



HAL
open science

Tonal noise generation in flows around aerofoils: a global stability analysis

Miguel Ángel Fosas de Pando

► **To cite this version:**

Miguel Ángel Fosas de Pando. Tonal noise generation in flows around aerofoils: a global stability analysis. Fluids mechanics [physics.class-ph]. Ecole Polytechnique X, 2012. English. NNT: . tel-00816987

HAL Id: tel-00816987

<https://theses.hal.science/tel-00816987>

Submitted on 23 Apr 2013

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



École Polytechnique
Laboratoire d'Hydrodynamique
Thèse présentée pour obtenir le grade de
DOCTEUR DE L'ÉCOLE POLYTECHNIQUE

Spécialité : Mécanique

par

Miguel Ángel Fosas de Pando

Tonal noise generation in flows around aerofoils: a global stability analysis*

*Génération d'ondes acoustiques à fréquences discrètes autour d'un profil d'aile :
analyse de stabilité globale

Thèse soutenue le 19 décembre 2012 devant le jury composé de

Colm-cille Caulfield	examineur	University of Cambridge, BP Institute
Carlo Cossu	président du jury	IMFT, CNRS
Patrick Huerre	examineur	École Polytechnique, CNRS
Nigel Peake	rapporteur	University of Cambridge
Maurice Rossi	rapporteur	Université Pierre et Marie Curie, CNRS
Richard D. Sandberg	examineur	University of Southampton
Peter J. Schmid	directeur de thèse	École Polytechnique, CNRS
Denis Sipp	directeur de thèse	ONERA

TONAL NOISE GENERATION IN FLOWS AROUND AEROFOILS:
A GLOBAL STABILITY ANALYSIS

Miguel Ángel Fosas de Pando

TONAL NOISE GENERATION
IN FLOWS AROUND AEROFOILS
A GLOBAL STABILITY ANALYSIS

Laboratoire d'Hydrodynamique (LadHyX)
École Polytechnique, France

Contents

Contents	iii
List of Figures	vii
List of Tables	xi
Remerciements	xiii
1 Introduction	1
1.1 Noise on aerofoils	1
1.2 Tonal noise on aerofoils	1
1.3 Goals of this work	5
1.4 Outline of the thesis	6
2 Direct numerical simulation code	7
2.1 Introduction	7
2.1.1 Governing equations: the compressible Navier–Stokes equations	8
2.2 Curvilinear coordinates: fields and geometry	9
2.2.1 Curvilinear coordinates	9
2.2.2 Discretization	11
2.2.3 Implementation	12
2.3 Pseudo-characteristics formulation	13
2.3.1 Introduction of pseudo-characteristics	13
2.3.2 Summary of equations	15
2.3.3 Discussion	16
2.3.4 Discretization	17
2.3.5 Implementation	17
2.4 Boundary conditions	18
2.4.1 Preliminaries	18
2.4.2 Wall boundary conditions	20
2.4.3 Inflow boundary condition	23
2.4.4 Outflow boundary condition	24
2.4.5 Sponge layers	24
2.4.6 Implementation	24
2.5 Spatial differentiation	25
2.5.1 Finite differences (FD)	25
2.5.2 Properties of finite-difference schemes: modified wave number	26

2.5.3	Choice of spatial schemes	27
2.5.4	Implementation	29
2.6	Temporal advancement	34
2.6.1	Stability region	35
2.6.2	Implementation	36
2.7	Validation cases	36
2.7.1	Exact solutions of the linearized Euler equations	37
2.7.2	Sound generated by a co-rotating vortex pair	38
2.A	Spatial differentiation and temporal advancement schemes: coefficients	41
2.B	Numerical grid used for the validation cases	42
3	Nonlinear simulations	43
3.1	Problem description	43
3.2	Computational set-up	44
3.2.1	Physical and computational domains: boundaries	44
3.2.2	Grid generation	45
3.2.3	Boundary and initial conditions	48
3.2.4	Wake connection	48
3.3	Simulations	48
3.3.1	Configuration	48
3.3.2	Estimation of spatial scales	49
3.3.3	Numerical grids	50
3.3.4	Simulation time	52
3.4	Temporal evolution	53
3.5	Statistics	56
3.5.1	Convergence	56
3.5.2	Mean flow	58
3.5.3	Standard deviation	60
3.5.4	Trailing edge	60
3.6	Frequency analysis	62
3.7	Conclusions from the nonlinear simulations	65
4	Efficient evaluation of the linearized dynamics	67
4.1	Introduction	67
4.2	Direct and adjoint linear operators from a nonlinear code	70
4.2.1	Direct numerical simulation (DNS) code	70
4.2.2	Practical implementation using a modular code structure	71
4.2.3	The linearization step	72
4.2.4	The direct operator	73
4.2.5	The adjoint operator	74
4.3	A general framework for the evaluation of direct and adjoint operators	76
4.3.1	Generalization	76
4.3.2	Linearization	77
4.3.3	Direct and adjoint operator evaluation	77
4.3.4	Extensions	78
4.4	Applications	81
4.4.1	Implementation and performance	81

4.4.2	Spatially developing compressible boundary layer	83
4.4.3	Sound generation by an aerofoil	85
4.4.4	Lid-driven cavity flow	86
4.5	Summary and conclusions	87
4.A	Derivation of the adjoint	88
5	Global linear stability analysis	91
5.1	Introduction	91
5.1.1	Mathematical tools	92
5.1.2	Singular value decomposition: optimal perturbations and optimal forcings . .	94
5.1.3	Sensitivity of eigenvalues and singular values	95
5.1.4	Evaluation of the direct and adjoint operators	95
5.2	Krylov subspace methods	96
5.2.1	Vector basis	96
5.2.2	Matrix exponential: time integration	97
5.2.3	Ritz values: eigenvalues and eigenvectors	98
5.3	Choice of base flow, domain sizes and inner product	99
5.3.1	Base flow	99
5.3.2	Domain sizes	99
5.3.3	Inner product	99
5.4	Impulse response analysis	100
5.4.1	Evolution of the norm: estimation of non-normal effects	100
5.4.2	Evolution of the flow field	102
5.4.3	Summary	107
5.5	Global modes	108
5.5.1	Computation: eigenvalue problem set-up	108
5.5.2	Spectrum	109
5.5.3	Leading modes: the coupled dynamics of the separation bubbles	109
5.5.4	High-frequency modes: suction-surface shear-layer instabilities	121
5.5.5	Low-frequency: separation-bubbles flapping and reattached flow dynamics . .	124
5.5.6	Convection effects and free-stream dynamics	126
5.5.7	Robustness of the global mode representation: effect of the domain size	129
5.5.8	Conclusions from the global mode analysis	130
5.6	Frequency response	134
5.6.1	Singular value problem as an eigenvalue problem: numerical set-up	134
5.6.2	Optimal gain	136
5.6.3	Optimal forcing and optimal response	136
6	Conclusions	141
6.1	Summary of the thesis	141
6.2	Future perspectives	143
	Bibliography	145

List of Figures

1.1	Flow conditions producing aerofoil blade-self noise	2
1.2	Typical experimental set-up in open-jet wind tunnels	3
1.3	Diagram showing aerofoil flow process	4
1.4	Smoke visualization at the pressure-surface trailing edge	5
2.1	Cartesian and curvilinear coordinates	11
2.2	Grid partitioning and fields storage	13
2.3	Pseudo-waves in the momentum equation	16
2.4	Block diagram of the Navier–Stokes right-hand-side evaluation	18
2.5	Vector basis and boundary conditions	21
2.6	Spatial schemes: modified wavelength and dissipation	28
2.7	Spatial schemes: spectrum	30
2.8	Spatial schemes: parallelization and explicit schemes	32
2.9	Spatial schemes: parallelization and implicit schemes	33
2.10	Temporal scheme: stability region and stability limits	36
2.11	Temporal evolution of small-amplitude perturbations	39
2.12	Co-rotating vortex pair: comparison	40
2.13	Generation of numerical grids for the validation cases	42
3.1	Sketch of the problem under consideration	43
3.2	C-grid (physical domain) and unit square (computational domain)	44
3.3	Grid stretching function	47
3.4	Normal and tangential grid spacing	51
3.5	Numerical grid, domain sizes and boundary conditions	52
3.6	Instantaneous flow fields (vorticity and instantaneous dilatation)	53
3.7	Instantaneous flow fields (sequence snapshots)	55
3.8	Stream-wise velocity recorded at the far-field probe	56
3.9	Low frequency flapping showing the instantaneous vorticity fields	56
3.10	Estimation of the mean and the standard deviation of the velocity	57
3.11	Mean flow pressure and skin-friction coefficient distribution over the aerofoil	59
3.12	Boundary layer thickness	60
3.13	Mean vorticity levels and boundary layer velocity profiles	61
3.14	Standard deviation of mean flow stream-wise velocity	61
3.15	Mean-flow streamlines at the trailing edge and rms levels	62
3.16	Magnitude of the discrete Fourier transform of the pressure signal in the far-field	63

3.17	Fourier component of the dominant tones, showing the real part of the dilatation and vorticity fluctuations	64
3.18	Low-frequency flapping of the separation bubble	65
4.1	Block diagram for the evaluation of the right-hand side of the compressible Navier–Stokes equations	71
4.2	Block diagram demonstrating the linearization step for a selected module of the code . .	72
4.3	Block diagram of the fully linearized direct operation	74
4.4	Block diagram of the fully linearized adjoint operation	74
4.5	Differentiation module and its transconjugate	75
4.6	Relative residual for the linearized direct operator for the application cases	82
4.7	Compressible boundary-layer flow: global spectrum	84
4.8	Compressible boundary-layer flow: TS waves, free-stream and Orr modes	85
4.9	Compressible flow around an airfoil: leading direct and adjoint modes	86
4.10	Lid-driven square cavity: representative modes	87
5.1	Impulse response of the linearized direct operator: set-up	101
5.2	Impulse response of the linearized direct operator: temporal evolution of the amplitude of the initial disturbance	102
5.3	Impulse response: initial growth of disturbances in the vicinity of the leading edge	103
5.4	Impulse response: convective growth of instabilities on the pressure-surface and suction-surface boundary layers	104
5.5	Impulse response: feedback-loop	106
5.6	Impulse response: entropy signal at the suction-side and pressure-side probes	107
5.7	Global spectrum	110
5.8	Global spectrum: least-stable modes	111
5.9	Least-stable direct global modes: spatial structure	112
5.10	Least-stable direct global modes: evolution of the stream-wise velocity peak	113
5.11	Least-stable direct global modes: evolution of the pressure peak	113
5.12	Comparison between the leading global modes and the frequential snapshots	114
5.13	Least-stable adjoint global modes: spatial structure	118
5.14	Least-stable adjoint global modes: evolution of the adjoint stream-wise velocity peak . . .	119
5.15	Least stable global modes: wavemaker region	120
5.16	Least-stable global modes: evolution of the amplitude of the direct mode, adjoint mode and wavemaker	121
5.17	Global spectrum: high-frequency modes	122
5.18	High-frequency modes: spatial structure of the direct and adjoint modes	123
5.19	Global spectrum: low-frequency modes	124
5.20	Low-frequency modes: spatial structure of the direct and adjoint modes	125
5.21	Downstream advection effects	127
5.22	Downstream advection effects: exponential growth	128
5.23	Free-stream modes	128
5.24	Effect of the domain size on the global spectrum	129
5.25	Global spectrum: summary of families of modes	131
5.26	Families of direct global modes	132
5.27	Families of adjoint global modes	133
5.28	Evolution of the temporal residual for a random perturbation	135

5.29	Frequency response: amplitude gain	136
5.30	Frequency response: spatial structure of the optimal responses	137
5.31	Frequency response: spatial structure of the optimal forcings	138
5.32	Sensitivity of the optimal gain with respect to modifications of the linearized operator . .	139
6.1	Mechanism	143
6.2	Transfer-function model	143

List of Tables

2.1	Modified wave numbers	29
2.2	Butcher tableau with the coefficients for a general Runge–Kutta schemes	35
2.3	Butcher tableau for explicit Runge–Kutta scheme RK4(3)5[2R+]C	35
2.4	Coefficients for the CULD (compact upwind with low dissipation) scheme	41
2.5	Coefficients for the CULD (compact upwind with low dissipation) scheme (cont)	41
2.6	Coefficients for the centred compact scheme	41
2.7	Coefficients for the centred compact scheme (cont)	41
2.8	Runge–Kutta coefficients (cont)	41
2.9	Runge–Kutta coefficients (cont)	41
3.1	Details of the case under consideration	49
3.2	Numerical grids and simulation parameters	50
3.3	Numerical grids and simulation parameters (continuation)	51
3.4	Separation, reattachment points and maximum reverse flow	58
3.5	Values of selected frequency components and their amplitude in the pressure spectrum	65
4.1	Linearization technique benchmark	83
5.1	Size of the domains considered for the global analysis	99
5.2	Comparison between the tones in the acoustic spectrum and the frequency of the global modes with local maxima of temporal growth rate	115
5.3	Summary of the performed frequency responses	140

Remerciements

Je n’aurais jamais pensé que le plus difficile dans la rédaction de la thèse, soit de trouver les mots justes pour remercier tous ceux avec qui j’ai partagé ces années aussi riches en émotions.

Tout d’abord, je tiens à remercier mes directeurs de thèse Peter Schmid et Denis Sipp. Merci Denis pour m’avoir confié ce fascinant travail et pour ton optimisme. Merci Peter pour l’enthousiasme que tu m’as apporté tout au long de cette thèse et pour ton aide lors de moments difficiles. Je garderai toujours un bon souvenir de ta générosité, ta sympathie et ton sens de l’humour.

Je voudrais aussi remercier Nigel Peake et Maurice Rossi d’avoir accepté d’être rapporteurs ainsi que Carlo Cossu, Colm-cille Caulfield, Patrick Huerre et Richard D. Sandberg d’avoir participé au jury de thèse.

Ce travail n’aurait pas été possible sans la bonne ambiance et l’ouverture d’esprit qui règne au LadHyX, mais ce n’est pas un secret ! Merci à Jean-Marc Chomaz, Christophe Clanet ainsi qu’à tout les *LadHyXiens* que j’ai eu l’honneur de rencontrer. Je tiens à remercier également Laurent Jacquin et tout le personnel du DAFE de m’avoir accueilli chaleureusement à l’ONERA Meudon.

Je remercie ici en particulier mes collègues de bureau. Merci Xavier pour ton sens critique ainsi que tout ce que tu m’as appris tout de long de ces trois ans sur des sujets variés, comme par exemple les instabilité hydrodynamique ou la langue française. Merci aussi Cristóbal, Jon, Gianluca, et Fabien pour les bons moments tant au bureau comme en dehors du labo. Merci Daniel pour ton inestimable aide et ta capacité d’écoute ; j’ai eu beaucoup de plaisir à discuter avec toi sur des sujets très différents. Face au *trafic fortement perturbé*, venir vous voir au labo était souvent motivation suffisante pour attendre le RER B.

Durant ces années on m’a appris à faire de la recherche mais aussi à enseigner. Pour cela je souhaiterai remercier Sabine Ortiz et Thierry Schuller de m’avoir confié leurs élèves à l’ENSTA et à l’Ecole Centrale Paris.

Merci également à mes colocs pour leur amitié et leurs efforts d’apprendre l’espagnol : ¡Hola señor!—Franz, ¡Hola! ¿Qué tal?—Nicolas. Merci Franz pour les échanges lors du petit-déjeuner quotidien et le ravitaillement de café si nécessaire en fin de thèse. Merci Nicolas pour les diners ensemble, ou les brunchs du dimanche avec Sabine et Fabien ou encore les jeux de sociétés. Ce fut un véritable plaisir !

L’écriture de ces mots représente pour moi un évènement marquant d’un long voyage qui s’est commencé il y a quelque temps. Je suis très reconnaissant à Antonio Barrero et à José Manuel Gordillo de l’Escuela Superior de Ingenieros de la Universidad de Sevilla pour les cours que j’ai eu la chance de suivre et pour m’avoir transmis votre passion pour la recherche.

Je voudrais saluer mes amis Isra et Juanpe. Avec eux, j’ai fréquemment eu l’occasion de partager les inquiétudes du thésard, aussi bien que l’émotion qu’on éprouve quand un calcul marche, ou quand on voit le soleil à Paris. Pour les rigolades ensemble et pour votre amitié je vous dit “Merci !”. Bonne courage avec vos thèses ! Qui sait, un jour on se retrouvera au *Caleta Institute of Technology* !

J'ai aussi été accompagné par d'autres *españoles en el mundo*. Sans vous ce n'aurait pas été pareil. C'est pour cela que remercie ici ceux que j'ai rencontrés d'abord à Toulouse (Enric et Francesco) et ensuite à Paris (Sonia, David, Jessy, Bea, Cath et Alba). Merci pour les soirées, les pique-niques et les voyages ensembles. Merci également à mes amis de la Fondation Suisse et le Colegio de España.

Toda mi vida he tenido la suerte de contar con el cariño y el apoyo de mis seres más queridos. Por ello les doy las gracias a todos. Gracias especialmente a mis padres, por todo lo que me enseñaron, a mis hermanos y a mi abuelo, por estar cada día ahí, y a la familia de Cris, por recibirme siempre con los brazos abiertos.

Las últimas palabras son para tí, Cris; gracias por tu amor, por creer en mí y por estar siempre a mi lado pese a la distancia. Sin tí no sería lo que soy.

Introduction

1.1 Noise on aerofoils

When an obstacle is immersed in a fluid with relative motion, the obstacle disrupts its course, and as a consequence, the flow field in the vicinity of the obstacle is modified. As it passes the body, depending on the relative speed, the shape of the obstacle, and its orientation, the flow displays different degrees of unsteadiness which are perceived even at far distances: hydrodynamic fluctuations are noticed in the wake, and acoustic radiation is detected in the far-field. The understanding of the physical mechanisms involved in the generation of these disturbances is of crucial interest for engineering applications, and more precisely for the transport and the energy industries, as high levels of both hydrodynamic and acoustic fluctuations are commonly linked to performance losses in technical devices; in addition, they negatively impact the surrounding environment. In the case of aerofoil sections—even if they are comparatively less intrusive than other types of obstacles—the unsteadiness of the perturbed flow leads to substantial levels of noise radiation.

Brooks et al. (1989) introduced the term aerofoil *self-noise*, to refer to the ensemble of physical mechanisms, depicted in figure 1.1, that may account for noise radiation in isolated aerofoil flows. The term *self-noise* points at the aerofoil and the disturbances that it introduces in the flow, namely boundary layers, separation regions and the wake, as the sources of noise. These mechanisms must not be mistaken with leading edge noise, which is related to the interaction between incoming free-stream turbulence and the aerofoil surface.

The focus of the present study is the noise generation mechanism under circumstances where the aerofoil flow consists of laminar boundary layers supporting instability waves; see figure 1.1b. Typical values for the non-dimensional flow parameters are in the range of moderate Reynolds numbers $10^5 \leq Re \leq 10^6$. In this case, the acoustic spectrum is then dominated by strong tones that rise above the background noise. This phenomenon is typically referred to as *tonal noise*, *boundary-layer instability noise* or *laminar boundary layer–vortex shedding noise*, and it is commonly observed on wind-tunnel test models, sailplanes and wind turbines.

1.2 Tonal noise on aerofoils

The work of Paterson et al. (1973) is commonly considered to be the first experimental study that specifically addressed the tonal noise phenomenon on aerofoils. In their analysis, they characterized

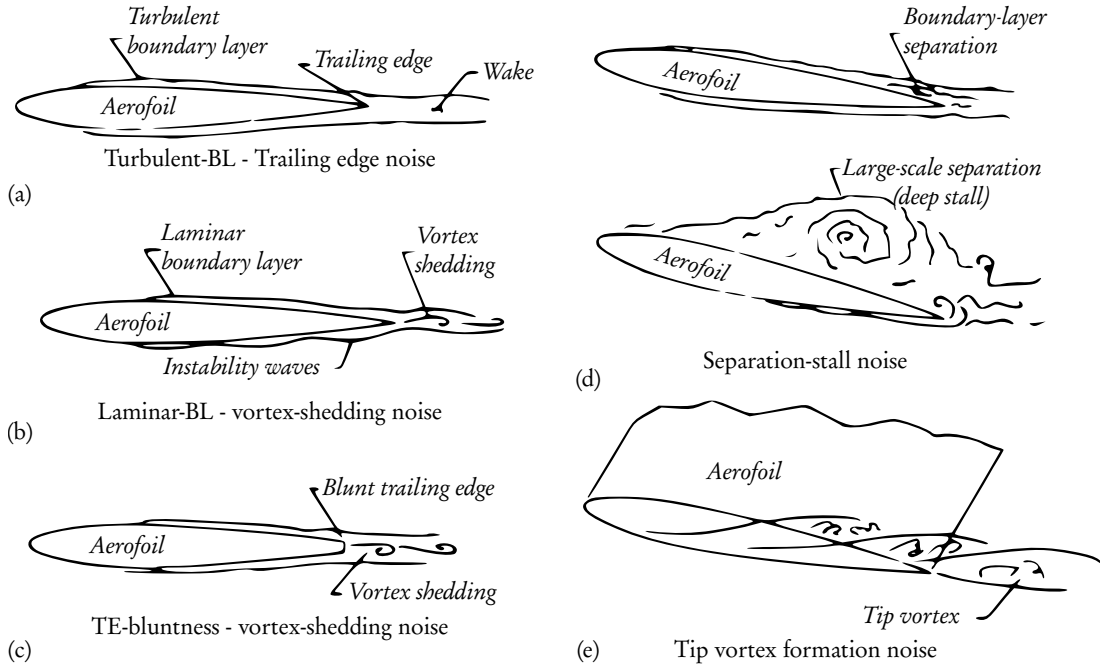


Figure 1.1: Flow conditions producing aerofoil blade-self noise (original caption, figure adapted from Brooks et al. (1989)).

the flow regimes where strong acoustic tones were present in the sound spectrum. A significant number of theoretical (Fink, 1974; Tam, 1974; Kingan and Pearse, 2009), experimental (Wright, 1976; Longhouse, 1977; Arbey and Bataille, 1983; Brooks et al., 1989; McAlpine et al., 1999; Nash et al., 1999; Nakano et al., 2006; Takagi et al., 2006; Arcondoulis et al., 2009; Chong and Joseph, 2012), and, more recently, numerical investigations (Desquesnes et al., 2007; Jones et al., 2008; Le Garrec, 2008; Sandberg et al., 2009; Jones et al., 2010; Jones and Sandberg, 2011; Tam and Ju, 2011) have followed the initial work of Paterson et al. (1973). However, a complete and widely accepted description of the physical mechanisms for the occurrence of discrete tones is still missing. In view of the above studies, we provide here a brief summary of the current understanding of the tonal noise generation mechanism. For a presentation in chronological order, the reader is referred to Desquesnes et al. (2007).

In figure 1.2a, we depict a typical experimental set-up consisting of an open-jet wind tunnel. Common aerofoil sections have chord-lengths in the range $8\text{cm} \leq c \leq 30\text{cm}$, small angles of incidence $0^\circ \leq \alpha \leq 8^\circ$ and are subjected to flow speeds in the range $10\text{ms}^{-1} \leq U \leq 50\text{ms}^{-1}$. At typical values of atmospheric pressure and air kinematic viscosity, the corresponding values for the Reynolds number $\text{Re} = Uc/\nu$ and the Mach number $M = a/U$ are in the ranges $10^5 \leq \text{Re} \leq 2 \cdot 10^6$ and $M < 0.15$.

Under such circumstances, the spectrum of a pressure probe in the far field (figure 1.2b) shows the appearance of a set of strong, equally-spaced peaks, denoted by $f_{n \max}$ and f_n , above the wind-tunnel background noise that are commonly perceived by the human ear as a strong whistle. For small variations in the free-stream velocity, the frequency of the largest amplitude peak $f_{n \max}$ scales with $U^{0.8}$. However, at given velocities, secondary peaks f_n rise in amplitude and overtake the previous maximum. Since the tones are discrete, the effect is a sudden jump in the main frequency,

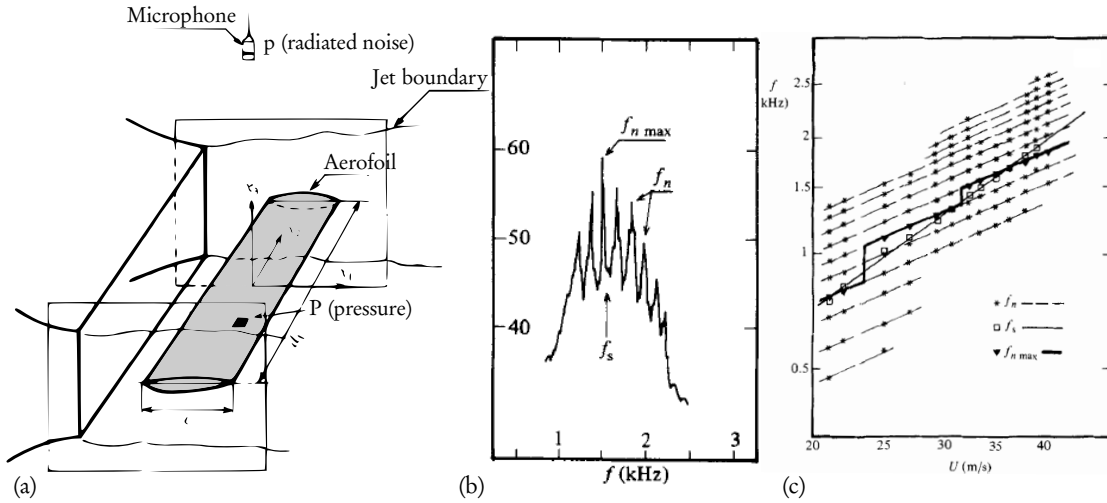


Figure 1.2: (a) Typical experimental set-up in open-jet wind tunnels: (b) the spectrum showing multiple equally spaced frequencies and (c) a ladder structure showing the evolution of the highest peak-amplitude for varying free-stream velocity (original figures from Arbey and Bataille (1983)).

which amounts to a *ladder* structure illustrated in figure 1.2c. The overall trend for the maximum frequency evolution $f_{n \max}$ is $U^{1/2}$. The acoustic levels of the discrete tones increase with the velocity, then reach a maximum, until they become imperceptible as the boundary layers undergo transition to turbulence. In Paterson et al. (1973), an empirical law for the overall evolution of the main tone was derived: $f = KU^{3/2}(c\nu)^{-1/2}$, where $K = 0.011$, ν is the kinematic viscosity, and c is the aerofoil chord. In terms of the Reynolds number Re and the Strouhal number $St = fc/U$, the above law reads $St = KRe^{1/2}$. Pressure measurements indicate that the acoustic radiation originates as a dipolar acoustic source located in the vicinity of the trailing edge.

The occurrence of multiple equally-spaced peaks following a ladder evolution has been observed experimentally by Arbey and Bataille (1983); Nakano et al. (2006); Arcondoulis et al. (2009); Takagi and Konishi (2010); Chong and Joseph (2012). In numerical studies, multiple peaks are commonly present in the pressure spectrum (Desquesnes et al., 2007; Le Garrec, 2008; Jones and Sandberg, 2011); however, the ladder structure has not been investigated numerically as the calculations are computationally expensive. In contrast, the experiments carried out by Nash et al. (1999) and the simulations of Tam and Ju (2011) displayed a single peak. Interestingly, the experimental cases that featured multiple peaks are correlated to experiments carried out in open-jet wind tunnels (figure 1.2a), whereas the case where only one peak is observed, corresponds to an experiment in a closed-section wind tunnel. In Le Garrec (2008), it is shown by means of numerical simulations that the observation of a single tone can possibly be related to the effect of the wind-tunnel walls on the potential flow.

It is widely accepted that the acoustic features presented above are strongly linked to flow configurations where the pressure surface displays a laminar boundary layer: the tones in the spectrum disappear as the boundary layer undergoes transition to turbulence, either naturally or induced by a trip wire. On the contrary, the boundary layer on the suction surface is usually turbulent, and the main features of the tonal noise remain insensitive to changes on the suction-surface boundary layer. In figure 1.3, typical features of the flow field are sketched.

For common ranges of Reynolds numbers and angles of attack, owing to the adverse pressure




Figure 1.3: Diagram showing aerofoil flow process (original caption, figure adapted from Nash et al. (1999)).

gradient downstream of the maximum thickness of the aerofoil, the boundary layer detaches over the last quarter of the chord, leading to a large laminar separation bubble (Longhouse, 1977; Nash et al., 1999; Desquesnes et al., 2007) that extends up to the trailing edge. Furthermore, the frequencies of the tones fall within the range of frequencies at which the pressure-surface boundary layer can sustain instability waves.

This latter fact, that the tone of the dominant frequency in the spectrum falls within the range of frequencies where instability waves are amplified along the laminar boundary layer, has been confirmed by most studies of tonal noise. Furthermore, by means of spatial stability theory, it is observed that the frequency that is most amplified along the pressure surface corresponds, within reasonable error bounds, to the frequency of the acoustic tones, thus suggesting the fundamentally important role of pressure-surface layer instabilities.

However, a physical mechanism based solely on boundary-layer instabilities cannot explain a frequency-selection process with multiple discrete peaks and, more importantly, cannot account for self-sustained oscillations in the absence of external noise: despite the separation bubbles on the pressure-surface, the reverse flow is typically small, approximately 10% of the free-stream speed, and it is thus unlikely that absolute instabilities are present (Hammond and Redekopp, 1998). Both experimental (Nash et al., 1999) and numerical studies (Jones et al., 2010; Tam and Ju, 2011) have confirmed that the boundary layers and the near wake do not display absolute instabilities.

Based on these observations, different types of *ad-hoc* mechanisms have been suggested in the last four decades, involving mostly aeroacoustic feedback loops (Tam, 1974; Wright, 1976; Arbey and Bataille, 1983; Desquesnes et al., 2007; Chong and Joseph, 2012, among others). The proposed mechanisms involve convective instabilities on the boundary layers and separated-flow regions, amplifying incoming disturbances and producing, as a result of the trailing edge geometry, significant acoustic scattering. The acoustic waves that further propagate in the far-field would excite in turn the boundary layers upstream, reinforcing discrete frequency components following an integer wave-number phase condition.

The above observations indicate that the tonal noise-generation mechanism is intrinsically global, as it involves instability mechanisms of different flow phenomena coupled with boundary-layer receptivity to free-stream disturbances. Indeed, flow receptivity is likely to take place at regions in the flow with significant spatial inhomogeneities, such as separation points or abrupt changes in the aerofoil geometry, where an energy transfer mechanism between free-stream disturbances and boundary layer instabilities takes place. For a general review of boundary-layer receptivity theory, the reader is referred to Saric et al. (2002).

Nonetheless, there is debate about the existence of aeroacoustic feedback loops (Nash et al., 1999; Tam and Ju, 2011) and their importance (Jones and Sandberg, 2011). For this reason, Jones

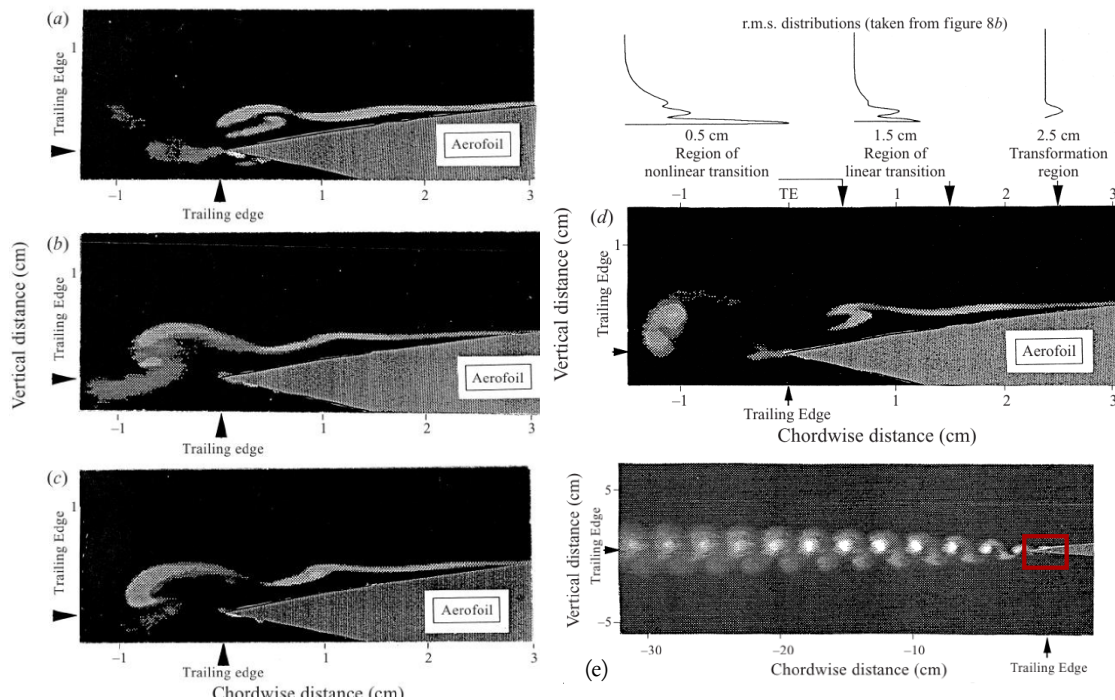


Figure 1.4: Smoke visualization at the pressure-surface trailing edge. (a) Instability structures begin to roll up near the trailing edge, (b) rolls up into a vortex at the trailing edge and a second instability begins to form. (d) A vortex is shed from the trailing edge, as the second instability rolls up into a vortex. (e) Smoke visualization of the wake (original caption, figure adapted from Nash et al. (1999)).

et al. (2010); Jones and Sandberg (2011) investigated by means of direct numerical simulations the evolution of a small perturbations in aerofoil flows; it was found that aeroacoustic feedback loops are indeed present and that they may play a significant role in the frequency selection.

Flow visualization shows that the flow organizes at the trailing edge into coherent structures (see figure 1.4), displays the ringing of strongly non-parallel hydrodynamic structures, and involves the dynamics of the separated boundary layers including vortex shedding at the trailing edge and acoustic radiation.

It appears fair to state that a clear analysis of the instability processes as well as the role of the different flow features during the tonal noise-generation process has been limited in the past by the narrow scope and shortcomings of classical local stability theory and the use of this tool to address non-parallel flows and long-distance feedback interactions.

1.3 Goals of this work

In view of the previous work, it is suggested that the occurrence of discrete tones in flow around aerofoils arises from a complex interaction between hydrodynamic instabilities in the boundary layers and the near wake, with acoustic radiation at the trailing edge.

To address this problem, global and non-modal analyses (Schmid and Henningson, 2001; Chomaz, 2005; Schmid, 2007) constitute a promising theoretical framework to address complex flow problems such as the one presented here: the non-parallelism of the flow field is considered without further

simplifications, the receptivity of the flow can be analysed by using adjoint techniques (Hill, 1995; Giannetti and Luchini, 2006) and the effect of non-normality of the underlying equations can directly be investigated. Furthermore, the identification of regions in the flow responsible for the global dynamics can be addressed by a structural sensitivity analysis (Hill, 1992; Bottaro et al., 2003; Chomaz, 2005; Giannetti and Luchini, 2006; Marquet et al., 2008). This principally motivates the application of global stability theory to the flow under investigation.

It should be mentioned that there exists a great variety of flow regimes that produce discrete tones in flows around aerofoils. In this work, we focus on a flow regime similar to the one analysed in the numerical investigation of Desquesnes et al. (2007), since their study is closely aligned with the experiments performed by Nash et al. (1999). The goal of the present thesis is to gain further insight into the coupled dynamics of the boundary-layers on both surfaces, the near wake and the subsequent acoustic radiation. More precisely, we address the question whether multiple tones can be explained using global stability theory and how this type of analysis can capture the dominant observed structures and describe long-distance global effects.

The importance of applying global stability theory to this problem is twofold: more specifically, it can provide a more complete understanding of the tonal-noise phenomenon, but more generally and more importantly, it serves as an illustration (and testbed) of this type of analysis in view of more complex flow configurations exhibiting similarly coupled hydrodynamic and acoustic features, such as the instabilities arising in turbomachinery and airframe flows.

1.4 Outline of the thesis

The present work is structured as follows. In chapter 2, we develop and outline a direct numerical simulation code that will serve as a powerful and flexible tool for the simulation of compressible flows in complex domains. This numerical code is based on previous experience with a direct numerical simulation code available in our group. In chapter 3, the numerical set-up and the specifics of the flow case are described. The nonlinear simulations are discussed and compared to the case described in Desquesnes et al. (2007).

In order to carry out global stability analyses, access to the linearized direct and adjoint operators must be established. In the case of compressible flow solvers, the associated derivation is challenging, as attention must be paid to computational efficiency. For this reason, we have developed a procedure, presented in chapter 4, that allows the evaluation of the direct and adjoint linearized dynamics efficiently from the nonlinear numerical code presented in chapter 2.

Once all the necessary tools are available, we perform a global analysis, presented in chapter 5, featuring an impulse response analysis, the computation and discussion of the global spectrum as well as the frequency response of the flow to optimal forcings. Concluding remarks and future perspectives are offered in chapter 6.

Direct numerical simulation code

2.1 Introduction

The physics of sound generation around an aerofoil features complex interaction between hydrodynamic and acoustic phenomena. The ultimate goal of the present study is to analyze such features via nonlinear and linearized simulations of the compressible Navier–Stokes equations. In a first stage, a computational code for the numerical integration of the nonlinear equations has been implemented and validated.

At the beginning of this work, the author had access to a computational code that has been previously used for the simulation of high-speed compressible flows (see for instance Mack, 2009). Unfortunately, important limitations arose in its adaptation for the use in the present study, concerning in particular the computational efficiency, the maintainability of the source code, the handling of more complex geometries and the implemented boundary conditions.

In order to overcome such limitations, a new numerical code has been developed during the course of this thesis, using the same theoretical approach but featuring a completely new implementation.

In this chapter we give a brief but self-contained description of both the theoretical foundations and the practical implementation of the numerical methods included in our direct numerical simulation (DNS) code. The presentation of theoretical and practical aspects is carried out in parallel in order to highlight the modularity of the code and nontrivial issues of its implementation. Nevertheless, the sections containing implementation details may be safely omitted according to the interests of the reader.

From a theoretical point of view, our code computes the approximate *temporal evolution* of a discretized flow *field* subjected to the *compressible Navier–Stokes equations* with prescribed *boundary conditions*. Spatial derivatives are evaluated using *finite differences* on *curvilinear multi-block structured grids*.

From an implementation point of view, our numerical code consists of several numerical libraries (using an object-oriented design) that implement the above *abstract objects* in terms of classes. The classes rely extensively on template techniques to allow maximum code reuse without hindering specialization¹. For instance, setting up a simulation for a given flow case requires writing a custom application that uses the above libraries for its purposes. The code is written in C++ and has been

¹In this respect, special care has been taken in the design of the template classes, in order to avoid significant performance penalties associated with the abstraction overhead in portions of the code where high-performance is crucial.

parallelized using the Message Passing Interface (MPI) library v2.2 (Message Passing Interface Forum, 2009). The input/output operations are performed using the HDF5 library v1.8.9 (The HDF Group, 2012).

2.1.1 Governing equations: the compressible Navier–Stokes equations

The compressible Navier–Stokes equations are commonly presented as a set of partial differential equations describing the temporal evolution of the density ρ , velocity field \mathbf{u} and internal energy e , which are obtained from the application of the mass, momentum and energy conservation principles to a fluid particle. These equations read

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

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

and

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

where \mathbf{f}_v is a volume forcing, and Q_r is a heat source. The above equations are augmented by the constitutive relations for the viscous stress tensor $\boldsymbol{\tau}$ and the heat flux \mathbf{q} given by

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

and

$$\mathbf{q} = -k \nabla T, \quad (2.5)$$

together with the state equations for a calorically perfect gas:

$$p = \rho r T \quad \text{and} \quad e = \frac{r}{\gamma - 1} T. \quad (2.6)$$

Note that μ is the viscosity, μ_v is the volume viscosity, k is the heat conductivity, r is the gas constant and γ is the heat capacity ratio.

The above equations form a closed system of partial differential equations, first order in time and second order in space, that can be solved provided that a well-posed set of initial and boundary conditions is given. For the present purposes, the continuity and energy conservation laws need to be expressed in terms of the pressure p and entropy s , instead. Using the fundamental thermodynamic relation $de = T ds - p d(1/\rho)$ and the definition of the speed of sound $a^2 = (\partial p / \partial \rho)_s$, we have

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

and

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

Furthermore, the following non-dimensional variables (denoted by the superscript $'$) are introduced:

$$\mathbf{x}' = \frac{\mathbf{x}}{x_0}, \quad t' = \frac{u_0 t}{x_0}, \quad \mathbf{u}' = \frac{\mathbf{u}}{u_0}, \quad p' = \frac{p}{p_0}, \quad s' = \frac{s}{r}, \quad \rho' = \frac{\rho}{\rho_0}, \quad \mu' = \frac{\mu}{\mu_0}, \quad \text{and} \quad k' = \frac{k}{k_0}. \quad (2.9)$$

The choice of the characteristic scales (denoted by the subscript $_0$) is usually motivated by the initial and/or boundary conditions, such as the length scale of an obstacle, the free-stream values of fluid variables, or the reference value of physical properties of the fluid.

Finally, our governing equations are summarized in equations (2.10) to (2.12), where the primes have been dropped. The introduced non-dimensional parameters, namely Reynolds number Re , Prandtl number Pr and Mach number M , are given in equation (2.13).

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

$$\rho \left(\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} \right) = -\frac{1}{\gamma \text{M}^2} \nabla p + \frac{1}{\text{Re}} \nabla \cdot \boldsymbol{\tau} + \mathbf{f}_v, \quad (2.10b)$$

$$p \left(\frac{\partial s}{\partial t} + \mathbf{u} \cdot \nabla s \right) = \frac{\gamma \text{M}^2}{\text{Re}} \boldsymbol{\tau} : \nabla \mathbf{u} - \frac{\gamma}{(\gamma - 1) \text{PrRe}} \nabla \cdot \mathbf{q} + Q_r, \quad (2.10c)$$

$$\boldsymbol{\tau} = \mu (\nabla \mathbf{u} + \nabla \mathbf{u}^t) + \left(\mu_v - \frac{2}{3} \mu \right) (\nabla \cdot \mathbf{u}) \mathbf{I}, \quad (2.11a)$$

$$\mathbf{q} = -k \nabla T, \quad (2.11b)$$

$$p = \rho T, \quad (2.12a) \quad (\gamma - 1) s = \log \frac{p}{\rho^\gamma}, \quad (2.12b) \quad a^2 = \frac{p}{\rho}, \quad (2.12c)$$

$$\text{Re} = \frac{\rho_0 u_0 x_0}{\mu_0}, \quad (2.13a) \quad \text{Pr} = \frac{\gamma r}{(\gamma - 1)} \frac{\mu_0}{k_0}, \quad (2.13b) \quad \text{M} = \frac{u_0}{a_0}, \quad (2.13c)$$

2.2 Curvilinear coordinates: fields and geometry

2.2.1 Curvilinear coordinates

We are interested in simulating the flow in a physical domain \mathcal{D} . Although in open flows the domain is formally unbounded, it is common practice to consider a sufficiently large bounded domain, in the hope that its particular extent will not affect the observations. The flow at a location given by its Cartesian coordinates \mathbf{x} is described by scalar, vector and tensor fields, such as the previously defined pressure $p(\mathbf{x})$, velocity $\mathbf{u}(\mathbf{x})$ and the viscous stress tensor $\boldsymbol{\tau}(\mathbf{x})$ fields. In an effort to alleviate the notation, the dependence on the spatial coordinates will be omitted when there is no ambiguity in the interpretation.

Except for the simplest geometries, the Cartesian coordinates \mathbf{x} are not amenable for practical calculations. Let us consider for instance the aerofoil surface: in such a coordinate system, the spatial

derivatives and the boundary conditions are cumbersome to write, since the boundary of the domain is usually given by a somewhat complicated set of equations. In contrast, it is preferred to define curvilinear coordinates ξ , tailored to the particular geometry, where the calculations are simplified. This usually amounts to choosing a *boundary conforming* coordinate system.

Given the boundaries of the physical domain $\partial\mathcal{D}$, we shall assume that a time-independent boundary-conforming coordinate system can readily be defined via the bijective transformation $\mathbf{x}(\xi)$ that maps the physical domain \mathcal{D} into a computational domain Ω and vice versa. The fields in the latter domain are recovered by using the transformation $f(\mathbf{x}(\xi)) = f(\xi)$.

For instance, a transformation may be chosen so that the boundary $\partial\Omega$ corresponds to the boundaries of the unit cube given by $0 \leq \xi^k \leq 1$. On the other hand, expressing equations (2.10) and (2.11) in the domain Ω requires the introduction of a more complex induced metric, that nevertheless can be dealt with systematically.

In the following, the index notation is used, and the summation convention for repeated indices applies unless they are parenthesized. It is useful to introduce the Jacobian of the transformation j_{ik} and its inverse k^{ki} . Hence

$$j_{ik} = \frac{\partial x_i}{\partial \xi^k}, \quad (2.14a)$$

$$k^{ki} = \frac{\partial \xi^k}{\partial x_i}. \quad (2.14b)$$

In addition to the Cartesian basis \mathbf{x}_i , two vector bases are defined: the contravariant basis ξ_k and the covariant basis ξ^k given by

$$\xi_k = j_{ik} \mathbf{x}_i, \quad (2.15a)$$

$$\xi^k = k^{ki} \mathbf{x}_i. \quad (2.15b)$$

Here we use the indices k and l for the curvilinear system, while the indices i and j are reserved for the Cartesian counterpart. It is also helpful to define the covariant metric tensor g_{kl} and the contravariant metric tensor g^{kl} as

$$g_{kl} = \xi_k \cdot \xi_l, \quad (2.16a)$$

$$g^{kl} = \xi^k \cdot \xi^l, \quad (2.16b)$$

as well as the corresponding scale factors h_k and h^k given by

$$h_k = \sqrt{g_{(kk)}}, \quad (2.17a)$$

$$h^k = \sqrt{g^{(kk)}}. \quad (2.17b)$$

A vector \mathbf{a} may equally be represented in the Cartesian basis $a_i \mathbf{x}_i$, contravariant basis $a^k \xi_k$ or covariant basis $a_k \xi^k$. From the above definitions, it can be shown that the following relations hold:

$$a^k = k^{ki} a_i, \quad (2.18a) \quad a_k = j_{ik} a_i, \quad (2.18b) \quad a^k = g^{kl} a_l, \quad (2.18c) \quad a_k = g_{kl} a^l. \quad (2.18d)$$

In addition, note that in curvilinear coordinates the operator ∇ reads

$$\nabla = \xi^k \frac{\partial}{\partial \xi^k}. \quad (2.19)$$

An illustration of the different vector bases is shown in figure 2.1. Geometrically, the direction of the vectors ξ_k is given by the intersection of the tangent planes with the constant- ξ^l surfaces ($l \neq k$), the direction of the vectors ξ^k is orthogonal to the corresponding constant- ξ^k plane.




Figure 2.1: Curvilinear coordinates and vector bases: two-dimensional representation showing the physical domain \mathcal{D} and the computational domain Ω . The domain corresponds to $0 \leq \xi^k \leq 1$. Two bi-orthogonal vector bases are defined: the contravariant vectors basis ξ^k , contained in the intersection of the constant- ξ^l planes ($l \neq k$), and the covariant vectors basis ξ_k , normal to the corresponding constant- ξ^k plane. A vector \mathbf{a} can be represented in any of the vector bases using equation (2.18). Note that although the vectors ξ_1 and ξ_2 are perpendicular in the drawing (b), in general we have $\xi_1 \cdot \xi_2 \neq 0$.

2.2.2 Discretization

The fields are discretized in the computational domain using an equally-spaced rectilinear grid.

Let $\mathbf{a}(\xi)$ be a scalar, vector or tensor field with a total of P components, then its discrete representation is given by the values of each component of the field at every grid point. For one-dimensional fields, the discretized field reads

$$a_{p i} = a_p \left(\frac{i-1}{I-1} \right) \quad \text{with} \quad p \in \{1, \dots, P\}, \quad (2.20)$$

and $i \in \{1, \dots, I\}$.

In two dimensions it is given by

$$a_{p i j} = a_p \left(\frac{i-1}{I-1}, \frac{j-1}{J-1} \right) \quad \text{with} \quad p \in \{1, \dots, P\}, \quad (2.21)$$

$i \in \{1, \dots, I\}$,
and $j \in \{1, \dots, J\}$.

Finally, for three dimensions we have

$$a_{p\,ijk} = a_p \left(\frac{i-1}{I-1}, \frac{j-1}{J-1}, \frac{k-1}{K-1} \right) \quad \text{with} \quad \begin{aligned} p &\in \{1, \dots, P\}, \\ i &\in \{1, \dots, I\}, \\ j &\in \{1, \dots, J\}, \\ \text{and } k &\in \{1, \dots, K\}. \end{aligned} \quad (2.22)$$

The index p will be omitted in the notation for scalar fields. Note that the same discretization applies for the above-introduced metric terms.

Hereinafter we define the column vector \mathbf{a} as the discrete representation of the field $\mathbf{a}(\boldsymbol{\xi})$ given by the sequence of values of the field at every grid point, following an increasing index order. For instance, in three dimensions, the components of \mathbf{a} read

$$\mathbf{a} = [a_l] \quad \text{with} \quad a_{P \times I \times J \times (k-1) + P \times I \times (j-1) + P \times (i-1) + p} = a_{p\,ijk}, \quad (2.23)$$

The vector \mathbf{a} has $P \times I \times J \times K$ components.

2.2.3 Implementation

2.2.3.1 Fields

In the numerical code, we provide a general class for the storage and manipulation of fields distributed among several processing units.² In figure 2.2 we show an illustration of the grid partitioning in two dimensions, and the distribution of data among the processes for a grid consisting of a single block. The processing units for the illustrated case are arranged into an MPI Cartesian topology, and according to their position in the topology, they are assigned a given subset of the field values.

In the example given in figure 2.2, a processor (m, n) stores the values of a given vector field \mathbf{a} in its subdomain, following the order given by the indices $p\,ij$, i.e. first the P components of the vector, then following the ξ_1 -axis from i_m to $i_{m+1} - 1$, and finally, following the ξ_2 -axis from j_n to $j_{n+1} - 1$. Note that maximum data locality is achieved for the components of the field at a single point. Scalar and tensor fields are managed analogously, and the extension to three dimensions and grids consisting of multiple blocks is conceptually straightforward although considerably more involved.

Several operations have been implemented on the field class, such as the scalar product $\langle \cdot, \cdot \rangle$ given by

$$\langle \mathbf{b}, \mathbf{a} \rangle = \sum_{ij} b_{p(ij)} M_{pq(ij)} a_{q(ij)}. \quad (2.24)$$

Parallel input/output operations are implemented using the HDF5 library (The HDF Group, 2012) which allows the storage of a CPU-number independent representation of the field in a given storage device. The post-processing of the resulting .h5 file is performed using the Python wrapper for the HDF library h5py (Collette, 2012) and the computer program ParaView (Henderson, 2007).

2.2.3.2 Geometry

A class for handling the geometry is provided. It stores the given grid mapping $\mathbf{x}(\boldsymbol{\xi})$, the Jacobians $\mathbf{j}(\boldsymbol{\xi})$ and $\mathbf{k}(\boldsymbol{\xi})$, etc. Note that although the grid mapping $\mathbf{x}(\boldsymbol{\xi})$ is usually generated in a preprocessing step

² In this document, we use the terms *processor* and *CPU* indifferently to refer to a single process unit, and not to the physical device that may consist of several processing units.




Figure 2.2: Grid partitioning and fields storage: representation of the two-dimensional single-block case. The computational domain is partitioned in rectangular subdomains, and each processing unit stores the respective field components in its subdomain.

and loaded from a storage device, the Jacobian is not available and needs to be computed explicitly. This operation requires the calculation of spatial derivatives which will be introduced in section 2.5.

2.3 Pseudo-characteristics formulation

In this section, the governing equations presented above are rearranged in a form amenable to their discretization and subsequent implementation. We focus on formulations that are suitably discretized using high-order finite differences in space and explicit integration schemes in time.

The approach followed here is known as the pseudo-characteristics formulation (see Sesterhenn, 2000), and it deals with the simulation of high-speed *smooth* compressible flows. The key idea of the formulation is the rearrangement of the inviscid part of the compressible Navier–Stokes equations to explicitly generate advection terms whose speed of propagation are given in a convenient form for their numerical evaluation.

2.3.1 Introduction of pseudo-characteristics

With the previous definitions for curvilinear grids (equations (2.14) to (2.18)), we can now focus on the derivation of the pseudo-characteristics (or pseudo-waves). Although, in general, the inviscid part of the compressible Navier–Stokes equations cannot be decomposed into a set of travelling plane waves, it can still be rearranged into transport terms with the transport speed of information appearing explicitly. Such a decomposition is not unique and furthermore is only practical when the transport terms are obtained along the curvilinear coordinates (see Sesterhenn, 2000). Three different types of transport terms can be defined, namely acoustic pseudo-waves, cross-transport terms and entropy pseudo-waves.

