



INSTITUT
POLYTECHNIQUE
DE PARIS

NNT : 20XXIPPAXXXX

Thèse de doctorat



On the flutter bifurcation in laminar flows: linear and nonlinear modal methods

Thèse de doctorat de l'Institut Polytechnique de Paris
préparée à l'École polytechnique et
à l'Office National d'Études et de Recherches Aérospatiales

École doctorale n°626
École Doctorale de l'Institut Polytechnique de Paris (ED IP Paris)
Spécialité de doctorat : Mécanique des fluides et des solides

Thèse présentée et soutenue à Meudon, le Date, par

JOHANN MOULIN

Composition du Jury :

Jacques Magnaudet Directeur de recherche, CNRS (Institut de Mécanique des Fluides de Toulouse)	Président
Grigorios Dimitriadis Professeur, Université de Liège (Département d'Aérospatiale et Mécanique)	Rapporteur
Ardeshir Hanifi Professeur, KTH Royal Institute of Technology in Stockholm (Linné FLOW Centre)	Rapporteur
Xavier Amandolèse Maître de conférences, Conservatoire National des Arts et Métiers (Laboratoire de Mécanique des Structures et des Systèmes Couplés)	Examineur
Flavio Giannetti Professeur associé, Università degli Studi di Salerno (Dipartimento di Ingegneria Industriale)	Examineur
Pierre Jolivet Chargé de recherche, CNRS (Institut de Recherche en Informatique de Toulouse)	Examineur
Denis Sipp Directeur de recherche, ONERA (Département Aérodynamique, Aéroélasticité, Acoustique)	Directeur de thèse
Olivier Marquet Maître de recherche, ONERA (Département Aérodynamique, Aéroélasticité, Acoustique)	Directeur de thèse

ABSTRACT

The flutter instability has been the focus of numerous works since the middle of the twentieth century, due to its critical application in aeronautics. Flutter is classically described as a linear instability using potential flow models, but viscous and non-linear fluid effects may both crucially impact this aeroelastic phenomenon.

The first part of this thesis is devoted to the development of theoretical and numerical methods for analyzing the linear and nonlinear dynamics of a typical aeroelastic section — a heaving and pitching spring-mounted plate — immersed in a two-dimensional laminar flow modeled by the incompressible Navier–Stokes equations. A semi-analytical weakly nonlinear analysis (WNL) is first developed in order to derive an amplitude equation for the flutter bifurcation. In order to bypass the inherent limitations of this method to weak nonlinearities, we then develop a harmonic balance type method, known as the Time Spectral Method (TSM), allowing to efficiently compute — possibly unstable — highly-nonlinear periodic flutter solutions. The challenging task of solving the TSM equations, especially when large numbers N of Fourier harmonics are considered, is tackled *via* a time-parallel Newton–Krylov approach in combination with a new, so-called block-circulant preconditioner, for which N -robustness is numerically demonstrated.

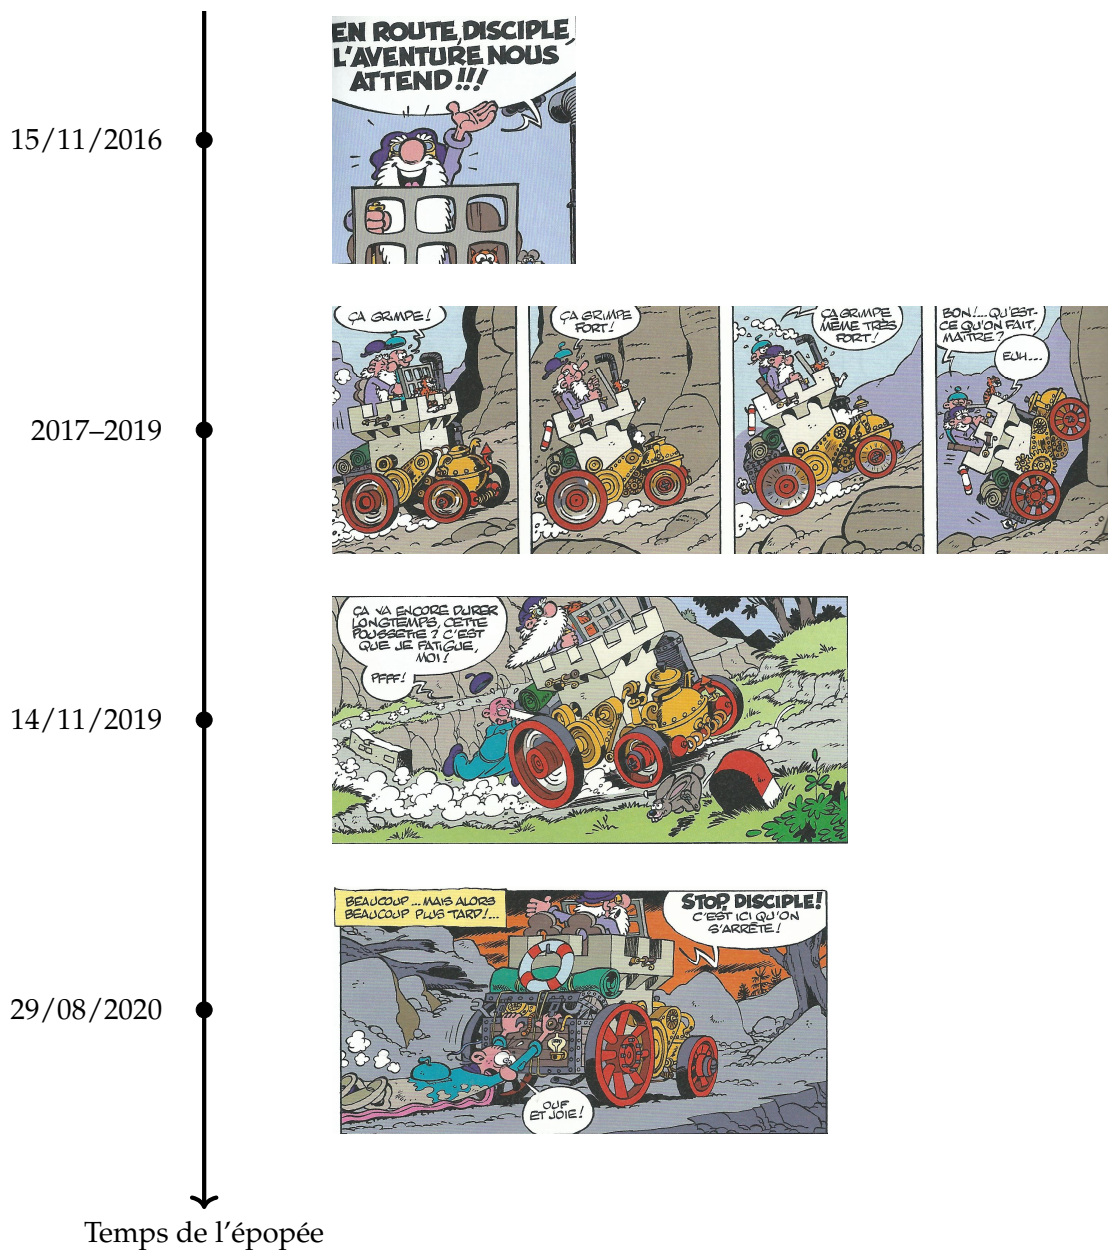
The second part of this thesis focuses on the physical investigation of the flutter bifurcation of the spring-mounted plate. We start by revisiting the linear stability problem using a Navier–Stokes fluid model allowing to highlight, in particular, the effect of viscosity. Comparisons to classical quasi-steady and unsteady potential flow (Theodorsen model) theories are performed. Contrary to what happens in potential flows, the flutter instability is shown to re-stabilize at very high reduced velocities in viscous flows. We continue our route on the flutter bifurcation by investigating the effect of fluid nonlinearities. Low solid-to-fluid mass ratios and increasing Reynolds numbers foster subcritical bifurcations. The role of leading-edge shear layers is pointed out. For intermediate mass ratios, an unusual bifurcation scenario that combines a supercritical bifurcation and the existence of subcritical high-amplitude flutter solutions is discovered. We conclude our study of the flutter bifurcation by investigating the appearance of low-frequency amplitude modulations on top of a previously established periodic flutter solution. Using an original TSM-based Floquet stability analysis, we explain this behavior by the destabilization of the periodic solutions by a pair of complex-conjugate Floquet modes. An analysis of the latter shows that the physical mechanism governing the instability borrows elements from the classical flutter instability arising on steady solutions.

The last part of this thesis aims at initiating the extension of the different methods previously evoked to large-scale three-dimensional configurations. As a first step towards this long-term goal, we develop an open-source massively parallel tool, based on the FreeFEM library and its PETSc/SLEPc interface, able to compute the nonlinear steady-state flow and subsequently solve the linear stability eigenproblem, for three-dimensional flows (the structure is fixed) possessing several tens of millions of degrees of freedom.

REMERCIEMENTS

Pour la version finale ...

RÉSUMÉ INFORMEL



Petit résumé informel illustré. Toute ressemblance avec des personnes ou des situations existantes ou ayant existé *ne saurait être* fortuite. Vignettes reproduites avec autorisation de TURK & DE GROOT (1982), "Génie en balade". *Léonard*, tome 6, Éditions du Lombard.

CONTENTS

Abstract	i
Remerciements	iii
Chronologie informelle	v
Contents	vii
Introduction	1
I Theoretical and numerical methods for nonlinear flutter analysis	17
1 Mathematical models for a “typical section” in viscous flows	19
1.1 Equations of motion	20
1.1.1 Solid model: spring-mounted rigid solid	21
1.1.2 Fluid model: incompressible Navier–Stokes	23
1.1.2.1 The “absolute velocity - rotating axis” formulation	24
1.1.2.2 The “reference configuration ALE” formulation	25
1.1.2.3 From the “reference configuration ALE” to the “ab- solute velocity - rotating axis”	27
1.1.3 Interface velocity continuity and interface forces	29
1.1.3.1 “Absolute velocity - rotating axis”	29
1.1.3.2 “Reference configuration ALE”	29
1.2 Spatial discretization with the finite element method	30
1.2.1 A weak formulation of the typical section problem	30
1.2.2 Finite elements discretized problem	33
1.3 Time-integration algorithm	35
1.3.1 Temporal scheme	36
1.3.2 Fluid-structure pressure-segregation algorithm	36
2 A weakly nonlinear solver for flutter bifurcation analysis	39
2.1 Introduction	40
2.2 Equations of motion: typical section in viscous flow	41
2.3 Weakly nonlinear analysis	42
2.3.1 Zero-th order: steady solution	42
2.3.2 First order: linear stability	43
2.3.3 Second order	43
2.3.4 Third order: Stuart–Landau equation	45
2.3.5 Solutions of the Stuart–Landau equation	46
2.3.6 Continuous adjoint equations	48
2.3.7 Summary	50
2.4 Numerical solution	51
2.4.1 SUPG-stabilized finite elements discretization	51
2.4.2 Solution method	53

2.4.3	Mesh adaptation loop	54
2.4.4	Validation	58
2.5	Conclusion	59
Appendices		61
Appendix 2.A	Development in ϵ series of trigonometric nonlinearities	61
Appendix 2.B	Operators of the SUPG-discretized eigenproblem eq. (2.30b)	63
3	A Time Spectral solver for periodic flutter solutions using block-circulant preconditioning	65
3.1	Introduction	66
3.2	Harmonic Balance methods	68
3.2.1	Galerkin approach: frequency-domain HBM	69
3.2.1.1	Analytical Harmonic Balance Method (AHBM)	70
3.2.1.2	Pseudo-spectral Harmonic Balance Method	71
3.2.2	Collocation approach : time-domain HBM	73
3.2.3	Comparison of the Galerkin and collocation approaches	74
3.2.3.1	Equivalence between the pseudo-spectral Galerkin and collocation approaches	74
3.2.3.2	Aliasing error	75
3.2.3.3	Synthetic view	76
3.3	Numerical solution of the Time Spectral Method	77
3.3.1	Newton method	77
3.3.2	Solution of the linearized TSM equations	78
3.3.3	A block-circulant preconditioner for the TSM Jacobian	80
3.3.4	Parallel implementation	82
3.4	Numerical results	83
3.4.1	NACA0012 in forced heaving motion	83
3.4.1.1	Reference DNS solution and TSM time-convergence	84
3.4.1.2	Assessment of the TSM solver performance	86
3.4.2	Flutter instability of a 2DOF spring-mounted plate: a case with unknown frequency	95
3.5	Conclusion	97
Appendices		101
Appendix 3.A	Discretized operators for the numerical examples of section 3.4	101
3.A.1	Coupled problem: spring-mounted plate	101
3.A.1.1	Nonlinear equations	101
3.A.1.2	Linearized equations	102
3.A.2	Forced problem: solid with imposed heaving motion	104
3.A.2.1	Nonlinear equations	104
3.A.2.2	Linearized equations	104
II	Flutter bifurcation analysis in viscous flows	105
4	Linear stability of a typical section in viscous incompressible flows	107
4.1	Introduction	108
4.2	Problem settings and methods	109
4.3	Different types of flow-induced vibrations	111
4.3.1	Leading eigenmodes	111

4.3.2	Varying the reduced velocity: from VIV to coupled-mode flutter and divergence	112
4.3.3	Increasing the steady angle of attack: from coupled-mode to single-mode flutter	117
4.4	Parametric explorations	120
4.4.1	Effect of mass ratio	121
4.4.2	Effect of Reynolds number	124
4.4.3	Full parametric exploration of flutter thresholds in the (\tilde{m}, Re) plane	126
4.5	Comparison to simplified fluid models	128
4.5.1	Theodorsen's model	130
4.5.2	Quasi-steady models	131
4.5.3	Numerical comparison of the flutter predictions	132
4.6	Conclusion	134
Appendices		137
	Appendix 4.A Computation of the steady lift and moment slopes with Navier–Stokes model	137
5	Effect of fluid nonlinearity on the flutter bifurcation of a typical section	139
5.1	Introduction	140
5.2	Governing equations and numerical methods	141
5.2.1	Time-marching simulations of nonlinear solutions	143
5.2.2	Methods for periodic nonlinear solutions	143
5.2.2.1	Weakly nonlinear analysis	143
5.2.2.2	Time Spectral Method	147
5.2.3	Cross-validation of numerical methods	148
5.3	Hard and soft flutter with weakly nonlinear analysis	148
5.3.1	Effects of mass ratio and Reynolds number on the bifurcation	151
5.3.2	Fluid and geometric nonlinearities	153
5.4	A mean flow approach for the bifurcation analysis at high Reynolds number and mass ratio	155
5.5	Bifurcation scenarios at low Reynolds number $Re = 500$	159
5.5.1	Supercritical bifurcation at high mass ratio	160
5.5.2	Subcritical bifurcation at low mass ratio	161
5.5.3	Double-fold bifurcation scenario at intermediate mass ratio	167
5.5.4	Experimental evidence of the double-fold scenario	168
5.6	Conclusion	170
Appendices		173
	Appendix 5.A Grid convergence study	173
6	Low-frequency modulation of flutter solutions: a Floquet analysis based on the Time Spectral Method	175
6.1	Introduction	176
6.2	Governing equations and TSM-based numerical approach for Floquet stability analysis	178
6.2.1	Governing equations	178
6.2.2	Linear stability analysis of Limit Cycle Oscillations	179
6.2.3	A Time Spectral Method-based approach for Floquet stability analysis	180

6.2.4	Time-marching simulations	185
6.3	Results	185
6.3.1	From periodic to quasi-periodic solutions	186
6.3.2	Floquet stability analysis of flutter LCO's	190
6.3.3	Analysis of the quasi-periodic perturbation: a "generalized flutter" instability	195
6.3.4	Comparison with time-marching results	199
6.4	Conclusion	203
III Towards large-scale linear stability analysis		205
7	Augmented Lagrangian Preconditioner for Large-Scale Hydrodynamic Stability Analysis	207
7.1	Introduction	207
7.2	Methods for linear stability analysis in hydrodynamics	210
7.2.1	Governing equations	210
7.2.2	Nonlinear steady-state solver	212
7.2.3	Linear eigensolver	213
7.3	An augmented Lagrangian approach for the shifted Jacobian matrix	213
7.4	Parallel implementation with FreeFEM and its interface to PETSc/SLEPc	215
7.4.1	Outer solvers	215
7.4.2	Inner mAL-preconditioned linear solvers	216
7.4.3	Innermost velocity and pressure linear solvers	218
7.5	Numerical results	219
7.5.1	Two- and three-dimensional test cases	219
7.5.2	Influence of numerical and physical parameters	221
7.5.2.1	Effect of the augmentation parameter	222
7.5.2.2	Effect of the shift parameter	223
7.5.2.3	Effect of the mesh refinement and Reynolds number	224
7.5.3	Comparison with other block preconditioners	225
7.5.4	Performance of the parallel implementation	228
7.5.4.1	Comparison with a direct solver on a small-scale 3D configuration	228
7.5.4.2	Parallel performance on a large-scale 3D configuration	230
7.6	Conclusion	234
Appendices		235
Appendix 7.A	Reproducibility	235
Appendix 7.B	Definition of other block preconditioners	235
Appendix 7.C	Linear solver tolerance and eigenvalue convergence criterion	236
8	Conclusion & perspectives	239
8.1	Conclusion	239
8.2	Perspectives	243
Bibliography		247

INTRODUCTION

The interaction of fluid flows and flexible solids is ubiquitous in a wide variety of natural systems. The vibrations of the soft palate at the origin of snoring, the oscillations of the vocal folds that produce voice, insects and birds flight or fish locomotion, all these phenomena are driven by the interaction of moving solids and fluid flows. Man-made systems, for their part, are also largely impacted by such *fluid-structure interactions*. Most of the time, they lead to undesirable dangerous oscillatory behaviors of elongated flexible-enough structures like aircraft wings, tall towers or long bridges. In some cases, on the contrary, these same interactions are deliberately triggered by engineers, for example for energy harvesting purposes or for designing bio-inspired drones. In both cases, a deep understanding of fluid-structure interactions is required in order to accurately predict them and ultimately propose adequate designs, adapted to these multiple applications. The present manuscript is dedicated to the study of the emergence of *fluid-structure instabilities*, *i.e.* oscillatory behaviors of coupled fluid-structure systems that emerge naturally without externally applied forcing.

Fluid-structure instabilities: a brief overview

Systems that couple fluids and structures are prone to a wide variety of instabilities. In this thesis, we are mainly focused on the so-called coupled-mode flutter, that we first introduce. Then, we very briefly present a few other fluid-structure instabilities, some of which will be occasionally met across the manuscript.

Coupled-mode flutter

The *coupled-mode flutter* instability occurs when two eigenmodes of a structure interact through the action of fluid forces so that one of the two modes becomes unstable. This instability typically occurs on airplane wings at high velocities. From the early years of aeroelastic studies [Bisplinghoff et al. 1955], it was understood that the core mechanisms at the origin of wing flutter could be described using a simplified model, referred to as a *typical aeroelastic section* (fig. 1). It consists in a thin airfoil section — or more simply a thin plate in this thesis — attached to flexion (heaving) and torsion (pitching) springs. This system is meant to mimic a two-dimensional cross section of a three-dimensional cantilever wing that possesses flexion and torsion modes. To physically understand the very basic mechanism of flutter, no complex aerodynamic model is necessary. Typically, it can be assumed that the instantaneous flow around the typical section is simply the flow around the same section, *frozen* at the instantaneous angle of attack. As it ignores any effect of flow unsteadiness this hypothesis is referred to as the *steady flow hypothesis*. Using the latter, it can be shown [E. H. Dowell et al. 1989, §3.3.5], that flutter occurs when the velocity of the incoming flow increases above the *critical flutter velocity*. The instability is linked to the progressive coupling of the flexion and torsion modes as velocity increases. With a steady flow model, this materializes as a frequency coalescence of the two modes (fig. 2(a)). As a consequence of aeroelastic coupling, the modes progressively acquire coupled flexion-torsion dynamics that are crucial for the emergence of the instability.




FIGURE 1: Typical aeroelastic section




FIGURE 2: Coupled-mode flutter instability. (a) growth rate (top) and frequency (bottom) of the heaving (red) and pitching (black) eigenmodes as a function of reduced velocity, computed with steady aerodynamics. The exact formulation is the one presented in [Hodges et al. 2011, fig. 5.3-5.4]. The vertical dashed line represents the critical reduced velocity (b-c) Illustration of the typical dynamics for flutter in (b) (red disk in (a)) and anti-flutter in (c) (black disk in (a)). The vertical blue arrows represent the lift force, applied at the quarter-chord point, according to thin airfoil aerodynamics.

The unstable mode (red disk) is associated to a particular swimming-like motion [De Langre 2002], illustrated in fig. 2(b), where the instantaneous angle of attack signal¹ ($-\theta(t)$, dashed line) precedes the heaving signal ($h(t)$, solid line) of about a quarter period. It is easy to see from the vertical blue arrows representing the lift force that the flow actually provides energy to the plate. On the contrary, the stabilized mode (black disk) illustrated in fig. 2(c) corresponds to a motion where the instantaneous angle of attack signal ($-\theta(t)$, dashed line) lags behind the heaving signal ($h(t)$, solid line) of about a quarter period. In this case, the flow works against the movement of the plate. In fact, it is also a swimming motion, but for a plate moving in the opposite direction (*i.e.* to right, here). As it looks like “the opposite of flutter”, such a motion is called *anti-flutter* in this manuscript.

¹Across this manuscript, we adopt the trigonometric convention for orienting angles, *i.e.* $\theta > 0$ nose-down. However, most of the time in the figures, we will represent the angle of attack $-\theta$ (“positive, nose-up”), as more classically done in aerodynamic studies.

Divergence, stall flutter, vortex-induced vibrations, etc

During the early years of monoplane aviation (*e.g.* the Fokker D-8 combat airplane), the passage from biplane to monoplane designs resulted in a number of failures during high-speed dives due to the so-called *divergence instability*. Contrary to flutter, this instability is static in the sense that it does not involve any oscillations of the wing. Divergence is the consequence of a loss of total stiffness — the sum of structural and aerodynamically-induced stiffnesses — at high velocities [Hodges et al. 2011, Chapter 4]. Because of their low torsional stiffness, early monoplanes were particularly prone to divergence in flight conditions that biplanes could withstand.

Stall flutter refers to aeroelastic instabilities that involve a stall phenomenon, or more generally, significant flow separation. It is the consequence of a loss of total damping — the sum of structural and aerodynamically-induced dampings — as velocity increases. Such a situation may occur for example on a single flexion mode, if the lift is a decreasing function of angle of attack. This typically occurs for airfoils positioned at large angles of attack, close to stall [E. H. Dowell et al. 1989, Chapter 5], or for certain bluff-bodies like square cylinders [Païdoussis et al. 2011, Chapter 2]

All instabilities mentioned until now emerge despite a fluid that is completely free of unsteadiness. The coupling between the structure motion and the flow is necessary for these instabilities to occur. However, fluid on their own are known to exhibit a wide variety of instabilities [Schmid et al. 2001; Charru 2011]. Some of them may interact with structural modes, triggering new types of fluid-structure instabilities. In *Vortex-induced vibrations (VIV)*, the hydrodynamic wake instability of bluff-bodies [Zebib 1987; Sipp et al. 2007] couples to a solid mode (the heaving mode typically). In the nonlinear regime, such a coupling typically leads to high-amplitude oscillations and a lock-in of the wake frequency on the solid frequency [Williamson et al. 2004; Navrose et al. 2016]. Since a fluid unsteadiness pre-exists to any structural motion, VIV are often thought of as the response of the solid to an imposed *external forcing* of the fluid, with negligible retro-action of the solid motion onto the flow. This vision has been proved wrong by a series of studies [Cossu et al. 2000; Mittal et al. 2005] where a spring-mounted cylinder was shown to get unstable at Reynolds numbers as low as half the critical Reynolds number for the fixed cylinder. Other purely fluid instabilities, like transonic shock buffet result in similar interactions with moving structures [Gao et al. 2017; Gao et al. 2020].

The challenges of aeroelasticity in laminar flows

The pioneering works in aeroelasticity by Glauert (1930), Theodorsen (1935) or Von Kármán et al. (1938) made the hypothesis of inviscid flows. In the inviscid world, the flow remains fully attached to the airfoil and the viscous effects that appear in the boundary layers are only taken into account *via* a Kutta condition. For the high-Reynolds flows typical of airplanes ($Re \sim 10^7 - 10^8$, see fig. 3(d-e)), the boundary layers are so thin that these assumptions are valid. In addition, they have the advantage of allowing analytical calculations, making these inviscid theories highly valuable for the engineer².

In the past decades, interest has grown in the design of smaller and slower aerial vehicles, often globally categorized as *Micro Aerial Vehicles (MAVs)*. Their range of application is wide, as summarized by Mueller et al. (2003): “surveillance, communication relay links, ship decoys, and detection of biological, chemical, or nuclear

²To this day, they are still used as a template to propose new models that are valid well outside the potential flow limit, *e.g.* [Brunton et al. 2013]




FIGURE 3: Typical Reynolds numbers for a variety of aerial vehicles (e.g. [Lissaman 1983]). (a) Robobee³(b) DelFly Micro[Croon et al. 2016] (b) palm-sized gliding MAV [Wood et al. 2007] (c) quadrotor Drone Volt Hercules 2⁴(d) Cessna 172⁵(e) Airbus A350⁶

materials". In addition to these critical, often military, applications, civilian use of MAVs has exploded in the past few years, for example for agricultural, commercial or recreational purposes. Due to their low cruise speed and small dimensions, these vehicles fly in regimes characterized by low- to moderate-Reynolds numbers, *i.e.* $Re = 10,000 - 500,000$ (see fig. 3(c)). Among the many challenges brought by these new flying devices is the need for understanding the specifics of low-Reynolds number aerodynamics and their repercussion on the aeroelastic behavior. For example, new types of aeroelastic instabilities may appear due to transitional flow features like Laminar Separation Bubbles [Poirel et al. 2008; Yuan et al. 2013]. The current trend in MAVs design is to further miniaturization, leading to very small unmanned aerial vehicles of the order of a few centimeters that operate in the so-called "ultra-low Reynolds number" regime, $Re \sim 10 - 10,000$ [Shyy et al. 2010]. These are often designed for flapping flight, by mimicking insects (see fig. 3(a-b)). The very light and flexible materials used imply strong fluid-structure interactions. Some researchers have even hypothesized that flutter-type instabilities may help various animals sustain flapping flight at reduced energetic cost [Michelin et al. 2009; Curet et al. 2013].

In the field of energy harvesting, fully passive devices that exploit coupled-mode flutter have been proposed [Peng et al. 2009; Pigolotti et al. 2017; Boudreau et al. 2018]. For these applications, the goal is not anymore to prevent instabilities to occur but rather to trigger vibrations at the lowest wind speed possible and yielding the highest energy extraction possible. Here also, the relevant Reynolds numbers are much lower than the ones encountered in classical aeronautics, typically in the low- to moderate-Reynolds number range.

Overall, it is seen that, in addition to their purely scientific interest, aeroelastic phenomena occurring at low Reynolds numbers are of significant interest for various modern applications. This is the global context that motivates the studies in this manuscript.

³Reproduced from <http://www.aboutkevinma.com>

⁴Reproduced from <https://www.coavmi.com>

⁵Reproduced from <https://www.dronevolt.com>

⁶Reproduced from https://fr.wikipedia.org/wiki/Airbus_A350_XWB

Linear stability

The first question relative to aeroelastic instabilities is to determine the threshold at which they first appear *i.e.* the parameters for which a given steady-state of an aeroelastic system is first destabilized. This threshold is usually expressed in terms of a non-dimensional wind velocity, called *reduced velocity*. Mathematically, the most classical indicator of the stability of a steady-state is called *linear stability* (or asymptotic stability) and consists in assessing the long-term behavior (growth or decay) of infinitesimal perturbations superposed onto the steady-state [Schmid et al. 2001]. In practice, this comes down to computing the rightmost eigenvalues of the coupled fluid-structure equations, linearized around the steady-state [Badcock et al. 2005].

The linear stability of spring-mounted airfoils and wings in high-Reynolds flows has been the focus of a vast amount of works. So much that every textbook on aeroelasticity contains chapters dedicated to it [Bisplinghoff et al. 1955, Ch. 8-9], [Hodges et al. 2011, Ch. 4-5]. Due to the high Reynolds number, these studies are usually performed using potential flow models, like the Theodorsen model [Theodorsen 1935] for thin airfoils or the Doublet Lattice Method (DLM) [Albano et al. 1969] for more general geometries. These analyses also form a key component of industrial practice to ensure that instability thresholds (for divergence and flutter, mainly) are outside the flight envelope of an airplane [Garrigues 2018]. During the last decade, several research groups have focused on taking into account viscous effects so as to improve the prediction of flutter stability in high-Reynolds transonic flows. These configurations are characterized by the so-called *transonic dip* which corresponds to a decrease of the flutter velocity threshold at transonic Mach numbers. For a heaving and pitching spring-mounted NACA64A010 (Isogai) airfoil, comparisons of Euler-based and RANS-based flutter stability boundaries were presented in [Badcock et al. 2011; Güner et al. 2018] and showed that viscous effects significantly affect the shape of the transonic dip and tend to delay flutter.

When it comes to the linear stability of spring-mounted airfoils at low Reynolds numbers, the literature is much less abundant. Some work has been done in [Le Maître et al. 2003; Bruno et al. 2008; Brunton et al. 2013] to numerically assess the effect of the Reynolds number on the flutter derivatives for heaving and pitching motion, but did not extend up to examining the repercussions on the flutter velocity threshold. Experimental and numerical investigations have been performed in [Chae et al. 2013] at moderate Reynolds number, showing that viscous effects are important for assessing coupled-mode flutter thresholds, in particular for low solid-to-fluid mass ratios where potential flow models overestimate the threshold. However, the Reynolds number $Re \sim 10^6$ in their study is still quite high. Recently, the aeroelastic stability of a spring-mounted pitching NACA0012 airfoil was investigated at transitional Reynolds numbers ($Re \sim 50000$) in [Negi 2019, Paper 3] in an effort to investigate the linear regime of the so-called “laminar separation flutter” originally reported in the experiments of Poirel et al. (2008). Perhaps an illustrating example of the need for more studies on the effect of viscosity on flutter can be found in [E. H. Dowell et al. 1989, p. 117] where the authors present the effect of some nondimensional parameters on the reduced velocity threshold for flutter on a typical aeroelastic section. Regarding the evolution of the threshold as a function of the solid-to-fluid mass ratio, it is noticed that the two-dimensional potential theory (Theodorsen model) predicts an infinite threshold for light enough airfoils, *i.e.* no flutter is observed. It is mentioned that no clear explanation for this infinite threshold has been provided. Among other hypothesis, the neglected effect of viscosity is proposed as a possible source for this singular behavior.

Apart from the typical aeroelastic section, the linear stability of several fluid-structure systems has however been studied in low-Reynolds viscous flows. The most classical of them might be the destabilization of a spring-mounted circular cylinder due to vortex-induced vibrations [Cossu et al. 2000]. Other than that, Cisonni et al. (2017) investigated the flutter of an elastic filament in viscous channel flow. Goza et al. (2018) studied the flow-induced oscillations of an inverted flag. The development of oscillations in the vertical path of freely falling or rising rigid objects have been considered by Assemat et al. (2012), Tchoufag et al. (2014a), and Tchoufag et al. (2014b). A flexible splitter plate attached to a rigid circular cylinder has been studied by J. L. Pfister et al. (2020).

In light of this bibliography, we propose in this manuscript to **revisit the linear stability of a typical aeroelastic section** — precisely, a heaving and pitching spring-mounted flat plate — **immersed in a low-Reynolds incompressible flow**. This study will be realised in the framework of global stability analysis, similar to the work of Assemat et al. (2012).

Nonlinear effects: bifurcation study

Once the critical velocity threshold is crossed, linear stability analyses predict an unbounded — hence unphysical — exponential growth of the unstable eigenmode. As the amplitude of oscillations increases, the nonlinearities of the system progressively come at play, saturate the exponential growth and settle the solution on a periodic orbit. This is a qualitative description of a *Hopf bifurcation* scenario [Nayfeh et al. 1995, §2.3].

Hopf bifurcations may be of two types. In a supercritical Hopf bifurcation, the periodic solutions exist only after the critical threshold and the amplitude of oscillations smoothly increases from the critical point (blue line in fig. 4). On the contrary, in a subcritical Hopf bifurcation, the branch of periodic solutions that emerges at the critical point progresses towards lower velocities. These solutions are unstable but subsequently stabilize *via a fold bifurcation* (red line in fig. 4). Such a scenario, often referred to as “hard” or “catastrophic” flutter, is particularly dangerous for two reasons. First, high-amplitude oscillations can be observed *below* the critical flutter threshold. Second, the systems response may be different when increasing or decreasing the wind velocity, due to the presence of a large hysteresis loop.

Despite the large predominance of linear aeroelasticity in industrial practice and undergraduate courses, the importance of nonlinear effects in aeroelastic phenomena has been pointed out since the early 1950’s. The work of Woolston et al. (1955) is often reported as the first detailed study on nonlinear flutter. The authors used a two-degrees-of-freedom typical section set-up with a nonlinear pitching spring and immersed in a linear unsteady potential flow. Different types of spring nonlinearities were considered such as a free-play, hysteresis and cubic stiffness, yielding unstable flutter responses below the linear threshold. A vast amount of literature followed that initial impulse [B. H. K. Lee et al. 1999a], confirming the possibility of subcritical flutter responses due to structural nonlinearities. More realistic structure models were then considered in an attempt to explain experimental wind-tunnel or flight-test observations. The investigation of nonlinear beam models [M. J. Patil et al. 2004] was motivated by the study of high-aspect ratio wings and their large static deflections. Conversely, nonlinear plate models were used for low-aspect ratio wings [D. Tang et al. 1999]. The aforementioned studies — and most of the reference




FIGURE 4: Two typical bifurcation scenarios for flutter: supercritical (blue) and subcritical (red) Hopf bifurcation. Dashed lines materialize unstable solutions. U_c is the (linear) critical velocity.

cited therein — focus on modelling structural (due to nonlinear stress-strain dependencies) and geometric (due to large displacements) nonlinearities, while retaining linear potential flow models like the Theodorsen model [Yang et al. 1988; B. H. K. Lee et al. 1999b] or the doublet-lattice method [M. J. Patil et al. 2004; D. Tang et al. 1999].

Mostly two types of aerodynamic nonlinearities were considered. First, stall-induced nonlinearities have been investigated in relation to high-aspect ratio wings and helicopter blades, that operate at large angles of attack. Semi-empirical fluid models were dominantly used, like static [M. Patil et al. 2001] or dynamic [D. M. Tang et al. 1992; D. M. Tang et al. 2004; Stanford et al. 2013] stall models. Second, transonic nonlinearities have received growing attention since the beginning of the 2000s, in particular thanks to the development of efficient CFD tools [Thomas et al. 2002; Schewe et al. 2003; Kholodar et al. 2004; Bendiksen 2011]. Several parametric studies have been performed, typically regarding the effect of the Mach number, the position of the elastic axis and the structural frequency ratio [Van Rooij et al. 2017b].

In parallel to these application-oriented studies, a series of fundamental experimental works have focused on the occurrence of coupled-mode flutter in moderate-Reynolds ($Re \sim 10^4 - 10^5$) flows, around zero [Dimitriadis et al. 2009; Amandolese et al. 2013] or nonzero [Razak et al. 2011] static angle of attack. These studies — in particular with zero static angle of attack — reported the existence of subcritical high-amplitude periodic solutions involving dynamic stall [McCroskey 1982; Šidlof et al. 2016]. The occurrence of such non-divergent oscillations despite the use of mostly linear structures showed that subsonic flow nonlinearities alone are able to saturate the exponential growth of the flutter instability. Due to the observed importance of the pitching motion in these oscillating regimes, some experimental [Onoue et al. 2015; Zhu et al. 2020] and numerical [Amiralaei et al. 2010; Menon et al. 2019] studies have been dedicated to pitching-only airfoils, with similar observations of subcritical high-amplitude periodic solutions. Recently, the transitional Reynolds number regime has been the focus of several works [Poirel et al. 2008; Poirel et al. 2010; Yuan et al. 2013; Barnes et al. 2018; Negi et al. 2018] showing the highly-nonlinear behavior of the transitional unsteady boundary layers and its connection to the occurrence of amplitude-bounded aeroelastic oscillations, specific to the transitional regime.

To our knowledge however, no work has been dedicated to a fundamental understanding of the nonlinear dynamics of coupled-mode flutter at low-Reynolds numbers (fully laminar flows). Typically, no systematic investigation of the effects of key parameters, like the mass ratio or the Reynolds number, on the type of the flutter bifurcation has been reported. It is thus one of the objectives of this thesis to **propose a numerical exploration of the bifurcation following the coupled-mode flutter instability in low-Reynolds incompressible flows**, governed by the two-dimensional Navier–Stokes equations.

Methods for nonlinear flutter analysis

To achieve such a goal, the most straightforward approach is to march in time the system of coupled nonlinear equations that govern the structure and the fluid, *via* some classical time-integration scheme (Euler, Runge-Kutta, Backward Differences, ...). Open-source or commercial codes implementing these methods are widely available and well validated. If this approach has the undeniable advantage of generality, it also suffers from several theoretical and practical issues. On the theoretical side, time-marching approaches, because they mimic the sequential nature of time, are not able to remain on unstable orbits. In the context of bifurcation study, this is a significant shortcoming since unstable periodic solutions are involved as soon as the bifurcation is subcritical or, if it is supercritical, when secondary instabilities are triggered (see below). On the more practical side, bifurcation study is mostly interested in permanent regimes of oscillations (*e.g.* periodic or quasi-periodic solutions) that a time-marching approach typically visits only after (possibly long) transient regimes. Large amounts of unexploited data (the transients) are thus produced, leading to large computational costs. For these reasons, more specific approaches are required in order to successfully address the objective stated above. We introduce here two of them, namely weakly nonlinear analyses and harmonic balance methods, that have been used in this thesis.

Weakly nonlinear analyses A so-called *Weakly NonLinear analysis* (WNL) aims at deriving the normal form associated to a bifurcation [Nayfeh 2011], which is a simple low-dimensional nonlinear equation that represents the behavior of the full system in the vicinity of the bifurcation point. Normal forms have been obtained for different aeroelastic instabilities like coupled-mode flutter of a typical section with linear aerodynamics [Sedaghat et al. 2000] or a galloping oscillator [Vio et al. 2005]. Nonlinear beams were considered in [Nayfeh et al. 2012] and a dynamic stall model in [Stanford et al. 2013]. In parallel, weakly nonlinear analyses have also been used to compute normal forms of different hydrodynamic bifurcations occurring in laminar flows [Sipp et al. 2007; Meliga et al. 2009; Meliga et al. 2012]. These studies used the framework of global stability where the flow model directly consists in a spatial discretization of the incompressible Navier–Stokes equations. These analyses were then extended to fluid-structure problems like a spring-mounted cylinder subjected to VIV's [Meliga et al. 2011] or freely falling objects [Tchoufag et al. 2015]. Towards more industry-oriented applications, weakly nonlinear models, based on Center Manifold Reduction, have been successfully developed for the transonic flutter of different wings in [Woodgate et al. 2007]. In this thesis, we will use an approach similar to Tchoufag et al. (2015) and adapt it to the particular case of the coupled-mode flutter of a typical section.

Harmonic Balance Methods Though very well-suited for the study of a bifurcation in the close vicinity of the instability threshold, weakly nonlinear approaches rapidly yield erroneous predictions as the bifurcation parameter deviates from the threshold [Gallaire et al. 2016]. More nonlinear methods are then required.

As a trade-off between weakly nonlinear analysis and fully nonlinear transient time-marching, one can choose an alternative approach based on the prior knowledge that only periodic solutions are sought close enough to the bifurcation point. In this case, the solution can be searched under the form of a Fourier series, which leads to the concept of *Harmonic Balance Methods* (HBM). These methods rely on transforming a time-dependent problem into a system of stationary equations for the Fourier coefficients of the unknown solution. This approach is often traced back to the work of Krylov et al. (1949) where it was used as an analytical tool to compute periodic solutions to nonlinear ordinary differential equations. Recently, it was re-popularized as a time-discretization scheme for CFD-based aeroelastic computations [Hall et al. 2002; Gopinath et al. 2005; McMullen et al. 2006]. In this thesis, we will use the particular variant by Gopinath et al. (2005) called the *Time Spectral Method* (TSM).

Behind the above mentioned advantages of HBM/TSM approaches hides a darker side: the increasing complexity of numerically solving the corresponding system for large frequencies or numbers of harmonics. Typically, the solution method of the first works on TSM was to march the equations in artificial pseudo-time, using explicit schemes. However, Van der Weide et al. (2005) showed that the associated CFL restriction depends on the number of harmonics such that the allowed timestep drastically reduces — and the convergence to the periodic orbit is increasingly difficult — as more harmonics are taken into account. The same reasoning is true for large oscillation frequencies. Several improvements of the solution method were achieved using implicit time-marching [Sicot et al. 2008] or Newton–Krylov solvers [Mundis et al. 2014]. Still, the issues associated to large numbers of harmonics were displaced to the linear solver and to the Krylov method’s preconditioner, respectively. It is only with the works of Mundis et al. (2015) and Mundis et al. (2017) that solution methods with iteration counts independent of the number of harmonics were reported. If the Newton–Krylov approach seems to be the most promising, the question of finding efficient preconditioners for the linearized TSM equations, *i.e.* that are robust in particular to the number of harmonics, is an active area of research. An objective of this thesis is thus to contribute to this effort by **proposing an efficient solution method for the TSM equations based on a Newton–Krylov approach and on an adequate preconditioning of the linearized system.**

Secondary flutter bifurcations

Just like steady solutions may destabilize through (primary) instabilities, periodic solutions may also destabilize through (secondary) instabilities. Secondary instabilities of periodic flutter solutions are rarely studied in the literature. The main reason for that is to consider that secondary bifurcations are irrelevant from the engineer’s point of view, because an airplane should never cross — nor even get too close to the critical threshold — thus making pointless the study of the solutions that appear above it. To that argument, one may oppose two. The first is given in by E. Dowell et al. (2003) in the conclusion to his review on nonlinear aeroelasticity: “because of nonlinear aeroelastic effects, finite amplitude oscillations can in some cases replace what would otherwise be the rapidly growing and destructive oscillations

of classical flutter behavior. A careful consideration and design of favorable nonlinearities offers a new opportunity for improved performance and safety of valuable wind-tunnel models, flight vehicles, and their operators and passengers.” In other words, a detailed knowledge of the finite-amplitude oscillating solutions that occur beyond the flutter linear threshold is a pre-requisite for the assessment of the dangers associated to flying in these regimes, or if needed, to design adequate control strategies [Livne 2018]. The second reason is a very practical consequence of the first. Indeed, the current trend to use harmonic balance approaches for performing nonlinear aeroelastic computations raises the question of the “observability” (*i.e.* the stability) of the so-computed solutions. Since harmonic balance methods are designed to compute periodic solutions, independently of their stability, how can one assess if the computed solution is the observed (stable) behavior, in reality ?

For the flutter instability, some works on secondary instabilities of flutter LCO’s can be found, for a typical section mounted on nonlinear springs with cubic stiffness [L. Liu et al. 2004; B. H. Lee et al. 2005; G. Liu et al. 2018]. In all cases, a transition from periodic to quasi-periodic regimes was observed for high-enough velocity. The investigations of secondary instabilities with nonlinear fluid models are very scarce and, to our knowledge, always based on empirical stall models [X. Li et al. 1997]. In this thesis, we thus propose to follow a complementary path and **investigate the stability of the nonlinear periodic flutter solutions of a structurally linear typical section, in a low-Reynolds flow modelled by the Navier–Stokes equations.**

The adequate mathematical framework for assessing the stability of periodic solutions is found in the theoretical work by Floquet (1883) on linear differential equations with periodic coefficients. Classically, the so-called Floquet stability of a periodic solution is assessed by analyzing the eigenvalues of the monodromy matrix (*e.g.* [Peletan et al. 2013] for a review). On the other hand, in an harmonic balance framework, the same analysis may be performed by scrutinizing the eigenvalues of the harmonic balance equations (equivalently the TSM equations), linearized around the periodic orbit under examination. In the logical continuation of the efforts announced above relatively to the development of efficient TSM solution methods, we will adopt this latter approach.

The computational burden of 3D flows

From the computation of steady solutions and the analysis of their linear stability to the TSM-based computation of periodic orbits and their Floquet stability, all numerical methods evoked in the previous paragraphs heavily rely on solving the linearized coupled fluid-structure equations, around some given solution, and for some prescribed right-hand side. For most part of this thesis, we will use as a model problem a typical aeroelastic section immersed in a two-dimensional Navier–Stokes flow. Using modern computers, the linear systems arising from the spatial discretization of the two-dimensional Navier–Stokes equations can be easily handled through direct sparse solvers (*e.g.* MUMPS [Amestoy et al. 2001]), based on the lower-upper (LU) factorization of the system’s matrix. These solvers have the double advantage of being very robust and mostly operated as black-boxes from the point of view of a non-specialist user. The counterpart of that robustness lies in the rapid growth of memory and cpu-time costs, associated with constructing and storing the LU factorization, as the number of degrees of freedom rises. As a consequence, this approach quickly becomes out-of-reach for three-dimensional configurations where the number of degrees of freedom easily exceeds several tens of millions.

For such cases, a widely used approach in hydrodynamic stability studies is to resort to the so-called *matrix-free* (or *time-stepping*) approach [Bagheri et al. 2009a; Loiseau et al. 2019]. It consists in re-using subroutines that are usually available in CFD solvers (time-steppers) in order to compute steady-states (or periodic solutions) and then assess their stability. Taking the example of the stability analysis of steady solutions, the latter are first computed by time-marching the nonlinear equations, either with the classical CFD solver if the searched solution is stable or with a stabilized version of it if the searched solution is unstable (e.g. [Shroff et al. 1993; Åkervik et al. 2006]). Second, the eigenvalue of the Navier–Stokes Jacobian matrix are deduced from the leading eigenvalues of an exponential-based transformation of it, the action of which is relatively easily obtained from existing subroutines⁷. An advantage of that approach is that CFD solvers are often highly optimized for large-scale parallel computations and thus offer very efficiently implemented subroutines. Two downsides however may be pointed out. First, stabilization techniques are not always trivial to use as they sometimes require prior knowledge of the searched solution [Åkervik et al. 2006], nor do they typically yield fast (quadratic) convergence to the solution, as a Newton method. Second, the exponential-based transformation implies a trade-off in the value of the timestep: small timesteps yield accurate application of the exponential transformation (and thus accurate eigenvalues), but at the cost of a slow convergence towards the eigenvalues [Tuckerman et al. 2000b]. The opposite is true for large timesteps.

In order to circumvent these drawbacks, one may seek to extend the methods that are used efficiently for two-dimensional configurations — *i.e.* Newton methods for the steady-state computation and shift-invert spectral transformation for the eigenvalue computations — to three-dimensional configurations. Hence the necessity to efficiently solve high-dimensional linear problems, involving the Navier–Stokes Jacobian matrix. To handle such high-dimensional linear problems, iterative methods are the method of choice [Saad 2003]. In particular, Krylov subspace methods are today recognized as the most efficient ones for poorly conditioned, non-normal systems. These methods rely on projecting the original high-dimensional linear system onto a low-dimensional subspace — the so-called Krylov subspace — where the system can be solved with basic linear solvers for dense matrices. The most emblematic algorithm from this class of methods is the Generalized Minimal Residual algorithm (GMRES) [Saad et al. 1986] that applies to generic non-Hermitian matrices. Some of the most appealing properties of Krylov methods lie in their superior robustness with respect to classical, so-called stationary, iterative methods (Jacobi, Gauss–Seidel, SOR, etc). For example, in exact arithmetic, GMRES is known to possess a *finite termination property*⁸ and a *non-increasing residual property*⁹. Those are clear progresses with respect to stationary methods that typically need diagonal dominance to ensure convergence. However, for large, poorly conditioned, nonnormal matrices — like the Jacobian matrix of the linearized Navier–Stokes equations —, the convergence behavior of Krylov methods can be very slow. To remedy that, preconditioning techniques must be used in order to reduce the condition number of the system, that is known to (partially, for nonnormal matrices) determine the convergence speed of the method [Liesen et al. 2004].

The research of efficient preconditioners for the *incompressible* linearized Navier–Stokes is an active field of applied mathematics research [Segal et al. 2010]. One of its particularities lies in the necessity to handle the saddle-point structure [Benzi et

⁷Precisely, from a subroutines that advances in time the *linearized* equations, of one timestep

⁸The exact solution is obtained in at most a number of iterations equal to the size of the system

⁹The residual is a decreasing function of the iteration number

al. 2005] due to the incompressibility constraint. Significant progresses were made through a series of works consecutive to the proposition of Benzi et al. (2006) to use augmented Lagrangian techniques in order to improve the problematic approximation of the pressure Schur complement. The application of this technique to solve the linearized Navier–Stokes problem for hydrodynamic stability purposes was first proposed by M. A. Olshanskii et al. (2008). The authors performed a theoretical study that they successfully verified on some two-dimensional numerical examples. However, to our knowledge, the practical use of augmented Lagrangian preconditioners for hydrodynamic stability eigenvalue computations was never attempted. Our objective is thus to first **assess the practical efficiency of the modified Augmented Lagrangian preconditioner [Benzi et al. 2011b] for performing complete linear stability analyses (i.e. for steady-state and eigenvalue computations) of incompressible laminar flows**. In a second step, we aim at proposing a parallel implementation of this approach, able to **handle stability analyses of large-scale three-dimensional flows on high-performance parallel computers**.

Organization of the manuscript

The manuscript organizes in three parts, that detail as follows:

1. The first part is devoted to the presentation of the different linear and nonlinear methods that will be used subsequently for analyzing the flutter bifurcation of the aeroelastic section.
 - Chapter 1 introduces in details the typical section model used in this thesis. The dimensional and non-dimensional parameters are defined as well as the governing equations. Two variants of the Arbitrary Lagrangian Eulerian formalism, used to handle the moving fluid domain, are introduced. The technical details on the finite element spatial discretization of the Navier–Stokes equations are provided. Finally, a classical time-marching approach for integrating the fluid-structure equations is presented, as it will serve as a reference across the manuscript. This chapter is mostly meant to offload the rest of the manuscript from the cumbersome technical details inherent to fluid-structure interaction.
 - In chapter 2, we present in extensive details the weakly nonlinear approach for the typical aeroelastic section. We also introduce a Hessian-based mesh adaptation framework that is used in the subsequent sections to efficiently explore the parameter space. The chapter ends with a validation of the weakly nonlinear results against reference fully nonlinear time-marching solutions.
 - Chapter 3 is dedicated to the Time Spectral Method. As a (long) preamble, TSM is replaced in the more global zoology of harmonic balance methods. A classification borrowed from spectral methods in space is used to organize the discussion. The core of the chapter presents the Newton—Krylov solution method and introduces the so-called block-circulant preconditioner that is at the core of the efficiency of the method when large number of harmonics are used. A parallelization in time is also presented. Several numerical experiments are reported to assess the robustness of the preconditioner. The method is applied first on cases where the solution frequency is known (forced oscillations) and then extended to the case of unknown frequency (self-sustained oscillations).

2. The second part of the thesis is dedicated to the physical analysis of the bifurcation associated to the coupled-mode flutter instability of the aeroelastic section in laminar flows:
 - We start our explorations in chapter 4 by revisiting the linear stability analysis of the typical section in low-Reynolds number flows. In addition to coupled-mode flutter, divergence and vortex-induced vibrations are also obtained. The thresholds for these different instabilities are computed while varying the solid-to-fluid mass ratio, thus allowing to build neutral curves for each type of instability. The effect of the Reynolds number is explored and comparisons are proposed with classical low-dimensional flow models such as quasi-steady models or the Theodorsen model.
 - Chapter 5 focuses on exploring the effects of low-Reynolds number fluid nonlinearities on the type of flutter bifurcation (subcritical or supercritical). The weakly nonlinear analysis is first used to explore parametrically the type of bifurcation for different solid-to-fluid mass ratios and Reynolds numbers. A decomposition of the cubic coefficient of the normal form of the flutter bifurcation is introduced so as to identify the contributions of the different nonlinearities and spatial locations with respect to the type of the bifurcation. Then, the highly nonlinear regime is explored by using results from both the Time-Spectral Method and reference time-marching solutions. For certain parameters, a unusual bifurcation scenario is obtained and is discussed with respect to unexplained experimental results in the literature.
 - We conclude our exploration of the flutter bifurcation in chapter 6 by investigating the solutions obtained up to 15% above the critical velocity, in a supercritical case. The bifurcation is first explored with reference time-marching solutions. Then a Floquet stability analysis is performed in order to explain the time-marching results. To that end, an original method based on the analysis of the spectrum of the linearized TSM operator is presented. The results from the Floquet analysis are analyzed and compared to the time-marching solutions, allowing to highlight the mechanism responsible for the observed time-marching results.
3. The third part of this thesis aims at initiating the extension of the solution methods used in the previous chapters for two-dimensional flows, to the three-dimensional case. It is composed of only one chapter.
 - Chapter 7 focuses on performing large-scale purely hydrodynamic (the structure is fixed) linear stability analysis of incompressible flows. To this end, the methods used in the previous chapters for two-dimensional configurations are extended. The core issue of the linear solver for high-dimensional configurations is tackled by using a Krylov subspace method, preconditioned by the modified Augmented Lagrangian preconditioner. The performance of the preconditioner is first assessed on a two-dimensional test case and benchmarked with respect to other state-of-the-art alternatives. Then, a parallel implementation using the FreeFEM language and its PETSc/SLEPc interface is described and tested on a large-scale three-dimensional case.

The results are finally summarized in the general conclusion, chapter 8, leading to several propositions for future directions of research.

Communications

Here are listed the different past and foreseen communications this thesis has lead to.

Papers

1. J.Moulin, P.Jolivet, O.Marquet, *Augmented Lagrangian Preconditioner for Large-Scale Hydrodynamic Stability Analysis*, published in *Computer Methods in Applied Mechanics and Engineering*, vol. 351, pp. 718 - 743. Code available at <https://github.com/prj-/moulin2019a1>.
2. J.Moulin, O.Marquet, *Flow-induced instabilities of springs-mounted plates in viscous flows: a global stability approach*, in preparation for *Physics of Fluids*.
3. J.Moulin, O.Marquet, *Hard and soft flutter of thin plates in laminar incompressible flows: weakly and strongly nonlinear analyses*, in preparation for *Journal of Fluid Mechanics*.
4. J.Moulin, O.Marquet, *Low-frequency amplitude modulation of flutter oscillations: a Floquet analysis based on the Time Spectral Method*, in preparation for *Journal of Fluid Mechanics*.

Conferences

1. J. Moulin, J-L. Pfister, M.Carini, O. Marquet. *Assessment of Fictitious Domain Method for Linear Stability Analysis of Fluid-Structure Systems*, ERCOFTAC-SIG33, Sienne, 19-21 juin 2017
2. J. Moulin, O. Marquet, D. Sipp. *Prediction of Flutter Instability in Turbulent Flow Based on Linear Stability Analysis*, BIFD, Houston, 11-14 juillet 2017
3. J. Moulin, J-L. Pfister, O. Marquet, P. Jolivet. *Augmented Lagrangian Preconditioner for Linear Stability Analysis of Incompressible Fluid Flows on Large Configurations*, FreeFem++Days, Paris, 14-15 décembre 2017
4. J.Moulin, O.Marquet, *Effect of Fluid Nonlinearities on Flutter Instability in Laminar Flows*, EFMC12, Vienne, 8-12 septembre 2018
5. J.Moulin, O.Marquet, P.Jolivet, *Augmented Lagrangian Approach for Large-Scale Linear Stability Analysis of Incompressible Flows and its Extension to Fluid-Structure Interaction*, FreeFem++Days, Paris, 13 décembre 2018

I

THEORETICAL AND NUMERICAL METHODS FOR NONLINEAR FLUTTER ANALYSIS

1

MATHEMATICAL MODELS FOR A “TYPICAL SECTION” IN VISCOUS FLOWS

This chapter introduces in details the typical section model used in this thesis. The dimensional and non-dimensional parameters are defined as well as the governing equations. Two variants of the Arbitrary Lagrangian Eulerian formalism, used to handle the moving fluid domain, are introduced. The technical details on the finite element spatial discretization of the Navier–Stokes equations are provided. Finally, a classical time-marching approach for integrating the fluid-structure equations is presented, as it will serve as a reference across the manuscript. This chapter is mostly meant to offload the rest of the manuscript from the cumbersome technical details inherent to fluid-structure interaction.

1.1 Equations of motion

The highly complex problem of full three dimensional wing flutter is modeled through a “typical airfoil section” [E. H. Dowell et al. 1989, §3.2.1]. It consists in a two-dimensional section (fig. 1.1) that translates along the negative horizontal axis $-\mathbf{e}_x$ at the constant airplane cruising speed U_∞ . Due to the wing flexibility the section also moves along the vertical direction \mathbf{e}_y (flexion mode) and rotates around the elastic center (torsion mode), noted A . The movement along \mathbf{e}_y is measured by the *heaving displacement*, noted h , whereas the rotation is measured by the *pitching angle*, noted θ . Note that h is counted positively upwards and θ is counted positively nose-down¹. Thus the classical definition of the angle of attack (positive nose-up) is simply $-\theta$.

Across this manuscript, the particular section studied is a rigid plate of thickness e and chord c , with a chord-to-thickness aspect ratio equal to $c/e = 20$. The plate mass (per unit length, as all subsequent constants) is noted m and its moment of inertia around the elastic center I_{ea} . The plate is mounted on heaving and pitching springs that model the wing flexibility. Both are attached at the elastic axis A , with respective stiffnesses K_h, K_θ . The center of gravity G is located at a distance x_{cg} from the elastic axis (with $x_{cg} > 0$ when G is downstream of A). The structural damping coefficients are noted C_h and C_θ . The plate moves in a two-dimensional viscous fluid of density ρ_f and dynamic viscosity μ .

In the following, all equations are nondimensionalized with respect to the characteristic length c , mass $\rho_f c^3$ and time c/U_∞ . This yields a set of eight nondimensional parameters (in addition to the chord-to-thickness ratio) governing the coupled problem. First, a list of purely solid parameters characterize the plate behavior: the heaving-to-pitching frequency ratio Ω , the heaving and pitching damping ratios ζ_h and ζ_p , the radius of gyration r_θ and the non-dimensional position of the center of gravity is noted x_θ . Their precise definitions are given in table 1.1 together with their values, that remain fixed throughout the manuscript. The reduced velocity

$$U^* = \frac{U_\infty}{c \sqrt{K_\theta/I_{ea}}}$$

characterizes the ratio of the horizontal plate velocity U_∞ to a velocity typical of the plate pitching mode, which has a natural frequency $\sqrt{K_\theta/I_{ea}}$. Alternatively it can be seen as the inverse of the non-dimensional pitching natural frequency. On the fluid side, the Reynolds number

$$\mathcal{R}e = \frac{\rho_f U_\infty c}{\mu}$$

characterizes the relative magnitude of the fluid inertial and viscous forces. Finally, the level of fluid-structure coupling is quantified by the the mass ratio²

$$\tilde{m} = \frac{m}{1/2 \rho_f c^2}$$

that compares the typical solid and fluid mass. The strongest coupling usually arise when $\tilde{m} \simeq 1$, *i.e.* when the solid is light enough so that fluid inertial can significantly

¹The usual aeroelastic convention is to take θ positive nose-up, to remain consistent with the classical aerodynamic definition of the angle of attack. However, the vertical displacement then must be taken positive downwards, in order to remain in a correctly positively oriented system of axes.

²On a side note, let us mention that \tilde{m} was defined as *twice* the non-dimensional solid mass in order to avoid an annoying factor $1/2$ in the right-hand side of the non-dimensional solid equations, eq. (1.3)

perturb it. On the contrary, when $\tilde{m} \gg 1$ the fluid momentum is negligible and couplings are weaker.

All these nondimensional numbers are summarized in table 1.1, along with typical values used in this work.




FIGURE 1.1: “Typical airfoil section” model.

	U^*	Ω	ζ_h	ζ_p
definition	$\frac{U_\infty}{c \sqrt{K_\theta/I_{ea}}}$	$\sqrt{\frac{K_h/m}{K_\theta/I_{ea}}}$	$\frac{C_h}{2\sqrt{mK_h}}$	$\frac{C_\theta}{2\sqrt{I_{ea}K_\theta}}$
typical values	$10^{-2} - 10^1$	0.8	0	0.05
	r_θ	x_θ	\tilde{m}	$\mathcal{R}e$
definition	$\sqrt{\frac{I_{ea}}{mc^2}}$	$\frac{x_{cg}}{c}$	$\frac{m}{1/2\rho_f c^2}$	$\frac{\rho_f U_\infty c}{\mu}$
typical values	0.290	0	$10^1 - 10^4$	$10^1 - 10^4$

TABLE 1.1: Non-dimensional parameters governing the flutter instability

1.1.1 Solid model: spring-mounted rigid solid

Let us start by describing the kinematics of the plate motion (*cf* fig. 1.2). Considering a point \mathbf{x} of the plate at the initial time $t = 0$. At time t , point \mathbf{x} has moved to a new position noted $\mathbf{x}_t(\mathbf{x}, t)$ such that

$$\mathbf{x}_t(\mathbf{x}, t) = \mathbf{x} + \boldsymbol{\zeta}_s(\mathbf{x}, t) \quad \text{with} \quad \boldsymbol{\zeta}_s(\mathbf{x}, t) = -1 t \mathbf{e}_x + h(t) \mathbf{e}_y + (\mathbf{R}(\theta(t)) - \mathbf{I})\mathbf{x} \quad (1.1)$$

and $\mathbf{R}(\theta)$ the rotation matrix of angle θ :

$$\mathbf{R}(\theta) = \begin{pmatrix} \cos(\theta) & -\sin(\theta) \\ \sin(\theta) & \cos(\theta) \end{pmatrix}$$

The displacement field $\boldsymbol{\zeta}_s(\mathbf{x}, t)$ is composed of three components respectively corresponding to the horizontal constant (unit, in non-dimensional form) velocity motion,

the heaving motion and the pitching motion. The corresponding velocity field is:

$$\mathbf{u}_s(\mathbf{x}, t) = \frac{\partial \boldsymbol{\zeta}_s}{\partial t} = -1 \mathbf{e}_x + \dot{h} \mathbf{e}_y + \dot{\theta} \frac{\partial \mathbf{R}}{\partial \theta} \mathbf{x}$$

where the dot symbol is the classical representation of the time-derivative.

The equations governing the motion of the plate, submitted to the elastic restoring forces, structural damping and fluid forces, are obtained using the Lagrange equations. The corresponding derivations are classical and the reader is reported to [Bisplinghoff et al. 1955, §3.8] (linearized version) or [Malher 2016, §1.1.2] (fully nonlinear version) for more details. After non-dimensionalizing and introducing the parameters of table 1.1, we obtain:

$$\begin{aligned} \frac{\partial^2 h}{\partial t^2} + x_\theta \left(\cos \theta \frac{\partial^2 \theta}{\partial t^2} - \sin \theta \left(\frac{\partial \theta}{\partial t} \right)^2 \right) + 2\zeta_h \left(\frac{\Omega}{U^*} \right) \frac{\partial h}{\partial t} + \left(\frac{\Omega}{U^*} \right)^2 h &= \frac{1}{\tilde{m}} \mathcal{C}_L \\ \frac{\partial^2 \theta}{\partial t^2} + \frac{x_\theta}{r_\theta^2} \cos \theta \frac{\partial^2 h}{\partial t^2} + 2\zeta_p \left(\frac{1}{U^*} \right) \frac{\partial \theta}{\partial t} + \left(\frac{1}{U^*} \right)^2 \theta &= \frac{1}{\tilde{m} r_\theta^2} \mathcal{C}_M \end{aligned} \quad (1.2)$$

where \mathcal{C}_L and \mathcal{C}_M are the lift and moment coefficients, defined as:

$$\mathcal{C}_L = \frac{L^\diamond}{1/2\rho_f U_\infty^2 c}, \quad \mathcal{C}_M = \frac{M_z^\diamond}{1/2\rho_f U_\infty^2 c^2}$$

with L^\diamond and M_z^\diamond the (dimensional) lift and moment (about the elastic center) of the fluid forces, per unit span. Note that eq. (1.2) naturally couples heaving and pitching due to nonlinear inertial terms. In most classical texts, those terms are linearized for flutter analysis [Bisplinghoff et al. 1955; E. H. Dowell et al. 1989; Hodges et al. 2011; Dimitriadis 2017]. If this is a natural approximation for studying linear flutter dynamics, it is not the case anymore for nonlinear flutter investigations.

In the rest of this manuscript however, the nonlinear inertial terms do vanish because we will only consider plates pitching around their center of gravity, *i.e.* $x_\theta = 0$, yielding structurally³ decoupled heaving and pitching equations:

$$\frac{\partial^2 h}{\partial t^2} + 2\zeta_h \left(\frac{\Omega}{U^*} \right) \frac{\partial h}{\partial t} + \left(\frac{\Omega}{U^*} \right)^2 h = \frac{1}{\tilde{m}} \mathcal{C}_L \quad (1.3a)$$

$$\frac{\partial^2 \theta}{\partial t^2} + 2\zeta_p \left(\frac{1}{U^*} \right) \frac{\partial \theta}{\partial t} + \left(\frac{1}{U^*} \right)^2 \theta = \frac{1}{\tilde{m} r_\theta^2} \mathcal{C}_M \quad (1.3b)$$

Notice finally that the natural heaving and pitching frequencies can be identified as

$$\omega_{0h} = \frac{\Omega}{U^*} \quad \text{and} \quad \omega_{0\theta} = \frac{1}{U^*}$$

³We will see later that a coupling is re-introduced *via* the fluid forces in the right-hand sides




FIGURE 1.2: Kinematic description of the plate undergoing a constant unit velocity translation along $-\mathbf{e}_x$ and time-dependent heaving and pitching motions, as described by eq. (1.1).

1.1.2 Fluid model: incompressible Navier–Stokes

The fluid dynamics are governed by the incompressible two-dimensional Navier–Stokes equations, written in a moving domain $\Omega_{f,t}$:

$$\begin{aligned} \frac{\partial \mathbf{u}_t}{\partial t} \Big|_{\mathbf{x}_t} + \nabla \mathbf{u}_t \mathbf{u}_t - \nabla \cdot \left[-p_t \mathbf{I} + \mathcal{R}e^{-1} \left(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^T \right) \right] &= 0 & \mathbf{x}_t \in \Omega_{f,t} \\ -\nabla \cdot \mathbf{u}_t &= 0 & \mathbf{x}_t \in \Omega_{f,t} \end{aligned} \quad (1.4)$$

Having adopted the convention where the plate moves in a fluid at rest (see fig. 1.1), a homogeneous Dirichlet boundary condition is imposed on the boundary Γ_{in} . At the outflow boundary Γ_{out} , a natural “do-nothing” condition is imposed:

$$\mathbf{u}_t = 0 \quad \text{on } \Gamma_{\text{in}} \quad (1.5a)$$

$$\left[-p_t \mathbf{I} + \mathcal{R}e^{-1} \left(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^T \right) \right] \mathbf{n}_f = 0 \quad \text{on } \Gamma_{\text{out}} \quad (1.5b)$$

with \mathbf{n}_f the outward normal to the fluid domain boundaries. The condition imposed at the fluid-structure interface is detailed later, in section 1.1.3.

The formulation eq. (1.4) of the Navier–Stokes equations presents two shortcomings. First, the time derivative, taken at \mathbf{x}_t constant, is not well defined. Indeed, imagine having to compute the time derivative of the fluid velocity at a point \mathbf{x}_t , for example with finite differences:

$$\frac{\partial \mathbf{u}_t}{\partial t} \Big|_{\mathbf{x}_t} \simeq \frac{\mathbf{u}_t(\mathbf{x}_t, t + \Delta t) - \mathbf{u}_t(\mathbf{x}_t, t)}{\Delta t} \quad \Delta t \ll 1$$

The ill-posedness arises in the possible scenario where the location \mathbf{x}_t is close enough to the interface so that \mathbf{x}_t belongs to the fluid domain at instant t but is in the solid domain at instant $t + \Delta t$ (or oppositely). Thus, the time-derivative has to be handled differently for moving domains. The second shortcoming lies in the fact that the equations are defined on a moving domain $\Omega_{f,t}$, that is typically dependent on the solid variables. As a consequence, a proper linearization (*a fortiori* “weakly nonlinearization”, *cf.* chapter 2) of those equations would require non-trivial linearization with respect to the domain.

To address these shortcomings, two formulations are introduced and will be used in the next chapters. The first formulation, referred to as the “absolute velocity - rotating axis” formulation, is taken from [Mougin et al. 2002] where it was proposed

for studying falling rigid bodies. It will be used for the weakly nonlinear developments of chapter 2. The second formulation, labeled as the “reference configuration ALE” formulation, is adapted from J.-L. Pfister (2019) where it was considered to study the dynamics of flexible structures. It is used in this manuscript for the Harmonic Balance approach developed in chapter 3. In the two following paragraphs, we provide the detailed equations for the two formulations. In addition, we show in section 1.1.2.3 that the “absolute velocity - rotating axis” equations can be derived as a particular case of the more general “reference configuration ALE” formulation.

1.1.2.1 The “absolute velocity - rotating axis” formulation

In the “absolute velocity - rotating axis” formulation by Mougín et al. (2002), the position of any point M in the fluid domain is written in the system of coordinates $(A, \mathbf{e}_X, \mathbf{e}_Y)$ of which the origin is the elastic axis A and the axes are the principal axes of the plate (see fig. 1.3). More precisely, we note \mathbf{X} the position vector of a point M :

$$\mathbf{X} = [X, Y]^T = \overrightarrow{AM}_{(\mathbf{e}_X, \mathbf{e}_Y)}$$

where we have precised as an index the system of axis the coordinates X and Y refer to. In $(A, \mathbf{e}_X, \mathbf{e}_Y)$, the fluid-structure interface is fixed, and is noted Γ_{fs} , where the index reference to time t has been dropped. For unbounded flows, the fluid domain is thus independent of the solid movement, and is consecutively noted Ω_f .

Following [Mougín et al. 2002], the Navier–Stokes equations for the *absolute* velocity $\mathbf{U}(\mathbf{X}, t) = [U, V]^T$ written in the rotating axes and the pressure $p(\mathbf{X}, t)$, read:

$$\frac{\partial \mathbf{U}}{\partial t} + \dot{\theta} \mathbf{e}_z \wedge \mathbf{U} + \nabla \mathbf{U} (\mathbf{U} - \mathbf{W}) - \nabla \cdot \sigma_f(p, \mathbf{U}) = 0 \quad (1.6a)$$

$$\nabla \cdot \mathbf{U} = 0 \quad (1.6b)$$

with the fluid stress tensor, $\sigma_f(p, \mathbf{U}) = -p\mathbf{I} + \mathcal{R}e^{-1} (\nabla \mathbf{U} + \nabla \mathbf{U}^T)$, and the “rigid-body velocity” (or “extension velocity”⁴), also written in the rotating axes:

$$\mathbf{W}(\dot{h}, \theta, \dot{\theta}) = \mathbf{R}(\theta)^T (-1, \dot{h})^T + \dot{\theta} \mathbf{e}_z \wedge \mathbf{X} \quad (1.7)$$

The rigid-body velocity is composed of two terms: the first term accounts for translation whereas the second accounts for rotation. Note in particular that the rotation matrix $\mathbf{R}(\theta)$ is used in order to project the translation velocity $-1\mathbf{e}_x + \dot{h}\mathbf{e}_y$ onto the rotating axes.

⁴This denomination will become clearer in ???




FIGURE 1.3: Illustration of several quantities used in the “absolute velocity - rotating axis” Navier–Stokes equations. In particular, are shown in red the vectors that are written in the rotating system of axes ($A, \mathbf{e}_X, \mathbf{e}_Y$).

1.1.2.2 The “reference configuration ALE” formulation

In the Arbitrary Lagrangian Eulerian (ALE) framework [Donea et al. 2004] a *fixed* reference domain, noted Ω_f , is introduced in addition to the physical domain $\Omega_{f,t}$. A one-to-one mapping $\mathcal{A}(\mathbf{x}, t)$ links any position \mathbf{x} in Ω_f to the corresponding physical position \mathbf{x}_t at instant t in $\Omega_{f,t}$.

$$\mathbf{x}_t = \mathcal{A}(\mathbf{x}, t) \quad \mathbf{x}_t \in \Omega_{f,t}, \quad \mathbf{x} \in \Omega_f$$

In addition, the so-called *extension displacement field* $\boldsymbol{\zeta}_e$ is defined as:

$$\boldsymbol{\zeta}_e(\mathbf{x}, t) = \mathbf{x}_t - \mathbf{x} = \mathcal{A}(\mathbf{x}, t) - \mathbf{x}$$

In this manuscript, the reference domain is chosen to be the fluid domain in the initial configuration, which typically corresponds to $h = \theta = 0$. The mapping is arbitrary as long as it conforms to the solid displacement $\boldsymbol{\zeta}_s(\mathbf{x}, t)$ on the (reference configuration) interface Γ_{fs} , *i.e.* $\boldsymbol{\zeta}_e(\mathbf{x}, t) = \boldsymbol{\zeta}_s(\mathbf{x}, t)$ on Γ_{fs} .

Using the chain-rule, the time-derivative at constant \mathbf{x}_t in eq. (1.4) can be rewritten at constant \mathbf{x} :

$$\left. \frac{\partial \mathbf{u}_t}{\partial t} \right|_{\mathbf{x}_t} = \left. \frac{\partial \mathbf{u}_t}{\partial t} \right|_{\mathbf{x}} - (\nabla \mathbf{u}_t) \mathbf{w}_t \quad (1.8)$$

where the *extension velocity* \mathbf{w}_t , defined on the deformed domain, is introduced:

$$\mathbf{w}_t = \left. \frac{\partial \mathcal{A}}{\partial t} \right|_{\mathbf{x}} \circ \mathcal{A}^{-1} \quad (1.9)$$

Using eq. (1.8) in eq. (1.4) yields the classical ALE form of the Navier–Stokes equations:

$$\begin{aligned} \left. \frac{\partial \mathbf{u}_t}{\partial t} \right|_{\mathbf{x}} + \nabla \mathbf{u}_t (\mathbf{u}_t - \mathbf{w}_t) - \nabla \cdot \left[-p_t \mathbf{I} + \mathcal{R} e^{-1} \left(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^T \right) \right] &= 0 & \mathbf{x}_t \in \Omega_{f,t} \\ -\nabla \cdot \mathbf{u}_t &= 0 & \mathbf{x}_t \in \Omega_{f,t} \end{aligned} \quad (1.10)$$

A new convection term has appeared where the convection velocity is \mathbf{w}_t . If the solid is fixed, then $\mathbf{w}_t = 0$ and the classical Eulerian Navier–Stokes is retrieved. On the contrary, imagine the mapping \mathcal{A} exactly follows the motion of the fluid particles; then $\mathbf{w}_t = \mathbf{u}_t$ and a Lagrangian description of the Navier–Stokes equations

is retrieved. For an arbitrary mapping, the description lies *arbitrarily* between the Eulerian and Lagrangian representations.

As stated previously, we wish to obtain a formulation on a fixed domain. To that aim, alternative fluid unknowns \mathbf{u} and p , defined in Ω_f , are introduced such that

$$\mathbf{u}(\mathbf{x}, t) = \mathbf{u}_t(\mathbf{x}_t(\mathbf{x}, t), t) \quad \text{and} \quad p(\mathbf{x}, t) = p_t(\mathbf{x}_t(\mathbf{x}, t), t) \quad \mathbf{x} \in \Omega_f$$

Using the above, eq. (1.10) may be entirely written in Ω_f whereas the domain motion is handled by the extension displacement field (e.g. [J.-L. Pfister et al. 2019]):

$$J(\boldsymbol{\zeta}_e) \left. \frac{\partial \mathbf{u}}{\partial t} \right|_{\mathbf{x}} + \nabla \mathbf{u} \Phi(\boldsymbol{\zeta}_e) (\mathbf{u} - \mathbf{w}) - \nabla \cdot \boldsymbol{\Sigma}_f(\mathbf{u}, p, \boldsymbol{\zeta}_e) = 0 \quad \text{in } \Omega_f \quad (1.11a)$$

$$-\nabla \cdot (\Phi(\boldsymbol{\zeta}_e) \mathbf{u}) = 0 \quad \text{in } \Omega_f \quad (1.11b)$$

where the extension velocity in Ω_f is

$$\mathbf{w} = \left. \frac{\partial \mathcal{A}}{\partial t} \right|_{\mathbf{x}} = \left. \frac{\partial \boldsymbol{\zeta}_e}{\partial t} \right|_{\mathbf{x}} \quad (1.12)$$

The deformation gradient $\mathbf{F} = \mathbf{I} + \nabla \boldsymbol{\zeta}_e$, its determinant $J = \det \mathbf{F}$ and the deformation operator $\Phi = J\mathbf{F}^{-1}$ have been introduced. The fluid stress tensor, transported in the reference configuration writes $\boldsymbol{\Sigma}_f = [-p\mathbf{I} + \mathcal{R}e^{-1}J^{-1}(\nabla \mathbf{u} \Phi + \Phi^T \nabla \mathbf{u}^T)] \Phi^T$.

Remark. The passage from eq. (1.10) to eq. (1.11) involves several transport theorems for spatial differential operators. Those may be found for example in appendix C of [J.-L. Pfister 2019].

Choice of the extension In the general case of an elastic solid, one can typically define $\boldsymbol{\zeta}_e$ as the solution of a *pseudo-elasticity* problem. If only rigid body motion is considered, more practical analytical extensions may be easily defined. Here, we use an analytical extension borrowed from [Persson et al. 2009]. Starting from the rigid-body displacement field $\boldsymbol{\zeta}_s(\mathbf{x}, t)$ defined by eq. (1.1), we define a “blending extension” (fig. 1.4)

$$\boldsymbol{\zeta}_e(\mathbf{x}, t) = -1t \mathbf{e}_x + b(\mathbf{x} - \mathbf{x}_G) (h\mathbf{e}_y + (\mathbf{R}(\theta) - \mathbf{I}) \mathbf{x}) \quad (1.13)$$

with $b(\mathbf{x})$ the *blending function*:

$$b(\mathbf{x}) = \begin{cases} 1 & |\mathbf{x}| < r_1 \\ 1 - f\left(\frac{|\mathbf{x}| - r_1}{r_2 - r_1}\right) & r_1 < |\mathbf{x}| < r_2 \\ 0 & |\mathbf{x}| > r_2 \end{cases} \quad \text{with} \quad f(s) = 10s^3 - 15s^4 + 6s^5 \quad (1.14)$$

With this extension, a different treatment is applied to the horizontal movement on the one hand, and to the heaving and pitching motions on the other. For the horizontal motion with constant velocity $-1\mathbf{e}_x$, it is simply propagated to the whole mesh. On the contrary, for the heaving and pitching motions the plate kinematics is applied to the mesh only inside a disk of radius r_1 (inner circle in fig. 1.4). Far enough from the solid, *i.e.* beyond a radius $r_2 > r_1$, the mesh is kept fixed (outer circle in fig. 1.4). Between r_1 and r_2 the fifth order polynomial $f(s)$ is used to connect smoothly the two zones.




FIGURE 1.4: Reference configuration ALE formalism with blending extension. The reference domain is illustrated in (a), where the extension displacement field ζ_e is represented by orange arrows. In (b), we show the corresponding deformed domain. The dashed circles mark the radii r_1 and r_2 of the blending function definition. For simplicity, the horizontal displacement component in eq. (1.13) is set to zero.

1.1.2.3 From the “reference configuration ALE” to the “absolute velocity - rotating axis”

In this section we demonstrate how the “absolute velocity - rotating axis” formulation, eq. (1.6) can be derived from eq. (1.11), by adding two simple ingredients. First, a “rigid-body extension” displacement is substituted for the blending extension, and second (ii) eq. (1.11) are projected along the system of axes $(\mathbf{e}_x, \mathbf{e}_y)$ that rotate with the plate. These developments constitute an alternative to the original derivation of Mougín et al. (2002).

Rigid body extension The rigid extension is the simplest analytical extension possible for rigid solid in unbounded flows. It simply propagates the rigid-body motion of the solid eq. (1.1) to the whole mesh (see fig. 1.5):

$$\zeta_e(\mathbf{x}, t) = \zeta_s(\mathbf{x}, t) \quad (1.15)$$

Using eq. (1.15), we immediately have $J = 1$ and $\mathbf{F} = \mathbf{R}(\theta)$ such that eq. (1.11) writes:

$$\begin{aligned} \left. \frac{\partial \mathbf{u}}{\partial t} \right|_{\mathbf{x}} + \nabla \mathbf{u} \mathbf{R}^T (\mathbf{u} - \mathbf{w}) - \nabla \cdot \left[-p \mathbf{R} + \mathcal{R} e^{-1} (\nabla \mathbf{u} + \mathbf{R} \nabla \mathbf{u}^T \mathbf{R}) \right] &= 0 \quad \text{in } \Omega_f \\ -\nabla \cdot (\mathbf{R}^T \mathbf{u}) &= 0 \quad \text{in } \Omega_f \end{aligned}$$




FIGURE 1.5: Reference configuration ALE formalism with rigid-body extension. The reference domain is illustrated in (a), where the extension displacement field ζ_e is represented by orange arrows. In (b), we show the corresponding deformed domain. We mark in red some vectors that are written in the rotating system of axis $(\mathbf{e}_x, \mathbf{e}_y)$, so as to yield the “absolute velocity - rotating axis” equations, eq. (1.6). As in fig. 1.4, the horizontal displacement component in eq. (1.15) is set to zero for simplicity.

Projection along the rotating axes First, we consider the system of axes $(\mathbf{e}_x, \mathbf{e}_y)$ that rotate with the plate, as shown in red in fig. 1.5(a). We then introduce the projections $\mathbf{U} = \mathbf{R}(\theta)^T \mathbf{u}$ and $\mathbf{W} = \mathbf{R}(\theta)^T \mathbf{w}$ of \mathbf{u} and \mathbf{w} on these axes and left-multiply eq. (1.11a) by \mathbf{R}^T , yielding

$$\begin{aligned} \mathbf{R}^T \left. \frac{\partial \mathbf{u}}{\partial t} \right|_{\mathbf{x}} + \nabla \mathbf{U} (\mathbf{U} - \mathbf{W}) - \nabla \cdot \left[-p \mathbf{I} + \mathcal{R} e^{-1} (\nabla \mathbf{U} + \nabla \mathbf{U}^T) \right] &= 0 \quad \text{in } \Omega_f \\ -\nabla \cdot \mathbf{U} &= 0 \quad \text{in } \Omega_f \end{aligned}$$

Finally noticing that

$$\frac{\partial \mathbf{U}}{\partial t} = \frac{\partial (\mathbf{R}^T \mathbf{u})}{\partial t} = \dot{\theta} \frac{\partial \mathbf{R}^T}{\partial \theta} \mathbf{u} + \mathbf{R}^T \frac{\partial \mathbf{u}}{\partial t} = \underbrace{\dot{\theta} \frac{\partial \mathbf{R}^T}{\partial \theta} \mathbf{R} \mathbf{U}}_{= -\dot{\theta} \mathbf{e}_z \wedge \mathbf{U}} + \mathbf{R}^T \frac{\partial \mathbf{u}}{\partial t}$$

one obtains exactly the “absolute velocity - rotating axis” formulation eq. (1.6):

$$\left. \frac{\partial \mathbf{U}}{\partial t} \right|_{\mathbf{x}} + \dot{\theta} \mathbf{e}_z \wedge \mathbf{U} + \nabla \mathbf{U} (\mathbf{U} - \mathbf{W}) - \nabla \cdot \left[-p \mathbf{I} + \mathcal{R} e^{-1} (\nabla \mathbf{U} + \nabla \mathbf{U}^T) \right] = 0 \quad \text{in } \Omega_f \quad (1.16a)$$

$$-\nabla \cdot \mathbf{U} = 0 \quad \text{in } \Omega_f \quad (1.16b)$$

where the extension velocity projected in the rotating axes, $\mathbf{W} = \mathbf{R}(\theta)^T \mathbf{w}$, may be developed to retrieve eq. (1.7):

$$\begin{aligned} \mathbf{W} &= \mathbf{R}(\theta)^T \mathbf{w} \\ &= \mathbf{R}(\theta)^T \left(\begin{pmatrix} -1 \\ \dot{h} \end{pmatrix} + \dot{\theta} \frac{\partial \mathbf{R}}{\partial \theta} \mathbf{x} \right) \\ &= \mathbf{R}(\theta)^T (-1, \dot{h})^T + \dot{\theta} \mathbf{e}_z \wedge \mathbf{x}. \end{aligned} \quad (1.17)$$

1.1.3 Interface velocity continuity and interface forces

To complete the solid and fluid models presented in the previous sections, interface conditions must be provided. First, the continuity of velocities is enforced on the time-dependent fluid-structure interface.

$$\mathbf{u}_t(\mathbf{x}_t, t) - \mathbf{w}_t(\mathbf{x}_t, t) = 0 \quad \text{on } \Gamma_{fs,t} \quad (1.18)$$

Second, the detailed expression for the lift and moment coefficients in the right-hand side of eq. (1.3) can now be given as a function of the fluid variables:

$$\mathcal{C}_L = \left\{ 2 \int_{\Gamma_{fs,t}} \left[-p_t \mathbf{I} + \mathcal{R} e^{-1} \left(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^T \right) \right] \mathbf{n}_{s,t} d\mathbf{x}_t \right\} \cdot \mathbf{e}_y \quad (1.19a)$$

$$\mathcal{C}_M = \left\{ 2 \int_{\Gamma_{fs,t}} [\mathbf{x}_t - \mathbf{x}_t(A_t)] \wedge \left[-p_t \mathbf{I} + \mathcal{R} e^{-1} \left(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^T \right) \right] \mathbf{n}_{s,t} d\mathbf{x}_t \right\} \cdot \mathbf{e}_z \quad (1.19b)$$

where the moment is specifically computed at the elastic axis A_t . The factor 2 in the above expressions is due to the classical definition of the lift (resp. moment) coefficient which is *twice* the non-dimensional lift (resp. moment). Depending on the particular Navier–Stokes formulation (“absolute velocity - rotating axis” or “reference configuration ALE”), the interface conditions are slightly remodeled.

1.1.3.1 “Absolute velocity - rotating axis”

The velocity continuity, projected on the rotating axes, becomes:

$$\mathbf{U}(\mathbf{X}, t) - \mathbf{W}(\mathbf{X}, t) = 0 \quad \text{on } \Gamma_{fs} \quad (1.20)$$

whereas the lift and moment coefficients become:

$$\mathcal{C}_L = \left\{ 2 \int_{\Gamma_{fs}} \sigma_f(\mathbf{U}, p) \mathbf{N}_s d\mathbf{X} \right\} \cdot \underbrace{\mathbf{R}^T(\theta) \mathbf{e}_y}_{\text{projection of } \mathbf{e}_y \text{ on the rotating axis}} \quad (1.21a)$$

$$\mathcal{C}_M = \left\{ 2 \int_{\Gamma_{fs}} \mathbf{X} \wedge \sigma_f(\mathbf{U}, p) \mathbf{N}_s d\mathbf{X} \right\} \cdot \mathbf{e}_z \quad (1.21b)$$

where the fluid force has been appropriately projected back along the vertical direction in order to retrieve the lift coefficient.

1.1.3.2 “Reference configuration ALE”

The velocity continuity becomes:

$$\mathbf{u}(\mathbf{x}, t) - \mathbf{w}(\mathbf{x}, t) = 0 \quad \text{on } \Gamma_{fs} \quad (1.22)$$

whereas the lift and moment coefficients become:

$$C_L = 2 \left\{ \int_{\Gamma_{fs}} \boldsymbol{\Sigma}_f(\mathbf{u}, p, \boldsymbol{\zeta}_e(h, \theta)) \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{e}_y \quad (1.23a)$$

$$C_M = 2 \left\{ \int_{\Gamma_{fs}} (\mathbf{R}(\theta)\mathbf{x}) \wedge \boldsymbol{\Sigma}_f(\mathbf{u}, p, \boldsymbol{\zeta}_e(h, \theta)) \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{e}_z \quad (1.23b)$$

Remark. As done in section 1.1.2.3 for the momentum and mass balance, the expression of the lift and moment coefficients eq. (1.21) can be deduced from eq. (1.23) by using the rigid-body extension eq. (1.15) and introducing $\mathbf{U} = \mathbf{R}^T \mathbf{u}$. For example, we have for the lift:

$$\begin{aligned} C_L &= 2 \left\{ \int_{\Gamma_{fs}} \boldsymbol{\Sigma}_f(\mathbf{u}, p, \boldsymbol{\zeta}_s) \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{e}_y \\ &= 2 \left\{ \int_{\Gamma_{fs}} \left[-p\mathbf{R} + \mathcal{R}e^{-1}(\nabla \mathbf{u} + \mathbf{R}\nabla \mathbf{u}^T \mathbf{R}) \right] \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{e}_y \\ &= 2 \left\{ \int_{\Gamma_{fs}} \left[-p\mathbf{R} + \mathcal{R}e^{-1}(\mathbf{R}\nabla \mathbf{U} + \mathbf{R}\nabla \mathbf{U}^T) \right] \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{e}_y \quad (1.24) \\ &= 2 \left\{ \int_{\Gamma_{fs}} \mathbf{R} \boldsymbol{\sigma}_f(p, \mathbf{U}) \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{e}_y \\ &= 2 \left\{ \int_{\Gamma_{fs}} \boldsymbol{\sigma}_f(p, \mathbf{U}) \mathbf{n}_s \, d\mathbf{x} \right\} \cdot \mathbf{R}^T \mathbf{e}_y \end{aligned}$$

The last step to recover exactly eq. (1.21a) consists in noticing that \mathbf{x} and \mathbf{X} on the one hand, and \mathbf{n}_s and \mathbf{N}_s on the other, are described by the same coordinates. Indeed for \mathbf{x} and \mathbf{X} we have $\mathbf{x} = \overrightarrow{AM}_{(\mathbf{e}_x, \mathbf{e}_y)} = \mathbf{R}(\theta)^T \left(\mathbf{R}(\theta) \overrightarrow{AM}_{(\mathbf{e}_x, \mathbf{e}_y)} \right) = \mathbf{R}(\theta)^T \overrightarrow{A_t M_t}_{(\mathbf{e}_x, \mathbf{e}_y)} = \overrightarrow{A_t M_t}_{(\mathbf{e}_x, \mathbf{e}_y)} = \mathbf{X}$. The same steps can be followed to show that $\mathbf{n}_s = \mathbf{N}_s$. Using $\mathbf{x} = \mathbf{X}$ and $\mathbf{n}_s = \mathbf{N}_s$ in the last line of the above equation yields eq. (1.21a).

Finally a similar reasoning (not shown here) can be followed for the moment coefficient.

1.2 Spatial discretization with the finite element method

In this section, we present the finite element spatial discretization of the Navier–Stokes equations that will be used throughout this manuscript. For sake of conciseness, details are provided only for the “absolute velocity - rotating axis formulation” (1.6). The same procedure (not reported) is followed for the “ALE reference configuration” equations⁵.

1.2.1 A weak formulation of the typical section problem

Weak imposition of the interface velocity continuity Before deriving the weak formulation itself, we briefly discuss the strategy used to enforce the velocity continuity boundary condition eq. (1.22) on Γ_{fs} . A classical choice is to embed the Dirichlet boundary conditions in the functional spaces the unknowns belong to. Numerically (*i.e.* after finite element discretization), this is done by modifying *a posteriori*

⁵The interested reader may find some details for this particular formulation, in the context of an elastic solid, in [J.-L. Pfister 2019, §1.1.5]

the discrete operators. Typically, suppose the finite element discretization of some linear problem writes $\mathbf{Ax} = \mathbf{b}$ *without taking into account the Dirichlet boundary conditions*. And suppose also that one wants to impose the condition $x_i = \bar{x}_i$ on some subset $i \in \Gamma_D$ of the degrees of freedom. Then, to impose that the finite element trial function respects the latter boundary condition, one simply sets the lines $i \in \Gamma_D$ of \mathbf{A} to zero, except for the diagonal terms a_{ii} equal to 1 and then sets $b_i = \bar{x}_i$. This procedure locks the degrees of freedom x_i so that the obtained solution \mathbf{x} verifies the desired boundary condition.

An alternative approach is to incorporate the Dirichlet condition in the weak formulation *via* an additional constraint equation [Deparis et al. 2016; J.-L. Pfister et al. 2019]. That additional constraint comes with the introduction of the corresponding Lagrange multiplier unknown, noted Λ , which can be conveniently defined as the interface fluid stress exerted by the fluid on the solid:

$$\Lambda = \sigma_f(p, \mathbf{U})\mathbf{N}_s \quad (1.25)$$

This approach has two advantages that make it well-suited for fluid-structure applications. First, thanks to the introduction of Λ , the interface stress is a natural unknown of the problem, thus avoiding having to compute the spatial derivatives of \mathbf{U} and p to evaluate the interface fluid stress (eq. (1.21)):

$$\begin{aligned} \mathcal{C}_L(\theta, \Lambda) &= \left\{ 2 \int_{\Gamma_{fs}} \Lambda d\mathbf{X} \right\} \cdot \mathbf{R}^T(\theta) \mathbf{e}_y \\ \mathcal{C}_M(\Lambda) &= \left\{ 2 \int_{\Gamma_{fs}} \mathbf{X} \wedge \Lambda d\mathbf{X} \right\} \cdot \mathbf{e}_z \end{aligned} \quad (1.26)$$

Second, it has been shown in purely fluid computations that a weak imposition of the no-slip boundary condition tends to yield better convergence properties of the solution, in particular close to the boundary [Bazilevs et al. 2007].

