac{(x-x_0)^2 + (y-y_0)^2}{0.005}}. \quad (5.4)$$

Note that if the adjoint variables are initialised with the null condition, $Q^* = Q^{*'}$ holds.

5.1.1 Multi-block Periodic Case

The first verification test case consists of a two-dimensional six-block setup with periodic boundary conditions in x and y directions, where the block interfaces and periodic boundary conditions are treated as characteristic interface conditions. The domain size is 1 length unit in the vertical direction and 4 length units in the horizontal direction, with 200 and 800 uniformly spaced grid points respectively. The multi-block distribution consists of 3 and 2 blocks of the same size in the x and y directions. The forward case is initialised as a laminar free-stream flow and both simulations run for ten complete forcing periods, with a time step of $5 \cdot 10^{-4}$, which is long enough for the disturbances to go through the periodic boundaries. The Reynolds number is 500 and is based in one length unit, and the Mach number is 0.5. As both simulations run in opposite time directions, the flow patterns keep a quasi-symmetric behaviour with respect to the forcing location² (see figure 5.1).

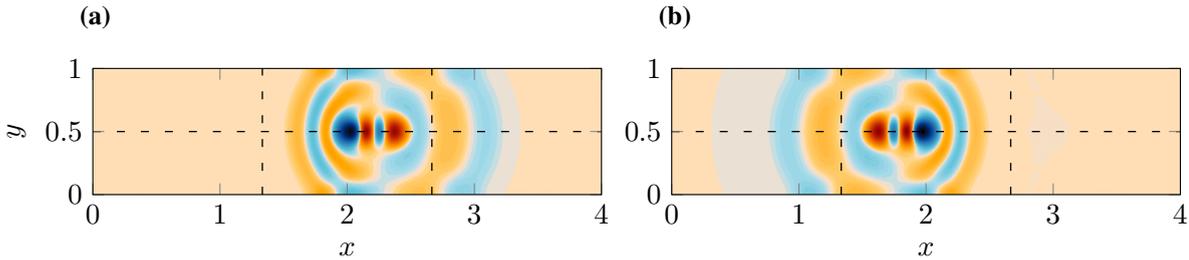


Figure 5.1: Left and right figures show contours of ρ' and p^* respectively, both after 2 forcing periods. Contour ranges from $4.5 \cdot 10^{-4}$ to $-4.5 \cdot 10^{-4}$. The dashed lines represent the location of the block characteristic interfaces.

The verification test has been repeated twice for this case using always the forcing term indicated in (5.3), forcing density (ρ) and adjoint pressure (p^*) in the first test, and u-momentum (ρu) and adjoint u-momentum (m_u^*) in the second one. Since the forcing applied in both forward

¹Note that for the adjoint simulation time is marching backwards in time, starting from the final time of the forward simulation.

²Note the different sign of the convective terms from the Navier-Stokes equations and the Adjoint Navier-Stokes equations (see chapter 3)

and adjoint simulations is the same, the spatial integration should match at every time step. As shown in figure 5.2, the agreement is excellent for both p^* and m_u^* , so (5.1) is satisfied, proving the correct behaviour for the adjoint governing equations and the CIC.

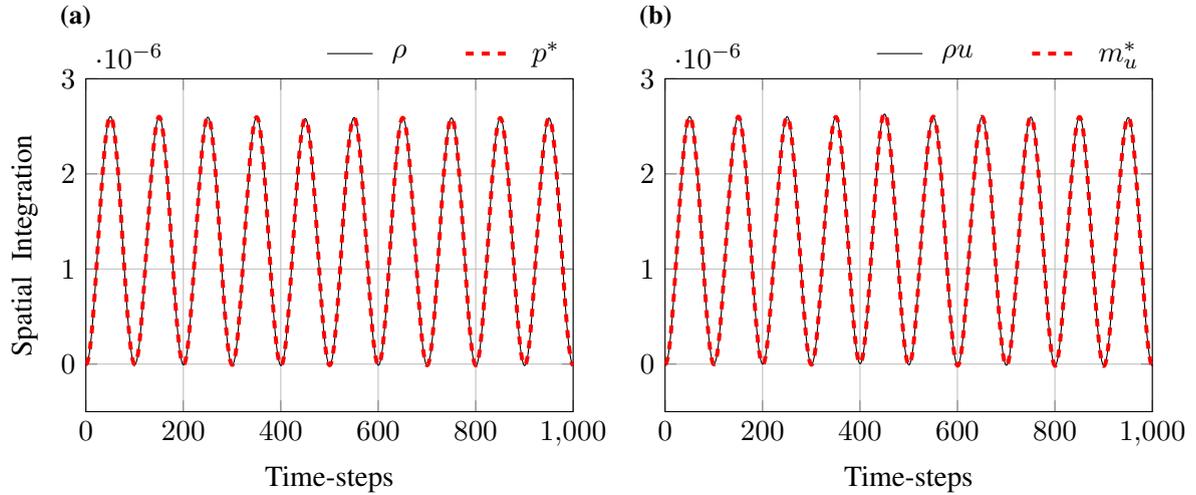


Figure 5.2: Spatial integration as a function of time for ρ' and p^* (a), and $(\rho u)'$ and m_u^* (b).

5.1.2 Laminar Free-stream Case

The second verification test case is a single-block setup, where the dimensions of the domain are 4 length units in x and y directions, with 800 points on each direction. The Reynolds number and Mach number are 500 and 0.5 respectively, where the x and y boundaries have been set as soft inflow/outflow and characteristic non-reflective boundaries, respectively. The initialisation of the forward simulation is as well a laminar free-stream flow, and both simulations run for ten complete forcing periods with a time step of $1 \cdot 10^{-3}$.

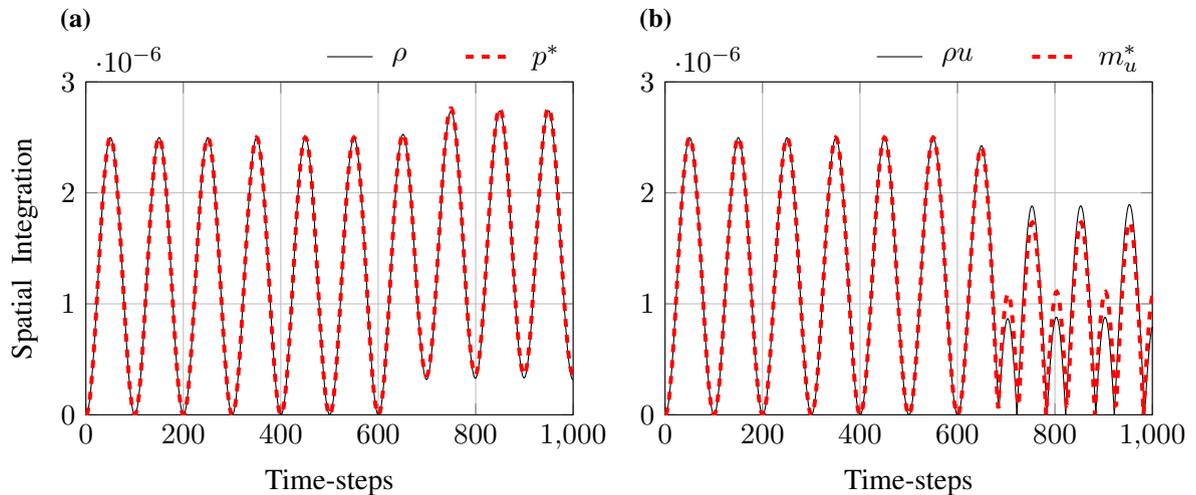


Figure 5.3: Spatial integration as a function of time for ρ' and p^* (a), and $(\rho u)'$ and m_u^* (b).

The same procedure for verification as for the previous case has been followed, forcing ρ and p^* , and ρu and m_u^* . From figure 5.3, it is observed that the spatial integrations suffer an amplitude shift (more pronounced when forcing ρu and m_u^*) as the waves start to cross the computational boundaries. Nevertheless, this behaviour was expected since the domain for the present case is open and the adjoint system of equations is not in conservative form. Note that for (5.1) to be satisfied, it would imply that $Q' = Q^*$, which is not true. Hence, this verification method is not suitable for open flows.

5.2 Finite Difference Gradient

An alternative method to verify the adjoint equations relies on the comparison of the gradient of the cost functional ($\nabla \mathcal{J}$) obtained using the adjoint method (2.11) with the one from a finite difference scheme (e.g. Schulze et al., 2011). If a first order forward finite difference scheme is chosen, the approximate gradient can be expressed as

$$\nabla_{\phi_n} \mathcal{J} \approx \frac{\mathcal{J}(\phi_n + \delta\phi_n) - \mathcal{J}(\phi_n)}{\delta\phi_n}, \quad (5.5)$$

where n is the control parameter index and $\delta\phi_n$ is a relatively small perturbation to the control parameter ϕ_n . To use the above equation it is required to establish a flow control problem with its corresponding cost functional. This expression implies that each of the control parameters has to be perturbed individually, which dictates that it is necessary to carry out $n_{cp} + 1$ function evaluations, being n_{cp} the number of control parameters.

5.2.1 Free-stream Noise Cancelling Case

As a test case, a simple two-dimensional noise-cancelling problem is defined, where the domain has the same dimensions as specified in subsection 5.1.1. The boundary conditions in the x direction are soft inflow/outflow CBC and free-stream CBC in the boundaries in y direction. The grid resolution consists in 400 and 100 uniformly spaced points in x and y directions respectively, and the time step employed for the current case is $1 \cdot 10^{-3}$. The flow is initialised as a laminar free-stream flow with $Re = 500$ and $M = 0.5$, where the noise source f is a time-periodic volume force in the streamwise direction located in the centre of the domain ($x_0 = 2$, $y_0 = 2$) defined as in (5.3). The sensing region $W_s(x, y)$ is located 1 length unit further downstream from the centre of the domain ($x_0 = 3$, $y_0 = 0.5$) and has the same spatial distribution as the noise source $W_a(x, y)$. The control is modelled as a time dependent volume force $\phi(t)$, which is initially turned off and is superimposed with the noise source, having the same spatial distribution $W_a(x, y)$. Hence, adding the contribution of the forcing terms to the governing

equations, the streamwise momentum and energy equations are rewritten as

$$\frac{\partial \rho u}{\partial t} = RHS + \rho f W_a(x, y) + \rho \phi(t) W_a(x, y) \quad (5.6)$$

$$\frac{\partial \rho E}{\partial t} = RHS + \rho u f W_a(x, y) + \rho u \phi(t) W_a(x, y). \quad (5.7)$$

To achieve a noise reduction with respect to the uncontrolled case, a convenient cost function that penalizes the density fluctuations at the sensing region can be written as

$$\mathcal{J} = \int_T \int_{\Omega} \frac{1}{2} (\rho - \rho_0)^2 W_s(x, y) d\Omega dt, \quad (5.8)$$

where ρ_0 is the value of the non-perturbed density for the non-dimensionalised Navier-Stokes equations. Including the cost function and the forcing terms in the Lagrangian functional and then substituting them into (3.11) leads to their equivalent adjoint terms, which for the current case they can be written as source terms of the adjoint pressure equation (3.15)

$$-\frac{\partial p^*}{\partial t} = RHS - \underbrace{(\rho - \rho_0) W_s(x, y)}_{\text{Cost Function}} + \underbrace{m_u^* f W_a(x, y) + m_u^* \phi(t) W_a(x, y)}_{\text{Forcing Terms}}. \quad (5.9)$$

Note that the forcing terms in (5.7) are cancelled out as the total energy equation is transformed into the pressure equation by subtracting the kinetic energy equation (3.9) from the total energy equation (3.3).

The flow actuation or control is defined as a time-dependent signal, which implies that the number of control parameters or degrees of freedom of the problem increases proportionally with the time duration of the simulation. For the present case, the simulation runs for ten complete forcing periods of the noise source (1000 time-steps), which is long enough for the disturbances to travel through the characteristic boundaries. Using (5.5) to approximate the gradient at each time-step requires a total of $1000 + 1$ function evaluations. The perturbation chosen for the controller is $\delta\phi = 0.01$, which corresponds to the maximum amplitude of the noise source. Figure 5.5 shows the comparison between the exact gradient of the functional provided by the adjoint

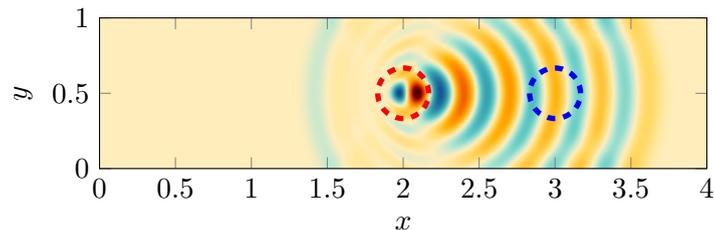


Figure 5.4: Illustration of the case setup. The red dashed circle represents the location of the noise source and controller, whereas the blue dashed circle represents the location of the sensing region. The contours show pressure fluctuations which range from $-0.5 \cdot 10^{-4}$ to $0.5 \cdot 10^{-4}$.

Chapter 5 Verification

simulation and the approximated gradient using (5.5), where the agreement of both approaches is excellent. The current setup can be considered as an extreme case since the top and bottom boundaries are relatively close to the noise source, which will potentially introduce reflections as the waves travel outside the domain. Ideally, the boundaries should be placed further away from the region of interest, preferably stretching the grid resolution to dissipate the outgoing waves and reduce the spurious reflections.

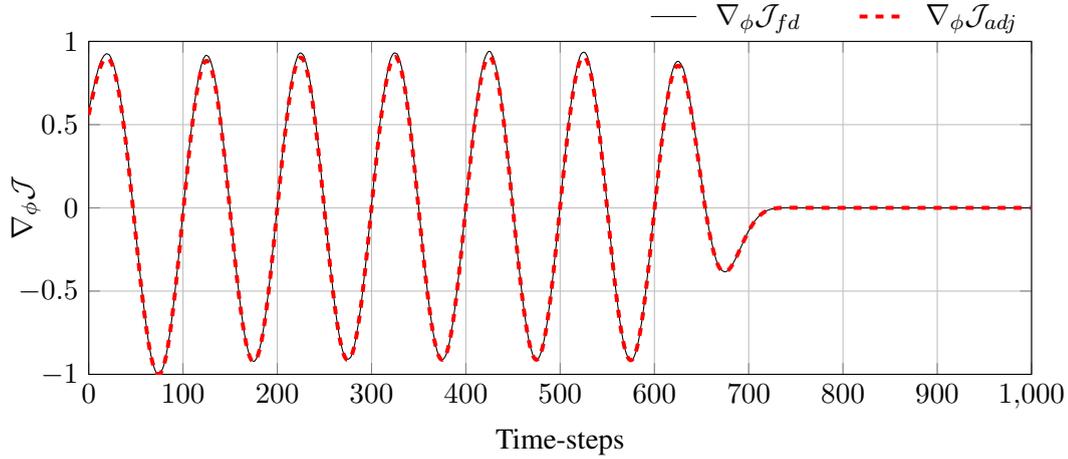


Figure 5.5: Time dependent gradient of the cost functional obtained with a finite difference approach ($\nabla_{\phi} \mathcal{J}_{fd}$) and adjoint method ($\nabla_{\phi} \mathcal{J}_{adj}$).

5.2.2 Laminar Boundary Layer Noise Cancelling Case

To conclude the verification of all the adjoint boundary conditions introduced in chapter 4, a noise cancelling problem on a laminar boundary layer can be regarded as a strict case to ensure the correct behaviour of the adjoint framework under wall-bounded flows. The case setup is the same as for the free-stream noise cancelling case, except for the lower boundary condition in the y direction, which has been replaced by an isothermal no-slip wall. The flow is initialised with a Blasius solution for a laminar boundary layer with its origin placed 20 length units upstream from the inflow boundary, and the time step has been increased to $2 \cdot 10^{-3}$.

The simulation is initially run for 1000 time steps to minimise the effect from the initial solution, preventing any spurious behaviour, where the noise source is always switched on. Similarly to the previous case, the exact and approximated gradient of the functional is computed with the adjoint simulation and a finite difference approach for the following 1000 time steps, which requires again of $1000 + 1$ forward simulations. As shown in figure 5.6, the agreement of both approaches is satisfactory, which concludes the successful verification of the adjoint framework.

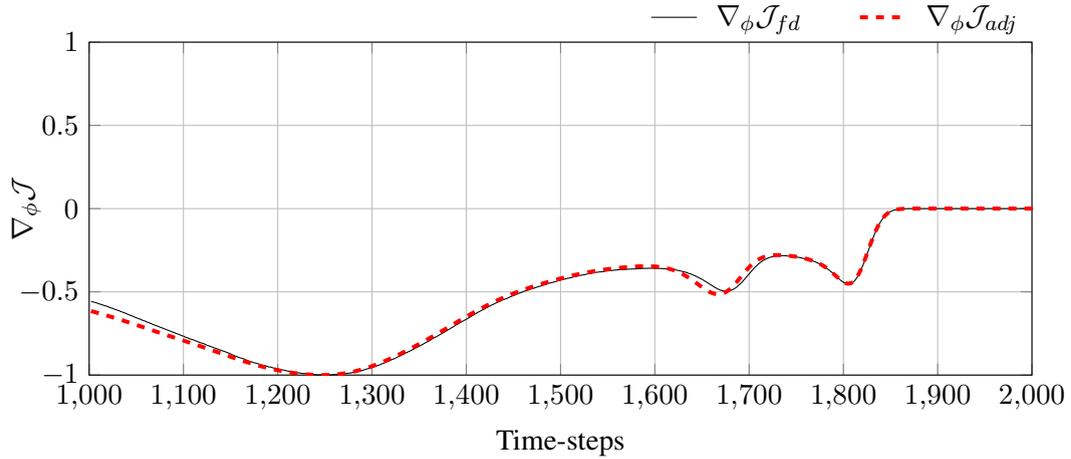


Figure 5.6: Time dependent gradient of the cost functional obtained with a finite difference approach ($\nabla_{\phi} \mathcal{J}_{fd}$) and adjoint method ($\nabla_{\phi} \mathcal{J}_{adj}$).

5.3 Summary

The verification of the novel adjoint framework was broken down into several simple test-cases. The CIC were successfully tested using the Green-Lagrange identity (3.12), where this method was shown to be valid only for closed flows. For this reason, all the characteristic boundary conditions, and also the no-slip wall condition, were verified comparing the adjoint gradients to a finite difference approach. Furthermore, an additional, and perhaps more thorough verification case is shown in chapter 7, where forward and adjoint-based stability analysis of a two-dimensional cavity basic flow yield the same leading eigenvalues.

Chapter 6

Periodic Orbits & Equilibria in a 2D Cavity Flow

This chapter reports for the first time a family of exact periodic solutions in compressible flows with the added difficulty of a complex geometry, which is the leading contribution from this chapter. In addition, we also present a novel framework which permits the computation of such solutions in any other complex geometry. Furthermore, a family of equilibrium flow solutions associated to the periodic orbits is also introduced, where it is shown that both families meet at the quasi-incompressible regime, forming a subcritical Hopf bifurcation (see also chapter 7). This proves the destabilising effect of compressibility in cavity flows, which is analysed in detail, covering an existing gap in the literature. Moreover, the flow-acoustic interaction present in cavity flows is generally overlooked by literature, where here we provide a thorough description of its evolution across Mach number. For the aspect ratio used, these results extend the data available in the literature of open cavity flows.

6.1 Introduction

The study of cavity flows has been tackled by numerous researchers for several decades. In particular, the constant flow separation and reattachment make this special case very interesting from a fundamental research perspective in fluid dynamics. Additionally, the complexity of the interior dynamics can be increased almost as desired by just raising the Reynolds number. The limited computational resources available to the early open cavity flow investigators forced them to focus exclusively on the flow mechanisms from an incompressible flow perspective. Hence, insights into the compressible events present in cavity flows were only accessible through experimental research, which was mainly motivated by military applications (aircraft bomb bays and so on). Precisely in one of these military motivated investigations, [Rossiter \(1964\)](#) documented and studied in detail for the first time the self-sustained flow oscillations of compressible origin, which are commonly known as Rossiter modes. The origin of these periodic events resides in the

growth of Kelvin-Helmholtz instabilities which grow along the separated shear layer, impinging onto the cavity's trailing edge. This flow impingement radiates an acoustic wave which also propagates upstream and, due to the high receptivity of the leading edge (see later chapter 7), it fuels further the appearance of new shear layer instabilities. Far from just being responsible for triggering these new instabilities, this flow-acoustic interaction is a Mach number dependent and synchronised process which governs the overall sound directivity in cavity flows. These complex phenomena are usually overlooked by literature and later on we analyse in detail. One of the first relevant computational studies which simulated two-dimensional cavity flows using compressible DNS was carried out by [Rowley et al. \(2002\)](#). Of special interest, they performed a large parametric study (changing Re_D , M , L/D , etc.) and documented in thorough detail the dynamics in each case. In particular, for low aspect ratio cavities (such as the one investigated herein) they found that the system's dynamics were governed by a shear layer (Rossiter) mode. For higher aspect ratios, the cavity flow is known to abandon the shear layer mode and undergo a transition towards a wake mode type of motion. In a parametric sense, the results shown here extend their database further for cavities of aspect ratio $L/D = 3$.

In addition, this chapter presents a novel framework which permits the computation of steady and periodic exact compressible flow solutions in complex geometries. These type of solutions are also referred to as 'invariant solutions' and 'recurrent' or 'exact coherent' flows. The aim behind the calculation of such exact flow solutions consists in reaching the understanding of complex flow phenomena by breaking down the state-space into simpler flow solutions without the use of any mathematical modelling (i.e., Fourier analysis, POD, DMD, and so on). Following this approach, if such solutions are unstable, highly chaotic systems can be thought of a continuous transition from the proximity of one solution to others. Consequently, these chaotic systems may be approximated as an expansion over periodic orbits (e.g. [Cvitanović et al., 2016](#)). Contrarily, stable exact solutions can act as a flow attractor, where the system continuously decays towards it. Examples of previous investigations searching steady states in incompressible flows used methods such as continuation ([Keller, 1977](#)), selective frequency damping ([Åkervik et al., 2006](#)), and even an adjoint-based approach ([Farazmand, 2016](#)). The numerical computation of exact steady solutions was only expanded to compressible flows by [Yamouni et al. \(2013\)](#), where they used the iterative Newton method for its calculation. On the other hand, this type of work was first carried out with the aim of seeking non-steady states by [Nagata \(1990\)](#), and since then, the search of such flow solutions remained exclusive of incompressible flows. For further details of similar investigations, the reader is referred to the reviews from [Kawahara et al. \(2012\)](#) and [Cvitanović \(2013\)](#). More recently, [Farazmand \(2016\)](#) extended the available methods to obtain steady and travelling wave solutions using an adjoint-based approach. With his framework, he guarantees convergence to an exact solution from any initial condition. Note that this is a remarkable advantage over more classic methods (e.g. Newton-Raphson iterations), which require of a starting point within the proximity of the sought flow solution to yield satisfactorily converged results. Although it was not his main objective, the reader is also encouraged to go through this article for an interesting introduction/review of invariant flow solutions. In the context of this thesis, one potential application of finding these periodic flow solutions, uses

the aforementioned possibility of approximating chaotic systems as an expansion over periodic orbits. Thus, in the future, we may be able to carry out control of chaotic flows by evaluating the sensitivities of the periodic orbits to the actuation. Therefore, this would remove the need of calculating the sensitivities of chaotic systems with the adjoint method, which as discussed earlier in subsection 2.1.1, it presents some major limitations for such scenarios.

In the current investigation we apply our novel framework on a two-dimensional open cavity flow at $Re_D = 2000$ and $M = 0.5$. As shown later on this chapter, this particular flow solution decays very slowly towards a periodic limit cycle. Hence, we use our method to drive the flow state directly to the periodic orbit, skipping this long transient. After this reference periodic solution is computed, we use it as the initial guess to compute the equivalent periodic solution at neighbouring Mach numbers, keeping the Reynolds number constant. The computation of this family of periodic solutions across Mach numbers aims to cover an existing gap in the literature, examining in detail the evolution of the cavity flow dynamics from a strictly compressible flow perspective. Previous studies have tried to understand the behaviour of Rossiter modes in compressible cavity flows with a considerable Reynolds number (e.g. [Keizer et al., 2004](#)). Under such conditions, the convective instabilities dominate and exert a strong modulation over the Rossiter modes, which makes the understanding of these shear layer modes extremely difficult. For example, in the work of [Yamouni et al. \(2013\)](#), the fact that their cavity flow became more unstable with an increasing Mach number appeared as a surprise to them. Hence, in this chapter, we have chosen the Reynolds number below the threshold where convective instabilities commence to appear, permitting us to show the physical mechanisms which are exclusive from compressible flows. In addition, as discussed in further detail later in chapter 7, our setup would present instabilities in the spanwise direction (of convective nature) if the full three-dimensional scenario would be considered ([Meseguer-Garrido et al., 2014](#)). Moreover, these 3D instabilities would exert as well a frequency modulation on the purely 2D shear layer modes ([Brès and Colonius, 2008](#)). For this reason, we restrict our simulations to be two-dimensional, focusing our analysis on the Mach number evolution of the aforementioned shear layer or Rossiter mode. Moreover, all the simulations discussed in this chapter (and also chapters 7 and 8) were carried out using the same domain discretisation, whose convergence was tested up to $Re_D = 5000$ (see appendix D). For further details about the simulations' grid and domain, the reader is referred to chapter 8.

6.2 Numerical Method

The numerical framework used to obtain the exact flow solutions presented in this chapter couples the in-house compressible DNS solver (HiPSTAR) and the optimisation algorithm L-BFGS, both described in chapter 4. Essentially, the process of finding these exact flow solutions boils down to a simple optimisation exercise, where the optimisation parameters are just the state variables at the initialisation $Q(\vec{x}, t_0)$. Starting from an initial condition, the DNS provides the value of a chosen cost function to be optimised, but also the gradients to all the control

parameters (see, for example, subsection 6.2.1). The value of the cost function alongside the gradients are fed to the L-BFGS algorithm, which returns a better estimate of initial state variables. This simple algorithm is then repeated until some stopping criterion is met. The nature of the solution found by the algorithm (steady or time-periodic) depends exclusively on the chosen cost function. The two subsections below gather the cost functions, stopping criterion and other modifications to the algorithm, necessary to successfully obtain both steady and periodic solutions.

6.2.1 Steady Solutions

To find a steady flow solution, we require of a cost function which penalises the change of the flow-field throughout the time integration respect to the initial state. Such function can be written as

$$\mathcal{J}(Q(\vec{x}, t), Q_0(\vec{x}, t_0)) = \int_T \int_{\Omega} \frac{1}{2} [Q(\vec{x}, t) - Q_0(\vec{x}, t_0)]^2 d\vec{x}dt, \quad (6.1)$$

where T is the duration of the time integration and Ω is the computational domain. The reference state or initial state is represented by $Q_0(\vec{x}, t_0)$. As described in chapter 2, following equation (2.2), we obtain the gradients necessary to drive the cost function towards its minimum as

$$\frac{D\mathcal{J}}{DQ_0} = \frac{\partial\mathcal{J}}{\partial Q} \frac{dQ}{dQ_0} + \frac{\partial\mathcal{J}}{\partial Q_0}, \quad (6.2)$$

where, in this particular case, the first term of the right hand side drops out. For further clarity, we could use the adjoint method, also described in chapter 2, to express the above equation. Using the equation (2.11) on the present case gives

$$\frac{D\mathcal{J}}{DQ_0} = -(Q^*)^T \frac{\partial\mathcal{N}}{\partial Q_0} + \frac{\partial\mathcal{J}}{\partial Q_0}. \quad (6.3)$$

Similarly to equation (6.2), the first term of the right hand side cancels out because the initial conditions Q_0 are not present in the governing equations \mathcal{N} , which leads to $\partial\mathcal{N}/\partial Q_0 = 0$. Hence, the term $D\mathcal{J}/DQ_0$ can be expressed as

$$\frac{D\mathcal{J}}{DQ_0} = \int_T \int_{\Omega} -[Q(\vec{x}, t) - Q_0(\vec{x}, t_0)] d\vec{x}dt, \quad (6.4)$$

and gives the gradients of the control parameters. Note that these gradients can be computed straight from the DNS, without any further cost. On a preliminary study, it was found that longer horizons provide ‘smarter’ gradients with further information about how the flow-field evolves away from the initial state. Hence, increasing the length of this horizon resulted in a flow-field matching the initial state for a longer time. Note that this would only be the case when the steady solutions are unstable. Contrarily, stable solutions should remain in the vicinity of the initial state regardless of the horizon, or even get closer to a complete steady flow with time-marching. Finally, the stopping criterion used in the current study consists in the variations of the cost function \mathcal{J} being of the order of the numerical precision used.

The reader should also note that compressible codes are very susceptible to potentially spurious changes in the state variables. This is because our optimisation routine is not being constrained by the governing equations nor discretised with the same numerical scheme. This can potentially lead to non-valid initial conditions resulting from the optimisation software, which would crash our DNS code. Implicitly and initially, the optimisation code is constrained by the discretised governing equations (i.e., the gradients are computed directly from the discretised flow quantities). But in following iterations, the numerics behind the optimisation algorithm might force an inconsistent drift in the new initial state, away from the from a solution of the Navier-Stokes equations discretised with our numerical scheme. If such catastrophic situation would occur, our algorithm would refresh the optimisation algorithm, using as the new ‘initial guess’ the last initial state given by the optimisation algorithm which did not produce the fatal NaNs.

6.2.2 Periodic Solutions

In order to search for a time-periodic flow solution, we use a more unconstrained version of the above cost function for the steady solutions, such as

$$\mathcal{J}(Q(\vec{x}, T), Q_0(\vec{x}, t_0)) = \int_{\Omega} \frac{1}{2} [Q(\vec{x}, T) - Q_0(\vec{x}, t_0)]^2 d\vec{x}. \quad (6.5)$$

Note that in comparison with equation (6.1) we have dropped the time integral, where the cost is now just the difference amongst the initial and final flow states. Following the same reasoning as detailed in the previous subsection, the gradients also follow directly from the DNS, where this time are computed as

$$\frac{D\mathcal{J}}{DQ_0} = \int_{\Omega} -[Q(\vec{x}, T) - Q_0(\vec{x}, t_0)] d\vec{x}. \quad (6.6)$$

Differently from the steady solutions, the time horizon T is not fixed and it also has to be optimised. One possible way to find a good initial estimate of T is seeking a globally periodic pattern in the flow-field with distributed monitor points across the entire flow domain. After performing Fourier transforms to each individual time signals, all monitor points should show an energy peak at a common frequency, where its inverse can be used as the first initial guess of T . After that, to optimise the horizon length, instead of computing a time gradient respect to the cost function and running it through the L-BFGS, this horizon length is adjusted by our algorithm every time the L-BFGS code returns a new iterate. This is achieved by overrunning the horizon T by a minimal amount (in our case it was 100 time-steps), and finding which neighbouring time-step has the minimum cost (see figure 6.1). Note that for this approach to work, the initial time horizon guess must be close to the period from the final flow periodic orbit. The convergence criterion used for these periodic solutions was $\mathcal{J} / \|Q_0\| < 10^{-12}$.

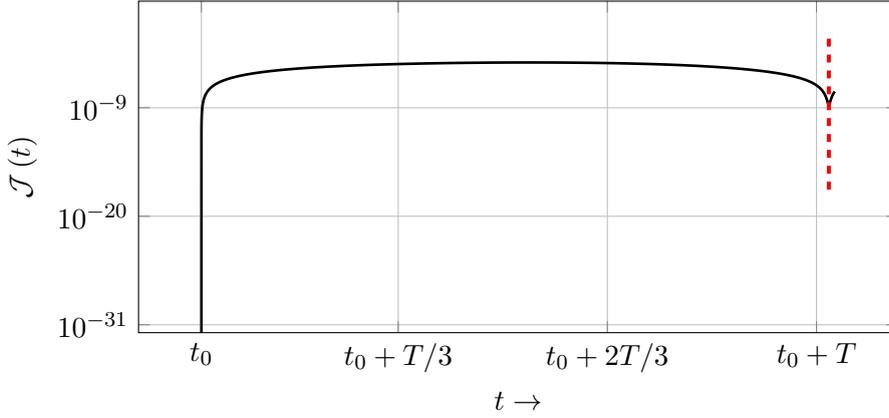


Figure 6.1: Illustration of the time horizon T optimisation. After L-BFGS returns a new iterate, the algorithm updates the old value of T (indicated by the grid line at $t_0 + T$), by the new value (highlighted with the vertical red dashed line).

6.3 Periodic Solution at $M=0.5$

Before delving into the analysis of the evolution of periodic solutions across Mach number, we first describe the periodic orbit found for $M = 0.5$ and $Re = 2000$. This solution, in particular, is the one which was found first and will be used as the main reference when comparing other orbits later in this chapter. The analysis carried out in this section breaks down the periodic trajectory into more fundamental intervals. These splits aim to enclose simpler physical mechanisms which are easier to understand and occur sequentially throughout the period. This will also facilitate the analysis of the differences amongst the several solutions found across Mach number.

To characterise the periodic orbits, we define the norm

$$\|\alpha\| = \int_{\Omega} \alpha(\vec{x}, t)^2 W_i(\vec{x}) d\Omega, \quad (6.7)$$

where α is a space-time dependent flow quantity and Ω is the integration domain. As the flow presents a time-periodic behaviour, the time evolution of this norm will show a closed trajectory of the global (rather than local) behaviour of the flow. Note that the spatial resolution of the simulations is only kept within DNS range in the vicinity of the cavity. Hence, the poor grid resolution near the boundaries of the simulation domain, also in collaboration with the boundary conditions, are likely to produce spurious flow behaviour which would be accounted for in (6.7) if the entire domain were considered. For this particular reason, the spatial function $W_i(\vec{x})$ is included in (6.7), which excludes from the integration these areas which might contain non-physical flow behaviour. $W_i(\vec{x})$ is defined as

$$W_i(\vec{x}) = \begin{cases} 1 & \text{if } \vec{x} \in [(-1.5, 0), (4.5, 4)] \\ 0 & \text{if } \vec{x} \notin [(-1.5, 0), (4.5, 4)] \end{cases}. \quad (6.8)$$

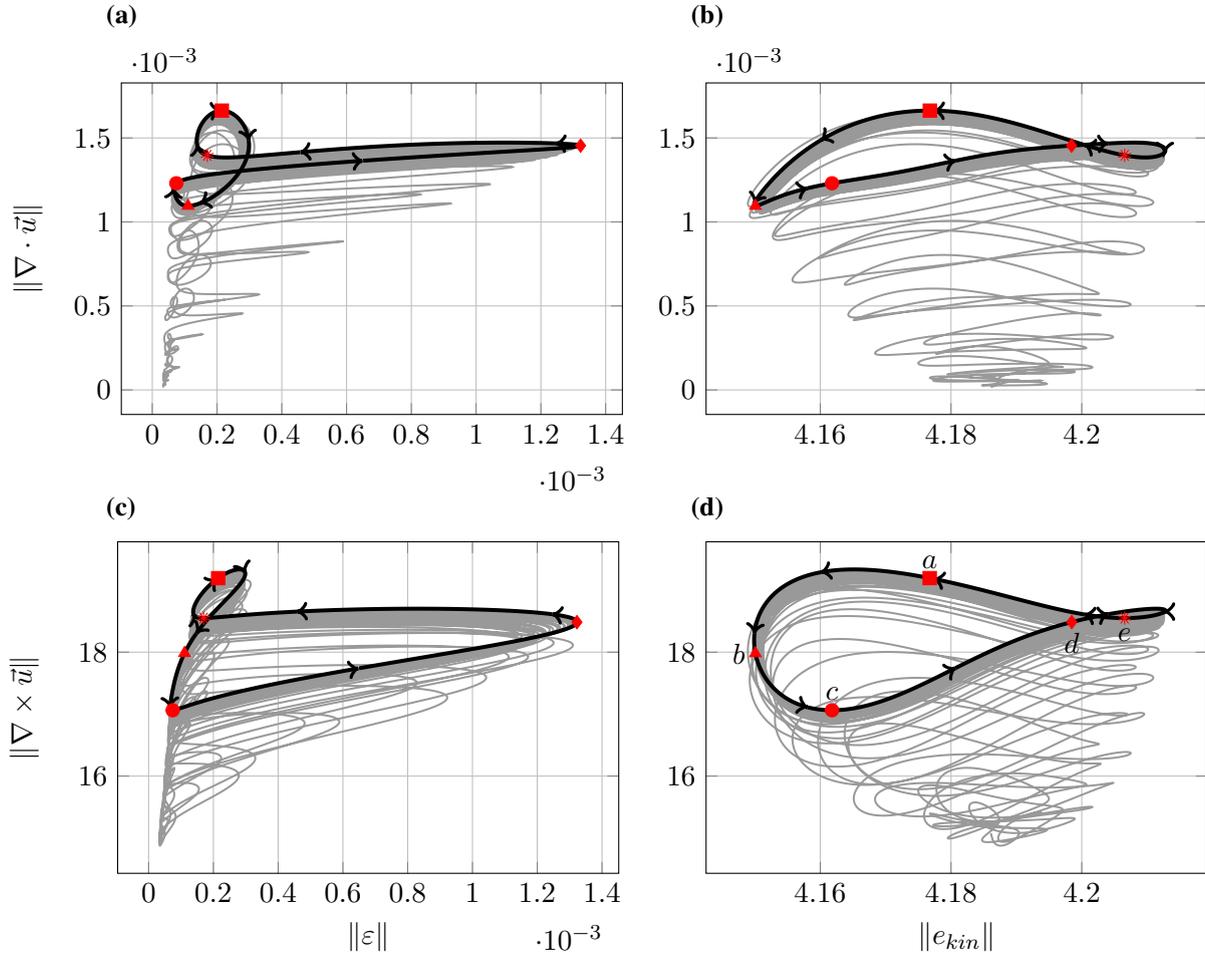


Figure 6.2: 4D representation of the periodic orbit at $M = 0.5$. The trajectory is projected over the norm of dilatation ($\|\nabla \cdot \vec{u}\|$), vorticity ($\|\nabla \times \vec{u}\|$), dissipation rate ($\|\varepsilon\|$) and kinetic energy ($\|e_{kin}\|$). The red marks represent key physical phenomena occurring along the periodic orbit. For the sake of clarity, the symbols are labelled with letters from a to e in (d). The grey line shows the natural flow evolution from the vicinity of the initial condition towards the periodic orbit.

To represent the periodic orbits in their integral form as defined in (6.7), we have chosen dilatation ($\nabla \cdot \vec{u}$), which highlights the compressible events, and viscous dissipation rate (ε - see Kundu et al., 2012), which should emphasise the flow phenomena with strong shear, such as vortex merging. Additionally, we also use vorticity ($\nabla \times \vec{u}$) and kinetic energy (e_{kin}), which are often used in literature to characterise periodic orbits in incompressible flows. Figure 6.2 shows a four-dimensional phase portrait of the periodic solution for $Re = 2000$ and $M = 0.5$, projected onto the four variables mentioned above. The locations of the key physical events occurring in the solution are highlighted with red symbols. These symbols are labelled chronologically from a to e in figure 6.2d, where a is the instant where compressible effects reach their maximum. This point corresponds to the vortex impingement on the trailing edge of the cavity, which will be used as the relative ‘starting point’ of the periodic orbit. Note that the criterion used to choose the origin is arbitrary, and if desired, any other starting point would be equally as

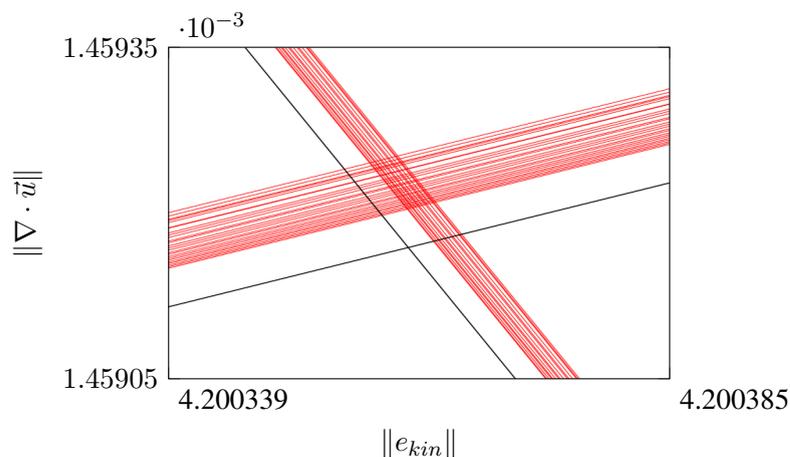


Figure 6.3: Zoomed-in area at the orbit cross-point from figure 6.2b. Black and red lines represent the trajectories of the periodic orbit and fully developed flow, respectively.

valid. Figure 6.2 also shows how the flow-field evolves from shortly after the flow initial condition towards the periodic trajectory. This periodic orbit is an attractor, where for this particular case, it almost represents the complete behaviour of the flow. This follows as a consequence of the relatively low Reynolds number alongside with, as explained further in this section, a strong acoustic shear layer feedback mechanism which stabilises the flow on this periodic behaviour. Hence, the flow lacks the sufficient energy to ‘jump’ to another state and sits indefinitely on the close vicinity of this stable orbit. Figure 6.3 shows a zoomed-in area at the orbit cross-point from figure 6.2b. The trajectory described by the developed flow follows indeed the periodic solution very closely, where the gap to this periodic orbit appears to decay asymptotically¹.

To further illustrate the flow behaviour, snapshots of the density, dilatation and vorticity fields are shown in figures 6.6, 6.7 and 6.8, respectively. In these three figures, the subfigures are also ordered chronologically, where each sub-caption indicates the same time instant as labelled in 6.2d. At the starting point *a*, a low-density spot (high velocity) is located on top of the vortex located in the downstream end of the cavity (figure 6.6a). This vortex was also observed in, for example, Brès and Colonius (2008), and remains in that location throughout the entire periodic orbit. In the present chapter, it will be referred to as the stationary vortex. The above-mentioned low-density area resulted from the merging of an incoming vortex from the shear layer with the stationary vortex. Additionally, another vortex is attached to the leading edge of the cavity (figure 6.8a). This is the point in the periodic orbit where the compressible effects are higher, mostly due to the low-density area impinging on the trailing edge of the cavity. Figure 6.7a shows how the sound radiation is mainly generated in the upstream direction. The physical phenomenon behind this is the cooperation amongst two aligned dipoles in perfect synchronisation (figure 6.4). The strongest dipole is located at the trailing edge of the cavity and its origin is the impingement of the low-density spot on the cavity’s trailing edge. The

¹Note that this figure shows the developed trajectory approximately 440 time units after the flow initialisation. Considering the domain shown in figures 6.6 and assuming a convecting velocity of $0.7U_\infty$, this time span represents a total of 51 flow-through periods.

weaker dipole is caused by the clockwise rotation of the vortex currently attached to the leading edge of the cavity. Note that this dipole interaction partially cancels the sound radiation in the downstream direction.

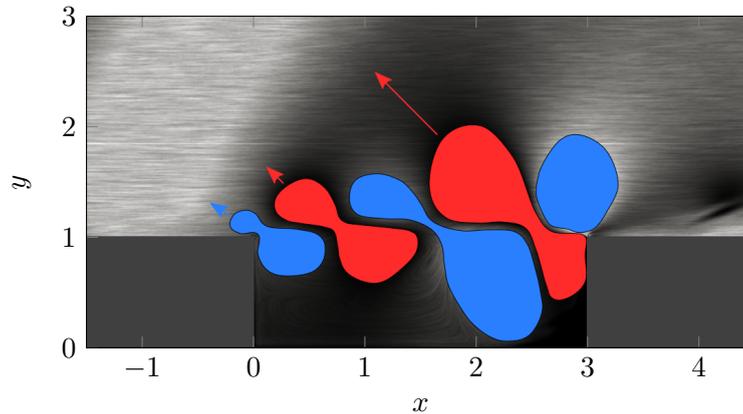


Figure 6.4: Illustration of the trailing edge - shear layer dipole interaction.

a to b The intensity of the impingement decreases rapidly, which causes the trailing edge dipole to dissipate while the shear layer dipole gains in intensity (figure 6.7b). Additionally, the upstream vortex is detached from the leading edge of the cavity. As this vortex moves downstream, a new vortex is formed in the upstream lower corner of the cavity. This vortex rotates counter-clockwise and induces a downwash along the upstream vertical wall of the cavity, which expands the flow as it interacts with the incoming boundary layer. This results into a new dipole attached to the leading edge, which also contributes to the amplification of the low-density acoustic wave reflected from the trailing edge at *a* (figure 6.5). When the flow reaches *b*, dilatation and kinetic are at their minimum values (figure 6.2b). At this point, the trailing edge dipole has vanished, leaving temporarily as leading noise sources the shear layer and leading edge dipoles.

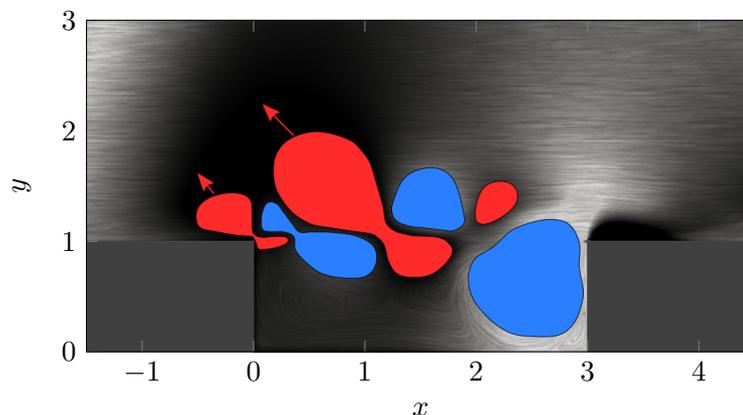


Figure 6.5: Illustration of the leading edge - shear layer dipole interaction.

b to c The vortex in the shear layer keeps moving forward, with its corresponding dipole now decaying in intensity, but still further amplifying the upstream propagating sound wave.

The main contributor to this amplification is now the dipole at the leading edge. The instant b can be seen as the start of a compressible interaction between the shear layer and stationary vortices. Their counter-rotating behaviour creates a high-density spot located amongst the two vortices. Note that the norm of the dilatation field increases due to this phenomenon (figure 6.2a). Also, the norm of vorticity (figure 6.2c) reaches minimum values due to the low interaction of vortical structures with the trailing edge (figure 6.8c).

- c to d* The shear layer and stationary vortices collide and start a merging process. This interaction gains in intensity continuously until the orbit reaches d . This phenomenon is clearly visible in figure 6.2a, where the viscous dissipation rate increases suddenly and reaches maximum values at the instant d . The high-density spot located downstream from the shear layer vortex has grown further and impinges on the trailing edge at d . This impingement reflects a high-density wave that propagates upstream. At the leading edge of the cavity, a new vortex started forming at c due to the suction caused by the vortex at the upstream lower corner of the cavity. Again, the counter-rotation between the leading edge vortex and the shear layer vortex create another high-density spot in between them, which further amplifies the upstream propagating high-density wave. The instantaneous dilatation field shows again two dipoles in the shear layer and trailing edge in perfect synchronisation radiating noise in the upstream direction (figure 6.7d). Note that this time, the dipoles show opposite sign respect to a . Hence, d could be seen as the phase counterpart of a .
- d to e* Soon after d , the core of the shear layer vortex gets absorbed by the stationary vortex. As a result of this merging, the low-density area starts orbiting the stationary vortex. This causes the trailing edge dipole to slowly invert in sign. As the vortex merging completes, the norm of the viscous dissipation rate experiences a sudden drop (figure 6.2a). Meanwhile, the leading edge vortex continues growing in size. This amplifies the shear layer dipole which keeps contributing to the upstream sound radiation. Also, the clockwise rotation of the leading edge vortex pushes the flow upwards along the upstream vertical cavity wall compressing the flow. This phenomenon is observed in the dilatation field (figure 6.7e) as a weak dipole located at the leading edge, which keeps amplifying the upstream travelling high-density wave. Note that this same amplification mechanism occurred for the low-density wave in c (figure 6.5).
- e to a* At this point, the low-density spot starts interacting with the trailing edge as it orbits the stationary vortex. This interaction is also observed in figure 6.2a, where the norm of the dilatation field increases monotonically from e to a . As the low-density area keeps moving downstream, compressibility effects reach again maximum values in a . During this interval, the leading edge vortex has grown in size considerably, up to about half cavity length, just before it detaches from the leading edge again in a . These two physical phenomena occur in perfect synchronisation, leading to the already mentioned double dipole, which radiates the low-density wave upstream.

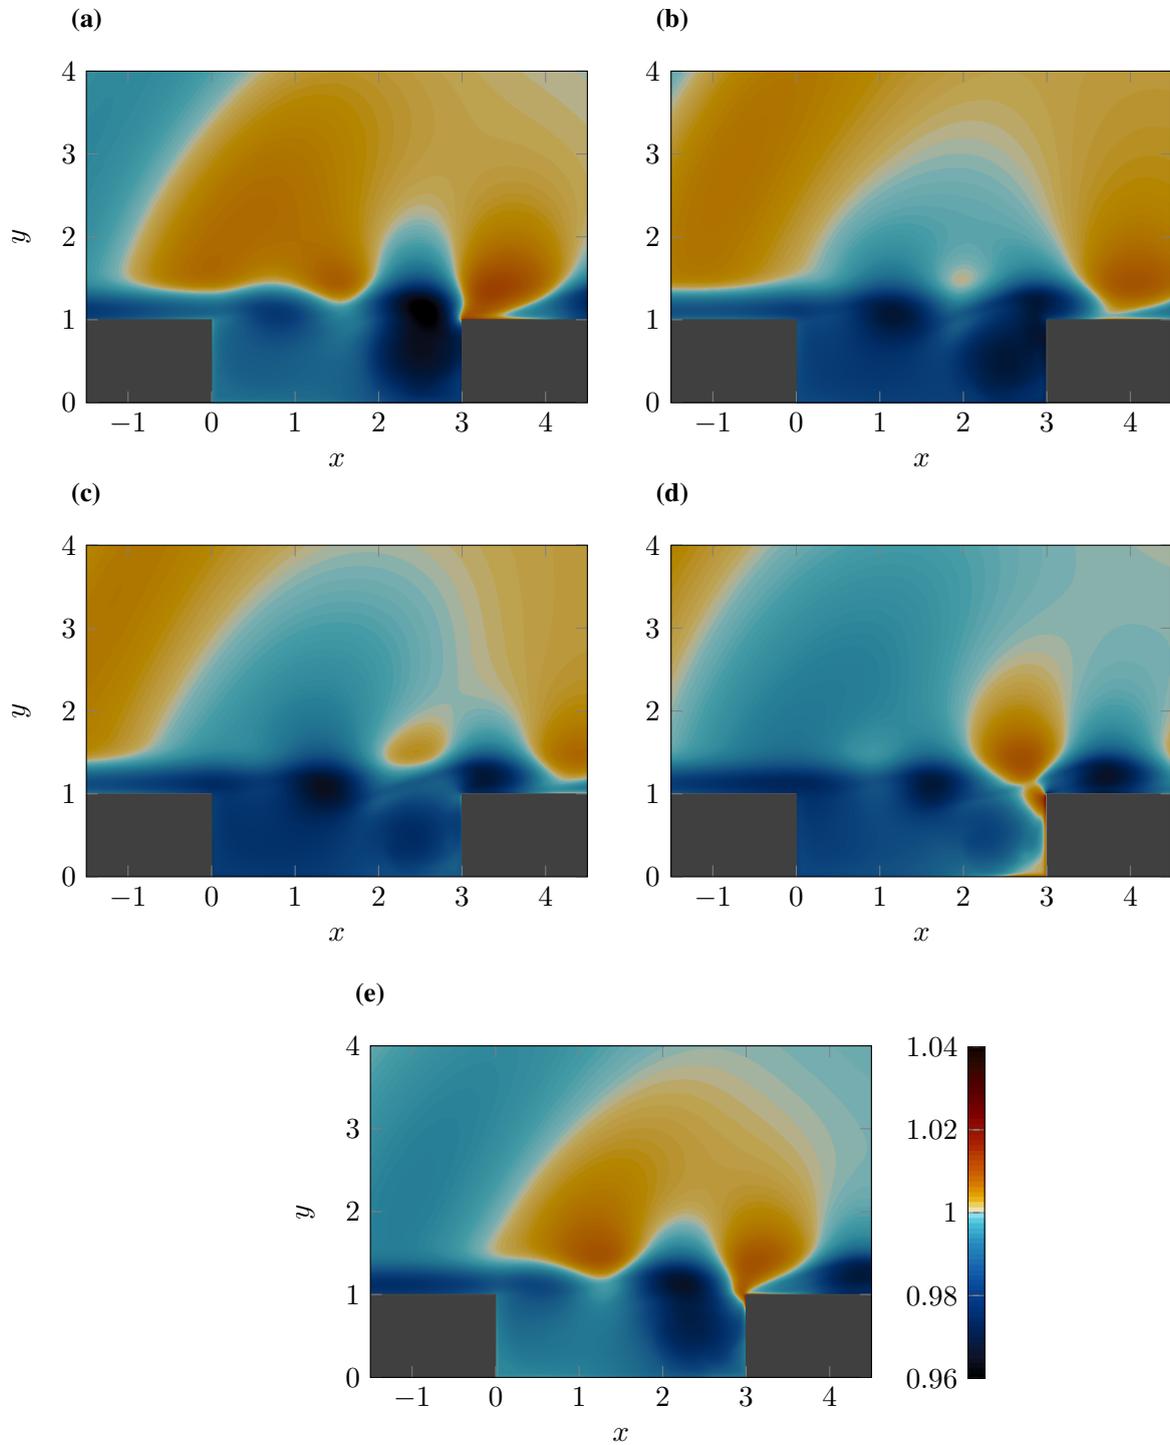


Figure 6.6: Contours of instantaneous density at the time steps represented by the red marks in figure 6.2. Subfigures (a) to (e) show the density field at the instants represented in figure 6.2d with the same letter. This is, (a) corresponds to the mark labelled with a , (b) with the mark labelled with b and so on. Density contours range from 0.96 to 1.04, where blue and red represent low and high density respectively.