2.3.1.1 Acoustic pseudo-waves

We consider now the pressure equation and the projection of the momentum equation onto the vector ξ^k , which yields an equation for u^k . The inviscid terms read

$$\frac{\partial p}{\partial t} + u^l \frac{\partial p}{\partial \xi^l} + \gamma \rho a^2 \frac{\partial u_i}{\partial \xi^l} k^{li} = \dots, \quad (2.25a)$$

and

$$\frac{\partial u_i}{\partial t} k^{ki} + u^l \frac{\partial u_i}{\partial \xi^l} k^{ki} + \frac{1}{\gamma M^2} \frac{1}{\rho} \frac{\partial p}{\partial \xi^l} g^{kl} = \dots \quad (2.25b)$$

We can divide the first equation by $\gamma \rho a$, and add or subtract the second equation multiplied by M/h^k . After some manipulation, the terms for $l = k$ read³

$$\left(\frac{1}{\gamma \rho a} \frac{\partial p}{\partial t} \pm M \frac{\partial u_i}{\partial t} \frac{k^{(k)i}}{h^{(k)}} \right) + \underbrace{\left(u^{(k)} \pm \frac{a}{M} h^{(k)} \right) \left(\frac{1}{\gamma \rho a} \frac{\partial p}{\partial \xi^{(k)}} \pm M \frac{\partial u_i}{\partial \xi^{(k)}} \frac{k^{(k)i}}{h^{(k)}} \right)}_{\Theta_k^\pm} = \text{RHS}_k^\pm, \quad (2.26)$$

which suggests the introduction of the acoustic pseudo-waves Θ_k^\pm . Note that the present definition differs conceptually from the classical definition for the Riemann invariants which do not include the transport velocity in their definition. The inviscid terms in the pressure equation (2.25a) can be written in terms of the acoustic pseudo-waves by noticing that

$$u^l \frac{\partial p}{\partial \xi^l} + \gamma \rho a^2 \frac{\partial u_i}{\partial \xi^l} k^{li} = \gamma \rho a \sum_l \frac{\Theta_l^+ + \Theta_l^-}{2}. \quad (2.27)$$

2.3.1.2 Cross-transport terms

Similarly, the acoustic pseudo-waves can be inserted back into the equation 2.25b by using the relation

$$u^l \frac{\partial u_i}{\partial \xi^l} k^{ki} + \frac{1}{\gamma M^2} \frac{1}{\rho} g^{lk} \frac{\partial p}{\partial \xi^l} = \sum_l \frac{\Theta_{(l)}^+ - \Theta_{(l)}^-}{2M} \frac{g^{k(l)}}{h^{(l)}} + \underbrace{\sum_l u^{(l)} \frac{\partial u_i}{\partial \xi^{(l)}} \left(k^{ki} - k^{(l)i} \frac{g^{k(l)}}{g^{(ll)}} \right)}_{\Theta^{kl}}. \quad (2.28)$$

From the latter expression we define the cross-transport terms Θ^{kl} . Note that Θ^{kl} is identically zero for $l = k$.

2.3.1.3 Entropy pseudo-waves

Finally, the equation for the entropy reads

$$\frac{\partial s}{\partial t} + \underbrace{u^{(k)} \frac{\partial s}{\partial \xi^{(k)}}}_{\Theta_s^k} = \dots, \quad (2.29)$$

which suggests the definition of the entropy pseudo-waves Θ_s^k .

³The terms for $l \neq k$ and viscous terms are grouped into the right-hand-side term RHS_k^\pm .

2.3.2 Summary of equations compressible Navier–Stokes equations in pseudo-characteristics form

The following pseudo-waves and transport terms are defined:

$$\text{acoustic pseudo-waves} \quad \Theta_k^\pm = \left(u^{(k)} \pm \frac{a}{M} h^{(k)} \right) \left(\frac{1}{\gamma \rho a} \frac{\partial p}{\partial \xi^{(k)}} \pm M \frac{\partial u_i}{\partial \xi^{(k)}} \frac{k^{(k)i}}{h^{(k)}} \right), \quad (2.30a)$$

$$\text{cross-transport terms} \quad \Theta^{kl} = u^{(l)} \frac{\partial u_i}{\partial \xi^{(l)}} \left(k^{ki} - k^{(l)i} \frac{g^{k(l)}}{g^{(ll)}} \right), \quad (2.30b)$$

$$\text{entropy pseudo-waves} \quad \Theta_s^k = u^{(k)} \frac{\partial s}{\partial \xi^{(k)}}. \quad (2.30c)$$

The compressible Navier–Stokes equations can be written using curvilinear coordinates as

$$\text{pressure equation} \quad \frac{\partial p}{\partial t} + \gamma \rho a \sum_l \frac{\Theta_l^+ + \Theta_l^-}{2} = (\gamma - 1) p \text{RHS}_s, \quad (2.31a)$$

$$\text{momentum equation} \quad \frac{\partial u^k}{\partial t} + \sum_l \frac{\Theta_{(l)}^+ - \Theta_{(l)}^-}{2M} \frac{g^{(l)k}}{h^{(l)}} + \sum_l \Theta^{kl} = \text{RHS}^k, \quad \frac{\partial u_i}{\partial t} = j_{ik} \frac{\partial u^k}{\partial t}, \quad (2.31b)$$

$$\text{entropy equation} \quad \frac{\partial s}{\partial t} + \sum_l \Theta_s^l = \text{RHS}_s, \quad (2.31c)$$

with the shorthands

$$\text{RHS}^k = \frac{1}{\text{Re}} \frac{1}{\rho} k^{ki} \frac{\partial \tau_{ji}}{\partial \xi^l} k^{lj} + \frac{1}{\rho} k^{ki} f_{vi}, \quad (2.31d)$$

$$\text{RHS}_s = \frac{1}{\rho T} \left(\frac{\gamma M^2}{\text{Re}} \tau_{ji} \frac{\partial u_i}{\partial \xi^l} k^{lj} - \frac{\gamma}{(\gamma - 1) \text{PrRe}} \frac{\partial q_i}{\partial \xi^l} k^{li} + Q_v \right). \quad (2.31e)$$

The constitutive laws (equation (2.11)) in curvilinear coordinates read

$$\text{heat flux} \quad q_i = -\chi \frac{\partial T}{\partial \xi^l} k^{li}, \quad (2.32a)$$

$$\text{viscous stresses} \quad \tau_{ij} = \mu \left(\frac{\partial u_i}{\partial \xi^l} k^{lj} + \frac{\partial u_j}{\partial \xi^l} k^{li} \right) + \left(\mu_v - \frac{2}{3} \mu \right) \frac{\partial u_m}{\partial \xi^l} k^{lm} \delta_{ij}. \quad (2.32b)$$

The equations presented above are completed by the thermodynamic relations given in equation (2.12), the non-dimensional parameters in equation (2.13) and the grid metrics defined in equations (2.14) to (2.18). Given a state vector \mathbf{v}

$$\mathbf{v} = [p \quad s \quad u_1 \quad \dots]^t, \quad (2.33)$$

we can symbolically write the above equations as

$$\frac{\partial \mathbf{v}}{\partial t} = \mathbf{f} \left(\mathbf{v}, \frac{\partial \mathbf{v}}{\partial \xi}, \frac{\partial^2 \mathbf{v}}{\partial \xi^2} \right), \quad (2.34)$$

where the dependence on the computational domain metric is implied.

2.3.3 Discussion

2.3.3.1 Interpretation of the pseudo-waves and cross-transport terms

We can now give an interpretation of the previously defined terms. Using vector notation, the momentum equation reads

$$\frac{\partial \mathbf{u}}{\partial t} + \underbrace{\sum_l \frac{\Theta_{(l)}^+ - \Theta_{(l)}^-}{2M} \frac{\boldsymbol{\xi}^{(l)}}{b^{(l)}} + \sum_l \Theta^{kl} \boldsymbol{\xi}_k}_{(\mathbf{u} \cdot \nabla) \mathbf{u} + \frac{1}{\rho} \nabla p} = \underbrace{\text{RHS}^k \boldsymbol{\xi}_k}_{\frac{1}{\rho} (\nabla \cdot \boldsymbol{\tau} + \mathbf{f}_v)}. \quad (2.35)$$

This means that the temporal change of the velocity vector \mathbf{u} at a given location is caused by the combined action of the acoustic pseudo-waves, cross-transport terms and the constraints grouped on the right-hand side. The quantities $\pm \Theta_{(l)}^\pm / (2M b^{(l)})$ and Θ^{kl} can be interpreted as the covariant (contravariant) components of an *acoustic pseudo-waves (cross-transport terms) vector*. Moreover, from the sign of the contravariant components of the velocity vector, we can infer the advection direction of the transported quantities.

In other words, the acoustic pseudo-waves for l have been defined such that all the terms with partial derivatives over ξ^l are simultaneously grouped into the pressure equation and the momentum equation for the covariant component of the velocity u_l . The respective cross-transport terms are the remaining transport terms, so that finally we have $(\mathbf{u} \cdot \nabla) \mathbf{u} + 1/\rho \nabla p$.

The reader will note that this decomposition is not unique in the sense that different curvilinear coordinates lead to different set of pseudo-waves. Nonetheless, the above observations will be particularly useful for the enforcement of boundary conditions and the application of the proper *upwind biasing* in the evaluation of spatial derivatives.




Figure 2.3: Pseudo-waves in the momentum equation: (a) acoustic pseudo-waves Θ_l^\pm and (b) cross-transport terms Θ^{kl} .

2.3.3.2 Remarks on spatial differentiation schemes and boundary conditions

Although the derivation of the spatial differentiation schemes and the determination of the boundary conditions will be addressed later, it is convenient to make some minor remarks here.

It is well-known in numerical analysis that numerical instabilities can arise in the time integration of the advection terms if centred schemes are used. A procedure to overcome such problems is to apply *upwind biased* schemes according to the sign of the advection speed. Note that the present approach allows us to obtain the advection speed explicitly. In contrast, the evaluation of diffusion terms using centred schemes does not cause numerical instabilities provided that a stability condition is satisfied.

Furthermore, we encounter at the boundaries the issue of imposing a certain number of physical boundary conditions. Numerically, on the boundary it is impossible to apply proper *upwind biasing* in the evaluation of spatial derivatives over the coordinate off the boundary surface, if the velocity has a component in the normal direction. However, the number of (inviscid) boundary conditions that can be imposed coincides with the number of pseudo-waves that cannot be properly evaluated. The procedure that will be detailed below consists of setting the values of the appropriate pseudo-waves so that the desired (inviscid) boundary conditions are satisfied.

2.3.4 Discretization

The equations are discretized on the numerical grid using finite differences. We shall assume that the spatial differentiation schemes are given and that the grid metric is known. The temporal evolution of the flow reduces to the evaluation of the right-hand-side of equation (2.34). As we have mentioned before, we use central, positive upwind and negative upwind biased schemes. Consequently, in two-dimensions the resulting system of equations reads

$$\frac{\partial \mathbf{v}_{ij}}{\partial t} = \mathbf{f}_{ij} \left(\mathbf{v}_{ij}, \left. \frac{\partial \mathbf{v}}{\partial \xi} \right|_{ij}, \left. \frac{\partial \mathbf{v}}{\partial \xi} \right|_{ij}^+, \left. \frac{\partial \mathbf{v}}{\partial \xi} \right|_{ij}^-, \left. \frac{\partial^2 \mathbf{v}}{\partial \xi^2} \right|_{ij} \right), \quad (2.36)$$

where \mathbf{f}_{ij} is nonlinear and the functional dependence has been extended to take into account the different schemes: central (no super-index), positive upwind (+ super-index) and negative upwind (- super-index).

2.3.5 Implementation

The evaluation of the right-hand-side of equation (2.36) for one, two and three dimensions has been implemented in an equation class. Note that such a class requires the specification of the spatial discretization schemes, which will be introduced in section 2.5, and a *modifier* class. The purpose of the latter class is to allow the user to add arbitrary source terms and specify boundary conditions. This procedure avoids the customization of the equations class for each flow case. Provided that the required classes are available, the equations module evaluates the right-hand-side of equation (2.36) and calls the spatial differentiation and modifier methods when required.

2.3.5.1 Auxiliary functions

The evaluation of equation (2.36) can be split into several sub-procedures.

We adopt here the convention that the subscript $_{,k}$ indicates the spatial derivative in the k th spatial direction in the computational domain. Furthermore, the subscript can also take the value $k = 0$, in which case, no differentiation is performed at all.

Let us first assume that the flow field and its centred spatial derivatives $\mathbf{v}_{0,k}$ are available. Then, the evaluation at a grid point ξ of the heat flux \mathbf{q} and the viscous stress tensor $\boldsymbol{\tau}$ can be performed




Figure 2.4: Block diagram illustrating the evaluation of the temporal derivative. The routines that perform the spatial differentiation (represented in grey blocks) have not been defined yet.

based on evaluating the flow field and its derivatives $\mathbf{v}_{0,k}$ at that grid point ξ ; see equations (2.32a) and (2.32b). We define the auxiliary vector \mathbf{v}_1 as

$$\mathbf{v}_1 = [q_i \quad \tau_{ij}]^t, \quad (2.37a)$$

and implement the above procedure as

$$\mathbf{v}_1 = \mathbf{f}_1(\mathbf{v}_{0,k}). \quad (2.37b)$$

Analogously, from equations (2.31a) to (2.31c) we observe that if the values of the upwind-biased spatial derivatives of the flow field $\mathbf{v}_{0,k}^\pm$ and the centred spatial derivatives $\mathbf{v}_{1,k}$ are known, then the evaluation can be performed using exclusively the values of those vectors at the given point. We introduce the vector \mathbf{v}_2 :

$$\mathbf{v}_2 = \frac{\partial \mathbf{v}_0}{\partial t}, \quad (2.38a)$$

and implement the above procedure as

$$\mathbf{v}_2 = \mathbf{f}_2(\mathbf{v}_{0,k}, \mathbf{v}_{0,k}^+, \mathbf{v}_{0,k}^-, \mathbf{v}_{1,k}). \quad (2.38b)$$

An illustration of the above procedure is given in figure 2.4 in block-diagram form. For simplicity, no distinction has been made in the above diagram between centred and upwind-biased spatial derivatives. Note that spatial differentiation (in grey) has not been defined yet, but we have managed to divide the evaluation of the right-hand side of the equations into (i) the evaluation of nonlinear point-wise local functions and (ii) the computation of spatial derivatives. Note that these observations are fundamental for the linearization technique presented in chapter 4.

2.4 Boundary conditions

2.4.1 Preliminaries

Given an initial flow field, the subsequent temporal evolution is dependent on boundary conditions. There is a vast body of literature on the determination and implementation of proper boundary conditions for the compressible Navier–Stokes equations. An in-depth study of the different choices is beyond the scope of this work, but it shall be mentioned that the numerical implementation must

be performed judiciously as numerical instabilities would occur otherwise. The boundary conditions derived here are inspired by those described in Poinso and Lele (1992), known as the Navier–Stokes Characteristics Boundary Conditions (NSCBC), and by several of the enhancements proposed by Lodato et al. (2008). The reader is referred to the given references for a detailed description and a comparison with alternative formulations.

2.4.1.1 Open flows

As it has been alluded to before, the analysis of open flows requires the truncation of the domain for practical calculations. Although the boundary conditions are formally given as the asymptotic value of the governing flow variables for an infinite domain, the precise dynamics at the boundary have to be prescribed in the calculation. In this respect, we encounter the problem of imposing *artificial* (Colonus, 2004) dynamics that will capture the physical behaviour of the flow.

To illustrate the latter point, let us consider, for instance, sound generation by an obstacle in an open flow. On the one hand, the inlet boundary condition could be chosen to impose the asymptotic values of the unperturbed flow. On the other hand, an incident perturbation—such as an acoustic wave—would be reflected at the inlet, leading to an unphysical downstream-propagating perturbation. Furthermore, in two- and three-dimensional flows, spurious reflections are unavoidable, but their effect can nevertheless be controlled. The compromise solution is to let the value of the flow field at the boundary vary on a short time scale to allow incident perturbations to leave the computational domain, but still impose the values on a longer time scale using a restoring term such that the reference values are set. This procedure amounts to the construction of *partially non-reflecting* boundary conditions whose specific properties and parameters have to be fine-tuned.

2.4.1.2 Effect on hydrodynamic instabilities

It is well-known that the boundary conditions may potentially change the nature of the instability mechanisms that lead to unsteady flow. For instance, Buell and Huerre (1988) show an example of an absolute instability triggered by an inaccurate treatment of the outflow boundary condition. In this regard, great care needs to be exercised in the design of effective boundary conditions to minimize the artificial effect of spurious reflections.

2.4.1.3 Inviscid and viscous boundary conditions

We shall distinguish inviscid- and viscous-type boundary conditions, according to whether they are applicable in the inviscid limit ($Re \rightarrow \infty$) or not. The inviscid boundary conditions receive the treatment analogous to boundary conditions for hyperbolic systems of equations. At the boundary, a certain number of waves are leaving the computational domain whereas some are entering. The waves entering the domain are chosen in such a way that the boundary displays the desired behaviour.

The finite-Reynolds number regime increases the order of the spatial derivatives, thus increasing the number of boundary conditions to be specified. The viscous boundary conditions typically consist of constraints on the viscous stress tensor and the heat fluxes. These are commonly treated by replacing the given value in the equations during their evaluation.

As pointed out in Poinso and Lele (1992), a requirement for the implementation of consistent boundary conditions exists in pairing each type of boundary condition with the corresponding type of term in the equations.

2.4.1.4 Boundary conditions in the pseudo-characteristics formulation

In the pseudo-characteristics formulation, the arrangement of the equations makes the determination of incoming and outgoing waves quite straightforward.

In section 2.2 we have increased the complexity of the governing equations with respect to their formulation in Cartesian coordinates by introducing a boundary-conforming curvilinear coordinate system. The advantages of such coordinate systems are now becoming clear: the components of the velocity in the covariant and contravariant components alleviate the task of determining which pseudo-wave enters or leaves the domain; the geometric relations of orthogonality or the tangency relations that the vector basis fulfils with respect to the boundary plane (see figure 2.1) will accomplish this task.

For the following derivations, we shall assume that we are treating a boundary condition on $\xi^n = 0$ or $\xi^n = 1$ and $0 \leq \xi^t \leq 1$ with $t \neq n$. This situation is depicted in figure 2.5.

We focus on boundary conditions for the subsonic regime, i.e. $|u| \leq a/M$ at edges (in two dimensions) or faces (in three dimensions). The treatment of corners and edges—where several different types of boundary conditions apply simultaneously—requires additional compatibility conditions. Such considerations will not be treated here, and the reader is referred to Poinso and Lele (1992) for further details.

2.4.2 Wall boundary conditions

Let us first consider the boundary conditions for a wall at rest. The velocity at an arbitrary point on the surface is then zero. The reasoning that follows can be easily extended for momentum injection.

Regarding the velocity, two different types of wall boundary conditions can be considered: (i) no-slip, with zero velocity at the wall $\mathbf{u} = 0$, and (ii) slip, with a zero wall-normal velocity component $\mathbf{u} \cdot \boldsymbol{\xi}^n = 0$. The specification of the wall-normal and wall-tangent velocity components is conceptually different, as in the inviscid limit only the former applies. In either case, provided that the initial flow field satisfies the boundary condition $\mathbf{u} = 0$, the boundary condition can be expressed as

$$\frac{\partial \mathbf{u}}{\partial t} = 0 \quad \text{or} \quad \frac{\partial \mathbf{u} \cdot \boldsymbol{\xi}^n}{\partial t} = 0. \quad (2.39)$$

In addition, in the case of slip boundary condition, the viscous stresses tangential to the wall are set to zero.

The thermal boundary condition is of viscous type. The temperature at the wall T_w or the heat flux \mathbf{q}_w may be prescribed. As before, provided that the initial flow field satisfies the temperature boundary condition, we enforce instead

$$\frac{\partial T}{\partial t} = \frac{\partial T_w}{\partial t}. \quad (2.40)$$

2.4.2.1 Inviscid boundary condition: wall-normal velocity

The natural basis for expressing the velocity components and the momentum equation is the contravariant basis. We recall that we have obtained such equations by projecting the momentum equation onto the covariant basis and that only contravariant components of the velocity u^n appear in setting the wall-normal velocity as $\boldsymbol{\xi}_t \cdot \boldsymbol{\xi}^n = 0$. Considering equation (2.31b) for $k = n$, and due to the

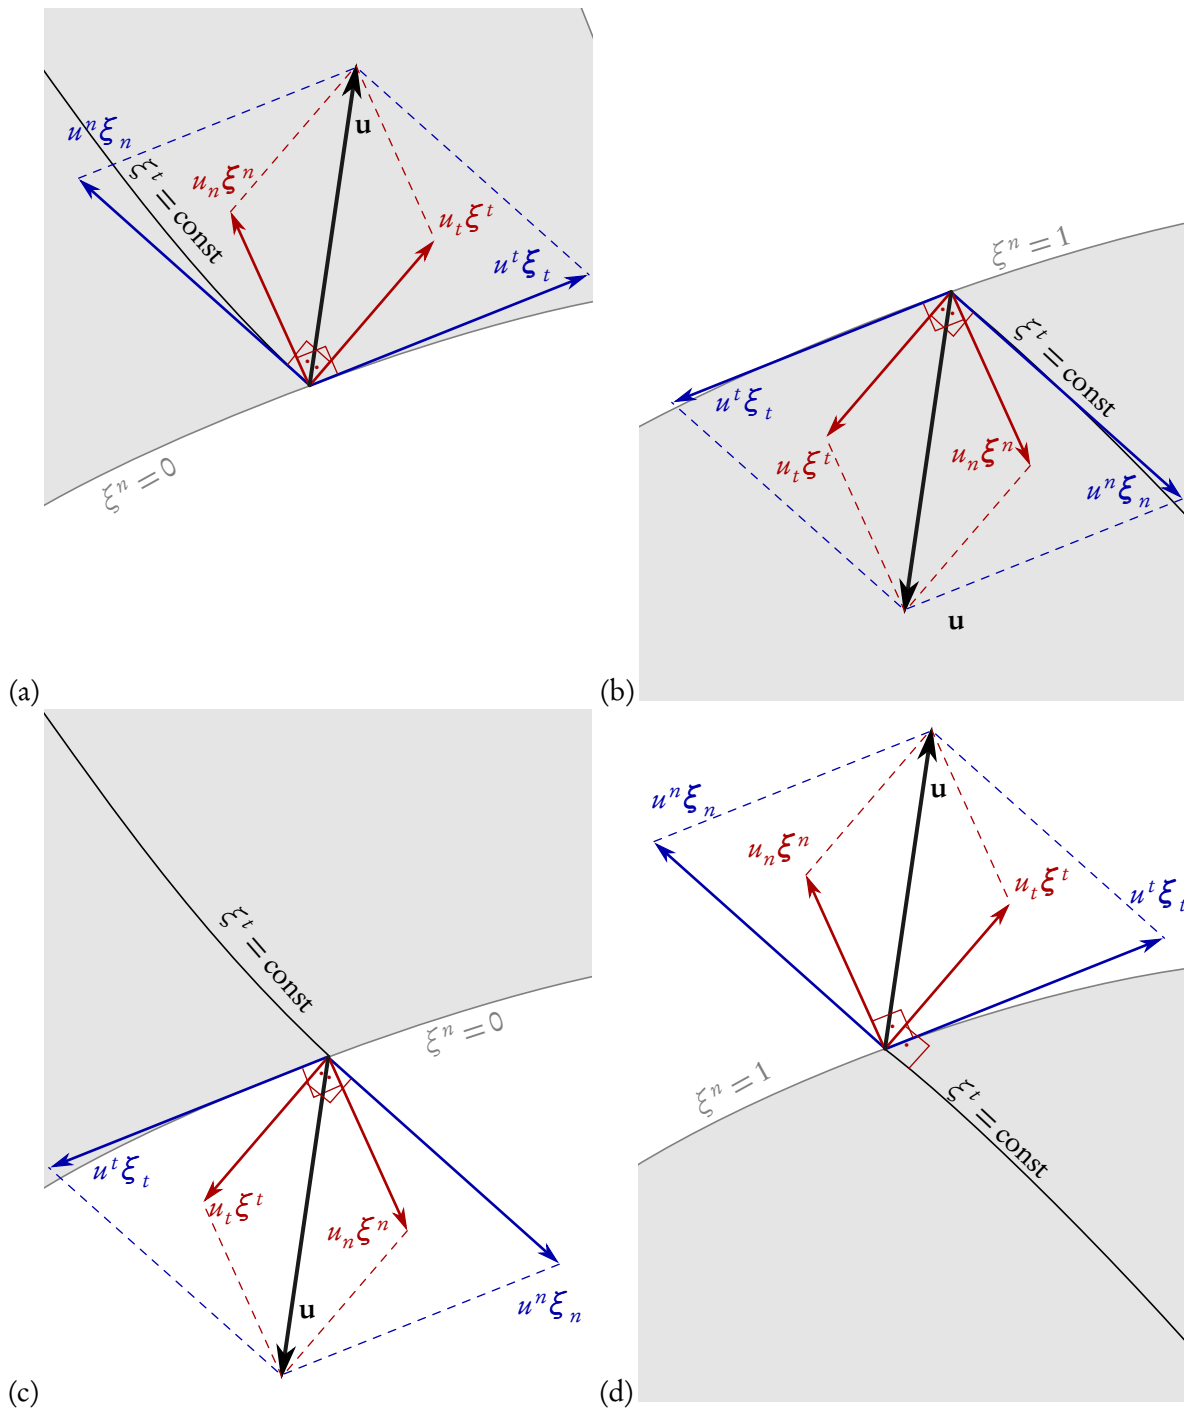


Figure 2.5: Vector basis and boundary conditions showing (a) inflow at $\xi^n = 0$, (b) inflow at $\xi^n = 1$, (c) outflow at $\xi^n = 1$, (d) inflow at $\xi^n = 1$. Note that ξ_t is contained in the tangent plane to the boundary surface and ξ^n is normal to the boundary plane.

location of the boundary, we solve for the incoming pseudo-wave. Hence

$$\Theta_n^\pm = \Theta_n^\mp \mp \frac{2M}{h^n} \left(\frac{\partial u^{\vec{n}}}{\partial t} + \sum_{l \neq n} \frac{\Theta_{(l)}^+ - \Theta_{(l)}^-}{2M} \frac{g^{(l)n}}{h^{(l)}} + \Theta^n - \text{RHS}^n \right) \quad (2.41)$$

It is legitimate to wonder if the same effect could be achieved by directly setting the temporal derivative to zero. In fact, the obtained acoustic pseudo-wave appears in the momentum equations for the tangential components u^t and the equation for the pressure, thus not yielding the same behaviour. The subtle difference is that if only the temporal derivative is set to zero, the incoming acoustic pseudo-wave in the remaining equations would be evaluated numerically using an incorrect upwind biasing, leading to an unstable calculation.

2.4.2.2 Viscous boundary condition: velocity

The wall-tangent velocity components are readily written in the contravariant basis. It is tempting to proceed similar to the previous development. However, in this case the viscous constraints that enforce the boundary condition are sought instead, since the present boundary condition is of viscous type. In practical terms, as these do not appear in the remaining equations, the same result can be achieved by directly setting the temporal derivative to zero. Hence

$$\frac{\partial u^t}{\partial t} = 0. \quad (2.42)$$

Alternatively, a slip boundary condition is equivalent to setting the wall-tangent viscous stresses to zero, i.e.

$$\xi_t \cdot (\boldsymbol{\tau} \cdot \xi^n) = 0 \quad \text{and} \quad \xi^n \cdot (\boldsymbol{\tau} \cdot \xi_t) = 0. \quad (2.43)$$

This can be achieved by redefining the viscous stress tensor in equation (2.32b) as

$$\boldsymbol{\tau} - \sum_{t \neq n} \xi_t \cdot (\boldsymbol{\tau} \cdot \xi^n) (\xi^t \otimes \xi_n + \xi_n \otimes \xi^t). \quad (2.44)$$

2.4.2.3 Viscous boundary conditions: heat transfer

The temporal variations of the temperature can be related to changes in the pressure and the entropy. It can be shown that

$$\frac{\partial T}{\partial t} = \frac{\gamma - 1}{\gamma} T \left(\frac{1}{p} \frac{\partial p}{\partial t} + \frac{\partial s}{\partial t} \right). \quad (2.45)$$

Using the previous relation and equations (2.31a) and (2.31c), we can set RHS_s such that the above equation is satisfied:

$$\text{RHS}_s = \frac{1}{\gamma - 1} \frac{1}{T} \frac{\partial T_w}{\partial t} + \frac{1}{a} \sum_l \frac{\Theta_l^+ + \Theta_l^-}{2} + \frac{1}{\gamma} \sum_l \Theta_s^l. \quad (2.46)$$

If the heat flux is specified (such as for an adiabatic boundary condition), the natural basis to express the boundary condition for the wall-normal heat flux \mathbf{q}_n is the covariant basis:

$$\mathbf{q}_w = q_{(n)} \xi^{(n)}. \quad (2.47)$$

This condition can be directly set in equation (2.32a) by redefining the heat flux as:

$$\mathbf{q} + (q_{(n)} - \xi_{(n)} \cdot \mathbf{q}) \xi^{(n)}. \quad (2.48)$$

2.4.3 Inflow boundary condition

2.4.3.1 Incoming pseudo-characteristics

At a subsonic inlet, the reference values for the velocity and the entropy are \mathbf{u}_{ref} and s_{ref} , respectively. The inflow is characterized by $u^n > 0$ if $\xi^n = 0$, or $u^n < 0$ if $\xi^n = 1$; see figure 2.5a,b. The sign of the pseudo-waves advection speed allows us to determine the pseudo-characteristic. For subsonic flow, the incoming pseudo-characteristics are

$$\Theta_n^+ \text{ or } \Theta_n^-, \Theta^{tn}, \text{ and } \Theta_s^n \quad (2.49)$$

We follow the natural extension to curvilinear coordinates of the procedure in Lodato et al. (2008). Let us reconsider again equation (2.26). The key idea is to choose the appropriate acoustic pseudo-wave to satisfy the following equation at the boundary:

$$\left(\frac{1}{\gamma \rho a} \frac{\partial p}{\partial t} \pm M \frac{\partial u_i}{\partial t} \frac{k^{(n)i}}{b^{(n)}} \right) \pm \frac{\eta_{\pm}^{(n)}}{L} (1 - M^2) \frac{a}{b^{(n)}} (u^{(n)} - u_{\text{ref}}^{(n)}) = 0, \quad (2.50)$$

which leads to

$$\Theta_{\pm}^n = \pm \frac{\eta_{\pm}^n}{L} (1 - M^2) \frac{a}{b^n} (u^n - u_{\text{ref}}^n) + \text{RHS}_{\pm}^n. \quad (2.51)$$

Similarly, the incoming cross-transport terms Θ^{tn} are chosen so that

$$\frac{\partial u^t}{\partial t} + \frac{\eta^{(t)}}{L} \frac{a}{M} (u^{(t)} - u_{\text{ref}}^{(t)}) = 0, \quad (2.52)$$

which leads to

$$\Theta^{tn} = \frac{\eta^{(t)}}{L} \frac{a}{M} (u^{(t)} - u_{\text{ref}}^{(t)}) + \text{RHS}^{tn}. \quad (2.53)$$

Finally, the incoming entropy pseudo-wave Θ_s^n is chosen so that we have

$$\frac{\partial s}{\partial t} + \frac{\eta_s}{L} \frac{a}{M} (s - s_{\text{ref}}) = 0 \quad (2.54)$$

which results in

$$\Theta_s^n = \frac{\eta_s}{L} \frac{a}{M} (s - s_{\text{ref}}) + \text{RHS}_s^n. \quad (2.55)$$

In the above, L refers to a characteristic non-dimensional length of the domain. The various relaxation parameters η_{\pm}^n , η^t and η_s are of unit order, and the optimal values depend on the specifics of the problem.

The resulting equations at the boundary can be interpreted as a system that is forced by the advection terms with a restoring force towards the reference values (denoted by the subscript $_{\text{ref}}$).

2.4.3.2 Viscous boundary conditions

Following Poinot and Lele (1992), additional viscous boundary conditions must be imposed to obtain a numerically stable computation. In particular, we impose:

$$k^{ni} \frac{\partial \tau_{ij}}{\partial \xi^n} k^{nj} = 0. \quad (2.56)$$

2.4.4 Outflow boundary condition

2.4.4.1 Outgoing pseudo-characteristics

The case of an outflow boundary condition is treated similarly to the inflow boundary condition. In such a case, the velocity vector points out of the domain, see figure 2.5c,d, and consequently only one pseudo-characteristic needs to be specified. Ideally, the suitable pseudo-wave may be set to zero, but as seen in Poinso and Lele (1992), this choice is inappropriate for long-time calculations as it leads to a pressure drift.

It is rather preferred to introduce a restoring term that avoids substantial deviations from the reference pressure. Following Lodato et al. (2008), the incoming acoustic pseudo-wave is chosen to satisfy

$$\left(\frac{1}{\gamma \rho a} \frac{\partial p}{\partial t} \pm \frac{k_{ki}}{\sqrt{g^{(kk)}}} M \frac{\partial u_i}{\partial t} \right) + \frac{\sigma}{L} \frac{1-M^2}{\gamma M} \frac{p-1}{\rho} - \beta \left(\text{RHS}_k^\pm - \text{RHS}_{k,\text{ref}}^\pm \right) = 0. \quad (2.57)$$

Hence

$$\Theta_\pm^k = \frac{\sigma}{L} \frac{1-M^2}{\gamma M} \frac{p-1}{\rho} + (1-\beta) \text{RHS}_k^\pm + \beta \text{RHS}_{k,\text{ref}}^\pm. \quad (2.58)$$

Two relaxation coefficients, namely σ and β , have been introduced. The values $\sigma = 0.528$ and $\beta = M$ have been suggested in Poinso and Lele (1992); Lodato et al. (2008).

2.4.4.2 Viscous boundary conditions

Additional viscous boundary conditions have to be specified. We have

$$\frac{\partial q_n}{\partial \xi^n} = 0, \quad \text{and} \quad \frac{\partial \tau_{tn}}{\partial \xi^n} = 0. \quad (2.59)$$

2.4.5 Sponge layers

In some cases, the non-reflecting outflow boundary condition presented above are nonetheless insufficient to avoid spurious reflections. A common example is the passage of strong vortical structures at the outlet. In this situation, sponge layers (Bodony, 2006) are typically used to further reduce the reflections. Note that formally sponge layers cannot be considered as boundary conditions.

A sponge layer consists of a restoring term towards a reference state \mathbf{v}_{ref} that is added to the governing equations. The implementation is performed by adding a spatially-dependent forcing term $-\sigma(\mathbf{x})(\mathbf{v} - \mathbf{v}_{\text{ref}})$ to equation (2.34).

2.4.6 Implementation

The above boundary conditions have been implemented in the code using the already mentioned *modifier* class. Such a class is used at two stages. In the equations module, the equations are assembled by evaluating first the viscous stress tensor and the heat flux, then the *modifier* is given the chance to modify the computed values in order to impose the various viscous boundary conditions given above. Similarly, after the evaluation of the viscous forces and pseudo-waves, the *modifier* is allowed to reassign a different value to the pseudo-characteristics, replace an equation at the boundary or add a source term. Note that the incorporation of boundary conditions does not alter the decomposition into auxiliary functions presented in section 2.3.

2.5 Spatial differentiation

In our numerical code, the integration in time of the compressible Navier–Stokes equations is performed using the method of lines: the governing equations are first discretized in space, yielding a high-dimensional system of ordinary differential equations that will subsequently be integrated in time.

High resolution is crucial for the analysis of wave propagation phenomena, and jet noise is a classical example that illustrates this fact (see Crighton, 1975, and references therein). In this type of flow, the energy associated with the radiated sound is several orders of magnitude smaller than the typical energy fluctuations. Nevertheless, the generated noise is the main feature of interest, and numerical calculations that target such feature require proper spatio-temporal resolution to capture the subtle energy transfer mechanisms. The adequacy of spatio-temporal discretization schemes is typically characterized by the order of accuracy of the discretization and by the effect on the dispersion relation of model equations, such as the convection–diffusion equation.

A common approach to achieve high resolution while containing the total computational cost is the use of high-order schemes. In this respect, the numerical properties of the spatial differentiation schemes featured in the previous (original) numerical code are satisfactory for our purposes, and they thus have been retained here: upwind-biased compact schemes for the advection terms and centred compact schemes for the diffusive terms. We mentioned, however, that additional changes, inspired by the manner in which explicit schemes are computed in practice, have been introduced to allow for a higher computational efficiency for parallel calculations.

We present in this section the characteristics of both spatial schemes—explicit and compact—and discuss implementation and parallelization issues. It is important to point out that efficient evaluation of the spatial derivatives is essential for the overall performance of the numerical code.

2.5.1 Finite differences (FD)

2.5.1.1 Explicit schemes

In the finite-difference approach, the n th order spatial derivative of a function at a given location x_i is approximated by a linear combination of the values of the function at a set of nodes x_{i+k} . For a given function $f(x)$, this approximation has the form

$$f'_i = \sum_k b_{i,i+k} f_{i+k}, \quad (2.60)$$

or, equivalently in matrix notation

$$\mathbf{f}' = \mathbf{B}\mathbf{f}, \quad (2.61)$$

where the superscript $'$ denotes spatial differentiation. Finite-difference schemes that are naturally written in the previous form are known as *explicit* schemes. In practice, the coefficients $b_{i,i+k}$ are only non-zero for k smaller than the *stencil width* k_{\max} . The approximation is then said to be of compact support, since the domain of dependence of the scheme is restricted to a few neighbouring nodes. The former expression can alternatively be interpreted as the exact derivative at x_i of a polynomial that interpolates the function at the nodes x_{i+k} . In this respect, it is well known that high-order interpolation on uniform grids may suffer from Runge’s phenomenon, which makes the derivation of effective high-order explicit schemes a difficult task.

2.5.1.2 Implicit schemes: compact schemes

A far more general family of finite-difference schemes known as *implicit* schemes is given by:

$$\sum_l a_{i,i+l} f'_{i+l} = \sum_k b_{i,i+k} f_{i+k}, \quad (2.62)$$

or, equivalently in matrix form:

$$\mathbf{A}f' = \mathbf{B}f. \quad (2.63)$$

In this case, the value of the approximated derivative is not explicitly given, and its calculation requires the solution of a linear system of equations that stems from writing the above approximation at each node. In contrast to explicit schemes, implicit schemes have global support since the value of the derivative at a given node depends, in general, on the function values at all nodes. There is a wide range of schemes that fall into this category, such as the spectral (e.g. Fourier and Chebyshev) and the pseudo-spectral methods.

The compact schemes (Lele, 1992) used in our code fall into the latter category, and they feature near-spectral resolution at a reasonable computational cost. The coefficient matrices \mathbf{A} and \mathbf{B} are banded (typically tri- or penta-diagonal), and consequently, the solution of the above system of equations can be obtained efficiently.

2.5.2 Properties of finite-difference schemes: modified wave number

A common procedure to benchmark the spatial resolution properties of finite-difference schemes is the analysis of the modified wave number. We consider the function $f(x)$ given by

$$f(x) = \exp(i\alpha x), \quad (2.64)$$

where α is a real wave number. In the following, the analysis is restricted to discretizations in unbounded domains with uniform grid spacing Δx .

The exact and the approximate value of the derivative computed with the spatial scheme read

$$f'_{i,\text{exact}} = i\alpha \exp(i\alpha x_i), \quad (2.65)$$

and

$$f'_{i,\text{fd}} = i\tilde{\alpha} \exp(i\alpha x_i), \quad \text{with} \quad \tilde{\alpha} = \frac{1}{i} \frac{\sum_k b_{i,i+k} \exp(ik\alpha\Delta x)}{\sum_l a_{i,i+l} \exp(il\alpha\Delta x)}, \quad (2.66)$$

respectively. The quantity $\tilde{\alpha} = \tilde{\alpha}_r + i\tilde{\alpha}_i$ is commonly referred to as the modified wave number, and it provides insight into the influence of the discretization scheme on the computed spatial derivatives. At first sight, the comparison suggests that it would be desirable to have spatial differentiation schemes with $\tilde{\alpha}_r = \alpha$ and $\tilde{\alpha}_i = 0$ at our disposal. In practice, we have $|\tilde{\alpha}_r - \alpha|/\alpha \leq \epsilon$ for wave numbers $\alpha \leq C\Delta x^{-1}$, where C is a constant depending on the scheme. The range of wave numbers that satisfies the above inequality defines *resolved* wave numbers of the problem.

In contrast, the relation $\tilde{\alpha}_i = 0$ can be easily satisfied: it suffices to provide symmetric coefficients $a_{i,i+l} = a_{i,i-l}$ and antisymmetric coefficients $b_{i,i+k} = -b_{i,i-k}$. The numerical schemes with these symmetry properties are known as centred schemes. However, from a practical point of view, and as will be seen later, under some circumstances it is desirable to relax condition $\tilde{\alpha}_i = 0$ in the design of numerical schemes in order to achieve numerical stability for the temporal integration. In

this case, the schemes are referred to as *upwind*-biased schemes, and, depending on whether more relative weight is given to nodes x_{i+k} for $k > 0$ or $k < 0$, they represent *positive* or *negative* upwinding, respectively.

For a given α , the associated real part $\tilde{\alpha}_r$ and imaginary part $\tilde{\alpha}_i$ of the modified wave number $\tilde{\alpha}$ allow us to quantify the effect of the discretization on the convection speeds and diffusion times. To illustrate this point, let us first consider the advection and diffusion equations

$$\frac{\partial f}{\partial t} + U \frac{\partial f}{\partial x} = 0 \quad (2.67) \quad \text{and} \quad \frac{\partial f}{\partial t} = \frac{1}{\text{Re}} \frac{\partial^2 f}{\partial x^2}, \quad (2.68)$$

respectively, and restrict the analysis to functions of the form

$$f(x, t) = \hat{f}(t) \exp(i\alpha x). \quad (2.69)$$

On one hand, by replacing the exact value of the spatial derivative of $f(x, t)$ into equations (2.67) and (2.68), we have

$$\frac{d\hat{f}}{dt} = -i\alpha U \hat{f} \quad \text{and} \quad \frac{d\hat{f}}{dt} = -\frac{1}{\text{Re}} \alpha^2 \hat{f}, \quad (2.70)$$

while, on the other hand, the finite-difference approximation yields

$$\frac{d\hat{f}}{dt} = (-i\tilde{\alpha}_r + \tilde{\alpha}_i) U \hat{f} \quad \text{and} \quad \frac{d\hat{f}}{dt} = -\frac{1}{\text{Re}} (\tilde{\alpha}_r^2 - \tilde{\alpha}_i^2 + 2i\tilde{\alpha}_r \tilde{\alpha}_i) \hat{f}. \quad (2.71)$$

It is illustrative to compare the exact and the modified advection equation: we notice that the real part of the modified wavelength can be interpreted as a change in the advection speed for a given wavelength; similarly, the imaginary part of the modified wavelength can be interpreted as an amplification or an attenuation in time. Equivalently, in the modified diffusion equation, the quantity $\tilde{\alpha}_r^2 - \tilde{\alpha}_i^2$ corresponds to a change in the characteristic dissipation time; the imaginary part can be interpreted as an advection term.

In the above, the effect of non-uniform grids can be accounted for in a straightforward manner. However, the effect of finite domains complicates the assessment of numerical stability and accuracy; see for instance the stability theory of Gustafsson, Kreiss and Sundström (Trefethen, 1983).

2.5.3 Choice of spatial schemes

It is well known that the discretization of the advection equation using centred finite differences and explicit time integration typically leads to unstable calculations due to the amplification of unresolved waves. A technique to overcome this difficulty is the use of upwind-biased schemes, which are characterized by an appreciable amount of artificial dissipation ($\tilde{\alpha}_i < 0$) to ensure stability of the numerical integration. The upwind spatial scheme used here is known as Compact Upwind with Low Dissipation (CULD) and was designed by Adams and Shariff (1996).

The discretization of the diffusion equation using centred schemes does not suffer from numerical instabilities, provided that a sufficiently small time-step is used in the time integration. The reader familiar with incompressible flow calculations will notice that in such a case, the diffusion terms are often integrated implicitly in time to overcome time-step limitations. For the flow cases considered here, however, the time-step restriction for the advection terms is generally more restrictive than the time-step restriction for the diffusion terms, thus justifying the explicit time integration.

The numerical coefficients for the 5th-order CULD scheme and a 3th-order centred spatial schemes used in our numerical code are given in section 2.A. In figure 2.6 we represent the real and imaginary parts of the modified wave numbers for such spatial schemes. As a reference, the modified wave number for the first-order upwind and centred explicit finite differences are included as well.

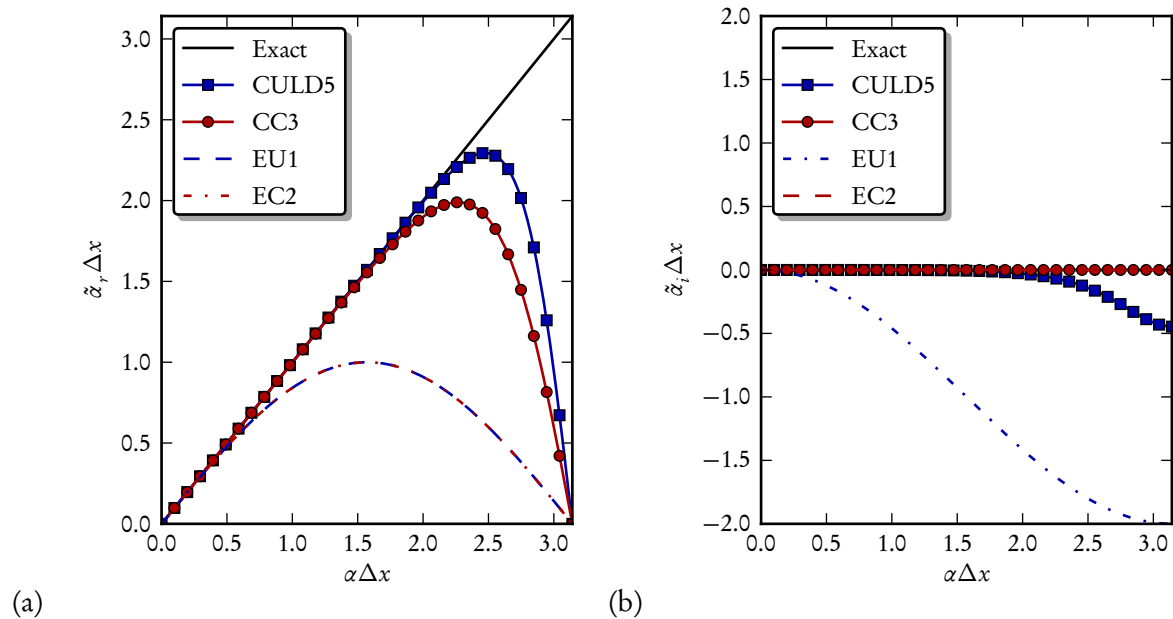


Figure 2.6: Spatial-scheme characteristics showing the modified wave number $\tilde{\alpha}$. (a) real part $\tilde{\alpha}_r$, and (b) imaginary part $\tilde{\alpha}_i$. Explicit schemes are shown for reference. EC1: Explicit centred (1st order), CC3: Compact centred (3rd order), EU1: Explicit upwind (1st order), CULD5: Compact upwind low dissipation (5th order).

The real part of the modified wave number $\tilde{\alpha}_r$ shows that for small wave numbers, the agreement with the exact value is excellent, whereas discrepancies appear for wave numbers around Δx^{-1} . The imaginary part of the modified wave number $\tilde{\alpha}_i$ shows that centred schemes are not afflicted by artificial dissipation. Upwind schemes exhibit artificial dissipation for a rather large range of wave numbers (short wavelengths).

From a practical point of view, this observation allows a rule-of-thumb for the number of points required per wavelength, in order to accurately represent the smallest spatial scales in our simulation. If the wavelength is given by $2\pi/\alpha$, the number of points per wavelength (ppw) reads

$$\text{ppw} = \left\lfloor \frac{2\pi}{\alpha\Delta x} \right\rfloor, \quad (2.72)$$

where $\lfloor \cdot \rfloor$ denotes the floor function. We present in table 2.1 the values of the real and imaginary parts of the modified wave-lengths for increasing number of points per wavelength.

It is readily confirmed that the centred schemes (CC3 and EC) are not artificially dissipative. The CULD5 scheme displays a reduced amount of dissipation only for the range of unresolved wave numbers. In contrast, the first-order upwind scheme (EU1) shows significant artificial dissipation over a wide range of wave numbers, which discourages its use for wave propagation problems.

ppw	CULD5		CC3		EU1		EC1	
	$ \tilde{\alpha}_r - \alpha /\alpha$	$\tilde{\alpha}_i$	$ \tilde{\alpha}_r - \alpha /\alpha$	$\tilde{\alpha}_i$	$ \tilde{\alpha}_r - \alpha /\alpha$	$\tilde{\alpha}_i$	$ \tilde{\alpha}_r - \alpha /\alpha$	$\tilde{\alpha}_i$
2	$1.00 \cdot 10^{+0}$	$-4.46 \cdot 10^{-1}$	$1.00 \cdot 10^{+0}$	0	$1.00 \cdot 10^{+0}$	$-2.00 \cdot 10^{+0}$	$1.00 \cdot 10^{+0}$	0
4	$3.20 \cdot 10^{-4}$	$-5.64 \cdot 10^{-3}$	$9.70 \cdot 10^{-3}$	0	$3.63 \cdot 10^{-1}$	$-1.00 \cdot 10^{+0}$	$3.63 \cdot 10^{-1}$	0
8	$1.56 \cdot 10^{-5}$	$-7.00 \cdot 10^{-5}$	$1.20 \cdot 10^{-4}$	0	$9.97 \cdot 10^{-2}$	$-2.93 \cdot 10^{-1}$	$9.97 \cdot 10^{-2}$	0
12	$1.45 \cdot 10^{-6}$	$-5.89 \cdot 10^{-6}$	$1.01 \cdot 10^{-5}$	0	$4.51 \cdot 10^{-2}$	$-1.34 \cdot 10^{-1}$	$4.51 \cdot 10^{-2}$	0
16	$2.64 \cdot 10^{-7}$	$-1.03 \cdot 10^{-6}$	$1.78 \cdot 10^{-6}$	0	$2.55 \cdot 10^{-2}$	$-7.61 \cdot 10^{-2}$	$2.55 \cdot 10^{-2}$	0
24	$2.34 \cdot 10^{-8}$	$-8.96 \cdot 10^{-8}$	$1.55 \cdot 10^{-7}$	0	$1.14 \cdot 10^{-2}$	$-3.41 \cdot 10^{-2}$	$1.14 \cdot 10^{-2}$	0

Table 2.1: Modified wave numbers for selected spatial differentiation schemes as a function of the number of points per wavelength. EC1: Explicit centred (1st order), CC3: Compact centred (3rd order), EU1: Explicit upwind (1st order), CULD5: Compact upwind low dissipation (5th order).

As stated earlier, centred compact schemes are used for the dissipative terms of the equations, and second-order derivatives are computed by applying the differentiation operator twice. As a consequence, the numerical scheme does not present dissipation for $\alpha\Delta x = \pi$. This fact was also observed by Lele (1992), who suggested the use of second-order finite-difference schemes as the method of choice. However, for practical reasons, the differentiation operator will be applied twice in our case.

A complementary view of the numerical discretization schemes can be obtained by inspecting the spectrum of the discrete differentiation operator. The result for the compact schemes used in this present study is shown in figure 2.7; it will be useful in determining the maximum allowable time-step for a given integration scheme (more details later). As before, we have omitted a discussion on the effect of boundary closures for finite domains.

2.5.4 Implementation

2.5.4.1 Numerical evaluation of spatial derivatives of a scalar function in one dimension

We consider a vector containing the values of the function $f(x)$ on a one-dimensional grid with I uniformly distributed nodes, and a scheme that approximates the spatial derivatives on such nodes determined by the coefficient matrices \mathbf{A} and \mathbf{B} . The value of the spatial derivatives can be obtained by solving the linear system of equations given in equation (2.63).

In the particular case of explicit schemes $\mathbf{A} \equiv \mathbf{I}$, and the calculation of spatial derivatives is performed by repeated matrix-vector product evaluations with different right-hand-sides \mathbf{f} .

In the general case, the system of equations has to be solved, which is performed in two steps: (i) at the beginning of the computation, an LU-decomposition of the coefficient matrix \mathbf{A} is performed⁴, (ii) the above decomposition is used repeatedly for different right-hand-sides \mathbf{f} . The decomposition reads

$$\mathbf{LUf}' = \mathbf{Bf}. \quad (2.73)$$

If the coefficient matrix \mathbf{A} is banded, which is the case for the spatial schemes considered here, then the matrices \mathbf{L} and \mathbf{U} are banded lower and upper triangular, respectively. Moreover, since \mathbf{A} is

⁴For periodic domains, \mathbf{A} contains nonzero components off the bands, in which case the Sherman–Morrison formula is used.