Weak formulation Let us start by defining the functional spaces in which the fluid velocity, pressure and Lagrange multiplier are sought:

$$\begin{aligned} \mathcal{V}^U &= \left\{ \mathbf{f} \in \mathcal{H}^1(\Omega_f) \quad \text{s.t.} \quad \mathbf{f} = 0 \text{ on } \Gamma_{in} \right\} \\ \mathcal{V}^p &= \left\{ f \in \mathcal{L}^2(\Omega_f) \right\} \\ \mathcal{V}^\Lambda &= \left\{ \mathbf{f} \in \mathcal{H}^{-1/2}(\Gamma_{fs}) \right\} \end{aligned}$$

By multiplying the fluid momentum equation eq. (1.6a) by some test function, $\boldsymbol{\psi}^U \in \mathcal{V}^U$, and integrating over Ω_f , we have:

$$\begin{aligned} \underbrace{\int_{\Omega_f} \frac{\partial \mathbf{U}}{\partial t} \cdot \boldsymbol{\psi}^U}_{:= \left\langle \frac{\partial \mathbf{U}}{\partial t}, \boldsymbol{\psi}^U \right\rangle} + \underbrace{\int_{\Omega_f} \left\{ u_\theta \mathbf{e}_z \wedge \mathbf{U} + \nabla \mathbf{U} (\mathbf{U} - \mathbf{W}(u_h, h, u_\theta)) - \nabla \cdot \sigma_f(p, \mathbf{U}) \right\} \cdot \boldsymbol{\psi}^U}_{:= \left\langle \mathcal{R}^U(\mathbf{U}, p, u_h, \theta, u_\theta), \boldsymbol{\psi}^U \right\rangle} = 0 \end{aligned} \quad (1.27)$$

Then, the divergence term is integrated by part where the inflow and outflow contributions respectively disappear due to the test functions being null on Γ_{in} and to the “do-nothing” outflow condition and where the definition eq. (1.25) of the interface

Lagrange multiplier is used:

$$\begin{aligned} \langle \mathcal{R}^{\mathbf{U}}(\mathbf{U}, p, u_h, \theta, u_\theta), \boldsymbol{\psi}^{\mathbf{U}} \rangle &= \int_{\Omega_f} \{u_\theta \mathbf{e}_z \wedge \mathbf{U} + \nabla \mathbf{U} (\mathbf{U} - \mathbf{W}(u_h, h, u_\theta))\} \cdot \boldsymbol{\psi}^{\mathbf{U}} \\ &+ \int_{\Omega_f} \boldsymbol{\sigma}_f(p, \mathbf{U}) : \nabla \boldsymbol{\psi}^{\mathbf{U}} + \int_{\Gamma_{fs}} \boldsymbol{\Lambda} \cdot \boldsymbol{\psi}^{\mathbf{U}} \end{aligned}$$

Note that the interface integral comes with a positive sign due to the definition of $\boldsymbol{\Lambda}$ that uses the *inward* normal to the fluid domain, \mathbf{N}_s .

The weak formulations corresponding to the continuity equation eq. (1.6b) and to the interface velocity constraint eq. (1.20) are directly obtained by projecting the strong equations along $\boldsymbol{\psi}^p \in \mathcal{V}^p$ and $\boldsymbol{\psi}^\Lambda \in \mathcal{V}^\Lambda$:

$$\begin{aligned} \langle \mathcal{R}^p(\mathbf{U}), \boldsymbol{\psi}^p \rangle &= - \int_{\Omega_f} \nabla \cdot \mathbf{U} \boldsymbol{\psi}^p = 0 \\ \langle \mathcal{R}^\Lambda(\mathbf{U}, u_h, \theta, u_\theta), \boldsymbol{\psi}^\Lambda \rangle_{\Gamma_{fs}} &= \int_{\Gamma_{fs}} \{\mathbf{U} - \mathbf{W}(u_h, h, u_\theta)\} \cdot \boldsymbol{\psi}^\Lambda = 0 \end{aligned} \quad (1.28)$$

In order to complete the weak formulation of the coupled problem, we introduce the heaving and pitching velocities, $u_h = \partial h / \partial t$ and $u_\theta = \partial \theta / \partial t$ thanks to which the solid equations (1.3) can be formulated only with first order in time derivatives. Then, the above definitions of u_h, u_θ and the solid equations are simply multiplied by adequate test functions $\psi^h, \psi^\theta, \psi^{u_h}, \psi^{u_\theta}$ that are numbers in \mathbb{R} . Gathering together the fluid (eq. (1.27) and (1.28)) and solid weak formulations finally yields the coupled fluid-structure weak form for the typical section problem:

$$\begin{aligned} \left\langle \boldsymbol{\psi}, \mathcal{M} \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}) \right\rangle &= \left\{ \frac{\partial h}{\partial t} - u_h \right\} \psi^h \\ &+ \left\{ \frac{\partial \theta}{\partial t} - u_\theta \right\} \psi^\theta \\ &+ \left\{ \frac{\partial u_h}{\partial t} + 2\zeta_h \frac{\Omega}{U_c^*} u_h + \frac{\Omega^2}{U_c^{*2}} h - \frac{1}{\tilde{m}} C_L(\theta, \boldsymbol{\Lambda}) \right\} \psi^{u_h} \\ &+ \left\{ \frac{\partial u_\theta}{\partial t} + 2\zeta_p \frac{1}{U_c^*} u_\theta + \frac{1}{U_c^{*2}} \theta - \frac{1}{r_\theta^2 \tilde{m}} C_M(\boldsymbol{\Lambda}) \right\} \psi^{u_\theta} \quad (1.29) \\ &+ \left\langle \frac{\partial \mathbf{U}}{\partial t}, \boldsymbol{\psi}^{\mathbf{U}} \right\rangle + \langle \mathcal{R}^{\mathbf{U}}(\mathbf{U}, p, u_h, \theta, u_\theta), \boldsymbol{\psi}^{\mathbf{U}} \rangle \\ &+ \langle \mathcal{R}^p(\mathbf{U}), \boldsymbol{\psi}^p \rangle \\ &+ \langle \mathcal{R}^\Lambda(\mathbf{U}, u_h, \theta, u_\theta), \boldsymbol{\psi}^\Lambda \rangle_{\Gamma_{fs}} \\ &= 0 \end{aligned}$$

where we have introduced the vector $\mathbf{q} = [h, \theta, u_h, u_\theta, \mathbf{U}, p, \boldsymbol{\Lambda}]^T$ of all unknowns, the corresponding test function $\boldsymbol{\psi} = [\psi^h, \psi^\theta, \psi^{u_h}, \psi^{u_\theta}, \boldsymbol{\psi}^{\mathbf{U}}, \boldsymbol{\psi}^p, \boldsymbol{\psi}^\Lambda]^T$, the symbolic operators \mathcal{M} that accounts for terms proportional to a time-derivative and \mathcal{R} that gathers all remaining terms. The first two lines correspond to the definition of the heaving and pitching velocities, the third and fourth lines are the momentum balances along the heaving and pitching degrees of freedom. The fifth and sixth lines are the classical Galerkin weak form of the Navier–Stokes equations. The last line is the constraint imposing the continuity of velocity on Γ_{fs} .

By an abuse of notation, we will often refer to the strong equations corresponding to the weak form eq. (1.29) using:

$$\mathcal{M} \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}) = 0 \quad (1.30)$$

where

$$\mathcal{M} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix} \quad \mathcal{R}(\mathbf{q}) = \begin{pmatrix} -u_h \\ -u_\theta \\ 2\zeta_h \Omega U^{*-1} u_h + \Omega^2 U^{*-2} h - \tilde{m}^{-1} \mathcal{C}_L(\theta, \Lambda) \\ 2\zeta_p U^{*-1} u_\theta + U^{*-2} \theta - \tilde{m}^{-1} r_\theta^{-2} \mathcal{C}_M(\Lambda) \\ \mathcal{R}^U(\mathbf{U}, p, u_h, \theta, u_\theta) \\ \mathcal{R}^p(\mathbf{U}) \\ \mathcal{R}^\Lambda(\mathbf{U}, u_h, \theta, u_\theta) \end{pmatrix}$$

1.2.2 Finite elements discretized problem

To obtain the discrete problem corresponding to eq. (1.29), a Delaunay triangulation $\mathcal{T}_\Delta = \{K\}$ of Ω_f is built, with typical element size Δ . The infinite dimensional functional spaces $\mathcal{V}^U, \mathcal{V}^p, \mathcal{V}^\Lambda$ are approximated by corresponding finite element spaces $\mathcal{V}_\Delta^U, \mathcal{V}_\Delta^p$ and $\mathcal{V}_\Delta^\Lambda$ defined below. For velocity and pressure \mathbb{P}_2 and \mathbb{P}_1 elements, respectively, are used so that the classical inf-sup Ladyženskaja–Babuška–Brezzi (LBB) condition for the incompressibility constraint is fulfilled [Fortin et al. 1991]. For the interface Lagrange multiplier, \mathbb{P}_1 elements are used.

$$\begin{aligned} \mathcal{V}_\Delta^U &= \{f_\Delta \in \mathcal{C}^0(\Omega_f) \quad \text{s.t.} \quad f_\Delta|_K \in \mathbb{P}_2(K), \forall K \in \mathcal{T}_\Delta, \quad f_\Delta = 0 \text{ on } \Gamma_{fs}\} \\ \mathcal{V}_\Delta^p &= \{f_\Delta \in \mathcal{C}^0(\Omega_f) \quad \text{s.t.} \quad f_\Delta|_K \in \mathbb{P}_1(K), \forall K \in \mathcal{T}_\Delta\} \\ \mathcal{V}_\Delta^\Lambda &= \{f_\Delta \in \mathcal{C}^0(\Gamma_{fs}) \quad \text{s.t.} \quad f_\Delta|_K \in \mathbb{P}_1(K), \forall K \in \mathcal{T}_\Delta\} \end{aligned}$$

Each of these spaces is generated by the corresponding finite element basis

$$\begin{aligned} \mathcal{V}_\Delta^U &= \text{span} \left(\psi_0^U, \psi_1^U, \dots, \psi_{n_{\text{dof},U}-1}^U \right) \\ \mathcal{V}_\Delta^p &= \text{span} \left(\psi_0^p, \psi_1^p, \dots, \psi_{n_{\text{dof},p}-1}^p \right) \\ \mathcal{V}_\Delta^\Lambda &= \text{span} \left(\psi_0^\Lambda, \psi_1^\Lambda, \dots, \psi_{n_{\text{dof},\Lambda}-1}^\Lambda \right) \end{aligned}$$

such that the finite dimensional approximations $\mathbf{U}_\Delta \in \mathcal{V}_\Delta^U, p_\Delta \in \mathcal{V}_\Delta^p$ and $\Lambda_\Delta \in \mathcal{V}_\Delta^\Lambda$ to \mathbf{U}, p and Λ can be decomposed as:

$$\begin{aligned} \mathbf{U}_\Delta(\mathbf{X}, t) &= \sum_{i=0}^{n_{\text{dof},U}-1} \mathbf{U}_i(t) \psi_i^U(\mathbf{X}) \\ p_\Delta(\mathbf{X}, t) &= \sum_{i=0}^{n_{\text{dof},p}-1} p_i(t) \psi_i^p(\mathbf{X}) \\ \Lambda_\Delta(\mathbf{X}, t) &= \sum_{i=0}^{n_{\text{dof},\Lambda}-1} \Lambda_i(t) \psi_i^\Lambda(\mathbf{X}) \end{aligned} \quad (1.31)$$

where the $\mathbf{U}_i(t), p_i(t)$ and $\Lambda_i(t)$ are the (time-dependent) finite element degrees of freedom. Finally, for the solid part, the unknowns h, u_h, θ and u_θ are discrete by

nature. As a consequence, the corresponding test and trial functions are simply the trivial base function of \mathbb{R} , *i.e.*, the scalar "one".

Using eq. (1.31) in eq. (1.29) the discrete problem is then obtained and writes:

$$\mathbf{M}(\mathbf{q}_\Delta) \frac{\partial \mathbf{q}_\Delta}{\partial t} + \mathbf{R}(\mathbf{q}_\Delta) = 0 \quad (1.32)$$

where the solid and fluid degrees of freedom are stored in

$$\mathbf{q}_\Delta = \left(h, \theta, u_h, u_\theta, \mathbf{U}_0, \dots, \mathbf{U}_{n_{\text{dof},U}-1}, p_0, \dots, p_{n_{\text{dof},p}-1}, \Lambda_0, \dots, \Lambda_{n_{\text{dof},\Lambda}-1} \right)^\top$$

$$\mathbf{M}(\mathbf{q}_\Delta) = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & \mathbf{M}_{U,\tau}(\mathbf{q}_\Delta) & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

$$\mathbf{R}(\mathbf{q}_\Delta) = \begin{pmatrix} -u_h \\ -u_\theta \\ 2\zeta_h \Omega U^{*-1} u_h + \Omega^2 U^{*-2} h - \tilde{m}^{-1} \mathcal{C}_L(\theta, \Lambda_\Delta) \\ 2\zeta_p U^{*-1} u_\theta + U^{*-2} \theta - \tilde{m}^{-1} r_\theta^{-2} \mathcal{C}_M(\Lambda_\Delta) \\ \mathbf{R}^U(\mathbf{q}_\Delta) \\ \mathbf{R}^p(\mathbf{q}_\Delta) \\ \mathbf{R}^\Lambda(\mathbf{q}_\Delta) \end{pmatrix}$$

The discrete residuals for the incompressibility and interface velocity constraints are vectors defined by

$$[\mathbf{R}^p(\mathbf{q}_\Delta)]_i = \left\langle \mathcal{R}^p(\mathbf{U}_\Delta), \psi_i^p \right\rangle \quad 0 \leq i < n_{\text{dof},p}$$

$$[\mathbf{R}^\Lambda(\mathbf{q}_\Delta)]_i = \left\langle \mathcal{R}^\Lambda(\mathbf{U}_\Delta, u_h, \theta, u_\theta), \psi_i^\Lambda \right\rangle_{\Gamma_{\text{fs}}} \quad 0 \leq i < n_{\text{dof},\Lambda}$$

For the fluid momentum equation, we slightly modify the classical Galerkin formulation in eq. (1.29), as seen in the next paragraph.

Streamlined Upwind Petrov–Galerkin (SUPG) stabilization of the fluid momentum equation In the advection-dominated regime, it is known that the classical Galerkin finite element discretization eq. (1.27) of advection-diffusion equations generates numerical oscillations in the solution. In order to avoid such oscillations — that are likely to appear for the highest Reynolds numbers investigated in this manuscript ($\mathcal{O}(10^4)$) — we stabilize the finite element method using the Streamline Upwind Petrov–Galerkin (SUPG) [Brooks et al. 1982] technique. The latter has been used through the last four decades to stabilize the Navier–Stokes equations for both steady–state and time-dependent nonlinear solutions. It is based on adding to the classical Galerkin test function for the fluid momentum equation, $\boldsymbol{\psi}^U$, a convected component $\nabla \boldsymbol{\psi}^U(\mathbf{U} - \mathbf{W})$. This new test function introduces artificial diffusion along the streamlines of the field $(\mathbf{U} - \mathbf{W})$, which provides some damping of the numerical oscillations. The amount of artificial diffusion is governed by a

space-dependent parameter noted τ .

$$\tau = \frac{\zeta(\text{Pe}_\Delta) \Delta}{2\|\mathbf{U} - \mathbf{W}\|} \quad \text{with} \quad \text{Pe}_\Delta = \frac{\|\mathbf{U} - \mathbf{W}\| \Delta}{\mathcal{R}e^{-1}} \quad \text{and} \quad \zeta(x) = \begin{cases} x/3 & \text{if } x \leq 3 \\ 1 & \text{if } x > 3 \end{cases}$$

where Δ is a typical mesh size⁶ and Pe_Δ is the element Péclet number.

Using the SUPG test function, $\boldsymbol{\psi}^{\mathbf{U}} + \tau \nabla \boldsymbol{\psi}^{\mathbf{U}}(\mathbf{U} - \mathbf{W})$, instead of the pure Galerkin test function $\boldsymbol{\psi}^{\mathbf{U}}$, the vector $\mathbf{R}^{\mathbf{U}}$ and matrix $\mathbf{M}_{\mathbf{U},\tau}$ that define the discretized Navier–Stokes momentum equation read

$$\begin{aligned} [\mathbf{R}^{\mathbf{U}}(\mathbf{q}_\Delta)]_i &= \left\langle \mathcal{R}^{\mathbf{U}}(\mathbf{U}_\Delta, p_\Delta, u_h, \theta, u_\theta), \boldsymbol{\psi}_i^{\mathbf{U}} \right\rangle \\ &+ \left\langle \mathcal{R}^{\mathbf{U}}(\mathbf{U}_\Delta, p_\Delta, u_h, \theta, u_\theta), \tau \nabla \boldsymbol{\psi}_i^{\mathbf{U}}(\mathbf{U}_\Delta - \mathbf{W}(u_h, \theta, u_\theta)) \right\rangle \quad 0 \leq i < n_{\text{dof},\mathbf{U}} \end{aligned}$$

and

$$[\mathbf{M}_{\mathbf{U},\tau}(\mathbf{q}_\Delta)]_{ij} := \left\langle \boldsymbol{\psi}_j^{\mathbf{U}}, \boldsymbol{\psi}_i^{\mathbf{U}} \right\rangle + \left\langle \boldsymbol{\psi}_j^{\mathbf{U}}, \tau \nabla \boldsymbol{\psi}_i^{\mathbf{U}}(\mathbf{U}_\Delta - \mathbf{W}(u_h, \theta, u_\theta)) \right\rangle \quad 0 \leq i, j < n_{\text{dof},\mathbf{U}}$$

In both cases, the first term on the right-hand side corresponds to the classical Galerkin formulation, eq. (1.27), whereas the second term is the SUPG contribution. If $\tau = 0$, the pure Galerkin formulation is retrieved. Note that due to SUPG, $\mathbf{M}_{\mathbf{U},\tau}$ now depends on the solution \mathbf{q}_Δ , through the SUPG test function. For the same reason, an additional level of nonlinearity in \mathbf{q}_Δ is introduced in $\mathbf{R}^{\mathbf{U}}$. As a consequence, it is seen that additional nonlinearities appear in the discretized problem due to the SUPG formulation, that are not present in the original Galerkin weak problem.

1.3 Time-integration algorithm

In this section, we present the algorithm used to obtain general time-dependent solutions of the typical section problem. Let us start by recalling below the solid model eq. (1.3) (formulated at first order in time by introducing $u_h = \partial h / \partial t$ and $u_\theta = \partial \theta / \partial t$):

$$\underbrace{\frac{\partial}{\partial t} \begin{pmatrix} h \\ \theta \\ u_h \\ u_\theta \end{pmatrix}}_{:= \mathbf{q}_s} + \underbrace{\begin{pmatrix} -u_h \\ -u_\theta \\ 2\zeta_h \Omega U_c^{*-1} u_h + \Omega^2 U_c^{*-2} h - \tilde{m}^{-1} \mathcal{C}_L \\ 2\zeta_p U_c^{*-1} u_\theta + U_c^{*-2} \theta - r_\theta^{-2} \tilde{m}^{-1} \mathcal{C}_M \end{pmatrix}}_{:= \mathbf{R}_s(\mathbf{q}_s, (\mathcal{C}_L, \mathcal{C}_M))} = 0 \quad (1.33)$$

For convenience, we introduce in eq. (1.33) the vector \mathbf{q}_s that gathers the solid degrees of freedom. As for the fluid model, we use the Navier–Stokes equations with ALE formalism *written in the deformed domain* $\Omega_{f,t}$ eq. (1.10) and accompanied with

⁶In practice, the size of a (triangular) mesh element is taken as the diameter of its circumscribed circle, given by the `hTriangle` keyword in the FreeFEM language.

the velocity continuity condition at the fluid-structure interface eq. (1.18):

$$\left. \frac{\partial \mathbf{u}_t}{\partial t} \right|_{\mathbf{x}} + \nabla \mathbf{u}_t (\mathbf{u}_t - \mathbf{w}_t(\mathbf{q}_s)) - \nabla \cdot \left[-p_t \mathbf{I} + \mathcal{R}e^{-1} \left(\nabla \mathbf{u}_t + \nabla \mathbf{u}_t^T \right) \right] = 0 \quad \mathbf{x}_t \in \Omega_{f,t}(\mathbf{q}_s) \quad (1.34a)$$

$$-\nabla \cdot \mathbf{u}_t = 0 \quad \mathbf{x}_t \in \Omega_{f,t}(\mathbf{q}_s) \quad (1.34b)$$

$$\mathbf{u}_t - \mathbf{w}_t(\mathbf{q}_s) = 0 \quad \mathbf{x}_t \in \Gamma_{fs,t}(\mathbf{q}_s) \quad (1.34c)$$

Notice that we have precised in eq. (1.34) the dependence of the fluid problem on \mathbf{q}_s , through the interface condition and the shape of the fluid domain.

1.3.1 Temporal scheme

Time-discretization is performed using the Backward Differences Formula of order q (BDF q) that approximates the time-derivative of a quantity ϕ at instant $n + 1$ by

$$\left. \frac{\partial \phi}{\partial t} \right|_{t=t_{n+1}} \simeq D\phi^{n+1} = \frac{1}{\Delta t} \left(\alpha_0 \phi^{n+1} + \sum_{i=1}^q \alpha_i \phi^{n+1-i} \right) \quad (1.35)$$

where ϕ^n the short notation for $\phi(t_n)$. In the following, we use the BDF2 scheme which defines by $\alpha_0 = 3/2$, $\alpha_1 = -2$ and $\alpha_2 = 1/2$. This leads to the time-discretized versions of eq. (1.33)

$$D\mathbf{q}_s^{n+1} + \mathbf{R}_s(\mathbf{q}_s^{n+1}, (\mathcal{C}_L, \mathcal{C}_M)^{n+1}) = 0 \quad (1.36)$$

and eq. (1.34)

$$\begin{aligned} D\mathbf{u}_t^{n+1} + \nabla \mathbf{u}_t^{n+1} \left(\mathbf{u}_t^{n+1} - \mathbf{w}_t(\mathbf{q}_s^{n+1}) \right) - \nabla \cdot \left[-p_t^{n+1} \mathbf{I} + \mathcal{R}e^{-1} \left(\nabla \mathbf{u}_t^{n+1} + \nabla \mathbf{u}_t^{n+1T} \right) \right] &= 0 \\ &\text{in } \Omega_{f,t}(\mathbf{q}_s^{n+1}) \\ -\nabla \cdot \mathbf{u}_t^{n+1} &= 0 \quad \text{in } \Omega_{f,t}(\mathbf{q}_s^{n+1}) \\ \mathbf{u}_t^{n+1} - \mathbf{w}_t(\mathbf{q}_s^{n+1}) &= 0 \quad \text{on } \Gamma_{fs,t}(\mathbf{q}_s^{n+1}) \end{aligned} \quad (1.37)$$

1.3.2 Fluid-structure pressure-segregation algorithm

The coupled nonlinear problem constituted by eq. (1.36) and (1.37) is solved using the pressure segregation method proposed in [Badia et al. 2007]. It consists in mixing fixed-point iterations for handling the fluid-structure coupling with a pressure-correction scheme [Guermond et al. 2006] for the incompressible Navier-Stokes problem. In the following, we provide some insights on this method.

Fixed-point iterations for fluid-structure coupling Consider the fluid *Dirichlet-to-Neumann Steklov-Poincaré* operator \mathcal{S}_f

$$\mathcal{S}_f : \mathbf{q}_s^{n+1} \rightarrow (\mathcal{C}_L, \mathcal{C}_M)^{n+1}$$

that associates to the solid state \mathbf{q}_s^{n+1} , the lift and moment coefficients $(\mathcal{C}_L, \mathcal{C}_M)^{n+1}$ obtained by solving eq. (1.37). Symmetrically, define the solid *Neumann-to-Dirichlet*

Steklov-Poincaré operator \mathcal{S}_s^{-1}

$$\mathcal{S}_s^{-1} : (\mathcal{C}_L, \mathcal{C}_M)^{n+1} \rightarrow \mathbf{q}_s^{n+1}$$

that associates the solid state \mathbf{q}_s^{n+1} verifying eq. (1.36) to the lift and momentum coefficients $(\mathcal{C}_L, \mathcal{C}_M)^{n+1}$.

Using the Steklov-Poincaré operators, the coupled problem eq. (1.36)–(1.37) is equivalently described as a fixed point of the operator $\mathcal{S}_s^{-1} \circ \mathcal{S}_f$:

$$\mathbf{q}_s^{n+1} = \mathcal{S}_s^{-1}(\mathcal{S}_f(\mathbf{q}_s^{n+1}))$$

Then, a classical method to obtain \mathbf{q}_s^{n+1} is to perform fixed-point iterations, indexed by k :

$$\mathbf{q}_s^{n+1,k+1} = \mathcal{S}_s^{-1}(\mathcal{S}_f(\mathbf{q}_s^{n+1,k})) \quad (1.38)$$

Application of \mathcal{S}_f The first step of the fixed-point iteration eq. (1.38) is to compute $(\mathcal{C}_L, \mathcal{C}_M)^{n+1,k+1} = \mathcal{S}_f(\mathbf{q}_s^{n+1,k})$. This is a nonlinear problem due to the Navier–Stokes convection term. It can be solved by using pressure-correction iterations, indexed by $0 \leq l < l_{\max}$, that read:

- Step 1: Advection-diffusion

$$\begin{aligned} & D\mathbf{u}_t^{n+1,k+1,l+1} + \nabla \mathbf{u}_t^{n+1,k+1,l+1} \left(\mathbf{u}_t^{n+1,k+1,l} - \mathbf{w}_t(\mathbf{q}_s^{n+1,k}) \right) \\ & - \nabla \cdot \left[-p_t^{n+1,k+1,l} \mathbf{I} + \mathcal{R}e^{-1} \left(\nabla \mathbf{u}_t^{n+1,k+1,l+1} + \nabla \mathbf{u}_t^{n+1,k+1,l+1\text{T}} \right) \right] = 0 \quad (1.39) \\ & \hspace{15em} \text{in } \Omega_{f,t}(\mathbf{q}_s^{n+1,k}) \\ & \mathbf{u}_t^{n+1,k+1,l+1} - \mathbf{w}_t(\mathbf{q}_s^{n+1,k}) = 0 \quad \text{on } \Gamma_{fs,t}(\mathbf{q}_s^{n+1,k}) \end{aligned}$$

- Step 2: Pressure-correction

$$\begin{aligned} \nabla^2 \left(p_t^{n+1,k+1,l+1} - p_t^{n+1,k+1,l} \right) &= \frac{\alpha_0}{\Delta t} \nabla \cdot \mathbf{u}_t^{n+1,k+1,l+1} \quad \text{in } \Omega_{f,t}(\mathbf{q}_s^{n+1,k}) \\ \nabla \left(p_t^{n+1,k+1,l+1} - p_t^{n+1,k+1,l} \right) \cdot \mathbf{n}_f &= 0 \quad \text{on } \Gamma_{fs,t}(\mathbf{q}_s^{n+1,k}) \end{aligned} \quad (1.40)$$

Thanks to the pressure-correction approach, velocity and pressure are computed in a decoupled manner. First, the new velocity estimate $\mathbf{u}_t^{n+1,k+1,l+1}$ is obtained in step 1 by solving a linear (note indeed that the advection velocity is taken at the previous step, l) advection-diffusion problem. Then, in step 2, the pressure estimate is corrected by solving a simple Poisson problem. With this approach, at a general iteration l , the velocity field is not necessarily divergence-free nor verifies the *nonlinear* Navier–Stokes momentum equation. However, at convergence of the pressure-correction iterations $l = l_{\max}$, velocity and pressure do not evolve anymore: $\mathbf{u}_t^{n+1,k+1,l_{\max}} = \mathbf{u}_t^{n+1,k+1,l_{\max}-1}$ and $p_t^{n+1,k+1,l_{\max}} = p_t^{n+1,k+1,l_{\max}-1}$. As a consequence, the left-hand side of eq. (1.40) is zero, which guarantees that incompressibility is enforced. In addition, the advection-diffusion problem eq. (1.39) has converged to the nonlinear Navier–Stokes momentum equation since the velocity is everywhere the same.

Application of \mathcal{S}_s^{-1} The second step of the fixed-point iteration eq. (1.38) is to compute $\mathbf{q}_s^{n+1,k+1} = \mathcal{S}_s^{-1}((\mathcal{C}_L, \mathcal{C}_M)^{n+1,k+1})$. This is a linear⁷ problem for $\mathbf{q}_s^{n+1,k+1}$ that is straightforwardly solved with a direct method.

Mixed iterations In the end, solving the coupled problem eq. (1.36) and (1.37) requires two nested loops: an outer fixed-point loop (eq. (1.38)) to handle the fluid-structure coupling and an inner pressure-correction loop (eq. (1.39) and (1.40)) to handle the Navier–Stokes nonlinearity and the incompressibility constraint. To avoid that and hopefully obtain a more efficient algorithm, we follow [Badia et al. 2007] and perform only one loop. This is equivalent to using only $l_{\max} = 1$ pressure-correction iteration per application of \mathcal{S}_t . The mixed iteration is stopped when the difference between two consecutive iterates of the pressure and solid state are below some specified relative tolerances:

$$\frac{\|p_t^{n+1,k+1} - p_t^{n+1,k}\|_{\mathcal{L}^2}}{\|p_t^{n+1,k}\|_{\mathcal{L}^2}} < \text{tol}_p \quad \text{and} \quad \frac{\|\mathbf{q}_s^{n+1,k+1} - \mathbf{q}_s^{n+1,k}\|_{\mathbb{R}^4}}{\|\mathbf{q}_s^{n+1,k}\|_{\mathbb{R}^4}} < \text{tol}_s$$

In practice, we used $\text{tol}_p = 10^{-4}$ and $\text{tol}_s = 10^{-6}$.

⁷It is linear because we use $x_\theta = 0$, else it would be nonlinear according to eq. (1.2)

2

A WEAKLY NONLINEAR SOLVER FOR FLUTTER BIFURCATION ANALYSIS

This chapter presents in extensive details the weakly nonlinear approach for the typical aeroelastic section. We also introduce a Hessian-based mesh adaptation framework that is used in the subsequent chapters to efficiently explore the parameter space. The chapter ends with a validation of the weakly nonlinear results against reference fully nonlinear time-marching solutions.

2.1 Introduction

One of the goals of this manuscript is to explore the possible scenarios for the flutter bifurcation for different combinations of non-dimensional parameters, like the Reynolds number or the solid-to-fluid mass ratio. To do so efficiently, the classical time-integration approach (section 1.3) is not well-suited, not only due to its high computational cost, but also because its typical output is a general transient solution, not particularly relevant to bifurcation analysis where permanent regimes (limit cycle oscillations) are searched. As an alternative, methods have been developed to derive nonlinear reduced-order models that yield lower computational cost and possibly more adapted outputs [Lucia et al. 2004; Badcock et al. 2011; Fabre et al. 2018].

Close enough to the bifurcation threshold, a general mathematical form that describes the nonlinear dynamics of a system is known as the *normal form* [Nayfeh et al. 1995, §2.3.2]. For a codimension-one bifurcation like flutter, it is a simple nonlinear equation for a complex amplitude variable that governs the amplitude of the oscillations. As a consequence, once the normal form is known, the system's response can be computed at negligible computational cost. In addition, key properties of the system, like the sub- or supercritical nature of the bifurcation (*cf* section 2.3.5), can be deduced directly from the normal form coefficients, without solving it.

Various methods may be applied to derive the normal form. In the context of hydrodynamic global stability studies, several researchers have used the method of multiple scales [Nayfeh et al. 1995, §2.3.6], often referring to it as *weakly nonlinear analysis*. In this method, the solution is expanded as a power series of a (small) deviation from the linear instability threshold, yielding a sequence of equations governing each order. For a Hopf bifurcation, the normal form typically emerges at third order as a result of a compatibility condition that needs to be enforced to ensure a well-posed problem. Such analysis have proved an interesting tool for exploring the bifurcations occurring in a variety of viscous flows around bluff-bodies: *e.g.* the classical two-dimensional cylinder flow [Sipp et al. 2007; Gallaire et al. 2016] or three-dimensional flows around disks [Meliga et al. 2009] and rotating spheres [Citro et al. 2016]. This approach has been extended by Tchoufag et al. (2015) for a fluid-structure application, namely, the study of the fluttering and spiraling motion of buoyancy-driven disks in viscous flows. For aeroelastic applications, nonlinear reduced-order models¹ have been derived for the transonic flutter of commercial airplane wings [Woodgate et al. 2007] and for the roll instability of a delta wing [Badcock et al. 2008], using an Euler flow model. These models were derived using center manifold reduction [Nayfeh et al. 1995, §2.3.4]. With this strategy, the analytical aspects are drastically reduced in comparison to the method of multiple scales, so that the approach becomes mostly numerical. If this is well-suited for practical integration in a previously developed CFD code, it however provides less insights in the mechanisms that govern the solution at the different orders. Note finally that in [Carini et al. 2015] the normal forms for the cylinder flow bifurcation, obtained through the multiple scales and center manifold approaches, were compared and yielded consistent results.

In this manuscript, we opt for the method of multiple scales due to its more analytical nature. The present chapter is dedicated to the construction and validation of the weakly nonlinear solver that will ultimately give the normal form representation of the flutter bifurcation for the typical section model introduced in the previous

¹Strictly speaking, they are not normal forms ...

chapter. In section 2.2, we start by briefly recalling the physical models. Then, the weakly nonlinear approach itself is detailed in section 2.3 leading to a sequence of partial differential equations to solve. Finally, in section 2.4, some numerical aspects of the solver are detailed such as the SUPG-stabilized finite element discretization and a Hessian-based mesh adaptation framework. This numerical section ends with a validation of the solver implementation.

2.2 Equations of motion: typical section in viscous flow

Before diving into the weakly nonlinear analysis of the typical section problem, we briefly recall the governing equations, that have been presented in details in chapter 1. The plate movement is defined by the heaving displacement h and pitching angle θ . A linear model is considered for the structure so that h and θ are both governed by classical damped linear oscillator equations (with a fluid forcing term). The heaving and pitching velocities, u_h and u_θ respectively, are also introduced so that the problem is kept first order in time. The fluid is described by the incompressible Navier–Stokes equations, formulated in the “absolute velocity - rotating axis” formalism, introduced in section 1.1.2.1. This particular formalism is used because it combines the following double advantage. First, it is written on a *fixed* fluid domain, which allows exact linearization (and “weakly nonlinearization”) with respect to the solid degrees of freedom. Second, it presents a moderate algebraic complexity in the sense that the movement of the physical fluid domain only materializes through a grid velocity contribution $-\mathbf{W}(\theta, u_h, u_\theta)$ in the convection velocity and a rotation term $u_\theta \mathbf{e}_z \wedge \mathbf{U}$. As a consequence, the cumbersome analytical manipulations involved in the proposed weakly nonlinear analysis remain of reasonable complexity²

Re-using the symbolic operators introduced in eq. (1.30), the governing equations for the typical section problem write, in strong form:

$$\mathcal{M} \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}; U^*) = 0 \quad (2.1)$$

with

$$\mathbf{q} = \begin{pmatrix} h \\ \theta \\ u_h \\ u_\theta \\ \mathbf{U} \\ p \\ \Lambda \end{pmatrix}, \quad \mathcal{M} = \begin{pmatrix} 1 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 1 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

and the nonlinear evolution operator:

$$\mathcal{R}(\mathbf{q}; U^*) = \begin{pmatrix} -u_h \\ -u_\theta \\ 2\zeta_h \Omega U^{*-1} u_h + \Omega^2 U^{*-2} h - \tilde{m}^{-1} \mathcal{C}_L(\Lambda, \theta) \\ 2\zeta_p U^{*-1} u_\theta + U^{*-2} \theta - \tilde{m}^{-1} r_\theta^{-2} \mathcal{C}_M(\Lambda) \\ u_\theta \mathbf{e}_z \wedge \mathbf{U} + \nabla \mathbf{U} (\mathbf{U} - \mathbf{W}(u_h, \theta, u_\theta)) - \nabla \cdot \sigma_f(p, \mathbf{U}) \\ -\nabla \cdot \mathbf{U} \\ \mathbf{U} - \mathbf{W}(u_h, \theta, u_\theta) \end{pmatrix}$$

²In other words, the task would be much more painful with the “reference configuration ALE” formalism of section 1.1.2.2 ...

where the dependence of \mathcal{R} on the reduced velocity U^* is recalled for clarity. As a consequence, eq. (2.1) is a system of seven equations that detail as follows. The two first lines correspond to the definition of the heaving and pitching velocities and the third and fourth lines to the solid momentum equations. The fifth and sixth lines are the Navier–Stokes momentum equation and incompressibility constraint. The last line enforces the continuity of velocity on the fluid–structure interface Γ_{fs} . All the following weakly nonlinear developments are performed using the strong form eq. (2.1) of the governing equations.

2.3 Weakly nonlinear analysis

In this section, we present a weakly nonlinear analysis of eq. (2.1), using the method of multiple scales [Nayfeh et al. 1995]. Noting U_c^* the linear critical velocity threshold at which the flutter instability occurs, we seek to reconstruct the (weakly) nonlinear solution for U^* close to U_c^* . The reduced velocity is thus developed as

$$\frac{1}{U^{*2}} = \frac{1}{U_c^{*2}} + \epsilon^2 \Delta_{U^*} \quad (2.2)$$

with $0 < \epsilon \ll 1$ quantifying the deviation from the threshold. Note that the parameter $\Delta_{U^*} = \pm 1$ is only introduced to handle the two possible bifurcation scenarios. Indeed, if the bifurcation is supercritical, then we seek solutions for $U^* \geq U_c^*$, which requires $\Delta_{U^*} = -1$. On the contrary, if the bifurcation is subcritical, we seek solutions for $U^* \leq U_c^*$, which requires $\Delta_{U^*} = +1$. Note also that we have opted for developing directly $U^{*-2} = U_c^{*-2} + \epsilon^2 \Delta_{U^*}$ instead of $U^* = U_c^* + \epsilon^2 \Delta_{U^*}$, for example. This choice is motivated by the work of Gallaire et al. (2016) who showed that, in the case of the vortex shedding bifurcation behind a circular cylinder, developing the parameter as it naturally appears in the equation (*i.e.* Re^{-1} , instead of Re) improves the accuracy of the weakly nonlinear analysis. On the practical side, this choice also has the advantage of avoiding additional calculatory steps that arise when expanding U^{*-2} from the expansion of U^* .

The multiple scales approach consists in seeking for a solution \mathbf{q} under a series form:

$$\mathbf{q} = \mathbf{q}_0 + \epsilon \mathbf{q}_1(t_0, t_1, t_2, t_3, \dots) + \epsilon^2 \mathbf{q}_2(t_0, t_1, t_2, t_3, \dots) + \epsilon^3 \mathbf{q}_3(t_0, t_1, t_2, t_3, \dots) + \dots \quad (2.3)$$

where the \mathbf{q}_i are functions of space (implicit here) and of a series of timescales t_i . In addition, the physical time is decomposed along the different timescales:

$$\partial_t = \partial_{t_0} + \epsilon \partial_{t_1} + \epsilon^2 \partial_{t_2} + \epsilon^3 \partial_{t_3} + \dots \quad (2.4)$$

By injecting eq. (2.3) and eq. (2.4) in the governing equations eq. (2.1), one can determine, at each order ϵ^i a system of equations for $\epsilon^i \mathbf{q}_i$. In order to facilitate the analytical developments, the nonlinear terms that involve trigonometric functions, (*i.e.* $\mathbf{R}(\theta)$, $\mathbf{W}(u_h, \theta, u_\theta)$ and $\mathcal{C}_L(\Lambda, \theta)$) are pre-developed in series of ϵ powers in section 2.A.

2.3.1 Zero-th order: steady solution

Gathering all terms at order ϵ^0 , it is immediately obtained that \mathbf{q}_0 is a steady solution at $U^* = U_c^*$, *i.e.*

$$\mathcal{R}(\mathbf{q}_0; U_c^*) = 0 \quad (2.5)$$

2.3.2 First order: linear stability

At order ϵ^1 , the linearized equations at $U^* = U_c^*$ are obtained, which write symbolically:

$$\mathcal{M}\partial_{t_0}\mathbf{q}_1 + \mathcal{J}(\mathbf{q}_0; U_c^*)\mathbf{q}_1 = 0 \quad (2.6)$$

where the Jacobian operator $\mathcal{J}(\mathbf{q}_0; U_c^*) = \partial\mathcal{R}/\partial\mathbf{q}|_{U^*=U_c^*}$ is defined by:

$$\mathcal{J}(\mathbf{q}_0; U_c^*)\mathbf{q}_1 = \begin{pmatrix} -u_{h1} \\ -u_{\theta 1} \\ 2\zeta_h\Omega U_c^{*-1}u_{h1} + \Omega^2 U_c^{*-2}h_1 - \tilde{m}^{-1} \left[\frac{\partial\mathcal{C}_L}{\partial\theta}\theta_1 + \frac{\partial\mathcal{C}_L}{\partial\Lambda}\Lambda_1 \right] \\ 2\zeta_p U_c^{*-1}u_{\theta 1} + U_c^{*-2}\theta_1 - \tilde{m}^{-1}r_\theta^{-2} \frac{\partial\mathcal{C}_M}{\partial\Lambda}\Lambda_1 \\ u_{\theta 1}\mathbf{e}_z \wedge \mathbf{U}_0 + \nabla\mathbf{U}_1(\mathbf{U}_0 - \mathbf{W}_0) + \nabla\mathbf{U}_0 \left(\mathbf{U}_1 - \frac{\partial\mathbf{W}}{\partial\mathbf{q}}\mathbf{q}_1 \right) - \nabla \cdot \sigma_f(p_1, \mathbf{U}_1) \\ -\nabla \cdot \mathbf{U}_1 \\ \mathbf{U}_1 - \frac{\partial\mathbf{W}}{\partial\mathbf{q}}\mathbf{q}_1 \end{pmatrix} \quad (2.7)$$

where the derivative of the grid velocity with respect to the state \mathbf{q} is such that

$$\frac{\partial\mathbf{W}}{\partial\mathbf{q}}\mathbf{q}_1 = \frac{\partial\mathbf{W}}{\partial u_h}u_{h1} + \frac{\partial\mathbf{W}}{\partial\theta}\theta_1 + \frac{\partial\mathbf{W}}{\partial u_\theta}u_{\theta 1}$$

The first order solution \mathbf{q}_1 to eq. (2.6) can be viewed as a superimposition of the so-called *global modes* $\hat{\mathbf{q}}^k$

$$\mathbf{q}_1 = \sum_k A_c^k(t_1, \dots, t_3)\hat{\mathbf{q}}^k e^{(\lambda_k + i\omega_k)t_0} + \text{c.c.} \quad (2.8)$$

with the amplitudes $A_c^k(t_1, \dots, t_3)$ depending only on the slow timescales ($t_i, i > 1$) and the $\hat{\mathbf{q}}^k$ defined as the eigenvectors of the linear stability eigenproblem [Sipp et al. 2010]:

$$\sigma_k \mathcal{M}\hat{\mathbf{q}}^k + \mathcal{J}(\mathbf{q}_0; U_c^*)\hat{\mathbf{q}}^k = 0 \quad (2.9)$$

where $\sigma_k = \lambda_k + i\omega_k$ is the corresponding eigenvalue.

By definition, at $U^* = U_c^*$, all global modes are stable ($\lambda < 0$) except for the *critical mode*, noted $\hat{\mathbf{q}}_c$, which has a purely imaginary eigenvalue $0 + i\omega_c$. As a consequence, only the critical mode is kept in the sum eq. (2.8), such that

$$\mathbf{q}_1 = A_c(t_1, \dots, t_3)\hat{\mathbf{q}}_c e^{i\omega_c t_0} + \text{c.c.} \quad (2.10)$$

2.3.3 Second order

At order ϵ^2 , the following forced linear problem is obtained:

$$\mathcal{M}\partial_{t_0}\mathbf{q}_2 + \mathcal{J}(\mathbf{q}_0; U_c^*)\mathbf{q}_2 = -\mathcal{M}\partial_{t_1}\mathbf{q}_1 + \mathcal{F}_2(\mathbf{q}_1, \mathbf{q}_1) \quad (2.11)$$

with the right-hand side forcing term:

$$\mathcal{F}_2(\mathbf{q}_a, \mathbf{q}_b) = \begin{pmatrix} 0 \\ 0 \\ \tilde{m}^{-1} \mathcal{C}_{L2,11}(\mathbf{q}_a, \mathbf{q}_b) \\ 0 \\ -u_{\theta_a} \mathbf{e}_z \wedge \mathbf{U}_b - \nabla \mathbf{U}_a \left(\mathbf{U}_b - \frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q}_b \right) + \nabla \mathbf{U}_0 \mathbf{W}_{2,11}(\mathbf{q}_a, \mathbf{q}_b) \\ 0 \\ \mathbf{W}_{2,11}(\mathbf{q}_a, \mathbf{q}_b) \end{pmatrix}$$

Note that the order ϵ^2 contributions from the lift coefficient $\mathcal{C}_{L2,11}$ and grid velocity $\mathbf{W}_{2,11}$ are bilinear functions, defined in section 2.A.

Due to the particular shape of \mathbf{q}_1 , eq. (2.10), the first forcing term in the right-hand side of eq. (2.11) is *resonant*, meaning that it vibrates at the critical eigenfrequency ω_c which is an eigenvalue of the linearized operator in the left-hand side:

$$\mathcal{M} \partial_{t_1} \mathbf{q}_1 = \mathcal{M} \partial_{t_1} A_\epsilon \hat{\mathbf{q}}_c e^{i\omega_c t_0} + \text{c.c.}$$

In general, such a resonant excitation yields a diverging response, which is not acceptable for a bounded physical weakly nonlinear solution. In order to avoid that, the Fredholm alternative [Sipp et al. 2007; Meliga et al. 2009] is invoked by stating that the resonant forcing term must be orthogonal to the kernel of the adjoint linear operator, *i.e.* to the adjoint critical mode \mathbf{q}_c^\dagger . For now, we report the exact definition of the adjoint modes to section 2.3.6 but we precise already the definition of the scalar product used to define orthogonality:

$$\langle \mathbf{q}_a, \mathbf{q}_b \rangle = h_a^* h_b + u_{h_a}^* u_{h_b} + \theta_a^* \theta_b + u_{\theta_a}^* u_{\theta_b} + \int_{\Omega_f} \{ \mathbf{U}_a^* \cdot \mathbf{U}_b + p_a^* p_b \} + \int_{\Gamma_{fs}} \{ \boldsymbol{\Lambda}_a^* \cdot \boldsymbol{\Lambda}_b \}$$

where the $*$ sign indicates the complex conjugate. Exploiting the bi-orthogonality property of the direct-adjoint basis, the Fredholm alternative gives:

$$\langle \mathbf{q}_c^\dagger, \mathcal{M} \partial_{t_1} \mathbf{q}_1 \rangle = 0 \implies \partial_{t_1} A_\epsilon = 0 \quad (2.12)$$

which turns out to simply cancel the resonant forcing term.

Plugging eq. (2.10) into the second (non-resonant) forcing term of eq. (2.11) yields four contributions:

$$\begin{aligned} \mathcal{F}_2(\mathbf{q}_1, \mathbf{q}_1) &= |A_\epsilon|^2 (\mathcal{F}_2(\hat{\mathbf{q}}_c, \hat{\mathbf{q}}_c^*) + \mathcal{F}_2(\hat{\mathbf{q}}_c^*, \hat{\mathbf{q}}_c)) \\ &\quad + \left(A_\epsilon^2 \mathcal{F}_2(\hat{\mathbf{q}}_c, \hat{\mathbf{q}}_c) e^{2i\omega_c t_0} + \text{c.c.} \right) \end{aligned}$$

Due to linearity, the second-order solution can then be written as the sum of the individual responses to each forcing term:

$$\mathbf{q}_2 = |A_\epsilon|^2 \hat{\mathbf{q}}_2^{|A|^2} + \left(A_\epsilon^2 \hat{\mathbf{q}}_2^{A^2} e^{2i\omega_c t_0} + \text{c.c.} \right) \quad (2.13)$$

with $\hat{\mathbf{q}}_2^{|A|^2}$ and $\hat{\mathbf{q}}_2^{A^2}$, two time-independent functions verifying

$$\mathcal{J}(\mathbf{q}_0; U_c^*) \hat{\mathbf{q}}_2^{|A|^2} = \mathcal{F}_2(\hat{\mathbf{q}}_c, \hat{\mathbf{q}}_c^*) + \mathcal{F}_2(\hat{\mathbf{q}}_c^*, \hat{\mathbf{q}}_c) \quad (2.14a)$$

$$[2i\omega_c \mathcal{M} + \mathcal{J}(\mathbf{q}_0; U_c^*)] \hat{\mathbf{q}}_2^{A^2} = \mathcal{F}_2(\hat{\mathbf{q}}_c, \hat{\mathbf{q}}_c) \quad (2.14b)$$

2.3.4 Third order: Stuart–Landau equation

At order ϵ^3 , we obtain another forced linear system:

$$\begin{aligned} \mathcal{M}\partial_{t_0}\mathbf{q}_3 + \mathcal{J}(\mathbf{q}_0; U_c^*)\mathbf{q}_3 = & -\mathcal{M}\partial_{t_2}\mathbf{q}_1 - \mathcal{M}\partial_{t_1}\mathbf{q}_2 - \Delta_{U^*}\mathcal{L}\mathbf{q}_1 \\ & + \mathcal{F}_{3,21}(\mathbf{q}_2, \mathbf{q}_1) + \mathcal{F}_{3,111}(\mathbf{q}_1, \mathbf{q}_1, \mathbf{q}_1) \end{aligned} \quad (2.15)$$

where

$$\mathcal{L} = \begin{pmatrix} 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ \Omega^2 & 0 & \zeta_h \Omega U_c^* & 0 & 0 & 0 & 0 \\ 0 & 1 & 0 & \zeta_p U_c^* & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \\ 0 & 0 & 0 & 0 & 0 & 0 & 0 \end{pmatrix}$$

with the following forcing terms:

$$\mathcal{F}_{3,21}(\mathbf{q}_a, \mathbf{q}_b) = \begin{pmatrix} 0 \\ 0 \\ \tilde{m}^{-1} \mathcal{C}_{L3,21}(\theta_a, \theta_b) \\ 0 \\ -u_{\theta_a} \mathbf{e}_z \wedge \mathbf{U}_b - u_{\theta_b} \mathbf{e}_z \wedge \mathbf{U}_a + \dots \\ \dots - \nabla \mathbf{U}_a \left(\mathbf{U}_b - \frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q}_b \right) - \nabla \mathbf{U}_b \left(\mathbf{U}_a - \frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q}_a \right) + \nabla \mathbf{U}_0 \mathbf{W}_{3,21}(\mathbf{q}_a, \mathbf{q}_b) \\ 0 \\ \mathbf{W}_{3,21}(\mathbf{q}_a, \mathbf{q}_b) \end{pmatrix}$$

and

$$\mathcal{F}_{3,111}(\mathbf{q}_a, \mathbf{q}_b, \mathbf{q}_c) = \begin{pmatrix} 0 \\ 0 \\ \tilde{m}^{-1} \mathcal{C}_{L3,111}(\theta_a, \theta_b, \theta_c) \\ 0 \\ \nabla \mathbf{U}_a \mathbf{W}_{2,11}(\mathbf{q}_b, \mathbf{q}_c) + \nabla \mathbf{U}_0 \mathbf{W}_{3,111}(\mathbf{q}_a, \mathbf{q}_b, \mathbf{q}_c) \\ 0 \\ \mathbf{W}_{3,111}(\mathbf{q}_a, \mathbf{q}_b, \mathbf{q}_c) \end{pmatrix}$$

As for the second order forcing, the order ϵ^3 contributions from the lift coefficient $\mathcal{C}_{L3,21}$ and $\mathcal{C}_{L3,111}$, and from the grid velocity $\mathbf{W}_{3,21}$ and $\mathbf{W}_{3,111}$ are defined in section 2.A.

Using eq. (2.12) in eq. (2.13), we obtain $\partial_{t_1}\mathbf{q}_2 = 0$, which directly cancels the second term in the right-hand side of eq. (2.15). In the four remaining forcing terms, we use the particular forms of \mathbf{q}_1 (eq. (2.10)) and \mathbf{q}_2 (eq. (2.13)). It is then observed that the forcing may be split between resonant and non-resonant terms. In particular, the resonant part writes

$$-\mathcal{M}\partial_{t_2}A_\epsilon \hat{\mathbf{q}}_c - \Delta_{U^*}A_\epsilon \mathcal{L}\hat{\mathbf{q}}_c + |A_\epsilon|^2 A_\epsilon \mathcal{F}_{\text{res}}(\hat{\mathbf{q}}_c, \hat{\mathbf{q}}_2^{|A|^2}, \hat{\mathbf{q}}_2^{A^2})$$

with

$$\begin{aligned} \mathcal{F}_{\text{res}}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_2^{A^2}, \widehat{\mathbf{q}}_2^{A^2}) &= \mathcal{F}_{3,21}(\widehat{\mathbf{q}}_2^{A^2}, \widehat{\mathbf{q}}_c) + \mathcal{F}_{3,21}(\widehat{\mathbf{q}}_2^{A^2}, \widehat{\mathbf{q}}_c^*) \\ &+ \mathcal{F}_{3,111}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c^*) + \mathcal{F}_{3,111}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c^*, \widehat{\mathbf{q}}_c) + \mathcal{F}_{3,111}(\widehat{\mathbf{q}}_c^*, \widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c) \end{aligned} \quad (2.16)$$

As already done at second order, the Fredholm alternative must again be invoked to ensure a well-behaved expansion despite the presence of the above resonant forcing. It leads to the following compatibility condition:

$$\left\langle \mathbf{q}_c^\dagger, -\mathcal{M} \partial_{t_2} A_\epsilon \widehat{\mathbf{q}}_c - \Delta_{U^*} A_\epsilon \mathcal{L} \widehat{\mathbf{q}}_c + |A_\epsilon|^2 A_\epsilon \mathcal{F}_{\text{res}}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_2^{A^2}, \widehat{\mathbf{q}}_2^{A^2}) \right\rangle = 0 \quad (2.17)$$

The Stuart–Landau equation — also referred to as *amplitude equation* — of the flutter instability is finally deduced from eq. (2.17):

$$\partial_t A = \alpha \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right) A + \beta |A|^2 A \quad (2.18)$$

after defining a rescaled amplitude $A = \epsilon A_\epsilon$ and using eq. (2.2) to re-introduce U^* as the bifurcation parameter. The normal form coefficients α and β only depend on the solutions previously computed at orders ϵ^0 to ϵ^2 and on the adjoint critical mode. They read:

$$\alpha = -\frac{1}{U_c^{*2}} \frac{\langle \mathbf{q}_c^\dagger, \mathcal{L} \widehat{\mathbf{q}}_c \rangle}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \widehat{\mathbf{q}}_c \rangle} \quad \beta = \frac{\langle \mathbf{q}_c^\dagger, \mathcal{F}_{\text{res}}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_2^{A^2}, \widehat{\mathbf{q}}_2^{A^2}) \rangle}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \widehat{\mathbf{q}}_c \rangle} \quad (2.19)$$

2.3.5 Solutions of the Stuart–Landau equation

To solve the Stuart–Landau equation, we introduce the time-dependent module $R(t)$ and phase $\phi(t)$ of the complex amplitude $A(t) = R(t)e^{i\phi(t)}$. The latter expression is injected in eq. (2.18), the real and imaginary part are separated and we obtain — supposing R is different from the trivial zero solution — a coupled system of equations for $R(t)$ and $\phi(t)$:

$$\begin{aligned} \partial_t R &= \Re \epsilon (\alpha) \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right) R + \Re \epsilon (\beta) R^3 \\ \partial_t \phi &= \Im \epsilon (\alpha) \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right) + \Im \epsilon (\beta) R^2 \end{aligned} \quad (2.20)$$

Linear solutions Let us first consider the case where R is sufficiently small so that the nonlinear term in eq. (2.18) (or eq. (2.20)) is negligible. From eq. (2.20), we immediately obtain:

$$R(t; U^*) = C_1^{\text{te}} \exp \left[\underbrace{\Re(\alpha) \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right)}_{:= \lambda_{\text{lin}}(U^*)} t \right]$$

$$\phi(t; U^*) = \underbrace{\Im(\alpha) \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right)}_{:= \omega_{\text{lin}}(U^*)} t + C_2^{\text{te}}$$

with C_1^{te} and C_2^{te} two unspecified (they would be specified by initial conditions) real-valued constants. Note in passing the dependence of R and ϕ in the parameter U^* . Using the above expressions in combination with eq. (2.3) and eq. (2.10), the first-order response — in the pitching degree of freedom, for example — reads:

$$\theta(t; U^*) = 2C_1^{\text{te}} |\hat{\theta}_c| \cos [(\omega_c + \omega_{\text{lin}}(U^*)) t + C_2^{\text{te}} + \phi_{\theta_c}] e^{\lambda_{\text{lin}}(U^*) t} \quad (2.21)$$

where $|\hat{\theta}_c|$ and ϕ_{θ_c} are the modulus and phase of the pitching degree of freedom in the critical mode, $\hat{\mathbf{q}}_c$. First, we note that, C_1^{te} being unspecified, the amplitude of the oscillation is free, and would only be fixed by additional initial conditions. Then, in the cos function, we see that ω_{lin} is a correction of the oscillation frequency when U^* departs from the critical flutter threshold U_c^* . As for λ_{lin} , it is seen to predict the growth rate of the flutter mode when $U^* \neq U_c^*$. More precisely, for $\Re(\alpha) < 0$ (resp. $\Re(\alpha) > 0$), λ_{lin} is positive — *i.e.* oscillations grow exponentially — when $U^* > U_c^*$ (resp. $U^* < U_c^*$). For an instability like flutter, which typically occurs when increasing the reduced velocity *i.e.* when $U^* > U_c^*$, we thus necessarily have

$$\Re(\alpha) < 0 \quad (2.22)$$

Nonlinear periodic solutions In this manuscript we are mostly interested in the permanent nonlinear flutter regimes, *i.e.* we do not study the transients. As a consequence we seek after periodic solutions of the Stuart–Landau equation. The periodicity hypothesis allows to write that the modulus R does not depend on time. The $\partial_t R$ term in eq. (2.20) is thus zero so that we directly obtain

$$R(U^*)^2 = \frac{\Re(\alpha)}{\Re(\beta)} \left(1 - \frac{U_c^{*2}}{U^{*2}} \right) \quad (2.23a)$$

$$\phi(t; U^*) = \underbrace{\left(\Im(\beta) \frac{\Re(\alpha)}{\Re(\beta)} - \Im(\alpha) \right)}_{:= \omega_{\text{nl}}(U^*)} \left(1 - \frac{U_c^{*2}}{U^{*2}} \right) t + C^{\text{te}} \quad (2.23b)$$

where C^{te} is an arbitrary real-valued constant that fixes the LCO phase. Using the above expressions in combination with eq. (2.3) and eq. (2.10), we again reconstruct

the pitching first-order response as:

$$\theta(t; U^*) = 2\sqrt{\frac{\Re(\alpha)}{\Re(\beta)} \left(1 - \frac{U_c^{*2}}{U^{*2}}\right)} |\hat{\theta}_c| \cos[(\omega_c + \omega_{nl}(U^*))t + C^{te} + \phi_{\theta_c}] \quad (2.24)$$

where $|\hat{\theta}_c|$ and ϕ_{θ_c} are the modulus and phase of the pitching degree of freedom in the critical mode, $\hat{\mathbf{q}}_c$. Similarly to the linear case detailed in the previous paragraph, the pulsation ω_{nl} represents a correction to the critical flutter frequency ω_c when departing from U_c^* . This time, the correction is nonlinear: it can be seen as the sum of the linear correction ω_{lin} (that depends on α only) introduced in the previous paragraph, and of a nonlinear contribution that depends on β .

Recalling that in the case of flutter we have $\Re(\alpha) < 0$ (eq. (2.22)), the sign of $\Re(\beta)$ is of particular interest. Indeed, according to eq. (2.23a), R can be properly defined for U^* above U_c^* only if $\Re(\beta)/\Re(\alpha) > 0$, i.e. $\Re(\beta) < 0$. In this case the bifurcation is supercritical, as shown in the upper left cell of table 2.1. On the contrary, if $\Re(\beta)/\Re(\alpha) < 0$ i.e. $\Re(\beta) > 0$, then R can be properly defined only for U^* below U_c^* , which means that the bifurcation is subcritical (lower left cell of table 2.1). For completeness, we have also represented in the right column of table 2.1 the two possible cases when $\Re(\alpha) > 0$. Overall, the sign of $\Re(\beta)$ alone gives the nature of the bifurcation.

	$\Re(\alpha) < 0$	$\Re(\alpha) > 0$
$\Re(\beta) < 0$	<p>supercritical</p>	<p>supercritical</p>
$\Re(\beta) > 0$	<p>subcritical</p>	<p>subcritical</p>

TABLE 2.1: Nature of the bifurcation depending on the signs of $\Re(\alpha)$ and $\Re(\beta)$. In each cell of the table, an amplitude- $vs-U^*$ diagram is shown. The central horizontal branch always represents the trivial $R = 0$ solution. By convention, the solid (resp. dashed) lines represent stable (resp. unstable) solutions.

2.3.6 Continuous adjoint equations

At order 2 and 3, the adjoint \mathbf{q}_c^\dagger of the critical eigenmode is projected onto the resonant forcings to obtain compatibility conditions. In this paragraph we provide the definition of the adjoint modes. First, we formally define the adjoint operators

$\mathcal{J}^\dagger(\mathbf{q}_0)$ and \mathcal{M}^\dagger of $\mathcal{J}(\mathbf{q}_0)$ and \mathcal{M} through the classical variational statement :

$$\left\langle \left(\sigma^* \mathcal{M}^\dagger + \mathcal{J}^\dagger(\mathbf{q}_0; U_c^*) \right) \mathbf{q}^\dagger, \mathbf{q} \right\rangle = \left\langle \mathbf{q}^\dagger, (\sigma \mathcal{M} + \mathcal{J}(\mathbf{q}_0; U_c^*)) \mathbf{q} \right\rangle \quad \forall \mathbf{q}, \mathbf{q}^\dagger \quad (2.25)$$

with $\mathbf{q}^\dagger = (h^\dagger, \theta^\dagger, u_h^\dagger, u_\theta^\dagger, \mathbf{U}^\dagger, p^\dagger, \boldsymbol{\Lambda}^\dagger)^\top$ the adjoint variable to \mathbf{q} . We recall that the scalar product defines as $\langle \mathbf{q}_a, \mathbf{q}_b \rangle = h_a^* h_b + u_{h_a}^* u_{h_b} + \theta_a^* \theta_b + u_{\theta_a}^* u_{\theta_b} + \int_{\Omega_f} \{ \mathbf{U}_a^* \cdot \mathbf{U}_b + p_a^* p_b \} + \int_{\Gamma_{fs}} \{ \boldsymbol{\Lambda}_a^* \cdot \boldsymbol{\Lambda}_b \}$. Then, the adjoint \mathbf{q}^\dagger of the (direct) eigenpair $(\sigma, \hat{\mathbf{q}})$ is defined as the eigenvector of the adjoint operator with eigenvalue σ^* :

$$\sigma^* \mathcal{M}^\dagger \mathbf{q}^\dagger + \mathcal{J}^\dagger(\mathbf{q}_0; U_c^*) \mathbf{q}^\dagger = 0 \quad (2.26)$$

Let us now derive an explicit strong form for eq. (2.26). To that end, we use the definition eq. (2.25) and explicitly develop its right-hand side which represents a weak formulation for the direct linearized eigenproblem eq. (2.9). It is obtained using the same steps³ as detailed in section 1.2.1 for the nonlinear problem such that only the result is stated here:

$$\begin{aligned} & \left\langle \left(\sigma^* \mathcal{M}^\dagger + \mathcal{J}^\dagger(\mathbf{q}_0; U_c^*) \right) \mathbf{q}^\dagger, \mathbf{q} \right\rangle = \left\langle \mathbf{q}^\dagger, (\sigma \mathcal{M} + \mathcal{J}(\mathbf{q}_0; U_c^*)) \mathbf{q} \right\rangle = \\ & \quad \{ \sigma h - u_h \} h^{\dagger*} \\ & + \{ \sigma \theta - u_\theta \} \theta^{\dagger*} \\ & + \left\{ \sigma u_h + 2\zeta_h \frac{\Omega}{U_c^*} u_h + \frac{\Omega^2}{U_c^{*2}} h - \frac{1}{\tilde{m}} \left[\frac{\partial \mathcal{C}_L}{\partial \theta} \theta + \frac{\partial \mathcal{C}_L}{\partial \boldsymbol{\Lambda}} \boldsymbol{\Lambda} \right] \right\} u_h^{\dagger*} \\ & + \left\{ \sigma u_\theta + 2\zeta_p \frac{1}{U_c^*} u_\theta + \frac{1}{U_c^{*2}} \theta - \frac{1}{r_\theta^2 \tilde{m}} \frac{\partial \mathcal{C}_M}{\partial \boldsymbol{\Lambda}} \boldsymbol{\Lambda} \right\} u_\theta^{\dagger*} \\ & + \int_{\Omega_f} \left\{ \sigma \mathbf{U} + u_\theta \mathbf{e}_z \wedge \mathbf{U}_0 + \nabla \mathbf{U} (\mathbf{U}_0 - \mathbf{W}_0) + \nabla \mathbf{U}_0 \left(\mathbf{U} - \left(\frac{\partial \mathbf{W}}{\partial u_h} u_h + \frac{\partial \mathbf{W}}{\partial \theta} \theta + \frac{\partial \mathbf{W}}{\partial u_\theta} u_\theta \right) \right) \right\} \cdot \mathbf{U}^{\dagger*} \\ & + \int_{\Omega_f} \sigma_f(\mathbf{U}, p) : \nabla \mathbf{U}^{\dagger*} - \int_{\Omega_f} \nabla \cdot \mathbf{U} p^{\dagger*} + \int_{\Gamma_{fs}} \boldsymbol{\Lambda} \cdot \mathbf{U}^{\dagger*} \\ & + \int_{\Gamma_{fs}} \left\{ \mathbf{U} - \left(\frac{\partial \mathbf{W}}{\partial u_h} u_h + \frac{\partial \mathbf{W}}{\partial \theta} \theta + \frac{\partial \mathbf{W}}{\partial u_\theta} u_\theta \right) \right\} \cdot \boldsymbol{\Lambda}^{\dagger*} \end{aligned}$$

where it is recalled that $\boldsymbol{\Lambda} = \sigma_f(p, \mathbf{U}) \mathbf{N}_s$ (cf eq. (1.25)). We need several ingredients to reorganize this weak formulation. First, we notice that $(\nabla \mathbf{U}_0 \mathbf{U}) \cdot \mathbf{U}^{\dagger*} = \nabla \mathbf{U}_0^\top \mathbf{U}^{\dagger*} \cdot \mathbf{U}$ and define $\mathbf{l}_\Lambda = 2 \mathbf{R}^\top(\theta_0) \mathbf{e}_y$ and $\mathbf{m}_\Lambda = 2 \mathbf{e}_z \wedge \mathbf{X}$ such that

$$\begin{aligned} \frac{\partial \mathcal{C}_L}{\partial \boldsymbol{\Lambda}} \boldsymbol{\Lambda} &= 2 \int_{\Gamma_{fs}} \left(\mathbf{R}^\top(\theta_0) \mathbf{e}_y \right) \cdot \boldsymbol{\Lambda} = \int_{\Gamma_{fs}} \mathbf{l}_\Lambda \cdot \boldsymbol{\Lambda} \\ \frac{\partial \mathcal{C}_M}{\partial \boldsymbol{\Lambda}} \boldsymbol{\Lambda} &= 2 \int_{\Gamma_{fs}} (\mathbf{X} \wedge \boldsymbol{\Lambda}) \cdot \mathbf{e}_z = \int_{\Gamma_{fs}} \mathbf{m}_\Lambda \cdot \boldsymbol{\Lambda} \end{aligned}$$

Then, we integrate by part the ‘‘perturbation convection by the baseflow’’ term

$$\int_{\Omega_f} \{ \nabla \mathbf{U} (\mathbf{U}_0 - \mathbf{W}_0) \} \cdot \mathbf{U}^{\dagger*} = - \int_{\Omega_f} \{ \nabla \mathbf{U}^{\dagger*} (\mathbf{U}_0 - \mathbf{W}_0) \} \cdot \mathbf{U} + \int_{\partial \Omega_f} \{ ((\mathbf{U}_0 - \mathbf{W}_0) \cdot \mathbf{N}_f) \mathbf{U}^{\dagger*} \} \cdot \mathbf{U}$$

where one notices that the boundary integral on $\partial \Omega_f$ can be reduced to Γ_{out} because it is zero both on Γ_{in} (due to $\mathbf{U} = 0$) and on Γ_{fs} (due to $\mathbf{U}_0 - \mathbf{W}_0 = 0$). Finally gathering

³Namely, integration by part of the viscous stress term, incorporation of the farfield boundary conditions and of the definition of $\boldsymbol{\Lambda}$

the different terms, we obtain a weak formulation where \mathbf{q} is viewed as the test function:

$$\begin{aligned}
& \left\langle \left(\sigma^* \mathcal{M}^\dagger + \mathcal{J}^\dagger(\mathbf{q}_0; U_c^*) \right) \mathbf{q}^\dagger, \mathbf{q} \right\rangle = \\
& \left\{ \sigma h^{\dagger*} + \frac{\Omega^2}{U_c^{*2}} u_h^{\dagger*} \right\} h \\
& + \left\{ \sigma \theta^{\dagger*} + \frac{1}{U_c^{*2}} u_\theta^{\dagger*} - \frac{1}{\tilde{m}} \frac{\partial \mathcal{C}_L}{\partial \theta} u_h^{\dagger*} - \int_{\Omega_f} \left(\nabla \mathbf{U}_0^T \mathbf{U}^{\dagger*} \right) \cdot \frac{\partial \mathbf{W}}{\partial \theta} - \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^{\dagger*} \cdot \frac{\partial \mathbf{W}}{\partial \theta} \right\} \theta \\
& + \left\{ \sigma u_h^{\dagger*} + 2\zeta_h \frac{\Omega}{U_c^*} u_h^{\dagger*} - h^{\dagger*} - \int_{\Omega_f} \left(\nabla \mathbf{U}_0^T \mathbf{U}^{\dagger*} \right) \cdot \frac{\partial \mathbf{W}}{\partial u_h} - \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^{\dagger*} \cdot \frac{\partial \mathbf{W}}{\partial u_h} \right\} u_h \\
& + \left\{ \sigma u_\theta^{\dagger*} + 2\zeta_p \frac{1}{U_c^*} u_\theta^{\dagger*} - \theta^{\dagger*} + \int_{\Omega_f} (\mathbf{e}_z \wedge \mathbf{U}_0) \cdot \mathbf{U}^{\dagger*} - \int_{\Omega_f} \left(\nabla \mathbf{U}_0^T \mathbf{U}^{\dagger*} \right) \cdot \frac{\partial \mathbf{W}}{\partial u_\theta} - \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^{\dagger*} \cdot \frac{\partial \mathbf{W}}{\partial u_\theta} \right\} u_\theta \\
& + \int_{\Omega_f} \left\{ \sigma \mathbf{U}^{\dagger*} - \nabla \mathbf{U}^{\dagger*} (\mathbf{U}_0 - \mathbf{W}_0) + \nabla \mathbf{U}_0^T \mathbf{U}^{\dagger*} \right\} \cdot \mathbf{U} + \int_{\Gamma_{out}} \left\{ ((\mathbf{U}_0 - \mathbf{W}_0) \cdot \mathbf{N}_f) \mathbf{U}^{\dagger*} \right\} \cdot \mathbf{U} \\
& + \int_{\Omega_f} \sigma_f(\mathbf{U}^{\dagger*}, p^{\dagger*}) : \nabla \mathbf{U} - \int_{\Omega_f} \nabla \cdot \mathbf{U}^{\dagger*} p + \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^{\dagger*} \cdot \mathbf{U} \\
& + \int_{\Gamma_{fs}} \left\{ \mathbf{U}^{\dagger*} - \frac{1}{\tilde{m}} \mathbf{1}_\Lambda u_h^{\dagger*} - \frac{1}{r_\theta^2 \tilde{m}} \mathbf{m}_\Lambda u_\theta^{\dagger*} \right\} \cdot \boldsymbol{\Lambda}
\end{aligned}$$

A last integration by part of the diffusion term, and an adequate cancellation of all boundary terms leads to the strong form of the continuous adjoint equations:

$$\begin{aligned}
& \sigma^* h^\dagger + \frac{\Omega^2}{U_c^{*2}} u_h^\dagger = 0 \\
& \sigma^* \theta^\dagger + \frac{1}{U_c^{*2}} u_\theta^\dagger - \frac{1}{\tilde{m}} \frac{\partial \mathcal{C}_L}{\partial \theta} u_h^\dagger - \int_{\Omega_f} \left(\nabla \mathbf{U}_0^T \mathbf{U}^\dagger \right) \cdot \frac{\partial \mathbf{W}}{\partial \theta} - \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^\dagger \cdot \frac{\partial \mathbf{W}}{\partial \theta} = 0 \\
& \sigma^* u_h^\dagger + 2\zeta_h \frac{\Omega}{U_c^*} u_h^\dagger - h^\dagger - \int_{\Omega_f} \left(\nabla \mathbf{U}_0^T \mathbf{U}^\dagger \right) \cdot \frac{\partial \mathbf{W}}{\partial u_h} - \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^\dagger \cdot \frac{\partial \mathbf{W}}{\partial u_h} = 0 \\
& \sigma^* u_\theta^\dagger + 2\zeta_p \frac{1}{U_c^*} u_\theta^\dagger - \theta^\dagger + \int_{\Omega_f} (\mathbf{e}_z \wedge \mathbf{U}_0) \cdot \mathbf{U}^\dagger - \int_{\Omega_f} \left(\nabla \mathbf{U}_0^T \mathbf{U}^\dagger \right) \cdot \frac{\partial \mathbf{W}}{\partial u_\theta} - \int_{\Gamma_{fs}} \boldsymbol{\Lambda}^\dagger \cdot \frac{\partial \mathbf{W}}{\partial u_\theta} = 0 \\
& \sigma^* \mathbf{U}^\dagger - \nabla \mathbf{U}^\dagger (\mathbf{U}_0 - \mathbf{W}_0) + \nabla \mathbf{U}_0^T \mathbf{U}^\dagger - \nabla \cdot \sigma_f(\mathbf{U}^\dagger, p^\dagger) = 0 \quad \text{in } \Omega_f \\
& -\nabla \cdot \mathbf{U}^\dagger = 0 \quad \text{in } \Omega_f \\
& \mathbf{U}^\dagger - \frac{1}{\tilde{m}} \mathbf{1}_\Lambda u_h^{\dagger*} - \frac{1}{r_\theta^2 \tilde{m}} \mathbf{m}_\Lambda u_\theta^{\dagger*} = 0 \quad \text{on } \Gamma_{fs}
\end{aligned} \tag{2.27}$$

with the adjoint inflow and outflow boundary conditions

$$\begin{aligned}
& \mathbf{U}^\dagger = 0 \quad \text{on } \Gamma_{in} \\
& \sigma_f(\mathbf{U}^\dagger, p^\dagger) \mathbf{n}_f + ((\mathbf{U}_0 - \mathbf{W}_0) \cdot \mathbf{N}_f) \mathbf{U}^\dagger = 0 \quad \text{on } \Gamma_{out}
\end{aligned}$$

and the definition of the adjoint Lagrange multiplier on Γ_{fs} : $\boldsymbol{\Lambda}^\dagger = \sigma_f(\mathbf{U}^\dagger, p^\dagger) \mathbf{N}_s$.

2.3.7 Summary

In the end, the weakly nonlinear solver detailed in this chapter allows to compute the approximate *nonlinear* behavior of the typical section at the cost of solving the

following sequence of — mostly *linear*, at the exception of eq. (2.28a) — problems:

$$\mathcal{R}(\mathbf{q}_0; U_c^*) = 0 \quad (2.28a)$$

$$(\mathbf{i}\omega_c \mathcal{M} + \mathcal{J}(\mathbf{q}_0; U_c^*)) \widehat{\mathbf{q}}_c = 0 \quad (2.28b)$$

$$(-\mathbf{i}\omega_c \mathcal{M}^\dagger + \mathcal{J}^\dagger(\mathbf{q}_0; U_c^*)) \mathbf{q}_c^\dagger = 0 \quad (2.28c)$$

$$\mathcal{J}(\mathbf{q}_0; U_c^*) \widehat{\mathbf{q}}_2^{|A|^2} - \mathcal{F}_2(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c^*) - \mathcal{F}_2(\widehat{\mathbf{q}}_c^*, \widehat{\mathbf{q}}_c) = 0 \quad (2.28d)$$

$$[2\mathbf{i}\omega_c \mathcal{M} + \mathcal{J}(\mathbf{q}_0; U_c^*)] \widehat{\mathbf{q}}_2^{A^2} - \mathcal{F}_2(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c) = 0 \quad (2.28e)$$

of which the Stuart–Landau equation is deduced:

$$\partial_t A = \alpha \left(\frac{U_c^{*2}}{U_c^{*2}} - 1 \right) A + \beta |A|^2 A$$

$$\alpha = -\frac{1}{U_c^{*2}} \frac{\langle \mathbf{q}_c^\dagger, \mathcal{L} \widehat{\mathbf{q}}_c \rangle}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \widehat{\mathbf{q}}_c \rangle} \quad \beta = \frac{\langle \mathbf{q}_c^\dagger, \mathcal{F}_{\text{res}}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_2^{|A|^2}, \widehat{\mathbf{q}}_2^{A^2}) \rangle}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \widehat{\mathbf{q}}_c \rangle}$$

Obviously, solving for a periodic flutter solution with this analysis comes at much lower numerical cost than performing a classical time-integration of the governing equations (see section 1.3) or even a Harmonic-Balance type calculation (see chapter 3). From a more theoretical point of view, the weakly nonlinear analysis provides some insights on how each order generates the next one through the forcing terms in the right-hand sides of eq. (2.11) and (2.15). Furthermore, we will see in chapter 5 that the semi-analytical expression of the normal form coefficient β allows to decompose it in individual contributions coming from different sources of nonlinearity, ultimately shedding some light on the physical origin of the sub- or supercritical nature of the flutter bifurcation.

2.4 Numerical solution

In this section, we present the numerical approach used to obtain approximate solutions to the sequence of problems, eq. (2.28). First, the equations are discretized in space using a SUPG-stabilized finite element method and the algorithms used to solve the discrete problems are briefly presented. Then, a mesh adaptation framework is proposed with the objective of avoiding a manual mesh design for each new Reynolds number investigated. This will come in particularly handy for the wide parametric explorations exposed in chapter 4 and chapter 5. The section ends with some validation results.

2.4.1 SUPG-stabilized finite elements discretization

The SUPG-stabilized finite element method [Brooks et al. 1982] is used to discretize in space the sequence of problems eq. (2.28). The weak statements that define the

discrete problems then read

$$\langle \boldsymbol{\psi} + \tau \mathcal{L}_{\text{adv}}[\mathbf{q}_0](\boldsymbol{\psi}), \mathcal{R}(\mathbf{q}_0) \rangle = 0 \quad (2.29a)$$

$$\langle \boldsymbol{\psi} + \tau \mathcal{L}_{\text{adv}}[\mathbf{q}_0](\boldsymbol{\psi}), (\mathbf{i}\omega_c \mathcal{M} + \mathcal{J}(\mathbf{q}_0)) \widehat{\mathbf{q}}_c \rangle = 0 \quad (2.29b)$$

$$\langle \boldsymbol{\psi} - \tau \mathcal{L}_{\text{adv}}[\mathbf{q}_0](\boldsymbol{\psi}), (-\mathbf{i}\omega_c \mathcal{M}^\dagger + \mathcal{J}^\dagger(\mathbf{q}_0)) \mathbf{q}_c^\dagger \rangle = 0 \quad (2.29c)$$

$$\langle \boldsymbol{\psi} + \tau \mathcal{L}_{\text{adv}}[\mathbf{q}_0](\boldsymbol{\psi}), \mathcal{J}(\mathbf{q}_0) \widehat{\mathbf{q}}_2^{|A|^2} - \mathcal{F}_2(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c^*) - \mathcal{F}_2(\widehat{\mathbf{q}}_c^*, \widehat{\mathbf{q}}_c) \rangle = 0 \quad (2.29d)$$

$$\langle \boldsymbol{\psi} + \tau \mathcal{L}_{\text{adv}}[\mathbf{q}_0](\boldsymbol{\psi}), [2\mathbf{i}\omega_c \mathcal{M} + \mathcal{J}(\mathbf{q}_0)] \widehat{\mathbf{q}}_2^{A^2} - \mathcal{F}_2(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c) \rangle = 0 \quad (2.29e)$$

with $\boldsymbol{\psi} = (\psi^h, \psi^{u_h}, \psi^\theta, \psi^{u_\theta}, \boldsymbol{\psi}^U, \psi^p, \boldsymbol{\psi}^\Lambda)^\top$ the (Galerkin) test function taken in $\mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathbb{R} \times \mathcal{V}^U \times \mathcal{V}^p \times \mathcal{V}^\Lambda$, where the functional spaces \mathcal{V}^U , \mathcal{V}^p and \mathcal{V}^Λ are defined in section 1.2.1. The SUPG stabilization is introduced thanks to the operator $\mathcal{L}_{\text{adv}}[\mathbf{q}_0](\boldsymbol{\psi}) = (0, 0, 0, 0, \nabla \boldsymbol{\psi}^U (\mathbf{U}_0 - \mathbf{W}_0), 0, 0)^\top$ which convects the fluid velocity test function with convection velocity $\mathbf{U}_0 - \mathbf{W}_0$. Obviously, the SUPG contribution is zero for all equations, except the fluid momentum equation. Note that in the adjoint problem, the convection velocity is $-(\mathbf{U}_0 - \mathbf{W}_0)$, as visible in the strong form eq. (2.27). This is the reason why a minus sign is used in front of the SUPG contribution in eq. (2.29c). More details about the SUPG formulation (definition of the stabilization parameter τ , etc) are provided in section 1.2.2 where it was applied to discretize the fully nonlinear problem.

Discrete problems The discrete problems corresponding to eq. (2.29) are obtained by decomposing the test and trial functions in a finite element basis, as already done in section 1.2.2. The Taylor–Hood ($\mathbb{P}_2, \mathbb{P}_1$) finite element pair is used for (\mathbf{U}, p) , and \mathbb{P}_1 elements are used for Λ . The discrete problems finally write

$$\mathbf{R}(\mathbf{q}_0; U_c^*) = 0 \quad (2.30a)$$

$$(\mathbf{i}\omega_c \mathbf{M}(\mathbf{q}_0) + \mathbf{J}(\mathbf{q}_0; U_c^*)) \widehat{\mathbf{q}}_c = 0 \quad (2.30b)$$

$$(-\mathbf{i}\omega_c \mathbf{M}^\dagger(\mathbf{q}_0) + \mathbf{J}^\dagger(\mathbf{q}_0; U_c^*)) \mathbf{q}_c^\dagger = 0 \quad (2.30c)$$

$$\mathbf{J}(\mathbf{q}_0; U_c^*) \widehat{\mathbf{q}}_2^{|A|^2} - \mathbf{F}_2(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c^*) - \mathbf{F}_2(\widehat{\mathbf{q}}_c^*, \widehat{\mathbf{q}}_c) = 0 \quad (2.30d)$$

$$(2\mathbf{i}\omega_c \mathbf{M}(\mathbf{q}_0) + \mathbf{J}(\mathbf{q}_0; U_c^*)) \widehat{\mathbf{q}}_2^{A^2} - \mathbf{F}_2(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_c) = 0 \quad (2.30e)$$

The discrete nonlinear residual \mathbf{R} and the mass matrix $\mathbf{M}(\mathbf{q}_0)$ were already defined in details — *i.e.* block-by-block — in eq. (1.32). For the Jacobian matrix $\mathbf{J}(\mathbf{q}_0; U_c^*)$, we provide in section 2.B the detailed block definition that results from the weak form eq. (2.29b). The block definitions of the other discrete operators can be similarly deduced from the corresponding weak forms and are not further detailed.

“Stabilize last” vs “stabilize first”: a brief discussion One may notice that a choice was made to stabilize the weakly nonlinear equations, while “weakly nonlinearizing” the stabilized nonlinear equations would have been a possible alternative. From now, let us refer to these two approaches as weakly-nonlinearize-then-stabilize and stabilize-then-weakly-nonlinearize, respectively. The choice of the weakly-nonlinearize-then-stabilize approach is motivated by the two following reasons. First, by stabilizing last, we avoid overburdening the already cumbersome weakly nonlinear derivations with additional highly nonlinear stabilization terms.




FIGURE 2.1: Adjoint critical eigenvectors obtained with stabilize-then-adjointize (a) *vs* adjointize-then-stabilize (b) approaches. Oscillations of the size of the mesh elements are observed with the first approach while totally avoided with the second.

Second, we also settle the issue of *adjoint consistency*. Indeed, it is known that the weak form resulting from SUPG stabilization is not adjoint consistent [Collis et al. 2002; Cyr et al. 2014], meaning that a stabilize-then-adjointize strategy⁴ does not produce a consistent discretization of the continuous adjoint operator, \mathcal{J}^\dagger . In contrast, the adjointize-then-stabilize strategy does. Several studies have reported the superiority of the adjointize-then-stabilize approach in terms of accuracy of the finite element approximation, for example in the context of optimal control problems [Collis et al. 2002] or *a posteriori* adjoint-based error estimation [Cyr et al. 2014]. For the flow of interest in this work, we illustrate adjoint inconsistency in fig. 2.1. Whereas the stabilize-then-adjointize approach (a) produces an adjoint critical eigenvector with element-sized oscillations close to the plate, the adjointize-then-stabilize approach (b) yields a perfectly smooth solution. Such observations are consistent, for example, with the work of Hicken et al. (2011, fig. 1) and motivate the use of the adjointize-then-stabilize approach for the present work. The main downside of this approach, however, is that a discretization of the adjoint problem is not simply obtained anymore by trans-conjugating the discretization of the direct problem. Instead, a dedicated routine must be implemented to build the discretization of the stabilized adjoint equations.

2.4.2 Solution method

The nonlinear problem eq. (2.30a) is solved with Newton iterations. The Jacobian matrix that is needed at each Newton iteration is computed by analytical differentiation of the stabilized weak formulation, except for the stabilization parameter τ which is kept constant. The linear system to be solved at each Newton iteration is handled with the sparse LU solver MUMPS [Amestoy et al. 2001], which is used from FreeFEM's interface to PETSc [Balay et al. 2019].

To achieve a quick convergence to the critical eigenmode, we directly solve for the critical triplet $(\hat{\mathbf{q}}_c, \omega_c, U_c^*)$, defined by:

$$\begin{aligned} (\mathbf{i}\omega_c \mathbf{M} + \mathbf{J}(\mathbf{q}_0; U_c^*)) \hat{\mathbf{q}}_c &= 0 \\ \hat{\mathbf{q}}_c^H \mathbf{M} \hat{\mathbf{q}}_c - 1 &= 0 \\ \Im(\hat{h}_c) &= 0 \end{aligned} \tag{2.31}$$

⁴More precisely, stabilize-then-adjointize consists in taking as a discrete approximation to \mathcal{J}^\dagger the Hermitian transpose of the SUPG-dcretized linear operator, \mathbf{J}^H

where eq. (2.30b) is re-stated in the first line, the second line is an (arbitrary) normalization of the eigenvector and the third line is a (arbitrary) phase condition that sets the imaginary part of the heaving degree of freedom to zero. Note the dependence of the Jacobian in U_c^* . The above problem is nonlinear and is solved again with Newton iterations. Due to the limited radius of convergence of the Newton iterations, a good enough initial guess must be provided. To that end, an extraction of the rightmost eigenvalues of $(\sigma \mathbf{M} + \mathbf{J}(\mathbf{q}_0; U^*)) \hat{\mathbf{q}} = 0$ is first performed with SLEPc's Krylov–Schur algorithm [Hernandez et al. 2005] (with shift-invert spectral transformation) for different U^* spanning a large enough range that contains U_c^* . Then, the triplet $(\hat{\mathbf{q}}, \omega, U^*)$ corresponding to the growth rate nearest to zero is taken as an initial guess. If one has previously obtained a critical triplet for physical parameters close enough to the current ones, it is also possible to take that former solution as an initial guess. In practice, this is the most frequent situation when we explore the (\tilde{m}, Re) parameter space (see chapter 4 and chapter 5).

Once the direct critical mode is known, the adjoint critical mode is obtained by performing some Krylov-Schur iterations on eq. (2.30c), using a shift close to $-i\omega_c$ for quick convergence.

Finally, the second order problems eq. (2.30d)–(2.30e) consist in simple linear solves that we perform again with MUMPS.

2.4.3 Mesh adaptation loop

The quantities involved in the weakly nonlinear analysis possess very different spatial supports. An illustration of that is provided in fig. 2.2. The steady solution (a), the direct eigenmode (b, top) and the second order forcings (c) are seen to have strong structures in the boundary layers and in the wake. Differently, the adjoint mode (b, bottom) has structures close to the plate and upstream. In addition to that, the spatial support of the function will significantly change when varying the Reynolds number across the range of interest in the subsequent chapters, $10 \leq Re \leq 10^4$ (cf chapter 4 and chapter 5). To cope with those difficulties and make the physical exploration as automatical as possible, we use mesh adaptation. A feature-based strategy [Alauzet et al. 2016] is adopted where the metrics is provided by the intersection⁵ of the individual metrics adapted to the Hessian of the following weakly nonlinear quantities:

$$U_0, V_0, |\hat{U}_c|, |\hat{V}_c|, |U_c^\dagger|, |V_c^\dagger|, |\hat{U}_2^{A^2}|, |\hat{V}_2^{A^2}|, |\hat{U}_2^{A^2}|, |\hat{V}_2^{A^2}|, |f_{\beta_f}|$$

The last quantity, f_{β_f} , corresponds to the volume distribution of β . Indeed, by developing the scalar product in the numerator of the definition of β , eq. (2.19), we have

$$\beta = \frac{\mathcal{F}_{\text{res}}^{u_h}(\hat{\mathbf{q}}_c, \mathbf{q}_2^{A^2}, \mathbf{q}_2^{A^2}) u_{hc}^{\dagger*}}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \hat{\mathbf{q}}_c \rangle} + \underbrace{\int_{\Omega_f} \frac{\mathcal{F}_{\text{res}}^{\mathbf{U}}(\hat{\mathbf{q}}_c, \mathbf{q}_2^{A^2}, \mathbf{q}_2^{A^2}) \cdot \mathbf{U}_c^{\dagger*}}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \hat{\mathbf{q}}_c \rangle}}_{= f_{\beta_f}(\mathbf{X})} + \int_{\Gamma_{fs}} \frac{\mathcal{F}_{\text{res}}^{\Lambda}(\hat{\mathbf{q}}_c, \mathbf{q}_2^{A^2}, \mathbf{q}_2^{A^2}) \cdot \Lambda_c^{\dagger*}}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \hat{\mathbf{q}}_c \rangle}$$

where $\mathcal{F}_{\text{res}}^{u_h}$, $\mathcal{F}_{\text{res}}^{\mathbf{U}}$ and $\mathcal{F}_{\text{res}}^{\Lambda}$ are the components of \mathcal{F}_{res} along u_h , \mathbf{U} and Λ , respectively. The second contribution is seen to be a volume integral of a quantity defined in the fluid domain. We identify its integrand, noted f_{β_f} , as the spatial distribution (or “density”) of β . As seen in fig. 2.2(d), this quantity is highly localized close to

⁵A definition of the intersection of metrics can be found, for example, in the PhD manuscript of Frazza (2018) (§5.2.3)

the plate leading edge and is believed to be an additional pertinent indicator about where the mesh should be refined to capture β at best. Note that f_{β_t} will be used again in chapter 5, for physical discussion on the origin of the nonlinearities.

In FreeFEM, the previously described mesh adaptation can be simply performed by using the interface to BAMG [Hecht 2006]:

```
thFadapted=adaptmesh(thF, [uFb,vFb], // U0,V0
                        [uFDmod,vFDmod], // |Ūc|,|V̂c|
                        [uFAmod,vFAmod], // |Uc†|,|Vc†|
                        [uF20mod,vF20mod], // |Ū2|A2||,|V̂2|A2||
                        [uF22mod,vF22mod], // |Ū2A2||,|V̂2A2||
                        betaspatialmod, // |fβt|
                        nbvx=1e6,ratio=1.5,nbsmooth=20,
                        hmin=1e-9,hmax=1,err=0.01);
```

Obviously, the metric used to obtain an adapted mesh is computed from the solution on a previous mesh. As a consequence, mesh adaptation is actually a non-linear process that we handle with the “fixed-point” loop described in algorithm 1.