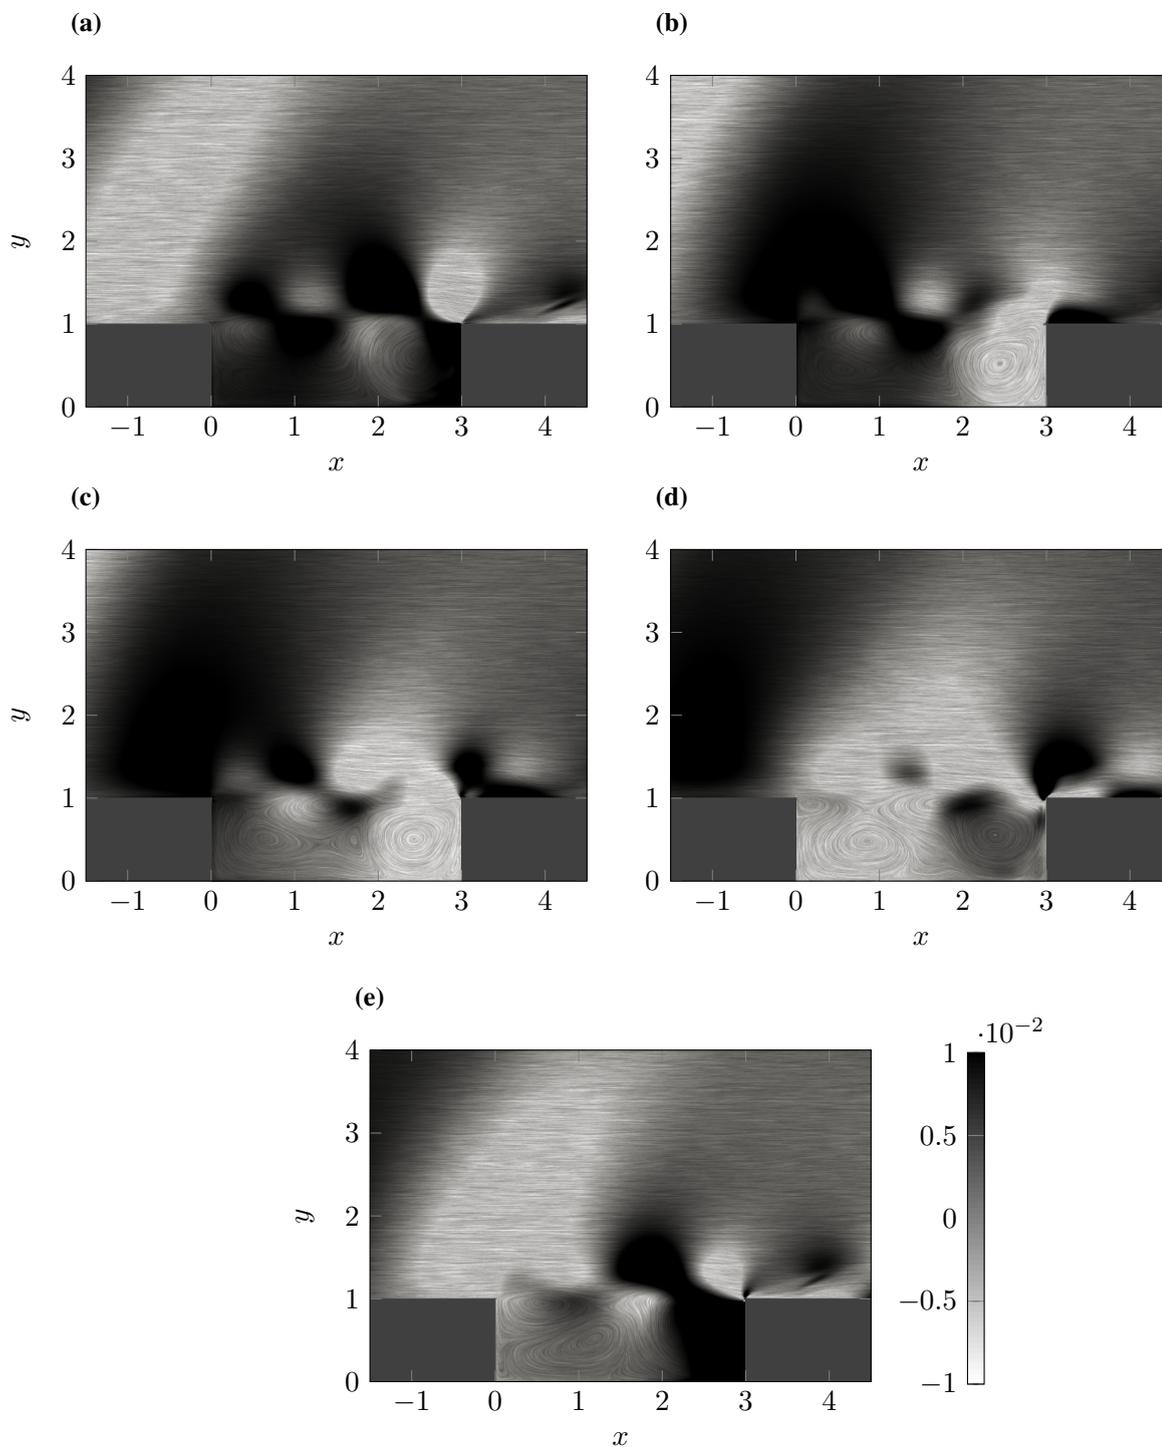


Figure 6.7: Contours of instantaneous dilatation at the time steps represented by the red marks in figure 6.2. Contour plots are overlapped with the instantaneous streamlines. Subfigures (a) to (e) show the dilatation field at the instants represented in figure 6.2d with the same letter. This is, (a) corresponds to the mark labelled with a , (b) with the mark labelled with b and so on. Dilatation contours range from -0.01 to 0.01 , where black and white represent flow expansion and compression respectively. This contour range has been chosen specifically to highlight the sound generation mechanisms.

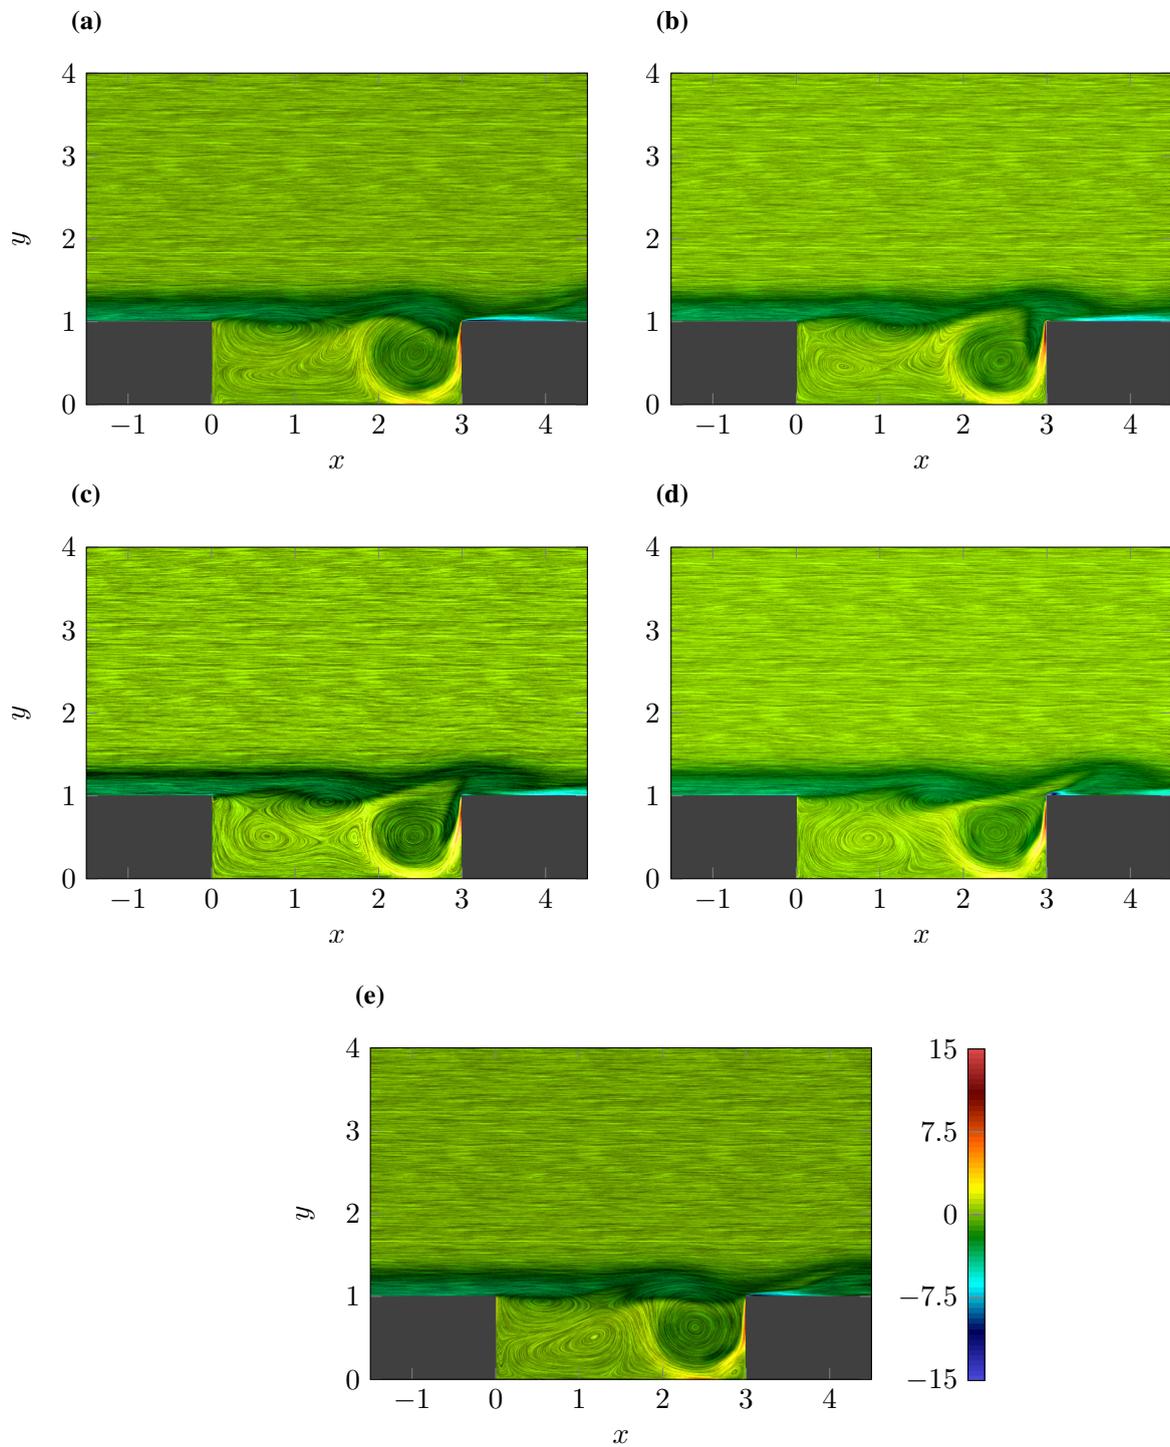


Figure 6.8: Contours of instantaneous vorticity at the time steps represented by the red marks in figure 6.2. Contour plots are overlapped with the instantaneous streamlines. Subfigures (a) to (e) show the vorticity field at the instants represented in figure 6.2d with the same letter. This is, (a) corresponds to the mark labelled with a , (b) with the mark labelled with b and so on. Vorticity contours range from -15 to 15, where blue and red represent negative and positive vorticity values respectively. This contour range has been chosen specifically to highlight the shear layer.

6.4 Family of Periodic Solutions across Mach Number

In this section, we show the evolution of the $M = 0.5$ orbit presented in the previous section across Mach number. The range of Mach numbers studied covers from $M = 0.25$ to $M = 0.8$. At the lower end of the range, the solution ceases to exist due to the low compressibility of the system and leads to a steady state. The collaboration amongst the compressible and convective phenomena shown previously for the $M = 0.5$ solution results in a family of periodic solutions which are stable. The non-dimensional character of our numerical framework permits the use of a periodic orbit with Mach number $M = M_0$ as an initial guess in the search of a new periodic orbit in neighbouring Mach numbers $M = M_0 \pm \delta$. Note that in flows with relevant compressible events, such as the one regarded in this study, a minor instantaneous change in Mach number will result into a non-physical flow-field. This is mainly due to the acoustic reflections created as the flow adapts to the new travelling speed of sound. Effectively, a decrease (or increase) in Mach number originates a sudden acceleration (or deceleration) of the acoustic waves, which creates artificial compressible phenomena in the vicinity of these waves. Hence, it might be convenient to let the flow evolve through the transient (i.e. spurious acoustic waves to leave the flow domain) before beginning the search of a new orbit. Initially, the orbits were continued from the periodic trajectory found at $M = 0.5$, both in ascending and descending order with increments of 0.05 in Mach number. Given the large variations in the flow quantities plotted in figure 6.9 for Mach numbers above 0.65, the step size in Mach number was reduced to 0.01. Hence, with this higher resolution, we are able to ensure that all the periodic orbits shown in this chapter belong to the same family of flow solutions.

Figure 6.9 shows the period T and the time-averaged norms of dilatation, kinetic energy, vorticity and dissipation rate as functions of Mach number. The periods shown in figure 6.9a are compared with the predictions calculated using Rossiter's semi-empirical formula (Rossiter, 1964)

$$T_M = \frac{L}{U_\infty} \frac{M + 1/\kappa_M}{n - \gamma_M}, \quad (6.9)$$

where L is the cavity length and κ_M and γ_M are empirical constants². These results correspond to the second cavity mode³ ($n = 2$). The agreement of the periods from the orbits found in the current investigation and the predictions calculated using Rossiter's formula is remarkable in the central section of figure 6.9a ($0.35 \leq M \leq 0.65$). This particular range of Mach numbers shows a smooth and monotonic behaviour across all the plots shown in figure 6.9. Note that the solutions at $M = 0.25$ and $M = 0.3$ are steady flow solutions, where the periods shown in 6.9a are purely numerical outputs obtained by the algorithm without physical interpretation. As mentioned in the previous section for the $M = 0.5$ orbit, this family of periodic solutions arise from the self-sustained compressible feedback mechanism characteristic of cavity flows. The vortex impingement on the trailing edge of the cavity radiates an upstream travelling acoustic

²These empirical constants were calculated based on the periods obtained for the periodic orbits at Mach numbers 0.5 and 0.55. After solving the system of two equations with κ_M and γ_M as unknowns, we arrive to $\kappa_M = 0.6096$ and $\gamma_M = 0.5003$.

³See section 6.5 for comments on the cavity mode selection mechanism.

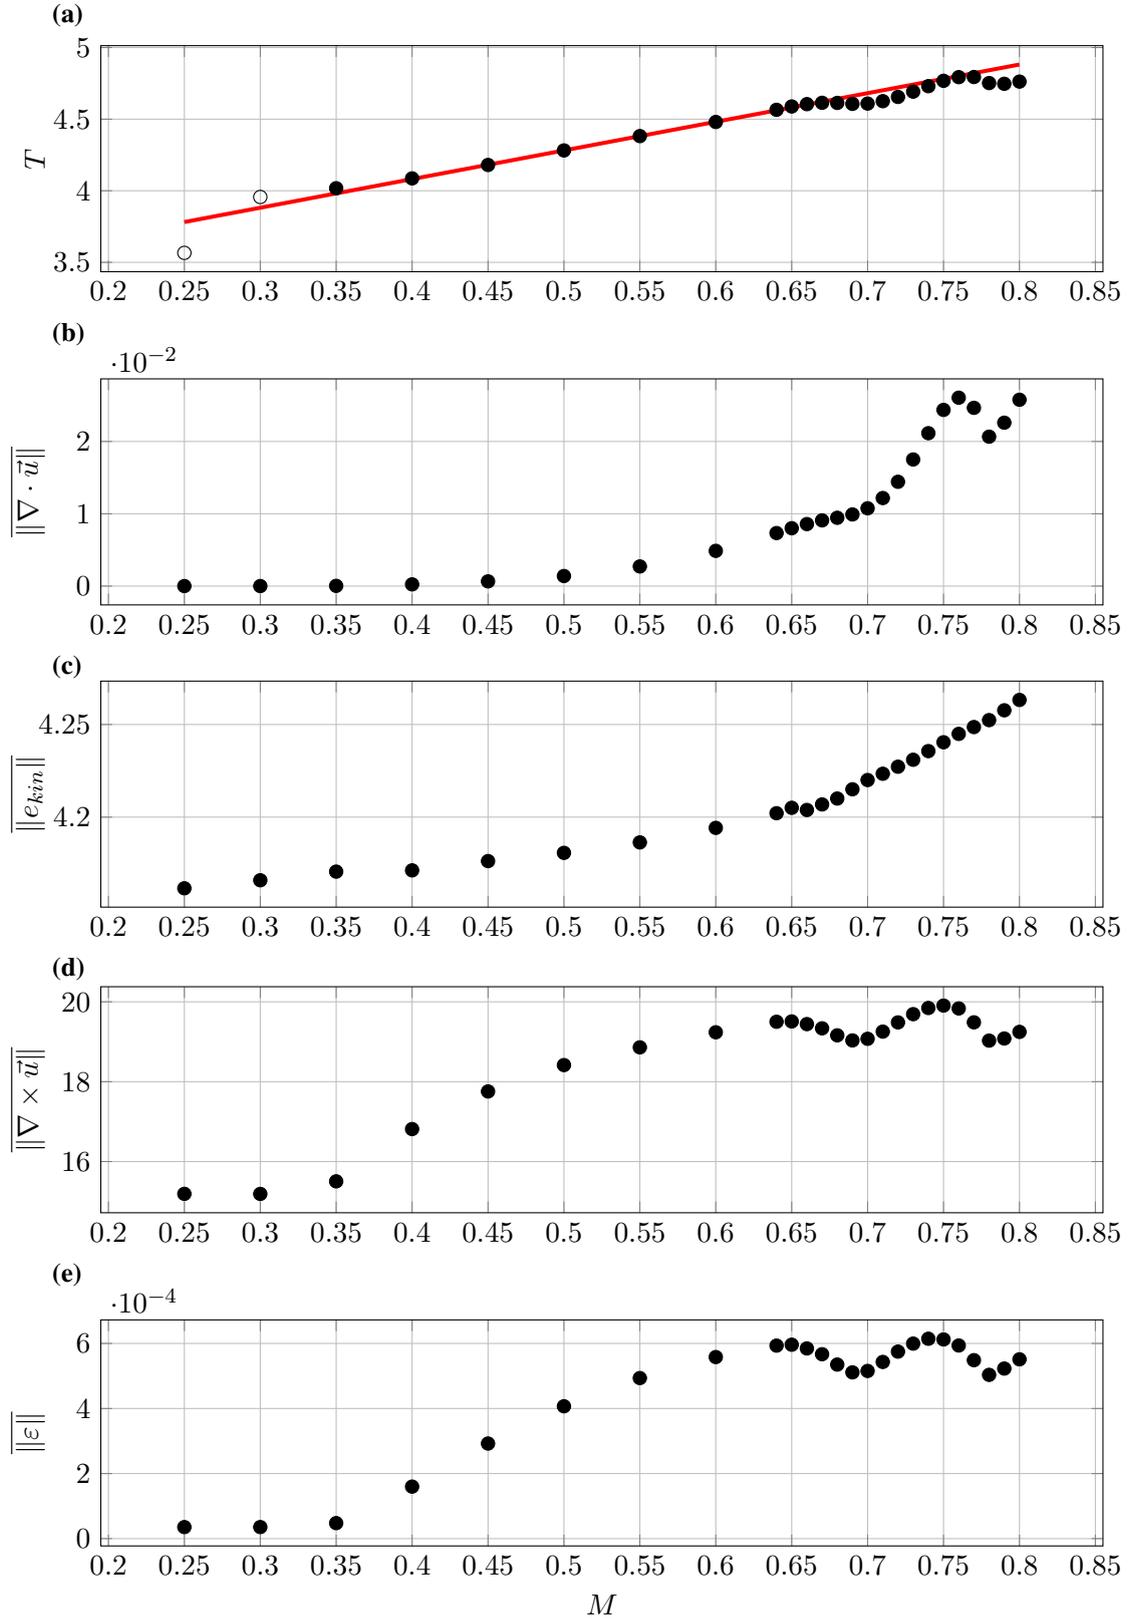


Figure 6.9: Evolution of the period (T), and time-averaged norms of dilatation ($\|\nabla \cdot \vec{u}\|$), kinetic energy ($\|e_{kin}\|$), vorticity ($\|\nabla \times \vec{u}\|$) and dissipation rate ($\|\varepsilon\|$) as a function of Mach number. Red solid line in (a) shows the periods predicted by Rossiter's formula.

wave that interacts with the oncoming new shear layer vortex. These two phenomena mutually benefit from each other due to their phase synchronisation along the entire orbit. When moving away from the reference $M = 0.5$ solution in Mach number, the change in the propagating speed of the upstream travelling acoustic wave results into a phase modification of the acoustic wave shear layer interaction. For clarity in the following analysis, we refer to a synchronised interaction as ‘in phase’, where the dipoles associated with the shear layer vortex and vortex impingement at the trailing edge have the same sign and enhance the upstream travelling acoustic wave (for example figures 6.7a or 6.7d). Contrarily, we say that the interaction occurs ‘out of phase’ when these two dipoles have opposite sign and lessen the intensity of the radiated sound wave.

The smooth and monotonic increase of the quantities shown in figure 6.9, suggest that the higher flow compressibility associated with a higher Mach number dominates the average behaviour of the periodic orbits up to $M = 0.65$. From this point onwards, the phase of the interaction appears to become a dominant phenomenon, leading to the oscillatory behaviour of the mean quantities observed mainly in figures 6.9d and 6.9e. Additionally, the rate of increase in average dilatation reduces from $M = 0.65$ to $M = 0.7$ due to the opposite phase of the acoustic wave and the shear layer, which partially cancels out the compressible phenomena. Note that this phase coupling might also vary as a function of Reynolds number (number of vortices in the shear layer) and cavity length (travelling distance of the upstream propagating acoustic wave). As the acoustic wave keeps decreasing its propagating speed (increasing the Mach number), the phase of the acoustic event becomes favourable again for the two physical mechanisms to work in synchronisation. This behaviour is reflected as a pronounced increase in the average dilatation from $M = 0.7$ to $M = 0.76$. As discussed further on, the higher flow compressibility allows the shear layer to not only modulate the upstream travelling acoustic wave but also radiate sound of comparable magnitude. This phenomenon, alongside its phase synchronisation with the already existing acoustic wave, is responsible for the sudden and steep changes in the periodic orbits above $M = 0.75$ observed in figure 6.9. Moreover, it is also worth tracking the evolution across Mach number of relevant local quantities (rather than integral) which are important to characterise the flow, such as the boundary layer’s momentum thickness Θ at the flow separation point. As well, we later show in section 6.5 that the variations in the momentum thickness produce considerable changes in the amplitude of the above-described shear layer oscillations, and also slight modulations in the system’s characteristic frequency. In particular, the amplitude of these fluctuations appears to increase with rising L/Θ , expressing its maximum differences when the phase portraits are compared in terms of $\overline{\|e_{kin}\|}$. Hence, the monotonic decrease in Re_{Θ} for $M > 0.5$ shown in figure 6.10 agrees with the steep increase in $\overline{\|e_{kin}\|}$ shown in figure 6.9c for the highest Mach number orbits. Furthermore, the trajectory’s period also shows a close relation with Re_{Θ} , where a decrease in L/Θ (increase in Re_{Θ}) yields a longer period (see section 6.5 and also appendix C). Especially for $M > 0.65$, an opposite oscillatory behaviour to the orbit’s period as a function of Mach number (figure 6.9a) is reflected in figure 6.10. More precisely, Mach number ranges which present a steep increase in the orbit’s period correspond to a steep decrease in Re_{Θ} at that same Mach number range. Hence, it appears that

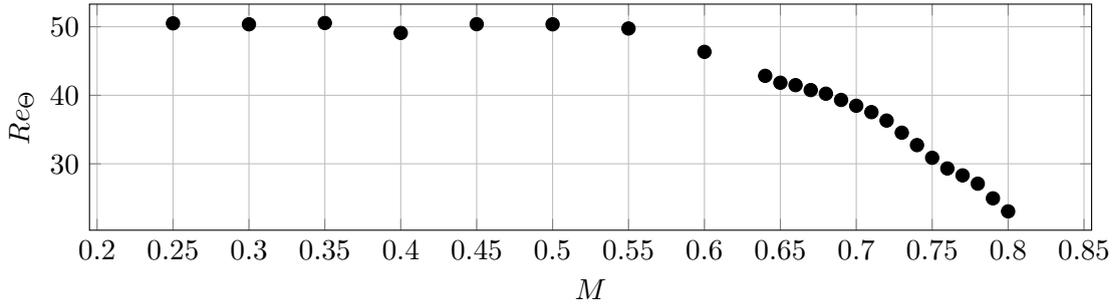


Figure 6.10: Evolution of Re_{Θ} as a function Mach number. The Re_{Θ} values are computed respect to the averaged flow quantities through one complete period at each period T .

the optimisation algorithm uses the incoming boundary layer thickness to balance the frequency of each flow trajectory, maintaining it within the proximity of the predictions from Rossiter's formula (6.9).

Figure 6.11 shows the phase portraits of the periodic orbits up to $M = 0.65$. Note that three of the main features described in the previous section can be identified in figure 6.11a. Instants a and b (vortex impingement, and shear layer vortex moving forward before interacting with the stationary vortex, respectively) correspond to the local maximum and minimum values of dilatation on the low viscous dissipation zone of the orbit. The instant named as c (beginning of the merging of the shear layer and stationary vortices) in figure 6.2, merges with b for Mach numbers above 0.5, as the difference in flow and sound velocity reduces. On the other hand, the instant d (maximum intensity of the vortex merging) always follows the absolute maximum value of viscous dissipation rate, reaching maximum intensity at $M = 0.6$. At Mach numbers 0.35 and 0.4, the strongest compressible event (maximum value of dilatation) is initially the vortex merging. After further increasing the Mach number, the maximum value of the norm of dilatation shifts to the instant corresponding with the vortex impinging on the trailing edge. This transition is strongly related to the physical mechanisms that cause the bifurcation of this family of periodic solutions from the steady solutions. As shown in figure 6.12a, the solution at $M = 0.30$ sits at perfect equilibrium thanks to the low Reynolds number and the laminar boundary layer. This steady solution becomes unstable for immediately higher Mach numbers, which leads to the periodic trajectories represented in figure 6.11. Figure 6.12b shows the instantaneous vorticity field of the $M = 0.35$ orbit where the norm of the dilatation field is highest. When the steady solution becomes unstable, the weak leading edge vortex gets absorbed by the stationary vortex. During this merging process, similarly to the $M = 0.5$ solution, the counter rotation of the two vortices originates a high-density spot located downstream from the shear layer vortex⁴, slightly on top of the stationary vortex. This phenomenon is the responsible for the highest norm of the dilatation field and coincides in time with the also highest norm of the viscous dissipation. Note that the flow-field shown in figure 6.12b is the $M = 0.35$ equivalent of the one shown in figure 6.8d for the $M = 0.5$ case. As the merging continues, this high-density

⁴We refer to the leading edge vortex as shear layer vortex once it detaches from the leading edge of the cavity.

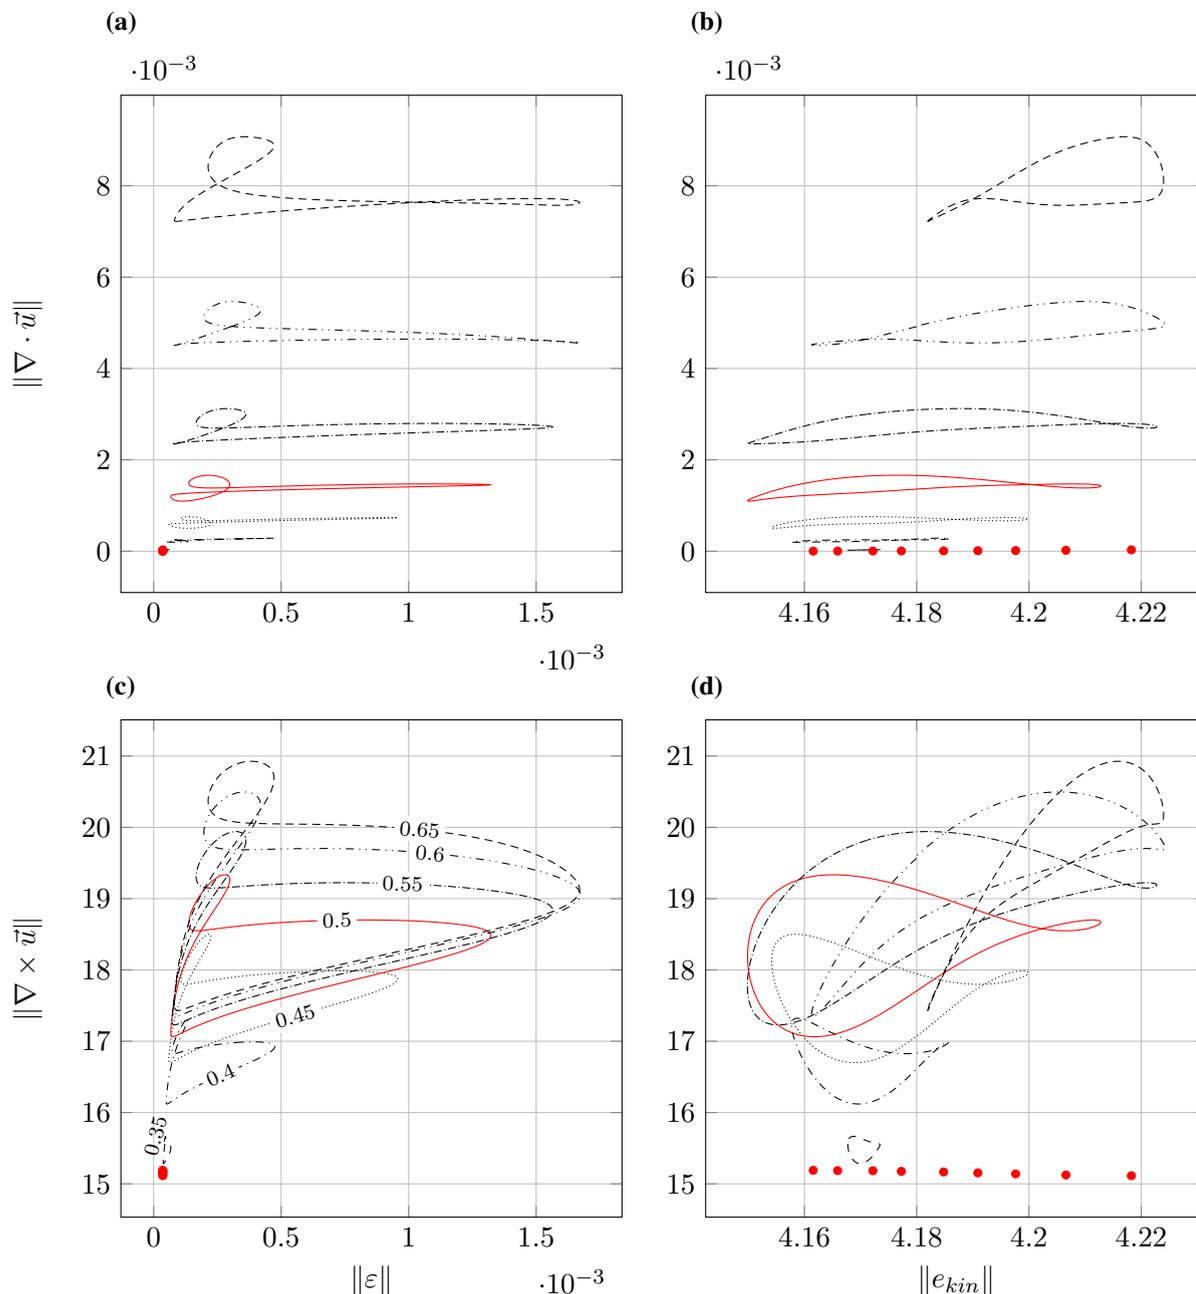


Figure 6.11: 4D representation of the family of solutions across Mach number, ranging from 0.25 to 0.65 with increments of 0.05. As a reference, the periodic orbit that corresponds to $M = 0.5$ is coloured in red. The trajectories are projected over the norm of dilatation ($\|\nabla \cdot \vec{u}\|$), vorticity ($\|\nabla \times \vec{u}\|$), dissipation rate ($\|\varepsilon\|$) and kinetic energy ($\|e_{kin}\|$). The periodic orbits are labelled in (c). The equilibrium solutions (see section 6.6) from Mach numbers 0.25 to 0.65 (ordered with increasing kinetic energy) are represented as red dots. Note how both families of solutions collapse in the quasi-incompressible regime at $M = 0.30$, which can be seen as a bifurcation point.

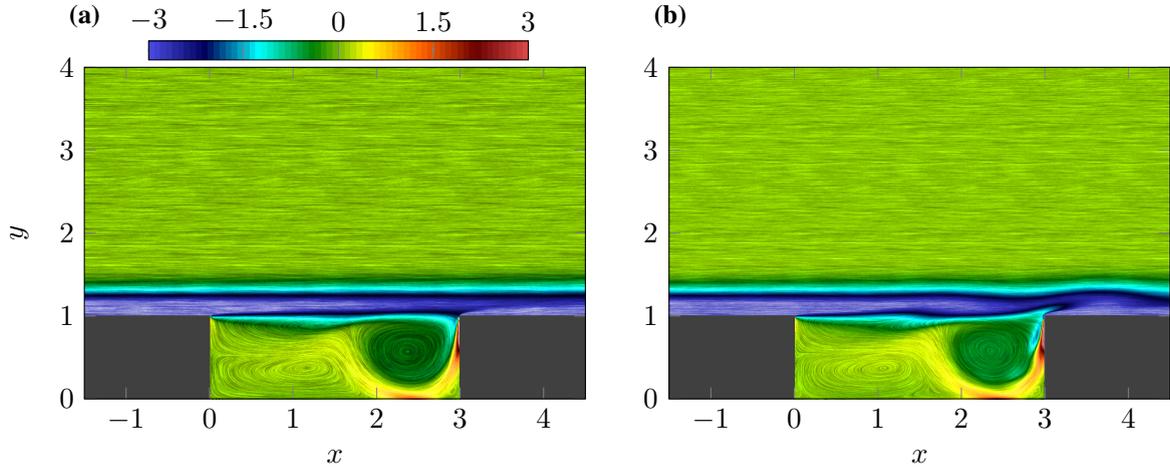


Figure 6.12: Instantaneous contours of vorticity at Mach numbers 0.30 **(a)** and 0.35 **(b)**.

area orbits the stationary vortex and impinges onto the trailing edge of the cavity. The reason why the vortex impingement is not as relevant from the compressible point of view is that the shear layer vortex is not strong enough to endure the orbit of the stationary vortex. This results in the low-density area, originally located on top of the shear layer vortex, to rapidly dissipate. Hence, when the low-density area impinges on the trailing edge, the intensity of the compressible events is not as high as with the impingement of its high-density counterpart. However, as the Mach number increases, the leading edge vortex gains in strength, which eventually leads to the low-density area impingement as the flow event with the highest compressibility. This same phase modification is observed in figure 6.11b. For the low Mach numbers, the periodic orbit reaches maximum kinetic energy in the vicinity of the maximum dilatation point. When the Mach number is increased, the intensification of the shear layer vortex rotation induces a greater flow velocity magnitude at the above-mentioned low-density area. Hence, once the vortex merging is completed, the high velocity (low density) zone orbits the stationary vortex, which further enhances the velocity magnitude in that area. This phenomenon displaces the section of the periodic trajectory from e to a towards higher kinetic energy values. Figure 6.11c shows the norm of the viscous dissipation rate against the norm of the vorticity field. The reader should note that these two quantities do not depend directly on the acoustic or compressible events occurring in the flow field. Hence, the plot keeps the same shape across Mach number, showing a monotonic increase in the amplitude of the periodic trajectory as the Mach number increases. Additionally, this figure shows how the periodic orbits bifurcate from a family of steady solutions in the quasi-incompressible flow regime. Given the steady character of the flow solution at $M = 0.25$ and $M = 0.3$, these two flow solutions are represented as overlapped points in the bottom left corner of the plot. The absolute minima of vorticity represent the start of the merging amongst the shear layer and stationary vortices (instant c). Figure 6.11d shows how the kinetic energy varies significantly with Mach number. The local vorticity minima at the right side of the flow trajectory (high kinetic energy zone) correspond to the instant where the shear layer vortex has merged completely with the stationary vortex. Furthermore, the overall shape of the periodic

trajectory suffers significant changes from $M = 0.55$ to $M = 0.65$, which is an indication that the acoustic wave shear layer interaction starts dominating the flow physics of the orbit. Particularly, at the absolute minimum of vorticity, the kinetic energy increases considerably from $M = 0.60$ to $M = 0.65$, induced by the drop of Θ . Also, the maximum value of the kinetic energy norm appears to reach saturation from $M = 0.6$ to $M = 0.65$ which shows that this interaction is shifting towards an out of phase synchronisation.

Figure 6.13 shows the phase-dominated Mach number range ($0.65 < M < 0.8$). Differently from the lower Mach number orbits represented in figure 6.11, the averaged norms of vorticity and viscous dissipation rate do not increase monotonically with Mach number. Instead, they now register an oscillatory behaviour that is related to the phase of the acoustic wave shear layer interaction. Figure 6.13a highlights how the phase of the above-mentioned interaction modifies the orbit at the vortex impingement on the trailing edge. As mentioned in section 6.3, this phenomenon is represented as the top left loop observed in the periodic orbits shown in figure 6.13a, which slowly unfolds as the Mach number increases. For the orbit at $M = 0.74$, the beginning and the maximum intensity of the vortex impingement are distinctively represented as two sharp corners. To have an insight on the phase of the interaction, figure 6.14 reveals the dilatation field at the maximum of the norm of dilatation (instant a) for Mach numbers 0.65 and 0.74. The upstream propagating acoustic wave is located further downstream on the $M = 0.74$ case, as its propagating velocity is lower. Also, note that the dipole associated with the shear layer vortex gains in intensity as the flow becomes more compressible. For comparative purposes, these dilatation contours from these figures show a remarkable agreement with the ones shown in Rowley et al. (2002), despite the fact their cavity configuration had $L/D = 2$. At the end of the Mach number range represented in figure 6.13, the intensity of the shear layer dipole is such that the radiated sound is of comparable magnitude to the already existing acoustic wave. As a consequence of this stronger interaction, the sound directivity changes towards a more vertical direction (see section 6.4). Figure 6.13b highlights the rapid change in the behaviour of the periodic orbit as it approaches $M = 0.75$. In the trajectory, three symbols (ordered chronologically with increasing kinetic energy) show the location of the flow events that cause these sharp bends in the phase portrait of the orbit. The first point shows the location right after the maximum intensity of the merging process between the shear layer and stationary vortices. The counter rotation of the new leading edge vortex with the stationary vortex causes a flow compression (at almost constant kinetic energy) in the centre of the cavity. This phenomenon also coincides with the impingement of the high-density spot on the trailing edge. After reaching the second point, despite the leading edge vortex keeps growing and compressing the flow (as seen in section 6.3 for the periodic orbit $M = 0.5$ from points e to a), the impingement of the high-density zone has finished, which is reflected as a drop in the flow compressibility. This occurs as the low-density spot further accelerates when orbiting the stationary vortex, causing an overall increase of kinetic energy, just before the impingement of the low-density area commences. Differently from the lower Mach number range shown in figure 6.11, the vortex impingement now sits as the highest value of dilatation and also in the vicinity of the highest kinetic energy. Now similarly to the previous Mach number range, the

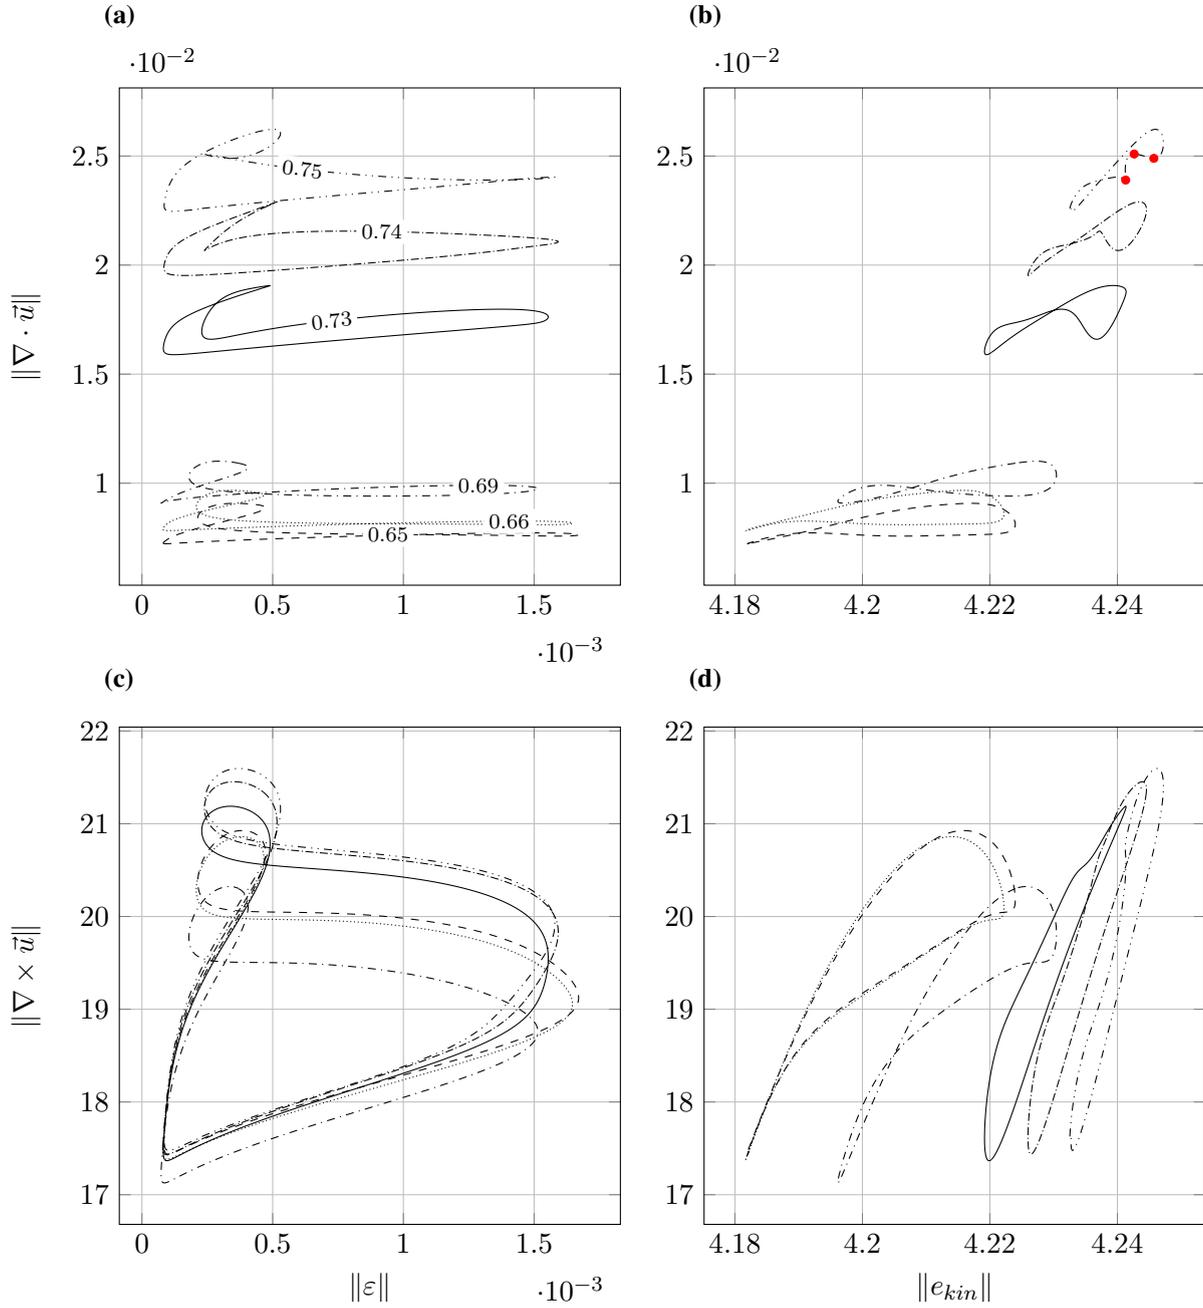


Figure 6.13: 4D representation of the family of solutions across Mach number, ranging from 0.65 to 0.75. As a reference, the periodic orbit that corresponds to $M = 0.73$ is coloured in red. Similarly to figure 6.2, the trajectories are projected over the norm of dilatation ($\|\nabla \cdot \vec{u}\|$), vorticity ($\|\nabla \times \vec{u}\|$), dissipation rate ($\|\varepsilon\|$) and kinetic energy ($\|e_{kin}\|$). The periodic orbits are labelled in (a).

maximum intensity of the vortex merging is also identified as the instant of maximum viscous dissipation rate in figures 6.13a and 6.13c. These two figures also show the shear layer vortex travelling downstream, right after detaching from the leading edge, as the minimum point in dilatation, vorticity and viscous dissipation. Shortly before that instant, the projected trajectory over the viscous dissipation and vorticity appears to briefly become independent from Mach

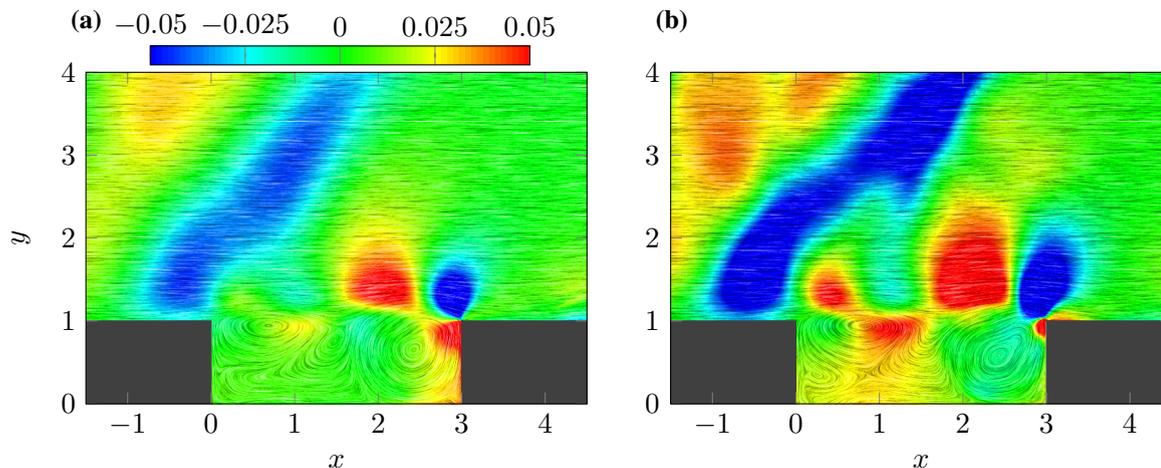


Figure 6.14: Snapshots of the dilatation field at the maximum value of the norm of dilatation for Mach numbers 0.65 and 0.74. Similarly to figure 6.7, dilatation contours range from -0.01 to 0.01, where black and white represent flow expansion and compression respectively.

number, where all the trajectories collapse. On the other hand, the phase synchronisation of the acoustic wave shear layer interaction modulates the behaviour of the flow-field, leading to the oscillatory stretching of the orbits in viscous dissipation and vorticity axis, which was also reflected in figure 6.9. As seen before in figure 6.9c, the average kinetic energy of the periodic orbits drops slightly from Mach numbers 0.65 to 0.66. The origin of this reduction resides in the lower kinetic energy carried from start to finish of the vortex impingement at $M = 0.66$ (figure 6.13d). For the following Mach numbers, the average kinetic energy follows an increasing quasi-linear trend up to $M = 0.8$. This is also reflected in figures 6.13b and 6.13d, where the horizontal displacement of the orbits as a function of Mach number is considerably higher than observed in figure 6.11d. Further, the variations of kinetic energy respect to the mean value appear to decrease, leading to more vertically aligned orbits in figure 6.13d.