Figure 2.7: Spectra of the spatial discretization schemes. Explicit schemes are shown for reference. EC1: Explicit centred (1st order), CC3: Compact centred (3rd order), EU1: Explicit upwind (1st order), CULD5: Compact upwind with low dissipation (5th order).

usually a diagonally dominant matrix, no partial pivoting strategy is needed, and the matrices \mathbf{L} and \mathbf{U} have the same bandwidth as \mathbf{A} . Finally, the resulting system of equations can be solved by simple forward/backward substitution.

In what follows, it is convenient to symbolically write the differentiation operation as

$$\mathbf{D}_I = \mathbf{A}_I^{-1} \mathbf{B}_I \quad (2.74)$$

where the subscript I indicates the number of nodes.

It is important to point out that in both the explicit schemes and the implicit compact schemes, the complexity of the algorithm is $\mathcal{O}(n)$. At first sight, it might seem that implicit schemes require longer computations since they require solving an additional linear system. Nevertheless, for a fixed resolution, the total calculation time is often shorter as explicit schemes require a larger stencil width and/or a larger number of points.

2.5.4.2 Numerical evaluation of spatial differentiation in n dimensions

In the above, we have only considered spatial differentiation of one-dimensional scalar functions. However, in our numerical code, spatial differentiation is commonly performed on a vector \mathbf{v} with

several components in n dimensions as defined in section 2.2. The sequential case, where a process-unit has access to the entire field, is considered first.

We have implemented a general class that provides methods for performing spatial differentiation on a sequence of \tilde{K} vectors of length \tilde{J} with \tilde{I} components and banded coefficient matrices \mathbf{A} and \mathbf{B} . Symbolically, this algorithm performs the operation

$$\mathbf{D}_{\tilde{K},\tilde{J},\tilde{I}}\mathbf{v} = (\mathbf{I}_{\tilde{K}} \otimes \mathbf{D}_{\tilde{J}} \otimes \mathbf{I}_{\tilde{I}}) \mathbf{v}, \quad (2.75)$$

where $\cdot \otimes \cdot$ denotes the Kronecker product, and \mathbf{I} is the identity matrix.

For a given grid ordering, the implementation of spatial differentiation in this form has several advantages. First, by choosing the appropriate values of the parameters \tilde{K} , \tilde{J} and \tilde{I} , spatial differentiation can be performed along any direction with the same code. Second, the resulting algorithm is cache-friendly and maximizes contiguous memory access, thus potentially allowing the compiler to perform platform-specific optimizations (for example, the use of SIMD instructions).

In order to illustrate the first point, let us consider the vector \mathbf{v} that contains the discrete representation of the field $\mathbf{v}(\xi)$ in three dimensions. The spatial derivative with respect to ξ^2 reads

$$\mathbf{v}_{,2} = \mathbf{D}_{,2}\mathbf{v} = (\mathbf{I}_K \otimes \mathbf{D}_J \otimes \mathbf{I}_I \otimes \mathbf{I}_P) \mathbf{v}. \quad (2.76)$$

The value of the parameters \tilde{K} , \tilde{J} and \tilde{I} can be deduced by rewriting the right-hand side of the previous expression as

$$(\mathbf{I}_K \otimes \mathbf{D}_J \otimes \mathbf{I}_{I \times P}) \mathbf{v}, \quad (2.77)$$

from where we obtain $\tilde{K} = K$, $\tilde{J} = J$ and $\tilde{I} = I \times P$.

Note that the LU decomposition is performed once during a setup phase of the calculation, whereas the matrix-vector product, and the forward/backward substitutions are repeatedly performed during the time-advancement of the equations. In this respect, computational efficiency of the latter operations is crucial.

2.5.4.3 Parallelization: explicit schemes

The algorithm presented above is not straightforward to apply in parallel: the fields are distributed among the process units according to the partitioning displayed in figure 2.2. In general, the value of the spatial derivative at a grid point in a sub-domain may depend on the field values at grid points in other sub-domains. In this case, the spatial derivatives cannot be computed since they are not directly accessible, and field values need to be transferred between process units.

In the case of explicit schemes, the finite-difference approximation has compact support; therefore only the value of the spatial derivative at grid points at a distance to the sub-domain limits smaller than the stencil width k_{\max} depend on the field values that belong to an adjacent sub-domain, which are not directly accessible. The parallel calculation of spatial derivatives for explicit schemes is illustrated in figure 2.8 and can be decomposed into the following steps:

1. Prior to computing the spatial derivatives, each process unit exchanges field values with adjacent process units (or neighbours) in order to be able to compute spatial derivatives in its sub-domain (shown in red) along the differentiation direction.
2. Once the exchange of information has been completed, each process can now compute the spatial derivatives on the grid points in the respective sub-domain.

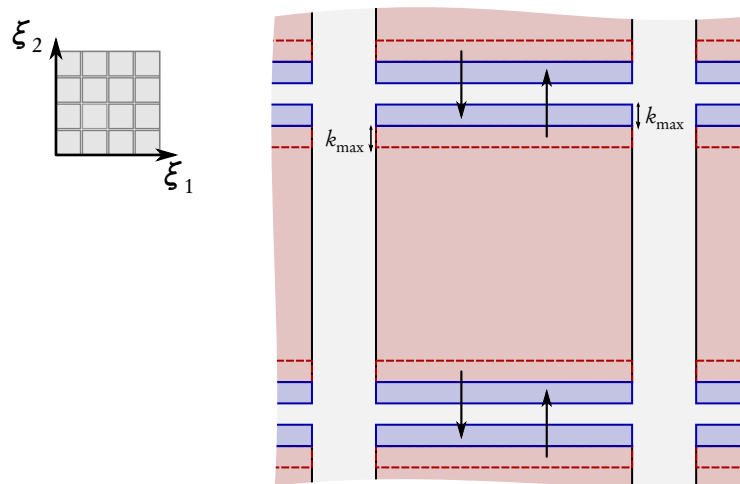


Figure 2.8: Diagram showing the parallelization strategy for the calculation of spatial derivatives in the case of explicit differentiation schemes.

In the above procedure, the operations required for step 1 involve two point-to-point communications per processing unit and per spatial direction; these communications are performed using MPI calls.

Two possibilities arise as to the moment when the information is exchanged. On one hand, the routine `MPI_Sendrecv` performs the exchange of information at the moment that it is called, returning only after the communication has been completed: in this case, the communication is said to be synchronous. On the other hand, using the routines `MPI_Isend` and `MPI_Irecv`, the exchange of information is initiated at the moment the routines are called. However, the transfer continues in the background, and the completion can be later verified by calling the routine `MPI_Wait`. In this case, the communication is said to be asynchronous.

2.5.4.4 Parallelization: implicit schemes

In the case of implicit schemes, the finite-difference approximation has global support, and therefore the value of the derivative at a grid point along a given direction depends on all the values of the field along the differentiation direction; these values are not directly accessible to the process unit.

The procedure for the parallel calculation of spatial derivatives for implicit schemes is sketched in figure 2.9. It is convenient to introduce a new domain partitioning, shown in blue, which is defined such that each process stores *slices* with all the field values along the differentiation direction. We distinguish the following steps:

1. Each group of processes aligned in the differentiation direction perform data exchanges in order to copy the data from the original partitioning to the partitioning defined above.
2. Once the exchange of information has been completed, each process can now compute the spatial derivatives along the differentiation direction.
3. The derivatives have been computed in the partitioning defined before, and the processes have to perform data exchanges again to copy the data back to the original partitioning.

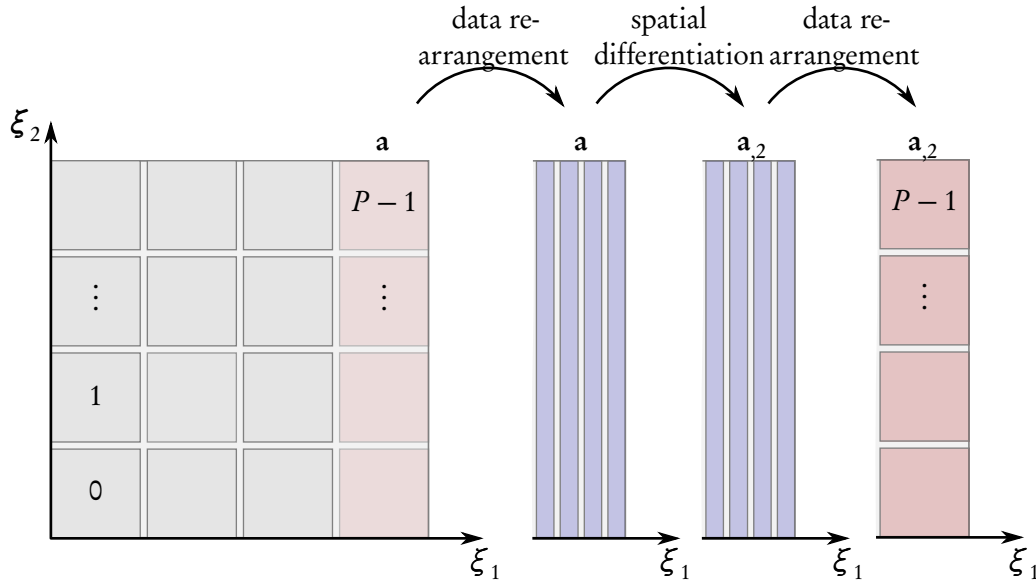


Figure 2.9: Diagram showing the parallelization strategy of the calculation of spatial derivatives in the case of implicit schemes.

If the number of processors along the differentiation direction is P_i , a total of $2 \times P_i$ communications are performed for step 1 and step 3, per spatial derivative; these communications can be implemented either point-to-point or collectively. Since $P = \prod_{i=1}^n P_i$, topology dimensions close to $P_i = \sqrt[n]{P}$, with n as the number of spatial directions, lead to a minimum number of communications.

Similarly to explicit schemes, there exist two possibilities as to the moment when the information is exchanged. On one hand, the collective routine `MPI_Alltoall` and its variants are synchronous. On the other hand, the operation can only be performed asynchronously using the point-to-point operations presented before.

2.5.4.5 Computational performance comparison: explicit vs implicit schemes

It is instructive to compare the computational performance of parallel calculations with the procedures given before for explicit and implicit schemes. Here, we focus on the differences that arise from the different parallelization strategy and the underlying communication patterns.

The time it takes to transfer a certain amount of information N over the network is, in a very simplified manner, given by $t = t_0 + bN$, where t_0 is the time it takes to establish the link and b is the network bandwidth⁵. These quantities are not only dependent on the network architecture, but also on the number of destinations, the traffic in the network at the moment of sending and the amount of information N to transfer, since vendors implement different routing algorithms depending on the size of the information packets.

⁵not to be confused with the matrix bandwidth

From the previous considerations, implicit schemes are in clear disadvantage with respect to explicit schemes, due to both a higher number of communications and a higher amount of data that is to be transferred over the network.

Although communication time cannot be avoided, there exist several solutions to mitigate the associated overheads in the calculation. First, it is desirable to group network transfers into as few as possible to avoid excessive penalization due to t_0 . Second, several computer architectures provide hardware mechanisms to allow transmission of data over the network without requiring the intervention of the central process-unit (Direct Memory Access or DMA).

2.5.4.6 Modification of the parallel calculation of implicit schemes

Even in this case, the performance of implicit schemes rapidly degrades for a large number of processing units. For this reason, we have modified the parallelized evaluation of implicit schemes, by truncating the stencil at the boundaries between the sub-domains, including an overlap region with adjacent sub-domains. Then, the evaluation of the modified compact schemes can be performed using the procedure sketched in figure 2.8 for explicit schemes. Note that the compact support is no longer the entire domain but the processor sub-domain together with the overlap region that needs to be communicated.

Such a modification obviously affects the accuracy of the scheme at the boundaries of the sub-domains. Nevertheless, from the analysis that have been performed (not shown here), it can be concluded that possible negative effects for the flow cases considered here can be mitigated by considering a sufficiently large overlap region. The use of this modification has to be used carefully when considering different flow cases.

2.6 Temporal advancement

In the previous section we have presented the spatial discretization of the equations. At this point, the original system of partial differential equations is reduced to a system of ordinary differential equations given by:

$$\frac{d\mathbf{v}}{dt} = \mathbf{F}(\mathbf{v}, t) \quad \text{with} \quad \mathbf{v}(0) = \mathbf{v}_0. \quad (2.78)$$

A popular family of schemes for integrating the above system of equations in time is the family of Runge–Kutta schemes. In these methods, a time-step from \mathbf{v}_k to \mathbf{v}_{k+1} , with the subscript k denotes the k th time step, is performed by applying a formula of the following form

$$\mathbf{v}_{k+1} = \mathbf{v}_k + \sum_{i=1}^s b_i \mathbf{k}_i, \quad (2.79)$$

with

$$\mathbf{k}_i = \Delta t \mathbf{f} \left(t_k + c_i \Delta t, \mathbf{v}_k + \sum_{j=1}^s a_{ij} \mathbf{k}_j \right). \quad (2.80)$$

It is conventional to present the above coefficients in the so-called Butcher tableau; see table 2.2. We focus on the *explicit Runge–Kutta schemes*, which satisfy $a_{ij} = 0$ for $i \leq j$. The coefficients depend on the particular Runge–Kutta method, and they are chosen according to the accuracy of the method although different criteria can also be applied.

c_1	a_{11}	a_{12}	\cdots	a_{1s}
c_2	a_{21}	a_{22}	\cdots	a_{2s}
\vdots	\vdots	\vdots	\ddots	\vdots
c_s	a_{s1}	a_{s2}	\cdots	a_{ss}
	b_1	b_2	\cdots	b_s

Table 2.2: Butcher tableau with the coefficients for a general Runge–Kutta schemes. For explicit Runge–Kutta schemes $a_{ij} = 0$ for $i \leq j$.

For our numerical code we have chosen the 4th-order accurate method RK4(3)5[2R+]C derived in Kennedy et al. (2000) and used in the previous numerical code. This method requires five function evaluations per time step; it can be classified as a low-storage method since only two temporary vectors are required for performing one time step, thus reducing the required amount of memory for the calculation. The Butcher tableau that corresponds to the above scheme is given in table 2.3.

0					
c_2	a_{21}				
c_3	b_1	a_{32}			
c_4	0	b_2	a_{43}		
c_5	0	0	b_3	a_{54}	
	b_1	b_2	b_3	b_4	b_5

Table 2.3: Butcher tableau for the explicit low-storage Runge–Kutta scheme RK4(3)5[2R+]C derived in Kennedy et al. (2000).

2.6.1 Stability region

The accuracy and stability of the Runge–Kutta methods is typically analysed by considering the time integration of the linear first-order differential equation

$$\frac{df}{dt} = \lambda f. \quad (2.81)$$

Given f_k , the approximate solution at time $t_k + \Delta t$ reads

$$f_{k+1} = (1 + p(\lambda\Delta t))f_k. \quad (2.82)$$

The above recursion remains bounded if $|1 + p(\lambda\Delta t)| \leq 1$; in figure 2.10a, we present, for the current scheme, the region of the complex plane where this inequality is satisfied for the present scheme.

In the case of an advection–diffusion equation (see section 2.5), it is required that the spectrum of the spatially discretized equation fall within the stability region. From this stability condition, a suitable time-step that leads to a stable calculation can readily be obtained. Hence

$$\frac{U\Delta t}{\Delta x} \leq \frac{3.38}{2.29} \quad \text{and} \quad \frac{\Delta t}{\text{Re}\Delta x^2} \leq \frac{4.82}{3.96} \quad (2.83)$$

which constitute the advection and viscous CFL condition, respectively. This situation is summarized in figure 2.10b.

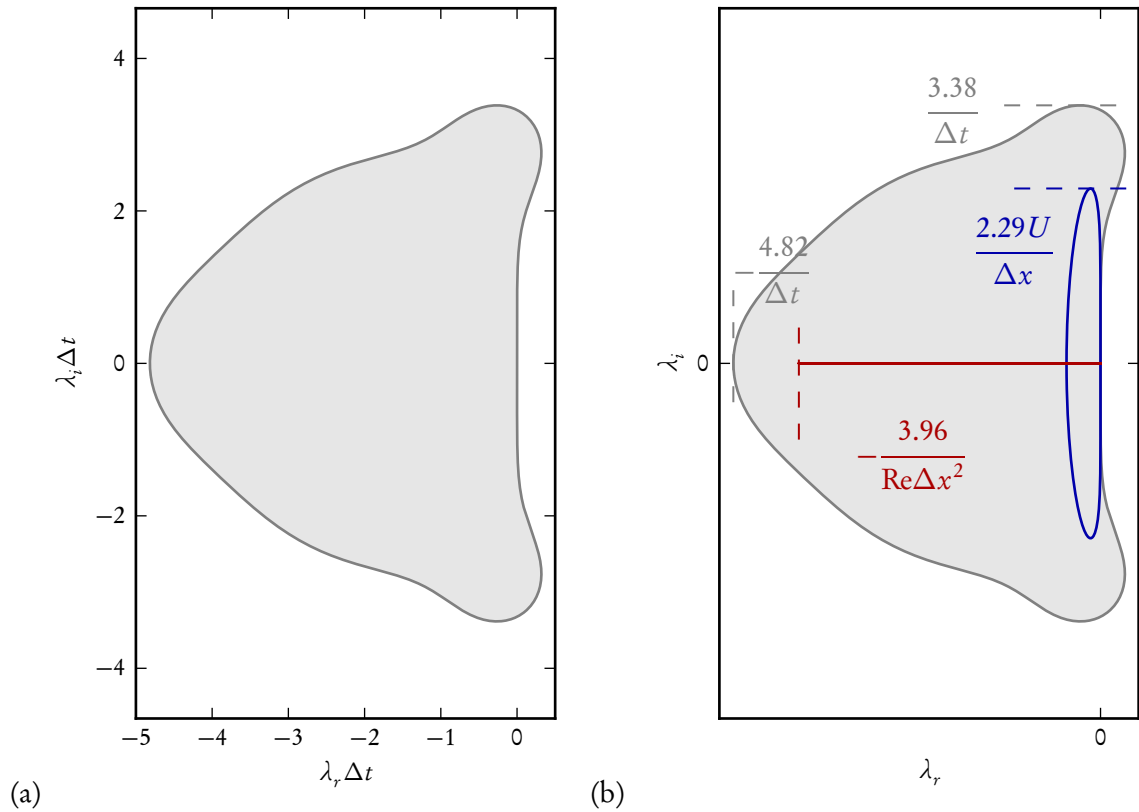


Figure 2.10: (a) stability region for the Runge-Kutta scheme RK4(3)5[2R+]C, (b) stability limits of the Runge-Kutta scheme RK4(3)5[2R+]C for an advection-diffusion process.

Note that for the compressible Navier-Stokes equations the advection speeds are given by $u \pm a/M$. Since we are interested in high-speed flow with moderate-to-high Reynolds numbers, the time step condition stemming from the advection terms is usually more restrictive than the equivalent one from the dissipative terms, thus justifying explicit time integration of the latter.

2.6.2 Implementation

The time integration using the formula given in equation (2.79) for the explicit case has been implemented in our numerical code in a time-integration class that takes as inputs the equation class as well as a suitable time step (which is estimated from the initial flow field). Note that the implementation in parallel is straightforward as it does not require exchanges between the processing units.

2.7 Validation cases

The computational code described before has been validated with the aid of exact and approximate analytic solutions of the compressible Navier-Stokes equations. In this section, we present several validation cases for two-dimensional compressible flow solvers.

Our aim is to verify the correctness of the implemented governing equations and its discretization in the general case of curvilinear grids with spatially-varying grid spacing and deformation.

In the following, we consider physical domains with a rectangular shape and dimensions L_x and L_y , discretized using a curvilinear numerical grid. The procedure followed for the construction of these numerical grid is briefly described in section 2.B. At the boundaries of the domain, non-reflecting inflow/outflow boundary conditions are imposed.

2.7.1 Exact solutions of the linearized Euler equations

We first address the evolution of small perturbations in pressure, velocity and entropy variables of infinitesimally small amplitude ϵ in a uniform flow in the inviscid limit, i.e. $\epsilon \ll 1$ and $\text{Re} \rightarrow \infty$. In this case, an analytic expression for the temporal evolution of Gaussian-shaped perturbations can readily be obtained (see Tam and Webb, 1993).

The following cases are considered:

1. Temporal evolution of a pressure perturbation.

$$p(x, y, t) = 1 + p' \left(\sqrt{(x-t)^2 + y^2}, t \right) \quad (2.84)$$

$$s(x, y, t) = 0 \quad (2.85)$$

$$u(x, y, t) = 1 + \frac{x-t}{\sqrt{(x-t)^2 + y^2}} u'_r \left(\sqrt{(x-t)^2 + y^2}, t \right) \quad (2.86)$$

$$v(x, y, t) = \frac{y}{\sqrt{(x-t)^2 + y^2}} u'_r \left(\sqrt{(x-t)^2 + y^2}, t \right) \quad (2.87)$$

with

$$p'(r, t) = \frac{\epsilon}{2} \int_0^\infty \exp\left(-\frac{\xi^2}{4}\right) \cos\left(\xi \frac{t}{M}\right) J_0(\xi r) \xi \, d\xi \quad (2.88)$$

$$u'_r(r, t) = \frac{\epsilon}{2\gamma M} \int_0^\infty \exp\left(-\frac{\xi^2}{4}\right) \sin\left(\xi \frac{t}{M}\right) J_1(\xi r) \xi \, d\xi \quad (2.89)$$

2. Temporal evolution of a divergence-free velocity perturbation.

$$p(x, y, t) = 1 \quad (2.90)$$

$$s(x, y, t) = 0 \quad (2.91)$$

$$u(x, y, t) = 1 - \frac{y}{\sqrt{(x-t)^2 + y^2}} u'_\theta \left(\sqrt{(x-t)^2 + y^2} \right) \quad (2.92)$$

$$v(x, y, t) = \frac{x-t}{\sqrt{(x-t)^2 + y^2}} u'_\theta \left(\sqrt{(x-t)^2 + y^2} \right) \quad (2.93)$$

with

$$u'_\theta(r) = r \exp(-r^2) \quad (2.94)$$

3. Temporal evolution of an entropy perturbation.

$$p(x, y, t) = 1 \quad (2.95)$$

$$s(x, y, t) = \epsilon \exp\left(-((x - t)^2 + y^2)\right) \quad (2.96)$$

$$u(x, y, t) = 1 \quad (2.97)$$

$$v(x, y, t) = 0 \quad (2.98)$$


In the above, we have taken $M = 0.4$, $\gamma = 1.4$, $\text{Re} = 10^{10}$, $\text{Pr} = 1$; the amplitude of the perturbations is $\epsilon = 10^{-4}$. The computational domain extends over $L = L_x = L_y = 64$ unit lengths. The grid spacing is uniform equal to $\Delta x = \Delta y = 0.5$. Finally, the deformation of the grid is given by $A_x = A_y = 0.375$, and $\alpha = \beta = 8\pi/L$.

A comparison between the numerical result and the analytical solution for the validation cases presented before is displayed in figure 2.11a-c.


2.7.2 Sound generated by a co-rotating vortex pair

Typical aeroacoustic investigations involve flow phenomena with very different spatio-temporal scales: the amplitude and the wavelength of both hydrodynamic and acoustic fluctuations differ by several orders of magnitude. As an example of validation case that present these features, we have considered the sound generated by a co-rotating vortex pair. The details from the validation case set up are detailed in Kim (2013), and an analytical estimate for the amplitude of the pressure fluctuations in the far-field are given in (Mitchell et al., 1995).


We have taken $M = 1$, $\gamma = 1.4$, $\text{Re} = 10^{10}$ and $\text{Pr} = 1$; the circulation of the co-rotating vortices $\Gamma = \pi^2/10$ and the rotation frequency is $\Omega = \pi\Gamma M/4$. The computational domain extends over $L = L_x = L_y = 766$ unit lengths. The grid spacing close to the origin is $\Delta x_O = \Delta y_O = 0.05$, and in the far-field is $\Delta x = \Delta y = 4$. Finally, the deformation of the grid is given by $A_x = A_y = 0.375$, and $\alpha = \beta = 16\pi/L$. In addition to non-reflecting boundary conditions, a sponge layer extending over $L_s = 40$ has been introduced at the boundaries of the domain. The amplitude of the sponge layer has been taken $\sigma = 0.1$. A typical instantaneous snapshot of the pressure is shown in figure 2.12a-b, and a comparison between the numerical solution and the above analytical estimate is displayed in 2.12.



(a)



(b)



(c)

Figure 2.11: Comparison between the analytical and numerical solution for the temporal evolution of small-amplitude disturbances in (a) pressure, (b) entropy and (c) vorticity.

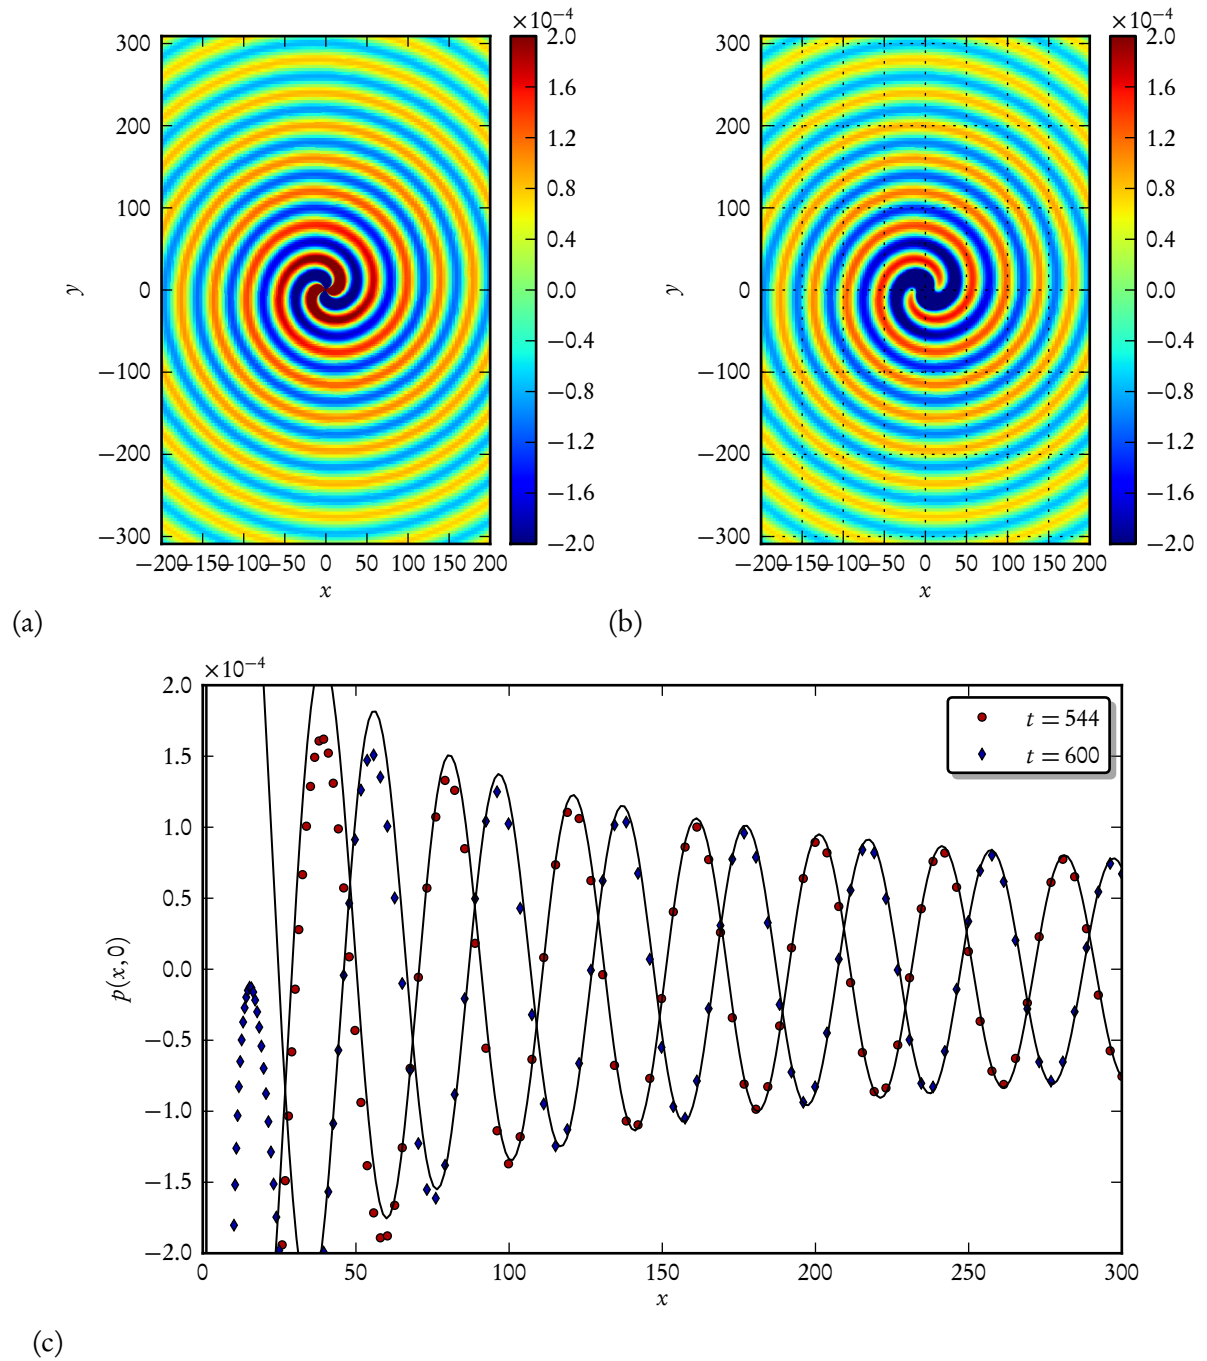


Figure 2.12: Sound radiated by a co-rotating vortex pair showing comparison between the (a) analytical and (b) the numerical solution. (c) Comparison along the x-axis for different times.

2.A Spatial differentiation and temporal advancement schemes: coefficients

k	-2	-1	0	1	2
$a_{1,1+k}$			1	-3.0607612480196	-9.2825687875052
$a_{2,2+k}$		-1.3841732050877193	1	6.620636765789473684	-2.12293497649122807
$a_{i,i+k}$	0.027844800835977185	0.45586882948901069	1	0.477479190809459099	0.038825537571973596
$a_{n-1,n-1+k}$	15.74839510877193	27.04931558947368	1	-0.3211204877192982	
$a_{n,n+k}$	4.074311507053876	6.583618303429768	1		

Table 2.4: Coefficients for the CULD (compact upwind with low dissipation) scheme.

k	-2	-1	0	1	2
$b_{1,3+k}$	-2.0916904202872	12.7390135650198	-7.5911418720294	-3.324665234327	0.2684839616238
$b_{2,3+k}$	3.3623251557894737	-11.84525148578947	8.83692208	0.29908109	-0.65307684
$b_{i,i+k}$	-0.118727423955	-0.7482059425881687	0.00052666822731662	0.7184558284234524	0.1479508698923997
$b_{n-1,n+k-2}$	-1.8466026	-30.228379	21.1592312	12.27649478947368	-1.36074438947368
$b_{n,n+k-2}$	0.040891100302009	-0.7577318529877	-6.875427455144652	4.202556257560724	3.389711950269619

Table 2.5: Coefficients for the CULD (compact upwind with low dissipation) scheme (cont).

k	-1	0	1
$a_{1,1+k}$		2	4
$a_{2,2+k}$		4	1
$a_{i,i+k}$	1	3	1
$a_{n-1,n+k-1}$	1	4	
$a_{n,n+k}$	4	2	

Table 2.6: Coefficients for the centred compact scheme.

k	-2	-1	0	1	2
$b_{1,1+k}$			-5	4	1
$b_{2,2+k}$		-3	0	3	
$b_{i,i+k}$	$-1/12$	$-7/3$	0	$7/3$	$1/12$
$b_{n-1,n+k-1}$		-3	0	3	
$b_{n,n+k}$	-1	-4	5		

Table 2.7: Coefficients for the centred compact scheme (cont).

a_{21}	$\frac{970286171893}{4311952581923}$	b_1	$\frac{1153189308089}{22510343858157}$
a_{32}	$\frac{6584761158862}{12103376702013}$	b_2	$\frac{1772645290293}{4653164025191}$
a_{43}	$\frac{2251764453980}{15575788980749}$	b_3	$\frac{-1672844663538}{4480602732383}$
a_{54}	$\frac{26877169314380}{34165994151039}$	b_4	$\frac{2114624349019}{3568978502595}$
		b_5	$\frac{5198255086312}{14908931495163}$

Table 2.8: Runge-Kutta coefficients (cont).

Table 2.9: Runge-Kutta coefficients (cont).

2.B Numerical grid used for the validation cases

The numerical grids used for the validation cases described in section 2.7 is generated in two steps.

First, an intermediate rectilinear grid $\mathbf{x}'(\xi)$ with non-uniform grid spacing (see figure 2.13a) is generated with the aid of the stretching functions defined in section 3.2.2.1. Close to the origin, the spacing along the x -direction and the y -direction is denoted by Δx_O and Δy_O , respectively. Far from the origin, the maximum grid spacing is Δx and Δy . The domain size is given by L_x and L_y ; the maximum stretching has been set to $f_{\max} = 0.04$. Then, the final numerical grid $\mathbf{x}(\xi)$ is obtained from the previous rectilinear grid using the following relations:

$$x(\xi, \eta) = x'(\xi, \eta) + A_x \sin(\alpha\xi) \sin(\beta\eta), \quad (2.99)$$

$$y(\xi, \eta) = y'(\xi, \eta) + A_y \sin(\alpha\xi) \sin(\beta\eta). \quad (2.100)$$

In the above relations, the parameters A_x, A_y, α, β characterize the deformation with respect to the intermediate grid, and are chosen on a case by case basis. An example of the numerical grid used in the validation cases is shown in figure 2.13b.




Figure 2.13: Generation of numerical grids for the validation cases showing (a) the intermediate rectilinear grid and (b) the final curvilinear grid.

Nonlinear simulations

3.1 Problem description

Our study concerns an aerofoil section immersed in an unbounded domain and subjected to a uniform flow, as sketched in figure 3.1. The geometry of the problem is readily defined using an inertial reference frame with respect to which the aerofoil remains at rest: the origin of the coordinate system coincides with the leading edge, the x_1 -axis is aligned with the trailing edge, and the x_2 -axis is perpendicular to the former following the right-hand rule.




Figure 3.1: Sketch of the problem under consideration.

The aerofoil flow obeys the two-dimensional compressible Navier–Stokes equations in non-dimensional form; the reference values for the flow variables are those of the unperturbed flow¹, the reference length-scale is the aerofoil chord, and the temporal scale is the passage time over one unit length. The free-stream velocity \mathbf{u} has norm one, and the free-stream values for the pressure and the entropy are $p = 1$ and $s = 0$, respectively. In this frame of reference, the free-stream velocity vector forms an angle α , or angle of attack, with the x_1 -axis. The aerofoil surface is considered adiabatic and a no-slip wall boundary condition applies, i.e. $\mathbf{q} \cdot \mathbf{n} = 0$ and $\mathbf{u} = 0$.

¹except for the entropy, where the ideal gas constant r has been taken as the reference value

3.2 Computational set-up

The aim of this section is to describe the set-up for the nonlinear simulations. Although the attention is focused on two-dimensional configurations, the numerical set-up can easily be extended to three dimensions.

There are several prerequisites for the simulation of the above-described flow configuration with the numerical code presented in chapter 2: the physical domain under consideration has to be truncated at a *far* distance from the aerofoil; furthermore, the mapping between the physical and computational domain has to be established.

3.2.1 Physical and computational domains: boundaries

We have considered a C-type grid for the physical domain, whose boundary is determined by the curves \mathbf{x}_{E_i} , with $i = 1, \dots, 4$, and a unit-square computational domain given by $0 \leq \xi^k \leq 1$. Both domains and their corresponding boundaries are illustrated in figure 3.2.

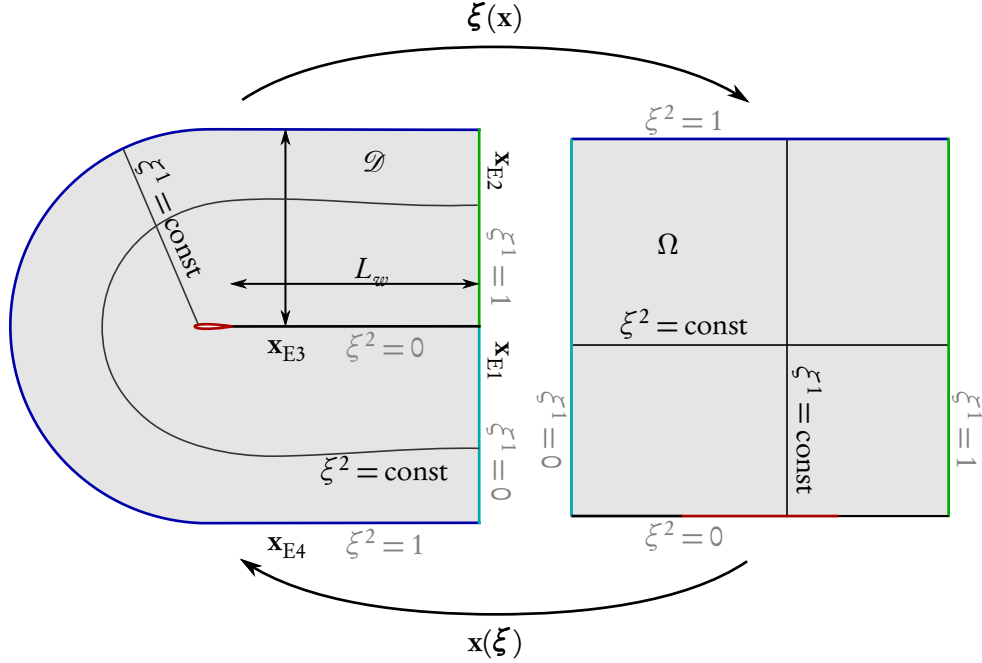


Figure 3.2: C-grid (physical domain, left) and unit square (computational domain, right) showing the correspondence between the boundaries. The C-grid is delimited by the curves \mathbf{x}_{E_i} . The downstream outlet given by \mathbf{x}_{E1} and \mathbf{x}_{E2} , at a distance L_w from the trailing edge, is mapped to the unit-square edges given by $\xi^1 = 0$ and $\xi^1 = 1$, respectively. The aerofoil surface together with the wake are described by the curve \mathbf{x}_{E3} , which corresponds to $\xi^2 = 0$. Note that the wake is traversed twice. Finally, the far-field boundary, at a distance L_n from the aerofoil is given by \mathbf{x}_{E4} ; it is mapped to $\xi^2 = 1$.

The outlet is located downstream at a distance L_w from the aerofoil trailing edge and is represented by the curves \mathbf{x}_{E1} and \mathbf{x}_{E2} defined by

$$(x_{E1}, y_{E1}) = (1 + L_w, -s), \quad \text{with } 0 \leq s \leq L_n, \quad (3.1)$$

and

$$(x_{E2}, y_{E2}) = (1 + L_w, s), \quad \text{with } 0 \leq s \leq L_n, \quad (3.2)$$

respectively. The curve \mathbf{x}_{E1} (\mathbf{x}_{E2}) corresponds to the computational-domain boundary given by $\xi^1 = 0$ ($\xi^1 = 1$) and $0 \leq \xi^2 \leq 1$.

The aerofoil section, placed at $0 \leq x \leq 1$, and the segment joining the trailing edge with both \mathbf{x}_{E1} and \mathbf{x}_{E2} , called *wake connection*, are described by the curve \mathbf{x}_{E3} :

$$(x_{E3}, y_{E3}) = \begin{cases} (1 + L_w - s, 0) & 0 \leq s < L_w \\ (x_a(s - L_w), y_a(s - L_w)) & 0 \leq s - L_w \leq L_a \\ (1 + s - (L_w + L_a), 0) & 0 \leq s - (L_w + L_a) \leq L_w \end{cases}. \quad (3.3)$$

This curve is mapped to the computational-domain boundary given by $\xi^2 = 0$ and $0 \leq \xi^1 \leq 1$. In the above, $(x_a(t), y_a(t))$ is a parametrization of the aerofoil section using the arc-length t measured from the trailing edge, and L_a is the length of the curve that describes the aerofoil surface. Note that the wake and the trailing edge are included *twice*; we refer to the first and second segments as the *lower* and *upper* wake connection, respectively. It is important to remark that the wake connection requires a particular treatment in the code that will be addressed later.

Finally, the far-field boundary \mathbf{x}_{E4} is given by a sufficiently smooth curve at a distance L_n from \mathbf{x}_{E3} . It is defined as

$$(x_{E4}, y_{E4}) = \begin{cases} (1 + L_w - s, -L_n) & 0 \leq s < 1 + L_w \\ \left(-L_n \sin\left(\frac{s - (1 + L_w)}{L_n}\right), -L_n \cos\left(\frac{s - (1 + L_w)}{L_n}\right) \right) & 0 \leq s - (1 + L_w) \leq \pi L_n \\ (s - (1 + L_w + \pi L_n), L_n) & 0 < s - (1 + L_w + \pi L_n) \leq 1 + L_w \end{cases}, \quad (3.4)$$

and it is mapped to the computational domain boundary $\xi^2 = 1$ and $0 \leq \xi^1 \leq 1$.

In the above, the variable s refers to the arc-length parameter. The distances at which the far field is truncated in the normal and in the downstream direction are given by the parameters L_n and L_w , respectively.

3.2.2 Grid generation

We consider the physical and the computational domain in discrete form. The computational domain is discretized uniformly along the ξ^1 - and ξ^2 -axis using I and J points, respectively.

In the construction of the mapping $\mathbf{x}(\xi)$, we distinguish two steps. First, we determine the grid mapping between the grid points on the *boundaries* of the physical and of the computational domain. Second, the grid points in the *interior* of the physical domain are mapped to the interior grid points of the computational domain. Note that in contrast to the computational domain, the grid point distribution on the physical domain is non-uniform.

3.2.2.1 Boundary grid points: stretching function

At the boundary of the unit square, each grid node, given by: $\xi_i^n = 0$ or $\xi_i^n = 1$, and $\xi_i^t = (i-1)/(I-1)$, is mapped to the C-grid boundary node $\mathbf{x}_E(s_i)$, with $s_i = s(\xi_i^t)$. The function s_i is commonly referred to as a stretching function and allows a user-specified grid-point density adapted to the expected gradients of the fluid variables.

It is important to note that the variation of the grid spacing has to be performed cautiously to avoid a decline in the spatial resolution properties of the numerical schemes. For this reason, we define the stretching function in terms of the relative rate of increase (or decrease) of the grid spacing between two consecutive grid points:

$$f_i = \frac{\Delta s_i - \Delta s_{i-1}}{\Delta s_{i-1}}, \quad (3.5)$$

from where we obtain the grid spacing $\Delta s_i = s_i - s_{i-1}$:

$$\Delta s_i = (1 + f_i)\Delta s_{i-1}, \quad (3.6)$$

and finally, the arc-length parameter values s_i :

$$s_{i+1} = s_i + f_i (s_i - s_{i-1}). \quad (3.7)$$

The above recursion relation is fully determined by specifying $s_1 = 0$ as the starting point of the curve and Δs_1 as the initial grid spacing. For $f_i > 0$ ($f_i < 0$) we have progressively diverging (clustered) grid points.

In particular, we use for the present case a linear combination of stretching rates of the form

$$f_i \equiv f(\xi_i, \xi_s, \xi_e, w_s, w_e, f_{\max}) = \frac{f_{\max}}{4} \left[1 + \operatorname{erf} \left(6 \frac{\xi_i - \xi_s}{w_s} \right) \right] \left[1 - \operatorname{erf} \left(6 \frac{\xi_i - \xi_e}{w_e} \right) \right], \quad (3.8)$$

which allows for smooth variations of the grid spacing.

The stretching rate, the corresponding grid spacing and the final grid point distribution are illustrated in figure 3.3: the parameters ξ_s and ξ_e determine the extent of the stretching region, and the parameters w_s and w_e give the number of points over which the stretching passes from 0 to f_{\max} and from f_{\max} to zero, respectively. Following the recommendations of Mitchell et al. (1995), we set the maximum stretching rate to $-0.04 \leq f_{\max} \leq 0.04$.

The determination of the grid spacings relies on estimates of the spatial scale of the main flow features that need to be resolved for a given flow case. A discussion for the present case is postponed to section 3.3, and, for the time being, it shall be assumed that the stretching functions have already been determined and that the stretching function for the downstream outlets (\mathbf{x}_{E1} and \mathbf{x}_{E2}), denoted by $s_{n,i}$, is the same.

3.2.2.2 Interior grid points: linear interpolation

The ξ^1 -constant grid-lines in the interior of the domain are chosen as straight segments connecting the curves \mathbf{x}_{E3} and \mathbf{x}_{E4} . The grid stretching along such segments is equal to the grid stretching at the downstream outlets $s_{n,i}$. Hence

$$\mathbf{x}_{ij} = \mathbf{x}_{E3,i} + (\mathbf{x}_{E4,i} - \mathbf{x}_{E3,i}) \frac{s_{n,j}}{L_n}. \quad (3.9)$$

Unfortunately, since \mathbf{x}_{E3} is non-differentiable at the trailing edge, the induced metric is discontinuous. While the former cannot be avoided, the numerical grid can be *smoothed* at a certain distance from the trailing edge so as to render the grid metric continuous and differentiable. In practice, this is performed by applying a second-order filter to the grid mapping $\mathbf{x}(\xi)$.

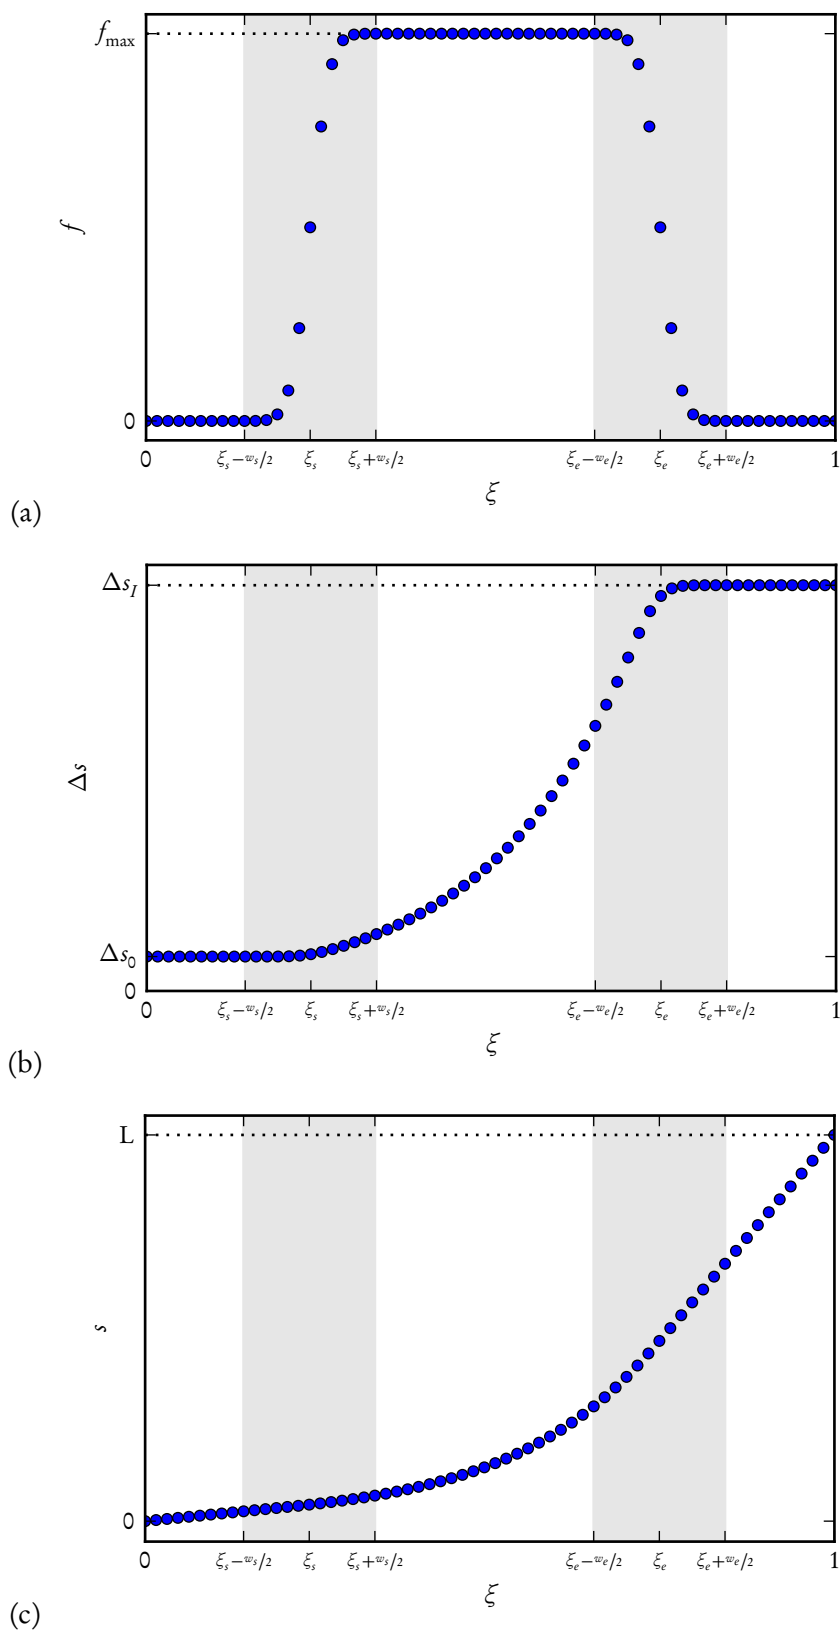


Figure 3.3: Grid stretching function showing (a) stretching rate, equations (3.5) and (3.8), the resulting (b) grid spacing, equation (3.6), and (c) the distribution of the arc-length parameter equation (3.7).

3.2.3 Boundary and initial conditions

The simulation of the flow requires the specification of the boundary conditions and an appropriate initial flow field.

The practical implementation of the boundary conditions follows the formalism presented in section 2.4. On the aerofoil surface, an adiabatic no-slip boundary condition applies. At the far-field and downstream outlets, the boundary conditions correspond to the non-reflecting inflow or outflow boundary condition depending on the sign of the boundary-normal velocity component.

3.2.3.1 Sponge layer

Since sound generation is the main feature of interest in our flow case, it is imperative to ensure that the imposed boundary conditions do not reflect spurious acoustic waves back into the computational domain. In this respect, the non-reflecting outflow boundary conditions prove insufficient at the downstream outlets, where vortical structures exit the domain; for this reason, a sponge layer (see section 2.4) had to be included to attenuate the reflections.

The value of the sponge-layer forcing-term strength σ is zero for $x < L_w + 1 - L_s$; for $x > L_w + 1 - L_s$ it reads:

$$\sigma(x) = A \left\{ 1 + \exp \left(\frac{L_s}{x - (L_w + 1)} + \frac{L_s}{x - (L_w + 1 - L_s)} \right) \right\}^{-1}, \quad (3.10)$$

where A is a free parameter and L_s is the length of the sponge layer measured from the downstream outlet. The specified initial condition is taken as the reference state \mathbf{v}_{ref} .

3.2.3.2 Initial flow field

Finally, an appropriate initial condition has to be chosen in order to start the numerical computation. We typically consider high Reynolds-number and moderate Mach-number flows; In this case, a reasonable guess for the initial velocity field can be obtained from the incompressible inviscid limit using a panel method. The pressure is obtained from the conservation of the total enthalpy, and the entropy is set to zero. Furthermore, the resulting velocity field has been gradually decreased to zero close to the aerofoil wall to comply with the no-slip boundary condition.

3.2.4 Wake connection

For the practical calculation of spatial derivatives across the wake, i.e. along the ξ^2 -direction, we have adapted the procedure described in section 2.5 to correctly take into account the values of the flow field on the opposite side of the wake connection. At every time step, the points in the wake and at the trailing edge are replicated on the upper and lower side of the wake line, and their value is averaged during the temporal evolution.

3.3 Simulations

3.3.1 Configuration

We consider a symmetric NACA 0012 aerofoil section at 2 degrees of incidence subjected to a uniform flow at Reynolds number 200 000 and at Mach number 0.4. The Prandtl number is taken as 0.71 and

	Angle of attack	Mach number	Reynolds number
present case	2 deg	0.4	200 000
Desquesnes et al. (2007)	2 deg	0.1	200 000

Table 3.1: Details of the case under consideration. The aerofoil section for all cases is a NACA 0012. The flow parameters from Desquesnes et al. (2007) correspond to the case that exhibits tonal noise (case 1).

the heat capacity ratio is 1.4; the flow has constant viscosity and constant heat conductivity. In addition, the volume viscosity is set to zero $\mu_v = 0$ (following Stokes' assumption).