Algorithm 1: Mesh adaptation for weakly nonlinear analysis

Data: (Re, \tilde{m}, \dots) , thF, \mathbf{q}_0 , $\hat{\mathbf{q}}_c$, U_c^* , ω_c , $\hat{\mathbf{q}}_2^{|A|^2}$, $\hat{\mathbf{q}}_2^{A^2}$, α , β , f_{β_t}

while $\{(\Delta U_c^* > tol_{U_c^*}) \text{ or } (\Delta \alpha > tol_\alpha) \text{ or } (\Delta \beta > tol_\beta)\}$ **do**

$U_{c,old}^* = U_c^*$, $\alpha_{old} = \alpha$, $\beta_{old} = \beta$;

thF=AdaptMesh(thF, \mathbf{q}_0 , $\hat{\mathbf{q}}_c$, \mathbf{q}_c^\dagger , $\hat{\mathbf{q}}_2^{|A|^2}$, $\hat{\mathbf{q}}_2^{A^2}$, f_{β_t});

\mathbf{q}_0 =SolveOrder0(thF, \mathbf{q}_0);

$[\hat{\mathbf{q}}_c, U_c^*, \omega_c, \mathbf{q}_c^\dagger]$ =SolveOrder1(thF, \mathbf{q}_0 , $\hat{\mathbf{q}}_c$, U_c^* , ω_c);

$[\hat{\mathbf{q}}_2^{|A|^2}, \hat{\mathbf{q}}_2^{A^2}]$ =SolveOrder2(thF, \mathbf{q}_0 , $\hat{\mathbf{q}}_c$);

$[\alpha, \beta, f_{\beta_t}]$ =BuildStuartLandau(thF, \mathbf{q}_0 , $\hat{\mathbf{q}}_c$, \mathbf{q}_c^\dagger , $\hat{\mathbf{q}}_2^{|A|^2}$, $\hat{\mathbf{q}}_2^{A^2}$);

$\Delta U_c^* = \frac{|U_c^* - U_{c,old}^*|}{|U_{c,old}^*|}$, $\Delta \alpha = \frac{|\alpha - \alpha_{old}|}{|\alpha_{old}|}$, $\Delta \beta = \frac{|\beta - \beta_{old}|}{|\beta_{old}|}$;

end

For a given set of physical parameters, an adequate initial guess for the mesh and different weakly nonlinear quantities are provided. The loop starts by adapting the mesh to the current solution, given the set of sensors. Then, the sequence of weakly nonlinear problems eq. (2.30) is solved. First, a new \mathbf{q}_0 is computed on the adapted mesh, given an initial guess. Then, the first order direct/adjoint problems are solved with the current \mathbf{q}_0 and given an initial guess for $\hat{\mathbf{q}}_c$ ⁶. The second order problems are solved given the newly computed \mathbf{q}_0 and $\hat{\mathbf{q}}_c$. The weakly nonlinear solve completes with the building of the Stuart–Landau equation, *i.e.* its coefficients α and β . The spatial distribution of β is also an output, as it is used for the next mesh adaptation iteration. The loop is stopped when the mesh and the solution are considered to be converged. Here, we used a criteria based on the relative variations of U_c^* , α and β . The tolerances were set to 0.01. In practice, the number of mesh adaptation iterations needed for convergence depends on the initial guess. Most of the time however, a

⁶An initial guess is needed for the direct mode since it is computed with a Newton method, cf section 2.4.2

weakly nonlinear run is performed starting from the solution at a previous nearby parameter set. In this case, two iterations were enough⁷.

For a prescribed interpolation error of 0.01 and a given set of physical parameters, the final adapted mesh is composed of $3.31 \cdot 10^4$ triangles and is illustrated in fig. 2.2(e). The mesh is seen to well adapt to the features from the different sensors. In the wake, streamwise elongated elements are obtained so as to efficiently capture the variation of \mathbf{U}_0 , $\widehat{\mathbf{U}}_c$, $\widehat{\mathbf{U}}_2^{|\mathbf{A}|^2}$ and $\widehat{\mathbf{U}}_2^{A^2}$. In the upstream part, the mesh is similarly refined but, this time, to capture the adjoint mode structure (bottom half of fig. 2.2(b)). Close to the plate, both the leading and trailing edge are particularly refined. The leading edge refinement is the consequence of the sharp variations of the steady-state, direct mode and second-order responses in that area. On the contrary, the trailing edge refinement is due to the adjoint mode. Now zooming on the (upper) leading edge in fig. 2.3, we observe several highly localized regions of extreme mesh refinement. If these appear largely over-refined with respect to the steady-state solution (fig. 2.3(a)), they are essential for capturing the extremely sharp variations of f_{β_i} (fig. 2.3(b)).

⁷In the case the initial guess is taken at a nearby different set of physical parameters, we *forced* at least 2 mesh adaptations in order for the mesh to be adapted, at least once, to a solution at the target physical parameters.

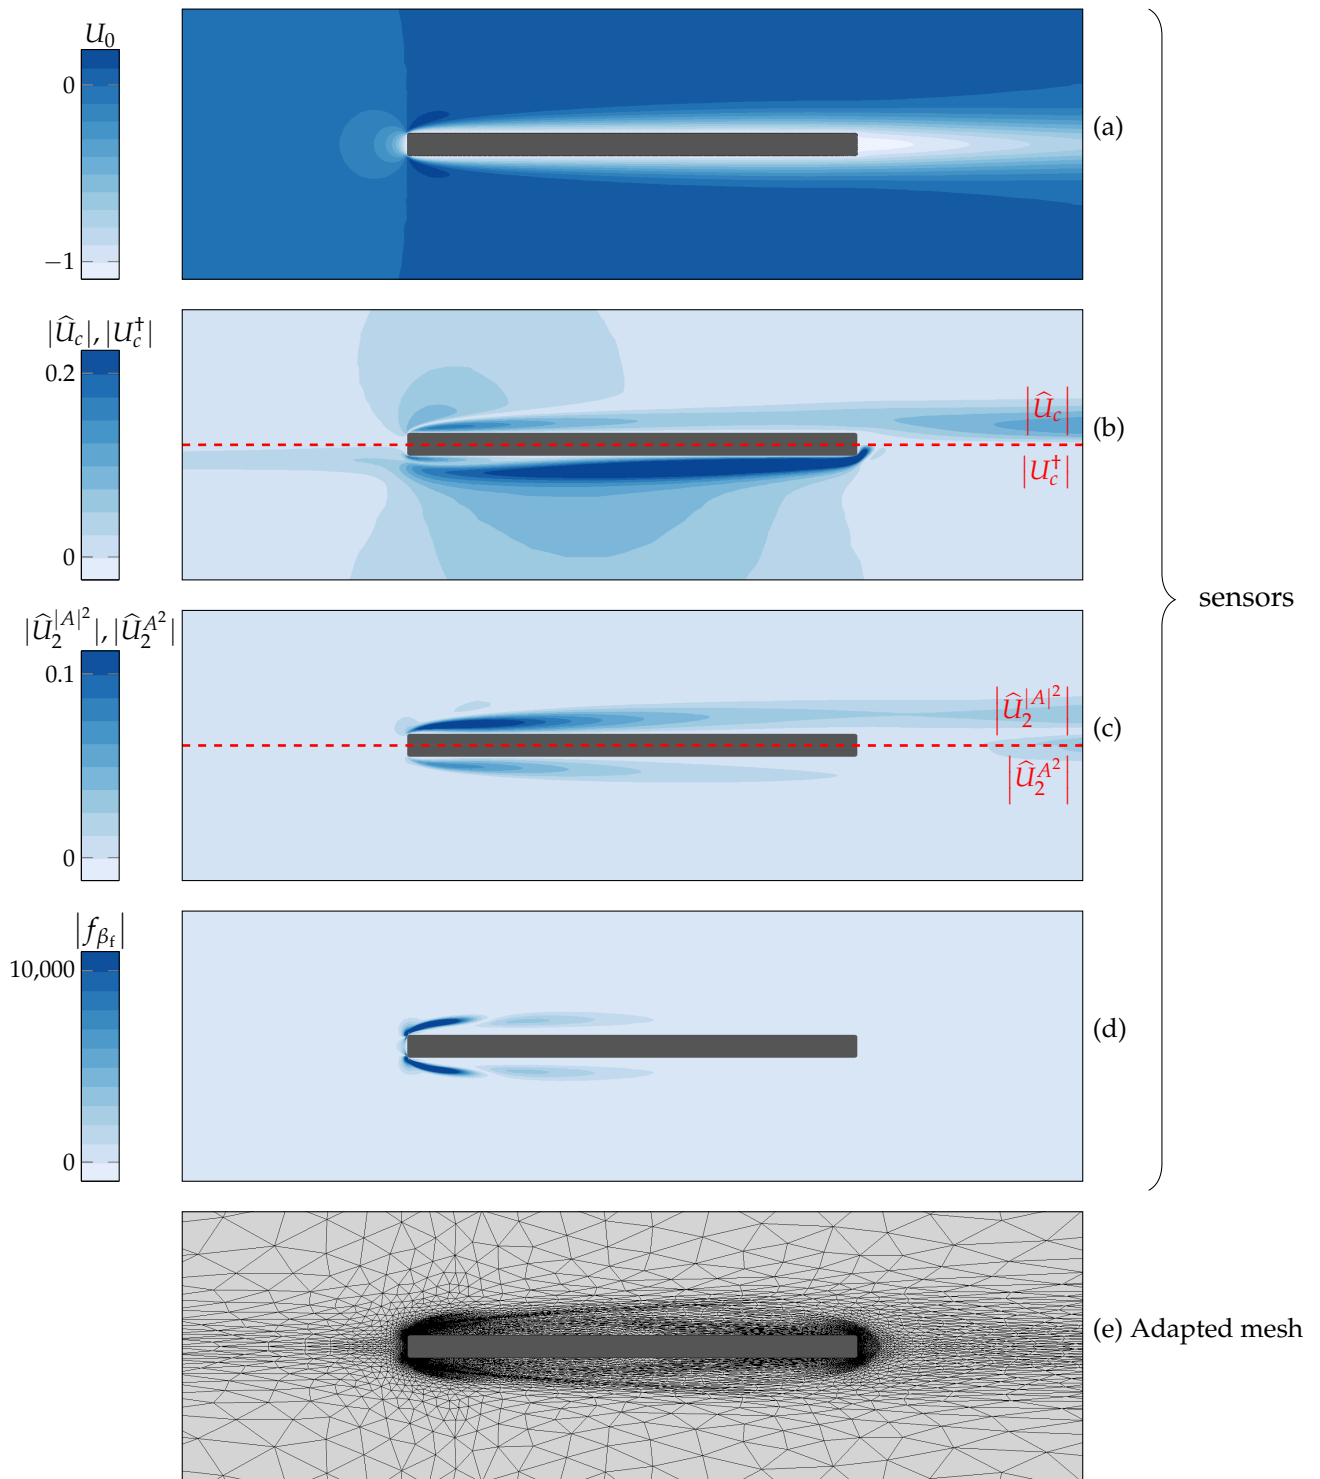


FIGURE 2.2: Hessian-based adaptation of the computational mesh to various flow sensors arising in the weakly nonlinear analysis: the steady-state solution U_0 (a), the direct and adjoint critical modes $|\hat{u}_c|, |u_c^+|$ (b), the second order responses $|\hat{u}_2^{A^2}|, |\hat{u}_2^{A^2}|$ (c) and the spatial distribution $|f_{\beta_f}|$ of the normal form coefficient β (d). The resulting adapted mesh is shown in (e). This example is given for $Re = 2682.81$ and $\bar{m} = 1000$.




FIGURE 2.3: Zoom on the leading edge of the adapted mesh of fig. 2.2(e). The mesh is superimposed in (a) on the steady solution and in (b) on the spatial distribution $|f_{\beta_f}|$ of the normal form coefficient, β .

2.4.4 Validation

Validating the weakly nonlinear solver itself is not an easy task as no established benchmark exists for weakly nonlinear flutter solutions in laminar incompressible flows. However, it is possible to compare the weakly nonlinear results to time-accurate simulations obtained with a classical time-stepping approach (see section 1.3), that we consider as reference. In fig. 2.4, the weakly nonlinear predictions (gray lines) are compared with permanent regime time-accurate simulations (black disks) on a range of U^* . On the upper figure (a), we show the (nonlinear) flutter frequency. In the weakly nonlinear framework, the latter is reconstructed as $\omega_c + \omega_{nl}(U^*)$ whereas it is simply extracted from the DNS time signal, using Fast Fourier Transform. On the lower figure (b), the amplitude of the pitching degree of freedom is presented, where the weakly nonlinear result is obtained through the first order reconstruction eq. (2.24).

Notice first that in both figures, the linear stability threshold, predicted from the eigenproblem eq. (2.9), is in agreement with the DNS results that start oscillating precisely at $U_c^* = 4.96$. This gives confidence in the stability analysis part of the weakly nonlinear analysis. Focusing now the part $U^* > U_c^*$, we observe in fig. 2.4(a) a perfect agreement of the flutter frequencies on the whole range of reduced velocities. The situation is slightly different for the pitching amplitude which are in good agreement only for U^* close enough to U_c^* . We point out the fact that this increasing gap between WNL and DNS as we move away from the stability threshold is expected. It is classically the result of higher order nonlinear effects coming into play and that are not captured with the present third-order weakly nonlinear approach. Similar trends can be found for example in [Gallaire et al. 2016] for the study of the vortex shedding bifurcation around a circular cylinder, in [Vio et al. 2007] for the galloping instability, in [Tchoufag et al. 2015] for the study of buoyancy-driven bodies or in [Woodgate et al. 2007; Gai et al. 2016] for the flutter instability.

Overall, the weakly nonlinear solver — and as part of it, the linear stability solver — is shown to be consistent with reference DNS solutions. In passing, we

note that the DNS code is implemented without any SUPG stabilization, and in addition, with an Arbitrary Lagrangian Eulerian framework that is different from (but consistent with) the “absolute velocity-rotating axis” formulation of the weakly nonlinear solver. As a consequence, the good agreement between these two different approaches is considered a satisfying sign of the correctness of the weakly nonlinear solver and its implementation. Further evidence of that will be given in chapter 5 for different sets of physical parameters and with additional comparison to Harmonic Balance calculations.




FIGURE 2.4: Validation of the weakly nonlinear analysis. The weakly nonlinear results (gray lines) for frequency (a) and pitching amplitude (b) are compared with reference time-accurate (DNS) solutions (disks) ($Re = 500$, $\bar{m} = 1000$).

2.5 Conclusion

In this chapter, we presented a weakly nonlinear approach based on the method of multiple scales that allows to derive a normal form for the flutter bifurcation, using a typical aeroelastic section immersed in an incompressible Navier–Stokes flow. The numerical solution of this semi-analytical approach has been addressed using a SUPG-stabilized finite element spatial discretization. A particular focus was put on the discretization of the adjoint problem for which it was shown that stabilizing the continuous adjoint equations leads to significantly smoother solutions than “adjointizing” the discrete stabilized direct equations. Then, a Hessian-based mesh adaptation framework was proposed where the adaptation metric is built as the intersection of various quantities involved in the weakly nonlinear analysis. In particular, a scalar field, noted $f_{\beta_r}(\mathbf{X})$, that reflects the spatial distribution of the nonlinear normal form coefficient β was taken into account. Its small wavelength variations close to the plate leading edge could then be accurately captured, hopefully leading to more

efficient computation of β . Finally, the weakly nonlinear predictions were verified in comparison to results from a classical fully nonlinear time-marching solver, yielding consistent results.

APPENDIX

Appendix 2.A Development in ϵ series of trigonometric nonlinearities

This appendix provides the decomposition in ϵ power series of quantities used in the “absolute velocity - rotating axis” formulation and that involve trigonometric nonlinearities: the rotation matrix, $\mathbf{R}(\theta)$ and the quantities that depend on it, $\mathbf{W}(u_h, \theta, u_\theta)$ and $\mathcal{C}_L(\boldsymbol{\Lambda}, \theta)$.

First, we seek for a ϵ -power series for $\cos \theta$ and $\sin \theta$. To ease calculations, we consider, as an intermediate step, the case $\theta_0 = 0$. A nonzero θ_0 component is introduced back a few lines below. Let us approximate \cos and \sin functions with a Taylor development⁸ around 0:

$$\begin{aligned}\cos(\theta) &= 1 - \frac{\theta^2}{2} + o(\theta^3) \\ \sin(\theta) &= \theta - \frac{\theta^3}{6} + o(\theta^3)\end{aligned}$$

Injecting in the above expressions the ϵ -series decomposition of θ (with $\theta_0 = 0$ here)

$$\theta = \epsilon\theta_1 + \epsilon^2\theta_2 + \epsilon^3\theta_3 + \dots$$

we obtain ϵ -series for $\cos \theta$ and $\sin \theta$:

$$\begin{aligned}\cos(\theta) &= 1 - \epsilon^2\frac{1}{2}\theta_1^2 - \epsilon^3\theta_1\theta_2 + \dots \\ \sin(\theta) &= \epsilon\theta_1 + \epsilon^2\theta_2 + \epsilon^3\left(\theta_3 - \frac{1}{6}\theta_1^3\right) + \dots\end{aligned}$$

From the above, it is easy to define the ϵ -series for $\mathbf{R}(\theta)$, including a nonzero θ_0 through:

$$\begin{aligned}\mathbf{R}(\theta) &= \mathbf{R}\left(\theta_0 + (\epsilon\theta_1 + \epsilon^2\theta_2 + \epsilon^3\theta_3 + \dots)\right) \\ &= \mathbf{R}(\theta_0)\mathbf{R}(\epsilon\theta_1 + \epsilon^2\theta_2 + \epsilon^3\theta_3 + \dots) \\ &= \mathbf{R}(\theta_0) \left[\begin{pmatrix} 1 & 0 \\ 0 & 1 \end{pmatrix} + \epsilon \begin{pmatrix} 0 & -\theta \\ \theta & 0 \end{pmatrix} + \epsilon^2 \begin{pmatrix} -\frac{1}{2}\theta_1^2 & -\theta_2 \\ \theta_2 & -\frac{1}{2}\theta_1^2 \end{pmatrix} + \epsilon^3 \begin{pmatrix} -\theta_1\theta_2 & -(\theta_3 - \frac{1}{6}\theta_1^3) \\ \theta_3 - \frac{1}{6}\theta_1^3 & -\theta_1\theta_2 \end{pmatrix} + \dots \right] \\ \mathbf{R}(\theta) &= \mathbf{R}(\theta_0) + \epsilon \frac{\partial \mathbf{R}}{\partial \theta} \theta_1 + \epsilon^2 \left[\frac{\partial \mathbf{R}}{\partial \theta} \theta_2 + \mathbf{R}_{2,11}(\theta_1, \theta_1) \right] + \epsilon^3 \left[\frac{\partial \mathbf{R}}{\partial \theta} \theta_3 + \mathbf{R}_{3,21}(\theta_2, \theta_1) + \mathbf{R}_{3,111}(\theta_1, \theta_1, \theta_1) \right] + \dots\end{aligned}$$

⁸In practice, order 3 is sufficient in the Taylor expansion since the weakly nonlinear analysis is pushed up to order 3 only

with

$$\begin{aligned} \frac{\partial \mathbf{R}}{\partial \theta} \theta &= \mathbf{R}(\theta_0) \begin{pmatrix} 0 & -\theta \\ \theta & 0 \end{pmatrix} & \mathbf{R}_{2,11}(\theta_a, \theta_b) &= -\frac{1}{2} \mathbf{R}(\theta_0) \begin{pmatrix} \theta_a \theta_b & 0 \\ 0 & \theta_a \theta_b \end{pmatrix} \\ \mathbf{R}_{3,21}(\theta_a, \theta_b) &= \mathbf{R}(\theta_0) \begin{pmatrix} -\theta_a \theta_b & 0 \\ 0 & -\theta_a \theta_b \end{pmatrix} & \mathbf{R}_{3,111}(\theta_a, \theta_b, \theta_c) &= \mathbf{R}(\theta_0) \frac{1}{6} \begin{pmatrix} 0 & \theta_a \theta_b \theta_c \\ -\theta_a \theta_b \theta_c & 0 \end{pmatrix} \end{aligned}$$

Replacing $\mathbf{R}(\theta)$ by its ϵ -series, we can derive the series for the solid/grid velocity $\mathbf{W}(u_h, \theta, u_\theta) = \mathbf{R}(\theta)^\top (-U_\infty, u_h)^\top + u_\theta \mathbf{e}_z \wedge \mathbf{X}$:

$$\begin{aligned} \mathbf{W}(u_h, \theta, u_\theta) &= \mathbf{W}(0, \theta_0, 0) + \epsilon \frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q}_1 + \epsilon^2 \left[\frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q}_2 + \mathbf{W}_{2,11}(\mathbf{q}_1, \mathbf{q}_1) \right] \\ &+ \epsilon^3 \left[\frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q}_3 + \mathbf{W}_{3,21}(\mathbf{q}_2, \mathbf{q}_1) + \mathbf{W}_{3,111}(\mathbf{q}_1, \mathbf{q}_1, \mathbf{q}_1) \right] + \dots \end{aligned}$$

where the following functions are defined:

$$\begin{aligned} \frac{\partial \mathbf{W}}{\partial \mathbf{q}} \mathbf{q} &= \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta (-U_\infty, 0)^\top + \mathbf{R}(\theta_0)^\top (0, u_h)^\top + u_\theta \mathbf{e}_z \wedge \mathbf{X}, \\ \mathbf{W}_{2,11}(\mathbf{q}_a, \mathbf{q}_b) &= \mathbf{R}_{2,11}(\theta_a, \theta_b)^\top (-U_\infty, 0)^\top + \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta_a (0, u_{hb})^\top, \\ \mathbf{W}_{3,21}(\mathbf{q}_a, \mathbf{q}_b) &= \mathbf{R}_{3,21}(\theta_a, \theta_b)^\top (-U_\infty, 0)^\top + \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta_a (0, u_{hb})^\top + \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta_b (0, u_{ha})^\top, \\ \mathbf{W}_{3,111}(\mathbf{q}_a, \mathbf{q}_b, \mathbf{q}_c) &= \mathbf{R}_{3,111}(\theta_a, \theta_b, \theta_c)^\top (-U_\infty, 0)^\top + \mathbf{R}_{2,11}(\theta_a, \theta_b)^\top (0, u_{hc})^\top \end{aligned}$$

Similarly for the lift coefficient $\mathcal{C}_L(\boldsymbol{\Lambda}, \theta) = 2 \int_{\Gamma_{fs}} \{ \boldsymbol{\Lambda} \cdot \mathbf{R}^\top(\theta) \mathbf{e}_y \}$, we can write:

$$\begin{aligned} \mathcal{C}_L(\boldsymbol{\Lambda}, \theta) &= \mathcal{C}_L(\boldsymbol{\Lambda}_0, \theta_0) + \epsilon \frac{\partial \mathcal{C}_L}{\partial \mathbf{q}} \mathbf{q}_1 + \epsilon^2 \left[\frac{\partial \mathcal{C}_L}{\partial \mathbf{q}} \mathbf{q}_2 + \mathcal{C}_{L2,11}(\mathbf{q}_1, \mathbf{q}_1) \right] \\ &+ \epsilon^3 \left[\frac{\partial \mathcal{C}_L}{\partial \mathbf{q}} \mathbf{q}_3 + \mathcal{C}_{L3,21}(\mathbf{q}_2, \mathbf{q}_1) + \mathcal{C}_{L3,111}(\mathbf{q}_1, \mathbf{q}_1, \mathbf{q}_1) \right] + \dots \end{aligned}$$

where the following functions are defined:

$$\begin{aligned} \frac{\partial \mathcal{C}_L}{\partial \mathbf{q}} \mathbf{q} &= 2 \int_{\Gamma_{fs}} \left\{ \boldsymbol{\Lambda} \cdot \mathbf{R}^\top(\theta_0) \mathbf{e}_y + \boldsymbol{\Lambda}_0 \cdot \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta \mathbf{e}_y \right\} \\ \mathcal{C}_{L2,11}(\mathbf{q}_a, \mathbf{q}_b) &= 2 \int_{\Gamma_{fs}} \left\{ \boldsymbol{\Lambda}_0 \cdot \mathbf{R}_{2,11}(\theta_a, \theta_b)^\top \mathbf{e}_y + \boldsymbol{\Lambda}_a \cdot \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta_b \mathbf{e}_y \right\} \\ \mathcal{C}_{L3,21}(\mathbf{q}_a, \mathbf{q}_b) &= 2 \int_{\Gamma_{fs}} \left\{ \boldsymbol{\Lambda}_0 \cdot \mathbf{R}_{3,21}(\theta_a, \theta_b)^\top \mathbf{e}_y + \boldsymbol{\Lambda}_a \cdot \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta_b \mathbf{e}_y + \boldsymbol{\Lambda}_b \cdot \frac{\partial \mathbf{R}^\top}{\partial \theta} \theta_a \mathbf{e}_y \right\} \\ \mathcal{C}_{L3,111}(\mathbf{q}_a, \mathbf{q}_b, \mathbf{q}_c) &= 2 \int_{\Gamma_{fs}} \left\{ \boldsymbol{\Lambda}_0 \cdot \mathbf{R}_{3,111}(\theta_a, \theta_b, \theta_c)^\top \mathbf{e}_y + \boldsymbol{\Lambda}_a \cdot \mathbf{R}_{2,11}(\theta_b, \theta_c)^\top \mathbf{e}_y \right\} \end{aligned}$$

Appendix 2.B Operators of the SUPG-discretized eigenproblem eq. (2.30b)

In this appendix, we provide the detailed expressions of the matrices involved in the SUPG-discretized eigenproblem eq. (2.30b). The weak formulation defining the discrete operators is given by eq. (2.29b), whereas the finite element test and trial spaces have been defined in section 1.2.2. Here we simply recall that the fluid velocity, pressure and interface Lagrange multiplier are searched in the finite element spaces \mathcal{V}_Δ^U , \mathcal{V}_Δ^p and $\mathcal{V}_\Delta^\Lambda$ respectively. The corresponding basis are noted:

$$\begin{aligned}\mathcal{V}_\Delta^U &= \text{span} \left(\boldsymbol{\psi}_0^U, \boldsymbol{\psi}_1^U, \dots, \boldsymbol{\psi}_{n_{\text{dof},U}-1}^U \right) \\ \mathcal{V}_\Delta^p &= \text{span} \left(\boldsymbol{\psi}_0^p, \boldsymbol{\psi}_1^p, \dots, \boldsymbol{\psi}_{n_{\text{dof},p}-1}^p \right) \\ \mathcal{V}_\Delta^\Lambda &= \text{span} \left(\boldsymbol{\psi}_0^\Lambda, \boldsymbol{\psi}_1^\Lambda, \dots, \boldsymbol{\psi}_{n_{\text{dof},\Lambda}-1}^\Lambda \right)\end{aligned}$$

Let us now detail the matrices $\mathbf{J}(\mathbf{q}_0, U^*)$ and $\mathbf{M}(\mathbf{q}_0)$ that appear in eq. (2.30b). For convenience, we split the solid and fluid unknowns such that eq. (2.30b) writes:

$$\mathbf{i}\omega_c \underbrace{\begin{pmatrix} \mathbf{M}_s & 0 \\ 0 & \mathbf{M}_f(\mathbf{q}_0) \end{pmatrix}}_{=\mathbf{M}(\mathbf{q}_0)} \begin{pmatrix} \widehat{\mathbf{q}}_{\text{sc}} \\ \widehat{\mathbf{q}}_{\text{fc}} \end{pmatrix} + \underbrace{\begin{pmatrix} \mathbf{J}_{\text{ss}}(\mathbf{q}_0, U_c^*) & \mathbf{J}_{\text{sf}}(\mathbf{q}_0) \\ \mathbf{J}_{\text{fs}}(\mathbf{q}_0) & \mathbf{J}_{\text{ff}}(\mathbf{q}_0) \end{pmatrix}}_{=\mathbf{J}(\mathbf{q}_0, U_c^*)} \begin{pmatrix} \widehat{\mathbf{q}}_{\text{sc}} \\ \widehat{\mathbf{q}}_{\text{fc}} \end{pmatrix} = 0$$

where the solid unknowns are gathered in $\widehat{\mathbf{q}}_{\text{sc}} = (\widehat{h}_c, \widehat{\theta}_c, \widehat{u}_{h_c}, \widehat{u}_{\theta_c})^\text{T}$ and the fluid finite element degrees of freedom in $\widehat{\mathbf{q}}_{\text{fc}} = (\widehat{\mathbf{U}}_c, \widehat{p}_c, \widehat{\Lambda}_c)^\text{T}$.

The linearized solid problem is defined by

$$\begin{aligned}\mathbf{M}_s &= \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \\ \mathbf{J}_{\text{ss}}(\mathbf{q}_0, U^*) &= \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \\ \Omega^2 U^{*-2} & -\frac{1}{\tilde{m}} \frac{\partial \mathcal{C}_L}{\partial \theta} \Big|_{\mathbf{q}_0} & 2\zeta_h \Omega U^{*-1} & 0 \\ 0 & U^{*-2} & 0 & 2\zeta_p U^{*-1} \end{pmatrix} \\ \mathbf{J}_{\text{sf}}(\mathbf{q}_0) &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{\tilde{m}} \frac{\partial \mathcal{C}_L}{\partial \Lambda} \Big|_{\mathbf{q}_0} \\ 0 & 0 & -\frac{1}{\tilde{m}r\theta^2} \frac{\partial \mathcal{C}_M}{\partial \Lambda} \end{pmatrix}\end{aligned}$$

Note that $\partial \mathcal{C}_L / \partial \Lambda|_{\mathbf{q}_0}$ and $\partial \mathcal{C}_M / \partial \Lambda$ are row matrices respectively defined by $[\partial \mathcal{C}_L / \partial \Lambda|_{\mathbf{q}_0}]_j = 2 \int_{\Gamma_{\text{fs}}} \{ \boldsymbol{\psi}_j^\Lambda \cdot \mathbf{R}^\text{T}(\theta_0) \mathbf{e}_y \}$ and $[\partial \mathcal{C}_M / \partial \Lambda]_j = 2 \int_{\Gamma_{\text{fs}}} \{ \mathbf{X} \wedge \boldsymbol{\psi}_j^\Lambda \} \cdot \mathbf{e}_z$.

The linearized fluid problem is defined by

$$\mathbf{M}_f(\mathbf{q}_0) = \begin{pmatrix} \mathbf{M}_{\mathbf{U},\tau}(\mathbf{q}_0) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}, \quad \mathbf{J}_{\text{ff}}(\mathbf{q}_0) = \begin{pmatrix} \mathbf{A}_\tau(\mathbf{q}_0) & \mathbf{B}_\tau^\text{T}(\mathbf{q}_0) & \mathbf{I}_\Lambda^\text{T} \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ \mathbf{I}_\Lambda & \mathbf{0} & \mathbf{0} \end{pmatrix}$$

$$\mathbf{J}_{\text{fs}}(\mathbf{q}_0) = \begin{pmatrix} \mathbf{0} & \mathbf{D}_\theta(\mathbf{q}_0) & \mathbf{D}_{u_h}(\mathbf{q}_0) & \mathbf{D}_{u_\theta}(\mathbf{q}_0) \\ \mathbf{0} & \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{W}_\theta(\mathbf{q}_0) & -\mathbf{W}_{u_h}(\mathbf{q}_0) & -\mathbf{W}_{u_\theta}(\mathbf{q}_0) \end{pmatrix}$$

where $\mathbf{M}_{\mathbf{U},\tau}(\mathbf{q}_0)$ is the mass matrix on the velocity space that incorporates the SUPG stabilization term. It was already defined in eq. (1.32) and is recalled below for convenience:

$$[\mathbf{M}_{\mathbf{U},\tau}(\mathbf{q}_0)]_{ij} = \int_{\Omega_f} \left\{ \boldsymbol{\psi}_j^\text{U} \cdot (\boldsymbol{\psi}_i^\text{U} + \tau \nabla \boldsymbol{\psi}_i^\text{U} (\mathbf{U}_0 - \mathbf{W}_0)) \right\}$$

Matrices \mathbf{A}_τ and \mathbf{B}_τ^T correspond to the linearized convection-diffusion and pressure gradient operators (with stabilization):

$$[\mathbf{A}_\tau(\mathbf{q}_0)]_{ij} = \int_{\Omega_f} \left\{ \nabla \boldsymbol{\psi}_j^\text{U} (\mathbf{U}_0 - \mathbf{W}_0) + \nabla \mathbf{U}_0 \boldsymbol{\psi}_j^\text{U} \right\} \cdot \boldsymbol{\psi}_i^\text{U} + \int_{\Omega_f} \nu \nabla_s \boldsymbol{\psi}_j^\text{U} : \nabla \boldsymbol{\psi}_i^\text{U} + \dots$$

$$+ \int_{\Omega_f} \tau \left\{ \nabla \boldsymbol{\psi}_j^\text{U} (\mathbf{U}_0 - \mathbf{W}_0) + \nabla \mathbf{U}_0 \boldsymbol{\psi}_j^\text{U} - \mathcal{R}e^{-1} \nabla^2 \boldsymbol{\psi}_j^\text{U} \right\} \cdot \nabla \boldsymbol{\psi}_i^\text{U} (\mathbf{U}_0 - \mathbf{W}_0)$$

$$[\mathbf{B}_\tau^\text{T}(\mathbf{q}_0)]_{ij} = \int_{\Omega_f} \left\{ -\boldsymbol{\psi}_j^p \nabla \cdot \boldsymbol{\psi}_i^\text{U} + \tau \nabla \boldsymbol{\psi}_j^p \cdot \nabla \boldsymbol{\psi}_i^\text{U} (\mathbf{U}_0 - \mathbf{W}_0) \right\}$$

whereas \mathbf{B} is the classical (Galerkin) divergence matrix:

$$[\mathbf{B}]_{ij} := - \int_{\Omega_f} \nabla \cdot \boldsymbol{\psi}_j^\text{U} \boldsymbol{\psi}_i^p$$

Matrix \mathbf{I}_Λ is a mass matrix defined on the interface Lagrange multiplier space. The column matrices \mathbf{W}_θ , \mathbf{W}_{u_h} and \mathbf{W}_{u_θ} transmit the solid velocity to the fluid at the interface, whereas the *shape derivatives* \mathbf{D}_θ , \mathbf{D}_{u_h} and \mathbf{D}_{u_θ} represent the effect of a solid movement on the Navier–Stokes momentum equation:

$$[\mathbf{W}_\theta(\mathbf{q}_0)]_i = \int_{\Gamma_{\text{fs}}} \frac{\partial \mathbf{W}}{\partial \theta} \cdot \boldsymbol{\psi}_i^\Lambda$$

$$[\mathbf{W}_{u_h}(\mathbf{q}_0)]_i = \int_{\Gamma_{\text{fs}}} \frac{\partial \mathbf{W}}{\partial u_h} \cdot \boldsymbol{\psi}_i^\Lambda$$

$$[\mathbf{W}_{u_\theta}(\mathbf{q}_0)]_i = \int_{\Gamma_{\text{fs}}} \frac{\partial \mathbf{W}}{\partial u_\theta} \cdot \boldsymbol{\psi}_i^\Lambda$$

$$[\mathbf{D}_\theta(\mathbf{q}_0)]_i = \int_{\Omega_f} \left\{ -\nabla \mathbf{U}_0 \frac{\partial \mathbf{W}}{\partial \theta} \right\} \cdot (\boldsymbol{\psi}_i^\text{U} + \tau \nabla \boldsymbol{\psi}_i^\text{U} (\mathbf{U}_0 - \mathbf{W}_0))$$

$$[\mathbf{D}_{u_h}(\mathbf{q}_0)]_i = \int_{\Omega_f} \left\{ -\nabla \mathbf{U}_0 \frac{\partial \mathbf{W}}{\partial u_h} \right\} \cdot (\boldsymbol{\psi}_i^\text{U} + \tau \nabla \boldsymbol{\psi}_i^\text{U} (\mathbf{U}_0 - \mathbf{W}_0))$$

$$[\mathbf{D}_{u_\theta}(\mathbf{q}_0)]_i = \int_{\Omega_f} \left\{ -\nabla \mathbf{U}_0 \frac{\partial \mathbf{W}}{\partial u_\theta} + \mathbf{e}_z \wedge \mathbf{U}_0 \right\} \cdot (\boldsymbol{\psi}_i^\text{U} + \tau \nabla \boldsymbol{\psi}_i^\text{U} (\mathbf{U}_0 - \mathbf{W}_0))$$

3

A TIME SPECTRAL SOLVER FOR PERIODIC FLUTTER SOLUTIONS USING BLOCK-CIRCULANT PRECONDITIONING

This chapter is dedicated to the presentation of a harmonic-balance-type approach, known as the Time Spectral Method (TSM). As a (long) preamble, TSM is considered in the more global context of harmonic balance methods. A classification borrowed from spectral methods in space is used to organize the discussion. The core of the chapter presents a Newton—Krylov solution method and introduces the so-called block-circulant preconditioner that is at the core of the efficiency of the proposed method, when large number of harmonics are used. A strategy for a time-parallel implementation is presented. Several numerical experiments are reported in order to assess the robustness of the preconditioner, in particular for large numbers of harmonics. The method is applied first on a case where the solution frequency is known (forced oscillations) and then extended to the case of unknown frequency (self-sustained oscillations), using a Schur complement approach.

3.1 Introduction

The preliminary results obtained with weakly nonlinear analysis in chapter 2 show that when nonlinearity becomes strong enough — typically when U^* deviates significantly from the critical threshold U_c^* — the weakly nonlinear analysis predicts less and less accurately the fully nonlinear dynamics. To overcome that shortcoming — and yet, avoid resorting to classical, but costly, time-stepping algorithms (cf section 1.3) — we propose to use a Harmonic Balance approach.

Harmonic Balance Methods (HBM) are a class of theoretical/numerical methods devoted to the research of periodic solutions to nonlinear systems of ordinary (ODE) or partial (PDE) differential equations. The mathematical foundation of the method lies in the ability to search for periodic solutions under the form of Fourier series. The first reference to the name “Harmonic Balance” is often traced back to the work of Krylov et al. (1949) where it was used to compute analytical solutions to generic (weakly) nonlinear equations¹. In its original form, the method solves for the unknown Fourier coefficients of the solution. This may lead to possibly cumbersome analytical aspects as the system of nonlinear equations verified by the Fourier coefficients must be derived, before it is eventually solved by any means available.

The method has been actively used in various fields of physics, such as electrical engineering [Gilmore et al. 1991] and nonlinear vibrations mechanics [Krack et al. 2019], where the study of periodic solutions to nonlinear equations is ubiquitous. In the aeroelasticity community, Harmonic Balance analyses have been performed to study low-dimensional systems of ODE’s, such as pitching an plunging rigid airfoils with linear analytical fluid models and structural nonlinearities [Shen 1959; Laursen et al. 1980; Price et al. 1995; D. Tang et al. 1998; L. Liu et al. 2004; L. Liu et al. 2005]. With the progress of computers, moderate size systems such as arising from spatial discretization of reasonable two-dimensional fluid flows were considered in theoretical fluid mechanics works [Carte et al. 1995; Fabre et al. 2018; Rigas et al. 2020].

Consistently with the growing use of CFD-based aeroelasticity for industrial purposes, HBM has received renewed interest starting with the work of Hall et al. (2002) who relied on the use of discrete Fourier transforms in order to formulate a time-domain HBM, later christened *High-Dimensional Harmonic Balance Method* [L. Liu et al. 2006]. In this approach, the analytical manipulations required to derive the equations verified by the Fourier coefficients are avoided so that the method is well-suited to the highly nonlinear and high-dimensional systems arising from the spatial discretization of the Euler or Navier–Stokes equations on industrial configurations. In addition, this formulation has the advantage of being more easily implementable into pre-existing CFD codes. From there, different names have been introduced for similar methods. With the *Nonlinear Frequency Domain Method* McMullen et al. (2006) searched for the solution in the frequency domain using a *pseudo-spectral approach* to compute the nonlinear terms. With the *Time Spectral Method*, [Gopinath et al. 2005] used the same formulation as [Hall et al. 2002], but proposed an alternative way of computing the spectral time-derivative operator which avoids the use of any Fourier transform. The different variant of Hall’s Harmonic Balance Method, have been used to compute solutions to various periodically forced two- and three-dimensional problems: pitching/plunging airfoils [Hall et al. 2002; Gopinath et al.

¹The method is also known as the *Equivalent Linearization Method* or also *Describing Function Method* as it allows to derive, thanks to the hypothesis harmonic oscillation, an equivalent linear representation for the nonlinear terms of the equation (cf. [Dimitriadis 2017, §4.3])

2005; Leffell et al. 2014], helicopter rotors [Leffell et al. 2014], wind turbines [Leffell et al. 2016], turbomachines [Van der Weide et al. 2005], etc. Also, in order to handle self-sustained limit cycle oscillations, several strategies have been proposed to deal with the unknown frequency [Thomas et al. 2002; McMullen et al. 2006; Ekici et al. 2011; Yao et al. 2015; Yao et al. 2018]. Recently, HBM has been nested in outer algorithms used to determine directly flutter thresholds [Thomas et al. 2018; Prasad et al. 2018; H. Li et al. 2018; S. He et al. 2018], or to perform unsteady shape optimisation [Thomas et al. 2005; Mader et al. 2012; Choi et al. 2014; Blondeau et al. 2019].

In Harmonic Balance Methods, all Fourier coefficient — or equivalently, all time instants for time-domain HBM — are searched together, in a coupled fashion. As a consequence, the number of degrees of freedom grows not only with spatial discretization, but also with time discretization. Thus, high-dimensional problems arise, even with reasonable spatial discretizations, when one takes into account growing numbers of harmonics. The question of the efficient and robust numerical solution of HBM formulations is thus a critical, on-going research focus, in the perspective of large-scale applications. Originally, the HBM equations were explicitly marched in pseudo-time until a steady-state was reached. However, the CFL number used for pseudo-time marching depends on the maximal resolved frequency [Van der Weide et al. 2005] so that the pseudo timestep must be decreased as the number of harmonics and/or the physical frequency increase, thus yielding very slow convergence to steady-state. To overcome this issue, implicit pseudo-timestepping has been proposed by Sicot et al. (2008), who solve the linear system at each pseudo-time step through block-Jacobi fixed-point iterations. As the number of harmonics increases, the diagonal dominance of the implicit matrix decreases so that the pseudo timestep is limited in this approach also, but this time to guarantee proper convergence of stationary block-Jacobi iterations. To circumvent that issue, several authors considered the use of linear solvers that do not require diagonal dominance for convergence, such as the Generalized Minimal RESidual method (GMRES). The performance of the approach is then entirely driven by one's ability to find a good preconditioning strategy. In [Woodgate et al. 2009] a block incomplete lower upper factorization was proposed, in [Su et al. 2010] a multigrid strategy with Gauss-Seidel smoother, in [Mundis et al. 2014] (Block-Colored) Gauss-Seidel was used as a preconditioner, either with or without the temporal coupling terms and defect correction iterations. In [Mundis et al. 2015], the authors proposed to rewrite the unknowns so as to group together all instants at one spatial location². Then, block Gauss-Seidel iterations were used as a preconditioner, with redefined blocks of size [number of steady-state unknowns per cell \times number of instants]. With this strategy, wave-number independent GMRES iteration counts were obtained, despite a preconditioner of complexity $\mathcal{O}(N^3)$, where N is the number of harmonics. In [Mundis et al. 2017; Leffell et al. 2016], approximate factorization-based preconditioners have been proposed, showing similar wave-number independence but with a complexity brought back to $\mathcal{O}(N^2)$.

The objectives of the present chapter are two-fold. First, we propose an unified overview of the different Harmonic Balance Methods that appeared in the aeroelasticity literature through the years. By adopting a generic classification historically used for spectral methods (in space) [Canuto et al. 1988], we rigorously compare the different alternatives. This first step ultimately motivates our choice to opt for the Time Spectral Method, as formulated by Gopinath et al. (2005). The second objective of the chapter is to tackle the challenging task of the numerical solution of the

²in the way a new physical equation would be added to the model ...

TSM problem. To this end, we adopt a Newton–Krylov approach, for which a new preconditioning strategy, based on circulant matrix properties, is introduced for the efficient solution of the linearized TSM equations. A parallel-in-time implementation is presented and its performance is assessed through two numerical examples: (i) a NACA0012 profile forced in plunging motion and (ii) the self-sustained flutter LCO developing on the spring-mounted plate studied in chapter 4. For the second case an original way of dealing with the additional frequency unknown, through Schur complements, is proposed.

3.2 Harmonic Balance methods

In this section, we present the Harmonic Balance approach and compare the different variants that have emerged in the literature through the years. For simplicity, we consider some generic system of K equations, of which a T -periodic solution vector $q(t) \in \mathbb{R}^K$ is sought:

$$M \frac{\partial q}{\partial t} + R(q) = f(t) \quad (3.1)$$

with $R(q)$ the stationary residual and $f(t)$ a T -periodic forcing term. The forcing is here meant in a very general sense and may consist in practice in a volume source term or in imposed boundary conditions (cf section 3.4.1 for the latter). In eq. (3.1), the frequency $\omega = 2\pi/T$ is known, as it is imposed by f . For sake of clarity, we remain in that case in all section 3.2. The case with unknown ω and zero forcing term is handled in section 3.4.2.

Fourier decomposition of the solution Due to the T -periodicity, the unknown exact solution q and the imposed forcing term f can be written as their Fourier series

$$q(t) = \sum_{k=-\infty}^{\infty} \hat{q}_k e^{ik\omega t} \quad f(t) = \sum_{k=-\infty}^{\infty} \hat{f}_k e^{ik\omega t} \quad (3.2)$$

with Fourier coefficients:

$$\hat{q}_k = \frac{1}{T} \int_0^T q(t) e^{-ik\omega t} \quad \hat{f}_k = \frac{1}{T} \int_0^T f(t) e^{-ik\omega t} \quad (3.3)$$

As any numerical method, Harmonic Balance relies on approaching the exact solution q by a finite dimensional approximation noted q_N . Because of its spectral convergence property [Canuto et al. 1988, §2.1.1] the truncated Fourier series offers a good *ansatz* for q_N . The Harmonic Balance approach thus relies on searching solutions of the form:

$$q_N(t) = \sum_{k=-N}^N \hat{q}_k e^{ik\omega t} \quad (3.4)$$

where the coefficients \hat{q}_k , $k = -N \dots N$ form the finite set of unknowns, with N the truncation order.

Finding the \hat{q}_k can be done *via* two main approaches. In the first approach, known as the *Galerkin formulation*, the equations governing the \hat{q}_k are derived by requiring that the residual of eq. (3.1) be orthogonal to any of the $2N + 1$ Fourier basis functions. In the second approach, known as the *collocation formulation*, the governing equations eq. (3.1) are enforced in a point-wise manner on a predefined grid of $2N + 1$ time instants. This classification, though not widely used when referring to

(time) Harmonic Balance methods, is commonly used in the literature using spectral methods for space discretization (Canuto et al. 1988, Chapter 3). It is adopted here because it provides a rigorous and general framework to present and compare the wide variety of Harmonic Balance methods that have emerged through the years.

3.2.1 Galerkin approach: frequency-domain HBM

Introducing the Fourier decompositions 3.2 in eq. (3.1), testing against the Fourier basis functions and invoking the orthogonality property of the Fourier basis³, one can derive a system of equations for the unknowns \hat{q}_k :

$$\mathbf{i}k\omega M \hat{q}_k + \hat{R}_k = \hat{f}_k \quad \forall k \in [-N \dots N] \quad (3.5)$$

where \hat{f}_k are the (known) Fourier coefficients of the forcing term eq. (3.3) and \hat{R}_k are the unknown Fourier coefficient of $R(q_N)$:

$$\hat{R}_k = \frac{1}{T} \int_0^T R(q_N(t)) e^{-\mathbf{i}k\omega t} \quad (3.6)$$

For future convenience, this system is rewritten in matrix form:

$$\omega \hat{\mathbf{D}} \hat{\mathbf{q}} + \hat{\mathbf{R}} = \hat{\mathbf{f}} \quad (3.7)$$

where

$$\hat{\mathbf{D}} = \mathbf{i} \text{diag}(-NM, \dots, NM) \quad (3.8)$$

is a block-diagonal matrix with diagonal blocks equal to $\mathbf{i}kM$, $k \in [-N, N]$ and

$$\hat{\mathbf{q}} = (\hat{q}_{-N}, \dots, \hat{q}_N)^T \quad \hat{\mathbf{R}} = (\hat{R}_{-N}, \dots, \hat{R}_N)^T \quad \hat{\mathbf{f}} = (\hat{f}_{-N}, \dots, \hat{f}_N)^T$$

The system of equations eq. (3.7) is not well-posed, because it possesses only $2N + 1$ equations for the $2 \times (2N + 1)$ unknowns that are the solutions harmonics $(\hat{q}_{-N}, \dots, \hat{q}_N)$ and the residual harmonics $(\hat{R}_{-N}, \dots, \hat{R}_N)$. All the methods introduced hereinafter allow specifying the dependence of the residual harmonics with the solution harmonics, that can be formally denoted $\hat{\mathbf{R}}(\hat{\mathbf{q}})$. Once this dependence is specified, the new system of equations writes

$$\omega \hat{\mathbf{D}} \hat{\mathbf{q}} + \hat{\mathbf{R}}(\hat{\mathbf{q}}) = \hat{\mathbf{f}} \quad (3.9)$$

We now introduce the different methods that allow to specify the dependence of residual harmonics with the solutions harmonics. In the Analytical Harmonic Balance Method developed in section 3.2.1.1, this dependence is obtained analytically in the spectral domain. In pseudo-spectral method introduced in section 3.2.1.2 the Fourier transform is used so that this dependence can be more easily evaluated in the temporal domain.

³The orthogonality of the Fourier basis, noted here $\{\phi_k\}_{k \in [-N, N]}$, $\phi_k(t) = e^{\mathbf{i}kt}$, with respect to the $\mathcal{L}^2(0, T)$ scalar product, reads as follows, with the star denoting the complex conjugate:

$$\frac{1}{T} \int_0^T \phi_k(t) \phi_l(t)^* dt = \delta_{kl} = \begin{cases} 1 & \text{if } k = l \\ 0 & \text{else} \end{cases}$$

3.2.1.1 Analytical Harmonic Balance Method (AHBM)

A first choice is to express $\widehat{\mathbf{R}}(\widehat{\mathbf{q}})$ analytically by injecting q_N into $R(q)$ and then invoking the orthogonality property of the Fourier basis to isolate the residual Fourier coefficients \widehat{R}_m .

For a simple example, suppose the residual is bilinear (typically, the Navier–Stokes convection term) $R(q) = \mathcal{N}(q, q)$. We have:

$$\begin{aligned} R(q_N(t)) &= \mathcal{N} \left(\left(\sum_{k=-N}^N \widehat{q}_k e^{ik\omega t} \right), \left(\sum_{l=-N}^N \widehat{q}_l e^{il\omega t} \right) \right) \\ &= \sum_m \sum_{\substack{-N \leq k, l \leq N \\ k+l=m}} \mathcal{N}(\widehat{q}_k, \widehat{q}_l) e^{i(k+l)\omega t} \\ &= \sum_{|m| \leq N} \sum_{\substack{-N \leq k, l \leq N \\ k+l=m}} \mathcal{N}(\widehat{q}_k, \widehat{q}_l) e^{i(k+l)\omega t} + \sum_{|m| > N} \sum_{\substack{-N \leq k, l \leq N \\ k+l=m}} \mathcal{N}(\widehat{q}_k, \widehat{q}_l) e^{i(k+l)\omega t} \end{aligned} \quad (3.10)$$

where we have separated in the last line the contributions that vibrate at frequencies lower (first sum) and higher (second sum) than the truncation frequency $N\omega$. Plugging the above expression in the definition of \widehat{R}_m eq. (3.6) and using the Fourier basis orthogonality gives the analytical expression of the residual harmonics as a function of the solutions harmonics \widehat{q}_k :

$$\widehat{R}_m = \sum_{\substack{-N \leq k, l \leq N \\ k+l=m}} \mathcal{N}(\widehat{q}_k, \widehat{q}_l) \quad (3.11)$$

In particular, we note that the high-frequency terms in eq. (3.10) (second sum) vanish naturally in the process, thanks to orthogonality. The full nonlinear residual

$$\widehat{\mathbf{R}}(\widehat{\mathbf{q}}) = \left(\widehat{R}_{-N}, \dots, \widehat{R}_N \right)^T$$

is then explicitly specified as a function of $\widehat{\mathbf{q}}$, using eq. (3.11).

In the rest of this manuscript, this method is referred to as the *Analytical Harmonic Balance Method* (AHBM). It is attractive because it allows to formulate the problem completely in the frequency domain and brings an analytical understanding of the generation of contributions at the different orders due to the nonlinear terms. Such an approach was used first by [Carte et al. \(1995\)](#) and more recently by [Fabre et al. \(2018\)](#) for computing periodic solutions of the laminar incompressible flow developing in the wake of a circular cylinder. In addition, [Rigas et al. \(2020\)](#) used it to study the nonlinear transition scenarios in the boundary layer over a flat plate. Also, the interested reader will find pedagogical examples using this approach in the context of aero-elasticity in the book of [Dimitriadis \(2017\)](#).

However, AHBM is tractable only for systems with “simple” enough nonlinearities, *i.e.* typically polynomial and of low order. Indeed, as the order of the nonlinearity increases, more and more complex convolutions are needed leading to cumbersome analytical expressions. In addition, from the numerical point of view, the cost of evaluating $\widehat{\mathbf{R}}(\widehat{\mathbf{q}})$ scales with N^p where p is the order of the polynomial nonlinearity. Contrary to the other approaches introduced below, this cost depend on p and quickly becomes large for highly nonlinear problem. Finally, nonpolynomial nonlinearities, as often encountered in turbulence models (e.g. the RANS Spalart-Allmaras

model) or in fluid-structure interaction (e.g. [Richter 2015]), cannot be handled directly. Of course, some tricks can be used to deal with this issue, like the definition of additional unknowns [Hall et al. 2002] or — as was considered temporarily in this thesis — the use of power series to reformulate the problem with only polynomial nonlinearities. In any case, large computational and analytical costs are still expected.

3.2.1.2 Pseudo-spectral Harmonic Balance Method

In order to overcome the drawbacks of the Analytical Harmonic Balance Method exposed above, $\widehat{\mathbf{R}}(\widehat{\mathbf{q}})$ can be evaluated using a *pseudo-spectral approach*. This relies on evaluating the residual in the time domain at a number of discrete time instants, say the $2N + 1$ equidistant points $t_n = Tn/(2N + 1)$ with $n = 0, \dots, 2N$, and then recast it back to the frequency domain, using Discrete Fourier Transform. Let us first introduce Discrete Fourier Transform, and then use it to define the pseudo-spectral HBM.

3.2.1.2.a Discrete Fourier Transform

The *discrete Fourier series* of q is defined as:

$$I_N[q](t) = \sum_{k=-N}^N \tilde{q}_k e^{ik\omega t} \quad (3.12)$$

where the \tilde{q}_k are the *discrete Fourier coefficients*, defined as:

$$\tilde{q}_k = \frac{1}{2N + 1} \sum_{n=0}^{2N} q(t_n) e^{-ik\omega t_n} \quad (3.13)$$

Let us recall that, in general, the discrete Fourier coefficients of a signal q are different from the continuous ones defined by eq. (3.3) (e.g. [Canuto et al. 1988, eq. (2.1.29)]). They become identical if q contains only frequencies lower than the truncation frequency $N\omega$ ⁴.

The Discrete Fourier Transform (DFT) can be viewed as a linear mapping between the $2N + 1$ instantaneous values $\mathbf{q} = (q(t_0), \dots, q(t_{2N}))^T$ and the $2N + 1$ discrete Fourier coefficients $\tilde{\mathbf{q}} = (\tilde{q}_{-N}, \dots, \tilde{q}_N)^T$ through eq. (3.13). For convenience, this mapping is written in matrix form:

$$\tilde{\mathbf{q}} = \mathbf{F}\mathbf{q} \quad (3.14)$$

with $(\mathbf{F})_{kn} = \frac{1}{2N + 1} e^{-ik\omega t_n} \quad k \in [-N, N], n \in [0, 2N]$

In addition, the inverse Discrete Fourier Transform \mathbf{F}^{-1} is defined as:

$$\mathbf{q} = \mathbf{F}^{-1}\tilde{\mathbf{q}} \quad (3.15)$$

with $(\mathbf{F}^{-1})_{nk} = e^{ik\omega t_n} \quad n \in [0, 2N], k \in [-N, N]$

⁴Or, in other words, if the grid of the $\{t_n\}_{n \in [0, 2N+1]}$ respects the Shannon criteria for the signal q .

For more clarity, \mathbf{F} and \mathbf{F}^{-1} are also written in developed form:

$$\mathbf{F} = \frac{1}{2N+1} \begin{pmatrix} 1 & e^{iN\omega t_1} & \dots & e^{iN\omega t_{2N}} \\ \vdots & \vdots & \ddots & \vdots \\ 1 & 1 & \dots & 1 \\ \vdots & \vdots & \ddots & \vdots \\ 1 & e^{-iN\omega t_1} & \dots & e^{-iN\omega t_{2N}} \end{pmatrix} \quad (3.16)$$

and

$$\mathbf{F}^{-1} = \begin{pmatrix} 1 & \dots & 1 & \dots & 1 \\ e^{-iN\omega t_1} & & 1 & & e^{iN\omega t_1} \\ \vdots & & \vdots & & \vdots \\ e^{-iN\omega t_{2N}} & \dots & 1 & \dots & e^{iN\omega t_{2N}} \end{pmatrix} \quad (3.17)$$

Remark. With the convention adopted in this manuscript, the \mathbf{F} is *almost* a unitary matrix, since $\mathbf{F}^{-1} = (2N+1)\mathbf{F}^H$, where the H superscript denotes the Hermitian conjugate.

3.2.1.2.b Pseudo-spectral HBM

Having introduced the Discrete Fourier Transform in the previous paragraph, the pseudo-spectral approach consists in approximating the *continuous* Fourier coefficients of the residual $\widehat{\mathbf{R}} = (\widehat{R}_{-N}, \dots, \widehat{R}_N)^T$ by its *discrete* Fourier coefficients $\widetilde{\mathbf{R}} = (\widetilde{R}_{-N}, \dots, \widetilde{R}_N)^T$:

$$\widehat{\mathbf{R}}(\widehat{\mathbf{q}}) \simeq \widetilde{\mathbf{R}}(\widehat{\mathbf{q}}) = \mathbf{FR}(\mathbf{F}^{-1}\widehat{\mathbf{q}}) \quad (3.18)$$

The rightmost term of eq. (3.18) defines a three-steps process to evaluate $\widetilde{\mathbf{R}}(\widehat{\mathbf{q}})$ that details as follows:

1. Evaluate the approximate solution q_N at the instants $t_n = Tn/(2N+1)$ with $n = 0, \dots, 2N$, using inverse Discrete Fourier Transform:

$$\mathbf{q} = \mathbf{F}^{-1}\widehat{\mathbf{q}} \quad \mathbf{q} = (q_N(t_0), \dots, q_N(t_{2N}))^T \quad (3.19)$$

2. Compute the residual at each instant from the instantaneous values gathered in \mathbf{q} :

$$\mathbf{R}(\mathbf{q}) = (R(q_N(t_0)), \dots, R(q_N(t_{2N})))^T \quad (3.20)$$

3. Transfer the residual back in the frequency domain using a Discrete Fourier Transform:

$$\widetilde{\mathbf{R}} = \mathbf{FR} \quad (3.21)$$

Let us now make some comments. First, note that in step 1, we directly link the *continuous* (instead of the discrete) Fourier coefficients $\widehat{\mathbf{q}}$ to \mathbf{q} through DFT. This is valid because, by definition, the approximate solution has a bandwidth equal to N . Thus, the Shannon criteria is respected, by construction, on the $2N+1$ instants sampling grid, making the discrete and continuous Fourier coefficients equal. On the contrary, in step 3., the notation $\widetilde{\mathbf{R}}$ is used instead of $\widehat{\mathbf{R}}$, to insist on the fact that the *discrete* Fourier coefficients of $R(q_N)$ are obtained, not the continuous ones. Indeed, in general, nothing guarantees that the Shannon criteria is enforced for the residual $R(q_N)$. Aliasing errors are then introduced and $\widetilde{\mathbf{R}}$ is only an approximation

of $\widehat{\mathbf{R}}$, as specified in eq. (3.18). The approach presented here is thus an approximate Galerkin method, hence the name *pseudo-spectral*. The interested reader is reported to section 3.2.3.2 for further discussion about aliasing error.

To our knowledge, in the Harmonic Balance literature, the pseudo-spectral approach has been first used by McMullen et al. (2006) who labeled it *Nonlinear Frequency Domain method*. Compared to AHBM presented above, this method does not require any analytical development and is readily applicable to any $R(q)$. In addition, the cost of evaluating $\widehat{\mathbf{R}}(\widehat{\mathbf{q}})$ is dominated by the DFT and thus scales with N^2 (the evaluation of the $2N + 1$ instantaneous residuals costs only N and is neglected here). In practice, further acceleration is possible [McMullen et al. 2006] by using Fast Fourier Transform instead of DFT, bringing back the cost to $N \log N$.

3.2.2 Collocation approach : time-domain HBM

Differently from the Galerkin approach which is obtained by orthogonal projection of the residual on the Fourier basis (eq. (3.5)), the collocation approach consists in enforcing eq. (3.1) in a pointwise manner, at the nodes of the temporal grid $t_n = Tn/(2N + 1)$, $n = 0, \dots, 2N$:

$$M \left. \frac{\partial q_N}{\partial t} \right|_{t=t_n} + R(q_N(t_n)) = f(t_n) \quad \forall n \in [0, 2N + 1] \quad (3.22)$$

By time-differentiating the Fourier series form of q_N eq. (3.4), and evaluating its Fourier coefficients *via* DFT (eq. (3.13)), a spectral approximation to the time derivative in the above equation can be obtained (see [Trefethen 1996; Gopinath et al. 2005] for a detailed derivation):

$$\left. \frac{\partial q_N}{\partial t} \right|_{t=t_n} \simeq \omega \sum_{k=0}^{2N} d_k q_N(t_{n+k}) \quad \text{with} \quad d_k = \begin{cases} \frac{1}{2}(-1)^{k+1} \csc\left(\frac{\pi k}{2N+1}\right) & \text{if } k \neq 0 \\ 0 & \text{if } k = 0 \end{cases} \quad (3.23)$$

Note that here, the index $n + k$ must be understood “modulo $2N + 1$ ” due to periodicity.

For convenience, eq. (3.22) is re-written in matrix form:

$$\omega \mathbf{D} \mathbf{q} + \mathbf{R}(\mathbf{q}) = \mathbf{f} \quad (3.24)$$

with

$$\begin{aligned} \mathbf{q} &= (q_N(t_0), \dots, q_N(t_{2N}))^T \\ \mathbf{R} &= (R(q_N(t_0)), \dots, R(q_N(t_{2N})))^T \\ \mathbf{f} &= (f(t_0), \dots, f(t_{2N}))^T \end{aligned}$$

and where $\omega \mathbf{D}$ is the spectral time-derivative operator. Due to eq. (3.23), it possesses a particular block-circulant structure:

$$\mathbf{D} = \begin{pmatrix} 0 & d_1 M & \dots & d_{2N} M \\ d_{2N} M & 0 & \ddots & \\ & \ddots & \ddots & d_1 M \\ d_1 M & & d_{2N} M & 0 \end{pmatrix} \quad (3.25)$$

The obtained formulation, eq. (3.24) and (3.25), is often referred to as the *Time Spectral Method* (TSM) [Gopinath et al. 2005]. However, a variant can be obtained by

using an alternative and equivalent definition of \mathbf{D} , based on Fourier transforms:

$$\mathbf{D} = \mathbf{F}^{-1} \widehat{\mathbf{D}} \mathbf{F} \quad (3.26)$$

With that definition, the time derivative is applied by (i) computing the Fourier coefficient of q_N with DFT, (ii) applying the time derivative in the spectral domain⁵ thanks to matrix $\widehat{\mathbf{D}}$ (see the definition, eq. (3.8)) and (iii) casting the result back in the time domain with inverse DFT. Authors that use eq. (3.26) instead of eq. (3.25) often refer to the collocation formulation as the *High-Dimensional Harmonic Balance method* [Hall et al. 2002; L. Liu et al. 2006; LaBryer et al. 2009].

A particularity of the operator \mathbf{D} is that it is global, in the sense that the value of the derivative at instant t_n depends on the value of the function at all other instants. Thus, the cost of applying \mathbf{D} in the form given by eq. (3.25) scales with N^2 . With the alternative definition eq. (3.26), the cost seems at first sight even higher, since a direct and an inverse DFT are needed, with a cost of N^2 each. However, it is well-known that costly DFT can always be replaced by much more economical Fast Fourier Transforms, that only cost $N \log N$. Here appears the main interest of evaluating the time-derivative *via* eq. (3.26): it decreases the application cost from N^2 to $N \log N$.

3.2.3 Comparison of the Galerkin and collocation approaches

The three Harmonic Balance formulations introduced previously are now compared. First, we show that the pseudo-spectral Galerkin and the collocation approaches are equivalent. Second, we discuss the only difference between the Analytical Harmonic Balance Method and the other formulations, which is the existence of aliasing errors in the latter.

3.2.3.1 Equivalence between the pseudo-spectral Galerkin and collocation approaches

By combining eq. (3.5) and (3.18), the pseudo-spectral Galerkin reads:

$$\omega \widehat{\mathbf{D}} \widehat{\mathbf{q}} + \mathbf{FR} \left(\mathbf{F}^{-1} \widehat{\mathbf{q}} \right) = \widehat{\mathbf{f}} \quad (3.27)$$

Since the Shannon criteria is respected for q_N on the temporal grid $\{t_n\}_{n \in [0, 2N]}$, the continuous Fourier coefficients $\widehat{\mathbf{q}}$ are equal to the discrete ones, $\widetilde{\mathbf{q}}$, and thus can be evaluated with DFT from the instantaneous values, *i.e.* $\widehat{\mathbf{q}} = \mathbf{F} \mathbf{q}$. Using also the fact that $\widehat{\mathbf{f}} = \mathbf{F} \mathbf{f}$ ⁶ yields

$$\omega \widehat{\mathbf{D}} \mathbf{F} \mathbf{q} + \mathbf{FR} (\mathbf{q}) = \mathbf{F} \mathbf{f} \quad (3.28)$$

Then, left-multiplying the above equation by \mathbf{F}^{-1} , the Time Spectral Method eq. (3.24) is retrieved:

$$\omega \underbrace{\mathbf{F}^{-1} \widehat{\mathbf{D}} \mathbf{F}}_{=\mathbf{D}} \mathbf{q} + \mathbf{R} (\mathbf{q}) = \mathbf{f} \quad (3.29)$$

where the spectral differentiation operator \mathbf{D} is identified using eq. (3.26). Thus, the pseudo-spectral Harmonic Balance Method and Time Spectral Method are simply the frequency domain and time domain versions of the same method.

⁵It simply corresponds to multiplying each harmonic \widetilde{q}_k by ik

⁶Again, the direct link between the instantaneous values of f and its continuous Fourier coefficients can only be computed with DFT if the Shannon criteria is respected for f . We suppose it is ...

3.2.3.2 Aliasing error

The aliasing phenomena arises when a signal is sampled on a grid that fails to ensure the Shannon criteria. In the context of Harmonic Balance methods, aliasing introduces a particular error, called *aliasing error*, that adds to the more classical truncation error due to neglecting harmonics higher than N .

As we shall see next, the aliasing error can be easily identified in the context of the pseudo-spectral Galerkin approach presented in section 3.2.1.2. We mention here that, due to the equivalence, demonstrated in section 3.2.3.1, between the Galerkin pseudo-spectral and collocation approaches, the latter is, as the former, subjected to aliasing errors.

With the Galerkin pseudo-spectral approach, we recall that the *discrete* Fourier coefficients of the residual are used as an approximation of the continuous ones, as stated in eq. (3.18). However, from classical textbooks, e.g. [Canuto et al. 1988, eq. (2.1.29)], it is known that the discrete and continuous Fourier coefficients are linked by the following relation:

$$\tilde{R}_k = \hat{R}_k + \underbrace{\sum_{\substack{m=-\infty \\ m \neq 0}}^{+\infty} \hat{R}_{k+(2N+1)m}}_{\text{aliasing error}} \quad k \in [-N, N] \quad (3.30)$$

where the aliasing error naturally appears as the difference between the \tilde{R}_k and the \hat{R}_k . In other words, the k -th harmonics of the residual computed with a pseudo-spectral approach is a mixture of the “true” continuous Fourier coefficient \hat{R}_k and an infinite number of contributions at frequencies $k + (2N + 1)m$, $m \in \mathbb{Z} - \{0\}$. This is linked to the fact that, once sampled on the $2N + 1$ -points grid $\{t_n\}_{n \in [0, 2N]}$, the Fourier basis functions of the form $k + (2N + 1)m$ are indistinguishable, $e^{ik\omega t_n} = e^{i(k+(2N+1)m)\omega t_n}$; they are *aliases*. In the end, due to aliasing errors, the pseudo-spectral approach does not solve the original Galerkin Harmonic Balance equations eq. (3.5) but a perturbed version of it that includes the aliased contributions:

$$ik\omega M \hat{q}_k + \hat{R}_k + \sum_{\substack{m=-\infty \\ m \neq 0}}^{+\infty} \hat{R}_{k+(2N+1)m} = \hat{f}_k \quad \forall k \in [-N \dots N] \quad (3.31)$$

In eq. (3.30), it is clear that, if the bandwidth of the residual signal $R(q_N(t))$ is higher than N , then there exist nonzero \hat{R}_k for $k > N$, and the aliasing error term is nonzero. In practice this is necessarily the case as soon as the residual is nonlinear since nonlinear interactions between harmonics of q_N lower than N generate harmonics of $R(q_N)$ higher than N . An example of that in the case of a quadratic nonlinearity is given below.

Example. Let us go back to the example introduced in section 3.2.1.1 where $R(q) = \mathcal{N}(q, q)$ is bilinear. In that case, the continuous Fourier coefficients of $R(q_N)$, obtained in eq. (3.11), are:

$$\hat{R}_m = \sum_{\substack{-N \leq k, l \leq N \\ k+l=m}} \mathcal{N}(\hat{q}_k, \hat{q}_l)$$

where it is seen that the highest possible non-zero harmonics is $m = 2N > N$. Then, from equation eq. (3.30) we have, for example for harmonics $k = 0, 1, N$:

$$\begin{aligned}\tilde{R}_0 &= \hat{R}_0 + \underbrace{0}_{\text{aliasing error}} \\ \tilde{R}_1 &= \hat{R}_1 + \underbrace{\hat{R}_{-2N}}_{\text{aliasing error}} \\ \tilde{R}_N &= \hat{R}_N + \underbrace{\hat{R}_{-(N+1)}}_{\text{aliasing error}}\end{aligned}$$

Finally, note that in comparison to either the Galerkin pseudo-spectral or collocation approaches, the Analytical Harmonic Balance Method (section 3.2.1.1) is aliasing-free because the aliased contributions are removed by construction, see eq. (3.11). In other words, due to aliasing errors, the Analytical Harmonic Balance Method actually solves a different system of equations than the two other approaches.

3.2.3.3 Synthetic view

We conclude the section by providing a synthetic view in table 3.2.1. On the one hand, the Galerkin methods (first line) are formulated in the frequency domain where the unknowns are the Fourier coefficients. The nonlinear terms $\hat{\mathbf{R}}$ are either computed analytically yielding a truly spectral aliasing-free method (first column), or computed in time-domain and brought back to frequency domain with DFT yielding an aliasing-prone pseudo-spectral method (second column). On the other hand, the collocation methods (second line) are formulated in the time domain where the unknowns are a set of instantaneous values spanning the period. By construction, the collocation methods are pseudo-spectral in the sense that they suffer from aliasing errors. Aliasing-free — or at least less aliasing-prone — time-domain methods have not been reviewed in this work but some strategies have been considered by different researchers, such as the use of Fourier filtering in [LaBryer et al. 2009] or the addition of spectral viscosity in [Huang et al. 2014].

	Spectral (aliasing-free)	Pseudo-spectral (aliasing-prone)
Galerkin (frequency domain)	<ul style="list-style-type: none"> Analytical Harmonic Balance Method (AHBM) 	<ul style="list-style-type: none"> Nonlinear Frequency Domain Method (NLFD)
Collocation (time domain)	<ul style="list-style-type: none"> HDHBM + Fourier filtering HDHBM + spectral viscosity 	<ul style="list-style-type: none"> High-Dimensional Harmonic Balance Method (HDHBM) Time Spectral Method (TSM)

TABLE 3.2.1: Summary of the different Harmonic Balance-based methods

3.3 Numerical solution of the Time Spectral Method

In the previous sections, we have presented the Harmonic Balance framework. Different formulations have been introduced and compared. In the present section, we focus on the Time Spectral Method (TSM) that will be used in practice in the rest of this manuscript (see chapter 5 and 6), because of its generality and ease of implementation for highly nonlinear equations. We now expose the algorithmic details for the numerical solution of the TSM equation eq. (3.24), that is recalled below for convenience:

$$\omega \mathbf{D} \mathbf{q} + \mathbf{R}(\mathbf{q}) = \mathbf{f} \quad (3.32)$$

Because TSM yields a stationary system of equations, Newton iterations (section 3.3.1) are a natural choice to reach fast convergence to the solution. However, the Newton iterations mostly rely on the solution of the linearized TSM system, which attains very large dimensions for typical fluid-structure problems. The delicate issue of its solution is tackled in section 3.3.2 and 3.3.3 where we introduce a new preconditioning strategy based on the almost circulant pattern of the TSM Jacobian matrix. The section ends with section 3.3.4 where a (time-)parallelization strategy is described. As in section 3.2, we mainly discuss the forced case $f(t \neq 0)$ while the extension to autonomous systems is addressed in section 3.4.2.

3.3.1 Newton method

For convenience, let us start by introducing the residual $\mathbf{R}^{TSM}(\mathbf{q})$ of the TSM system, eq. (3.32), in the case of an imposed frequency:

$$\mathbf{R}^{TSM}(\mathbf{q}) = 0 \quad \text{with} \quad \mathbf{R}^{TSM}(\mathbf{q}) = \omega \mathbf{D} \mathbf{q} + \mathbf{R}(\mathbf{q}) - \mathbf{f} \quad (3.33)$$

The above nonlinear system is solved with a Newton-like method, the i -th iteration of which reads:

$$\begin{aligned} \left(\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM}(\mathbf{q}^i) \right) \delta \mathbf{q} &= -\mathbf{R}^{TSM}(\mathbf{q}^i) \\ \mathbf{q}^{i+1} &= \mathbf{q}^i + \delta \mathbf{q} \end{aligned} \quad (3.34)$$

where $\mathbf{M} = \text{diag}(M, \dots, M)$ is a block-diagonal matrix and \mathbf{J}^{TSM} the TSM Jacobian defined by

$$\mathbf{J}^{TSM}(\mathbf{q}) = \left. \frac{\partial \mathbf{R}^{TSM}}{\partial \mathbf{q}} \right|_{\mathbf{q}} = \omega \mathbf{D} + \mathbf{J} \quad \text{with} \quad \mathbf{J} = \left. \frac{\partial \mathbf{R}}{\partial \mathbf{q}} \right|_{\mathbf{q}} = \text{diag}(J(q(t_0)), \dots, J(q(t_{2N}))) \quad (3.35)$$

The TSM Jacobian can also be written in block form, which reads:

$$\mathbf{J}^{TSM}(\mathbf{q}) = \begin{pmatrix} J(q(t_0)) & \omega d_1 M & \dots & \omega d_{2N} M \\ \omega d_{2N} M & J(q(t_1)) & \ddots & \\ & \ddots & \ddots & \omega d_1 M \\ \omega d_1 M & & \omega d_{2N} M & J(q(t_{2N})) \end{pmatrix} \quad (3.36)$$

where one can identify the diagonal blocks that are provided by \mathbf{J} , and the off-diagonal blocks that are provided by the time-derivative contribution $\omega \mathbf{D}$.

The pseudo timestep τ introduced in eq. (3.34) has no physical meaning. It is numerical parameter only used to improve the robustness of the nonlinear iterations when the system is particularly stiff, see *e.g.* [Crivellini et al. 2011]. It is defined from

a CFL number:

$$\tau = \frac{h_K CFL}{U_\infty} \quad (3.37)$$

with h_K the local mesh size. For the first iterations, a small CFL is generally used as it avoids over-corrections when the Newton iterate \mathbf{q}^i is far from convergence and the likely blow-up of the algorithm. As $\|\mathbf{R}^{TSM}(\mathbf{q}^i)\|_2$ decreases, CFL is progressively increased until it reaches infinity. In that limit, the classical Newton method is retrieved and, with it, the quadratic convergence properties that can drastically accelerate convergence with respect to the case of a constant finite CFL . Note finally that in the TSM literature, it is common to use a CFL definition that depends on the maximum resolvable frequency $N\omega$: $\tau = h_K CFL / (U_\infty + N\omega h_K)$ [Sicot et al. 2008; Antheaume et al. 2011; Gopinath et al. 2006; Mundis et al. 2014]. With this definition, the pseudo timestep naturally decreases as N or ω increase, thus changing the properties of the linear system eq. (3.34) and making assessment of N -robustness less natural. Recently, researchers targeting wavenumber independent TSM preconditioning [Mundis et al. 2015; Mundis et al. 2017] have switched back to using the classical definition of CFL eq. (3.37), that we adopt as well.

3.3.2 Solution of the linearized TSM equations

At each Newton iteration eq. (3.34), a $K \times (2N + 1)$ unknowns linearized TSM system must be solved:

$$\left(\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM} \right) \mathbf{x} = \mathbf{y} \quad (3.38)$$

Sparse LU solvers A first robust and simple alternative to numerically solve eq. (3.38) is to use sparse direct linear solvers based on LU factorization (MUMPS, SuperLU, etc). However, direct solvers are known to be highly memory intensive for large number of unknowns, and even more so when the sparsity of the matrix decreases. In the case of the TSM Jacobian, both effects are combined. Indeed, not only the total number of unknowns increases while refining the spatial (K large) and/or time (N large) discretization, but in addition, the sparsity pattern fills up as N increases. These trends can be better grasped with the numerical example provided in table 3.3.1, where some features of the mono-instant and TSM Jacobians are reported, for a simple 2D flow. It is seen that both the size of the TSM system and the average number of nonzero coefficients per row linearly increases with N . A direct consequence is that the LU solver (MUMPS here) requires more and more memory and time to factorize \mathbf{J}^{TSM} . More importantly, these increasing costs scale approximately with N^2 . As a result, when TSM is applied to space discretizations of the Navier–Stokes equations, even in 2D, the memory requirements quickly become out of reach for practical computations.

	size	# of nnz per row	memory [MB]	time [s]
$J(q(t_i))$	47,510	28	90	1
$\mathbf{J}^{TSM}(N=1)$	142,530	48	1,011	23
$\mathbf{J}^{TSM}(N=2)$	237,550	68	3,055	103
$\mathbf{J}^{TSM}(N=3)$	332,570	88	6,199	323

TABLE 3.3.1: Computational cost of using the direct sparse solver MUMPS on the full TSM linearized problem eq. (3.38). The size, average number of nonzero coefficients per row, peak factorization memory and factorization time are shown for the one-instant Jacobian and the TSM Jacobian with different N . The system considered in this example is the finite element Taylor–Hood discretization of the classical 2D laminar flow around a circular cylinder at $Re = 55$. The (spatial) mesh is made of 10,385 triangles, yielding $K = 47,510$ degrees of freedom per instant.

Preconditioned Krylov subspace methods To circumvent the untractable memory requirements associated with the direct linear solvers highlighted in the previous paragraph, we propose to use an iterative Krylov subspace method to solve eq. (3.38), namely, the Generalized Minimal RESidual algorithm (GMRES) [Saad 1993]. In brief, GMRES searches the solution of eq. (3.38) in a low-dimensional subspace of dimension $m \ll K \times (2N + 1)$, known as a Krylov subspace:

$$\mathcal{K}_m \left(\left(\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM} \right), \mathbf{r}_0 \right) = \text{span} \left\{ \mathbf{r}_0, \left(\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM} \right)^1 \mathbf{r}_0, \dots, \left(\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM} \right)^{m-1} \mathbf{r}_0 \right\}$$

with \mathbf{r}_0 some vector. Building \mathcal{K}_m thus only requires to apply $1/\tau \mathbf{M} + \mathbf{J}^{TSM}$ to a given sequence of vectors. One of the most powerful results for Krylov methods is that they are known to converge in at most in $K \times (2N + 1)$ iterations (*e.g.* Liesen et al. 2004 for a review). However, the convergence rate remains highly dependent on the particular problem. For problems as complex as the linearized TSM system, preconditioning is required to obtain solutions in reasonable cpu-time [Mundis et al. 2014; Mundis et al. 2015; Mundis et al. 2017]. In the following, we use right preconditioning, *i.e.*, GMRES is applied to the following preconditioned linear system, instead of directly to eq. (3.38):

$$\left(\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM} \right) \mathbf{P}^{-1} \mathbf{z} = \mathbf{y} \quad (3.39)$$

where \mathbf{P} is the preconditioner. The wanted solution \mathbf{x} is finally retrieved by solving

$$\mathbf{P} \mathbf{x} = \mathbf{z}$$

A good preconditioner is an operator \mathbf{P} that combines the two following properties: (i) the spectrum of $(1/\tau \mathbf{M} + \mathbf{J}^{TSM}) \mathbf{P}^{-1}$ is well clustered around 1 and (ii) applying \mathbf{P}^{-1} to a given vector has a low cost. A qualitative rule of thoughts to obtain such a preconditioner is that \mathbf{P}^{-1} should be a good enough but cheap approximation of $1/\tau \mathbf{M} + \mathbf{J}^{TSM}$.

3.3.3 A block-circulant preconditioner for the TSM Jacobian

We now introduce the so-called *block-circulant* preconditioner, designed in this thesis. First, let us recall the block structure of the linearized TSM system eq. (3.38):

$$\frac{1}{\tau} \mathbf{M} + \mathbf{J}^{TSM} = \begin{pmatrix} \frac{1}{\tau} M + J(t_0) & \omega d_1 M & \dots & \omega d_{2N} M \\ \omega d_{2N} M & \frac{1}{\tau} M + J(t_1) & \ddots & \vdots \\ \vdots & \ddots & \ddots & \omega d_1 M \\ \omega d_1 M & \dots & \omega d_{2N} M & \frac{1}{\tau} M + J(t_{2N}) \end{pmatrix} \quad (3.40)$$

The preconditioner we propose is based on the observation that the above system matrix possesses a pattern very close to a circulant matrix. More precisely, only the diagonal blocks break the circulant pattern, because the stationary Jacobian J is evaluated at a different time instant on each line. In order to retrieve a circulant matrix, we need to replace all the $J(t_i)$ on the diagonal by a same matrix. We propose to use the Jacobian evaluated on the mean flow $J(\tilde{q}_0)$, so that the block-circulant preconditioner \mathbf{P}_C is formed:

$$\mathbf{P}_C = \begin{pmatrix} \frac{1}{\tau} M + J(\tilde{q}_0) & \omega d_1 M & \dots & \omega d_{2N} M \\ \omega d_{2N} M & \frac{1}{\tau} M + J(\tilde{q}_0) & \ddots & \vdots \\ \vdots & \ddots & \ddots & \omega d_1 M \\ \omega d_1 M & \dots & \omega d_{2N} M & \frac{1}{\tau} M + J(\tilde{q}_0) \end{pmatrix} \quad (3.41)$$

Here, we choose to use the Jacobian *evaluated at the mean flow* $J(\tilde{q}_0)$ as the diagonal but we could opt instead, for example, for the *mean Jacobian*, \tilde{J}_0 . This alternative is currently explored by Diogo Ferreira Sabino (PhD student at ONERA-DAAA/MAPE) who is testing the circulant preconditioner for compressible flow applications, using ONERA's discontinuous Galerkin code, AGHORA. Note that, if the Jacobian depends linearly on the state — as in the incompressible Navier–Stokes case in fixed domain — then the two alternatives are mathematically equivalent. However, their algorithmic cost strongly differs. On the one hand, using $J(\tilde{q}_0)$ requires to (i) form \tilde{q}_0 and (ii) build the Jacobian at \tilde{q}_0 , where step (i) is cheap (linear combination of vectors) and step (ii) is more costly. On the other hand, using \tilde{J}_0 does not require any new construction of Jacobian matrix from a solution vector, but requires to combine linearly the $2N + 1$ mono-instant Jacobian matrices, $J(t_i)$. These being matrices of already large dimensions (size of the spatial discretization), manipulating them can reveal cumbersome. This issue is expected to be particularly stringent in the perspective of a parallel implementation, as proposed in section 3.3.4, where assembling \tilde{J}_0 would require to communicate the $J(t_i)$ between processes. It is probable that passing such large objects between processes deteriorates the parallel performance of the solver. In any case, quantitative comparisons would be needed in order to further compare these alternatives.

Application of \mathbf{P}_C^{-1} In a preconditioned GMRES algorithm, one needs to provide a routine that efficiently applies the preconditioner's inverse \mathbf{P}_C^{-1} to a given vector, say \mathbf{x}_{in} , and stores it in \mathbf{x}_{out} . This comes down to solving a system $\mathbf{P}_C \mathbf{x}_{\text{out}} = \mathbf{x}_{\text{in}}$. Despite the block-circulant form of \mathbf{P}_C , such systems still cannot be handled with

direct solvers (cf section 3.3.2), as \mathbf{P}_C possesses the same size and sparsity pattern as the original system matrix, $1/\tau \mathbf{M} + \mathbf{J}^{TSM}$. Instead, we exploit the fact that block-circulant matrices can be block-diagonalized using the discrete Fourier transform:

$$\mathbf{P}_C = \mathbf{F}^{-1} \widetilde{\mathbf{P}}_C \mathbf{F}$$

where the block-diagonal matrix $\widetilde{\mathbf{P}}_C$ is defined as

$$\widetilde{\mathbf{P}}_C = \begin{pmatrix} \widetilde{P}_{C-N} & & & & \\ & \ddots & & & \\ & & \widetilde{P}_{C0} & & \\ & & & \ddots & \\ & & & & \widetilde{P}_{CN} \end{pmatrix} \quad \text{and} \quad \widetilde{P}_{Ck} = \left(\frac{1}{\tau} + i\omega k \right) M + J(\tilde{q}_0) \quad (3.42)$$

The above property simply follows from the fact that

$$\begin{aligned} \mathbf{P}_C &= \left(\frac{1}{\tau} M + J(\tilde{q}_0) \right) \mathbf{I} && + \omega \mathbf{D} \\ \mathbf{P}_C &= \left(\frac{1}{\tau} M + J(\tilde{q}_0) \right) \mathbf{F}^{-1} \mathbf{I} \mathbf{F} && + \omega \mathbf{F}^{-1} \widehat{\mathbf{D}} \mathbf{F} \\ \mathbf{P}_C &= \mathbf{F}^{-1} \left[\left(\frac{1}{\tau} M + J(\tilde{q}_0) \right) \mathbf{I} + \omega \widehat{\mathbf{D}} \right] \mathbf{F} \end{aligned}$$

where $\mathbf{I} = \text{diag}(I, \dots, I)$ is the $(2N+1) \times (2N+1)$ block-identity matrix (I is the $K \times K$ identity). Going from the first to the second line, we used eq. (3.26) to express \mathbf{D} as a function of $\widehat{\mathbf{D}}$. Also, let us stress the fact that $1/\tau M + J(\tilde{q}_0)$ is not a $(2N+1) \times (2N+1)$ block-matrix but only one “instantaneous” block, of dimension $K \times K$. This is the reason why we can permute $1/\tau M + J(\tilde{q}_0)$ and \mathbf{F}^{-1} when going from the second to the third line⁷.

In the end, the preconditioner inverse simply reads:

$$\mathbf{P}_C^{-1} = \mathbf{F}^{-1} \widetilde{\mathbf{P}}_C^{-1} \mathbf{F} \quad (3.43)$$

Thus applying \mathbf{P}_C^{-1} to some vector only requires: one DFT, $2N+1$ solves with the complex shifted one-instant Jacobian matrix at the mean flow \widetilde{P}_{Ck} and finally an inverse DFT.

Equivalent preconditioner in the Fourier space To gain some insights into the expected performance of the circulant preconditioner, it is helpful to consider the preconditioned linearized TSM system in the Fourier space. For simplicity, we take $\tau = +\infty$ here. Left-multiplying eq. (3.39) by \mathbf{F} , and using eq. (3.43) we have:

$$\left(\mathbf{F} \mathbf{J}^{TSM} \mathbf{F}^{-1} \right) \widetilde{\mathbf{P}}_C^{-1} \tilde{\mathbf{x}} = \tilde{\mathbf{y}}$$

with, $\tilde{\mathbf{x}} = \mathbf{F} \mathbf{x}$ and $\tilde{\mathbf{y}} = \mathbf{F} \mathbf{y}$. The above equation is a linear system where we identify the system matrix $\mathbf{F} \mathbf{J}^{TSM} \mathbf{F}^{-1}$ and the preconditioner $\widetilde{\mathbf{P}}_C$. On the one hand, the matrix

⁷This is most simply seen in the case $K = 1$ where $1/\tau M + J(\tilde{q}_0)$ is a scalar ...

$\mathbf{FJ}^{TSM}\mathbf{F}^{-1}$ develops as:

$$\mathbf{FJ}^{TSM}\mathbf{F}^{-1} = \omega\hat{\mathbf{D}} + \mathbf{FJ}\mathbf{F}^{-1} = \begin{pmatrix} -\mathbf{i}\omega NM + \tilde{J}_0 & \tilde{J}_{-1} & \dots & \tilde{J}_{-N} & \tilde{J}_N & \dots & \tilde{J}_1 \\ \tilde{J}_1 & -\mathbf{i}\omega(N-1)M + \tilde{J}_0 & \tilde{J}_{-1} & \dots & \dots & \dots & \vdots \\ \vdots & \dots & \dots & \dots & \dots & \dots & \tilde{J}_N \\ \tilde{J}_N & \dots & \dots & \dots & \dots & \dots & \tilde{J}_{-N} \\ \tilde{J}_{-N} & \dots & \dots & \dots & \dots & \dots & \vdots \\ \vdots & \dots & \dots & \dots & \tilde{J}_1 & \mathbf{i}\omega(N-1)M + \tilde{J}_0 & \tilde{J}_{-1} \\ \tilde{J}_{-1} & \dots & \tilde{J}_{-N} & \tilde{J}_N & \dots & \tilde{J}_1 & \mathbf{i}\omega NM + \tilde{J}_0 \end{pmatrix}$$

where the \tilde{J}_k are the discrete Fourier coefficient of $J(q(t))$. On the other hand, from its definition eq. (3.42), the preconditioner $\tilde{\mathbf{P}}_C$ is seen to correspond to the diagonal of the above matrix where $J(\tilde{q}_0)$ substitutes for \tilde{J}_0 . As a consequence, the block-circulant preconditioner for TSM is close to a block-diagonal preconditioner in the Fourier space. It is thus expected that the preconditioner will perform greatly in cases where \tilde{J}_0 dominates over the \tilde{J}_k for $|k| \geq 1$, or by extension for the cases where the solution around which the system is linearized verifies " $\tilde{q}_0 \gg \tilde{q}_k$ ".

3.3.4 Parallel implementation

To handle efficiently large number of harmonics, the TSM solution method is implemented in parallel. Each one of the $2N + 1$ unknown instants $q(t_n)$ is handled by a different MPI process, so that the number of degrees of freedom per process remains constant for all N . The Jacobian matrix \mathbf{J}^{TSM} (cf eq. (3.36)) is distributed by rows: each processor possesses the *entire* row of the matrix corresponding to the degrees of freedom it handles. In practice this comes down to storing $J(q(t_n))$, M and the coefficients d_n .

Due to the global nature of the spectral derivative operator, the time derivative at each instant depends on the solution at all other instants. If the instants are distributed in parallel, then a "brute force" way of computing those time derivatives is to make each process send its local solution to all $2N$ other processes, and then compute the derivatives locally. This strategy has two downsides, that may become cumbersome as N increases. First, it means that each process stores all $2N + 1$ instants, at least during the time derivative computation. Second, it requires $2N \times (2N + 1)$ parallel communications of a vector of size K . In order to bypass those shortcomings, we rather adopt a "round-robin" strategy (fig. 3.3.1), as proposed by Mundis et al. (2017). At the first round-robin iteration, each process \mathcal{P}_n passes its local solution $q(t_n)$ only to its neighbor \mathcal{P}_{n+1} , which stores it in the same memory slot as it used to store its own local solution $q(t_{n+1})$ in. Each process \mathcal{P}_{n+1} is then able to add the contribution proportional to $q(t_n)$ to its local time derivative, i.e. $\omega d_{2N} M q(t_n)$. By performing a total of $2N + 1$ such iterations, each process has been given access to all instants (and finally gets back its own local instant). Contrary to the brute force approach, the round-robin strategy requires to store only one vector of the size K of the spatial discretization per process⁸. In addition, only $2N + 1$ communications to a direct neighbor are performed.

⁸In practice we store the received solution in a temporary second container, so we actually store two vectors of size K ...




FIGURE 3.3.1: Computation of global operators through a “round-robin” strategy. Each line corresponds to a process and each column to a round-robin iteration. In the circles, we indicate the instant possessed by the process at the *beginning* of the current round-robin iteration. The arrows materialize the parallel communications. The last column gives the distribution of the instants at the end of the round-robin loop.

Note that an identical round-robin strategy is used to perform the direct and inverse DFT’s needed in the application of the circulant preconditioner.

Finally, it is worth noticing that this implementation of the spectral derivative operator leads to a cost of $2N$ matrix vector product of the type $Mq(t_k)$ and as much vector summations per process. As a consequence, even supposing infinitely fast communications between processes, this implementation cannot scale better than linearly in terms of wall-clock time. The same is true in the preconditioner application, where DFT is used.

3.4 Numerical results

In this section, we study the performance of the circulant preconditioner on two numerical examples. The first configuration is a NACA0012 profile forced in heaving motion at different frequencies. The second configuration is the spring-mounted flat plate introduced in chapter 1, undergoing a self-sustained flutter instability. For this second case, the solution method presented in section 3.3 is extended to the case where the frequency is unknown.