The trajectories corresponding to the highest Mach numbers from the family of solutions shown in figure 6.9 are gathered in figure 6.15. At this regime, the compressibility of the shear layer is high enough to radiate sound independently, which now interferes with the well known upstream propagating sound wave originated with the vortex impingement on the trailing edge of the cavity. This new acoustic wave will be referred to as the shear layer acoustic wave. The averaged quantities of these orbits show significant changes respect to the trend followed by lower Mach number solutions. The most noticeable event is the sudden drop in the averaged dilatation from 0.76 to 0.78. Differently from lower Mach number solutions, the averaged dilatation decreases with increasing Mach number. Precisely at $M = 0.78$, it appears to reach a local minimum, which is followed by a continuous increase until the upper bound of our Mach number range. Furthermore, the averaged values in vorticity and viscous dissipation rate follow the previously mentioned oscillatory behaviour, decreasing from $M = 0.75$ to $M = 0.78$. Similarly to the averaged dilatation, these two quantities also show a sharp change in trend at $M = 0.78$. On the other hand, these sudden changes are not reflected in the period of the solution (it carries an

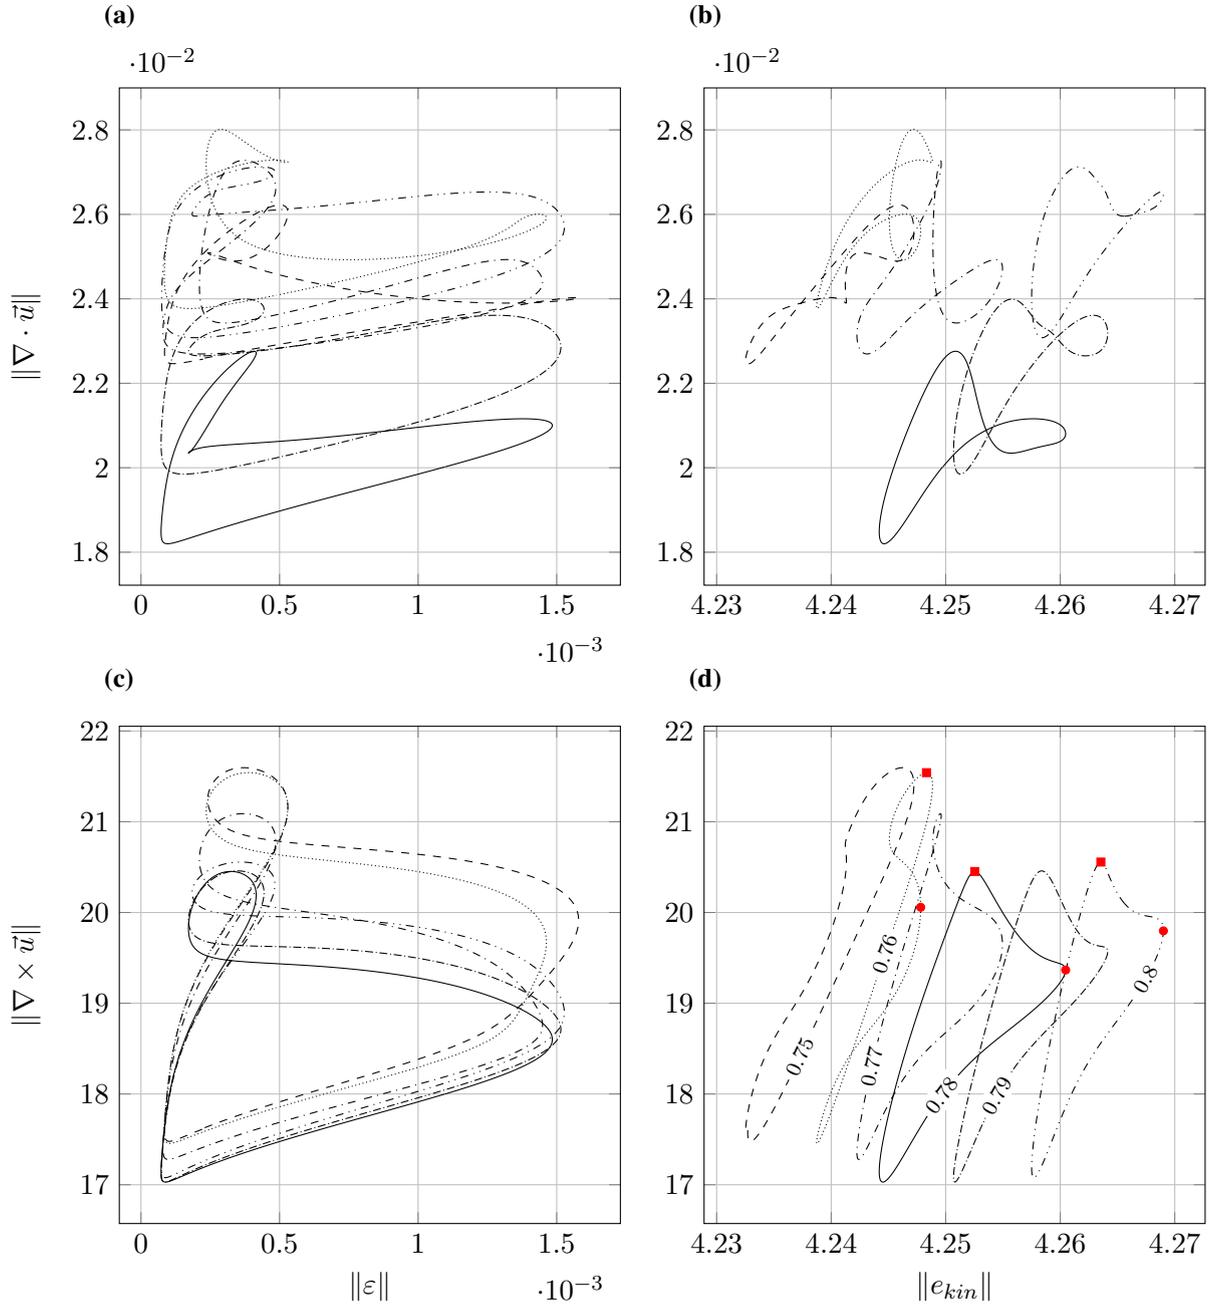


Figure 6.15: 4D representation of the family of solutions across Mach number, ranging from 0.75 to 0.8. As a reference, the periodic orbit that corresponds to $M = 0.78$ is coloured in red. Similarly to figure 6.2, the trajectories are projected over the norm of dilatation ($\|\nabla \cdot \vec{u}\|$), vorticity ($\|\nabla \times \vec{u}\|$), dissipation rate ($\|\varepsilon\|$) and kinetic energy ($\|e_{kin}\|$). The periodic orbits are labelled in (d).

oscillatory behaviour from $M = 0.65$ to $M = 0.75$), nor in the average kinetic energy, which follows an increasing quasi-linear trend.

Figure 6.14 revealed how the phase of the interaction plays a significant role in the sound radiation. There we showed how with increasing Mach number, the dipoles associated with the leading edge and leading edge vortex gain in strength (higher compressible effects). This trend

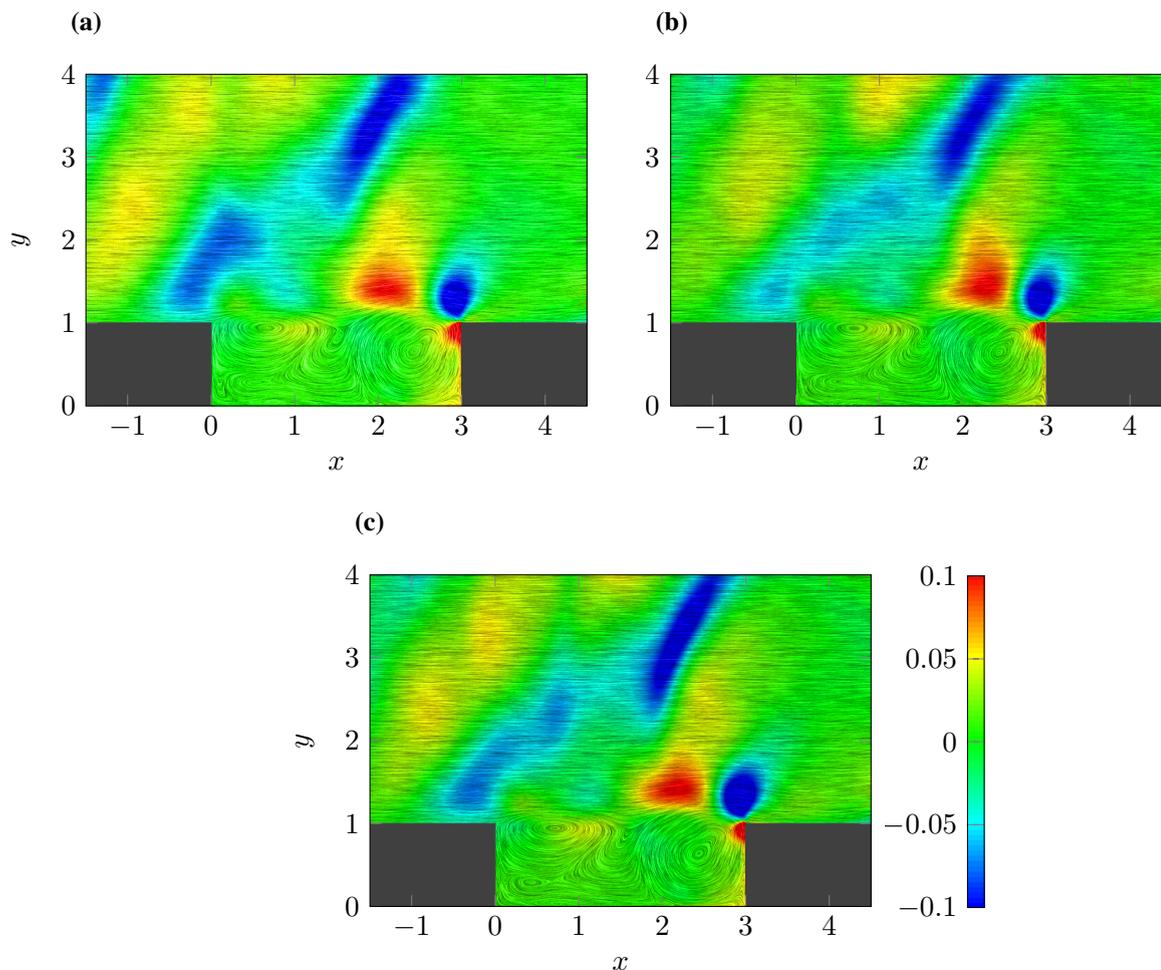


Figure 6.16: Instantaneous contours of the dilatation field at Mach numbers 0.76 **(a)**, 0.78 **(b)** and 0.8 **(c)**. These instants correspond to the maximum norm of the vorticity field along the periodic orbit.

continues with higher Mach numbers as the flow velocity approaches the propagating velocity of sound. Figure 6.16 shows the snapshots of the dilatation field at Mach numbers 0.76, 0.78 and 0.8 at the point of maximum norm of the vorticity field (figure 6.15). The position of these snapshots are highlighted in figure 6.15d with red square symbols at the flow orbit. In the same figure, the red circles show the opposite phase to the plots from 6.16. The interference amongst the upstream propagating acoustic wave and shear layer acoustic wave, which we observed earlier on in figure 6.14b, grows continuously through $M = 0.76$ (figures 6.16a), reaching the anti-phase of the interaction at $M = 0.78$. At this point, both acoustic waves cancel each other out in the vicinity of the cavity (figure 6.16b). This explains the sudden drop in the average dilatation field from $M = 0.76$ to $M = 0.78$. From this Mach number onwards, the shear layer acoustic wave keeps increasing its energy, and also the phase of the interaction slowly becomes favourable again. Note that the features of the purely convecting events (i.e. vortices) are barely altered throughout this interval. Additionally, despite the considerable difference in Mach number, the shape and alignment of the dipoles related to convective phenomena are remarkably

similar to the ones illustrated in figure 6.4 and shown in figure 6.6a for the $M = 0.5$ case. Furthermore, the qualitative agreement of figure 6.16c with the data shown by Rowley et al. (2002) (their figure 6) is excellent.

Overall Sound and Directivity

So far in the present section, we have seen how the overall behaviour of the periodic orbits changes as a consequence of the phase modification of the interaction amongst the shear layer and the trailing edge acoustic radiation. In particular, for the higher Mach numbers ($M > 0.65$) the sound radiation of the shear layer is of comparable magnitude to the radiation from the trailing edge. This leads to an enhanced interaction which, as seen in the above dilatation plots, has an effect on the energy and directivity of the overall sound radiation. To characterise the sound (or noise) radiation for the current family of periodic orbits we define the overall sound pressure level as

$$OSPL = 10 \cdot \log \left(\frac{\int_{-\infty}^{\infty} |\mathcal{F}(p'(t))|^2 df}{(p_{ref})^2} \right), \quad (6.10)$$

where $p'(t)$ are the pressure fluctuations at the measurement point, p_{ref} is a reference pressure⁵ and \mathcal{F} indicates a Fourier transform. To evaluate the sound directivity, the OSPL was computed in 35 equally distributed monitor points along an arc of radius 5, each of them separated by 5 degrees⁶. The centre of the arc is located at coordinates 1.5 and 1 in the streamwise and vertical directions, respectively. In order to have a higher resolution in the low-frequency range, the pressure signal at every measurement point for each Mach number was repeated 100 times. Note that this will not have an effect on the overall result of the integration since all frequency content below the periodic orbit's frequency has been removed during the optimisation. Also, the use of a window function is not necessary due to the periodicity of the time series. The OSPL of the most representative periodic orbits are gathered in figure 6.17. Angles lower than 90 degrees and greater than 90 degrees show, respectively, an upstream and downstream sound radiation. The OSPL values are normalised with the maximum upstream propagating OSPL value from the orbit at $M = 0.5$.

In section 6.3 we observed how the shear layer and leading edge dipoles enhance the acoustic wave originally radiated by the trailing edge at $M = 0.5$. At the same time, the interaction of the shear layer and trailing edge dipoles partially cancels out the sound radiation in the downstream direction (figure 6.4). This behaviour agrees with the content of figure 6.17a, where the OSPL values reach higher magnitudes at the upstream propagating angles, and the downstream propagation is severely mitigated by the shear layer dipole. The quasi-circular shape of the plots at Mach numbers 0.4 and 0.5 indicate that this upstream sound radiation is almost entirely governed by the cavity's trailing edge noise. Note that this is not the case for $M = 0.6$ and $M = 0.65$, where the shear layer is compressible enough to start radiating sound. As the

⁵The reference pressure used is $p_{ref} = 2 \cdot 10^{-5}$.

⁶Note that the location of each of the monitor points has been approximated to the nearest grid point.

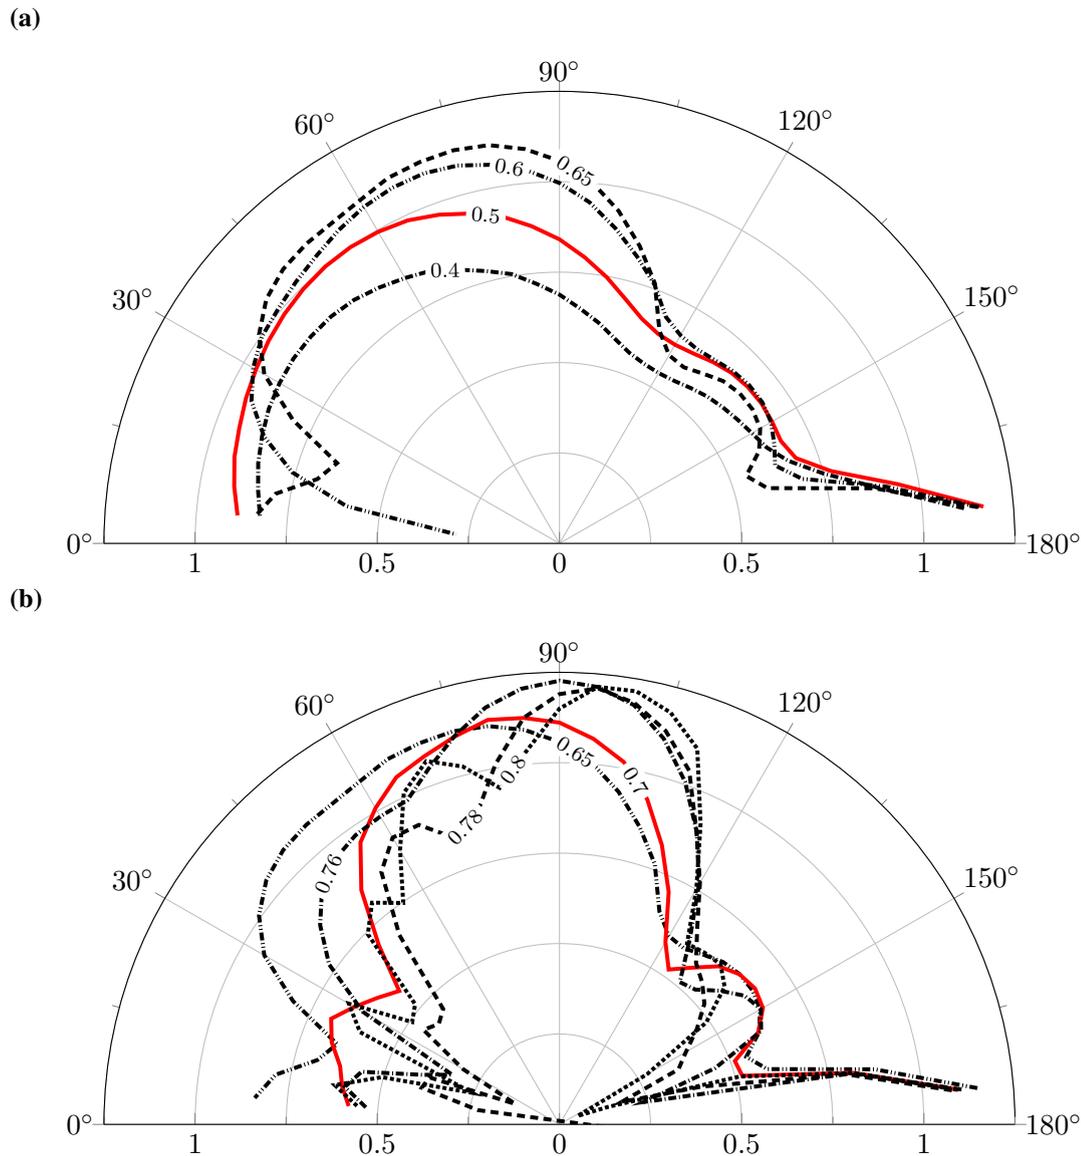


Figure 6.17: OSPL from 35 equally distributed monitor points along an arc of radius 5, each of them separated by 5 degrees. The centre of the arc is located at coordinates 1.5 and 1 in the streamwise and vertical directions, respectively.

Mach number is further increased, the higher flow compressibility leads to stronger dipoles at the leading edge and shear layer. This phenomenon in combination with a lower speed of sound forces the above-mentioned interaction to slowly get out of synchronisation. The OSPL values for Mach numbers 0.6 and 0.65 show how the dipole located at the leading edge cancels out the upstream sound propagation in the vicinity of the wall. Furthermore, the high intensity of the shear layer dipole is reflected in the more pronounced lobe at about 140° . Note that as the Mach number is increased, the leading sound radiation direction shifts monotonically towards a higher angle. This trend also applies for the higher Mach numbers (figure 6.17b). At Mach numbers from 0.65 to 0.70, the dipole interaction continues out of synchronisation which leads to a substantial reduction of the acoustic radiation in the upstream direction. As the interaction becomes favourable again from $M = 0.70$ to $M = 0.76$, the upstream sound radiation increases

in the $30^\circ - 60^\circ$ range. This is solely caused by the higher radiation of the shear layer dipole at this Mach number range. As shown in figure 6.16, for Mach numbers above $M = 0.76$ the sound radiation is split in two different waves, where now the shear layer dipole is the responsible from the upstream propagating acoustics. This is observed in figure 6.17b as two main distinct lobes for $M > 0.76$. The sound radiation from the cavity's trailing edge is cancelled out almost completely in the upstream direction as it interacts with the shear layer acoustics, which forces its propagation in the vertical direction.

6.5 Other Periodic Solutions and Stability

In the section above we presented the family of periodic solutions which arise from continuing the periodic orbit at $M = 0.5$ across Mach number. But, yet at this point, one question remains unknown: *are the solutions continued from $M = 0.50$ the same as if they were obtained from their own fully developed flow at other Mach numbers?* A priori, given the considerable drop in Re_Θ with rising M shown in earlier in figure 6.10, which was caused by the optimisation algorithm, one would guess that the answer is *no*. In order to provide a more rigorous answer to this question, we computed the periodic orbits from independent fully developed flow-fields at Mach numbers 0.65 and 0.8. Figure 6.18 compares these new solutions with the orbits continued from the $M = 0.5$ trajectory at the aforementioned Mach numbers. Observing this figure, we can readily and certainly answer the question posed above: *no*. Moreover, these new orbits are considerably less energetic than the orbits continued from the $M = 0.5$ solution, exhibiting substantially inferior values in dilatation and kinetic energy. On the other hand, the shape of the trajectories described by both periodic solutions at each Mach number are identical. This suggests that the flow mechanisms which govern both flow solutions (with equal M) are the same, with the only difference of superior energy in the continued orbits. In addition, the periods of these new orbits are almost identical to the ones from their respective continued orbits, which confirms that both sets of solutions correspond to the same Rossiter mode.

To achieve a further insight into the differences amongst the periodic trajectories, we have computed the Re_Θ for each solution at $M = 0.8$, giving values of 23.06 and 64.37 for the continued and new orbits, respectively. One of the primary effects of this smaller Re_Θ in the continued flow trajectory (caused only by the decrease of Θ) is a more unstable character of the shear layer (Brès and Colonius, 2008). This enhanced instability yields a stronger leading edge vortex, which following the flow mechanisms described earlier in section 6.3, it eventually travels downstream and impinges onto the cavity's trailing edge, radiating a stronger acoustic wave. Due to the relatively high Mach number and the strong upstream noise radiation, all these physical events are also reflected in the velocity magnitude contours shown in figure 6.19. In particular, these plots correspond to the point in the flow trajectory of the maximum norm of dilatation for each case. For instance, these plots reveal significantly stronger (wider) upstream-travelling acoustic waves than figure 6.19b, and also, the leading edge vortex shows a slightly higher rotation velocity for the continued orbit case. All these differences indicate that there is a relevant

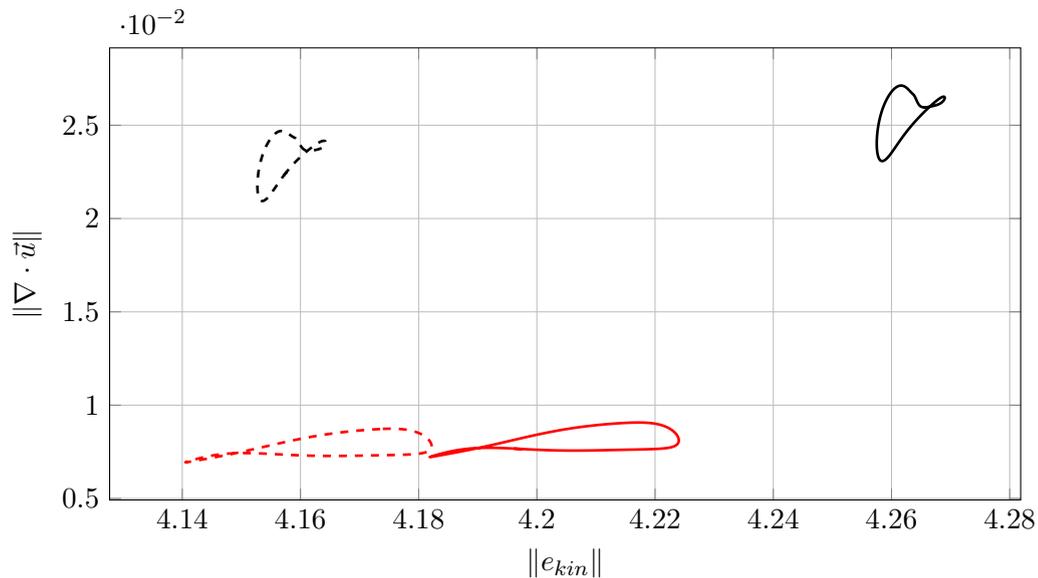


Figure 6.18: Comparison of the continued and original orbits at Mach numbers 0.65 and 0.8. The continued orbits are represented with solid lines. The trajectories found from the developed flows are shown as dashed lines. The solutions corresponding to Mach numbers 0.65 and 0.8 are coloured with red and black, respectively.

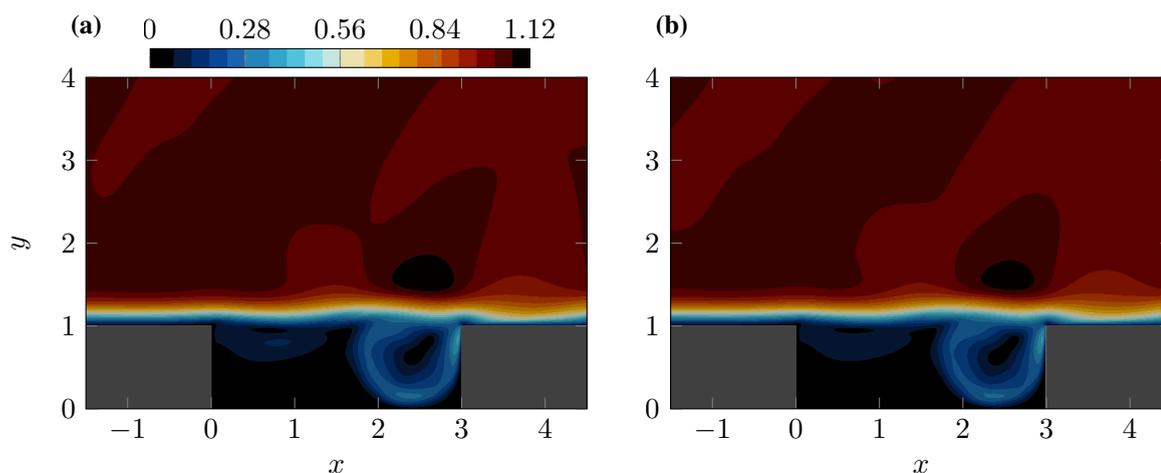


Figure 6.19: Snapshot of velocity magnitude at the point of maximum dilatation for the continued (a) and original (b) orbits at $M = 0.8$.

inheritance present in orbits continued from the $M = 0.5$ solution. As mentioned in section 6.4, when continuing an orbit across Mach number, the optimisation algorithm alters the incoming boundary layer thickness to presumably slightly adjust the characteristic frequency of the periodic solution. Hence, in essence, these different orbits arise from different initial conditions to the optimisation problem, which lead to different solutions.

In order to assess the stability of the periodic solutions presented in this chapter, we introduced random noise perturbations across the entire flow-field for several flow orbits (both continued from $M = 0.5$ and directly calculated from their developed flow). Despite the amplitude of these disturbances ranging up to 10% of Q in some situations, the flow-field eventually adjusted

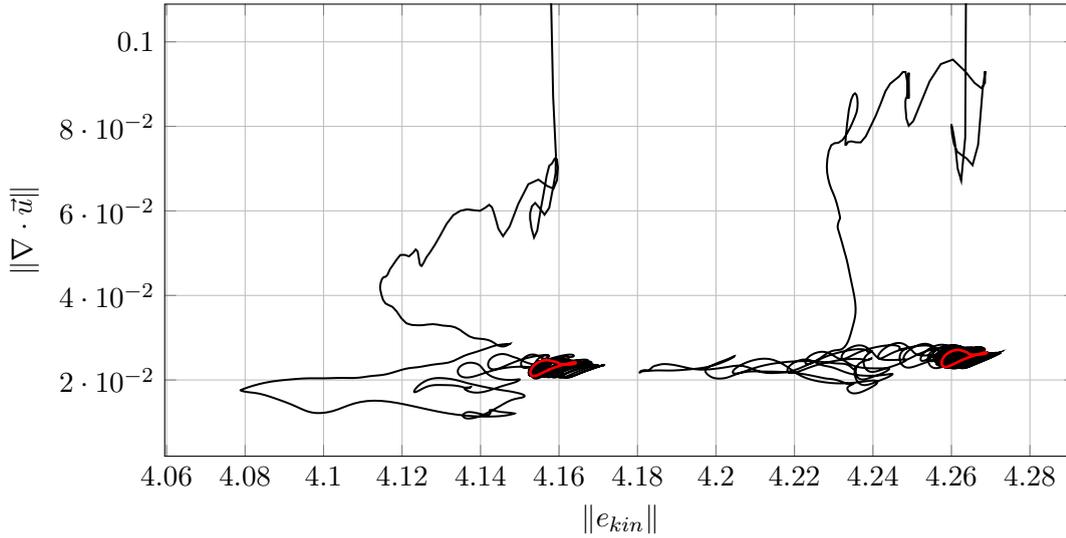


Figure 6.20: Evolution of both periodic orbits at $M = 0.8$ represented earlier in this section, following a large amplitude body force perturbation at the shear layer (black lines). Both solutions return back to their respective limit cycles (red lines) showing an stable character.

back to the unperturbed trajectory in every case. If any of these periodic orbits were unstable, a flow-field perturbation would have triggered a transition towards a different flow state. In any case, such transitioning scenario did not happen due to the remarkable stability provided by the flow-acoustic feedback mechanism. In addition, bear in mind that random noise perturbations are essentially strong numerical point-to-point oscillations in the flow-field. For this reason, they might be interpreted as spurious oscillations by the high-order explicit filter (see section 4.1.1), which partially removes these disturbances, contributing to the stability of the periodic orbits. Alternatively, to agitate the periodic solutions more efficiently with the intention of seeking their transition to another state, we now introduce an initial body force disturbance in the streamwise and vertical momentum state variables. The spatial domain of activity for this perturbation is defined as

$$W_f(\vec{x}) = e^{-\frac{(x-x_0)^2 + (y-y_0)^2}{0.05}}, \quad (6.11)$$

where both streamwise and vertical momentum components experience an external forcing up to $-0.1\rho_\infty u_\infty$ at the centre of the Gaussian function, at $\vec{x}_0 = (1, 1)$. The exact location of this forcing has been carefully chosen to alter the flow-field at the shear layer and the proximity of the cavity's leading edge. If any existing, an unstable shear layer mode would exhibit its highest receptivity values at these particular spatial regions (see chapter 7). For clarification purposes, at this point, it is worth naming the flow trajectories studied under this new perturbation. We refer to the flow solution at $M = 0.8$ continued from the original periodic trajectory at $M = 0.5$ as M080, whereas the one obtained straight from the developed flow is designated as M080-fd. Additionally, the flow trajectories representing perturbed orbits are appended with the prefix 'p-' (i.e. p-M080 for the perturbed M080 and so on).

Figure 6.20 shows the evolution of the trajectories p-M080 and p-M080-fd from the introduction

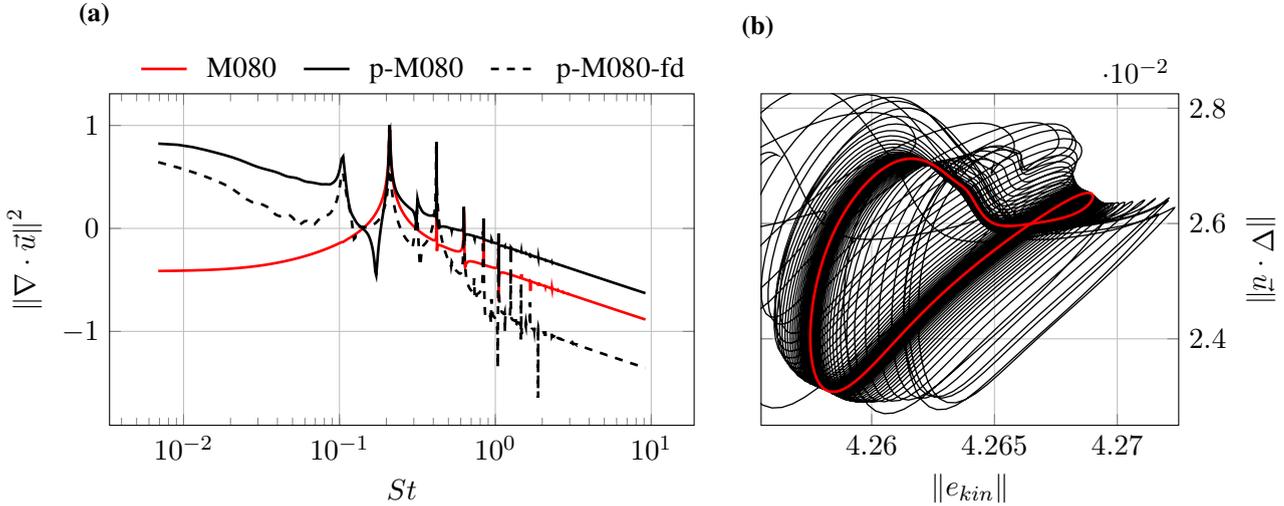


Figure 6.21: **(a)** Frequency spectrum of the perturbed flow trajectories at $M = 0.8$ and the exact (continued) solution also at $M = 0.8$. **(b)** Final stage of p-M080 decaying to M080.

of the body force until they decay to their respective M080 and M080-fd orbits. This shows that both orbits are stable, where despite that p-M080 gets closer (from a kinetic energy perspective) to M080-fd than M080 shortly after the perturbation is introduced, it eventually returns to its original state M080. Also, this particular disturbance appears to cause a slight phase shift in the trajectory, which impedes the direct comparison of both perturbed and unperturbed states⁷. Moreover, figure 6.21a displays the Fourier transformed norms of the dilatation field across time from M080, p-M080 and p-M080-fd. For simplicity, M080-fd has not been represented as it yields almost identical results to M080. Indeed, due to the high receptivity of the area surrounding the cavity's leading edge, this perturbation initially triggers the first Rossiter mode for both p-M080 and p-M080-fd, which appears as a peak at $St \approx 0.103$. Taking now a closer look on how M080-p decays to M080, as shown in figure 6.21b, we can spot the influence of the first Rossiter mode on the trajectory described by p-M080. Once the perturbed state p-M080 reaches the vicinity of M080, it commences an oscillation about the M080 flow solution with both frequency and exponential decaying rate associated with this stable first Rossiter mode. The stability of these Rossiter modes is related to the mode selection phenomenon also discussed by Brès and Colonius (2008). This flow mechanism consists of the cavity flow 'choosing' one Rossiter mode to govern the shear layer dynamics, where all the remaining modes below experience an exponential decay. This mode selection is not fully understood yet, and according to the literature, it appears to be dependent on parameters such as Reynolds number or the cavity's aspect ratio. Brès and Colonius (2008) reported for their 2M06 case that the shear layer oscillated with the first Rossiter mode. Similarly to the present investigation, in the initial stages of time marching, they also observed an additional Rossiter mode (in their case the second mode) with relevant activity, which eventually decayed completely. Furthermore, on a potential

⁷The result of subtracting the perturbed and unperturbed fields will never reach zero due to this phase modification.

three-dimensional scenario, they suggested that the interaction with the spanwise modes might affect the selection of the dominating Rossiter mode.

6.6 Family of Equilibrium Solutions across Mach Number

Section 6.4 revealed how the nature of the periodic solutions presented herein is purely dominated by compressible events. Further, as the Mach number was decreased down to the incompressible regime ($M \approx 0.30$), the amplitude of this periodic behaviour decayed to zero, giving rise to a steady state. Hence, this incompressible limit could be seen as a bifurcation point amongst the family of periodic solutions seen in the sections above, and a family of steady or equilibrium solutions. Away from the lower Mach number regime ($M \gtrsim 0.35$), these steady solutions are unstable by nature (see chapter 7), where any minimal flow perturbation will bring the equilibrium off balance, triggering a transition towards the stable periodic orbit at the corresponding Mach number. The instability of these periodic orbits is such that, just the contribution of the numerical errors will be enough to break the equilibrium of this family of flow solutions. Additionally, the steady solutions are originally unobservable from a naturally evolving flow at ‘compressible Mach numbers’; Mach numbers where the compressible effects are not negligible ($M \gtrsim 0.30$). In this section, we show the evolution of the steady solutions across Mach number from 0.25 to 0.65, with increments of 0.05. These solutions are computed using their corresponding periodic orbit as the starting point. The integration time for each optimisation was set to ten times the period of the periodic trajectory at each particular Mach number. Note that, as discussed earlier in section 6.2, the integration horizon has a major effect on the result. A longer horizon will provide further insight into how the flow instability evolves in time. Additionally, the instability of the solutions is expected to increase with Mach number due to a higher flow compressibility. Despite the flow solutions for $M \lesssim 0.30$ presented in previous sections are regarded as steady, we have run these two solutions further (as indicated in subsection 6.2.2) to ensure they satisfy the same convergence criterion than the other steady solutions⁸.

Figure 6.22 presents the family of steady solutions across Mach number, which is characterised by the norms of dilatation, kinetic energy, vorticity, viscous dissipation, alongside with $Re_{e\theta}$. Similarly to the behaviour of the periodic orbits observed earlier in figure 6.9, the norm of the dilatation of the steady solutions appears to increase exponentially with Mach number. On the other hand, the kinetic energy norm presents a continuously increasing trend proportional to M , where this time occurs regardless of the variations in the momentum thickness. Before continuing any further, it is convenient to describe first the flow topology of these solutions, which is revealed with the streamlines shown in figure 6.23. This flow configuration is the same as for the steady solution at $M = 0.30$, shown above in figure 6.12a. It consists of a weak shear layer vortex which counter-rotates on top of another vortex located at the upstream bottom end of the cavity. Furthermore, the vortex referred to as the stationary vortex in the previous sections

⁸This is the reason of the minimal disagreement in the simulation parameters presented in appendix C at $M = 0.25$ and $M = 0.30$.

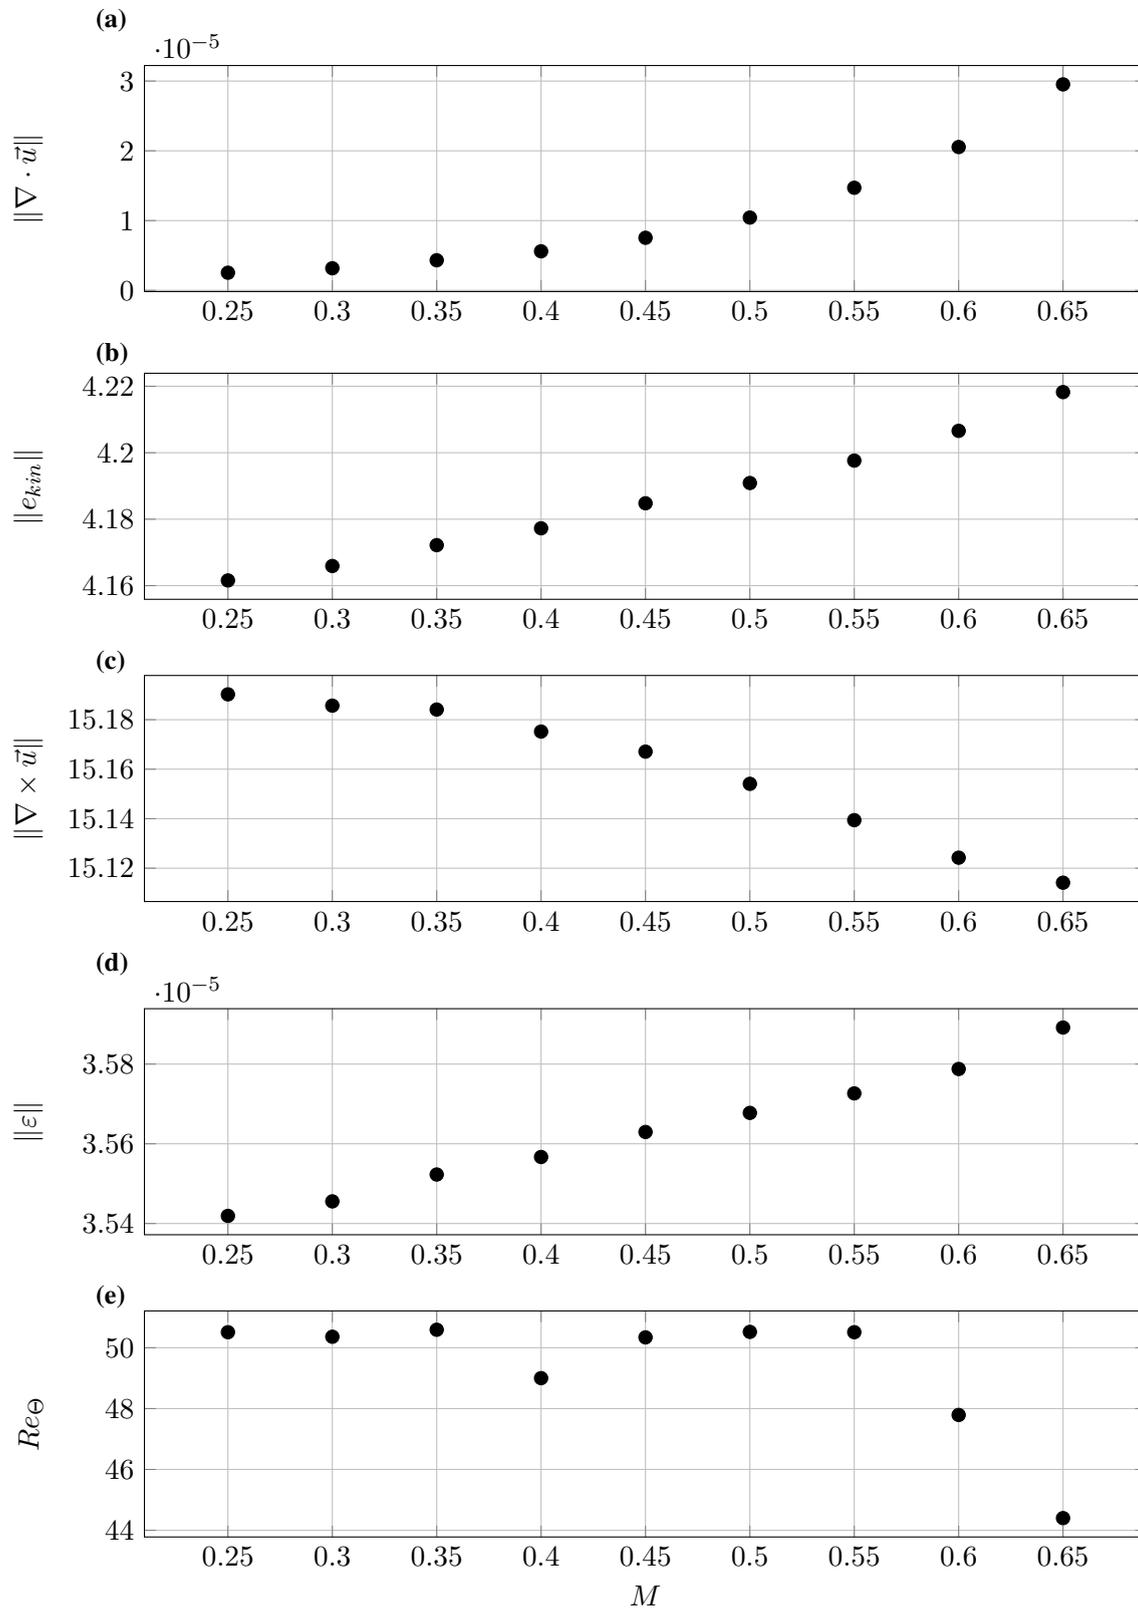


Figure 6.22: Evolution of the norms of dilatation ($\|\nabla \cdot \vec{u}\|$), kinetic energy ($\|e_{kin}\|$), vorticity ($\|\nabla \times \vec{u}\|$), dissipation rate ($\|\varepsilon\|$) and Reynolds number based on the boundary layer's momentum thickness (Re_{θ}) as a function of Mach number.

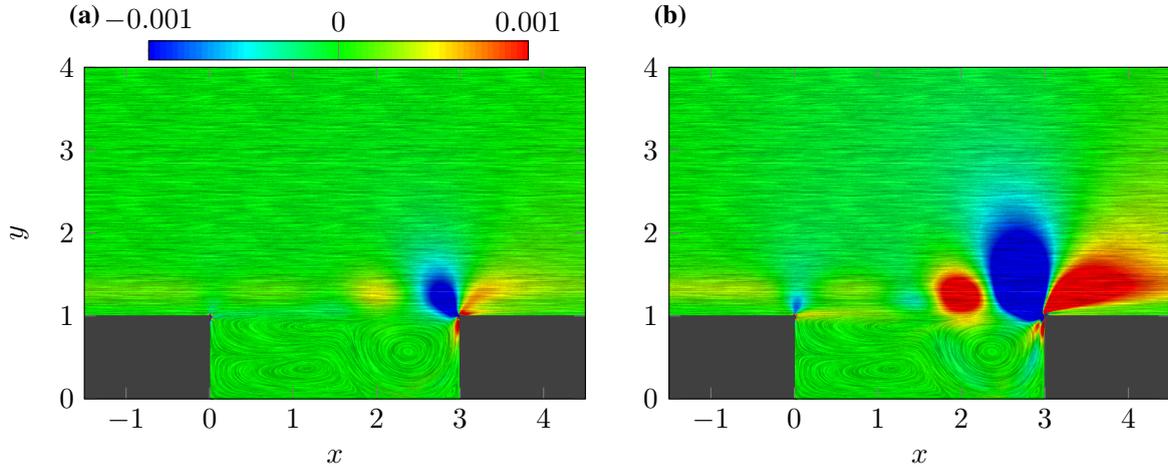


Figure 6.23: Contour plots of dilatation on the steady solutions at Mach numbers 0.25 (a) and 0.5 (b).

also takes part in the steady solutions, and it is located at the downstream end of the cavity. Additionally, from the streamlines illustrated in figure 6.23, we also perceive that the stationary and shear layer vortices are interacting amongst themselves in what it appears to be a fragile equilibrium. A small perturbation crossing the shear layer could destabilise this balance point, which would trigger the merging process between these two vortices as we saw earlier in section 6.3. Going back to the description of the compressible phenomena in the equilibrium solutions, figure 6.23 shows contours of dilatation for the steady solution at $M = 0.25$ and $M = 0.5$. This figure shows how the higher flow compressibility associated with a higher Mach number results into stronger compressible events in the vicinity of the cavity's trailing edge and shear layer. This phenomenon arises from the interaction amongst the shear layer vortex, the stationary vortex and the trailing edge. In a similar fashion as for the periodic orbits, the counter-rotating character of the shear layer and stationary vortices gives rise to a flow compression at the mid-point of the two vortices. Just downstream from this location, the stationary vortex further accelerates the flow on top of the shear layer, which causes a flow expansion. However, the strongest compressible phenomenon in this family of equilibrium solutions is the flow compression in the vicinity of the trailing edge of the cavity, which emanates from the constant flow impingement (this time steady impingement) onto the cavity's trailing edge. The increase in kinetic energy as a function of Mach number (figure 6.22b) arises from the stronger compressible effects above the stationary vortex. The greater flow expansion, shown in figure 6.23b as the red area above the stationary vortex, causes a higher acceleration on the shear layer flow on top of this vortex. Note that this higher flow velocity is the responsible of the stronger flow compression occurring at the cavity's trailing edge. Furthermore, this higher magnitude impingement also enhances the norm of the viscous dissipation, where the greater flow velocities at the impingement location produce higher values of viscous dissipation rate. Differently from all the other monitored integral quantities, the norm of vorticity shows a slight decaying trend as a function of Mach number. Due to the very low variations in vorticity, we have been unable to identify the source of this phenomenon.

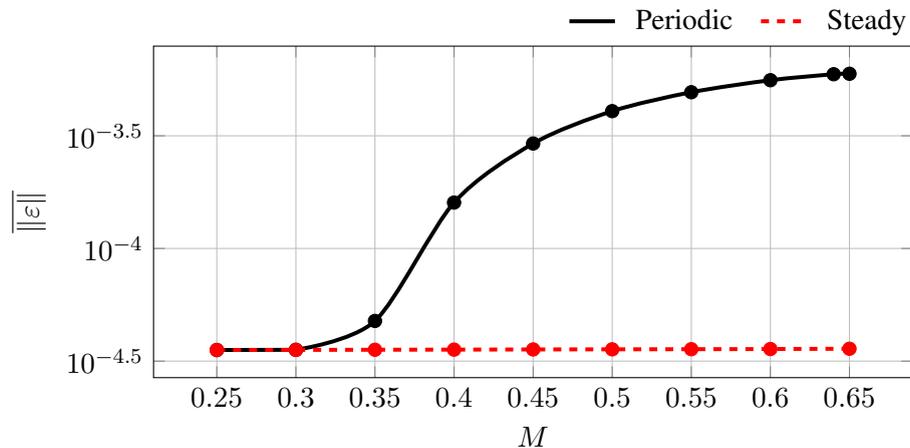


Figure 6.24: Bifurcation of the periodic and steady families of solutions across Mach number. The red and black marks represent the data points corresponding to the steady and periodic solutions, respectively. The black solid and red dashed lines show the expected trend of both families of solutions in between the points.

The bifurcation amongst the families of steady and periodic solutions is illustrated in figure 6.11. Bear in mind that this particular bifurcation appears to occur in a continuous and progressive manner, which differs from other similar scenarios reported such as the supercritical bifurcation of a 2D cylinder's wake (Provansal et al., 1987). In any case, in order to determine the bifurcation's nature, it is required a detailed stability analysis of these equilibrium solutions, which is shown later in chapter 7. The equilibrium solutions presented only show comparable changes respect to the periodic solutions in the kinetic energy. More precisely, figures 6.11b and 6.11d show how the two families of solutions meet at $M = 0.30$ (the lower left corner of both pictures) when descending in Mach number, where the other two plots show all the steady solutions overlapped into one single point. As well, figure 6.24 also shows this bifurcation between these two families of flow solutions, this time as a function of the averaged norm of viscous dissipation across Mach number. In this figure, we can readily observe how the periodic and steady families get indeed progressively closer as we descend in Mach number. In addition, this plot shows matching values for the flow solutions from both families with $M \leq 0.3$. As mentioned earlier in this section and shown later in chapter 7, this set of equilibrium solutions is unstable away from the lower Mach number regime ($M \gtrsim 0.35$). Figure 6.25 shows the evolution of the unstable steady solution at $M = 0.5$ through the equivalent of 100 periods of the corresponding periodic solution at the same Mach number. Shortly after the beginning of the simulation, the flow instability experiences low amplitude oscillations which grow exponentially in time. This amplitude growth continues until the flow-field approaches the vicinity of the periodic solution. There, the amplitude's growth rate decays rapidly, where the flow saturates and follows closely the trajectory described by the periodic orbit. For a more thorough stability analysis of these steady flow solutions, the reader is referred to chapter 7. Also, the parametric details of these equilibrium solutions can be found in appendix C.