This particular choice of symmetric aerofoil section and flow parameters is motivated by the numerical study of Desquesnes et al. (2007), see table 3.1, where it was found that two-dimensional simulations exhibit similar features to the experiments reported in Nash et al. (1999). Note, however, that a higher Mach number has been chosen here to allow for larger time steps in our numerical calculations.

It is instructive to determine the physical parameters that would correspond to an experiment under atmospheric conditions. Considering the gas constant r_g equal to $287 \text{ J kg}^{-1} \text{ K}^{-1}$, an ambient temperature of 298.15 K and a kinematic viscosity equal to $15.72 \cdot 10^{-6} \text{ m}^2 \text{ s}^{-1}$, the governing parameters would describe an aerofoil configuration of 2.27 cm chord subjected to a free stream speed of 148.44 ms^{-1} .

3.3.1.1 Modified NACA 4-digit aerofoil sections

The aerofoil sections of the NACA 4-digit family have finite thickness at the trailing edge. However, our attention is focused on aerofoils with sharp trailing edges, and for this reason a modified profile has been considered instead. The thickness of the modified aerofoil section reads

$$y_{\pm}(x) = \pm \frac{t}{0.2} \left(0.296375\sqrt{x} - 0.12635x - 0.35195x^2 + 0.283775x^3 - 0.10185x^4 \right),$$

with $0 \leq x \leq 1$.

(3.11)

In the above, the parameter $t = 0.12$ is the maximum aerofoil thickness, which is reached at $x = 0.3$, and the radius of curvature at the leading edge is $R_{le} = 1.1t^2$.

3.3.2 Estimation of spatial scales

In order to define the numerical grid, it is crucial to identify the flow features that determine the length scales that need to be resolved in the calculation: the boundary-layer thickness at the leading edge, the smallest spatial scale in the boundary layer at the trailing edge and the acoustic wave-length in the far field of the highest tone in the sound spectrum determine the grid spacing at each respective location.

The leading-edge boundary-layer thickness at the stagnation point is approximately equal to that of a hypothetical cylinder with the same radius of curvature as the aerofoil. For the present case, the radius of curvature R_{le} and the thickness δ_{le} (Schlichting and Gersten, 2004) are

$$R_{le} = 0.016, \quad \text{and} \quad \delta_{le} \approx 2R_{le}^{1/2} \text{Re}^{-1/2} \approx 5 \cdot 10^{-4},$$
(3.12)

respectively. They determine the spatial scales to be resolved at the leading edge in the tangential and in the normal direction.

At the trailing edge, a rough estimation of the lower bound of the boundary layer thickness is given by:

$$\delta_{te} \approx 4.91\text{Re}^{-1/2} \approx 10^{-2}. \quad (3.13)$$

It is known from experiments that the non-dimensional frequency of the main tone can be roughly estimated by the empirical relation $f \approx 0.011\text{Re}^{1/2} \approx 5$; in fact, Desquesnes et al. (2007) found in their simulation that the acoustic spectrum is dominated by discrete frequencies centred around $f \approx 7$. From the relation $\lambda f = c$, with c as the phase speed, we can estimate the wavelength associated with the propagating acoustic waves in the far-field as well as the wavelength associated with the hydrodynamic instabilities on the aerofoil surface. The wavelengths of acoustic waves propagating upstream and downstream follow as

$$\lambda_{ac}^- = \left| 1 - \frac{1}{M} \right| f^{-1} \approx 0.21, \quad \text{and} \quad \lambda_{ac}^+ = \left| 1 + \frac{1}{M} \right| f^{-1} \approx 0.5. \quad (3.14)$$

Typical phase speeds of wall-bounded hydrodynamic instabilities are in the range 0.3–0.5; by considering the lower bound $c = 0.3$, we obtain

$$\lambda_{bl} = c_{bl} f^{-1} \approx 0.04. \quad (3.15)$$

3.3.3 Numerical grids

With the above considerations, we have defined four numerical grids with increasing grid resolution and domain size. The main parameters are given in table 3.2 and in table 3.3: coarse grid (G1), refined grid (G3), intermediate resolution (G2) and large domain size (G1S2).

Case	Grid spacing					Domain size	
	w_{le}	w_{te}	$h_{le} = h_{te}$	w_{far}	h_{far}	L_w	L_n
G1	$1.62 \cdot 10^{-3}$	$1.50 \cdot 10^{-3}$	$5.00 \cdot 10^{-4}$	$6.00 \cdot 10^{-3}$	$2.70 \cdot 10^{-2}$	7.55	7.34
G1S2	$1.62 \cdot 10^{-3}$	$1.50 \cdot 10^{-3}$	$5.00 \cdot 10^{-4}$	$6.00 \cdot 10^{-3}$	$2.70 \cdot 10^{-2}$	12.62	11.41
G2	$1.08 \cdot 10^{-3}$	$1.00 \cdot 10^{-3}$	$3.33 \cdot 10^{-4}$	$4.00 \cdot 10^{-3}$	$1.80 \cdot 10^{-2}$	7.63	7.06
G3	$7.22 \cdot 10^{-4}$	$6.67 \cdot 10^{-4}$	$2.22 \cdot 10^{-4}$	$2.67 \cdot 10^{-3}$	$1.20 \cdot 10^{-2}$	7.68	7.66

Table 3.2: Numerical grids and simulation parameters.

The grid-point density along the curves \mathbf{x}_{E1} and \mathbf{x}_{E2} determines the grid resolution in the normal direction of the boundary layers and the near wake. In figure 3.4, we display the normal grid spacing Δs_n versus the arc-length s_n . The gradients of the flow variables are concentrated in a layer of thickness $\delta_{te} \approx 10^{-2}$ from the wall: in this zone, we have included 20 points for the coarsest grid (G1), and 45 points for the finest grid (G3). Far from the wall, the grid spacing is restricted by the wavelength of the acoustic waves propagating upstream. The grid resolution in the far-field ranges from 8 points per wavelength for the coarsest grid (G1) to 18 points for the finest grid (G3). The transition from the near-wall spacing to the far-field spacing is carried out from $s = 10^{-2}$ to $s = 1$. As previously indicated, the maximum stretching rate is 0.04.

The normal grid spacing may seem insufficient when compared to the boundary-layer thickness near the stagnation point: 1 to 2 points are included in this region for the coarsest and the finest grid,

	Relaxation			Timestepping			Spatial discretization		
	L_s	A	η^\pm	CFL	Δt	T	n	m_a	m_w
G1	2	3	5.0	1.2	$1.79 \cdot 10^{-4}$	150	384	1280	1280
G1S2	2	3	2.5	1.2	$1.79 \cdot 10^{-4}$	50	576	1280	1920
G2	2	3	5.0	1.2	$1.19 \cdot 10^{-4}$	50	512	1920	1920
G3	2	3	5.0	1.2	$7.87 \cdot 10^{-5}$	50	768	2880	2880

Table 3.3: Numerical grids and simulation parameters (continuation).

respectively. In fact, the boundary layer at the leading edge does not play a role in the phenomenon under investigation: the Reynolds number $\text{Re}_{\delta_{le}}$ is 150, and it can be anticipated that in that region the flow is stable. The time step for the entire calculation is proportional to the smallest grid spacing, which occurs at precisely this location. Since our goal is to achieve the largest possible time step, we have kept the number of grid points at this location to a minimum. During the simulation, numerically-induced waves may appear, but they are rapidly dissipated inside the boundary layer and do not propagate further downstream, thus ensuring that the accuracy of the relevant parts of the simulation is not degraded.

The tangential grid spacing at the leading edge is determined by the radius of curvature $R_{le} = 0.016$; in this region, we have included 10 points (within a distance R_{le}) for the coarsest grid (G1) and 22 points for the finest grid (G3). Downstream of the leading-edge region, the tangential grid spacing is limited by the boundary layer instabilities at the above-mentioned frequency and possibly higher harmonics. For this reason, we have included near the trailing edge 26 points per wavelength for the coarsest grid (G1) and 60 points for the finest grid (G3). Note that with this resolution, wavelengths three times shorter than that of the hydrodynamic instabilities would still be well resolved on the coarsest grid.




Figure 3.4: (a) Normal and (b) tangential grid spacing.

In the near wake, vortex shedding occurs, and the grid requirements in the normal and tangential direction are comparable. In view of this, the grid spacing has been kept constant from the trailing edge $x = 1$ to $x = 1.75$. Further downstream, the grid spacing is increased. Note that the loss of resolution downstream does not negatively influence the quality of the simulation.

With respect to the far-field boundary \mathbf{x}_{E4} there is no preferred zone where grid points need to be clustered. Therefore, we have specified the grid points, by convenience, from the grid-point distribution along the curve \mathbf{x}_{E3} as follows: for $x > 0.3$, downstream of the maximum aerofoil thickness, the grid points on \mathbf{x}_{E4} are given by the intersection between the vertical line that passes by each grid point on \mathbf{x}_{E3} and the curve \mathbf{x}_{E4} ; for $x < 0.3$, the grid points are given by the intersection of a line normal to the curve \mathbf{x}_{E3} and the far-field boundary \mathbf{x}_{E4} .

An illustration of the final grid is given in figure 3.5.

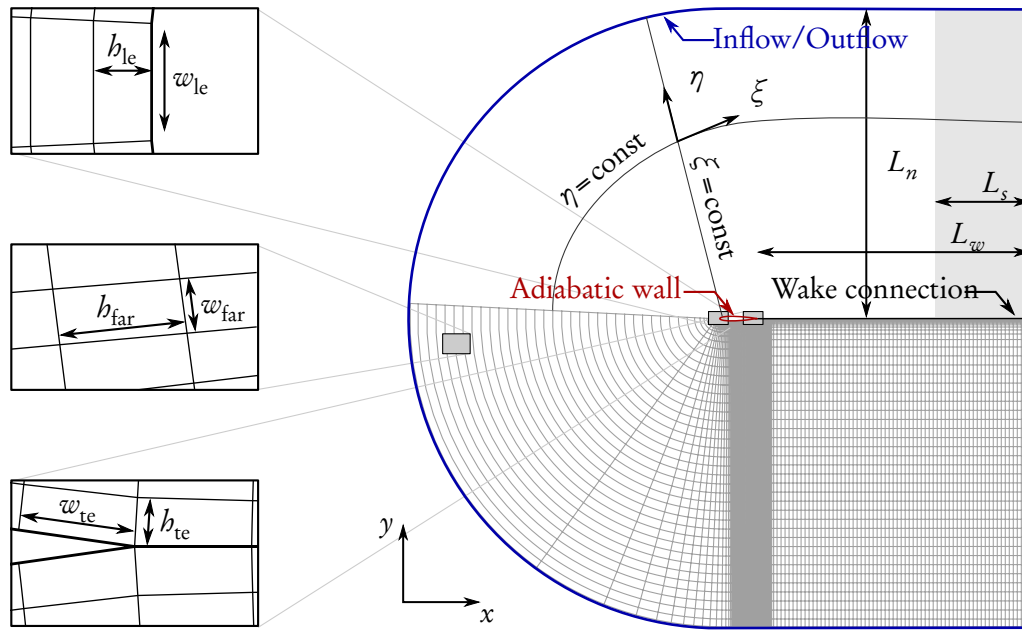


Figure 3.5: Numerical grid used in the present calculations. For a given aerofoil section, the domain is defined by the length of the wake L_w and the normal direction L_n ; a sponge layer of length L_s is added at the downstream outflow boundary. The boundary is discretized using m_a points along the aerofoil, m_w points along the wake and n points along the normal direction. The grid spacing at the leading edge (h_{le} , h_{le}), trailing edge (h_{te} , w_{te}) and far field (h_{far} , w_{far}) characterize the grid point distribution.

For all the simulations, the sponge layer strength parameter A is taken as 3, which is found to be sufficiently high to ensure that vortical structures are damped as they reach the boundary and small enough to avoid the generation of spurious radiation during the damping process.

3.3.4 Simulation time

Our flow case has been simulated using the previously defined numerical grids. The time step is given by the CFL number with respect to the smallest grid spacing in the physical domain: the grid size at the leading edge. In view of the results given in figure 2.10, a CFL number equal to 1.2 is found to be sufficient for a stable and accurate integration in time.

The computations have been run from the initial condition described in section 3.2 until a statistical steady-state has been established. For each simulation, the initial transient dynamics, representing

approximately 10 time units, have been discarded. The subsequent flow evolution has been integrated over $T_s = 150$ time units for case G1, and over $T_s = 50$ time units for the remaining cases. We have stored snapshots of the flow field evolution every $\Delta t = 0.01$, yielding a total of 15000 snapshots for case G1, and 5000 snapshots for the remaining cases. The purpose of using different numerical grids and different integration times is to analyse their influence in the observations made here.

In the following (unless explicitly stated otherwise) the presented results correspond to case G3.

3.4 Temporal evolution

The aim of this section is to provide a qualitative description of the flow features that are observed in the nonlinear simulations. In figure 3.6, the flow field at $t = 30$ is illustrated by the instantaneous vorticity $\omega_z = \partial_x v - \partial_y u$, and the instantaneous dilatation $\partial_x u + \partial_y v$.




Figure 3.6: Instantaneous flow field showing (a) vorticity levels and (b) instantaneous dilatation contours once a quasi-periodic regime is established.

The vorticity field (figure 3.6a) allows us to visualize the boundary layers that develop around the aerofoil as well as a quasi-periodic vortex shedding into the wake. The vorticity field displays the growth and the separation of the boundary layer on the suction surface with later reattachment at approximately mid-chord. Downstream of the reattachment point, vortices with a negative vorticity component are continuously emerging, driven by the outer flow, before they are ultimately shed into the wake at the trailing edge. These vortices will be labelled SSV (suction-surface vortices). We also notice that the boundary layer that develops on the pressure surface detaches over the last quarter of the chord, leading to a thin region of reverse flow that extends to the trailing edge, where vortices with a positive vorticity component appear. These latter vortices will be labelled PSV (pressure-surface vortices). In addition, secondary vortical structures with a considerably smaller spatial extent, and alternating positive/negative vorticity, arise at the trailing edge.

The dilatation field (figure 3.6b) reveals the presence of acoustic waves in the far-field at a characteristic wavelength² $\lambda \approx 0.5$, emanating from the near-wake; they are in counter-phase between the upper and lower half-plane of the aerofoil. In addition, secondary acoustic-waves in the far-field with a significantly smaller wave-length of $\lambda \approx 0.05$ and a considerably smaller amplitude can be observed.

Although the flow in this study is not strictly periodic, it is informative to visualize a representative flow-field sequence over one vortex-shedding period. The flow evolution from $t = 30.00$ to $t = 30.20$, shown every 0.05 time units, is depicted in figure 3.7. We focus on the flow evolution near the trailing edge (top and middle row) and the corresponding far-field (bottom row).

The evolution of the vorticity field in the near wake, figure 3.7 (top row) shows the passage of a SSV by the trailing edge, and the detachment of a PSV from the separation bubble on the pressure side. Although the succession of snapshots presented here suggests that the passage of a SSV by the trailing edge entails the shedding of a PSV, a detailed analysis of the temporal series affirms that this may not necessarily be the case. Although the shedding of either vortex occurs at approximately the same frequency, their phase varies during the simulation. The visualization of the vorticity field suggests that the secondary vortices arise from the interaction of SSV and PSV with the geometric singularity of the trailing-edge.

From the inspection of the dilatation field near the trailing edge, figure 3.7 (middle row), we observe that the emission of acoustic waves occurs at the same frequency as the shedding of vortical structures. The emission of the secondary acoustic waves is observed immediately after the formation of the secondary vortices; see figure 3.7(c,h).

The far-field dilatation field, figure 3.7 (bottom row), displays the propagation of acoustic waves emitted from the near-wake region; the analysis of the entire temporal series shows that the amplitude of the emitted wave critically depends on the precise phase relation between the passage of SSV and the shedding of PSV. However, this visualization does not allow to identify the contribution of each hydrodynamic feature to the sound radiation.

We present in figure 3.8 the stream-wise velocity signal recorded at $x = 1$ and $y = 0.5$. In addition to the high-frequency oscillations produced by the acoustic waves described above, we also observe a noticeable low-frequency oscillation. This latter oscillation is intrinsic to the flow as it has been reported in all numerical simulations that have been performed previously; moreover, it persists for longer integration times.

The time at which the pressure signal reaches a maximum or minimum is correlated with a low-frequency change in the pressure-surface separation bubble; no substantial changes have been observed in the separation bubble on the suction surface. These features are illustrated by displaying

²The measure has been made along the vertical line that passes through the trailing edge.

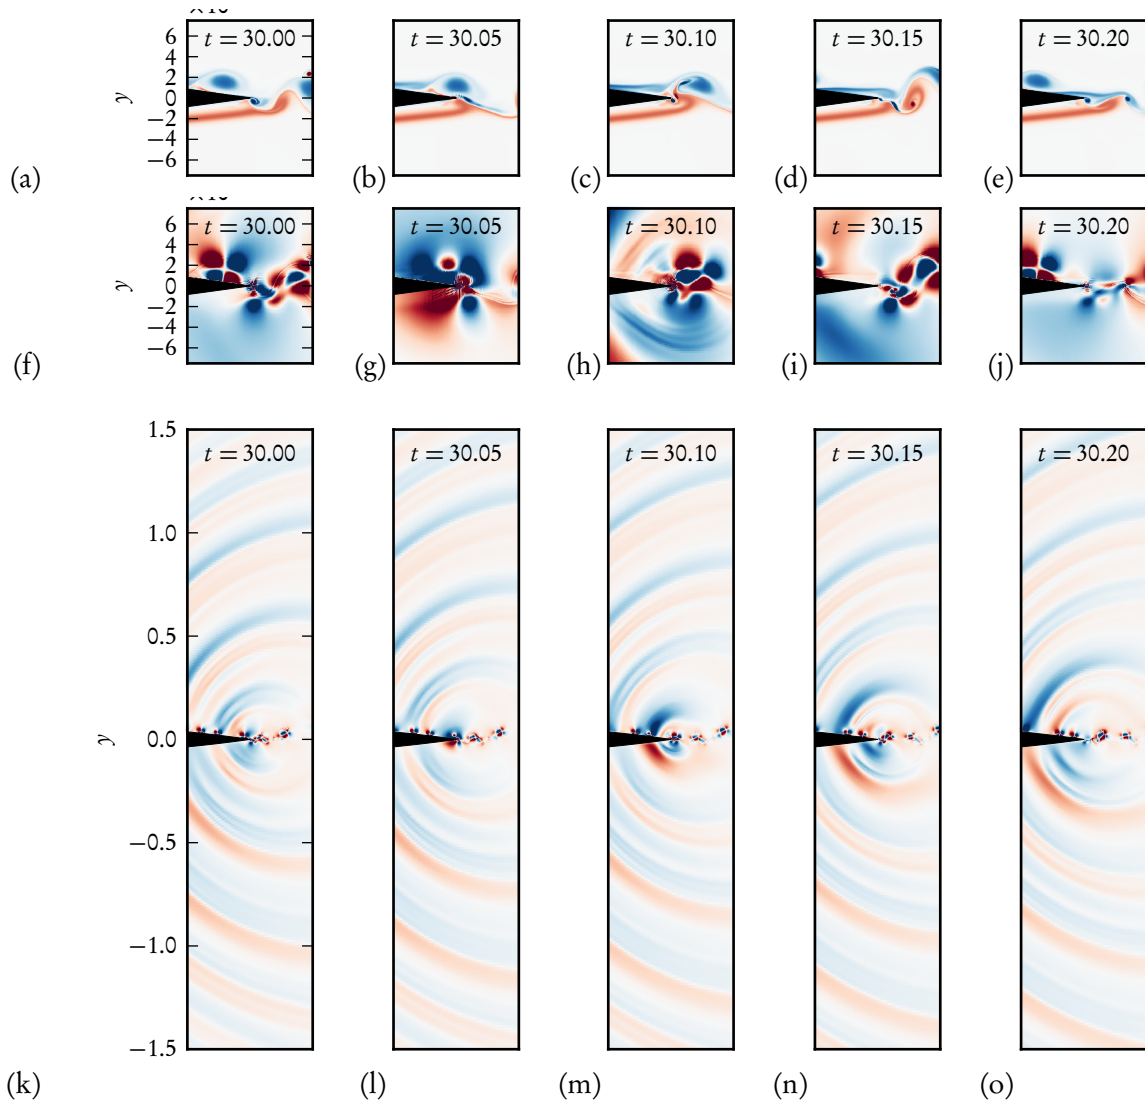


Figure 3.7: Representative flow field sequence showing (top row) instantaneous vorticity levels and (middle row) dilatation contours near the trailing edge, and (bottom row) dilatation contours of the far-field, over one vortex shedding period.

the vorticity fields for both separation bubbles at $t = 43.50$ and $t = 56.19$ (the times at which the pressure signal maximum and minimum is reached); see figure 3.9.

The low-frequency motion of the separation bubble affects the potential velocity field around the aerofoil. This effect is captured by the stream-wise velocity signal shown in figure 3.8.

This phenomenon seems to have been overlooked in previous calculations and experiments that address the tonal-noise generation mechanism. Nonetheless, it is well known (see Dovgal et al., 1994, and references therein) that separation bubbles exhibit a low-frequency flapping associated with their strong susceptibility to long-wave disturbances. The features described in the above-mentioned reference correspond to the behaviour observed here.




Figure 3.8: Stream-wise velocity recorded at a probe in the far field ($x = 1$ and $y = 0.5$).




Figure 3.9: Low frequency flapping showing the instantaneous vorticity fields. (a) the suction-surface separation bubble at $t = 43.50$, (b) the suction-surface separation bubble at $t = 56.10$, (c) the pressure-surface separation bubble at $t = 43.50$, and (d) the pressure-surface separation bubble at $t = 56.19$.

3.5 Statistics

We present in this section a quantitative description of the flow features in terms of the mean flow and standard deviation.

3.5.1 Convergence

Let u_i be an arbitrary quantity collected at time $i\Delta t$, such as the velocity signal at a certain grid point. An estimate of the mean at time $i\Delta t$, denoted by μ_i and the associated standard deviation, denoted by σ_i , of $\{u_j\}$ with $j = 1, \dots, i$ is given by

$$\mu_i = \frac{1}{i} \sum_{j=1}^i u_j, \quad \text{and} \quad \sigma_i = \sqrt{\frac{1}{i} \sum_{j=1}^i (u_j - \mu_i)^2}. \quad (3.16)$$

For this part of the study, we consider the temporal series of the velocity component u at the closest grid point to the following locations: (i) near the trailing edge at $x = 1$ and $y = 0.005$, (ii) near the reattachment of the suction side separation bubble at $x = 0.765$ and $y = 0.033$, (iii) in the far field at $x = 1$ and $y = 0.5$. The estimated value of the mean and the standard deviation at the different locations as a function of the time is displayed in figure 3.10.




Figure 3.10: Estimation of the mean and the standard deviation of the velocity signals recorded (a,b) near the trailing edge at $x = 1$ and $y = 0.005$, (c,d) near the reattachment of the suction side separation bubble at $x = 0.765$ and $y = 0.033$, and finally, (e,f) the far field at $x = 1$ and $y = 0.5$.

The following observations are made for the probes near the aerofoil (figure 3.10a–d). The results obtained for case G1 show that for simulation times $T > 50$, the estimates of the statistical moments are sufficiently converged in time. In addition, the values of the mean and the standard deviation for case G1S2 show small differences compared to case G1, indicating that the influence of the domain size and relaxation parameters are small. The mean and standard deviation for the probes near the aerofoil (i) and (ii) are within reasonable proximity. The comparison for numerical grids with increasing spatial resolution indicate that small differences persist in the values of the mean and the amplitude of the fluctuations. The discussion of the importance of these differences will be postponed to section 3.6.

The mean and the standard deviation of the velocity signal at the probe in the far field (iii), figure 3.10e–f, is more sensitive to the simulation time and the size of the physical domain. It is possible that this sensitivity arises from a difference in amplitude of the low-frequency oscillation previously discussed, since the mean oscillates on a comparable large time-scale.

Case	$x_{s,p}$	$x_{r,p}$	$u_{\min,p}$	$x_{s,s}$	$x_{r,s}$	$u_{\min,s}$
G1	0.717	0.974	-0.11	0.411	0.662	-0.05
G1S2	0.714	0.970	-0.12	0.411	0.659	-0.04
G2	0.715	0.974	-0.11	0.400	0.684	-0.06
G3	0.728	0.977	-0.10	0.405	0.686	-0.07

Table 3.4: Separation, reattachment points and maximum reverse flow.

3.5.2 Mean flow

The pressure coefficient distribution over the aerofoil $-c_p(x)$ (figure 3.11a) with

$$c_p(x) = \frac{2}{\gamma M^2} (p - 1) \quad (3.17)$$

shows that the boundary layers on both surfaces of the aerofoil are subjected to an adverse pressure gradient: from $x = 0.04$ to the trailing edge for the suction surface, and from $x = 0.24$ to the trailing edge for the pressure surface.

The separation point and zones of reverse flow can further be determined by the zero crossing and negative values of the skin-friction coefficient $c_f(x)$, respectively. This coefficient is defined as

$$c_f(x) = 2 \frac{\mathbf{t} \cdot \bar{\boldsymbol{\tau}} \cdot \mathbf{n}}{\text{Re}}, \quad (3.18)$$


where \mathbf{n} is the wall-normal vector, and \mathbf{t} is the tangential vector to the aerofoil surface in the direction of the free-stream velocity. The skin-friction coefficient (figure 3.11b) confirms the presence of separation bubbles on the pressure and suction surfaces. The mean separation and reattachment points, together with the maximum reverse flow are given in table 3.4.

It is worth highlighting that no suction-surface separation was found in the case studied by Desquesnes et al. (2007) where tonal noise is present, and marginal separation was found instead. The accuracy of the separation and reattachment point of the separation bubbles has been analysed by performing grid and domain-size dependence tests. The differences due to finer grids and larger domain sizes fall within 2%; moreover, the values reported here are in good agreement with the results computed by the Xfoil software (Drela, 1989). It is conceivable that the differences can be attributed to the difference choice in the Mach number.

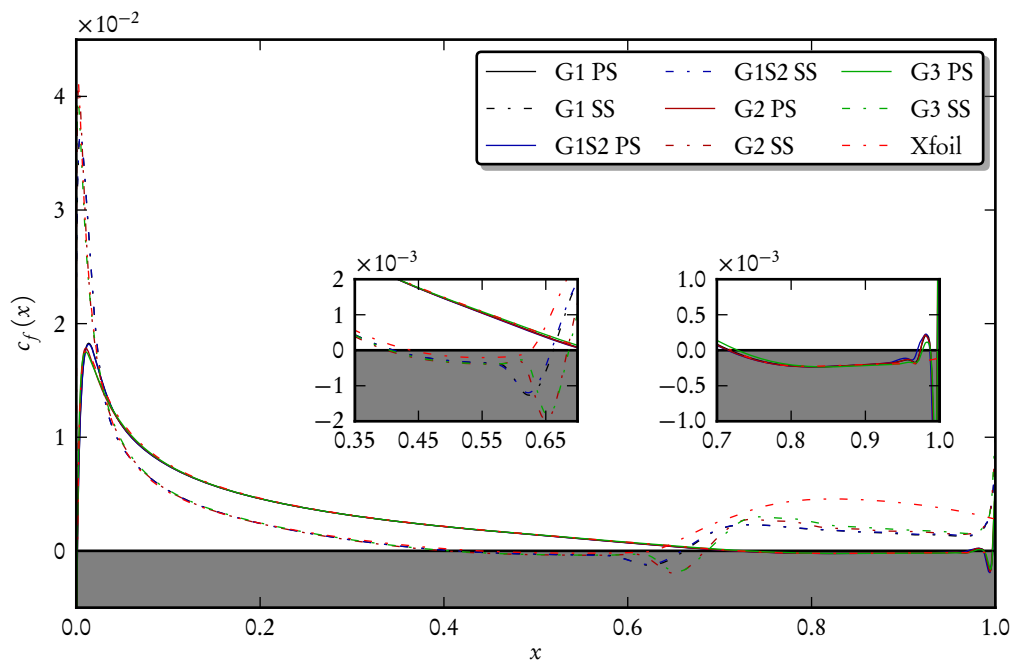
The boundary-layer displacement-thickness δ^* , shown in figure 3.12, has been defined in terms of the vorticity ω_z as

$$\delta^*(x) = \frac{\int_0^\infty n \omega_z(x, n) dn}{\int_0^\infty \omega_z(x, n) dn}, \quad (3.19)$$

where n is the wall-normal coordinate with the origin placed on the aerofoil surface for a given location x . The Reynolds number based on the displacement thickness Re_{δ^*} , also given, confirms that boundary-layer instabilities are certainly present at about one third of the aerofoil chord. The displacement thickness increases significantly after the separation points—especially on the pressure surface—due to reverse flow and the associated shear layer. On the suction surface, the vortex shedding after the reattachment averages out, decreasing the boundary-layer displacement thickness before it starts to grow again further downstream.



(a)



(b)

Figure 3.11: Mean flow (a) pressure coefficient (b) and skin-friction coefficient distribution over the aerofoil.




Figure 3.12: Boundary layer displacement thickness and Reynolds number based on the boundary layer displacement thickness.

Both on the suction and pressure surfaces, the velocity profiles are characterized by inflection points due to the negative pressure gradients, thus indicating that hydrodynamic instabilities may grow by inviscid mechanisms; see figure 3.13b,c. This observation is in agreement with the results of Nash et al. (1999) and Desquesnes et al. (2007); it is also a well-known feature of long laminar separation-bubbles (Dovgal et al., 1994).

3.5.3 Standard deviation

Previous studies—experimental and numerical—have drawn the attention to the standard deviation of the mean flow in the separation bubbles, which for the present calculations is shown in figure 3.14. The standard-deviation profiles on the suction side consist of two peaks that appear after the mean reattachment point, indicating an increased amount of unsteadiness caused by vortex growth and advection. On the pressure surface, only a single peak in the standard deviation appears at the separated shear layer. This fact has led previous studies to conclude that hydrodynamic instabilities play an important role. However, these profiles cannot be attributed by themselves to a specific travelling instability wave.

3.5.4 Trailing edge

We focus on the velocity profile at $x = 0.95$ on the pressure surface. At this location, the profile exhibits two inflection points and displays a positive stream-wise velocity close to the wall. From inspection of the averaged streamlines close to the trailing edge (figure 3.15), we conclude that the




Figure 3.13: (a) Mean vorticity levels and boundary layer velocity profiles on (b) the suction surface and (c) the pressure surface. The red circles indicate the inflection point; the blue circles indicate zero-velocity point.




Figure 3.14: Standard deviation of mean flow stream-wise velocity.




Figure 3.15: Mean-flow streamlines at the trailing edge and rms levels.

pressure-surface separation bubble gives rise to two counter-rotating trailing vortices that continue into the wake.

3.6 Frequency analysis

The nonlinear simulation can readily be described in terms of frequency content. For each numerical grid, the pressure signal at a point probe located at $x = 1$ and $y = 0.5$ has been extracted, and we have computed the Discrete Fourier Transform (DFT), which is given in figure 3.16. The spectral resolution is $\Delta f = 0.0067$ for the case G1 and $\Delta f = 0.02$ for the remaining cases. The Nyquist frequency is $F_s = 50$.

A set of equally spaced peaks is easily identified in the pressure spectrum. These discrete components are clearly related to the acoustic emission observed in the flow evolution; see figure 3.6. It is also noticed that low-frequency oscillations are present, which could not be seen directly from the inspection of the instantaneous flow. The precise value of the discrete frequencies and amplitudes of the acoustic tones and the largest low-frequency component are included in table 3.5. Although the amplitude of the pressure fluctuations of the acoustic tones can not be considered converged for the present calculations, the frequency at which the tones occur is a robust feature.

Additional information about the spatial structures related to the above-mentioned features can be obtained by analysing the DFT of the entire data set that has been collected during the nonlinear simulations. The frequency snapshots corresponding to the acoustic tones indicated in table 3.5 are represented in figure 3.17a–c and figure 3.18.

The spatial structures associated with the acoustic frequencies exhibit, in addition to the sound propagation in the far field, hydrodynamic features. Instabilities are observed to grow and saturate in the separated shear layer on the suction surface and in the wake.

At low frequency, $f = 0.06$, the associated spatial structure (figure 3.18) indicates the presence of a slow motion of the separation bubble, changing its size by displacing the separation point. This feature seems to not have been reported before in the context of the tonal noise problem although it is




Figure 3.16: Magnitude of the discrete Fourier transform of the pressure signal extracted from the probe at $x = 1$ and $y = 0.5$ for different numerical grids. (a) G1, (b) G1S2, (c) G2, and (d) G3.




Figure 3.17: Fourier component of the dominant tones, showing the real part of the dilatation and vorticity fluctuations. (a,b) $f = 6.28$, (c,d) $f = 6.66$ and (e,f) $f = 7.05$.

Case	Frequency bands							
	5.65–6.07		6.07–6.50		6.50–6.91		6.91–7.33	
	f_{peak}	p_{peak}	f_{peak}	p_{peak}	f_{peak}	p_{peak}	f_{peak}	p_{peak}
G1	5.88	$1.1 \cdot 10^{-4}$	6.28	$1.0 \cdot 10^{-4}$	6.66	$2.2 \cdot 10^{-4}$	7.05	$1.4 \cdot 10^{-4}$
G1S2	5.86	$1.1 \cdot 10^{-4}$	6.28	$2.4 \cdot 10^{-4}$	6.66	$3.9 \cdot 10^{-4}$	7.04	$1.8 \cdot 10^{-4}$
G2	5.92	$2.2 \cdot 10^{-4}$	6.36	$9.7 \cdot 10^{-5}$	6.74	$1.1 \cdot 10^{-4}$	7.14	$2.0 \cdot 10^{-4}$
G3	5.90	$1.6 \cdot 10^{-4}$	6.34	$1.8 \cdot 10^{-4}$	6.72	$2.6 \cdot 10^{-4}$	7.12	$1.9 \cdot 10^{-4}$

Table 3.5: Values of selected frequency components and their amplitude in the pressure spectrum shown in figure 3.16.




Figure 3.18: Real part of the vorticity component for $f = 0.0625$ showing low-frequency flapping of the separation bubble.

clearly observed in experimental and numerical standard-deviation profiles. Indeed, a spatio-temporal scale separation between this phenomena and the acoustic radiation seems to exist.

3.7 Conclusions from the nonlinear simulations

At this point, the presented calculations suggest that a strong interaction occurs between the dynamics of the separated flow on both sides, vortex shedding into the wake and the scattered acoustic waves. Our aim is then to understand the different degrees of participation of these flow features in the generation of sound and the establishment of a self-sustained process. These goals will be addressed in the remaining chapters.

Efficient evaluation of the linearized dynamics from compressible flow solvers

Abstract

The direct and adjoint operators play an undeniably important role in a vast number of theoretical and practical studies that range from linear stability to flow control and nonlinear optimization. Based on an existing nonlinear flow solver, the design of efficient and straightforward procedures to access these operators is thus highly desirable. In the case of compressible solvers, the use of high-order numerical schemes combined with complicated governing equations makes the derivation of efficient procedures a challenging and often tedious undertaking. In this work, a novel technique for the evaluation of the direct and adjoint operators directly from compressible flow solvers is presented and extended to include nonlinear differentiation schemes and turbulence models. The application to the incompressible counterpart is also discussed. The presented method requires minimal additional programming effort and automatically takes into account subsequent modifications in the governing equations and boundary conditions. The introduced methodology is demonstrated on existing numerical codes, and direct and adjoint global modes are calculated for three typical flow configurations. Implementation issues and the performance measures are also discussed. The proposed algorithm presents an easy-to-implement and efficient technique to extract valuable information for the quantitative analysis of complex flows.¹

4.1 Introduction

Even though most physical processes in fluid flows are most aptly described by a nonlinear mathematical model, our tools for analysing them often rely on a linear approximation of the underlying dynamics (Drazin and Reid, 2004). For this reason, the Jacobian matrix \mathbf{A} (or direct operator) of a temporal evolution process together with its adjoint \mathbf{A}^* play an undeniably important role in theoretical and practical studies (Schmid and Henningson, 2001). The direct operator describes the dynamics of small perturbations around an equilibrium state; the associated adjoint operator contains gradient information about changes in the system and their influence on its dynamics (Hill, 1992; Strykowski and Sreenivasan, 1990; Giannetti and Luchini, 2007; Marquet et al., 2008). Together, they lay a theoretical foundation for stability analyses and are integral parts in many gradient-based numerical

¹Published as an article in Journal of Computational Physics; see Fosas de Pando et al. (2012)

methods, e.g. Newton’s method (Knoll and Keyes, 2004), pseudo-arc-length continuation (Seydel, 2010), steepest-descent optimization (Kelley, 1999; Nocedal and Wright, 2006), nonlinear adjoint looping (Zuccher et al., 2006; Juniper, 2011), and many more. In a more physical context, since the Jacobian represents an approximation of the nonlinear dynamics around an equilibrium state, it plays a central part in the extraction of physical mechanisms for instabilities and, in the case of fluid dynamics, in the study of transition scenarios. In the latter case, theoretical studies within the linear framework—such as global stability (Chomaz, 2005) and transient growth analysis (Schmid, 2007)—have demonstrated their usefulness in the physical description of the onset of unsteadiness and the determination of its cause. Furthermore, weakly nonlinear analyses (Sipp and Lebedev, 2007) and many efficient techniques for active and passive flow control (Kim and Bewley, 2006) also depend on the linear dynamics.

Gradient-based analysis of fluid systems have concentrated predominantly on incompressible flow configurations because of the relative ease in deriving efficient procedures for the evaluation and extraction of the Jacobian. Despite some commendable attempts (Crouch et al., 2007; Brès and Colonius, 2008; Meliga et al., 2010; Mack and Schmid, 2011), many interesting issues of compressible flows are yet to be explored: for example, linear stability theory could shed some light on the mechanisms for sound generation and their role in the onset of instabilities (Desquesnes et al., 2007; Brès and Colonius, 2008; Mohseni et al., 2002). Although the structure of the compressible Navier–Stokes equations is algorithmically simpler because of the absence of the incompressibility constraint, the derivation of the direct linear operator and its adjoint for general compressible flows can be an arduous task and susceptible to mistakes due to a significantly larger number of terms. This situation is aggravated by the necessity of more complex boundary conditions (Poinsot and Lele, 1992; Lodato et al., 2008). An approach that has been successfully applied to spatial schemes with relatively short stencils is the explicit extraction of the linearised operator \mathbf{A} by a quasi-linearisation technique using a sequence of unit vectors \mathbf{e}_j , thus yielding, column by column, the Jacobian \mathbf{A} . In this case, given a base state \mathbf{V} , we have

$$\mathbf{A}_{ij} \approx \frac{\mathbf{F}_i(\mathbf{V} + \epsilon \mathbf{e}_j) - \mathbf{F}_i(\mathbf{V})}{\epsilon} \quad (4.1)$$

with \mathbf{F}_i denoting the i th component of the right-hand-side of the nonlinear equations and ϵ given as a small parameter. If the stencil is short and the grid ordering is known, this approach constitutes a very efficient method. This *a priori* knowledge will then allow a tiling-technique that determines multiple columns of \mathbf{A} with one evaluation.

However, the disparity of length and time scales (such as in sound generation problems) makes the use of high-order spatial schemes a convenient approach to achieve high resolution (Colonius and Lele, 2004; Pirozzoli, 2011), leading to low-sparsity or dense matrices whose explicit construction for its use in practical calculations is intractable, even for configurations with small numbers of degrees of freedom. Numerical methods relying on matrix-free evaluations of $\mathbf{A}\mathbf{v}$ and $\mathbf{A}^*\mathbf{w}$ for linear stability appear to be the best option to tackle this problem. This approach is reflected in the success of Krylov subspace techniques for large-scale linear algebra problems (Saad, 2003, 2011). A first attempt to overcome the above-mentioned sparsity problem is the numerical evaluation of the direct operator \mathbf{A} by quasi-linearisation, a technique that has been widely applied (see Mack and Schmid (2010) for an example of compressible flows). We have

$$\widetilde{\mathbf{A}}\mathbf{v} = \frac{\mathbf{F}(\mathbf{V} + \epsilon\mathbf{v}) - \mathbf{F}(\mathbf{V})}{\epsilon} \quad (4.2)$$

which provides an approximation for the action of the Jacobian (direct operator) \mathbf{A} on a given vector \mathbf{v} . This technique, however, fails to provide direct access to the adjoint operator \mathbf{A}^* .

The derivation of the direct and adjoint operators can follow two different strategies. In the *continuous* approach, the direct operator is derived from the continuous form of the nonlinear equations and implemented in a numerical code; the adjoint operator is then derived from the continuous form of the direct operator before it is discretized appropriately and implemented in a numerical code. In the *discrete* approach, the direct operator is obtained by linearization of the already discretized equations and then implemented; the discrete adjoint is then easily written in terms of the transconjugate of the discretized direct operator. In the context of flow stability studies, the continuous approach is more common. On the contrary, abundant examples of the discrete approach can be found in the optimization community. It is important to note that these two strategies do not yield the same operators, since the discretization and linearization steps generally do not commute (Sirkes and Tziperman, 1997). On the one hand, the continuous approach provides the numerical approximation of the sensitivity of the continuous equations; on the other hand, the discrete approach represents the sensitivity of the discretized approximation of the nonlinear equations.

The discrete approach offers important advantages over the continuous one, both from a practical and numerical point of view. The most relevant property is that, given the scalar product $\langle \cdot, \cdot \rangle$, \mathbf{A} and \mathbf{A}^* satisfy the fundamental relation $\langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle = \langle \mathbf{A}^*\mathbf{w}, \mathbf{v} \rangle$ up to machine precision rather than up to the discretization error of the numerical scheme on a given mesh, which is the case for the continuous approach. This property is very desirable as it avoids convergence problems of iterative gradient-based algorithms due to amplification of errors stemming from this discrepancy. In addition, the above-mentioned relation applied to the continuous case generates boundary terms, commonly by integration by parts, that have to be dealt with separately. For compressible flow simulations, this is not a trivial task. On the contrary, in the case of the discrete operators, if the discretized nonlinear equations together with a well-posed set of boundary conditions are available, then the linearized operators are in principle straightforward to derive and no additional thought has to be given to discretizations or boundary conditions for either operator, which includes, for example, the use of artificial numerical techniques such as sponge layers (Spagnoli and Airiau, 2008). This feature is also central to automatic differentiation (AD) (Carle and Fagan, 2000; Hascoët and Pascual, 2004; Giles and Pierce, 2000) which has become a powerful tool for the derivation of direct and adjoint code; the nonlinear source code is interpreted in a non-standard fashion to automatically derive new source code for the evaluation of the direct and adjoint operators. This technique has been successfully applied to a wide range of applications; however, issues related to parallelization or efficiency remain an active field of research (Müller and Cusdin, 2005).

In this article, we present a novel technique for the efficient evaluation of the linearized direct operator and its adjoint *directly* from a nonlinear compressible flow solver using the discrete approach. This technique benefits from all the advantages of the discrete approach while circumventing most of the difficulties mentioned above. Most importantly, it can be classified as a matrix-free algorithm since the operators are evaluated directly from a nonlinear solver without explicitly forming the resulting matrices. Moreover, it requires minimal storage and can be applied to a general class of spatial discretizations. We further contend that black-box strategies may not necessarily lead to optimal procedures; we thus prefer customization (without loss of generality), allowing for potential code reuse and avoiding restrictions to a specific geometry or set of boundary conditions. Only minor assumptions about the code, in particular its modularity, will be necessary; these assumptions should readily be met by exercising good programming practices and standards. In particular, it is important to notice that only spatial differentiation couples the temporal evolution of flow variables at each grid point to the flow variables at neighbouring grid points. As will be shown below, this fact can be exploited, and the general structure of the discrete direct and adjoint operators can be determined.

Even though our technique will be demonstrated on an explicit compressible code, we also mention generalizations of the algorithm to incompressible flow solvers. Once the direct and adjoint operations have been extracted from the nonlinear simulation code, they can be easily incorporated into any quantitative study that requires them. We present (in section 4.4) a representative application of our technique: we will illustrate the efficiency of the algorithm by computing direct and adjoint global modes for a spatially developing compressible boundary layer and for flow around an aerofoil.

4.2 Direct and adjoint linear operators from a nonlinear code: demonstration on the compressible Navier–Stokes equations

Before launching into a generalized formulation of the extraction of direct and adjoint operators from a modular nonlinear code, we will focus on a specific case and present a step-by-step algorithm based on a typical setup.

4.2.1 Direct numerical simulation (DNS) code

A compressible flow solver developed by the authors will be used as a demonstration example. In particular, we will only discuss details of the code as they pertain to the extraction technique. The program is based on the three-dimensional Navier–Stokes equations for compressible flow, formulated in terms of the pressure p , entropy s and velocity field \mathbf{u} . The equations are augmented by a material law, specifying the heat flux \mathbf{q} , the viscous stress tensor $\boldsymbol{\tau}$ and the state equation for an ideal gas, given by $p = \rho R_g T$. The heat capacity at constant volume c_v , the thermal conductivity λ and the viscosities μ and μ_v are taken as constants but a dependence on arbitrary state variables, such as the temperature T , can be readily considered. We have

$$\mathbf{q} = -\lambda \nabla T, \quad (4.3a)$$

$$\boldsymbol{\tau} = \mu (\nabla \mathbf{u} + \nabla \mathbf{u}^T) + \left(\mu_v - \frac{2}{3} \mu \right) (\nabla \cdot \mathbf{u}) \mathbf{I}, \quad (4.3b)$$

$$\frac{\partial p}{\partial t} + \mathbf{u} \cdot \nabla p + \rho c^2 \nabla \cdot \mathbf{u} = \frac{p}{c_v} \left(\frac{\partial s}{\partial t} + \mathbf{u} \cdot \nabla s \right), \quad (4.4a)$$

$$\frac{\partial s}{\partial t} + \mathbf{u} \cdot \nabla s = \frac{1}{\rho T} (-\nabla \cdot \mathbf{q} + \boldsymbol{\tau} : \nabla \mathbf{u}), \quad (4.4b)$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\frac{1}{\rho} \nabla p + \frac{1}{\rho} \nabla \cdot \boldsymbol{\tau}. \quad (4.4c)$$

Typically, time integration of this set of equations can be performed using the method of lines. An appropriate spatial discretization is applied, together with the treatment of the boundary conditions for the computational domain. The implementation of the boundary conditions can be thought of as the substitution of the governing equations at boundary grid points to impose a desired behaviour, i.e. inflow, outflow, solid wall, etc. A common procedure consists in the decomposition of the inviscid part of the equations into incoming and outgoing *characteristics* and specifying the value of the ingoing characteristics and viscous fluxes (Poinsot and Lele, 1992; Lodato et al., 2008). Sometimes it is preferred to implement Neumann boundary conditions or symmetry conditions directly in the




Figure 4.1: Block diagram for the evaluation of the right-hand side of the compressible Navier–Stokes equations.

differentiation schemes. In either case, the problem is reduced to the temporal integration of a system of nonlinear ordinary differential equations, symbolically written as

$$\frac{d\mathbf{v}}{dt} = \mathbf{F}(\mathbf{v}), \quad (4.5)$$

where \mathbf{v} represents the state vector containing all the variables at every grid point, and the right-hand-side $\mathbf{F}(\mathbf{v})$ comprises the governing equations, including the boundary conditions. We will use hereafter \mathbf{v} indistinctly for the state vector and for the small perturbations around a base state \mathbf{V} when there is no ambiguity. A suitable temporal integration scheme is used, starting from a specified initial field.

4.2.2 Practical implementation using a modular code structure

The evaluation of the right-hand-side $\mathbf{F}(\mathbf{v})$ of equation (4.5) can be performed by algorithmically processing the following steps:

1. Computation of the spatial derivatives of the state vector $\mathbf{v}_0 = [p \quad s \quad \mathbf{u}]^T$; the result will be called $\mathbf{v}_{0,k}$, where k stands for the respective spatial derivatives and can take on the values $k \in \{x, y, 0\}$ with $_0$ denoting no differentiation.
2. Evaluation of $\mathbf{f}_1(\mathbf{v}_{0,k})$, the discrete form of equation (4.3) including boundary conditions.
3. Computation of the spatial derivatives of $\mathbf{v}_1 = [\mathbf{q} \quad \boldsymbol{\tau}]$; the result will be called $\mathbf{v}_{1,k}$.
4. Evaluation of $\mathbf{f}_2(\mathbf{v}_{0,k}, \mathbf{v}_{1,k})$, the discrete form of equation (4.4), including boundary conditions; finally, the result (\mathbf{v}_2) is identical to the time derivative of \mathbf{v} .

In the above, \mathbf{v}_0 and \mathbf{v} are used, for convenience and indistinctly, as the state vector containing all the flow variables that are integrated in time. This four-step procedure is sketched in form of a block diagram in figure 4.1. In general, the access to each step is available by practising a modular layout of the code. The implementation of the boundary conditions fits this layout, as they are commonly enforced by replacing the governing equations on the domain boundaries and, most commonly, can be expressed in terms of the variables introduced above (see Poinso and Lele (1992); Lodato et al. (2008)). In the above diagram the differentiation routine, denoted by D_k , is used as a short form to obtain all the necessary derivatives of the state-vector \mathbf{v} , with the notational convention that $\mathbf{v}_{,x} = D_x \mathbf{v}$, $\mathbf{v}_{,y} = D_y \mathbf{v}$, and $\mathbf{v}_{,0} = D_0 \mathbf{v} = \mathbf{v}$. The subscript k thus denotes differentiation with

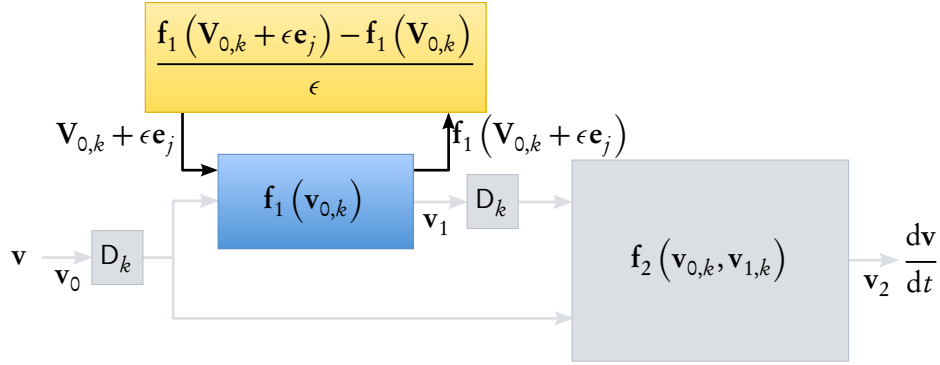


Figure 4.2: Block diagram demonstrating the linearization step for a selected module of the code.

respect to x and y , or (in the case of 0) no differentiation at all (see figure 4.5). The extension to three dimensions is obvious.

The modular structure of the code suggests the introduction of auxiliary variables, which will aid in the description of the subsequent linearization algorithm, as well as in the derivation of the adjoint operator. A new variable for each spatial derivative of the state vector is introduced. We recall that $\mathbf{v}_0 = [p \quad s \quad \mathbf{u}]^T$ is the vector state and $\mathbf{v}_1 = [\mathbf{q} \quad \boldsymbol{\tau}]^T$ is an auxiliary variable. Additionally, we have $\mathbf{v}_2 = \frac{\partial \mathbf{v}}{\partial t}$. We can then write symbolically

$$\begin{aligned}
 \mathbf{v}_0 &\longrightarrow \mathbf{v}_{0,k} = D_k \mathbf{v}_0, \\
 \mathbf{v}_1 = \mathbf{f}_1(\mathbf{v}_{0,k}) &\longrightarrow \mathbf{v}_{1,k} = D_k \mathbf{v}_1, \\
 \mathbf{v}_2 = \mathbf{f}_2(\mathbf{v}_{0,k}, \mathbf{v}_{1,k}) &\longrightarrow \mathbf{v}_2 = \frac{d\mathbf{v}}{dt}.
 \end{aligned} \tag{4.6}$$

In analogy to the block diagram (figure 4.1), this procedure can be interpreted as two distinct steps. First, the viscous tensor and the heat flux (\mathbf{v}_1) are computed from the state vector, together with its derivatives ($\mathbf{v}_{0,k}$). Second, the temporal derivative of the state vector (\mathbf{v}_2) is calculated from the flow variables and their derivatives ($\mathbf{v}_{0,k}$), as well as from the heat flux and viscous tensor and their derivatives ($\mathbf{v}_{1,k}$). The proper boundary conditions are incorporated into each respective step.

4.2.3 The linearization step

It is important to realize that, in the above scheme, only the calculation of the spatial derivatives links the flow field variables at a given point to their neighbouring points; all other operations are local (i.e., grid-point-wise). This fact can be exploited to derive a computationally efficient algorithm for the extraction and evaluation of the linearized dynamics.