3.4.1 NACA0012 in forced heaving motion

A NACA0012 profile, immersed in a flow with uniform upstream velocity $U_\infty \mathbf{e}_x$, is forced to oscillate along the vertical direction (heaving motion):

$$h(t) = h_1 \sin(\omega t)$$

where $h_1 = 0.08$ is the non-dimensional⁹ plunging amplitude and ω the reduced frequency. The Reynolds number based on the chord is $Re = 1850$. The numerical values specified here are taken from [Leffell et al. 2014]. However, contrary to those authors who considered a slightly compressible flow ($Ma = 0.2$), we use the incompressible Navier–Stokes equations. In addition, we vary the forcing frequency between $\omega = 0.1$ and 3.6. In particular, we will see that a high-frequency like $\omega = 3.6$ provides a challenging test case from the point of view of the numerical solution of the TSM equations, because a high number of harmonics must be taken into account to capture its temporal behavior.

In all the numerical examples of section 3.4, we handle the moving fluid domain with the “Reference configuration ALE” formalism introduced in section 1.1.2.2. The spatial discretization is taken care of, as in section 1.2.2, with the (Galerkin) finite element method. It yields a discretized system of equations similar to eq. (3.1):

$$M(q) \frac{\partial q}{\partial t} + R(q) = f(t)$$

We note that the above system is not strictly identical to eq. (3.1) since the mass matrix depends here on q . This is due to the fact that, in the “Reference configuration ALE” fluid momentum equation eq. (1.11a), the mesh deformation Jacobian appears in front of the time-derivative. From the TSM solution method point of view, the main consequence is that $M(q)$ must also be “averaged” when building the block-circulant preconditioner \mathbf{P}_C :

$$\mathbf{P}_C = \begin{pmatrix} \frac{1}{\tau}M(\tilde{q}_0) + J(\tilde{q}_0) & \omega d_1M(\tilde{q}_0) & \dots & \omega d_{2N}M(\tilde{q}_0) \\ \omega d_{2N}M(\tilde{q}_0) & \frac{1}{\tau}M(\tilde{q}_0) + J(\tilde{q}_0) & \ddots & \vdots \\ \vdots & \ddots & \ddots & \omega d_1M(\tilde{q}_0) \\ \omega d_1M(\tilde{q}_0) & \dots & \omega d_{2N}M(\tilde{q}_0) & \frac{1}{\tau}M(\tilde{q}_0) + J(\tilde{q}_0) \end{pmatrix}$$

Other than that, the solution method applies as presented in section 3.3.

3.4.1.1 Reference DNS solution and TSM time-convergence

Let us start by describing a reference solution that has been computed using the time-stepping algorithm presented in section 1.3, with a number of time steps per period equal to 1400. Vorticity snapshots of the obtained periodic solution are presented in fig. 3.4.1 for three values of ω . The wake is seen to strongly evolve between (a) $\omega = 0.4$ where it is quasi-steady and (c) $\omega = 3.6$ where it is highly unsteady with strong vortices being shed. For each frequency, time and Fourier representations of the vertical velocity at point $P = (2, 0)$ are shown in rightmost part of the figure. At the low-frequency $\omega = 0.4$, the velocity signal is strongly dominated by the fundamental frequency. As ω increases, the third, fifth and seventh harmonics increase in amplitude, so that the high-frequency case $\omega = 3.6$ presents a more complex time behavior featuring abrupt variations of the velocity at point P when a vortex goes through.

⁹We recall (cf section 1.1) that in this manuscript, we non-dimensionalize using the airfoil chord c , far-field velocity U_∞ and fluid density ρ_f as reference length, velocity and density, i.e. $h = h^\diamond/c$ and $\omega = \omega^\diamond c/U_\infty$, where the dimensional quantities are marked with the \diamond symbol.

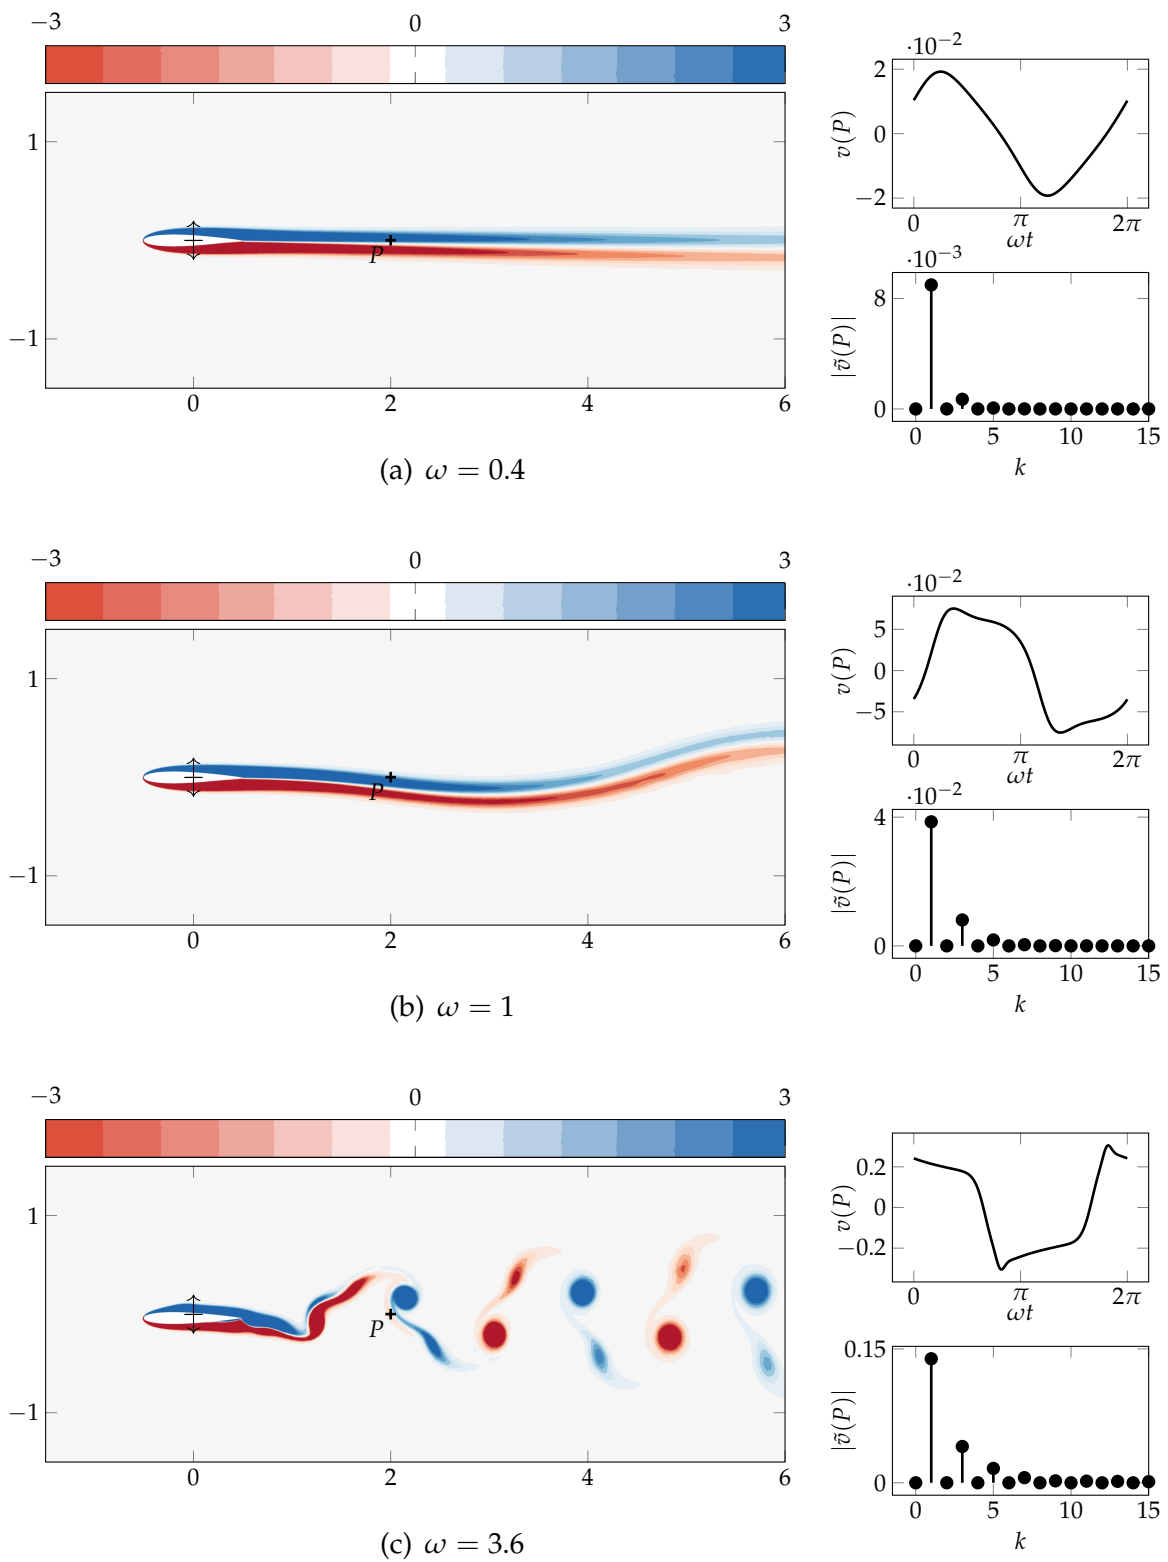


FIGURE 3.4.1: NACA0012 airfoil forced in plunging motion at different frequencies ω . Vorticity snapshots are shown on the left whereas the time signal of vertical velocity at point P is shown on the right (top: temporal series, bottom: Fourier spectrum).

In fig. 3.4.2, we present a study of the convergence of the TSM solution towards the previously described reference solution, as the number of harmonics N increases. Only the case with forcing frequency $\omega = 3.6$ is shown as it is the hardest to converge, due to the presence of higher harmonics. In fig. 3.4.2a, the vertical velocity signal at the probe point (P) is shown for $N = 5$ (squares) and $N = 10$ (triangles). The symbols correspond to the TSM collocation points t_n whereas the thin solid lines show the Fourier series reconstruction on the continuous time interval. Finally, the reference solution is marked by a thick black solid line. Clearly, for $N = 5$, the velocity signal is polluted by point-to-point oscillations of significant amplitude. These oscillations reduce for $N = 10$, though still visible, for example just before $\omega t = \pi$. For $N = 20$ (not shown) they are not observed anymore. In fig. 3.4.2b, the time convergence of the TSM solution is analyzed by looking at the vorticity field at $t = 0$ for increasing N . In agreement with what was observed on the velocity probe, the $N = 5$ solution contains small spatial wavelength oscillations in the wake. Those oscillations, are still present for $N = 10$, though with smaller wavelength and lower amplitude. Above $N = 20$ the oscillation are not visible anymore and a smooth wake, similar to the one the reference solution fig. 3.4.1c is retrieved everywhere, showing that TSM has converged.


The strong non-physical high-frequency oscillations described above for $N = 5$ may seem unexpected. Indeed, according to the reference time-stepping computation, fig. 3.4.1c, the harmonics higher than the fifth are negligible in the vertical velocity signal at the probe P . As a consequence, one would expect that with $N = 5$ the truncation error due to neglecting harmonics higher than the fifth would be small and the obtained solution should be close to the reference. However, we have seen in section 3.2.3.2 that TSM, as any pseudo-spectral method, suffers from aliasing errors in addition to classical truncation errors. As a consequence, despite $N = 5$ is likely to yield small truncation error, it may still present large aliasing error. Another argument in favor of explaining the oscillations by aliasing error is that they are high-frequency: this is in qualitative agreement with the aliasing phenomenon that tends to pollute the low-harmonics equations with high-frequency residual contributions, see eq. (3.31).

3.4.1.2 Assessment of the TSM solver performance


We now assess the numerical performance of the proposed TSM solution method. To do so, we perform several TSM computations for different forcing frequencies and numbers of harmonics. The nonlinear iterations are provided by the Newton method presented in section 3.3.1, where each linear system of the form eq. (3.34) is solved by the GMRES algorithm, preconditioned by the block-circulant preconditioner proposed in ???. The numerical pseudo time-step $\tau = h_K CFL / U_\infty$, introduced in the Newton iteration eq. (3.34), is governed by the following CFL law:

$$CFL = \max \left(\min \left(\frac{CFL_{\min}}{r_i^{1.5}} (i + 1 - i_0), CFL_{\max} \right), CFL_{\min} \right)$$

where $r_i = \|\mathbf{R}^{TSM}(\mathbf{q}_i)\|_2 / \|\mathbf{R}^{TSM}(\mathbf{q}_{i_0})\|_2$ is the residual reduction factor between the current iteration i and the iteration i_0 at which the CFL starts to vary. For $i < i_0$, we keep $CFL = CFL_{\min}$. In the following, we use $i_0 = 2$. The Newton iterations are stopped when the l_2 norm of the nonlinear residual $\|\mathbf{R}^{TSM}(\mathbf{q})\|_2$ is lower than 10^{-9} . The GMRES relative tolerance is set to 10^{-4} and the Krylov subspace size is set to $m = 200$, with no restart allowed. The initial guess for both the Newton and



(a) Vertical velocity at the probe point P as a function of time for $N = 5$ (squares) and $N = 10$ (triangles). The symbols indicate the TSM collocation points, whereas the thin solid lines shows the Fourier series reconstruction on the continuous time interval. The reference solution is shown with thick black line.



(b) Vorticity snapshots for different N

FIGURE 3.4.2: Effect of the number of harmonics N on the TSM solutions.

the GMRES iterations is the zero vector. Note that for such a rough initial guess, the use of a small enough CFL — here $CFL_{\min} = 100$ — for the first Newton iterations is necessary so as to avoid blow-up of the Newton method.

Linear solver performance First we want to measure the performance of the linear solver alone. To do so, we report in table 3.4.1 the number of GMRES iterations performed at the last Newton step. The advantage of taking the last Newton step is that the maximal CFL ($CFL_{\max} = 10^{10}$ here) is always reached, yielding the *true* linearized TSM system, without any additional term. Thus, the performance of the preconditioner can be assessed without it “getting help” from a potentially large diagonally dominant contribution (the $1/\tau M$ term in eq. (3.40)). We note that another possibility would have been to report the average number of GMRES iterations per Newton step across the whole Newton run. We however think this measure is less faithful, as the CFL evolution — *i.e.* the magnitude of $1/\tau M$ term — is not a parameter of the method. It rather depends on the path taken by the Newton algorithm to converge. For example, if the algorithm stays for a while in a region where the CFL is low, the average number of GMRES iterations per Newton step will decrease because those low CFL iterations benefit from a better conditioned system.

Coming back to table 3.4.1, we observe first that for low forcing frequencies $\omega \leq 0.4$, the GMRES iterations count is always around 10 and is completely independent of N . For higher frequencies ($\omega = 0.8$ and 1), a mild N -dependence is observed, particularly visible when comparing $N = 20$ and $N = 40$. This difference of behavior between full N -robustness at low ω and mild N -dependence at higher ω is not explained. For the highest forcing frequency $\omega = 3.6$, the maximum number of 200 iterations is reached. This deterioration could be expected due to the fact that substituting $J(\tilde{q}_0)$ to the instantaneous Jacobian $J(q(t_n))$, as done for building the circulant preconditioner (see eq. (3.41)), is likely to be a less and less good approximation as ω increases. Indeed, it is shown in fig. 3.4.3 that, in the wake, the instantaneous solution becomes very different from the mean solution as ω increases.

$\omega \backslash N$	3	5	10	20	40
0.1	9	9	10	9	9
0.2	9	9	9	9	9
0.4	9	9	9	9	9
0.8	22	21	22	25	37
1	38	34	42	50	86
3.6	200	200	200	200	200

TABLE 3.4.1: Performance of the block-circulant preconditioner for different forcing frequencies and number of harmonics. For each couple (ω, N) , we report the number of GMRES iterations at the last Newton iteration (for which infinite CFL is always reached).

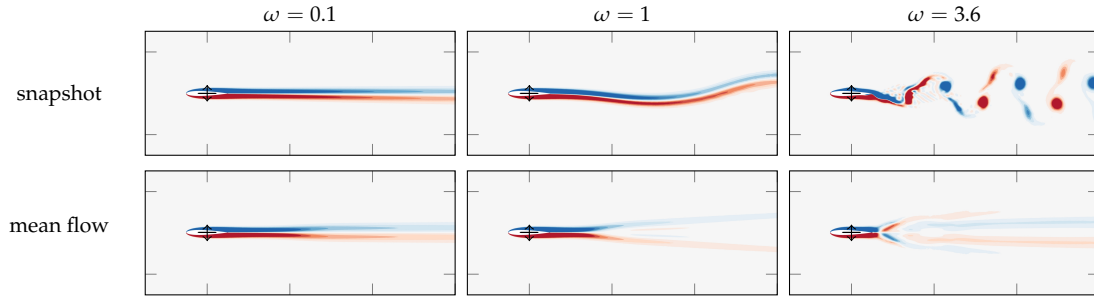


FIGURE 3.4.3: Instantaneous solution *vs* mean flow for different reduced frequencies. ($Re = 1850$, $h_0 = 0.08$, $N = 10$).

Finally, we study the efficiency of the linear solver in terms of wall-clock time in fig. 3.4.4. For different numbers of harmonics, the total wall-clock time spent in solving the linearized TSM system at the last Newton iteration is marked by cross symbols. Despite the N -independent GMRES iteration counts (see numbers between parenthesis) for the chosen forcing frequency ($\omega = 0.1$), it is observed that the wall-clock time scales linearly with N . To understand this behavior, the linear solution process is split into two phases that are: preconditioner set-up (red) and GMRES routine (blue). The preconditioner set-up is then further split into two contributions. The first contribution, shown in light red, consists in the building of \mathbf{P}_C (*i.e.* computing the mean solution \tilde{q}_0 , constructing $J(\tilde{q}_0)$ and setting-up all distributed PETSc objects). The second contribution, shown in dark red, consists in the LU factorization of the diagonal blocks of $\tilde{\mathbf{P}}_C$. The GMRES phase is also split into two contributions. First, the time spent in applying the matrix vector product (light blue) and second, the time spent in applying the preconditioner (dark blue). At this level of detail, we observe that the two phases scale differently. On the one hand, the preconditioner set-up phase is roughly independent of N . This is because the most time-consuming operations of this phase, that are (i) the assembly of $J(\tilde{q}_0)$, and (ii) the LU factorization of the diagonal blocks of $\tilde{\mathbf{P}}_C$, are both purely sequential operations. They are thus ideally parallelized. On the other hand, the GMRES phase scales linearly with N . This is in fact the best expected scenario, given the implementation of the spectral derivative operator through eq. (3.25) and our DFT-based preconditioner (*cf* section 3.3.4 for details). In other words, we stress here the fact that this linear scaling is not a consequence of a poor implementation of parallelism, but an intrinsic limitation of using DFT-based time-global operators. Still, a parallel implementation is significantly better than a purely sequential one, that would scale quadratically with N .




FIGURE 3.4.4: Wall-clock time spent in the solution of the linear system at the last Newton iteration, for different N . The forcing frequency is $\omega = 0.1$. The total time (cross symbols) is split into two phases: preconditioner set-up (red) and GMRES routine (blue). The preconditioner set-up is further split into the building of the preconditioner (light red) and the LU factorization of its diagonal blocks (dark red). The GMRES phase is split into the time spent applying the matrix vector product (light blue) and applying the preconditioner (dark blue). Finally, the GMRES iteration counts are recalled between parenthesis.

Nonlinear solver performance Having assessed the efficiency of the circulant preconditioner for solving the linearized TSM system, we now focus on the performance of the full Newton–Krylov method. In fig. 3.4.5 we report the evolution of the l_2 norm of the nonlinear residual $\|\mathbf{R}^{TSM}(\mathbf{q})\|_2$, the number GMRES iterations and the CFL as a function of the nonlinear iterations. Three forcing frequencies are compared. Let us focus first on the lowest frequency case, $\omega = 0.4$ (circles). We observe two phases in the evolution of the nonlinear residual: in the first phase, the residual slowly decreases as the CFL number is low, whereas in the second phase a fast quadratic convergence typical of the Newton method is retrieved, thanks to the CFL reaching high enough values. Finally, the GMRES iteration count shows that convergence of the linear solver is reached (up to the chosen tolerance of 10^{-4}) at each Newton iteration. The case $\omega = 1$ (triangles) behaves similarly though with higher GMRES iteration counts, which is consistent with the results of table 3.4.1. For $\omega = 3.6$, the GMRES algorithm is unable to converge to the required tolerance above some critical CFL value ($\simeq 10^3$ here). As a consequence, only approximate Newton increments eq. (3.34) are computed, yielding a degraded nonlinear convergence even with high CFL.

In fig. 3.4.6, we investigate the robustness of the Newton–Krylov approach when increasing the number of harmonics. At low enough forcing frequency ($\omega = 0.4$, fig. 3.4.6a), all indicators reported exactly superimpose, showing the robustness of the approach, for N between $N = 3$ and $N = 40$. For a higher forcing frequency $\omega = 1$, similar robustness is observed in terms of nonlinear residuals. However, as already mentioned in table 3.4.1, the GMRES iterations count mildly degrades for high N . This does not impair nonlinear convergence as GMRES still converges, but surely increases¹⁰ the wall-clock time necessary for convergence.

¹⁰... in addition to the “incompressible” linear increase of wall-clock time with N due to the linear solver (see fig. 3.4.4)




FIGURE 3.4.5: Performance of the Newton–Krylov approach different reduced frequencies: $\omega = 0.4$ (circles), $\omega = 1$ (triangles), $\omega = 3.6$ (squares). The number of harmonics is $N = 20$.



(a) $\omega = 0.4$ (b) $\omega = 1$

FIGURE 3.4.6: Performance of the Newton-Krylov approach when increasing the number of harmonics: $N = 3$ (crosses), $N = 5$ (circles), $N = 10$ (squares) and $N = 40$ ("plus" signs). The results are presented for $\omega = 0.4$ (a) and $\omega = 1$ (b).

Comparison of \mathbf{P}_C to the block-Jacobi preconditioner We conclude our study by comparing the proposed approach to a more classical one based on block-Jacobi iterations. In order to allow straightforward comparison, we use here block-Jacobi as a preconditioner for GMRES like in [Mundis et al. 2014; Mundis et al. 2015; Mundis et al. 2017] (referred to as “GMRES-EX” there), and not as fixed-point iteration as in [Sicot et al. 2008]. The block-Jacobi preconditioner writes:


$$\mathbf{P}_{BJ} = \begin{pmatrix} \frac{1}{\tau}M(q(t_0)) + J(q(t_0)) & 0 & \dots & 0 \\ 0 & \frac{1}{\tau}M(q(t_1)) + J(q(t_1)) & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \frac{1}{\tau}M(q(t_{2N})) + J(q(t_{2N})) \end{pmatrix} \quad (3.44)$$

yielding the following advantages over \mathbf{P}_C . First, all blocks involved are already assembled as parts of the TSM Jacobian. Second, no going back and forth between physical and Fourier spaces are needed. Finally, the one-instant linear systems to be solved are real-valued and can be directly handled by pre-existing (linearized) steady-state solvers.


In fig. 3.4.7, we propose a comparison of the nonlinear convergences obtained with \mathbf{P}_C (circles) and \mathbf{P}_{BJ} (crosses). The same indicators as in figures 3.4.5 and 3.4.6 are shown again. However, due to the difference of cost in applying \mathbf{P}_C^{-1} and \mathbf{P}_{BJ}^{-1} , the indicators are monitor *versus* wall-clock time, instead of Newton iterations.

For the one harmonic case $N = 1$ (fig. 3.4.7a), the block Jacobi approach is seen to converge about 30% faster than the circulant preconditioner. For such a low N , the GMRES algorithm always converges with both \mathbf{P}_{BJ} and \mathbf{P}_C , thus yielding the fast Newton convergence properties. Despite a much higher GMRES iteration count, the block Jacobi strategy is more efficient in this case, due to lower set-up and preconditioner application costs (not shown).

Now turning to the case $N = 5$ (fig. 3.4.7b), it is seen that with block-Jacobi, GMRES is mostly unable to converge and always reaches the maximum 200 iterations after 400 seconds. As a consequence, the Newton increments are not computed precisely and the nonlinear convergence stalls around $\|\mathbf{R}^{TSM}(\mathbf{q})\|_2 \simeq 10^{-3}$. On the contrary, the circulant preconditioner, thanks to its N -robustness conserves proper GMRES convergence, and thus yields the expected quadratic Newton convergence.



(a) $N = 1$



(b) $N = 5$

FIGURE 3.4.7: Comparison of the Newton–Krylov solver convergence using preconditioner P_C (circles) or P_{BJ} (crosses). The cases $N = 1$ is shown in (a) whereas $N = 5$ is shown in (b). The forcing frequency is $\omega = 0.4$.

3.4.2 Flutter instability of a 2DOF spring-mounted plate: a case with unknown frequency

Having demonstrated the efficiency of the Newton–Krylov approach in the previous paragraphs, we are now interested in using it in order to compute the self-sustained periodic solutions that emerge from the flutter instability of the 2DOF spring-mounted plate, presented in chapter 1 (and further studied in chapter 4 to 6). Contrary to the previous test case, the movement of the solid is not imposed anymore, but is an unknown of the problem.

When dealing with self-sustained LCO's the frequency ω is an additional unknown of the problem. As a consequence, an additional equation must be added to the TSM equations, in order to ensure a well-posed problem. Typically, this additional equation sets the phase of the researched LCO. We choose to impose that the pitching angle reaches an extremum at t_0 , which yields using the spectral derivation operator, eq. (3.23):

$$\dot{\theta}(t_0) \simeq \sum_{n=0}^{2N} d_n \theta(t_n) = 0$$

In the following, we refer to this equation as the *phase constraint*. From there, a Newton–Krylov method could be applied, for some given reduced velocity U^* , to the TSM system augmented by the phase constraint.

In this work, we rather opt for a slightly different approach, inspired from [S. He et al. 2018], where the imposed parameter is not the bifurcation parameter (U^* here) but rather an observable of the researched solution. Typically, this observable may be the pitching angle amplitude, noted \mathcal{A}_θ . Thus, in this strategy U^* is yet another unknown of the problem. The accompanying additional scalar equation, referred to as the *amplitude constraint*, fixes the (extremum, cf the phase constraint) value of θ at t_0 :

$$\theta(t_0) = \mathcal{A}_\theta$$

As a consequence, the nonlinear system to be solved writes:

$$\begin{aligned} \mathbf{R}^{TSM, self}(\mathbf{q}, [\omega, U^*]) &= \begin{pmatrix} \mathbf{R}^{TSM}(\mathbf{q}, [\omega, U^*]) \\ \mathbf{C}(\mathbf{q}) \end{pmatrix} = \mathbf{0} \\ \text{with } \mathbf{C}(\mathbf{q}) &= \begin{pmatrix} \sum_{n=0}^{2N} d_n \theta(t_n) \\ \theta(t_0) - \mathcal{A}_\theta \end{pmatrix} \end{aligned} \quad (3.45)$$

where $\mathbf{C}(\mathbf{q})$ gathers the two additional constraints introduced above.

Remark. In the case U^* is a parameter (and not an unknown), the amplitude constraint is not introduced and \mathbf{C} would be a simply: $\mathbf{C}(\mathbf{q}) = \sum_{n=0}^{2N} d_n \theta(t_n)$.

Preconditioning of the bordered Jacobian matrix The Jacobian of eq. (3.45) presents a typical *bordered* structure:

$$\begin{aligned} \mathbf{J}^{TSM, self}(\mathbf{q}, [\omega, U^*]) &= \begin{pmatrix} \mathbf{J}^{TSM}(\mathbf{q}, [\omega, U^*]) & \mathbf{b}(\mathbf{q}, [\omega, U^*]) \\ \mathbf{c}(\mathbf{q})^T & \mathbf{0} \end{pmatrix} \\ \text{with } \mathbf{b}(\mathbf{q}, [\omega, U^*]) &= \frac{\partial \mathbf{R}^{TSM}}{\partial [\omega, U^*]^T} \Big|_{\mathbf{q}, [\omega, U^*]} \quad \text{and} \quad \mathbf{c}(\mathbf{q}) = \frac{\partial \mathbf{C}}{\partial \mathbf{q}} \Big|_{\mathbf{q}} \end{aligned} \quad (3.46)$$

An optimal preconditioner for the bordered Jacobian can be built from the \mathbf{U} block of the following block-LU factorisation of $\mathbf{J}^{TSM, self}$:

$$\mathbf{J}^{TSM, self} = \begin{pmatrix} \mathbf{J}^{TSM} & \mathbf{b} \\ \mathbf{c}^T & \mathbf{0} \end{pmatrix} = \underbrace{\begin{pmatrix} \mathbf{I} & \mathbf{0} \\ \mathbf{c}^T \mathbf{J}^{TSM-1} & \mathbf{I} \end{pmatrix}}_{\mathbf{L}} \underbrace{\begin{pmatrix} \mathbf{J}^{TSM} & \mathbf{b} \\ \mathbf{0} & \mathbf{S} \end{pmatrix}}_{\mathbf{U}}$$

with $\mathbf{S} = -\mathbf{c}^T \mathbf{J}^{TSM-1} \mathbf{b}$ the 2-by-2 Schur complement of the additional unknowns block $[\omega, U^*]$. The preconditioned matrix $\mathbf{J}^{TSM, self} \mathbf{U}^{-1}$ has only ones as eigenvalues and would yield GMRES to converge in at most two iterations [Silvester et al. 2001].

Applying \mathbf{U}^{-1} requires first to build \mathbf{S} . This can be done, only once per Newton iteration, by computing the image of the canonical basis of \mathbb{R}^2 by \mathbf{S} , using its definition. The cost of that operation amounts to two solves for \mathbf{J}^{TSM-1} . Then, at each GMRES iteration, one must (i) apply \mathbf{S}^{-1} and (ii) solve for \mathbf{J}^{TSM-1} . Overall, the costs of one solution for the bordered linearized system amounts to at most four solves for \mathbf{J}^{TSM-1} : two to build \mathbf{S} and two, at most, to converge GMRES. However, the question of the required accuracy for those four solves is then posed, which leads to cumbersome empirical studies.

Instead, we choose to use as the preconditioner an approximate version of \mathbf{U} where each occurrence of \mathbf{J}^{TSM} is replaced by the circulant preconditioner \mathbf{P}_C , yielding:

$$\mathbf{P}_C^{auto} = \begin{pmatrix} \mathbf{P}_C & \mathbf{b} \\ \mathbf{0} & \mathbf{S}_C \end{pmatrix} \quad \text{with} \quad \mathbf{S}_C = -\mathbf{c}^T \mathbf{P}_C^{-1} \mathbf{b}$$

The cost of this strategy amounts to:

- two applications of \mathbf{P}_C^{-1} , performed only once during preconditioner set-up, in order to build \mathbf{S}_C .
- computing the dense inverse of \mathbf{S}_C
- at each GMRES iterations: one application of \mathbf{S}_C^{-1} , one matrix-vector product with \mathbf{b} and one application of \mathbf{P}_C^{-1} .

The dominant cost is obviously to apply \mathbf{P}_C^{-1} . Since this is done only once per GMRES iteration, the cost of the iteration for the banded system is roughly equivalent to that of the forced case.

Results Flutter LCO's are searched at $Re = 500$ and $\tilde{m} = 10^3$. According to the physical investigations reported later in this manuscript (see fig. 4.3.7b), the critical velocity is $U_c^* = 4.96$ and the critical frequency $\omega_c = 0.167$. Following the approach exposed above, we seek for the triplet $[\mathbf{q}, \omega, U^*]$ for two imposed maximum pitching angles $\mathcal{A}_\theta = 5^\circ$ and $\mathcal{A}_\theta = 10^\circ$. The unstructured mesh is composed of 15,274 triangles, yielding $K = 71,496$ degrees of freedom per instant. The nonlinear tolerance is set to 10^{-9} , the GMRES tolerance to 10^{-4} and the Krylov space size to 200 (no restart allowed). An initial guess for the Newton iterations is set by superimposing the steady solution and the critical flutter eigenmode obtained through linear stability analysis in chapter 4, multiplied by a small arbitrary amplitude. Thanks to this reasonable initial guess, the solution for $\mathcal{A}_\theta = 5^\circ$ could be converged with a constant infinite $CFL = +\infty$ (exact Newton method). The latter solution then served as an initial guess for the case $\mathcal{A}_\theta = 10^\circ$, which again converged with $CFL = +\infty$.

The two LCO solutions are shown in fig. 3.4.1. The vorticity snapshots on the left part of the figure are taken at $t = 0$, which corresponds to maximal instantaneous

angle of attack due to the phase constraint. At $\mathcal{A}_\theta = 5^\circ$, the flow remains attached to the plate all along the period whereas at $\mathcal{A}_\theta = 10^\circ$ a recirculation region, marked by the black solid line, forms over the whole plate. In the wake, the shear layers remain “stable” in the sense that no vortex formation is observed. This is expected due to the low frequency of the LCO $\omega \simeq 10^{-1}$ in both cases. On the right part of fig. 3.4.1, it is seen that for both amplitudes, the angle of attack (solid line) has a nearly harmonic behavior. On the contrary, the lift coefficient (dashed) contains higher harmonics at $\mathcal{A}_\theta = 10^\circ$. This is linked to the more important nonlinearity of the flow, consecutive to the appearance of the recirculation region.

In table 3.4.2, the number of Newton steps and average GMRES iteration counts per Newton step are reported for the two amplitudes and different N . The number of Newton iterations is low for both amplitudes, which is expected from the “good enough” initial guess we use. In addition, the Newton and GMRES iteration counts are both independent of N , showing that the robustness properties of the circulant preconditioner observed in forced case are conserved in the self-sustained case.

$\mathcal{A}_\theta \backslash N$	1	3	5	10	20
5°	7/28	7/27	7/27	7/26	7/26
10°	7/74	7/64	7/66	8/62	8/62

TABLE 3.4.2: Performance of the TSM solver for the self-sustained flutter LCO test case. Each cell of the table is formatted like: “number of Newton iterations / average number of GMRES iterations”.

3.5 Conclusion

Using the established classification of spectral methods in space from [Canuto et al. 1988], we proposed a synthetic view of the different HBM methods used in the aeroelasticity literature through the last twenty years. First, the direct application of a Galerkin procedure, using the Fourier basis, leads to what we refer to as the Analytical Harmonic Balance Method. In this method, the unknowns are the Fourier coefficients of the solution, and possibly cumbersome analytical manipulations are required to derive the equations verified by each of them. For more complex systems (highly nonlinear and/or high-dimensional), those analytical manipulations can be advantageously avoided by adopting a pseudo-spectral approach. This is the Nonlinear Frequency Domain Method. Another possibility is to directly formulate the problem in the time-domain, through a collocation approach. The High-Dimensional Harmonic Balance Method and the Time Spectral Method both use this strategy while they only differ in the way they evaluate the spectral time-derivative term. All three methods pseudo-spectral approaches (NLFD, HDHBM and TSM) were shown to be strictly equivalent so that they can simply be considered as different *implementations* of the pseudo-spectral scheme. The Analytical Harmonic Balance Method on the other hand remains a different method as it is, by construction, a truly spectral (aliasing-free) approach.

In the rest of the chapter, we detailed a solution method for the TSM equations based on a Newton–Krylov approach. The high-dimensional linear system arising at each Newton iteration is solved using the GMRES algorithm for which we presented a new preconditioner called circulant preconditioner. The latter is obtained




FIGURE 3.4.8: Spring-mounted thin plate undergoing coupled mode flutter LCO's of different amplitudes \mathcal{A}_θ . Vorticity snapshots at $t = 0$ are shown on the left where solid lines mark the recirculation area. The temporal series of angle of attack (solid line) and lift coefficient (dashed) are shown on the right. ($Re = 500, \tilde{m} = 10^3$)

by replacing the diagonal blocks of the TSM Jacobian matrix with the Jacobian evaluated at the mean flow. The resulting preconditioner being circulant, its inverse is efficiently computed in the frequency domain. The proposed approach has been implemented in a parallel-in-time manner in the FreeFEM library and then tested on two numerical examples.

The first example consisted in a NACA0012 airfoil forced in plunging motion at different frequencies. The performance of the linear solver alone was first investigated. For forcing frequencies up to $\omega \simeq 0.8$, nearly N -independent (for $1 \geq N \geq 40$) GMRES iteration counts were obtained for solving the linearized TSM problem without any pseudo time-step. For higher frequencies, a mild N -dependence of the GMRES iteration count was observed, which could not be explained. From a wall-clock time point of view, the linear solve was decomposed into two phases: preconditioner set-up and GMRES. The preconditioner set-up was seen to yield N -independent wall-clock time, due to its mostly sequential nature. On the contrary, the GMRES phase, which actively relies on operators with a dense-in-time pattern (spectral derivative operator and discrete Fourier transforms), showed linear scaling with N , as expected. On the downside, the circulant preconditioner was seen to exhibit a significant dependence on ω . This dependence is intrinsic to its construction as the mean flow approximation, that yields the circulant structure, is less and less verified for high ω . Then, the performance of the nonlinear solver was exposed. First, the Newton iteration count logically increases with ω , due to the more and more nonlinear nature of the solution. Second, because the nonlinear solver is fully independent of N (typically, the pseudo CFL number is only function of the spatial discretization), the N -robustness of the linear solver directly transfers to the nonlinear solver, thus yielding robust path to the LCO solution (in terms of iteration counts). The importance of using a robust linear solver was even more obvious when replacing the circulant preconditioner with a block-Jacobi preconditioner. Despite its lower setup and application costs, the block-Jacobi preconditioner was shown to be quickly outperformed when N was greater than 2 – 3. This is due to the quick degradation of GMRES iteration counts, ultimately yielding highly inaccurate linear solves, as N increases.

The second example, which consisted in computing self-sustained flutter LCO's, was the opportunity to extend our strategy to the case where a few additional unknowns have to be solved for (the frequency typically, but not only). To handle the associated scalar constraints, the circulant preconditioner was coupled with a Schur complement approach, at nearly no additional cost per GMRES iterations. The obtained preconditioner showed the same N -robust features as in the forced case.

Overall, the Newton–Krylov approach with block-circulant preconditioning was shown to provide an efficient, yet imperfect, algorithm for computing complex nonlinear LCO solutions, such as the one obtained for $\omega = 3.6$ in our forced NACA0012 test case. In order to further assess its performance, we propose, as a future work, to compare the present preconditioner with the approximate factorization preconditioner of [Mundis et al. 2017], that appears to be one of the current references for aeroelastic applications.

As it is implemented, the main drawback of our solver is its linear scaling in wall-clock time with respect to N . A straightforward solution to that issue is to follow Ramezani et al. (2017), by implementing both the spectral derivative and the circulant preconditioner through parallel-in-time *fast* Fourier transforms. With an efficient implementation, this should bring the scaling down to $\log N$.

Another possible direction of research is to improve the solver/preconditioner performance for highly nonlinear cases exhibiting strong harmonic contents. A suggestion is to investigate the effect of adding numerical viscosity in time to the classical TSM scheme. This may be motivated, for example, by the fact that [Leffell et al. \(2014\)](#) noted as a side remark that very high forcing frequencies (like $\omega = 6$) could only be computed by adding artificial viscosity. On the positive side, we mention that the circulant preconditioner could be easily adapted to the addition of a viscous second order spectral derivative in time, due to the similar structure of the linearized operator.

Finally, we should mention that, in the perspective of larger scale three-dimensional applications, it is important that efficient solvers exist to handle the complex-valued mono-instant problems that arise on the diagonal of eq. (3.42). In the framework of laminar incompressible Navier–Stokes, this challenge is partially addressed in chapter 7.

APPENDIX

Appendix 3.A Discretized operators for the numerical examples of section 3.4

In this section we provide the expressions of the discrete (in space) operators that are substituted for the generic problem eq. (3.1) in the numerical examples of section 3.4. We first present the coupled fluid-structure case used in section 3.4.2. Then, the particular case of section 3.4.1 where the solid is forced with heaving motion is deduced. In all cases, the “reference configuration ALE” formalism, introduced in section 1.1.2.2, is used. The weak formulations are obtained by following steps exactly similar to the ones detailed for the “absolute velocity – rotating axis” formalism, in section 1.2.1. In particular, the velocity continuity condition on the fluid interface is imposed weakly, by introducing the corresponding Lagrange multiplier, noted $\lambda = \Sigma_f \mathbf{n}_s$, where we recall Σ_f is the viscous stress tensor, written in the reference configuration.

Finite elements are used for spatial discretization of the Navier–Stokes equations: $(\mathcal{P}_2, \mathcal{P}_1)$ for (\mathbf{u}, p) and \mathcal{P}_2 for λ . All operators defining the finite element discretization of a differential operator are defined by their weak formulations.

3.A.1 Coupled problem: spring-mounted plate

The coupled problem for the numerical example of section 3.4.2 is detailed in this section.

3.A.1.1 Nonlinear equations

The solid and fluid equations are grouped under the generic form:

$$\begin{pmatrix} \mathbf{M}_s & 0 \\ 0 & \mathbf{M}_f(\mathbf{q}_s) \end{pmatrix} \frac{\partial}{\partial t} \begin{pmatrix} \mathbf{q}_s \\ \mathbf{q}_f \end{pmatrix} + \begin{pmatrix} \mathbf{R}_s(\mathbf{q}_s, \mathbf{q}_f) \\ \mathbf{R}_f(\mathbf{q}_f, \mathbf{q}_s) \end{pmatrix} = 0 \quad (3.47)$$

with the solid $\mathbf{q}_s = (h, \theta, u_h, u_\theta)^\top$ and fluid $\mathbf{q}_f = (\mathbf{u}, p, \lambda)^\top$ unknowns, and where

$$\mathbf{M}_s = \begin{pmatrix} 1 & 0 & 0 & 0 \\ 0 & 1 & 0 & 0 \\ 0 & 0 & 1 & 0 \\ 0 & 0 & 0 & 1 \end{pmatrix} \quad \mathbf{M}_f(\mathbf{q}_s) = \begin{pmatrix} \mathbf{M}_u(\mathbf{q}_s) & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & \mathbf{0} & \mathbf{0} \end{pmatrix}$$

with $\mathbf{M}_u(\mathbf{q}_s)$ is similar to a mass matrix on the velocity space, but incorporates a dependence in \mathbf{q}_s through the Jacobian of the extension deformation:

$$\mathbf{M}_u(\mathbf{q}_s) = \int_{\Omega_f} J(\xi_e(\mathbf{q}_s)) \mathbf{u} \cdot \ddot{\mathbf{u}}$$

The discrete stationary residuals write

$$\mathbf{R}_s = \begin{pmatrix} -u_h \\ -u_\theta \\ 2\zeta_h \frac{\Omega}{U_c^*} u_h + \frac{\Omega^2}{U_c^{*2}} h - \frac{1}{\tilde{m}} \mathcal{C}_L(\lambda) \\ 2\zeta_p \left(\frac{1}{U^*}\right) u_\theta + \left(\frac{1}{U^*}\right)^2 \theta - \frac{1}{\tilde{m}r_\theta^2} \mathcal{C}_M(\theta, \lambda) \end{pmatrix}, \quad \mathbf{R}_f = \begin{pmatrix} \mathbf{R}_u \\ \mathbf{R}_p \\ \mathbf{R}_\lambda \end{pmatrix}$$

where the three components of the \mathbf{R}_f are

$$\begin{aligned} \mathbf{R}_u(\mathbf{u}, p, \lambda, \mathbf{q}_s) &:= \int_{\Omega_f} \{ \nabla \mathbf{u} \Phi(\zeta_e(\mathbf{q}_s)) (\mathbf{u} - \mathbf{w}(\mathbf{q}_s)) \} \cdot \check{\mathbf{u}} \\ &\quad + \int_{\Omega_f} \boldsymbol{\Sigma}_f(\mathbf{u}, p, \zeta_e(\mathbf{q}_s)) : \nabla \check{\mathbf{u}} + \int_{\Gamma_{fs}} \boldsymbol{\lambda} \cdot \check{\mathbf{u}} \\ \mathbf{R}_p(\mathbf{u}, \mathbf{q}_s) &:= - \int_{\Omega_f} \{ \Phi(\zeta_e(\mathbf{q}_s))^T : \nabla \mathbf{u} \} \check{p} \\ \mathbf{R}_\lambda(\mathbf{u}, \mathbf{q}_s) &:= \int_{\Gamma_{fs}} (\mathbf{u} - \mathbf{w}(\mathbf{q}_s)) \cdot \check{\boldsymbol{\lambda}} \end{aligned}$$

3.A.1.2 Linearized equations

Linearizing the “reference configuration ALE” formalism is a cumbersome process, previously considered in the thesis of [J.-L. Pfister \(2019\)](#). Here, for sake of brevity, we limit ourselves to giving the final result. Consider a base (time-dependent) state, noted $\mathbf{q}_{s0} = (h_0, \theta_0, u_{h0}, u_{\theta0})^T$ and $\mathbf{q}_{f0} = (\mathbf{u}_0, p_0, \lambda_0)^T$. The linearization of eq. (3.47) around $(\mathbf{q}_{s0}, \mathbf{q}_{f0})$ yields:

$$\begin{pmatrix} \mathbf{M}_s & 0 \\ 0 & \mathbf{M}_f(\mathbf{q}_{s0}) \end{pmatrix} \frac{\partial}{\partial t} \begin{pmatrix} \mathbf{q}_s \\ \mathbf{q}_f \end{pmatrix} + \begin{pmatrix} \mathbf{J}_{ss} & \mathbf{J}_{sf} \\ \mathbf{J}_{fs} & \mathbf{J}_{ff} \end{pmatrix} \begin{pmatrix} \mathbf{q}_s \\ \mathbf{q}_f \end{pmatrix} = 0 \quad (3.48)$$

where the linearized solid problem is defined by

$$\begin{aligned} \mathbf{J}_{ss} &= \begin{pmatrix} 0 & 0 & -1 & 0 \\ 0 & 0 & 0 & -1 \\ \Omega^2 U^{*-2} & 0 & 2\zeta_h \Omega U^{*-1} & 0 \\ 0 & U^{*-2} - \frac{1}{\tilde{m}r_\theta^2} \frac{\partial \mathcal{C}_M}{\partial \theta} & 0 & 2\zeta_p U^{*-1} \end{pmatrix} \\ \mathbf{J}_{sf} &= \begin{pmatrix} 0 & 0 & 0 \\ 0 & 0 & 0 \\ 0 & 0 & -\frac{1}{\tilde{m}} \frac{\partial \mathcal{C}_L}{\partial \lambda} \\ 0 & 0 & -\frac{1}{\tilde{m}r_\theta^2} \frac{\partial \mathcal{C}_M}{\partial \lambda} \end{pmatrix} \end{aligned}$$

and the linearized fluid problem by

$$\mathbf{J}_{ff} = \begin{pmatrix} \mathbf{A} & \mathbf{B}^T & \mathbf{I}_\lambda^T \\ \mathbf{B} & \mathbf{0} & \mathbf{0} \\ \mathbf{I}_\lambda & \mathbf{0} & \mathbf{0} \end{pmatrix} \quad \mathbf{J}_{fs} = \begin{pmatrix} \mathbf{D}_{uh} & \mathbf{D}_{u\theta} & \mathbf{D}_{uu_h} & \mathbf{D}_{uu_\theta} \\ \mathbf{D}_{ph} & \mathbf{D}_{p\theta} & \mathbf{0} & \mathbf{0} \\ \mathbf{0} & -\mathbf{W}_\theta & -\mathbf{W}_{u_h} & -\mathbf{W}_{u_\theta} \end{pmatrix}$$

Matrices \mathbf{A} and \mathbf{B} are the finite element discretizations of the linearized Navier–Stokes convection–diffusion and divergence operators:

$$\begin{aligned}\mathbf{A} &:= \int_{\Omega_f} \left\{ \nabla \mathbf{u}' \Phi(\xi_e(\mathbf{q}_{s0})) (\mathbf{u}_0 - \mathbf{w}(\mathbf{q}_{s0})) + \nabla \mathbf{u}_0 \Phi(\xi_e(\mathbf{q}_{s0})) \mathbf{u}' \right\} \cdot \check{\mathbf{u}} \\ &\quad + \int_{\Omega_f} \left\{ \mathcal{R}e^{-1} \mathbf{G}(\mathbf{u}', \xi_e(\mathbf{q}_{s0})) \Phi(\xi_e(\mathbf{q}_{s0}))^T \right\} : \nabla \check{\mathbf{u}} \\ \mathbf{B} &:= - \int_{\Omega_f} \left\{ \Phi(\xi_e(\mathbf{q}_{s0}))^T : \nabla \mathbf{u}' \right\} \check{p}\end{aligned}$$

whereas \mathbf{I}_λ is a mass matrix defined on the interface Lagrange multiplier space. Matrices \mathbf{W}_θ , \mathbf{W}_{u_h} and \mathbf{W}_{u_θ} transmit the solid velocity to the fluid at the interface,

$$\mathbf{W}_\theta := \int_{\Gamma_{fs}} \frac{\partial \mathbf{w}}{\partial \theta} \theta' \cdot \check{\lambda} \quad \mathbf{W}_{u_h} := \int_{\Gamma_{fs}} \frac{\partial \mathbf{w}}{\partial u_h} u'_h \cdot \check{\lambda} \quad \mathbf{W}_{u_\theta} := \int_{\Gamma_{fs}} \frac{\partial \mathbf{w}}{\partial u_\theta} u'_\theta \cdot \check{\lambda}$$

whereas the *shape derivatives* \mathbf{D}_{u_h} , \mathbf{D}_{u_θ} , $\mathbf{D}_{u u_h}$, $\mathbf{D}_{u u_\theta}$ represent the effect of a solid movement on the Navier–Stokes momentum equation:

$$\begin{aligned}\mathbf{D}_{u_h} &:= \int_{\Omega_f} \left\{ \nabla \mathbf{u}_0 \Phi' \left(\frac{\partial \xi_e}{\partial h} h' \right) (\mathbf{u}_0 - \mathbf{w}(\mathbf{q}_{s0})) \right\} \cdot \check{\mathbf{u}} \\ &\quad + \int_{\Omega_f} \Sigma_f' \left(\frac{\partial \xi_e}{\partial h} h' \right) : \nabla \check{\mathbf{u}} \\ &\quad + \int_{\Omega_f} J' \left(\frac{\partial \xi_e}{\partial h} h' \right) \frac{\partial \mathbf{u}_0}{\partial t} \cdot \check{\mathbf{u}} \\ \mathbf{D}_{u_\theta} &:= \int_{\Omega_f} \left\{ \nabla \mathbf{u}_0 \Phi' \left(\frac{\partial \xi_e}{\partial \theta} \theta' \right) (\mathbf{u}_0 - \mathbf{w}(\mathbf{q}_{s0})) - \nabla \mathbf{u}_0 \Phi \left(\xi_e(\mathbf{q}_{s0}) \right) \frac{\partial \mathbf{w}}{\partial \theta} \theta' \right\} \cdot \check{\mathbf{u}} \\ &\quad + \int_{\Omega_f} \Sigma_f' \left(\frac{\partial \xi_e}{\partial \theta} \theta' \right) : \nabla \check{\mathbf{u}} \\ &\quad + \int_{\Omega_f} J' \left(\frac{\partial \xi_e}{\partial \theta} \theta' \right) \frac{\partial \mathbf{u}_0}{\partial t} \cdot \check{\mathbf{u}} \\ \mathbf{D}_{u u_h} &:= \int_{\Omega_f} \left\{ -\nabla \mathbf{u}_0 \Phi \left(\xi_e(\mathbf{q}_{s0}) \right) \frac{\partial \mathbf{w}}{\partial u_h} u'_h \right\} \cdot \check{\mathbf{u}} \\ \mathbf{D}_{u u_\theta} &:= \int_{\Omega_f} \left\{ -\nabla \mathbf{u}_0 \Phi \left(\xi_e(\mathbf{q}_{s0}) \right) \frac{\partial \mathbf{w}}{\partial u_\theta} u'_\theta \right\} \cdot \check{\mathbf{u}}\end{aligned}$$

and \mathbf{D}_{ph} , $\mathbf{D}_{p\theta}$ represent the effect of a solid movement on the Navier–Stokes continuity equation:

$$\begin{aligned}\mathbf{D}_{ph} &:= - \int_{\Omega_f} \left\{ \Phi' \left(\frac{\partial \xi_e}{\partial h} h' \right)^T : \nabla \mathbf{u}_0 \right\} \check{p} \\ \mathbf{D}_{p\theta} &:= - \int_{\Omega_f} \left\{ \Phi' \left(\frac{\partial \xi_e}{\partial \theta} \theta' \right)^T : \nabla \mathbf{u}_0 \right\} \check{p}\end{aligned}$$

Remark. The last term in \mathbf{D}_{u_h} and \mathbf{D}_{u_θ} comes from the linearization of the term $J(\xi_e(\mathbf{q}_s)) \frac{\partial \mathbf{u}}{\partial t}$ in the fluid momentum equation, with respect to solid variables. Because it is proportional to $\frac{\partial \mathbf{u}_0}{\partial t}$, this term vanishes if \mathbf{u}_0 is a steady state, like in global stability analysis. Here, \mathbf{u}_0 is a periodic orbit, thus the term should be kept for exact linearization.

The different linearized operators that appear in the shape derivatives are detailed below. First, the linearization of $\Sigma_f = (-p_0\mathbf{I} + \mathcal{R}e^{-1}\mathbf{G})\Phi^T$ with respect to ξ_e :

$$\Sigma_f'(\xi_e') = \mathcal{R}e^{-1}\mathbf{G}'(\xi_e')\Phi(\xi_e(\mathbf{q}_{s0}))^T + \left(-p_0\mathbf{I} + \mathcal{R}e^{-1}\mathbf{G}(\mathbf{u}_0, \xi_e(\mathbf{q}_{s0}))\right)\Phi'(\xi_e')^T$$

Second, the linearization of $\mathbf{G} = J^{-1}(\nabla\mathbf{u}\Phi + \Phi^T\nabla\mathbf{u}^T)$ with respect to ξ_e :

$$\mathbf{G}'(\xi_e') = \frac{1}{J(\xi_{e0})} \left(\nabla\mathbf{u}_0\Phi'(\xi_e') + \Phi'(\xi_e')^T\nabla\mathbf{u}_0^T \right) - \frac{J'(\xi_e')}{J^2(\xi_{e0})} \left(\nabla\mathbf{u}_0\Phi(\xi_{e0}) + \Phi^T(\xi_{e0})\nabla\mathbf{u}_0^T \right)$$

Third, the linearization of $J = \det(\mathbf{F})$ with respect to ξ_e :

$$J'(\xi_e') = \Phi(\xi_{e0})^T : \nabla\xi_e'$$

Finally, for two-dimensional flows, the deformation operator $\Phi = J\mathbf{F}^{-1}$ is linear:

$$\Phi(\xi_e) = \begin{pmatrix} 1 + \frac{\partial\xi_y}{\partial y} & -\frac{\partial\xi_x}{\partial y} \\ -\frac{\partial\xi_y}{\partial x} & 1 + \frac{\partial\xi_x}{\partial x} \end{pmatrix} \quad \text{and} \quad \Phi'(\xi_e') = \Phi(\xi_e')$$

3.A.2 Forced problem: solid with imposed heaving motion

The forced fluid problem for the numerical example of section 3.4.1 is detailed in this section. The NACA0012 profile is forced along the heaving direction, as described by the following solid vector:

$$\mathbf{q}_s^{\text{forc}}(t) = \begin{pmatrix} h_1 \sin(\omega t) \\ \omega h_1 \cos(\omega t) \\ 0 \\ 0 \end{pmatrix}$$

3.A.2.1 Nonlinear equations

Only the fluid equations are kept from eq. (3.47):

$$\mathbf{M}_f(\mathbf{q}_s^{\text{forc}}(t)) \frac{\partial\mathbf{q}_f}{\partial t} + \mathbf{R}_f(\mathbf{q}_f, \mathbf{q}_s^{\text{forc}}(t)) = \mathbf{0} \quad (3.49)$$

with $\mathbf{q}_f = (\mathbf{u}, p, \lambda)^T$ gathering the fluid unknowns.

Remark. Due to the highly nonlinear form of the ALE equations, most contributions involving to the solid forcing $\mathbf{q}_s^{\text{forc}}(t)$ cannot in practice be transferred to the right-hand side, as an *external* forcing. Instead, they are kept in operator \mathbf{R}_f which is thus a nonlinear operator in \mathbf{q}_f , with time-dependent periodic coefficients.

3.A.2.2 Linearized equations

The linearized problem around the fluid state \mathbf{q}_{f0} reduces to

$$\mathbf{M}_f(\mathbf{q}_s^{\text{forc}}(t)) \frac{\partial\mathbf{q}_f}{\partial t} + \mathbf{J}_{ff}(\mathbf{q}_{f0}, \mathbf{q}_s^{\text{forc}}(t))\mathbf{q}_f = \mathbf{0} \quad (3.50)$$

II

FLUTTER BIFURCATION ANALYSIS IN VISCOUS FLOWS

4

LINEAR STABILITY OF A TYPICAL SECTION IN VISCOUS INCOMPRESSIBLE FLOWS

We revisit, for low-Reynolds laminar incompressible flows, the linear stability analysis of a typical aeroelastic section consisting in a thin plate mounted on flexion and torsion springs. A global stability framework is adopted for the coupled fluid-structure system. Depending on the reduced velocity U^* and steady angle of attack, we show that the global stability framework captures four types of flow-induced instabilities — vortex-induced vibrations at low U^* and coupled-mode/single-mode flutter and divergence at high U^* — that are often studied with separate flow models. Neutral curves for these instabilities are presented in the parametric plane composed of the solid-to-fluid mass ratio and the reduced velocity. At a particular, low, value of the mass ratio, noted \tilde{m}^{II} , a codimension-two point is found where both flutter and divergence occur at the same critical reduced velocity. For mass ratios higher than \tilde{m}^{II} , flutter occurs prior to divergence whereas divergence precedes flutter (if it exists at all) for mass ratios below \tilde{m}^{II} . In addition, for very high U^* , a restabilization of the flutter mode is observed. The effect of the Reynolds number is then investigated on the range $10 \leq Re \leq 10^4$, showing in particular that this high- U^* restabilization of flutter is caused by viscosity. The global stability results are finally compared to more classical fluid modelizations based on quasi-steady models or on the Theodorsen theory, both calibrated using the static aerodynamic coefficients. If all approaches converge for the low-frequency flutter that occurs at large mass ratios ($\tilde{m} > 10^3$), they significantly differ for the higher-frequency flutter occurring at low mass ratios ($\tilde{m} \sim 10$), where quasi-steady models are overconservative while the (recalibrated) Theodorsen model is unconservative. For some low value of the mass ratio, the Theodorsen model predicts a neutral curve asymptotically tending towards infinity. With the full Navier–Stokes approach, no such asymptote is observed as the neutral curve “folds back” so as to allow the high- U^* restabilization of the flutter mode. Overall, our findings show that accurate predictions of the flutter thresholds, in particular for low mass ratios, require taking into account the viscous effects.

4.1 Introduction

The assessment of flutter linear stability is often performed using potential flow models, like the Theodorsen model [Theodorsen 1935] or the Doublet Lattice Method (DLM) [Albano et al. 1969]. These elegant and low-computational-cost methods have encountered a great success for the computation of aeroelastic stability of high-speed aircraft in industrial practice [Garrigues 2018]. However, since the middle of the twentieth century, it has been pointed out that classical aeroelastic (potential flow) methods are not consistent with experimental results for low mass ratio applications, typical of hydrofoils [Woolston et al. 1951; Abramson 1969; Besch et al. 1971; Chae et al. 2013]. For example, for a two-dimensional foil section mounted on heaving and pitching springs (a “typical aeroelastic section”), it is well-known that the Theodorsen model predicts a critical flutter velocity asymptotically tending towards infinity, for some low value of the mass ratio (e.g. [E. H. Dowell et al. 1989, fig. 3.27] or [Hodges et al. 2011, fig. 5.14]). This often yields an overestimated (unconservative) prediction of the experimental flutter thresholds for low mass ratios, in the region of the asymptote [Chae et al. 2013]. No uniform theory that would accurately capture fluid-structure instabilities at these low mass ratio seems to be generally accepted in the literature. As a consequence, the most widely used method to this day consists in time-marching the coupled nonlinear fluid-structure equations [Young et al. 2012], using a Reynolds Average Navier–Stokes approach for turbulence modeling. In several of these studies, viscosity is cited as the main source of difficulty, despite the high-enough Reynolds numbers used in hydrofoil applications, $Re \sim 10^6$. If viscous effects are already problematic for these high-Reynolds configurations, their impact is expected to be even stronger for modern applications, like Micro-Aerial Vehicles [Mueller et al. 2003; Shyy et al. 2010], that combine (very) low-Reynolds numbers ($10 \leq Re \leq 10^5$), low mass ratios and very flexible wings.

The effect of Reynolds number on the fluid forces exerted on a forced heaving and pitching thin streamlined plate was investigated by Bruno et al. (2008), using an indicial function approach. In the large range they studied, from $Re = 10$ to $Re = 10^5$, the flutter derivatives were shown to significantly depend on the Reynolds number. As a consequence, large discrepancies with respect to Theodorsen’s potential theory were reported at low Re , leading the authors to recommend not to use the latter for such flows. A possible modelling approach for fluid forces at low-Reynolds flows was proposed by Brunton et al. (2013). The authors considered a generalized Theodorsen’s model by identifying empirical versions of the analytical potential flow parameters in the original model, like the added mass, damping and stiffness coefficients or the Theodorsen function. This generalization allowed the authors to significantly improve the lift prediction on a pitch oscillating flat plate at $Re = 100$, for a wide range of reduced frequencies. However, in these studies, the impact of the viscous flow model on the aeroelastic stability was not investigated. Recently, the aeroelastic stability of a spring-mounted pitching NACA0012 airfoil was investigated at transitional Reynolds numbers ($Re \sim 50000$) in [Negi 2019, Paper 3] in an effort to investigate the linear regime of the so-called “laminar separation flutter” originally reported in the experiments of Poirel et al. (2008).

For systems other than a typical airfoil section, the fluid-structure linear stability has been successfully studied in several (very) low-Reynolds viscous flows. The most classical of them might be the destabilization of a spring-mounted circular cylinder due to vortex-induced vibrations [Cossu et al. 2000]. In addition, Cisonni et al. (2017) investigated the flutter of an elastic filament in viscous channel flow.

Goza et al. (2018) studied the flow-induced oscillations of an inverted flag. The development of oscillations in the vertical path of freely falling or rising rigid objects has been considered by Assemat et al. (2012), Tchoufag et al. (2014a), and Tchoufag et al. (2014b). A flexible splitter plate attached to a rigid circular cylinder has been studied by J. L. Pfister et al. (2020). In these studies, the fluid-structure stability is assessed with a so-called global stability approach [Sipp et al. 2007] where the leading (*i.e.* rightmost) eigenvalues of the linearized fluid-structure operator, including the incompressible Navier–Stokes equations, are directly computed. A more detailed presentation of this approach for fluid-structure instabilities can be found in [Assemat et al. 2012] in the case of rigid structures and in [J.-L. Pfister et al. 2019; Negi et al. 2019] for various approaches suitable for generic elastic solids.

In this study, we propose to revisit the linear stability of a “typical aeroelastic section” consisting in a rigid, heaving and pitching spring-mounted thin plate, immersed in two-dimensional laminar incompressible flows, using a global stability approach. A second goal of this work is to assess how differently a full Navier–Stokes fluid modeling predicts flutter, in comparison to simpler, less computationally intensive models. The chapter is organized as follows. In section 4.2, we very briefly recall the models and methods used to investigate the linear stability of a spring-mounted plate. In section 4.3, we present the different types of instabilities that appear depending on the reduced velocity and steady angle of attack. Then, in section 4.4, we parametrically explore the effect of the solid-to-fluid mass ratio and Reynolds number for the plate at zero steady angle of attack. The chapter ends in section 4.5 with a comparison of the present viscous analysis with simplified models such as quasi-steady approaches or the classical Theodorsen model.

4.2 Problem settings and methods

We consider in this chapter the system presented in section 1.1 which is composed of a rigid plate mounted on heaving and pitching springs and immersed in a two-dimensional incompressible viscous flow. Eight non-dimensional parameters, summarized in table 1.1, define the system. Five of them — the heaving-to-pitching frequency ratio Ω , the structural dampings ζ_h, ζ_p , the radius of gyration r_θ and the position of the elastic axis x_θ — are kept fixed to the values already specified in table 1.1. In particular, the elastic axis is positioned at the mid-chord ($x_\theta = 0$). The three other parameters, that are the solid-to-fluid mass ratio $\tilde{m} = m / (1/2\rho_f c^2)$, the Reynolds number $\mathcal{R}e = \rho_f U_\infty c / \mu$ and the reduced velocity $U^* = U_\infty / (c \sqrt{K_\theta / I_{ea}})$, vary in the present chapter.

The dynamics of the spring-mounted plate is governed by two (damped) linear oscillator equations eq. (1.3), whereas the flow verifies the incompressible Navier–Stokes equations. This system of coupled equations, described in detail in chapter 1, is formally written here as the following first-order in time evolution equation

$$\mathcal{M} \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}) = 0 \quad (4.1)$$

where the variable $\mathbf{q} = (h, \theta, u_h, u_\theta, \mathbf{U}, p, \mathbf{\Lambda})^T$ gathers all variables necessary to describe the fluid-solid interaction. These first four scalar variables allow describing the dynamics of the rigid plate. They are the heaving h and pitching θ displacements, as well as the corresponding velocities u_h and u_θ . The incompressible flow is described with the velocity \mathbf{U} and pressure p fields. Finally, the variable $\mathbf{\Lambda}$ defined at the fluid-solid interface represents the local stress exerted by the fluid onto the

solid. The exact definition of the linear operator \mathcal{M} in front of the time-derivative and of the nonlinear residual vector \mathcal{R} are given in eq. (1.30). In this residual, the first four components corresponds to the two coupled damped harmonic oscillators, written as a first-order problem in time. The fifth and sixth component correspond to the Navier–Stokes momentum and mass conservation equation, written using the “absolute velocity - rotating axis” formalism described in section 1.1.2.1. The last equation corresponds to the equality of fluid and solid velocities at the fluid-solid interface. Again, we refer to chapter 1 for more details.

Linear stability analysis To explore the occurrence of flutter, we investigate the linear stability of steady solutions $\mathbf{q}_0(\mathbf{X})$ of eq. (4.1), *i.e.* we search the conditions in which small amplitude time-dependent perturbations, noted $\epsilon \mathbf{q}_1(\mathbf{X}, t)$, are able to develop around \mathbf{q}_0 in a self-sustained manner. To that end, the solution is decomposed as the sum of the steady solution and the small perturbation:

$$\mathbf{q}(\mathbf{X}, t) = \mathbf{q}_0(\mathbf{X}) + \epsilon \mathbf{q}_1(\mathbf{X}, t) \quad \epsilon \ll 1 \quad (4.2)$$

In addition, the perturbation is further decomposed in the form of global modes [Sipp et al. 2010]:

$$\mathbf{q}_1(\mathbf{X}, t) = \hat{\mathbf{q}}(\mathbf{X}) e^{\sigma t} \quad \sigma = \lambda + i\omega$$

Plugging the above in eq. (4.1) yields the steady nonlinear problem eq. (2.5) for \mathbf{q}_0 and the eigenproblem eq. (2.9), that are recalled for convenience:

$$\mathcal{R}(\mathbf{q}_0, U_c^*) = 0 \quad (4.3a)$$

$$\sigma \mathcal{M} \hat{\mathbf{q}} + \mathcal{J}(\mathbf{q}_0, U_c^*) \hat{\mathbf{q}} = 0 \quad (4.3b)$$

where σ is complex eigenvalue and $\hat{\mathbf{q}}$ the corresponding complex-valued eigenvector. The long term stability of the steady solution is assessed by scrutinizing the eigenvalue spectrum of eq. (4.3b). If all eigenvalues have negative real parts ($\lambda < 0$) the system is stable in the sense that any infinitesimal perturbation will eventually decay to zero. If at least one eigenvalue has positive real part ($\lambda > 0$) then the system is unstable. The case $\lambda = 0$ describes a situation where a perturbation will neither be amplified nor damped, and is said neutrally stable.

A Newton method for neutral curves computation From a physical point of view, sets of parameters for which an eigenmode is neutrally stable are of particular interest since they define boundaries, called *neutral curves*, between stable and unstable regions of the parametric space. For the present study, we are typically interested (*e.g.* fig. 4.4.1) in computing neutral curves in the (\tilde{m}, U^*) parametric space for different modes of interest. In other words, we search curves of the form $U_c^* = f(\tilde{m})$, where U_c^* is a critical value of U^* , *i.e.* a value of U^* for which the mode of interest is neutrally stable. A naive way to do so is, for each \tilde{m} , to vary U^* in a wide enough range, compute the spectrum of eq. (4.3b) and find the particular value(s) of U^* for which a mode switches from $\lambda < 0$ to $\lambda > 0$. This naive approach is particularly inefficient as it requires to compute the full spectrum on a whole interval of U^* in order to extract in the end only one critical mode $\hat{\mathbf{q}}_c$, the corresponding critical frequency ω_c and U_c^* . In order to bypass those shortcomings, we solve directly for the

critical triplet $(\hat{\mathbf{q}}_c, \omega_c, U_c^*)$ that verifies:

$$(\mathbf{i}\omega_c \mathcal{M} + \mathcal{J}(\mathbf{q}_0, U_c^*)) \hat{\mathbf{q}}_c = 0 \quad (4.4a)$$

$$\hat{\mathbf{q}}_c^H \mathcal{M} \hat{\mathbf{q}}_c - 1 = 0 \quad (4.4b)$$

$$\mathcal{C}(\hat{\mathbf{q}}_c) = 0 \quad (4.4c)$$

The first equation ensures that the mode is a critical eigenmode of eq. (4.3b), whereas the second and third equations are arbitrary scalar conditions that are required to obtain a well-posed problem. In practice, they respectively fix the norm and phase of the eigenvector. For dynamic instabilities (*i.e.* $\omega_c \neq 0$: flutter or VIV in our case), we impose as a phase condition that the heaving degree of freedom has a zero real part: $\mathcal{C}(\hat{\mathbf{q}}_c) = \Re(\hat{h}_c)$. For static instabilities (*i.e.* $\omega_c = 0$, divergence in our case), a phase condition is irrelevant, and we can simply degenerate it to $\omega_c = 0$. The obtained system eq. (4.4) is nonlinear due to the dependence of the Jacobian operator in U_c^* and to the presence of ω_c in front of the mass matrix term, in the first equation. We solve it with Newton iterations. Note in addition that we had to wrap the resolution of eq. (4.4) in a pseudo arc-length algorithm, in order to properly handle the occurrence of turning points in the different computed neutral curves (see fig. 4.4.1).

We end this section by mentioning that all the problems mentioned above are discretized in space using the SUPG-stabilized finite element method presented in section 1.2.2.

4.3 Different types of flow-induced vibrations

In this section, we expose the linear stability results obtained by solving eq. (4.3). First, the leading eigenmodes are described in section 4.3.1 for a particular set of non-dimensional parameters. Then, by varying the reduced velocity in section 4.3.2 and the steady angle of attack in section 4.3.3, four types of instabilities are observed.

4.3.1 Leading eigenmodes

Consider the flow around the plate at $Re = 2900$. The steady solution \mathbf{q}_0 , at which the fluid-structure Jacobian is evaluated, is simply the steady flow around the plate fixed at $h = \theta = 0$. Such a flow, shown in fig. 4.3.1a, is mostly attached to the plate, even at the leading edge, due to the low Reynolds number. At the blunt trailing edge, a small recirculation area, characterized by backwards flow, is observed inside the region marked by the thin black line.

The rightmost eigenvalues¹ of eq. (4.3b) are shown in fig. 4.3.1b for particular values of the mass ratio $\tilde{m} = 1000$ and reduced velocity $U^* = 4.7$. First, one unstable mode, marked by a red dot, is observed at a low frequency $\omega = 0.17$. This frequency is of the order of $\mathcal{O}(1/U^*)$, which corresponds to the natural (non-dimensional) solid frequencies, $\omega_{0\theta} = 1/U^*$ (pitching) and $\omega_{0h} = \Omega/U^*$ (heaving). The structure of the corresponding eigenvector is shown in fig. 4.3.2a. The pressure field of the eigenmode, represented by the blue-white colormap, is mainly concentrated near the leading edge. As can be seen from the displacement field of the plate (orange arrows), this unstable mode couples heaving and pitching motions. By comparing the minimums of the heaving and pitching signals presented in the right part of the figure, we observed that their phase shift is such that pitching (dashed line) precedes

¹*i.e.* the eigenvalues with largest real part λ

heaving (solid line). This feature is a typical signature of the coupled-mode flutter instability, as mentioned in the introduction to this manuscript (see the discussion of fig. 2(b)). This mode thus corresponds to a classical *flutter mode*.

Close to the flutter mode, another low-frequency mode (black dot, $\omega = 0.15$), stable this time, stands out in fig. 4.3.1b due to its relatively small damping in comparison with the rest of modes marked by the black “plus” signs. Again, the corresponding eigenmode is shown in fig. 4.3.2b and presents very similar features to the flutter mode with a coupling between heaving and pitching motions and a pressure field concentrated near the leading edge. However, contrary to the flutter mode, the phase shift is now such that pitching (dashed) lags behind heaving (solid). Due to those properties “similar and opposite” to the flutter mode, this mode is an *anti-flutter mode* (see also the discussion of fig. 2(c) in the general introduction to this manuscript).

All other modes in fig. 4.3.1b are stable, most of them being strongly damped (notice the different scales on the left and right of the “zigzag” sign in the figure). They are qualified as fluid modes here, in the sense that they are already observed when the solid is fixed, *i.e.* they are present in the spectrum of the fluid-only Jacobian matrix. Among these modes, the one marked by the blue disk on the figure is remarkable, due to its nearly zero growth rate and high frequency $\omega_{\text{wake}} = 11.2$. The eigenvector structure presented in fig. 4.3.2c shows completely different features from the flutter and anti-flutter modes previously described. First, the solid contribution to the mode is negligible, meaning that the flow perturbation generated by this mode does not produce significant efforts on the plate. This is indeed observed in the pressure field representation where the fluctuations develop only in the wake of the plate in the form of a succession of downstream-convected, small-wavelength patterns. This feature is typical of the eigenmode associated with the instability of the recirculation region in the wake of a bluff body (*e.g.* [Jackson 1987; Sipp et al. 2007]). In the following, it is thus referred to as the *wake mode*. Finally, we mention that the critical Reynolds number for the wake mode, in the case of a fixed plate, is about $\mathcal{R}e_{c,\text{wake}} \simeq 2925$ (or equivalently $\mathcal{R}e_{c,\text{wake}}^H \simeq 145$ if the Reynolds is based on the plate’s thickness, H) and the thickness-based Strouhal number is $St_e \simeq 0.08$.

4.3.2 Varying the reduced velocity: from VIV to coupled-mode flutter and divergence

When varying the reduced velocity U^* , the three leading modes marked by disks in fig. 4.3.1b move in the complex plane, giving rise to different types of instabilities. In fig. 4.3.3, we mark by solid lines the paths of those three leading modes as U^* is varied between 0.05 and 8. For comparison purposes, the position of the modes at $U^* = 4.7$ (case of fig. 4.3.1b) is recalled with the disks symbols. The direction of increasing U^* is materialized on each path by an arrow. Finally, the path of the uncoupled (solid-only) heaving and pitching modes are shown with thin dashed lines.

First, it is observed that the path followed by the red eigenvalue is mostly close to the real (vertical) axis, which is also the locus for the uncoupled heaving mode. For that reason, we refer to the red mode as the heaving mode. For most values of U^* , the heaving mode is slightly damped but gets unstable at two different locations in the complex plane. On the one hand, at high-frequency $\omega \simeq 11$, the heaving mode interacts with the wake mode (blue) and destabilizes, as emphasized in the zoomed-in view proposed in the upper-right part of the figure. The interaction is even more obvious when visualizing the eigenmode structure (fig. 4.3.4a): the mode’s pressure

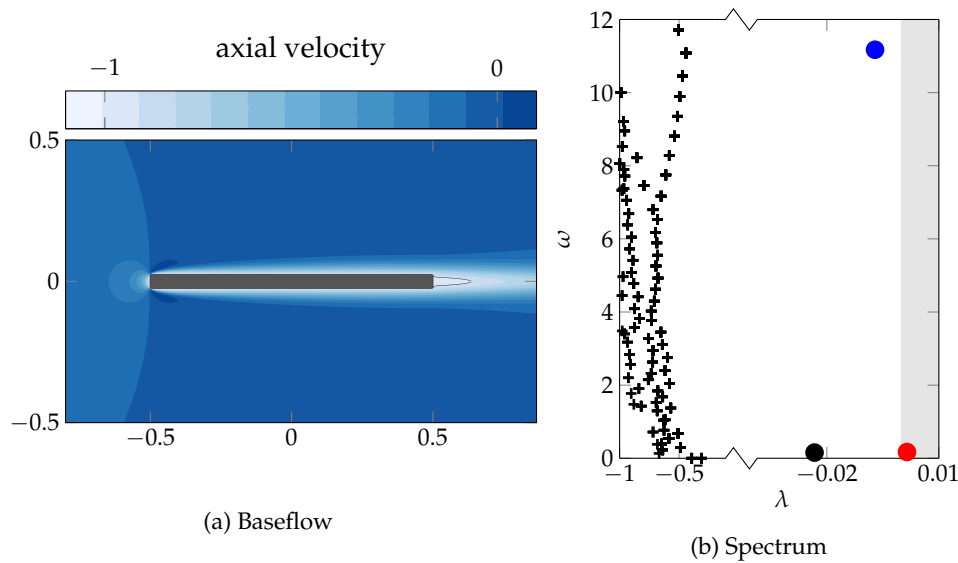


FIGURE 4.3.1: Linear stability analysis of the steady baseflow illustrated in 4.3.1a. The spectrum of the fluid-structure operator \mathcal{J} at ($U^* = 4.7$, $Re = 2900$, $\tilde{m} = 1000$) is shown in 4.3.1b. The leading mode, marked with a red disk, is slightly unstable and is responsible for the flutter instability. It is depicted in fig. 4.3.2a. All modes marked by crosses correspond to the fluid modes that are present without any fluid structure coupling. In particular, the high-frequency wake mode (blue disk) is presented in fig. 4.3.2c.

field has the typical features of the stable wake mode (*cf* fig. 4.3.2c), whereas the solid movement is largely dominated by heaving motion. Note that the destabilization of a fluid-structure system due to the interaction of a solid mode and a wake mode has been extensively documented in the circular cylinder flow configuration, using both global stability analysis [Cossu et al. 2000; Meliga et al. 2011; Navrose et al. 2016] and nonlinear time-accurate simulations [Mittal et al. 2005; Singh et al. 2005; Navrose et al. 2016] in the past twenty years. The second destabilization observed occurs in a low-frequency zone around $\omega \simeq 0.1$ and results in the flutter mode already described in section 4.3.1.

Focusing now on the black mode, it stays close to the solid-only pitching mode locus at low U^* , before it suddenly deviates to reach zero frequency and finally destabilizes, as show in the zoomed-in view of the lower-right part of fig. 4.3.3. The corresponding eigenmode is presented in fig. 4.3.4b. The destabilization of the pitching mode by a static instability ($\omega = 0$) is an example of the classical divergence instability [Bisplinghoff et al. 1955, Chapter 8]. Note that here, the divergence mode includes both heaving and pitching displacements. However, it is known that divergence is possible with only the pitching degree of freedom, for example in the thin airfoil theory limit [E. H. Dowell et al. 1989, §2.1].

For sake of completeness, we provide in fig. 4.3.5 a complementary (and more practical) representation for the flutter and divergence instabilities, by following the red and black eigenvalues as a functions of U^* . The growth rate and frequency are both normalized by the pitching natural frequency $\omega_{0\theta} = 1/U^*$, in order to avoid large variations in frequency. It is clear in fig. 4.3.5 that for $U^* \rightarrow 0$, the red and black

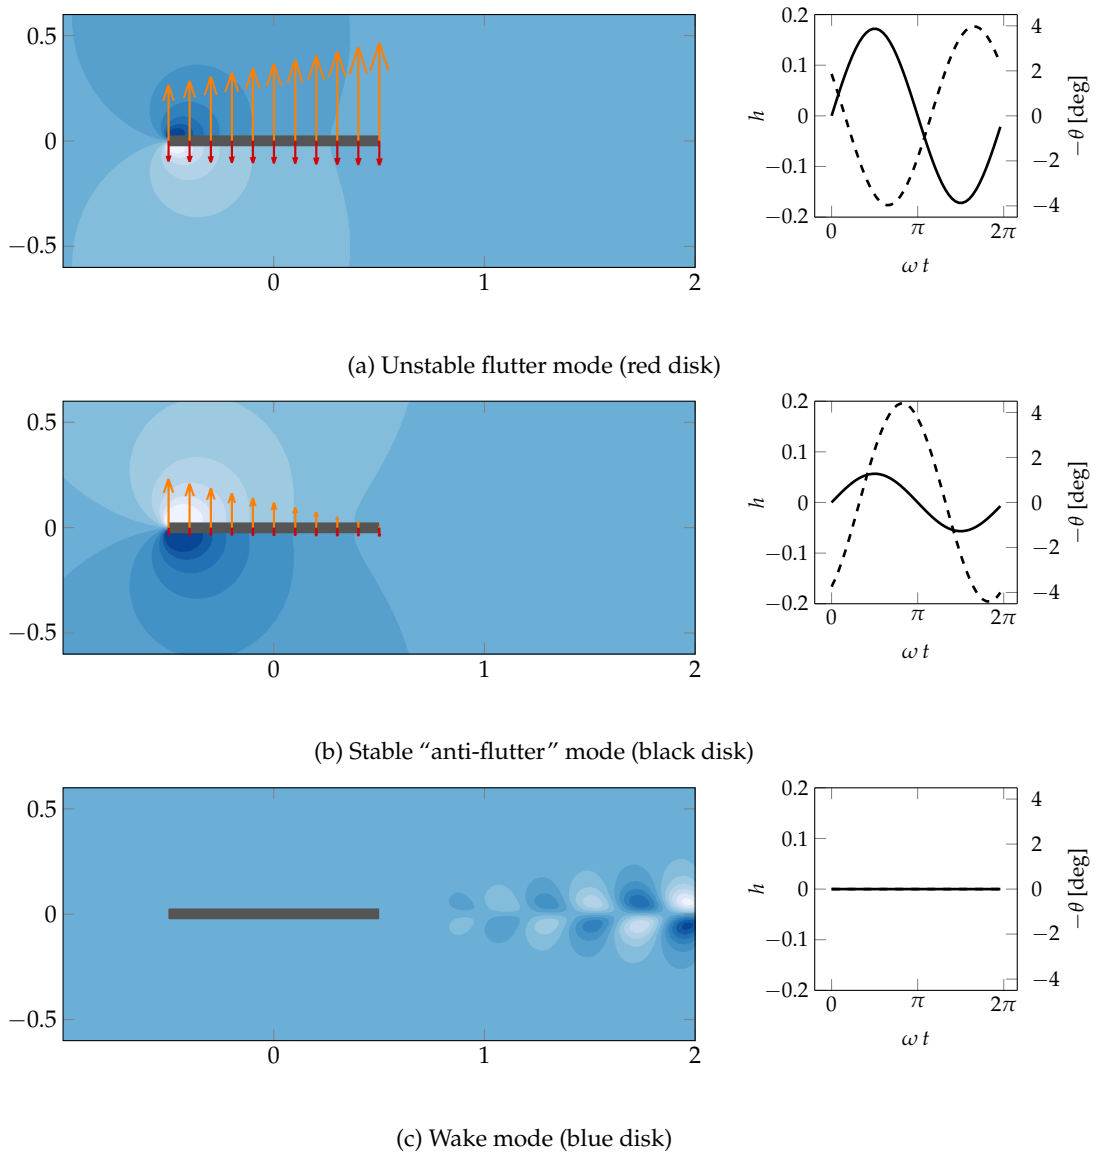


FIGURE 4.3.2: Visualisation of the three leading eigenmodes of fig. 4.3.1b. For each subfigure, the left part represents a snapshots of the mode at $t=3\pi/4$, where the blue colormap represents the pressure field, while the arrows materialise the solid displacement (orange) and velocity (red) vectors. On the right part, we show the time signals for the vertical displacement (solid line) and pitching angle (dashed line). The modes are normalized to unit solid mechanical energy.

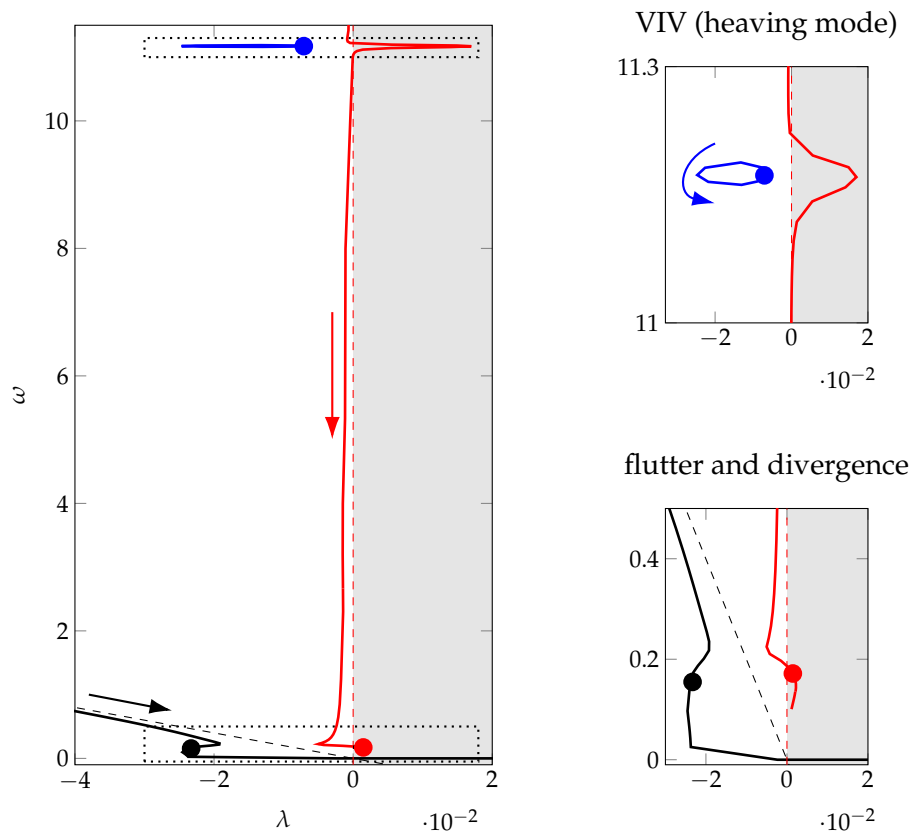


FIGURE 4.3.3: Loci of the three leading eigenvalues represented by disks in fig. 4.3.1b, as the reduced velocity varies ($U^* \in [0.05, 8]$). The dashed straight lines indicate the loci of the uncoupled (solid-only) solid modes (heaving in red, pitching in black). The arrows indicate the direction of increasing U^* . The dotted rectangles materialize the zones of VIV type and flutter/divergence type instabilities, of which zoomed views are exposed in the right part of the figure. The filled gray area represents the unstable half plane.

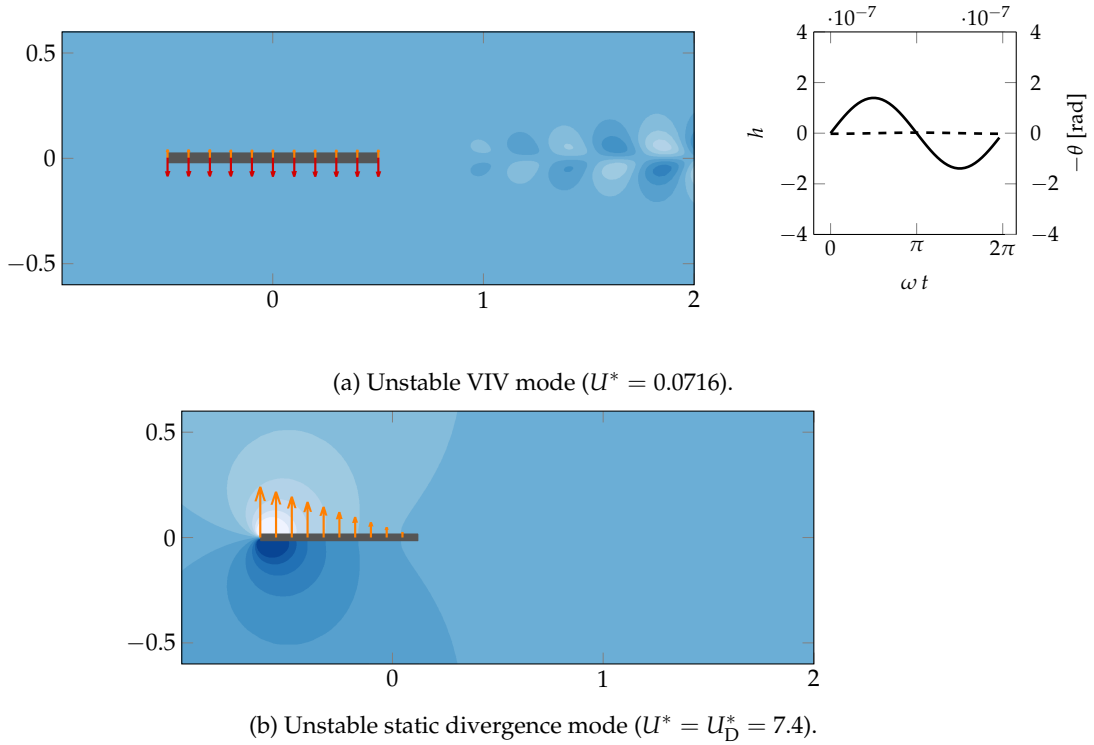


FIGURE 4.3.4: Eigenmodes corresponding to the (heaving) VIV and divergence instabilities observed in fig. 4.3.3. The orange (resp. red) arrows represent the displacement (resp. velocity) vectors whereas the pressure field is visualized by the blue colormap.

modes respectively tend towards the solid-only eigenmodes:

$$\frac{\lambda_h + \mathbf{i}\omega_h}{\omega_{0\theta}} = (-\Omega\zeta_h + \mathbf{i}\Omega\sqrt{1 - \zeta_h^2}) = 0 + 0.8 \times \mathbf{i}$$

$$\frac{\lambda_\theta + \mathbf{i}\omega_\theta}{\omega_{0\theta}} = (-\zeta_p + \mathbf{i}\sqrt{1 - \zeta_p^2}) = -0.05 + 0.999 \times \mathbf{i}$$

where we recall Ω is the heaving-to-pitching natural frequency ratio. As U^* increases, the frequencies of the branches first approach each other while growth rates slightly decrease. Around $U^* = 4$, both branches have equal frequency and a brutal increase of the growth rate of the heaving branch is observed so that the heaving branch goes unstable at $U_c^* = 4.5$. As a consequence, flutter is linked to the coalescence of the heaving and pitching frequencies. Hence the name *coupled-mode flutter*, usually associated to that instability [De Langre 2002]. The reduced velocity at which flutter is triggered is called the critical flutter velocity and is denoted U_c^* in the following. For even higher U^* , the pitching branch finally gets unstable for a reduced velocity $U_D^* = 7.4$, through the static divergence instability reported before. U_D^* is called the divergence velocity.

Let us finally mention that in the particular case studied here, it is the heaving branch that gets unstable *via* the flutter instability. However, by changing the damping parameters ζ_h and ζ_p , we could observe different scenarios where it is the pitching branch that gives rise to flutter. In all cases however, the common feature is the coalescence of the heaving and pitching frequencies, prior to flutter.




FIGURE 4.3.5: Evolution of the flutter and anti-flutter eigenmodes, marked respectively by ?? and ?? in fig. 4.3.1b, as a function of the reduced velocity ($\mathcal{Re} = 2900$, $\tilde{m} = 1000$). The flutter critical velocity U_c^* is defined as the reduced velocity for which the flutter mode represented in fig. 4.3.2a becomes unstable. For even higher U^* , a static divergence instability occurs for $U^* > U_D^*$. As $U^* \rightarrow 0$ both branches tend towards the solid-only eigenmodes. For the particular parameters used in the figure, those modes reduce to $\sigma = i0.8\omega_{0\theta}$ (heaving mode) and $\sigma = (-0.05 + i0.999)\omega_{0\theta}$ (pitching mode).