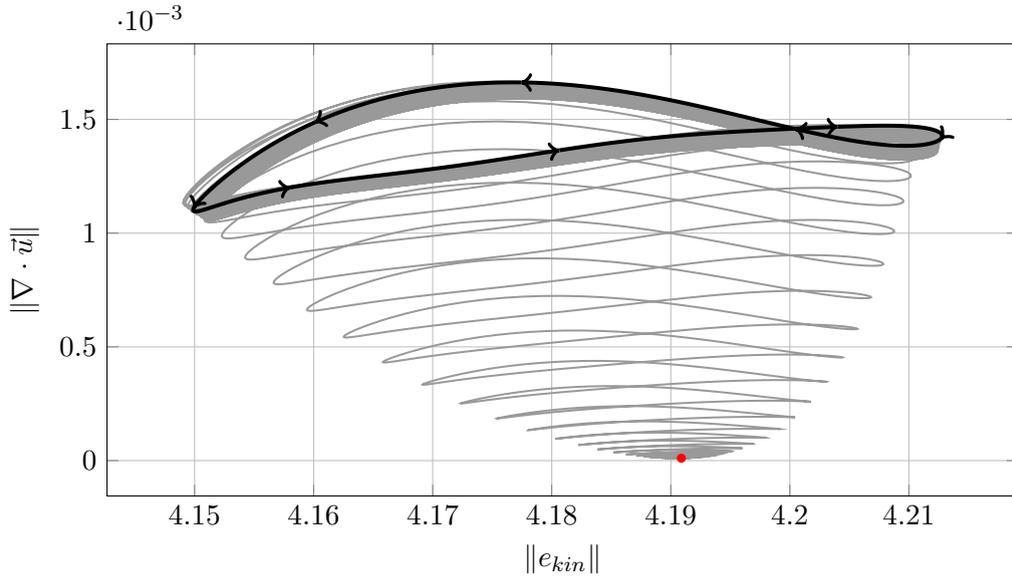


Figure 6.25: Transition from the steady solution to a fully developed flow at $M = 0.5$ (grey line). The steady or equilibrium solution is represented with a red dot, whereas the periodic solution at this Mach number is illustrated with a black line.

6.7 Summary and Conclusions

A newly developed framework to compute steady and periodic compressible flow solutions has been successfully applied on a two-dimensional open cavity flow. With this method, we have computed a family of compressible periodic flow solutions across Mach number for the first time. In addition, we have also calculated a family of equilibrium solutions using the same open cavity configuration. The Reynolds number based on the cavity depth was carefully chosen to avoid any convective instabilities in the flow-field, restricting the system to be only driven by purely compressible self-sustained oscillations. With this setup, we were able to show how the two families of periodic and steady solutions collapse in the quasi-incompressible regime ($M \approx 0.30$), which proves that the flow compressibility has a destabilising effect in cavity flows. To shed light into this not fully understood phenomenon by literature (e.g. [Yamouni et al., 2013](#)), a thorough analysis on the evolution of the compressible events across Mach number has been carried out. In particular, especial emphasis was put on the flow-acoustic shear layer interaction, which completely dominates the system's dynamics. In fact, this peculiar interaction was shown to have a major effect on the overall sound radiation in cavity flows. Furthermore, the stability of the periodic orbits was assessed by perturbing the system and tracking the growth of the perturbations, showing that these periodic flow solutions are stable even under large disturbances. On the other hand, this stable character does not seem to apply to the steady solutions, where by just time-marching the steady solution at $M = 0.5$, it commenced a transition towards its corresponding periodic orbit at the same Mach number. Moreover, the bifurcation between the families of steady and periodic orbits appears to occur in a continuous and progressive manner unlike other reported bifurcations, which adds further interest to its investigation. In any case, in order to provide a more rigorous description of the bifurcation, a more detailed stability

Chapter 6 Periodic Orbits & Equilibria in a 2D Cavity Flow

analysis of these equilibrium solutions is required, which is presented in chapter 7. Due to time-permitting limitations, we were not able to apply this framework to compute a periodic orbit with the presence of convective instabilities (at $Re_D = 5000$, for example). Hence, this task remains as future work.

Chapter 7

Stability Analysis of 2D Cavity Flow Equilibria

This chapter describes the linear stability analysis of the equilibrium flow solutions computed in chapter 6. The leading contribution presented herein is the insight over the compressible flow mechanisms which trigger an instability in these steady solutions. In addition, this analysis sheds light on the origin and nature of the bifurcation amongst the periodic and equilibrium families of solutions. Furthermore, this chapter aims to establish a connection with previous studies regarding the stability of 2D and 3D cavity flows, isolating the effect of purely compressible two-dimensional flow phenomena across Mach number. Also, the effect of the choice of base flow over the transition dynamics is briefly discussed. Moreover, these results extend the available data about stability analysis of cavity flows in terms of both aspect ratio and Reynolds number.

7.1 Introduction

In the previous chapter, we showed how the families of periodic and steady solutions of a two-dimensional cavity flow with aspect ratio $L/D = 3$ at $Re = 2000$ meet in the incompressible regime, approximately at $M = 0.30$. In ascending direction in terms of Mach number, this bifurcation occurs product of the separated shear layer becoming unstable as the flow gains in compressibility, abandoning the steady state and giving rise to a stable periodic orbit. Hence, at the bifurcation point, two branches of flow solutions appear, each of them corresponding to the families of periodic and steady solutions. In particular, away from the lower Mach number regime after the bifurcation has occurred, the steady solutions are unstable and, with time marching, they eventually undergo transition towards their corresponding stable limit cycle¹. Additionally, in chapter 6 we also observed how this divergence amongst the two branches occurs in a continuous and progressive manner, where the amplitude of the time-periodic solutions

¹See section 6.6 for a graphical representation of the transition from a steady state solution to a periodic orbit at $M = 0.5$.

grows proportionally to the Mach number, initially following Rossiter's semi-empirical formula (Rossiter, 1964). However, the reader must note that this gradual bifurcation differs from other scenarios reported in the literature, such as the bifurcation of a 2D cylinder's wake flow across Reynolds number. For that particular case, the initially steady and axis-symmetric recirculation bubble becomes unstable above the critical point $Re = 47$, triggering a transition towards radically different flow (Provansal et al., 1987), commonly known as the Von Kármán vortex street. Hence, the fact that the bifurcation studied here does not show such an abrupt behaviour, adds further interest to the investigation of the underlying physical mechanisms which cause such instability in the steady solutions presented.

The stability of several open cavity flow configurations has been studied in past investigations. Theofilis (2011) highlights this well-known flow geometry in his review about global linear instability. For example, Brès and Colonius (2008) carried out BiGlobal instability analysis (Theofilis and Colonius, 2003) on various compressible 2D and 3D cavity flows, where for the three-dimensional cases they used the two-dimensional mean flow as the base state. There they discovered that the spanwise instabilities were independent of the Mach number, where their origin was of convective (rather than acoustic) nature. This finding motivated the three-dimensional stability analysis in cavity flows assuming incompressible flow, such as de Vicente et al. (2014) or Meseguer-Garrido et al. (2014). In particular, from the neutral stability curves shown in Meseguer-Garrido et al. (2014) (their figure 7), for an aspect ratio of $L/D = 3$, the critical Reynolds number above which the three-dimensional cavity flow becomes unstable is $Re_D \approx 800$. Hence, this implies that if all the steady solutions studied here were extended to a three-dimensional space, they would certainly exhibit an unstable spanwise mode. If that were the case, the 2D and 3D modes would interact causing a frequency modulation respect to their isolated behaviour (Brès and Colonius, 2008). In that regard, we restrict our analysis to a two-dimensional space to exclusively investigate the onset of the instability related to the 2D Rossiter mode. One of the few stability analysis which used an exact solution as the basic flow was carried out by Åkervik et al. (2007). Despite that their two-dimensional cavity flow was modelled as incompressible with a relatively low Reynolds number ($Re = 350$), this equilibrium solution was unstable due to the large cavity's aspect ratio ($L/D \approx 25$), which caused the separated flow to undergo transition towards a wake mode type (Rowley et al., 2002) periodic cycle. Further, a similar study was carried out by Giannetti and Luchini (2007), this time investigating the above-mentioned instability in a 2D cylinder's wake, and also using several incompressible steady solutions as the base flow for the stability problem. The use of steady flow solutions to carry out stability analysis was extended to compressible flows by Yamouni et al. (2013). Using a finite element solver with artificial diffusion, they studied the evolution of the unstable eigenvalues of a 2D cavity flow across Mach number. With the aspect ratio $L/D = 1$ and the Reynolds number $Re_D = 7500$ used in their study, the flow was already unstable in the incompressible regime. This prevented them from isolating the origin of the instability as a function of the Mach number. Similarly to the 3D case described above, the constant presence of the unstable modes from purely convective origin is likely to cause a modulation over the onset and evolution of compressible events as a function of Mach number.

In this chapter we evaluate and analyse the global linear stability of the compressible 2D steady solutions presented previously, unveiling the flow physics which cause these solutions to become unstable as the Mach number is increased. To shed light into the stability of these exact solutions, we make use of both forward (non-linear) and adjoint (linear) Navier-Stokes solvers, thoroughly detailed in chapter 3. Respectively, these two ‘points of view’ yield information about the system’s response and receptivity, which can be combined to obtain the so-called *wave maker* or, more formally, the core of the instability. For a more elaborated introduction to global stability analysis and the use of forward and adjoint global modes, the reader is referred to section 2.3. To perform such stability analysis, we use dynamic mode decomposition, as explained in section 2.2.3, to get the eigenvalues and eigenvectors (modes) of the system straight from the data produced by our numerical framework without any further modifications. According to Schmid (2010), DMD recovers the global modes of a linear process, whereas for the non-linear case “it identifies the dominant frequencies and their associated spatial structures”. As described earlier in section 2.3, the development of a linearised forward Navier-Stokes solver is beyond the scope of the present project. For this reason, the forward global modes can only be approximated assuming linear behaviour of the full non-linear Navier-Stokes equations in the initial stages of time marching. Also, note that in this particular case the flow is undergoing a transition from a steady state to a periodic orbit. Consequently, the need for a compromise arises, since extending the dataset considerably might result in the DMD framework accounting for non-linear dynamics related to the limit cycle, and a short dataset might not produce converged results. On the other hand, the adjoint equations are linear, which allows extending the adjoint dataset as required to obtain converged adjoint dynamic modes. Hence, since the eigenvalues recovered from the forward and adjoint simulations should coincide, we use the adjoint eigenvalues as the reference and also as a measure of the veracity of the approximated forward modes.

7.2 Global Stability Analysis at $M=0.5$

In the present section, we show a full stability analysis of the steady state solution found at $M = 0.5$ and $Re = 2000$. As observed earlier in section 6.6, figure 6.25 revealed the unstable character of this particular steady solution, where it undergoes a transition from a steady state towards a periodic trajectory. For the reasons stated above, in order to approximate the forward global modes using dynamic mode decomposition, we must establish a time interval where the global behaviour of the system can be approximated as linear. Unfortunately, such information cannot be found in figure 6.25, where the projection of the transition is not represented as a function of time. Therefore, a more convenient manner to obtain information about the duration this linear transition is as displayed by figure 7.1.

The value of the cost function (6.1) is now shown across time. The initial steep increase in the cost function is the result of the flow solver reacting to an initial condition which was computed externally. For instance, the boundary conditions and derivative routines adapt the steady solution to satisfy the physical conditions imposed at the boundaries, and also fit the flow-field onto

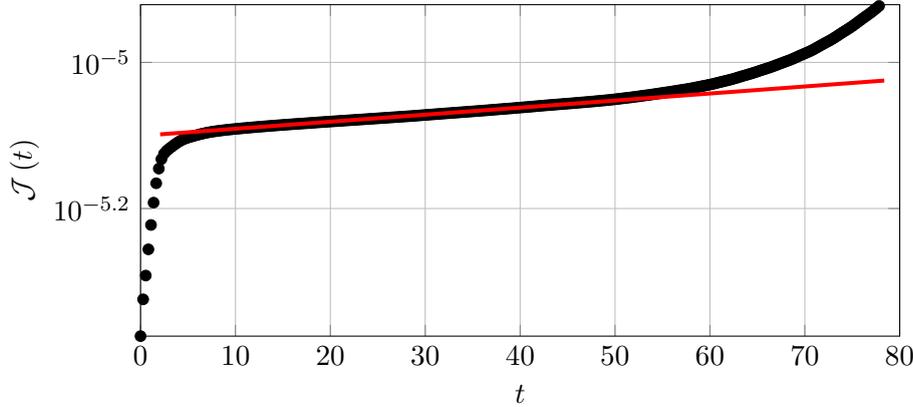


Figure 7.1: Steady-state cost function evolution as a function of non-dimensional time units. The red line represents the best fit of the linear flow behaviour during the transition.

a polynomial function of the same order as the derivative routines. After this initial ‘rejection’ is overcome, one would expect that the divergence of the flow-field from the initial steady solution would be much more drastic once the non-linear effects start taking place. Unfortunately for the present investigation, there is no way to tell exactly when this happens due to the lack of a linearised Navier-Stokes solver. Alternatively, we can estimate the start of the non-linear flow behaviour from figure 7.1, as the point where the trajectory of the cost function abandons the linear trend exhibited approximately from 7 to 48 time units. Furthermore, the converged eigenvalue spectrum obtained through the adjoint simulations can help to select the dataset which yields the most accurate eigenvalues. The adjoint dynamic mode decomposition was carried out using 801 snapshots sampled equidistantly across a time span of 220 time units². To avoid spurious behaviour accounted for by DMD from the boundary conditions, the snapshot sequences (both adjoint and forward datasets) were cropped down to the vicinity of the cavity. Finally, the forward dataset from 12.375 to 46.75 time units produced the leading eigenvalue closest to the one calculated using adjoint data.

Figure 7.2 shows the eigenspectrum calculated using the forward and adjoint simulations. These eigenvalues have been transformed to a more convenient form using the relations $\mu_r = \Re\{\log(\lambda_n)\}/\Delta t$ and $\mu_i = \Im\{\log(\lambda_n)\}/(2\pi\Delta t)$ (see section 2.2.3), indicating growth rate and Strouhal number³, respectively. The resulting leading (most unstable) eigenvalues yielded by both forward and adjoint approaches are very similar. The exact values produced by the adjoint DMD for this particular mode are $\mu_r = 0.042644$ and $\mu_i = 0.246062$, where the forward DMD results are just 3.33% off. Also note the closeness of both approaches on the secondary eigenvalue at $\mu_r = -0.007837$ and $\mu_i = 0.153670$. The remaining unmatched eigenvalues predicted by the forward DMD appear as a direct consequence of an insufficiently long snapshot sequence to produce a fully converged spectrum. This phenomenon was also observed when studying the

²The sampling frequency has been kept constant through all the DMD realisations presented in this chapter. Additionally, the convergence of the adjoint DMD results has been ensured using a shorter data subset which produced the same results.

³Strouhal number also defined as $St = \frac{fL_\infty}{U_\infty}$.

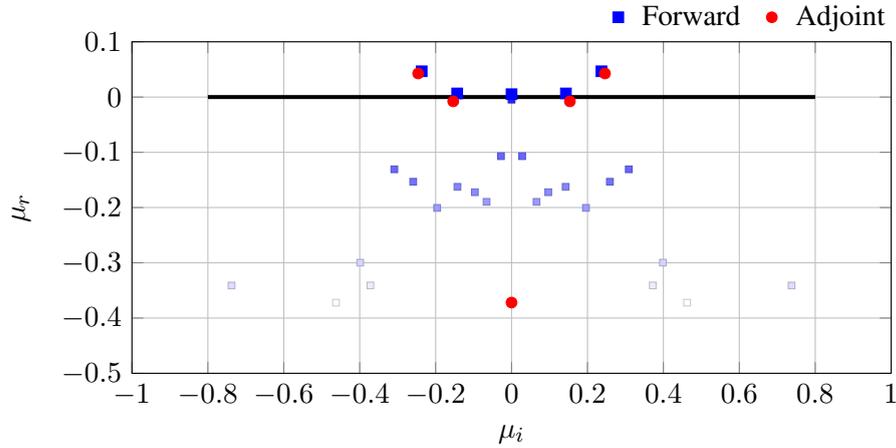


Figure 7.2: Eigenvalue spectrum for the steady solution at $M = 0.50$. The red circular markers show the eigenvalues calculated through the adjoint simulations. The blue squared symbols represent the approximate eigenvalues obtained using the non-linear forward simulations. The smaller squared symbols show purely numerical modes (noise) product of the insufficiently large forward snapshot sequence, where their colour intensity gives a measure of each mode's amplitude.

convergence of the adjoint spectrum, where most of the damped eigenvalues ($\mu_r < 0$) further decayed as the snapshot sequence was incremented. These stable eigenvalues (with $\mu_r < -0.5$) have been excluded from the above plot due to their lack of dynamical relevance. Overall, considering the assumption of linear flow dynamics from a non-linear Navier-Stokes solver, added to the limited forward dataset, the results achieved by the forward DMD can be categorised as satisfactory. Furthermore, figure 7.2 can be seen as the verification of both forward and adjoint compressible Navier-Stokes solvers, alongside with the DMD post-processing tool.

Figure 7.3 presents the forward and adjoint modes corresponding to the most unstable eigenvalue. From top to bottom, this figure shows the forward (left) and adjoint (right) components of this mode for the four state variables⁴ in $Q = [p, \rho u, \rho v, \rho]$ and $Q^* = [\rho^*, m_u^*, m_v^*, p^*]$. As introduced previously in chapter 2, the forward and adjoint modes can be seen, respectively, as spatial maps of the system's response and receptivity, both relative to their associated eigenvalue. Hence, since the current equilibrium solution is globally unstable, the forward mode reveals the laminar shear layer and the reattached laminar boundary layer as the leading amplifiers of this instability. In particular, the streamwise and vertical momentum variables (figures 7.3c and 7.3e, respectively) show the largest amplitudes of this mode. Contrarily, the density (figure 7.3g) and pressure (figure 7.3a) components appear not to be as active. On the other hand, it is expected that the density component will gain in amplitude as the Mach number increases due to a larger flow compressibility. Above all, note that all four forward mode components exhibit a gradual increase in their activity along the shear layer in the downstream direction. Apart from the mode amplitude, the locations where the mode switches sign reveal essential information to understand the underlying physical mechanisms which drive the mode dynamics. In

⁴See definition of Q and Q^* in chapter 3

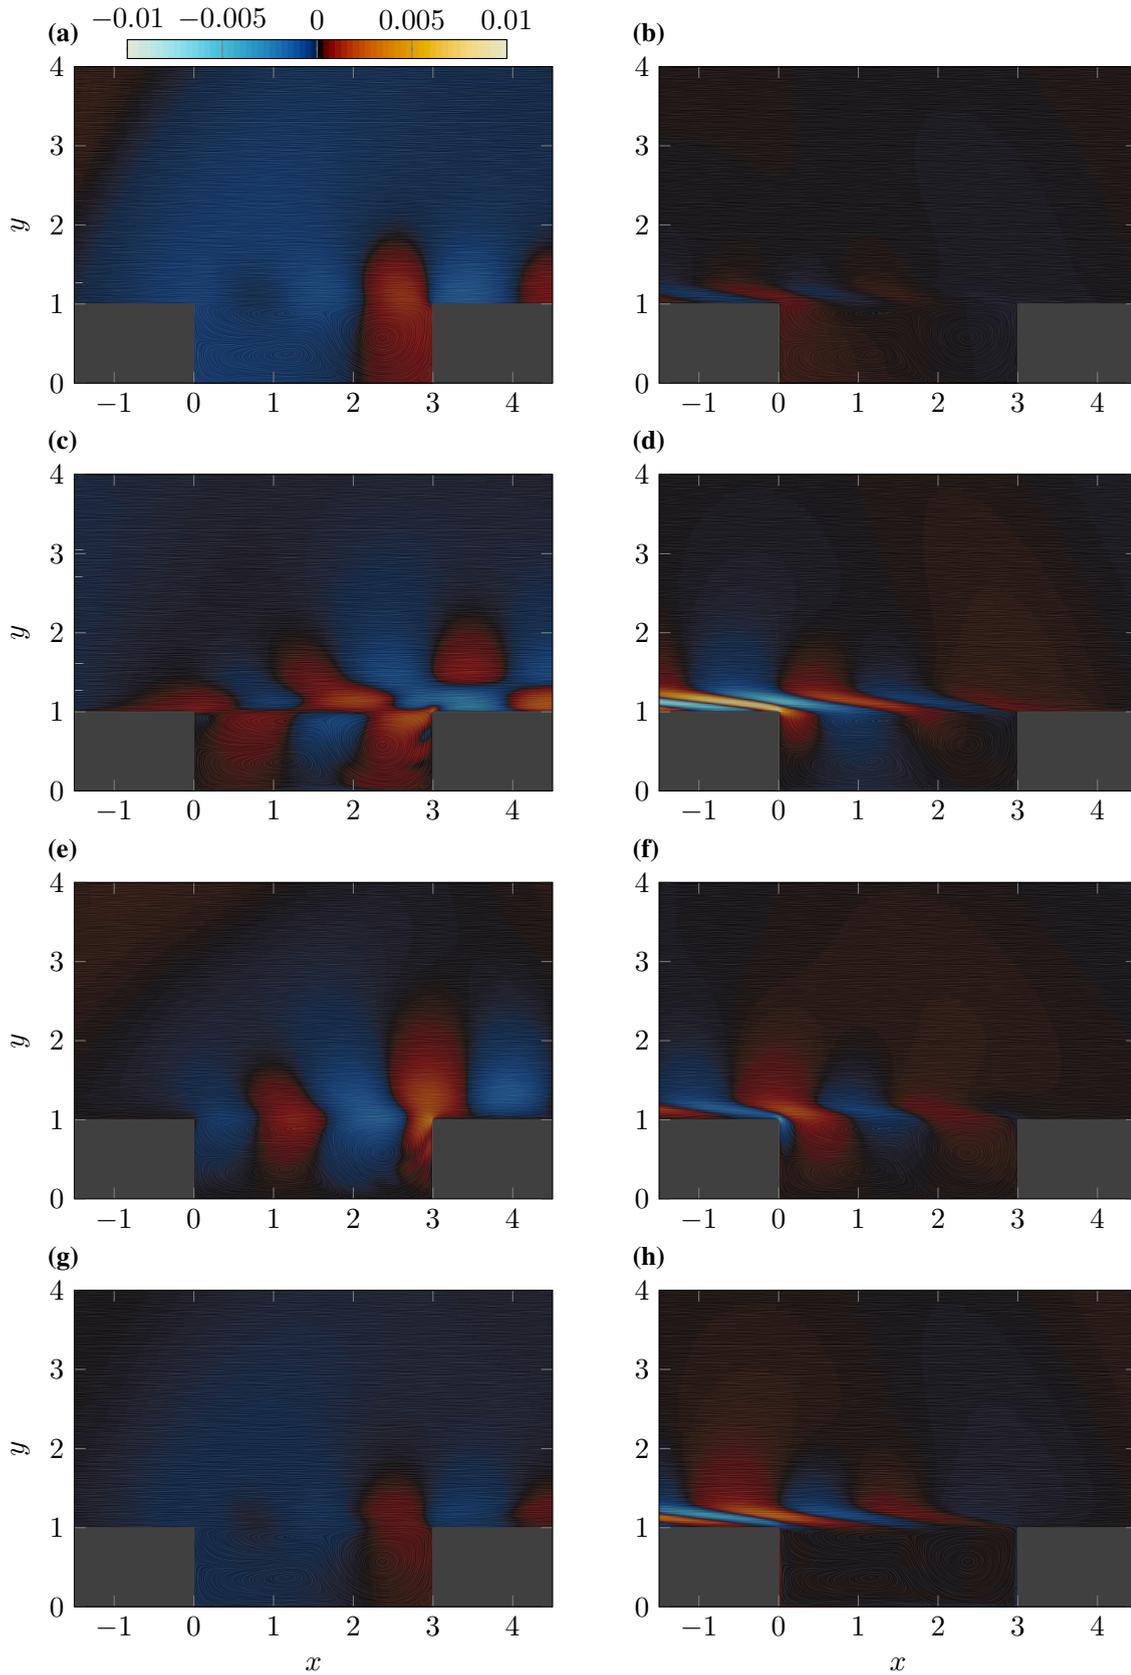


Figure 7.3: Forward ((a), (c), (e), (g)) and adjoint ((b), (d), (f), (h)) modes associated with the leading eigenvalue for the equilibrium solution at $M = 0.5$. From top to bottom, the figures show the components of Q and Q^* . Contour plots are overlapped with the streamlines from the steady solution.

substance, loci with contrary sign represent opposite flow behaviour, such as increase-decrease of pressure, momentum, density, etc. For reference purposes, the streamlines calculated from the velocity field of the steady flow solution are superimposed to the contour plots in figure 7.3. Curiously, the momentum mode shapes highlight the close relationship between the vortices located in the shear layer and also at the downstream end of the cavity⁵. If we analyse in detail the streamwise momentum component, we observe that it switches signs in the proximity of the vertical centreline of the vortices located in the shear layer and at the downstream end of the cavity. This phenomenon suggests that both vortices are compressed and expanded in the streamwise direction at the eigenvalue's frequency. Further, as also seen in figure 7.3c, similar sign changes occur across the laminar shear layer in the vertical direction, revealing its unstable character. Additionally, the vertical momentum component also shows relevant activity along the laminar shear layer. Figure 7.3e uncovers patches of opposite increase in vertical momentum, where one of the sign changes takes place just in between the above-mentioned shear layer and downstream vortices. Hence, the combination of the streamwise and vertical momentum components induces fluctuations in the laminar shear layer which affect both shear layer and downstream vortices. These two vortices begin an oscillating motion with increasing amplitude, which eventually breaks this fragile equilibrium, resulting in the downstream vortex absorbing the shear layer vortex. As detailed in chapter 6, the extra streamwise momentum resulting from this merging process impinges onto the cavity's trailing edge, which radiates upstream an acoustic wave, setting off the Rossiter's mode. From this point onwards, the system can no longer be approximated as linear and begins the last stage of its transition towards the stable limit cycle (see figure 6.25).

So far with the analysis above, we know the physical mechanisms associated with the unstable eigenvalue, which trigger a transition from a steady state towards a periodic orbit in the present 2D cavity flow. In other words, we only know that our steady flow solution is unstable and how it reacts to that instability. In order to delve into deeper questions, such as the reason or the origin of the instability, we require the adjoint mode of that unstable eigenvalue. As stated above the adjoint modes can be interpreted as receptivity maps which indicate how to most efficiently trigger their respective forward mode. Hence, the adjoint mode shown in figure 7.3 (right) reveals the receptivity of its forward mode (figure 7.3 - left). At this point, it is worth recalling that the adjoint formulation used links pressure p with adjoint density ρ^* and density ρ with adjoint pressure p^* (see section 3.2). Similarly to the forward mode, the streamwise and vertical adjoint momentum components (figures 7.3d and 7.3f, respectively) present a considerable activity, where this time, the adjoint pressure (figure 7.3h) also shows an amplitude of comparable magnitude. On the other hand, opposite to the behaviour exhibited by the forward mode, the adjoint mode shows increasing amplitude in the upstream direction, a product of the backwards time marching, highlighting the vicinity of the cavity's leading edge as a receptivity 'hot spot'. This high receptivity of the leading edge of the cavity, the separated shear layer and also the incoming boundary layer, is responsible for the global instability of the flow solution studied here. Any perturbation in the state vector Q (primarily density and momentum perturbations),

⁵For a detailed description of the flow topology of the equilibrium solution, the reader is referred to chapter 6.

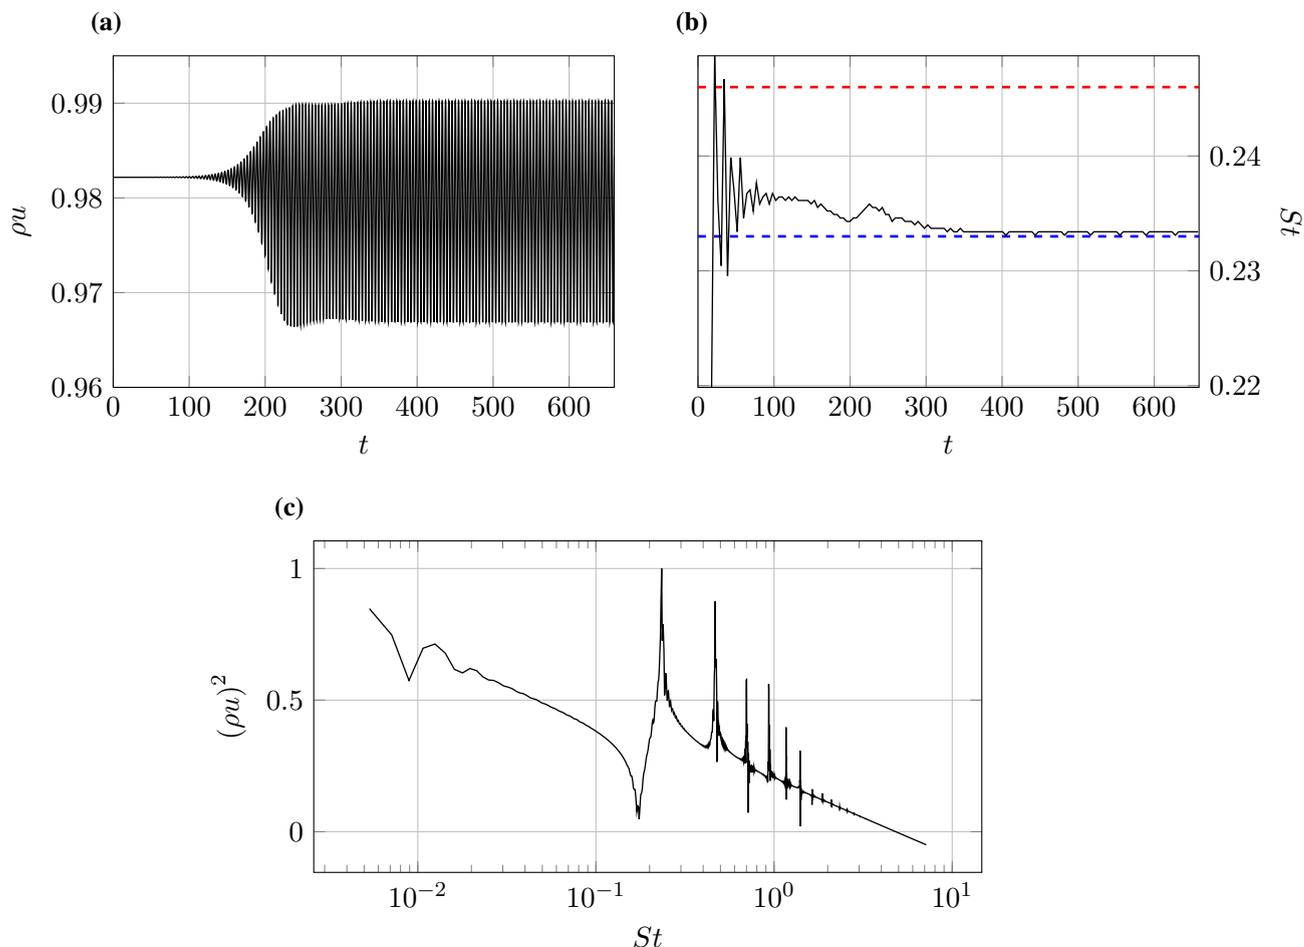


Figure 7.4: **(a)** Time signal of the streamwise momentum at $\vec{x} = (2.25, 1)$. **(b)** Estimated flow-acoustic feedback loop's frequency evolution across time. Red dashed and blue dashed horizontal lines indicate the frequencies associated the unstable mode and periodic limit cycle, respectively. **(c)** Frequency spectrum of the time signal show in **(a)**.

where its corresponding adjoint mode component is not zero will set off the instability growth. For the current equilibrium solution, the numerics from our compressible Navier-Stokes solver, alongside with the numerical precision of the flow solution, suffice to trigger this unstable mode. Analogously to the acoustic-convective feedback loop occurring in periodic cavity solutions (see chapter 6), any minimal flow imbalance which encounters the cavity's trailing edge will reflect upstream a weak acoustic wave. This upstream travelling acoustic wave eventually reaches the proximity of the leading edge of the cavity, setting in motion the unstable mode. Hence, once the unstable mode is active, it will further disturb the flow-field as shown in figure 7.3, producing a more energetic acoustic wave than the precedent in the following impingement on the trailing edge. This feedback loop continues fuelling the flow instability until the linear flow dynamics collapse. Once this occurs, the natural frequency from the flow-acoustic feedback loop progressively evolves from the frequency of the (linear) unstable mode of $St \approx 0.246$, to the characteristic frequency of its corresponding limit cycle (Rossiter mode) of $St \approx 0.233$.

A similar disagreement amongst the frequencies associated with the unstable eigenmode and the non-linear periodic solution was also observed by [Sipp and Lebedev \(2007\)](#) in an incompressible flow over a square cavity. Figure 7.4a shows the streamwise momentum time signal through this transition, captured by a monitor point located at the shear layer. At this precise location, the unstable forward mode is considerably active, which permits the estimation of the system's frequency by tracking the local minima (or maxima) of this signal across time. Since there is only one single unstable mode, the system's frequency should asymptotically approach the frequency of this unstable eigenvalue, which governs the long time behaviour of the system. Figure 7.4b reveals this frequency estimation as a function of time, where it appears to become stable after approximately 100 time units. The fact that the initial strong frequency oscillations ceased is an indication that all the stable modes have decayed significantly, leaving the unstable mode as the only one active in the system. At this point in time, we readily observe that the approximated frequency has already abandoned the frequency associated with the unstable mode, which confirms that the system's dynamics have stopped being governed by linear flow mechanisms ([de Vicente et al., 2014](#)). Furthermore, bear in mind that the linear behaviour ceases even before the vortex merging of the shear layer and downstream vortices begins, which occurs approximately after 135 time units. After this, the flow undergoes the last stage of the transition, where the system's frequency slowly decays to the Rossiter mode's frequency. Contrary to [Brès and Colonius \(2008\)](#), the transition of this case does not seem to trigger other Rossiter modes (figure 7.4c). The origin of this different behaviour is believed to reside in the different base flows employed. Note that here we make use of a steady exact flow solution as base flow, whereas, in their study, they employed an averaged flow-field, which is not itself a flow solution. Hence, from there on, the flow reacts to that artificial flow as if it were a perturbation, triggering other leading Rossiter modes as seen also in chapter 6.

Following the concept of structural sensitivity introduced back in section 2.3, we can now combine the forward (approximated) and adjoint modes to reveal the origin of the instability or wave maker. Such analysis reveals the area of the flow-field which acts as 'the driver of the oscillation' ([Luchini and Bottaro, 2014](#)), highlighting the regions with both high receptivity and response. These spatial maps are also referred to as sensitivity to a localised feedback ([Giannetti and Luchini, 2007](#)). Figure 7.5 shows the wave maker regions for the streamwise and vertical momentum components. Pressure and density plots have not been included as their activity is orders of magnitude lower. These structural sensitivity maps expose predominately both the shear layer and the incoming laminar boundary layer as the origin of the instability. Especially, the areas showing a greater intensity are located in the proximity of the cavity's leading edge, where the adjoint mode peaks. In addition, as observed in figure 7.5a, the laminar boundary layer only contributes to the onset of the instability in the streamwise direction. In fact, the point of highest magnitude is located just above the cavity's leading edge, before this incoming laminar boundary layer separates from the wall. Moreover, the high structural sensitivity patches exhibited along the shear layer in figure 7.5b are also strongly linked to the flow structures located underneath. The two weakest patches (yet non-negligible) are located in the shear layer, precisely above the cores from the downstream and shear layer vortices. Overall, for this

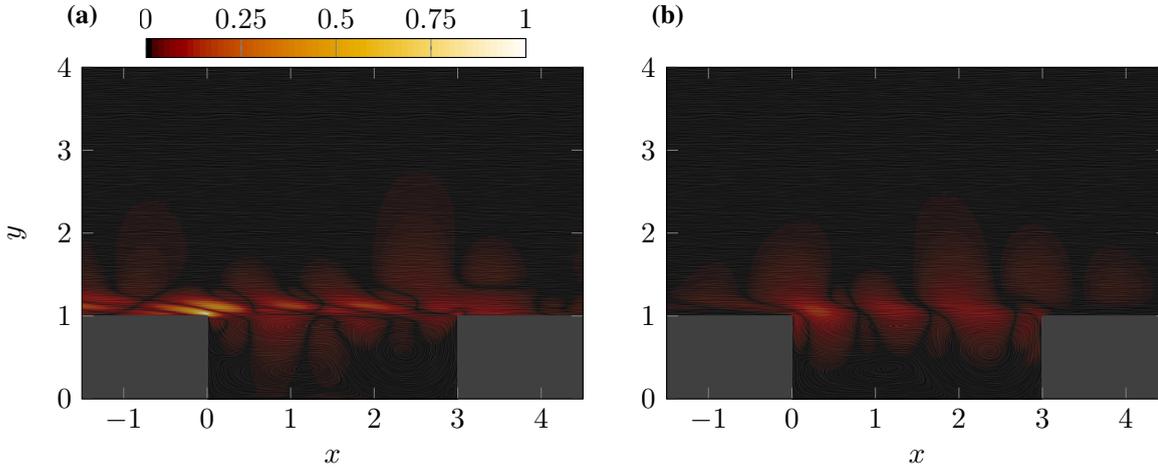


Figure 7.5: Structural sensitivity maps of the streamwise **(a)** and vertical **(b)** momentum components, calculated as $|\Psi_{m_u}^*| \cdot |\Psi_{\rho u}|$ and $|\Psi_{m_v}^*| \cdot |\Psi_{\rho v}|$, respectively. Contour levels are normalised with $6 \cdot 10^{-6}$. Figures are overlapped with the streamlines from the steady solution.

distinct geometry, the wave maker region is distributed along the incoming boundary layer and shear layer, peaking in intensity at the cavity's leading edge, to then decay towards the trailing edge. In particular, the activity shown in figure 7.5 vanishes rapidly after the cavity's trailing edge, a result of the aforementioned convective-acoustic feedback mechanism not holding past the trailing edge, and thus yielding an almost negligible receptivity.

7.3 Stability Evolution across Mach Number

The Mach number dependency of both families of periodic and steady solutions was thoroughly studied in chapter 6. There, the Mach number changes produced alterations in the flow compressibility and propagating velocity of the acoustics, which modified the flow mechanisms which drive the Rossiter feedback loop occurring in the periodic solutions. Additionally, in the family of equilibrium solutions, these Mach number modifications introduced alterations in the density gradients, which increased proportionally with M_∞ . On the other hand, the flow topology of these equilibrium solutions remained unaffected by the Mach number changes. Hence, in the present section, we evaluate how these changes in flow compressibility and the associated acoustic speed translate into the stability of these equilibrium solutions. In order to carry out this stability analysis across Mach number, we just refer to the eigenspectrum and eigenmodes resulting from our adjoint Navier-Stokes framework. Performing such analysis with the forward non-linear solver might lead to erroneous conclusions for the reasons stated above. Hence, obtaining and studying the forward eigenmodes will remain an item on the to-do list, focusing the present analysis on a receptivity perspective.

Figure 7.6 shows the eigenvalue spectrum for the steady solutions from Mach numbers 0.30 to 0.65. To not overpopulate the picture, the eigenspectrum from the $M = 0.25$ solution has

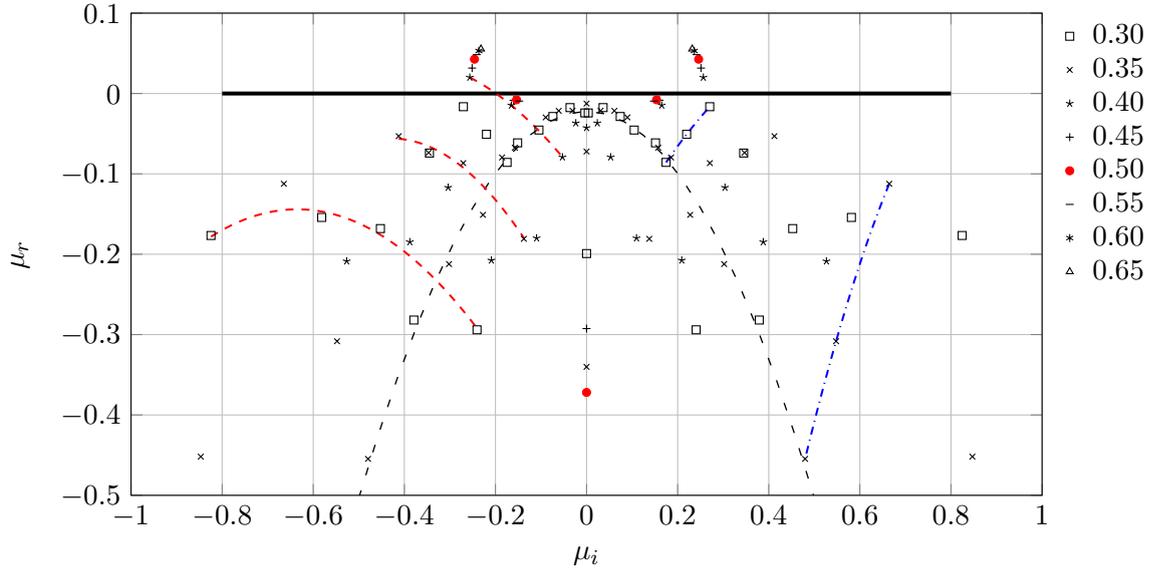


Figure 7.6: Eigenspectrum of the family of equilibrium solutions across Mach number. The red dashed lines show the rise of the unstable branch at Mach numbers 0.3, 0.35 and 0.4. The blue dash-dotted lines represent the evolution of a decaying branch as the Mach number is increased. The thin black dashed line highlights the branch appearing due to noise in the stable steady states.

not been included since all the eigenvalues corresponding to that Mach number lack dynamical relevance ($\mu_r \lesssim -20$). Being practical and stepping away from the numerics, this flow solution is, by all means, steady and stable. In this figure 7.6, it is straightforward to notice that the steady solutions at $M = 0.3$ and $M = 0.35$ are stable, where all their eigenvalues are $\mu_r < 0$. Contrarily, all the other steady states are globally unstable since they have at least one eigenvalue with $\mu_r > 0$. Similarly, Brès and Colonius (2008) also observed that their 2D base flow was stable at $M = 0.35$, but equivalently to the present investigation, it became unstable when raising the Mach number up to $M = 0.6$, leaving the rest of parameters unchanged⁶. Recalling the periodic orbits documented in chapter 6, a stable and periodic flow solution was found also at $M = 0.35$ when descending in Mach number from $M = 0.4$. This shows that both steady and periodic flow solutions are stable at this particular Mach number, where the flow will evolve to one or the other depending on the initial condition chosen. Additionally, this suggests that the transition from steady to periodic solutions does not occur at the same Mach number as the transition in the opposite direction. At the same time, this can be interpreted as the cavity flow just being compressible enough to maintain the self-sustained oscillations (when descending in Mach number), but not as compressible to trigger them (in ascending direction). Moreover, the coexistence of these two stable solutions at the same Mach number is known as bistability, and it is characteristic from subcritical Hopf bifurcations. Figure 7.7 shows a detailed illustration of this type of bifurcation fitted with the data points from the families of periodic and steady flow solutions. In ascending direction in terms of Mach number, the

⁶Additionally, the same base flow at $M = 0.35$ became unstable when slightly increasing the Reynolds number (based on the cavity depth) from 1500 to 1585.

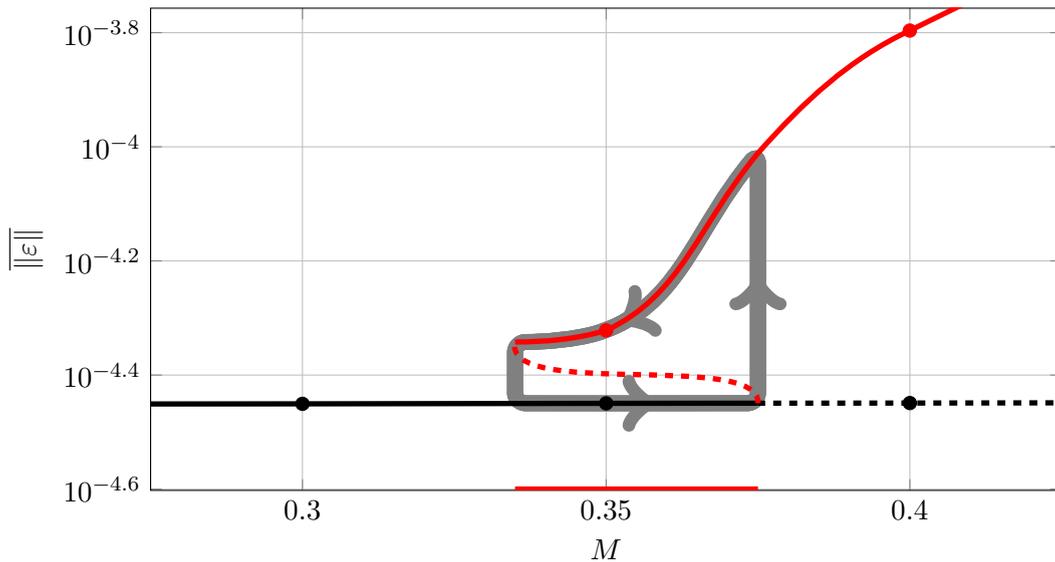


Figure 7.7: Illustration of the subcritical Hopf bifurcation of the periodic and steady families of solutions across Mach number. The solid lines represent stable flow states, whereas the dashed lines show unstable flow configurations. The red and black coloured lines indicate the periodic and steady families of solutions, respectively. Also with the same colour coding, the circle symbols show the actual flow solutions computed in this investigation. The hysteresis loop is depicted with the thick grey solid lines, where the arrows show the loop direction. For this illustration, the Mach number range where bistability occurs is represented with the red coloured horizontal axis.

family of equilibrium solutions is stable up to the subcritical bifurcation point, located between Mach numbers 0.35 and 0.40. At this bifurcation, two unstable branches emerge, giving rise to unstable equilibrium and periodic solutions. Hence, far from following any of the unstable branches past the bifurcation point, the system undergoes a transition towards the corresponding stable limit cycle at the same Mach number. Once the flow reaches this periodic orbit, further variations in Mach number will displace the fluid system along the family of periodic solutions as described earlier in chapter 6. However, if the Mach number is reduced below the lower limit of the bistability range, the periodic orbit ceases from being stable, which leads to a rapid decay towards its corresponding steady state, also at the same Mach number. Thus, the system's bistability leads to a hysteresis loop around this bistable range (see figure 7.7), which results into the two different transitions described above, from a steady state to a periodic orbit and vice versa. Unfortunately, in order to determine the exact Mach number at which these transitions occur, a finer sampling of flow solutions would be required.

Going back to figure 7.6, in an attempt to shed light on the origin of the instability as a function of Mach number, the leading eigenvalue can be traced back from the unstable to the stable regime, and there its eigenmode can be analysed. Note that the eigenvalues with a positive growth rate from the solutions ranging from $M = 0.4$ to $M = 0.65$ are remarkably aligned in the unstable plane. Intuitively, it might be tempting to relate these unstable eigenvalues to the leading (stable) eigenvalue from the $M = 0.3$ solution, which is placed just below. On the other hand, if that were the case, the leading eigenvalue corresponding to the $M = 0.35$ equilibrium solution

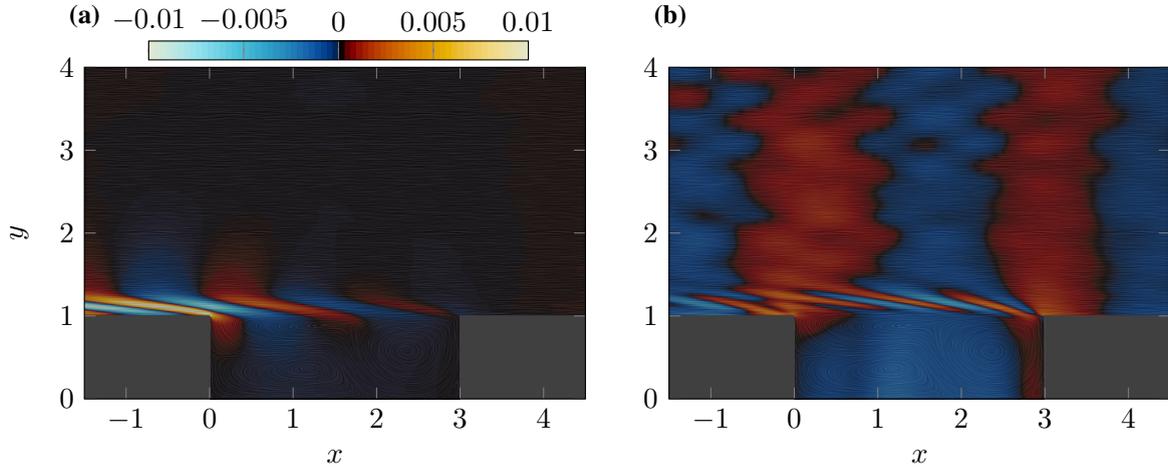


Figure 7.8: Streamwise momentum component of the adjoint mode corresponding to the leading eigenvalue of the decaying branch (highlighted with blue dash-dotted lines in figure 7.6) at $M = 0.3$ (a) and $M = 0.35$ (b).

should be placed somewhere in between. But this is not the case. Studying figure 7.6 in more detail, we readily observe the branch containing this leading eigenvalue at $M = 0.3$ (highlighted by the blue dash-dotted line) has its analogous branch at $M = 0.35$, which exhibits a further damped character, with also a much higher frequency. This fading behaviour agrees with figure 7.8, where the respective adjoint streamwise momentum components of the leading mode from this branch reflects a substantial loss in receptivity, particularly at the cavity's leading edge and incoming boundary layer. Furthermore, the appearance of large vertical patterns in the free-stream in figure 7.8b, alongside the above-mentioned phenomena, advocate that this particular branch is decaying towards the high-frequency and low-amplitude acoustic regime. In addition, at this point, it is worth highlighting that the eigenvalues aligned in a parabolic shape at Mach numbers 0.3 and 0.35 are believed to be associated with the noise present in those simulations⁷ (Bagheri, 2014).

Moreover, figure 7.6 also shows a rising eigenvalue branch from $M = 0.3$ to $M = 0.4$, highlighted with red dashed lines. In contrast to the preceding blue dash-dotted decaying branch, this new branch exhibits a least stable nature as the Mach number of the steady solution is increased, moving towards a lower frequency regime. Especially, at $M = 0.4$ its leading eigenvalue presents a positive growth rate, sitting in the unstable plane. For reference purposes in the current chapter, this branch is also denoted as the unstable branch, where its eigenvalues are sorted numerically with decreasing characteristic frequency. Figure 7.9 displays the evolution of some of the eigenvalues from this unstable branch across Mach number from a streamwise momentum receptivity point of view. Respectively, the first and fourth modes at $M = 0.3$ are shown in figures 7.9a and 7.9b. There, the areas with high receptivity appear to displace continuously from the incoming laminar boundary layer towards the shear layer as the mode's frequency increases along the branch. Additionally, as seen in figure 7.6, the eigenvalues from this branch get closer together as the Mach number is raised. As reflected in figures 7.9c and

⁷In fact, this parabola disappears for immediately higher Mach numbers.