To illustrate this point, we consider typical nonlinear advection terms $(\mathbf{u} \cdot \nabla) \mathbf{u}$ of a two-dimensional model problem. Following the above-mentioned procedure, we consider the spatial derivatives as independent variables and thus introduce a function $\mathbf{f}(\mathbf{u}, \mathbf{r}, \mathbf{s}) = [\mathbf{r} \mid \mathbf{s}] \mathbf{u}$ with $\mathbf{r} = D_x \mathbf{u}$ and $\mathbf{s} = D_y \mathbf{u}$. This puts further emphasis on the fact that spatial derivatives are calculated separately from the linearization procedure. Continuous differentiation is a linear operation and, for the time being, we will assume that the discrete differentiation scheme is also linear (see below for extensions to nonlinear

differentiation schemes). Consequently, only the nonlinear (grid-local) terms need to be linearized. In our particular example of the nonlinear advection term, we have

$$\mathbf{f}(\mathbf{u}, \mathbf{r}, \mathbf{s}) = \begin{pmatrix} u_1 r_1 + u_2 s_1 \\ u_1 r_2 + u_2 s_2 \end{pmatrix} \quad (4.7)$$

where the subscripts indicate the respective components of the variables. The linearization of this expression has the following structure

$$\left. \frac{\partial \mathbf{f}}{\partial \mathbf{u}} \right|_0 \mathbf{u} + \left. \frac{\partial \mathbf{f}}{\partial \mathbf{r}} \right|_0 \mathbf{r} + \left. \frac{\partial \mathbf{f}}{\partial \mathbf{s}} \right|_0 \mathbf{s} = \mathbf{A}_u \mathbf{u} + \mathbf{A}_r \mathbf{r} + \mathbf{A}_s \mathbf{s} \quad (4.8)$$

where \mathbf{A}_u , \mathbf{A}_r and \mathbf{A}_s are block-diagonal matrices depending only on the base-flow quantities of the associated variables \mathbf{u} , \mathbf{r} and \mathbf{s} (denoted by \mathbf{U} , \mathbf{R} and \mathbf{S} , respectively) and the subscript $_0$ denotes evaluation at the base state. Although in this simple example the coefficients can be readily obtained, they can be evaluated in a more systematic way. For instance, the matrix \mathbf{A}_r may be obtained numerically from a nonlinear code by two evaluations of our function $\mathbf{f}(\mathbf{u}, \mathbf{r}, \mathbf{s})$ according to

$$\mathbf{A}_r = \left. \frac{\partial \mathbf{f}}{\partial \mathbf{r}} \right|_0 \approx \left[\begin{array}{c|c} \frac{\mathbf{f}(\mathbf{U}, \mathbf{R} + \epsilon \mathbf{e}_1, \mathbf{S}) - \mathbf{f}(\mathbf{U}, \mathbf{R}, \mathbf{S})}{\epsilon} & \frac{\mathbf{f}(\mathbf{U}, \mathbf{R} + \epsilon \mathbf{e}_2, \mathbf{S}) - \mathbf{f}(\mathbf{U}, \mathbf{R}, \mathbf{S})}{\epsilon} \end{array} \right] \quad (4.9)$$

up to an accuracy of $\mathcal{O}(\sqrt{\epsilon_m})$ (with ϵ_m as the machine precision) which suffices for most applications (Knoll and Keyes, 2004). Other high-order approximations could be used without additional effort, however, attention must be paid to avoid cancellation errors. The same procedure applies for the determination of \mathbf{A}_u and \mathbf{A}_s . In summary, the complete linearization of the nonlinear advection term produces the following three matrices

$$\mathbf{A}_u = \begin{pmatrix} R_1 & S_1 \\ R_2 & S_2 \end{pmatrix} \quad \mathbf{A}_r = \begin{pmatrix} U_1 & 0 \\ 0 & U_1 \end{pmatrix} \quad \mathbf{A}_s = \begin{pmatrix} U_2 & 0 \\ 0 & U_2 \end{pmatrix}. \quad (4.10)$$

The exact linearization of the discretized equation is recovered by replacing \mathbf{r} and \mathbf{s} by $D_x \mathbf{u}$ and $D_y \mathbf{u}$, respectively. This process can be extrapolated to the case of the compressible Navier–Stokes equations, and it is illustrated in figure 4.2: the coefficient matrices $\mathbf{A}_{i,j,k}$ symbolize the linearized functions according to

$$\mathbf{A}_{i,j,k} = \left. \frac{\partial \mathbf{f}_i}{\partial \mathbf{v}_{j,k}} \right|_0. \quad (4.11)$$

In the above expression, the indices i and j represent components of the nonlinear function \mathbf{f} or the variable \mathbf{v} , while the index k indicates the derivative (i.e., $k \in \{x, y, 0\}$).

4.2.4 The direct operator

After each module of the nonlinear code has been linearized following the previous section, we can now assemble the various parts to explicitly derive the algorithm for the evaluation of the linearized direct operator. The linearization leads to

$$\begin{aligned} \mathbf{v}_0 &= \mathbf{v} \\ \mathbf{v}_1 &= \mathbf{A}_{1,0,k} D_k \mathbf{v}_0 \\ \mathbf{v}_2 &= \mathbf{A}_{2,0,k} D_k \mathbf{v}_0 + \mathbf{A}_{2,1,k} D_k \mathbf{v}_1 \quad \longrightarrow \quad \frac{d\mathbf{v}}{dt} = \mathbf{v}_2 \end{aligned} \quad (4.12)$$




Figure 4.3: Block diagram of the fully linearized direct operation.




Figure 4.4: Block diagram of the fully linearized adjoint operation.

where Einstein's summation convention applies for the repeated subscript k . By combining the various terms, a composite linearized operator \mathbf{A} can finally be determined and the temporal evolution of the field \mathbf{v} reads

$$\frac{d\mathbf{v}}{dt} = \underbrace{(\mathbf{A}_{2,0,k} \mathbf{D}_k + \mathbf{A}_{2,1,k} \mathbf{D}_k \mathbf{A}_{1,0,l} \mathbf{D}_l)}_{\mathbf{A}} \mathbf{v}. \quad (4.13)$$

This operator then forms the basis for enquiries into the linearized dynamics of the flow, such as its global stability properties or its response to external forcing or noise. It further includes the linearization of the boundary conditions implemented in the nonlinear code, and its form does not depend of any specific geometry. The algorithm for the evaluation of the linearized dynamics has exactly the same structure as the nonlinear code, but replaces the nonlinear blocks by the linearized ones. We can thus benefit from substantial code reuse and parallelization efforts already implemented in the nonlinear code.

4.2.5 The adjoint operator

Substantially more information about the linearized dynamics of the flow can be gained by considering the adjoint linearized operator or products of direct and adjoint variables. It is thus desirable to extract this operator from the nonlinear code. We first introduce the inner product $\langle \mathbf{w}, \mathbf{v} \rangle = \mathbf{w}^H \mathbf{M} \mathbf{v}$, with a Hermitian and positive-definite weight matrix \mathbf{M} . The matrix \mathbf{M} represents the discretization of the continuous inner product (containing, e.g., metric terms stemming from a non-uniform grid), but can also be used to give more weight to prescribed regions of the computational domain or selected components of the state vector. The adjoint operator is then trivially obtained from the




Figure 4.5: Differentiation module (a) and its transconjugate (b).

definition $\langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle = \langle \mathbf{A}^* \mathbf{w}, \mathbf{v} \rangle$, leading to $\mathbf{A}^* = \mathbf{M}^{-1} \mathbf{A}^H \mathbf{M}$. The temporal evolution of the adjoint field \mathbf{w} reads

$$\frac{d\mathbf{w}}{dt} = \mathbf{M}^{-1} \underbrace{\left(D_k^H \mathbf{A}_{2,0,k}^H + D_l^H \mathbf{A}_{1,0,l}^H D_k^H \mathbf{A}_{2,1,k}^H \right)}_{\mathbf{A}^*} \mathbf{M} \mathbf{w}. \quad (4.14)$$

Due to the modular nature of the extraction procedure for the direct operator, the adjoint operator can be computed by reversing the procedure in figure 4.3 and by conjugate transposing all involved linear modules. A detailed mathematical proof of this procedure is given in appendix 4.A. It shall suffice here to demonstrate the reversal of the direct methodology based on the block-diagram, which is given in figure 4.4. Starting at the right-hand side of the diagram with the variable \mathbf{w} , which is adjoint to the direct variable \mathbf{v} , we determine in a similar modular fashion the time-derivative of the adjoint variable.

$$\begin{aligned} \mathbf{w}_2 &= \mathbf{M} \mathbf{w} \\ \mathbf{w}_1 &= D_k^H \mathbf{A}_{2,1,k}^H \mathbf{w}_2 \\ \mathbf{w}_0 &= D_l^H \mathbf{A}_{1,0,l}^H \mathbf{w}_1 + D_k^H \mathbf{A}_{2,0,k}^H \mathbf{w}_2 \quad \longrightarrow \quad \frac{d\mathbf{w}}{dt} = \mathbf{M}^{-1} \mathbf{w}_0 \end{aligned} \quad (4.15)$$

The various matrices $\mathbf{A}_{i,j,k}$ which have been determined during the linearization step for the direct operator are simply transconjugated (symbolized by H). The differentiation matrices D_k need to be transconjugated as well. In one dimension, most of the spatial-derivative approximations using finite differences can symbolically be written as $M_L \mathbf{v}_{,x} = M_R \mathbf{v}$ leading to $D = M_L^{-1} M_R$. The matrix-vector multiplication $D^H \mathbf{w}$ is simple to implement since it only needs the transconjugation of the matrices M_L and M_R and its respective application in reverse order. The extension to spatial derivatives in two and three dimensions is straightforward, since derivatives in higher dimensions can easily be written as Kronecker products of one-dimensional differentiation operators. For instance, considering two dimensions and lexicographical ordering of the grid, the spatial derivatives along x and y are $D_x = \mathbf{I} \otimes D$ and $D_y = D \otimes \mathbf{I}$, respectively.

The adjoint of the differentiation matrices D_k is displayed in figure 4.5 and involves an additional summation over all components. On the one hand, the forward operation D_k takes a single variable \mathbf{v} and produces multiple spatial derivatives ($\mathbf{v}_{,k}$) which are subsequently used and summed in a linearized function block (note that \mathbf{f}_2 has been replaced by *two* matrices $\mathbf{A}_{2,1,k}, \mathbf{A}_{2,0,k}$); on the other hand, by reversing this sequence, the adjoint of the function block produces multiple variables from

a single-variable input which are in turn “adjoint differentiated” and summed in the D_k^H -block. This is consistent with the summation of repeated indices in the above equations. Adjoint first-order differentiation can be thought of as negative differentiation, since for continuous derivatives we have $(\partial/\partial x)^* = -\partial/\partial x$. The incorporation of boundary closures into the discrete differentiation matrices D_k , however, makes this analogy only true in an approximate or interpretive sense. We shall remark here that the linearized boundary conditions were automatically included in the direct operator, and therefore they are also taken into account in the adjoint operator without any additional effort.

By transconjugation of the procedural algorithm for the evaluation of the direct operator, we have thus obtained the adjoint linearized operator directly from our nonlinear code. We can then use both the direct and adjoint linearized operator to address and quantify various issues related to the linear dynamics of our flow.

4.3 A general framework for the evaluation of direct and adjoint operators

In this section we extend the evaluation technique of the direct and adjoint operators to an arbitrarily complex nonlinear code that uses an explicit discretization in time. At the end of the section, potential extensions of this procedure to take into account incompressibility or nonlinear differentiation schemes are analysed.

4.3.1 Generalization

The state vector \mathbf{v} that describes the evolution of a physical state is supposed to be governed by a system of partial differential equations, nonlinear and first order in time, which can be stated as

$$\frac{\partial \mathbf{v}}{\partial t} = \mathbf{f}(\mathbf{v}) \quad (4.16)$$

with appropriate initial and boundary conditions. In this equation, an explicit dependence on the spatial coordinates \mathbf{x} and time t will be excluded, even though it could be incorporated in the formulation that follows. When discretized explicitly in time, we arrive at a system of autonomous ordinary differential equations via the method of lines, and a numerical code that implements the right-hand side and advances the state vector in time is assumed to be available. In a continuous formulation, the right-hand side $\mathbf{f}(\mathbf{v})$ is generally a complicated function of not only the state vector \mathbf{v} but also of its spatial derivatives (up to arbitrary order). The link between changes in one of the state variables at a given grid point and associated changes in neighbouring grid points is caused by spatial derivatives.

We reconsider the continuous version of the equations and introduce auxiliary variables and local functions. In the discretized equations, nonlinearity arises from expressions involving those auxiliary variables. This latter fact is easily exploited as has been shown before. Following the procedure introduced in the previous section, equation (4.16) can be split into as many auxiliary functions as the order of the highest spatial derivative of the state vector in $\mathbf{f}(\mathbf{v})$. More formally, this can be written

as

$$\begin{aligned}
 \mathbf{v}_0 = \mathbf{v} & \longrightarrow \mathbf{v}_{0,k} = \frac{\partial \mathbf{v}_0}{\partial x_k}, \\
 \mathbf{v}_1 = \mathbf{f}_1(\mathbf{v}_{0,k}) & \longrightarrow \mathbf{v}_{1,k} = \frac{\partial \mathbf{v}_1}{\partial x_k}, \\
 \vdots & \vdots \\
 \mathbf{v}_i = \mathbf{f}_i(\mathbf{v}_{0,k}, \dots, \mathbf{v}_{i-1,k}) & \longrightarrow \mathbf{v}_{i,k} = \frac{\partial \mathbf{v}_i}{\partial x_k}, \\
 \vdots & \vdots \\
 \mathbf{v}_n = \mathbf{f}_n(\mathbf{v}_{0,k}, \dots, \mathbf{v}_{n-1,k}) & \longrightarrow \mathbf{v}_n = \frac{\partial \mathbf{v}}{\partial t}.
 \end{aligned} \tag{4.17}$$

As before, it is convenient to define $\mathbf{v}_{j,0} = \mathbf{v}_j$. Since our attention is only focused on the extraction of the linearized dynamics from already discretized nonlinear equations, only the discrete formulation will be considered. The structure of the above equations imposes weak constraints on the manner in which expressions are implemented; but a sufficiently well-written simulation code easily satisfies these constraints.

4.3.2 Linearization

The linearization of an arbitrary auxiliary function \mathbf{f}_i is carried out by introducing a base state \mathbf{V} plus a small perturbation \mathbf{v}' , i.e. $\mathbf{v} = \mathbf{V} + \epsilon \mathbf{v}'$. The values of the auxiliary variables at the base state can be obtained by subsequently evaluating $\mathbf{f}_i(\mathbf{V}_{j,k})$ using the nonlinear code. The equation for the perturbation is then

$$\mathbf{v}'_i = \sum_{j=0}^{i-1} \sum_k \left. \frac{\partial \mathbf{f}_i}{\partial \mathbf{v}_{j,k}} \right|_0 \mathbf{v}'_{j,k} \tag{4.18}$$

where the subscript $_0$ stands for an evaluation at the base state. As before, the submatrices of the Jacobian of \mathbf{f}_i can be determined numerically choosing a sufficiently small value ϵ and evaluating

$$[\mathbf{A}_{i,j,k}]_{lm} = \left. \frac{\partial \mathbf{f}_i}{\partial \mathbf{v}_{j,k}} \right|_0 \approx \left[\frac{\mathbf{f}_i(\mathbf{V}_{0,k}, \dots, \mathbf{V}_{j,k} + \epsilon \mathbf{e}_m, \dots, \mathbf{V}_{i-1,k}) - \mathbf{f}_i(\mathbf{V}_{0,k}, \dots, \mathbf{V}_{i-1,k})}{\epsilon} \right]_l. \tag{4.19}$$

This expression presumes that $\mathbf{f}_i(\mathbf{v}_{j,k})$ is easily accessible. Since all the elements of $\mathbf{A}_{i,j,k}$ can be obtained in very few operations (the auxiliary variables are defined at a grid point), the value of ϵ can be chosen adaptively in order to ensure that the numerical differentiation remains within an acceptable error tolerance. Furthermore, the primes (indicating the linearized variables) are omitted hereafter.

4.3.3 Direct and adjoint operator evaluation

The linearized version of the discretized system is easily determined by replacing the nonlinear functions by their linearized counterparts. The adjoint (equation 4.21) is obtained by explicitly

forming the direct operator (equation 4.20) that stems from this derivation and applying the definition of the adjoint operator. Details of the derivation of the adjoint operator are given in 4.A, but the main results are included below.

$$\begin{aligned}
 \mathbf{v}_0 &= \mathbf{v} & \mathbf{w}_n &= \mathbf{M}\mathbf{w} \\
 \vdots & & \vdots & \\
 \mathbf{v}_i &= \sum_{j=0}^{i-1} \sum_k \mathbf{A}_{i,j,k} \mathbf{D}_k \mathbf{v}_j & \mathbf{w}_i &= \sum_{j=i+1}^n \sum_k \mathbf{D}_k^H \mathbf{A}_{j,i,k}^H \mathbf{w}_j \\
 \vdots & & \vdots & \\
 \mathbf{v}_n &= \sum_{j=0}^{n-1} \sum_k \mathbf{A}_{n,j,k} \mathbf{D}_k \mathbf{v}_j \rightarrow \frac{d\mathbf{v}}{dt} = \mathbf{v}_n & \mathbf{w}_0 &= \sum_{j=1}^n \sum_k \mathbf{D}_k^H \mathbf{A}_{j,n,k}^H \mathbf{w}_j \rightarrow \frac{d\mathbf{w}}{dt} = \mathbf{M}^{-1} \mathbf{w}_0
 \end{aligned} \tag{4.20} \tag{4.21}$$

We see that the procedure follows the outline introduced previously for the specific case of the compressible Navier–Stokes equations. In particular, we recognize the reversed processing and transconjugation of the linearized modules for the adjoint operation.

4.3.4 Extensions

The above methodology is, so far, applicable to any flow solver provided that the differentiation schemes are linear and the temporal advancement is performed explicitly. In this section we relax these limitations and consider several extensions such as the use of nonlinear differentiation schemes (upwinding, WENO) and mixed discretizations, as is the case for incompressible flow solvers.

4.3.4.1 Nonlinear differentiation schemes

Weighted essentially non-oscillatory (WENO) schemes In WENO schemes (Shu, 1998) the differentiation operator is nonlinear and must be included in the linearization procedure. Several stencils are computed using different sets of neighbouring points, and a convex nonlinear combination of them is formed to avoid spurious oscillations close to discontinuities. The differentiation operator can be expressed as $\mathbf{D}(\mathbf{v})$, where the value of the derivative at x_i has compact support, i.e. it only depends on the values at $x_{i \pm k}$, where the largest value of k determines the overall stencil width which is commonly very small. If the stencil width is known *a priori*, an efficient linearization can be performed using a tiling technique. In this case, the linearized operator can be computed using $2k_{\max}$ function calls and can be stored in sparse format.

Upwinding In upwind schemes, two different stencils are considered for the computation of the spatial derivatives and the appropriate one is chosen depending on the sign of the advection velocity. The choice of spatial derivatives using upwinded stencils are $\mathbf{v}_{j,k}^+ = \mathbf{D}_k^+ \mathbf{v}_j$ and $\mathbf{v}_{j,k}^- = \mathbf{D}_k^- \mathbf{v}_j$, and a decision is made based on a generally nonlinear function which can be written as $\mathbf{f}_i(\mathbf{v}_{j,k}^+, \mathbf{v}_{j,k}^-)$. If the latter function is differentiable, one can proceed according to the method above; however, this function is usually non-differentiable and thus the linearization has to be performed cautiously. For instance, it is common to switch from one derivative (+) to the other (−) using the sign of the

advection speed. In numerical implementations this amounts to an *if-statement* which renders the function non-differentiable and the linearized operator ill-defined. Nevertheless, the *if-statement* can be replaced by a smooth representation such as the logistic function $L(u) \equiv 1/(1 + e^{-u/\delta})$, where δ is chosen sufficiently small to achieve a fast transition but sufficiently large to avoid numerical difficulties during the linearization. In practice, a choice of $\delta \sim O(\sqrt{\epsilon_m})$ (with ϵ_m as the machine precision) is typically satisfactory. If we reconsider the example given in section 4.2.3 and introduce upwinding in the differentiation along the x direction, we obtain

$$\mathbf{f}(\mathbf{u}, \mathbf{r}^+, \mathbf{r}^-, \dots) = H(u_1)u_1\mathbf{r}^+ + H(-u_1)u_1\mathbf{r}^- + \dots \approx \frac{u_1}{1 + e^{-u_1/\delta}}\mathbf{r}^+ + \frac{u_1}{1 + e^{u_1/\delta}}\mathbf{r}^- + \dots \quad (4.22)$$

with $H(\cdot)$ denoting the Heaviside function. Performing a linear stability analysis of the resulting discretization, it can be easily verified that such a modification does not introduce numerical instabilities or artefacts. Once the above function is linearized around a base state, the corresponding linearized equation reads

$$\mathbf{A}_r^+\mathbf{r}^+ + \mathbf{A}_r^-\mathbf{r}^- + \dots \quad (4.23)$$

and the thus linearized module can be incorporated in the overall procedure for the direct and adjoint operator. In summary, the above procedure to treat upwind schemes aims at avoiding problems arising from numerical differentiation.

4.3.4.2 Turbulence models

Turbulence models such as Spalart–Allmaras, k - ϵ , k - ω , Reynolds stress model (RSM) and their variants add one or several differential evolution equations for the modelled variables such as turbulent eddy viscosity, turbulent kinetic energy, dissipation, etc., to our governing equations. These new variables and their associated equations can be readily accounted for by expanding the state vector to include the respective quantities. However, it is important to verify that the chosen turbulence model is differentiable or can be regularized similar to the procedure given in the previous section. Also noteworthy is the fact that the linearized system includes the effect of small perturbations in the resolved structures on the modelled turbulent quantities (Crouch et al., 2007).

4.3.4.3 Incompressible flow solvers

As a further extension, we next consider the incompressible Navier–Stokes equations which, written in terms of primitive variables (p, \mathbf{u}) , read

$$\nabla \cdot \mathbf{u} = 0, \quad (4.24)$$

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} = -\nabla p + \nu \nabla^2 \mathbf{u}. \quad (4.25)$$

When integrating the above set of equations in time, it is common practice to consider implicit time-stepping for the diffusion term in order to achieve a reasonable time step. This fact, together with the incompressibility constraint, complicates the straightforward application of our linearization technique. For instance, the time evolution can no longer be written in the form $\frac{d\mathbf{v}}{dt} = \mathbf{F}(\mathbf{v})$ since the incompressibility constraint leads to a separate Poisson equation for the pressure. Nevertheless,

we observe that the system of equations resulting from the implicit treatment of diffusion and the incompressibility constraint is already linear and thus does not need to be linearized. Moreover, the coefficient matrix stemming from the discretization of the Laplacian is typically symmetric or can be transconjugated easily to form the corresponding adjoint operation.

There exists a great variety of numerical methods to integrate the above system of equations in time, and the derivation of a general linearization strategy for all of them is beyond the scope of this article. Rather, we consider an extension of our technique to the pressure-less fractional-step method (Kim and Moin, 1985). The discretization is performed on a staggered-grid, where the pressure is defined in the cell centres and the components of the velocity field \mathbf{u} at the cell edges. The advancement over one time step can be decomposed into the following steps:

1. nonlinear advection at interior grid points

$$\frac{\mathbf{u}^a - \mathbf{u}^n}{\Delta t} + (\mathbf{u}^n \cdot \nabla) \mathbf{u}^n = 0, \quad (4.26)$$

2. diffusion using backward Euler

$$\frac{\mathbf{u}^{ad} - \mathbf{u}^a}{\Delta t} = \nu \nabla^2 \mathbf{u}^{ad} \quad \text{with} \quad \mathbf{u}_{\partial\Omega}^{ad} = \mathbf{u}_{\partial\Omega}^n, \quad (4.27)$$

3. computation of ϕ and update of the final velocity field

$$\nabla^2 \phi^{n+1} = \frac{1}{\Delta t} \nabla \cdot \mathbf{u}^{ad} \quad \text{with} \quad \frac{\partial \phi^{n+1}}{\partial n} = 0, \quad (4.28)$$

$$\mathbf{u}^{n+1} = \mathbf{u}^{ad} - \Delta t \nabla \phi^{n+1}. \quad (4.29)$$

Due to the mixed discretization in space and time, it is desirable to derive a linearization procedure for the operator that propagates a given flow field over one time step, i.e. $\mathbf{v}^{n+1} = \mathbf{A}_{\Delta t} \mathbf{v}^n$, with $\mathbf{v} = [u_i, \phi]^T$. We consider hereafter the discretized equations and introduce the parameter $\alpha = \nu \Delta t$. The temporal advancement of the Stokes operator is readily identified in steps 2–3 and denoted by the linear operator $\mathbf{A}_{\mathcal{S}}$. As far as our linearization procedure is concerned, the Stokes operator $\mathbf{A}_{\mathcal{S}}$ plays a role similar to spatial differentiation and thus can be treated analogous to the latter. The propagation over one time step can now be stated as

$$\begin{aligned} \mathbf{v}_0 = \mathbf{v}^n &\longrightarrow \mathbf{v}_{0,k} = \frac{\partial \mathbf{v}_0}{\partial x_k} \\ \mathbf{v}_1 = \mathbf{f}_1(\mathbf{v}_{0,k}) &\longrightarrow \mathbf{v}^{n+1} = \mathbf{A}_{\mathcal{S}} \mathbf{v}_1, \end{aligned} \quad (4.30)$$

where $\mathbf{v}_1 = [u_i^a]^T$. Only the function $\mathbf{f}_1(\mathbf{v}_{0,k})$ needs to be linearized. As before, this can be performed using the numerical approximation of quasi-linearization. Once the matrices $\mathbf{A}_{1,0,k}$ are determined, the direct and/or adjoint operators for the propagation over one time step read

$$\begin{aligned} \mathbf{v}_0 = \mathbf{v}^n & & \mathbf{w}_1 = \mathbf{A}_{\mathcal{S}}^H \mathbf{M} \mathbf{w}^n \\ \mathbf{v}_1 = \mathbf{A}_{1,0,k} \mathbf{D}_k \mathbf{v}_0 &\longrightarrow \mathbf{v}^{n+1} = \mathbf{A}_{\mathcal{S}} \mathbf{v}_1, & \mathbf{w}_0 = \mathbf{D}_k^H \mathbf{A}_{1,0,k}^H \mathbf{w}_1 &\longrightarrow \mathbf{w}^{n+1} = \mathbf{M}^{-1} \mathbf{w}_0. \end{aligned} \quad (4.31) \quad (4.32)$$

What remains is the determination of the explicit expression for the transconjugate of the Stokes operator $\mathbf{A}_{\mathcal{S}} = \mathbf{A}_{\mathcal{D}} \mathbf{A}_{\mathcal{D}}$ (steps 2 and 3). In the latter expression, $\mathbf{A}_{\mathcal{D}}$ is the operator for the diffusion (step 2) and $\mathbf{A}_{\mathcal{P}}$ the computation of the divergence-free velocity field (step 3):

$$\mathbf{A}_{\mathcal{D}} = (1 - \alpha L_d)^{-1}, \quad (4.33) \quad \mathbf{A}_{\mathcal{P}} = [1 - GL_n^{-1}D], \quad (4.34)$$

where L stands for the discretized Laplacian, D for the discretized divergence and G for the discretized gradient. Subscripts $_d$ and $_n$ refer to Dirichlet and homogeneous Neumann boundary conditions, respectively. Using the previous relations we can write the transconjugate of $\mathbf{A}_{\mathcal{S}}$ as $\mathbf{A}_{\mathcal{S}}^H = \mathbf{A}_{\mathcal{D}}^H \mathbf{A}_{\mathcal{P}}^H$ with

$$\mathbf{A}_{\mathcal{D}}^H = (1 - \alpha L_d)^{-1}, \quad (4.35) \quad \mathbf{A}_{\mathcal{P}}^H = [1 - D^H L_n^{-1}G^H]. \quad (4.36)$$

Although the above derivation of the discrete adjoint operator for the linearized incompressible solver is rather simple and can be performed systematically, it may not be as straightforward to implement as for the compressible flow solver. The continuous Stokes operator is self-adjoint, but the use of the pressure projection and the discretization render this property only approximately true for the numerical implementation. For this reason, the transconjugate of the involved operators needs to be performed judiciously. Nevertheless, most of these operators are simple to transconjugate; for instance, L_n and L_d are usually symmetric, and, in the case considered here, D^H and G^H are equal to $-G$ and $-D$, respectively, which renders the projection operator $\mathbf{A}_{\mathcal{P}}$ and the diffusion operator $\mathbf{A}_{\mathcal{D}}$ symmetric.

4.4 Applications

In this section, we focus our attention on the application to compressible and incompressible flow solvers to demonstrate the efficiency and capabilities of the method introduced above. More precisely, the computation of direct and adjoint global modes – a central component in many quantitative flow analyses – will be performed on several flow configurations, namely a spatially developing compressible boundary layer, the compressible flow around an aerofoil, and the incompressible flow in a lid-driven cavity. Implementation details and the obtained performances will be discussed in particular.

4.4.1 Implementation and performance

The compressible flow solver under consideration implements the three-dimensional Navier–Stokes equations in simple multiblock-structured grids using the so-called *pseudo-characteristics* formulation (Sesterhenn, 2000) in curvilinear coordinates. The numerical code can be used to address typical flow configurations in aeroacoustics, and, consequently, high-order numerical methods have been chosen in order to accurately resolve all flow features at a reasonable computational cost. In particular, compact upwind low-dissipative (CULD) schemes (Adams and Shariff, 1996) are used for the advection terms and central compact schemes (Lele, 1992) for the computation of viscous and heat fluxes. The temporal advancement is carried out using a 4th-order low-storage Runge-Kutta scheme (Kennedy et al., 2000), and appropriate boundary conditions are implemented by extending the Navier–Stokes characteristics boundary conditions, or NSCBC (Poinsot and Lele, 1992; Lodato et al., 2008), to curvilinear coordinates.




Figure 4.6: Relative residual for the linearized direct operator for the application cases.

The nonlinear solver is written in C++ and conveniently parallelized using the message-passing library MPI. The evaluation of the direct and adjoint operators is implemented in two distinct modules: (a) the numerical linearization of the auxiliary nonlinear functions around a base state, equation (4.19), and (b) the evaluation of the linearized direct and adjoint operators, equations (4.12) and (4.15), respectively. Both modules can straightforwardly be incorporated into quantitative investigations that rely on them. It is worth pointing out that, in the case under consideration, both modules represent only 9% of the lines-of-code of the entire program: 4% for the linearization module and 5% for the evaluation module. The latter module shares most of the routines with its nonlinear counterpart and thus inherits its performance characteristics and reuses already parallelized code. Moreover, the linearization module automatically takes into account modifications in the boundary conditions and even in the governing equations.

More precisely, the linearization process is performed by evaluating equation (4.19) with a value of ϵ sufficiently small, typically $\sqrt{\epsilon_m}$, where ϵ_m is the machine precision. The optimal value can be determined by repeated evaluation using progressively smaller values of ϵ . This procedure yields the coefficients of $\mathbf{A}_{i,j,k}$ at an expected convergence rate, e.g., $\mathcal{O}(\epsilon)$ for our first-order Jacobian approximation. By further reducing ϵ , cancellation errors begin to dominate until they contaminate the approximate coefficients. The memory cost associated with the storage of the matrices $\mathbf{A}_{i,j,k}$ is usually small. For example, in our case, the storage of over one hundred coefficients per grid point are typically required for two-dimensional problems.

In order to verify the accuracy of the linearization procedure and the implementation of the direct operator, the norm of the difference between the linearized direct operator obtained using equation (4.2) and our technique is presented in figure 4.6. A local minimum in relative error between the two techniques is reached for a parameter $\epsilon \approx 4 \cdot 10^{-8}$; for values above, monotonic first-order convergence is observed, whereas for values below, round-off errors start to dominate. The imple-

mentation of the adjoint operator has been verified using the identity $\langle \mathbf{w}, \mathbf{A}\mathbf{v} \rangle = \langle \mathbf{A}^* \mathbf{w}, \mathbf{v} \rangle$, which for the cases presented here, is satisfied up to round-off error.

A comparison between the time spent in the evaluation of the nonlinear and linearized operators, together with the time spent in the linearization process, is shown in table 4.1. As expected, the nonlinear and linearized direct and adjoint codes show nearly identical performances; the time spent in the linearization module is comparable to the time it takes to perform one evaluation, since the computation of matrix coefficients can be performed very efficiently.

Case	Nonlinear		Linearization		Direct		Adjoint	
	Time (s)	Ratio	Time (s)	Ratio	Time (s)	Ratio	Time (s)	Ratio
Boundary layer	0.0096	1.0	0.1575	16.4	0.0103	1.07	0.0109	1.17
Airfoil	0.0428	1.0	0.4690	11.0	0.0604	1.41	0.0624	1.45

Table 4.1: Comparison between the time spent on the evaluation of the different operators and the linearization procedure. The ratio to the time spent on the nonlinear operator evaluation, shown in bold, is also presented.

The minimal additional programming effort that is required for the implementation of this technique, specially in the case of compressible flow solvers, and its performance result in a substantial advantage over linearized codes obtained from Automatic Differentiation. In the latter case, the *differentiated* code is usually as long as the original code, and sustained performance can only be achieved via careful optimization.

The above-mentioned modules have also been implemented for an incompressible flow solver. The advancement in time of the Stokes operator is performed using the pressure-less fractional-step method described in Kim and Moin (1985). The nonlinear advection terms are discretized in space using upwinded, central finite differences and advanced in time using the forward Euler method.

Finally, we illustrate the potential applications of this technique by computing direct and adjoint global modes of several flow configurations. The temporal advancement of the linearized operators (Edwards et al., 1994) has been coupled with the eigenvalue solver SLEPc (Hernandez et al., 2005). The calculations have been performed using the Krylov-Schur algorithm (Stewart, 2002), together with the harmonic-extraction method (Morgan and Zeng, 2006) to select different parts of interest of the spectrum.

4.4.2 Spatially developing compressible boundary layer

We first consider a two-dimensional compressible boundary layer over a flat plate. The boundary-layer displacement thickness δ_1 at the inlet of the domain is taken as the reference length; the Reynolds number is $\text{Re}_{\delta_1} = 1000$ and the Mach number is $M = 0.8$. The domain extends over 800 and 40 unit lengths along the tangential and wall normal directions, respectively. The numerical grid is refined in the vertical direction in the vicinity of the wall. The velocity and entropy profiles obtained from the self-similar solution of the compressible boundary-layer equations are imposed at the inlet of the domain using the characteristics-based approximate non-reflecting inflow boundary condition given in Lodato et al. (2008). A no-slip adiabatic boundary condition is imposed at the wall, and a characteristics-based approximate non-reflecting outflow boundary condition is implemented at the free-stream boundary and at the outlet of the domain. For this first example demonstrating the

matrix extraction technique, we will concentrate on a classical global stability analysis, determine the direct spectrum and focus on the least stable global modes.




Figure 4.7: Global spectrum of compressible boundary-layer flow for $Re_{\delta_1} = 1000$ and $M = 0.8$, displaying three families of spectral branches.

Given a suitable initial condition, the flow is advanced in time until the norm of the time derivative falls below 10^{-8} . The linearization of the equations around this steady solution is performed, and the temporal propagator related to the linearized direct problem is used to obtain the global modes. In the global spectrum (figure 4.7), three types of modal branches can be recognized, linked to Tollmien–Schlichting waves, Orr modes and free-stream modes.

Selected eigenfunctions corresponding to each of these branches are displayed in figure 4.8. The Tollmien–Schlichting branch is characterized by low phase velocities on the order of 30% of the free-stream value. Its associated structures show compact support inside the boundary layer and exponential decay towards the free-stream (see figure 4.8a). Exponential spatial growth in the streamwise direction is also observed for this type of modes. The disturbance dynamics in the free-stream is represented by the free-stream modes (see figure 4.8b) which show a phase velocity comparable to the free-stream value. We notice the typical weak exponential decay towards the free-stream as well as substantially stronger exponential decay towards the wall inside the boundary layer. These types of modes are particularly important for receptivity studies which address the transfer of disturbance energy between the free-stream and the boundary layer. The final type of mode is known as the Orr modes which commonly occur in predominantly two-dimensional configurations. They represent a specific dynamics whereby structures extract energy from the background base-flow via a tilting process induced by the mean shear. Similar to the Tollmien–Schlichting waves, Orr modes are confined to the boundary layer.




Figure 4.8: Three different types of modes selected from the global spectrum for the compressible boundary layer with $\text{Re}_{\delta_1} = 1000$ and $M = 0.8$. (a) Tollmien–Schlichting (TS) waves, (b) free-stream modes and (c) Orr modes.

4.4.3 Sound generation by an aerofoil

In this example, the compressible flow around a NACA0012 airfoil at 2° angle of attack is computed. The Reynolds number based on the chord length is $\text{Re} = 2 \cdot 10^5$ and the Mach number is $M = 0.4$. The Navier–Stokes equations are solved on a C-grid that extends over 7 chord lengths along the wake and wall-normal direction. The numerical grid consists of three curvilinear blocks with a total of 3840×384 points. In the chosen parameter regime, the flow exhibits a substantial level of acoustic noise that appears as a sharp peak in the frequency spectrum. Experimental and theoretical studies show that the generation of this type of sound can be linked to the linear stability of the flow. In order to avoid contamination of the acoustic field with spurious reflections caused by vortical structures as they leave the domain, a sponge layer is added at the outlet downstream of the aerofoil.

A nonlinear simulation of this configuration is conducted until the flow reaches a quasi-periodic state, and the linearization of the governing equations is performed around the mean flow. Figure 4.9 presents the most unstable direct and adjoint global mode extracted by our algorithm. Each panel depicts the pressure field and the streamwise velocity field. For the direct mode we observe a strong spatial growth in the chordwise direction that culminates near the trailing edge. In addition, amplified structures are visible near the separation bubble on the suction (upper) side. The pressure field displays the associated sound field scattered by the fluid structures. Again, the strongest emission of sound stems from the trailing edge of the aerofoil. The corresponding adjoint global mode identifies structures on the pressure (lower) side of the aerofoil and can be used to identify regions where the associated direct global mode (figure 4.9a) is particularly sensitive to perturbations. Both direct and adjoint modes play an important role in pinpointing localized areas that may be the source of




Figure 4.9: Leading direct (a) and adjoint (b) global modes ($\lambda \approx 43.8i$) for compressible flow around a NACA0012 aerofoil at 2° angle of attack, $\text{Re} = 2 \cdot 10^5$ and $M = 0.4$. In both cases pressure field and streamwise velocity contours are shown.

self-sustained oscillations.

4.4.4 Lid-driven cavity flow

We conclude this section by presenting results obtained using the incompressible flow solver. The flow in a lid-driven square cavity at $\text{Re} = 10000$ is considered. A typical high-order cavity mode is displayed in figure 4.10. It shows a typical vortical structure superimposed on the mean-flow cavity vortex. The adjoint mode (figure 4.10b) shows additional features near the downstream edge of the lid which indicates increased sensitivity in this region for the excitation of the associated direct mode (figure 4.10a).

This display of modes has been included here for demonstration purposes only, to illustrate the capability of the matrix-extraction technique even for incompressible flow solvers. In true stability calculations, a less dissipative spatial discretization scheme has to be chosen in order to isolate the physical dissipation from the dissipation of the numerical scheme. In this sense, the upwind scheme chosen for this demonstration is unsuitable to determine proper growth rates even though the qualitative features of the modes are readily captured.




Figure 4.10: Representative direct (a) and adjoint (b) global mode in a lid-driven square cavity at $Re = 10000$.

4.5 Summary and conclusions

In this article, we have presented a novel technique for the efficient evaluation of the linearized discrete direct and adjoint operators directly from existing compressible nonlinear simulation codes. The discrete approach contains several advantages over the continuous one: if the nonlinear solver is assumed to accurately describe the underlying physics, the associated linearized operators will consequently represent the linear dynamics and characterize their sensitivities. The derivation of these linearized operators can be obtained in a systematic manner without giving specific thoughts to boundary conditions or numerical discretizations. Nevertheless, the derivation of efficient procedures for discrete linearized operators, especially for the compressible case, can be cumbersome, error-prone and challenging. However, the apparent difficulties can be overcome by carefully analysing the structure of the discretized governing equations. A crucial observation is the fact that (i) variables at different grid points are only linked by spatial discretization, which is commonly a linear operation, and (ii) nonlinearities arise from expressions that relate local quantities at a single grid point. This observation also holds for the boundary conditions considered in this article. By introducing auxiliary variables and assuming a modular structure of the code, the nonlinear modules for the evaluation of the nonlinear terms can be linearized with a small number of function evaluations and minimal memory requirements. The adjoint of the linearized direct operator can be obtained by transconjugating the direct operator, or equivalently, *reversing* the block-diagram that represents the evaluation of direct operator. The evaluation technique consists of two steps: first, the numerical linearization of the nonlinear local expressions and, second, the evaluation of the direct and adjoint operators using a sequence of all linearizations. We have shown that the structure of the linearized direct and adjoint operators allows for potential code reuse and parallelization efforts from the nonlinear solver. For the compressible case shown here, the total added code represents only 9% of the entire program. The routines for the evaluation of the linearized operators inherit the performance characteristics of the nonlinear code. The evaluation technique has been derived and demonstrated first to a compressible flow solver and has then been extended to also treat non-

linear differentiation schemes, turbulence models and incompressible flow solvers. The presented algorithm has been applied to and showcased on three selected flow configurations: a compressible spatially growing boundary layer, compressible flow around an aerofoil and the incompressible flow in a lid-driven square cavity. Global spectra and global direct and adjoint eigenfunctions have been presented and discussed.

In particular, in the field of computational aeroacoustics (CAA) this technique could have substantial benefits in the analysis, optimization and control of noise-generating mechanisms. Even though earlier attempts at direct-adjoint optimizations have been made using a continuous formulation (Collis et al., 2003; Wei and Freund, 2005), the above technique based on the discrete adjoint offers additional advantages.

The efficient extraction of direct and adjoint operators has numerous and important applications in the quantitative analyses of complex fluid flows. Modal solutions of the direct operator give insight into stability properties, receptivity characteristics, and physical transition mechanisms. Combined with the associated adjoint modes, they allow for the quantification and localization of sensitivity measures, feedback mechanisms, and gradient information for optimization schemes. Moreover, direct and adjoint information is prevalent – but, in general, difficult to obtain – in the design of active and passive control strategies and in the reduction of high-dimensional models. For this reason, an efficient technique that provides this necessary information directly from nonlinear simulation codes with a moderate amount of effort is valuable and welcome addition to the currently available methods to analyze complex configurations arising in multi-physics and multi-scales flow applications.

Acknowledgments

The first author wishes to thank Xavier Garnaud for fruitful discussions. This work was performed using HPC resources from GENCI-CINES (Grant 2011-026451).

4.A Derivation of the adjoint

In order to proceed with the derivation of the adjoint of equation (4.20), we seek an expression for the time derivative, denoted by \mathbf{v}_n , of the state vector \mathbf{v} in terms of \mathbf{v} . We start by expressing the dependence of an intermediate variable \mathbf{v}_i on $\mathbf{v}_0, \dots, \mathbf{v}_{i-1}$ in matrix form:

$$\underbrace{\begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \vdots \\ \mathbf{v}_{i-1} \\ \mathbf{v}_i \end{bmatrix}}_{\mathbf{v}_{1:i}} = \underbrace{\begin{bmatrix} 0 & 0 & \cdots & 0 & 0 \\ \mathbf{A}_{2,1,k} D_k & 0 & \cdots & 0 & 0 \\ \vdots & \vdots & \ddots & \vdots & \vdots \\ \mathbf{A}_{i-1,1,k} D_k & \mathbf{A}_{i-1,2,k} D_k & \cdots & 0 & 0 \\ \mathbf{A}_{i,1,k} D_k & \mathbf{A}_{i,2,k} D_k & \cdots & \mathbf{A}_{i,i-1,k} D_k & 0 \end{bmatrix}}_{\mathbf{L}_i} \underbrace{\begin{bmatrix} \mathbf{v}_1 \\ \mathbf{v}_2 \\ \vdots \\ \mathbf{v}_{i-1} \\ \mathbf{v}_i \end{bmatrix}}_{\mathbf{v}_{1:i}} + \underbrace{\begin{bmatrix} \mathbf{A}_{1,0,k} D_k \\ \mathbf{A}_{2,0,k} D_k \\ \vdots \\ \mathbf{A}_{i-1,0,k} D_k \\ \mathbf{A}_{i,0,k} D_k \end{bmatrix}}_{\mathbf{I}_{1:i,0}} \mathbf{v}_0, \quad (4.37)$$

where $\mathbf{v}_{1:i}$ is the composite vector containing the intermediate variables $\mathbf{v}_1, \dots, \mathbf{v}_i$. The lower-triangular matrix \mathbf{L}_i and the matrix $\mathbf{I}_{1:i,0}$ describe in shorthand the relation between the auxiliary variables and the vector state \mathbf{v} . Using the above definitions, we write

$$\left[\begin{array}{c|c} \mathbf{I} - \mathbf{L}_{n-1} & \mathbf{0} \\ \hline -\mathbf{I}_{n,1:n-1} & \mathbf{I} \end{array} \right] \left[\begin{array}{c} \mathbf{v}_{1:n-1} \\ \mathbf{v}_n \end{array} \right] = \left[\begin{array}{c} \mathbf{I}_{1:n-1,0} \\ \mathbf{I}_{n,0} \end{array} \right] \mathbf{v}_0. \quad (4.38)$$

This system of equations can readily be reformulated by eliminating the composite vector $\mathbf{v}_{1:n-1}$. We obtain

$$\frac{d\mathbf{v}}{dt} = \mathbf{v}_n = \left(\mathbf{I}_{n,1:n-1} (\mathbf{I} - \mathbf{L}_{n-1})^{-1} \mathbf{I}_{1:n-1,0} + \mathbf{I}_{n,0} \right) \mathbf{v} \quad (4.39)$$

and, given the scalar product $\langle \mathbf{w}, \mathbf{v} \rangle = \mathbf{w}^H \mathbf{M} \mathbf{v}$. The equation adjoint the one above reads

$$\frac{d\mathbf{w}}{dt} = \mathbf{M}^{-1} \left(\mathbf{I}_{1:n-1,0}^H (\mathbf{I} - \mathbf{L}_{n-1}^H)^{-1} \mathbf{I}_{n,1:n-1}^H + \mathbf{I}_{n,0}^H \right) \mathbf{M} \mathbf{w}. \quad (4.40)$$

The evaluation of the above adjoint can be expressed in an equivalent form to equation (4.20), which is more amenable for practical implementation. After introducing the auxiliary adjoint variables $\mathbf{w}_n = \mathbf{M} \mathbf{w}$, $\mathbf{w}_{1:n-1} = (\mathbf{I} - \mathbf{L}_{n-1}^H)^{-1} \mathbf{I}_{n,1:n-1}^H \mathbf{w}_n$, we have $\frac{d\mathbf{w}}{dt} = \mathbf{M}^{-1} \mathbf{w}_0$. Combining these definitions, we arrive at an expression analogous to equation (4.38). It reads

$$\left[\begin{array}{c|c} \mathbf{I} & -\mathbf{I}_{1:n-1,0}^H \\ \hline \mathbf{0} & \mathbf{I} - \mathbf{L}_{n-1}^H \end{array} \right] \left[\begin{array}{c} \mathbf{w}_0 \\ \mathbf{w}_{1:n-1} \end{array} \right] = \left[\begin{array}{c} \mathbf{I}_{n,0}^H \\ \mathbf{I}_{n,1:n-1}^H \end{array} \right] \mathbf{w}_n. \quad (4.41)$$

Finally we obtain the familiar sequential procedure given by

$$\begin{aligned} \mathbf{w}_n &= \mathbf{M} \mathbf{w} \\ &\vdots \\ \mathbf{w}_i &= \sum_{j=i+1}^n \sum_k \mathbf{D}_k^H \mathbf{A}_{j,i,k}^H \mathbf{w}_j \\ &\vdots \\ \mathbf{w}_0 &= \sum_{j=1}^n \sum_k \mathbf{D}_k^H \mathbf{A}_{j,0,k}^H \mathbf{w}_j \quad \longrightarrow \quad \frac{d\mathbf{w}}{dt} = \mathbf{M}^{-1} \mathbf{w}_0. \end{aligned} \quad (4.42)$$

Global linear stability analysis

5.1 Introduction

Most studies that address the tonal-noise problem on aerofoils relate the appearance of acoustic tones in the sound spectrum to the laminar boundary layer on the pressure surface. The results from the nonlinear calculations (chapter 3) confirm that for our chosen parameter regime the tonal noise phenomenon is present. Furthermore, in agreement with the flow visualization carried out by Nash et al. (1999), the frequential snapshots suggest that a strong interaction occurs between the separated boundary layers, the acoustic radiation at the trailing edge and the wake dynamics.

In both experimental (Nash et al., 1999) and numerical (Desquesnes et al., 2007) investigations, it has been demonstrated that the frequency of the dominant acoustic tone is correlated, within reasonable bounds, to the frequency of the most amplified instability wave along the pressure-surface boundary layer. This seems to suggest that pressure-surface boundary-layer instabilities play a central role in the tonal-noise generation mechanism. It is important to point out, however, that boundary-layer instabilities are only convectively unstable and, thus, are not sufficient to explain self-sustained oscillations.

This latter fact has led previous researchers to introduce *ad-hoc* mechanisms, such as aeroacoustic feedback-loops—beyond the scope and validity of local stability theories—to explain a self-sustained process that selects discrete frequencies which in turn satisfy an appropriate phase condition between acoustic waves and hydrodynamic instabilities. In Jones et al. (2010), the tonal-noise mechanism has recently been investigated using numerical simulations of small disturbances superimposed on the mean flow, and it was found that aeroacoustic feedback loops are indeed present and can play an important role in the selection of acoustic tones. This fact then suggests that a global stability approach may shed light onto the tonal-noise problem, as it circumvents the limitations and shortcomings of local stability theory.

In this chapter, we start our analysis by presenting the principal features of an impulse response, which is then followed by an analysis of the global spectrum and a study of the response behaviour to optimal external forcings.

5.1.1 Mathematical tools

Even though most global stability concepts are formally defined in the continuous time domain based on the linearized Navier–Stokes equations, we directly present the framework for the spatially discretized equations. The reason for this choice is a practical one: the computations that follow are necessarily performed in finite-dimensional spaces in which many of the continuous concepts relate to familiar linear algebra tools.

Once the nonlinear compressible Navier–Stokes equations have been spatially discretized, the calculation of the temporal evolution is reduced to solving a set of ordinary differential equations given by

$$\frac{d\mathbf{v}}{dt} = \mathbf{F}(\mathbf{v}) \quad (5.1)$$

where $\mathbf{F}(\mathbf{v})$ is a nonlinear function. Linear global stability is then based on the linearized dynamics of this nonlinear function around a base flow \mathbf{V} . The study therefore relies on the Jacobian given by

$$\mathbf{A} = \left(\frac{\partial \mathbf{F}}{\partial \mathbf{v}} \right)_{\mathbf{v}}. \quad (5.2)$$

Provided that the physical mechanisms under investigation are well represented by the nonlinear model and its linearized version, the operator \mathbf{A} will form the foundation for our further study.

The analysis of \mathbf{A} is greatly aided by access to the adjoint operator \mathbf{A}^* . The definition of the adjoint operator is based on a specified inner product. In our case we have $\langle \mathbf{w}, \mathbf{v} \rangle = \mathbf{w}^H \mathbf{M} \mathbf{v}$, where \mathbf{M} is a Hermitian, positive definite matrix. Then, the adjoint operator \mathbf{A}^* is given as the operator that satisfies the duality relation $\langle \mathbf{w}, \mathbf{A} \mathbf{u} \rangle = \langle \mathbf{A}^* \mathbf{w}, \mathbf{u} \rangle$ for arbitrary \mathbf{w} and \mathbf{v} . From this fundamental relationship, the operator \mathbf{A}^* can be easily expressed in terms of the transconjugate \mathbf{A}^H . We have

$$\mathbf{A}^* = \mathbf{M}^{-1} \mathbf{A}^H \mathbf{M}. \quad (5.3)$$

At this point, it is important to emphasize that the linearized Navier–Stokes operator is non-normal, i.e. the direct and adjoint operator do not commute $\mathbf{A} \mathbf{A}^* \neq \mathbf{A}^* \mathbf{A}$, which has fundamental implications on the perturbation dynamics. It should also be mentioned that the adjoint operator \mathbf{A}^* is of critical importance for the understanding of the dynamical properties of the operator \mathbf{A} . It has countless and interesting applications of which we only consider a few here.