4.3.3 Increasing the steady angle of attack: from coupled-mode to single-mode flutter

In this paragraph, we vary the steady angle of attack of the plate and study how this impacts the occurrence of the flutter instability. In other words, we want to impose the value of the baseflow pitching angle $\theta_0 \neq 0$ as a parameter, and then study the linear stability of this steady solution. Mathematically however, this cannot be done freely since θ_0 , h_0 and the corresponding fluid variables $(\mathbf{U}_0, p_0, \Lambda_0)$ are already fully constrained by the steady-state problem, eq. (4.3a). As a consequence, an additional artificial “degree of freedom” must be added to the system in order to add the desired constraint on θ_0 . This degree of freedom is introduced as a *pre-stress* term in the pitching spring, which yields a “pre-stressed” version of the solid model, eq. (1.3):

$$\begin{aligned} \frac{\partial^2 h}{\partial t^2} + 2\zeta_h \left(\frac{\Omega}{U^*} \right) \frac{\partial h}{\partial t} + \left(\frac{\Omega}{U^*} \right)^2 h &= \frac{1}{\tilde{m}} C_L(\Lambda, \theta) \\ \frac{\partial^2 \theta}{\partial t^2} + 2\zeta_p \left(\frac{1}{U^*} \right) \frac{\partial \theta}{\partial t} + \left(\frac{1}{U^*} \right)^2 (\theta - \theta_{ps}) &= \frac{1}{r_\theta^2 \tilde{m}} C_M(\Lambda) \end{aligned} \quad (4.5)$$

where θ_{ps} is the parameter that tunes the magnitude of the pre-stress. A steady solution $(\mathbf{U}_0, p_0, \mathbf{\Lambda}_0, h_0, 0, \theta_0, 0)$ of the fluid-structure problem, using the above pre-stressed solid model then verifies:

$$\begin{aligned} \left(\frac{\Omega}{U^*}\right)^2 h_0 &= \frac{1}{\tilde{m}} \mathcal{C}_L(\mathbf{\Lambda}_0, \theta_0) \\ \left(\frac{1}{U^*}\right)^2 (\theta_0 - \theta_{\text{ps}}) &= \frac{1}{r_\theta^2 \tilde{m}} \mathcal{C}_M(\mathbf{\Lambda}_0) \end{aligned} \quad (4.6)$$

Since we impose θ_0 , the unknowns are now $(\mathbf{U}_0, p_0, \mathbf{\Lambda}_0, h_0, \theta_{\text{ps}})$ and verify eq. (4.6) (and the Navier–Stokes equations). The solution is trivially obtained in two steps. First, a fluid-only steady solution $(\mathbf{U}_0, p_0, \mathbf{\Lambda}_0)$ is computed for a plate with a pitching angle equal to the imposed θ_0 . Second, θ_{ps} and h_0 are obtained from eq. (4.6). With that procedure, we obtain a steady solution of the coupled problem that possess the desired θ_0 . Note that h_0 will vary depending on θ_0 : the higher θ_0 , the higher the lift $\mathcal{C}_L(\mathbf{\Lambda}_0, \theta_0)$, and thus the higher h_0 . However, it is clear that the value of h_0 has no physical meaning (and no impact on stability) since the steady flow is invariant by a vertical translation, and the linearized fluid-structure problem eq. (4.3b) does not depend on h_0 .

Turning now to the presentation of the results, let us fix the value of the Reynolds number to $Re = 500$. As shown in fig. 4.3.6b, the flow remains attached to the plate until $-\theta_0 = 7^\circ$. For these low angles of attack, the lift and moment coefficient are mostly linear functions of θ_0 . For $-\theta_0 = 8^\circ$ a large recirculation area, delimited by a thin black line, appears and covers the whole plate. At this angle, the lift and moment coefficient presented in fig. 4.3.6a start to plateau. For $-\theta_0 > 10^\circ$, the recirculation region keeps growing (not shown) and a slight decrease of the magnitude of the lift and moment coefficients is observed. This behaviour is similar to a classical stall phenomenon. However stall is here less brutal than it is in high-Reynolds aerodynamic flows. This “lighter” stall is a known feature of low-to-moderate Reynolds flows around plates [Amandolese et al. 2013] or airfoils [Mahbub Alam et al. 2009]. Finally, note that for $-\theta_0 \geq 9.8^\circ$, the steady baseflow has such a large recirculation region that it is hydrodynamically unstable (dashed line in fig. 4.3.6a) due to a classical wake mode (not shown) becoming unstable with frequency $\omega \simeq 4$. Thus, for $-\theta_0 \geq 9.8^\circ$, the baseflow lift and moment coefficients reported here are to be used with caution, as wake unsteadiness is likely to make them progressively drift away from their mean-flow counterparts, observed experimentally. In any case, the subsequent discussion on the flutter instability will be restricted to $0^\circ \leq -\theta_0 \leq 10^\circ$, thus avoiding any unsteady fluid effects.

In fig. 4.3.7a, we present the evolution of the heaving (red lines) and pitching (black lines) modes as the reduced velocity is increased, and for different θ_0 . For the case $\theta_0 = 0^\circ$ (dashed lines), the typical coupled-mode flutter (CMF) scenario detailed in section 4.3.2 is observed, with the frequencies of the two modes approaching each other and giving rise to the instability. For $0 < -\theta_0 < 7$, the same coupled-mode flutter behavior is still observed. Note that, as $-\theta_0$ increases, the branch giving rise to CMF transitions from the heaving branch (case $-\theta_0 = 0^\circ$, dashed) to the pitching branch (case $-\theta_0 = 6^\circ$, dash-dotted). In addition, the critical threshold for CMF is progressively delayed to higher values: from $U_c^* = 4.95$ at $-\theta_0 = 0^\circ$ to $U_c^* = 6.05$ at 6.8° . For $-\theta_0 \geq 7$, the scenario qualitatively changes as the instability is triggered on the pitching mode (black solid curve, $-\theta_0 = 10^\circ$) despite the absence of any frequency coalescence. This type of flutter which only involves a one-mode pitching

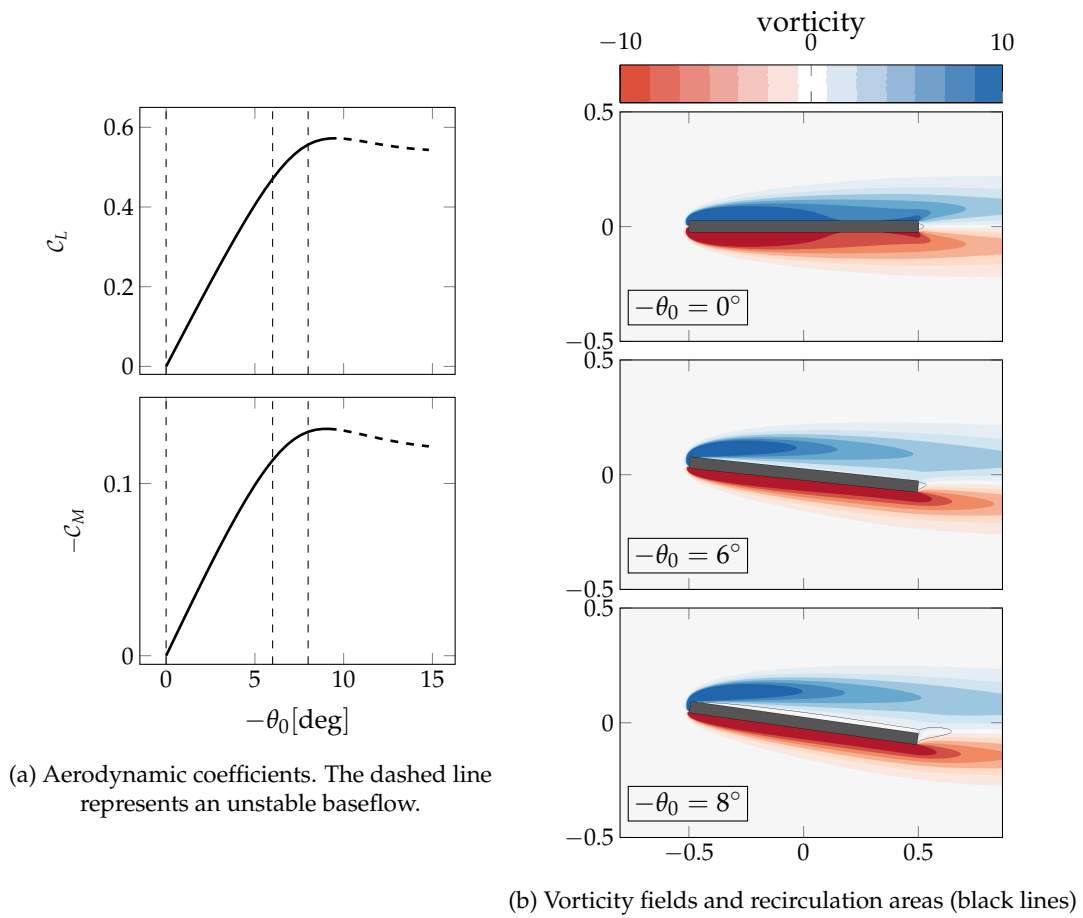


FIGURE 4.3.6: Features of the steady flow around the fixed plate at different angles of attack ($Re = 500$).

motion (see the eigenvector represented in fig. 4.3.8c) is referred to as *pitching single-mode flutter* (pSMF). The transition from CMF to pSMF around $-\theta_0 \simeq 7$ is accompanied by a brutal change in the variation of U_c^* with θ_0 , as shown in fig. 4.3.7b. In the pSMF range, the critical velocity decreases with increasing angle of attack. Overall, the critical eigenmodes for different angles of attack represented in fig. 4.3.8, show clearly the transition between a coupled heaving and pitching flutter (CMF) at low incidence and a pitch-dominated flutter (pSMF) at higher incidence.

It is interesting to note that the features reported above have been observed experimentally, for example by Razak et al. (2011), where the flutter instability on a heaving and pitching NACA0012 profile was shown to transition from CMF to pSMF as the static angle of attack is increased. The higher Reynolds number ($Re \simeq 10^5$) and different airfoil profile used in this study does not allow a quantitative comparison with our results, but brings the idea that the mechanisms at play in those two different configurations are similar. However, further investigations of that statement are needed, in particular with respect to the cohabitation of coupled-mode flutter and laminar separation flutter [Poirel et al. 2008; Poirel et al. 2012] for these transitional Reynolds numbers.

For $-\theta_0 = 10^\circ$, we notice in fig. 4.3.7a that, in addition to the pitching mode (solid black line) becoming unstable around $U_c^* = 4.8$, the heaving mode (solid red line) is also very close to unstable. It actually does become (slightly) unstable around $U_c^* = 5.2$, giving rise to a (secondary) *heaving single-mode flutter* (hSMF). The corresponding critical velocity is materialized in fig. 4.3.7b by the red disk. We should mention that hSMF could have been anticipated already by observing the shape of the steady lift curve fig. 4.3.6a. Indeed, using a quasi-steady approach (see section 4.5.2) that approximates the heaving motion *via* an effective angle of attack $\theta_{\text{eff}} = \dot{h}$, we can approximate the linearized fluid force for a heaving motion around $-\theta_0 = 10^\circ$ as

$$\left. \frac{dC_L^{\text{stat}}}{d\theta} \right|_{-\theta_0=10^\circ} \dot{h}_1$$

with C_L^{stat} the steady lift coefficient shown in fig. 4.3.6a and where we recall the index 1 refers to the perturbation around the steady-state (*cf* eq. (4.2)). At $-\theta_0 = 10^\circ$, the slope $dC_L^{\text{stat}}/d\theta$ is *positive*². Using the above expression in the linearized heaving momentum equation yields a negative damping term that can destabilize the system, if it counterbalances the structural damping (*i.e.* if U^* is high enough):


$$\ddot{h}_1 + \left(2\zeta_h \frac{\Omega}{U_c^*} - \frac{1}{\tilde{m}} \left. \frac{dC_L^{\text{stat}}}{d\theta} \right|_{-\theta_0=10^\circ} \right) \dot{h}_1 + \frac{\Omega^2}{U_c^{*2}} h_1 = 0$$

Note finally that since we use $\zeta_h = 0$ across this thesis, the system should be immediately unstable, already at $U^* = 0$. The (small) damping that prevents that — and that is not represented in the above equation — is likely to be due to viscous dissipation.

4.4 Parametric explorations

In the previous section, we exhibited four types of flow-induced instabilities, depending on the reduced velocity and steady angle of attack. For the rest of the study,

²Again, recall our sign convention for pitching angles: positive in the trigonometric orientation, *i.e.* nose-down



(a) Evolution of the pitching (black) and heaving (red) modes for increasing reduced velocity at different steady angles of attack : $\alpha = 0^\circ$ (dashed), 6° (dash-dotted), 8° (dotted) and 10° (solid).

(b) Critical flutter speed.

FIGURE 4.3.7: Transition from coupled-mode flutter (CMF) to pitching single-mode flutter (pSMF) as the steady angle of attack increases. When the steady solution is stalled ($-\theta_0 > 9^\circ$), a single-mode flutter on the heaving mode (hSMF) is also triggered. ($Re = 500, \tilde{m} = 10^3$)

we focus on the case with zero steady angle of attack, and parametrically explore the effect of mass ratio and Reynolds number.

4.4.1 Effect of mass ratio

When varying the mass ratio \tilde{m} , similar destabilization scenarios to those depicted in fig. 4.3.3 and fig. 4.3.5 are found for the VIV, flutter and divergence modes, but with different values of the critical thresholds. In this paragraph, we aim at studying the effect of \tilde{m} on the instability thresholds. In other words, we wish to compute neutral curves for each instability, in the plane (\tilde{m}, U^*) .

In fig. 4.4.1, we show the evolution of the critical flutter, divergence and VIV thresholds as a function of \tilde{m} . Three different regions are then defined in the plane (\tilde{m}, U^*) corresponding to regions where the flutter mode (red area), the divergence mode (grey area) or the the VIV mode (blue area) are unstable. Conversely, in the non-shaded area of the plane, all eigenvalues have negative real parts and the system is linearly stable. First, we observe that the divergence boundary $U_D^* = f(\tilde{m})$ is a straight line with slope 1/2 in the log-log plane. This tendency can be analytically explained in the hypothesis of pure pitching motion and thin airfoil theory, for which the analytical expression for the divergence threshold is $U_D^* = \sqrt{r_\theta^2 / (\pi/2) \tilde{m}}$ [E. H. Dowell et al. 1989, §2.1]³, where we recall r_θ is the radius of gyration (*cf* table 1.1).

The flutter boundaries behave similarly to the divergence boundary at high \tilde{m} with a slope close to 1/2. As \tilde{m} decreases, the slope attenuates until the critical velocity reaches a minimum around $\tilde{m} \simeq 20$. For even lower \tilde{m} , *i.e.* around $\tilde{m} \simeq 7$, the neutral curve encounters a turning point. As a consequence, no unstable flutter mode is observed for very low $\tilde{m} \leq 7$, in the investigated velocity range. A second consequence of the “turning” of the flutter neutral curve is that in cases where flutter can develop (*i.e.* $\tilde{m} > 7$), there exists a reduced velocity higher than U_c^* for which the flutter mode restabilizes. To our best knowledge, this restabilization of the flutter

³This expression is obtained by stating that, at the critical divergence velocity, the torsional spring elastic force is balanced by the fluid moment from thin airfoil theory: $U_D^{*-2} \theta = 1 / (r_\theta^2 \tilde{m}) \pi / 2 \theta$.

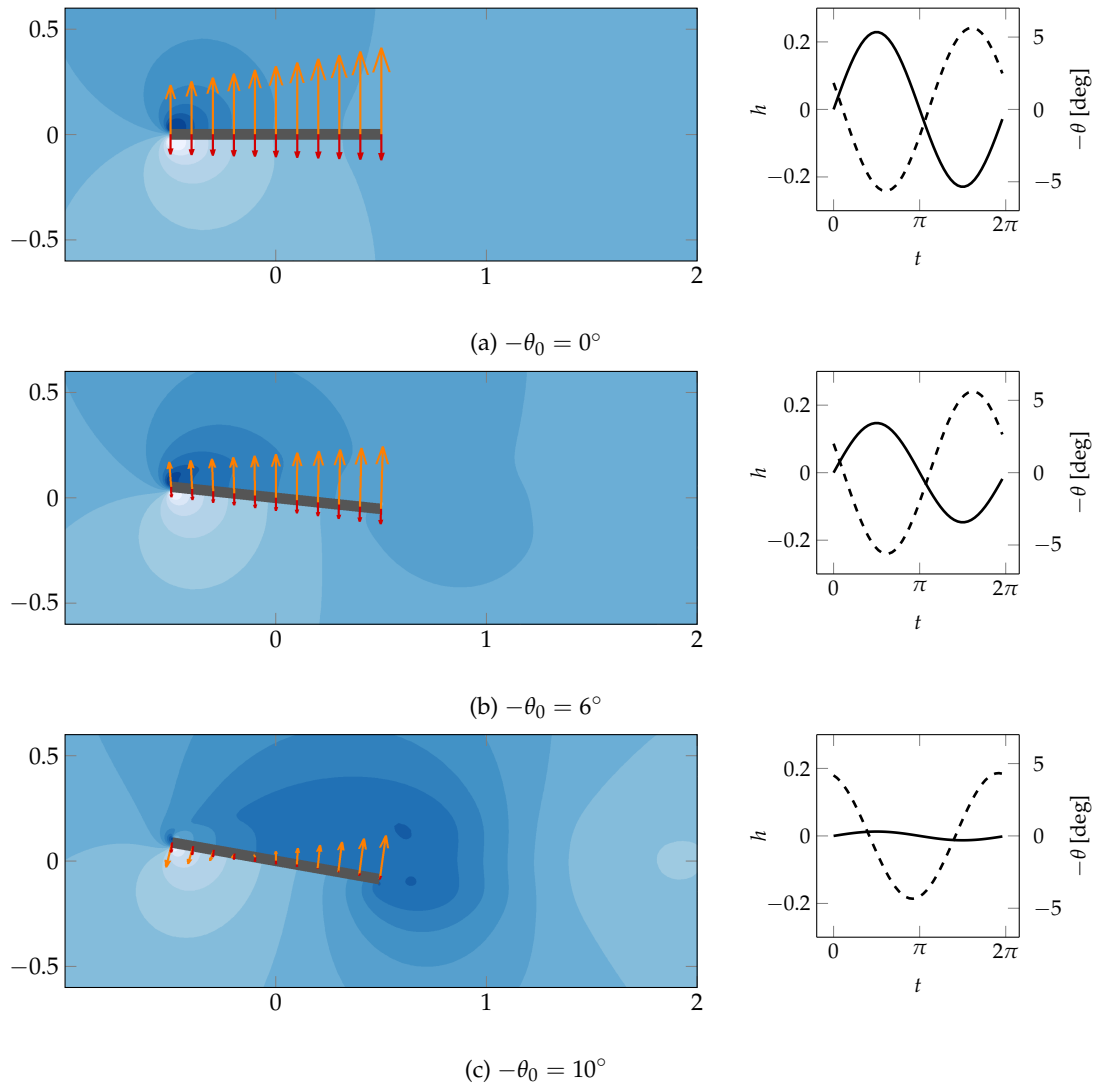


FIGURE 4.3.8: Visualisation of the critical flutter eigenmodes at three different angles of attack. For each subfigure, the left part represents a snapshots of the mode at $t=3\pi/4$, where the blue colormap represents the pressure field, while the arrows materialise the solid displacement (orange) and velocity (red) vectors. On the right part, we show the time signals for the vertical displacement (solid line) and pitching angle (dashed line).

mode was never reported in the literature and the underlying physical mechanism is to be precised. For now, we continue with the description of fig. 4.4.1 and will come back later on that issue (*cf* section 4.4.2 and fig. 4.5.2).

At the intersection of the flutter and divergence neutral curves, a codimension-two point can be identified for a particular mass ratio $\tilde{m}^{II} \simeq 11$ and $U^* \simeq 0.9$. At this point, both the flutter and divergence modes are neutrally stable and can be involved in the observed nonlinear dynamics. Typically, the simultaneous presence of these two modes may explain the occurrence of so-called “dynamic divergence” [Chae et al. 2013; Chae et al. 2016] where a combination of oscillations and mean deviation of the wing/plate may be observed in the nonlinear regime. These studies were performed at higher Reynolds numbers than ours ($Re \simeq 10^6$), on a different 2D section (NACA0016 airfoil) and with different structural parameters. However, the range of dynamic divergence reported there is roughly $\tilde{m} \in [1.6, 6.3]$, which is the same order of magnitude as the values found in the present work. Note that the effect of Re on the value of \tilde{m}^{II} is studied in the next section (fig. 4.4.5).

Turning finally to the VIV instability region, represented in blue, we observe first that no matter the mass ratio, VIV always occur at much lower values of U^* than the flutter and divergence instabilities. The shape of the VIV area strongly depends on the mass ratio. At high enough \tilde{m} , two separated VIV regions are observed. The lower one is a nearly constant interval of reduced velocities between $\tilde{m} = 20$ and $\tilde{m} = 10^3$, whereas the upper one vanishes around $\tilde{m} = 10^2$. The existence of two separated regions of VIV is related to the existence of the two solid modes that resonate with the vortex-shedding mode at different U^* . In the upper region, the resonance of the vortex-shedding mode occurs with the pitching mode when $\omega_{0\theta} = 1/U^* \simeq \omega_w$, *i.e.* $U^* \simeq 0.09$ (upper dashed line). In the lower region, it occurs with the heaving mode when $\omega_{0h} = \Omega/U^* \simeq \omega_w$, *i.e.* $U^* \simeq 0.07$ (lower dashed line). Towards the low \tilde{m} , the two VIV regions widen and ultimately merge in one large unique region. Let us mention that the widening of the VIV region at low mass ratio has also been observed in the case of an only-heaving spring-mounted cylinder [Navrose et al. 2016]⁴.

We end that section by proposing in fig. 4.4.2 an alternative visualization for the neutral curves already shown in fig. 4.4.1. Indeed, it may seem inconsistent to the reader that the neutral curve for the static divergence instability depends on mass ratio \tilde{m} : considering a fluid with fixed density, \tilde{m} is simply proportional to the solid mass, which should have no effect on a static (divergence) instability. In fact, the $\propto \sqrt{\tilde{m}}$ behavior of U_D^* is only an artifact of the definition of the reduced velocity, that is: dimensional velocity U_∞ divided by natural pitching frequency, which depends itself on solid mass. A more sound choice of non-dimensional parameter is thus to replace the reduced velocity by the non-dimensional dynamic pressure, noted q^* , that represents the ratio between the fluid inertial forces and solid elastic forces⁵:

$$q^* = \frac{U^{*2}}{\tilde{m}r_\theta^2} = \frac{\frac{1}{2}\rho_f U_\infty^2}{K_\theta/c^2} \quad (4.7)$$

This is done in fig. 4.4.2 where the neutral curves are shown in the (\tilde{m}, q^*) plane. The expected constant divergence threshold (in terms of dynamic pressure) is retrieved, together with the idea that flutter boundaries (in terms of dynamic pressure) are also independent of \tilde{m} , at high enough \tilde{m} . Notice finally that against intuition, low

⁴Note however that their bifurcation parameter was Re instead of the reduced velocity.

⁵This number is also referred to as the *Cauchy number* in the literature




FIGURE 4.4.1: Neutral curves in the plane (\tilde{m}, U^*) for flutter (red area, thick line), divergence (gray area, dashed thick line) and VIV (blue area, thin line) instabilities ($Re = 2900$). The thin horizontal dashed lines materialize the reduced velocities for which the heaving (lower line) and pitching (upper line) natural frequencies equal the VIV mode frequency ω_w .

\tilde{m} require higher q^* to trigger flutter. In other words, the lighter the structure, the higher must be the typical fluid pressure forces *with respect* to the solid elastic forces.

To conclude, fig. 4.4.1 and 4.4.2 can be thought of as representations of two different “dimensional” behaviors. In the one hand, if one changes solid mass while maintaining constant solid *natural frequencies*, then fig. 4.4.1 faithfully traduces the trend for the dimensional critical wind velocity as a function of solid mass⁶. On the other hand, if one changes solid mass while maintaining constant solid *stiffness*, then fig. 4.4.2 is the more natural representation. Obviously, in practical wing design, the engineer usually navigates somewhere in between those two “idealized” situations due to the inter-dependence of stiffness (Young modulus) and mass (density) for real-life materials (*cf* Ashby’s material property charts [Ashby 1999]). In the remaining of this chapter, we use U^* or q^* depending on the features we wish to highlight.

4.4.2 Effect of Reynolds number

In the previous sections we explored the linear behavior of the spring-mounted plate by varying some classical aeroelastic parameters that are the reduced velocity and the mass ratio. We now exploit the full potential of a global stability approach by studying how viscosity affects flutter and divergence. Its effect on the VIV instability is left aside and we refer the interested reader to dedicated works on this aspect [Meliga et al. 2011; Navrose et al. 2016].

Influence of Reynolds on the steady baseflow As a first step to understand the effect of the Reynolds number on flutter and divergence, we give in fig. 4.4.3 some insights into the steady baseflow solution at $Re = 10$, $Re = 110$, $Re = 500$ and $Re = 2900$. First, the corresponding solutions (at zero incidence) are presented in

⁶Rigorously, we note that the Reynolds number is kept fixed in fig. 4.4.1 and 4.4.2. Experimentally however, if the dimensional wind velocity is varied in a given fluid and for a given size of the plate, Re will vary. The effects of Reynolds are studied in the next section.




FIGURE 4.4.2: Alternative representation, using dynamic pressure q^* , of the neutral curves of fig. 4.4.1. Same color conventions are used.

fig. 4.4.3c to illustrate the clear evolution of flow features as Re increases. The baseline case $Re = 2900$ possesses the classical features of a high-speed flow with uniform velocity far enough from the plate and sharp velocity gradients concentrated in a thin enough boundary layer⁷. As Reynolds decreases, the boundary layer thickens so much that diffusion effects finally dominate the whole represented domain at $Re = 10$. In fig. 4.4.3a, we present the lift and moment (evaluated at the elastic axis, which is the mid-chord in this manuscript) coefficients at those four Reynolds numbers, as a function of angle of attack⁸. For $Re = 10$, C_L and C_M evolve almost linearly with angle of attack, in the range of angles investigated. Indeed, for such a low Reynolds, the nonlinear Navier–Stokes equations tend towards the linear Stokes problem. The slope of the lift coefficient is around 3.22, which is comparable to the results of Sun et al. (2001) and Sun et al. (2004), who reported a slope of 2.8 at $Re = 13.57$. The comparison between our work and Sun’s is however only qualitative since his studies took into account compressible ($Ma = 0.2$) and rarefaction effects that may be non-negligible at such low Reynolds numbers [Sun et al. 2001]. As the Reynolds number increases, the aerodynamic coefficients depend more and more nonlinearly on the angle of attack. At $Re = 500$, a plateau appears for $|\theta| > 10^\circ$. At $Re = 2900$, the plateau transforms into a stronger drop of both C_L and C_M , and occurs at lower $|\theta| \simeq 3^\circ$. The increase of Reynolds also leads to a significant increase of the slope at $\theta = 0^\circ$, as observed in fig. 4.4.3a. More quantitatively, from fig. 4.4.3b we see that between $Re = 10$ and the baseline case $Re = 2900$ investigated in the previous section, the slopes are multiplied by two.

To conclude, we like to remind that the results presented here are for steady Navier–Stokes solutions that may not necessarily be observed in practice, because unstable. More precisely, it was previously checked that the zero incidence solution is stable for all $Re \leq 2900$. However, it is likely to not be the case for every angle of attack considered in fig. 4.4.3. The systematic check of linear stability was not performed here for each angle of attack, and the reader is reported to section 4.3.3 for such an investigation, at $Re = 500$. In any case, our purpose in this section is to

⁷Obviously, due to the low Reynolds number, compared to classical aerodynamics flows (planes, ...), the thickness is still of the order of the plate thickness.

⁸Due to the sign convention used in this work, the angle of attack as traditionally defined in aerodynamics — *i.e.* positive nose-up — is equal to $-\theta$. Hence the choice to represent $C_L = f(-\theta)$ and $-C_M = f(-\theta)$ in fig. 4.4.3a.

study viscosity effects on flutter and divergence instabilities around $\theta = 0^\circ$. These only depends on the linearized aerodynamics around $\theta = 0^\circ$, which are characterized by the slopes of the aerodynamic coefficients at zero.

Flutter and divergence thresholds for different Reynolds In fig. 4.4.4, we show the flutter (solid lines) and divergence (dashed lines) neutral curves in the (\tilde{m}, q^*) plan and for three different Reynolds numbers spanning three orders of magnitude: $Re = 10$ (brown), $Re = 110$ (gray) and $Re = 2900$ (black). For each Reynolds number, the divergence instability occurs above the corresponding dashed lines, while the flutter instability occurs inside the corresponding solid curve. First, it is observed that the critical dynamic pressure for divergence decreases as Re increases. By classical arguments [E. H. Dowell et al. 1989, §2.1], this change can be linked to the increase in the slope of the steady lift and moment coefficients, previously reported in fig. 4.4.3b. A similar trend is observed for the flutter threshold in the high \tilde{m} limit. In addition, the restabilization of flutter for high q^* is more and more delayed when Re increases. This suggests that this restabilization is typically a viscous effects, which would explain why it was never reported — to the author’s knowledge — since most existing flutter studies use potential flow models. More insights on that aspect are provided in section 4.5 where we compare the present approach with the classical (potential) Theodorsen theory. Finally, we note a large change in the position of the codimension-two point between $Re = 10$ and $Re = 110$. Indeed, between those two Reynolds, its position drops from $\tilde{m} \simeq 300$ to $\tilde{m} \simeq 20$. Between $Re = 110$ and the baseline case used in the previous section at $Re = 2900$, further shift is observed but in much less important proportion (from $\tilde{m} \simeq 20$ to $\tilde{m} \simeq 10$).

Due to the wide variations of the high- q^* limit of the flutter region, two scenarios are possible, depending on Re . At high-enough Re , the “classical scenario” is observed, where, as one increases q^* , the flutter mode destabilizes first, remains unstable, and finally the divergence mode destabilizes. At low Reynolds ($Re = 10$), the flutter mode destabilizes first, restabilizes — leaving the system stable on some range of q^* — and finally divergence is triggered.

4.4.3 Full parametric exploration of flutter thresholds in the (\tilde{m}, Re) plane

In order to provide an intermediate conclusion to our parametric explorations, we propose a synthetic view of the results obtained in sections 4.3.2 to 4.4.2 for the flutter instability. To that aim, we represent critical flutter quantities in the (\tilde{m}, Re) parameter space. By “critical” we refer to the low U^* (or equivalently low q^*) end of the flutter area in fig. 4.4.1, as it is the threshold of interest for typical aeronautical applications.

In fig. 4.4.5(a), we start with showing the critical flutter velocity U_c^* in the (\tilde{m}, Re) plane. With this representation, the mass ratio clearly has a dominant effect on the value of U_c^* , in comparison to the Reynolds number. This is the $U_c^* \propto \sqrt{\tilde{m}}$ trend already highlighted in fig. 4.4.1. To see more clearly the effect of Re , it is better to switch to the representation of the flutter threshold in terms of the critical dynamic pressure q_c^* , as proposed in fig. 4.4.5(b). Indeed, with this representation, the “automatic” $\propto \sqrt{\tilde{m}}$ of the flutter threshold that dominates fig. 4.4.5(a) is eliminated, leaving place to a clearer visualization of the destabilizing effect of increasing Re . By comparing fig. 4.4.5(a) and (b), we clearly see that depending on the chosen representation of the flutter threshold (U_c^* or q_c^*), decreasing the mass ratio can either be seen, for all Re , as destabilizing (a) or slightly stabilizing (b). We refer the reader to the end of section 4.4.1 for a discussion on the “dimensional” interpretation of

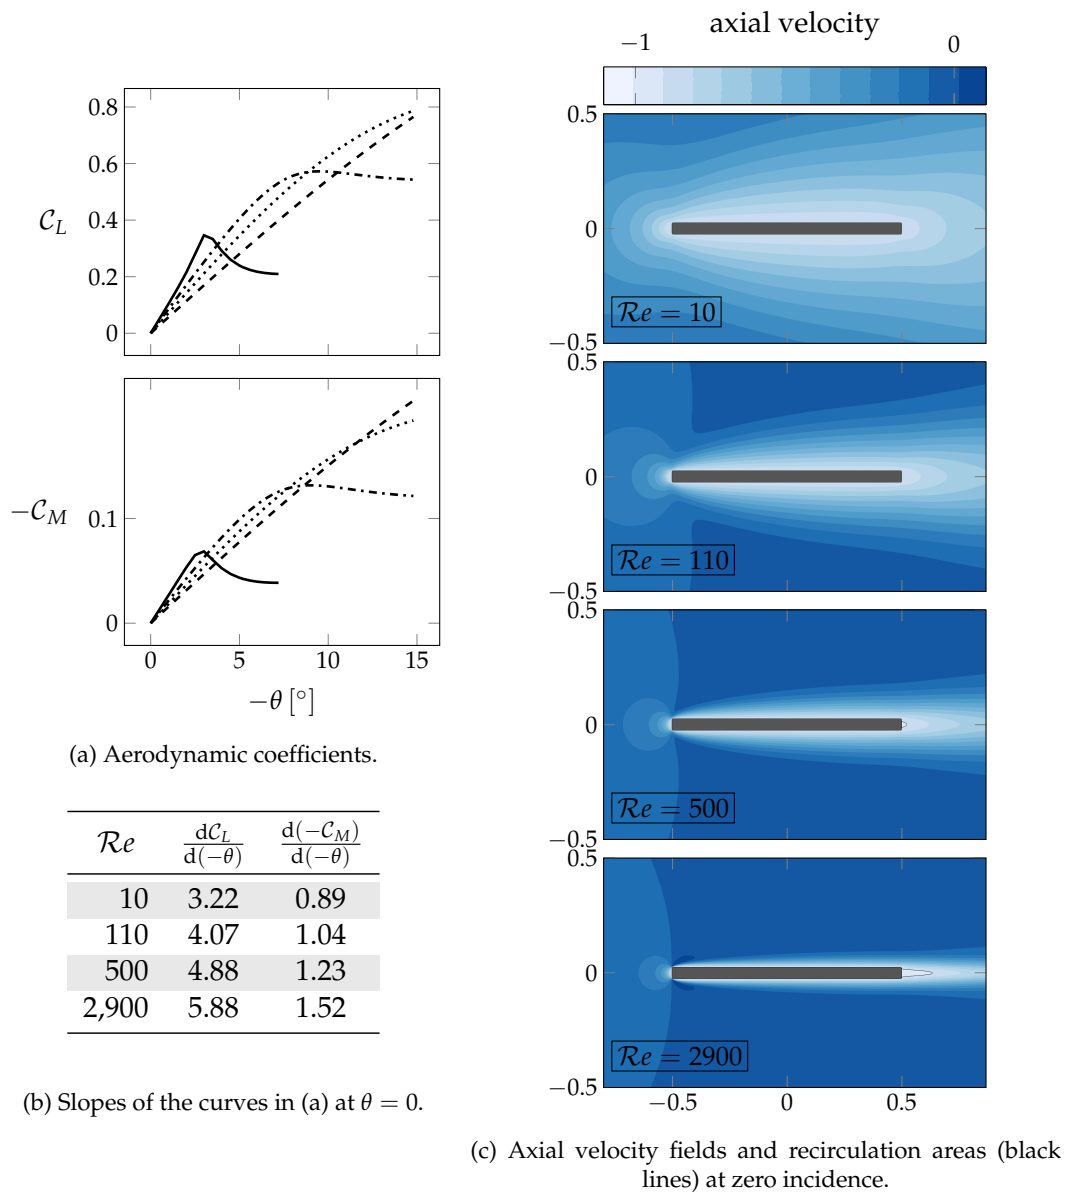


FIGURE 4.4.3: Features of the steady flow for different Reynolds numbers. In (a), the $C_L(\theta)$ and $C_M(\theta)$ curves for $Re = 10$ (dashed), $Re = 110$ (dotted), $Re = 500$ (dashdotted) and $Re = 2900$ (solid) are shown. In (b), the slopes of those curves at $\theta = 0$ are tabulated. In (c), we present the steady solutions associated to the different Reynolds at zero incidence.




FIGURE 4.4.4: Linear stability neutral curves for flutter (solid lines) and divergence (dashed lines) instabilities in the plane (\tilde{m}, q^*) for different Reynolds numbers: $Re = 10$ (brown), $Re = 110$ (gray), $Re = 2900$ (black). For a clear identification of the stable/unstable regions, the reader is reported to fig. 4.4.2.

using either U_c^* or q_c^* as a flutter threshold. The last map, fig. 4.4.5(c) provides additional information on the frequency of the flutter mode at criticality, ω_c . Similarly as in fig. 4.4.5(a), the flutter frequency is dominantly governed by the mass ratio with the general trend that heavier (resp. lighter) structures yield low-frequency (resp. high-frequency) flutter.

We conclude this paragraph by discussing the boundaries of the flutter domain represented in the three maps of fig. 4.4.5. At low mass ratios, the flutter domain is limited by the codimension-two point evidenced in fig. 4.4.1 and marked here by the thin dashed curve. Indeed, for $\tilde{m} < \tilde{m}^{II}$ the flutter threshold, if it exists at all, is triggered after divergence occurs. For clarity, we prefer to remove this scenario from fig. 4.4.5, which is the reason why, the region $\tilde{m} < \tilde{m}^{II}$ has been shaded in gray. At high Reynolds number, we choose to limit the exploration to $Re \leq 10^4$. Also, the horizontal dashed line materializes the critical Reynolds number, $Re_{c,wake} \simeq 2925$, for the wake instability developing on the fixed plate. Indeed, for $Re > Re_{c,wake}$, the steady baseflow is not stable anymore so that flutter thresholds obtained with stability analyses of steady-states, as performed here, are questionable. A stability analysis based on the mean flow (time-averaged over the vortex-shedding period) or even on the Floquet stability of the vortex-shedding periodic solution may be more appropriate.

4.5 Comparison to simplified fluid models

In this section, we wish to compare the flutter stability results from the global stability approach with simplified fluid models like the classical Theodorsen theory [Theodorsen 1935] or a quasi-steady model [E. H. Dowell et al. 1989, §3.4.2.2]. Obviously, vortex-induced vibrations are out of the scope of this section as none of those models can capture a wake instability. We start by briefly recapitulating the basics of Theodorsen's theory in section 4.5.1, and the quasi-steady approach in section 4.5.2. The procedure used to calibrate both models with steady viscous flow data is explained in section 4.A. In section 4.5.3, numerical comparisons for the flutter thresholds prediction are exposed.

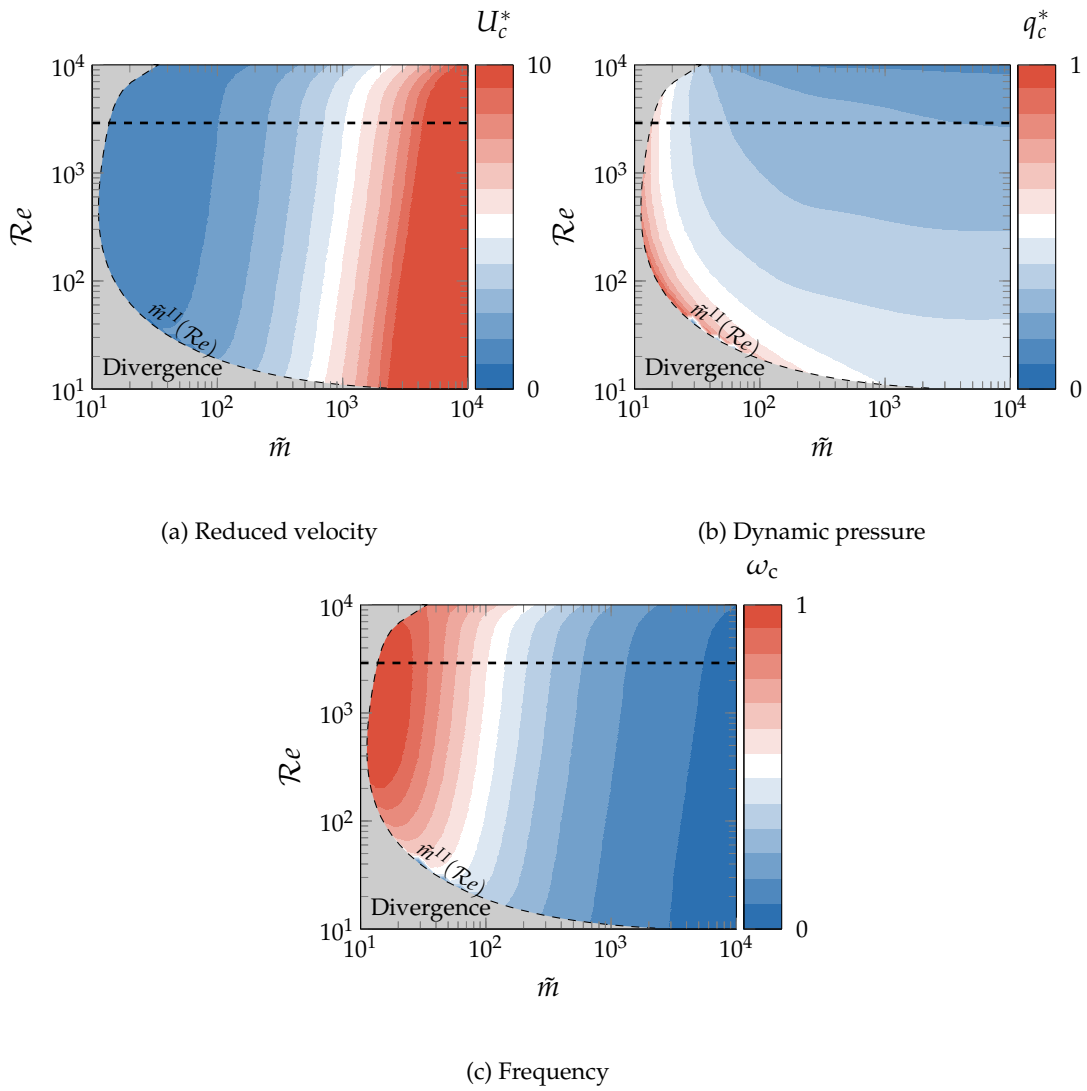


FIGURE 4.4.5: Influence of mass ratio and Reynolds number on different critical quantities for flutter instability.

4.5.1 Theodorsen's model

In his seminal paper, [Theodorsen \(1935\)](#) proposed analytical expressions for the fluid lift and moment applied on an infinitely thin flat plate, force to oscillate in pitch and plunge motions. In the original theory, the effect of an additional trailing edge flap was included, but is omitted here. The theory uses the framework of (unsteady) potential flows and is derived by separating the solution into two contributions: a non-circulatory and a circulatory part. The non-circulatory part is obtained with the classical steps of (steady) potential flow theory: (i) Joukowski's conformal transformations, (ii) distribution of sources and sinks on the plate so as to find a potential that satisfies the boundary condition (continuity of normal velocity) and (iii) (unsteady) Bernoulli theorem in order to finally retrieve a pressure distribution from the potential field. As the non-circulatory part does not satisfy the Kutta condition at the trailing edge, Theodorsen adds a circulatory contribution. It consists in a distribution of bound vortices on the plate, with the corresponding distribution of counter-rotating vortices in the wake. With such a set-up, the Kutta condition can be satisfied, while still enforcing the boundary condition on the plate and so that the full total circulation is maintained to zero. The workplan sketched above yields the following expressions for the lift and moment coefficients. The interested reader is reported to the enlightening work of [Bisplinghoff et al. \(1955, §5.6\)](#) for a step-by-step derivation of those.

$$\begin{aligned}
 C_L &= \underbrace{-\frac{\pi}{2} [\dot{\theta} + \dot{h} - a\ddot{\theta}]}_{\text{non-circulatory}} + \underbrace{C(k) \frac{dC_L^{\text{stat}}}{d\theta} \Big|_{\theta_0} \left[\theta + \dot{h} + \left(\frac{1}{4} - a\right) \dot{\theta} \right]}_{\text{circulatory}} \\
 C_M &= \underbrace{\frac{\pi}{4} \left[a\dot{h} - \left(\frac{1}{16} + 2a^2\right) \ddot{\theta} \right]}_{\text{non-circulatory}} + \underbrace{\frac{\pi}{4} \left(2a - \frac{1}{2}\right) \dot{\theta} + C(k) \frac{dC_M^{\text{stat}}}{d\theta} \Big|_{\theta_0} \left[\theta + \dot{h} + \left(\frac{1}{4} - a\right) \dot{\theta} \right]}_{\text{circulatory}}
 \end{aligned} \tag{4.8}$$

where a is the nondimensional distance between the elastic axis and the mid-chord⁹ (counted positively when the elastic axis is downstream of the mid-chord), and where $C(k = \omega/2)$ is the Theodorsen function, which can be analytically expressed using Hankel functions (*c.f.* [\[Bisplinghoff et al. 1955\]](#)). The circulatory terms in eq. (4.8) have a typical form composed of the steady lift and moment slopes at zero incidence, $dC_L^{\text{stat}}/d\theta|_{\theta_0}$ and $dC_M^{\text{stat}}/d\theta|_{\theta_0}$, multiplied by an *effective* angle of attack $\theta + \dot{h} + (1/4 - a)\dot{\theta}$ and the Theodorsen function. The second circulatory term in C_M is more difficult to interpret physically. As for the non-circulatory term, they mainly provide added mass to the system.

In the circulatory parts, the slopes $dC_L^{\text{stat}}/d\theta|_{\theta_0}$ and $dC_M^{\text{stat}}/d\theta|_{\theta_0}$ of the *steady* lift and moment curves are used to define the unsteady response. In the original theory, the thin airfoil theory values naturally arise:

$$\frac{dC_L^{\text{stat}}}{d\theta} \Big|_{\theta_0} = -2\pi \qquad \frac{dC_M^{\text{stat}}}{d\theta} \Big|_{\theta_0} = \frac{\pi}{2} (1 + 4a)$$

However, several authors (*e.g.* [\[Amandolese et al. 2013; Brunton et al. 2013\]](#)) have shown that calibrating the original Theodorsen model with “true values” — *i.e.* values extracted from experimental or numerical steady $C_L^{\text{stat}} - \theta$ and $C_M^{\text{stat}} - \theta$ curves —

⁹Given our particular plate set-up, the mid-chord is also the center of gravity. In addition, we place the elastic axis at the center of gravity ($x_\theta = 0$), thus $a = 0$

yields much better agreement for non-potential flows.

4.5.2 Quasi-steady models

In the limit of low frequency, a usual approximation is to consider that the fluid “adapts instantaneously to the solid movement”. Then, the fluid forces and moments can be derived from computations of *steady* forces and moments, using appropriate boundary conditions. This type of approximation is usually referred to as the *quasi-steady* hypothesis. Multiple variants of quasi-steady models have been proposed through the years [E. H. Dowell et al. 1989, §3.4.2.2]. We use two of them in the following.

It is expected that the quasi-steady hypothesis will have a restricted range of validity. However, the obvious ease of implementation of that analysis, which relies exclusively on steady measurements and/or theory, makes it a worthwhile first guess for flutter analysis. For that reason, we think it is useful comparing its results to the global stability approach.

Quasi-Steady Model 1 In this first quasi-steady model (QSM1) one considers that the fluid reacts on such a fast time scale to solid movement that it sees the solid as moving at a constant velocity. The instantaneous lift and moment are then equal to the *steady* lift and moment, taken at an effective angle of attack $\theta + \dot{h} + (1/4 - a)\dot{\theta}$. This effective angle of attack is composed of the “geometric” angle of attack θ , plus two contributions proportional to \dot{h} and $\dot{\theta}$. The \dot{h} contribution is easily justified by considering a heaving only motion. Indeed, in that case, the instantaneous solid velocity is simply $-1 \mathbf{e}_x + \dot{h} \mathbf{e}_y$. This configuration is identical to the one where the plate moves at velocity $-1 \mathbf{e}_x$ with incidence $\theta \simeq \dot{h}$, hence the \dot{h} contribution. Giving an interpretation for the term $(1/4 - a)\dot{\theta}$ in terms of effective angle of attack is less trivial because a constant angular velocity motion yields a different local velocity at each location on the surface of the plate. Thus it cannot be taken easily into account through a unique effective angle of attack. As a consequence, this second contribution is mainly seen here as an inheritance from Theodorsen’s theory. Further physical interpretation is discussed for example in [Hodges et al. 2011, §5.5] or [Païdoussis et al. 2011, §2.7.2]. Finally, it is supposed that the motion is of small amplitude, hence, only the linearized dependency of the steady lift and moment are kept, yielding QSM1:

$$\begin{aligned} C_L &= \left. \frac{dC_L^{\text{stat}}}{d\theta} \right|_{\theta_0} \left[\theta + \dot{h} + \left(\frac{1}{4} - a \right) \dot{\theta} \right] \\ C_M &= \left. \frac{dC_M^{\text{stat}}}{d\theta} \right|_{\theta_0} \left[\theta + \dot{h} + \left(\frac{1}{4} - a \right) \dot{\theta} \right] \end{aligned} \quad (4.9)$$

Quasi-Steady Model 2 The second quasi-steady model (QSM2) we use in this section is simply the limit of the Theodorsen model eq. (4.8) when the frequency is small

$\omega \rightarrow 0$. In that case, the Theodorsen function tends towards $C(k) \rightarrow 1$, yielding:

$$\begin{aligned}
 C_L &= \underbrace{-\frac{\pi}{2} [\dot{h} - a\ddot{\theta}]}_{\text{added mass}} - \underbrace{\frac{\pi}{2} \dot{\theta} + \frac{dC_L^{\text{stat}}}{d\theta} \Big|_{\theta_0} \left[\dot{h} + \left(\frac{1}{4} - a \right) \dot{\theta} \right]}_{\text{added damping}} + \underbrace{\frac{dC_L^{\text{stat}}}{d\theta} \Big|_{\theta_0} \theta}_{\text{added stiffness}} \\
 C_M &= \underbrace{\frac{\pi}{4} \left[a\dot{h} - \left(\frac{1}{16} + 2a^2 \right) \ddot{\theta} \right]}_{\text{added inertia}} + \underbrace{\frac{\pi}{4} \left(2a - \frac{1}{2} \right) \dot{\theta} + \frac{dC_M^{\text{stat}}}{d\theta} \Big|_{\theta_0} \left[\dot{h} + \left(\frac{1}{4} - a \right) \dot{\theta} \right]}_{\text{added damping}} + \underbrace{\frac{dC_M^{\text{stat}}}{d\theta} \Big|_{\theta_0} \theta}_{\text{added stiffness}}
 \end{aligned} \tag{4.10}$$

where it is now easy to identify *constant* (i.e. independent of frequency) added mass, damping and stiffness coefficients. QSM2 can be seen as an extension of QSM1 where an additional damping term and — more importantly — added mass effects have been added.

4.5.3 Numerical comparison of the flutter predictions

Before numerically comparing the Theodorsen and quasi-steady models to the global stability approach, we stress the fact that to calibrate these models, need to compute the slope of the lift and moment curves at zero incidence. This may of course be done by fitting linear functions to the lift and moment curves fig. 4.4.3. However, the process of building the *nonlinear* curve for each value of the Reynolds number and fitting a line to it is cumbersome and overkill (we only need the *linear* behavior). As an alternative, we rather use a linearized approach, detailed in section 4.A, that allows to extract directly the slopes of interest from the linearized response of the fluid to an increment of angle of attack.

In fig. 4.5.1 (a) and (b), we show how the slopes of the lift and moment curves at $\theta_0 = 0$ vary with the Reynolds number. In comparison, the theoretical values from thin airfoil theory are materialized by thin horizontal dashed lines. Wide variations of slopes are obtained: for example the lift slope goes from 3.2 at $Re = 10$ to more than 12 at $Re = 10^4$. The very large slopes observed for $Re > Re_{c,\text{wake}}$ (on the right of the vertical dashed line) are however to be taken cautiously, as we remind that they are computed from steady Navier–Stokes solutions, which are unstable at those high Reynolds numbers. In particular, the abrupt increase of the slopes around $Re = 7000$ is likely to be unphysical. Further investigations are needed to precisely assess that assumption.

We compare in fig. 4.5.1 (c) the (lower) critical dynamic pressure q_c^* for $10 \leq Re \leq 10^4$ and a large mass ratio $\tilde{m} = 10^4$ predicted by the three different flow models: full Navier–Stokes (solid line), Theodorsen (dotted red line) and QSM1 (dotted gray line) and QSM2 (dotted blue line). For completeness, the original (i.e. non-calibrated) Theodorsen prediction is shown with red horizontal dashed line. It is observed that the calibration with viscous slopes yields a good agreement between both the Theodorsen and quasi-steady models and the global stability results. The fact that the Theodorsen and quasi-steady models are very close to each other is expected, due to the low frequency of the flutter instability arising at such a high mass ratio ($\omega_c = \mathcal{O}(10^{-2})$). On the contrary, the non-calibrated Theodorsen theory is only valid for a particular Reynolds number, that happens to present steady slopes close to the thin airfoil theory predictions. Note that in the present case, it turns out that this Reynolds is very close to the critical Reynolds for the wake destabilization, $Re_{c,\text{wake}}$. Whether this is a coincidence or a more broadly shared feature remains an opened question.




FIGURE 4.5.1: Calibration of Theodorsen and quasi-steady models with the slopes of the steady $C_L^{\text{stat}}(\theta)$ and $C_M^{\text{stat}}(\theta)$ curves, evaluated at $\theta_0 = 0$. The slopes of the lift and moment coefficients are shown as a function of Re in (a) and (b) respectively. The critical dynamic pressure computed with full Navier–Stokes (solid line), Theodorsen (dotted red) and quasi-steady models (QSM1 in dotted gray and QSM2 in dotted blue) are presented in (c) for a large mass ratio $\tilde{m} = 10^4$. The horizontal dashed line shows the threshold predicted by original (non-calibrated) Theodorsen theory. The vertical dashed line materializes the critical Reynolds for the wake instability.

Having verified that all four models perform similarly in the case of a high \tilde{m} , low ω_c flutter, we now assess whether the simplified models accurately reproduce the particular shape of the flutter neutral curve in the (q^*, \tilde{m}) space, shown in section 4.4.1. In fig. 4.5.2, the neutral curves obtained for the four different models are presented, at $Re = 2900$. The Navier–Stokes curve (solid line), Theodorsen curve (dotted red) and quasi-steady curves (QSM1 in dotted gray and QSM2 in dotted blue) are presented. It is observed that, in the limit of large \tilde{m} , all three curves converge as expected. As \tilde{m} decreases, QSM1 and QSM2 quickly diverge from the Navier–Stokes and Theodorsen predictions, and both predict a lower threshold. In particular, QSM1 is over-conservative as q_c^* strongly decreases with \tilde{m} . The fact that QSM2 is significantly better than QSM1 shows that the added mass effects — that are the major difference between QSM1 and QSM2 — should not be neglected. The Theodorsen model remains close enough to the Navier–Stokes prediction down to $\tilde{m} \simeq 10^2$ whereas, for lower mass ratios, q_c^* is overestimated. Finally, we notice that Theodorsen and QSM2 predict an abrupt increase of q_c^* to infinity for $\tilde{m} \simeq 10$, whereas Navier–Stokes predicts the neutral curve should turn back. To shed some light onto that behavior, let us recall the results of fig. 4.4.4. It was shown there that as Re increases, the point where the neutral curve turns back to higher \tilde{m} is shifted to higher and higher q^* . If the Theodorsen model is now viewed as an asymptotic model for Navier–Stokes as $Re \rightarrow \infty$, then it should possess a turning point shifted to $q^* \rightarrow \infty$. In other words, the Theodorsen neutral curve should never turn back, which is exactly what we observe in fig. 4.5.2.




FIGURE 4.5.2: Comparison of the neutral curves in the (\tilde{m}, q^*) plane at $Re = 2900$ given by the Navier–Stokes (solid line), Theodorsen (dotted red) and quasi-steady (QSM1 in dotted gray and QSM2 in dotted blue) flow models.

4.6 Conclusion

In this chapter, we revisited the linear stability of a typical aeroelastic section with an incompressible Navier–Stokes flow model. The linear stability of this fluid-structure system was assessed by using a global fluid-structure stability approach, which consists in scrutinizing the leading eigenvalues of the linearized fluid-structure operator.

Four types of instabilities were observed through extensive parametric explorations. First, a classical coupled-mode flutter occurred at high-enough reduced velocity (or equivalently high dynamic pressure) due to the convergence of the heaving and pitching frequencies. For even higher velocity, a static divergence instability was found. For low reduced velocities, vortex induced vibrations instabilities occur, as the hydrodynamic wake mode interacts with the heaving (or pitching) mode. By varying the angle of attack of the steady solution, we showed how the flutter instability progressively transitions from a coupled-mode flutter involving both heaving and pitching motions to a single-mode flutter involving only the pitching mode. This transition is concomitant with the appearance of a large recirculation region on the plate suction side. Overall, the main interests of the global stability approach used in this work is its ability to investigate the stability of steady-state solutions that present nonlinear flow features, like recirculation regions, and thus to capture a variety of fluid-structure instabilities with one uniform flow model.

After having identified these different instabilities, a parametric exploration of the effect of mass ratio and Reynolds number was proposed, in the case of zero steady angle of attack. Neutral curves were computed in the (\tilde{m}, U^*) (or alternatively (\tilde{m}, q^*)) plane. For all mass ratios, vortex-induced vibrations occurred on a range of reduced velocities centered around two well-defined values that correspond to the coincidence of the pitching (resp. heaving) frequencies with the hydrodynamic wake instability frequency. Decreasing the mass ratio tends to widen the range of reduced velocity where vortex-induced vibrations occur. At a particular, low, mass ratio \tilde{m}^{II} (about $\tilde{m}^{II} \sim 10$ at $Re = 500$) a codimension-two point is found where both flutter and divergence occur at the same critical reduced velocity. For mass ratios higher than \tilde{m}^{II} , flutter occurs prior to divergence whereas divergence precedes flutter (if

it exists at all) for mass ratios below \tilde{m}^{II} . In addition, it was shown that for low mass ratios slightly below \tilde{m}^{II} , the flutter neutral curve “folds back” towards higher reduced velocities and mass ratios. For a fixed mass ratio, this implies that flutter restabilizes for (very) high reduced velocities. By varying the Reynolds number, it was shown that high Reynolds postpone this high- U^* restabilization of the flutter mode, leading us to attribute this unexpected behavior to viscosity effects. This trend is consistent with the potential Theodorsen model ($Re \rightarrow \infty$) that predicts a vertical asymptote of the flutter neutral curve for some mass ratio slightly below \tilde{m}^{II} .

The chapter was concluded by comparing the stability predictions obtained from the global Navier–Stokes stability approach to several variants of quasi-steady approaches and to the Theodorsen model. After a proper calibration of these models using the slopes of the viscous steady $C_L^{\text{stat}} - \theta$ and $C_M^{\text{stat}} - \theta$ curves, it was shown that all approaches provide accurate enough prediction of flutter thresholds for the low-frequency flutter, $\omega_c = \mathcal{O}(10^{-2})$, which typically occurs at high mass ratios. As mass ratio is decreased, the quasi-steady models become quickly irrelevant and significantly underestimate the flutter threshold. The Theodorsen model remains close enough to the Navier–Stokes results down to $\tilde{m} \simeq 10^2$. Contrary to quasi-steady approaches, Theodorsen tends to overestimate the threshold, which makes it a more accurate but unconservative stability criteria. Typically, at $\tilde{m} = 10^2$ the flutter threshold is overestimated by about 10% by Theodorsen with respect to Navier–Stokes. As a consequence, if accuracy is required in the evaluation of flutter thresholds, it is believed that the modelization of the unsteady viscous effects through the global stability approach is more adequate, in particular at low mass ratios close to \tilde{m}^{II} .

APPENDIX

Appendix 4.A Computation of the steady lift and moment slopes with Navier–Stokes model

The slopes $d\mathcal{C}_L^{\text{stat}}/d\theta$ and $d\mathcal{C}_M^{\text{stat}}/d\theta$ at $\theta_0 = 0$ are computed using the following linearized approach. Suppose a steady solution $\mathbf{q}_{f_0} = (\mathbf{U}_0, p_0, \Lambda_0)^T$ is known for the flow at incidence θ_0 . We recall that the aerodynamic coefficients can be computed directly from the interface stress variable (eq. (1.26)):

$$\begin{aligned}\mathcal{C}_L(\theta, \Lambda) &= \left\{ 2 \int_{\Gamma_{\text{fs}}} \Lambda d\mathbf{X} \right\} \cdot \mathbf{R}^T(\theta) \mathbf{e}_y \\ \mathcal{C}_M(\Lambda) &= \left\{ 2 \int_{\Gamma_{\text{fs}}} \mathbf{X} \wedge \Lambda d\mathbf{X} \right\} \cdot \mathbf{e}_z\end{aligned}$$

Differentiating the above lift coefficient with respect to θ , the slope of the lift coefficient $\mathcal{C}_L(\theta, \Lambda)$ at θ_0 writes:

$$\left. \frac{d\mathcal{C}_L^{\text{stat}}}{d\theta} \right|_{\theta_0} = \left. \frac{\partial \mathcal{C}_L^{\text{stat}}}{\partial \theta} \right|_{\theta_0, \Lambda_0} + \left. \frac{\partial \mathcal{C}_L^{\text{stat}}}{\partial \Lambda} \right|_{\theta_0, \Lambda_0} \frac{\partial \Lambda}{\partial \theta} \quad (4.11)$$

with

$$\begin{aligned}\left. \frac{\partial \mathcal{C}_L^{\text{stat}}}{\partial \theta} \right|_{\theta_0, \Lambda_0} \delta\theta &= \left\{ 2 \int_{\Gamma_{\text{fs}}} \Lambda_0 \right\} \cdot \left. \frac{\partial \mathbf{R}^T}{\partial \theta} \right|_{\theta_0} \delta\theta \mathbf{e}_y \\ \left. \frac{\partial \mathcal{C}_L^{\text{stat}}}{\partial \Lambda} \right|_{\theta_0, \Lambda_0} \delta\Lambda &= \left\{ 2 \int_{\Gamma_{\text{fs}}} \delta\Lambda \right\} \cdot \mathbf{R}^T(\theta_0) \mathbf{e}_y\end{aligned}$$

A similar derivation for the moment derivative leads to:

$$\left. \frac{d\mathcal{C}_M^{\text{stat}}}{d\theta} \right|_{\theta_0} = \underbrace{\left. \frac{\partial \mathcal{C}_M^{\text{stat}}}{\partial \theta} \right|_{\Lambda_0}}_{=0} + \left. \frac{\partial \mathcal{C}_M^{\text{stat}}}{\partial \Lambda} \right|_{\Lambda_0} \frac{\partial \Lambda}{\partial \theta} = 2 \int_{\Gamma_{\text{fs}}} \mathbf{X} \wedge \frac{\partial \Lambda}{\partial \theta} d\mathbf{X} \quad (4.12)$$

where the first term is zero because \mathcal{C}_M does not explicitly depend on θ .

The only unknown quantity in eq. (4.11) and eq. (4.12) is the interface stress increment $\partial\Lambda/\partial\theta$ due to a variation in θ . An equation for that quantity can be derived by differentiating the nonlinear steady-state fluid residual $\mathbf{R}_f = (\mathbf{R}^U, \mathbf{R}^p, \mathbf{R}^\Lambda)^T$ — composed of the fluid momentum equation, mass balance equation and the interface velocity continuity condition, cf section 1.2.2 for the definitions) — with respect to θ :

$$\left. \frac{d\mathbf{R}_f}{d\theta} \right|_{\theta_0} = \left. \frac{\partial \mathbf{R}_f}{\partial \mathbf{q}_f} \right|_{\theta_0} \frac{\partial \mathbf{q}_f}{\partial \theta} + \left. \frac{\partial \mathbf{R}_f}{\partial \mathbf{q}_s} \right|_{\theta_0} \frac{\partial \mathbf{q}_s}{\partial \theta} = \mathbf{0}$$

with $\mathbf{q}_f = (\mathbf{U}, p, \Lambda)^T$ the fluid variables and $\mathbf{q}_s = (h, u_h, \theta, u_\theta)^T$ the solid variables. In the above equation appear the linearization of the fluid residual with respect to fluid variables $\partial\mathbf{R}_f/\partial\mathbf{q}_f$ (*i.e.* the fluid-only Jacobian matrix) and the linearization

of the fluid residual with respect to the solid variables $\partial \mathbf{R}_f / \partial \mathbf{q}_s$. These operators have already been used and noted \mathbf{J}_{ff} and \mathbf{J}_{fs} , respectively, when deriving the linear stability eigenproblem in chapter 2 (more precisely in the appendix detailing their discrete version, section 2.B). Using the latter, the above equation rewrites:

$$\begin{aligned} \mathbf{J}_{ff} \frac{\partial \mathbf{q}_f}{\partial \theta} &= -\mathbf{J}_{fs} \frac{\partial \mathbf{q}_s}{\partial \theta} \\ &= -\mathbf{J}_{fs} (0, 0, 1, 0)^T \\ \mathbf{J}_{ff} \frac{\partial \mathbf{q}_f}{\partial \theta} &= \begin{pmatrix} -\mathbf{D}_\theta \\ \mathbf{0} \\ \mathbf{W}_\theta \end{pmatrix} \end{aligned}$$

where \mathbf{D}_θ and \mathbf{W}_θ respectively are the derivatives of the fluid momentum equation and interface velocity continuity condition, with respect to θ . This linear system can be solved for $\partial \mathbf{q}_f / \partial \theta$, which finally yields the slopes $dC_L^{\text{stat}} / d\theta$ and $dC_M^{\text{stat}} / d\theta$ through eq. (4.11) and eq. (4.12).

5

EFFECT OF FLUID NONLINEARITY ON THE FLUTTER BIFURCATION OF A TYPICAL SECTION

We investigate numerically the role of incompressible flow nonlinearities on the periodic flutter of a thin plate mounted on a system of bending/torsion linear springs located at its center of mass. The steady flow solution gets unstable to linear flutter eigenmodes at a critical reduced velocity. Close to that threshold, limit cycle oscillations of the plate appear for lower or higher reduced velocities, depending on the nature of the bifurcation. A weakly nonlinear analysis is first developed to compute the coefficients of the cubic amplitude equation that determine the subcritical or supercritical nature of the bifurcation. A parametric investigation of the effect of the solid-to-fluid mass ratio and Reynolds number shows that the bifurcation is supercritical (soft flutter) at very low Reynolds numbers $Re < 90$ independently of the mass ratio. At intermediate Reynolds numbers, $90 < Re < 2000$, it gets subcritical (hard flutter) on a range of low-to-moderate mass ratios. For larger values of the Reynolds number $2000 < Re$, the bifurcation is (almost) always subcritical independently of the mass ratio. The bifurcation scenarios are further investigated at the Reynolds number $Re = 500$ with the Time Spectral Method (TSM) allowing to compute accurately periodic solutions with large amplitude oscillations. The transition from a supercritical to a subcritical bifurcation when decreasing the mass ratio from high values is scrutinized, revealing a double-fold bifurcation scenario at intermediate mass ratio. In this scenario, the bifurcation is supercritical, as shown by the weakly nonlinear analysis, but a fold bifurcation of limit cycle solutions occurs slightly above the critical reduced velocity, leading to (unstable) limit cycle oscillations at lower reduced velocity. The second fold of periodic solutions finally leads to the branch of large amplitude oscillations that are observed in time-marching simulations. This double-fold scenario is discussed in light of experimental results by [Amandolese et al. \(2013\)](#).

5.1 Introduction

The flutter phenomenon is well understood as a linear instability in the sense that, above some critical velocity threshold, the coupled fluid-structure system possesses a pair of complex conjugate eigenvalues with a positive growth rate. These aspects were revisited in chapter 4 with a particular focus on the role of fluid viscosity.

However, if the governing equations are nonlinear — and they should be, because Nature is —, the unbounded exponential growth predicted by linear theory is saturated at some point, leading to finite amplitude oscillations *via* a Hopf bifurcation. Depending on the nature of the nonlinearity at play, the flutter bifurcation may be of a supercritical or subcritical nature (often denominated respectively as “soft flutter” and “hard flutter”). For a supercritical bifurcation, linear analysis usually provides a conservative stability criteria since the LCO amplitude grows smoothly from the bifurcation point. In the subcritical case however, large amplitude stable solutions exist below the linear threshold. Finite amplitude perturbations can then suddenly transition the system from a stable steady-state to a high-amplitude solution.

The occurrence of subcritical flutter was first reported by [Woolston et al. 1955; Woolston et al. 1957] using a nonlinear pitching spring with cubic stiffness and freeplay and an indicial function approach for the incompressible aerodynamics. In particular they showed that a soft torsion spring¹ leads to flutter oscillations well below the linear threshold. In contrast a hard spring preserves a supercritical flutter. Numerous investigations on the role of nonlinear spring stiffness have confirmed and enriched those findings (e.g. [B. H. K. Lee et al. 1999b] and references therein, or more recently [Padmanabhan et al. 2018]). The common hypothesis that a hard pitching spring leads to supercritical flutter was then nuanced by Collier et al. (2004) who exhibited some counter examples, using center manifold theory. In all above mentioned studies, the aerodynamics were modeled through variants of unsteady thin airfoil theory [Theodorsen 1935; Bisplinghoff et al. 1955] which is linear.

More recently, efforts to model nonlinear aerodynamics have been made. In the transonic regime, [Schewe et al. 2003] showed experimental evidence of strong fluid nonlinear effects that limit the amplitude of flutter oscillation. In particular, coexisting flutter LCO's were found for a supercritical² NLR 7310 airfoil. The coexistence of these solutions was later reproduced with Euler-based numerical simulations and attributed to different dynamical behaviors of shocks on the upper and lower surface of the (asymmetrical) NLR 7310 airfoil (see §7.2 in the review by Bendiksen (2011)). Numerical investigations by [Thomas et al. 2002] reported both super- and subcritical bifurcations on a transonic symmetrical NACA64A010A airfoil, using Harmonic Balance solutions of the Euler equation. In [Van Rooij et al. 2017a], parametric investigations of the flutter bifurcation of a supercritical NLR7301 airfoil in inviscid transonic flows were reported. In particular, the effect of the heaving-to-pitching natural frequency ratio was investigated. In [Thomas et al. 2004], evidence was provided on the role of viscosity in the transonic regime by comparing aeroelastic Harmonic Balance solutions of the Euler and RANS Spalart Allmaras models. In this work, viscosity tends to have a “hardening” effect, i.e. drives the bifurcation towards supercritical.

¹A *soft spring* (resp. *hard spring*) is a spring with decreasing (resp. increasing) stiffness as deformation increases

²Here, supercritical has nothing to do with the nature of flutter, but refers to these airfoils shapes that delay the apparition of supersonic regions on the suction side.

Turning now to the subsonic regime, nonlinear fluid effects on flutter have often been investigated in the context of airfoils pitching around a nonzero static angle of attack, close to the static stall angle, by either using reduced order models (*e.g.* the ONERA stall model in [D. Tang et al. 2002; Stanford et al. 2013]), experiments [Razak et al. 2011; Bhat et al. 2013] or numerical simulations [Menon et al. 2019]. Recently, the transitional Reynolds number regime has been the focus of several experimental and numerical works [Poirel et al. 2008; Poirel et al. 2010; Yuan et al. 2013; Barnes et al. 2018; Negi et al. 2018] that aimed at explaining the occurrence of amplitude-bounded aeroelastic oscillations — the so-called “laminar separation flutter” — that are particular to the transitional regime. For airfoils at zero static angle of attack undergoing classical coupled-mode flutter, nonlinear fluid effects were mostly investigated experimentally. In [Dimitriadis et al. 2009; Šidlof et al. 2016], the flutter of spring-mounted NACA0012 airfoil was investigated, whereas a similar study was conducted in [Amandolese et al. 2013], replacing the airfoil profile with a rigid flat plate. Both studies were performed at a moderate Reynolds, $Re \simeq 10^4$. Each time, hysteresis and subcritical LCO’s were reported, showing that not only transonic nonlinearities can saturate the flutter oscillation growth. In addition, Amandolese et al. (2013) showed the co-existence of low- and high-amplitude LCO solutions in a range of post-critical flow velocities. By ensuring that the experimental set-up yielded an “as linear as possible” solid model, they suggested that nonlinear aerodynamic effects are the origin of the observed behavior. However, the exact nature of the physical phenomenon involved remains an open question. Finally, we note that due to the mostly experimental nature of these investigations, the effect of the governing non-dimensional parameters is not easy to grasp, as they are often varied simultaneously. For example, increasing the wind velocity increases the reduced velocity and the Reynolds number together. Knowing that the Reynolds number typically governs the amount of nonlinearity in the fluid model, a separate evaluation of its effect would nevertheless be valuable. Similarly the effect of solid-to-fluid mass ratio is rarely investigated alone despite its known major effect on the critical flutter frequency (see chapter 4), and hence, on the amount of fluid unsteady effects.

In this work, we aim at isolating and understanding the role of incompressible fluid nonlinearities on the flutter bifurcation of a spring-mounted heaving and pitching plate. To that aim we consider a purely linear solid model immersed in a two-dimensional Navier–Stokes flow. The chapter is organized as follows. In section 5.2, the physical model and computational approaches are briefly described. In section 5.3, the effect of mass ratio and Reynolds number on the type of flutter is investigated, using a weakly nonlinear approach. In addition, we take advantage of the weakly nonlinear formalism to exhibit a scalar field that quantifies how much the different regions of the flow contribute to the type of bifurcation. Finally, in section 5.5, the fully nonlinear regime is visited where radically different flow features are observed for supercritical and subcritical cases. In the process, a particularly interesting scenario is discovered, where two stable solutions coexist for a given reduced velocity. This scenario enlightens the experimental results of Amandolese et al. (2013).

5.2 Governing equations and numerical methods

We investigate the non-linear dynamics of a rigid plate mounted on heaving and pitching springs and immersed in a two-dimensional incompressible viscous flow. As discussed in section section 1.1, eight non-dimensional parameters defined in

table 1.1, govern that fluid-structure interaction problem. As always in this thesis, five of them — the heaving-to-pitching frequency ratio Ω , the structural dampings ζ_h, ζ_p , the radius of gyration r_θ and the position of the elastic axis x_θ — are kept fixed to the values already specified in table 1.1. The parameters that vary in the present chapter are the solid-to-fluid mass ratio $\tilde{m} = m/(1/2\rho_f c^2)$, the Reynolds number $\mathcal{Re} = \rho_f U_\infty c/\mu$ and the reduced velocity $U^* = U_\infty/(c\sqrt{K_\theta/I_{ea}})$

The dynamics of the spring-mounted plate is governed by two (damped) linear oscillator equations eq. (1.3), whereas the flow verifies the incompressible Navier–Stokes equations. This system of coupled equations, described in detail in chapter 1, is formally written here as the following first-order in time evolution equation

$$\mathcal{M} \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}) = 0 \quad (5.1)$$

where the variable $\mathbf{q} = (h, \theta, u_h, u_\theta, \mathbf{U}, p, \mathbf{\Lambda})^T$ gathers all variables necessary to describe the fluid-solid interaction. The first four scalar variables allow describing the dynamics of the rigid plate. They are the heaving h and pitching θ displacements, as well as the corresponding velocities u_h and u_θ . The incompressible flow is described with the velocity \mathbf{U} and pressure p fields. Finally, the variable $\mathbf{\Lambda}$ defined at the fluid-solid interface represents the local stress exerted by the fluid onto the solid. The exact definition of the linear operator \mathcal{M} in front of the time-derivative and of the nonlinear residual vector \mathcal{R} are given in eq. (1.30). In this residual, the first four components corresponds to the two coupled damped harmonic oscillators, written as a first-order problem in time. The fifth and sixth component correspond to the Navier–Stokes momentum and mass conservation equations. The last equation corresponds to the equality of fluid and solid velocities at the fluid-solid interface. Again, we refer to chapter 1 for more details.

Several (numerical) methods are used in the following to compute the unsteady nonlinear solutions that emerge close to the critical reduced velocity U_c^* where steady solutions of the above equation ($\mathcal{R}(\mathbf{q}) = 0$) get linearly unstable (see chapter 4 for more detailed investigations on the linear stability of the system). These nonlinear unsteady solutions often consist in time-periodic solutions, also referred to as Limit Cycle Oscillations (LCO), that can be computed using either classical time-marching method (see section 5.2.1) or nonlinear modals methods specifically developed to compute periodic solutions (see section 5.2.2). In this work, the following modal methods are used: the weakly-nonlinear analysis (WNL) introduced in details in chapter 2 and the Time Spectral Method (TSM) presented in chapter 3. The weakly-nonlinear analysis is a semi-analytical method based on the assumption of weak nonlinear effects and thus tends to be valid close to the critical reduced velocity U_c^* . It is mainly used in the following to determine the supercritical or subcritical nature of the flutter bifurcation. On the other hand, the Time Spectral Method allows computing nonlinear periodic solutions independently of their amplitude (strength of the nonlinearity). In both cases, those modals methods give access to unstable LCO, which can not be determined with a time-marching method, unless stabilization techniques are implemented [Jallas et al. 2017; Shaabani-Ardali et al. 2019].

5.2.1 Time-marching simulations of nonlinear solutions

A Backward Differences Formula of order two (BDF2) is used to approximate the time derivative:

$$\left. \frac{\partial \mathbf{q}}{\partial t} \right| (t_n) \simeq \frac{3\mathbf{q}(t_n) - 4\mathbf{q}(t_{n-1}) + \mathbf{q}(t_{n-2})}{2\Delta t} \quad (5.2)$$

yielding the following implicit nonlinear system at each timestep

$$\mathcal{M} \mathbf{q}(t_n) + \frac{2\Delta t}{3} \mathcal{R}(\mathbf{q}(t_n)) = \frac{4}{3} \mathbf{q}(t_{n-1}) - \frac{1}{3} \mathbf{q}(t_{n-2})$$

The latter is solved using the pressure segregation method proposed in [Badia et al. 2007](#) which consists in mixing a pressure correction approach [[Guermond et al. 2006](#)] to handle the incompressibility constraint with Dirichlet–Neumann fixed-point iterations for handling the fluid–structure coupling. Solving the implicit nonlinear system is thus decomposed in solving a sequence of simpler linear problems that are: (i) a linear advection–diffusion equation for the fluid velocity, (ii) a Poisson problem for the pressure increment and (iii) a four-by-four³ linear (the solid dynamics are fully linear due to the elastic axis being at the center of mass) solid problem. The fluid problems are space-discretized using the well-known Taylor–Hood ($\mathbb{P}_2, \mathbb{P}_1$) finite element pair for (velocity, pressure) *via* the finite element library FreeFEM [[Hecht 2012](#)]. The discrete problems are solved in parallel with preconditioned Krylov subspace methods from the PETSc library [[Balay et al. 2019](#)], accessed through its FreeFEM interface. More details about the algorithm can be found in section [1.3](#).

5.2.2 Methods for periodic nonlinear solutions

We now describe the weakly-nonlinear analysis in [§5.2.2.1](#) and the time spectral method in [§5.2.2.2](#), that both allow to determine periodic nonlinear solutions. Compared to the time-marching simulations, these methods can be used to compute both stable and unstable Limit Cycle Oscillations.

5.2.2.1 Weakly nonlinear analysis

In the vicinity of the critical reduced velocity U_c^* , the amplitude of the LCO are expected to be small. That property can be used to approximate the periodic solution and thus to efficiently determine the supercritical or subcritical nature of the bifurcation. As detailed in chapter [2](#), the weakly nonlinear approximation is based on a development of the solution in the vicinity of the critical velocity. The departure from the critical velocity is written

$$\frac{1}{U^{*2}} = \frac{1}{U_c^{*2}} + \epsilon^2 \Delta_{U^*} \quad (5.3)$$

with ϵ is a small positive parameter and the integer $\Delta_{U^*} = \pm 1$ is introduced to choose the departure side. According to the methods of multiple scales [[Nayfeh et al. 1995](#)], the solution \mathbf{q} is then searched in the form of a series of ϵ powers:

$$\mathbf{q} = \mathbf{q}_0 + \underbrace{\epsilon \left(A_\epsilon \hat{\mathbf{q}}_c e^{i\omega_c t} + \text{c.c.} \right)}_{\mathbf{q}_1} + \epsilon^2 \mathbf{q}_2 + \epsilon^3 \mathbf{q}_3 + \dots \quad (5.4)$$

³Recall the heaving and pitching velocities have been introduced to formulate the problem at first order in time...

where \mathbf{q}_0 is solution of the steady-state equation eq. (2.5) at the critical velocity U_c^* and the first-order solution \mathbf{q}_1 is proportional to the marginally stable eigenmode $\widehat{\mathbf{q}}_c$ oscillating at frequency ω_c . The small but finite-size complex amplitude $A = \epsilon A_\epsilon$ of this critical eigenmode is solution of the well-known Stuart–Landau equation:

$$\frac{dA}{dt} = \alpha \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right) A + \beta |A|^2 A \quad (5.5)$$

where the complex coefficient α drives the linear term while the complex coefficient β drives the cubic nonlinear term. As explained in chapter 2, this amplitude equation results from a compatibility condition obtained at third-order in (5.4), which gives the following definition of the linear α and nonlinear β coefficients,

$$\alpha = -\frac{1}{U_c^{*2}} \frac{\langle \mathbf{q}_c^\dagger, \mathcal{L} \widehat{\mathbf{q}}_c \rangle}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \widehat{\mathbf{q}}_c \rangle} \quad \text{and} \quad \beta = \frac{\langle \mathbf{q}_c^\dagger, \mathcal{F}_{\text{res}}(\widehat{\mathbf{q}}_c, \widehat{\mathbf{q}}_2^{A^2}, \widehat{\mathbf{q}}_2^{A^2}) \rangle}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \widehat{\mathbf{q}}_c \rangle}. \quad (5.6)$$

The linear coefficient α only depends on the critical eigenmode \mathbf{q}_c while the nonlinear coefficient β also depends on the second-order solution \mathbf{q}_2 . Both coefficients depends on the critical adjoint eigenmode \mathbf{q}_c^\dagger . The exact definition of the scalar product $\langle \cdot, \cdot \rangle$, linear operators \mathcal{L} and \mathcal{M} and residual vector \mathcal{F}_{res} is given in section 2.3.

Once these coefficients are determined, the dynamics of the Limit Cycle Oscillations is obtained by injecting the polar decomposition of the amplitude $A = R e^{i\omega_{\text{nl}} t + \phi}$ in (5.5), yielding first that the real amplitude R satisfies

$$\frac{dR}{dt} = \Re(\alpha) \left(\frac{U_c^{*2}}{U^{*2}} - 1 \right) R + \Re(\beta) R^3 \quad (5.7)$$

Searching for periodic solutions, we have $dR/dt = 0$, such that eq. (5.7) yields either the trivial solution $R = 0$ (which gives the base flow solution \mathbf{q}_0) or the non-zero amplitude solution

$$R^2 = \frac{\Re(\alpha)}{\Re(\beta)} \left(1 - \frac{U_c^{*2}}{U^{*2}} \right) \quad (5.8)$$

The existence of the LCO thus clearly depends on the sign of $\Re(\alpha)/\Re(\beta)$. If $\Re(\alpha)/\Re(\beta) > 0$ (resp. $\Re(\alpha)/\Re(\beta) < 0$), then the LCO exists for flutter velocities above (resp. below) the critical value. To determine the supercritical or subcritical nature of the bifurcation, we also need to know the stability of the trivial solution $R = 0$ on both sides of the critical value. This is given by the sign of the linear term in (5.7). In the following, we will always consider the case where $\Re(\alpha) < 0$, so that the trivial solution gets unstable for $U^* > U_c^*$. Summarizing in fig. 5.2.1, the bifurcation is then supercritical if $\Re(\beta) < 0$ while it is subcritical if $\Re(\beta) > 0$.

The frequency of the LCO $\omega = \omega_c + \omega_{\text{nl}}$ is the sum of the linear critical frequency ω_c and the non-linear correction ω_{nl} given by

$$\omega_{\text{nl}} = \left(\Im(\beta) \frac{\Re(\alpha)}{\Re(\beta)} - \Im(\alpha) \right) \left(1 - \frac{U_c^{*2}}{U^{*2}} \right). \quad (5.9)$$

Knowing this frequency and the real amplitude R from (5.8), we can reconstruct analytically the first-order approximation for any components of the LCO. In particular,




FIGURE 5.2.1: Nature of the bifurcation depending on the signs of $\Re(\beta)$ and given $\Re(\alpha) < 0$. The amplitude is represented as a function of U^* . The central horizontal branch always represents the trivial $R = 0$ solution that becomes unstable for $U^* > U_c^*$. By convention, the solid (resp. dashed) lines represent stable (resp. unstable) solutions.

the temporal evolution of the pitching angle is given by

$$\theta(t) = 2\sqrt{\frac{1}{\tilde{\beta}} \left(1 - \frac{U_c^{*2}}{U^{*2}}\right)} \cos[(\omega_c + \omega_{nl})t + \phi + \phi_{\theta_c}] \quad (5.10)$$

where we have introduced the normalized quantity $\tilde{\beta}$

$$\tilde{\beta} = \frac{\Re(\beta)}{\Re(\alpha)|\hat{\theta}_c|^2} \quad (5.11)$$

From eq. (5.10), the amplitude of the pitching motion is seen to depend only on $\tilde{\beta}$ and the deviation to the threshold given by the factor $1 - (U_c^*/U^*)^2$. As a consequence $\tilde{\beta}$ — its inverse, exactly — quantifies how fast the amplitude of the pitching oscillations grows when moving away from U_c^* . More precisely, a large positive $\tilde{\beta}$ corresponds to a supercritical bifurcation with slowly increasing amplitude whereas a very negative $\tilde{\beta}$ indicates a strongly subcritical bifurcation. When $\tilde{\beta} = 0$, the cubic nonlinear term in the amplitude equation (5.5) vanishes. This means that the nonlinearities at stake order ϵ^3 are not strong enough to saturate the linear growth of the flutter instability. In such cases, the weakly nonlinear analysis must typically be pushed at higher order. We chose not to pursue this path and simply note that results of the weakly nonlinear analysis around $\tilde{\beta} = 0$ should be interpreted with caution. To confirm them, we will rather consider the strongly nonlinear analysis developed below, in section 5.2.2.2.

Identifying the fluid and geometric nonlinearities The nature of the bifurcation is entirely driven by the scalar parameter $\tilde{\beta}$, defined in eq. (5.11). We now propose a decomposition of this parameter into two contributions that quantify the fluid and geometric nonlinearity.

Coming back to the definition (5.6) of the nonlinear coefficient β , the resonant forcing \mathcal{F}_{res} (exactly defined in eq. (2.16)) is formally written as

$$\mathcal{F}_{\text{res}}(\hat{\mathbf{q}}_c, \hat{\mathbf{q}}_2^{|\Lambda|^2}, \hat{\mathbf{q}}_2^{A^2}) = (0, 0, \mathcal{F}_{\text{res}}^{u_h}, 0, \mathcal{F}_{\text{res}}^U, 0, \mathcal{F}_{\text{res}}^\Lambda)^T \quad (5.12)$$

to highlight that it has non-zero components only for the heaving momentum equation (third component), the Navier–Stokes momentum equation (fifth) and the velocity continuity equation at the fluid–solid interface (seventh). The terms $\mathcal{F}_{\text{res}}^{u_h}$ and $\mathcal{F}_{\text{res}}^\Lambda$ are due to the trigonometric nonlinearities inherent to the rotation movement of the plate. Indeed, in the heaving equation, the lift obtained in the rotating frame ($\mathbf{e}_x, \mathbf{e}_y$) must be projected back onto the laboratory vertical axis \mathbf{e}_y (cf eq. (1.21a)). In the interface boundary condition, it is the vertical velocity of the plate that must be projected back in the rotating axis (cf eq. (1.20)). These terms thus correspond to the so-called *geometric nonlinearities*. On the contrary, the term $\mathcal{F}_{\text{res}}^U$ is generated by *fluid nonlinearities* of the Navier–Stokes momentum equation which are the convection term, including the modification of the convective velocity by the domain’s motion. Based on that decomposition of the resonant forcing, and recalling that the adjoint mode also decomposes as $\mathbf{q}^\dagger = (h^\dagger, \theta^\dagger, u_h^\dagger, u_\theta^\dagger, \mathbf{U}^\dagger, p^\dagger, \Lambda^\dagger)^\top$, we can thus develop the nonlinear coefficient β as the sum

$$\beta = \beta_f + \beta_g$$

of the fluid β_f and geometric β_g contributions defined as

$$\beta_f = \int_{\Omega_f} f_{\beta_f}(\mathbf{X}) d\Omega \quad \text{and} \quad \beta_g = \frac{\mathcal{F}_{\text{res}}^{u_h} u_h^{\dagger*}}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \hat{\mathbf{q}}_c \rangle} + \int_{\Gamma_{fs}} \frac{\mathcal{F}_{\text{res}}^\Lambda \cdot \Lambda^{\dagger*}}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \hat{\mathbf{q}}_c \rangle} d\Gamma.$$

The scalar field $f_{\beta_f}(\mathbf{X})$ in the definition of the fluid contribution is defined as

$$f_{\beta_f}(\mathbf{X}) = \frac{\mathbf{U}^{\dagger*} \cdot \mathcal{F}_{\text{res}}^U(\hat{\mathbf{q}}_c, \mathbf{q}_2^{|A|^2}, \mathbf{q}_2^{A^2})}{\langle \mathbf{q}_c^\dagger, \mathcal{M} \hat{\mathbf{q}}_c \rangle}$$

and will be used to determine flow regions where the nonlinearity is at play. This quantity can be assimilated to different local quantities previously introduced in the literature to assess the regions of the flow that are most decisive to linear instabilities. The so-called *wavemaker* [Huerre et al. 1990; Gianetti et al. 2007], for example, writes as the product of the local norms of the adjoint and direct modes. It quantifies the drift of an eigenvalue of the linear stability problem to a spatially localized feedback force. Closer to our expression of f_{β_f} , Marquet et al. (2015b) have introduced a way of decomposing an eigenvalue of the (fluid) stability problem as the integral of a local quantity, called *endogeneity*, that writes $\mathbf{U}^{\dagger*} \cdot (\mathcal{J}^U \hat{\mathbf{U}})$ with \mathcal{J}^U the Jacobian of the fluid problem (formulated here only with velocity, for simplicity). Both these approaches have been applied to the vortex shedding instability behind a circular cylinder in (see [Paladini et al. 2019] for a direct comparison) and evidenced the sensitive regions of the flow. In our study, it is the fluid contribution to the nonlinear coefficient β_f that is decomposed as the integral of a local quantity, f_{β_f} . The latter may thus be seen as a (weakly) nonlinear extension of the endogeneity concept, where the resonant forcing is projected onto the adjoint mode instead of the linear operator.

Obviously, the decompositions proposed above hold also for the normalized quantity $\tilde{\beta}$ that we use in the following instead of β , for reasons explained before. Specifically,

$$\tilde{\beta} = \tilde{\beta}_f + \tilde{\beta}_g, \quad \tilde{\beta}_f = \frac{\Re(\beta_f)}{\Re(\alpha)|\theta_c|^2} \quad \tilde{\beta}_g = \frac{\Re(\beta_g)}{\Re(\alpha)|\theta_c|^2} \quad (5.13)$$

and

$$\tilde{\beta}_f = \int_{\Omega_f} \tilde{f}_{\beta_f}(\mathbf{x}), \quad \tilde{f}_{\beta_f} = \frac{\Re(f_{\beta_f})}{\Re(\alpha)|\theta_c|^2} \quad (5.14)$$

5.2.2.2 Time Spectral Method

To investigate LCOs with large amplitudes, and in particular the hard flutter occurring via a subcritical bifurcation, we consider the Time Spectral Method (TSM), belonging to the family of harmonic balance methods, that is designed to compute periodic solutions. A short explanation is provided here, and the reader is referred to chapter 3 for a detailed derivation.

In the Time Spectral Method, a spectral approximation of the time derivative in eq. (5.1) is considered, i.e.

$$\left. \frac{\partial \mathbf{q}}{\partial t} \right| (t_n) \simeq \omega \sum_{k=0}^{2N_h} d_k \mathbf{q}(t_{n+k}) \quad (5.15)$$

where $\omega = 2\pi/T$ is the (unknown) frequency of the T -periodic solution and the period is discretized by the grid composed of the $2N_h + 1$ instants $t_n = n/(2N_h + 1)T$, $n = 0, \dots, 2N_h$. The coefficients d_k are given by

$$d_k = \begin{cases} \frac{1}{2}(-1)^{k+1} \csc\left(\frac{\pi k}{2N_h + 1}\right) & \text{if } k \neq 0 \\ 0 & \text{if } k = 0 \end{cases} \quad (5.16)$$

Note that the index $n + k$ must be understood modulo $2N_h + 1$ due to periodicity. Compared to the second-order BDF formula eq. (5.2) used in time-marching simulations that depends only on three consecutive time instants, the above time-spectral derivative depends on all $2N_h + 1$ instants. Using eq. (5.15), the time-dependent governing equations can be recast into a system of $2N_h + 1$ coupled time-independent equations that writes:

$$\omega \mathbf{D} \mathbf{Q} + \mathbf{R}(\mathbf{Q}) = 0 \quad (5.17)$$

where the solution and residual vector are

$$\mathbf{Q} = \begin{pmatrix} \mathbf{q}(t_0) \\ \vdots \\ \mathbf{q}(t_{2N_h}) \end{pmatrix} \text{ and } \mathbf{R}(\mathbf{Q}) = \begin{pmatrix} \mathcal{R}(\mathbf{q}(t_0)) \\ \vdots \\ \mathcal{R}(\mathbf{q}(t_{2N_h})) \end{pmatrix} \quad (5.18)$$

The time-spectral derivation matrix \mathbf{D} is defined (cf eq. (3.25)) from the coefficients (5.16). As the LCO frequency ω is unknown, one additional scalar equation is required to close the system. This equation may be arbitrary chosen and we here impose that the pitching velocity is zero, $u_\theta = 0$, at instant $t = 0$. The time-spectral solutions thus satisfy the system of equation (6.9) augmented with one scalar constraint. Its total size is $(2N_h + 1) \times K + 1$ where K is the number of degrees of freedom used for the spatial discretization of \mathbf{q} . To obtain accurate solutions, the number of harmonics N_h (equivalently, the number of instants $2N_h + 1$) and the number of spatial degrees of freedom K may quickly grow. Efficient numerical methods are then required to obtain time-spectral solutions with adequate discretizations. To address this challenging task, a solver was proposed in chapter section 3.3 based on the combination of a Newton–Krylov strategy with the so called “block-circulant”

preconditioner. This numerical method has been used to obtain results shown in the next section.