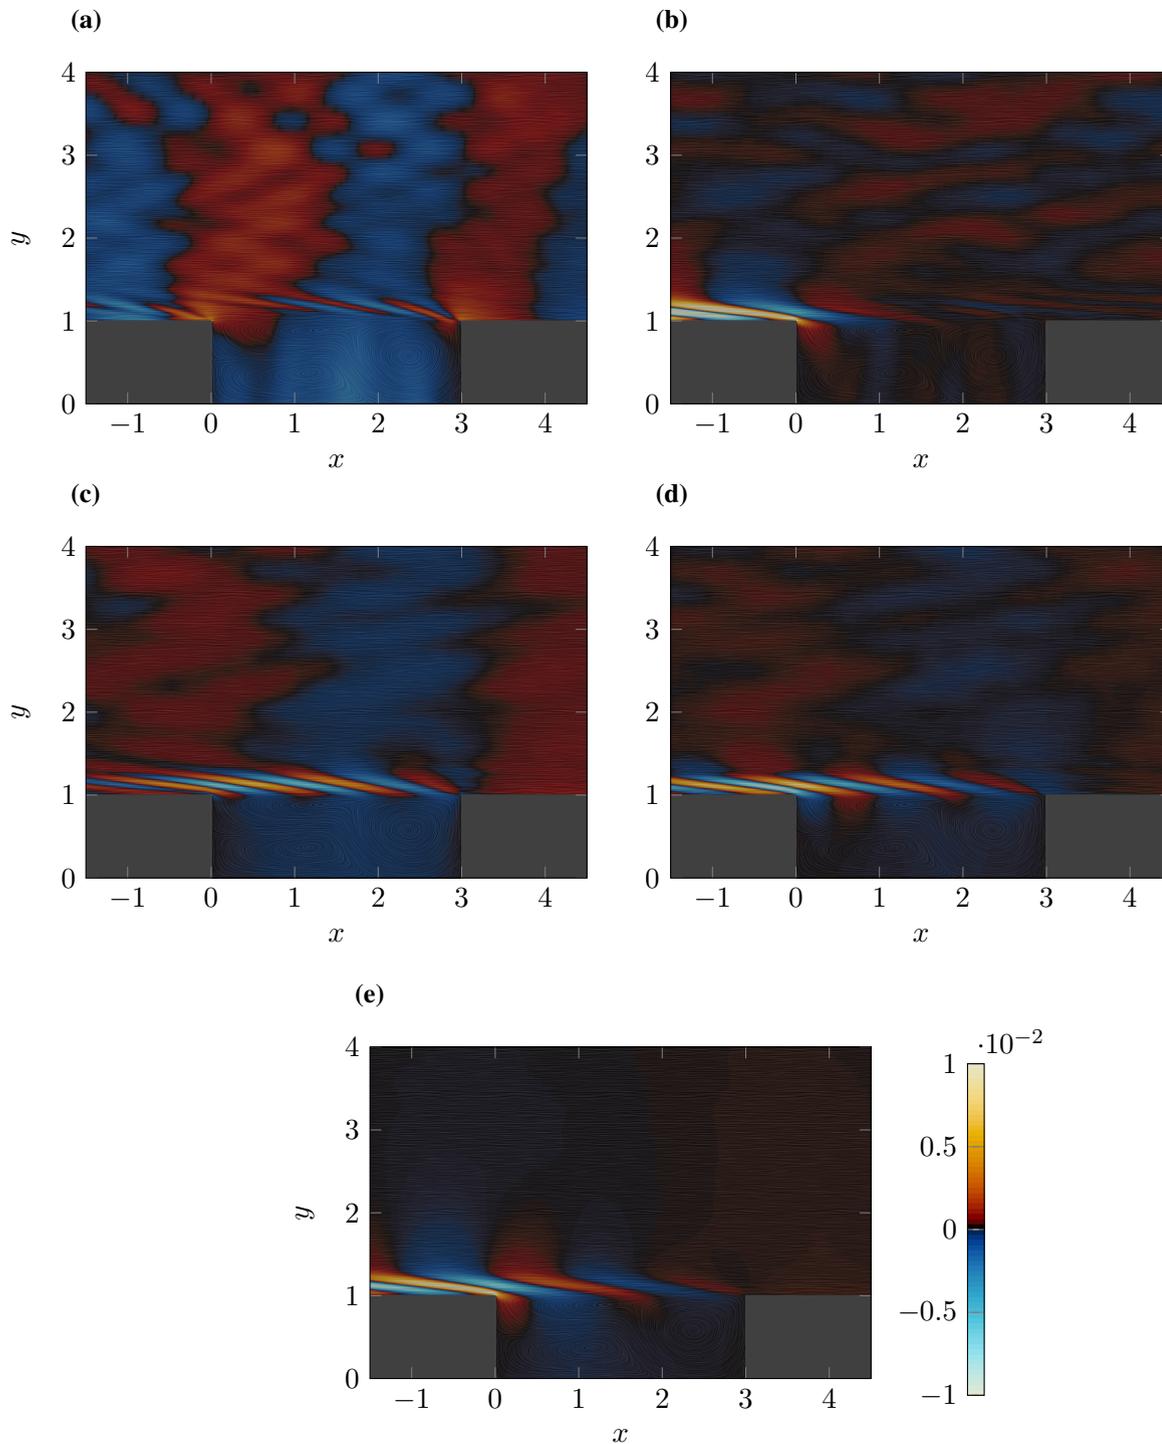


Figure 7.9: Evolution of the streamwise momentum receptivity of some of the eigenvalues from the unstable branch across Mach number. The first mode is shown in figures (a), (c) and (e) at Mach numbers 0.30, 0.35 and 0.4, respectively. The fourth mode at $M = 0.3$ and the second mode at $M = 0.35$ are also shown in (b) and (d).

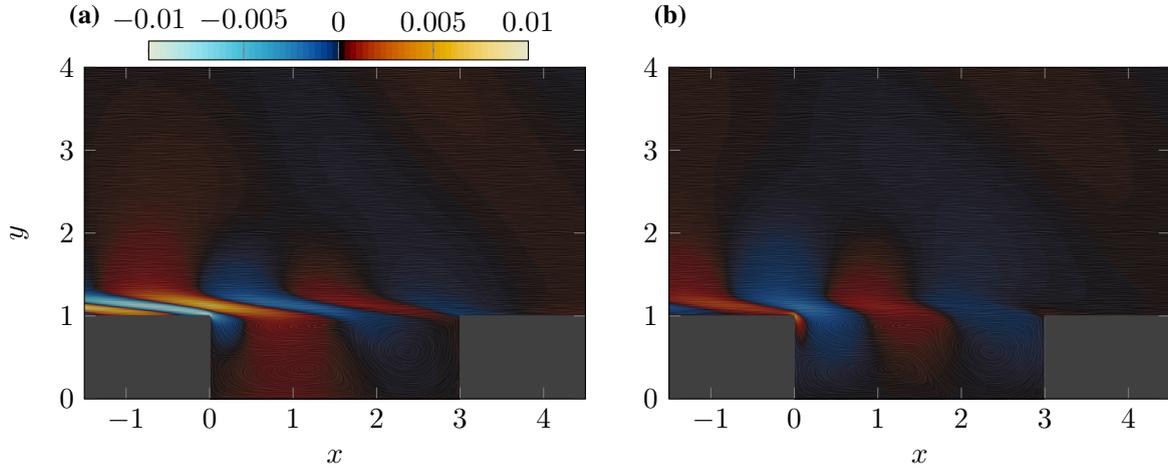


Figure 7.10: Streamwise **(a)** and vertical **(b)** adjoint momentum component of the unstable mode at $M = 0.65$.

7.9d, this phenomenon also occurs in their respective eigenvectors, where the first and second adjoint modes at $M = 0.35$ show a very similar receptivity (both amplitude and shape-wise) in the streamwise momentum component. Hence, if this approaching behaviour would continue, the first and second modes would merge eventually. This would result into only three eigenvalues forming the unstable branch, which is precisely the case observed at $M = 0.4$. Unfortunately, a finer sampling of equilibrium solutions across Mach number would be required in order to rigorously demonstrate this potential eigenvalue merging, which is believed to have a relevant influence on the onset of the unstable eigenvalue.

Lastly, figure 7.9e reveals the adjoint streamwise momentum component of the unstable mode at $M = 0.4$. Inspecting the progression of this particular adjoint mode across Mach number, we detect how the dominant activity of this mode displaces gradually from the shear layer towards the proximity of the cavity's leading edge and incoming boundary layer. Once there, this 'receptivity migration' appears to come to a halt. This phenomenon is confirmed by the adjoint momentum components of the unstable mode at $M = 0.65$ as shown in figure 7.10, exhibiting a remarkably similar receptivity to its corresponding modes at Mach numbers 0.4 and 0.5. Perhaps the only noticeable difference between these three plots resides in the content of the free-stream, where a more distinct downstream⁸ reflected wave appears with increasing flow compressibility. Note that an analogous Mach number effect was shown in the Rossiter modes taking place in time-periodic flow solutions (see chapter 6). In the same manner, once the unstable mode of the steady solution is active, a higher flow compressibility results into a stronger upstream-travelling acoustic wave reflected off the cavity's trailing edge. This phenomenon boosts the gain of the acoustic-convective feedback loop, yielding a larger growth rate μ_r . In addition, the close relation amongst the periodic orbits and the stability of the steady solutions is manifested as well in the evolution of the frequency of this unstable mode. Similarly to these flow periodic trajectories, the non-dimensional frequency μ_i of the unstable mode

⁸Downstream in the sense of the forward simulation.

M_∞	$St_{periodic}$	μ_i	μ_r
0.30	—	0.824514	-0.176693
0.35	0.24871	0.412341	-0.053106
0.40	0.24472	0.256365	0.019984
0.45	0.23923	0.251265	0.031612
0.50	0.23355	0.246062	0.042644
0.55	0.22822	0.241150	0.048355
0.60	0.22319	0.236504	0.052891
0.65	0.21792	0.231924	0.055347

Table 7.1: Comparison across Mach number of the non-dimensional frequency associated with the periodic solutions and the first eigenvalue of the unstable branch of the corresponding steady solution.

decreases as the Mach number raises (see table 7.1). The reason behind this behaviour is simply the reduction in the propagating speed of sound relative to the convective flow velocity, which increases the lag of the aforementioned acoustic-convective feedback loop. Relating these results to previous studies, the decreasing trend in the unstable mode’s frequency as a function of Mach number was also observed in Yamouni et al. (2013). Also analogously to our findings, they observed the appearance of new unstable modes as the Mach number was raised (five extra modes at $M = 0.5$ and seven at $M = 0.9$). As advanced earlier in section 7.1, in their investigation, the constant presence of unstable modes from purely convective origin is prone to exert a frequency and growth rate modulation over the exclusively compressible unstable modes, and vice versa. In fact, the proportional increase in growth rate as a function of Mach number came as a surprise to them, when comparing their results with literature. Hence, in the present investigation, we have shed light into the physical mechanisms which cause both aforementioned increase in growth rate, and appearance of additional unstable modes with Mach number from an exclusively compressible receptivity point of view, covering an existing gap in the literature.

7.4 Summary and Conclusions

The linear global stability of the equilibrium flow solutions computed in chapter 6 was studied in depth. DMD was used to recover the adjoint global modes and eigenspectrum. In addition, the forward global modes were also approximated for the steady solution at $M = 0.5$. This was accomplished using DMD over a non-linear forward dataset, assuming linear flow behaviour in the initial stages of time marching. The resulting eigenvalue spectrum was compared with the adjoint results for verification purposes. The forward approximated modes revealed the laminar shear layer as the leading amplifier of the instability. Further, the respective adjoint modes highlighted the cavity’s leading edge and incoming boundary layer as the regions with the highest receptivity. Additionally, the evolution of the frequency corresponding to the unstable mode was tracked throughout the steady-to-periodic transition at $M = 0.5$. This showed that the system quickly abandons the linear behaviour, even before the shear layer and downstream

vortices begin the merging process. The wave maker region was also computed showing activity in the boundary layer and shear layer, with maximum intensity at the cavity's leading edge, to then decay towards the trailing edge. The evolution of the eigenvalue instability was also analysed across Mach number. This showed that the transitions from steady to periodic flow and vice versa do not occur at the same Mach number, product of the subcritical Hopf bifurcation of the steady and periodic families. This proves that periodic or steady solutions are not necessarily exclusive in cavity flows, where it has been demonstrated that they both can be stable over a short Mach number range. Moreover, the eigenvalue branch which eventually becomes unstable as the Mach number is raised was also identified. The lack of both forward modes and a finer sampling of solutions across Mach number permits us only to hypothesise on the onset of the unstable character, and analyse such modes from a receptivity point of view. Hence, a forward stability analysis and a finer sampling of steady flow solutions (in particular from $M = 0.35$ to $M = 0.40$) remain future work. Overall, this chapter has provided a new insight into the compressible physical phenomena which makes exact steady solutions more unstable as the Mach number increases.

Chapter 8

Noise Radiation and Directivity Control of a 2D Cavity Flow

This chapter presents a new application of adjoint-based optimisation. In particular, we make use of this method to perform noise radiation and directivity control of a two-dimensional open cavity flow, which is the main contribution of this chapter. In addition, we present the resulting optimal actuation, showing the ideal forcing patterns for such control in cavity flows. These results are valuable information which can be used in future investigations to model other flow actuation strategies.

8.1 Introduction

The control of separated flows has caught the attention of numerous researchers for several decades. The current computational capabilities permit the use of more advanced techniques for their study, which give a new perspective to develop flow control schemes. The use of the adjoint method to fluid flows is gaining in popularity amongst the flow control community to carry out gradient-based optimal flow control. Examples of successful applications of adjoint-based optimal flow control are turbulence reduction in a channel flow ([Bewley et al., 2001](#)), aero-acoustic jet noise reduction ([Cerviño et al., 2002](#); [Kim et al., 2014](#)), or optimal porous media distribution in an aerofoil's trailing edge to minimise aero-acoustic radiation ([Schulze and Sesterhenn, 2013](#)). Note that due to the high computational demands of the method, it is a common practice to make simplifications of the flow, such as assuming incompressibility ([Blackburn et al., 2008](#)), constant viscosity ([Vishnampet et al., 2015](#)) or even linearising the dynamics of the flow ([Fosas de Pando et al., 2012](#)).

In the current chapter we make use of the previously introduced adjoint-based optimal flow control framework for compressible DNS (see chapters 2, 3 and 4). The code is applied to a two-dimensional separated flow scenario to explore the capabilities and limitations of the method.

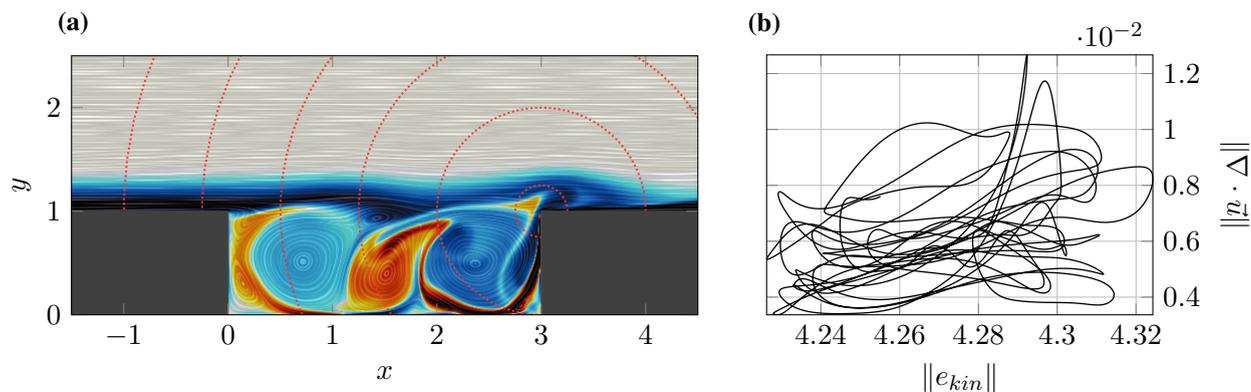


Figure 8.1: **(a)** Illustration of the 2D cavity flow. Contours show instantaneous z -vorticity. Red dashed lines represent the sound radiation from the trailing edge. **(b)** Phase portrait of the trajectory followed by the non-actuated cavity flow-field.

Similarly to the past two chapters, the flow configuration chosen for the present investigation is an open cavity flow with aspect ratio $L/H = 3$ (figure 8.1a). This time, the Reynolds number is boosted to $Re = 5000$, where the Mach number is kept as $M = 0.5$. Despite its simple geometry, this flow configuration gives rise to complex non-linear flow phenomena with constant flow separation and constant reattachment. Kelvin-Helmholtz type instabilities grow in the separated shear layer, creating non-linear structures which impinge on the trailing edge of the cavity. This unsteady interaction radiates acoustic waves that also propagate upstream (see chapter 6). Due to the high receptivity of the leading edge, these compressible flow events further excite some of the unstable shear layer modes (see chapter 7). The concatenation of these phenomena results in a feedback loop that is responsible for self-sustained oscillations, which are also known as Rossiter modes (Rossiter, 1964). Note that in comparison with the cavity flows studied in chapter 6, the increase in Re provokes a more chaotic flow behaviour (figure 8.1b). For further reading about open cavity flows, a summary of the recent computations and experimental techniques can be found in Lawson and Barakos (2011). Regardless of being a widely studied flow, the amount of literature referring to the active control of cavity flows is rather scarce. Amongst these few published articles, without getting into the details, the most popular method to actuate such flow is through linear closed-loop control (for example, Rowley et al., 2005, 2006; Henningson and Åkervik, 2008; Barbagallo et al., 2009; S. J. Illingworth and Rowley, 2012; Nagarajan et al., 2013). For further details and other control methodologies in open cavity flows, the reader is referred to the review articles from Rowley and Williams (2006) and Cattafesta et al. (2008). In addition, one of the more recent contributions to passive control in cavity flows, using for the first time an adjoint-based approach, is the work of Liu et al. (2016). Yet so far, to the best of the author's knowledge, the active control of open cavity flows using the aforementioned adjoint-based method remains unexplored by literature. Hence, in the present study, we employ for the first time active adjoint-based optimal flow control on a two-dimensional cavity flow. The aim of this work is to minimise the acoustic radiation at the sensor location actuating on the upstream flow, where in particular, we focus on the importance of the sensor and actuator locations.

8.2 Computational Setup

The cavity flow considered in this chapter has a laminar inflow condition with Reynolds number $Re_D = 5000$, based on the cavity depth and free-stream velocity. The cavity depth is also chosen to non-dimensionalise the streamwise and vertical directions. With the coordinate origin at the lower-left corner of the cavity, the domain ranges from -20 to 20 in x and from 0 to 10 in y . The grid consists of a total of 480000 points, where the DNS resolution was only kept in the vicinity of the cavity. The Reynolds number based on the momentum thickness at the separation point is $Re_\Theta \approx 33.8$, where the non-dimensional free-stream velocity is $U_\infty = 1$. The Prandtl and Mach numbers are respectively $Pr_\infty = 0.72$ and $M_\infty = 0.5$. To carry out the adjoint simulations, the flow field is stored every 250 time-steps (at intervals of 0.1375 time units) and it is linearly interpolated in between for every sub-step of the Runge-Kutta scheme. Note that this implicitly decreases the Nyquist frequency of our sensor, acting as a low-pass filter.

To penalise the noise radiation from the trailing edge of the cavity, a convenient cost function for this particular problem is

$$\mathcal{J}(Q, \phi) = \frac{1}{T} \int_0^T \int_\Omega \frac{1}{2} (p - \bar{p})^2 W_s(\vec{x}) \, d\Omega dt, \quad (8.1)$$

where T represents the optimisation horizon, \bar{p} is the averaged pressure over this time period and $W_s(\vec{x})$ is a spatial weighting function that defines the sensing region. In the present investigation, we use three different functions $W_s(\vec{x})$ to study their effect over the resulting control. The first two sensing regions are defined as

$$W_s(\vec{x}) = e^{-\frac{(x-x_0)^2 + (y-y_0)^2}{0.05}}, \quad (8.2)$$

located downstream (at $\vec{x}_0 = (4, 3)$) and upstream (at $\vec{x}_0 = (-1, 3)$) of the cavity. According to Åkervik et al. (2007), a downstream located sensor will efficiently capture the shear layer response, since it is where the corresponding forward global mode has a large amplitude (see also chapter 7). On the other hand, in compressible cavity flows, the leading acoustic radiation occurs in the upstream direction as part of the convective-acoustic feedback mechanism which constitutes the Rossiter modes (see chapter 6). Hence, the sensor placed upstream of the cavity will adequately gather these dominant acoustic-related pressure fluctuations. Furthermore, a third sensing scenario is studied, where the overall sound radiation from the cavity is accounted for. Defined as

$$W_s(\vec{x}) = e^{-\frac{(r_s - \sqrt{(x-1.5)^2 + (y-1)^2})^2}{0.15}} - e^{-\frac{(y-1)^2}{0.2}}, \quad (8.3)$$

this sensor is an arc of radius $r_s = 5$ and centred at $\vec{x}_0 = (1.5, 1)$, which effectively covers the sound radiation in all directions. Note that a ramp-down function has been added to switch off the sensor in the near wall region ($y = 1$), contributing to the stability of the adjoint simulations.

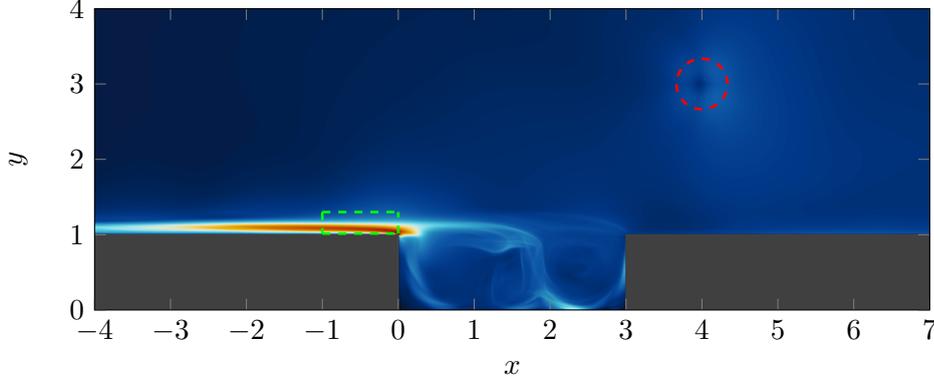


Figure 8.2: Root mean square of the magnitude of the adjoint momentum \vec{m}^* when using the downstream sensor. Red and blue contours indicate high and low values, respectively. The green dashed rectangle indicates the actuating sub-volume. The red-dashed circle represents the downstream sensing region.

To exploit the advantages of the adjoint method, the actuation is chosen to be a time-dependent body forcing sub-domain $\vec{f}(\vec{x}, t)$, where every grid point within the sub-domain has two control parameters (the streamwise and vertical components of the forcing vector \vec{f}). Including both cost function and forcing terms in the Lagrangian functional (2.3), which are then substituted into (3.11), leads to their respective adjoint terms which can be written as

$$-\frac{\partial \rho^*}{\partial t} = RHS + \underbrace{(p - \bar{p}) W_s(\vec{x})}_{\text{Cost Function}} \quad (8.4)$$

$$-\frac{\partial p^*}{\partial t} = RHS + \underbrace{(\vec{m}^* \cdot \vec{f}) W_a(\vec{x})}_{\text{Forcing Terms}}. \quad (8.5)$$

From (2.11) the gradient for each of the control parameters follows as

$$\frac{D\mathcal{J}}{D\phi_x} = \rho m_u^* \quad \text{and} \quad \frac{D\mathcal{J}}{D\phi_y} = \rho m_v^*. \quad (8.6)$$

To choose an appropriate location for the actuation, an initial statistical study of the adjoint quantities is carried out for an entire time horizon T , which gives a measure of the flow sensitivities. Given the time-dependence of the chosen actuation, its control authority should be proportional to the temporal fluctuations of \vec{m}^* . Figure 8.2 reveals that the area with the highest root mean square values in \vec{m}^* is in the boundary layer before the flow separation point, reaching maximum values right after the leading edge of the cavity. In addition, this information is also supported by the results discussed earlier in chapter 7. There, adjoint mode and wave maker plots relative to the cavity's unstable shear-layer mode marked the vicinity of the leading edge as a both receptivity and structural sensitivity 'hot-spot'. On that account, the proximity of the cavity's leading edge is the locus of the flow-field where we can most effectively modify the response of such unstable mode, which is the responsible for the sound radiation. Hence, the

location of our actuation is defined as

$$W_a(\vec{x}) = \begin{cases} 1 & \text{if } \vec{x} \in [(-1, 0), (1.015, 1.3)] \\ 0 & \text{if } \vec{x} \notin [(-1, 0), (1.015, 1.3)] \end{cases}. \quad (8.7)$$

The actuating sub-domain contains a total of 8000 grid points. In a similar fashion as for the flow variables in the adjoint simulations, the control parameters are updated every 250 time iterations and they are linearly interpolated in between every capture. Note that this also removes implicitly the high frequencies from the controller. For optimisation using the downstream sensor, the receding horizon is chosen to be $T = 19.25$ time units, based on a preliminary study on the growth of the adjoint field at the controller location. The leap forward after the optimisation over the time horizon has converged is $T_a = 8.25$ time units. Each optimisation is run for 5 complete horizons, which covers a total of 52.25 time units. This results in a total of 2.256 million control parameters per receding horizon T , and 6.096 million control parameters over the five total optimisation periods. The fact that this sensing region is located downstream of the cavity, in a region of the flow with strong convection, reduces the communication time between the sensor and the actuation. On the other hand, this is not the case for both upstream located and arc sensing regions, where a longer communication time is required. Hence, for these two sensors, the optimisation horizon is increased to $T = 38.5$ time units.

8.3 Optimal Control Results

The entire optimisation required a total of 333 function and gradient evaluations (76 downstream sensor, 77 upstream sensor and 180 arc sensor). In addition, 4 extra forward simulations were required for each case, in order to advance the optimisation between receding horizons once the previous one reached convergence. For the present study, the adjoint computations required a similar amount of computational effort as the forward simulations. As an example, figure 8.3 illustrates the significant reduction in cost for all receding horizons, registered at the downstream sensor. The local sudden increases in the cost function value (especially in T_1) are caused by a poor estimation of the control parameters by the control update algorithm (L-BFGS), which are rapidly corrected in the next iteration. The convergence criterion consisted in the last two function evaluations of every horizon to be equal down to three significant figures, and also being the lowest value of the series. Due to limitations in computational resources, the maximum number of function evaluations per receding horizon was limited to 50. Once the iterative optimisation has concluded for all three scenarios, the effect of the three controllers is evaluated over the entire optimisation horizon by performing controlled and uncontrolled simulations for 52.25 time units each. The resulting values of the cost function for each case are summarised in table 8.1.

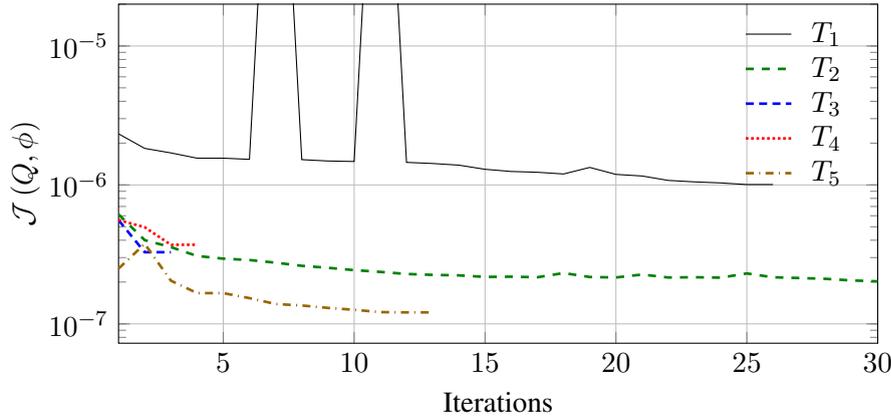


Figure 8.3: Downstream sensor case. Evolution of the cost function for the five receding horizons.

sensor	no control	control	$\Delta\mathcal{J}$
downstream	$2.679 \cdot 10^{-6}$	$5.458 \cdot 10^{-7}$	-79.6%
upstream	$3.356 \cdot 10^{-5}$	$2.673 \cdot 10^{-5}$	-20.4%
arc	$2.489 \cdot 10^{-3}$	$4.472 \cdot 10^{-4}$	-82.0%

Table 8.1: Cost function values for the three different sensors with and without flow actuation.

Before analysing the quantitative results, figure 8.4 shows the contours of the root mean square of pressure (p_{rms}) for the non-actuated (figure 8.4a) and actuated simulations. Ordered alphabetically, figures 8.4b, 8.4c and 8.4d represent the p_{rms} contours for the downstream, upstream and arc sensor cases. For reference purposes, the locations of the upstream and downstream sensors are also represented in the plots. Bear in mind that for all four cases we use a laminar inflow, which forces the activity of these p_{rms} contours in the free-stream flow to be exclusive from the acoustic perturbations. Hence, these figures provide a direct impression of the magnitude and directivity of the cavity's sound radiation. Moreover, the activity of the shear layer can also be estimated through the p_{rms} contour levels at that location. Similarly to what we saw before in chapter 6, the overall sound radiation from a cavity flow is a combination of the flow impingement onto the trailing edge which radiates upstream an acoustic wave, and even more importantly, the interaction of the shear layer vortices with this upstream-travelling acoustic wave. This complex phenomenon was broken down (from a fully compressible flow perspective) as a synergy of several dipoles, where their alignment completely determines not only the leading direction where the sound is emitted but also the regions with insignificant sound. Therefore, comparing the uncontrolled case with figure 8.4b, we readily observe how the locus with the weakest sound levels, originally in the proximity of the downstream sensor, has spread further covering the sensor, as a consequence of the actuation. Far from moving the flow-field towards a complete steady state (the ideal flow solution where no sound is emitted), our adjoint-based controller has opted to change the noise directivity to optimally cancel out the

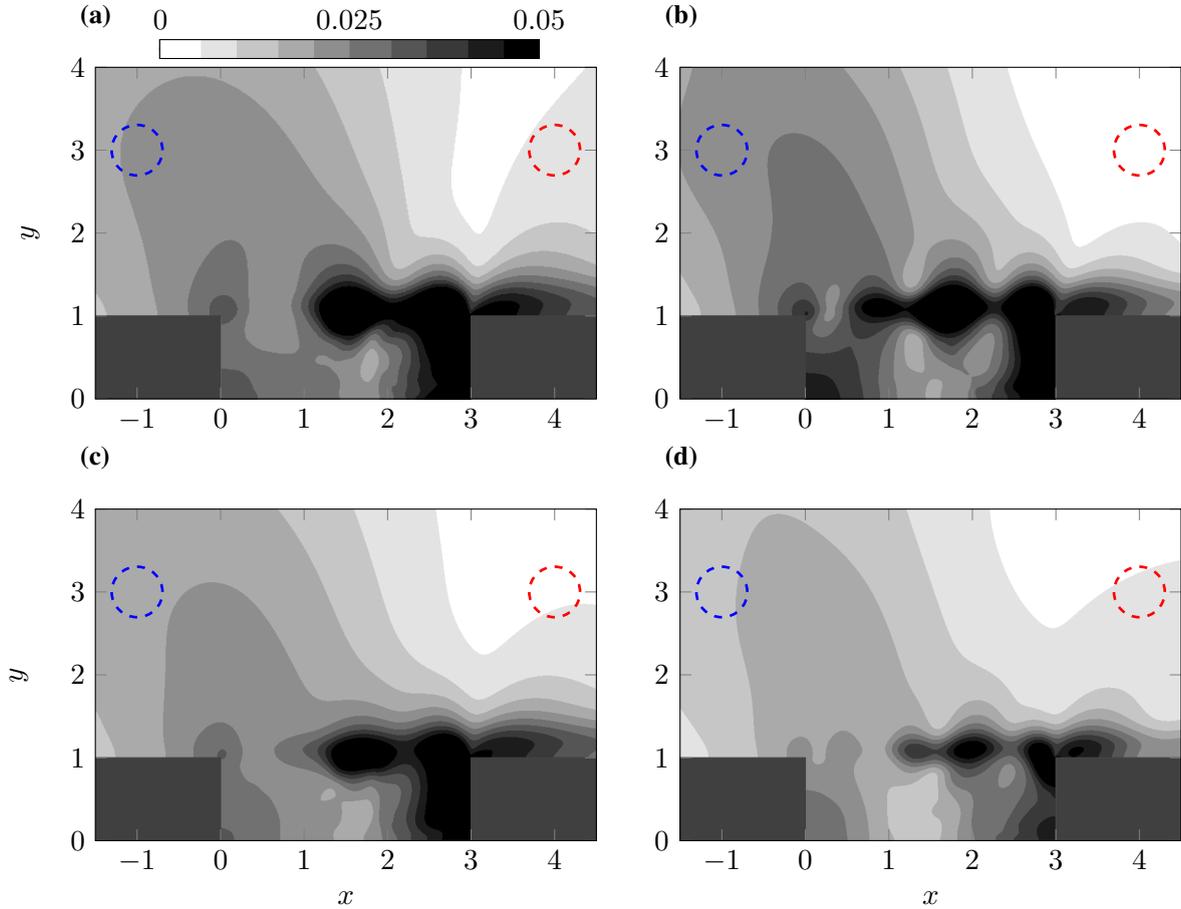


Figure 8.4: Root mean square pressure (p_{rms}) contours for uncontrolled **(a)** and controlled cases. Ordered alphabetically, **(b)**, **(c)** and **(d)** represent the p_{rms} contours for the downstream, upstream and arc sensor cases, respectively. Locations of upstream and downstream regions are indicated with blue-dashed and red-dashed circles.

noise levels just at the sensor location, by altering the aforementioned dipole interaction. As a side effect, both upstream sound radiation and shear layer now show enhanced p_{rms} contours¹. On the other hand, with the sensor located upstream (figure 8.4c), the actuation now reduced significantly the sound radiation in the upstream direction. Incidentally, the p_{rms} contours also show a notorious sound reduction at the downstream sensor, despite it was not active during this optimisation. From this particular case, it is worth highlighting the different p_{rms} patterns exhibited in the shear layer in comparison with the previous optimisation. This suggests that the flow actuator triggers different physical mechanisms to reduce the sound levels at both sensors. Lastly, the case optimised to reduce the pressure fluctuations at the arc sensor appears to be the most successful reducing the overall sound radiation. Also, given the low p_{rms} contours across the entire figure 8.4d, we can say that this actuation yields the flow-field closest to the steady state. Bear in mind that this larger noise reduction comes at the price of more than doubling the computational effort, respect to the upstream or downstream sensor cases. Moreover, the sound reduction at the upstream sensor is much greater than the one obtained just optimising for

¹In any case, the flow actuation has certainly decreased the pressure fluctuations at the sensor location, which is all we wanted for this particular case.

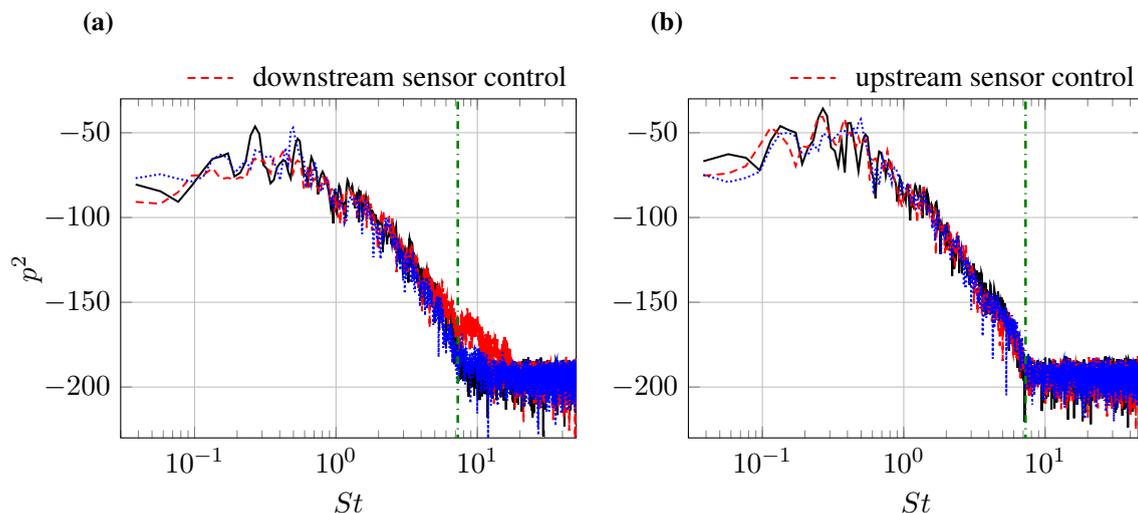


Figure 8.5: Power spectrum of pressure at the centre of the downstream **(a)** and upstream **(b)** monitor points. The black solid and blue dotted lines represent the non-actuated and arc cases. The red dashed lines show the results from the downstream case in **(a)** and upstream case in **(b)**. The green dash dotted vertical line indicates the sampling frequency of the flow sensors used for the adjoint-based optimisation.

this particular location. The reason behind these two different results resides in the ‘quality’ of the information (adjoint-field) with which the gradients are computed. The arc sensor provides sensitivity data from all directions, which represents much better the flow mechanisms which give birth to the sound generation than the information from just the upstream sensor. Hence, sensing the pressure fluctuations with the arc sensor shows the optimisation framework that the optimal way to achieve the desired sound reduction is driving the flow-field as close as possible to a steady state. Contrarily, the poor sensitivity data provided by the upstream sensor drives the solution to a local minimum. On the other hand, comparing the results from figures 8.4b and 8.4d, we observe how the downstream sensor case is more effective reducing the sound levels at this location. This should not come as a surprise since the arc sensor seeks an overall sound reduction rather than just focusing the control efforts on that specific spatial region.

To obtain an insight into the frequency content of the pressure fluctuations radiated by the cavity in each case, we have placed two monitor points at the centre of the downstream and upstream sensing regions. Furthermore, since this acoustic radiation is directly related to the shear layer impinging onto the cavity’s trailing edge, the analysis of these signals in the frequency domain can be used as an indication of the shear layer dynamics. The monitor points capture data with a sampling frequency 25 times higher than the flow sensors used for the adjoint-based optimisation, which permits obtaining the full frequency range. Figure 8.5 shows the frequency spectrum of various cases at the downstream (figure 8.5a) and upstream (figure 8.5b) monitor points. The solid black and blue dotted lines represent the pressure fluctuations for the non-actuated and arc cases, respectively. Especially, the red dashed lines show the results from the downstream case in figure 8.5a and upstream case in 8.5b. These two figures show how the sound radiation is dominated by the cavity’s Rossiter modes, which appear as distinct peaks in the spectrum. In

agreement with the p_{rms} contours from figure 8.4, figure 8.5 shows how the acoustic radiation associated with these resonant frequencies is considerably more energetic in the upstream direction than when they are advected downstream. This phenomenon was reflected previously in figure 8.4 as higher p_{rms} contours, whereas now in figure 8.5, it is manifested with higher peaks in figure 8.5b than in 8.5a. In addition, due to the above-mentioned interaction with the shear layer² (which has an amplifier or damping role over the acoustic radiation), the first and third Rossiter modes are mitigated in the pressure signal captured by the downstream monitor point. Contrarily, these modes are amplified by the shear layer as they propagate upstream. Bear in mind that, in comparison with the results shown in chapter 6, the characteristic frequencies of the Rossiter modes have experienced a modulation as a direct consequence of the Reynolds number increase³. In particular, this Reynolds number variation has displaced the second Rossiter mode's frequency from $St \approx 0.233$ (at $Re = 2000$) to $St \approx 0.268$ (at $Re = 5000$). Also in comparison with the results presented in chapter 6, the first Rossiter mode appears to gain in activity as the Reynolds number is raised. Despite this fact, the second mode still dominates the sound radiation, presenting the largest amplitude in both spectra for the non-actuated case shown in figure 8.5. For this reason, the second Rossiter mode is the main target of the flow actuation in every case studied herein. The reduction achieved in the value of the cost function for the downstream case is 79.6%. The diminution of energy in the lower resonant frequencies is evident in figure 8.5a, where the amplitude of the second Rossiter mode lessens from -46.28dB to -65.23dB . Moreover, the first and fourth modes also experience a mitigation in their amplitude, where only the third Rossiter mode appears to be enhanced. This event agrees with the p_{rms} contours from figure 8.4b, where the shorter flow patterns along the shear layer indicate that the actuation triggers higher Rossiter modes, which we now know that this corresponds to the third mode. Further, note that there is a considerable increase in the energy present in the high frequencies above the flow-field sampling frequency ($St \approx 7.272$). As mentioned earlier, these high-frequency fluctuations cannot be captured by the cost function, which acts as a low-pass filter. Regarding the upstream sensor case, the overall reduction in the cost function is of 20.4%, which is the lowest of the three optimisation cases studied in the present chapter. As illustrated by figure 8.5b, this flow actuation fails to mitigate the first and second Rossiter modes, where it only causes a slight reduction in their amplitude and characteristic frequencies. Additionally, this frequency modulation also occurs to the third mode, whose amplitude is raised a minimal amount. Linking this spectrum with the results shown earlier in figure 8.4c, they both show how the second Rossiter mode remains dominant in the cavity dynamics. Far from cancelling out this second mode, the flow actuation appears to marginally redirect the shear layer, altering the flow acoustic interaction which briefly diminishes the sound levels at the upstream sensor. Lastly, as shown in table 8.1, the arc sensor case achieves the greatest reduction in the cost function. Similarly to the downstream sensor case, this actuation causes significant reductions in the amplitude of the second Rossiter mode (from -46.28dB to -60.32dB and from -35.78dB to -50.50dB at both downstream and upstream sensors, respectively). As well, the amplitude of the third

²See chapter 6 for more details.

³This frequency modulation can be accounted for by the Rossiter's semi-empirical formula (6.9) adjusting the coefficients κ and γ .

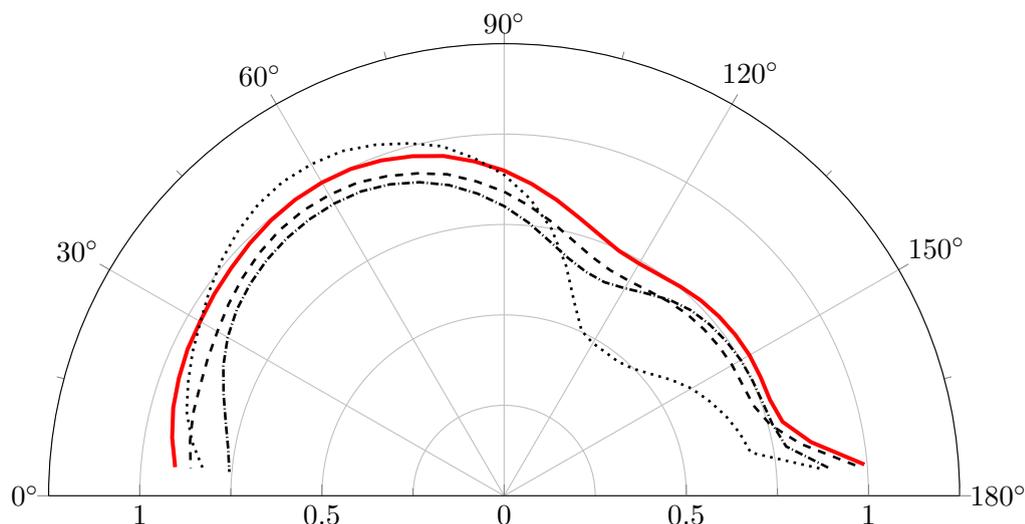


Figure 8.6: Overall sound pressure level plots, with 0° and 180° as the upstream and downstream directions. The red solid line shows the non-actuated case. Dotted and dashed lines represent the downstream and upstream sensor cases, respectively. The dash-dotted line indicates the result of the arc sensor case. The OSPL levels are normalised with the maximum value of the non-actuated case.

mode it is also lessened, where this time the actuation enhances the fourth Rossiter mode as the dominant mode. On the other hand, the first mode presents minimal reduction respect to the non-actuated case. These results agree with the p_{rms} contours from figure 8.4d, where the flow patterns along the shear layer are even smaller than in 8.4b, since the fourth Rossiter mode dominates for this case the sound radiation.

Figure 8.6 shows the overall sound pressure levels for all cases as defined in (6.10), computed along the arc sensor. This analysis sheds light into the cavity's sound directivity and how it changes depending on the actuation obtained with each sensor. In addition, figure 8.6 proves that several flow control strategies are possible, depending on the direction where we desire the sound to be reduced. The non-actuated case is used as the baseline case and is represented by a red solid line. Additionally, the magnitude of all the curves is normalised with the maximum value of the baseline case. In particular, this plot confirms the arc sensor setup (dash-dotted line) as the most successful reducing the overall cavity's sound radiation. This fact should not come as a surprise since the downstream (dotted line) and upstream (dashed line) setups had as their only objective to reduce the sound levels at their respective locations, rather than an overall sound reduction. In fact, the downstream case is the one which most efficiently mitigates the sound propagation in the downstream direction. Additionally, this analysis highlights the importance of the sensor location, not only for adjoint-based optimisation but for every study interested in cavity sound. Due to the sound amplifying character of the shear layer, the leading sound radiation direction is approximately at 60° , and it strongly depends on the Mach number (see chapter 6). Hence, at the current Mach number, a flow control strategy based on a sensor located downstream is prone to lead to higher sound levels radiated upstream, as it is the case here for our downstream sensor-based optimisation. On the other hand, that is not the case for

Case	μ_r	μ_i	Mode
non-actuated	-0.008	0.267	II
downstream	-0.007	0.400	III
upstream	+0.012	0.135	I
arc	-0.025	0.481	IV

Table 8.2: Growth rate μ_r and frequency (Strouhal number) μ_i of the dominant dynamic modes for each case. Last column indicates the Rossiter mode number associated to these leading dynamic modes.

the upstream sensor, where a reduction in the sound levels at the sensor location also yields a global sound reduction.

With the aim of seeking a more global description of the dominant flow dynamics for each flow actuation, we can apply DMD over the snapshot sequence generated for each case. As observed earlier in figure 8.5, the time horizon is not long enough to provide a fine detail on the lower resonant frequencies. This is generally an indication that carrying out DMD over such dataset might not yield converged results. Nonetheless, DMD results still provide a reasonable estimate of the system's dynamics, which can be used to unveil the physical mechanisms triggered by each flow actuation. Similarly to Schmid (2010), we determine the dominant or leading mode projecting the POD base over each dynamic mode, which ranks them by their the energy content. Hence, the dynamic mode referred to as dominant is simply the one with the highest energy. Using this criterion, the corresponding eigenvalues to the dominant dynamic mode on each case are summarised in table 8.2. In addition, this table also includes the Rossiter mode number associated to each of the leading dynamic modes. These modes' number can be readily estimated comparing the frequencies exhibited by the Rossiter modes in figure 8.5 with the eigenvalue's frequency. Bear in mind that the associated frequency of each Rossiter mode for each case is subject to a modulation from the flow actuation. Especially, this frequency modulation is directly related to the location where the flow actuation is applied. As discussed earlier and also in chapter 7, in the present setup this particular spatial region shows a high receptivity, but perhaps more importantly, it also presents high values of structural sensitivity. Hence, an actuation located in the core of flow instability or wave maker is prone to alter the system's eigenvalues. Evidence of such frequency modulation were already observed in figure 8.5, where the characteristic frequency of the fourth Rossiter mode suffered a slight decrease through the influence of the flow actuation in the arc case. This event agrees with the frequencies summarised in table 8.2. There, the frequencies from the second and third modes from the non-actuated ($\mu_i \approx 0.267$) and downstream cases ($\mu_i \approx 0.4$) are almost exact multiples from the frequency of the first mode from the upstream case ($\mu_i \approx 0.135$). Contrarily, the frequency from the dominant dynamic mode in the arc case, associated to the fourth Rossiter mode, is lower than expected when following the trend from the previous three cases⁴. Moreover, the streamwise and vertical momentum com-

⁴Without any frequency modulation from the actuation, one would expect the the fourth Rossiter mode to exhibit a characteristic frequency of twice the frequency of the second mode of the non-actuated case ($\mu_i \approx 2 \times 0.267$). Furthermore, the dynamic modes corresponding to the three first Rossiter modes in the arc case also present alterations in their characteristic frequency - $\mu_{i_1} \approx 0.153$, $\mu_{i_2} \approx 0.241$ and $\mu_{i_3} \approx 0.344$.