5.1.1.1 Linearized dynamics: impulse response and frequency response

The ultimate goal of stability studies is the assessment of the linear flow behaviour, i.e., the response of the dynamical system \mathbf{A} to initial perturbations \mathbf{u} or to external forcings $f e^{-i\omega t}$. Hence

$$\frac{d\mathbf{v}}{dt} = (\mathbf{A} + i\omega) \mathbf{v} + \mathbf{f} \quad \text{with} \quad \mathbf{v}(0) = \mathbf{u}, \quad (5.4)$$

where the term $e^{-i\omega t}$ has been factored out.

At first, linear mechanisms can be investigated by the analysis of the temporal evolution of a given initial condition. In this case, we have $\mathbf{f} = 0$ and $\omega = 0$, and the evolution of the perturbation $\mathbf{v}(t)$ reads

$$\mathbf{v}(t) = e^{t\mathbf{A}} \mathbf{u}. \quad (5.5)$$

This type of analysis will be considered in section 5.4. It enables us to probe transient mechanisms in the flow as well as its asymptotic global stability. Even in the case of a globally stable flow, the

non-normality of the operator \mathbf{A} leads to important transient growth, and the temporal evolution of perturbations has important implications on the triggering of non-linearities in the flow.

The observed physical features of the flow can also be related to a sustained response to external perturbations. A special and generic case is the flow response to a harmonic forcing. In that case, regardless of the initial condition, the asymptotic solution is given by

$$\mathbf{v} = -(\mathbf{A} + i\omega\mathbf{I})^{-1}\mathbf{f}. \quad (5.6)$$

Both impulse response and harmonic response are means to gain more insight into linear amplification mechanisms of the flow under investigation.

5.1.1.2 Eigenvalue decomposition: global modes

A more quantitative picture of the flow behaviour arises from the eigenvalue spectrum of the linearized operator \mathbf{A} . For this we assume a modal decomposition into elementary solutions of the form $\mathbf{v}_i e^{-i\omega_i t}$. The mathematical problem for the eigenpairs (ω_i, \mathbf{v}_i) is given by

$$-i\omega_i \mathbf{v}_i = \mathbf{A}\mathbf{v}_i. \quad (5.7)$$

In the terminology of global stability analyses, the eigenvalues ω_i and the eigenvectors \mathbf{v}_i are referred to as the global frequencies and the direct global modes, respectively. The eigenvectors are undetermined up to a multiplicative constant. Unless explicitly stated, the direct global modes considered here have been scaled such that $\|\mathbf{v}_i\|^2 = 1$. The above decomposition allows us to determine and categorize the physical mechanisms that are present in the flow.

Analogously, for each global frequency ω_j , the adjoint operator has an eigenpair $(-\bar{\omega}_j, \mathbf{w}_j)$, where the adjoint global mode \mathbf{w}_j needs to be calculated by solving

$$i\bar{\omega}_j \mathbf{w}_j = \mathbf{A}^* \mathbf{w}_j. \quad (5.8)$$

The duality property has immediate implications for the direct \mathbf{v}_i and adjoint \mathbf{w}_j global modes. Hence

$$\langle \mathbf{A}^* \mathbf{w}_j, \mathbf{v}_i \rangle = \langle \mathbf{w}_j, \mathbf{A}\mathbf{v}_i \rangle \longrightarrow (\bar{\omega}_j + \omega_i) \langle \mathbf{w}_j, \mathbf{v}_i \rangle = 0, \quad (5.9)$$

which indicates that the direct \mathbf{v}_i and adjoint \mathbf{w}_j global modes are bi-orthogonal unless $i = j$. As before, the adjoint global modes are defined up to a multiplicative constant. Hereinafter, the adjoint global modes are normalized such that $\langle \mathbf{w}_i, \mathbf{v}_i \rangle = 1$.

The above property permits us to project an arbitrary field \mathbf{v} and \mathbf{w} onto a linear combination of direct and adjoint global modes. Hence

$$\mathbf{v}_{\parallel} = \sum_i a_i \mathbf{v}_i \longrightarrow a_i = \langle \mathbf{w}_i, \mathbf{v} \rangle, \quad (5.10) \quad \text{and} \quad \mathbf{w}_{\parallel} = \sum_i b_i \mathbf{w}_i \longrightarrow b_i = \langle \mathbf{v}_i, \mathbf{w} \rangle. \quad (5.11)$$

The response of the flow field to an impulse response \mathbf{u} or to a harmonic forcing \mathbf{f} , given in terms of an expansion into global modes, reads for \mathbf{v}_{\parallel} as

$$\mathbf{v}_{\parallel}(t) = \sum_i \langle \mathbf{w}_i, \mathbf{u} \rangle \mathbf{v}_i e^{-i\omega_i t}, \quad (5.12) \quad \text{and} \quad \mathbf{v}_{\parallel} = i \sum_i \frac{\langle \mathbf{w}_i, \mathbf{f} \rangle}{\omega - \omega_i} \mathbf{v}_i, \quad (5.13)$$

respectively. Although direct and adjoint global modes are helpful in providing insight into the physical mechanisms that are present in the flow, modal compositions are known to be rather ill-conditioned for highly non-normal operators. This limitation can be partially overcome by considering the singular value decomposition, which provides a basis for a robust representation.

5.1.2 Singular value decomposition: optimal perturbations and optimal forcings

Different choices of initial conditions \mathbf{u} and different forcings \mathbf{f} generally lead to rather different flow behaviour. For the type of analysis considered here, we are particularly interested in situations that induce the maximum response in the flow, i.e., we seek initial conditions \mathbf{u}_{\max} and external harmonic forcings \mathbf{f}_{\max} that yield the largest possible gain. In the first case, the analysis is known as *transient growth analysis*, whereas the second case is typically referred to as *optimal frequency-response analysis* (Trefethen et al., 1993; Schmid and Henningson, 2001). In order to quantify the effect, we define the gain as the ratio of the norm of the output (response) to the norm of the input. For each type of analysis we have

$$G_{\max} = \max_{\mathbf{u}} \frac{\|e^{\mathbf{A}t}\mathbf{u}\|}{\|\mathbf{u}\|}, \quad (5.14) \quad \text{and} \quad G_{\max} = \max_{\mathbf{f}} \frac{\|(\mathbf{A} + i\omega\mathbf{I})^{-1}\mathbf{f}\|}{\|\mathbf{f}\|}. \quad (5.15)$$

For our purpose, it is more convenient to rewrite the gains in terms of the 2-norm. By introducing the Cholesky decomposition of the matrix $\mathbf{M} = \mathbf{L}^H\mathbf{L}$ that defines the inner product $\langle \mathbf{w}, \mathbf{v} \rangle$, we arrive at

$$G_{\max} = \max_{\mathbf{u}} \frac{\|\mathcal{P}(t; \mathbf{A})\mathbf{L}\mathbf{u}\|_2}{\|\mathbf{L}\mathbf{u}\|_2}, \quad (5.16) \quad \text{and} \quad G_{\max} = \max_{\mathbf{f}} \frac{\|\mathcal{R}(\omega; \mathbf{A})\mathbf{L}\mathbf{f}\|_2}{\|\mathbf{L}\mathbf{f}\|_2} \quad (5.17)$$

where we have introduced the matrices

$$\mathcal{P}(t; \mathbf{A}) = \mathbf{L}e^{t\mathbf{A}}\mathbf{L}^{-1}, \quad (5.18) \quad \text{and} \quad \mathcal{R}(\omega; \mathbf{A}) = \mathbf{L}(\mathbf{A} + i\omega\mathbf{I})^{-1}\mathbf{L}^{-1}. \quad (5.19)$$

The solution to the above optimization problems is given by three leading singular triplet $(\sigma_1, \mathbf{u}_1, \mathbf{v}_1)$ of each respective matrix and reads

$$\mathbf{u}_{\max} = \mathbf{L}^{-1}\mathbf{v}_1, \quad (5.20) \quad \mathbf{f}_{\max} = \mathbf{L}^{-1}\mathbf{v}_1, \quad (5.23)$$

$$\mathbf{v}_{\max} = \sigma_1\mathbf{L}^{-1}\mathbf{u}_1, \quad (5.21) \quad \mathbf{v}_{\max} = \sigma_1\mathbf{L}^{-1}\mathbf{u}_1, \quad (5.24)$$

$$G_{\max} = \sigma_1, \quad (5.22) \quad G_{\max} = \sigma_1. \quad (5.25)$$

The reader should note that, in contrast to the eigenvalue decomposition, the bases associated with the singular value decomposition are orthonormal. This means that for a given initial condition \mathbf{u} or an arbitrary forcing \mathbf{f} , the respective response at time T or frequency ω , respectively, is far better represented by singular vectors than by eigenvectors. We have

$$\mathbf{v}_{\parallel}(t) = \sum_i \langle \mathbf{u}_{\max, i}, \mathbf{u} \rangle \mathbf{v}_{\max, i}, \quad (5.26) \quad \text{and} \quad \mathbf{v}_{\parallel} = \sum_i \langle \mathbf{f}_{\max, i}, \mathbf{f} \rangle \mathbf{v}_{\max, i}. \quad (5.27)$$

However, in contrast to the eigenvalue decomposition, the SVD-based vector basis has to be constructed anew for changing target times T and frequencies ω .

5.1.3 Sensitivity of eigenvalues and singular values

In both the eigenvalue and singular value decompositions, it is important to assess the sensitivity of eigenvalues and singular values to small perturbations, respectively, which is related to the structural sensitivity of the flow. We proceed first with the sensitivity of the eigenvalues.

Given the eigenpair (λ, \mathbf{v}) of an arbitrary matrix \mathbf{A} , the effect of an operator perturbation $\delta \mathbf{A}$ on the eigenvalues can be determined by considering the perturbed eigenvalue problem and by neglecting higher order terms. We have

$$(\mathbf{A} + \delta \mathbf{A})(\mathbf{v} + \delta \mathbf{v}) = (\lambda + \delta \lambda)(\mathbf{v} + \delta \mathbf{v}) \longrightarrow \delta \lambda = \langle \mathbf{w}, \delta \mathbf{A} \mathbf{v} \rangle. \quad (5.28)$$

The above expression is central for the analysis of structural sensitivity. In the case of global modes, the eigenvalue sensitivity reads

$$-i \delta \omega_i = \langle \mathbf{w}_i, \delta \mathbf{A} \mathbf{v}_i \rangle. \quad (5.29)$$

This procedure can also be extended to a singular value problem, as the latter can be expressed as an equivalent eigenvalue problem. For instance, in the case of the optimal frequency-response, it is convenient to express the sensitivity of the optimal gain in terms of the sensitivity of the smallest-magnitude eigenvalue given by

$$\begin{pmatrix} 0 & \mathcal{R}^{-H} \\ \mathcal{R}^{-1} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} = \sigma^{-1} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}. \quad (5.30)$$

As before, we consider small perturbations to the operator $\delta \mathcal{R}^{-1}$ and seek first-order effects on $\delta \sigma^{-1}$. Hence

$$\begin{pmatrix} 0 & \delta \mathcal{R}^{-H} \\ \delta \mathcal{R}^{-1} & 0 \end{pmatrix} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix} = \delta \sigma^{-1} \begin{pmatrix} \mathbf{u} \\ \mathbf{v} \end{pmatrix}, \quad (5.31)$$

from where we obtain

$$\delta \sigma = -\sigma^2 \operatorname{Re} \left(\mathbf{v}^H \delta \mathcal{R}^{-1} \right) \mathbf{u}. \quad (5.32)$$

As a result, the sensitivity of the optimal gain to small variations $\delta \mathbf{A}$ of the operator \mathbf{A} and to small variations in the frequency $\delta \omega$ reads (Brandt et al., 2011)

$$\frac{1}{G_{\max}} \delta G_{\max} = -\operatorname{Re} \langle \mathbf{f}_{\max}, \delta \mathbf{A} \mathbf{v}_{\max} \rangle, \quad (5.33)$$

and

$$\frac{1}{G_{\max}} \frac{\delta G_{\max}}{\delta \omega} = \operatorname{Im} \langle \mathbf{f}_{\max}, \mathbf{v}_{\max} \rangle, \quad (5.34)$$

respectively.

5.1.4 Evaluation of the direct and adjoint operators

Although conceptually simple, the above linear algebra problems that arise from global stability calculations are not straightforward to solve: the high-dimensionality of \mathbf{A} and \mathbf{A}^* put severe restrictions on the numerical methods that can be chosen for practical implementations. The above choice is further restricted by the fact that the operators \mathbf{A} and \mathbf{A}^* are in general dense, owing to our choice

of discretization. The limitation is twofold: the storage required for the full operators is prohibitive large, and the computational time required for solving the entire problem using dense algorithms far exceeds typical computational resources.

In this context, the only numerical methods that provide a solution to our problem are necessarily based on iterative matrix-free algorithms. This type of algorithms solely relies on matrix-vector products $\mathbf{A}\mathbf{v}$ and $\mathbf{A}^*\mathbf{w}$. Algorithms based on Krylov subspaces techniques are thus the methods of choice. For our case, we have augmented the direct numerical simulation code presented in chapter 2 by an evaluation procedure of matrix-vector products for both linearized operators \mathbf{A} and \mathbf{A}^* (see chapter 4).

5.2 Krylov subspace methods

In this section we provide a brief summary of the Krylov subspace methods used in the present study: the exponential Krylov time integration method and the Arnoldi algorithm for eigenvalue problems (Edwards et al., 1994). An in-detail theoretical description of Krylov methods is given in Saad (2003, 2011); Van der Vorst (2003).

5.2.1 Vector basis

Probably the most well-known numerical algorithm for the iterative solution of eigenvalue problems is the power iteration: the largest-magnitude eigenpair (λ, \mathbf{v}) of a matrix \mathbf{A} is obtained from the recursive relation $\mathbf{v}_{n+1} = \mathbf{A}\mathbf{v}_n$ with \mathbf{v}_1 an arbitrary vector. After a sufficient number of iterations $\lambda = \|\mathbf{v}_{n+1}\|/\|\mathbf{v}_n\|$ approximates the largest-magnitude eigenvalue, and the eigenvector is given by $\mathbf{v} \approx \mathbf{v}_{n+1}$. Note that in the final result only the last matrix-vector product $\mathbf{A}\mathbf{v}_n$ is considered. A similar situation occurs in iterative algorithms for the solution of linear system of equations such as the Jacobi method. Krylov methods arise from the observation that more valuable information about the solution can be obtained by considering the complete series of matrix-vector products, and seeking the solution in the subspace that they define. The Krylov subspace of a matrix \mathbf{A} and vector \mathbf{v} of size m is then defined as

$$\mathcal{K}_m(\mathbf{A}, \mathbf{v}) \equiv \text{span} \{ \mathbf{v}, \mathbf{A}\mathbf{v}, \mathbf{A}^2\mathbf{v}, \dots, \mathbf{A}^{m-1}\mathbf{v} \}. \quad (5.35)$$

For practical reasons, a vector basis for the above subspace is not constructed directly from its definition, as the resulting basis is not orthogonal and round-off errors induced by finite-accuracy arithmetic render the basis ill-conditioned. Instead, a superior alternative to construct a basis for the Krylov subspace $\mathcal{K}_m(\mathbf{A}, \mathbf{v})$ is the Arnoldi algorithm with modified Gram–Schmidt reorthogonalizations (algorithm 1).

In matrix notation, the Arnoldi algorithm can be written as

$$\mathbf{A}\mathbf{V}_m = \mathbf{V}_m\mathbf{H}_m + h_{m+1,m}\mathbf{v}_{m+1}\mathbf{e}_m^H, \quad (5.36)$$

where the matrix \mathbf{V}_m is formed by m column vectors generating the subspace $\mathcal{K}_m(\mathbf{A}, \mathbf{v})$. The square matrix \mathbf{H}_m is upper Hessenberg with leading dimension m , and \mathbf{v}_{m+1} is by construction orthogonal to \mathbf{V}_m .

Krylov methods consist then in seeking in a Krylov subspace the *optimal* element that approximates the solution to a specific problem. Typical estimates of the error are obtained as a function of \mathbf{v}_{m+1} . The most important application of Krylov methods is the iterative solution of linear system of equations, where the inversion of \mathbf{A} is typically replaced by the inversion of \mathbf{H}_m . Hence

$$(\lambda\mathbf{I} - \mathbf{A})^{-1}\mathbf{v} \approx \beta\mathbf{V}_m(\lambda\mathbf{I} - \mathbf{H}_m)^{-1}\mathbf{e}_1. \quad (5.37)$$

Algorithm 1 Arnoldi with Modified Gram–Schmidt

```

Choose initial vector  $\mathbf{v}$ 
 $\beta \leftarrow \|\mathbf{v}\|$ 
 $\mathbf{v}_1 \leftarrow \mathbf{v}/\beta$ 
for  $i = 1, m$  do
   $\mathbf{w} \leftarrow \mathbf{A}\mathbf{v}_i$ 
  for  $j = 1, i$  do
     $h_{ji} = \mathbf{v}_j^H \mathbf{v}_i$ 
     $\mathbf{w} \leftarrow \mathbf{w} - h_{ji}\mathbf{v}_j$ 
  end for
   $h_{i+1,i} \leftarrow \|\mathbf{w}\|$ 
  if  $h_{i+1,i} \leq \text{tol}$  then
    break
  end if
   $\mathbf{v}_{i+1} \leftarrow \mathbf{w}/h_{i+1,i}$ 
end for

```

The above relation can be used in general to approximate matrix functions (Hochbruck and Lubich, 1997). We have

$$\begin{aligned}
f(\mathbf{A})\mathbf{v} &= \frac{1}{2\pi i} \int_{\Gamma} f(\lambda)(\lambda\mathbf{I} - \mathbf{A})^{-1} \mathbf{v} d\lambda \\
&\approx \frac{1}{2\pi i} \int_{\Gamma} f(\lambda)\beta\mathbf{V}_m(\lambda\mathbf{I} - \mathbf{H}_m)^{-1} \mathbf{e}_1 d\lambda = \beta\mathbf{V}_m f(\mathbf{H}_m)\mathbf{e}_1.
\end{aligned} \tag{5.38}$$

Depending of the function $f(\mathbf{A})$ whose approximation is to be computed, the convergence properties of the Krylov approximation vary. It should be noted that Krylov subspaces are predominantly oriented towards the direction of the eigenvectors that have associated the largest-magnitude eigenvalues. This fact is of great importance when assessing the convergence of Krylov subspace techniques.

5.2.2 Matrix exponential: time integration

In the case of the exponential matrix $e^{t\mathbf{A}}\mathbf{v}$, the Krylov approximation converges with superlinear error decay (Hochbruck and Lubich, 1997), and leads to a very efficient technique for the integration in time of large-scale linear systems of ordinary differential equations. For an application for compressible flows, the reader is referred to Schulze et al. (2009). Following Sidje (1998), we consider the initial value problem

$$\frac{d\mathbf{v}}{dt} = \mathbf{A}\mathbf{v} + \mathbf{f} \quad \text{with} \quad \mathbf{v}(0) = \mathbf{u}. \tag{5.39}$$

The solution at time t reads

$$\mathbf{v}(t) = e^{t\mathbf{A}}\mathbf{u} + t\varphi(t\mathbf{A})\mathbf{f} \quad \text{with} \quad \varphi(z) = z^{-1}(e^z - 1), \tag{5.40}$$

or equivalently,

$$\mathbf{v}(t) = \mathbf{u} + t\varphi(t\mathbf{A})(\mathbf{A}\mathbf{u} + \mathbf{f}). \tag{5.41}$$

In the present work, Krylov subspace methods are used to approximate the matrix-vector product involving the matrix $t\varphi(t\mathbf{A})$ and the vector $(\mathbf{A}\mathbf{u} + \mathbf{f})$. Hence

$$t\varphi(t\mathbf{A})(\mathbf{A}\mathbf{u} + \mathbf{f}) \approx \beta \mathbf{V}_m t\varphi(t\mathbf{H}_m)\mathbf{e}_1. \quad (5.42)$$

The evaluation of $t\varphi(t\mathbf{H}_m)\mathbf{e}_1$ is computationally more attractive, and it can be performed using typical dense algorithms such as Padé approximations. In our numerical code, exponential Krylov time integration for the linearized operators \mathbf{A} and \mathbf{A}^* has been motivated by the previous work of Schulze et al. (2009). The implementation of the Arnoldi algorithm (algorithm 1) has been parallelized with the aid of the computational library PETSc (Balay et al., 1997, 2011), and the evaluation of the exponential of the matrix \mathbf{H}_m has been performed using the routine `zgpadm` provided by EXPOKIT Sidje (1998).

From our experience, optimal sizes for Krylov subspaces are in the range $30 < m < 64$, leading to a CFL condition of $\text{CFL} = 16$ for relative errors of 10^{-10} , thus representing about two function evaluations for an equivalent time step with $\text{CFL} = 1$ for the Runge–Kutta method described in chapter 2.

It is important to mention that the high-accuracy of intermediate transformations for the calculation of eigenvalues, presented below, is known to be critical in order to assess reasonable accuracy in the computed eigenvalues. In this respect, exponential Krylov time-integration schemes present a clear advantage owing to the superlinear convergence properties, which permits us to achieve small residuals at reasonable computational cost.

5.2.3 Ritz values: eigenvalues and eigenvectors

Following a similar strategy, eigenvalues and eigenvectors of \mathbf{A} are computed by considering the approximation to the large-scale eigenvalue problem given by

$$(\mathbf{H}_m - \tilde{\lambda}\mathbf{I})\mathbf{v} = 0. \quad (5.43)$$

The solution of the eigenvalue problem involving \mathbf{H}_m is typically obtained by using the QZ method. The approximated eigenvectors are then given by $\mathbf{v} = \mathbf{V}_m\boldsymbol{\xi}$, and the eigenvalues λ are approximated by the Ritz values $\tilde{\lambda}$.

It should be noted that the above technique provides good approximations of the eigenpairs (λ, \mathbf{x}) located in the periphery of the eigenvalue spectrum. For global stability calculations, the eigenpairs of physical importance are those with largest real part (or largest temporal growth-rate). In that case, different techniques such as the time stepping technique Edwards et al. (1994) consisting in the application of the above procedure to the operator $e^{t\mathbf{A}}$ allows to ensure convergence to the least stable eigenvalues.

In our numerical code, the calculation of eigenvalues is carried out using the Krylov–Schur technique (Stewart, 2002) as implemented in the numerical library SLEPc (Hernandez et al., 2005). The time-stepping method presented before has been used in order to favour the convergence of the least stable eigenvalues, and in addition, the harmonic extraction technique implemented in SLEPc has been used to improve convergence on the eigenvalues corresponding to the different frequencies of interest.

Case	L_n	L_w	Description	Analysis
G1	7.34	7.55	Mean flow	Global modes
G1N2W3	2.17	3.58	Mean flow (truncated domain)	Impulse and frequency response

Table 5.1: Size of the domains considered for the global analysis. The global mode analysis has been carried out in the original domain, used for the nonlinear simulations. For the impulse response and the frequency response, the above domain has been truncated in order to reduce the time required for the calculations.

5.3 Choice of base flow, domain sizes and inner product

5.3.1 Base flow

Prior to launching into the global analysis, a base flow must be chosen for the linearization. In stability analyses the base flow is usually taken as the steady-state solution to the Navier-Stokes equations. The typical Reynolds number for the first bifurcation of our aerofoil flow takes place at about $Re = 9000$ (Akbar, 2010).

The steady state for our flow configuration (not shown here) has been computed using the selective frequency damping method (Åkervik et al., 2006), and it presents boundary-layer detachment close to the aerofoil nose, leading to a massive separation bubble that extends over dozens of chord lengths downstream. In the context of aerofoil noise, identical observations have been made by Jones and Sandberg (2011) for a similar flow regime. A theoretical justification for this steady solution at high Reynolds numbers in the asymptotic limit is given in Cheng and Smith (1982).

Despite this traditional approach, it seems reasonable to expect that global stability characteristics of the flow can be assessed by considering the mean flow as the basic state. This hypothesis will be revisited later in view of the results.

5.3.2 Domain sizes

We have defined two cases for the linearized simulations: G1 and G1N2W3; see table 5.1 for details. The case G1 has been used for the calculation of global modes. The case G1N2W3 derives from the case G1, where the difference lies in the fact that the computational domain has been truncated in order to reduce the number of degrees of freedom and accelerate the calculations.

As regards the global mode calculation, the domain size is known to have a significant influence on the computed global spectrum and the associated modes; for this reason we have performed the calculation of the global modes for the case G1N2W3 as a test case to probe the influence of the domain size.

5.3.3 Inner product

The inner product used in the present study is based on the small-perturbation energy norm derived by Chu (1965); Hanifi et al. (1996) for compressible flows. The energy \mathcal{E} contained in a volume \mathcal{V} , written in terms of the variables (p, s, \mathbf{u}) , reads

$$2\mathcal{E} = \int_{\mathcal{V}} \left(\frac{1}{\gamma^2 M^2} \frac{p^2}{\bar{p}} + \frac{\gamma-1}{\gamma^2 M^2} \bar{p} s^2 + \bar{\rho} \|\mathbf{u}\|^2 \right) d\mathcal{V}, \quad (5.44)$$

where $\bar{\rho}$ and \bar{p} stand for the base-flow density and pressure, respectively. The dependence on the physical coordinates \mathbf{x} has been omitted, and all quantities are nondimensionalized. It is then natural to introduce the inner product given by

$$\langle \mathbf{w}, \mathbf{v} \rangle = \int_{\mathcal{V}} \mathbf{w}^H \mathbf{m} \mathbf{v} d\mathcal{V}, \quad \text{with } \mathbf{m} = \begin{pmatrix} \frac{1}{\gamma^2 M^2} \frac{1}{\bar{p}} & 0 & 0 & 0 \\ 0 & \frac{\gamma-1}{\gamma^2 M^2} \bar{p} & 0 & 0 \\ 0 & 0 & \bar{\rho} & 0 \\ 0 & 0 & 0 & \bar{\rho} \end{pmatrix}. \quad (5.45)$$

Once the above expression is spatially discretized in the computational domain, we finally have

$$\langle \mathbf{w}, \mathbf{v} \rangle = \mathbf{w}^H \mathbf{M} \mathbf{v}, \quad \text{with } \text{diag}\{\mathbf{m}_i\} \text{ and } \mathbf{m}_i = \begin{pmatrix} \frac{1}{\gamma^2 M^2} \frac{1}{\bar{p}} & 0 & 0 & 0 \\ 0 & \frac{\gamma-1}{\gamma^2 M^2} \bar{p} & 0 & 0 \\ 0 & 0 & \bar{\rho} & 0 \\ 0 & 0 & 0 & \bar{\rho} \end{pmatrix}_i |\mathbf{x}(\xi)|_i \Delta \xi_i \Delta \eta_i. \quad (5.46)$$

5.4 Impulse response analysis

In this section, we describe the features of the impulse response of the linearized direct operator \mathbf{A} . The impulse response provides insight into both the linearized transient dynamics (a physical point of view) and the properties of the operator \mathbf{A} (a mathematical point of view). More importantly, it enables us to determine, at first glance, whether the linearized dynamics are able to capture the phenomena under investigation; it thus provides evidence for the pertinence of a linear study. We consider the initial value problem (IVP) given by

$$\frac{d\mathbf{v}}{dt} = \mathbf{A} \mathbf{v} \quad \text{with} \quad \mathbf{v}(0) = \mathbf{u}. \quad (5.47)$$

The initial condition \mathbf{u} consists of a random perturbation, with norm one, localized at the leading edge of the aerofoil as illustrated in figure 5.1. We integrate the above IVP using the exponential Krylov time-integration scheme described in section 5.2. The subsequent flow evolution has been computed over $T = 30$ time units, and the snapshots of the flow field have been collected every $\Delta t = 0.01$.

5.4.1 Evolution of the norm: estimation of non-normal effects

In figure 5.2, we present the evolution of the norm of the perturbation (or amplitude) $\|\mathbf{v}(t)\| = \|e^{t\mathbf{A}} \mathbf{u}\|$ as a function of time. It is important to highlight that only the contribution of perturbations within the circle sketched in figure 5.1 is considered. Formally, the above measure is a semi-norm, and it has been chosen such that the features described below become independent of the domain size. In addition, the individual contributions of the pressure, the entropy and the velocity terms to the amplitude are represented.

Two regimes can clearly be distinguished. For $t \leq 5$, transient effects amount to a large increase in amplitude, reaching the maximum $\|\mathbf{v}(t_{\max})\| \approx 2.5 \cdot 10^2$ for $t_{\max} \approx 2.5$. For $t > 2.5$, the amplitude




Figure 5.1: Impulse response of the linearized direct operator \mathbf{A} showing the base flow. The initial condition, also shown, consists of a random perturbation localized at the leading edge. The selected norm includes only the contribution from perturbations within a circle of radius 1.25 centred at mid-chord. Two probes are defined: on the suction-surface (labelled a), and on the pressure surface (labelled b).

decays exponentially at the approximate rate given by

$$\|\mathbf{v}(t)\| \approx \|\mathbf{v}(t_{\max})\| e^{N(t-t_{\max})}, \quad (5.48)$$

with $N \approx -0.18$.

According to the previous expression, it is obvious that the amplitude decreases asymptotically to zero. An estimate of the characteristic time t over which the amplitude reaches a certain fraction of the initial amplitude is given by

$$t = t_{\max} + \frac{1}{N} \log \frac{\|\mathbf{v}(t)\|}{\|\mathbf{v}(t_{\max})\|}. \quad (5.49)$$

Large values of t are given either by: (i) a value N close to zero, or (ii) large value for $\|\mathbf{v}_{\max}\|$ or t_{\max} . For the present case, at about $t \approx 46$, the amplitude of the perturbation is order one. It is important to point out that the above value is considerably larger than the typical period of the oscillations that are observed in the nonlinear simulation, which indicates that even though the operator is stable, perturbations are highly amplified by linear mechanisms.

From a mathematical point of view, the above observations can be related to the properties of the operator \mathbf{A} , in particular, to its non-normality and the non-orthogonality of its eigenvectors. We have the following bounds for the 2-norm $\|\cdot\|_2$ of the exponential matrix $e^{t\mathbf{A}}$ (Trefethen, 1997):

$$e^{t\alpha(\mathbf{A})} \leq \|e^{t\mathbf{A}}\|_2 \leq \|\mathbf{V}\|_2 \|\mathbf{V}^{-1}\|_2 e^{t\alpha(\mathbf{A})}, \quad (5.50)$$




Figure 5.2: Impulse response of the linearized direct operator \mathbf{A} showing the temporal evolution of the amplitude of the initial disturbance. Two regimes can be distinguished. For short times $t \leq 2.5$, the linearized operator exhibits important transient growth of the disturbances. For longer integration times $t > 2.5$, the amplitude decays exponentially, and therefore the operator is found to be stable.

where $\alpha(\mathbf{A})$ is the spectral abscissa of \mathbf{A} , i.e. the largest real part in the eigenvalue spectrum, and the matrix \mathbf{V} is the matrix whose columns are given by the normalized eigenvectors of \mathbf{A} . In this case, it is straightforward to prove that

$$\|e^{t\mathbf{A}}\mathbf{u}\|_2 \leq \|\mathbf{V}\|_2 \|\mathbf{V}^{-1}\|_2 e^{t\alpha(\mathbf{A})}. \quad (5.51)$$

The above indicates that if the spectral abscissa is negative, i.e. the operator is stable, and the transient effects are large, then the condition number of \mathbf{V} is large, or equivalently, the eigenvectors are highly non-orthogonal.

5.4.2 Evolution of the flow field

The aim of this section is to understand the flow features responsible for both the transient growth and the asymptotic regime. The transient-growth mechanisms are analysed in terms of the instantaneous snapshots of the flow field and the evolution of the amplitude from $t = 0$ to $t = 5.5$; see figures 5.3–5.5. For each snapshot, the colour levels have been normalized by the ∞ -norm of $\mathbf{v}(t)$. The flow dynamics in the asymptotic regime are assessed in terms of the entropy signal recorded at the probes shown in figure 5.1: near the reattachment point of the suction-surface separation bubble and near the reattachment point of the pressure-surface separation bubble; see figure 5.6. Note that the exponential decay rate has been factored out.

5.4.2.1 Initial growth of disturbances: from the leading edge to the trailing edge

Let us focus first on the transient-growth mechanisms that appear as the initial perturbation is advected along the pressure- and suction-surface boundary-layer of the aerofoil. The decrease in amplitude at early stages $t < 0.5$, figure 5.3a, can be interpreted in terms of the flow field evolution shown in figure 5.3b–c.

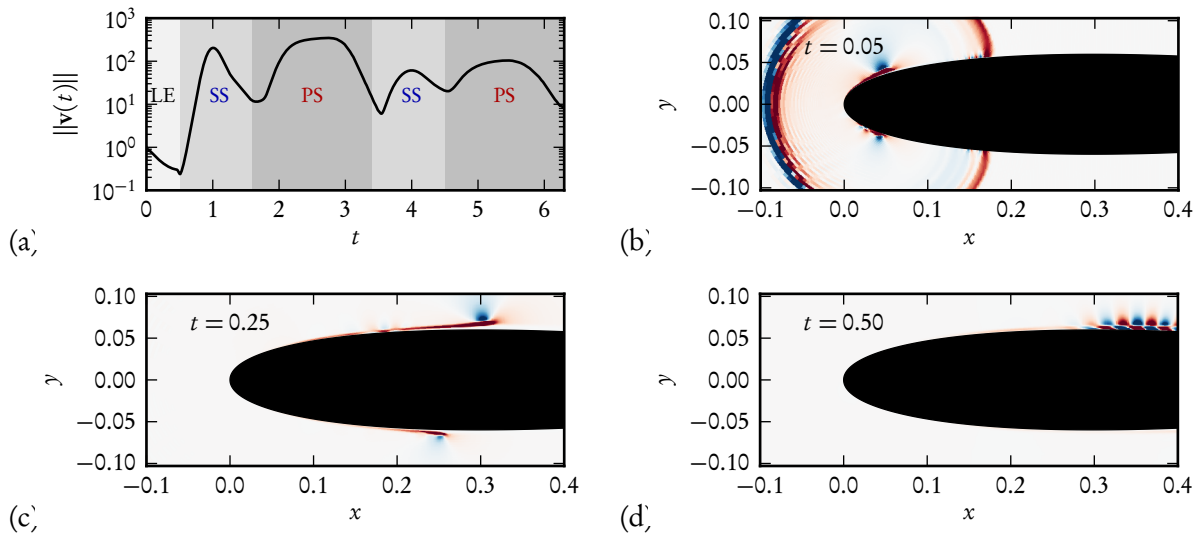


Figure 5.3: Initial growth of disturbances in the vicinity of the leading edge showing the dominant effect of the suction surface.

At $t = 0.05$ (figure 5.3b), a transient acoustic pulse is generated by the initial perturbation, and it propagates radially according to the local speed of sound in the far-field. We can attribute this pressure wave to the fact that the chosen initial condition is not divergence-free. Close to the leading edge, the perturbation separates into travelling perturbations along the suction and the pressure surface owing to the transportation of the disturbance by the base-flow velocity field. From inspection of figure 5.3c, we readily observe that the perturbation does not propagate at uniform speed across the boundary layer: close to the wall, the advection speed is close to zero, whereas far from the wall, the propagation speed reaches the local free-stream velocity.

In detail, and within the limits of a parallel-flow approximation, the perturbation can be interpreted as a combination of acoustic, free-stream and boundary-layer modes that are obtained in the spatial stability analysis of the local velocity profiles. As the perturbation is advected, it evolves according to the local dispersion relation: each single mode is characterized by its respective phase speed, and grows or decays in space according not only to the spatial growth rate, but also to the projection onto the modes of the immediate downstream profile. Note that the advection speed of the emerging wave packets is given by the group velocity, since they consist of a superposition of local modes of the same branch.

We observe the propagation in the free-stream of the initial perturbation at the free-stream velocity (which is higher on the suction surface than on the pressure surface) and without spatial decay; see figure 5.3c. Similar observations can be made about the acoustic wave within the boundary layer as observed in figure 5.3b.

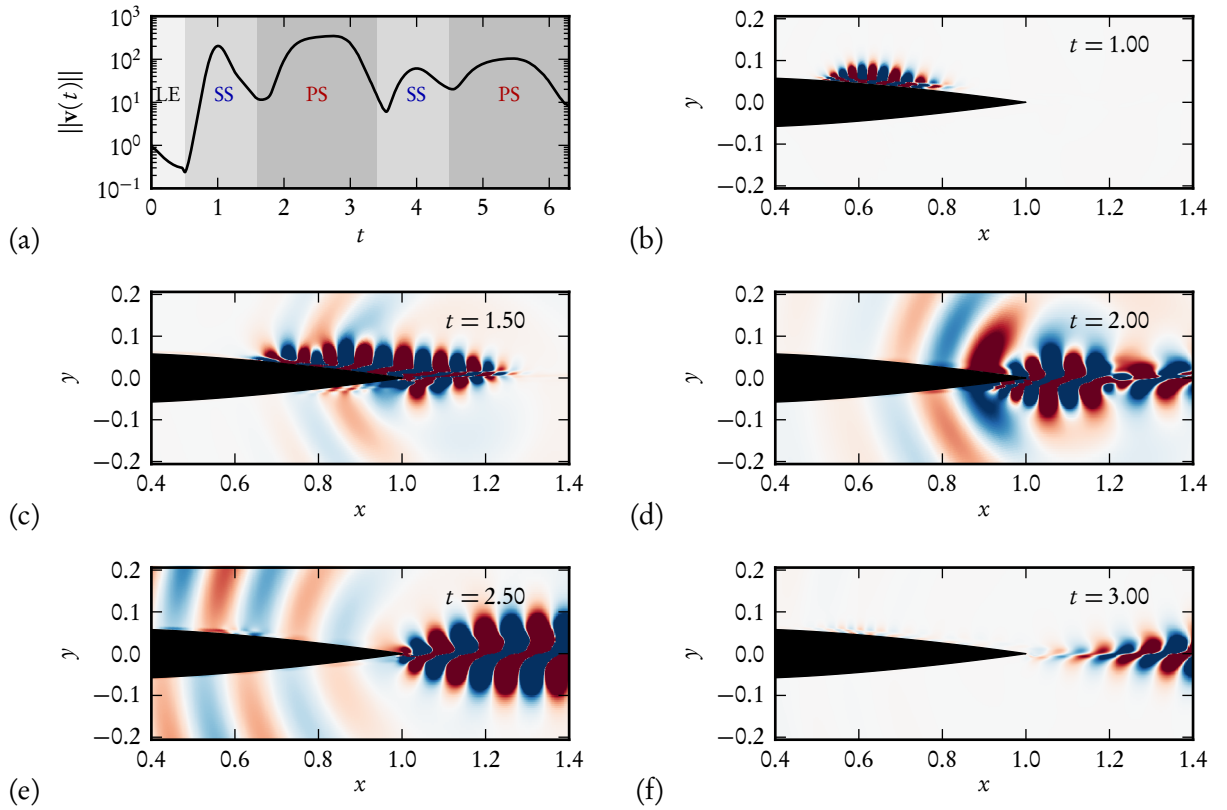


Figure 5.4: Impulse response: convective growth of instabilities on the pressure-surface and suction-surface boundary layers. Substantial acoustic radiation is observed after the passing of disturbances at the trailing edge.

However, a fraction of the perturbation in the boundary layer projects onto boundary-layer modes, that may present spatial growth or decay. We recall that close to the leading edge, the boundary layers are subjected to positive pressure gradients ($x < 0.04$ for the suction surface and $x < 0.24$ for the pressure surface, see figure 3.11a), and for the given flow regime, this indicates that the boundary-layer modes are stable. The fraction of the perturbation that travels inside the boundary layer is then damped as it travels along convectively stable profiles, from where we can explain the initial decrease in amplitude of the perturbation.

In figure 5.3d, we observe the growth of a wave packet on the suction surface that can be attributed to inviscid-type instability mechanisms. We recall that the stability characteristics of the velocity profiles is different for both surfaces: the suction-surface velocity profiles exhibit inflection points substantially closer to the leading edge than the pressure-surface velocity profiles; see figure 3.13. On the pressure surface, the boundary layer is expected to be convectively stable for $x < 0.24$. The above reasons indicate why the global linear dynamics for $0.5 < t < 1$ is predominantly localized on the suction surface.

The subsequent flow evolution is represented in figure 5.4. For $0.5 < t < 1$, the dominant mechanisms are related to Kelvin-Helmholtz instabilities within the shear layer of the suction-surface separation bubble. A local maximum in amplitude (figure 5.4a) is reached at $t = 1$, and it coincides with the instant where the wave packet approaches the reattachment point $x_{r,s} = 0.66$; see figure 5.4b

and table 3.4. For $1 < t < 1.6$, we observe a decrease in amplitude as the wave packet passes the separation bubble and reaches the reattached flow. This behaviour can be related to a mismatch in the stability characteristics between separation bubble and the reattached flow: the wave packet does not project efficiently into convectively unstable waves in the reattached-flow region. After passing the trailing edge, the wave packet produces scattering of acoustic waves with a phase-shift of 180° between the upper and lower half-plane (figure 5.4c).

For $1.5 < t < 2.7$, the amplitude increases due to the growth of a wave packet on the pressure-side separation bubble; see figure 5.4c–e. Note that although perturbations are damped close to the leading edge, for $0.3 < x < 0.7$ the velocity profiles are inflectional (figure 3.13b) and sustain spatial growth. In the temporal evolution, we observe a slow spatial amplification for $0.3 < x < 0.7$. However, most of the growth takes place in the shear layer of the separation bubble $0.7 < x < 1$. It is reasonable to expect that the spatial growth rates of the attached (but inflectional, see figure 3.13c) profiles are comparatively smaller than the spatial growth rates in the separation bubble where there is reverse flow.

A more intense acoustic radiation is observed as the pressure-surface wave packet passes the trailing edge, compared to the suction surface wave-packet. Although in both cases the radiation has the same qualitative characteristics, it is found that: first, the frequency is slightly lower for the acoustic waves generated at the pressure surface than for the ones from the suction surface, and second, and more importantly, the acoustic field produced by the waves from the pressure surface is more intense than the one from the suction surface. For $2.7 < t < 3.5$, the wave packet is located in the wake, where the dynamics are dominated by convection.

At this point, it can be concluded that no region of absolute instability, i.e. regions where the growth of wave packets with zero group velocity takes place, exists in the flow. Further confirmation has been obtained by analysing the local velocity profiles: all velocity profiles exhibit negative temporal growth-rates at the pinch point in the complex plane.

5.4.2.2 Regeneration of wave packets: feedback effects

The flow evolution for $3.5 < t < 6$, depicted in figure 5.5, displays the same sequence of phenomena that has been observed for $0.5 < t < 3.5$. A first wave packet, issued from the boundary layer on the suction surface, emerges (figure 5.5b), and it scatters acoustic waves as it passes by the trailing edge (figure 5.5c). A second wave packet emerges on the pressure surface (figure 5.5d); at its passage by the trailing edge, it radiates substantial acoustic energy into the far-field (figure 5.5e).

It is important to notice that for $0.5 < t < 3.5$, the instabilities were triggered by the initial perturbation at the leading edge. In contrast, for $3.5 < t < 5$, the triggering mechanism can be attributed to the receptivity of both boundary layers to free-stream disturbances originated downstream. This idea is supported by the observation that the period between consecutive wave packets is of the same order of magnitude as the convective time over the aerofoil length. The above observation raises the question whether the receptivity mechanism can mainly be attributed to the acoustic waves or to the hydrodynamic disturbances induced by the wave packet at the trailing edge.

On the one hand, the hydrodynamic features would excite the boundary layers upstream by a mechanism of forced receptivity due to a match in frequency and wavelength. On the other hand, the acoustic waves have the same frequency but a different wavelength than the hydrodynamic features and would then require a mechanism of natural receptivity. Although the forced receptivity mechanisms are more efficient than the natural receptivity, it should be noted that the acoustic radiation decays, with the distance d to the source as $d^{-1/2}$, whereas the hydrodynamic perturbations decay exponentially with d . Consequently, as far as the receptivity mechanism is concerned, the

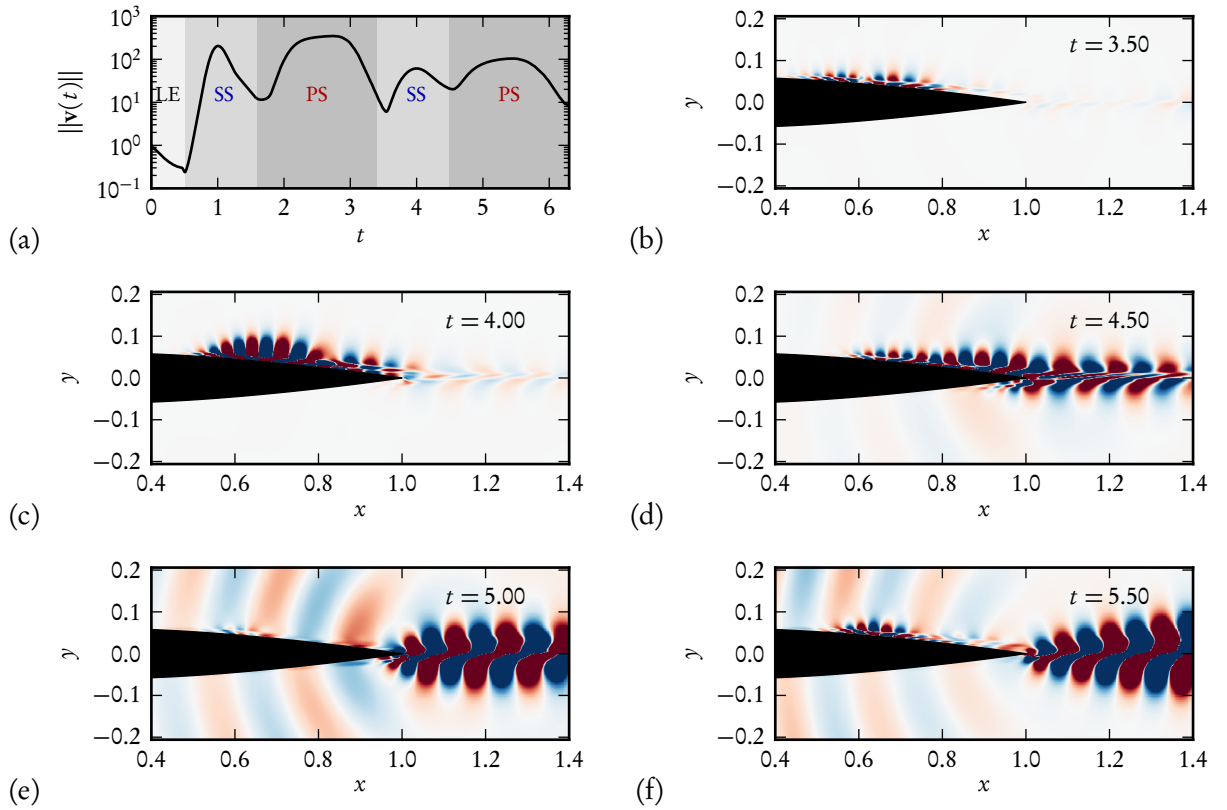


Figure 5.5: Impulse response: feedback-loop showing instantaneous perturbation evolution displaying pressure levels and entropy contours.

far-field acoustic radiation dominates over the far-field hydrodynamic disturbances (see for instance figure 5.5f).

It is thus found that, for the present case, there exists a complex feedback loop involving the hydrodynamic features on the (i) pressure- and (ii) suction surface of the aerofoil and (iii) the near wake. The above features are related by feedback-loop effects due to the natural receptivity of the boundary layers to the upstream-propagating acoustic waves.

5.4.2.3 Cross-talk between the suction and pressure surface

At this point, it is instructive to analyse the long-time behaviour in order to gain insight into the features of the feedback loop. To this end, we present in figure 5.6 the entropy signal, where the temporal decay has been factored out, recorded at the probe locations indicated before: on the suction surface at the reattachment point (figure 5.6a), and on the pressure surface near the trailing edge (figure 5.6b). Since we focus on the hydrodynamic features in the boundary layer that propagate downstream and not on the upstream-propagating acoustic waves, we have chosen the entropy as the measure.

For the transient regime observed before, we readily identify the period of wave packet shedding on the suction surface as $\Delta T_{ss} \approx 2.96$, and that on the pressure surface as $\Delta T_{ps} \approx 2.78$.




Figure 5.6: Impulse response: entropy signal at (a) the suction-side probe, (b) the pressure-side probe, displaying the instantaneous signal in grey, the envelope in black and the low-frequency component in green.

For longer times, $t > 6$, there is an increase in complexity of the behaviour as the acoustic radiation, emitted by wave-packets from either surface, is susceptible to interact with the opposite boundary layer. Indeed, cross-talk effects are observed, since the shedding period of wave packets on the pressure side ΔT_{ps} appears spontaneously on the suction surface signal, thus indicating that the boundary layer on the pressure surface exerts, *a priori*, a dominant effect.

The attention is now directed towards the average given by a time-moving window of the recorded signals, shown in figure 5.6 and displayed in green. We can readily identify a low-frequency long-period oscillation that seems to correspond to the low-frequency oscillation presented in figure 3.18.

5.4.3 Summary

In this section, we have presented a detailed description of the impulse response. The linearized compressible Navier–Stokes equations operator \mathbf{A} has been found stable. However, and as in most high-Reynolds number flows, important transient-growth effects can be observed.

An analysis of the perturbation evolution reveals the different hydrodynamic features of the flow under investigation. The acoustic radiation at the trailing edge is more intense for wave packets issued from the pressure surface than for those issued from the suction surface.

No absolute instability, which can be interpreted as a *local feedback*, has been found. In contrast, the regeneration of hydrodynamic instabilities upstream strongly suggests the existence of *global feedback* effects. The feedback effects link the dynamics of the pressure surface and the suction surface, and the above results indicate that the pressure-surface boundary layer has a dominant effect.

5.5 Global modes

In the previous section, we have described the dynamics of the linearized Navier–Stokes operator \mathbf{A} following the temporal evolution of an initial impulse. As we were interested in the dynamics of the flow near the aerofoil surface, the initial condition was intentionally specified as a localized disturbance at the leading edge. In the ensuing temporal evolution, important transient effects have been observed as the perturbation travels downstream along the aerofoil chord until it is ultimately shed into the wake. As a result, we have found that, in the long-time behaviour, a weakly-damped feedback loop appears with a characteristic period involving the dynamics of both boundary layers and the acoustic radiation at the trailing edge.

In this section we present a complementary approach to gain further insight into the description of the linearized dynamics: within the global-modes framework, the temporal evolution of the impulse response can be interpreted as that of a linear combination of *direct* global modes whose relative coefficients are given by the projection of the initial condition onto the corresponding *adjoint* global modes.

From that perspective, an alternative interpretation of the impulse response can be provided. For the short time scale, all direct modes (regardless of their growth-rate) whose corresponding adjoint has significant spatial support at the location of the initial impulse are important for the short-time dynamics of the flow. On a long time scale, the dynamics are mainly related to the least-damped global modes.

In this section, we focus on the features of the least-stable direct and adjoint global modes, which may be responsible for the observed feedback loop, and more importantly, for the structure of the acoustic spectrum in the nonlinear simulations as well as in experimental investigations.

5.5.1 Computation: eigenvalue problem set-up

Prior to launching into the global modes analysis, we summarize here the parameters of the numerical calculation that permit us to obtain the global spectrum.

As already discussed, we are interested in computing the least-stable (or leading) eigenpairs of the direct operator \mathbf{A} and its adjoint \mathbf{A}^* . To this end, we have applied a Krylov–Schur technique (presented earlier), as implemented by the eigensolver SLEPc, to the operator that represents the temporal advancement over time Δt for the direct and the adjoint operators, i.e. $e^{\Delta t \mathbf{A}}$ and $e^{\Delta t \mathbf{A}^*}$, respectively.

Our goal is to obtain a large portion of the global spectrum. The parameters of the eigenvalue solver are set as follows: first, for given computational resources, we determine the maximum temporal advancement T that can be performed and, second, the available amount of memory determines the size m of the largest Krylov subspace that can be constructed; the time-step is then given by $\Delta t = T/(m - 1)$. Note that in the opposite case where only a few eigenvalues are sought but higher accuracy is required, the considerations are different: the considered Krylov subspace is smaller and relies on restarting techniques.

In the present calculations, the typical size of Krylov subspaces is 2048, and a typical time step is taken as $\Delta t = 0.015$. The results from the nonlinear simulations (table 3.5) hint at $\omega = 2\pi f \approx 44$ as the most probable location for the global modes of interest. In order to improve the convergence of the eigenvalues at this location in the complex plane, we have enabled the use of the harmonic extraction technique in the eigensolver with the shift parameter σ as $e^{i\Delta t \omega_r}$.