5.2.3 Cross-validation of numerical methods

The implementation of the weakly-nonlinear and time-spectral methods are validated by comparison with results of the time-marching simulations performed at the mass ratio $\tilde{m} = 1000$ and the Reynolds number $Re = 500$.

The frequency and the maximal pitching angle are depicted in Figure fig. 5.2.2a and fig. 5.2.2b, respectively, as a function of the reduced velocity. The reference results, corresponding to time-marching simulations performed with about 1400 timesteps per period, are reported with black dots in the figures. Results of the weakly-nonlinear analysis, shown with the grey curve, matches fairly well the time-marching results especially close to the critical reduced velocity. For the maximal value of the pitching angle (Figure fig. 5.2.2b), the disagreement clearly increases with the reduced velocity. Such a disagreement is expected and was already reported, for instance by [Gallaire et al. 2016] for the vortex shedding behind a circular cylinder. Interestingly, the frequency prediction is much better. Turning now to results of the time-spectral method with $N_{\eta} = 20$ (squares), they compare very well with time-marching results, for all values of the reduced velocity considered here. The critical flutter velocity U_c^* corresponds to that predicted by the linear stability analysis. In addition, the periodic evolution of the lift coefficient (Figure fig. 5.2.2c) and pitching amplitude (Figure fig. 5.2.2d) obtained with the time-spectral method and time-marching simulation also perfectly match. Similar agreements were obtained for other sets of parameters (\tilde{m}, Re).

5.3 Hard and soft flutter with weakly nonlinear analysis

The first objective of this section is to understand the effect of the mass ratio and Reynolds number on the sub- or supercritical nature of the flutter bifurcation, using the weakly nonlinear analysis (§5.3.1). Results of the decomposition approach, previously introduced in eq. (5.13), are then presented to quantify the distinct effect of the hydrodynamic and geometric nonlinearities on the bifurcation (§5.3.2).

Before addressing those objectives, we start by presenting a typical coupled-mode flutter LCO, obtained at $Re = 500$, $\tilde{m} = 1000$ and $U^* = 1.016U_c^*$. The solution is reconstructed in time (at second order) using eq. (5.4). In fig. 5.3.1a, the time signals for the heaving displacement (solid line) and pitching angle (dashed line) are purely sinusoidal. This is consistent with the signals obtained with TSM and time-marching computations using similar parameters in fig. 5.2.2. The heaving and pitching motion oscillate with a phase shift such that the instantaneous angle of attack $-\theta$ is “in advance” with respect to the heaving motion — here by slightly less than half a period. This phase shift is a typical feature of the coupled-mode flutter instability, as recalled in the introduction to this manuscript (*cf* fig. 2). The lift (solid line) and moment (dashed line) coefficient signals, presented in the central part of fig. 5.3.1a, vibrate almost in phase with respect to each other, and with respect to the pitching angle shown above. Over one period of oscillation, this motion allows the extraction of a small, but positive, mean power from the flow to the solid (equal to the power dissipated by the structural damping). In fig. 5.3.1b, four snapshots of the vorticity field are presented. The flow features fully attached shear layers that evolve *quasi-steadily* in time, consistently with the low frequency $\omega \simeq 0.17$ of the flutter instability at the chosen set of parameters. One notice that in these snapshots,




FIGURE 5.2.2: Comparison of numerical results obtained with time-marching simulations and modal methods for $\bar{m} = 1000$ and $\mathcal{R}e = 500$. The frequency (a) and the maximal pitching angle (b) of the LCO are presented as a function of reduced velocity for the weakly nonlinear analysis (gray solid line), the time spectral method ($N_h = 20$, squares) and the time-marching simulations (1400 timesteps per period, black dots). The vertical dashed line indicates the critical flutter velocity. Periodic evolution of (c) the lift coefficient and (d) the pitching angle are depicted for the time-spectral method and time-marching simulations (using the same symbols) for the reduced velocity $U^* = 5.22$.




FIGURE 5.3.1: Weakly nonlinear flutter LCO at $Re = 500$, $\tilde{m} = 1000$ and $U^*/U_c^* = 1.016$. In (a) we present the time signals of the heaving displacement (solid) and pitching angle (dashed) (top part), the lift (solid) and moment (dashed) coefficients (central part) and the power transmitted by the fluid forces to the solid (bottom part). In (b), vorticity snapshots corresponding to four instants in the period are shown.

the wake seems to almost “rotate as a block” with the plate. Close to the plate, this is indeed the correct physical flow (as will be seen later in fig. 5.5.2). Far from the plate, this block-rotation becomes increasingly unphysical and is entirely due to the *weakly nonlinear* nature of the method. This is not problematic however since the nature of the flutter bifurcation, that we explore in this section, is an asymptotic feature that is fully determined by the cubic coefficient of the normal form, eq. (5.5).

5.3.1 Effects of mass ratio and Reynolds number on the bifurcation




FIGURE 5.3.2: Nature of the flutter bifurcation at different mass ratio and constant $\mathcal{R}e = 500$. The value of the criticality indicator $\tilde{\beta}$, built from weakly nonlinear analysis, is shown as a function of mass ratio. In the insert the amplitude of θ is represented as a function of U^*/U_c^* for the different mass ratio marked with solid symbols on the main curve ($\tilde{m} = 20, 30, 60, 80, 120, 200$). Note that we do not investigate the range $\tilde{m} < 12$ materialized by the grey area, since for those low mass ratio a static divergence mode is already unstable when the flutter threshold U_c^* is reached (cf fig. 4.4.1). This scenario is out of the validity assumptions of the present weakly nonlinear analysis.

We first investigate the effect of the solid-to-fluid mass ratio on the bifurcation for the fixed Reynolds number $\mathcal{R}e = 500$. For each mass ratio, the critical reduced velocity for flutter is first found. Then, the non-dimensional quantity $\tilde{\beta}$ that governs the supercritical or subcritical nature of the bifurcation, and defined in eq. (5.11), is computed. The evolution of $\tilde{\beta}$ as a function of the mass ratio \tilde{m} is shown in fig. 5.3.2. Three regions are observed in the figure. For high mass ratio $\tilde{m} > \tilde{m}_u$, the bifurcation is supercritical ($\tilde{\beta} > 0$). For intermediate values in the range $\tilde{m}_1 < \tilde{m} < \tilde{m}_u$, the bifurcation is subcritical ($\tilde{\beta} < 0$). And finally, for lower mass ratio $\tilde{m} < \tilde{m}_1$, the bifurcation gets back to supercriticality. The transition values are $\tilde{m}_1 = 13.5$ and $\tilde{m}_u = 70$ for the Reynolds number $\mathcal{R}e = 500$ considered here. We recall that those values should be considered with caution, since they correspond to $\tilde{\beta} = 0$ for which the cubic amplitude equation (5.5) degenerates, as explained previously. To account for the nonlinear effects in that case, the quintic amplitude equation should be considered but this is out of the scope of the present paper. Instead, the effect of higher-order nonlinearities close to \tilde{m}_u will be investigated in section 5.5 using time-spectral method or time-marching simulations. We note also that the supercritical region at low mass ratio is of very small extension since it is bounded on its lower end by a region (grey in the figure) where divergence instability occurs for reduced velocity lower than the critical flutter velocity.

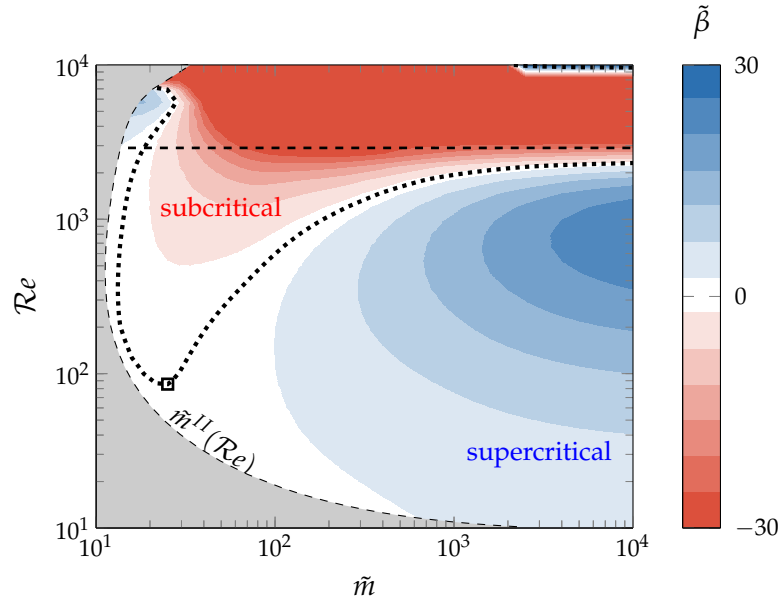


FIGURE 5.3.3: Effect of mass ratio and Reynolds number on the type of the flutter bifurcation. The colormap represents the criticality indicator $\tilde{\beta}$ in the (\tilde{m}, Re) plane, whereas the dotted line indicates the contour $\tilde{\beta} = 0$ which separates subcritical and supercritical cases. The gray area on the left represents mass ratios $\tilde{m} < \tilde{m}^{II}(Re)$, for which a divergence is the primary instability (see fig. 4.4.1). The upper horizontal dashed line materializes the critical Reynolds $Re_{c,wake}$ above which the wake of the fixed plate becomes unstable.

Let us now focus on the inset of fig. 5.3.2 which shows the bifurcation diagrams corresponding to several values of \tilde{m} . Since the critical reduced velocity varies with the mass ratio, we show the amplitude of θ as a function of the ratio between the reduced velocity and its critical value. The color of the different curves corresponds to the color of the solid circles in the main figure. As expected, the LCO branches exists above (resp. below) the threshold for supercritical (resp. subcritical) bifurcations displayed with orange (blue) colors. If the sign of $\tilde{\beta}$ gives the nature of the bifurcation, its absolute value $|\tilde{\beta}|$ indicates the strength of the nonlinearity and consequently the growth of the LCO amplitude when moving away from the threshold $U^* = U_c^*$. This is clearly observed in fig. 5.3.2 by comparing the growth of the pitching amplitude when moving away from the threshold (inset) with the corresponding values of $\tilde{\beta}$ (main figure). For large values of $|\tilde{\beta}|$, the maximal pitching amplitude of the LCO grows slowly. The nonlinear effects are strong compared to the linear growth and the saturation occurs for small amplitudes. On the other hand, for small values of $|\tilde{\beta}|$, the maximal pitching amplitude of the LCO grows abruptly close to the threshold. The nonlinearity being weak, the saturation occurs for larger amplitudes

By additionally varying the Reynolds number, we determine the transition between supercritical and subcritical bifurcations in the parameter space (\tilde{m}, Re) . Figure 5.3.3 displays the isocontours of $\tilde{\beta}$ in that parameter space, with positive values (in blue) indicating a supercritical bifurcation and negative values (in red) indicating a subcritical bifurcation. The black dotted curve highlights the transition between supercritical and subcritical regions. Supercritical bifurcations occurs predominantly for large mass ratio and low Reynolds number. For very low Reynolds

number (for instance $Re = 20$), decreasing the mass ratio rapidly leads to the divergence instability occurring prior to the flutter instability. This region is shaded in grey and the dashed curve $\tilde{m}^{II}(Re)$ corresponds to codimension-two points where the flutter and divergence modes are both neutrally stable at the same reduced velocity. For Reynolds number above $Re \sim 90$ (square symbol) and below $Re \sim 2000$, the bifurcation changes from supercritical to subcritical when decreasing the mass ratio. This is the scenario previously described in fig. 5.3.2 ($Re = 500$) around $\tilde{m} = \tilde{m}_u$. For very low mass ratios, the opposite transition from subcritical to supercritical again is also retrieved ($\tilde{m} = \tilde{m}_l$). The low mass ratio supercritical region is very narrow as we quickly reach the codimension-two curve. When increasing the Reynolds number, one notices that the lower transition mass ratio \tilde{m}_l only slightly varies. On the contrary, the upper transition mass ratio \tilde{m}_u strongly increases. As a consequence, for Reynolds numbers higher than $Re \sim 2000$, the bifurcation is subcritical for all mass ratios reported in the figure (and above \tilde{m}_l). Interestingly, the upper transition mass ratio \tilde{m}_u seems to asymptotically tend to infinity when Re approaches some value just below the critical Reynolds for the vortex-shedding instability of the fixed plate, $Re_{c,wake} \simeq 2925$ (horizontal dashed line). For Reynolds numbers above $Re_{c,wake}$ (horizontal dashed line), the results of the weakly nonlinear analysis should be interpreted cautiously since the vortex-shedding mode is unstable in addition to the flutter mode, thus requiring in theory a two-mode weakly nonlinear analysis [Meliga et al. 2012]. Still, we observe an abrupt transition from subcritical to supercritical for $Re \simeq 10^4$ and large mass ratio. This point is further discussed in section 5.4.

5.3.2 Fluid and geometric nonlinearities

The model considered for the elastic motion of the rigid plate is purely linear. Still, the rotational motion of the plate induces a geometric nonlinearity in addition to the inherent fluid nonlinearity. The weakly-nonlinear analysis easily allows identifying those two contributions, as explained in §5.2.2.1 with the decomposition (5.13).

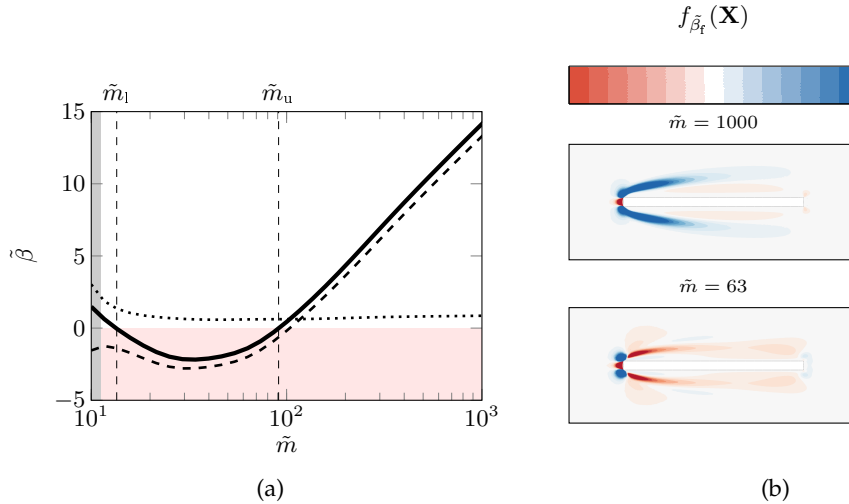


FIGURE 5.3.4: Fluid and geometric nonlinearities. (a) Fluid $\tilde{\beta}_f$ (dashed line) and geometric $\tilde{\beta}_g$ (dotted line) contribution to the nonlinear coefficient $\tilde{\beta}$ (solid line) as a function of the mass ratio, (b) Spatial distribution of $\tilde{f}_{\tilde{\beta}_f}(\mathbf{x})$ (its integral over the domain gives the value of the fluid contribution $\tilde{\beta}_f$) for two values of the mass ratio corresponding to a supercritical bifurcation (top, $\tilde{m} = 1000$) and a subcritical bifurcation (bottom, $\tilde{m} = 63$). The Reynolds number is $Re = 540$.

Figure 5.3.4a displays the evolution of the fluid $\tilde{\beta}_f$ and geometric $\tilde{\beta}_g$ contributions to the nonlinear coefficient $\tilde{\beta}$ as a function of the mass ratio for the fixed Reynolds number $Re = 540$. In that case, we recall that the bifurcation is supercritical for large ($\tilde{m} > \tilde{m}_u$) and low ($\tilde{m} < \tilde{m}_1$) values of the mass ratio, while it is subcritical for intermediate values ($\tilde{m}_1 < \tilde{m} < \tilde{m}_u$). For all values of the mass ratio, the contribution of the geometric nonlinearity (dotted curve) is positive and, in most cases, it is small compared to the fluid contribution (dashed line), that evolves very similarly to the total contribution (solid line). In particular, the upper transition at \tilde{m}_u from a super- to sub-critical bifurcation is clearly driven by the fluid nonlinearity. On the other hand, around the lower transition at \tilde{m}_1 , the geometric and fluid nonlinearity have opposite contributions of equal strength. For small enough values of the mass ratio, the total contribution is even positive. Thus, the transition from sub-critical to super-critical around \tilde{m}_1 is driven by the geometric nonlinearity.

In cases where the bifurcation is driven by the fluid nonlinearity, we can determine the flow regions that mostly contribute by examining the spatial distribution of the scalar field $\tilde{f}_{\tilde{\beta}_f}(\mathbf{X})$, defined in (5.14). Its integral over space gives the fluid contribution $\tilde{\beta}_f$. Figure 5.3.4b displays $\tilde{f}_{\tilde{\beta}_f}$ for two mass ratios corresponding to a supercritical bifurcation (top figure, $\tilde{m} = 1000$) and a subcritical bifurcation (bottom figure, $\tilde{m} = 63$). Largest values of $\tilde{f}_{\tilde{\beta}_f}$ are obtained close to the plate, and more specifically close to the leading-edge. For the largest mass ratio (top), the symmetric shear layers emerging from the leading-edge corners strongly contribute to the supercritical nature of the bifurcation. For the lowest mass ratio (bottom), this effect is (almost) entirely opposite and the bifurcation gets subcritical. The shear-layers around the leading-edge are thus the flow region responsible for the upper transition around \tilde{m}_u .




FIGURE 5.4.1: Instantaneous vorticity field around the *fixed plate* for (a) a steady solution at $Re = 2682$ and unsteady solutions at (b) $Re = 8325$ and (c) $Re = 10^4$.

5.4 A mean flow approach for the bifurcation analysis at high Reynolds number and mass ratio

For the highest mass ratios, $\tilde{m} \simeq 10^4$, the weakly nonlinear analysis predicts a brutal transition from subcritical to supercritical for large Reynolds numbers, $Re \simeq 10^4$ (see fig. 5.3.3). After discussing the physical relevance of the linear and weakly nonlinear flutter analysis for such high Reynolds numbers, an alternative mean flow-based approach is proposed and applied.

First, we recall that linear stability analysis predicts the destabilisation of a vortex-shedding mode for Reynolds number above $Re_{c,wake} \simeq 2925$ (cf horizontal dashed line in fig. 5.3.3), with a frequency $\omega_{wake} \simeq 11.2$ much higher than the flutter frequency. Therefore, the weakly nonlinear analysis developed around the base flow is questionable as the latter may not represent the physical flow seen by the plate below the critical flutter velocity. To better illustrate that discrepancy, time-marching simulations of the flow around the fixed rigid plate (not mounted on the two springs) have been performed for various Reynolds numbers. Snapshots of the vorticity fields are displayed in fig. 5.4.1. For $Re_{c,wake} < Re \lesssim 8500$, the flow unsteadiness is mainly visible in the wake (see fig. 5.4.1-b)) with vortices of opposite signs that are alternatively shed at the trailing edge. The flow unsteadiness at the leading edge and around the plate is very weak and barely visible. For higher values of the Reynolds number, as $Re = 10^4$ corresponding to fig. 5.4.1(c), large fluctuations are clearly visible around the plate. More specifically, small vortices are shed from




FIGURE 5.4.2: Spatial distribution of the streamwise velocity for the base (upper halves) and mean (lower halves) flow at Reynolds numbers $Re = 8325$ and $Re = 10^4$. The black curves delimit the recirculation regions.

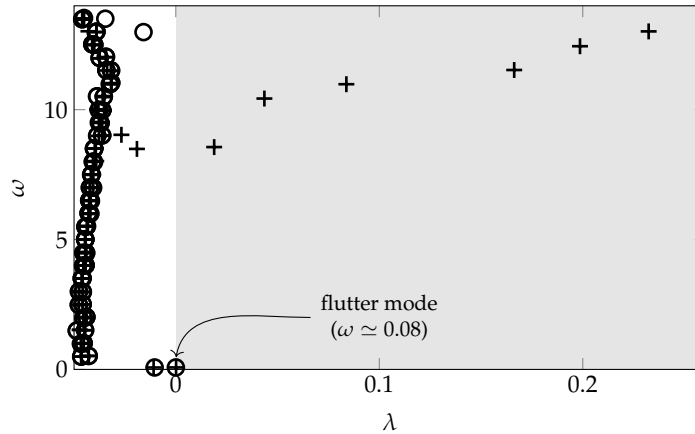


FIGURE 5.4.3: Comparison of the baseflow (crosses) and mean flow (circles) linear stability spectra at $Re = 10^4$, $\tilde{m} = 10^4$ and $U^* = 10.7$.

the leading-edge shear layers, travel downstream along the plate and interact with the trailing-edge vortices. For a thorough investigation of the interaction between those leading-edge and trailing-edge vortices when varying the plate's length, we refer to [Hourigan et al. 2001]. Here, we focus on the effect of the flow unsteadiness on the (time-averaged) mean flow, that are shown and compared to the base flow in fig. 5.4.2. For $Re = 8325$ (top figure), the difference between the base (upper half) and mean (lower half) flows is large in the wake. The main effect of the fluctuations is to decrease the recirculation region at the trailing edge. The two symmetric recirculation regions located around the leading-edge corners are very similar in the base and mean flows. For $Re = 10^4$ (bottom figure), the stream-wise extent of the leading-edge recirculation regions clearly decreases in the mean flow.

In the following, we propose to take into account this mean-flow distortion in the nonlinear analysis of the flutter instability. Close to the critical Reynolds number $Re_{c,wake}$ for vortex-shedding, a rigorous weakly nonlinear analysis may be performed around a critical set of parameters (mass ratio and Reynolds number) where the vortex-shedding and flutter modes get simultaneously unstable. This type of two-mode expansion was successfully applied, for example, by Meliga et al. (2012) for studying mode selection during vortex breakdown in swirling jets. Far from the




FIGURE 5.4.4: Linear flutter analysis performed with the base and mean flows for $\tilde{m} = 10^4$. (a) Critical reduced velocity as a function of the Reynolds number obtained with base (solid line) and mean (circles) flow analyses. The critical Reynolds number $Re_{c,wake}$ for the onset of vortex-shedding is depicted with the vertical dashed line. (b) Spatial distribution of the pressure field for the (real part of) flutter eigenmode obtained at $Re = 8325$ and $Re = 10^4$ using the base flow (upper halves) and the mean flow (lower halves). Black curves delimit the recirculation regions in the base and mean flows.

critical Reynolds number, the two-mode weakly nonlinear analysis may not be accurate. Rather than focusing on this rigorous analysis, we propose a simpler alternative that replaces the base flow \mathbf{q}_0 by the hydrodynamic mean-flow $\bar{\mathbf{q}}$ in the weakly nonlinear expansion eq. (5.4). With this approach the weakly nonlinear analysis can be performed taking into account only the flutter mode, since all hydrodynamic modes of the meanflow are stable (see fig. 5.4.3). By doing so, we partially retain the nonlinear interaction between purely-hydrodynamic fluctuations and flutter instability, since the hydrodynamic mean-flow accounts for the nonlinear interaction of flow harmonics. However, we neglect the dynamic interaction between the purely-hydrodynamic harmonics and the flutter harmonics. That assumption is reasonable as the vortex-shedding frequency ($\omega_{wake} = 11.2$ at $Re_{c,wake}$) is much higher than the flutter frequency ($\omega \simeq 0.1$ for $\tilde{m} = 10^4$). In other words, the hydrodynamic vortex-shedding is time-averaged at the slow time scale of the flutter phenomenon. In the weakly nonlinear analysis around the mean flow, we thus analyze the nonlinear behavior of the low-frequency flutter instability that develops on the mean hydrodynamic flow

Let us first examine results of the linear flutter analysis of a plate of mass ratio $\tilde{m} = 10^4$. Figure 5.4.4(a) shows the critical reduced velocity as a function of the Reynolds number obtained with the base (solid curve) and mean (circles) flow analyses. They are very similar in both analyses, for all values of the Reynolds number. The marginally stable flutter modes obtained for the Reynolds numbers $Re = 8325$ and $Re = 10^4$ are shown in fig. 5.4.4(b) using the real part of the pressure fields. In both analyses and at both Reynolds numbers, they display larger values and gradients around the reattachment of the leading-edge recirculation regions. At $Re = 8325$, the baseflow and mean flow modes are very similar in the leading-edge region. This is in agreement with the similar velocity threshold reported in 5.4.4(a). For $Re = 10^4$, the baseflow and mean flow modes are still qualitatively similar, but the latter is slightly shifted upstream, consistently with the shortening of the leading-edge recirculation region in the mean flow. This shift however appears to

have only a small effect on the critical velocity shown in 5.4.4(a).

Let us now consider results of the weakly nonlinear analysis by examining more specifically the nonlinear coefficient $\tilde{\beta}$ defined in 5.11 and related to the strength of the cubic nonlinear term in the amplitude equation. Figure 5.4.5(a) displays that quantity, computed with the base (solid line) and mean (circles) flow, as a function of the Reynolds number. Results for the base flow indicate that, when increasing the Reynolds number up to $Re \approx 8500$, the bifurcation is more and more subcritical since $\tilde{\beta}$ is more and more negative. Above that Reynolds number, $\tilde{\beta}$ abruptly increases until it gets positive. For $Re = 10^4$, the weakly nonlinear analysis of the base flows thus predicts a supercritical bifurcation. This is the behaviour already observed in fig. 5.3.3 for high Reynolds numbers and mass ratios. Results for the mean flow are very similar up to $Re \approx 7000 - 8000$. Above $Re = 8000$, the mean flow results start to deviate from the baseflow analysis, but keep the same trend. The difference becomes striking for Re above ~ 9500 where the bifurcation remains subcritical in the mean flow analysis whereas it becomes brutally supercritical in the baseflow analysis. To further understand that behaviour, we display in 5.4.5(b) the spatial distribution of \tilde{f}_{β_f} whose integral over the computational domain is almost ⁴ equal to $\tilde{\beta}$. The upper plot displays the scalar field obtained with the base (upper) and mean (lower) flow at $Re = 8325$. In both cases, largest values are localized in the shear layer of the leading-edge recirculation regions. As noticed previously, the flow unsteadiness in those regions is very weak for that Reynolds number, and the mean flow is very similar to the base flow. This explains that $\tilde{\beta}$ are almost equal for the base and mean flow analyses. The wake of the plate, where the base and mean flow are very different, does not contribute at all to $\tilde{\beta}$ which explains why the shortening of the recirculation bubble in the mean flow does not impact $\tilde{\beta}$. The lower plot in 5.4.5(b) displays \tilde{f}_{β_f} for $Re = 10^4$. In this case, the mean flow distribution is clearly shifted upstream in comparison to the baseflow and the variations of \tilde{f}_{β_f} are less important than in the baseflow analysis. Overall, we note that the large impact of the mean flow on the value of $\tilde{\beta}$ is only due to mild differences in the balance between the supercritical (blue lobes) and subcritical (red lobes) regions.

⁴It was checked that geometric part $\tilde{\beta}_g$ almost does not contribute to $\tilde{\beta}$ for the chosen mass ratio ($\tilde{m} = 10^4$).




FIGURE 5.4.5: Weakly nonlinear flutter analysis performed around the base or mean flows for $\tilde{m} = 10^4$. (a) Nonlinear coefficient $\tilde{\beta}$ as a function of the Reynolds number obtained with the base (solid line) and mean (circles) flow analyses. Negative and positive values correspond to subcritical and supercritical bifurcation, respectively. (b) Spatial distribution of the nonlinear fluid contribution $\tilde{f}_{\tilde{\beta}_i}$, (cf eq. (5.14)) is presented for both the base and mean flow analyses at $Re = 8325$ and $Re = 10^4$. The integral of that contribution gives the dominant contribution to $\tilde{\beta}$.

In the end, the proposed mean flow-based weakly nonlinear analysis shows that the type of the flutter bifurcation is strongly impacted by switching from baseflow to mean flow analysis for Re higher than ≈ 8000 . This is linked to the fact that for such Reynolds numbers, unsteadiness occurs close to the leading edge and changes the topology of the leading edge shear layers, which are the most contributing regions to $\tilde{\beta}$. On the contrary, vortex shedding in the trailing edge area does not affect $\tilde{\beta}$. In any case, the present analysis, is not meant as a rigorous analysis, but rather as a way of stressing that results in fig. 5.3.3 must be taken with caution for high Reynolds numbers. In addition, it underlines again the critical role of the leading edge shear layers in deciding the type of the flutter bifurcation, and hence the need for accurate fluid modeling.

5.5 Bifurcation scenarios at low Reynolds number $Re = 500$

In this section, we focus on the bifurcation scenarios occurring at the low Reynolds number $Re = 500$ when varying the mass ratio and the reduced velocity. As seen in fig. 5.3.3, for that Reynolds number, the weakly nonlinear analysis predicts a transition from a supercritical bifurcation at high mass ratio ($\tilde{m} \sim 1000$) to a subcritical bifurcation at lower mass ratio ($\tilde{m} \sim 60$). We further explore those bifurcation scenarios using, in addition to the weakly nonlinear analysis valid close to the critical reduced velocities, the time spectral method (TSM) and time-marching simulations. The supercritical and subcritical bifurcation scenarios are described in §5.5.1 and §5.5.2, respectively. For intermediate mass ratios (when the nonlinear coefficient of the amplitude equation is close to zero), a more complex scenario is examined in §5.5.3 and is discussed in light of experimental results by [Amandolese et al. 2013](#).




FIGURE 5.5.1: Supercritical bifurcation for $\hat{m} = 1000$ and $Re = 500$. (a) Maximal pitching angle and (b) angular frequency as a function of the reduced velocity ratio U^*/U_c^* (with $U_c^* = 4.95$). (c) Phase diagram of the periodic solution $U^*/U_c^* = 1.016$ in the map (θ, C_M) . The black and grey lines correspond to a branch of solutions obtained with time spectral and weakly nonlinear methods, respectively, while black circles correspond to results of time-marching simulations. The cross in (c) is the steady solution.

5.5.1 Supercritical bifurcation at high mass ratio

Diagrams corresponding to a supercritical bifurcation are displayed in fig. 5.5.1 for the mass ratio $\hat{m} = 1000$. The maximal pitching angle is shown as a function of the reduced velocity in fig. 5.5.1(a) where the solid curve and black dots correspond to results of TSM (with $2N_h + 1 = 41$ instants) and time-marching simulations (with about 1400 timesteps per period), respectively. The thin gray lines, materializing the weakly nonlinear results, are superimposed for comparison. For all reduced velocities, TSM and time-marching simulations superimpose perfectly so that we can not only conclude that TSM is accurate enough with this number of instants, but also that the LCO solutions are stable⁵. The solution amplitude smoothly increases when increasing the reduced velocity above U_c^* . The maximum pitching angles is about 8° at $U^*/U_c^* = 1.05$. This is a relatively low amplitude, in comparison with those achieved for subcritical bifurcations, presented in the following. A consequence of these low amplitudes is that the weakly nonlinear approach (thin gray line) yields satisfying pitching amplitude predictions up to, say, $U^* = 1.05U_c^*$. In fig. 5.5.1(b), we show the bifurcation diagram for the oscillation frequency. On the considered interval of reduced velocity, the frequency only slightly decreases, in fact following a $1/U^*$ trend which is the typical frequency of the natural solid modes. In particular, we remark that the weakly nonlinear analysis (thin gray line) perfectly captures the frequency evolution. For a given reduced velocity, $U^*/U_c^* = 1.016$ (filled circles in (a)), the solution obtained with TSM (solid black line) is represented in in fig. 5.5.1(c),

⁵Here, we evaluate the stability of TSM solutions by simply restarting a time-marching computation from the TSM solution: if the time-marching solution remains on the periodic orbit predicted by TSM, then we consider the TSM solution stable, if not, we consider it unstable.

in the $\theta - \mathcal{C}_M$ map⁶. The elliptic shape of the obtained orbit shows that both the plate displacement and moment coefficient follow a mostly harmonic evolution.




FIGURE 5.5.2: Supercritical bifurcation. Snapshots of vorticity for the low-amplitude LCO at $\tilde{m} = 1000$ and $U^*/U_c^* = 1.016$.

The spatial distribution of the vorticity is depicted in fig. 5.5.2 at four instants of the period. At all instants, the flow remains attached to the plate and closely resembles the weakly nonlinear solution already shown in fig. 5.3.1. This further confirms the mildly nonlinear nature of this LCO. Furthermore, the flow also resembles the steady solution (see fig. 4.3.6b), which is in agreement with the low-frequency $\omega \simeq 0.17$ characterizing this supercritical LCO. In other words, this LCO is a mildly nonlinear quasi-steady periodic solution.

5.5.2 Subcritical bifurcation at low mass ratio

We now consider the lower mass-ratio $\tilde{m} = 60$ for which the weakly-nonlinear analysis predicts a subcritical bifurcation. The bifurcation and phase diagrams are displayed in fig. 5.5.3 as for the supercritical bifurcation presented above. At the critical reduced velocity ($U^* = U_c^*$), an unstable LCO branch emerges and progresses towards lower reduced velocities with increasing amplitude. The branch predicted by the weakly-nonlinear analysis (dashed grey curve) is clearly of too high amplitude compared to the branch obtained with the TSM (dashed black line), showing that strong nonlinear effects are at play, that can be only captured with the latter method. As the reduced velocity is further decreased, the amplitude of the pitching angle drastically increases when getting close to $U^* \simeq 0.93U_c^*$. This is observed for the TSM branch but not for the weakly-nonlinear branch. In fact, the weakly nonlinear analysis cannot approximate the turning point observed at $U^* \simeq 0.93U_c^*$, unless considering high-order terms in the expansion eq. (5.4), so as to obtain a quintic amplitude equation. When the branch turns and folds back with increasing U^* , the

⁶We recall that, throughout the thesis, we use the trigonometric convention for orienting angles, *i.e.* $\theta > 0$ nose-down. Thus, we chose to represent $-\mathcal{C}_M$ as a function of $-\theta$ in fig. 5.5.1(c), in order to obtain a figure respecting the more classical "positive, nose-up" convention of aerodynamic studies.




FIGURE 5.5.3: Subcritical bifurcation for $\tilde{m} = 60$ and $Re = 500$. Same legend as in fig. 5.5.1. The critical reduced velocity is $U_c^* = 1.33$. The dashed curves corresponds to unstable LCO. The LCO displayed in (c) corresponds to $U^* = 0.95U_c^*$.

pitching amplitude increases drastically in a small range of reduced velocity. For the reduced velocity $U^* = 0.95U_c^*$, two solutions coexist: a low-amplitude LCO of maximal pitching angle $\theta_{max} \sim 20^\circ$ and a high-amplitude LCO of maximal pitching angle $\theta_{max} \sim 40^\circ$. Note that the nonlinearity is so strong that we could not compute the high-amplitude LCO branch using the TSM method, because the performance of the block-circulant preconditioner deteriorates too much for very large fluctuations, as explained in table 3.4.1. Still it is successfully used to capture the unstable branch of periodic solution that cannot be obtained with time-marching simulations, while the latter are reserved for computing the high-amplitude branch of solutions (marked with black symbols), that are stable and still periodic. For the largest reduced velocity considered here ($U^* \simeq 1.08U_c^*$), the maximal pitching angle is very large, around 60° . Contrary to the pitching amplitude, the oscillation frequency, shown in fig. 5.5.3(b), does not vary significantly. The large-amplitude branch possesses a slightly higher frequency than the unstable low-amplitude branch, but the global trend is, as in the supercritical case, dictated by the natural solid frequencies (not shown).

The unstable and stable LCO solutions co-existing at $U^*/U_c^* = 0.95$ are displayed in fig. 5.5.3(c) in the $\theta - C_M$ map, with dashed and solid lines, respectively. The central cross materializes the steady (stable here) solution and the thin gray dashed line the weakly nonlinear prediction. In contrast to the supercritical case (see fig. 5.5.1(c)), both LCO orbits are far from ellipses. In particular, the large amplitude stable LCO (solid line) presents multiple local maxima of the moment coefficient. When the reduced velocity is decreased, the two solutions get closer and closer, until they collide and disappear for $U^* < 0.93$. This is a saddle-node bifurcation of periodic orbits.

Snapshots of vorticity are shown in fig. 5.5.4 for the (a) unstable low-amplitude and (b) stable high-amplitude solutions. In comparison with the periodic solution shown in fig. 5.5.2 for the supercritical bifurcation, the flow unsteadiness is stronger in both cases. For the low-amplitude solution in fig. 5.5.4(a), weak leading-edge vortices of opposite signs are alternatively formed in the shear layers on the plate's suction side, when the extremal values of the pitching angle are reached at $t = 0$

and $t = T/2$. These leading-edge vortices are then shed in the wake during the upstroke and downstroke phases, respectively. This vortex-shedding was not observed for the lower amplitude solution of the supercritical bifurcation since the shear-layers evolve quasi-steadily in this case. For the high-amplitude solution shown in fig. 5.5.4(b), the vortex-shedding pattern is more complex. First, we note that the vorticity field satisfies the spatio-temporal symmetry $\omega_z(x, -y, t + T/2) = -\omega_z(x, y, t)$ *i.e.* the pattern of the vorticity fields during the first and second half-periods are symmetric and of opposite sign. Thus, we can focus on describing only the first half-period.

At $t = 0$, the plate is at maximal pitching angle. A large recirculation region, delimited by the thin black line on the figure, (almost) covers the plate. A strong clockwise-oriented (blue) leading edge vortex is being fed by the detached upper leading-edge shear layer. In addition, an elongated counter-rotating (red) vorticity layer occupies the plate's upper surface. At $t = T/8$, the leading-edge vortex detaches from the feeding shear layer. Due to its large size, its centroid is already located above the plate's trailing edge. Concomitantly, a weaker secondary clockwise leading-edge vortex is growing in the upper shear layer. Under the influence of the primary vortex, the counter-rotating upper side shear layer progressively detaches from the plate's surface. At $t = 2T/8$, the primary vortex has been shed in the wake, whereas the secondary vortex detaches from the feeding shear layer. In addition, the counter-rotating upper side shear layer has evolved into a better defined weak counter-rotating vortex. At $t = 3T/8$, the primary clockwise vortex has paired together with the weak counter-rotating vortex in the wake, forming a dipole. The secondary clockwise vortex is also shed in the wake so that the flow is fully re-attached on the upper side of the plate. Notice also the detached flow on the lower side of the plate which constitutes the early stage of the sequence of events in the lower shear layer symmetric to the one we are currently describing in the upper shear layer. At the plate's trailing edge, a tertiary weak clockwise vortex is being formed due the nose-down pitching motion of the plate that induces an upstream movement of the trailing edge. At $t = T/2$, this tertiary clockwise vortex appears more clearly as it is shed in the wake. For completeness, we can even guess a quaternary (very weak) clockwise vortex being formed at the trailing edge, and later released in the wake at $t = 5T/8$.

In order to get more detailed insights into the role of the different vortical structures generated during the high-amplitude LCO presented in fig. 5.5.4b, we propose in fig. 5.5.5 a more detailed description of the generation of the three dominant structures of that solution — *i.e.* the primary and secondary clockwise vortices and the counter-rotating vortex — and their connection to lift and moment forces. The LCO solution is presented in the $h - C_L$ and $\theta - C_M$ maps in the center of the figure. On a side note, we notice that, contrary to the $\theta - C_M$ orbit which is composed of one clockwise oriented loop, the $h - C_L$ orbit presents a more complex pattern of five loops (the three larger are well visible in the figure, the two small ones are found at the extremities of the orbit). Because these loops have opposite orientations, the average power provided by the fluid to the solid over one period along the heaving direction is negligible (by three orders of magnitude) with respect to the work along the pitching direction. Resuming now the description of the vortical structures and their connection the aerodynamic forces, six points marked by letters from (a) to (f) are chosen on the orbit and the corresponding snapshots are presented in fig. 5.5.5. For comparison with fig. 5.5.4b, the vorticity field is represented, together with the pressure isocontours (black lines, negative contours with dashed line) and

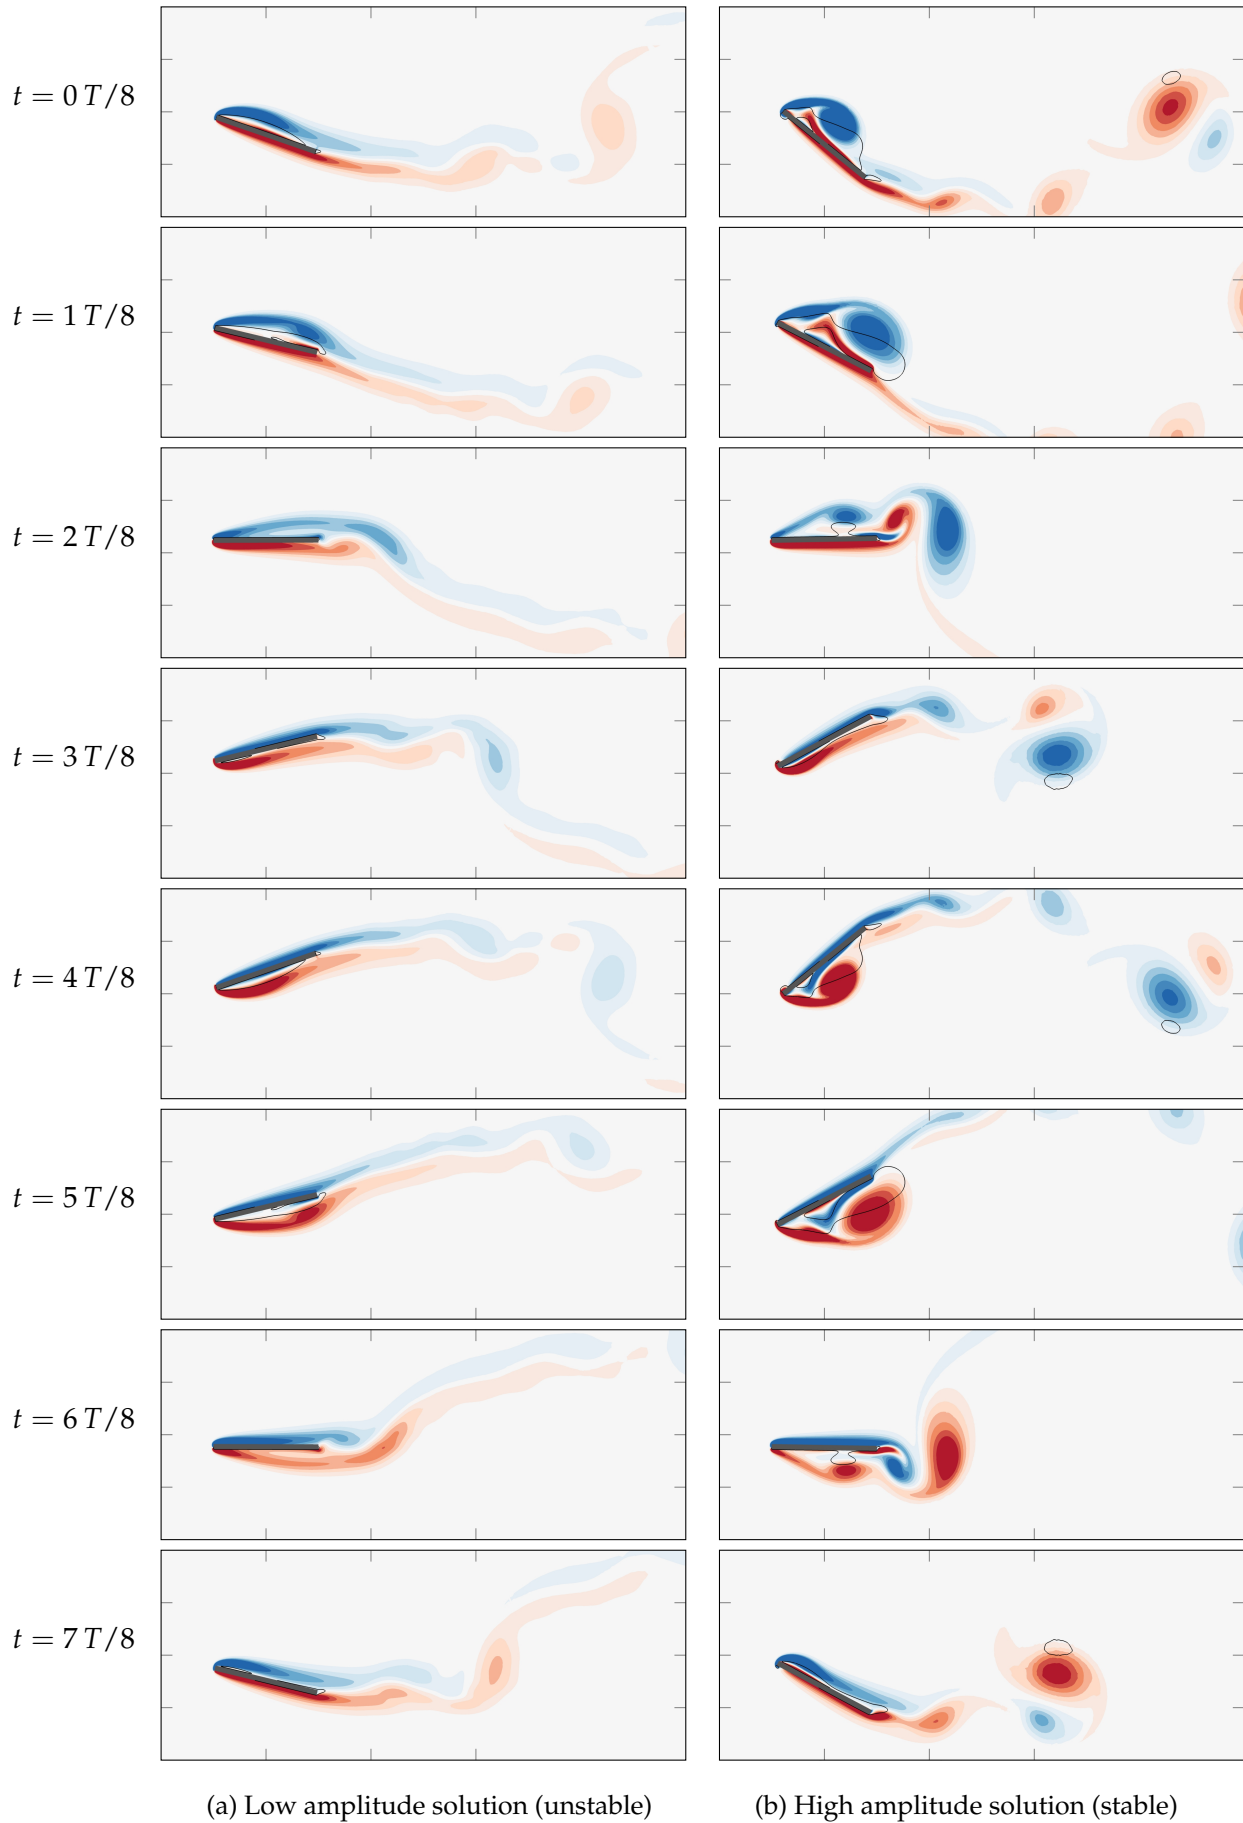


FIGURE 5.5.4: Vorticity snapshots of the two coexisting flutter LCO's in the subcritical case ($\tilde{m} = 60$, $U^*/U_c^* = 0.95$).




FIGURE 5.5.5: Description of the high-lift generation process of the high-amplitude stable solution in subcritical flutter ($\bar{m} = 60$, $U^*/U_c^* = 0.95$). For each instant from (a) to (f), the vorticity (colored contours) and pressure (black isolines, negative values are dashed) fields are presented, along with a few streamlines in orange.

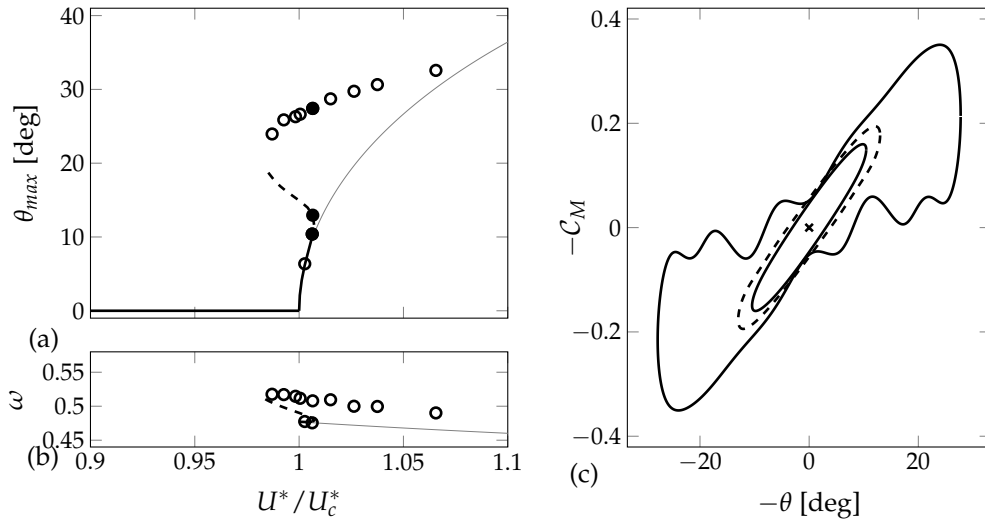


FIGURE 5.5.6: Bifurcation for intermediate mass ratio $\tilde{m} = 120$ ($Re = 500$). Same legend as in fig. 5.5.1. The critical reduced velocity is $U_c^* = 1.78$. The three LCO displayed in (c) corresponds to $U^*/U_c^* = 1.006$.

some streamlines (orange). Starting at point (a), we are just before the instant of maximum pitching angle. The upper leading edge shear layer is progressively detaching, thus creating a recirculation region down to the midchord. The low-pressure area at the leading edge induces a lift and strongly clockwise moment (recall the moment is computed at the elastic axis, which is the midchord). At (b), the primary clockwise vortex introduced before has formed and locates approximately above the midchord. The low-pressure area is still on the upper side of the plate, thus conserving a high lift. Despite it has displaced downstream with the vortex, the existence of a leading-edge high-pressure area on the lower side of the plate (due to incoming flow) still generates a strong clockwise moment. In (c), the primary vortex and the corresponding low-pressure area are above the trailing edge. A high-lift is still obtained but the moment brutally drops. Besides, the secondary clockwise vortex and the counter-rotating vortex are clearly visible at the leading edge. Between (c) and (d), the lift drops to negative values due to the upper side low-pressure area leaving the plate with the primary vortex. In the same interval of time, the moment slightly mitigates its drop thanks to the low-pressure leading edge regions associated to the secondary vortices. In (e) and (f), the lift and moment continue decreasing (*i.e.* anti-clockwise for the moment) due to the fact that all the upper side low-pressure area has left the surface of the plate, with the secondary vortices.

Overall, the previous descriptions of the supercritical (fig. 5.5.2) and subcritical (fig. 5.5.4(b) and 5.5.5) show that the super and sub-critical stable solutions are linked to radically different flows. In the supercritical case, the quasi-steady shear layers remain attached to the plate across the cycle, whereas in the subcritical case, leading edge shear layers dynamically roll-up into a series of vortices of various strength and signs that generate strongly nonlinear lift and moment coefficients. We mention that this crucial importance of the shear layers was announced already in the analysis of the spatial decomposition of $\tilde{\beta}_f$ proposed in section 5.3.2 where these regions were shown to decide the type of the bifurcation.

5.5.3 Double-fold bifurcation scenario at intermediate mass ratio

For intermediate mass ratio, there exists a transition between a supercritical and a subcritical bifurcation, that is identified, in the weakly nonlinear analysis, when the nonlinear coefficient vanishes, i.e. $\tilde{\beta} = 0$. However, results of the weakly nonlinear analysis are then of limited interest since, at third-order, all nonlinear terms vanish. Rather than considering the effect of higher-order terms in the development, we again use time spectral and time-marching simulations to investigate the nonlinear saturation of the LCOs at those intermediate mass ratio.

As for the previous bifurcation scenarios, we display the maximal pitching amplitude and frequency as a function of the reduced velocity in fig. 5.5.6(a) and (b), respectively. At the critical reduced velocity U_c^* , the bifurcation is supercritical in agreement with the weakly nonlinear prediction (grey curve) giving birth to a branch of stable low-amplitude LCO's. The amplitude growth away from U_c^* is very quick as the solution already reaches about 10° when the reduced velocity only exceeds U_c^* of less than one percent. This is the consequence of the small positive value of $\tilde{\beta}$ for this set of parameter (see fig. 5.3.2). On this branch, the solution is mostly harmonic as shown by the corresponding elliptic orbit in fig. 5.5.6c (inner solid line orbit). The vorticity field (not shown) is similar to fig. 5.5.2 despite the higher oscillation frequency, $\omega \sim 0.5$, shown in (b). Then, instead of the expected smooth increase of the LCO amplitude with increasing U^* away from the threshold, a first fold bifurcation is observed for $U^*/U_c^* \simeq 1.007$, leading to an unstable medium amplitude LCO branch. Traveling further along the medium amplitude branch, a second fold point is guessed, but as for the case $\tilde{m} = 60$, could not be reached with TSM for numerical reasons. Beyond the second fold, a high-amplitude stable branch of LCO's is retrieved using DNS. On this branch, the solutions resembles fig. 5.5.4(b). On the range of reduced velocities between U_c^* and $1.007U_c^*$, the system thus possesses three nested LCO solutions, as shown in fig. 5.5.6c for $U^* = 1.006U_c^*$. The small and large amplitude LCO's are stable (solid lines) whereas the medium amplitude LCO is unstable (dashed line).

As a consequence, the case $\tilde{m} = 120$ gives rise to a non-standard bifurcation scenario where a supercritical Hopf is followed by two folds. Two important remarks then need to be made. First, because the second turning point is located for $U^* < U_c^*$, stable high-amplitude periodic solutions exist below the linear velocity threshold, despite the supercritical nature of the primary bifurcation. Second, if we suppose the system to be on the low amplitude branch, then, as one increases U^* , a brutal change of the system response will be observed around $U^*/U_c^* \simeq 1.007$, when the low amplitude branch disappears at the first fold, and the system consequently jumps to the large amplitude branch. In addition, if one now decreases U^* from the high amplitude branch, an hysteretic behavior is expected, as the system will go back to the steady baseflow solution only for $U^*/U_c^* < 0.98$ (where the second fold point is located), instead of $U^*/U_c^* = 1$.

We conclude this section by proposing in fig. 5.5.7 a synthetic view of the three different bifurcation scenarios observed before at $Re = 500$. In this figure, we gather the results for $\tilde{m} = 60, 120$ and 1000 previously shown in fig. 5.5.3, 5.5.6 and 5.5.1, respectively. The evolution of the pitching angle and oscillation frequency as a function of reduced velocity are reproduced, together with an additional representation of the heaving amplitude. The horizontal axis represents the ratio of U^* to U_c^* so that all bifurcations conveniently start at the same point. However, it should be kept in mind that the critical reduced velocities for these three mass ratios are different (from $U_c^* = 1.33$ at $\tilde{m} = 60$ to $U_c^* = 4.95$ at $\tilde{m} = 1000$). The pitching amplitude plot

(a) is meant as a summary and the reader is reported to the figures corresponding to the individual mass ratios for a complete description. For the heaving amplitude (not shown until now), we notice that varying the mass ratio does not affect much the heaving amplitude: for example, at $U^* = 1.05U_c^*$, the heaving amplitude simply goes from about 0.3 at $\tilde{m} = 1000$ and $\tilde{m} = 120$ to 0.4 at $\tilde{m} = 60$. In comparison, the pitching amplitude goes from 10° to about 60° . This induces the idea that the difference between the high-amplitude subcritical LCO's and the low-amplitude supercritical LCO's is mainly in the pitching motion. Finally, in (c) we first observe that for all mass ratios, the frequency evolves along a $1/U^*$ trend imposed by the natural solid frequencies that are materialized by the thin dotted (heaving) and dash-dotted (pitching) lines. The low-amplitude solutions at the highest mass ratios ($\tilde{m} = 1000$, orange) are rather associated to the heaving frequency, whereas the high-amplitude solutions at lower mass ratios shift towards the pitching frequency. Overall, the occurrence of large amplitude LCO's (and subcritical flutter) is associated with an increase of the oscillation frequency. It can be interpreted that high enough frequencies are necessary to generate the highly unsteady vortex shedding pattern associated to the large amplitude LCO's (see fig. 5.5.4b).

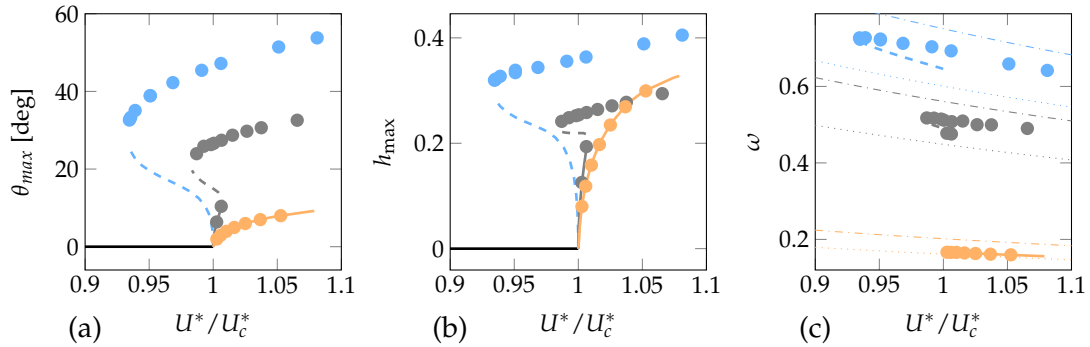


FIGURE 5.5.7: Transition from subcritical to supercritical bifurcations at low Reynolds number ($Re = 500$) when increasing the mass ratio ($\tilde{m} = 60$ in blue, $\tilde{m} = 120$ in grey and $\tilde{m} = 1000$ in orange). (a) Maximal pitching angle, (b) maximal heaving displacement and (c) frequency of the LCO as a function of U^*/U_c^* . Dots and curves correspond to results of time-marching and time-spectral methods, respectively, while solid and dashed curves distinguish stable from unstable LCOs. In (c), the natural heaving and pitching frequencies are represented by dotted and dash-dotted curves, respectively.

5.5.4 Experimental evidence of the double-fold scenario

In [Amandolese et al. 2013], the flutter response of a system similar to the one used in this work was investigated experimentally, in a wind tunnel. In particular, the assumption of a fully linear solid model that is used throughout the present work, was verified, at least for angles of attack lower than 25° . The characteristics of their spring-mounted plate, summarized in table 5.5.1, are similar to ours though not identical. In this paragraph we aim at comparing the bifurcation diagram obtained by these authors with the present work. If fitting our solid parameters to the experimental ones would have been possible, our 2D Navier–Stokes fluid model is not reliable at Reynolds as high as the one used experimentally ($Re \sim 3 \cdot 10^4$). Therefore, the following comparison is only qualitative.

	c/e	x_θ	Ω	ζ_h	ζ_p	r_θ
exp. [Amandolese et al. 2013]	23	0.08	0.78	0.002	0.015	0.35
present work	20	0	0.8	0	0.05	0.29

TABLE 5.5.1: Comparison of the parameters of the plate used by [Amandolese et al. 2013] with the present work.

By varying the wind velocity, the authors obtained the bifurcation diagrams that are reproduced in fig. 5.5.8a. High-amplitude subcritical responses and strong hysteresis, similar to the results of fig. 5.5.3 were observed. If those features are compatible with a classical subcritical Hopf scenario, the presence of the low amplitude response for $U^*/U_c^* \simeq 1.08$, marked by a red disk, remained unexplained by the authors. In fig. 5.5.8b, the bifurcation diagram for $(\tilde{m}, Re) = (120, 500)$ is reproduced. To ease the comparison with fig. 5.5.8a, a red disk symbol is added on the low-amplitude branch. By comparing the pitching (left) and heaving (right) amplitude diagrams, we see that the major difference between the low-amplitude and the high-amplitude branches is in the pitching dynamics. Indeed, whereas the pitching amplitude is almost tripled at $U^*/U_c^* \simeq 1.007$ when going from the low-amplitude to the high-amplitude branch, the heaving amplitude only increases of about 40%. These qualitative features are present in both the experimental results and in our numerical example. As a consequence, it can be argued that the supercritical Hopf bifurcation followed by two folds offers a possible explanation for the experimentally observed co-existence of (stable) low-amplitude supercritical and large-amplitude subcritical solutions.

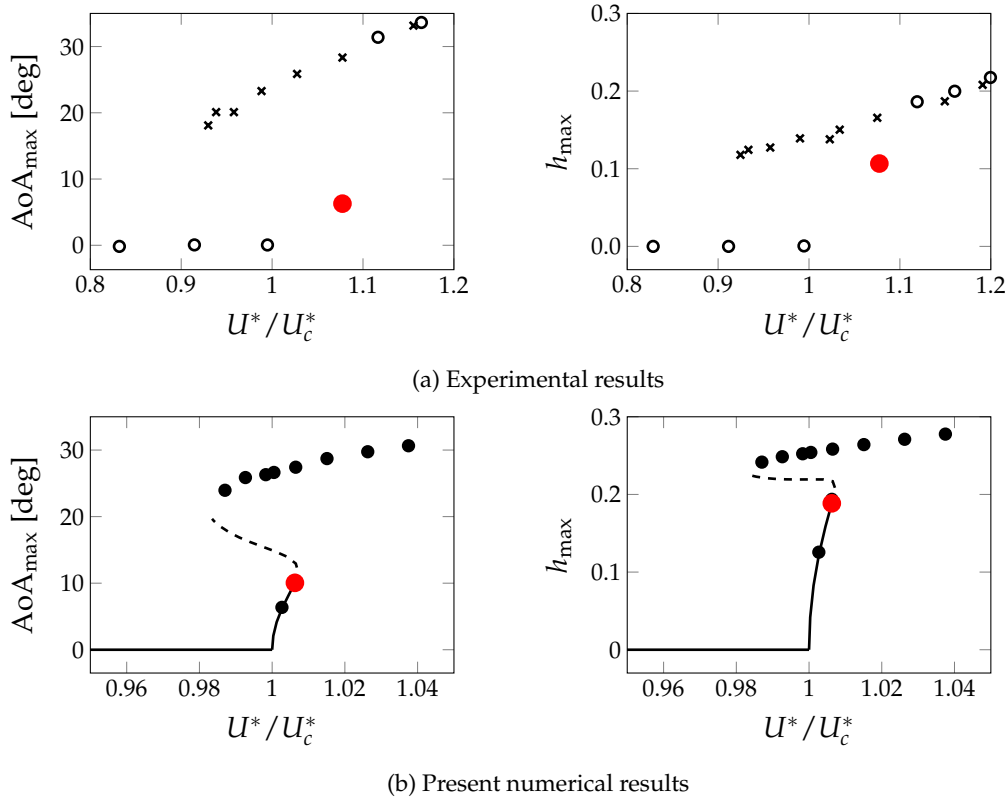


FIGURE 5.5.8: Qualitative comparison of the supercritical double-fold scenario exposed in fig. 5.5.7 with wind-tunnel data from [Amandolese et al. 2013]. Pitching (left) and heaving (right) amplitude results are given for (a) the experimental results ($\tilde{m} = 2100$, $\mathcal{Re} \simeq 30000$) and (b) the present numerical results ($\tilde{m} = 120$, $\mathcal{Re} = 500$). For the description of (a), see fig. 5.5.7. In (b), circles (resp. crosses) mark increasing (resp. decreasing) wind velocity.

5.6 Conclusion

Using a weakly nonlinear analysis, we showed that the flutter bifurcation is strongly affected by the solid-to-fluid mass ratio and the Reynolds number. For very low $\mathcal{Re} < 90$, the bifurcation is always supercritical. For intermediate Reynolds number, $90 < \mathcal{Re} \lesssim 2000$, the bifurcation is supercritical at high mass ratios $\tilde{m} > \tilde{m}_u$, transitions to subcritical on a range of mass ratios $\tilde{m}_l < \tilde{m} < \tilde{m}_u$ and finally goes back to supercritical for very low mass ratios below \tilde{m}_l . For the highest Reynolds numbers considered in this study, $2000 \lesssim \mathcal{Re} < 10000$, the bifurcation is subcritical for all mass ratios. In the rest of the study we mostly focused on an intermediate Reynolds number of $\mathcal{Re} = 500$. We proposed a decomposition of the cubic coefficient in the amplitude equation that allows to determine how the different nonlinearities of the model contribute to the sub- or super-critical behavior. It was shown that geometric nonlinearities — the ones that are due to the rotation of the plate — mostly drive the low mass ratio transition, at $\tilde{m} = \tilde{m}_l$ whereas, the fluid nonlinearities coming from the Navier–Stokes momentum equation clearly drive the upper mass ratio transition, at $\tilde{m} = \tilde{m}_u$. By further decomposing the fluid nonlinear contribution in space, we demonstrated that the nature of the upper mass ratio transition is mostly decided by the shear layers at the plate leading edge. In order to investigate the highly-nonlinear solutions that appear away from the threshold we combined

results from the Time Spectral Method and reference time-marching simulations. For high mass ratio ($\tilde{m} = 1000$), where the bifurcation is supercritical, stable periodic solutions with low amplitude appear above the critical reduced velocity U_c^* . These solutions oscillate at a low frequency leading to quasi-steady flow features with well-attached shear layers. On the contrary, for a low mass ratio ($\tilde{m} = 60$), where the bifurcation is subcritical, high-amplitude periodic solutions are observed down to $0.94U_c^*$. These solutions possess a high oscillation frequency and are characterized by strong unsteady fluid effects involving the shedding of multiple vortical structures. The largest vortical structure corresponds to a strong leading edge vortex that is generated when the plate reaches high angles of attack. Finally, at an intermediate mass ratio ($\tilde{m} = 120$), we presented an unusual scenario where the bifurcation is supercritical but still allows high-amplitude solutions below the critical velocity. This situation is made possible by the succession of two fold bifurcations of LCO's that follow the original supercritical Hopf bifurcation. We ended the study with a discussion of the double-fold scenario in light of previous experimental results by [Amandolèse \(2016\)](#) where both low-amplitude supercritical solutions and high-amplitude subcritical solutions were observed. It is proposed that the double-fold scenario provides a consistent explanation for these previously unexplained results.

APPENDIX


Appendix 5.A Grid convergence study

The sensitivity of the weakly nonlinear results with respect to computational domain size and grid refinement are presented in fig. 5.A.1a, for a subcritical case ($\tilde{m} = 1000, \mathcal{R}e = 2682.81$). In the first column, the radius of the computational domain R is reported. In the second column, the interpolation error for the mesh adaptation process (cf section 2.4.3) is given whereas the corresponding number of triangles is reported in the third column. The last three columns present respectively the critical flutter velocity, the critical frequency and the normalized quantity $\tilde{\beta}$ defined in eq. (5.11). From the first to the fourth row, the computational mesh is refined. It is observed that the linear quantities, U_c^* and ω_c , are both almost converged already with the coarsest mesh. On the contrary, $\tilde{\beta}$ converges slower with the grid. This is easily linked to the small scale variations of its spatial distribution, as illustrated in fig. 5.4.5(b).

In the last three rows, the dependence of the results as a function of the computational domain size is explored. It is seen again that the linear quantities are very robust, while $\tilde{\beta}$ more significantly depends on R . With the smallest domain, $R = 10$, the obtained $\tilde{\beta}$ deviates about 7% from the reference value obtained at $R = 50$. In any case, it should be kept in mind that these variations of a few percent are negligible in regards to the physical variations observed in fig. 5.3.3.

R	err	n_{tri}	U_c^*	ω_c	$\tilde{\beta}$
20	$3 \cdot 10^{-2}$	16,662	4.576	0.1808	-18.94
20	$1 \cdot 10^{-2}$	36,750	4.576	0.1809	-19.17
20	$5 \cdot 10^{-3}$	65,434	4.575	0.1809	-19.18
20	$3 \cdot 10^{-3}$	100,362	4.575	0.1809	-19.18
10	$1 \cdot 10^{-2}$	33,444	4.530	0.1828	-20.93
30	$1 \cdot 10^{-2}$	43,194	4.571	0.1811	-19.65
50	$1 \cdot 10^{-2}$	58,280	4.565	0.1813	-19.59

(a) Grid sensitivity table



(b) Computational domain

FIGURE 5.A.1: Grid sensitivity of the weakly nonlinear analysis at ($\tilde{m} = 1000, \mathcal{R}e = 2682.81$).

6

LOW-FREQUENCY MODULATION OF FLUTTER SOLUTIONS: A FLOQUET ANALYSIS BASED ON THE TIME SPECTRAL METHOD

The flutter oscillations of a thin plate mounted on a system of bending/torsional springs and immersed in a laminar incompressible flow are investigated numerically for a low Reynolds number $Re = 500$ and a large solid-to-fluid mass ratio $\tilde{m} = 1000$. Time-marching simulations of the fluid-solid interaction show that, when increasing the reduced velocity above the critical value U_c^ , periodic oscillations of the plate are first observed as a consequence of the classical (primary) flutter instability of the steady solution. These oscillations, of moderate frequency $\omega_0 \sim 0.16$ are characterized by a well-behaved flow that remains mostly attached to the plate during the motion. For slightly larger reduced velocity, a low-frequency modulation of the oscillations is observed, where the plate reaches higher angles of attack and the flow more massively detaches. Drawing a Poincaré map clearly indicates that this quasi-periodic solution is a torus attractor. To explain the emergence of the quasi-periodic solutions, a Floquet stability analysis is performed. The latter fully relies on the Time Spectral Method, not only to compute the unstable periodic limit cycle oscillations, but also to determine the leading Floquet modes. We show that an asynchronous Floquet mode gets unstable for values of the reduced velocity where quasi-periodic solutions are observed. The low modulation frequency ω is well predicted by the Floquet analysis in the vicinity of the critical velocity for the secondary instability. An analysis of the pitching and heaving components of the Floquet mode and the associated perturbation reveals that the low-frequency modulation is linked to the continuous drift of the phase difference between pitching and heaving across the modulation period. This is made possible by the fact that the pitching angle predominantly oscillates at a slightly higher frequency $\omega_0 + \omega$ than the heaving displacement which dominantly oscillates at $\omega_0 - \omega$. When the pitching motion precedes the heaving motion, energy is extracted from the flow with perturbation dynamics that are reminiscent of the classical flutter instability, and the perturbation amplitude increases. On the contrary, when heaving precedes pitching, the solid energy is dissipated into the flow via so-called “anti-flutter” dynamics, and the perturbation amplitude decreases.*

6.1 Introduction

The flutter linear instability that occurs on a heaving and pitching spring-mounted typical section [E. H. Dowell et al. 1989, §3.3.5] gives a rise to a Hopf bifurcation, as the consequence of the destabilization of a pair of complex conjugate fluid-structure eigenmodes of the steady configuration where the section is horizontal. For reduced velocities above the critical threshold for flutter, this pair of eigenmodes grows exponentially until the system's nonlinearity come into play and saturate their growth. Then, periodic regimes with well defined amplitude, referred to as *Limit Cycle Oscillations* (LCO), are observed. In the case of a supercritical flutter bifurcation, that is the focus of this chapter, the amplitude of the LCO is usually expected to grow smoothly as the reduced velocity increases. However, LCO solutions — just like steady solutions — can be unstable and transition, through secondary instabilities, to richer temporal behaviors, like quasi-periodic oscillations or even chaos [Nayfeh et al. 1995].

Several researchers have reported the destabilization of flutter LCO on two degrees of freedom spring-mounted typical aeroelastic sections. The vast majority of these works investigated various types of nonlinearities in the structure, while keeping the fluid modeling linear (steady, quasi-steady or Theodorsen models). For example, in [L. Liu et al. 2004; B. H. Lee et al. 2005], a cubic stiffness nonlinearity in the pitch spring and coupled to Theodorsen's model for the fluid was investigated. In [Zhang et al. 2017; G. Liu et al. 2018], the cubic stiffness was put in the degree of freedom governing an added external store while steady thin airfoil theory was used for the fluid. These authors showed that for high enough reduced velocities, the flutter LCO undergoes low-frequency amplitude modulations, which are attributed to the destabilization of a pair of complex conjugate Floquet modes [G. Liu et al. 2018]. Using different types of structural nonlinearities, *e.g.* freeplay or hysteresis, a wide zoology of dynamics may be observed, as exposed in the comprehensive review by B. H. K. Lee et al. (1999a). Regarding the effect of fluid nonlinearities, the literature is much scarcer and — to the authors' best knowledge — always based on empirical dynamic stall models. Using such a fluid model, X. Li et al. (1997) showed that a structurally linear airfoil with nonzero static angle of attack transitions from stable static equilibrium, to flutter LCO, and finally to chaos, through a "quasi-periodic route to chaos". More recently, similar quasi-periodic oscillations were reported in [Jian et al. 2009], for a structurally nonlinear high-aspect-ratio wing.

Mathematically, the asymptotic linear stability of a periodic orbit is carried on using Floquet stability analysis [Floquet 1883]. This analysis typically consists in computing the so-called Floquet multipliers, that are complex numbers that quantify the growth/decay and frequency of infinitesimal perturbations that may develop on top of a previously established LCO. The most widely spread numerical approach is to obtain the Floquet multipliers as the eigenvalues of the *monodromy matrix* Φ , defined as the operator that propagates a perturbation $\mathbf{q}'(t)$ over one period T_0 of the LCO: $\mathbf{q}'(t + T_0) = \Phi \mathbf{q}'(t)$. Building Φ can be quite a cumbersome task, as it requires K (K being the number of degrees of freedom) time-integrations of the linearized governing equations, across one period T_0 . More details on different variants of this approach can be found, for example, in [Peletan et al. 2013]. When high-dimensional physical models — typically, the spatially discretized Navier–Stokes equations — are used, K becomes so large that building Φ is not feasible. In that case, only a few dominant Floquet multipliers are computed, using Arnoldi type algorithms [Barkley et al. 1996; Elston et al. 2004; Deng et al. 2016; Jallas et al. 2017; Shaabani-Ardali et al. 2019]. In this case, the linearized governing equations must be time-marched

only a number of time equal to the size of the Krylov subspace times the number of Arnoldi iterations, which is always much smaller than K . In these approaches, based on time-marching the linearized equations, another challenge is to first compute the base LCO solution of which Floquet stability is then computed. If the LCO is stable, this can be simply done by time-marching the nonlinear equations on a long enough interval of time. If the LCO is unstable however, stabilization techniques must be devised in order for the time-marching approach to settle on these naturally unstable LCO solutions. Several methods have been proposed to achieve such a task, but are often based on the prior knowledge of a property of the targeted LCO. For example, if the LCO is known to possess some spatial symmetry, the latter can be directly imposed in the time-marching algorithm [Elston et al. 2004; Deng et al. 2016]. When only a spatio-temporal symmetry property of the LCO is known, specific stabilization techniques are used to artificially damp the symmetry-breaking components from the time-marching solution [Jallas et al. 2017; Shaabani-Ardali et al. 2019]. In parallel to these time-domain approaches, some researcher [L. Liu et al. 2004; G. Liu et al. 2018] have opted for frequency-domain approaches, often referred to as Harmonic Balance Methods (HBM). These methods impose the temporal periodicity of the researched solution by decomposing it as a (truncated) Fourier series. As a consequence, no *a priori* knowledge on the LCO is required to compute it, which gives HBM a significant advantage over the aforementioned time-domain approaches. The reader is reported to chapter 3 for extensive details on Harmonic Balance Methods. In addition to its ability to arbitrary compute LCO's regardless of their stability, HBM also offers the capacity to assess Floquet stability by computing the Floquet exponents as the eigenvalues of the linearized HBM equations [Deconinck et al. 2006; Lazarus et al. 2010; Krack et al. 2019]. As a consequence, and despite some known shortcomings mostly related to sorting the Floquet exponents [Lazarus et al. 2010], Harmonic Balance Methods appear as an appealing alternative for LCO stability analysis.

In this work, we study the stability of flutter LCO solutions that develop on a two degrees of freedom spring-mounted plate, immersed in a viscous incompressible flow. Whereas the solid model is purely linear, we are interested in the effect of the fluid nonlinearities, contained in the Navier–Stokes equations. To that purpose, we use a particular type of HBM, referred to as the Time Spectral Method (TSM) [Gopinath et al. 2005] to efficiently compute flutter LCO's. The TSM equations are solved thanks to the Newton–Krylov solver previously developed in section 3.3. Then, we show how the Floquet stability of those LCO's can be assessed through an eigenvalue analysis of the Jacobian of the TSM equations. An advantage of this approach is that it naturally fits in the TSM framework and heavily relies on routines already developed for the Newton–Krylov TSM solver.

The rest of the chapter organizes as follows. Section 6.2 starts by a brief summary of the equations governing the motion of the spring-mounted plate. Then, after recalling some elements of Floquet theory (section 6.2.2), we introduce in section 6.2.3 a TSM-based approach for first computing periodic flutter solutions and then assessing their Floquet stability. Section 6.3 gathers the presentation and analysis of the numerical results. We start by reporting in section 6.3.1, using time-marching computations, the emergence of low-frequency modulations on top of small amplitude flutter LCO's, for high-enough reduced velocities. Then, in section 6.3.2, these modulations are explained by a linear instability of the flutter LCO, due to the destabilization of a pair of complex conjugate Floquet modes. In section 6.3.3, we reconstruct the physical perturbation associated to the pair of unstable modes, allowing us

to shed some light on the physical mechanism governing the appearance of the low-frequency modulation in the fully nonlinear time-marching solutions (section 6.3.4).

6.2 Governing equations and TSM-based numerical approach for Floquet stability analysis

6.2.1 Governing equations

We investigate the non-linear dynamics of a rigid plate mounted on heaving and pitching springs with respective stiffness K_h and K_θ and immersed in a two-dimensional incompressible viscous open flow with far-field velocity U_∞ . As discussed in section 1.1, eight non-dimensional parameters defined in table 1.1, govern that fluid-structure interaction problem, namely, the heaving-to-pitching frequency ratio Ω , the structural damping ratios ζ_h, ζ_p , the radius of gyration r_θ , the position of the elastic axis x_θ (zero here), the solid-to-fluid mass ratio \tilde{m} , the Reynolds number Re and the reduced velocity U^* . In this chapter, only the reduced velocity is varied while all other parameters are kept fixed to values specified later. The reduced velocity is defined as the ratio of the far-field fluid velocity to a velocity typical of the natural pitching mode, $U^* = U_\infty / (c \sqrt{K_\theta / I_{ea}})$, where c is the plate's chord and I_{ea} is its moment of inertia with respect to the elastic axis.

The dynamics of the spring-mounted plate is governed by two (damped) linear oscillator equations eq. (1.3), that are recalled here:

$$\frac{\partial^2 h}{\partial t^2} + 2\zeta_h \left(\frac{\Omega}{U^*} \right) \frac{\partial h}{\partial t} + \left(\frac{\Omega}{U^*} \right)^2 h = \frac{1}{\tilde{m}} C_L \quad (6.1a)$$

$$\frac{\partial^2 \theta}{\partial t^2} + 2\zeta_p \left(\frac{1}{U^*} \right) \frac{\partial \theta}{\partial t} + \left(\frac{1}{U^*} \right)^2 \theta = \frac{1}{\tilde{m} r_\theta^2} C_M \quad (6.1b)$$

with C_L the lift coefficient and C_M the moment coefficients about the elastic axis (which is the mid-chord here). The flow is modeled with the two-dimensional incompressible Navier–Stokes equations, for which an Arbitrary Lagrangian-Eulerian (ALE) framework (written in a reference domain) is used to handle fluid domain motion [J.-L. Pfister et al. 2019]. The equations have been previously stated in chapter 1 (see eq. (1.11)) to which the interested reader is reported for more details.

In the following, the coupled system constituted by the solid and fluid equations is written formally as a first-order in time evolution operator:

$$\mathcal{M}(\mathbf{q}) \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}) = 0 \quad (6.2)$$

where the variable $\mathbf{q} = (h, \theta, u_h, u_\theta, \mathbf{u}, p, \boldsymbol{\lambda})^T$ gathers all variables necessary to describe the fluid-solid interaction. The first four scalar variables allow describing the dynamics of the rigid plate. They are the heaving h and pitching θ displacements, as well as the corresponding velocities u_h and u_θ . The incompressible flow is described with the velocity \mathbf{u} and pressure p fields. Finally, the variable $\boldsymbol{\lambda}$, defined at the fluid-solid interface, represents the local stress exerted by the fluid onto the solid. System (6.2) is composed of seven equations: the first four correspond to the solid model eq. (6.1), written as a first-order problem in time. The fifth and sixth equations are the Navier–Stokes momentum and mass conservation equations. The last equation corresponds to the equality of fluid and solid velocities at the fluid-solid interface. Note that the mass matrix $\mathcal{M}(\mathbf{q})$ depends here on the unknown solution \mathbf{q} . This is

due to the particular reference configuration ALE framework used to formulate the Navier–Stokes equations, where the Jacobian of the mesh deformation field appears in front of the time derivative (see eq. (1.11a)).

6.2.2 Linear stability analysis of Limit Cycle Oscillations

Limit Cycle Oscillations (LCO) are periodic solutions $\mathbf{q}_0(t)$ of (6.2) that can loose stability *via* linear instabilities. In this work, the linear stability of LCO's is assessed through a Floquet stability analysis which consists in evaluating the asymptotic fate of perturbations \mathbf{q}' that develop around the LCO. To that aim, the solution is decomposed as the sum of the periodic solution $\mathbf{q}_0(t)$ and an infinitesimal perturbation $\mathbf{q}'(t)$:

$$\mathbf{q}(t) = \mathbf{q}_0(t) + \epsilon \mathbf{q}'(t) \quad \epsilon \ll 1 \quad (6.3)$$

By injecting this *ansatz* solution in eq. (6.2), we obtain at order ϵ^0 the equation governing the LCO, i.e.

$$\mathcal{M}(\mathbf{q}_0) \frac{\partial \mathbf{q}_0}{\partial t} + \mathcal{R}(\mathbf{q}_0) = 0 \quad \text{with} \quad \mathbf{q}_0(t + T_0) = \mathbf{q}_0(t) \quad (6.4)$$

and at order ϵ^1 the equation governing the perturbation, i.e.

$$\mathcal{M}(\mathbf{q}_0(t)) \frac{\partial \mathbf{q}'}{\partial t} + \mathcal{J}(\mathbf{q}_0(t)) \mathbf{q}' = 0 \quad (6.5)$$

where the Jacobian operator of the coupled problem, linearized around \mathbf{q}_0 , is defined as

$$\mathcal{J}(\mathbf{q}_0(t)) = \partial \mathcal{R} / \partial \mathbf{q} |_{\mathbf{q}_0(t)} + \partial \mathcal{M} / \partial \mathbf{q} |_{\mathbf{q}_0(t)} \partial \mathbf{q}_0 / \partial t$$

In the present fluid-structure formulation, since the mass matrix depends on the state variable, the derivative of $\mathcal{M}(\mathbf{q})$ with respect to \mathbf{q} appears in the definition of $\mathcal{J}(\mathbf{q})$. Obviously, for physical systems with formalism less intricate than the present one, \mathcal{M} is generally a constant mass matrix, and we simply retrieve $\mathcal{J}(\mathbf{q}_0(t)) = \partial \mathcal{R} / \partial \mathbf{q} |_{\mathbf{q}_0(t)}$.

Following Floquet's theory [Floquet 1883], the solutions to a linear system with periodic coefficients, like eq. (6.5), can be found as a superposition of *Floquet form* signals:

$$\mathbf{q}'(t) = \mathbf{q}^\circ(t) e^{\sigma t} + \text{c.c.} \quad (6.6)$$

where $\mathbf{q}^\circ(t)$ is a T_0 -periodic function, called *Floquet mode*, and $\sigma = \lambda + i\omega \in \mathbf{C}$ is the *Floquet exponent*. Inserting the Floquet form in eq. (6.5) we find that $(\sigma, \mathbf{q}^\circ(t))$ are (generalized) eigenpairs of a spatio-temporal linear operator:

$$\sigma \mathcal{M}(\mathbf{q}_0(t)) \mathbf{q}^\circ + \left(\mathcal{M}(\mathbf{q}_0(t)) \frac{\partial}{\partial t} + \mathcal{J}(\mathbf{q}_0(t)) \right) \mathbf{q}^\circ = 0 \quad (6.7)$$

By solving eq. (6.7), the linear stability of the base LCO $\mathbf{q}_0(t)$ is thus determined from the growth rate of the Floquet exponents λ . If all $\lambda < 0$, then any perturbation is asymptotically damped and the LCO is linearly asymptotically stable. On the contrary, if at least one exponent verifies $\lambda > 0$, the corresponding perturbation grows exponentially and the LCO is linearly unstable. In the next paragraph, we focus on the numerical solution of the base LCO problem eq. (6.4) and the Floquet eigenproblem eq. (6.7) using a type of space-time discretization, called Time Spectral Method, well-suited to periodic problems.

6.2.3 A Time Spectral Method-based approach for Floquet stability analysis

The Time Spectral Method [Gopinath et al. 2005] is used for the computation of the LCO (6.4) and of the Floquet exponents/modes (6.7). It consists in discretizing time on a periodic grid $t_n = T_0 n / (2N + 1)$, $n = 0, \dots, 2N$ of $2N + 1$ instants, where $T_0 = 2\pi / \omega_0$ is the (*a priori* unknown) period of the solution and N is the number of harmonics captured. The time derivative of any periodic field is approximated at the grid points using a spectral approximation:

$$\left. \frac{\partial \mathbf{q}}{\partial t} \right| (t_n) \simeq \omega_0 \sum_{k=0}^{2N} d_k \mathbf{q}(t_{n+k}) \quad (6.8)$$

where the coefficients d_k are given by

$$d_k = \begin{cases} \frac{1}{2} (-1)^{k+1} \operatorname{csc} \left(\frac{\pi k}{2N+1} \right) & \text{if } k \neq 0 \\ 0 & \text{if } k = 0 \end{cases}$$

LCO computation based on TSM Using the spectral time-derivative approximation eq. (6.8), the time-dependent equations (6.4) governing the LCO can be recast into a system of $2N + 1$ coupled time-independent equations that writes:

$$\omega_0 \mathbf{D} \mathbf{Q}_0 + \mathbf{R}(\mathbf{Q}_0) = 0 \quad (6.9)$$

where the solution and residual vector are

$$\mathbf{Q}_0 = \begin{pmatrix} \mathbf{q}_0(t_0) \\ \vdots \\ \mathbf{q}_0(t_{2N}) \end{pmatrix} \text{ and } \mathbf{R}(\mathbf{Q}_0) = \begin{pmatrix} \mathcal{R}(\mathbf{q}_0(t_0)) \\ \vdots \\ \mathcal{R}(\mathbf{q}_0(t_{2N})) \end{pmatrix} \quad (6.10)$$

and the time-spectral derivative matrix is defined as

$$\mathbf{D} = \begin{pmatrix} 0 & d_1 \mathcal{M}(\mathbf{q}_0(t_0)) & \dots & d_{2N} \mathcal{M}(\mathbf{q}_0(t_0)) \\ d_{2N} \mathcal{M}(\mathbf{q}_0(t_1)) & 0 & \ddots & \\ & \ddots & \ddots & d_1 \mathcal{M}(\mathbf{q}_0(t_{2N-1})) \\ d_1 \mathcal{M}(\mathbf{q}_0(t_{2N})) & & d_{2N} \mathcal{M}(\mathbf{q}_0(t_{2N})) & 0 \end{pmatrix} \quad (6.11)$$

Since the frequency ω_0 is unknown, an additional equation is necessary to close the system. In practice, this equation is an arbitrary condition that fixes the phase of the LCO. In this work, we chose impose that the pitching velocity is zero, $u_\theta = 0$, at instant $t = 0$. Finally, we mention that in our TSM solver, the reduced velocity U^* is not directly imposed as a parameter. Instead, it is a second additional unknown of the problem. As for the frequency, a new scalar equation is provided to close the system, and consists in imposing the amplitude θ_{\max} of the pitching motion.

The TSM problem eq. (6.9), augmented with the additional phase and amplitude constraints introduced above is a nonlinear system of equation with a large number of degrees of freedom, $(2N + 1) \times K + 2$ with K the number of degrees of freedom corresponding to the finite element spatial discretization of \mathbf{q} . This problem is solved using a Newton–Krylov approach that combines classical Newton iterations with a Krylov subspace linear solver [Saad 1993] to tackle the challenging high-dimensional

linear problems arising at each Newton iteration. An adequate preconditioner for the Krylov method, proposed in chapter 3, is built using the so-called block-circulant preconditioner for the TSM equations and a Schur complement approach for the two additional scalar (phase and amplitude) constraints. The interested reader will find extensive details regarding the formulation, implementation, validation and numerical performance of the above described TSM solver in chapter 3.

Floquet mode/exponent computation based on TSM The eigenproblem eq. (6.7) is discretized in time and space using again the Time Spectral and Finite Element methods, respectively. This yields the following fully discrete problem of dimension $(2N + 1) \times K$:

$$\sigma \mathbf{M} \mathbf{Q}^\circ + (\omega_0 \mathbf{D} + \mathbf{J}) \mathbf{Q}^\circ = 0 \quad (6.12)$$

where $\sigma = \lambda + i\omega$ is the complex Floquet exponent associated to the Floquet mode \mathbf{Q}° , defined as

$$\mathbf{Q}^\circ = \begin{pmatrix} \mathbf{q}^\circ(t_0) \\ \vdots \\ \mathbf{q}^\circ(t_{2N}) \end{pmatrix}$$

The block-diagonal mass \mathbf{M} and Jacobian \mathbf{J} matrices write

$$\mathbf{M} = \begin{pmatrix} \mathcal{M}(\mathbf{q}_0(t_0)) & 0 & \dots & 0 \\ 0 & \mathcal{M}(\mathbf{q}_0(t_1)) & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mathcal{M}(\mathbf{q}_0(t_{2N})) \end{pmatrix} \quad \text{and}$$

$$\mathbf{J} = \begin{pmatrix} \mathcal{J}(\mathbf{q}_0(t_0)) & 0 & \dots & 0 \\ 0 & \mathcal{J}(\mathbf{q}_0(t_1)) & \ddots & \vdots \\ \vdots & \ddots & \ddots & 0 \\ 0 & \dots & 0 & \mathcal{J}(\mathbf{q}_0(t_{2N})) \end{pmatrix}$$

At this stage, it must be mentioned that the proposed approach is found to be very similar to the Floquet–Fourier–Hill method, originally devised as an analytical tool by Hill (1886), and more recently re-popularized as a numerical method [Deconinck et al. 2006; Lazarus et al. 2010]. The major difference resides in the fact that in the Floquet–Fourier–Hill method, a frequency-domain Harmonic Balance discretization of eq. (6.7) is used, instead of a time-domain here.

Numerical solution of the Floquet eigenproblem The size of the matrices involved in eq. (6.12) is $(2N + 1)K \times (2N + 1)K$, which quickly grows with both the spatial discretization (N) and the temporal discretization (K). As a consequence — and contrary to previous works [Deconinck et al. 2006; Lazarus et al. 2010] who focused on low-dimensional systems, $K \simeq \mathcal{O}(1)$ —, the spectrum of eq. (6.12) cannot be obtained at reasonable cost with classical dense algebra algorithms, even for low N . Instead, we perform Krylov–Schur iterations on the *shifted-and-inverted* eigenproblem :

$$\tilde{\sigma} \mathbf{Q}^\circ + \mathbf{T} \mathbf{Q}^\circ = 0$$

$$\text{with } \tilde{\sigma} = (\sigma - s)^{-1} \quad \text{and} \quad \mathbf{T} = [s\mathbf{M} + (\omega_0 \mathbf{D} + \mathbf{J})]^{-1} \mathbf{M}$$

in order to converge a few eigenvalues of interest, near the user-defined shift $s \in \mathbb{C}$. By moving s along the imaginary axis ($\lambda \simeq 0$), one can monitor the destabilization

Krylov–Schur it.	Linear solver calls	total GMRES it.	Wall-clock time
6	714	27,540	1h06min

TABLE 6.2.1: Cost of a Krylov–Schur run for the following numerical parameters: $s = 1 + 0i$, $n_{ev} = 40$, $n_{cv} = 200$, $N = 5$, $\text{tol}_{\text{GMRES}} = \text{tol}_{\text{KS}} = 10^{-6}$. The corresponding spectrum is shown in fig. 6.2.1(a).

of Floquet modes. As only matrix-vector products with \mathbf{T} are needed in the Krylov–Schur iterations, \mathbf{T} is never assembled. Only its action $\mathbf{x} = \mathbf{T}\mathbf{z}$ on a given vector \mathbf{z} is computed through one matrix-vector product $\mathbf{y} = \mathbf{M}\mathbf{z}$ and one solve of the linear system $[s\mathbf{M} + (\omega_0\mathbf{D} + \mathbf{J})]\mathbf{x} = \mathbf{y}$. Finally, noticing that the latter system is very similar to a Newton step of the nonlinear TSM solver, we simply re-use the block-circulant preconditioned GMRES solver already used in the nonlinear solver (see chapter 3). The typical cost of a Krylov–Schur run is presented in table 6.2.1 where we compute the 40 Floquet multipliers closest to the (real-valued for this case) shift $s = 1$, using $N = 5$ harmonics (*i.e.* 11 temporal grid points). The obtained spectrum is the one presented in fig. 6.2.1(a). With a Krylov space of size 200, we need 6 Krylov–Schur iterations to converge all 40 eigenvalues to 10^{-6} . This corresponds to 714 linear solves (GMRES calls), for a total amount of 27,540 GMRES iterations (*i.e.* as much applications of the circulant preconditioner and matrix-vector product with $[s\mathbf{M} + (\omega_0\mathbf{D} + \mathbf{J})]$). This corresponds to about 40 GMRES iterations per linear solve. Using the time-parallel implementation detailed in section 3.3.4, the total wall-clock time is about an hour.

Overall, it appears that most of the matrices and routines involved in the building and subsequent solution of eq. (6.12) are bricks of the Newton-Krylov TSM solver used for computing LCO solutions. In particular, the matrix $\omega_0\mathbf{D} + \mathbf{J}$ is simply the Jacobian of the TSM system, eq. (6.9). As a consequence, the Floquet eigenproblem can be assembled and solved by mostly re-using the computational routines of the nonlinear solver. From a practical point of view, we point out however that all computational routines must be adapted to handle complex algebra if one is to use a shift with nonzero imaginary part.