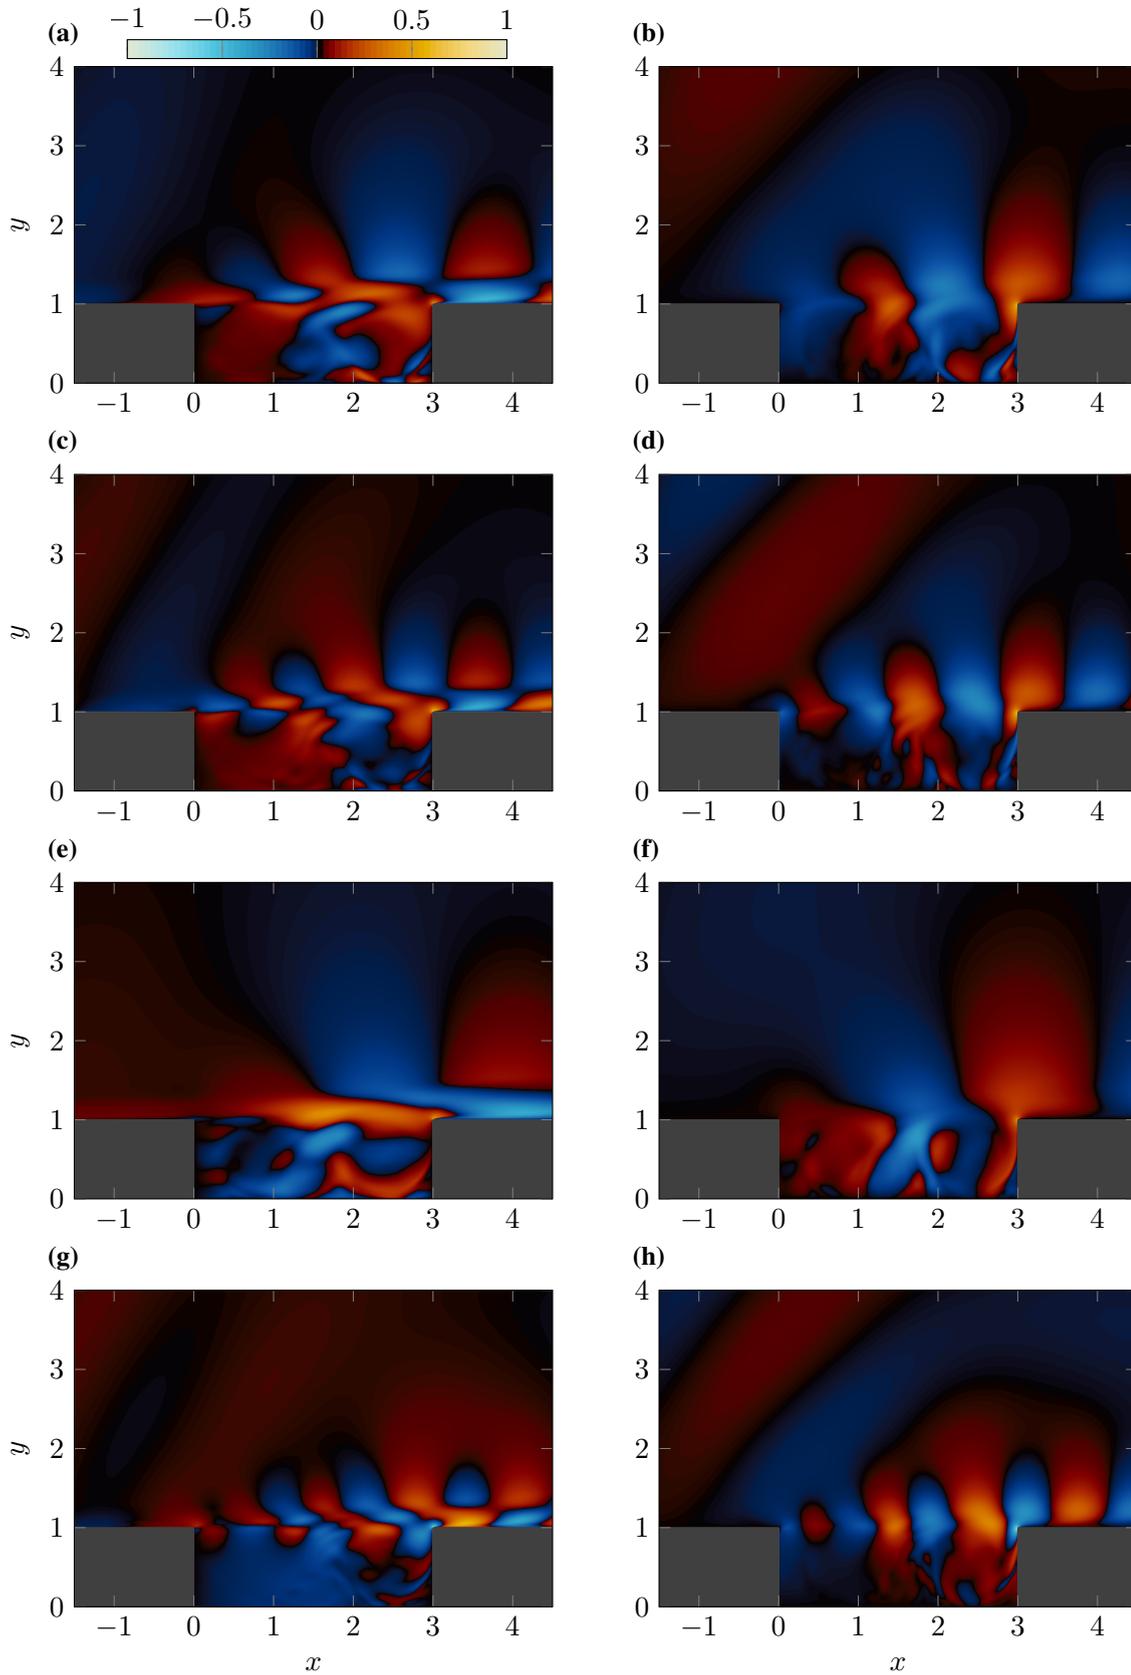


Figure 8.7: Streamwise (left) and vertical (right) momentum components of the leading dynamic modes for each case. From top to bottom are the non-actuated, downstream, upstream and arc cases. Contour levels are normalised with 0.007.

ponents of the dominant dynamic modes for each case are shown in figure 8.7. Equivalently to figure 8.4, the flow patterns present a shorter wavelength as the leading Rossiter mode's number is raised. Due to the laminar character of the inflow condition used, the sound radiation can also be estimated in these momentum components. In particular, these plots highlight how the sound generation and its directivity is strongly linked to the shear layer dynamics. As seen before in chapter 6, when the shear layer instabilities grow, they compress the flow around them forming a dipole. This dipole interacts with the upstream travelling acoustic wave reflected from the cavity's trailing edge, which can be either enhanced or mitigated, depending on the phase of the interaction⁵. As we move to higher Rossiter modes, the characteristic length of these instabilities decreases, permitting more simultaneous dipoles along the shear layer. Hence, by changing the dominant Rossiter mode in a cavity flow, we are able to change the way this flow-acoustic interaction occurs, which has an impact on the overall cavity sound radiation and its directivity. For example, at the downstream sensor location, the activity of the third Rossiter mode for the downstream case (figures 8.7c and 8.7d) is lower than any other mode show in figure 8.7. Also as a consequence of the shorter wavelengths in the shear layer (compared to the non-actuated case), the upstream travelling acoustic wave is further enhanced, yielding a greater sound radiation in this direction, as seen earlier in figures 8.4 and 8.6. Focusing now in the arc case, we see from figures 8.7g and 8.7h that the upstream sound radiation from the fourth Rossiter mode is much weaker principally compared with the dominant modes from the non-actuated case (figures 8.7a and 8.7b) and downstream case (figures 8.7c and 8.7d). The origin of this lower sound radiation is also exclusive of the flow dynamics governing the shear layer. With the fourth Rossiter mode as the dominant mode, the flow structures impinging onto the cavity's trailing edge have a shorter wavelength, which implies that the impingement is less energetic compared to the non-actuated case. This event radiates a weaker acoustic wave upstream, which despite the numerous dipoles placed along the shear layer, it does not get as amplified as occurred in the non-actuated or downstream sensor cases.

Optimal Flow Actuation

So far we have analysed and studied in detail the effect of each flow actuation over the system's dynamics and its associated sound generation. On the other hand, very little has been said about the flow actuation used. Given the high complexity of our controller, with above six million of degrees of freedom per case, a thorough description of the actuation is crucial to reach understanding on how to trigger the aforementioned physical mechanisms which reduce the cavity's sound radiation at the sensor location. Despite that the flow actuation used here is not entirely physical (in the sense that it cannot be reproduced in an experimental setup), it yet provides a valuable information on how to optimally perturb the flow-field, which cannot be obtained by other methods. Hence, the present section exposes these flow actuations as obtained from the adjoint-based optimisation, leaving their potential conversion to more real flow actuators as future work. Due to the limited success of the actuation in the upstream

⁵Note that this phase is strongly linked to the Mach number of the flow.

Case	#	μ_r	μ_i
downstream	D1	-0.436	0.152
	D2	-0.413	0.666
	D3	-0.559	0.560
	D4	-0.320	0.289
arc	A1	-0.315	0.105
	A2	-0.273	1.120
	A3	-0.265	0.728
	A4	-0.120	0.443

Table 8.3: Dominant actuation modes for the downstream and arc cases.

case, and also considering that a greater noise reduction at the upstream sensor was achieved by the arc case, we restrict our controller analysis to the downstream and arc cases. In order to reduce the complexity of the flow actuations and also seeking a better understanding of the flow forcing mechanisms, we break down the forcing $\vec{f}(\vec{x}, t)$ into spatial modes with associated growth rate and frequency by using DMD. Since the modes with the highest growth rate are the ones which govern the system's (controller's) dynamics in the long time (e.g., Brès and Colonius, 2008), we only show the fourth least stable modes for each case. Similarly to the criterion used above, these dynamic modes are ranked by their energy content by projecting them onto a POD basis. Incidentally, the four modes with the highest growth rate are also the ones with the highest energy in both downstream and arc cases. Table 8.3 summarises the growth rates and frequencies of the four dominant modes in each case. Also for these modes, figures 8.8 and 8.9 show the streamwise and vertical forcing components for the downstream and arc cases, respectively. From the characteristic frequencies gathered in table 8.3 relative to the downstream actuation, we can establish a direct connection between these frequencies and the more active Rossiter modes present in the non-actuated cavity flow. This occurs as a direct consequence of the strong flow convection between the actuation and sensor locations. By proximity in their oscillating frequencies, we readily observe how modes D1, D2, D3 and D4 are related to the first, fifth, fourth and second Rossiter modes. In particular, looking back to the spectrum shown in figure 8.5a, these four modes are the ones exhibiting the highest amplitudes, which immediately justifies why the flow actuation tries to mitigate them. Hence, as a result of the flow control strategy being successful, the third mode becomes dominant simply by elimination, rather than being enhanced by the flow actuation. Moreover, since the second Rossiter mode showed the largest amplitude in the non-actuated case, the actuation mode responsible for its control (D4) presents the least damped growth rate. In other words, the mode D4 is required to 'stay on' for longer to deal with this more energetic Rossiter mode. The reason behind all the controller modes exhibit a negative growth rate resides in the relatively short time-span of the optimisation. Initially, the amplitude of these modes is large, but it instantly commences a decay as soon as the flow actuation begins to mitigate its target. Hence, if this optimisation would be continued further in time, some of the actuation modes would slowly decay to Fourier modes (with zero growth rate and both constant amplitude and frequency). Regarding the arc case, the relation between the controller and Rossiter modes is not as obvious as in the downstream case. Perhaps,

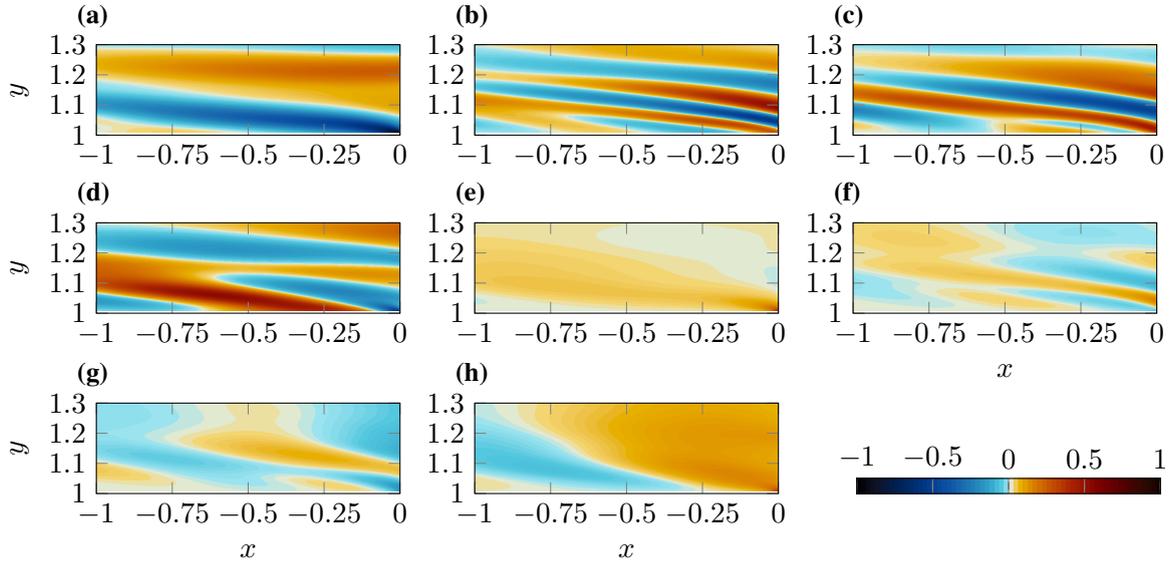


Figure 8.8: Dominant actuation modes for the downstream case. Ordered from mode D1 to D4, (a) to (d) and (e) to (h) show the streamwise and vertical forcing components. Contour levels are normalised with 0.3741.

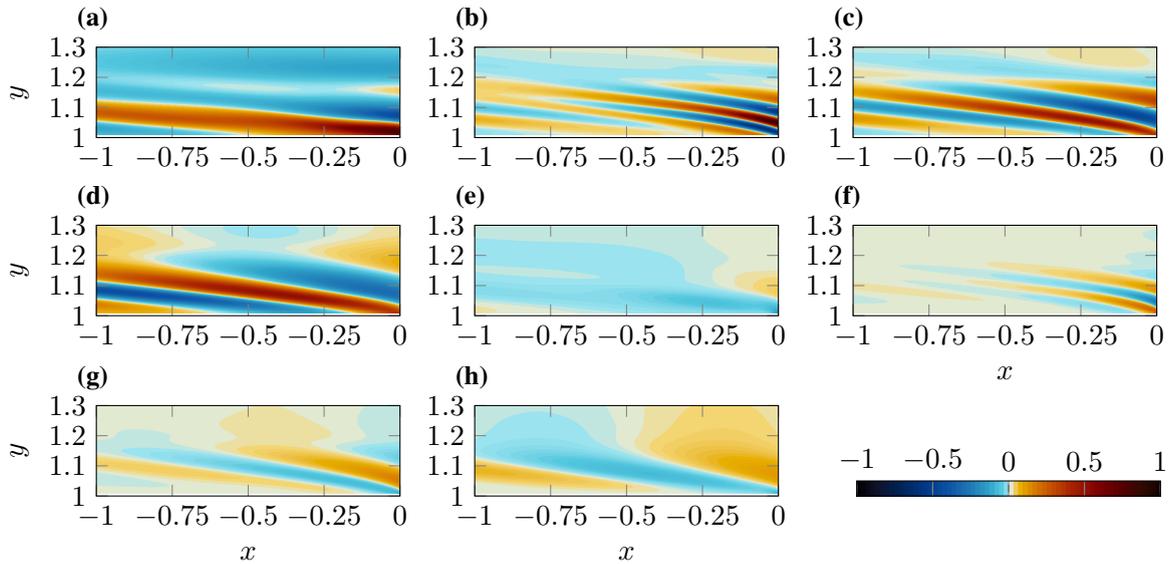


Figure 8.9: Dominant actuation modes for the arc case. Ordered from mode A1 to A4, (a) to (d) and (e) to (h) show the streamwise and vertical forcing components. Contour levels are normalised with 0.3741.

the much bigger sensing region in the arc case alongside the lack of strong convection between the sensor and the actuation, give rise to more complex events in the adjoint-field, which then lead to this final control strategy. The only modes where we can distinguish a relation with the cavity's leading resonant frequencies are modes A4, A1 and A3. The frequencies from A4 and A1 are slightly below the associated frequencies to the fourth and first Rossiter modes, whereas the oscillation presented by A3 is associated with the fifth Rossiter mode. Similarly to the downstream case, the less energetic actuating mode (A4) shows the highest growth rate,

this time to fuel the activity of the fourth Rossiter mode. Examining the mode shapes from figures 8.8 and 8.9, we quickly notice that the forcing is principally exerted in the streamwise direction, where the forcing in the vertical direction is approximately an order of magnitude lower. Further, the most important feature from these plots is that they reveal the ideal shape of how the forcing should be applied. In the present case, all forcing modes exhibit diagonal streaks, which is the result of the receptivity patterns from the Rossiter modes in this particular area (see chapter 7). In addition, these plots also show how the streamwise and vertical forcing components are generally in opposite phase. This means that an increment in the streamwise momentum has an associated decrease in vertical momentum at the same spatial location, and vice versa. Lastly, the high activity of these forcing modes towards the cavity's leading edge suggests that further extending the forcing around this corner might increase the efficiency of the actuation. This would also be supported by high amplitude at the vicinity of the cavity's leading edge exhibited in the wave maker and receptivity plots, as shown earlier in chapter 7.

8.4 Summary and Conclusions

Adjoint-based optimal flow control was successfully applied on a two-dimensional open cavity flow for the first time. The target of the actuation was to minimise the pressure fluctuations at the sensor location. To investigate the influence of the sensor on the resulting flow actuation, three different sensing scenarios were carried out, where amplitude reductions of 79.6%, 20.4% and 82.0% were achieved respect to the non-actuated case. A statistical study of the adjoint-field prior to the optimisation, alongside with the findings from previous chapters, indicated the ideal location for the actuator. To exploit the advantages of the adjoint method, a large number of control parameters was used. Decreasing the flow-field sampling frequency acted as an implicit low-pass filter in the sensor. For all three cases, p_{rms} contours revealed a decrease of pressure fluctuations in the area surrounding the sensors, by either changing the directivity of the noise radiation or reducing the overall sound emission. Furthermore, Fourier analysis of temporal signals in several monitor points, accompanied by dynamic mode decomposition, shed light into the variations in the dominant physical mechanisms which govern the flow-field on each scenario. As well, DMD analysis was carried out over the control subspaces to show the leading frequencies and forcing shapes. This provides a valuable insight into the ideal forcing patterns to achieve noise reduction in a cavity flow, which can be used in future investigations to model more realistic flow actuation strategies. As future work, it would be interesting to isolate single frequencies associated to Rossiter modes in the cost function (Schulze et al., 2011), forcing the controller to reveal the optimal actuation to mitigate each particular mode.

Chapter 9

Control of a 3D Backward-Facing Step Flow

This chapter presents a new application of adjoint-based optimisation. The investigation is oriented to cabin noise reduction, where we optimise an upstream flow actuation to reduce the flow impinging scales which present the largest impact on cabin noise. In addition, the complexity present in the flow studied pushes the adjoint-based approach at its limits. These intrinsic limitations of the method are discussed, where we also suggest and demonstrate a successful alternative which permits the use of adjoint-based optimisation methods in separated flows. If the reader is not familiar with the use of the adjoint method in chaotic systems, it is strongly recommended to read appendix [E](#) before continuing any further.

9.1 Introduction

When a vehicle is moving at relatively high speed, fluid-structure-acoustic interaction (FSAI) takes place everywhere around the vehicle's body. The pressure fluctuations present in the flow-field interact with the structure exciting some of its vibration modes that will radiate sound into the passenger's cabin, leading to a phenomenon known as 'aero-vibro-acoustics'. Hence, the structure can be seen as a filter or transfer function that only allows certain fluctuations to propagate through it. Depending on the material properties, the shape and how each structural panel is attached to the body, the vibrating properties will change, which makes some structural parts more prone to radiate noise than others. Recent advances reducing noise generation from the mechanical components of the vehicles, such as engine noise in cars and aeroplanes, have made aero-acoustic noise become the main contributor to cabin noise.

In aero-vibro-acoustics problems, the frequencies of the structural vibrations lie within the audible or acoustic range and the deflections are usually smaller than the smallest eddy present in the flow-field. Due to these low amplitude vibrations, it is very common to solve the fluid

and structural problems separately as the structural deformation barely influences the flow motion. [Hwang and Hambric \(2001\)](#) stated that given this situation, for large structures with fixed boundaries, high wavenumber pressure fluctuations on the structure have a small contribution to its vibration. Hence, it is essential to capture accurately the low wavenumber range of the spectrum (i.e. acoustic waves) to obtain a realistic excitation of the structure. In addition, according to [Bremner and Wilby \(2002\)](#), the acoustic content of the pressure field has a considerable influence on the overall noise transmission, despite the energy of these acoustic fluctuations is considerably lower than the convective ones. [Bremner \(2012\)](#) studied several “aero-acoustic loading mechanisms” that could have a direct repercussion on the wind noise transmission into a car cabin. In his study, wavenumber-frequency analysis showed that convection, propagating acoustics and diffuse acoustics contributed to wind noise transmission. [Graham \(1997\)](#) highlighted three factors that will determine if the coupling of the flow spectrum and the vibrating modes of the structure will be strong or not; the acoustic efficiency, resonance and excitation. The acoustic efficiency and the resonance are determined by the radiation efficiency and mode shapes of the structure, respectively, and the excitation exclusively depends on the near wall characteristics of the flow. [Van Herpe et al. \(2011\)](#) studied the wall pressure fluctuations of a flow past a car’s wing mirror with a wavenumber-frequency spectrum. The flow excitation was used in a second study to estimate the radiated noise by the side window of a car into the passenger’s cabin ([Van Herpe and Duarte, 2012](#)).

The aim of the present investigation is to apply adjoint-based optimal flow control on a three-dimensional backward-facing step, manipulating the upstream flow to reduce the flow excitation over the given sensing region. By targeting only the flow excitation, the study is kept independent from the structural properties. Previous studies in adjoint-based optimal flow control oriented to aero-acoustic noise reduction can be found, for example, in [Cerviño et al. \(2002\)](#) and [Kim et al. \(2014\)](#). Usually, these studies are carried out on flows over simple geometries with a strong flow convection, which facilitate the application of the adjoint method. Note that this is not the case for the present investigation, as the flow separation present in the BFS adds an extra level of complexity to the flow optimisation problem.

9.2 Computational Setup

We investigate the flow over a three-dimensional backward-facing step with a laminar inflow (figure 9.1). Of particular interest from this geometry are the constant-height flow separation and the flow impingement in the reattachment region ([Le et al., 1997](#)). This causes an energetic complex flow dominated by the chaotic behaviour of turbulence. The target or sensing region is located in the vicinity of the lower wall after the step, where the flow impingement occurs. This permits to penalise the flow unsteadiness that would interact with, for example, a vehicle’s window placed at the same location. The flow would excite some of the vibrating modes of the structure which would radiate noise into the passenger’s cabin. As mentioned in section 9.1, the main contributors to the structural vibration are the low wavenumber fluctuations. Hence, a

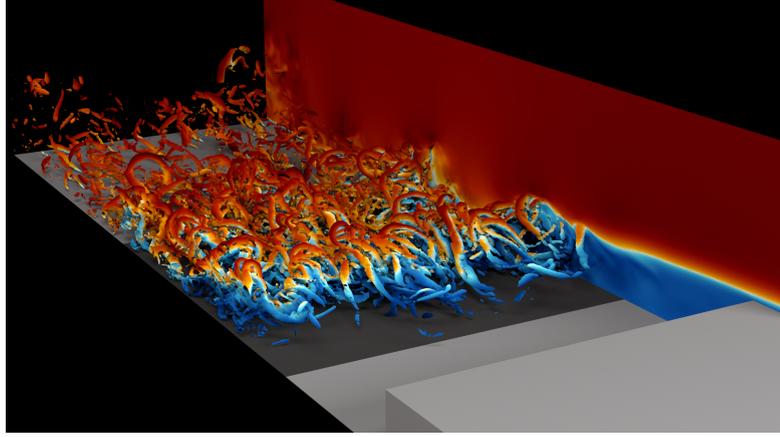


Figure 9.1: Illustration of the flow topology over a backward-facing step. The turbulent structures (represented as isosurfaces of Q-criterion coloured with streamwise velocity) impinge on the target region (dark grey surface) causing the pressure fluctuations that would generate flow-induced vibrations on a structure at that same location.

convenient cost function (based on [Schulze et al. \(2011\)](#)) which targets the pressure fluctuations that match the leading bending wavelengths of the structure is

$$\mathcal{J}(Q, \phi) = \frac{1}{T} \int_0^T \int_{k_L}^{k_U} \frac{1}{2} [\mathcal{F}(p'(\vec{x}, t) W(\vec{x}))]^2 d\vec{k} dt, \quad (9.1)$$

where k_L and k_U are the lower and upper cut-off wavenumbers respectively. Also, \mathcal{F} indicates a Fourier transform and W is a spatial function defining the sensing region.

The streamwise, vertical and spanwise directions are non-dimensionalised by the step height. With the coordinate origin at the lower corner of the step, the domain ranges from -20 to 80 in x , 0 to 20 in y and from 0 to 6.4 in the spanwise direction z ([Kopera, 2011](#); [von Terzi, 2004](#)). The Reynolds number based on the step height is $Re_\infty = 3000$ where the non-dimensional free-stream velocity is $U_\infty = 1$. The Prandtl and Mach numbers are respectively $Pr_\infty = 0.72$ and $M_\infty = 0.5$. In particular, this three-dimensional case is a continuation from the *medium* case from appendix E, where the boundary layer thickness at the separation point is $\delta \approx 0.382$. The time-step for the simulations is $\Delta t = 0.0008$. The computational domain is divided into three blocks around the step corner. The first block is discretised with 100 points in streamwise and wall-normal directions and blocks two and three use 600 by 100 points for the streamwise and wall-normal directions respectively. The flow information is interchanged through the block interfaces using halo points. The grid resolution is clustered towards the wall boundaries to make sure that all structures in the flow-field are resolved, but the resolution is coarsened towards the inflow, outflow and free-stream to avoid reflections in both forward and adjoint simulations. Note that the resolution is kept within the DNS range only up to 20 step heights downstream from the step. The spanwise dimension is discretised with 258 uniformly spaced collocation points, or 128 Fourier modes, with periodic boundary conditions in this direction. The wall boundaries are modelled as isothermal no-slip walls and the inflow, free-stream and outflow are

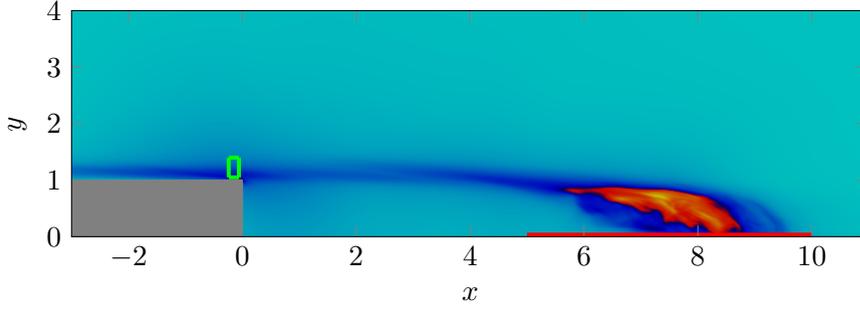


Figure 9.2: Root mean square of the magnitude of the adjoint momentum \vec{m}^* . Red and blue contours indicate high and low values, respectively. The green rectangle indicates the actuating sub-volume. The sensing region is represented with a red line at the flow impingement area.

of characteristic type, where the outflow zonal CBC is applied over the last 50 points. To carry out adjoint simulations the flow-field is stored every 250 time-steps and it is linearly interpolated for every sub-step of the Runge-Kutta scheme. A preliminary study showed that a finer flow-field sampling rate in time leads to the same results.

The sensing region ($W_s(\vec{x})$) extends from 5 to 10 step heights downstream of the step in the streamwise direction, covering the entire domain across the spanwise direction. For stability purposes, the sensing region ranges from the second to the eighth grid points in the wall normal direction, where a Gaussian function is used to smooth out the discontinuities in the adjoint simulations. To exploit the advantages of the adjoint method, the actuation is chosen to be a time-dependent body forcing sub-domain $\vec{f}(\vec{x}, t)$, where every grid point within the sub-domain has three control parameters (the streamwise, vertical and spanwise components of the forcing vector \vec{f}). Including both cost function and forcing terms in the Lagrangian functional (2.3), which are then substituted into (3.11), leads to their respective adjoint terms which can be written as

$$-\frac{\partial \rho^*}{\partial t} = RHS + \underbrace{\tilde{p} W_s(\vec{x})}_{\text{Cost Function}} \quad (9.2)$$

$$-\frac{\partial p^*}{\partial t} = RHS + \underbrace{(\vec{m}^* \cdot \vec{f}) W_a(\vec{x})}_{\text{Forcing Terms}}, \quad (9.3)$$

where \tilde{p} is the filtered pressure field. $W_a(x)$ is the spatial function representing the actuation sub-volume, which is defined as

$$W_a(\vec{x}) = \begin{cases} 1 & \text{if } \vec{x} \in [(-0.24, -0.06), (1.05, 1.4), (0, 6.4)] \\ 0 & \text{if } \vec{x} \notin [(-0.24, -0.06), (1.05, 1.4), (0, 6.4)] \end{cases}, \quad (9.4)$$

and encloses a total of 222,138 grid points. This actuating region was chosen as a result of a similar preliminary statistical study on the adjoint field, as shown earlier in chapter 8 (see figure

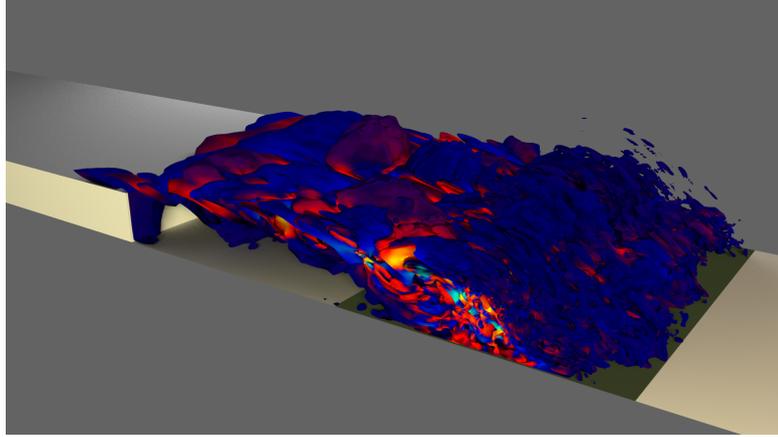


Figure 9.3: Illustration of the adjoint topology over a backward-facing step with the sensing region located in the flow impingement area. The adjoint structures (represented as isosurfaces of ρ^* coloured with p^*) travel upstream from the sensing region (dark green surface) towards the flow actuation.

9.2). From (2.11) the gradients for each of the control parameters follows as

$$\frac{D\mathcal{J}}{D\phi_x} = \rho m_u^* \quad , \quad \frac{D\mathcal{J}}{D\phi_y} = \rho m_v^* \quad \text{and} \quad \frac{D\mathcal{J}}{D\phi_z} = \rho m_w^*. \quad (9.5)$$

Due to the chaotic (turbulent) nature of the flow studied in the present chapter, the initial condition (terminal condition in the adjoint simulations) is forgotten by the flow past a given time. This means that the sensitivities of the system will diverge as the optimisation horizon increases. On the other hand, the horizon should be sufficiently long to allow communication between the sensor and the actuation (i.e. the actuation scales must have enough time to reach the sensing location in the forward simulations, and vice versa in the adjoint simulations - see figure 9.3). For further details on the limitations of adjoint-based optimisation in separated flows the reader is referred to subsection 2.1.1, section 9.4 and appendix E. For this particular case, we choose the optimisation horizon to be 24 time units, which according to figure 9.4a, is just over three shedding cycles. On a preliminary study, extending further the optimisation horizon resulted in an erratic behaviour of the flow actuation due to the low accuracy of the flow sensitivities. After each optimisation horizon has converged, the optimisation is advanced in time for T_a time units, which for the present case is 12 time units. In a similar manner as for the flow variables in the adjoint simulations, the control parameters are updated every 250 time iterations and they are linearly interpolated in between every capture. Note that this also removes implicitly the high frequencies from the controller. This setup results in approximately 80.6 million control parameters per time horizon. To better illustrate the effect of the exponentially growing adjoint field over the final actuation, figure 9.4b shows in advance the time signal of the resulting controller at an arbitrary location within the actuating sub-domain. This diverging adjoint field with backwards time marching leads to higher magnitudes in the control parameters the closer they are to the start of the optimisation horizon. The reason behind this behaviour resides in the respective gradient, which as indicated in (9.5), is proportional to the adjoint quantity \vec{m}^* . Note

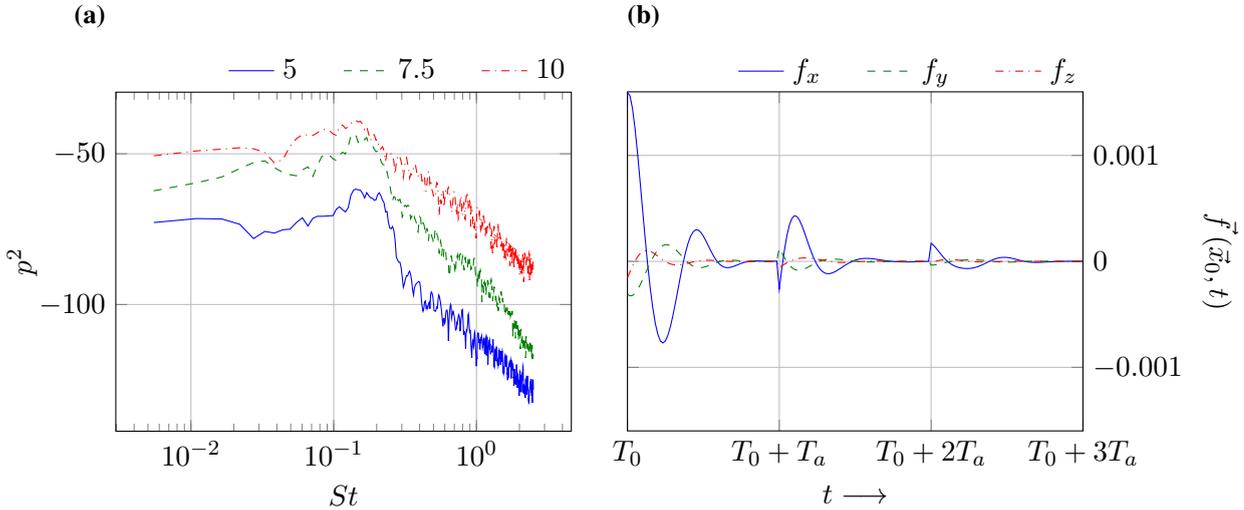


Figure 9.4: **(a)** Fourier transformed pressure signal at three different downstream locations of the step for the uncontrolled case. **(b)** Streamwise, vertical and spanwise components of the body-forcing vector $\vec{f}(\vec{x}_0, t)$ at an arbitrary location \vec{x}_0 within the actuation sub-domain. Note the pronounced increase in magnitude of the actuation parameters caused by the divergence of the adjoint field with backwards time marching.

that the divergence rate of the adjoint field is an indication on how ‘fast’ the system forgets the initial or terminal condition. As an alternative to circumvent this highly damped control signal, an intuitive method would consist in reducing the parameter T_a down to the updating period of the controller (250 time steps in the present case). Unfortunately, this approach would require the computation of a considerable larger amount of horizons to cover an equivalent time span, which makes this alternative unaffordable.

9.3 Optimal Flow Control

The flow optimisation has been run for three complete receding horizons, which is over 5 step shedding cycles. The cut-off wavenumbers were set to be the 10th mode in both streamwise and spanwise directions of the sensing region ($k_x = [-2, 2]$ and $k_z = [-1.5625, 1.5625]$). For this period of time, the flow actuation which resulted from the adjoint framework achieved a reduction in the value of the cost function of 7.05% with respect to the non-actuated simulation. Figure 9.5 shows graphical evidence of this cost reduction, reflected in the two-dimensional wavenumber spectrum of the pressure fluctuations gathered at the sensing region for the non-actuated (figures 9.5a and 9.5b) and actuated (figures 9.5c and 9.5d) scenarios. In addition, both figures 9.5c and 9.5a also show a sufficient grid resolution in a critical region such as the flow impingement area, as the energy decays considerably towards the highest wavenumbers. Moreover, the range of the spectrum accounted for in the cost function (9.1), delimited with a green dashed line in figures 9.5c and 9.5a, is shown in greater detail in figures 9.5b and 9.5d. From these zoomed-in pictures, we can readily observe how highest energy area is located along the $kz = 0$ axis, near the lowest wave numbers in the streamwise direction. This energy footprint is

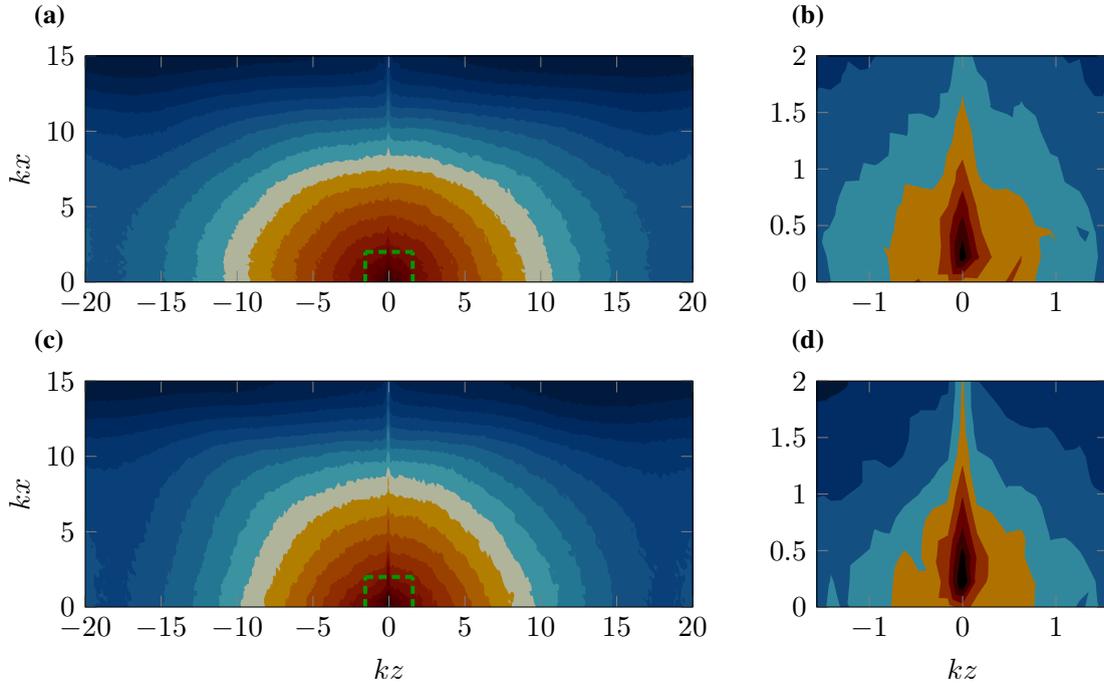


Figure 9.5: Wavenumber spectrum of the pressure fluctuations gathered at the sensing region for the non-actuated **(a)** and **(b)**, and actuated scenarios **(c)** and **(d)**. Contour levels range from -200 to -40 dB in **(a)** and **(c)**, and from -100 to -40 dB in **(b)** and **(d)**. In both cases the total number of contours is 22.

associated with the coherent structures in the spanwise direction, which are intrinsic from these type of flows. Especially, the incoming laminar boundary layer favours the appearance of structured roll-ups of the separated boundary layer in the spanwise direction, product of the growth of Kelvin-Helmholtz instabilities. In fact, the final flow actuation appears to enhance the energy present in these structures respect to the non-actuated case. On the other hand, the cost reduction is achieved by notably mitigating the energy content in the areas of the spectrum with non-zero kz . Note that this reduction in the cost function occurs despite the fact that the controller is effectively switched off most of the time (see figure 9.4b). Perhaps this highly damped in time behaviour of the controller is the reason why the actuation fails to reduce the energy content in the dominant spanwise coherent structures. In any case, the short energy reduction achieved in the low wavenumber regime cost by equation (9.1) has a much larger impact over the entire spectrum, where the overall energy reduction¹ is of 19%. This is a direct consequence of the flow actuation acting as a turbulence generator, which leads to a less structured and less energetic impingement of the shear layer on the lower wall. This phenomenon is also observed in figure 9.6, where the non-actuated DNS shows higher density spots and larger structures impinging on the wall compared to the actuated case. Also note the strong presence in both figures 9.6a and 9.6b of the spanwise coherent footprint of the shear layer roll-up. Despite the flow actuation, the streamwise location of these particular structures matches the ones from the non-actuated case. On the other hand, this is not the case for the flow impingement, which appears to have

¹Calculated extending the limits of the cost function (9.1) to the full wavenumber spectrum.

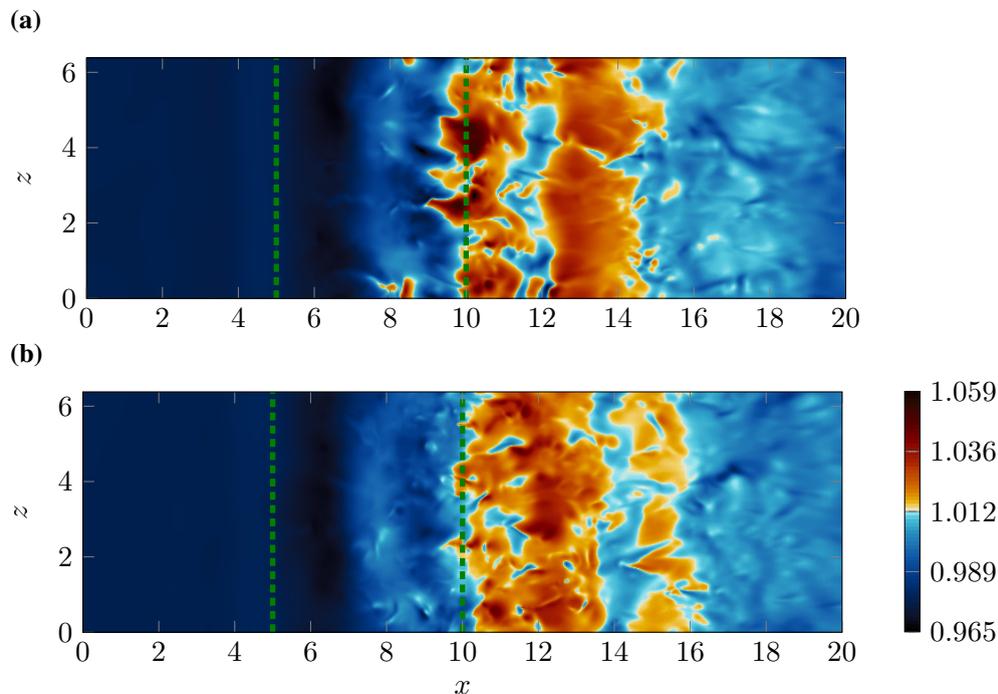


Figure 9.6: Instantaneous snapshot of the density-field in the lower wall of the step for the non-actuated **(a)** and actuated **(b)** scenarios. Both density-fields were captured at $t = T_0 + 3Ta$. Sensing region is delimited by the two vertical green dashed lines.

been slightly displaced further downstream by the controller. Note that this behaviour could be associated with the actuation acting as a flow trip or turbulent generator, which effectively leads to larger reattachment lengths (Adams and Johnston, 1988). Unfortunately, the overall optimisation period is not long enough to yield reliable statistics, which could help us to determine whether this different reattachment location is an instantaneous event or it prevails in time.

Moreover, with the time history of the pressure fluctuations at the sensing region, we can also compute the wavenumber-frequency spectrum. This analysis reveals the directivity of flow structures sorted by their frequency and wavenumber content (see chapter 2), which can shed light on what the actuation is doing to the flow. In addition, for subsonic Mach numbers, such as the one used herein, this flow decomposition is able to separate the energy content of the convective and acoustic events. Figure 9.7 shows the WFS at $St = 2.5$ for the non-actuated (figure 9.7a) and actuated (figure 9.7b) simulations, where the acoustic propagating speed in the $x - z$ plane is represented by the black dotted circle. Lower St planes did not show significant differences. Note that this St value is considerably higher than the shear layer roll-up frequency (see figure 9.4a), and corresponds to the upper limit of the frequency range captured by the sensor. The absence of a strong energy footprint in the streamwise propagating direction, suggests that the structures represented in this plane are associated with the spanwise oriented flow impingement, which is induced by the streamwise coherent rollers present in this type of flows (von Terzi, 2004). These structured rollers appear as the laminar shear layer collapses in the spanwise direction, and they present a counter-rotating behaviour respect each other. In particular, from these plots, we observe a change in the directivity of these structures associated with

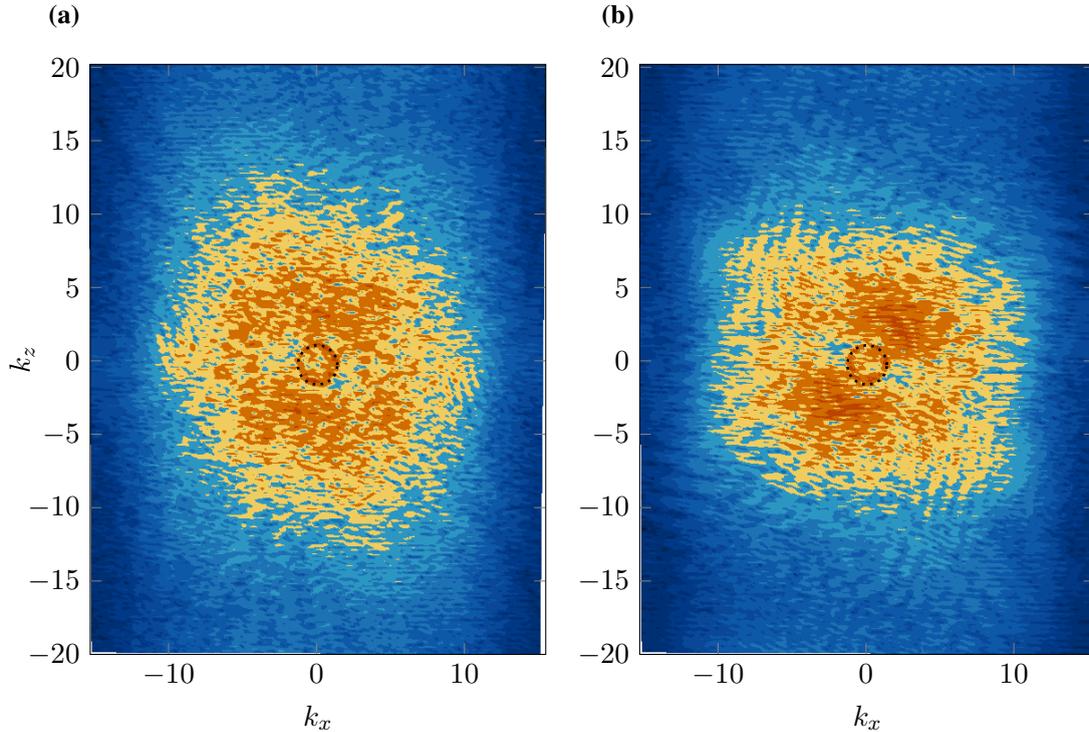


Figure 9.7: Wave number frequency spectrum at $St = 2.5$. Figures (a) and (b) correspond to the non-actuated and actuated cases, respectively. Contour levels range from -122 to 102. The acoustic circle (black dotted circle) represents the propagating speed of sound in the $x - z$ plane.

these streamwise rollers, caused by the flow actuation. The energy is now clustered forming a dipole shape which sits diagonally across the acoustic circle. Additionally, figure 9.7b shows a faster energy decay for the higher wavenumbers, especially in the spanwise direction. This event was also observed earlier in figure 9.5. The acoustic energy content (both propagating and diffuse acoustic energy) of these plots does not show significant changes when comparing the non-actuated and actuated simulations. Due to the slow convergence of this particular flow decomposition, perhaps further continuing the optimisation over more horizons would yield clearer information. In this present investigation, we are unable to execute such task due to our limited computational resources.

Despite that the flow actuation presents a highly damped behaviour, we can still extract valuable information on the forcing mechanisms by observing the initial stages of the actuation, where the control's amplitude has not yet decayed to zero. Figure 9.8 shows the instantaneous forcing contours at $t = T_0$, corresponding to the streamwise, vertical and spanwise forcing components. In particular, these three components display a spatial forcing oscillation in the spanwise direction, which appears to follow the second spanwise shear mode². More precisely, this fluctuation can be better observed by tracking the zero forcing contour along the spanwise direction, where we observe a strong resemblance in the forcing patterns exhibited by the streamwise (figure 9.8a) and vertical (figure 9.8b) actuating components. Similarly to the optimal control results

²Assuming the first shear layer mode in span to have a wavelength equal to the width of the domain.

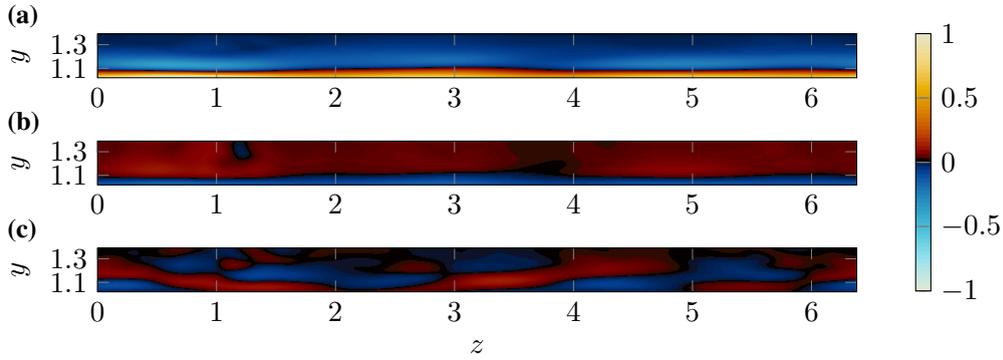


Figure 9.8: Instantaneous contours of the flow actuation at $t = T_0$, showing the streamwise (a), vertical (b) and spanwise (c) forcing components. Frontal view in the upstream direction. Contour levels are normalised by 0.0027.

presented in chapter 8, these two forcing components are opposite in phase, where a positive forcing in the streamwise direction has an associated negative forcing in the vertical direction (and vice versa). Note that this phenomenon is also reflected in figure 9.4b. In addition, this figure shows how all three forcing components oscillate at the same dominant frequency of approximately $St \approx 0.192$, which is within the vicinity of the shedding frequency of the step (see figure 9.4a). As well, since a preliminary study based on several other arbitrary locations showed the same dominant frequency and phase, we can also estimate the phase of the vertical and spanwise forcing respect to the streamwise component as $\psi_y \approx \pi 14/13$ and $\psi_z \approx \pi 7/13$.

9.4 Bypassing Adjoint-based Limitations: Gradient Normalisation

In previous sections from this chapter, we observed how standard adjoint-based optimisation methods fail to compute the correct flow sensitivities in chaotic systems. This phenomenon leads to an exponential divergence of the adjoint-field with backwards time marching, which yielded a very inefficient flow actuation. In addition, as shown with a practical example in appendix E, in separated flows this divergence of the adjoint field is further enhanced by the separated flow area. This flow recirculation region traps the errors arising from computing the sensitivities of the small chaotic scales of the system and leads to a cumulative error growth. On the other hand, this cumulative error growth would not occur in open chaotic (turbulent) flows with a strong convection, such as a turbulent boundary layer flow (Bewley et al., 2001) or a turbulent jet (Kim et al., 2014). In these scenarios, there is also an added error to the flow sensitivities emerging from the chaotic flow structures, but these errors get convected away by the flow, which prevents the aforementioned divergence of the adjoint field. Hence, when carrying out adjoint-based optimal control in this type of flows, the error present in the flow sensitivities will just depend on the distance between the sensor and the actuation, and it will remain approximately constant across time. Contrarily, in separated flows where the sensing region or flow actuation interact with the flow recirculation region, the error present in the sensitivities builds also as a function of time. For this particular reason, using an increasingly longer optimisation

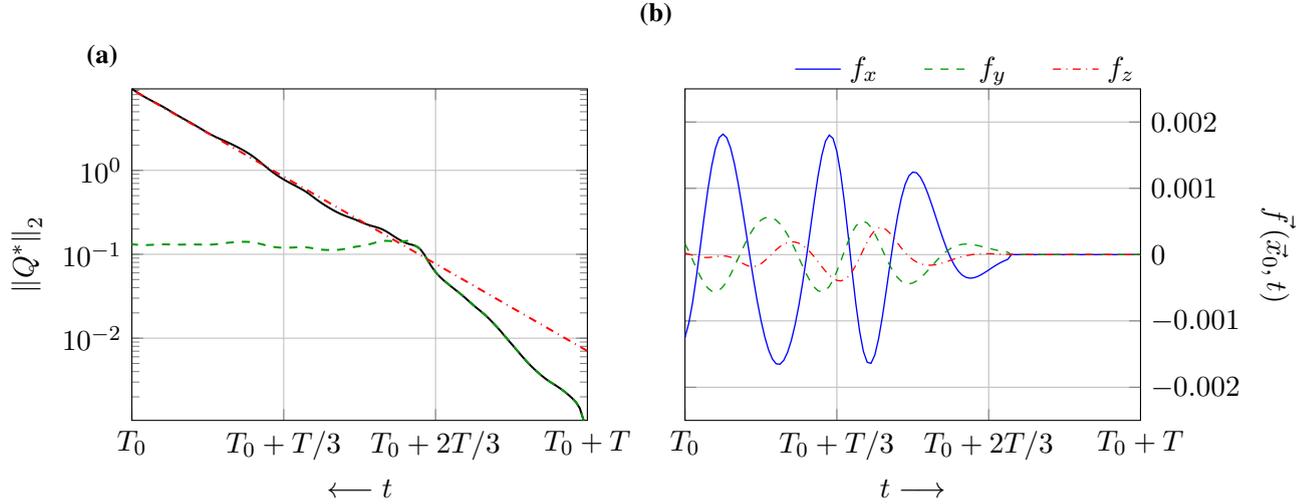


Figure 9.9: **(a)** Illustration of the sensitivity re-scaling method based on the norm of the adjoint field. Original and re-scaled norms are represented as black and green dashed lines, respectively. The red dash-dotted line shows the logarithmic line fit of the saturated norm of the adjoint field. **(b)** Streamwise, vertical and spanwise components of the preliminary body-forcing vector $\vec{f}(\vec{x}_0, t)$ when using the gradient normalisation method. Time signals correspond to an arbitrary location \vec{x}_0 within the actuation sub-domain.

horizon will eventually result in a completely degraded adjoint field, which will yield gradients that cause an erratic evolution of the cost function, rather than drive it to a minimum. Furthermore, this time-dependent error growth is the reason why the adjoint framework produces the inefficient control shown above in figure 9.4b. As mentioned earlier, the gradients are computed directly from the adjoint quantities, leading to a highly damped in time flow actuation. Thus, the need for a method that delivers a flow actuation which stays active throughout the optimisation arises. To this end, two intuitive potential alternatives would consist in restricting our controller to be steady or time periodic. On the other hand, the reader should bear in mind that the gradient computation for this type of flow actuations would account for the adjoint quantities throughout the optimisation horizon. Hence, since the adjoint field is diverging with backwards time marching, the result of such time integration would be strongly biased by the adjoint field close to T_0 . In other words, this optimisation would lead to a flow actuation which stays active throughout the simulation, but it is essentially optimised for an instantaneous flow snapshot, rather than the entire horizon. Not a good idea. Also note that, paradoxically, the adjoint field which dominates the result of the gradient integration is the one with the largest accumulated error. Consequently, in order to calculate a steady or periodic actuation which equally accounts for every instant of the optimisation horizon, we would require of a weighting function to counteract the adjoint-field divergence, product of the cumulative error growth. Note that this strategy could even be applied to the gradients of the time-dependent actuation used above, preventing the appearance of the damped character described earlier. Therefore, in the present investigation, we have opted to base such weighting function on the norm of the time-dependent adjoint field. As shown in

figure 9.9a, the raw norm of the adjoint field (black solid line) grows exponentially with backwards time marching from $T_0 + T_a$ to just after $T_0 + 2T_a/3$, where the adjoint field appears to reach some sort of saturation which reduces its growth rate. From this point backwards (to T_0), we carry out a line fit of the logarithmic-transformed norm (red dash-dotted line), obtaining the approximated growth rate associated with the cumulative error growth. Bear in mind that this growth rate is equivalent to the leading Lyapunov exponent of the system (Wolf et al., 1985). Finally, using the point after saturation as the reference, alongside the logarithmic fit, we can re-scale the adjoint field and remove the magnitude growth (green dashed line). Note that this process just normalises the adjoint quantities respect to the growth rate of the norm of the adjoint field, where the aforementioned errors associated with the computation of the sensitivities on the chaotic scales are still present. Figure 9.9b shows the preliminary actuation results from applying this normalisation to the control strategy described in previous sections, at an arbitrary location within the forcing sub-volume. Unfortunately, due to the limited available computational resources, we were just able to carry out 9 function-gradient evaluations. In any case, despite not having a converged solution, comparing the control signals exhibited in figures 9.4b and 9.9b we can already observe the benefits of this adjoint field normalisation, where the flow actuation stays active for much longer. In addition, these preliminary results produce a reduction in the value of cost function of 2.2%, whereas the final results from section 9.3 (evaluated just on the first horizon) reduced the value of the cost function a 5%. Also, at this point, it is fair to say that the results from the previous section needed more than twice the number of function-gradient evaluations to produce such result, in comparison to these preliminary results. Hence, for this reason, this gradient normalisation is a path worth exploring in the future, while a better method to compute the flow sensitivities in turbulent flows is not available (such as the method suggested by Wang, 2013; Wang et al., 2014). Moreover, other alternatives from past investigations are later discussed in appendix E. As well, the reader is encouraged to go through Blonigan (2016) for a recent review into these alternatives.