It is important to mention that the estimate of the relative error with respect to the propagator is commonly more optimistic than the actual error with respect to the linearized operator \mathbf{A} . For this

reason, we measure the degree of convergence according to the relative residual given by

$$\varepsilon = \frac{\|\mathbf{A}\mathbf{v} - \lambda\mathbf{v}\|_2}{\|\lambda\mathbf{v}\|_2}. \quad (5.52)$$

The computed spectrum of global modes is depicted in figure 5.7, showing the location of the eigenvalues in the complex plane according to the eigenfrequency $-i\omega = \lambda$, as is conventional in global stability analyses, and coloured according to their relative residual; see equation (5.52). The shift parameter $\omega = 44$ for the harmonic extraction method is indicated in the ω -plane by a black cross and the iso-contours of distance to the target (dashed lines); note that the eigenvalues with the smallest residual (near 10^{-6}) correspond to the ones closest to the target point.

A discussion about the robustness of the global spectrum with respect to the extent of the physical domain is postponed to the end of this section.

5.5.2 Spectrum

We present in figure 5.7 the global spectrum of the linearized compressible Navier–Stokes operator \mathbf{A} . Each direct global mode \mathbf{v} is represented in the complex plane by its associated complex frequency or eigenfrequency $\omega = \omega_r + i\omega_i$, where ω_r is the angular frequency and ω_i is the temporal growth-rate. The associated adjoint global mode is denoted by \mathbf{w} , and its eigenfrequency is $-\bar{\omega} = -\omega_r + i\omega_i$.

We recall that since the operator \mathbf{A} is real, the spectrum is symmetric with respect to the vertical axis ($\omega_r = 0$). More precisely, if (ω, \mathbf{v}) is an eigenpair and $(-\bar{\omega}, \mathbf{w})$ is the corresponding adjoint, then $(-\bar{\omega}, \bar{\mathbf{v}})$ is also an eigenpair and its adjoint is $(\omega, \bar{\mathbf{w}})$. As a consequence, it suffices to only focus on the features of the global modes in either the left or the right half-plane.

A first inspection of figure 5.7 confirms that the operator \mathbf{A} is stable, since all the eigenvalues fall into the lower half of the complex plane $\omega_i < 0$. No isolated branches of global modes can be observed in the spectrum. This suggests that the features of the direct and adjoint global modes vary, at least in principle, in a continuous manner.

For varying frequency, the growth rate associated with the least damped mode exhibits two maxima of temporal growth-rate: (i) at $\omega = 43.86 - 0.15i$ and (ii) close to the origin $\omega = 0$. These maxima correspond, within a reasonable approximation, to the main frequencies observed in the acoustic tones from the nonlinear simulations, namely, multiple acoustic tones centred near $\omega = 2\pi f \approx 41.84$ and low-frequency components (see table 3.5, case G1, for reference).

5.5.3 Leading modes: the coupled dynamics of the separation bubbles

In figure 5.8 we provide a zoomed view of the spectrum for the subdomain comprising the leading global modes.

Inspection reveals that the leading global modes consist of multiple local maxima of growth-rate for varying frequency, centred around the maximum growth-rate peak at $\omega_r = 43.86$ and $\omega_i = -0.15$. The frequency difference between consecutive maxima is constant and equal to $\Delta\omega_r = 2.44$, and the temporal growth for the local maxima is approximately $\omega_i \approx -0.15$. Between the peaks, the frequency difference of neighbouring eigenvalues is $\Delta\omega_r = 0.34$, and their minimum growth-rate is approximately $\omega_i = -0.25$.

Our investigation will now focus on the quantitative features of the leading direct and adjoint global modes. We thus consider the individual modes labelled M1–M4 in the spectrum (see figure 5.8) corresponding to the local maxima (modes M1 and M4) and the eigenvalues in-between (modes M2




Figure 5.7: Global spectrum. The colours indicate the relative error with respect to the operator \mathbf{A} as given in equation 5.52.

and M3). In particular, our aim is to quantify the similarities and the differences between the modes at and between the maxima, (M1, M4) and (M2, M3), respectively.

5.5.3.1 Direct modes

In figure 5.9 we present the direct global modes previously labelled as M1–M4. In order to facilitate a comparison, each global mode has been normalized by the maximum value of the velocity in the near wake $1 < x < 1.2$, which for the displayed modes occurs at about $x = 1.15$. For each mode, we display the real part of the stream-wise velocity component for the downstream half of the aerofoil (left column) and the real part of the associated pressure in the near-field of the aerofoil (right column).

From the velocity fields at the aerofoil, we conclude that the modes represent boundary-layer instabilities with predominant spatial support downstream of the reattachment point of the suction-surface separation bubble and in the wake. As expected, for increasing frequency the characteristic wavenumber of the oscillations increases.

Under the chosen scaling, the spatial growth becomes noticeable on the suction surface for $x > 0.5$ and on the pressure surface for $x > 0.95$. These locations fall within the respective regions given by the separated shear-layer and the reattachment point of the separation bubble. Therefore, we can deduce that the family of global modes in the range $40 < \omega < 50$ represent the coupled global dynamics of both separation bubbles.

In the pressure field we observe the acoustic footprint of the boundary-layer instabilities. In the near-field, the modes display substantial levels of acoustic radiation into the free-stream with higher wave-lengths. The radiation can be attributed to a dipolar acoustic source contained in the global




Figure 5.8: Detailed view of the global spectrum showing the least stable (or leading) global modes. The modes M1–M4 are considered, and the corresponding eigenvalues are $\omega_{M1} = 43.86 - 0.15i$, $\omega_{M2} = 43.86 - 0.15i$, $\omega_{M3} = 44.63 - 0.23i$ and $\omega_{M4} = 46.30 - 0.15i$. The relative residual is below 10^{-5} for both the direct and adjoint modes.

mode at the trailing edge, whose acoustic pattern corresponds to cylindrical waves with a shift of 180° between the upper and lower half-plane. Furthermore, the acoustic waves show a preferred upstream directionality.

We see that the chosen global modes exhibit nearly identical features in both the hydrodynamic instabilities and the associated acoustic radiation. We then focus on the evolution of the maximum peak of stream-wise velocity (figure 5.10) and of the pressure-surface (figure 5.11) on both surfaces of the aerofoil.

The evolution of the stream-wise velocity peak on the suction surface (figure 5.10) displays the exponential growth of the instability waves at $0.53 < x < 0.62$. This range coincides with the separated shear-layer at the edge of reattachment. As already hinted by the impulse response analysis, further downstream the mode decays exponentially in space until it reaches the trailing edge. Further upstream, the velocity values are approximately constant (in logarithmic scale) and they correspond to the levels of velocity fluctuations outside the boundary layer.

On the pressure surface, the exponential growth of stream-wise velocity takes place near the trailing edge $0.9 < x < 1$, where the base flow exhibits a separated shear-layer at the edge of the reattachment. Similarly to what has been observed for the pressure side, upstream values of the velocity are nearly constant (in logarithmic scale) and are related to velocity fluctuations outside the boundary layer.

The evolution of the pressure fluctuations on the suction and pressure surface reveals the appearance of standing waves on the suction surface and, to a lesser extent, on the pressure surface.

On the suction surface, the nodes of the standing wave are located at $x \approx 0.5$ and are uniformly distributed for $0.8 < x < 1$. Note that their location varies significantly for the different modes. On the pressure surface, however, the node is located at $x \approx 0.95$. A detailed inspection shows that the modes with the higher growth rates (M1 and M4) are correlated with the location of the suction-surface node at $x \approx 0.95$, which corresponds to the location of the node on the pressure surface.

It is important to point out that the identified region of dominant exponential growth does not necessarily imply that all growth takes place at this location: upstream, the amplitude of the hydrodynamic perturbations is far smaller than the induced level of fluctuations associated with

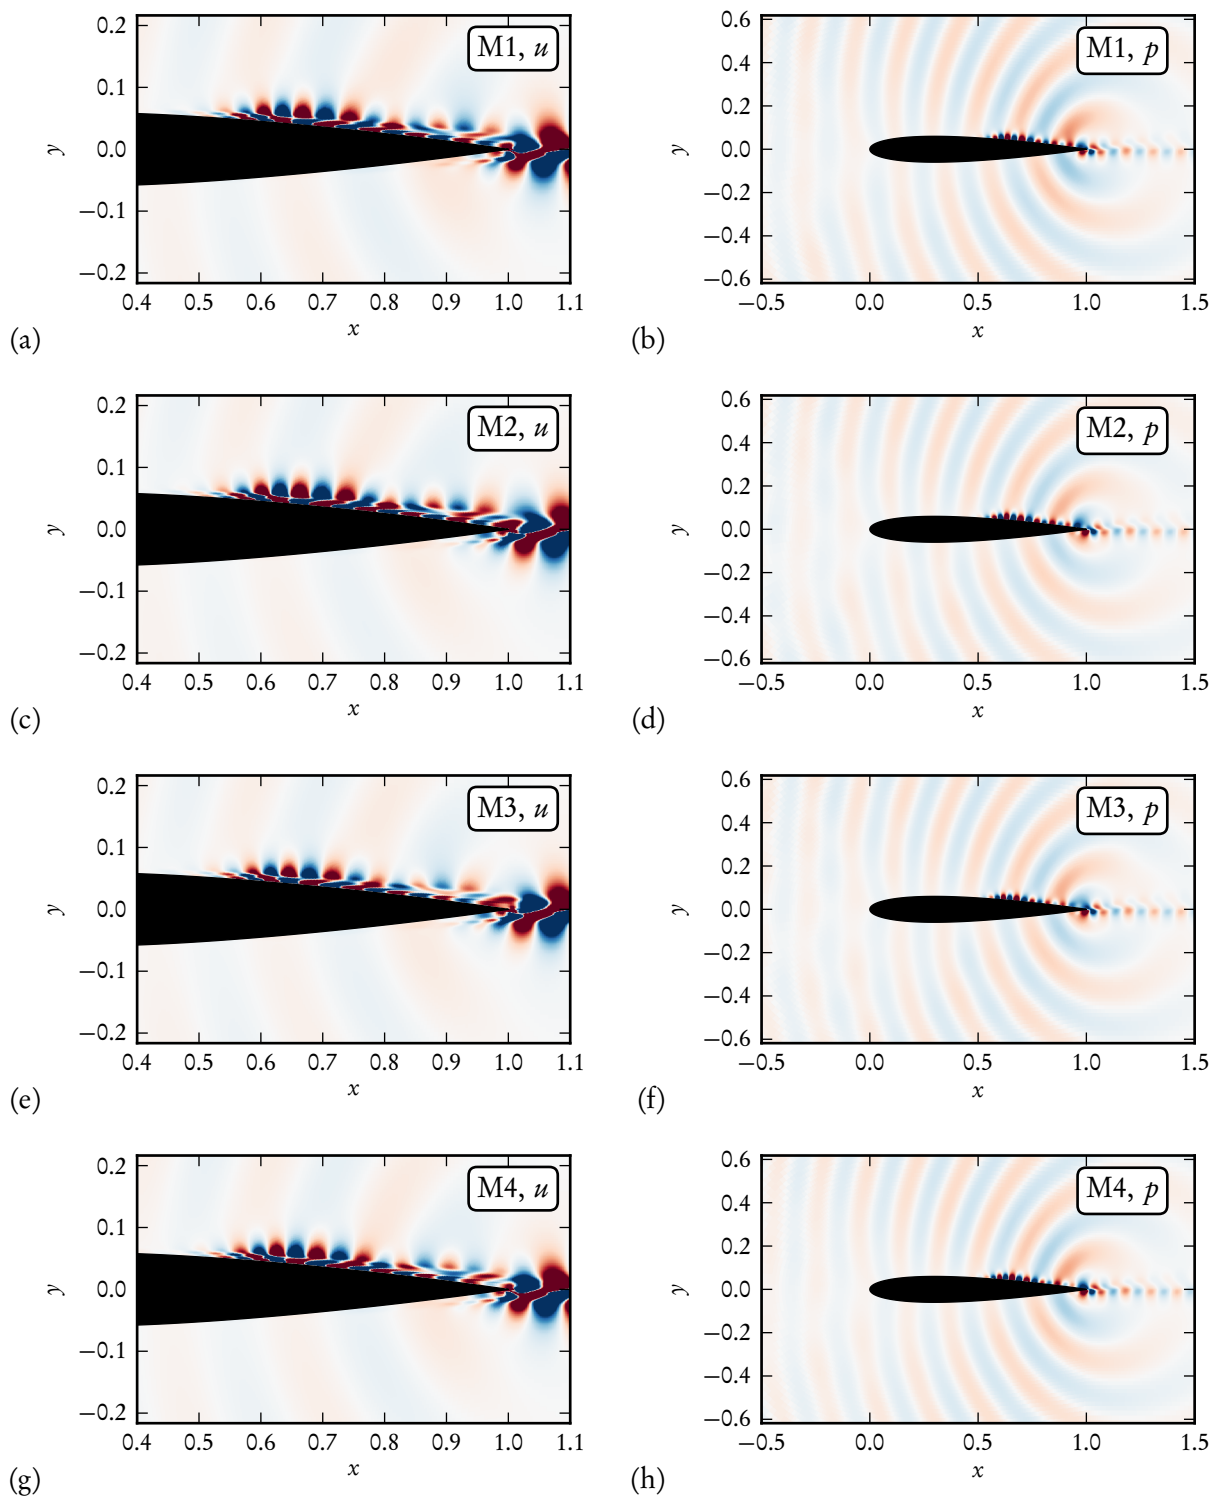


Figure 5.9: Spatial structure of the global modes labelled M1–M4 in figure 5.8, displaying the real part of the velocity levels in the vicinity of the aerofoil surface (left column), and the real part of the associated near-field pressure levels (right column). The modes presented here have been normalized by the maximum value of the velocity field in the near wake $1 < x < 1.2$.




Figure 5.10: Evolution of the stream-wise velocity peak along the aerofoil chord (a) on the suction- and (b) on the pressure-surface of the aerofoil for the modes labelled M1–M4 in figure 5.8.




Figure 5.11: Evolution of the pressure peak along the aerofoil chord (a) on the suction- and (b) on the pressure-surface of the aerofoil for the modes labelled M1–M4 in figure 5.8.

the acoustic waves. Since the acoustic waves and the hydrodynamic perturbations occur at the same frequency, we are unable to determine precisely the exact location of the origin of exponential growth.

5.5.3.2 Comparison to the frequential snapshots from the nonlinear simulation

The leading direct global (M1) mode can now be compared to the frequential snapshots; see figure 5.12. At first sight, the leading global mode (figure 5.12a) is in good agreement with the features that




Figure 5.12: Comparisons between (a) the leading global mode in the spectrum (M1) and (b) the frequential snapshot from the nonlinear simulation (see figure 3.17c,d). Both snapshots have been normalized by the maximum stream-wise velocity amplitude on the suction surface.

are observed in the nonlinear simulation: the growth of instabilities on the suction surface at the separation bubble and at the reattachment point of the pressure-surface boundary layer. The source of noise is readily captured by the linear global modes together with the radiation pattern.

However, the following qualitative differences are observed.

In the nonlinear simulations, downstream of the reattachment point on the suction surface, the flow field shows saturated levels of velocity fluctuations due to the nonlinearities that are triggered by finite-amplitude perturbation levels. At this location, the perturbation interacts nonlinearly with the outer flow, producing vortices that are advected downstream. Within the linearized framework, downstream of the reattachment point the hydrodynamic instabilities are still governed by the linear dynamics of the reattached mean-flow boundary layer. As already discussed in our analysis of the impulse response, the separation bubble at the reattachment point and the boundary-layer that subsequently forms downstream have rather dissimilar stability characteristics.

Likewise, vortices that are shed into the wake by both boundary layers also reach saturated levels and give rise to a quasi-periodic vortex street that is observed in the nonlinear simulations, including an interaction between the two-dimensional vortices. Within the linear framework, the instability waves are convected continuously into the wake, where convection effects dominate. Additionally, due to linearity perturbations cannot mutually interact.

Case G1	Frequency peaks				
	ω_{-2}	ω_{-1}	ω_0	ω_1	ω_2
linear: global modes	—	39.26	41.38	43.86	46.30
nonlinear: Fourier	36.94	39.46	41.84	44.30	—
relative difference	—	0.5%	1.1%	1%	—

Table 5.2: Comparison between the tones in the acoustic spectrum for the nonlinear simulation (see table 3.5, case G1) and the frequency of the global modes with local maxima of temporal growth rate in the global spectrum. The peak with maximum growth rate and the loudest tone in the nonlinear simulation are indicated in bold.

In spite of the limitations of linear stability theory, the resemblance between the direct leading mode and the frequential snapshot from the nonlinear simulation is remarkable.

It is hence tempting to compare the occurrence of equally-spaced local maxima in the global spectrum with the location of the equally spaced tones in the acoustic spectrum; see table 5.2. For frequencies where this comparison can be performed, the agreement between the acoustic tones and the spacing in the global spectrum is excellent: the maximum relative error is 1.1%.

The above observations suggest that the tonal noise phenomenon and the occurrence of multiple equally-spaced peaks in the acoustic spectrum have their physical origin reflected in the characteristics of the global spectrum, and, more importantly, the principal mechanism can be ascribed to the linearized dynamics of the flow. In light of these results, it seems legitimate to further employ linear-stability tools to gain additional insight into the noise generation mechanism.

5.5.3.3 Superposition of global modes: aeroacoustic feedback-loop

We observed in the impulse response analysis (in section 5.4) that, triggered by the passage of perturbations by the trailing edge of the aerofoil, an aeroacoustic feedback loop establishes itself and subsequently drives the disturbance evolution. The features of this feedback loop were then described phenomenologically in terms of the dynamics of hydrodynamic and acoustic wave packets, and their origin was linked to local properties of the base flow. Even though for short times we observed distinct wave packets, for longer times, we could no longer identify individual wave packets but rather high-frequency components modulated by low-frequency oscillations with a period of $T_f \approx 2.8$.

With the direct global modes extracted, we are now in a position to provide a complementary description of the feedback loop in terms of the dynamics of the least stable modes: for sufficiently long time, the temporal evolution is mainly given by a superposition of the least-stable global modes, since the remaining modes have comparatively smaller growth rates. As was discussed before, the global modes display the most relevant physical features that are also observed in the nonlinear simulations, matching in spatial structure as well as in frequency. The question remains, then, how the global modes can account for the above feedback loop.

To answer this question, we consider the superposition of two plane waves with similar spatio-temporal characteristics: their respective frequencies are taken as $\bar{\omega} + \Delta\omega$ and $\bar{\omega} - \Delta\omega$, and the

wavenumbers are $\bar{\alpha} + \Delta\alpha$ and $\bar{\alpha} - \Delta\alpha$. Hence their superposition is given by

$$A(x, t) = e^{i[(\bar{\alpha} + \Delta\alpha)x - (\bar{\omega} + \Delta\omega)t]} + e^{i[(\bar{\alpha} - \Delta\alpha)x - (\bar{\omega} - \Delta\omega)t]} = \quad (5.53)$$

$$= 2 \cos(\Delta\alpha x - \Delta\omega t) e^{i(\bar{\alpha}x - \bar{\omega}t)} \quad (5.54)$$

$$= 2 \cos(\Delta\alpha(x - c_g t)) e^{i\bar{\alpha}(x - c_p t)}, \quad (5.55)$$

with the phase velocity $c_p = \bar{\omega}/\bar{\alpha}$ and the group velocity $c_g = \Delta\omega/\Delta\alpha$. From this expression, it is readily observed that for $\Delta\omega \ll \bar{\omega}$ and for $\Delta\alpha \ll \bar{\alpha}$ the evolution can be interpreted as oscillations at the average frequency $\bar{\omega}$ and wavenumber $\bar{\alpha}$, modulated by a low-frequency component at half the frequency difference $\Delta\omega$ and half the wavenumber difference $\Delta\alpha$. Furthermore, the velocity of the low-frequency component is given by the group velocity c_g , which is different from the phase velocity of the high-frequency component c_p . In other words, the waves display *beating*, which arises from their linear superposition and is linked to their spatio-temporal resemblance.

In an analogous manner, global modes with similar spatial structures and nearly matched frequencies can also exhibit beating. The evolution of the flow can be described by a superposition of global modes $\mathbf{v}(t) = \sum_i a_i \mathbf{v}_i e^{-i\omega_i t}$. For each pair of modes \mathbf{v}_k and \mathbf{v}_l , we can write their superposition as

$$a_k \mathbf{v}_k e^{-i\omega_k t} + a_l \mathbf{v}_l e^{-i\omega_l t} = 2 [\cos(\Delta\omega t) \bar{\mathbf{v}} - i \sin(\Delta\omega t) \Delta\mathbf{v}] e^{-i\bar{\omega}t}, \quad (5.56)$$

where $\bar{\omega} = (\omega_l + \omega_k)/2$, $\Delta\omega = (\omega_l - \omega_k)/2$, $\bar{\mathbf{v}} = (a_l \mathbf{v}_l + a_k \mathbf{v}_k)/2$, and $\Delta\mathbf{v} = (a_l \mathbf{v}_l - a_k \mathbf{v}_k)/2$. The above expression shows that the same decomposition in terms of a high-frequency carrier wave and a low-frequency modulation can be performed. In the context of fluid dynamics, this beating between global modes has been observed in several fluid cases; see for instance Schmid and Henningson (2002) for the case of a falling liquid curtain and Ehrenstein and Gallaire (2008); Cherubini et al. (2010) for the beating of separation bubbles.

For the present case, we can estimate the typical period of the feedback loops from half the difference between the peaks at both extremes of the leading modes; we get $\Delta\omega = 2.44$. This estimate is in good agreement with the angular frequency associated with the period of the feedback loops that was observed in the impulse response: $\omega_{ps} = 2\pi/\Delta T_{ps} \approx 2.26$. The difference in frequency can be attributed to the fact that the full interaction involves multiple modes with similar frequencies, which complicates the mechanism.

The propagation velocity of the wave packets is given by the group velocity: $c_g = \Delta\omega/\Delta\alpha$, where $\Delta\alpha$ is obtained from the spatial structure of the global modes involved in the feedback loop. For the aerofoil boundary layers, the instability waves are dispersive, and therefore the group velocity differs from the phase velocity. In contrast, the propagation of the acoustic pulse in the far-field coincides with the propagation speed of the individual waves as the far-field flow is non-dispersive.

From a global perspective, it can be concluded that the feedback loop that appears in the impulse response (and that has also been observed in the nonlinear simulation) arises from the interaction of the leading global modes with similar frequencies. Their *modal* interaction describes the growth of wave packets on the suction and pressure surfaces and includes the radiation of acoustic waves. We emphasize that the feedback-loop phenomenon cannot be attributed to any individual global mode.

5.5.3.4 Adjoint modes

The description of the perturbation dynamics in terms of global modes is incomplete if no thoughts are given to the processes by which an initial perturbation (or equivalently, a given forcing) projects

onto the involved direct global modes. In more physical terms, such a projection is related to the receptivity mechanism.

As was discussed in the introduction of the present chapter (section 5.1), the adjoint global modes represent a powerful theoretical concept to assess receptivity mechanisms of the flow: an arbitrary perturbation can be expressed, in principle, as a linear combination of global modes $\mathbf{u}_{||} = \sum_j a_j \mathbf{v}_j$. In

order to obtain the coefficients a_i , we need to evaluate the inner product on both sides with respect to the corresponding adjoint global mode \mathbf{w}_i . By invoking the bi-orthogonality condition, and noting that we have chosen the adjoint global modes such that $\langle \mathbf{w}_i, \mathbf{v}_i \rangle = 1$, we then arrive at $a_i = \langle \mathbf{w}_i, \mathbf{u} \rangle$.

The above expression highlights that the individual components of the adjoint global mode, i.e. the adjoint pressure, entropy and velocity components at every grid point, can be interpreted as the projection of a unit impulse in the corresponding variable at that location onto the direct global modes, from where we can evaluate quantitatively the receptivity mechanisms of the flow. From an optimization point of view, the adjoint global mode corresponds to the perturbation that produces the largest projection onto the corresponding direct mode¹, and its norm $\|\mathbf{w}_i\|$ constitutes a useful measure to evaluate the overall receptivity.

The adjoint global modes corresponding to the direct global modes L1–L4 are illustrated in figure 5.13 showing the magnitude of the adjoint stream-wise velocity $|u^*|$ in logarithmic scale together with the representation (in linear scale) of the real part of the adjoint stream-wise velocity (see insets). The magnitude of the adjoint mode is a meaningful measure to compare the receptivity between different adjoint modes.

The spatial distribution of the four adjoint global modes for M1–M4 is very similar: the support of the adjoint mode is highly localized on the pressure-surface boundary layer, between the leading edge and upstream of the separation point, i.e., $0.1 < x < 0.75$, where the maximum value is approximately 10^9 . Considerable levels of receptivity are also observed in the free-stream in the region where the base flow convects perturbations to the pressure-surface boundary layer. Note that the levels of receptivity on the suction surface are about 10^3 times smaller than the receptivity measures on the pressure surface; see figure 5.14. The real part of the adjoint velocity, shown in the zoomed view in figure 5.13, displays features reminiscent of the Orr-mechanism, with flow structures tilted against the base-flow shear.

The difference in spatial support between the direct global modes and the adjoint global modes is striking and has implications on the understanding of the physical phenomena.

It is well-known that in convectively unstable flows, one of the common signs of non-normality of the Navier–Stokes operator is the spatial separation between the direct and the adjoint global mode in the downstream and upstream direction (Chomaz, 2005), respectively. The direct global mode typically displays exponential growth in the convectively unstable region, and if the flow is stable downstream, its amplitude reaches a constant value in that region. Conversely, from a global point of view, the adjoint shows increased spatial support in the region where a localized perturbation yields a large projection onto the direct global mode. From a local point of view, the adjoint shows the highest value where a localized point-wise perturbation undergoes maximum amplification by means of the instability mechanisms present in the flow.

In our case, however, the direct global mode is concentrated on the suction surface and in the near wake, whereas the adjoint global mode is predominantly located on the pressure surface. Paradoxically, the perturbation that causes the largest growth of a global mode with spatial support on the

¹The reader can easily verify that the solution to the optimization problem $\arg \max_{\mathbf{u}} \frac{|\langle \mathbf{w}_i, \mathbf{u} \rangle|^2}{\|\mathbf{u}\|^2}$ is precisely $\mathbf{u} = \mathbf{w}_i$.




Figure 5.13: Spatial structure of the adjoint global modes labelled M1–M4 in figure 5.8, displaying the magnitude of the stream-wise velocity levels, and the real part in the inset. The modes presented here have been normalized such that $\langle \mathbf{w}, \mathbf{v} \rangle = 1$. The corresponding direct modes are given in figure 5.9.




Figure 5.14: Evolution of the adjoint stream-wise velocity peak along the aerofoil chord (a) on the suction and (b) on the pressure surface of the aerofoil for the modes labelled M1–M4 in figure 5.8.

suction surface is not *convectively linked* to the location on the pressure surface that causes maximum receptivity.

The spatial structure of the adjoint modes suggests a different route for the excitation of the instabilities in the flow via (i) the growth of instabilities by convective mechanisms on the pressure surface, (ii) the scattering of acoustic wave at the trailing edge, (iii) the receptivity at the leading edge, where the suction-surface boundary layer displays convective instabilities. This receptivity mechanism prevails over the direct growth of instabilities on the suction surface as the total spatial growth along the pressure-surface is larger than that on the suction-surface. This observation is supported by previous experiments (Nash et al., 1999) and numerical calculations (Desquesnes et al., 2007): the frequency of the dominant tone in the spectrum coincides with the most amplified frequency along the pressure surface based on local stability theory.

Besides, not all modes seem to be equally receptive. In figure 5.14 we depict the magnitude of the adjoint mode on the suction-surface (figure 5.14a) and on the pressure surface (figure 5.14b). It can easily be seen that the receptivity on the pressure surface to stream-wise velocity perturbations is higher for the modes at the peaks in growth rate, M1 and M4, than for the modes M2 and M3. As discussed before, the frequencies of the modes M1 and M4 are in excellent agreement with the frequencies that appear in the acoustic spectrum of the nonlinear numerical simulation.

The occurrence of multiple peaks in the sound spectrum cannot solely be explained from the convective instability properties of the boundary-layer: the range of frequencies where the boundary layer displays convective growth is larger than the frequency difference between the peaks. In this respect, it is suggested that global effects further discretize the components that are most amplified.




Figure 5.15: Wavemaker region of the eigenvalue labelled M1 in figure 5.8, showing increased sensitivity on the pressure surface.

5.5.3.5 Wavemaker region

The direct global modes and the corresponding adjoint modes determine the structural sensitivity of the linearized operator \mathbf{A} (Hill, 1992; Bottaro et al., 2003; Chomaz, 2005; Giannetti and Luchini, 2007; Marquet et al., 2008). As discussed in the introduction, the sensitivity of the eigenvalues to perturbations in the operator \mathbf{A} can be addressed by considering the perturbed eigenvalue problem $(\mathbf{A} + \delta \mathbf{A})(\mathbf{v} + \delta \mathbf{v}) = (\lambda + \delta \lambda)(\mathbf{v} + \delta \mathbf{v})$ and retaining only first-order terms. By computing the inner product of both sides with respect to the adjoint mode, we obtain $\delta \lambda = \langle \mathbf{w}, \delta \mathbf{A} \mathbf{v} \rangle$ (note that we have chosen \mathbf{w} such that $\langle \mathbf{w}, \mathbf{v} \rangle = 1$). The above expression indicates that the effect of a unitary proportional feedback between flow variables at different locations on any eigenvalue can be quantified by the product between the adjoint at the actuator and the value of the direct global mode at the sensor. Note that in general the feedback affects many eigenvalues of the matrix \mathbf{A} .

For the case of a diagonal $\delta \mathbf{A}$ we restrict ourselves to a local feedback between each individual component at its given location. This approach leads to the concept of a wavemaker: the structural sensitivity of the flow to local feedback quantified by the product of the value of the direct global mode and the corresponding adjoint global mode at a given location.

In figure 5.15 we display the wavemaker of the leading mode M1. We observe that the spatial support of the wavemaker is closer to the adjoint than to the direct mode, since the compactness of the adjoint is more important than that of the direct mode (note that the adjoint mode displays significant acoustic levels in the far-field). Further confirmation is obtained by comparing the value of the wave maker, the direct global mode and the adjoint global mode on the aerofoil surfaces; see figure 5.16.

This fact may indicate that the individual leading eigenfrequency is more sensitive to operator perturbations \mathbf{A} on the pressure surface of the aerofoil (and, in particular, upstream of the separation point) than in the rest of the domain. The first implication of the wavemaker location on the pressure surface has been made by the previous numerical study of Desquesnes et al. (2007):

The computation of the amplification ratio for different chordwise stations and the same x_0 shows that the most amplified frequency is constant and equal to 845Hz from $x/c = 0.7$ to $x/c = 0.96$. The tone noise frequency is therefore selected before the boundary-layer separation. Desquesnes et al. (2007)

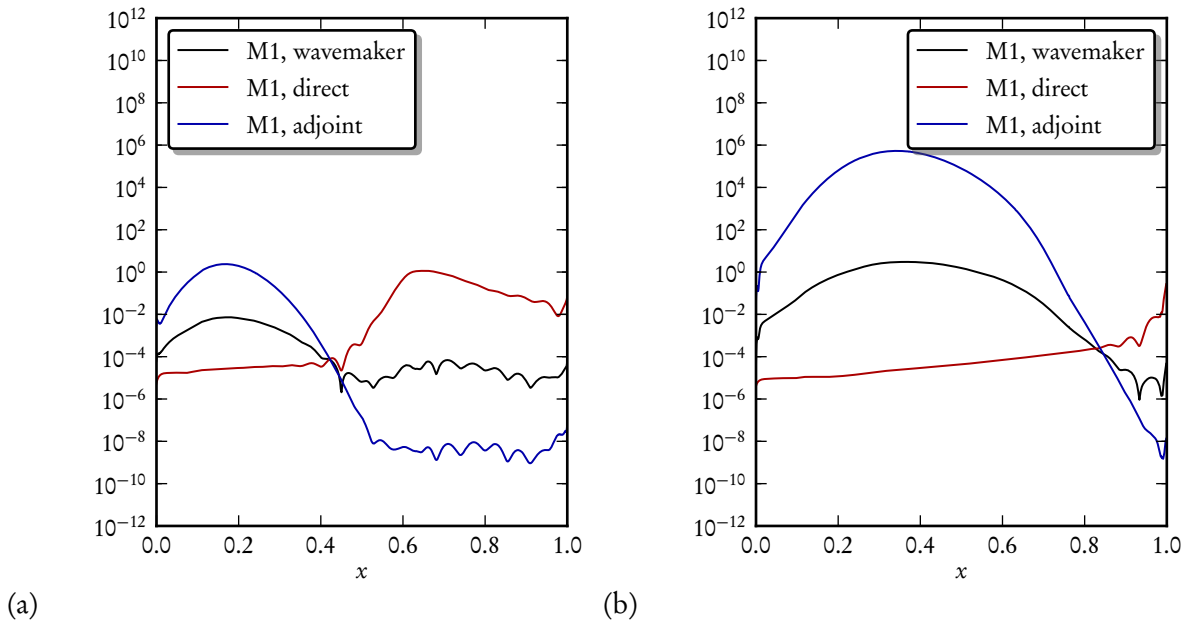


Figure 5.16: Evolution of the amplitude of the direct mode, adjoint mode and the wavemaker along the (a) suction surface and (b) pressure surface for the mode M1 shown in figure 5.8.

In this respect, the wavemaker describes the region in the flow that *imposes* its dynamics on the remaining flow. The results presented here are therefore in good agreement with both numerical and experimental studies:

A boundary-layer trip wire was found to have no effect on the tone when placed at various chord-wise positions on the suction surface of this airfoil but caused the tone to disappear when placed forward of 80% chord on the pressure surface.

Paterson et al. (1973)

The above observations, describing the sensitivity of the flow to changes in the base flow, are in excellent agreement with the spatial location of the wavemaker. However, it is doubtful whether the triggering of the pressure-surface boundary layer can be linked to the small operator perturbations considered above.

The second implication is related to the choice of base flow. Although the instantaneous flow and the mean flow are not steady solutions of the Navier–Stokes equations, the time residual at the pressure-surface boundary layer upstream of the separation bubble is very small. This indicates that the modes of interest exhibit sufficient robustness, provided that the pressure-surface boundary layer remains unaltered.

5.5.4 High-frequency modes: suction-surface shear-layer instabilities

Our attention is now directed toward the features of the remaining modes of the spectrum. We present in figure 5.17 a closer view of the global spectrum for the range of high frequencies (with respect to the leading eigenfrequencies).




Figure 5.17: Global spectrum displaying high-frequencies $\omega > 45$. The labels H1–H4 correspond to the eigenfrequency of the global modes (direct and adjoint) represented figure 5.18. The associated eigenfrequencies are: $\omega_{H1} = 50.04 - 0.29i$, $\omega_{H2} = 66.78 - 0.54i$, $\omega_{H3} = 83.19 - 0.61i$, $\omega_{H4} = 100.02 - 0.62i$. The residual is smaller than $5 \cdot 10^{-3}$.

This part of the global spectrum contains modes with increasingly lower temporal growth rates, representing short time-scale dynamics. Modes in this spectral subdomain thus describe the transient dynamics that were observed in the impulse response (section 5.4) for $t < 1$. Four representative modes, labelled H1–H4 in figure 5.17, are depicted in figure 5.18a–d, visualized by the real part of the free-stream velocity. The high-frequency direct modes have been scaled by the maximum value of the free-stream velocity on the suction surface. As before, the adjoint global modes are normalized to satisfy $\langle \mathbf{w}, \mathbf{v} \rangle = 1$.

The direct mode H1, depicted in figure 5.18a in terms of the stream-wise velocity, displays features similar to the coupled dynamics between the suction-surface and pressure-surface separation bubbles that was observed for the leading modes in the frequency range $30 < \omega < 50$. Although acoustic radiation is observed, it is comparatively smaller than that of the leading modes. Likewise, the corresponding adjoint stream-wise velocity field (see figure 5.18b–c) displays a dominant receptivity to perturbations on the pressure-surface. Owing to the higher-frequency, the spatial scales are noticeably shorter than those of the leading modes M1–M4.

For even higher frequencies (figure 5.18d) we observe a qualitative change from the mode H1 to the mode H2: the direct global mode reduces the spatial support on the pressure-surface separation bubble; on the suction surface, the front of the global mode moves upstream into the shear-layer region. In the reattached flow region, the mode describes the advection by the free stream, lacking significant interaction with the boundary layer. Furthermore, the acoustic radiation disappears for higher-frequency modes (H2–H4).

The transition from H1 to H2 can be explained by returning to the characteristic scales of the different separated shear-layers. In order to comply with the prevalent dispersion relations, the spatial support of the mode shows shorter wavelengths for increasing frequencies. Since the thickness of the separated shear layer is smaller on the suction side than on the pressure side, for short wavelengths (or higher frequencies) the suction-surface separation bubble supports instabilities whereas the growth on the pressure surface is less pronounced. The adjoint mode for H2 (figure 5.18e–f) has not changed qualitatively when compared to that of H1. In other words, rather than directly placing a perturbation on the suction surface, it is still more energy-efficient to excite the direct mode H2 by taking advantage of the spatial instability growth on the pressure surface combined with a scattering of acoustic waves

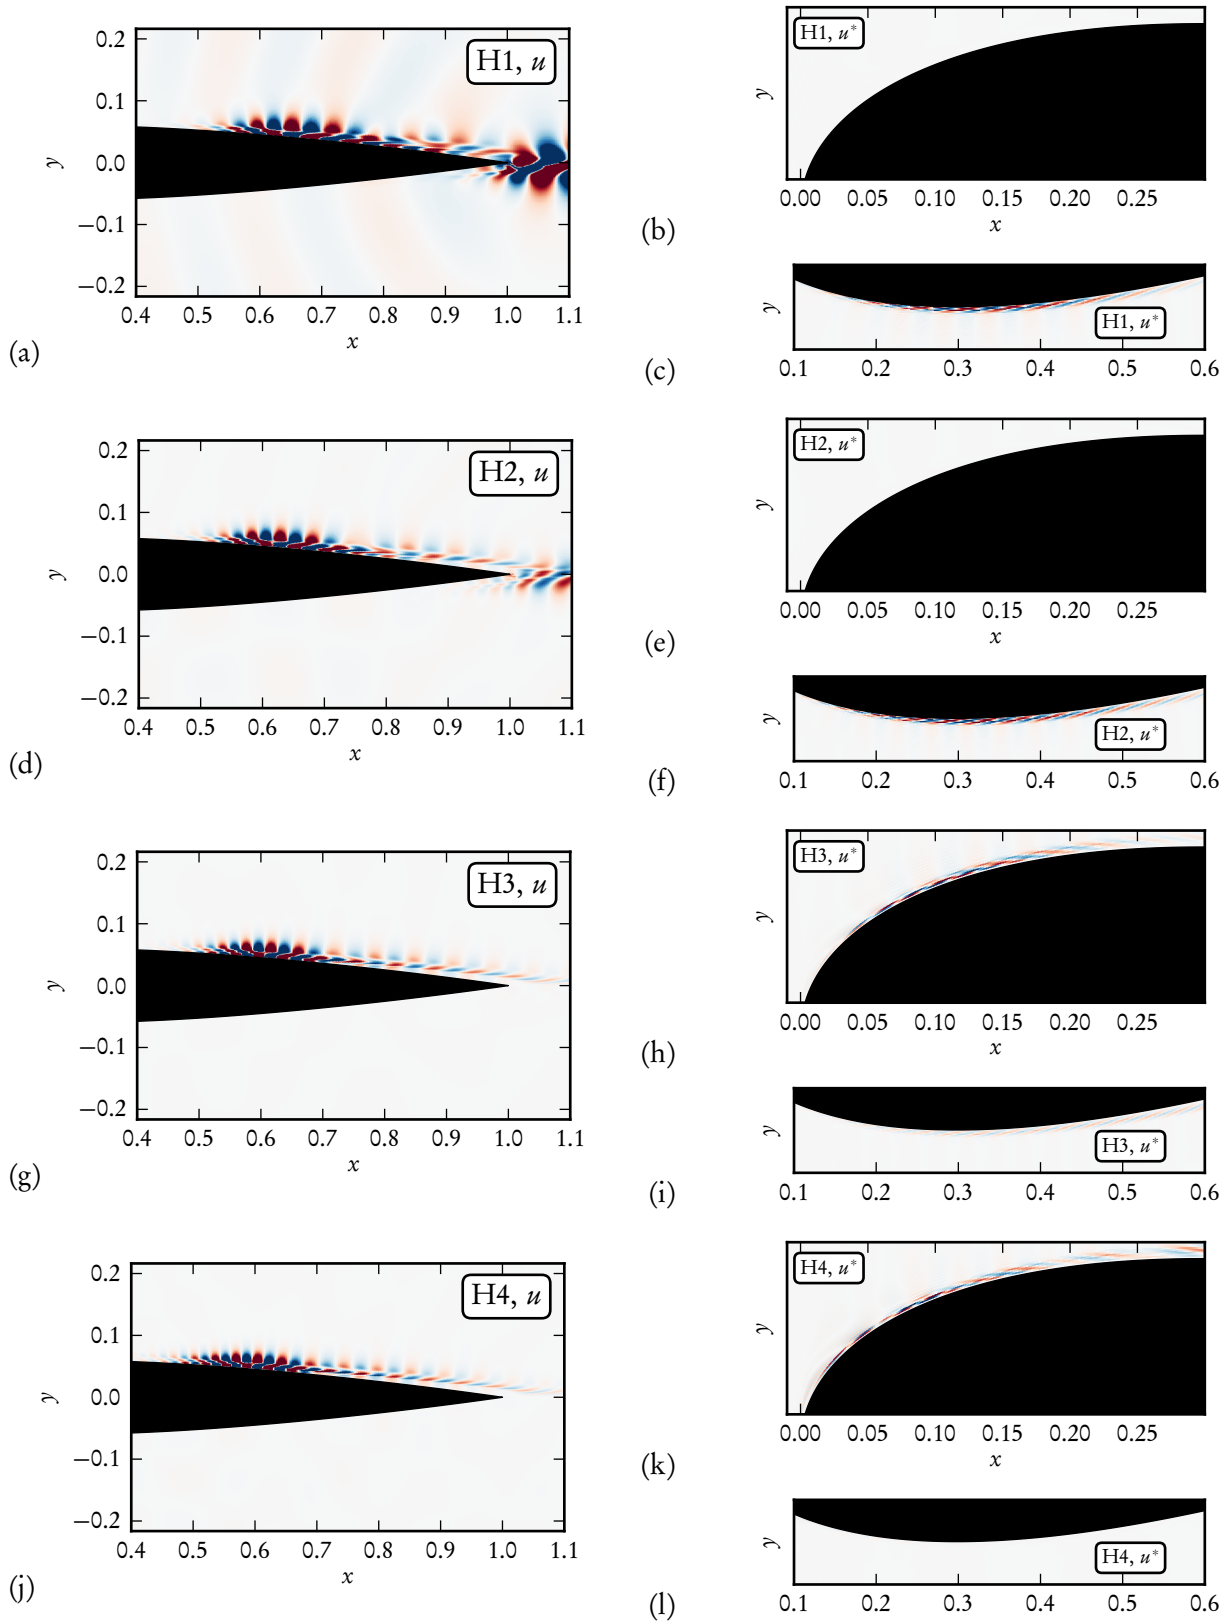


Figure 5.18: Direct and adjoint global modes corresponding to H1-H4 indicated in figure 5.17, visualized by the stream-wise velocity (left column) and adjoint stream-wise velocity (right column) at the leading edge and on the pressure surface. From top to bottom row, modes with increasing frequency are displayed.




Figure 5.19: Low-frequency modes in the spectrum. The modes labelled L1–L4 are depicted in figure 5.20, and the corresponding eigenvalues are $\omega_{L1} = 29.96 - 0.36i$, $\omega_{L2} = 20.08 - 0.38i$, $\omega_{L3} = 10.03 - 0.40i$, $\omega_{L4} = 0.03 - 0.04i$.

and a re-triggering on the leading edge via a receptivity mechanism.

We observe several qualitative changes in the adjoint global modes for H3 and H4. For the adjoint global mode H3 (see figure 5.18h–i), the receptivity on both the suction and the pressure surface is comparable. The growth of instabilities on the pressure surface is then equally important for the excitation of the direct modes. After a further increase in frequency (mode H4) the direct global mode moves further upstream (see 5.18j), and the adjoint global mode ceases to display spatial support on the pressure-surface; rather, it becomes highly localized on the suction surface near the leading edge.

5.5.5 Low-frequency: separation-bubbles flapping and reattached flow dynamics

The global spectrum for $\omega < 35$ is displayed in figure 5.19. As before, we have chosen four representative modes L1–L4 in order to investigate the features of the global spectrum for decreasing frequency. The normalization criterion for the direct and adjoint global modes is identical to the one described for the high-frequency modes.

The global mode L1 shown in figure 5.20a corresponds to the lower bound in frequency (and wavenumber) of the modes that represent the coupled dynamics of the separation bubble on both surfaces. As before, it displays increased receptivity on the pressure-surface (figure 5.20b–c).

Contrary to the previous case, for decreasing frequencies we expect an increase in the wavelength of the observed instabilities; the corresponding global modes should display structures in the shear layers that exceed that scale of the separation bubbles.

Modes L2 and L3 confirm that for lower frequencies (figure 5.20d,g) the global modes undergo a qualitative change: the direct global modes localize downstream of the separation bubble, in the region where the flow is reattached. Moreover, boundary-layer-type instabilities on the suction surface are then excited by the pressure-surface dynamics for the mode L2 (figure 5.20e–f), while for the mode L3 the receptivity on the suction surface becomes important.

For the lowest-frequency mode L4 the wavelength is comparable with the extent of the separation bubbles; this mode represents synchronized low-frequency oscillations, in opposite phase, between both separation bubbles (figure 5.20j). The corresponding adjoint (figure 5.20k) captures the receptiv-




Figure 5.20: Direct and adjoint global modes corresponding to L1-L4 indicated in figure 5.19, showing stream-wise velocity (left column) and adjoint stream-wise velocity (right column) at the leading edge and on the pressure surface. From top to bottom row, modes with decreasing frequencies are displayed.

ity to long-wave disturbances in the free-stream. It is important to recall that this type of oscillation was already observed, both in the nonlinear simulations and in the impulse response analysis.

5.5.6 Convection effects and free-stream dynamics

5.5.6.1 Effect of the convection on the computed global modes

In the description of the modal dynamics, we have so far focused on the dynamics in the vicinity of the aerofoil surface and the near wake. However, there also exist global modes that capture the advection of perturbations in the wake as well as the propagation of acoustic waves in the far-field.

In figure 5.21a–c we present the real part of the stream-wise velocity for representative modes from each of the families of global modes that have been identified before. The global modes have been chosen such that the temporal growth rate is approximately equal to $\omega_i \approx -0.43$. In figure 5.21d the evolution of the stream-wise velocity along the wake is represented. A common characteristic for the selected modes is the exponential growth in the wake. Contrary to intuition, this feature is associated with the advection dynamics, since the base flow does not exhibit a significant velocity deficit in the wake.

The global modes describe the dynamics of instability waves in the domain. These waves reach the trailing edge, after which they are subjected to advection by the base flow in the wake; this advective part is characterized by a dispersion relation of the form $\omega = \alpha$. As a consequence, the spatial structure of the global mode in the wake necessarily displays spatial, exponential growth. In figure 5.22a we display the magnitude of the stream-wise velocity in the wake using a logarithmic scale. Since $\omega_i = -0.43$, the spatial growth rate in the wake can be approximated as $-\alpha_i \approx 0.43$, as confirmed in the figure. Following the explanation given in Garnaud (2012) in the context of global modes in compressible jets, the growth in the free stream is physical. It represents features of temporally stable perturbations that were emitted earlier; due to pure advection, they have thus higher amplitudes, compared to perturbations that have been emitted at the current time.

A similar phenomenon is observed in figure 5.22b for the evolution of the pressure amplitude in the far-field along a vertical line through the trailing edge for the modes M1–M4.

5.5.6.2 Free-stream modes

In the previous sections we have presented the direct and adjoint global modes that describe instability mechanisms of the flow in the vicinity of the aerofoil. The description of the operator \mathbf{A} in terms of these global modes is, however, incomplete since a representation of the free-stream dynamics is missing.

In figure 5.23 we display two examples of global modes, labelled F1 and F2, that exhibit free-stream dynamics. The eigenvalues associated with this family of modes have negative temporal growth rates. They are furthermore characterized by a prominent spatial support in the free-stream. A noteworthy property is the fact that they show spatial exponential growth towards the outlet and well-defined real wavenumbers in the two spatial directions.

The features of the free-stream modes are reminiscent of the pseudospectral characteristics of both the advection-diffusion equation and the wave equation with nearly absorbing boundary conditions. The reader is referred to Trefethen (1997) for further details.

From a qualitative point of view, the free-stream global modes can be understood as the result of both the finite size of the domain and the dispersion relation of the one-dimensional advection equation $\omega = \alpha$.

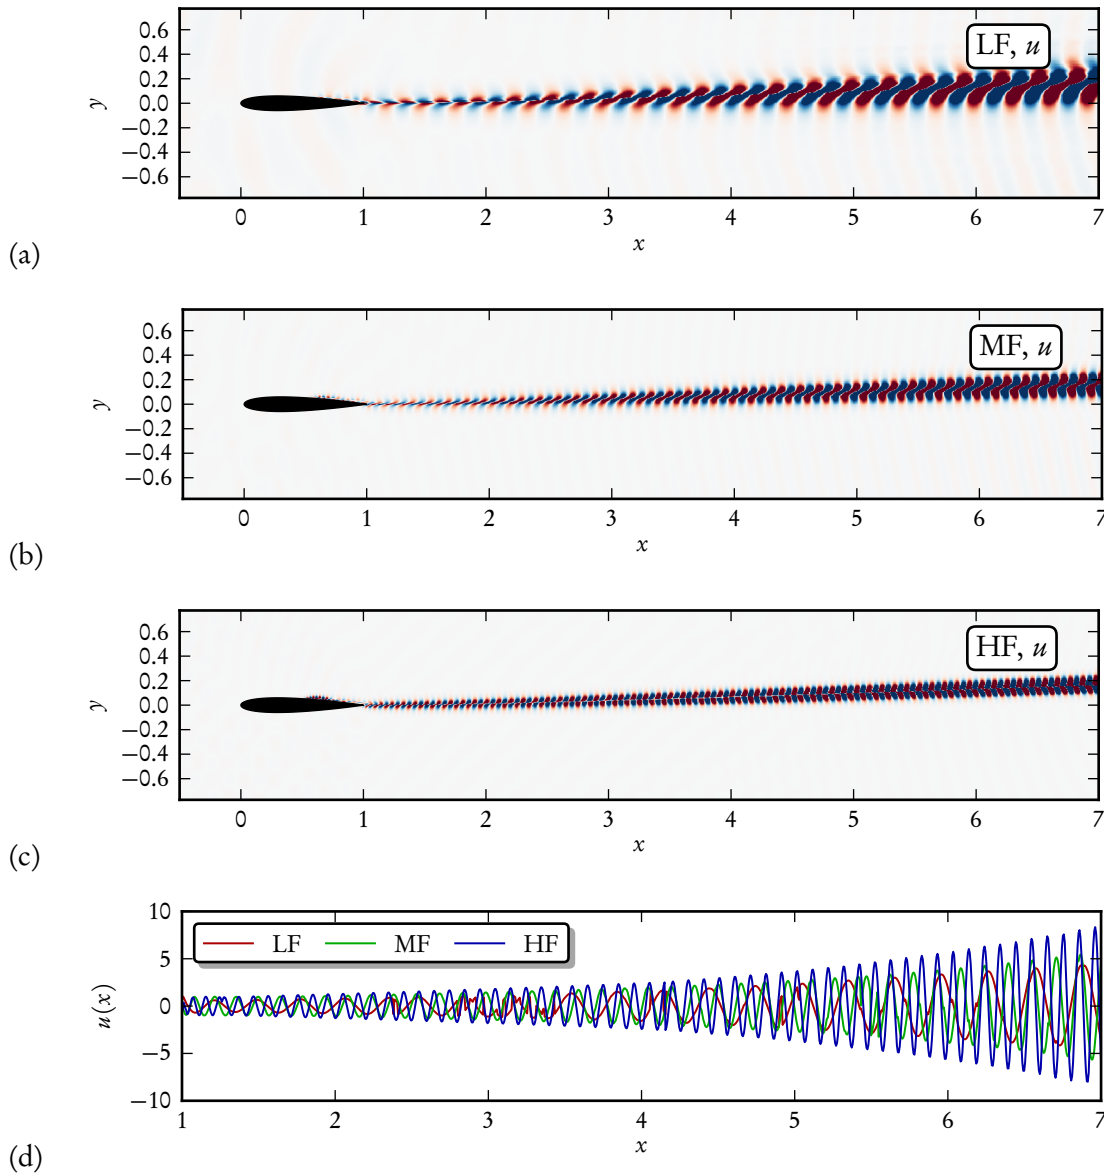


Figure 5.21: Downstream advection effects for the computed modes showing examples of (a) low-frequency modes $\omega_{\text{LF}} = 20.08 - 0.38i$, (b) leading modes $\omega_{\text{MF}} = 39.87 - 0.43i$ and (c) high-frequency modes $\omega_{\text{HF}} = 59.97 - 0.43i$. The real part of the stream-wise velocity is represented in (a-c); the evolution along the wake is shown in (d).