Periodicity of the Floquet exponents and choice of the relevant modes In this paragraph we point out the periodicity property of the Floquet exponents and explain how it affects the choice of the physically relevant eigenpairs amongst all eigenpairs $(\sigma, \mathbf{q}^\circ(t))$ of eq. (6.12).

In fig. 6.2.1(a), we show a typical spectrum of the leading Floquet exponents obtained from the numerical solution of eq. (6.12). The spectrum presents a repeating pattern — except at the upper and lower extremities, on which we will come back below — along the imaginary axis, with period ω_0 . The observed periodicity of the Floquet exponents originates from the fact that the Floquet form of the linear perturbation $\mathbf{q}'(t)$, introduced in eq. (6.6), in fact admits an infinity of equivalent representations, indexed by the integer $p \in \mathbb{Z}$:

$$\mathbf{q}'(t) = \mathbf{q}^{\circ,p}(t)e^{(\sigma + ip\omega_0)t} + \text{c.c.} \quad \text{with} \quad \mathbf{q}^{\circ,p}(t) = \sum_{n=-\infty}^{\infty} \widehat{\mathbf{q}}_n^{\circ} e^{i(n-p)\omega_0 t} \quad (6.13)$$

where the $\widehat{\mathbf{q}}_n^\circ$ are the Fourier coefficients of the ω_0 -periodic Floquet mode:

$$\mathbf{q}^\circ(t) = \sum_{n=-\infty}^{\infty} \widehat{\mathbf{q}}_n^\circ e^{in\omega_0 t}$$

Varying the index p simply “displaces” multiples of the base LCO frequency ω_0 from the Floquet mode to the Floquet exponent and *vice-versa*. However, the real physical perturbation $\mathbf{q}'(t)$ associated to each couple $(\sigma + ip\omega_0, \mathbf{q}^{\circ,p}(t))$ is left unchanged by changing p , as shown by eq. (6.13). As a direct consequence, it is easily verified that if $(\sigma, \mathbf{q}^\circ(t))$ is an eigenpair of eq. (6.7), then all the $(\sigma + ip\omega_0, \mathbf{q}^{\circ,p}(t))$, $p \in \mathbb{Z}$ are also eigenpairs. Thus, the spectrum of Floquet exponents is $i\omega_0$ -periodic, as observed in fig. 6.2.1(a).

In theory, this periodicity implies that only the eigenvalues in a section of the complex plane of extension ω_0 along the imaginary axis are representative of the whole spectrum. In practice however, truncation effects due to finite N (here $N = 5$) break that periodicity, as observed in the upper and lower extremities of the spectrum in fig. 6.2.1(a). The physically relevant eigenvalues (*i.e.* the ones that respect the periodicity) must then be carefully separated from the non-relevant ones (*i.e.* the ones that break the periodicity). In this work, the distinction is made by looking at the Fourier spectrum of the corresponding Floquet mode. As an example, we show in fig. 6.2.1(b) the Fourier spectrum of the dominant Floquet exponent (marked by a gray disk in (a)). If the harmonics of the mode well decrease to zero on the extremities of the computational spectrum $[-N\omega_0, N\omega_0]$ (*e.g.* the central Fourier spectrum in (b)), then the mode is well converged (in time), and thus physically relevant. On the contrary, if the harmonics are large at one of the extremities of the Fourier spectrum (see the upper or lower spectra in (b)), then the mode is ill-converged and thus discarded for physical analysis. Once all ill-converged exponents have been excluded, it is usually more convenient to represent LCO stability using the Floquet multipliers $\mu = e^{\sigma T_0}$ represented in fig. 6.2.1(c), instead of the exponents. Indeed, with the multiplier the aforementioned, purely mathematical, periodicity of the exponents naturally disappears due to $\mu = e^{(\sigma + ip\omega_0) T_0} = e^{\sigma T_0}$, for any $p \in \mathbb{Z}$.




FIGURE 6.2.1: Floquet stability of the base LCO $\mathbf{q}_0(t)$ computed at $U^* = 5.45$. The Floquet exponents obtained as a raw output of a Krylov–Schur run ($s = 1$, $N = 5$) are shown in (a). For a particular exponent marked by a gray disk (and its periodically repeating representations), the Fourier spectrum of the corresponding mode is shown in (b). Due to the periodicity of the exponents, only one pattern (with properly converged modes) is needed to assess Floquet stability (yellow area). In practice, the exponents representation is advantageously replaced by the Floquet multipliers representation (c), where the (mathematical) periodicity of the exponents naturally disappears.

6.2.4 Time-marching simulations

In order to compute solutions with unspecified time behavior (*e.g.* not only periodic solutions), we use a classical time-stepping approach to march in time the coupled equations eq. (6.2). A Backward Differences Formula of order two (BDF2) approximates the time derivative:

$$\left. \frac{\partial \mathbf{q}}{\partial t} \right| (t_n) \simeq \frac{3\mathbf{q}(t_n) - 4\mathbf{q}(t_{n-1}) + \mathbf{q}(t_{n-2})}{2\Delta t}$$

yielding the following nonlinear system at each timestep:

$$\frac{3}{2\Delta t} \mathcal{M}(\mathbf{q}(t_n)) \mathbf{q}(t_n) + \mathcal{R}(\mathbf{q}(t_n)) = \frac{4}{2\Delta t} \mathbf{q}(t_{n-1}) - \frac{1}{2\Delta t} \mathbf{q}(t_{n-2})$$

The latter is solved using the pressure segregation method proposed in [Badia et al. 2007](#) which consists in mixing a pressure correction approach [[Guermond et al. 2006](#)] to handle the incompressibility constraint with Dirichlet–Neumann fixed-point iterations for handling the fluid–structure coupling. Solving the implicit nonlinear system is thus decomposed in solving a sequence of simpler linear problems that are: (i) a linear advection-diffusion equation for the fluid velocity, (ii) a Poisson problem for the pressure increment and (iii) a four-by-four linear¹ solid problem. The fluid problems are space-discretized using the well-known Taylor–Hood ($\mathbb{P}_2, \mathbb{P}_1$) finite element pair for (\mathbf{u}, p) *via* the finite element library FreeFEM [[Hecht 2012](#)]. The discrete problems are solved in parallel with preconditioned Krylov subspace methods from the PETSc library [[Balay et al. 2019](#)], accessed through its FreeFEM interface. More details about the algorithm can be found in section 1.3.

6.3 Results

From now, we investigate the dynamics of typical section of mass ratio $\tilde{m} = 10^3$, immersed in an incoming uniform flow characterized by the Reynolds number $Re = 500$. As in chapters 4 and 5, the heaving-to-pitching frequency ratio Ω and the structural damping coefficients, ζ_h, ζ_p are set to $\Omega = 0.8$, $\zeta_h = 0$ and $\zeta_p = 0.05$, respectively. For this set of parameter, we recall that a flutter instability occurs at the critical reduced velocity $U_c^* = 4.96$ (see chapter 4). In chapter 5, we showed that the Hopf bifurcation that results from the flutter linear instability is supercritical. The branch of Limit Cycle Oscillation (LCO) emerging at U_c^* is characterized by a low fundamental frequency $\omega_0 \simeq 0.16 - 0.17$. Due to their low frequency, these LCO's are considered as quasi-steady solutions. Now, we investigate the unsteady nonlinear solutions that exist for a reduced velocity significantly above the critical flutter threshold, for U^* between U_c^* and $1.2U_c^*$. Using time-marching simulations, we first explore and characterize in section 6.3.1 the different solutions that naturally emerge on that velocity range. Depending on the reduced velocity, periodic or quasi-periodic solutions are found. The Time Spectral Method is then used in section 6.3.2, firstly to compute the LCO solutions that exist on the whole considered range of reduced velocities, and secondly to investigate their linear stability. In section 6.3.3, we reconstruct the real perturbation associated to the leading pair of Floquet multipliers and analyze it in order to understand the physical mechanism at play in the transition from periodic

¹The solid dynamics are fully linear due to the elastic axis being at the center of mass

to quasi-periodic solutions. We conclude in section 6.3.4 with a detailed comparison of the linear stability (Floquet) results and the fully nonlinear (time-marching) solutions.

6.3.1 From periodic to quasi-periodic solutions

Time-marching simulations are performed on a triangular grid of 15,274 triangles, starting from the steady solution as initial condition. The time step Δt is chosen so that the CFL number $CFL = U_\infty \Delta t / \Delta$ (with U_∞ the far-field velocity and Δ the typical size of the smallest triangle) is around 50. This corresponds to about 3000 timesteps per fast period² (T_0). The simulations are run for a time long enough so that a permanent regime is reached for each reduced velocity.

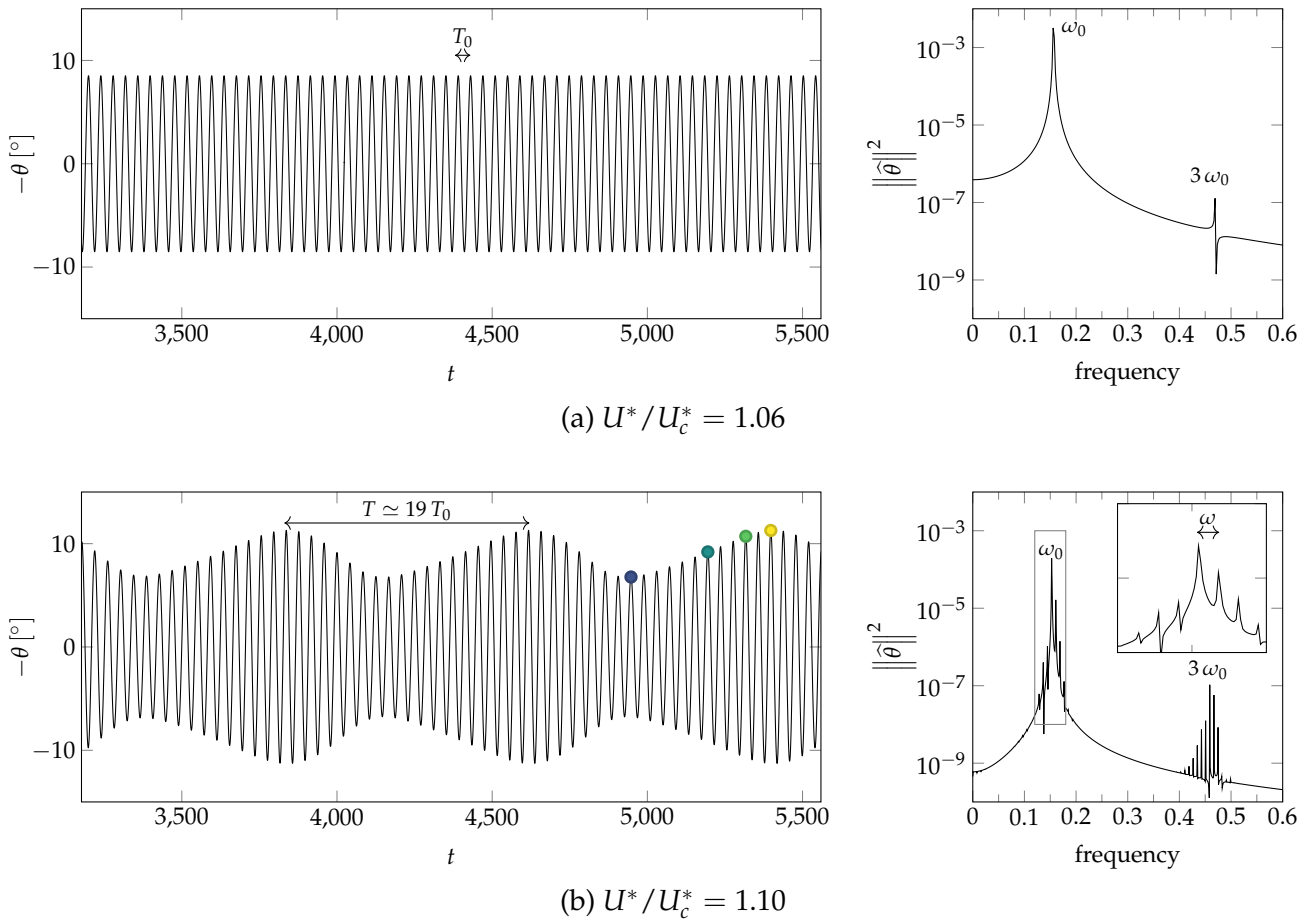


FIGURE 6.3.1: (a) Periodic and (b) quasi-periodic oscillations of the pitching angle displayed with time series (left column) and Fourier spectra (right column) for two reduced velocities above the critical flutter velocity U_c^* . ω_0 corresponds to the fundamental frequency of the periodic solution, while ω is the small frequency shift.

Figure 6.3.1 displays the temporal evolution of the pitching angle and its Fourier spectrum for two values of the reduced velocity. For $U^* = 1.06 U_c^*$, slightly above the critical reduced velocity, the solution is periodic with a frequency $\omega_0 = 2\pi/T_0 =$

²“Fast period” refers in this chapter to the period of the primary flutter instability at U_c^* , see below

0.16. The amplitude of the harmonic oscillating at $3\omega_0$ is very small when examining the pitching angle. However, the periodic LCO is not harmonic since higher amplitude harmonics, not shown here, are observed in the flow. Snapshots of the periodic evolution of the vorticity are shown in fig. 6.3.2(a-d). The dynamics consist in a coupled heaving and pitching motion of the plate, typical of the flutter instability. During the plate's oscillation, the flow is mainly attached to the plate, and it is slightly detached only for the highest angle of attack (around 9° here) reached during the plate's oscillation.

When increasing the reduced velocity to $U^* = 1.1U_c^*$, the fast flutter oscillations are still observed in fig. 6.3.1(b), but their amplitude is now modulated by a slow frequency $\omega = 2\pi/T$. The spectrum exhibits multiple peaks appearing around the fundamental frequency, at frequencies $\omega_0 - 2\omega, \omega_0 - \omega, \omega_0 + \omega, \omega_0 + 2\omega$, etc ... Since the two frequencies ω and ω_0 are not commensurable ($\omega \simeq \omega_0/19$), this solution is quasi-periodic. The flow evolution of this quasi-periodic solution is shown in fig. 6.3.2(e-h) at four instants of the slow period $T \simeq 19T_0$, corresponding to the local maxima of the pitching angle marked by colored disks in fig. 6.3.1(b). The flow separation is much more pronounced especially in the motion's phase where the plate reaches the largest pitching angle, close to 12° (see fig. 6.3.2(h)).




FIGURE 6.3.2: (a-d) Periodic and (e-h) quasi-periodic solutions displayed with the streamwise velocity field. The periodic solution obtained for $U^* = 1.073U_c^*$ is displayed at four instants of the period (a) $t = 0$, (b) $t = T_0/4$, (c) $t = T_0/2$ and (d) $t = 3T_0/4$. The slow evolution of the quasi-periodic solution obtained for $U^* = 1.1U_c^*$ is shown by displaying the solution at the four instants marked by colored disks in fig. 6.3.1(b). The thin black curves, corresponding to the contour of zero streamwise velocity, delimit regions of flow recirculation.

To gain some more insights in the periodic and quasi-periodic nature of these two solutions, we display in fig. 6.3.3(a-b) their temporal evolution in the map $\theta - \mathcal{C}_L$. The periodic solution (a) evolves on a simple orbit whereas the quasi-periodic solution (b) evolves on a *thick* orbit. Only one modulation period is depicted here, for visualization purpose, but with more slow periods, the orbit visits the whole thickness.




FIGURE 6.3.3: (a-b) Temporal evolution in the map $(-\theta, C_L)$ of the (a) periodic solution at $U^*/U_c^* = 1.060$ and (b) the quasi-periodic evolution at $U^*/U_c^* = 1.10$. (c-d) Poincaré section defined as the hyperplane $\{\dot{\theta} = 0, \ddot{\theta} > 0\}$ for the (c) periodic solution and (d) quasi-periodic solutions. Points are colored according to the time at which they were reached.

To show that this quasi-periodic solution is a torus attractor, the Poincaré sections is displayed in fig. 6.3.3(c-d). It is defined as the hyperplane $\{-\dot{\theta} = 0, \ddot{\theta} < 0\}$, which contains all local maxima of the instantaneous angle of attack. In practice, the Poincaré section is built by sampling the lift coefficient and pitching angle each time the latter passes through a local maximum. So, by construction, the Poincaré section is a discrete set of points. A total of about 800 fast periods (T_0) — *i.e.* of equivalently 42 slow periods (T) — is computed during the permanent regime in order to build a representative section. The periodic solution, seen in fig. 6.3.3(c), is a single point, since the lift coefficient has always the same value when the angle of attack reaches its maximal value. On the other hand, the quasi-periodic solution seen in fig. 6.3.3(d) is a closed curve, which is not continuously visited when time grows. This is highlighted by coloring each point of this curve with the time at which it is visited. As time grows, new points are added on that curve without superimposing to a previous one, thus progressively filling the closed curve. This shows that the quasi-periodic solution is a torus attractor (*e.g.* [Nayfeh et al. 1995]).




FIGURE 6.3.4: (a) Local maxima of the pitching angle for the periodic (single circle) and quasi-periodic (multiple circles) solutions as a function of the reduced velocity ratio. The filled black circle at $U^*/U_c^* = 1.06$ corresponds to the maximum pitching angle ($-\theta$) of the periodic solution (see fig. 6.3.1a) whereas the colored filled circles at $U^*/U_c^* = 1.10$ correspond to four local maxima of the pitching angle in the quasi-periodic solution (see fig. 6.3.1b). (b) Fundamental frequencies (circles) of the solutions as a function of the reduced velocity ratio. In the periodic range, only one fundamental frequency ω_0 exists, while in the quasi-periodic range two fundamental frequencies ω_0 and ω are visible. In (a) and (b), the bottom horizontal line corresponds to the steady solution with the plate horizontal (solid when stable, dashed when unstable).

Finally, we display in fig. 6.3.4 the solutions obtained when varying the reduced velocity in the range $1 < U^*/U_c^* < 1.2$. Maximal values of pitching angles are shown in fig. 6.3.4(a) as a function of the reduced velocity ratio. The stable and unstable steady solutions, corresponding to $\theta_{max} = 0^\circ$, are depicted with solid and dashed lines, respectively. We clearly see that a supercritical branch of periodic solutions, corresponding to a single circle for a reduced velocity, emerges above the critical reduced velocity U_c^* , where the steady solution gets unstable. Quasi-periodic solutions, characterized by multiple circles, are obtained when increasing the reduced velocity above $U^* = 1.073U_c^*$. The incommensurate frequencies ω_0 and ω characterizing those quasi-periodic oscillations are shown in fig. 6.3.4(b). A slight decrease of the fast frequency ω_0 is observed when increasing the reduced velocity, whereas the modulation frequency ω appears constant, at the scale of the representation (cf also fig. 6.3.10b)

Those diagrams clearly suggest that, like the periodic LCOs emerge from a linear instability of the steady solution, the quasi-periodic solutions result from a secondary instability, i.e. an instability of the periodic LCO. This scenario is investigated in the next paragraph by performing a Floquet analysis of the periodic LCO.

6.3.2 Floquet stability analysis of flutter LCO's

For reduced velocity $U^* > 1.073U_c^*$, only quasi-periodic solutions are obtained with time-marching simulations. To compute periodic solutions and investigate their linear stability, we use the time spectral methods described in §6.2.3.

We first examine the accuracy of the time-spectral methods for computing the periodic base solutions and its leading Floquet multiplier by varying the number of harmonics N (equivalently the number of instants $2N + 1$) and the size of the mesh

used for spatial discretisation. Up to $N = 20$ harmonics ($2N + 1 = 41$ instants) have been used for the computations of periodic solutions as well as Floquet modes. Let us recall that with our TSM solver, the reduced velocity is not imposed in practice as a parameter for computing the base LCO's. Instead, the reduced velocity is an unknown of the computation where an additional constraint is imposed by prescribing the pitching amplitude θ_{\max} .

N		1	3	5	10	20
ω_0	M_0	0.14859	0.15316	0.15315	0.15315	0.15315
	M_1	0.14859	0.15315	0.15314	0.15314	0.15314
$\mathcal{C}_{L\max}$	M_0	0.6031	0.6886	0.7047	0.7084	0.7083
	M_1	0.6030	0.6886	0.7048	0.7085	0.7083
U^*/U_c^*	M_0	1.12653	1.09980	1.10006	1.10002	1.10002
	M_1	1.12659	1.09984	1.10011	1.10006	1.10006
$\lambda(\cdot 10^{-5})$	M_0	-7.5015	1.7963	1.7866	1.7550	1.7548
	M_1	-7.0844	1.8014	1.7916	1.7559	1.7597
$\omega(\cdot 10^{-3})$	M_0	9.2168	9.5138	9.5379	9.5341	9.5341
	M_1	9.2187	9.5145	9.5388	9.5350	9.5350

TABLE 6.3.1: Effect of the number of harmonics N and mesh refinement M_i on the accuracy of the periodic base solutions (top) and leading Floquet exponent (bottom) both computed with time spectral methods. For all computations, the pitching amplitude is imposed to $\theta_{\max} = 10^\circ$ (U^* is an unknown), as demanded in practice by our TSM solver (cf section 6.2.3). The corresponding reduced velocity of the converged solution is $U^*/U_c^* = 1.10$. For the periodic base solution, we report the frequency ω_0 , the maximum lift coefficient, and the reduced velocity. The growth rate λ and frequency ω are shown for the Floquet exponents. The two meshes M_0 and M_1 are made of 15,274 and 28,394 triangles, respectively.

Table 6.3.1 displays results of this convergence study for an imposed pitching amplitude, $\theta_{\max} = 10^\circ$. In the upper part of the table, we report first some quantities related to the baseflow: its frequency ω_0 , the maximum lift coefficient $\mathcal{C}_{L\max}$ reached over a period and the reduced velocity corresponding to the imposed θ_{\max} . For the frequency ω_0 , spectral convergence is very fast since only three harmonics are required to capture it up to the fourth significant digits. In comparison, the case $N = 1$ clearly deviates. Convergence of the lift coefficient is significantly more difficult since the convergence of the fourth digit is only reached with $N = 20$. For the reduced velocity, convergence lies in between: easier than for lift but harder than for frequency. Then, we examine the sensitivity of the leading Floquet exponent to the number of harmonics N and grid size. The growth rate converges towards $1.755 \cdot 10^{-5}$ (four significant digits) using $N = 10$ harmonics. Except for the erroneous result obtained with $N = 1$, which predicts a stable Floquet mode, the growth rates obtained with a smaller number of harmonics are fairly good. For instance, with only $N = 3$ harmonics, the growth rate is predicted to 2 significant digits. The convergence of the frequency ω is quite similar. Finally, for all quantities considered, doubling the number of triangles in the mesh (from M_0 to M_1) barely modifies the results. As a consequence, all results shown hereinafter are obtained with the mesh M_0 .

The stability of the periodic base solution computed for $U^* = 1.073U_c^*$ is now addressed. This solution is very similar to the one displayed in fig. 6.3.1 and fig. 6.3.2. The main effect observed on the base LCO when increasing the reduced velocity (not

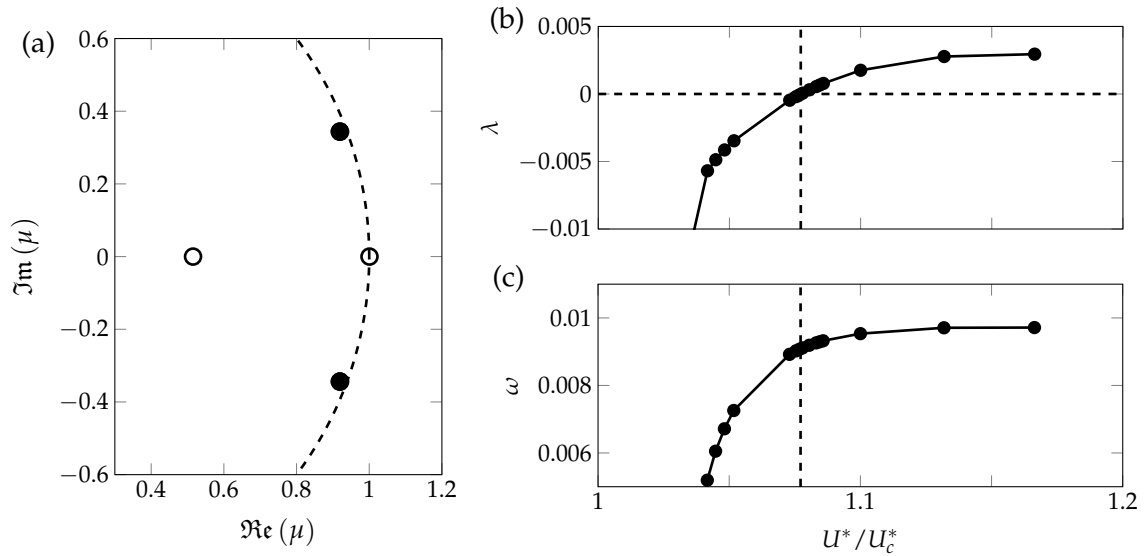


FIGURE 6.3.5: (a) Floquet multipliers computed for the periodic base flow shown in fig. 6.3.2 for the reduced velocity $U^*/U_c^* = 1.073$. The stability boundary (unit circle) is depicted by the dashed curve. (b-c) Evolution of the growth rate λ and frequency ω of the leading Floquet mode (black dots in figure a) as a function of the reduced velocity. The mode becomes unstable at a critical velocity $U^*/U_c^* \simeq 1.078$.

shown here) is the growing pitching and heaving amplitude as well as the widening of the recirculation region. The four leading Floquet multipliers computed for that periodic solution are shown in fig. 6.3.5(a). We recall first that the Floquet multiplier $\mu = 1$ is not of interest when addressing the stability of the periodic solution. Indeed, it is easily verified³ that the time-derivative of the periodic solution ($\partial \mathbf{q}_0 / \partial t$) is an eigenvector of eq. (6.7), associated to the exponent $\sigma = 0$ (or equivalently Floquet multiplier $\mu = 1$). The Floquet multipliers of interest when addressing the stability of the periodic base flow are marked with black dots in the figure. In the present case, this is a pair of complex conjugate that gets unstable when increasing the reduced velocity, as seen in fig. 6.3.5(b) and (c) that present the growth rate and frequency as a function of U^*/U_c^* . We recall that the frequency of the Floquet mode is associated to the argument ϕ of the Floquet multiplier according to $\omega = (\phi/2\pi)\omega_0$. As clearly visible in fig. 6.3.5(b), the argument of the Floquet multiplier is small, and the frequency associated to the asynchronous Floquet mode is also small. For $U^*/U_c^* = 1.100$, we have $(\phi/2\pi) \simeq 0.0622$, indicating that ω is approximately 6% of the fundamental frequency ω_0 . From the non-integer value of the ratio $\omega_0/\omega \simeq 16.07$, we see that the new period $T = 2\pi/\omega$ introduced by the Floquet mode is not a multiple of the fundamental period T_0 . The oscillating part of the perturbation, defined by neglecting the exponential growth as

$$\tilde{\mathbf{q}}(t) = \mathbf{q}^\circ(t) e^{i\omega t} + \text{c.c.} = 2 [\Re(\mathbf{q}^\circ(t)) \cos(\omega t) - \Im(\mathbf{q}^\circ(t)) \sin(\omega t)] \quad (6.14)$$

³take $\mathbf{q} = \mathbf{q}_0$ in eq. (6.2) and derive with respect to time

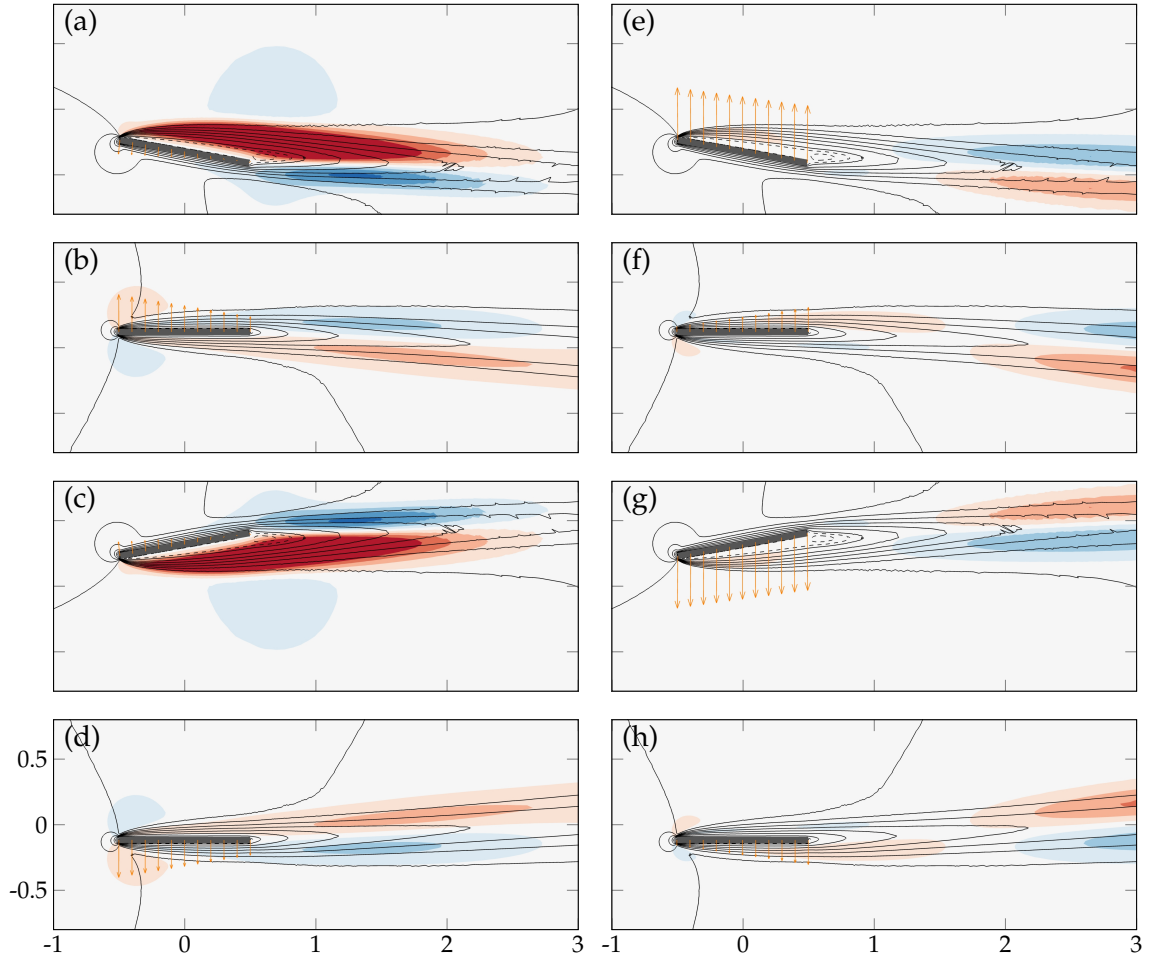


FIGURE 6.3.6: Periodic evolution of the flow and solid displacement fields for the unstable Floquet mode (with positive frequency $\omega > 0$) at $U^*/U_c^* = 1.10$. The streamwise velocity (red and blue) and solid displacement (orange arrows) fields are shown for the (a-d) real and (e-h) imaginary parts of the Floquet mode at instants (a,e) $t = 0$, (b,f) $t = T_0/4$, (c,g) $t = T_0/2$ and (d,h) $t = 3T_0/4$.

is thus a quasi-periodic function with two fundamental frequencies⁴: the low frequency ω introduced by the Floquet exponent and the high frequency $\omega_0 = 2\pi/T_0$ or the ω_0 -periodic mode, $\mathbf{q}^\circ(t)$.

The temporal evolution of the quasi-periodic perturbation $\tilde{\mathbf{q}}$ is directly linked to the periodic evolution of the real and imaginary parts of the Floquet mode, \mathbf{q}° . More precisely, because ω is much smaller than ω_0 , we have according to eq. (6.14) that $\tilde{\mathbf{q}}(t) \simeq 2\Re(\mathbf{q}^\circ(t))$ on an interval length T_0 centered on $\omega t \simeq 0$. Similarly, $\tilde{\mathbf{q}}(t) \simeq 2\Im(\mathbf{q}^\circ(t))$ on an interval centered on $\omega t \simeq \pi/2$. The real and imaginary parts of $\mathbf{q}^\circ(t)$ are displayed in fig. 6.3.6(a-d) and (e-h), respectively, at four instants of the fast period T_0 . The flow perturbations of the real and imaginary parts have similar spatial distribution, except at instants of the base flow period when the plate reaches its extrema positions (see figure a and e). At these instants, the flow perturbation

⁴The question whether these frequencies are commensurable or not cannot be answered with numerical computations. However, the continuous nature of the Poincaré sections obtained in the fully nonlinear regime with time-marching computations (see fig. 6.3.3(d)) tend to indicate that they are not.

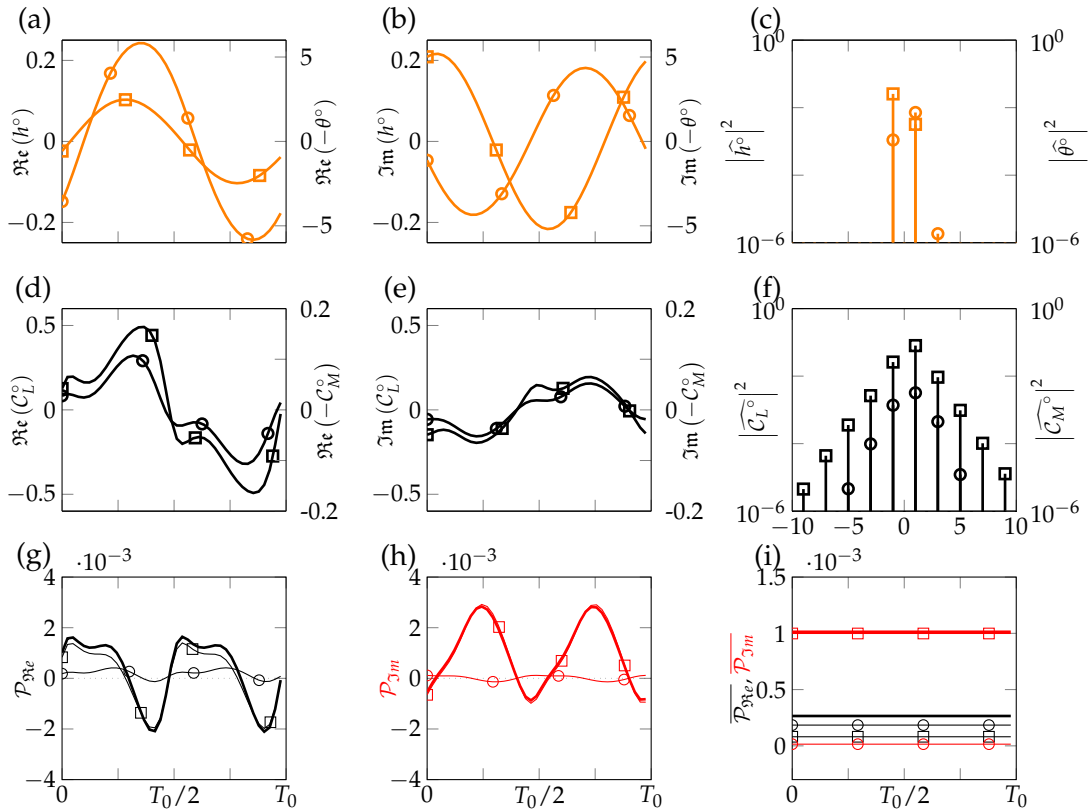


FIGURE 6.3.7: Periodic evolution of the (a-c) heaving (squares) and pitching (circles) components and (d-f) lift (squares) and moment (circles) coefficients for the unstable Floquet mode (with positive frequency $\omega > 0$) at $U^*/U_c^* = 1.10$. (g-h) Instantaneous power provided by the fluid to the solid: total power (thick line), heaving power (squares) and pitching power (circles). The real and imaginary parts of the time signals are displayed in (a,d,g) and (b,e,h), respectively. The modulus of the Fourier harmonics are presented in (c,f). (i) Mean power of the real (black) and imaginary (red) parts of the signal corresponding to the time series in (g) and (f) respectively.

is of much larger magnitude and located closer to the plate for the real part of the Floquet mode. We note that such a large flow perturbation is associated to a weak plate displacement (displayed by the orange arrows). This may be attributed to large solid inertial effects due to the high mass ratio used in this study, $\tilde{m} = 10^3$.

The periodic evolution of the plate's displacements is more visible in fig. 6.3.7(a-b). First, for both the real and imaginary parts displayed in (a) and (b) respectively, the heaving (squares) and pitching (circles) perturbation evolve harmonically in time. This is also assessed by the Fourier spectrum presented in fig. 6.3.7(c) where the first harmonic largely dominates. For the real part, heaving and pitching evolve almost in phase with pitching lagging slightly behind heaving. Oppositely, for the imaginary part, the signals are clearly out of phase with pitching preceding heaving of about $T_0/4$. Focusing on the heaving signal (squares), its magnitude is higher in the imaginary part than in the real part. The opposite is true for the pitching signal. The periodic evolution of the lift (squares) and moment (circles) coefficients are shown in fig. 6.3.7(d-e), with the corresponding Fourier spectrum displayed in fig. 6.3.7(f). Differently from the solid degrees of freedom, the aerodynamics loads

have a more complex time-behavior, typically involving higher harmonics, as shown in the Fourier spectra. Again, a clear difference is observed between the real and imaginary parts. However, this difference seems to be more involved than a simple phase shift (and rescaling), as for the solid components. For example, this can be seen by comparing the respective magnitude of the two local maxima of the lift coefficient in (d) and (e). To evaluate how the dynamics of the real/imaginary parts of the Floquet mode reflects in terms of energy transfer between the fluid and the plate, we show in fig. 6.3.7(e-f) the total power provided by the fluid to the solid (thick line) for each part of the signal:

$$\begin{aligned}\mathcal{P}_{\Re} &= \Re(C_L^\circ) \Re(u_h^\circ) + \Re(C_M^\circ) \Re(u_\theta^\circ) \\ \mathcal{P}_{\Im} &= \Im(C_L^\circ) \Im(u_h^\circ) + \Im(C_M^\circ) \Im(u_\theta^\circ)\end{aligned}$$

In addition the power associated to the individual heaving (squares) and pitching (circles) motions are represented separately. In both (g) and (h) we first observe that most of the variations in instantaneous power are provided by the heaving power, whereas the pitching power varies only marginally. In addition, the total power is roughly oscillating around zero in the real part whereas it is shifted towards positive values in the imaginary part. As a consequence, the mean power over one period T_0 significantly differs between the two signals. This is shown in (i) where we represent the mean power corresponding to the different signals in (g) and (h) (using the same conventions). For the real part (black), the total mean power is around $0.25 \cdot 10^{-3}$ and is constituted of heaving and pitching contributions of the same order of magnitude. Turning now to the imaginary part (red), the total mean power is about four times higher. In addition, it is entirely provided by the heaving degree of freedom, while the pitching degree of freedom produces zero mean power.

6.3.3 Analysis of the quasi-periodic perturbation: a “generalized flutter” instability

In the previous section, we described the spatio-temporal features of the Floquet mode $\mathbf{q}^\circ(t)$. In this section, we reconstruct the quasi-periodic perturbation $\tilde{\mathbf{q}}$ (see eq. (6.14)) and investigate its dynamics, allowing us to better understand the physical origin of the low-frequency modulation. These dynamics give rise to what we refer to as a *generalized flutter* instability.

In fig. 6.3.8, we present the temporal evolution of the oscillating part of the perturbation, $\tilde{\mathbf{q}}$, defined by eq. (6.14), for $U^*/U_c^* = 1.100$. The different signals are represented on an interval of $17 T_0$, which allows to visualize one full oscillation at the modulation frequency ω (the period of the modulation is $T = T_0/(\phi/2\pi) \simeq 16.07 T_0$ at $U^*/U_c^* = 1.100$). In fig. 6.3.8(a), the signal of \tilde{h} (solid line) and $\tilde{\theta}$ (dashed line) are shown. They both possess the typical shape of a high-frequency oscillation at a frequency about ω_0 , modulated by the low frequency ω . However, the latter modulation is not phased for both signals so that they do not reach their modulation maxima simultaneously. More than that, the phase shift between both signals continuously drifts. This can be seen, for example, by taking the heaving signal as a reference and realizing that the pitching signal is continuously “getting ahead” (*i.e.* shifting to the left) of heaving. Alternatively, we show in fig. 6.3.8(b) a measure “per fast period” of the phase difference between heaving and pitching. More precisely, the $17T_0$ long signal is divided into 17 subintervals of length T_0 . On each of these, the phase difference is measured as the interval of time between the two closest maxima of $-\tilde{\theta}$ and \tilde{h} . By convention, the phase shift is defined in $[-\pi, \pi]$ and is positive




FIGURE 6.3.8: Time representation of the oscillating part, $\tilde{\mathbf{q}} = \mathbf{q}^\circ(t)e^{i\omega t} + \text{c.c.}$, of the perturbation generated by the pair of unstable Floquet modes at $U^*/U_c^* = 1.10$. (a) Heaving (solid line) and pitching (dashed line) signals. (b) Average phase shift between pitching and heaving over each fast period T_0 . (c) Power transmitted by the fluid to the solid. In all figures, the vertical dashed lines materialize the periods of maximal fluid energy extraction by the solid and maximal solid energy dissipation by the fluid.

(resp. negative) when pitching (resp. heaving) precedes heaving (resp. pitching). In the figure, it is clearly seen that the phase difference visits the whole interval $[-\pi, \pi]$ (two times, precisely) during the slow frequency oscillation. This continuous drift of the phase shift between heaving and pitching is made possible by the fact that both signals do not dominantly oscillate at the same frequency. This can be seen first by simply counting the number of local maxima in the heaving and pitching signals across the $17 T_0$ interval of fig. 6.3.8(a): they amount respectively to 16 and 18. A second indication of that fact can be deduced from the Fourier spectra of the Floquet mode, presented in fig. 6.3.7(c). Indeed, the harmonics of the quasi-periodic perturbation, $\widehat{\mathbf{q}}$, are directly linked to the ones of the Floquet mode by:

$$\widehat{\mathbf{q}}(v) = \begin{cases} \widehat{\mathbf{q}}_n^\circ & \text{if } v = n\omega_0 + \omega \\ \widehat{\mathbf{q}}_{-n}^{\circ*} & \text{if } v = n\omega_0 - \omega \\ 0 & \text{else} \end{cases} \quad (6.15)$$

As a consequence, by observing that $|\widehat{h}_{-1}^\circ| > |\widehat{h}_{+1}^\circ|$ in fig. 6.3.7(c) (squares), we deduce that $|\widehat{h}(\omega_0 - \omega)| > |\widehat{h}(\omega_0 + \omega)|$. Conversely, from $|\widehat{\theta}_{+1}^\circ| > |\widehat{\theta}_{-1}^\circ|$ (circles), we obtain that $|\widehat{\theta}(\omega_0 + \omega)| > |\widehat{\theta}(\omega_0 - \omega)|$. Both these analyses reflect the fact that the pitching signal dominantly vibrates at a slightly higher frequency, $\omega_0 + \omega$, than the heaving signal which vibrates at $\omega_0 - \omega$.

In fig. 6.3.8(c), we monitor the energy exchange between fluid and solid through the instantaneous power $\widetilde{\mathcal{P}}_f$ transmitted by the fluid forces to the solid. It is seen that depending on the phase inside the slow oscillation, the fluid either provides ($\widetilde{\mathcal{P}}_f > 0$) or dissipates ($\widetilde{\mathcal{P}}_f < 0$) energy from the solid. For example, between $2T_0$ and $3T_0$, the fluid always provides energy to the solid, whereas between $6T_0$ and $7T_0$, the fluid almost always dissipates solid energy. In average over the modulation period, the fluid provides energy to the solid. By comparing fig. 6.3.8(b) and (c), we notice that the extracted power is closely correlated to the phase difference. Indeed, the maximum energy extraction from the fluid ($t \in [2T_0, 3T_0]$) corresponds to a phase difference around $+\pi/2$, whereas the maximum solid energy dissipation ($t \in [6T_0, 7T_0]$) corresponds to a phase difference close to $-\pi/2$.

Let us further explore that link by zooming on the two particular subintervals $[2T_0, 3T_0]$ and $[6T_0, 7T_0]$ in fig. 6.3.9(a) and fig. 6.3.9(b), respectively. The time signal for \widehat{h} (solid line) and $-\widehat{\theta}$ (dashed line) are reproduced from fig. 6.3.8, along with snapshots of the instantaneous perturbation pressure field \widetilde{p} (blue isocontours) and perturbation displacement vectors (orange arrows). The black isoline demarcates the recirculation area of the base LCO ($\mathbf{u}_0 \cdot \mathbf{e}_x < 0$). We note that at each snapshot, the fields are represented in the deformed configuration corresponding to the LCO only. In other words, the position of the plate corresponds to $\mathbf{q}_0(t)$ alone. To reconstruct the full displacement of the perturbed solution $\mathbf{q}_0(t) + \epsilon \widetilde{\mathbf{q}}(t)$, one has to add to the represented position of the plate, the perturbation displacement direction materialized by the orange arrows (multiplied by the arbitrary constant ϵ). For the subinterval $[2T_0, 3T_0]$ (a), the perturbation displaces the plate around the LCO position with a movement where the pitching motion precedes the heaving motion by a phase difference around $+\pi/2$ (phase quadrature). From the pressure field, we observe that the flow mostly works positively on the plate. For example at $t = 2 T_0$ (snapshot (i)), one easily deduces from the pressure distribution on the plate that the moment (calculated around the midchord) is clockwise-oriented. At the same time, the pitching velocity is also clockwise-oriented (*cf* dashed line in the central figure), thus yielding positive work on the pitching degree of freedom. Similarly,

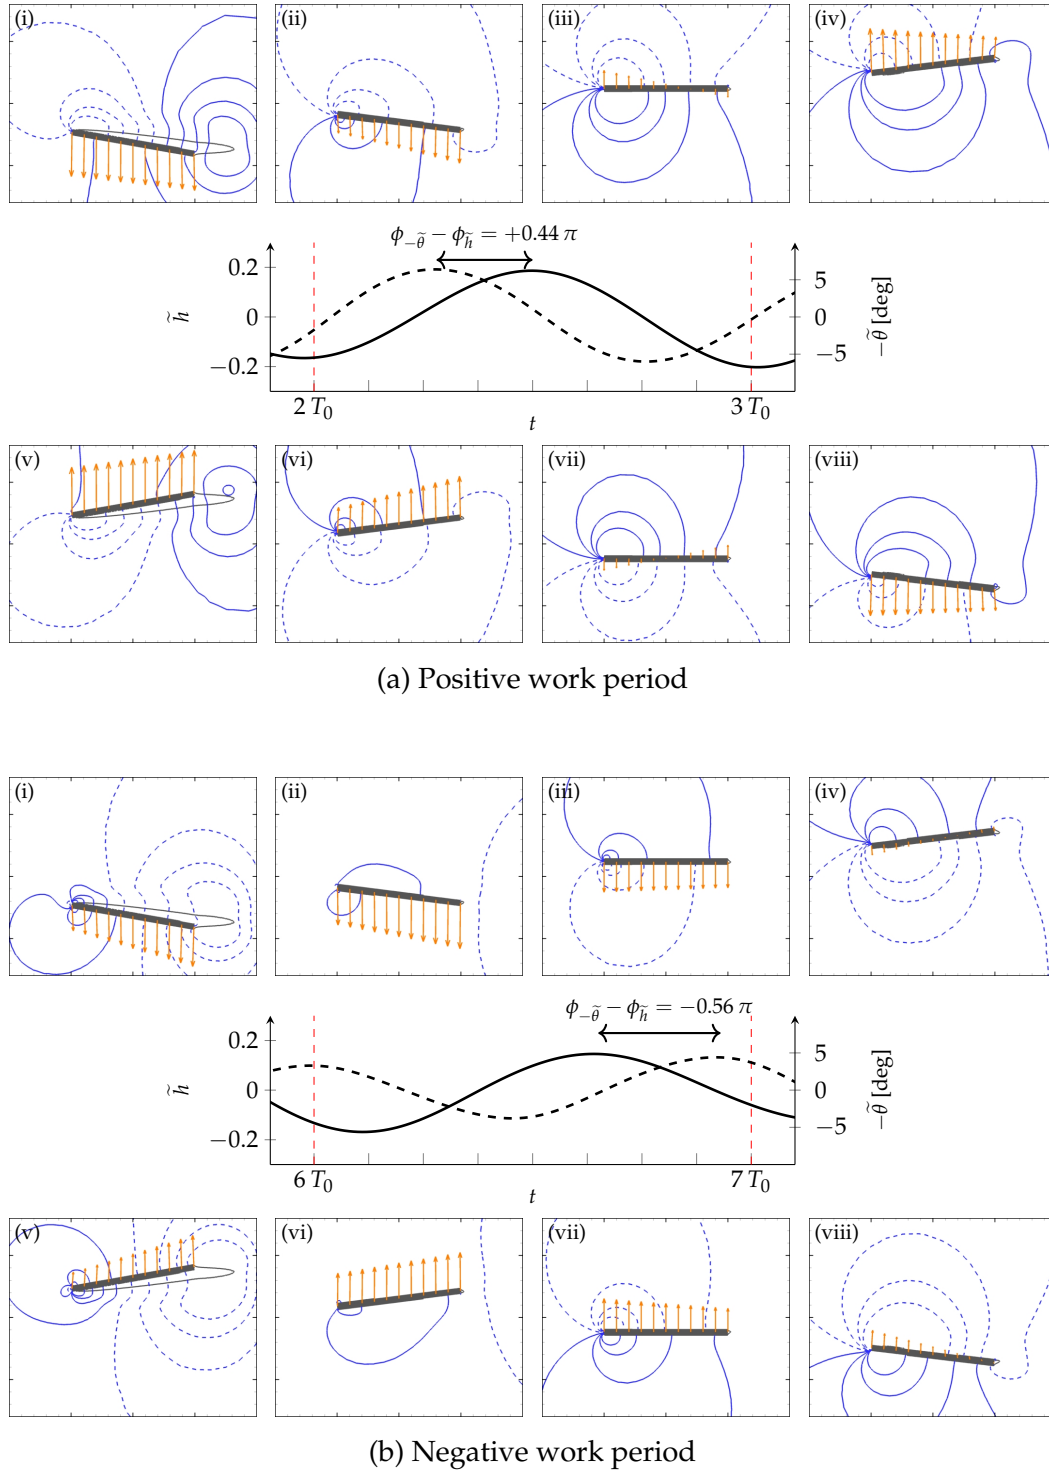


FIGURE 6.3.9: Temporal evolution of the quasi-periodic perturbation on two periods corresponding to (a) maximal fluid energy extraction by the solid and (b) maximal solid energy dissipation by the fluid, both indicated with vertical dashed lines in fig. 6.3.8. The heaving and pitching displacement are shown with solid and dashed lines, respectively. Snapshots (i) to (viii) are taken at the instants $t = nT_0$, $t = (n + 1/8)T_0$, ..., $t = (n + 7/8)T_0$ ($n = 2$ in (a) and $n = 6$ in (b)). They show the perturbation pressure field (blue contours, negative contours are dashed) and displacement vector (orange arrows). The thin black contour represents the recirculation zone of the LCO. The unperturbed position of the plate is displayed in these figures.

at $t = (2 + 1/8)T_0$ (snapshot (ii)) or $t = (2 + 2/8)T_0$ (snapshot (iii)), the lift force is clearly oriented upwards, while the heaving velocity is positive; the work on the heaving degree of freedom is then again positive. This intuitively explains why, on average, during the period $[2T_0, 3T_0]$ the solid is extracting energy from the fluid (*cf* fig. 6.3.8(c)). Overall, it is remarked the perturbation dynamics described here are highly similar to the one of the classical flutter mode that develops on steady solutions, and that we previously investigated in fig. 4.3.2a.

Turning now to the subinterval $[6T_0, 7T_0]$ in (b), the movement presents somehow opposite features. First, the heaving motion now precedes the pitching motion with a phase difference around $-\pi/2$. As a result, from the perturbation displacement vectors (orange arrows), we have the impression of a mode “the other way around” in comparison to (a), *i.e.* a classical flutter mode but where the flow would originate from the right. From the pressure fields we observe that the flow now mostly works against the plate’s movement. For example, at $t = (6 + 2/8)T_0$ (snapshot (iii)) or $t = (6 + 3/8)T_0$ (snapshot (iv)), the lift is clearly downwards, while the perturbation’s heaving velocity is oriented upwards (*cf* solid line in the central figure). The opposite situation occurs at $t = (6 + 7/8)T_0$ (snapshot (viii)) with upwards lift and downwards heaving motion. This analysis is in agreement with the average negative power of the fluid forces during the period $[6T_0, 7T_0]$ (*cf* fig. 6.3.8(c)). Here, we make the symmetric observation to the one made for (a): the perturbation dynamics on the interval $[6T_0, 7T_0]$ are very close to the (stable) so-called anti-flutter mode that was exhibited in the case of a steady baseflow in fig. 4.3.2b.

As a consequence of the above analysis of periods $[2T_0, 3T_0]$ and $[6T_0, 7T_0]$ we can argue that the dynamics of the unstable Floquet mode oscillates on the slow timescale between subintervals with flutter-like dynamics that provide energy to the solid from the flow, and subintervals with anti-flutter-like dynamics that dissipate solid energy in the flow. During the former, the solid motion amplitude increases, whereas it decreases during the latter, thus yielding a low-frequency amplitude modulation. In other words, the observed instability can be viewed as a kind of generalized coupled-mode flutter that, contrary to the classical one, develops on top of a previously established flutter LCO.

6.3.4 Comparison with time-marching results

After having studied in details the destabilization of the Floquet mode and the associated perturbation, we now wish to compare the Floquet predictions to the reference fully nonlinear time-marching solutions, previously presented in section 6.3.1.

In fig. 6.3.10(a), we start by comparing the pitching amplitude of the periodic orbits computed with TSM (thick lines) to the time-marching results (circles) previously exposed in fig. 6.3.4(a). As usual, we represent stable (*resp.* unstable) orbits with a solid (*resp.* dashed) line. The bottom horizontal line simply recalls the existence of an unstable steady solution (zero pitching amplitude). For low reduced velocities, where time-marching predicts periodic solutions, TSM and time-marching perfectly agree, as expected. For higher reduced velocities, the time-marching solutions become quasi-periodic in agreement with the destabilization of the TSM solution (dashed lines) occurring at $U^* = 1.078U_c^*$ according to the Floquet stability analysis (*cf* fig. 6.3.5). For such high velocities, the agreement between time-marching and TSM is necessary lost on the pitching amplitude as TSM cannot predict the global maximum reached by the plate across the slow modulation. For example, at $U^* = 1.165U_c^*$, the quasi-periodic solution visits the whole range $[5.5^\circ, 14.5^\circ]$, to be compared with the 12° predicted by TSM. Still, it can be considered that TSM




FIGURE 6.3.10: Comparison of Floquet analysis and time-marching results. (a) Bifurcation diagram for the pitching amplitude comparing the periodic orbits computed with TSM (thick lines) to the time-marching results (circles) previously shown in fig. 6.3.4(a). (b) Bifurcation diagram comparing the fundamental frequency ω_0 (upper curve) of the periodic orbits and the frequency of their leading Floquet mode ω (lower curve) to the dominant frequencies of the time-marching computations (circles). In both figures, stable (resp. unstable) LCO's are marked by solid (resp. dashed) lines.

predicts a “representative” pitching amplitude since the LCO amplitude is always included in the range visited by the quasi-periodic solution. In fig. 6.3.10(b), we compare the fundamental frequency ω_0 of the base LCO obtained with TSM (upper solid line) and the low-frequency ω introduced by the unstable Floquet mode (lower solid line) to the two fundamental frequencies of the time-marching solutions (circles). The latter are identified through a FFT analysis of the signal. First, it is observed that the high frequency in the time-marching solutions is in very good agreement with the TSM periodic orbits for all U^* , even when the periodic orbits are unstable (dashed line). This shows that the low-frequency modulation has a very small effect on the fast-frequency ω_0 that is mostly decided by the underlying periodic solution. This is in contrast with the pitching amplitudes shown in (a) that are clearly impacted by the modulation. The low-frequency ω is accurately predicted by the Floquet exponent for U^* close to the instability threshold, $U^* = 1.078U_c^*$, whereas the trend for growing U^* is not accurately captured. Similarly to what happens for the vortex shedding bifurcation around a circular cylinder (eg. [Sipp et al. 2007]), this deviation can be safely attributed to the nonlinear effects that come in play when moving away from the instability threshold.

Qualitative reconstruction of the time-marching solutions based on the Floquet mode We conclude this work by showing how superposing the base LCO \mathbf{q}_0 and the linear perturbation $\tilde{\mathbf{q}}$ generated by the unstable pair of Floquet modes compares to the fully nonlinear time-marching solution. First, we must point out an *a priori* inconsistency between the perturbation and the fully nonlinear quasi-periodic time-marching results. Indeed, whereas the perturbation $\tilde{\mathbf{q}}$ features two modulation amplitude maxima per slow period (see fig. 6.3.8(a)), the fully nonlinear solution has only one amplitude maximum per slow period (see fig. 6.3.1(b)). To understand and dissipate this apparent discrepancy, we need to reconstruct the perturbed solution $\mathbf{q}_0(t) + \epsilon\tilde{\mathbf{q}}(t)$, which is the sum of the base LCO and the (oscillating part of) the




FIGURE 6.3.11: Reconstruction of the full perturbed pitching solution $\theta_0(t) + \epsilon\theta'(t)$ (red line). For comparison, the LCO component $\theta_0(t)$ is shown with thin black solid line, whereas the perturbation component $\epsilon\theta'(t)$ is shown with thin black dashed line. The perturbation amplitude ϵ is arbitrary, for easy visualization.

perturbation. In theory, this quantity is — given an adequate amplitude ϵ — a first order approximation of the corresponding nonlinear quasi-periodic solution.

Results of that reconstruction are presented in fig. 6.3.11 where we show the pitching components for the LCO $\theta_0(t)$ (thin solid black line), the perturbation $\epsilon\tilde{\theta}(t)$ (thin dashed black line) and the reconstructed perturbed solution $\theta_0(t) + \epsilon\tilde{\theta}(t)$ (thick red line) over the same interval of length $17T_0$ as in fig. 6.3.8. Note that the perturbation amplitude ϵ is arbitrarily chosen for easy visualization. As a consequence, only the LCO solution amplitude has a physical meaning here. By comparing the reconstructed perturbed solution (red line) and the perturbation alone (dashed line) we immediately see that the reconstructed solution yields a signal with only one modulation amplitude maximum per slow period, in agreement with the time-marching results. Looking more closely at the signals, it is intriguing to observe that the maxima of modulation amplitude in the reconstructed solution (in red) do not necessarily match with the maxima of modulation amplitude in the perturbation alone (dashed). For example, around $t = 2T_0$, the perturbation is at a maximum of modulation amplitude while the reconstructed solution is close to a minimum. To better understand this feature, it is convenient to think of the reconstructed solution as the superposition of two oscillating signals, θ_0 and $\tilde{\theta}$, that dominantly oscillate at very close, but different, frequencies: ω_0 for θ_0 and $\omega_0 + \omega$ for $\tilde{\theta}$ (cf the discussion around eq. (6.15)). As a consequence, the phase difference of both signals continuously changes in time, which results in an interference-like phenomenon. When the waves are in-phase — as is the case at the seventh local maximum of the reconstructed solution (see vertical arrow in the figure) — the interference are constructive and the amplitude of the perturbed solution is higher than the amplitude of the LCO alone. On the contrary, when the waves are out-of-phase — as is the case at the fifteenth local maximum of the reconstructed solution (see vertical arrow in the figure) — the interference are destructive and the amplitude of the perturbed solution is lower than the amplitude of the LCO alone. In other words, the amplitude of the reconstructed solution depends only very partially on the amplitude of the perturbation alone; the key dependence is the phase between the base LCO and the perturbation.

In fig. 6.3.8(b), we showed that the perturbation $\tilde{\mathbf{q}}$ is characterized by a continuous drift of the phase difference between the heaving and pitching motions. This key feature was driven by the fact that pitching vibrates slightly faster than heaving.

We now wish to assess if these dynamics are retrieved in the fully nonlinear solutions. In fig. 6.3.12, we propose a comparison of the Fourier spectrum of the base LCO \mathbf{q}_0 (solid peaks) and Floquet perturbation $\tilde{\mathbf{q}}$ (dashed peaks) with respect to the fully nonlinear quasi-periodic time-marching solution (gray lines). We focus on a range of frequencies centered on ω_0 since the higher harmonics are negligible in the solid dynamics of the Floquet mode (see fig. 6.3.7(c)). Looking first at the spectra of the heaving motion (a), we make several observations. First, the amplitude of the ω_0 harmonic in the quasi-periodic time-marching solution is well predicted by the periodic TSM solution, despite the fact that it is unstable. Second, the dashed peaks $\omega_0 \pm \omega$ coming from the Floquet perturbation are slightly shifted with respect to the time-marching peaks. This is directly linked to the overestimation of the modulation frequency ω with Floquet analysis for reduced velocities as far from the threshold as $U^* = 1.10U_c^*$ (cf fig. 6.3.10(b)). Third, we notice that several additional peaks at frequencies $\omega_0 \pm n\omega$ are present in the time-marching solution while they are not predicted by the Floquet analysis. This is typically a manifestation of the nonlinear effects that saturate the growth of the Floquet mode away from the threshold. Similar observations can be made for the pitching signal presented in fig. 6.3.12(b). Finally, by comparing the $\omega_0 - \omega$ and $\omega_0 + \omega$ peaks in the time-marching heaving signal fig. 6.3.12(a), we observe that, aside from the ω_0 component inherited from the flutter LCO, the heaving dominantly vibrates on the $\omega_0 - \omega$ harmonic (the scale is logarithmic, the difference in amplitude between the $\omega_0 \pm \omega$ is about one order of magnitude). It is remarked that this hierarchy between the harmonics amplitudes is retrieved in the Floquet perturbation (squares). A symmetric observation is made in fig. 6.3.12(b) for the pitching signal for which both the time-marching solution and the Floquet perturbation agree on the fact that the dominant frequency is $\omega_0 + \omega$. These comparisons show that the difference in the heaving and pitching frequencies, that characterizes the linear Floquet perturbation $\tilde{\mathbf{q}}$, also characterizes the fully nonlinear quasi-periodic solutions. This arguments tends to show that the linear mechanism described in section 6.3.3, by which flutter-like and anti-flutter-like dynamics alternate during the low-frequency modulation, is present also in the fully nonlinear regime.




FIGURE 6.3.12: Comparison of the Fourier spectra of the base LCO \mathbf{q}_0 (solid peaks) and Floquet perturbation $\tilde{\mathbf{q}}$ (dashed peaks) with respect to the fully nonlinear quasi-periodic time-marching solution (gray lines) for the heaving (a) and pitching (b) degrees of freedom ($U^*/U_c^* = 1.10$). The amplitude ϵ of $\tilde{\mathbf{q}}$ is arbitrary.

6.4 Conclusion

This work was aimed at investigating the self-sustained motion of a spring-mounted plate in a laminar viscous flow, above the critical reduced velocity threshold U_c^* at which the classical flutter instability occurs. For the parameters of the study, the flutter bifurcation was supercritical (soft flutter) and gave rise to periodic solutions for $U_c^* < U^* < 1.078 U_c^*$. Above $1.078 U_c^*$, the solutions became quasi-periodic and were characterized by a slow amplitude modulation of the fast flutter oscillations. This transition from periodic to quasi-periodic solutions was investigated by analyzing the linear (asymptotic) Floquet stability of the periodic solutions that underlie the observed solution at all reduced velocities. In order to first compute these — possibly unstable — LCO's, and then assess their Floquet stability, we proposed an original approach based on the Time Spectral Method (TSM). For each reduced velocity, the periodic solution was first computed by solving the TSM problem (with unknown frequency). Then, the Floquet stability of the periodic solution was assessed by computing its leading Floquet exponents, that turn out to be the eigenvalues of the linearized TSM operator. Using this method, we showed that the appearance of quasi-periodic solutions is due to the destabilization of a pair of complex conjugate asynchronous Floquet modes. The analysis of the perturbation associated to this unstable pair of modes revealed that the low-frequency modulation was linked to the alternation of fast periods where the mean power of the fluid forces exerted on the solid is positive/negative. During the positive mean power periods, the linear perturbation dynamics were reminiscent of the unstable flutter mode, that develops classically on steady solutions (*cf* chapter 4), and that is characterized by a pitching-heaving phase shift around $+\pi/2$. On the contrary, during the negative mean power periods, the dynamics of the perturbation resembled the stable, so-called anti-flutter-like mode, for which the pitching-heaving phase shift is around $-\pi/2$. This alternation of flutter-like and anti-flutter-like dynamics was made possible by the fact that the heaving and pitching signals in the linear perturbation do not oscillate at the same frequency. As a consequence, they continuously drift with respect to each other, thus allowing a continuous change of the pitching-heaving phase shift. The last part of this chapter was dedicated to the comparison of the fully nonlinear results from time-marching simulations and the Floquet linear stability results. It was verified that Floquet stability accurately predicts both the reduced velocity threshold for the appearance of the quasi-periodic solutions and, close enough to the threshold, the frequency of the modulation. In addition, we observed in a fully nonlinear solution, far enough from the Floquet instability threshold, that the heaving and pitching signals present different dominant frequencies. This is in agreement with the analysis of the linear Floquet perturbation and tends to show that the mechanism described for the Floquet perturbation is also at play in the nonlinear regime and sustains the amplitude modulation.

III

TOWARDS LARGE-SCALE LINEAR STABILITY ANALYSIS

7

AUGMENTED LAGRANGIAN PRECONDITIONER FOR LARGE-SCALE HYDRODYNAMIC STABILITY ANALYSIS

7.1 Introduction

Over the past century, hydrodynamic linear stability theory was developed to understand the early stage of laminar-turbulence transition in parallel flows, such as boundary layers and shear flows [Drazin et al. 2004]. In the local stability theory, the growth or decay of perturbations developing on parallel flows, described with mono-dimensional velocity profiles, is investigated assuming a normal form decomposition. The resulting eigenproblem is of small size and does not require large computational resources to be solved. Although the local stability theory (mono-dimensional) was then extended to the description of spatially developing flows [Huerre et al. 1990], the linear stability analysis of truly two- and three-dimensional flows has gained in popularity since the beginning of the century [Theofilis 2003; Chomaz 2005; Bagheri et al. 2009b; Sipp et al. 2010; Theofilis 2011; Dijkstra et al. 2014; Loiseau et al. 2019] thanks to the development of computational resources and numerical tools allowing (a) to compute steady solutions of the governing equations and (b) to determine the most unstable eigenmodes of the linearized equations around this steady solution. An efficient and highly-parallel numerical tool is proposed in the present paper to achieve these two steps in the case of the incompressible Navier–Stokes equations.

Two main numerical approaches exist to carry out a linear stability analysis. The first one is the “time-stepping” [Loiseau et al. 2019] or “matrix-free” [Bagheri et al. 2009b] approach based on the use of existing unsteady nonlinear solvers, developed in Computational Fluid Dynamics (CFD). The “matrix-free” denomination indicates that the action of matrices onto vectors is obtained without assembling them. The unsteady solvers are adapted to compute steady solutions and to extract the eigenmodes of largest growth rate, relevant in linear hydrodynamic stability analysis. For computing steady (stable or unstable) solutions, stabilization procedures, such as the recursive projection method [Shroff et al. 1993], the selective frequency damping method [Åkervik et al. 2006], or more recently the BoostConv algorithm [Citro et al. 2017], are applied together with the unsteady nonlinear solver. The computation of leading eigenvalues is then achieved by noticing that the operations performed at

each iteration of the *linearized* time-stepping solver correspond to an exponential-based transformation of the Jacobian operator [Bagheri et al. 2009b; Loiseau et al. 2019]. Classical Krylov subspace-based methods like Arnoldi or Krylov–Schur are then commonly used to compute the eigenvalues of the exponential operator with largest magnitude, which are also the leading (rightmost) eigenvalues of the Jacobian operator. One of the advantages of this approach is the computational-time efficiency of applying one time-step of the unsteady solvers. Indeed, these solvers are often highly optimized, not only thanks to very scalable parallel implementation, but also because efficient numerical algorithms have been developed for solving the time-discretized problems (e.g., splitting [Karniadakis et al. 1991] or fractional step [Kim et al. 1985] methods). The drawback of this approach is the slow convergence of the Arnoldi method induced by the use of time-steppers. Indeed, small time-steps are required for an application of the linearized time-stepper to approximate accurately the exponential transformation [Tuckerman et al. 2000b]. This leads to a large number of so-called “outer” iterations (in the 10^3 – 10^4 range) to converge only a few eigenvalues. The efficiency of the “time-stepping” approach is thus mainly based on fast outer iterations at the expense of a large number of such iterations in the Arnoldi process. Note that other strategies for computing matrix exponential allow to relax the small time-step constraint and thus provide better convergence properties [Caliari et al. 2014; Rostami et al. 2018].

The second existing numerical approach to perform linear stability analysis in hydrodynamics [Sipp et al. 2010] is referred here to as the “matrix-based” approach. It relies on the assembly of sparse matrices resulting from the spatial discretization of the underlying problem and the solution of corresponding linear systems using existing parallel libraries that implement direct sparse *LU* factorization of those matrices (MUMPS [Amestoy et al. 2001], SuperLU [X. S. Li 2005]). The steady-state solutions are then computed by solving the steady nonlinear equations with a (quasi-)Newton method. An invert-based spectral transformation of the Jacobian operator, like the shift-and-invert [Christodoulou et al. 1988; Ehrenstein et al. 2005; Sartor et al. 2015; Sipp et al. 2010] or Cayley [Cliffe et al. 1993; Meerbergen et al. 1996; Mack et al. 2010] transformations, is then applied with a Krylov subspace-based method (typically, the Arnoldi method [Saad 1980]) to determine the leading eigenvalues. The Newton method and the shift-and-invert strategy allow, respectively, to achieve fast convergence towards the steady solution and the leading eigenvalues. Usually, the number of Newton iterations is around 10 or so to compute a steady solution, while it may require a few hundred outer iterations in the Arnoldi algorithm to compute a few eigenvalues. This reduced number of applications comes at a price: it requires the ability to invert the linearized steady (and generally shifted) Navier–Stokes equations. Consequently, it has mainly been used for linear stability analysis of two-dimensional flow configurations, for which the number of unknowns remains limited (not much greater than 10^5 unknowns) and the Jacobian matrices remain sparse, so that the *LU* factorization is affordable. The “matrix-based” strategy is particularly efficient for the eigenvalue computation, since the time-consuming *LU* factorization of the sparse Jacobian matrix is done once for all, while only the forward eliminations and back substitutions are repeated at each outer iteration. The main drawback of this approach is the large amount of memory needed to perform factorization, especially for three-dimensional flow configurations [Marquet et al. 2015a]. For large-scale hydrodynamics problems, the high cost of forming the Jacobian matrix explicitly, and the prohibitive memory requirements of direct solvers drove many authors [Bagheri et al. 2009a; Loiseau et al. 2014; Citro et al. 2015] towards the “matrix-free” strategy.

At the early beginning of nineties, Tuckerman [Tuckerman 1989; Mamun et al. 1994; Tuckerman et al. 2000b] proposed to improve the slow convergence of the “matrix-free” approach by using a Newton method (resp. a shift-and-invert strategy) for the steady-state (resp. eigenvalue) computation, while still using an existing unsteady solver. This method is based on the observation that one can adapt a (linearized) unsteady solver in order to apply, to some given vector, the steady Navier–Stokes Jacobian operator, left-preconditioned by the (unsteady) Stokes operator. Thus, this technique provides a cheap “matrix-free” way of preconditioning the Navier–Stokes Jacobian operator by the (unsteady) Stokes operator, for use inside Krylov subspace linear solvers typically. The method is nowadays known as the “Stokes” preconditioning technique and has been largely applied during the last decades for the computation of steady-state and leading eigenvalues [Barkley et al. 1997; Bergeon et al. 1998; Tuckerman et al. 2000a; Mercader et al. 2006; Tuckerman 2015]. Recently, it has been adapted and applied to the determination of resolvent modes in large-scale three-dimensional configurations [Brynjell-Rahkola et al. 2017]. In the Stokes preconditioning technique, the time-step of the linearized unsteady solver becomes a parameter of the preconditioner. Large time-steps usually provide better preconditioning, but make the application of one linearized time iteration harder. More details and improvements of the method can be found in [Beaume 2017]. In any case, the performance of this method remains limited by the efficiency of the (unsteady) Stokes operator to precondition the linearized steady Navier–Stokes operator.

In the present paper, we propose to develop a “matrix-based” specific solver for performing linear stability analysis, which relies on state-of-the-art preconditioners for the linearized steady incompressible Navier–Stokes equations, thus avoiding the use of direct solvers on the full problem. Over the last decades, various promising approaches have been developed aiming at overcoming the two main difficulties of this problem: the saddle-point structure of the equations deriving from the incompressibility constraint and the absence of a (small) time-step parameter that greatly enhances the convergence of iterative algorithms, due to the resulting diagonal dominance of the matrix. Among those steady preconditioners are the well-known SIMPLE preconditioner [Patankar et al. 1983], the more recent Pressure Convection–Diffusion (PCD) preconditioner proposed by [Kay et al. 2002], as well as the original augmented Lagrangian (AL) [Benzi et al. 2006; Benzi et al. 2011a; Heister et al. 2013] and modified augmented Lagrangian (mAL) [Benzi et al. 2011b; Benzi et al. 2011c; Benzi et al. 2011a; Benzi et al. 2013] preconditioners. Several authors showed the superiority of the modified augmented Lagrangian approach over other state-of-the-art alternatives for solving the Oseen and linearized incompressible Navier–Stokes equations [Segal et al. 2010; X. He et al. 2016]. Moreover, a very recent work [Farrell et al. 2018] proposed an efficient and highly scalable steady Navier–Stokes solver based on the original augmented Lagrangian preconditioner. If the augmented Lagrangian strategy has been regularly used for steady-state computations, it was never tested on practical case of eigenvalue computations. A work in that direction was however proposed by Olshanskii and Benzi [M. A. Olshanskii et al. 2008], who adapted the original augmented Lagrangian preconditioner to solve the shifted linearized Navier–Stokes equations. They showed theoretically and numerically that AL was robust to a real-valued shift on a variety of 2D flow configurations. Complex-valued shifts, as needed in practice to efficiently explore the complex plane with a shift-and-invert strategy, were not considered.

The first objective of the present paper is to assess the efficiency of the modified

augmented Lagrangian preconditioner for the computation of steady-state solutions with a Newton method and leading eigenvalues with a complex shift-and-invert strategy. The second objective is to describe, and test on a three-dimensional flow configuration, an open-source parallel implementation of the modified augmented Lagrangian preconditioner for linear stability analysis purposes, using the FreeFEM finite element library [Hecht 2012] interfaced with PETSc [Balay et al. 2019] and SLEPc [Hernandez et al. 2005]. The full code is made available at <https://github.com/prj-/moulin2019a1>.

The paper is organized as follows. The governing equations required to carry out the linear stability analysis of incompressible flows are introduced in section 7.2. The Newton method used to solve the steady nonlinear equations and the eigensolver based on the shift-and-invert strategy are also described. The preconditioning technique and the modified augmented Lagrangian preconditioner are introduced in section 7.3. The parallel implementation is detailed in section 7.4. Numerical results are given in section 7.5. First we examine, on a two-dimensional problem, the effect of various numerical and physical parameters on the performance of the mAL preconditioner for solving the complex shifted linearized Navier–Stokes problem. Then we compare the performance of mAL with other state-of-the-art preconditioners. Finally, we evaluate the performance of the proposed parallel implementation by first comparing it to a sparse direct solver on a small-scale three-dimensional test case and then, by testing its scalability on a large-scale configuration that cannot be afforded with a direct sparse solver.

7.2 Methods for linear stability analysis in hydrodynamics

7.2.1 Governing equations

Let us consider an incompressible flow, described by the two-dimensional (resp. three-dimensional) velocity field $\mathbf{u} = [u, v]^T$ (resp. $\mathbf{u} = [u, v, w]^T$) and the pressure field p , that satisfy the incompressible Navier–Stokes equations:

$$\frac{\partial \mathbf{u}}{\partial t} + (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p - \frac{1}{Re} \nabla^2 \mathbf{u} = 0, \quad -\nabla \cdot \mathbf{u} = 0$$

The Reynolds number is defined as $Re = U_\infty L / \nu$, where U_∞ and L are characteristic velocity and length used to make non-dimensional the velocity and pressure fields, and ν is the kinematic viscosity. For conciseness, the Navier–Stokes equations are rewritten in a state-space form as follows

$$\mathcal{M} \frac{\partial \mathbf{q}}{\partial t} + \mathcal{R}(\mathbf{q}) = 0, \quad \mathcal{M} = \begin{pmatrix} 1 & 0 \\ 0 & 0 \end{pmatrix}, \quad \mathcal{R}(\mathbf{q}) = \begin{pmatrix} (\mathbf{u} \cdot \nabla) \mathbf{u} + \nabla p - Re^{-1} \nabla^2 \mathbf{u} \\ -\nabla \cdot \mathbf{u} \end{pmatrix} \quad (7.1)$$

where $\mathbf{q} = (\mathbf{u}, p)^T$ is the state-space vector. Base flows, denoted hereinafter $\mathbf{q}_b(\mathbf{x})$, are time-independent (steady) solutions of the Navier–Stokes equations (7.1) and thus satisfy the nonlinear steady Navier–Stokes equations

$$\mathcal{R}(\mathbf{q}_b) = 0. \quad (7.2)$$

Linear stability of base flows is investigated by superimposing infinitesimal perturbations $\mathbf{q}'(\mathbf{x}, t)$ to the base flow solution, i.e. $\mathbf{q}(\mathbf{x}, t) = \mathbf{q}_b(\mathbf{x}) + \epsilon \mathbf{q}'(\mathbf{x}, t)$, where ϵ is an infinitesimal parameter. After inserting this decomposition into eq. (7.1), using the definition eq. (7.2) of the base flow and neglecting high-order terms in ϵ , one

obtains the linearized Navier–Stokes equations governing the temporal evolution of the infinitesimal perturbation,

$$\mathcal{M} \frac{\partial \mathbf{q}'}{\partial t} + \mathcal{J}(\mathbf{q}_b) \mathbf{q}' = 0, \quad \text{where} \quad \mathcal{J}(\mathbf{q}_b) = \begin{pmatrix} (\mathbf{u}_b \cdot \nabla)(\bullet) + (\bullet \cdot \nabla) \mathbf{u}_b - \mathcal{R}e^{-1} \nabla^2 \bullet & \nabla \bullet \\ -\nabla \cdot \bullet & 0 \end{pmatrix}$$

is the Jacobian operator defined around the base flow \mathbf{q}_b . The long-term evolution of any infinitesimal perturbation is conveniently described by assuming a spectral decomposition of perturbations as $\mathbf{q}' = \hat{\mathbf{q}}(\mathbf{x})e^{\sigma t} + \text{c.c.}$, where $\hat{\mathbf{q}}(\mathbf{x})$ is a complex spatial field whose temporal evolution is exponential and given by the complex number $\sigma = \lambda + i\omega$. λ is the growth rate and ω is the angular frequency. Inserting this modal decomposition into the above linearized equations shows that σ and $\hat{\mathbf{q}}$ are respectively eigenvalues and eigenmodes of the generalized eigenproblem:

$$\sigma \mathcal{M} \hat{\mathbf{q}} + \mathcal{J}(\mathbf{q}_b) \hat{\mathbf{q}} = 0. \quad (7.3)$$

The stability of the base flow is then determined by considering the leading eigenmode $\hat{\mathbf{q}}_0$ associated to the eigenvalue $\sigma_0 = \lambda_0 + i\omega_0$ with the largest real part λ_0 . When the growth rate of the leading eigenmode is negative ($\lambda_0 < 0$), all the eigenvalues have negative real parts, and the base flow is linearly stable since any perturbations superimposed to the base flow is damped at sufficiently large time. On the other hand, when the growth rate of the leading eigenmode is positive ($\lambda_0 > 0$), the perturbation will grow in time and the base flow is linearly unstable [Sipp et al. 2010].

A linear stability analysis thus consists first in computing a base flow, which is a solution of the steady Navier–Stokes eq. (7.2), and then in determining the leading eigenvalues/eigenmodes of the eigenproblem (7.3) with the largest growth rate.

In the present paper, a finite element method is used for the spatial discretization of the nonlinear steady equations (7.2) and of the linear eigenproblem (7.3) on a d -dimensional ($d = 2, 3$) domain Ω . A grad–div stabilized weak formulation [M. Olshanskii et al. 2009] of eq. (7.2) is used, which consists in finding \mathbf{u}_b in $\mathcal{V}_\Gamma = \{\mathbf{u} \in (\mathcal{H}^1(\Omega))^d, \text{ s.t. } \mathbf{u} = \mathbf{u}_\Gamma \text{ on } \Gamma\}$ and p_b in $\mathcal{Q} = \mathcal{L}^2(\Omega)$ such that:

$$\mathcal{R}_u(\mathbf{q}_b; \check{\mathbf{u}}) = \langle \mathbf{u}_b \cdot \nabla \mathbf{u}_b, \check{\mathbf{u}} \rangle + \langle \mathcal{R}e^{-1} \nabla \mathbf{u}_b, \nabla \check{\mathbf{u}} \rangle - \langle p_b, \nabla \cdot \check{\mathbf{u}} \rangle + \gamma \langle \nabla \cdot \mathbf{u}_b, \nabla \cdot \check{\mathbf{u}} \rangle = 0 \quad (7.4a)$$

$$\mathcal{R}_p(\mathbf{q}_b; \check{p}) = -\langle \nabla \cdot \mathbf{u}_b, \check{p} \rangle = 0 \quad (7.4b)$$

for all $(\check{\mathbf{u}}, \check{p})$ in $\mathcal{V}_0 \times \mathcal{Q}$, where $\langle \bullet, \bullet \rangle$ denotes the \mathcal{L}^2 inner-product and $\mathcal{V}_0 = \{\mathbf{u} \in (\mathcal{H}^1(\Omega))^d, \text{ s.t. } \mathbf{u} = 0 \text{ on } \Gamma\}$ is the velocity space with vanishing velocity on the boundary Γ . The weak residuals of the momentum and mass conservation equations are \mathcal{R}_u and \mathcal{R}_p , respectively. The last term in the momentum residual \mathcal{R}_u is the grad–div stabilization (also called augmentation) term that corresponds to the weak form of $-\gamma \nabla (\nabla \cdot \mathbf{u}_b)$, with $\gamma \geq 0$ a numerical parameter. In the above continuous weak formulation, the stabilization term strictly vanishes on the solution: $\langle \nabla \cdot \mathbf{u}_b, \nabla \cdot \check{\mathbf{u}} \rangle = 0$. Indeed, the divergence-free condition $\langle \nabla \cdot \mathbf{u}_b, \check{p} \rangle = 0$ is satisfied for all $\check{p} \in \mathcal{Q}$, and in particular for $\nabla \cdot \check{\mathbf{u}} \in \mathcal{Q}$.

A Delaunay triangulation of the domain $\mathcal{T}_h = \{K\}$, consisting in triangular ($d = 2$) or tetrahedral ($d = 3$) elements K , is used. In order to satisfy the inf–sup Ladyženskaja–Babuška–Brezzi (LBB) condition (see [Fortin et al. 1991]), the Taylor–Hood finite element pair is chosen, so that the discrete velocity \mathbf{u}_b^h and pressure p_b^h are sought respectively in $\mathcal{V}_\Gamma^h = \{\mathbf{u}^h \in \mathcal{C}^0(\Omega), \text{ s.t. } \mathbf{u}^h|_K \in \mathbb{P}_2(K), \forall K \in \mathcal{T}_h, \mathbf{u}^h =$

\mathbf{u}_Γ on Γ and $\mathcal{Q}^h = \{p^h \in \mathcal{C}^0(\Omega), \text{ s.t. } p^h|_K \in \mathbb{P}_1(K), \forall K \in \mathcal{T}_h\}$. Note that, with the Taylor–Hood finite element pair, the discrete divergence of the velocity test functions does not belong to the discrete pressure space, i.e. $\nabla \cdot \check{\mathbf{u}}^h \notin \mathcal{Q}^h$. Therefore, contrary to the continuous case, the stabilization term does not vanish from the discrete momentum equation ($\langle \nabla \cdot \mathbf{u}_b^h, \nabla \cdot \check{\mathbf{u}}^h \rangle \neq 0$), and the discrete solution depends on the value of the stabilization parameter γ . Here, the grad–div stabilization is mainly introduced to improve, thanks to an efficient preconditioner (see section 7.3), the iterative solution of linear systems involved when solving the nonlinear discrete equations (7.5). The question of whether the grad–div stabilized discrete solution is closer or further from the continuous weak solution is out of the scope of this paper. However, several studies (e.g. [M. A. Olshanskii 2002; M. A. Olshanskii et al. 2004; M. Olshanskii et al. 2009; Linke et al. 2011; Heister et al. 2013]) showed that the grad–div stabilization often improves the mass conservation property and the velocity error of the discrete solution, for adequate values of γ . Numerical experiments are performed in section 7.5.2.1 to assess the accuracy of the stabilized discrete solution and to determine adequate values of γ .

Such a discretization yields the following discrete version of the nonlinear base flow eq. (7.2):

$$\mathbf{R}(\mathbf{q}_b^h) = \mathbf{0} \quad (7.5)$$

where \mathbf{q}_b^h denotes now the vector of coefficients of \mathbf{u}_b^h and p_b^h in the finite elements basis.

The generalized eigenproblem (7.3) is discretized similarly, yielding

$$\sigma \mathbf{M} \hat{\mathbf{q}}^h + \mathbf{J}(\mathbf{q}_b^h) \hat{\mathbf{q}}^h = 0 \quad (7.6)$$

where \mathbf{M} and $\mathbf{J}(\mathbf{q}_b^h)$, the finite element matrices obtained after discretization of the mass \mathcal{M} and Jacobian operator $\mathcal{J}(\mathbf{q}_b^h)$, are respectively defined as

$$\mathbf{M} = \begin{pmatrix} \mathbf{M}_u & 0 \\ 0 & 0 \end{pmatrix} \quad \text{and} \quad \mathbf{J}(\mathbf{q}_b^h) = \begin{pmatrix} \mathbf{A}_\gamma & \mathbf{B}^T \\ \mathbf{B} & 0 \end{pmatrix}. \quad (7.7)$$

The rectangular matrix \mathbf{B} is the discretization of the divergence operator and its transpose \mathbf{B}^T represents the discrete gradient. The mass matrix on the velocity space \mathbf{M}_u can be written as a 3-by-3 block diagonal matrix corresponding to the three velocity components. The 3-by-3 block matrix $\mathbf{A}_\gamma = \mathbf{A} + \gamma\Gamma$ is the sum of \mathbf{A} , which represents the linearized diffusion and convection terms in the momentum conservation equation, and Γ , obtained after discretization of the grad–div stabilization term. They write:

$$\mathbf{A}_\gamma = \begin{pmatrix} \mathbf{A}_{uu} + \gamma\Gamma_{uu} & \mathbf{A}_{uv} + \gamma\Gamma_{uv} & \mathbf{A}_{uw} + \gamma\Gamma_{uw} \\ \mathbf{A}_{vu} + \gamma\Gamma_{vu} & \mathbf{A}_{vv} + \gamma\Gamma_{vv} & \mathbf{A}_{vw} + \gamma\Gamma_{vw} \\ \mathbf{A}_{wu} + \gamma\Gamma_{wu} & \mathbf{A}_{wv} + \gamma\Gamma_{wv} & \mathbf{A}_{ww} + \gamma\Gamma_{ww} \end{pmatrix}, \quad \mathbf{M}_u = \begin{pmatrix} \mathbf{M}_u & 0 & 0 \\ 0 & \mathbf{M}_v & 0 \\ 0 & 0 & \mathbf{M}_w \end{pmatrix}. \quad (7.8)$$

In the following, we will mostly refer to the discrete solutions. Therefore, the superscript h is dropped unless confusion is possible.

7.2.2 Nonlinear steady-state solver

The nonlinear solution \mathbf{q}_b of the discrete problem eq. (7.5) is obtained by the classical Newton method. The approximated solution at the k th iteration is obtained as

$$\mathbf{q}_b^k = \mathbf{q}_b^{k-1} + \delta \mathbf{q}_b^k, \quad (7.9)$$

where $\delta \mathbf{q}_b^k$ denotes the solution increment, obtained by solving the linear problem

$$\mathbf{J}(\mathbf{q}_b^{k-1}) \delta \mathbf{q}_b^k = -\mathbf{R}(\mathbf{q}_b^{k-1}) \quad (7.10)$$

where $\mathbf{J}(\mathbf{q}_b^{k-1})$ is the Jacobian matrix defined in eq. (7.7) with the known approximation of the steady solution \mathbf{q}_b^{k-1} . The solution of this linear system is repeated for each iteration of the Newton algorithm, that is considered to be converged when the l^2 norm of the residual $\|\mathbf{R}(\mathbf{q}_b)\|_2$ is below some numerical tolerance.

7.2.3 Linear eigensolver

The Krylov–Schur algorithm [Stewart 2002] is used in the present study to solve the generalized eigenproblem (7.6). In order to compute the leading eigenvalues, which lie in the complex plane close to the zero growth-rate axis ($\lambda = 0$) for any frequency ω , a shift-and-invert spectral transformation is first applied, yielding the transformed eigenproblem

$$\mu \hat{\mathbf{q}} + \mathbf{T} \hat{\mathbf{q}} = 0, \quad \mathbf{T} = (\mathbf{J}(\mathbf{q}_b) + s\mathbf{M})^{-1} \mathbf{M} \quad (7.11)$$

where $s = s_r + is_i$ is the so-called complex shift. The eigenvalues μ of the transformed problem are related to the eigenvalues σ of eq. (7.6) through $\mu = (\sigma - s)^{-1}$ while the eigenvectors are left unchanged by the spectral transformation. Like the classical power method, the Krylov–Schur algorithm allows to compute the eigenvalues of largest magnitude. When applied to the transformed problem, it gives the eigenvalues μ of largest magnitude, which correspond to the eigenvalues σ closest to the complex shift s . To determine the leading eigenvalue of eq. (7.6), the eigenproblem (7.11) is solved for several values of the complex s close to the real axis, spanning appropriately the imaginary axis. For each eigenvalue computation, the Krylov–Schur algorithm requires multiple “matrix–vector” applications of the matrix \mathbf{T} . In other words, repeated solutions of the linear system $(\mathbf{J}(\mathbf{q}_b) + s\mathbf{M}) \mathbf{q}_o = \mathbf{q}_i$ are required, where the right-hand side vectors \mathbf{q}_i are given by the Krylov–Schur algorithm.

In the present work, linear stability analysis is thus performed using a nonlinear steady-state solver and a linear eigensolver, that both rely on multiple solutions of linear systems involving the complex shifted Jacobian matrix $(\mathbf{J} + s\mathbf{M})$. For the steady-state solver, this matrix reduces to the real Jacobian matrix \mathbf{J} as the complex shift vanishes $s = 0$. The next section introduces a preconditioned iterative method used to solve efficiently such systems.

7.3 An augmented Lagrangian approach for the shifted Jacobian matrix

As explained in the previous section, the main challenge of an hydrodynamic stability analysis is to solve efficiently the following linear equation:

$$\begin{pmatrix} \mathbf{A}_{\gamma,uu} + s\mathbf{M}_u & \mathbf{A}_{\gamma,uv} & \mathbf{A}_{\gamma,uw} & \mathbf{B}_u^T \\ \mathbf{A}_{\gamma,vu} & \mathbf{A}_{\gamma,vv} + s\mathbf{M}_v & \mathbf{A}_{\gamma,vw} & \mathbf{B}_v^T \\ \mathbf{A}_{\gamma,wu} & \mathbf{A}_{\gamma,wv} & \mathbf{A}_{\gamma,ww} + s\mathbf{M}_w & \mathbf{B}_w^T \\ \mathbf{B}_u & \mathbf{B}_v & \mathbf{B}_w & 0 \end{pmatrix} \begin{pmatrix} u_o \\ v_o \\ w_o \\ p_o \end{pmatrix} = \begin{pmatrix} u_i \\ v_i \\ w_i \\ p_i \end{pmatrix} \quad (7.12)$$

where $\mathbf{A}_{\gamma,\alpha\beta} = \mathbf{A}_{\alpha\beta} + \gamma\Gamma_{\alpha\beta}$ ($\alpha, \beta = u, v, w$), $\mathbf{q}_o = (u_o, v_o, w_o, p_o)^T$ is the solution vector and $\mathbf{q}_i = (u_i, v_i, w_i, p_i)^T$ is a right-hand side vector. In the perspective of large-scale computations, we must avoid the use of direct solvers applied directly to eq. (7.12), due to their huge memory cost [Marquet et al. 2015a]. Instead, we use the flexible Generalized Minimal Residual algorithm (GMRES) [Saad 1993] for solving iteratively eq. (7.12). The shifted-Jacobian matrix being indefinite and ill-conditioned, the use of an iterative method without preconditioning is inefficient as it requires a very large number of iterations [Tuckerman 2015]. To improve the numerical efficiency of the iterative solution, the above linear system is replaced by the right-preconditioned linear system:

$$\begin{pmatrix} \mathbf{A}_{\gamma,uu} + s\mathbf{M}_u & \mathbf{A}_{\gamma,uv} & \mathbf{A}_{\gamma,uw} & \mathbf{B}_u^T \\ \mathbf{A}_{\gamma,vu} & \mathbf{A}_{\gamma,vv} + s\mathbf{M}_v & \mathbf{A}_{\gamma,vw} & \mathbf{B}_v^T \\ \mathbf{A}_{\gamma,wu} & \mathbf{A}_{\gamma,wv} & \mathbf{A}_{\gamma,ww} + s\mathbf{M}_w & \mathbf{B}_w^T \\ \mathbf{B}_u & \mathbf{B}_v & \mathbf{B}_w & 0 \end{pmatrix} \mathbf{P}^{-1} \begin{pmatrix} \tilde{u}_o \\ \tilde{v}_o \\ \tilde{w}_o \\ \tilde{p}_o \end{pmatrix} = \begin{pmatrix} u_i \\ v_i \\ w_i \\ p_i \end{pmatrix} \quad (7.13)$$

where the matrix \mathbf{P} is the so-called preconditioner and $(\tilde{u}_o, \tilde{v}_o, \tilde{w}_o, \tilde{p}_o)^T$ is an intermediate solution. The final solution is found by solving the following linear system

$$\mathbf{P} \begin{pmatrix} u_o \\ v_o \\ w_o \\ p_o \end{pmatrix} = \begin{pmatrix} \tilde{u}_o \\ \tilde{v}_o \\ \tilde{w}_o \\ \tilde{p}_o \end{pmatrix} \quad (7.14)$$

The GMRES algorithm is applied to the right-preconditioned eq. (7.13) which, in addition to matrix–vector products with the shifted-Jacobian matrix, requires the repeated application of \mathbf{P}^{-1} , i.e., the solution of eq. (7.14). A good preconditioner achieves a compromise between a fast application of the preconditioner and a small number of iterations to solve the preconditioned system.