9.5 Summary and Conclusions

The newly developed framework oriented to adjoint-based optimal flow control was applied to a three-dimensional compressible backward-facing step flow. The geometry used and the presence of the flow separation area resulted in an added complexity on the application of the method. The target of the flow actuation was to minimise the low wavenumber pressure fluctuations at the sensor location, which was placed in the near wall region where the flow reattachment occurs. A total cost reduction of 7.05% was obtained with respect to the non-actuated case. The flow control framework produced an inefficient time-dependent control, where the actuation was effectively switched off most of the optimisation period. It was observed that this was due to the adjoint-field divergence with backwards time marching. In order to circumvent some of the intrinsic limitations of the method which cause this damped flow actuation, a gradient normalisation based on the L2 norm of the adjoint field was proposed. The preliminary results shown

were satisfactory and produced a flow actuation which remained active for a much longer time. This suggests that the use of this re-scaling approach can be beneficial in future applications of adjoint-based optimisation.

Chapter 10

Summary and Conclusions

In the current thesis, we have presented the development and several applications of an adjoint-based optimal flow control framework for compressible DNS, which was appended to an existing parallel in-house FORTRAN code named HiPSTAR (Sandberg, 2015). The background theory supporting this investigation was introduced in chapter 2, whereas the governing equations and code numerics were detailed in chapters 3 and 4, respectively. Before delving into the main chapters of the report, several canonical tests were executed for verification purposes of the newly developed framework in chapter 5. With minor coding effort and exploiting the modular setup of the optimal flow control framework, we were able to extract a new feature of our code. This new capability permits the computation of both stable and unstable exact steady and periodic flow solutions in compressible flows over complex geometries.

A leading theory of chaotic systems shows how they may be approximated by an expansion over periodic orbits (Cvitanović et al., 2016). Thus, a future potential application of these periodic orbits (which we have not explored in this work) resides in carrying out optimal flow control in chaotic systems computing the sensitivities of the periodic orbits to the flow actuation. This would remove the need of computing the sensitivities of chaotic trajectories, which as discussed in chapters 2 and 9, it presents major limitations. Hence, in chapter 6 we showed for the first time the families across Mach number of periodic and equilibrium solutions in a two-dimensional compressible flow over a rectangular cavity. The Reynolds number based on the cavity depth ($Re = 2000$) was carefully chosen to avoid any convective instabilities in the flow-field, restricting the system to be only driven by purely compressible self-sustained oscillations. The evolution of the compressible events across Mach number was thoroughly documented for the two families of flow solutions. Especially, the flow-acoustic shear layer interaction present this type of flows was carefully studied, as it completely dominates the system's dynamics and also its sound radiation. In addition, we were able to show how the two families of periodic and steady solutions collapse in the quasi-incompressible regime ($M \approx 0.30$). This proved that the flow compressibility has a destabilising effect in cavity flows, a phenomenon which is not fully understood by the literature (e.g. Yamouni et al., 2013). Also in this chapter, the family of periodic solutions was shown stable, whereas a brief preliminary insight into the stability of the

Chapter 10 Summary and Conclusions

steady solution at $M = 0.5$ revealed its unstable character. A more thorough stability analysis on the family of steady solutions was carried out in chapter 7. In particular, DMD was used to perform a global linear stability analysis and recover the adjoint global modes and eigenspectrum of the steady solutions. In addition, the forward global modes were also approximated for the steady solution at $M = 0.5$. This was accomplished using DMD over a non-linear forward dataset, assuming linear flow behaviour in the initial stages of time marching. The forward approximated modes revealed the laminar shear layer as the leading amplifier of the instability. Further, the respective adjoint modes highlighted the cavity's leading edge and incoming boundary layer as the regions with the highest receptivity. Additionally, the evolution of the frequency corresponding to the unstable mode was tracked throughout the steady-to-periodic transition at $M = 0.5$. This showed that the system quickly abandons the linear behaviour, even before the shear layer and downstream vortices begin the merging process. The evolution of the eigenvalue instability was also analysed across Mach number, which showed that the transitions from steady to periodic flow and vice versa do not occur at the same Mach number. This phenomenon confirmed that the bifurcation of the steady and periodic families is of subcritical Hopf type, hence, proving that periodic or steady solutions are not necessarily exclusive in cavity flows, where it has been demonstrated that they both can be stable over a short Mach number range. Moreover, the eigenvalue branch which eventually becomes unstable as the Mach number is raised was also identified. The lack of both forward modes and a finer sampling of solutions across Mach number permitted us only to hypothesise on the onset of the unstable character, and analyse such modes from a receptivity point of view.

In chapters 8 and 9 we slightly moved away from this fundamental side of fluid dynamics towards a more applied field. More precisely, in these two chapters, we have shown two novel applications of adjoint-based optimal flow control, where both of them were motivated by practical engineering applications. Chapter 8 studied the optimal control of a two-dimensional open cavity flow, where the target of the actuation was to minimise the pressure fluctuations at the sensor location. Due to the higher Reynolds number of this case ($Re = 5000$) compared to the cavity flow regarded in chapter 6, the presence of convective instabilities prevents the flow from saturating into a stable periodic orbit. To investigate the influence of the sensor on the resulting flow actuation, three different sensing scenarios were carried out, where reductions of 79.6%, 20.4% and 82.0% in the cost function were achieved respect to the non-actuated case. For all three cases, p_{rms} contours revealed a decrease of pressure fluctuations in the area surrounding the sensors, by either changing the directivity of the noise radiation or reducing the overall sound emission. Furthermore, Fourier analysis of temporal signals in several monitor points, accompanied by DMD, shed light into the variations in the dominant physical mechanisms which govern the flow-field on each scenario. As well, DMD analysis was carried out over the control subspaces to show the leading frequencies and forcing shapes. This provides a valuable insight into the ideal forcing patterns to achieve noise reduction in a cavity flow, which can be used in future investigations to model more realistic flow actuation strategies. Lastly, chapter 9 showed the application of the newly developed adjoint framework to a three-dimensional compressible backward-facing step flow. Differently from the cavity flow studied in chapter 8, the presence

of a highly chaotic flow separation area resulted in an added complexity on the application of the method. The target of the flow actuation was to minimise the low wavenumber pressure fluctuations at the sensor location, which was placed in the near wall region where the flow reattachment occurs. Despite the flow control framework produced an inefficient time-dependent control, where the actuation was effectively switched off most of the optimisation period, a total cost reduction of 7.05% was obtained with respect to the non-actuated case. The origin of this inefficient flow actuation was found to reside in the chaotic behaviour of the flow. Under such conditions, the adjoint method fails to provide the correct flow sensitivities, leading to an exponential blow-up of the adjoint field with backwards time-marching. In an attempt to circumvent some of the intrinsic limitations of the method which cause this damped flow actuation, a gradient normalisation based on the L2 norm of the adjoint field was proposed. The preliminary results shown were satisfactory and produced a flow actuation which remained active for a much longer time. This suggests that the use of this re-scaling approach can be beneficial in future applications of adjoint-based optimisation.

Appendix A

Derivation of the Adjoint Euler Equations

The derivation of the adjoint equations studied in the present research is based on the formulation of the fluid governing equations where the energy equation is expressed in terms of pressure (see section 3.2). For the sake of clarity, only the inviscid part of the governing equations (Euler equations) is considered in this appendix. If the governing equations are written in the form $\mathcal{N}(Q) = 0$, results into

$$\begin{pmatrix} \frac{\partial \rho}{\partial t} + \nabla \cdot \vec{m} \\ \frac{\partial \vec{m}}{\partial t} + \nabla \cdot \left(\vec{m} \otimes \frac{\vec{m}}{\rho} \right) + \nabla p \\ \frac{\partial p}{\partial t} + \nabla \cdot \left(p \frac{\vec{m}}{\rho} \right) + (\gamma - 1) \left[p \left(\nabla \cdot \frac{\vec{m}}{\rho} \right) \right] \end{pmatrix} = 0, \quad (\text{A.1})$$

where the vector \vec{m} represents the momentum components ($\rho \vec{u}$). Additionally, the Lagrange functional is introduced as

$$\mathcal{L}(Q, Q^*, \phi) = \mathcal{J}(Q, \phi) - Q^{*T} \cdot \mathcal{N}(Q, \phi), \quad (\text{A.2})$$

where \mathcal{J} is the cost function or target function, which takes part in the adjoint equations as a source term. The derivation of the adjoint equations begins with the linearisation of the system of equations (A.1) around the state variables, which for the present case are $Q = [p, \vec{m}, \rho]$. This linearisation is achieved by setting the first variation of the system respect to the state variables to zero, or equivalently, taking Fréchet derivatives respect to each of the state variables. If the differentiation method is chosen, an arbitrary state \tilde{Q} must be chosen such that

$$\lim_{\epsilon \rightarrow 0} \left(\frac{\mathcal{L}(Q + \epsilon \tilde{Q}, Q^*, \phi) - \mathcal{L}(Q, Q^*, \phi)}{\epsilon} \right) = 0, \quad (\text{A.3})$$

where Q^* represents the Lagrange multipliers of the system, know as *co-state* variables or *adjoint* variables. \mathcal{L} represents the Lagrange functional (A.2) where the cost function \mathcal{J} can be just

Appendix A Derivation of the Adjoint Euler Equations

ignored, as it is case specific and its derivation follows the same procedure. The differentiation of the system of governing equations (A.1) following (A.3) results into

$$\mathcal{N}'(Q) \tilde{Q} = \left(\begin{array}{l} \frac{\partial \tilde{p}}{\partial t} + \nabla \cdot \tilde{\vec{m}} \\ \frac{\partial \tilde{\vec{m}}}{\partial t} + \nabla \cdot \left(\tilde{\vec{m}} \otimes \frac{\tilde{\vec{m}}}{\rho} + \tilde{\vec{m}} \otimes \left(\frac{\tilde{\vec{m}}}{\rho} - \tilde{\rho} \frac{\tilde{\vec{m}}}{\rho^2} \right) \right) + \nabla \tilde{p} \\ \frac{\partial \tilde{p}}{\partial t} + \nabla \cdot \left(\tilde{p} \frac{\tilde{\vec{m}}}{\rho} + p \frac{\tilde{\vec{m}}}{\rho} - \tilde{\rho} \frac{\tilde{\vec{m}}}{\rho^2} \right) + (\gamma - 1) \left[\tilde{p} \left(\nabla \cdot \frac{\tilde{\vec{m}}}{\rho} \right) + p \left(\nabla \cdot \left(\frac{\tilde{\vec{m}}}{\rho} - \tilde{\rho} \frac{\tilde{\vec{m}}}{\rho^2} \right) \right) \right] \end{array} \right) = 0, \quad (\text{A.4})$$

At this point is worth introducing the Green-Lagrange identity as

$$\langle Q^*, \mathcal{N}'(Q) \tilde{Q} \rangle = \langle \mathcal{N}^*(Q) Q^*, \tilde{Q} \rangle + BT, \quad (\text{A.5})$$

which is satisfied by the inner product $Q^{*T} \cdot \mathcal{N}(Q)$ present in (A.2). Hence, (A.4) can be rearranged, removing the as much as possible the derivatives from the arbitrary state \tilde{Q} by integration by parts, which introduces additional terms defined only at the boundaries, known as *boundary terms* (see appendix B). Performing the integration by parts and grouping the same \tilde{Q} variables, the adjoint Euler equations can be written in operator form as

$$\mathcal{N}^*(Q) Q^* = \left(\begin{array}{l} -\frac{\partial \rho^*}{\partial t} - \vec{u} \cdot \nabla \rho^* + (\gamma - 1) \rho^* \nabla \cdot \vec{u} - \nabla \cdot \vec{m}^* \\ -\frac{\partial \vec{m}^*}{\partial t} - \gamma \frac{p}{\rho} \nabla \rho^* - (\gamma - 1) \frac{p^*}{\rho} \nabla p - \vec{u} \left(\nabla \vec{m}^* + (\nabla \vec{m}^*)^T \right) - \nabla p^* \\ -\frac{\partial p^*}{\partial t} + \gamma \frac{p \vec{u}}{\rho} \cdot \nabla \rho^* + (\gamma - 1) \frac{p^*}{\rho} \vec{u} \cdot \nabla p + \vec{u} \cdot (\vec{u} \cdot \nabla) \vec{m}^* \end{array} \right) = -BT, \quad (\text{A.6})$$

where BT represents the boundary terms.

Appendix B

Adjoint Boundary Terms

The temporal and spatial boundary terms that arise due to integration by parts in the derivation of the adjoint Euler equations are

$$\begin{aligned}
 BT = & \int_{\Omega} (p^* \tilde{\rho} + m_u^* \tilde{m}_u + m_v^* \tilde{m}_v + m_w^* \tilde{m}_w + \rho^* \tilde{p})|_{t=0}^{t=T} d\Omega \\
 & + \int_0^T \int_{\Gamma} \vec{n} \left[\tilde{p} (\rho^* (u + v + w) + m_u^* + m_v^* + m_w^*) \right. \\
 & \quad + \tilde{m}_u \left(m_u^* (2u + v + w) + m_v^* v + m_w^* w + \gamma \frac{p}{\rho} \rho^* + p^* \right) \\
 & \quad + \tilde{m}_v \left(m_u^* u + m_v^* (u + 2v + w) + m_w^* w + \gamma \frac{p}{\rho} \rho^* + p^* \right) \\
 & \quad + \tilde{m}_w \left(m_u^* u + m_v^* v + m_w^* (u + v + 2w) + \gamma \frac{p}{\rho} \rho^* + p^* \right) \\
 & \quad - \tilde{\rho} \left(\frac{m_u^*}{\rho^2} (u^2 + uv + uw) + \frac{m_v^*}{\rho^2} (uv + v^2 + vw) \right. \\
 & \quad \left. + \frac{m_w^*}{\rho^2} (uw + vw + w^2) + \gamma \frac{p}{\rho} \rho^* (u + v + w) \right) \Big] d\Gamma dt + BT_{Sv} + BT_{BC},
 \end{aligned} \tag{B.1}$$

where BT_{Sv} accounts for the additional boundary terms due to the viscous terms and BT_{BC} represents the boundary terms from the forward boundary condition. Since the \tilde{Q} components are arbitrary, every term multiplying them has to hold individually. Hence, the correct boundary conditions will cancel out the boundary terms at the boundaries of the domain. The first integral in (B.1) imposes the terminal condition of the adjoint problem. Unless a specific term is included in the Lagrangian functional to impose an given terminal condition, from (B.1) follows that $\rho^*|_{t=T} = 0$, $\vec{m}^*|_{t=T} = 0$ and $p^*|_{t=T} = 0$. By imposing the forward boundary conditions in the Lagrange functional, additional adjoint variables (ψ) will appear in (B.1) through the BT_{BC} term, which enforce the adjoint boundary conditions. Note that these additional adjoint variables are only defined at the boundary and they are uncoupled from the adjoint governing equations.

Appendix B Adjoint Boundary Terms

If the full Navier-Stokes equations are considered, the additional boundary terms are

$$\begin{aligned}
 BT_{Sv} = \int_0^T \int_{\Gamma} & \left[\tilde{p} \left(\rho^* \left(\frac{3\mu_a}{PrReM^2} \frac{\partial T}{\partial n} + \frac{3\mu\gamma}{PrRe} \frac{\partial 1/\rho}{\partial n} \right) + \frac{\mu_a}{3Re} \vec{m}^* \cdot A - \frac{3\mu\gamma}{PrRe\rho} \frac{\partial \rho^*}{\partial n} \right) \right. \\
 & + \tilde{m}_u \left(-\frac{\mu}{3Re\rho} B_1 + 2(\gamma-1) \frac{\rho^*\mu}{3Re\rho} A_1 + \frac{\mu}{3Re} \frac{\partial 1/\rho}{\partial n} C_1 \right) \\
 & + \tilde{m}_v \left(-\frac{\mu}{3Re\rho} B_2 + 2(\gamma-1) \frac{\rho^*\mu}{3Re\rho} A_2 + \frac{\mu}{3Re} \frac{\partial 1/\rho}{\partial n} C_2 \right) \\
 & + \tilde{m}_w \left(-\frac{\mu}{3Re\rho} B_3 + 2(\gamma-1) \frac{\rho^*\mu}{3Re\rho} A_3 + \frac{\mu}{3Re} \frac{\partial 1/\rho}{\partial n} C_3 \right) \\
 & + \tilde{\rho} \left(\frac{\mu}{3Re\rho} (\vec{u} \cdot B - 2(\gamma-1) \vec{u} \cdot A) + \frac{\mu_b}{3Re} \vec{m}^* \cdot A - \frac{\mu}{3Re} \vec{m}^* \cdot D \right. \\
 & \quad \left. + \rho^* \left(\frac{3\mu_b}{PrReM^2} \frac{\partial T}{\partial n} - \frac{3\mu\gamma}{PrRe} \frac{\partial p/\rho^2}{\partial n} \right) + \frac{3\mu\gamma p}{PrRe\rho^2} \frac{\partial \rho^*}{\partial n} \right) \\
 & + \frac{\partial \tilde{p}}{\partial n} \left(\rho^* \frac{3\mu\gamma}{PrRe\rho} \right) \vec{n} + \frac{\partial \tilde{m}_u}{\partial n} \left(\frac{\mu}{3Re\rho} C_1 \right) \vec{n} + \frac{\partial \tilde{m}_v}{\partial n} \left(\frac{\mu}{3Re\rho} C_2 \right) \vec{n} \\
 & \left. + \frac{\partial \tilde{m}_w}{\partial n} \left(\frac{\mu}{3Re\rho} C_3 \right) \vec{n} - \frac{\partial \tilde{\rho}}{\partial n} \left(\rho^* \frac{3\mu\gamma p}{PrRe\rho^2} + \frac{\mu}{3Re\rho} \vec{m}^* \cdot E \right) \vec{n} \right] d\Gamma dt,
 \end{aligned} \tag{B.2}$$

where the vectors A , B , C , D and E are

$$\begin{aligned}
 A &= \begin{bmatrix} 10 \frac{\partial u}{\partial n} + \frac{\partial v}{\partial n} + \frac{\partial w}{\partial n} \\ \frac{\partial u}{\partial n} + 10 \frac{\partial v}{\partial n} + \frac{\partial w}{\partial n} \\ \frac{\partial u}{\partial n} + \frac{\partial v}{\partial n} + 10 \frac{\partial w}{\partial n} \end{bmatrix} & B &= \begin{bmatrix} 10 \frac{\partial m_u^*}{\partial n} + \frac{\partial m_v^*}{\partial n} + \frac{\partial m_w^*}{\partial n} \\ \frac{\partial m_u^*}{\partial n} + 10 \frac{\partial m_v^*}{\partial n} + \frac{\partial m_w^*}{\partial n} \\ \frac{\partial m_u^*}{\partial n} + \frac{\partial m_v^*}{\partial n} + 10 \frac{\partial m_w^*}{\partial n} \end{bmatrix} \\
 C &= \begin{bmatrix} 10m_u^* + m_v^* + m_w^* \\ m_u^* + 10m_v^* + m_w^* \\ m_u^* + m_v^* + 10m_w^* \end{bmatrix} & D &= \begin{bmatrix} 10 \frac{\partial u/\rho}{\partial n} + \frac{\partial v/\rho}{\partial n} + \frac{\partial w/\rho}{\partial n} \\ \frac{\partial u/\rho}{\partial n} + 10 \frac{\partial v/\rho}{\partial n} + \frac{\partial w/\rho}{\partial n} \\ \frac{\partial u/\rho}{\partial n} + \frac{\partial v/\rho}{\partial n} + 10 \frac{\partial w/\rho}{\partial n} \end{bmatrix} \\
 E &= \begin{bmatrix} 10u + v + w \\ y + 10v + w \\ u + v + 10w \end{bmatrix},
 \end{aligned}$$

and also with μ_a and μ_b as

$$\begin{aligned}
 \mu_a &= \frac{3}{2} \left(\frac{p\gamma M^2}{\rho} \right)^{1/2} \frac{\gamma M^2}{\rho} - \left(\frac{p\gamma M^2}{\rho} \right)^{3/2} \left(\frac{\frac{\gamma M^2}{\rho} (1+S)}{\left(\left(\frac{p\gamma M^2}{\rho} \right) + S \right)^2} \right) \\
 \mu_b &= -\frac{p}{\rho} \mu_a.
 \end{aligned}$$

B.1 Boundary Conditions

Euler Isothermal Slip Wall

The forward conditions for a slip isothermal wall¹ in an inviscid flow are

$$\frac{\partial u}{\partial n} = 0, \quad \frac{\partial w}{\partial n} = 0, \quad v_w = 0, \quad T_w = \text{const.} \quad (\text{B.3})$$

These conditions are taken into account through the Lagrangian functional, introducing four additional adjoint variables which are only defined at the boundaries. Expressing the conditions in terms of the state variables² leads to

$$\psi^T \cdot BC(Q, \phi) = \psi_1 \frac{\partial}{\partial n} \left(\frac{\rho u}{\rho} \right) + \psi_2 \frac{\partial}{\partial n} \left(\frac{\rho w}{\rho} \right) + \psi_3 \left(\frac{\rho v}{\rho} \right) + \psi_4 \left(\frac{p\gamma M^2}{\rho} - \text{const} \right). \quad (\text{B.4})$$

Applying the same procedure as indicated in equation 3.11, this time using the Lagrangian functional indicated in (4.4), the extra boundary terms are

$$\begin{aligned} BT_{BC} = \int_0^T \int_{\Gamma} & \left[\tilde{p} \left(-\psi_4 \frac{\gamma M^2}{\rho} \right) \vec{n} + \tilde{m}_u \left(\psi_1 \frac{\partial 1/\rho}{\partial n} \right) + \tilde{m}_v \left(-\psi_3 \frac{1}{\rho} \right) \vec{n} + \tilde{m}_w \left(\psi_2 \frac{\partial 1/\rho}{\partial n} \right) \right. \\ & + \tilde{\rho} \left(\psi_3 \frac{v}{\rho} \vec{n} + \psi_4 \frac{p\gamma M^2}{\rho^2} \vec{n} - \psi_1 \frac{\partial u/\rho}{\partial n} - \psi_2 \frac{\partial w/\rho}{\partial n} \right) + \frac{\partial \tilde{m}_u}{\partial n} \left(\psi_1 \frac{1}{\rho} \right) \vec{n} \\ & \left. + \frac{\partial \tilde{m}_w}{\partial n} \left(\psi_2 \frac{1}{\rho} \right) \vec{n} - \frac{\partial \tilde{\rho}}{\partial n} \left(\psi_1 \frac{u}{\rho} + \psi_2 \frac{w}{\rho} \right) \vec{n} \right] d\Gamma dt. \end{aligned} \quad (\text{B.5})$$

Regrouping (B.1) and (B.5) and solving the system of equations results in the following adjoint boundary conditions:

$$\rho^* = \frac{m_v^*}{(u+w)(\gamma-1)} \quad (\text{B.6})$$

$$m_u^* = 0 \quad (\text{B.7})$$

$$m_w^* = 0 \quad (\text{B.8})$$

$$p^* = -\rho^* \left(\frac{\gamma p}{\rho} \right), \quad (\text{B.9})$$

¹This derivation assumes a slip isothermal wall located at the boundary in the y direction. An analogous procedure should be followed for the x boundary.

²For simplicity, the subscript ϕ_w is omitted in this section.

Appendix B Adjoint Boundary Terms

with the ψ adjoint boundary variables as

$$\psi_1 = 0 \quad (\text{B.10})$$

$$\psi_2 = 0 \quad (\text{B.11})$$

$$\psi_3 = m_u^* \rho u + m_v^* \rho (u + w) + m_w^* \rho w + \rho^* \gamma p + p^* \rho \quad (\text{B.12})$$

$$\psi_4 = \rho^* \frac{\rho}{\gamma M^2} (u + w) + \frac{\rho}{\gamma M^2} (m_u^* + m_v^* + m_w^*). \quad (\text{B.13})$$

Isothermal No-Slip Wall

The necessary conditions in the forward simulation for an isothermal no-slip wall are

$$u_w = v_w = w_w = 0, \quad T_w = \text{const}, \quad (\text{B.14})$$

which are taken into account in the Lagrangian functional as

$$\psi^T \cdot BC(Q, \phi) = \psi_1 \left(\frac{\rho u}{\rho} \right) + \psi_2 \left(\frac{\rho v}{\rho} \right) + \psi_3 \left(\frac{\rho w}{\rho} \right) + \psi_4 \left(\frac{p \gamma M^2}{\rho} - \text{const} \right), \quad (\text{B.15})$$

where the extra boundary terms are

$$BT_{BC} = \int_0^T \int_{\Gamma} \vec{n} \left[-\tilde{p} \left(\psi_4 \frac{\gamma M^2}{\rho} \right) - \tilde{m}_u \left(\psi_1 \frac{1}{\rho} \right) - \tilde{m}_v \left(\psi_2 \frac{1}{\rho} \right) - \tilde{m}_w \left(\psi_3 \frac{1}{\rho} \right) + \tilde{\rho} \left(\psi_1 \frac{u}{\rho} + \psi_2 \frac{v}{\rho} + \psi_3 \frac{w}{\rho} + \psi_4 \frac{p \gamma M^2}{\rho^2} \right) \right] d\Gamma dt. \quad (\text{B.16})$$

After substituting (B.16) and (B.2) into (B.1), the system of ten equations can be solved, resulting into the adjoint boundary conditions for an isothermal no-slip wall on a viscous flow

$$(m_u^*)_w = (m_v^*)_w = (m_w^*)_w = (\rho^*)_w = 0. \quad (\text{B.17})$$

Appendix C

Exact Solutions - Simulation Parameters

This appendix gathers the simulation parameters corresponding to the steady (table C.1) and periodic (table C.2) exact flow solutions studied in chapters 6 and 7. The parameters describing the boundary layer's momentum thickness in the table C.2 are averaged values over the entire solution's period. Despite the computation of periodic flow solutions for $M \lesssim 0.30$ yielded steady states, we have run these two solutions further with the steady solution setup (as indicated in subsection 6.2.2) to ensure they satisfy the same convergence criterion than the other steady solutions. This is the reason of the slight differences in the simulation parameters presented tables C.1 and C.2 at $M = 0.25$ and $M = 0.30$.

M_∞	Re_D	Re_Θ	L/Θ	μ_r	μ_i
0.25	2000	50.51	118.77	-	-
0.30	2000	50.36	119.13	-0.176693	0.824514
0.35	2000	50.59	118.56	-0.053106	0.412341
0.40	2000	49.00	122.52	0.019984	0.256365
0.45	2000	50.34	118.77	0.031612	0.251265
0.50	2000	50.52	118.74	0.042644	0.246062
0.55	2000	50.51	118.77	0.048355	0.241150
0.60	2000	47.79	125.52	0.052891	0.236504
0.65	2000	44.40	135.12	0.055347	0.231924

Table C.1: Parameters and stability of the 2D cavity equilibrium solutions. D and Θ are the cavity depth and momentum thickness at the cavity's leading edge, respectively. The columns μ_r and μ_i contain the growth rate and frequency of the leading (first) eigenvalue of the unstable branch.

Appendix C Exact Solutions - Simulation Parameters

M_∞	Re_D	Re_Θ	L/Θ	St
0.25	2000	50.50	118.82	–
0.30	2000	50.36	119.14	–
0.35	2000	50.55	118.68	0.2489
0.40	2000	49.10	122.19	0.2447
0.45	2000	50.38	119.08	0.2393
0.50	2000	50.36	119.15	0.2335
0.55	2000	49.75	120.58	0.2282
0.60	2000	46.34	129.46	0.2231
0.64	2000	42.84	140.05	0.2190
0.65	2000	41.83	143.40	0.2179
0.66	2000	41.48	144.64	0.2172
0.67	2000	40.76	147.20	0.2167
0.68	2000	40.23	149.15	0.2167
0.69	2000	39.32	152.32	0.2170
0.70	2000	38.47	155.96	0.2170
0.71	2000	37.54	159.80	0.2162
0.72	2000	36.31	165.25	0.2162
0.73	2000	34.54	173.69	0.2131
0.74	2000	32.73	183.34	0.2114
0.75	2000	30.88	194.28	0.2097
0.76	2000	29.32	204.60	0.2086
0.77	2000	28.31	211.91	0.2086
0.78	2000	27.10	221.41	0.2104
0.79	2000	24.96	240.42	0.2107
0.80	2000	23.06	260.24	0.2100
0.80-fd	2000	64.37	93.20	0.2090

Table C.2: Simulation parameters of the 2D cavity periodic solutions. D and Θ are the cavity depth and momentum thickness at the cavity's leading edge, respectively. St shows the Strouhal number of each periodic trajectory.

Appendix D

Grid Independence of 2D Cavity Flows

This appendix shows a grid independence study on the two-dimensional open cavity flows studied in chapters 6, 7 and 8. Bear in mind that the domain discretisation used throughout these three chapters is the same, despite all the cavity flow setups investigated do not have the same Reynolds number. Hence, by carrying out such grid independence study on the case with the highest Reynolds number case will automatically ensure the grid independence for all the flow solutions with lower Re . For this reason, we restrict this grid study to the setup shown in chapter 8 (uncontrolled case), where the Reynolds number is the highest ($Re = 5000$). The additional grid required for this study referred to as the fine grid, contains twice the number of points of the DNS grid in each direction. Both fine and DNS cases were run for over 350 time units, where Favre averaged statistics were collected during each run. In particular, we compare these two flow fields by tracking the streamwise and vertical velocity components

$$\hat{u} = \frac{\overline{\rho u}}{\bar{\rho}} \quad \text{and} \quad \hat{v} = \frac{\overline{\rho v}}{\bar{\rho}}, \quad (\text{D.1})$$

and also the diagonal components of the Reynolds stress tensor σ

$$\sigma_{11} = \frac{\overline{\rho u u}}{\bar{\rho}} - \hat{u}\hat{u} \quad \text{and} \quad \sigma_{22} = \frac{\overline{\rho v v}}{\bar{\rho}} - \hat{v}\hat{v}. \quad (\text{D.2})$$

As shown in figure D.1, these four statistical quantities show an outstanding agreement for both DNS and fine grids, which confirms the grid independence of all the cavity flow solutions studied. Perhaps the largest disagreement between these two cases can be observed in figures D.1c and D.1d in the near wall region inside the cavity. In any case, these minor differences should not have any repercussion in the findings presented in this thesis, since the flow in that area is purely driven by the separated shear layer, which shows an excellent grid independence. Furthermore, these brief differences are expected to rapidly decrease with Re , confirming as well the grid independence of the lower Re setups studied in chapters 6 and 7.

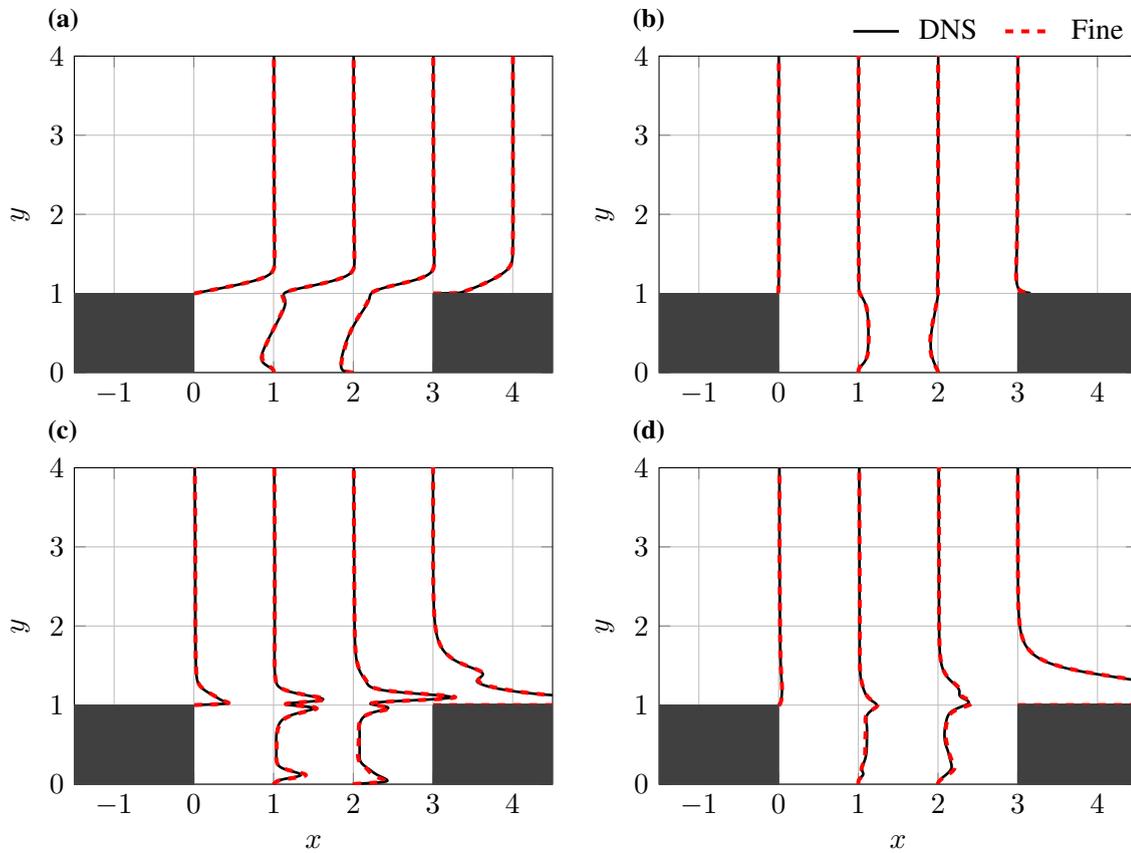


Figure D.1: Grid independence at $Re = 5000$ shown with four different Favre-averaged quantities at four different streamwise locations along the cavity. Subfigures (a) and (b) show profiles of the streamwise and vertical velocity components, respectively. On the other hand, subfigures (c) and (d) represent the σ_{11} and σ_{22} components of the Reynolds stress tensor σ . The magnitudes of σ_{11} and σ_{22} have been scaled up (x200) in their respective plots for clarity purposes.

Appendix E

Limitations of Adjoint-based Optimisation in Separated Flows

This appendix shows a preliminary sensitivity study of a two-dimensional BFS flow. The aim of this appendix is to illustrate the limitations of the adjoint method for separated flows, providing a context for the optimal flow control exercise shown in chapter 9. Given the preliminary character of this study, the results presented herein should only be regarded as qualitatively.

E.1 Introduction

The understanding of turbulent flows is one of the greatest scientific challenges. Yet to our current knowledge, turbulence is of a chaotic nature, which makes it extremely difficult to predict. In separated flows, the interaction between the large and small scales is accentuated, which adds an additional degree of difficulty to its prediction. A proof of this is, for example, the struggle of RANS models to simulate the correct behaviour of flows under large separation regions. Additionally, the presence of wall boundaries is crucial, as it forces the turbulent small scales to break down into smaller ones which widens the already large range of present scale sizes in the flow-field.

In a separated flow, the trajectories of two neighbouring fluid particles diverge exponentially with time due to the chaotic behaviour of turbulence. Consequently, the initial condition is ‘forgotten’ by the system and, the system’s sensitivities cannot be computed because the initial and final conditions are unrelated. This means that if the prediction time horizon is longer than the time which takes two neighbouring trajectories to diverge, the computed sensitivities will be incorrect. [Lea et al. \(2000\)](#) used the Lorenz attractor as an example for chaotic systems, showing how the adjoint analysis failed to give the correct time-averaged sensitivities. They suggested that for relatively long time horizons in relation to the predictability of the system’s time scales, the adjoint quantities diverge exponentially due to cumulative error growth. As an

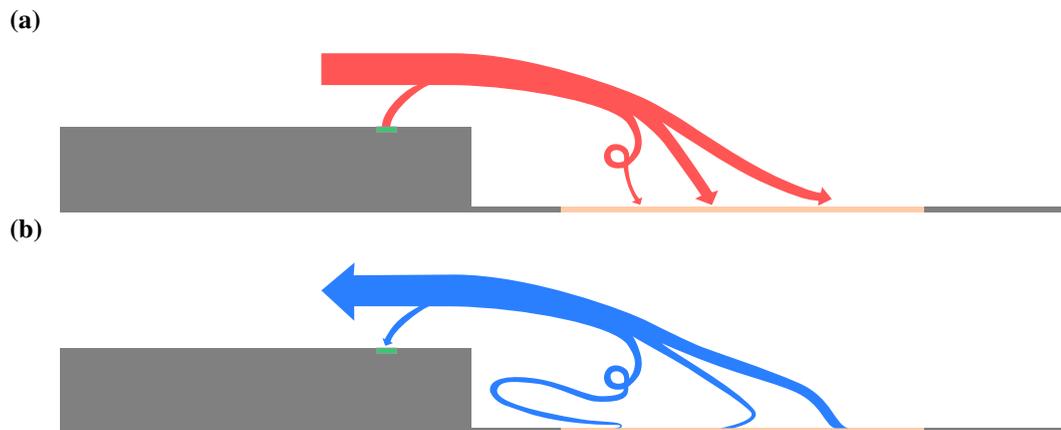


Figure E.1: Illustration of a forward **(a)** and adjoint **(b)** simulations in a flow control case. The green and orange boundaries represent the flow actuation location and the target region respectively.

alternative, they attempted to approximate the sensitivities by an ensemble average over several realisations during a short horizon for different initial conditions. [Eyink et al. \(2004\)](#) tested (also with the Lorenz attractor) two different ensemble averaging methods as a model to approximate time-averaged sensitivities. To achieve satisfactory results they required of the order of 10^6 realisations. Note that this large amount of simulations required by the ensemble average approach make this method unaffordable for most of the problems engineering interest. [Wang \(2013\)](#); [Wang et al. \(2014\)](#) redefined the method to compute the time-averaged sensitivities (for both forward and adjoint approaches) in chaotic systems based on the inverted shadow operator. This circumvents the exponential blow-up of the sensitivities and removes the dependency of the resulting sensitivity to the initial condition present in chaotic systems (*butterfly effect*- [Blonigan, 2016](#)). Unfortunately, due to the high computational demands of this method, it is not suitable yet to deal with turbulent flows.

If performing flow control, the correct election of the time horizon is essential. In the investigation carried out by [Bewley et al. \(2001\)](#), a larger time horizon was shown beneficial when controlling time-averaged flow quantities. On the other hand, they stated that as the prediction horizon is increased, the required effort to optimise the dynamical system for a given control strategy increases exponentially. An intrinsic limitation in the election of this time horizon resides in the travelling time of the control flow scales from the actuation location to the target region. Note that the same limitation occurs in adjoint simulations in the opposite direction. As illustrated in figure [E.1](#) for the case investigated herein (and also in chapter [9](#)), the travelling time will exclusively depend on the flow characteristics between the actuation and target locations.

Our objective in this section is to show on a practical flow control exercise how using the available tools to decompose the flow-field it is possible to identify the structures that contribute to the above-mentioned cumulative error growth. A full description of the case studied is given in section [E.2](#). On section [E.3](#) the influence of the relative location of the target region to the divergence of the adjoint simulations is shown with three different simulations. The identification of

the small scales as the main contributors to the exponential error growth is carried out in section E.4. A flow-scales separation is used to mitigate the divergence of the adjoint system in section E.5. Finally, the summary and conclusions from this chapter are gathered in section E.6.

E.2 Simulation Details

The case regarded in this chapter is a two-dimensional backward-facing step with a laminar inflow. Of particular interest from this geometry are the constant-height flow separation and the flow impingement in the reattachment region (e.g. [Le et al., 1997](#)). This results in an energetic complex flow interaction where the flow motion in the separated flow area is dominated by the chaotic behaviour of turbulence. This model is oriented towards the application of cabin noise reduction. A target or sensing region is located in the lower wall after the step, to cost the flow unsteadiness that would interact with, for example, a vehicle's window placed at the same location. The flow would excite some of the vibrating modes of the structure which would radiate noise into the passenger's cabin.

The step height is chosen to non-dimensionalise the streamwise and vertical directions. With the coordinate origin at the lower corner of the step, the domain ranges from -20 to 80 in x and from 0 to 20 in the y direction. The Reynolds number is $Re = 3000$ where the non-dimensional free-stream velocity is $U_\infty = 1$. The Prandtl and Mach numbers are respectively $Pr = 0.72$ and $M = 0.5$. The time-step for the three simulations is $\Delta t = 0.001$. The inflow laminar boundary layer thickness is specified in section E.3 as it is different for each of the three simulations regarded in this chapter. The computational domain is divided into three blocks as indicated in figure E.2, where the grid dimensions for the first block are 100 in both directions and 600 by 100 for the second and third blocks in the streamwise and normal directions. The flow information is interchanged through the block interfaces using halo points. The grid resolution is clustered towards the wall boundaries to make sure that all structures in the flow-field are resolved, but the resolution is coarsened towards the inflow, outflow and free-stream to avoid reflections in both forward and adjoint simulations. Note that the resolution is kept within the DNS range only up to 20 step heights downstream from the step. The wall boundaries are modelled as isothermal no-slip walls and the inflow, outflow and free-stream are of characteristic type. To carry out adjoint simulations the flow-field is stored every 250 time-steps and it is linearly interpolated in between for every sub-step of the Runge-Kutta scheme. A preliminary study showed that a finer in time flow-field sampling rate lead to the same results. In addition, the grid resolution showed to be sufficient on a preliminary grid convergence study.

To reduce the fluctuations over the structure, a convenient cost function for this particular problem is

$$\mathcal{J}(Q, \phi) = \frac{1}{T} \int_0^T \int_\Omega \frac{1}{2} (\rho - \rho_0)^2 W_s(\vec{x}) \, d\Omega dt, \quad (\text{E.1})$$

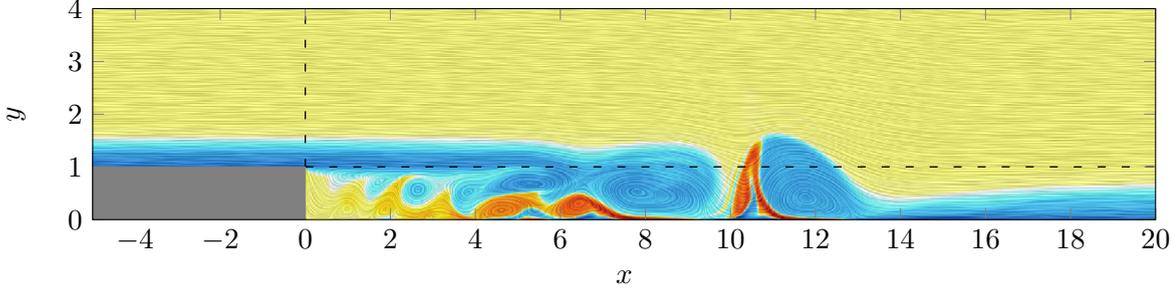


Figure E.2: Instantaneous contours of z -vorticity. Dashed lines represent the computational block interfaces.

where T represents the optimisation horizon and ρ_0 is the averaged density over this time period. $W_s(\vec{x})$ is a spatial weighting function that represents the sensing region, which is defined as

$$W_s(\vec{x}) = \begin{cases} e^{-\frac{(y-0.005)^2}{0.001}} & \text{if } x \in [3, 13] \\ 0 & \text{if } x \notin [3, 13] \end{cases}. \quad (\text{E.2})$$

Note that the sensing region is located just above the wall to avoid any modifications in the boundary terms from the adjoint equations.

E.3 Sensing Region Influence on Adjoint Divergence

The dependency of the cumulative error growth in adjoint simulations to the relative location of the sensing region in the flow-field is studied. In a two-dimensional case, the mean reattachment location of a flow over a backward-facing step can be altered modifying the inflow boundary layer (BL) thickness. If the sensing region remains unmodified, this can be interpreted as a change in the relative position of this sensing region with respect to the flow-field. The present study considers three different BL thickness scenarios, which are referred to as *thin*, *medium* and *thick*. The BL thickness at the step for each of the three cases are $\delta_{thin} \approx 0.162$, $\delta_{medium} \approx 0.382$ and $\delta_{thick} \approx 0.770$. The differences in the three cases are evident in figure E.3, as the location of the mean flow reattachment occurs further downstream as the BL thickens. The density fluctuations in the target region for the thin BL case are mostly dominated by the large vortices shed from the step, since the mean reattachment location occurs just before the sensing zone. Differently, the medium BL case has the mean impingement location close to the centre of the sensing region. This means that all the density fluctuations in this zone are driven by the shear layer flapping, alongside with the randomness of the high energy turbulent structures. Opposite from the thin BL setup, the mean reattachment location for the tick BL case occurs downstream from the target zone. Hence the density fluctuations at the sensing region are caused by the low-velocity structures from within the separated flow zone. The density at the centre of the sensing region is shown in figure E.4 for all three cases for a period of time. The higher amplitudes and frequencies present in the medium case suggest that the predictability of the

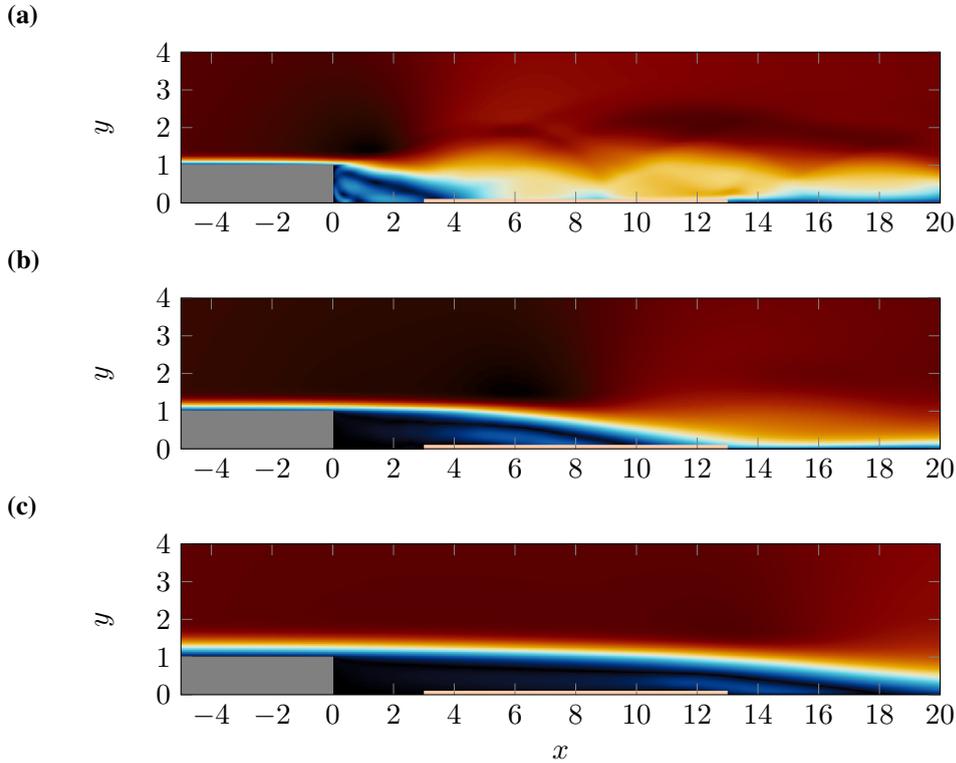


Figure E.3: Mean velocity field for the thin **(a)** medium **(b)** and thick **(c)** cases. The sensing region is represented with the orange line. Contours of velocity magnitude range from 0 to 1.1.

system's time-scales in the region of interest is the shortest from the three simulations. Hence, the divergence of the flow sensitivities for this case should be the highest respect to the other two. Additionally, for the medium and thick cases, the flow sensitivities will be driven by the flow motion from the separated flow area. The flow recirculation present in this region establishes a feedback loop that will enhance the error growth of the flow sensitivities (see figure E.1).