Figure 5.22: (a) Exponential growth in the free-stream to comply with the dispersion relation of the advection equation $\omega = \alpha$. (b) Amplitude of pressure fluctuations along a vertical line through the trailing edge.




Figure 5.23: Free-stream modes: (a) $\omega_{F1} = 0.94 - 0.21j$ and (b) F2 $\omega_{F2} = 5.02 - 0.38j$.

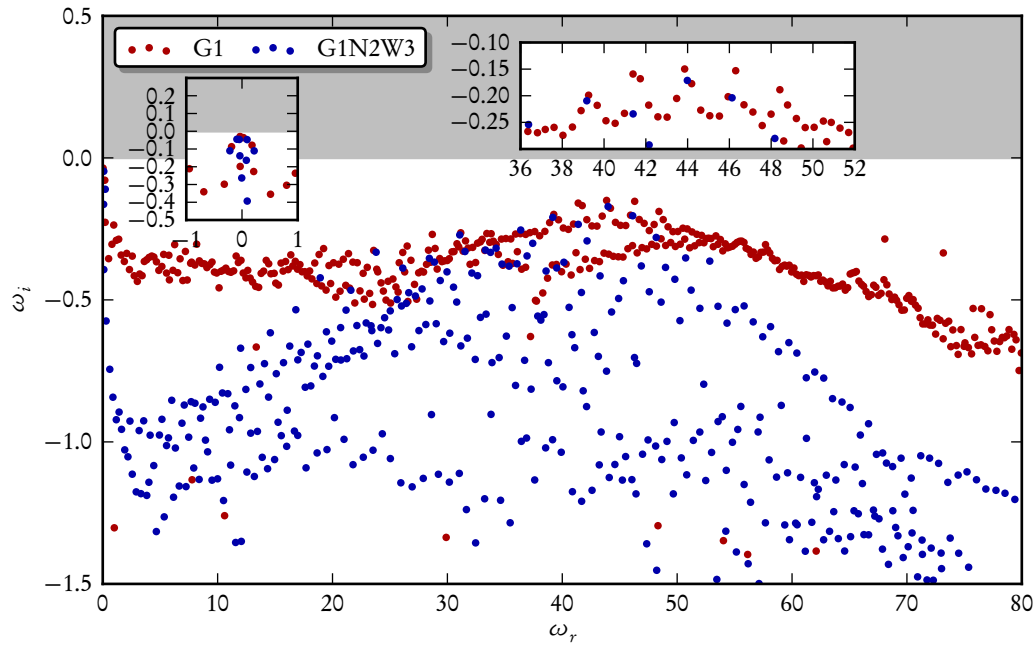


Figure 5.24: Effect of the domain size on the global spectrum.

Since the domain has a finite extent, any perturbation in the free-stream leaves the domain after a convective time period. As a consequence, the global modes that describe the free-stream dynamics are damped. In light of this fact, it is important to remark that the growth rates of the free-stream modes are inherently dependent on the size of the domain, and in the limit of an infinitely large domain they converge towards a zero temporal growth rate.

As before, for selected α (such that the diffusive effects are negligible), the physics in the free-stream can be approximated by pure advection, and therefore the free-stream eigenvalues with negative temporal growth rates ω have associated global modes with wavenumber $\alpha \approx \omega$: since $\omega_i < 0$, the free-stream global modes display structures with an associated exponential spatial growth of $-\alpha_i > 0$.

In the global spectrum there is no clear separation between free-stream modes and the global modes presented before. Nonetheless, our calculations show that they are predominant in the region of low temporal growth rates, where the computed eigenvalues have a rather large residual.

5.5.7 Robustness of the global mode representation: effect of the domain size

In figure 5.24 we display the effect of the domain size on the computed global spectrum for the cases G1 and G1N2W3 defined in table 5.1.

The temporal growth-rate and the frequency of the leading global modes are approximately the same. Even though some global modes can be observed at the peaks, the characteristic structure described before, with alternating local maxima and minima, is no longer evident. Note also that since the global modes have spatial support in the wake from the trailing edge to the domain outlet, the spacing of the modes is approximately given by $\Delta\omega \approx \pi/L$, and the convergence of the spectrum

cannot be evaluated as that of each individual eigenpair. In the case of domains G1 and G1N2W3, the frequency spacing between consecutive modes is $\Delta\omega \approx 0.43$ and $\Delta\omega \approx 0.95$, respectively.

Note that although the structure of the peaks can no longer be identified for small domains, the features observed in the impulse response (performed using the case G1N2W3) remained unchanged, which indicates that the physical features described earlier are intrinsic to the system. While these physical features are robust, their representation in terms of global modes, in the case of a stable system, is very sensitive to small modifications.

Nonetheless, for the present case the global modes have been shown to be useful in describing the dynamics of our system. A possible explanation for their usefulness in the present case lies in the fact that the leading modes describe a process with a feedback loop — possibly, a structure that can be efficiently described by global modes. This idea is supported by the fact that the global modes for higher frequencies (modes H·) and lower frequencies (modes L·), which describe pure, convectively unstable dynamics, are significantly affected by size of the domain.

5.5.8 Conclusions from the global mode analysis

In the global modes analysis presented before, it has been found that the frequency range of the least stable global modes correspond to that of the acoustic tones in the nonlinear simulation (chapter 3). Furthermore, a detailed inspection of the global spectrum reveals that a set of modes equally-spaced in frequency have associated comparatively higher temporal-growth rates.

The corresponding direct modes have been related to the individual tones that are present in the nonlinear simulations, and they describe, by superposition, the feedback loop observed in the impulse response analysis (section 5.4). The spatial support of the adjoint modes is concentrated on the pressure-surface, that has been related to strong non-normal effects involving regions in the flow that are not directly linked by the advection. The analysis of the wavemaker points at the boundary layer on the pressure surface as the region of increased sensitivity of the flow, thus providing a theoretical justification for the choice of the mean flow as the basic state.

Finally, different families of modes have been identified in the global spectrum (figure 5.25) and classified (figures 5.26–5.27) according to the physical mechanisms that they describe: low-frequency flapping of the separation bubbles, dynamics of the reattached-flow on the suction surface and high-frequency shear-layer instabilities on the separated-flow regions.




Figure 5.25: Global spectrum: summary of families of modes.

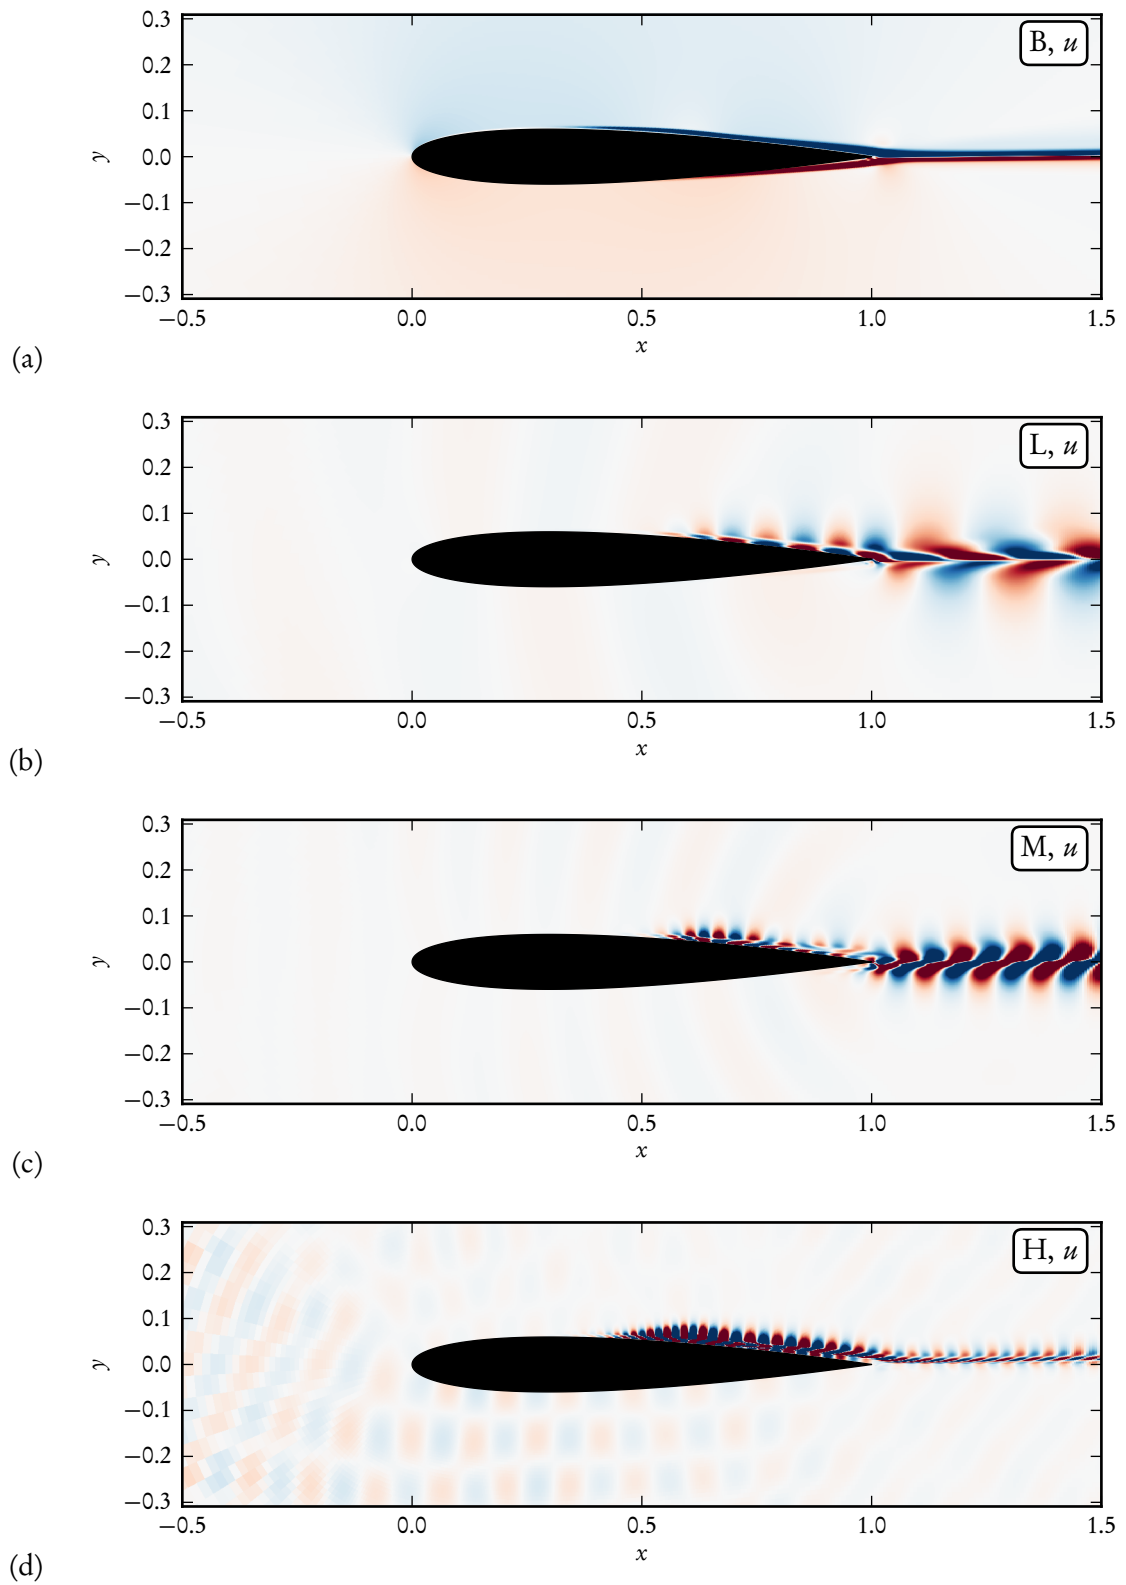


Figure 5.26: Families of direct global modes: (a) low-frequency oscillation of the separation bubbles, (b) spatial amplification along the suction-surface reattached boundary-layer (c) noise radiation and coupled dynamics of the separation bubbles (d) high-frequency shear-layer instabilities of the separation bubbles. The corresponding adjoint modes are displayed in figure 5.27.




Figure 5.27: Families of adjoint global modes: (a) low-frequency oscillation of the separation bubbles, (b) spatial amplification along the suction-surface reattached boundary-layer (c) noise radiation and coupled dynamics of the separation bubbles (d) high-frequency shear-layer instabilities of the separation bubbles. The corresponding direct modes are displayed in figure 5.26.

5.6 Frequency response

Our analysis found that the operator \mathbf{A} is stable. This observation has a major implication on the physical mechanisms that can be observed in experiments: it indicates that our fluid system behaves as a noise amplifier rather than as an oscillator.

So far, we have focused on the description of the flow features as they pertained to the *unforced* dynamics initiated by an external initial impulse. However, in a real environment, the perturbations are most aptly described in terms of a statistically-steady random forcing, and therefore the fluid system displays *forced* dynamics. In a linearized framework we can decompose the external forcing into a superposition of constant harmonic forcings $\mathbf{f}e^{-i\omega t}$, with ω real, and characterize the system response for different values of ω . To this end, we seek the forcing that maximizes, as described in 5.1.2, the response of the system for each frequency.

5.6.1 Singular value problem as an eigenvalue problem: numerical set-up

Prior to presenting the optimal forcings, responses and gains for our system, we briefly describe the numerical set-up. The calculation of singular values is easily expressed as an eigenvalue problem in terms of the resolvent $\mathcal{R}(\omega; \mathbf{A})$ defined in equation (5.19). One of the strategies implemented in the eigenvalue solver SLEPc consists of rewriting the singular value problem as an eigenvalue problem given in terms of the cyclic matrix $H(\mathcal{R})$:

$$H(\mathcal{R}) = \begin{pmatrix} 0 & \mathcal{R} \\ \mathcal{R}^H & 0 \end{pmatrix}. \quad (5.57)$$

The eigenvalues and eigenvectors of $H(\mathcal{R})$ are related to the singular values and singular vectors of \mathcal{R} . It is easily verified that the eigenvectors are

$$\begin{pmatrix} \pm \mathbf{u}_i \\ \mathbf{v}_i \end{pmatrix} \quad (5.58)$$

and that the associated eigenvalue is σ . Therefore, the singular triplet of interest is computed as the largest-in-magnitude eigenvalue of the cyclic matrix. Moreover, since the cyclic matrix is Hermitian, the convergence properties are, in this case, more favourable than for non-Hermitian problems, and thus only a few iterations are required.

We only need to provide the eigensolver with a routine that evaluates the matrix-vector products $\mathcal{R}\mathbf{v}$ and $\mathcal{R}^H\mathbf{u}$, i.e.

$$\mathcal{R}\mathbf{v} = \mathbf{L}_v (\mathbf{A} + i\omega\mathbf{I})^{-1} \mathbf{L}_f^{-1} \mathbf{v} \quad (5.59)$$

and

$$\mathcal{R}^H\mathbf{u} = \mathbf{L}_f^{-H} (\mathbf{A}^H - i\omega\mathbf{I})^{-1} \mathbf{L}_v^H \mathbf{u} \quad (5.60)$$

However, in the matrix-free framework presented in section 5.2, we have no direct means of evaluating the resolvent \mathcal{R} , since it involves the solution of a linear system of equations, which is not straightforward within the present matrix-free framework. A compromise to overcome this limitation consists of a long-time integration of equation 5.4 in the intention that for large times T




Figure 5.28: Evolution of the temporal residual for a random perturbation.

the solution is sufficiently close to the asymptotic regime. Then, the evaluation of matrix-vector products $\mathcal{R}\mathbf{v}$ and $\mathcal{R}^H\mathbf{u}$ read

$$\mathcal{R}\mathbf{v} \approx \mathbf{L}_v \left[(\mathbf{A} + i\omega\mathbf{I})^{-1} \left(e^{(\mathbf{A} + i\omega\mathbf{I})T} - \mathbf{I} \right) \right] \mathbf{L}_f^{-1}\mathbf{v}, \quad (5.61)$$

and

$$\mathcal{R}^H\mathbf{u} \approx \mathbf{L}_f^{-H} \left[(\mathbf{A}^H - i\omega\mathbf{I})^{-1} \left(e^{(\mathbf{A}^H - i\omega\mathbf{I})T} - \mathbf{I} \right) \right] \mathbf{L}_v^H\mathbf{u}. \quad (5.62)$$

with the advantage that it is now possible to use the exponential Krylov time integration. Note that in the above expressions the initial condition has been taken as zero. The obvious drawback of the above approach is the potential cost of the time integration, especially if reaching the steady regime requires long integration times. The time horizon T has been determined in our case from temporal simulations using a random forcing; see figure 5.28.

In view of the above results, it appears reasonable to conduct the time integration for $T = 15$ time units. Even though a longer integration time would seem desirable, we were limited by CPU time of the calculations to this value. Fortunately, the size of the Krylov spaces to determine the largest singular value is far smaller than in the case of global-mode calculations. For instance, a Krylov subspace size of $m = 4$ yielded leading singular values converged up to 10^{-5} for all cases presented here. Note that the second of the above equations corresponds to the temporal integration of the transconjugate of the operator. In the development of the technique for the efficient evaluation of the linearized dynamics, special attention has been paid to the fundamental relation of the adjoint up to the machine precision; this property will ensure the convergence of the iterative numerical method. It is important to verify the time integration techniques also conforms to the direct-adjoint relationship. As discussed in Garnaud (2012), for Runge-Kutta schemes this amounts to implementing the adjoint of the time integration scheme. For our case of Krylov-based exponential time integration, however, no such implementation is required, since the direct-adjoint property is automatically satisfied.




Figure 5.29: Frequency response showing gain in amplitude (also the leading singular value of the resolvent) versus frequency. Note that the continuous curve has been obtained by a piecewise interpolant constructed with sensitivity information.

5.6.2 Optimal gain

The global spectrum in the vicinity of the least stable eigenvalues, i.e. $34 < \omega < 54$, thus suggests that tonal noise is related to equally spaced eigenfrequencies at local maxima of the growth rate. This finding is further corroborated by a comparison of the acoustic frequencies from the nonlinear simulations. In our investigation, the linearized dynamics have been found stable, and the global modes are highly non-normal which indicates that caution must be exercised in the characterization of flow amplification mechanisms.

For this reason, we have computed the optimal responses for a range of frequencies of interest. The gain as a function of the frequency is depicted in figure 5.29, and a summary of the most relevant characteristics of the optimal frequency-response is provided in table 5.3.

In spite of a rather low number of frequencies to obtain a complete picture of the transfer function, it can be seen that the transfer function shows several maxima at locations that coincide with global modes of largest growth rates. This confirms the observation made during the global mode analysis that the multiple peaks in the spectrum have their origin in the linearized dynamics of the flow.

5.6.3 Optimal forcing and optimal response

We consider now the characteristics of the responses and the forcings for the frequencies at $\omega_1 = 44$, $\omega_2 = 45$ and $\omega_3 = 46.5$; see figure 5.30. The frequencies ω_1 and ω_3 are closer to a local maximum of the transfer function, whereas ω_2 is close to a local minimum.

We observe familiar features in the near-field of the aerofoil: boundary layer instabilities appear on the suction surface of the aerofoil and in the near wake. In our case, the forcing represents a fixed spatial structure harmonic in time, and therefore the response displays the complete behaviour of the advection in the wake and in the far-field. This enables a quantitative comparison between the optimal responses and the frequential snapshots from the nonlinear simulations.




Figure 5.30: Spatial structure of the optimal responses for selected frequencies of the transfer function displayed in figure 5.29; the real part of the velocity levels in the vicinity of the aerofoil surface (left column), and the real part the associated near-field pressure levels (right column) are visualized. The corresponding forcings are displayed in figure 5.31




Figure 5.31: Spatial structure of the optimal forcings for selected frequencies of the transfer function displayed in figure 5.29, showing the real part of the velocity levels in the vicinity of the aerofoil surface. The corresponding responses are displayed in figure 5.30.

In figure 5.31 we present the optimal forcings associated with the responses depicted in figure 5.30. Similarly to the adjoint global modes, the optimal forcings display compact support on the aerofoil surface and exponential decay towards the free-stream. The sensitivity of the optimal gain with respect to modifications in the linearized operator A is displayed in figure 5.32.




Figure 5.32: Sensitivity of the optimal gain for $\omega = 40$ with respect to modifications of the linearized operator \mathbf{A} .

Transfer function			Optimal forcing (maxima)					Optimal response (maxima)						
			Pressure surface		Suction surface			Suction surface			Wake		Near-field	
ω	$\sigma_1 \times 10^{-7}$	$\varepsilon_{\text{rel}} \times 10^4$	x	y	x	y	$ u^* \times 10^3$	x	y	$ u^* \times 10^{-5}$	x	y	$ u $	$p \times 10^{-5}$
34.00	0.42	1.96	0.37	-0.061	0.16	0.054	16.04	1.00	-0.002	1.45	1.04	-0.010	2.65	1.00
39.27	1.61	0.77	0.36	-0.061	0.14	0.053	5.92	0.64	0.043	5.86	1.04	-0.010	9.22	3.14
40.00	0.92	0.64	0.36	-0.061	0.14	0.053	3.97	0.63	0.043	3.42	1.05	-0.009	5.61	1.13
42.10	1.76	0.43	0.35	-0.061	0.14	0.053	2.83	0.64	0.043	6.61	1.13	-0.008	9.72	2.47
43.50	1.61	0.27	0.35	-0.061	0.14	0.053	2.98	0.63	0.044	6.46	1.05	-0.009	9.08	3.52
43.90	3.32	0.03	0.35	-0.061	0.14	0.053	2.82	0.63	0.044	15.15	1.14	-0.008	17.63	7.11
44.00	3.48	0.03	0.35	-0.061	0.14	0.053	2.65	0.63	0.044	15.67	1.14	-0.008	18.57	7.18
44.10	3.28	0.04	0.35	-0.061	0.14	0.053	2.44	0.63	0.044	14.55	1.14	-0.008	17.30	6.41
44.50	1.54	0.33	0.35	-0.061	0.14	0.053	1.31	0.63	0.044	5.22	1.14	-0.008	8.31	1.99
45.00	1.53	0.32	0.35	-0.061	0.14	0.052	1.40	0.63	0.044	5.65	1.14	-0.008	8.26	1.46
46.50	2.46	0.11	0.34	-0.061	0.13	0.052	1.19	0.63	0.044	9.90	1.14	-0.008	12.19	2.42
50.00	1.84	0.35	0.34	-0.061	0.13	0.051	0.36	0.62	0.044	6.43	1.14	-0.008	8.63	1.35
54.00	1.15	0.76	0.33	-0.061	0.13	0.051	0.19	0.62	0.044	3.29	1.14	-0.008	5.03	1.48

Table 5.3: Summary of the performed frequency responses

Summary of the performed frequency responses. The maximum gain has been obtained for $\omega = 44.0$, where the gain in amplitude is $\sigma_1 = 3.48 \cdot 10^7$. For the cases presented here, the maximum value of the stream-wise velocity forcing is placed on the pressure surface at $x = 0.35$, the value of the corresponding optimal forcing on the suction surface is $2.65 \cdot 10^{-3}$. The response is maximum in the near wake at approximately $x = 1.15$. The magnitude of the pressure at $x = 1$ and $y = 0.5$ is maximum for the frequency $\omega = 44.0$.

Conclusions

6.1 Summary of the thesis

In this work, the generation of discrete acoustic tones in flows around aerofoils has been addressed by numerical investigations featuring direct numerical simulations and global stability analyses. The primary purpose was to gain insight into the physical mechanisms underlying the occurrence of multiple peaks in the acoustic spectrum and isolating the principal features of aeroacoustic feedback loop.

To this end, we have designed and implemented a nonlinear simulation code, presented in chapter 2, that addresses two- and three-dimensional compressible flows in physical domains that can be described using multi-block grids. In the choice of numerical methods, special emphasis has been directed to the quality of the discretizations and the boundary conditions in order to analyse sound radiation mechanisms. The numerical code is parallelized to address flow problems that require substantial computational resources, and it has been extensively tested throughout this work.

Our numerical code has been used first in the simulation of the flow around an aerofoil, described in chapter 3. Despite the choice of a relatively high Mach number, the flow features are in qualitative agreement with previous experimental and numerical investigations. It has been shown that the pressure spectrum displays strong, equally spaced peaks, and their sensitivity to computational parameters has been assessed. The investigation of the hydrodynamic fields confirms that two separation bubbles are present: on the suction and pressure surfaces of the aerofoil. Based on their dynamics, a quasi-steady state establishes itself, consisting of vortex shedding and intense acoustic scattering at the trailing edge. The analysis of the frequential snapshots confirms that the radiation of noise is linked to the ringing of hydrodynamic coherent structures. Furthermore, the fluctuation levels observed near the trailing edge have been related to the well-known low-frequency flapping induced by the long laminar separation bubbles.

In order to tackle the above problem using global stability theory, a technique for the efficient extraction of the linearized direct and adjoint dynamics from compressible flow solvers has been developed (chapter 4). This technique is based in the observation that, owing to the underlying structure of the equations, the evaluation of nonlinear terms is only weakly coupled to the evaluation of spatial derivatives which themselves are linear. This fact allows us to evaluate the direct linearized dynamics in a modular fashion with a significantly reduced programming effort. The evaluation of the adjoint dynamics, a desirable component for quantitative flow analyses, is performed by evaluating

the direct code in reverse mode. The technique has been applied to spatially developing boundary layers and aerofoil flows, and the potential application of this technique to other flow solvers has been briefly discussed.

Once the above technique has been applied to the nonlinear numerical code, we present in chapter 5 a global stability analysis of the mean-flow dynamics of the aerofoil flow described in chapter 3. The resulting operator has been found globally stable. However, the impulse response of the flow to a localized initial perturbation at the leading edge displayed important transient growth effects that culminated in the establishment of aeroacoustic feedback loops, similar to ones reported in recent numerical studies. A novel result of this flow case is the observation of the competing dynamics of both suction- and pressure-surface separation bubbles, and their cross-interaction by means of acoustic radiation. Furthermore, the described feedback loop has been identified in the direct numerical simulations at the initial stages.

In order to gain insight into the properties of the linearized operator, we have conducted a global mode analysis. We have shown that the global spectrum is characterized by least stable global modes within the range of frequencies of acoustic tones. A more detailed analysis revealed that the temporal growth-rates of the leading modes follow a peak structure where the frequencies are in remarkably close agreement with the discrete frequencies observed in the nonlinear simulations. The spatial structure of the direct global modes captures the growth of hydrodynamic instabilities on the suction surface and the near wake. According to the frequency spacing and the spatial structure, these global modes are associated with the feedback loop observed in both the impulse response and the nonlinear simulations. Additional families of modes have been identified representing the high-frequency dynamics of the shear layer of the separation bubbles, the low-frequency dynamics of the reattached flow on the suction surface, the free-stream dynamics and the low-frequency flapping observed in the frequential snapshots.

The adjoint global modes display increased receptivity of the flow on the pressure surface upstream of the separation bubble, and the receptivity of the modes is more important for the modes at the acoustic frequencies. In particular, the spatial support of the adjoint modes is of particular importance for this flow, and it displays strong non-normal effects involving regions in the flow that are not directly linked by advection effects. Instead, the increased receptivity at this location can be linked to a combined effect of highly amplified instabilities, the acoustic scattering at the trailing edge and the leading-edge receptivity. The analysis of the wavemaker highlights, in agreement with previous investigations, the sensitivity of the flow to the pressure-surface boundary layer and the associated instability processes. In addition, it provides a theoretical justification for the choice of the mean flow as the basic state. For the remaining families of global modes, the importance of suction- and pressure-surface receptivity has been analysed.

Nevertheless, owing to the non-normality of the underlying equations, the global modes are known to form an ill-conditioned basis for quantitative flow investigations. For this reason, we have conducted a frequency response analysis, which constitutes a more robust representation of the dynamics of convection-dominated flows. Typical gains in amplitude have been found to be in the range of 10^8 – 10^9 . The above results suggest that although the base flow is stable, the frequency response displays a powerful and rather selective amplification mechanism at the frequencies of the spectral peaks. However, our study has been limited by computational resources that are required to determine the complete transfer function.

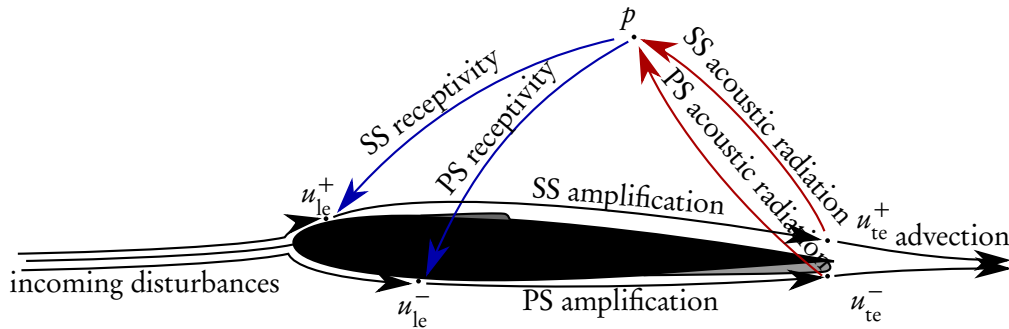


Figure 6.1: Mechanism.




Figure 6.2: Transfer-function model.

6.2 Future perspectives

Global stability analyses represent a promising tool for a detailed analysis of the tonal noise-generation mechanism on aerofoils and, in a broader sense, on complex aeroacoustic instabilities.

In the present case, it is reasonable to expect a global stability analysis to shed light on the quantitative importance of the flow features presented here. Based on our findings, we propose two simplified models that should form the basis of further investigations; see figure 6.1.

The first simplified model is inspired by a transfer-function model presented by Rowley et al. (2006) for cavity oscillations, and it consists of modelling the amplification processes that take place in the separation bubbles by delayed second-order linear oscillators, coupled by acoustic interactions which are in turn represented by a proportional, delayed feedback; see figure 6.2.

The second model consists of a set of partial differential equation involving Ginzburg–Landau equations with an upstream feedback from downstream flow conditions.

$$\frac{\partial \psi^+}{\partial t} + U^+(x) \frac{\partial \psi^+}{\partial x} = \mu^+(x) \psi^+ + \gamma \frac{\partial^2 \psi^-}{\partial x^2} \quad (6.1)$$

$$\frac{\partial \psi^-}{\partial t} + U^-(x) \frac{\partial \psi^-}{\partial x} = \mu^-(x) \psi^- + \gamma \frac{\partial^2 \psi^+}{\partial x^2} \quad (6.2)$$

with

$$\psi^+(0, t) = K^+(\psi^-(L, t - \tau) + \psi^+(L, t - \tau)) \quad (6.3)$$

$$\psi^-(0, t) = K^-(\psi^-(L, t - \tau) + \psi^+(L, t - \tau)) \quad (6.4)$$

$$\frac{\partial \psi^+}{\partial x} \Big|_L = \frac{\partial \psi^+}{\partial x} \Big|_L = 0 \quad (6.5)$$

From a global and local stability analysis, realistic parameters for those models can be obtained. It would be interesting to identify different flow regimes by varying the Reynolds number (affecting the transfer function and the dispersion relation of the Ginzburg–Landau oscillators), the Mach number (by changing the delay time in the proportional terms and the time delay in the delayed system of differential equation) and the effect of the angle of incidence (by changing the degree of asymmetry of the delayed second-order oscillators and the dispersion relations). With the aid of these model problems, it could be interesting to identify different regimes where features reminiscent of the ones observed in the above simulations persist. This type of study could further our understanding of the essentials of the tonal-noise generation mechanism, the importance of its various components, and its dependence on the governing parameters. In addition, a model-based approach using delay-differential oscillatory systems could provide guidance for a further investigation of the tonal-noise problem using the tools developed and presented in this thesis.

Bibliography

- N. A. Adams and K. Shariff. A high-resolution hybrid compact-ENO scheme for shock-turbulence interaction problems. *Journal of Computational Physics*, 127(1):27–51, 1996.
- T. Akbar. *Simulation of three-dimensional flows past infinite cylindrical bodies*. PhD thesis, Université de Strasbourg, 2010.
- E. Åkervik, L. Brandt, D. S. Henningson, J. Hoepffner, O. Marxen, and P. Schlatter. Steady solutions of the Navier–Stokes equations by selective frequency damping. *Physics of fluids*, 18:068102, 2006.
- H. Arbey and J. Bataille. Noise generated by airfoil profiles placed in a uniform laminar flow. *Journal of Fluid Mechanics*, 134(1):33–47, 1983.
- E. Arcondoulis, C. J. Doolan, and A. C. Zander. Airfoil noise measurements at various angles of attack and low Reynolds number. *Acoustics 2009*, 2009.
- S. Balay, W. D. Gropp, L. C. McInnes, and B. F. Smith. Efficient management of parallelism in object oriented numerical software libraries. In E. Arge, A. M. Bruaset, and H. P. Langtangen, editors, *Modern Software Tools in Scientific Computing*, page 163–202. Birkhäuser Press, 1997.
- S. Balay, J. Brown, K. Buschelman, W. D. Gropp, D. Kaushik, M. G. Knepley, L. C. McInnes, B. F. Smith, and H. Zhang. *PETSc Web page*. 2011. <http://www.mcs.anl.gov/petsc>.
- D. J. Bodony. Analysis of sponge zones for computational fluid mechanics. *Journal of Computational Physics*, 212(2):681–702, 2006.
- A. Bottaro, P. Corbett, and P. Luchini. The effect of base flow variation on flow stability. *Journal of Fluid Mechanics*, 476, 2003.
- L. Brandt, D. Sipp, J. O. Pralits, and O. Marquet. Effect of base-flow variation in noise amplifiers: the flat-plate boundary layer. *Journal of Fluid Mechanics*, 687:503–528, 10 2011.
- T. F. Brooks, D. S. Pope, and M. A. Marcolini. Airfoil self-noise and prediction. Reference Publication 1218, NASA, 1989.
- G.A. Brès and T. Colonius. Three-dimensional instabilities in compressible flow over open cavities. *Journal of Fluid Mechanics*, 599:309–339, 2008.
- J. C. Buell and P. Huerre. Inflow/Outflow boundary conditions and global dynamics of spatial mixing layers. *Proceedings of the Summer Program, Center for Turbulence Research*, 1988.
- A. Carle and M. Fagan. ADIFOR 3.0 overview. Technical Report CAAM-TR-00-02, Department of Computational and Applied Mathematics, Rice University, 2000.

- H.K. Cheng and F.T. Smith. The influence of airfoil thickness and Reynolds number on separation. *Zeitschrift für Angewandte Mathematik und Physik (ZAMP)*, 33(2):151–180, 1982.
- S. Cherubini, J.-Ch. Robinet, and P. De Palma. The effects of non-normality and nonlinearity of the Navier–Stokes operator on the dynamics of a large laminar separation bubble. *Physics of Fluids*, 22(1):014102, 2010.
- J.-M. Chomaz. Global instabilities in spatially developing flows: Non-normality and nonlinearity. *Annual Review of Fluid Mechanics*, 37:357–392, 2005.
- T. P. Chong and P. Joseph. "Ladder" structure in tonal noise generated by laminar flow around an airfoil. *Journal of the Acoustical Society of America*, 131(6):461–467, 2012.
- B. Chu. On the energy transfer to small disturbances in fluid flow (part i). *Acta Mechanica*, 1(3): 215–234, September 1965.
- A. Collette. h5py, 2012. URL <http://code.google.com/p/h5py/>.
- S. S. Collis, K. Ghayour, and M. Heinkenschloss. Optimal transpiration boundary control for aeroacoustics. *AIAA Journal*, 41(7):1257–1270, 2003.
- T. Colonius. Modeling artificial boundary conditions for compressible flow. *Annual Review of Fluid Mechanics*, 36:315–345, 2004.
- T. Colonius and S. K. Lele. Computational aeroacoustics: progress on nonlinear problems of sound generation. *Progress in Aerospace Sciences*, 40(6):345–416, 2004.
- D. G. Crighton. Basic principles of aerodynamic noise generation. *Progress in Aerospace Sciences*, 16(1):31–96, 1975.
- J.D. Crouch, A. Garbaruk, and D. Magidov. Predicting the onset of flow unsteadiness based on global instability. *Journal of Computational Physics*, 224(2):924–940, 2007.
- G. Desquesnes, M. Terracol, and P. Sagaut. Numerical investigation of the tone noise mechanism over laminar airfoils. *Journal of Fluid Mechanics*, 591:155–182, 2007.
- A.V. Dovgal, V.V. Kozlov, and A. Michalke. Laminar boundary layer separation: instability and associated phenomena. *Progress in Aerospace Sciences*, 30(1):61–94, 1994.
- P. G. Drazin and W.H. Reid. *Hydrodynamic Stability*. Cambridge Mathematical Library. Cambridge University Press, 2nd edition, 2004.
- M. Drela. XFOIL: an analysis and design system for low Reynolds number airfoils. University of Notre Dame, 1989.
- W. S. Edwards, L. S. Tuckerman, R. A. Friesner, and D. C. Sorensen. Krylov methods for the incompressible Navier–Stokes equations. *Journal of Computational Physics*, 110:82–102, 1994.
- U. Ehrenstein and F. Gallaire. Two-dimensional global low-frequency oscillations in a separating boundary-layer flow. *Journal of Fluid Mechanics*, 614:315, 2008.
- M. R. Fink. Prediction of airfoil tone frequencies. *Journal of Aircraft*, 12(2):118–120, 1974.

- M. Fosas de Pando, D. Sipp, and P. J. Schmid. Efficient evaluation of the direct and adjoint linearized dynamics from compressible flow solvers. *Journal of Computational Physics*, 231(23):7739–7755, October 2012.
- X. Garnaud. *Modes, transient dynamics and forced response of circular jets*. PhD thesis, École Polytechnique, 2012.
- F. Giannetti and P. Luchini. Leading-edge receptivity by adjoint methods. *Journal of Fluid Mechanics*, 547:21–53, 2006.
- F. Giannetti and P. Luchini. Structural sensitivity of the first instability of the cylinder wake. *Journal of Fluid Mechanics*, 581:167–197, 2007.
- M.B. Giles and N.A. Pierce. An introduction to the adjoint approach to design. *Flow, Turbulence and Combustion*, 65(3):393–415, 2000.
- D.A. Hammond and L.G. Redekopp. Local and global instability properties of separation bubbles. *European Journal of Mechanics-B/Fluids*, 17(2):145–164, 1998.
- A. Hanifi, P. J. Schmid, and D. S. Henningson. Transient growth in compressible boundary layer flow. *Physics of Fluids*, 8:826, 1996.
- L. Hascoët and V. Pascual. TAPENADE 2.1 user’s guide. Rapport technique 300, INRIA, Sophia Antipolis, 2004.
- A. Henderson. ParaView guide, a parallel visualization application, 2007.
- V. Hernandez, J. E. Roman, and V. Vidal. SLEPc: a scalable and flexible toolkit for the solution of eigenvalue problems. *ACM Transactions on Mathematical Software*, 31(3):351–362, 2005.
- D. C. Hill. A theoretical approach for analyzing the restabilization of wakes. Technical Report 103858, NASA, 1992.
- D. C. Hill. Adjoint systems and their role in the receptivity problem for boundary layers. *Journal of Fluid Mechanics*, 292:183–204, 1995.
- M. Hochbruck and C. Lubich. On Krylov subspace approximations to the matrix exponential operator. *SIAM Journal on Numerical Analysis*, 34(5):1911–1925, 1997.
- L. E. Jones and R. D. Sandberg. Numerical analysis of tonal airfoil self-noise and acoustic feedback-loops. *Journal of Sound and Vibration*, 2011.
- L. E. Jones, R. D. Sandberg, and N. D. Sandham. Direct numerical simulations of forced and unforced separation bubbles on an airfoil at incidence. *Journal of Fluid Mechanics*, 602, 2008.
- L. E. Jones, R. D. Sandberg, and N. D. Sandham. Stability and receptivity characteristics of a laminar separation bubble on an aerofoil. *Journal of Fluid Mechanics*, 648:257–296, 2010.
- M. P. Juniper. Triggering in the horizontal Rijke tube: non-normality, transient growth and bypass transition. *Journal of Fluid Mechanics*, 667:272–308, 2011.
- C. T. Kelley. *Iterative methods for optimization*, volume 18 of *Frontiers in Applied Mathematics*. Society for Industrial and Applied Mathematics, Philadelphia, 1999.

- C. A. Kennedy, M. H. Carpenter, and R. M. Lewis. Low-storage, explicit Runge–Kutta schemes for the compressible Navier–Stokes equations. *Applied Numerical Mathematics*, 35(3):177–219, 2000.
- J. Kim and T. R. Bewley. A linear systems approach to flow control. *Annual Review of Fluid Mechanics*, 39:383–417, 2006.
- J. Kim and P. Moin. Application of a fractional-step method to incompressible Navier–Stokes equations. *Journal of Computational Physics*, 59(2):308–323, 1985.
- J.W. Kim. Quasi-disjoint pentadiagonal matrix systems for the parallelization of compact finite-difference schemes and filters. *Journal of Computational Physics*, 241:168–194, 2013.
- M. J. Kingan and J. R. Pearse. Laminar boundary layer instability noise produced by an aerofoil. *Journal of Sound and Vibration*, 322(4-5):808–828, 2009.
- D.A. Knoll and D.E. Keyes. Jacobian-free Newton–Krylov methods: a survey of approaches and applications. *Journal of Computational Physics*, 193(2):357–397, 2004.
- T. Le Garrec. *Simulation directe du bruit de bord de fuite d’un profil par une méthode de multi domaines*. PhD thesis, École Nationale Supérieure d’Arts et Métiers, 2008.
- S. K. Lele. Compact finite difference schemes with spectral-like resolution. *Journal of Computational Physics*, 103(1):16–42, 1992.
- G. Lodato, P. Domingo, and L. Vervisch. Three-dimensional boundary conditions for direct and large-eddy simulation of compressible viscous flows. *Journal of Computational Physics*, 227(10):5105–5143, 2008.
- R. E. Longhouse. Vortex shedding noise of low tip speed, axial flow fans. *Journal of Sound and Vibration*, 53(1):25–46, 1977.
- C.J. Mack. *Global stability of compressible flow about a swept parabolic body*. PhD thesis, Ecole Polytechnique, 2009.
- C.J. Mack and P.J. Schmid. A preconditioned Krylov technique for global hydrodynamic stability analysis of large-scale compressible flows. *Journal of Computational Physics*, 229(3):541–560, 2010.
- C.J. Mack and P.J. Schmid. Global stability of swept flow around a parabolic body: the neutral curve. *Journal of Fluid Mechanics*, 678:589–599, 2011.
- O. Marquet, D. Sipp, and L. Jacquin. Sensitivity analysis and passive control of cylinder flow. *Journal of Fluid Mechanics*, 615:221–252, 2008.
- A. McAlpine, E. C. Nash, and M. V. Lowson. On the generation of discrete frequency tones by the flow around an aerofoil. *Journal of Sound and Vibration*, 222(5):753–779, 1999.
- P. Meliga, D. Sipp, and J.-M. Chomaz. Effect of compressibility on the global stability of axisymmetric wake flows. *Journal of Fluid Mechanics*, 660:499–526, 2010.
- Message Passing Interface Forum. MPI: a message-passing interface standard. version 2.2, 2009. URL <http://www.mpi-forum.org>.

- B.E. Mitchell, S.K. Lele, and P. Moin. *Direct computation of the sound generated by subsonic and supersonic axisymmetric jets*. PhD thesis, Stanford University, 1995.
- K. Mohseni, T. Colonius, and J.B. Freund. An evaluation of linear instability waves as sources of sound in a supersonic turbulent jet. *Physics of Fluids*, 14(10):3593, 2002.
- R. B. Morgan and M. Zeng. A harmonic restarted arnoldi algorithm for calculating eigenvalues and determining multiplicity. *Linear algebra and its applications*, 415(1):96–113, 2006.
- J.D. Müller and P. Cusdin. On the performance of discrete adjoint CFD codes using automatic differentiation. *International Journal for Numerical Methods in Fluids*, 47(8-9):939–945, 2005.
- T. Nakano, N. Fujisawa, and S. Lee. Measurement of tonal-noise characteristics and periodic flow structure around NACA0018 airfoil. *Experiments in Fluids*, 40(3):482–490, 2006.
- E. C. Nash, M. V. Lowson, and A. McAlpine. Boundary-layer instability noise on aerofoils. *Journal of Fluid Mechanics*, 382:27–61, 1999.
- J. Nocedal and S. J. Wright. *Numerical Optimization*. Springer, 2nd edition, 2006.
- R. W. Paterson, P. G. Vogt, M. R. Fink, and C. L. Much. Vortex noise of isolated airfoils. *Journal of Aircraft*, 10(5):296–302, 1973.
- S. Pirozzoli. Numerical methods for high-speed flows. *Annual Review of Fluid Mechanics*, 43(1):163–194, 2011.
- T. J. Poinso and S. K. Lele. Boundary conditions for direct simulations of compressible viscous flows. *Journal of Computational Physics*, 101(1):104–129, 1992.
- C. W. Rowley, D. R. Williams, T. Colonius, R. M. Murray, and D. G. Macmynowski. Linear models for control of cavity flow oscillations. *Journal of Fluid Mechanics*, 547:317, 2006.
- Y. Saad. *Iterative methods for sparse linear systems*. Society for Industrial and Applied Mathematics, 2nd edition, 2003.
- Y. Saad. *Numerical methods for large eigenvalue problems*. Society for Industrial and Applied Mathematics, revised edition, 2011.
- R. D. Sandberg, L. E. Jones, N. D. Sandham, and P. F. Joseph. Direct numerical simulations of tonal noise generated by laminar flow past airfoils. *Journal of Sound and Vibration*, 320:838–858, 2009.
- W.S. Saric, H.L. Reed, and E.J. Kerschen. Boundary-layer receptivity to freestream disturbances. *Annual Review of Fluid Mechanics*, 34(1):291–319, 2002.
- H. Schlichting and K. Gersten. *Boundary-layer theory*. Springer, 2004.
- P. J. Schmid. Nonmodal stability theory. *Annual Review of Fluid Mechanics*, 39:129–162, 2007.
- P. J. Schmid and D. S. Henningson. *Stability and transition in shear flows*, volume 142 of *Applied Mathematics Sciences*. Springer, 2001.
- P. J. Schmid and D. S. Henningson. On the stability of a falling liquid curtain. *Journal of Fluid Mechanics*, 463:163–171, 2002.

- J. C. Schulze, P. J. Schmid, and J. Sesterhenn. Exponential time integration using Krylov subspaces. *International Journal for Numerical Methods in Fluids*, 60(6):591–609, 2009.
- J. Sesterhenn. A characteristic-type formulation of the Navier–Stokes equations for high order upwind schemes. *Computers & fluids*, 30(1):37–67, 2000.
- R. Seydel. *Practical Bifurcation and Stability Analysis*. Interdisciplinary Applied Mathematics. Springer, third edition, 2010.
- C. W. Shu. Essentially non-oscillatory and weighted essentially non-oscillatory schemes for hyperbolic conservation laws. Technical Report NASA/CR-97-206253, NASA, 1998.
- R. B. Sidje. Expokit. a software package for computing matrix exponentials. *ACM Trans. Math. Softw.*, 24(1):130–156, 1998.
- D. Sipp and A. Lebedev. 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, 2007.
- Z. Sirkes and E. Tziperman. Finite difference of adjoint or adjoint of finite difference? *Monthly weather review*, 125(12):3373–3378, 1997.
- B. Spagnoli and C. Airiau. Adjoint analysis for noise control in a two-dimensional compressible mixing layer. *Computers & Fluids*, 37(4):475–486, 2008.
- G. W. Stewart. A Krylov–Schur algorithm for large eigenproblems. *SIAM Journal on Matrix Analysis and Applications*, 23(3):601–614, 2002.
- P. J. Strykowski and K. R. Sreenivasan. On the formation and suppression of vortex ‘shedding’ at low Reynolds numbers. *Journal of Fluid Mechanics*, 218:71–107, 1990.
- S. Takagi and Y. Konishi. Frequency selection mechanism of airfoil trailing-edge noise. *Journal of Aircraft*, 47(4):1111–1116, 2010.
- Y. Takagi, N. Fujisawa, T. Nakano, and A. Nashimoto. Cylinder wake influence on the tonal noise and aerodynamic characteristics of a NACA0018 airfoil. *Journal of Sound and Vibration*, 297(3-5): 563–577, 2006.
- C. K. W. Tam. Discrete tones of isolated airfoils. *Journal of the Acoustical Society of America*, 55(6): 1173–1177, 1974.
- C. K. W. Tam and H. Ju. Aerofoil tones at moderate Reynolds number. *Journal of Fluid Mechanics*, 2011.
- C.K.W. Tam and J.C. Webb. Dispersion-relation-preserving finite difference schemes for computational acoustics. *Journal of Computational Physics*, 107(2):262–281, 1993.
- The HDF Group. HDF5 library v1.8.9, 2012. URL <http://www.hdfgroup.org/HDF5/>.
- L. N. Trefethen. Group velocity interpretation of the stability theory of Gustafsson, Kreiss, and Sundström. *Journal of Computational Physics*, 49(2):199–217, 1983.
- L.N. Trefethen. Pseudospectra of linear operators. *Siam Review*, 39(3):383–406, 1997.

- L.N. Trefethen, A.E. Trefethen, S.C. Reddy, and T.A. Driscoll. Hydrodynamic stability without eigenvalues. *Science*, 261(5121):578–584, 1993.
- H.A. Van der Vorst. *Iterative Krylov methods for large linear systems*, volume 13. Cambridge University Press, 2003.
- M. Wei and J. B. Freund. A noise-controlled free shear flow. *Journal of Fluid Mechanics*, 546:123–152, 2005.
- S. E. Wright. The acoustic spectrum of axial flow machines. *Journal of Sound and Vibration*, 45(2): 165–223, 1976.
- S. Zuccher, A. Bottaro, and P. Luchini. Algebraic growth in a Blasius boundary layer: Nonlinear optimal disturbances. *European Journal of Mechanics-B/Fluids*, 25(1):1–17, 2006.

Résumé

La génération de fréquences discrètes dans l'écoulement autour d'un profil d'aile est étudiée dans cette thèse au moyen de simulations numériques non-linéaires et d'études de stabilité globale. À cette fin, un code numérique mettant en œuvre une nouvelle technique pour accéder à la dynamique linéaire, directe et adjointe, a d'abord été développé et ensuite appliqué à l'écoulement autour d'un profil. Les simulations non-linéaires confirment l'apparition de fréquences discrètes dans le spectre sonore, et pour le cas considéré, l'analyse de stabilité globale de la dynamique linéaire de l'écoulement moyen montre que celle-ci est stable. Cependant, la réponse de l'écoulement à des perturbations incidentes révèle de fortes croissances transitoires amenant à l'établissement de cycles de rétroaction aéroacoustique. Ces cycles comprennent la croissance des instabilités hydrodynamiques dans les couches limites sur l'intrados et l'extrados ainsi que leurs interactions avec les radiations acoustiques du bord de fuite. Les processus associés à l'apparition des fréquences discrètes dans le spectre sonore ainsi que les cycles de rétroaction sont ensuite mis en relation avec les modes globaux les moins stables : d'un côté, la structure spatiale des modes directs montre la croissance des instabilités hydrodynamiques sur l'extrados et la zone du sillage proche du bord de fuite; d'un autre côté, les modes adjoints associés présentent l'intrados comme la zone la plus réceptive à des perturbations externes. Finalement, l'analyse de la région dite du *wavemaker* indique, en accord avec les expériences, le rôle fondamental de la couche limite sur l'intrados.

Abstract

The generation of discrete acoustic tones in the compressible flow around an aerofoil is addressed in this thesis by means of nonlinear numerical simulations and global stability analyses. To this end, a nonlinear simulation code featuring a novel technique for gaining access to the linearized direct and adjoint dynamics has been developed and applied to the flow around an aerofoil. The nonlinear simulations confirm the appearance of discrete tones in the acoustic spectrum, and for the chosen flow case, the global stability analyses of the mean-flow dynamics reveal that the linearized operator is stable. However, the flow response to incoming disturbances exhibits important transient growth effects that culminate into the onset of aeroacoustic feedback loops, involving instability process on the suction- and pressure-surface boundary-layers together with their cross interaction by acoustic radiation at the trailing edge. The features of the aeroacoustic feedback loops and the appearance of discrete tones are then related to the features of the least stable modes in the global spectrum: on the one hand, the spatial structure of the direct modes display the growth of hydrodynamic instabilities on the suction surface and the near wake; on the other hand, the associated adjoint modes display increased receptivity of the flow on the pressure surface. Finally, the analysis of the wavemaker region highlights, in agreement with previous experimental investigations, the sensitivity of the flow to the pressure-surface boundary layer.