Every adjoint simulation regarded in this study was initialised with the "null" condition (M. D. Gunzburger, 2003) and they ran for $5 \cdot 10^5$ iterations. For non-chaotic systems, the sensitivities should initially grow until they saturate, where they become stable. On the other hand, in situations where the chaos dominates the fluid motion, the adjoint field (or sensitivities) will never reach this saturation due to the already mentioned error growth. This phenomenon can be observed looking at the temporal evolution of the norm of the adjoint field as shown in figure E.5. As expected, the medium BL case has the highest divergence ratio, followed by the thick BL. As the sensing region is placed outside the recirculation region for the thin BL case, only the coherent structures that pass over it are sensed. Hence, the sensitivities of the system are not affected by the chaotic behaviour of the flow inside the recirculation zone. This shows that the relative location of the target zone with respect to the flow-field has an influence on the flow sensitivities.

Note that the divergence of the sensitivities makes the computation of an optimal control an

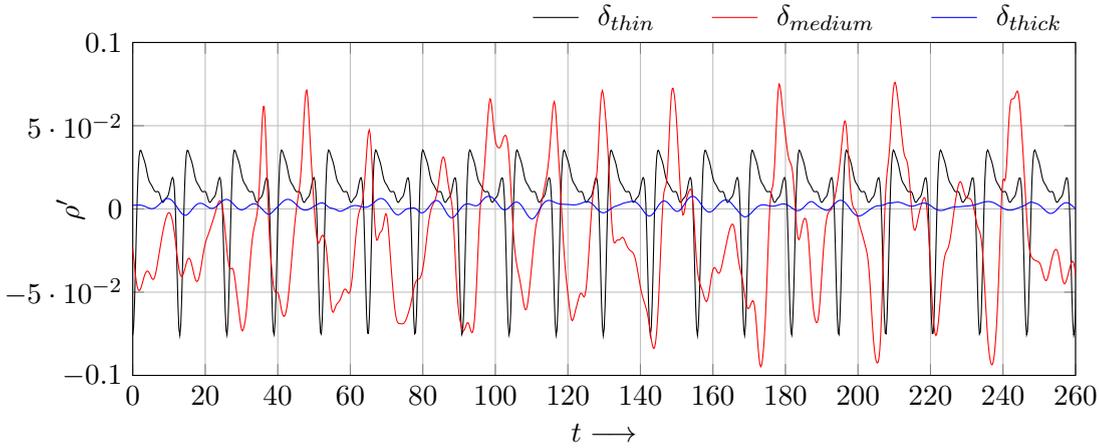


Figure E.4: Density fluctuations at the centre of the sensing zone. The fluctuations are computed as $\rho' = (\rho - \rho_0)$. The simulation time is relative to the start of the data capturing shown.

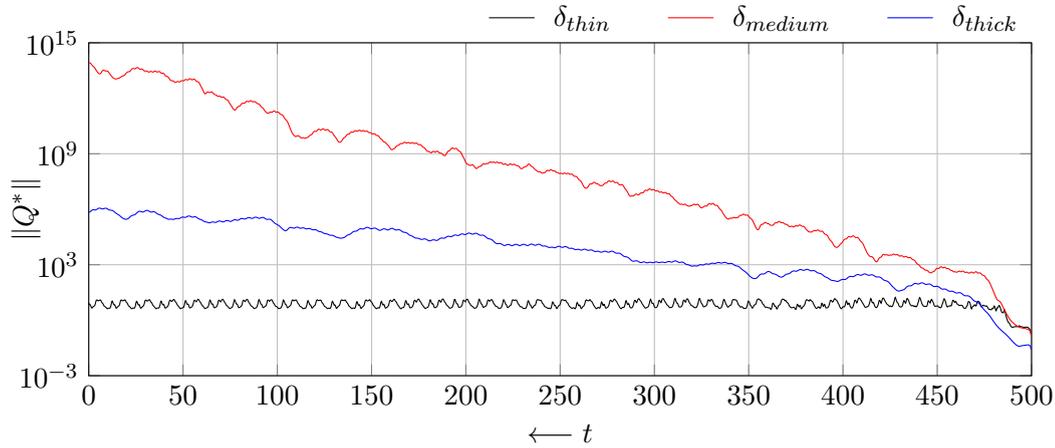


Figure E.5: Adjoint field norm for the three boundary layer cases. Note that the adjoint simulations march backwards in time.

impossible task. The errors introduced due to the unpredictability of turbulence contaminate the flow sensitivities. If performing time-dependent control, due to the exponential growth of the sensitivities, the gradients become unrealistic making the control diverge. Hence short time horizons must be chosen to limit these errors. On the other hand, as illustrated in figure E.1, the time horizon must be long enough to allow communication between the sensing and the actuation.

E.4 Adjoint-Field Decomposition

The flow-field decomposition methods introduced earlier in chapter 2 can be also used to achieve a deeper insight into the mechanisms that cause the divergence of the flow sensitivities. In a

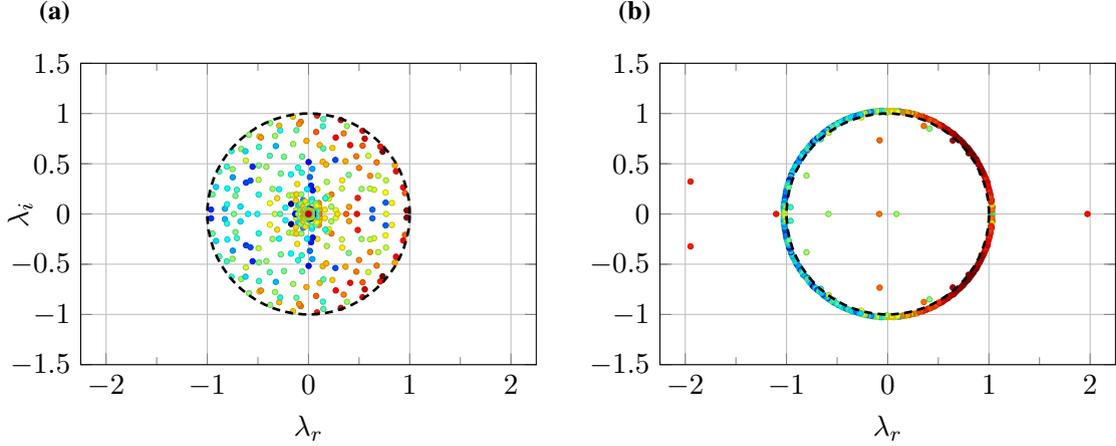


Figure E.6: DMD spectrum for the thin BL (a) and medium BL (b) cases. Each eigenvalue is coloured by its energy content by projecting the dynamic modes onto a POD basis, which also gives a measure of coherence (Schmid, 2010). The levels are independent for each case.

similar fashion as for (2.27), the adjoint simulation can be reformulated in state-space as

$$\mathbf{Q}^{*N} = \mathbf{A}\mathbf{Q}^{*N-1} + \mathbf{B}\mathbf{Y}^{N-1}, \quad (\text{E.3})$$

where the adjoint operator is split into the \mathbf{A} and \mathbf{B} matrices. The adjoint variables are represented by \mathbf{Q}^* and \mathbf{Y} is the input matrix, which in this case is formed by the forward flow-field and the sensing information. For such system, DMD with control (see section 2.2.3) can be applied on the adjoint field to obtain \mathbf{A} and \mathbf{B} . In particular, we are interested in the growth rates of the dynamic modes to identify which ones contribute to the exponential blow up. The mode shape of these unstable modes will highlight the structures that induce the error in the adjoint simulation.

The dynamic mode decomposition is carried out on the full adjoint field for all variables. To avoid any spurious influence of the near-boundary flow regions, the decomposition has been reduced to a domain of interest with the same dimensions as shown in figure E.2. For comparative purposes, the decomposition is carried out on the thin BL (stable) and medium BL (unstable) cases. The snapshots that form the datasets are captured every 500 iterations starting from the adjoint terminal condition during 250000 iterations (500 samples per dataset). Figure E.6 shows the approximated eigenvalues of the matrix \mathbf{A} where λ_r and λ_i are their real and imaginary components. The stability of their corresponding mode is determined by the location of the eigenvalue respect to the unitary circle. The modes with their eigenvalue sitting outside the unitary circle are considered unstable, whereas they are regarded as stable modes otherwise. As expected the approximated eigenvalues of \mathbf{A} for the thin BL case are inside the unitary circle indicating the stability of the flow sensitivities. On the other hand, the majority of the eigenvalues for the medium BL simulation represent the unstable modes that make the flow sensitivities to diverge. An alternative representation of the above plots is achieved calculating the (non-dimensional) frequency $\mu_i = \Im\{m\{\log(\lambda_n)\}\}/(2\pi\Delta t)$ and growth rate $\mu_r = \Re\{m\{\log(\lambda_n)\}\}/\Delta t$

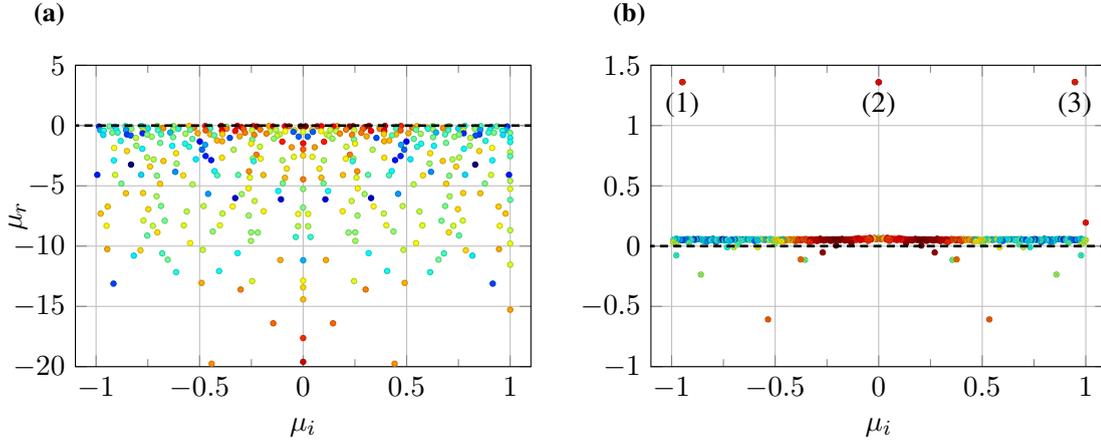


Figure E.7: DMD spectrum for the thin BL **(a)** and medium BL **(b)** cases. Labels (1), (2) and (3) indicate the three leading unstable modes.

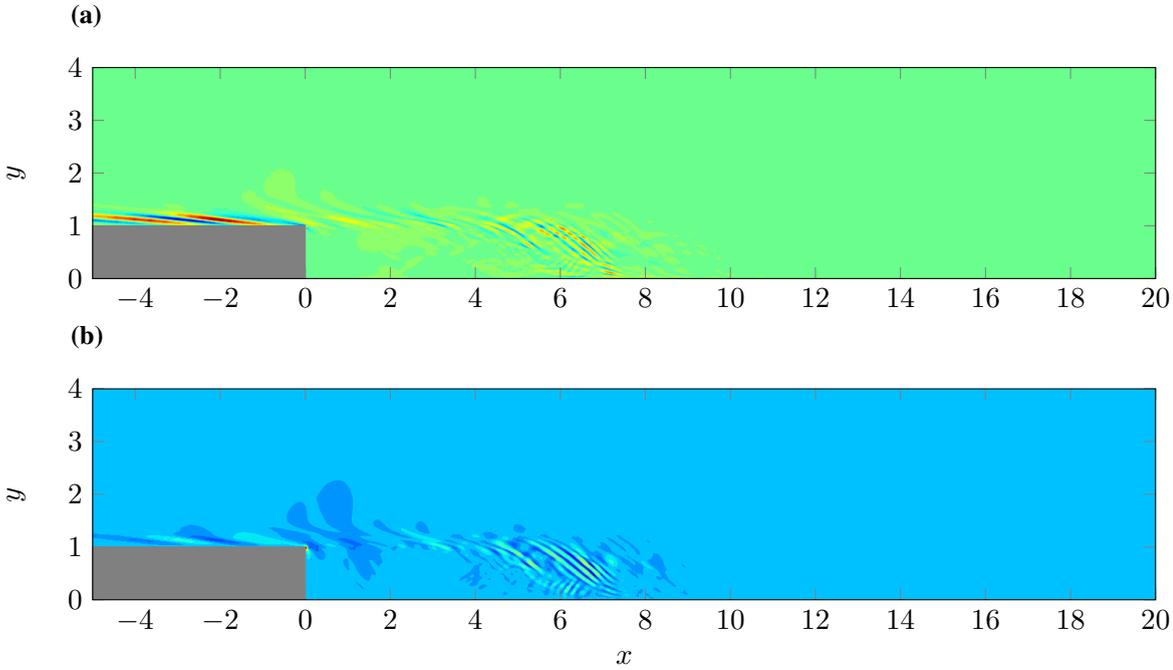


Figure E.8: Dynamic mode shapes of the mode (1) (see figure E.7b) for m_u^* **(a)** and m_v^* **(b)**. The high level of activity in the near flow reattachment zone shows the origin of the divergence in the adjoint simulation.

from each dynamic mode from its corresponding eigenvalue. We again observe from figure E.7 the stability of the flow sensitivities for the thin BL case as all the growth rates for the dynamic modes are lower or equal than zero. In contrast, most of the dynamic modes from the medium BL simulation have a positive growth rate. It is worth mentioning that the energy content of the dynamic modes decays as the frequency increases, which is an indication that the high energy dynamic modes correspond to large structures whereas the low energy modes capture the motion of the small scales. As a result of processing real valued data, the plots are symmetric with respect to $\lambda_i = 0$ and $\mu_i = 0$. From figure E.7 (right), given their high growth rate, the modes

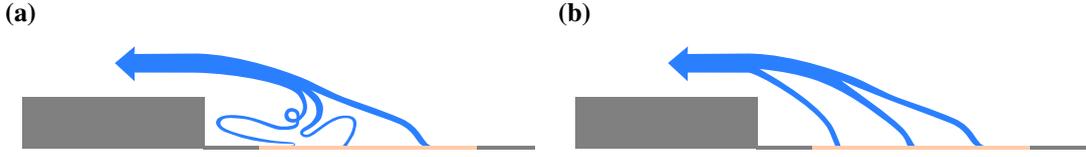


Figure E.9: Illustration of the propagation of the flow sensitivities for the adjoint simulations of the medium BL (a) and thin BL (b) cases.

1, 2 and 3 can be revealed as the main contributors to the divergence of the adjoint-field. Hence, following (2.33) the mode shapes can be calculated to identify the structures that cause the exponential growth of the sensitivities. Figure E.8 shows the dynamic modes for the stream-wise (m_u^*) and vertical (m_v^*) adjoint momentum components, where an evident activity of the modes takes place in the near flow reattachment region (see figure E.3). This finding supports the fact that the chaotic behaviour of the most energetic turbulent structures is principally responsible for the cumulative error growth in adjoint simulations. Far from being convected away, these errors propagate through the flow recirculation area resulting into an unstable feedback loop that contributes to the instability of the other dynamic modes (figure E.9a). Contrastingly, this feedback loop does not occur in the thin BL case, where the sensing is placed outside the separated flow area (figure E.9b).

E.5 Sensitivity Analysis on the Large Scales

From the adjoint system defined in E.3, the input matrix \mathbf{Y} has been identified as the responsible of the divergence of the flow sensitivities for chaotic systems. This matrix can be decomposed into

$$\mathbf{Y} = \begin{bmatrix} \mathbf{Q} & \mathbf{S} \end{bmatrix}, \quad (\text{E.4})$$

where \mathbf{S} contains the data from the sensing region. Additionally one could breakdown the forward state variables into

$$\mathbf{Q} = \tilde{\mathbf{Q}} + \mathbf{Q}', \quad (\text{E.5})$$

where the large scales are represented by $\tilde{\mathbf{Q}}$. The medium-small scales that present chaotic-like behaviour and cause the adjoint system to diverge are grouped into \mathbf{Q}' . A method which has been shown effective to carry out this decomposition is POD (see subsection 2.2.2). The flow-field is decomposed into several modes which are ranked by their energy content, where the large scales are accounted for in the high energy POD modes. Hence, an energy optimal approximation of the input space \mathbf{Y} is achieved dropping the low energy modes, which can be expressed as

$$\mathbf{Y} \approx \tilde{\mathbf{Y}} = \begin{bmatrix} \tilde{\mathbf{Q}} & \mathbf{S} \end{bmatrix}. \quad (\text{E.6})$$

The approximated flow field follows directly from (2.15), which leads to the expression

$$\tilde{\mathbf{Q}}(\vec{x}, t) = \sum_{n=1}^{N_{filt}} a_n(t) \Psi_n(\vec{x}), \quad (\text{E.7})$$

where N_{filt} is the number of POD modes used to reconstruct the flow-field. Figure E.10 shows how the divergence of the adjoint system decreases as the number of POD modes used to reconstruct $\tilde{\mathbf{Q}}$ is reduced. For the present case, the cumulative error growth is completely eliminated when the flow-field is reconstructed with only 3 POD modes. Therefore, this approach provides an observable approximation of the flow sensitivities on separated flows for long time horizons. Despite that the flow sensitivities are an approximation of those from the full system, they could potentially be used to perform flow control as the sensing information has not been approximated.

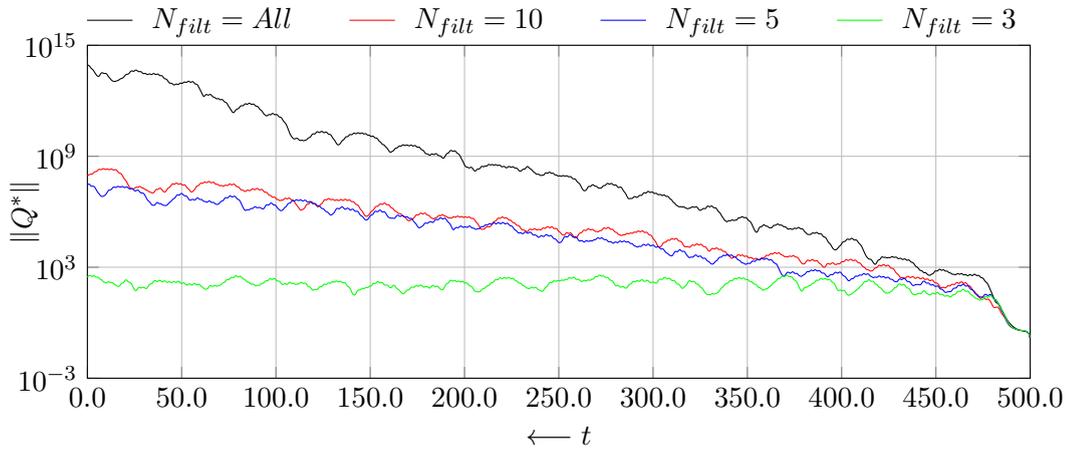


Figure E.10: Comparison of the temporal evolution of the adjoint field norm for the medium BL case using different number of POD modes for the flow-field reconstruction.

E.6 Summary and Conclusions

In the present chapter, the sensitivities of separated flows for long time horizons have been investigated. In particular, a backward-facing step was chosen as it provides a constant flow separation point. The cumulative error growth problem was shown to be dependent on the relative position of the sensing region respect to the flow recirculation zone. This was studied with three different inflow boundary layer thickness cases, which modified the average flow reattachment location. For the thin BL case, the adjoint field did not diverge as the flow sensitivities are not dominated by the chaotic behaviour of turbulence. The reason behind this is that the recirculation region does not interact with the sensing region. The medium BL simulation has the mean reattachment location occurring in the centre of the target zone, which leads to a rapid divergence of the flow sensitivities due to the high energy of the structures impinging on the

wall. The average flow reattachment location takes place further downstream of the sensing region for the thick BL case. This results into a slower divergence as the turbulent structures inside the recirculation region have a lower Reynolds number.

From a dynamic mode decomposition of the adjoint field, it was found that (as expected) the main contributors to the divergence of the flow sensitivities are the small scales present in the flow-field. Due to the flow recirculation, the error induced by this scales is not convected away, resulting in an unstable feedback loop that causes the blow up of the adjoint field. Using proper orthogonal decomposition, the small scales of the flow-field were filtered out. This circumvented the problem of the cumulative error growth and the method was shown to be successful for obtaining an approximation of the flow sensitivities for long time horizons.

References

- Adams, E. W. and Johnston, J. P. (1988). Effects of the separating shear layer on the reattachment flow structure part 2: Reattachment length and wall shear stress. *Experiments in Fluids*, 6(7):493–499.
- Åkervik, E., Brandt, L., Henningson, D. S., Höpfner, J., Marxen, O., and Schlatter, P. (2006). Steady solutions of the navier-stokes equations by selective frequency damping. *Physics of Fluids*, 18(6):068102.
- Åkervik, E., Höpfner, J., Ehrenstein, U., and Henningson, D. S. (2007). Optimal growth, model reduction and control in a separated boundary-layer flow using global eigenmodes. *Journal of Fluid Mechanics*, 579:305–314.
- Anderson, J. (1995). *Computational Fluid Dynamics, The Basics with Applications*. McGraw-Hill.
- Armaly, B. F., Durst, F., Pereira, J. C. F., and Schönung, B. (1983). Experimental and theoretical investigation of backward-facing step flow. *Journal of Fluid Mechanics*, 127:473 – 496.
- Badreddine, H., Vandewalle, S., and Meyers, J. (2014). Sequential quadratic programming (sqp) for optimal control in direct numerical simulation of turbulent flow. *Journal of Computational Physics*, 256:1 – 16.
- Bagheri, S. (2010). *Analysis and control of transitional shear flows using global modes*. PhD thesis, Department of Mechanics, Royal Institute of Technology, Stockholm.
- Bagheri, S. (2014). Effects of weak noise on oscillating flows: Linking quality factor, floquet modes, and koopman spectrum. *Phys. Fluids*, 26(9):094104.
- Barbagallo, A., Dergham, G., Sipp, D., Schmid, P. J., and Robinet, J.-C. (2012). Closed-loop control of unsteadiness over a rounded backward-facing step. *Journal of Fluid Mechanics*, 703:326–362.
- Barbagallo, A., Sipp, D., and Schmid, P. J. (2009). Closed-loop control of an open cavity flow using reduced-order models. *Journal of Fluid Mechanics*, 641:1–50.
- Barkley, D., Gomes, M. G. M., and Henderson, R. D. (2002). Three-dimensional instability in flow over a backward-facing step. *Journal of Fluid Mechanics*, 473:167 – 190.

References

- Basley, J., Pastur, L., Lusseyran, F., Faure, T. M., and Delprat, N. (2011). Experimental investigation of global structures in an incompressible cavity flow using time-resolved PIV. *Experiments in Fluids*, 50(4):905–918.
- Beaudoin, J., Cadot, O., Aider, J., and Wesfreid, J. E. (2004). Three-dimensional stationary flow over a backward-facing step. *European Journal of Mechanics-B/Fluids*, 23(1):147–155.
- Bergmann, M., Cordier, L., and Brancher, J. (2005). Optimal rotary control of the cylinder wake using proper orthogonal decomposition reduced-order model. *Physics of Fluids*, 17.
- Bewley, T. R., Moin, P., and Temam, R. (2001). DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms. *Journal of Fluid Mechanics*, 447:179–225.
- Bilanin, A. J. and Covert, E. (1973). Estimation of possible excitation frequencies for shallow rectangular cavities. *AIAA journal*, 11(3):347–351.
- Blackburn, H. M., Barkley, D., and Sherwin, S. J. (2008). Convective instability and transient growth in flow over a backward-facing step. *Journal of Fluid Mechanics*, 603:271–304.
- Blasius, H. (1908). Grenzschichten in Flüssigkeiten mit kleiner Reibung. *Z. Math. Phys.*, 56:1–37.
- Block, P. J. (1976). Noise response of cavities of varying dimensions at subsonic speeds. Technical report, NASA, TN D-8351.
- Blonigan, P. J. (2016). *Least Squares Shadowing for sensitivity analysis of large chaotic systems and fluid flows*. PhD thesis, Massachusetts Institute of Technology. Department of Aeronautics and Astronautics.
- Bogey, C., de Cacqueray, N., and Bailly, C. (2009). A shock-capturing methodology based on adaptative spatial filtering for high-order non-linear computations. *Journal of Computational Physics*, 228:1447 – 1465.
- Boujo, E. and Gallaire, F. (2014). Controlled reattachment in separated flows: a variational approach to recirculation length reduction. *Journal of Fluid Mechanics*, 742:618–635.
- Bremner, P. G. (2012). Vibroacoustic Source Mechanism under Aeroacoustic Loads. In *AIAA 2012-2206*, Colorado Springs, CO, USA.
- Bremner, P. G. and Wilby, J. F. (2002). Aero-Vibro-Acoustics: Problem Statement and Methods for Simulation-Based Design Solution. In *AIAA 2002-2551*, Breckenridge, Colorado.
- Bremner, P. G. and Zhu, M. (2003). Recent Progress using SEA and CFD to Predict Interior Wind Noise. In *SAE International 2003-01-1705*.
- Brès, G. A. and Colonius, T. (2008). Three-dimensional instabilities in compressible flow over open cavities. *Journal of Fluid Mechanics*, 599:309–339.

- Bunge, H., Richards, M. A., Lithgow-Bertelloni, C., Baumgardner, J. R., Grand, S. P., and Romanowicz, B. A. (1998). Time scales and heterogeneous structure in geodynamic earth models. *Science*, 280(5360):91–95.
- Carpenter, M. H., Nordström, J., and Gottlieb, D. (1999). A stable and conservative interface treatment of arbitrary spatial accuracy. *Journal of Computational Physics*, 148:341 – 365.
- Cattafesta, L. N., Song, Q., Williams, D. R., Rowley, C. W., and Alvi, F. S. (2008). Active control of flow-induced cavity oscillations. *Progress in Aerospace Sciences*, 44(7–8):479–502.
- Cerviño, L., Bewley, T., Freud, J., and Lele, S. (2002). Perturbation and adjoint analyses of flow-acoustic interactions in an unsteady 2D jet. In *Center for Turbulence Research. Proceedings of the Summer Program 2002*.
- Chang, K., Constantinescu, G., and Park, S. (2006). Analysis of the flow and mass transfer processes for the incompressible flow past an open cavity with a laminar and a fully turbulent incoming boundary layer. *Journal of Fluid Mechanics*, 561:113–145.
- Choi, H., Hinze, M., and Kunisch, K. (1999). Instantaneous control of backward-facing step flows. *Applied numerical mathematics*, 31(2):133–158.
- Chun, K. and Sung, H. (1996). Control of turbulent separated flow over a backward-facing step by local forcing. *Experiments in Fluids*, 21(6):417–426.
- Citro, V., Giannetti, F., Brandt, L., and Luchini, P. (2015). Linear three-dimensional global and asymptotic stability analysis of incompressible open cavity flow. *Journal of Fluid Mechanics*, 768:113–140.
- Cvitanović, P. (2013). Recurrent flows: the clockwork behind turbulence. *Journal of Fluid Mechanics*, 726:1–4.
- Cvitanović, P., Artuso, R., Mainieri, R., Tanner, G., and Vattay, G. (2016). *Chaos: Classical and Quantum*. ChaosBook.org, Niels Bohr Institute, Copenhagen.
- de Vicente, J., Basley, J., Meseguer-Garrido, F., Soria, J., and Theofilis, V. (2014). Three-dimensional instabilities over a rectangular open cavity: from linear stability analysis to experimentation. *Journal of Fluid Mechanics*, 748:189–220.
- Dinniman, M. S., Asay-Davis, X. S., Galton-Fenzi, B. K., Holland, P. R., Jenkins, A., and Timmermann, R. (2016). Modeling ice shelf/ocean interaction in antarctica: A review. *Oceanography*, 29(4):144–153.
- Douay, C. L., Lusseyran, F., and Pastur, L. R. (2016). The onset of centrifugal instability in an open cavity flow. *Fluid Dynamics Research*, 48(6):061410.

References

- Druault, P., Bouhoubeiny, E., and Germain, G. (2012). POD investigation of the unsteady turbulent boundary layer developing over porous moving flexible fishing net structure. *Experiments in Fluids*, 53(1):277–292.
- Druault, P., Hekmati, A., and Ricot, D. (2013). Discrimination of acoustic and turbulent components from aeroacoustic wall pressure field. *Journal of Sound and Vibration*, 332:7257 – 7278.
- Eaton, J. K. and Johnston, J. P. (1980). Turbulent flow re-attachment: an experimental study of the flow and structure behind a backward-facing step. Technical report, Stanford University, Rep. MD-39.
- Eyink, G. L., Haine, T. W. N., and Lea, D. J. (2004). Ruelle’s linear response formula, ensemble adjoint schemes and lévy flights. *Nonlinearity*, 17(5):1867.
- Farazmand, M. (2016). An adjoint-based approach for finding invariant solutions of Navier-Stokes equations. *Journal of Fluid Mechanics*, 795:278–312.
- Fichtner, A., Bunge, H., and Igel, H. (2006a). The adjoint method in seismology: I. theory. *Physics of the Earth and Planetary Interiors*, 157(1–2):86 – 104.
- Fichtner, A., Bunge, H., and Igel, H. (2006b). The adjoint method in seismology—: II. applications: traveltimes and sensitivity functionals. *Physics of the Earth and Planetary Interiors*, 157(1):105–123.
- Fosas de Pando, M., Sipp, D., and Schmid, P. J. (2012). Efficient evaluation of the direct and adjoint linearized dynamics from compressible flow solvers. *Journal of Computational Physics*, 231:7739–7755.
- Freund, J. B. (2011). Adjoint-based optimization for understanding and suppressing jet noise. *Journal of Sound and Vibration*, 330(17):4114 – 4122. {COMPUTATIONAL} AEROACOUSTICS (CAA) {FOR} {AIRCRAFT} {NOISE} {PREDICTION} – {PART} B.
- Frigo, M. and Johnson, S. (2005). The design and implementation of FFTW3. *Proceedings of the IEEE*, 93:216 – 231.
- Furuichi, N. and Kumada, M. (2002). An experimental study of a spanwise structure around a reattachment region of a two-dimensional backward-facing step. *Experiments in fluids*, 32(2):179–187.
- Gautier, N. and Aider, J. (2013). Control of the separated flow downstream of a backward-facing step using visual feedback. In *Proc. R. Soc. A*, volume 469, page 20130404. The Royal Society.
- Giannetti, F. and Luchini, P. (2007). Structural sensitivity of the first instability of the cylinder wake. *Journal of Fluid Mechanics*, 581:167–197.

- Giles, M. B. and Pierce, N. A. (2000). An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65:393–415.
- Graham, W. R. (1997). A comparison of models for the wavenumber-frequency spectrum of turbulent boundary layer pressures. *Journal of Sound and Vibration*, 206(4):541 – 565.
- Gunzburger, M. and Lee, H. C. (1994). Analysis, approximation and computation of a coupled solid/fluid temperature control problem. *Computational Methods and Applied Mechanical Engineering*, 118:133–152.
- H. Schlichting and K. Gersten (2000). *Boundary Layer Theory*. Springer, Berlin.
- Haruna, S., Nouzawa, T., Kamimoto, I., and Sato, H. (1990). An Experimental Analysis and Estimation of Aerodynamic Noise Using a Production Vehicle. In *SAE Technical Paper 900316*.
- Hasan, M. (1992). The flow over a backward-facing step under controlled perturbation: laminar separation. *Journal of Fluid Mechanics*, 238:73–96.
- Heenan, A. and Morrison, J. (1998). Passive control of pressure fluctuations generated by separated flow. *AIAA journal*, 36(6):1014–1022.
- Heimbach, P., Hill, C., and Giering, R. (2005). An efficient exact adjoint of the parallel MIT general circulation model, generated via automatic differentiation. *Future Generation Computer Systems*, 21(8):1356–1371.
- Heller, H. H., Holmes, D., and Covert, E. E. (1971). Flow-induced pressure oscillations in shallow cavities. *Journal of sound and Vibration*, 18(4):545IN9547–546IN10553.
- Henningson, D. S. and Åkervik, E. (2008). The use of global modes to understand transition and perform flow control. *Physics of Fluids*, 20:031302–15.
- Hernandez, V., Roman, J. E., and Vidal, V. (2005). SLEPc: A scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Trans. Math. Software*, 31(3):351–362.
- Holmes, P., Lumley, J. L., and Berkooz, G. (1996). *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*. Cambridge University Press.
- Huerre, P. and Monkewitz, P. A. (1990). Local and global instabilities in spatially developing flows. *Annu. Rev. Fluid Mech.*, 22:473 – 537.
- Hwang, Y. F. and Hambric, S. (2001). A numerical technique for computation of flow-induced vibration and noise. *The Journal of the Acoustical Society of America*, 110(5):2759.
- Jameson, A. (2003). Aerodynamic Shape Optimization Using the Adjoint Method. Technical report, von Karman Inst. Fluid Dyn., Brussels, Belgium.
- Jones, L. E. (2008). *Numerical studies of the flow around an airfoil at low Reynolds number*. PhD thesis, University of Southampton.

References

- Jović, S. (1998). Recovery of reattached turbulent shear layers. *Experimental Thermal and Fluid Science*, 17(1–2):57 – 62.
- Jović, S. and Driver, D. M. (1994). Backward-facing step measurements at low Reynolds number, $Re_h = 5000$. Technical Report NASA TM 108807, NASA.
- Kaiktis, L., Karniadakis, G. E., and Orszag, S. A. (1991). Onset of three-dimensionality, equilibria and early transition in flow over a backward-facing step. *Journal of Fluid Mechanics*, 231:505 – 538.
- Kaiktsis, L., Karniadakis, G. E., and Orszag, S. A. (1996). Unsteadiness and convective instabilities in two-dimensional flow over a backward-facing step. *Journal of Fluid Mechanics*, 321:157–187.
- Kang, S. and Choi, H. (2002). Suboptimal feedback control of turbulent flow over a backward-facing step. *Journal of Fluid Mechanics*, 463:201–227.
- Kawahara, G., Uhlmann, M., and van Veen, L. (2012). The significance of simple invariant solutions in turbulent flows. *Annual Review of Fluid Mechanics*, 44(1):203–225.
- Kegerise, M. A., Spina, E. F., Garg, S., and Cattafesta III, L. N. (2004). Mode-switching and nonlinear effects in compressible flow over a cavity. *Physics of Fluids (1994-present)*, 16(3):678–687.
- Keller, H. B. (1977). Numerical solution of bifurcation and nonlinear eigenvalue problems. *Applications of Bifurcation Theory*, page 359–84.
- Kennedy, C. A., Carpenter, M. H., and Lewis, R. M. (2000). Low-storage, explicit Runge-Kuta schemes for the compressible Navier-Stokes equations. *Applied Numerical Mathematics*, 35:177 – 219.
- Kennedy, C. A. and Gruber, A. (2008). Reduced aliasing formulations of the convective terms within the Navier-Stokes equations for a compressible fluid. *Journal of Computational Physics*, 227:1676 – 1700.
- Kim, J. and Bewley, T. R. (2007). A Linear Systems Approach to Flow Control. *Annu. Rev. Fluid Mech.*, 39:383–417.
- Kim, J., Bodony, D. J., and Freund, J. B. (2014). Adjoint-based control of loud events in a turbulent jet. *Journal of Fluid Mechanics*, 741:28–59.
- Kim, J. W. and Lee, D. J. (2000). Generalized characteristic boundary conditions for computational aeroacoustics. *AIAA*, 38:2040 – 2049.
- Kim, J. W. and Lee, D. J. (2003). Characteristic Interface Conditions for Multiblock High-Order Computation on Singular Structured Grid. *AIAA*, 41:2341 – 2348.

- Kolmogorov, A. N. (1991). The local structure of turbulence in incompressible viscous fluid for very large reynolds numbers. *Proceedings of the Royal Society of London A: Mathematical, Physical and Engineering Sciences*, 434(1890):9–13.
- Kopera, M. A. (2011). *Direct Numerical Simulation of Turbulent Flow over a Backward-Facing Step*. PhD thesis, University of Warwick.
- Kostas, J., Soria, J., and Chong, M. (2005). A comparison between snapshot pod analysis of PIV velocity and vorticity data. *Experiments in Fluids*, 38(2):146–160.
- Kumar, A. and Dhiman, A. K. (2012). Effect of a circular cylinder on separated forced convection at a backward-facing step. *International Journal of Thermal Sciences*, 52:176 – 185.
- Kundu, P. K., Cohen, I. M., and Dowling, D. R. (2012). *Fluid mechanics*. Waltham, MA Academic Press, 5th ed edition. Previous ed.: 2008.
- Lai, J., Yue, J., and Platzer, M. (2002). Control of backward-facing step flow using a flapping foil. *Experiments in Fluids*, 32(1):44–54.
- Lanzerstorfer, D. and Kuhlmann, H. C. (2012). Global stability of the two-dimensional flow over a backward-facing step. *Journal of Fluid Mechanics*, 693:1–27.
- Lawson, S. J. and Barakos, G. N. (2011). Review of numerical simulations for high-speed, turbulent cavity flows. *Progress in Aerospace Sciences*, 47(3):186–216.
- Le, H., Moin, P., and Kim, J. (1997). Direct Numerical Simulation of Turbulent Flow over a Backward-Facing Step. *Journal of Fluid Mechanics*, 330:349 – 374.
- Lea, D. J., Allen, M. R., and Haine, T. W. N. (2000). Sensitivity analysis of the climate of a chaotic system. *Tellus*, 52A:523–532.
- Lee, I. and Sung, J. (2000). Characteristics of wall pressure fluctuations in separated and reattaching flows over a backward-facing step: Part I. Time-mean statistics and cross-spectral analysis. *Experiments in Fluids*, 30(1):262 – 272.
- Lee, J. (2011). Model predictive control: Review of the three decades of development. *International Journal of Control, Automation and Systems*, 9(3):415–424.
- Lehoucq, R. B. and G., S. A. (2001). Large-scale eigenvalue calculations for stability analysis of steady flows on massively parallel computers. *Int. J. Numer. Methods Fluids*, 36:309 – 27.
- Lighthill, M. J. (1952). On Sound Generated Aerodynamically I. *Proceedings of the Royal Society of London. Series A*, 211:564 – 587.
- Lighthill, M. J. (1954). On Sound Generated Aerodynamically II. *Proceedings of the Royal Society of London. Series A*, 222:1 – 32.
- Liu, Q. (2016). *Global instability analysis and control of three-dimensional long open cavity flow*. PhD thesis, Technical University of Madrid (UPM).

References

- Liu, Q., Gómez, F., and Theofilis, V. (2016). Control of flow instabilities in an open aircraft bay model cavity. In *AIAA 2016-3317*, Washington, D.C.
- Luchini, P. and Bottaro, A. (2014). Adjoint Equations in Stability Analysis. *Annual Review of Fluid Mechanics*, 46:493–517.
- M. D. Gunzburger (2003). *Perspectives in Flow Control and Optimization*. SIAM, Ames, Iowa.
- Mader, C. and Martins, J. (2008). ADjoint: An Approach for the Rapid Development of Discrete Adjoint Solvers. *AIAA Journal*, 46:863 – 873.
- Mehrez, Z., Bouterra, M., El, C. A., Belghith, A., and Le, Q. P. (2011). Mass transfer control of a backward-facing step flow by local forcing-effect of reynolds number. *Thermal Science*, 15(2):367–378.
- Meliga, P., Sipp, D., and Chomaz, J.-M. (2010). Effect of compressibility on the global stability of axisymmetric wake flows. *Journal of Fluid Mechanics*, 660:499–526.
- Meseguer-Garrido, F., de Vicente, J., Valero, E., and Theofilis, V. (2014). On linear instability mechanisms in incompressible open cavity flow. *Journal of Fluid Mechanics*, 752:219–236.
- Meyer, K. E., Pedersen, J. M., and Özcan, O. (2007). A turbulent jet in crossflow analysed with proper orthogonal decomposition. *Journal of Fluid Mechanics*, 583:199–227.
- Mezić, I. (2005). Spectral properties of dynamical systems, model reduction and decompositions. *Nonlinear Dynamics*, 41(1-3):309–325.
- Mezić, I. (2013). Analysis of Fluid Flows via Spectral Properties of the Koopman Operator. *Annu. Rev. Fluid Mech.*, 45:57 –78.
- Miau, J., Chou, J., Lee, K., and Chen, M. (1991). Control of separated flow by a two-dimensional oscillating fence. *AIAA journal*, 29(7):1140–1148.
- Miles, J. W. (1958). On the disturbed motion of a plane vortex sheet. *Journal of Fluid Mechanics*, 4(05):538–552.
- Morales, J. L. and Nocedal, J. (2011). Remark on Algorithm 778: L-BFGS-B: Fortran Subroutines for Large-scale Bound Constrained Optimization. *ACM Transactions on Mathematical Software*, 38(1):7:1–7:4.
- Moret-Gabarro, L. (2009). *Aeroacoustic investigation and adjoint analysis of subsonic cavity flows*. PhD thesis, Université de Toulouse.
- Nagarajan, K. K., Cordier, L., and Airiau, C. (2013). Development and application of a reduced order model for the control of self-sustained instabilities in cavity flows. *Communications in Computational Physics*, 14:186–218.
- Nagata, M. (1990). Three-dimensional finite-amplitude solutions in plane couette flow: bifurcation from infinity. *J. Fluid Mech*, 217:519–527.

- Neumann, J. and Wengle, H. (2003). DNS and LES of passively controlled turbulent backward-facing step flow. *Flow, Turbulence and Combustion*, 71(1):297.
- Nocedal, J. and Wright, S. J. (2006). *Numerical Optimization*. Springer.
- Pavithran, S. and Redekopp, L. (1989). The absolute-convective transition in subsonic mixing layers. *Physics of Fluids A: Fluid Dynamics*, 1(10):1736–1739.
- Pereira, J. and Sousa, J. (1994). Influence of impingement edge geometry on cavity flow oscillations. *AIAA journal*, 32(8):1737–1740.
- Poinsot, T. J. and Lele, S. K. (1992). Boundary Conditions for Direct Simulations of Compressible Viscous Flows. *Journal of Computational Physics*, 101:104 – 129.
- Polak, E. and Ribiere, G. (1969). Note on the convergence of methods of conjugate directions. *Revue Francaise d'Informatique et de Recherche Operationnelle*, 3:35 – 43.
- Pope, S. (2000). *Turbulent Flows*. Cambridge University Press.
- Proctor, J. L., Brunton, S. L., and Kutz, J. N. (2014). Dynamic mode decomposition with control. *arXiv:1409.6358 [math.OC]*.
- Provansal, M., Mathis, C., and Boyer, L. (1987). Bénard-von kármán instability: transient and forced regimes. *Journal of Fluid Mechanics*, 182:1–22.
- Rossiter, J. (1964). Wind tunnel experiments on the flow over rectangular cavities at subsonic and transonic speeds. Technical report, Technical Report 64037. Royal Aircraft Establishment.
- Rowley, C., Juttijudata, V., and Williams, D. (2005). Cavity Flow Control Simulations and Experiments. In *AIAA 2005-292*, Reno, Nevada.
- Rowley, C. W. (2005). Model Reduction For Fluids Using Balanced Proper Orthogonal Decomposition. *Int. J. on Bifurcation and Chaos*, 13(3):997 – 1013.
- Rowley, C. W., Colonius, T., and Basu, A. J. (2002). On self-sustained oscillations in two-dimensional compressible flow over rectangular cavities. *Journal of Fluid Mechanics*, 455:315–346.
- Rowley, C. W., Mezić, I., Bagheri, S., Schlatter, P., and Henningson, D. S. (2009). Spectral analysis of nonlinear flows. *J. Fluid Mech.*, 641:115 – 127.
- Rowley, C. W. and Williams, D. R. (2006). Dynamics and control of high-reynolds-number flow over open cavities. *Annual Review of Fluid Mechanics*, 38(1):251–276.
- Rowley, C. W., Williams, D. R., Colonius, T., Murray, R. M., and Macmynowski, D. G. (2006). Linear models for control of cavity flow oscillations. *Journal of Fluid Mechanics*, 547:317–330.

References

- Ruisi, R., Zare-Behtash, H., Kontis, K., and Erfani, R. (2016). Active flow control over a backward-facing step using plasma actuation. *Acta Astronautica*, 126:354 – 363. Space Flight Safety.
- S. J. Illingworth, A. S. M. and Rowley, C. W. (2012). Feedback control of cavity flow oscillations using simple linear models. *Journal of Fluid Mechanics*, 709:223–248.
- Sampath, R. and Chakravarthy, S. (2014). Proper orthogonal and dynamic mode decompositions of time-resolved PIV of confined backward-facing step flow. *Experiments in Fluids*, 55(9).
- Sandberg, R. D. (2015). Compressible-Flow DNS with Application to Airfoil Noise. *Flow, Turbulence and Combustion*, 95:1–19.
- Sandberg, R. D. and Sandham, N. D. (2006). Nonreflecting Zonal Characteristic Boundary Condition for Direct Numerical Simulation of Aerodynamic Sound. *AIAA Journal*, 44:402 – 405.
- Sarohia, V. (1977). Experimental investigation of oscillations in flows over shallow cavities. *AIAA Journal*, 15(7):984–991.
- Schäfer, F., Breuer, M., and Durst, F. (2009). The dynamics of the transitional flow over a backward-facing step. *Journal of Fluid Mechanics*, 623:85 – 119.
- Schmid, P. and Henningson, D. (2000). *Stability and Transition in Shear Flows*. Applied Mathematical Sciences. Springer New York.
- Schmid, P. J. (2010). Dynamic mode decomposition of numerical and experimental data. *J. Fluid Mech.*, 656:5 – 28.
- Schmid, P. J. and Brandt, L. (2014). Analysis of Fluid Systems: Stability, Receptivity, Sensitivity. Lecture notes from the FLOW-NORDITA Summer School on Advanced Instability Methods for Complex Flows, Stockholm, Sweden, 2013. *Applied Mechanics Reviews*, 66:024803–024803–21.
- Schneider, R. (2006). *Applications of the Discrete Adjoint Method in Computational Fluid Dynamics*. PhD thesis, University of Leeds.
- Schulze, J., Schmid, P., and Sesterhenn, J. (2011). Iterative optimization based on an objective functional in frequency-space with application to jet-noise cancellation. *Journal of Computational Physics*, 230:6075 – 6098.
- Schulze, J. and Sesterhenn, J. (2011). Adjoint based noise minimization of a round supersonic jet. In *Journal of Physics: Conference Series*, volume 318, page 092005. IOP Publishing.
- Schulze, J. and Sesterhenn, J. (2013). Optimal distribution of porous media to reduce trailing edge noise. *Computers and Fluids*, 78:41 – 53.
- Selimefendigil, F. and Oztop, H. F. (2014). Control of laminar pulsating flow and heat transfer in backward-facing step by using a square obstacle. *Journal of Heat Transfer*, 136(8):081701.

- Selimefendigil, F. and Öztop, H. F. (2014). Effect of a rotating cylinder in forced convection of ferrofluid over a backward facing step. *International Journal of Heat and Mass Transfer*, 71:142 – 148.
- Silveira, A., Grand, D., Métais, O., and Lesieur, M. (1993). A numerical investigation of the coherent vortices in turbulence behind a backward-facing step. *Journal of Fluid Mechanics*, 256:1 – 25.
- Sipp, D. and Lebedev, A. (2007). Global stability of base and mean flows: a general approach and its applications to cylinder and open cavity flows. *Journal of Fluid Mechanics*, 593:333–358.
- Sirovich, L. (1987). Turbulence and the dynamics of coherent structures. parts i–iii. *Quarterly of Applied Mathematics*, 45:561–590.
- Smith, T. R. (2003). *Low-dimensional models of plane Couette flow using the proper orthogonal decomposition*. PhD thesis, Princeton University.
- Song, G., Alizard, F., Robinet, J. C., and Gloerfelt, X. (2013). Global and Koopman modes analysis of sound generation in mixing layers. *Physics of Fluids*, 25:124101–1 – 124101–29.
- Spagnoli, B. and Airiau, C. (2008). Adjoint analysis for noise control in a two-dimensional compressible mixing layer. *Computers and Fluids*, 37:475 – 486.
- Talagrand, O. and Courtier, P. (1987). Variational assimilation of meteorological observations with the adjoint vorticity equation. i: Theory. *Quarterly Journal of the Royal Meteorological Society*, 113(478):1311–1328.
- Tam, C. K. and Block, P. J. (1978). On the tones and pressure oscillations induced by flow over rectangular cavities. *Journal of Fluid mechanics*, 89(2):373–399.
- Theofilis, V. (2011). Global linear instability. *Annual Review of Fluid Mechanics*, 43(1):319–352.
- Theofilis, V. and Colonius, T. (2003). An algorithm for the recovery of 2- and 3-d global instabilities of compressible flow over 2-d open cavities. In *33 rd AIAA Fluid Dynamics Conference and Exhibit*.
- Thompson, K. W. (1987). Time Dependant Boundary Conditions for Hyperbolic Systems. *Journal of Computational Physics*, 68:1 – 34.
- Thompson, K. W. (1990). Time Dependant Boundary Conditions for Hyperbolic Systems, II. *Journal of Computational Physics*, 89:439 – 461.
- Tu, J. H., Rowley, C. W., and Luchtenburg, D. M. (2014). On dynamic mode decomposition: Theory and applications. *Journal of Computational Dynamics*, 1(2).
- Uruba, V., Jonáš, P., and Mazur, O. (2007). Control of a channel-flow behind a backward-facing step by suction/blowing. *International Journal of Heat and Fluid Flow*, 28(4):665–672.

References

- Van Herpe, F., Bordji, M., and Baresch, D. (2011). Wavenumber-Frequency Analysis of the Wall Pressure Fluctuations in the Wake of a Car Side Mirror. In *AIAA 2011-2936*, Portland, Oregon.
- Van Herpe, F. A. and Duarte, L. O. (2012). Sound vs. Pseudo-sound Contributions to the Wind Noise inside a Car Cabin: A Modal Approach. In *AIAA 2012-2207*, Colorado Springs, CO.
- Vishnampet, R., Bodony, D. J., and Freud, J. B. (2015). A practical discrete-adjoint method for high fidelity compressible turbulence simulations. *Journal of Computational Physics*, 285:173–192.
- von Terzi, D. A. (2004). *Numerical Investigation of Transitional and Turbulent Backward-Facing Step Flows*. PhD thesis, University of Arizona.
- Wang, Q. (2013). Forward and adjoint sensitivity computation of chaotic dynamical systems. *Journal of Computational Physics*, 235:1 – 13.
- Wang, Q., Hu, R., and Blonigan, P. (2014). Least squares shadowing sensitivity analysis of chaotic limit cycle oscillations. *Journal of Computational Physics*, 267:210 – 224.
- Wee, D., Yi, T., Annaswamy, A., and Ghoniem, A. F. (2004). Self-sustained oscillations and vortex shedding in backward-facing step flows: Simulation and linear instability analysis. *Physics of Fluids*, 16(9):3361–3373.
- Wei, M. and Freund, J. B. (2006). A noise-controlled free shear flow. *Journal of Fluid Mechanics*, 546:123–152.
- White, F. (1991). *Viscous Fluid Flow*. McGraw-Hill.
- Wolf, A., Swift, J. B., Swinney, H. L., and Vastano, J. A. (1985). Determining lyapunov exponents from a time series. *Physica D: Nonlinear Phenomena*, 16(3):285–317.
- Yamouni, S., Sipp, D., and Jacquin, L. (2013). Interaction between feedback aeroacoustic and acoustic resonance mechanisms in a cavity flow: a global stability analysis. *Journal of Fluid Mechanics*, 717:134–165.
- Yoshioka, S., Obi, S., and Masuda, S. (2001). Turbulence statistics of periodically perturbed separated flow over backward-facing step. *International Journal of Heat and Fluid Flow*, 22(4):393–401.