The augmented Lagrangian preconditioner allows to solve iteratively the Oseen [Benzi et al. 2006; Benzi et al. 2011b; Benzi et al. 2011a; Heister et al. 2013] and linearized Navier–Stokes equations [Benzi et al. 2011c; M. A. Olshanskii et al. 2008] in a very limited number of iterations, regardless of the mesh refinement and the Reynolds number value. Nevertheless, these interesting properties are counterbalanced by the difficulty of solving iteratively the coupled (two or three-dimensional) velocity subproblem arising in the application of the original preconditioner, as it requires highly specific multigrid solvers [Benzi et al. 2006; Farrell et al. 2018]. In order to circumvent this particular issue, the so-called modified augmented Lagrangian (mAL) preconditioner was introduced in [Benzi et al. 2011b]. It is derived from the original augmented Lagrangian preconditioner by neglecting either the lower block matrices [Benzi et al. 2011b] or the upper block matrices [X. He et al. 2016], as follows

$$\mathbf{P}_{\text{mAL}} = \begin{pmatrix} \mathbf{A}_{\gamma,uu} + s\mathbf{M}_u & 0 & 0 & 0 \\ \mathbf{A}_{\gamma,vu} & \mathbf{A}_{\gamma,vv} + s\mathbf{M}_v & 0 & 0 \\ \mathbf{A}_{\gamma,wu} & \mathbf{A}_{\gamma,wv} & \mathbf{A}_{\gamma,ww} + s\mathbf{M}_w & 0 \\ \mathbf{B}_u & \mathbf{B}_v & \mathbf{B}_w & \mathbf{S}_p \end{pmatrix}, \quad (7.15)$$

where \mathbf{S}_p is an approximation of the pressure Schur complement $-\mathbf{B}(\mathbf{A}_{\gamma} + s\mathbf{M}_u)^{-1}\mathbf{B}^T$.

Rather than being explicitly specified, this matrix is defined by the action of its inverse as

$$\mathbf{S}_p^{-1} = -(\gamma + \mathcal{R}e^{-1})\mathbf{M}_p^{-1} - s\mathbf{L}_p^{-1}, \quad (7.16)$$

where \mathbf{M}_p is the mass matrix and \mathbf{L}_p the Laplacian matrix, both defined on the discrete pressure space. Note that, for base flow computations $s = 0$, only the first term remains in the definition of the approximated Schur complement eq. (7.16). The lower block-triangular version of the preconditioner is chosen for practical reasons explained in section 7.4. In the original preconditioner proposed by [Benzi et al. 2006; Benzi et al. 2011b], the augmentation term of the Jacobian matrix $\mathbf{\Gamma}$ was defined algebraically as $\mathbf{B}^T\mathbf{M}_p^{-1}\mathbf{B}$, thus requiring two sparse matrix products to be constructed explicitly. As proposed in [Heister et al. 2013], the construction cost can significantly be reduced by building the matrix $\mathbf{\Gamma}$ from the finite element discretization of the grad-div stabilization term. They showed that $\mathbf{\Gamma}$ is then the sum of the algebraic augmentation $\mathbf{B}^T\mathbf{M}_p^{-1}\mathbf{B}$ and of a stabilization matrix. The efficiency of the mAL preconditioner is thus conserved while significantly reducing the construction costs. Finally, the grad-div augmentation matrix $\mathbf{\Gamma}$ is much more sparse than its algebraic counterpart, thus motivating our choice for this implementation of the mAL preconditioner, in the perspective of large-scale three-dimensional computations.

7.4 Parallel implementation with FreeFEM and its interface to PETSc/SLEPc

For realistic three-dimensional geometries, the approach derived in the previous sections requires the solution of nonlinear systems and generalized eigenproblems of large dimensions. Thus, high-performance computing becomes necessary. The goal of this section is to show how this is done using a finite element domain specific language, FreeFEM [Hecht 2012; Jolivet et al. 2012], interfaced with distributed linear algebra backends, PETSc [Balay et al. 2019] and SLEPc [Hernandez et al. 2005]. A thorough introduction of these libraries may be found in their respective manuals^{1,2,3}. Our implementation is openly available at <https://github.com/prj-/moulin2019a1> and the rest of this section follows the available source code.

7.4.1 Outer solvers

In this section, we describe how the outer solvers (i.e. the nonlinear steady-state solver and the eigensolver) are implemented.

The Newton method described in section 7.2.2 is implemented using FreeFEM. Only the inversion of the Jacobian matrix \mathbf{J} of eq. (7.10) is performed by PETSc. Given a FreeFEM distributed version of the Jacobian matrix $d\mathbf{J}$, PETSc options defining the linear solver for eq. (7.10) are set using the following FreeFEM syntax:

```
set(dJ, sparams = params, fields = vX[], names = names,
    schurPreconditioner = S, schurList = listX[]);
```

The keyword `sparams` is a string defined by the user gathering the PETSc runtime options for the Krylov subspace solver (KSP) and preconditioner (PC). The interested reader should refer to PETSc manual for details on the use of runtime options. The keywords `fields`, `names`, `schurPreconditioner`, and `schurList` allow to

¹<https://doc.freefem.org/pdf/FreeFEM-documentation.pdf>

²<http://www.mcs.anl.gov/petsc/petsc-current/docs/manual.pdf>

³<http://slepc.upv.es/documentation/slepc.pdf>

implement specific block preconditioners, like mAL, and their use is detailed in the next sections.

For the eigenvalue computation presented in section 7.2.3, only the finite element matrices are built by FreeFEM. Then, the Krylov–Schur algorithm is performed entirely by SLEPc through the use of the eigenvalue problem solver framework (EPS), which is called from within FreeFEM.

```
int k = EPSSolve(dJ, dM, vectors = vec, values = val, sparams = )
    ↵ params,
    fields = vX[], names = names, schurPreconditioner = S, )
    ↵ schurList = listX[]);
// solves the eigenvalue problem  $dJ\hat{q} = \sigma dM\hat{q}$ 
```

Contrary to the case of the linear solver interface, two matrices dJ and dM that define the generalized eigenproblem (7.6) must now be passed to SLEPc. In addition, sparams must also contain the SLEPc runtime options defining the eigensolver.

7.4.2 Inner mAL-preconditioned linear solvers

The inner linear solves of system eq. (7.12) with a mAL-preconditioned GMRES require the implementation of the block structure of the preconditioner (eq. (7.15)). This is done in PETSc by using the so-called fieldsplit structure that gives to the users a high-level of abstraction to define operators by blocks. The following PETSc runtime options define such a preconditioner:

```
string params = paramsXYZ + " " + paramsP + " " + paramsKrylov +
    "-pc_type fieldsplit -pc_fieldsplit_type multiplicative";
```

The desired lower block-triangular structure of the preconditioner is obtained by the use of PETSc keyword `multiplicative`. The strings `paramsXYZ` and `paramsP` respectively contain the innermost velocity and pressure block solvers options that will be detailed later on. The string `paramsKrylov` contains the definition of the Krylov subspace linear solver. For example, one should simply write `paramsKrylov ↵ ↵ = "-ksp_type fgmres"` to use the flexible GMRES. In order to implement the modified augmented Lagrangian preconditioner \mathbf{P}_{mAL} through the fieldsplit structure, in 3D, the four fields u, v, w , and p must be defined. Assuming the problem is formulated in the full vectorial finite element space \mathcal{W}_h , containing the velocities and pressure unknowns, one must be able to differentiate the degrees of freedom belonging to each field. To that aim a finite element function taking a different integer value for each one of the four fields is defined in FreeFEM and passed to PETSc/SLEPc through the keyword `fields`. Then, for simplicity, each field is attributed a name that will be used to identify it when defining the different innermost solvers associated to the diagonal blocks of \mathbf{P}_{mAL} , c.f. section 7.4.3. Those names are contained in an array of strings, that is provided to the solver through the keyword `names`.

```
Wh [vX, vY, vZ, p] = [1, 2, 3, 4]; // numbering of each field
string[int] names(4); // prefix of each field
names[0] = "vX"; // x-velocity
names[1] = "vY"; // y-velocity
names[2] = "vZ"; // z-velocity
names[3] = "p"; // pressure
```

Approximate Schur complements The default setting in PETSc, when using a multiplicative fieldsplit preconditioner, is to define the preconditioner as the lower block triangular part of the system matrix in eq. (7.12). Thus, on the block diagonal of such a preconditioner, one would have $\mathbf{A}_{\gamma,uu} + s\mathbf{M}_u$, $\mathbf{A}_{\gamma,vv} + s\mathbf{M}_v$, $\mathbf{A}_{\gamma,ww} + s\mathbf{M}_w$ and the null matrix. In order to implement \mathbf{P}_{mAL} , one must replace, in the preconditioner only, the default operator for the pressure field (the null matrix) by the ones necessary to implement the desired Schur complement approximation eq. (7.16). This is done in PETSc using `PCFieldSplitGetSubKSP` to retrieve the operators linked to each field of the fieldsplit structure and then `KSPSetOperators` to set the new operators that define eq. (7.16).

When the shift s is null, for base flow computations, the approximate Schur complement only requires the assembly of one operator: $(\gamma + \mathcal{R}e^{-1})^{-1}\mathbf{M}_p$. This is done in FreeFEM as shown below, and then passed to PETSc with the keyword `schurPreconditioner`.

```
matrix[int] S(1); // array with a single matrix
varf vSchur(p, q) = int3d(th, qforder = 3)
  (-1.0/(gamma + 1.0/Re) * p * q); // eq. (7.16) with s = 0
S[0] = vSchur(Qh, Qh); // matrix assembly
```

For eigenvalue computations, two auxiliary operators are now needed: $(\gamma + \mathcal{R}e^{-1})^{-1}\mathbf{M}_p$ and $s^{-1}\mathbf{L}_p$. The construction in FreeFEM is performed using the following lines, and then again passed to SLEPc with the keyword `schurPreconditioner`.

```
matrix<complex>[int] S(2); // array with two matrices
complex scale;
varf vMp(p, q) = int3d(th, qforder = 3)(scale * p * q); // eq. (7.7)
scale = 1.0/(gamma + 1.0/Re);
S[0] = vMp(Qh, Qh); // first matrix assembly
macro grad(p)[dx(p), dy(p), dz(p)]// macro for computing  $\nabla p$ 
varf vLp(p, q) = on(3, p = 1) // inlet boundary condition
  + int3d(th, qforder = 2)(scale * (grad(p)' * grad(q)));
// shift value s
complex s = getARGV("-shift_real", 1.0e-6) +
  getARGV("-shift_imag", 0.6) * 1i;
scale = 1.0/s;
S[1] = vLp(Qh, Qh); // second matrix assembly
```

Finally, we note that the operators needed for the Schur complement approximation are built on the pressure space Q_h . However, in FreeFEM, it is not possible to know a priori the correspondence between the numbering of W_h , where the full solution is defined, and Q_h . To circumvent this issue, we compute this correspondence in FreeFEM from an interpolation between Q_h and W_h , and then pass it to PETSc/SLEPc with the keyword `schurList`:

```
Qh pIdx; // function from the pressure space
pIdx[] = 1:pIdx[].n; // numbering of the unknowns of Qh
// renumbering into the complete space by doing an interpolation
  on Wh
Wh [listX, listY, listZ, listP] = [0, 0, 0, pIdx];
```

7.4.3 Innermost velocity and pressure linear solvers

The approximate inverse of the diagonal blocks in eq. (7.15) are defined using off-the-shelf iterative methods from PETSc. For each velocity field, the GMRES is used, right-preconditioned by an additive Schwarz method (ASM) with one-level of algebraic overlap, as well as exact *LU* factorizations for each subdomain solver. These factorizations are carried out by MUMPS [Amestoy et al. 2001]. A maximum Krylov dimension of 50 is prescribed and the GMRES is stopped when the relative unpreconditioned residual norm is lower than 10^{-1} . In our implementation, the PETSc runtime options defining the approximate inverse of the diagonal velocity blocks are contained in the string `paramsXYZ` detailed below:

```
real tolV = getARGV("-velocity_tol", 1.0e-1); // default to 10-1
// monodimensional velocity solver
string paramsV = "-ksp_type gmres -ksp_converged_reason )
    ⤵ -ksp_pc_side right " +
    "-ksp_rtol " + tolV + " -ksp_gmres_restart 50 -pc_type asm " +
    "-pc_asm_overlap 1 -sub_pc_type lu )
    ⤵ -sub_pc_factor_mat_solver_type mumps";
// each velocity component gets the same monodimensional solver
// defined by paramsV
string paramsXYZ = "-prefix_push fieldsplit_vX_ " + paramsV + " )
    ⤵ -prefix_pop"
    + " -prefix_push fieldsplit_vY_ " + paramsV + " )
    ⤵ -prefix_pop"
    + " -prefix_push fieldsplit_vZ_ " + paramsV + " )
    ⤵ -prefix_pop";
```

For the pressure Schur complement approximate inverse eq. (7.16), the PETSc runtime options defining the solver are contained in the string `paramsP`. We must distinguish the cases of the base flow and eigensolvers. For the former ($s = 0$), only the action of \mathbf{M}_p^{-1} has to be evaluated. For that purpose, we use at most five iterations of the Jacobi-preconditioned conjugate gradient algorithm:

```
string paramsP = "-prefix_push fieldsplit_p_ " +
    "-ksp_type cg -ksp_max_it 5 -pc_type jacobi -prefix_pop";
```

For the eigensolver ($s \neq 0$), the action of the inverse of the Schur complement is approximated by the sum of the action of $(\gamma + \mathcal{R}e^{-1})\mathbf{M}_p^{-1}$ and $s\mathbf{L}_p^{-1}$. This is done through PETSc composite preconditioner:

```
string paramsP = "-prefix_push st_fieldsplit_p_ " +
    "-ksp_type preonly -pc_type composite " +
    "-prefix_push sub_0_ " + // action of  $(\gamma + \mathcal{R}e^{-1})\mathbf{M}_p^{-1}$ 
    "-pc_type bjacobi -sub_pc_type icc -prefix_pop " +
    "-prefix_push sub_1_ " + // action of  $s\mathbf{L}_p^{-1}$ 
    "-pc_type gamg -pc_gamg_square_graph 10 -prefix_pop " +
    "-prefix_pop";
```

Here only one application of the block Jacobi preconditioner with ICC(0) sub-solvers [Chan et al. 1997] is used for approximating the mass matrix inverse while one V-cycle of GAMG [Adams et al. 2004] is used for the Laplacian term.

7.5 Numerical results

The efficiency of the modified augmented Lagrangian (mAL) preconditioner is investigated in this section by performing the linear stability analysis of two- and three-dimensional flow configurations described in section 7.5.1. The two-dimensional computations are always performed on one process as they are of limited size. For the three-dimensional case, the fully parallel implementation presented in section 7.4 is used. The influence of various numerical and physical parameters, such as the augmentation parameter, the mesh size and the Reynolds number, is first assessed in section 7.5.2 for the two-dimensional configuration, before comparing the performance of mAL preconditioner with other block preconditioners (PCD, SIMPLE, Stokes) in section 7.5.3. The efficiency of the parallel implementation is finally investigated in section 7.5.4 for the three-dimensional configuration.

7.5.1 Two- and three-dimensional test cases

The two-dimensional flow configuration is sketched in fig. 7.5.1a. A thin plate of height h and thickness $t = h/6$ is immersed in an incoming flow of uniform velocity U_∞ . The size of the computational box indicated in the figure and the flow variables are made non-dimensional using h as characteristic length and U_∞ as characteristic velocity, so that the Reynolds number is defined as $Re = U_\infty h/\nu$, where ν is the kinematic viscosity.

Triangulations of the computational domain are obtained with the internal mesh generator of FreeFEM. The no-slip boundary condition $u = v = 0$ is applied on the plate, symmetry boundary conditions ($\partial_y u = 0$ and $v = 0$) are applied at the top and bottom boundaries of the computational domain, and a stress-free boundary condition is applied at the outlet boundary.

A typical steady solution of the incompressible Navier–Stokes equation is displayed in fig. 7.5.2a for $Re = 40$. The flow recirculates in two symmetric regions in the wake of the plate, as indicated by the streamlines. The linear stability analysis of this base flow yields the eigenvalue spectrum shown in fig. 7.5.2c with circles. A pair of complex conjugate unstable eigenvalues is found ($\lambda \geq 0$) characterized by an angular frequency $\omega = 0.70$. For a lower value of the Reynolds number $Re = 30$, this eigenvalue is stable as shown by the square symbols. The real part of the eigenmode associated to this leading eigenvalue is depicted in fig. 7.5.2b with isocontours of the streamwise velocity. The spatial structure of this eigenmode breaks the symmetry of the steady solution and is responsible for the onset of the well-known Von Kármán vortex-street that becomes visible behind bluff bodies once the exponential growth of the linear instability saturates due to nonlinearities.




FIGURE 7.5.1: Two-dimensional and three-dimensional flow configurations. Sketch of the computational domains used for (a) the two-dimensional plate of height $h = 1$ and thickness $t = 1/6$ and (b) the three-dimensional plate of span $L = 2.5$ immersed in an upstream uniform streamwise flow U_∞ .

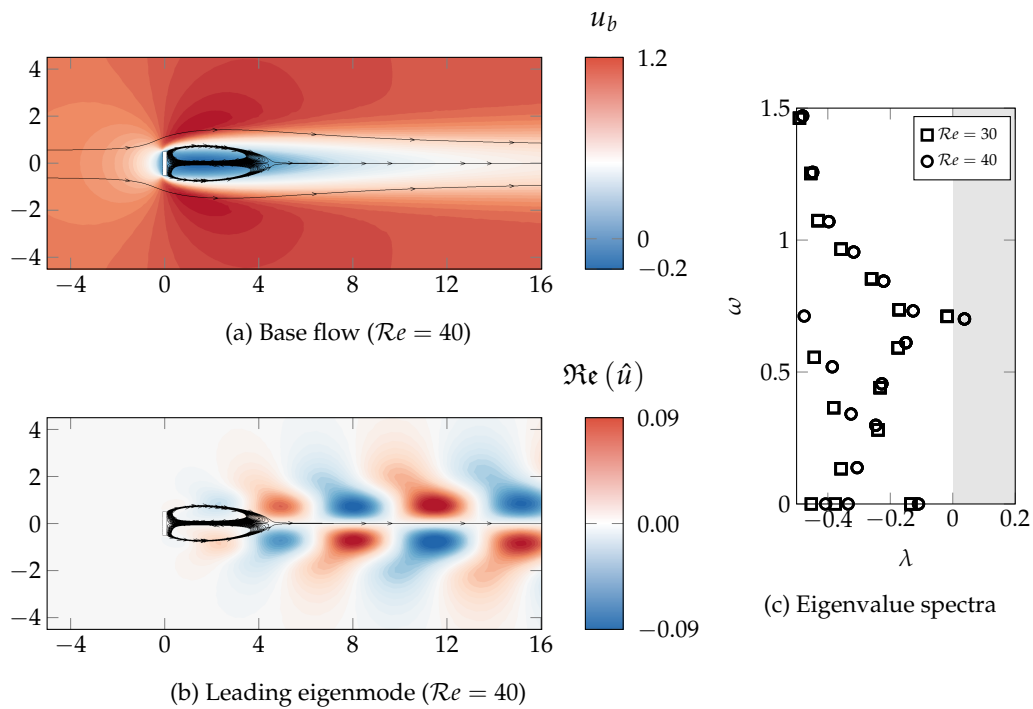


FIGURE 7.5.2: Results of the linear stability analysis for the two-dimensional configuration. Streamwise velocity u of (a) the steady-state solution and (b) the real part of the unstable eigenmode. (c) Eigenvalues are depicted with circles in the complex plane (growth rate λ and frequency ω). The unstable region is shown in gray. Only eigenvalues with positive frequencies are shown, the spectrum being symmetric.

The three-dimensional flow configuration is a plate of height and thickness identical to the two-dimensional plate, but of finite length L in the spanwise direction z , as sketched in fig. 7.5.1b. The computational domain is discretized using Gmsh [Geuzaine et al. 2009] by a Delaunay mesh composed of 17 million tetrahedra, which are then

partitioned between processes with ParMETIS [Karypis et al. 1998]. Using Taylor–Hood finite element pair, cf. ??, this leads to a total of 75 million unknowns. The boundary conditions are similar to those detailed above for the two-dimensional configuration.

The linear stability analysis of this flow configuration has been performed by [Marquet et al. 2015a] who determined the neutral curves of various unstable eigenmodes in the range of Reynolds number $40 \leq Re \leq 200$ and length $1 \leq L \leq 6$. Here, we specifically investigate the plate of length $L = 2.5$ for the Reynolds number $Re = 100$. The steady solution, depicted in fig. 7.5.3a, exhibits a large three-dimensional recirculation region in the wake of the plate. The stability analysis performed in [Marquet et al. 2015a] revealed that two pairs of complex eigenvalues get unstable above $Re \simeq 101$ for this parameter choice, with respective angular frequencies of $\omega \simeq 0.3$ and $\omega \simeq 0.57$. Hereinafter, we focus on the high-frequency eigenmode, depicted in fig. 7.5.3b. As shown in the figure, the three-dimensional eigenmode breaks the top/bottom symmetry of the steady-state solution, as for the two-dimensional plate.

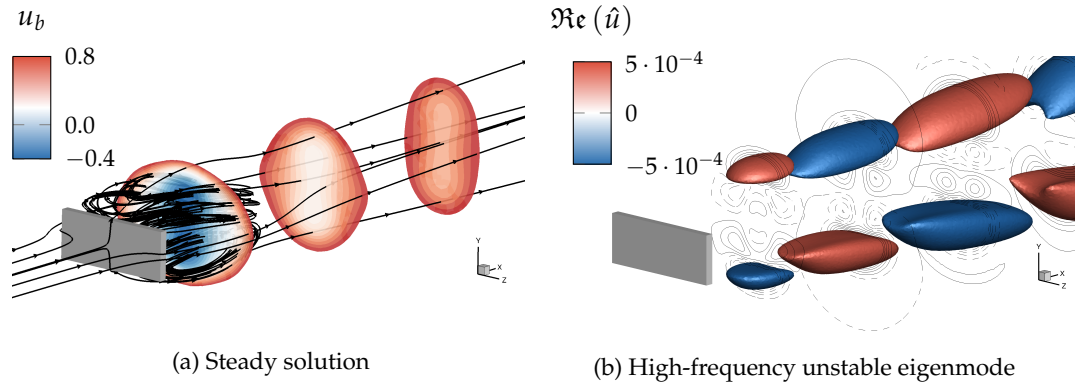


FIGURE 7.5.3: Linear stability analysis for the three-dimensional flow around of plate of spanwise length $L = 2.5$ and $Re = 100$. Streamwise velocity contours of (a) the steady solution and (b) the high-frequency unstable eigenmode ($\omega = 0.57$) are presented.

7.5.2 Influence of numerical and physical parameters

We investigate in this section the influence of various numerical and physical parameters on the performance of the mAL preconditioner. Tests are performed on the two-dimensional flow configuration previously introduced. The effect of the augmentation parameter on the preconditioner efficiency and solution accuracy is reported in section 7.5.2.1. The performance of the preconditioner is tested in section 7.5.2.2 for many values of the complex shift parameter used in the shift-and-invert strategy to compute the leading eigenvalues. The behavior of the preconditioner in regards to the mesh refinement and the Reynolds number is finally tested in section 7.5.2.3.

In all the numerical tests performed in this section, the full GMRES without restart is used in order to fairly assess the performance of the preconditioner. The diagonal blocks defined by the mAL preconditioner eq. (7.15) are here inverted using the sparse direct solver MUMPS.

7.5.2.1 Effect of the augmentation parameter

The effect of the augmentation parameter γ on the performance of the mAL preconditioner is first assessed by considering the number of GMRES iterations. Using the Newton method described in section 7.2.2, steady solutions are computed for the Reynolds number $Re = 40$ and several values of the augmentation parameters reported in the first columns of table 7.5.1a and table 7.5.1b, that correspond to results obtained with a coarse mesh (14,674 triangles) and a finer mesh ($1.29 \cdot 10^5$ triangles), respectively. The GMRES relative tolerance being fixed to 10^{-6} , the average numbers of inner (GMRES) iterations per outer (Newton) iteration are reported in the second columns of those tables. For both meshes, there exists an optimal value of the augmentation parameter, $\gamma \simeq 1$, for which a minimum number of iterations is reached. Similar observations are reported in other studies [Benzi et al. 2011b; Benzi et al. 2011c; Heister et al. 2013; Benzi et al. 2013; X. He et al. 2016] for different flow configurations such as the lid-driven cavity flow, the backward facing step or the flow over a flat plate. Note also that the number of iterations is quite similar for the coarse and fine meshes, regardless of the augmentation parameter value. For the optimal γ , the average number of inner iterations is around 50.

As briefly discussed in ??, the introduction of the grad-div stabilization term in the weak formulation 7.4a does not modify the conservation of momentum at the continuous level, since the continuous solution is divergence-free. However, with the spatial discretization chosen in the present study (Taylor-Hood finite element), the divergence of the velocity is only weakly satisfied and the grad-div stabilization term modifies the discrete momentum equation. The augmentation parameter has therefore an influence on the accuracy of the discrete solution. To assess this effect, a reference steady solution, denoted (\mathbf{u}_b^r, p_b^r) , is computed without stabilization parameter ($\gamma = 0$) on a very-fine mesh made of $5.13 \cdot 10^5$ triangles. The corresponding leading eigenvalue denoted σ^r is also computed. The two last columns of table 7.5.1a and table 7.5.1b report the relative errors of the steady velocity and the leading eigenvalue computed with the coarse and fine meshes, respectively, for several values of γ . Examining first the results obtained with the coarse mesh (see table 7.5.1a), a minimal error is obtained for $\gamma \simeq 1$, not only for the steady solution but also for the leading eigenvalue. When the mesh is refined (see table 7.5.1b), a minimal error is still obtained for $\gamma \simeq 1$, although less pronounced. Compared to results obtained with the coarse mesh, the relative error is decreased whatever the value of the augmentation parameter. As expected, the augmentation parameter less affects the accuracy of the discrete solution when the mesh is refined, since the discrete solution tends towards the continuous solution. It is worth noticing that the use of the stabilization term can significantly improve the accuracy of the solution. For instance, the accuracy of the eigenvalue obtained for the coarse mesh with $\gamma = 1$ is identical to the one obtained for the fine mesh without stabilization $\gamma = 0$. In other words, the same accuracy is obtained but with ten times fewer mesh elements.

The present results clearly indicate that $\gamma \simeq 1$ is an optimal value from both the solution accuracy point of view and the preconditioning efficiency point of view when considering not only steady solutions, as reported before [Heister et al. 2013], but also leading eigenvalues. As a consequence, in the following, we consider that γ can be chosen on preconditioning efficiency criteria only without compromising accuracy.

γ	# of GMRES iterations	$\frac{\ \mathbf{u}_b^h - \mathbf{u}_b^r\ _2}{\ \mathbf{u}_b^r\ _2}$	$\frac{ \sigma^h - \sigma^r }{ \sigma^r }$	# of GMRES iterations	$\frac{\ \mathbf{u}_b^h - \mathbf{u}_b^r\ _2}{\ \mathbf{u}_b^r\ _2}$	$\frac{ \sigma^h - \sigma^r }{ \sigma^r }$
0	860	$2.8 \cdot 10^{-4}$	$9.8 \cdot 10^{-4}$	873	$3.2 \cdot 10^{-5}$	$1.1 \cdot 10^{-4}$
10^{-1}	191	$2.1 \cdot 10^{-4}$	$7.9 \cdot 10^{-4}$	194	$2.7 \cdot 10^{-5}$	$1.1 \cdot 10^{-4}$
10^0	52	$8.4 \cdot 10^{-5}$	$1.1 \cdot 10^{-4}$	53	$1.4 \cdot 10^{-5}$	$5.4 \cdot 10^{-5}$
10^1	337	$4.4 \cdot 10^{-4}$	$1.5 \cdot 10^{-3}$	363	$2.7 \cdot 10^{-5}$	$8.5 \cdot 10^{-5}$
10^2	>1,000	$1.3 \cdot 10^{-3}$	$4.6 \cdot 10^{-3}$	>1,000	$6 \cdot 10^{-5}$	$2.1 \cdot 10^{-4}$

(A) Coarse mesh (14,674 triangles)

(B) Fine mesh ($1.29 \cdot 10^5$ triangles)

TABLE 7.5.1: Effect of the grad-div augmentation parameter γ on the mAL preconditioning efficiency and the solution accuracy. For both tables, the first column indicates values of γ . The second column represents the average number of mAL preconditioned GMRES iterations per Newton iteration. The last two columns give the relative errors of the steady solution and the leading eigenvalue compared with a reference solution $(\mathbf{u}_b^r, \sigma^r)$ computed on a very fine mesh without stabilization ($\gamma = 0$).

7.5.2.2 Effect of the shift parameter

The shift-and-invert strategy, adopted in the present study to compute the eigenvalues with largest real part, requires to specify the complex value $s = s_r + \mathbf{i}s_i$ that appears in the spectral transformation section 7.2.3. When investigating the transition of a steady solution from a stable to an unstable state, a common practice is to choose the shift parameter as a pure imaginary number, i.e. $s = \mathbf{i}s_i$, and to vary the imaginary part in order to compute complex eigenvalues with growth rates close to $\lambda = 0$. Depending on the flow configuration investigated, the steady solution may get unstable for eigenmodes characterized by very different frequencies. Ideally, the number of preconditioned GMRES iterations should be insensitive to the value of the complex shift, for the Krylov-Schur algorithm to converge rapidly whatever the eigenvalue of interest. To the best of our knowledge, only the case of a real-valued shift has been considered so far, either positive when solving the unsteady Oseen problem [Benzi et al. 2011b; Heister et al. 2013] or negative when solving the linearized Navier-Stokes equation [M. A. Olshanskii et al. 2008].

Here, we vary s in the whole complex plane and assess its effect on the efficiency of the mAL preconditioner by performing the following numerical experiment. The linear system eq. (7.12) is solved with right-hand side vectors whose coefficients are randomly generated in $[0, 1] + [0, 1]\mathbf{i}$, as done for instance in [Benzi 2008]. The Jacobian matrix \mathbf{J} of the linear system is computed with the steady solution at $Re = 40$ and the augmentation parameter $\gamma = 0.7$. In other words, only the inner solver is studied, no outer iteration (Newton or Krylov-Schur) is performed. The isocontours shown in fig. 7.5.4 in the complex plane (s_r, s_i) correspond to the number of inner (GMRES) iterations required to decrease the relative residual to 10^{-6} . The red circles are eigenvalues of the Jacobian matrix. First, the number of iterations increases when the shift gets closer to any eigenvalue. In that case, the matrix $\mathbf{J} + s\mathbf{M}$ involved in the spectral transformation (7.2.3) becomes singular, leading to a very ill-conditioned linear system and thus high iteration counts. Second, the number of iterations is reduced when increasing s_r . Solving the linear system for $s_r < 0$ is generally more expensive than for $s_r > 0$. According to [M. A. Olshanskii et al. 2008], this is due to the indefiniteness of the velocity block $\mathbf{A}_\gamma + s\mathbf{M}_u$ in eq. (7.12) for sufficiently large negative values of s_r . On the contrary, the velocity block is definite when $s_r > 0$. The




FIGURE 7.5.4: Influence of the complex shift. Isocontours represent the number of GMRES iterations needed to solve the linear system (7.13) with the mAL preconditioner, depicted in the complex plane (s_r, s_i) , and computed for $Re = 40$ and $\gamma = 0.7$. The red circles are the eigenvalues of the Jacobian matrix (7.6).

contribution of the shift can be interpreted as a (positive definite) time-step term, which reinforces the diagonal dominance of the problem and thus improve its spectral properties. For more details, the reader can refer to [Benzi et al. 2011b, section 2.6], where the mAL preconditioner is used to solve the unsteady Oseen problem. Finally, no particular trend is observed in the number of inner iterations when fixing the real part s_r and varying the imaginary part s_i , except when getting closer to an eigenvalue. For $s_r = 0$, the number of iterations is roughly constant for $s_i < 0.5$ and increases around $s_i = 0.6$ due to the proximity of the unstable eigenvalue marked by the red circle. By further increasing s_i , the number of iterations then decreases.

7.5.2.3 Effect of the mesh refinement and Reynolds number

The modified augmented Lagrangian preconditioner allows to compute steady solution in a number of GMRES iterations independent of the mesh refinement, as previously observed in table 7.5.1, and mildly dependent of the Reynolds number [Benzi et al. 2011b; Benzi et al. 2011a]. The influence of the mesh refinement and Reynolds number on the number of iterations needed to solve the complex-shifted linear system eq. (7.12) has not been investigated so far. The numerical experiment consists in solving the linear system to the relative tolerance of 10^{-6} for right-hand side vectors with randomly generated coefficients as explained before. First, the Reynolds number is fixed ($Re = 40$) as well as the augmentation parameter ($\gamma = 0.7$) while the mesh refinement changes. The number of inner GMRES iterations is reported in fig. 7.5.5a as a function of the number of triangles. The curves correspond to different values of the (purely imaginary) shift. Clearly, the iteration count is independent of the mesh refinement, regardless of the shift. Second, a fixed mesh refinement is chosen (14,828 triangles) and the linear system is solved for different values of the Reynolds number in the range $[10; 500]$ for different shift values. As reported in [Benzi et al. 2011c] when computing steady solutions, the optimal value of the augmentation parameter that minimizes the number of iterations depends on the

Reynolds number. The optimal value of γ has been first determined for each value of Re . For $Re = 10$, the optimal value is $\gamma \simeq 1.2$ and it decreases to $\gamma \simeq 0.4$ for $Re = 500$. These optimal values of γ have been determined for $s = 0$ but are used in the following regardless of the values of s . The number of iterations is depicted in fig. 7.5.5b as a function of the Reynolds number. The mAL preconditioner shows a mild degradation of its performance as Re increases, independently of the values of the shift. The increase of the number of inner iterations is proportional to $Re^{0.5}$ in this numerical experiment.

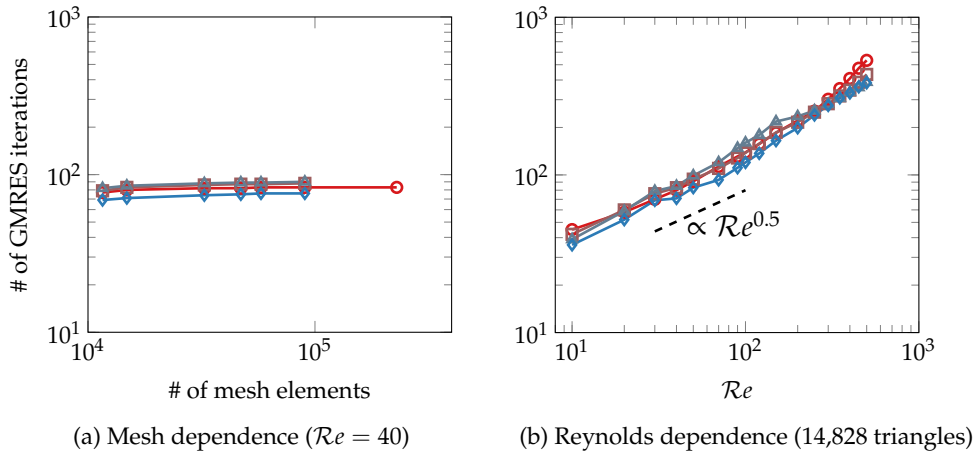


FIGURE 7.5.5: Effect of mesh refinement (a) and Reynolds number (b) on mAL preconditioning efficiency. The GMRES iteration count for solving eq. (7.12) to a relative tolerance of 10^{-6} is presented for different shifts: $s = 0$ (??), $s = 0.3i$ (??), $s = 0.6i$ (??) and $s = 1i$ (??)

As a conclusion, the modified augmented Lagrangian preconditioner exhibits interesting properties for performing efficiently a linear stability analysis using a shift-and-invert strategy: robustness with respect to a complex-valued shift, mesh independence, and a mild deterioration as Re increases.

7.5.3 Comparison with other block preconditioners

The mAL preconditioner is one of many other preconditioners developed to solve the steady incompressible Navier–Stokes equations. Among them, we select the Pressure Convection–Diffusion (PCD) [Kay et al. 2002] and the SIMPLE [Patankar et al. 1983] preconditioners, widely used and easily implemented, and compare their performance with those of the mAL preconditioner. In addition, we also test the unsteady Stokes preconditioner [Tuckerman 1989] which has gained in popularity in the hydrodynamic stability community [Tuckerman et al. 2019], as it can be easily implemented using existing time-steppers so as to compute base flows and leading eigenvalues [Tuckerman et al. 2000b]. In our implementation however, the Stokes preconditioner is itself applied using a nested Krylov subspace method instead of an existing time-stepper. More details on those preconditioners are given in section 7.B but it is worth recalling here that they are designed for classical Galerkin discretization of the Navier–Stokes equations. Therefore, their application to the iterative solution of eq. (7.12) is meant for $\gamma = 0$, i.e. without grad–div stabilization terms.

The numerical test case consists in solving iteratively eq. (7.12) with a random right-hand side vector, as detailed before, for a large relative tolerance equal to 10^{-3}

to limit the number of iterations. The shift is fixed to $s = 0$ since it was shown in section 7.5.2.3 that the mAL preconditioner depends very weakly on the shift parameter when the mesh is refined or the Reynolds number is increased. All innermost block solutions are performed using exact *LU* factorizations. The PCD and SIMPLE preconditioners are parameter-free, unlike the mAL and Stokes preconditioners. For the latter, the optimal values of the parameter (γ for mAL and a time-step like parameter for Stokes) are determined for each values of the Reynolds number.

The effect of the mesh refinement and Reynolds number on the number of inner iterations, studied in the previous paragraph for the mAL preconditioner, is assessed here for all the other preconditioners. Results are compared in fig. 7.5.6. This number of iterations is a good measure to compare the efficiency of the different preconditioners, in a first approximation, because for each inner iteration, the application of all preconditioners requires the solution of subproblems with similar complexities⁴. Therefore, the computational time of one inner iteration is roughly similar for all preconditioners.

All the preconditioners are independent of the mesh refinement, as shown in fig. 7.5.6a, except for the SIMPLE preconditioner for which the number of inner iterations slightly increases with the number of triangles. Interestingly, the number of iterations is significantly less for mAL and PCD (around 50) than for Stokes and SIMPLE (around 1,000). Note that for the two preconditioners depending on a parameter (mAL and Stokes), their optimal value was found to be independent of the mesh refinement.

The effect of the Reynolds number is reported in fig. 7.5.6b. For all tested preconditioners, the number of iterations increases with the Reynolds number, but with different slopes. The mAL preconditioner exhibits the best performance for all Reynolds numbers, except for low Reynolds number ($Re < 20$) where PCD is more efficient. However, the number of iterations obtained with the PCD preconditioner increases strongly for larger values of the Reynolds number ($Re > 80$). The mAL and SIMPLE preconditioners exhibit a similar trend: the number of iterations scales with the Reynolds number as $Re^{0.5}$. However, it is significantly larger with SIMPLE than with mAL, regardless of the Reynolds number. At low Re , the Stokes preconditioner behaves similarly to the SIMPLE preconditioner, but for larger Re , it degrades significantly and exhibits the same trend as the PCD preconditioner. Finally, when considering the number of iterations, the mAL preconditioner is undoubtedly the best preconditioner. We note that, contrary to the mesh dependence study, the optimal parameters of mAL and Stokes showed some variations with respect to Re .

⁴two scalar velocity solves and one pressure Schur complement solve for mAL and Stokes; one vectorial velocity solve and one pressure Schur complement solve for PCD and SIMPLE

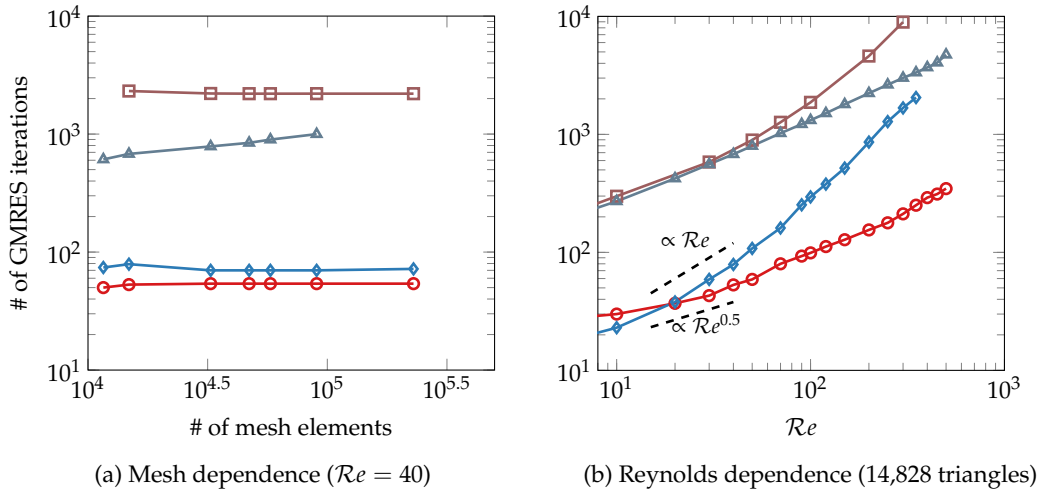


FIGURE 7.5.6: Influence of (a) the number of mesh elements and (b) the Reynolds number Re on the number of inner iterations required to solve eq. (7.12) with the mAL (??), PCD (??), SIMPLE (??), and Stokes (??) preconditioners. The relative tolerance is set to 10^{-3} and the shift to zero.

Let us now compare the computational time for applying the four preconditioners. To that aim, the direct LU factorizations used until now for the innermost velocity blocks solvers are replaced by GMRES right-preconditioned with an $ILU(2)$ method (as implemented in PETSc). The choice of an innermost iterative solution allows for a more comprehensive interpretation of the computational time, since it accounts for the various complexities in solving iteratively the velocity blocks involved in the different preconditioners. Moreover, such an innermost iterative solution is necessary when considering large-scale three-dimensional problems, as shown in the next section. The relative tolerance of the inner (resp. innermost) GMRES is fixed to 10^{-3} (resp. 10^{-2}). The computational times obtained with the four preconditioners are reported in fig. 7.5.7 for $Re = 40$ (left) and $Re = 100$ (right). The total time is split into the time spent in computing matrix–vector products, in applying the global preconditioner, and in constructing the global Krylov subspace. The inner iteration counts are given between parenthesis. Note that it may be slightly higher than what is presented in fig. 7.5.6 since the velocity blocks are now solved only approximately. All computations are run on a standard laptop computer.

For the low Reynolds number $Re = 40$, the mAL and PCD preconditioners are about ten times faster than the SIMPLE and Stokes preconditioners⁵. For this Reynolds number, the PCD preconditioner is comparable with the mAL preconditioner, as it is only 40% slower. However, when the Reynolds number is increased to $Re = 100$, the performance of PCD degrades significantly with respect to mAL, as it is now about five times slower. The computational times are not given for the SIMPLE and Stokes preconditioners because they largely exceed 2,100 seconds. The deterioration in the computational time of the PCD preconditioner when the Reynolds number is increased, is in agreement with the growth in the number of iterations observed before. For even higher Reynolds numbers $Re > 100$, the mAL preconditioner is expected to be increasingly more interesting than its competitors.

⁵note that the time for building the inner Krylov subspace is very small for the Stokes preconditioner at $Re = 40$, despite a large number of iterations. This is due to the fact that, in the implementation detailed in section 7.B, the Stokes preconditioner is itself applied with a nested iterative method. Therefore, no Krylov subspace of dimension 947 is actually built

As a conclusion, this benchmark shows that, compared to other widely used preconditioners, mAL provides a more efficient approach for solving eq. (7.12) on a configuration typical of two-dimensional external flows around bluff bodies. In particular, among the alternatives tested here, it is the only preconditioner combining a mesh-independent iteration count and a mild degradation with $Re^{0.5}$, making it the most efficient preconditioner for $Re \geq 20$.

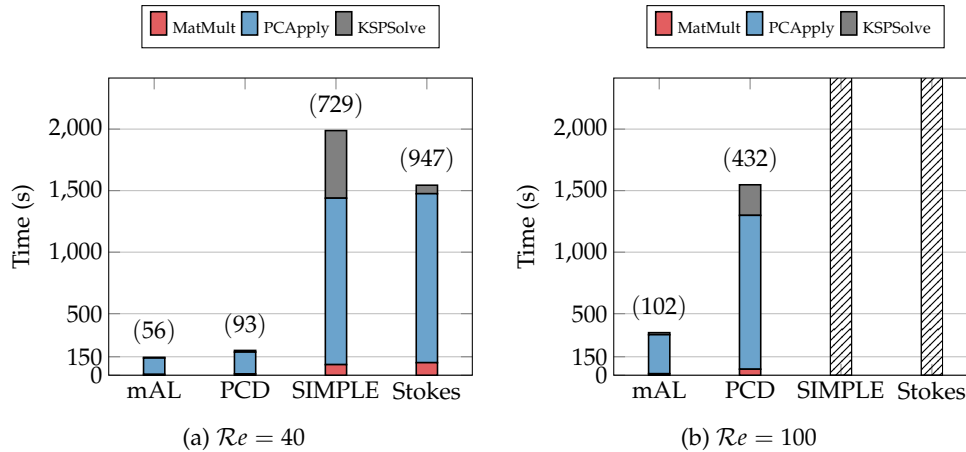


FIGURE 7.5.7: Comparison of the performance of various preconditioners on the two-dimensional test case (14,828 triangles). The total time is split between matrix–vector products, applying the preconditioner, and building the Krylov subspace. The number of global GMRES iterations is given between parenthesis. System eq. (7.12) is solved to a relative tolerance of 10^{-3} . For $Re = 100$, the hatched bars correspond to preconditioners for which the total time largely exceeded 2,100 seconds and is not reported in details. The velocity blocks in the preconditioners are solved iteratively to a relative tolerance of 10^{-2} using an innermost GMRES, right-preconditioned with $ILU(2)$. The pressure blocks are solved exactly with MUMPS. These times will depend on the particular preconditioners used for solving the diagonal blocks. Therefore, those results should be considered qualitatively.

7.5.4 Performance of the parallel implementation

In this section, the performance of the parallel implementation detailed in section 7.4 is tested on the three-dimensional configuration presented in section 7.5.1. First, a coarse mesh is used, in order to be able to compare our approach with the direct parallel solver MUMPS. Then, the full size 3D configuration presented before is considered to test the parallel performance of our approach on a problem that a direct solver could not handle at a reasonable memory cost.

7.5.4.1 Comparison with a direct solver on a small-scale 3D configuration

Despite its large memory requirements, some authors have used the “matrix-based” approach, combined with direct solvers for the arising linear systems, to perform the stability analysis of three-dimensional flows [Marquet et al. 2015a; Iorio et al. 2014]. In this section, we aim at comparing the performance of this approach to ours. To that end, the three-dimensional test case is considered using a coarse mesh of 1.1 million tetrahedra (4.8 million unknowns), in order to keep the memory consumption of the direct solver reasonable. The computations are performed on Sator,

an ONERA cluster composed of 620 nodes with two fourteen-core Intel Broadwell clocked at 2.4 GHz. The direct solver we compare ourselves to is MUMPS.

7.5.4.1.a Nonlinear solver

In this section, the Newton nonlinear tolerance and the GMRES relative tolerance are both set to 10^{-6} . In fig. 7.5.8a we report the average wall-clock time per Newton iteration for the mAL and MUMPS approaches, as a function of the number of processes. On the top x -axis, we report the amount of available memory corresponding to each number of processes. First, as expected, the memory requirements of our approach are lower than with MUMPS: we observe that MUMPS cannot be run on less than 224 processes, which corresponds to an available memory of 1,024 GB, whereas the mAL approach can be run on 28 processes (128 GB). Note that the memory requirements of the mAL approach could be even lower by using iterative methods for the subdomain solvers of the innermost ASM-preconditioned GMRES iterations. Moreover, thanks to good scalability properties and the absence of a full LU factorization at each Newton iteration, the mAL approach is clearly faster than MUMPS (about ten times with 448 processes).

7.5.4.1.b Eigensolver

For the eigensolver, the Krylov–Schur tolerance is 10^{-6} whereas the inner relative tolerance is 10^{-3} . Note that we use a larger tolerance for the inner linear solve than for the outer Krylov–Schur solver. Indeed, contrary to what is often recommended in the literature (e.g., [Roman et al. 2018, § 3.4.1]), we observed that it was not necessary to use a smaller tolerance for the inner solution of eq. (7.12) in order to keep a satisfying accuracy on the computed eigenvalues. More details on that aspect may be found in section 7.C.

We show in fig. 7.5.8b the total wall-clock time for computing 5 eigenvalues closest to the shift $s = 0.6i$, using mAL and MUMPS as inner solvers, as a function of the number of processes. The available memory is again reported on the top x -axis. Similar conclusions as for the nonlinear solver can be made for memory consumption with a multiplication factor of two, due to the use complex instead of real algebra. From a wall-clock time point of view, we observe, as for the nonlinear solver, that the mAL approach possesses much better scalability than MUMPS, which leads to a faster computation. We note however that MUMPS is harder to beat with an iterative approach when used in the eigensolver than in the nonlinear solver. The reason is that, in the Krylov–Schur method, the very high cost of forming the full LU factorization is greatly amortized by the many inner solves realized with it, whereas in the Newton method, each inner solve requires to build the factorization again.

As a conclusion, the mAL approach presents the double advantage of being much less memory-intensive than the direct solver and also faster, even for the unfavorable case of eigenvalue computations.

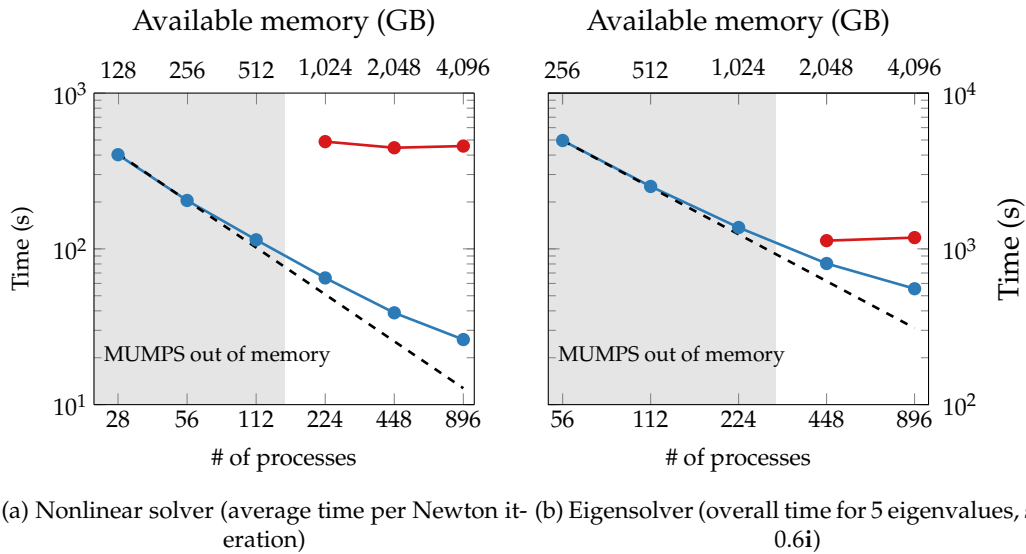


FIGURE 7.5.8: Comparison of mAL (??) and MUMPS (??) as inner solvers on (a) the nonlinear solver and (b) the eigensolver. The computation are performed on the three-dimensional test case using a coarse mesh (4.8 million unknowns) and $\gamma = 0.3$. The grey area indicates when the LU factorization is not feasible due to too large memory requirements. The dashed line represents ideal scalability.

7.5.4.2 Parallel performance on a large-scale 3D configuration

In this section, the 3D plate configuration is used, with a fine mesh, resulting in 75 million unknowns. The parallel performance of our implementation is investigated for the nonlinear base flow solver and eigenvalue solver.

Results were obtained on Curie, a system composed of 5,040 nodes with two eight-core Intel Sandy Bridge clocked at 2.7GHz. The interconnect is an InfiniBand QDR full fat tree and the MPI implementation exploited was bullxMPI version 1.2.9.2. All binaries and shared libraries were compiled with Intel compilers and Math Kernel Library support (for dense linear algebra computations) version 18.0.1.163. Recent releases of FreeFEM and PETSc/SLEPc were used (version 3.61 and 3.9.3 respectively). In all following plots and tables, the time spent in the finite element kernel is never accounted for because we are mostly interested in the performance of the preconditioner. Only the time spent in PETSc or SLEPc is reported.

7.5.4.2.a Nonlinear solver

In this paragraph, we investigate the parallel performance of the nonlinear steady-state solver. The inner Krylov solver is the flexible GMRES algorithm [Saad 1993], which is stopped when the relative unpreconditioned residual norm is lower than 10^{-1} . The Newton method is stopped when the l^2 norm of the residual is lower than 10^{-6} . As an initial guess for the computation at $Re = 100$, a solution at a lower Reynolds number $Re = 50$ is first computed using a higher nonlinear outer tolerance of 10^{-4} . Also note that in this preprocessing step, all the domain decomposition information obtained from ParMETIS partitioning is dumped and will be used in successive runs for the nonlinear and generalized eigenvalue solvers. In table 7.5.2, the numerical performance of the nonlinear solver are reported. One may notice that even if a high relative tolerance is used to stop the flexible GMRES, very few Newton iterations (second column) are needed for the solver to converge, independently

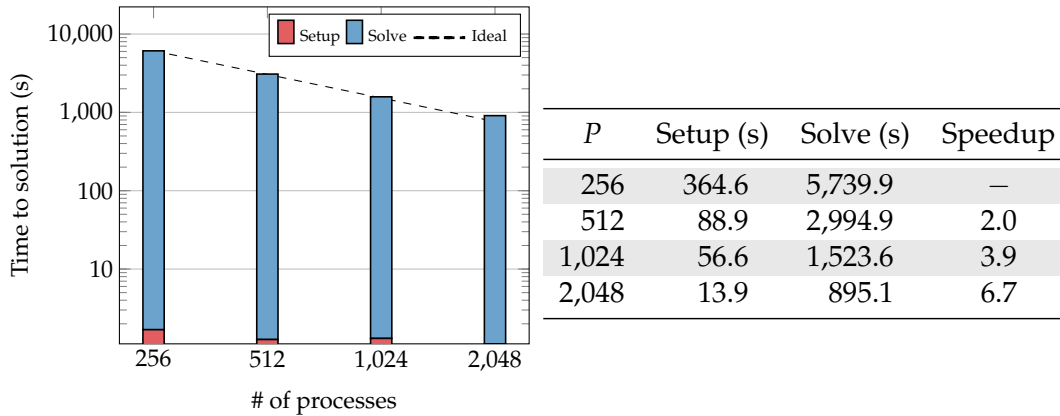


FIGURE 7.5.9: Scalability of the 3D nonlinear solver with respect to the number of processes.

of the number of subdomains (first column). The number of mAL-preconditioned inner iterations needed to reach convergence, averaged over all Newton iterations, is reported in the third column. It is seen not to depend on the number of processes. Eventually, in the last three columns, we show the average number of ASM-preconditioned innermost iterations needed for each velocity block of \mathbf{P}_{mAL} to reach the desired convergence tolerance of 10^{-1} (see section 7.4.3). A slight increase is observed with the number of processes. This is an expected feature of simple one-level domain decomposition methods, like the additive Schwarz method, that are known to not scale numerically [Dolean et al. 2015].

P	# of Newton iterations	# of iterations per Newton it.	# of iterations per field (x, y, z)		
256	6	83	30	12	19
512	5	81	31	13	20
1,024	5	84	35	15	21
2,048	5	84	44	17	27

TABLE 7.5.2: Numerical performance of the 3D nonlinear solver with respect to the number of processes ($Re = 100$). The second column represents the number of Newton outer iterations, the third is the number of mAL-preconditioned inner GMRES iterations per Newton step. The last three columns correspond to the average number of ASM-preconditioned innermost GMRES iterations for each velocity block per inner iteration.

In fig. 7.5.9, the scalability of our implementation is shown, using the run with 256 processes as the reference, and going up to 2,048 processes. The parallel efficiency of this approach remains above 83%. The fact that one additional Newton iteration is needed with 256 processes has to be highlighted, since it does improve the efficiency. Other than that, because exact LU factorizations are used as subdomain solvers in the additive Schwarz method used for each velocity field, the setup phase scales superlinearly (see the second column of the table in fig. 7.5.9). Moreover, because the number of iterations needed for the corresponding solvers only grows slightly, as shown in the three last columns from table 7.5.2, the solution phase also scales appropriately.

7.5.4.2.b Eigensolver

We now evaluate the parallel performance of the eigensolver. The tolerance on the Krylov–Schur algorithm is 10^{-6} whereas the relative tolerance for the inner linear solver is set⁶ to 10^{-4} . The main difference with the Newton method is that the multiple inner linear solves involve the very same shifted operator ($\mathbf{J} + s\mathbf{M}$) and preconditioner \mathbf{P}_{mAL} . To improve the performance of the eigensolver, let us show first the effect of using a recycled Krylov method for solving these systems, instead of the standard GMRES. This can be done by switching from the KSP objects of PETSc to the iterative methods of the HPDDM library [Jolivet et al. 2016] which handle subspace recycling. In particular, the flexible GCRO-DR method is used, with a recycled subspace between each linear solves of dimension five. The following lines allow to switch between PETSc and HPDDM Krylov methods from within FreeFEM:

```
int recycle = getARGV("-recycle", 0); // use GMRES by default
int restart = getARGV("-restart", 200); // default to 200
real innerTol = getARGV("-inner_tol", 1.0e-4); // default to 10-4
string paramsKrylov = (recycle == 0 ? "-st_ksp_type fgmres " +
  "-st_ksp_monitor -st_ksp_rtol " + innerTol +
  "-st_ksp_gmres_restart " + restart + "-st_ksp_max_it 1000"
  :
  "-st_ksp_type hpddm -hpddm_st_krylov_method gcrodr " +
  "-hpddm_st_recycle " + recycle + "-hpddm_st_max_it 1000" +
  "-hpddm_st_verbosity 4 -hpddm_st_gmres_restart " + restart +
  "-hpddm_st_tol " + innerTol + "-hpddm_st_variant flexible");
```

In fig. 7.5.10, the number of mAL-preconditioned inner linear iterations needed for each iterative method (FGMRES or FGCRO-DR) to solve the sequence of linear systems of the first iteration of the Krylov–Schur algorithm is reported. For this particular Krylov–Schur iteration, fifteen systems have to be solved. When using FGMRES, it corresponds to a total of 3,751 inner linear iterations. When using FGCRO-DR, it corresponds to only 2,209 inner linear iterations. Even though the solutions of all fifteen systems are not rigorously equal when switching from FGMRES to FGCRO-DR, after the first iteration of the eigensolver, convergence is reached for the two eigenpairs closest to the shift: $-1.03 \cdot 10^{-2} + 0.57i$ and $-7.81 \cdot 10^{-2} + 0.57i$.

In all the following runs, FGCRO-DR is used in order to reduce the number of inner iterations. The number of Krylov–Schur iterations needed to retrieve the requested eigenpairs is reported in the second column of table 7.5.3. In the third column is the average number of solved linear systems per eigensolver iteration. The average number of mAL-preconditioned FGCRO-DR iterations (inner iterations) per linear solve (outer iteration) is presented in the fourth column whereas the last three columns contain the average number of innermost iterations for each velocity field, per application of \mathbf{P}_{mAL} . It was not possible to have the code run on 256 processes due to memory requirements significantly higher than for the nonlinear solver. Indeed, we switch from a real-valued to a complex-valued problem and the additional operators \mathbf{M} and \mathbf{L}_p are assembled explicitly. In table 7.5.4, the scalability of our implementation is shown, using the run with 512 processes as the reference, and going up to 2,048 processes. The parallel efficiency of this approach is approximately the same as for the nonlinear solver, though on a narrower range of process counts, remaining above 82%.

⁶for the same reasons explained in 7.5.4.1.b, we use a larger tolerance for the inner linear solves than for the outer Krylov–Schur algorithm (see also section 7.C)




FIGURE 7.5.10: Effect of a recycled Krylov method on the performance of the eigensolver. The number of mAL-preconditioned inner iterations for the flexible GMRES and GCRO-DR algorithms is compared for each linear solve of a Krylov–Schur iteration.

P	# of eigensolver iterations	# of linear solves	# of iterations per linear solve	# of iterations per field (x, y, z)		
512	7	7	120	7	10	11
1,024	7	8	127	7	10	13
2,048	7	8	119	10	13	17

TABLE 7.5.3: Numerical performance of the 3D eigensolver with respect to the number of processes. The second and third columns represent respectively the number of Krylov–Schur iterations and the number of linear solves (outer iterations) per Krylov–Schur iteration. The fourth column is the number of mAL-preconditioned FGCRO-DR inner iterations per linear solve. The last three columns correspond to the average number of innermost ASM-preconditioned GMRES iterations for each velocity sub-block per inner iteration.

P	Setup (s)	Solve (s)	Speedup
512	55.3	39,160.7	–
1,024	25.7	24,508.3	1.6
2,048	27.1	11,849.9	3.3

TABLE 7.5.4: Scalability of the 3D eigensolver with respect to the number of processes.

7.6 Conclusion

The stationary base flow as well as the eigenvalue computations involved in hydrodynamic linear stability analysis require multiple solutions of linear systems based on the (shifted) Jacobian operator of the incompressible steady Navier–Stokes equations. To solve such systems on large-scale configurations involving hundreds of millions of unknowns, we proposed to use a Krylov subspace linear solver like the flexible GMRES algorithm, preconditioned by the modified augmented Lagrangian (mAL) preconditioner [Benzi et al. 2011b]. On a two-dimensional bluff-body flow, we studied numerically the performance of the mAL preconditioner for linear stability analysis purposes. We showed in particular that this approach handles efficiently complex-valued shifts and thus is well-suited for the computations of eigenvalues with possibly large frequencies, using the shift-and-invert spectral transformation. Then, the mAL preconditioner was tested against some other widely used steady-state (PCD, SIMPLE) and time-stepping-based (Stokes) preconditioners, all of them used in a sequential version. The mAL preconditioner was shown to require lower numbers of GMRES iterations and to be faster than all its competitors.

In order to perform large-scale three-dimensional stability analysis computations, a parallel implementation of the mAL preconditioner was presented and is made available online: <https://github.com/prj-/moulin2019a1>. The FreeFEM finite element language was used as a discretization kernel whereas PETSc and SLEPc were used as distributed linear algebra backends. First, a comparison with the parallel direct solver MUMPS was presented on a three-dimensional bluff-body flow configuration, using a coarse enough mesh to make the LU factorization possible. The mAL approach required about one tenth as much memory and had better strong scaling properties. Despite the attractiveness of a direct linear solver — when it can be afforded — for the eigenvalue computations (the factorization is done once and re-used multiple times), the mAL approach turned out to be a faster alternative than the direct approach, thanks to its much better parallel performance. Finally, the implementation was used on a fine mesh resulting in 75 million unknowns and showed satisfying strong scaling properties up to 2,048 processes, for both the base flow and eigenvalue computations. The role of subspace recycling between the multiple consecutive linear solves, inside the Krylov–Schur eigensolver, was tested and allowed significant performance gains.

Acknowledgments

This project has received funding from the European Research Council (ERC) under the European Unions Horizon 2020 research and innovation program (grant agreement number 638307). Moreover, this work was granted access to the HPC resources of TGCC@CEA under the allocation A0030607519 made by GENCI.

APPENDIX

Appendix 7.A Reproducibility

In addition to the few extracts of the code used in the paper, the interested reader can find the complete FreeFEM code in the following repository: <https://github.com/prj-/moulin2019a1>.

Appendix 7.B Definition of other block preconditioners

Here are defined the classical block preconditioners for incompressible Navier–Stokes that are compared to the modified augmented Lagrangian approach in section 7.5.3. Contrary to mAL, those preconditioners do not require an augmentation. Thus, they are used without grad–div stabilization ($\gamma = 0$). Versions that incorporate a complex shift s are proposed here.

Pressure convection–diffusion preconditioner The pressure convection–diffusion (PCD) preconditioner was proposed by [Kay et al. 2002]:

$$\mathbf{P}_{\text{PCD}} = \begin{pmatrix} \mathbf{A} + s\mathbf{M}_{\mathbf{u}} & \mathbf{B}^T \\ 0 & \mathbf{S}_p \end{pmatrix}, \quad (7.17)$$

with $\mathbf{S}_p^{-1} = -\mathbf{M}_p^{-1}(\mathbf{F}_p + s\mathbf{M}_p)\mathbf{L}_p^{-1}$, where \mathbf{F}_p is a convection–diffusion operator built on the pressure space. Compared to the classical PCD preconditioner for steady-state Navier–Stokes equations, the shift contribution $s\mathbf{M}_p$ is added to the pressure Schur complement approximation.

SIMPLE preconditioner The SIMPLE preconditioner was proposed as a solver by [Patankar et al. 1983]. We use its preconditioner version [Elman et al. 2008]:

$$\mathbf{P}_{\text{SIMPLE}} = \begin{pmatrix} \mathbf{A} + s\mathbf{M}_{\mathbf{u}} & 0 \\ \mathbf{B} & \mathbf{S}_p \end{pmatrix} \begin{pmatrix} \mathbf{I} & \text{diag}(\mathbf{A} + s\mathbf{M}_{\mathbf{u}})^{-1}\mathbf{B}^T \\ 0 & \mathbf{I} \end{pmatrix}, \quad (7.18)$$

and $\mathbf{S}_p^{-1} = -[\mathbf{B}\text{diag}(\mathbf{A} + s\mathbf{M}_{\mathbf{u}})^{-1}\mathbf{B}^T]^{-1}$.

Stokes preconditioner The Stokes preconditioning approach was popularized by [Tuckerman et al. 2000a] in the hydrodynamic stability community. Tuckerman’s idea is two-fold:

1. preconditioning the linearized Navier–Stokes problem eq. (7.12) by the Stokes problem, i.e.,

$$\mathbf{P}_{\text{Stokes}} = \begin{pmatrix} \mathbf{A}_{\text{Stokes}} + s\mathbf{M}_{\mathbf{u}} & \mathbf{B}^T \\ \mathbf{B} & 0 \end{pmatrix},$$

with $\mathbf{A}_{\text{Stokes}} = \mathbf{D} + \Delta t^{-1} \mathbf{M}_{\mathbf{u}}$ and \mathbf{D} contains only the diffusion terms. Note that the time-step contribution Δt has no physical meaning here: it only represents some numerical parameter of the preconditioner. A case-dependent optimal value may exist, as reported in [Beaume 2017]. We determined this optimal value numerically.

2. applying the preconditioner by adapting a pre-existing time-stepping code, which significantly reduces the development costs. In this work however, we prefer to apply $\mathbf{P}_{\text{Stokes}}^{-1}$ by using a few inner iterations of GMRES, preconditioned by:

$$\mathbf{P}_{\text{Stokes, inner}} = \begin{pmatrix} \mathbf{A}_{\text{Stokes}} + s\mathbf{M}_{\mathbf{u}} & \mathbf{B}^T \\ 0 & \mathbf{S}_p \end{pmatrix},$$

with $\mathbf{S}_p^{-1} = -\mathcal{R}e^{-1} \mathbf{M}_p^{-1} - (\Delta t + s) \mathbf{L}_p^{-1}$ [Cahouet A N et al. 1988]. A large relative tolerance of 10^{-2} is set for the inner iterations, as we observed that further convergence of the inner iterations did not improve the convergence of the outer GMRES iterations. Obviously, $\mathbf{A}_{\text{Stokes}}$ being block diagonal, applying $\mathbf{P}_{\text{Stokes, inner}}^{-1}$ naturally requires two scalar velocity solves (in 2D) and one pressure solve.

Note that, in section 7.5.3, the GMRES iteration count reported for the Stokes preconditioner corresponds to the total number of applications of $\mathbf{P}_{\text{Stokes, inner}}$ necessary to converge eq. (7.12) to the desired tolerance.

Appendix 7.C Linear solver tolerance and eigenvalue convergence criterion

In cases where an iterative linear solver is used, the action of $(\mathbf{J} + s\mathbf{M})^{-1}$, required when applying the spectrally transformed operator \mathbf{T} in the Krylov–Schur algorithm (see section 7.2.3), is approximated using some user-defined tolerance. As a consequence, matrix \mathbf{T} in eq. (7.11) is replaced by some approximation $\tilde{\mathbf{T}}$. It is usually recommended to set the tolerance for the linear solver lower than the one prescribed to the eigensolver, so that the imprecision of the linear solver does not pollute the eigensolver accuracy (see e.g. [Roman et al. 2018, § 3.4.1]). Here, we re-evaluate this statement numerically on the two-dimensional test case presented in section 7.5.1.

The following numerical experiment is performed. We solve the eigenproblem (7.11) using a mAL-preconditioned GMRES algorithm to apply $(\mathbf{J} + s\mathbf{M})^{-1}$. The relative tolerance of the GMRES algorithm ε_{lin} is varied between 10^{-8} and 10^{-1} while the tolerance of the Krylov–Schur algorithm ε_{eig} is kept constant to 10^{-6} . Only one eigenvalue, closest to the shift $s = 0.7i$, is demanded. In the second column of table 7.C.1, the number of GMRES iterations required to apply $(\mathbf{J} + s\mathbf{M})^{-1}$ is shown. The value is averaged over all applications of $(\mathbf{J} + s\mathbf{M})^{-1}$ to compute the demanded eigenvalue. In the third column, the total number of applications of $\tilde{\mathbf{T}}$ is shown. In the last three columns, we monitor the eigenvalue and the discrete l^2 norm of the eigenproblem residual. It is observed that one can in practice increase ε_{lin} well above $\varepsilon_{\text{eig}} = 10^{-6}$, without compromising significantly the accuracy of the computed eigenvalue. At least up to $\varepsilon_{\text{lin}} = 10^{-3}$, the computed eigenvalue is converged to satisfying accuracy, for a cost divided by two with respect to the “safe choice” $\varepsilon_{\text{lin}} = 10^{-6}$. Increasing ε_{lin} may thus allow some performance improvement.

Finally, the interpretation of the last column of table 7.C.1 deserves some further explanation. Indeed, one can observe that, for $\varepsilon_{\text{lin}} = 10^{-1}$ and $\varepsilon_{\text{lin}} = 10^{-2}$, despite

the fact that we kept $\varepsilon_{\text{eig}} = 10^{-6}$, the Krylov–Schur algorithm considered it had converged to an appropriate eigenvalue, while the residual was still above ε_{eig} . The reason to that observation is that most Arnoldi-based eigensolver packages, such as SLEPc, use convergence criteria based on the residual of the transformed problem eq. (7.11), not the original one eq. (7.6). As a consequence, the effect of using a large linear tolerance ε_{lin} is to apply an approximate operator $\tilde{\mathbf{T}}$ far from the exact one. But the eigenvalues of $\tilde{\mathbf{T}}$ can be computed to any prescribed tolerance by the Krylov–Schur algorithm. Note that, in SLEPc, a workaround is to use the option `-eps_true_residual` which forces the computation of the residual on the original problem eq. (7.6) and thus is free from any approximation linked to the underlying linear solver. In this case, the effect of using a large linear tolerance ε_{lin} would be to make the convergence of the Krylov–Schur algorithm increasingly slow (or even impossible). This option being more costly, it should however be avoided for large-scale computations.

ε_{lin}	# of GMRES iterations per linear solve	# of linear solves	λ	ω	$\ \sigma\mathbf{M}\hat{\mathbf{q}} + \mathbf{J}(\mathbf{q}_b)\hat{\mathbf{q}}\ _2$
10^{-1}	15	12	$4.48 \cdot 10^{-2}$	$6.90 \cdot 10^{-1}$	$1.5 \cdot 10^{-4}$
10^{-2}	28	10	$3.77 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$2.6 \cdot 10^{-6}$
10^{-3}	36	8	$3.75 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$9.5 \cdot 10^{-7}$
10^{-4}	45	8	$3.75 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$9.2 \cdot 10^{-7}$
10^{-5}	53	8	$3.75 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$9.2 \cdot 10^{-7}$
10^{-6}	62	8	$3.75 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$9.2 \cdot 10^{-7}$
10^{-7}	71	8	$3.75 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$9.2 \cdot 10^{-7}$
10^{-8}	80	8	$3.75 \cdot 10^{-2}$	$7.01 \cdot 10^{-1}$	$9.2 \cdot 10^{-7}$

TABLE 7.C.1: Effect of linear solver tolerance on eigenvalue computation ($\mathcal{R}e = 40$, $\gamma = 0.7$, $\varepsilon_{\text{eig}} = 10^{-6}$, $s = 0.7\mathbf{i}$)

8

CONCLUSION & PERSPECTIVES

8.1 Conclusion

This thesis represents an effort towards a better understanding of the bifurcation associated to the flutter instability in low-Reynolds viscous incompressible flows, using specifically designed linear and nonlinear methods such as global linear stability analysis, weakly nonlinear analysis, Harmonic Balance-type computations and Floquet stability analysis.

Theoretical and numerical methods for nonlinear flutter analysis

The first part of this thesis was dedicated to the introduction of different linear and nonlinear methods well-suited to the investigation of fluid-structure instabilities in viscous incompressible flows, and in particular to the flutter instability.

In chapter 1, we started by introducing the typical section model: a heaving and pitching spring-mounted rigid plate immersed in a viscous incompressible two-dimensional flow. Following the general spirit of this thesis, which is to focus on fluid modelling, we used a simple dynamically linear solid model composed of two degrees of freedom (heaving and pitching), each of them being governed by a damped oscillator equation. The flow, for his part, is modelled with the two-dimensional incompressible Navier–Stokes equations. Thanks to this approach, complex nonlinear flow features, like boundary layers, recirculations regions, dynamic stall vortices, etc, can be captured. To properly handle the movement of the fluid domain, two variants of the Arbitrary Lagrangian Eulerian framework were introduced. First, the so-called “absolute velocity - rotating axis” which is only valid for open flows and rigid-body motion but has the advantage of analytical simplicity. Second, the “reference configuration ALE” formulation, which has none of these limitations but at the cost of significant technical complications. It was shown that the “absolute velocity - rotating axis” formalism is in fact a special case of the “reference configuration ALE”, where a rigid-body extension is used and the momentum equations are projected onto the rotating axes. Finally, an algorithm for performing unsteady time-marching computations of the coupled system was presented, based on the combination of a pressure segregation approach to handle the incompressibility constraint and fixed-point Dirichlet-to-Neumann iterations for ensuring strong fluid-structure coupling.

In order to perform efficient investigations of the flutter bifurcation without resorting to fully nonlinear time-marching computations, a weakly nonlinear analysis was introduced in chapter 2. Using the method of multiple scales, a normal form

(amplitude equation) for the flutter bifurcation was derived in a semi-analytical fashion. Indeed, in this method, the nonlinear interactions between the different scales and the time dependence are handled analytically while a spatial discretization of the Navier–Stokes equations is still needed. With the objective to parametrically explore the effects of the Reynolds number, a Hessian-based mesh adaptation framework was proposed so as to adapt to the different flow features that appear when varying Reynolds in the range of interest of the subsequent chapters, $10 \leq Re \leq 10^4$. In addition, thanks to a spatial decomposition of the cubic coefficient of the normal form, it was realized that small-scale structures, localized around the plate leading edge, may participate to the determination of the nature of the flutter bifurcation. Mesh adaptation is then particularly well-suited for performing an adequate and automatic spatial discretization of these structures.

Anticipating the shortcomings of the weakly nonlinear analysis for computing highly-nonlinear flutter solutions, a harmonic balance-type approach known as the Time Spectral Method (TSM) was then introduced in chapter 3. With this approach, only periodic solutions are searched but without any hypothesis on the strength of the nonlinearity, *i.e.* on the amplitude of the oscillations. This yields several advantages for bifurcation study purposes that are: (i) the ability to compute possibly unstable periodic solutions and (ii) a direct targeting of the permanent regime of interest — the periodic orbits — without going through often long transient regimes. Before diving into the core of the chapter, which is the numerical solution of the TSM equations, we proposed a synthetic overview and consistent classification (table 3.2.1) of several variants of harmonic balance methods that have appeared through the years. We hope this effort may help the interested reader dive into this abundant and keen-on-acronyms literature. Solving the TSM problem mentioned above consists in solving a nonlinear system of $2N + 1$ “steady-state” equations, with N the number of harmonics taken into account, that are globally coupled *via* a source term that accounts for the time-derivative. A so-called Newton–Krylov solution method is adopted, which consists in solving the high-dimensional linear system arising at each Newton iteration using a Krylov subspace method, here the GMRES algorithm. For adequate convergence of GMRES, we presented a new preconditioner for the TSM Jacobian matrix, referred to as the “block-circulant preconditioner”. It consists in replacing the diagonal blocks of the original TSM Jacobian, which are the steady Navier–Stokes Jacobians evaluated at the $2N + 1$ instants of the temporal grid, by a unique matrix that is here chosen as the steady Navier–Stokes Jacobian evaluated at the mean flow. The resulting preconditioner possesses a block-circulant pattern, which implies that it can be block-diagonalized by the Fourier transform. Efficient inversion of the block-circulant preconditioner is thus possible in the Fourier space. Numerical testing of this solution method, both for imposed frequency (forced systems) and unknown frequency (autonomous systems) test cases, showed that it possesses good N -robustness properties. A comparison with a more classical block-Jacobi preconditioner showed largely superior performances of the block-circulant preconditioner for moderate and large numbers of harmonics ($5 \leq N \leq 40$). This N -robustness property makes the block-circulant preconditioner particularly well-suited for computing LCO’s with high harmonics.

Flutter bifurcation analysis in viscous flows

The second part of this manuscript mainly focused on the exploration of the flutter bifurcation. The succession of chapters took us on an increasingly nonlinear journey,

from the study of the linear stability of the typical section in chapter 4, to the exploration of the sub-/super-critical nature of the associated bifurcation in chapter 5, and finally to the investigation of a secondary instability of flutter LCO solutions in chapter 6.

In chapter 4, we revisited the linear stability of the typical aeroelastic section with an incompressible Navier–Stokes flow model. Four types of instabilities were observed through extensive parametric explorations: classical coupled-mode flutter and static divergence at high-enough reduced velocities, vortex-induced vibrations at low reduced velocities, and a transition from coupled-mode flutter to single-mode flutter as the steady angle of attack of the plate increases. Overall, the main interests of the global stability approach used in this work is its ability to investigate the stability of steady-state solutions that present nonlinear flow features (recirculation regions, etc) and thus to capture, with one flow model, a variety of fluid-structure instabilities that must usually be investigated with dedicated flow models. After having identified these different instabilities, a parametric exploration of the effect of mass ratio and Reynolds number was proposed, in the case of zero steady angle of attack. Neutral curves for coupled-mode flutter, divergence and vortex-induced vibrations were computed in the (\tilde{m}, U^*) (or alternatively (\tilde{m}, q^*)) plane. For all mass ratios, vortex-induced vibrations occurred on a range of reduced velocities centered around two well-defined values that correspond to the coincidence of the pitching (resp. heaving) frequencies with the hydrodynamic wake instability frequency. At a particular, low, mass ratio \tilde{m}^{II} , a codimension-two point was found where both flutter and divergence occur at the same critical reduced velocity. For mass ratios higher than \tilde{m}^{II} , flutter occurs prior to divergence whereas divergence precedes flutter (if it exists at all) for mass ratios below \tilde{m}^{II} . In addition, it was shown that for low mass ratios slightly below \tilde{m}^{II} , the flutter neutral curve “folds back” towards higher reduced velocities and mass ratios, thus allowing the flutter mode to restabilize at (very) high reduced velocities. By varying the Reynolds number, it was shown that high Reynolds postpone this restabilization of the flutter mode, leading us to attribute this unexpected behavior to viscosity effects. Comparisons of the stability predictions obtained from the global Navier–Stokes stability approach to several variants of quasi-steady approaches and to the Theodorsen model (previously calibrated with steady viscous aerodynamics coefficients). For the low-frequency flutter mass ratio occurring at large mass ratios, it was shown that all approaches provide accurate enough prediction of flutter thresholds. As mass ratio is decreased however, the critical flutter frequency decreases and the quasi-steady and Theodorsen models progressively deviate from the reference Navier–Stokes neutral curve. As a consequence, if accuracy is required in the evaluation of flutter thresholds, it is believed that the modelization of the unsteady viscous effects through the global stability approach is more adequate, in particular at low mass ratios where the Theodorsen model was found to be unconservative.

Having gained a parametric vision of the linear flutter instability in the previous chapter, an objective of chapter 5 was to enrich this vision with nonlinear elements. Using weakly nonlinear analysis, we determined the normal form associated to the flutter bifurcation on a large domain of the (\tilde{m}, Re) parametric plane, $10 \leq \tilde{m} \leq 10^4$ and $10 \leq Re \leq 10^4$. This allowed us to draw a “criticity map” (fig. 5.3.3) that represent, with one scalar indicator, the sub- or super-critical nature of the flutter bifurcation. A general trend could be observed that is: low mass ratios and high-Reynolds numbers foster subcritical bifurcations. In order to get further insights into the physical mechanisms at the origin of these transitions, we proposed a decomposition of the cubic normal form coefficient, allowing to separate contributions

from different physical origins (fluid *vs* geometric nonlinearities) and from different spatial locations in the flow. Applying this approach, we showed that the transition from super- to sub-critical when decreasing the mass ratio at fixed Reynolds number is primarily dictated by the fluid nonlinearity, and more specifically, by the behavior of the leading-edge shear layers. In a second step, we explored the highly nonlinear regime of the flutter bifurcation at a fixed Reynolds number ($Re = 500$), using a combination of TSM and time-marching computations. For high mass ratio, where the bifurcation is supercritical, stable periodic solutions with low amplitude appear above the critical reduced velocity U_c^* . These solutions oscillate at a low frequency leading to quasi-steady flow features with well-attached shear layers. On the contrary, for a low mass ratio ($\tilde{m} = 60$), where the bifurcation is subcritical, high-amplitude periodic solutions were observed below the critical velocity. These solutions possess a high oscillation frequency and are characterized by strong unsteady fluid effects involving the shedding of multiple vortical structures. The largest vortical structure corresponds to a strong leading edge vortex that is generated when the plate reaches high angles of attack. Finally, at an intermediate mass ratio ($\tilde{m} = 120$), we discovered an unusual scenario where the bifurcation is supercritical but still allows high-amplitude highly-unsteady solutions below the critical velocity. This situation is made possible by the succession of two-fold bifurcations of LCO's that follow the original supercritical Hopf bifurcation. We ended the study by noting that the double-fold scenario is consistent with previous unexplained experimental results by [Amandolese et al. \(2013\)](#), where both low-amplitude supercritical solutions and high-amplitude subcritical oscillations were observed.

In the case of the supercritical flutter bifurcations occurring at large mass ratios, it was observed that the periodic solutions that exist immediately above the critical velocity U_c^* transition to quasi-periodic oscillations with slow amplitude modulations, for reduced velocities higher than $1.078 U_c^*$. This transition was the focus of chapter 6 where we analyzed the linear (asymptotic) Floquet stability of the periodic solutions that underlie the observed quasi-periodic solutions. In order to first compute these — possibly unstable — LCO's, and then assess their Floquet stability, we proposed an original approach, entirely based on the Time Spectral Method. Using this method, we showed that the appearance of quasi-periodic solutions is due to the destabilization of a pair of complex conjugate asynchronous Floquet modes. An analysis of the pitching and heaving components of these modes and the associated perturbation revealed that the low-frequency modulation is linked to the continuous drift of the phase difference between pitching and heaving perturbation signals across the modulation period. This is made possible by the fact that the pitching angle predominantly oscillates at a slightly higher frequency than the heaving displacement. When the pitching motion precedes the heaving motion, energy is extracted from the flow with perturbation dynamics that are reminiscent of the classical flutter instability, and the perturbation amplitude increases. On the contrary, when heaving precedes pitching, the solid energy is dissipated into the flow *via* so-called “anti-flutter” dynamics, and the perturbation amplitude decreases.

Towards large-scale linear stability analysis

The last part of this thesis aimed at initiating the extension of the different numerical methods (global linear stability, WNL, TSM, Floquet-TSM) successfully used in the previous chapters for studying the flutter bifurcation of a 2D typical section, to large-scale 3D configurations.

A first step towards this long-term goal is presented in chapter 7, where we focused on performing purely hydrodynamic (the structure is frozen) global stability analysis of three-dimensional incompressible flows. We proposed to use preconditioned Krylov-subspace methods for solving the high-dimensional linearized Navier–Stokes system, which constitutes the critical building brick of both the steady flow computation *via* a Newton method and the eigenvalue computation *via* shift-invert Krylov–Schur iterations. With that strategy, the main issue is to find a good preconditioner for the linearized incompressible Navier–Stokes equations. After a performance study and a 2D benchmark against several other state-of-the-art preconditioners, we opted for the so-called modified Augmented Lagrangian (mAL) preconditioner. Then, a parallel implementation of the mAL preconditioner using the FreeFEM finite element language and its interface to PETSc/SLEPc was presented and made available online. Three-dimensional numerical experiments showed that the proposed approach generally outperforms the more “basic” approach of solving the linear systems with a (parallel) direct solver, like MUMPS. In particular, the mAL-preconditioned Krylov subspace solver featured much better parallel performance. Finally, we showed that the proposed implementation was able to perform the linear stability analysis of a high-dimensional problem with about 75 million unknowns, with satisfying strong scaling properties up to 2,048 processes.

8.2 Perspectives

Amplitude-robust TSM preconditioning

Though robust to the number of harmonics, the performance of the block-circulant TSM preconditioner, proposed in the present manuscript, was seen to significantly deteriorate for increasing LCO amplitudes. For some of the very-high-amplitude flutter LCOs encountered in this work (typically, fig. 5.5.4b), the block-circulant preconditioner was unable to sufficiently decrease the linear solver residual, thus causing a stagnation of the nonlinear residual and preventing the Newton method to converge in reasonable amounts of time. Here, we circumvented the issue by using time-marching computations for computing these particularly challenging solutions. However, a TSM solution robust also for these cases would be highly valuable. Towards this goal, we suggest several directions of research. First, from the point of view of the TSM preconditioner itself, it would be interesting to compare the performance of the block-circulant preconditioner to the current state-of-the-art in the transonic aeroelastic literature, which seems to be the work of [Mundis et al. \(2017\)](#). Another interesting preconditioning strategy may be to consider multigrid-in-time preconditioners [[Falgout et al. 2014](#)], where the linearized problem is solved subsequently on a series of temporal grids with different coarsening. Finally, as noted in the work of [Leffell et al. \(2014\)](#), some highly-nonlinear TSM solutions may need the addition of numerical viscosity (in time) in order to converge. Thus, it is possible that our inability to converge the highest-amplitude flutter solutions is partly due to insufficient stability properties of the original TSM scheme, for these particular cases. As a consequence, in our opinion, the interaction between the solver’s — in particular the preconditioner’s — performance and the TSM scheme (with or without numerical viscosity) should be carefully studied, even before trying to further improve the preconditioner.

Understanding the low-mass-ratio and high-Reynolds sub-/super-critical transitions

In chapter 5, we investigated in detail the high-mass-ratio sub-/super-critical transition ($\tilde{m} = \tilde{m}_u$ in fig. 5.3.2). However, two other transitions have been left mostly unexplored: the low-mass-ratio sub-/super-critical ($\tilde{m} = \tilde{m}_l$) at fixed Reynolds number and the high-Reynolds transition at fixed, high, mass ratio. Both this transition should be object of further investigations. In particular, the high-Reynolds transition occurs typically at very large mass ratio that yield low-frequency flutter. As a consequence, it is likely that (quasi-)steady flow models can be used to explain this transition. In other words, it is likely that the nature of the bifurcation could be predicted directly from the analysis of the curve representing the steady aerodynamic coefficient as a function of angle of attack.

Criticality maps for the flutter bifurcation: effect of the typical section profile

Throughout the present manuscript, the typical section investigated consisted in a thin rectangular plate. Naturally, the question of the robustness of the presented bifurcation scenarios with respect to the section shape is in order. Typically, an interesting continuation of the studies of chapter 5 would be to produce criticality maps, similar to fig. 5.3.3, but for various section shapes: airfoils in particular, but also other section representative of bridge decks, for example. To this end, the mesh adaptation strategy proposed in chapter 2 should come in particularly handy.

Flutter bifurcation at higher Reynolds: turbulence modeling and effect of coherent fluid fluctuations

Another intriguing perspective from chapter 5 is to push the parametric explorations to higher Reynolds numbers, so as to progressively extend the study to regimes typical of classical airplane flight. Indeed, knowing whether airplanes are likely to undergo subcritical flutter or not could have a major impact on the degree of confidence an engineer can put in flutter assessment *via* classical linear stability analyses, and as a consequence, in the way flutter margins are estimated.

On this route towards higher-velocity flows, significant extensions of the present work will be needed. First, the increase in Reynolds number supposes an adequate modelization of turbulence. Given the critical importance of the leading-edge shear layers and dynamic stall vortex, highlighted for the high-mass-ratio super-/sub-critical transition (\tilde{m}_u), the chosen turbulence model should be able to accurately predict such complex unsteady nonlinear behaviors at high-Reynolds numbers. For transitional Reynolds numbers, the issue is even more delicate as laminar separation bubbles should be appropriately modeled given their impact on the aeroelastic behavior [Poirel et al. 2008]. If the accurate modelization of laminar separation bubbles is known to be a challenging CFD problem, recent progresses in that area [Bernardos et al. 2019] should be reviewed for flutter bifurcation study purposes. The second extension consists in relaxing the incompressibility constraint used across this thesis and would be eventually required in order to tackle flows typical of airliners that fly in the transonic regime.

Tools for flutter bifurcation study with increasingly complex flow models

Generally speaking, the progressive complexification of the fluid model, whether it is due to turbulence or compressibility aspects, makes it increasingly cumbersome¹ to use an analytical technique such as the weakly nonlinear analysis proposed in this thesis. On the other hand, the Time Spectral Method for its part, is less impacted by increasing the flow model complexity. Indeed, in modern CFD codes, TSM is seen as an additional temporal scheme and is thus naturally compatible with any flow model. As a consequence, a possible approach for studying the flutter bifurcation in complex turbulent transonic flows is to only resort to TSM solutions (and to time-marching computations for reference). Still, it may be considered that some features of the weakly nonlinear analysis are of great practical interest: in particular its ability to represent the type of bifurcation with only one scalar value, determined from the normal form coefficients. One could then wonder if it is possible to extract, from a fully nonlinear TSM solver, only the information needed to establish the normal form of the bifurcation.

Large-scale linear stability analysis of elastic structures in viscous flows

A natural continuation of the developments in chapter 7 is to extend the proposed approach, and in particular the modified Augmented Lagrangian preconditioner, in order to perform large-scale parallel linear stability analysis of elastic structures in three-dimensional flows. In practice, this work has already started in the course of this thesis, using the “Lagrangian-based” approach for linear stability of fluid-structure systems proposed in [J.-L. Pfister et al. 2019]. Preliminary numerical tests of the parallel implementation have been successfully run on a two-dimensional configuration. A large-scale three-dimensional demonstrator will hopefully be soon achieved.

¹And in our opinion, highly unlikely and undesirable ...

BIBLIOGRAPHY

- ABRAMSON, H. N. (1969). "Hydroelasticity: a review of hydrofoil flutter". *Applied Mechanics Reviews* 22.2, pp. 115–121.
- ADAMS, M., BAYRAKTAR, H., KEAVENY, T., and PAPADOPOULOS, P. (2004). "Ultra-scalable Implicit Finite Element Analyses in Solid Mechanics with over a Half a Billion Degrees of Freedom". In: *Proceedings of the 2004 ACM/IEEE conference on Supercomputing*. SC04. IEEE Computer Society, 34:1–34:15.
- ÅKERVIK, E., BRANDT, L., HENNINGSON, D. S., HEPFFNER, J., MARXEN, O., and SCHLATTER, P. (2006). "Steady solutions of the Navier-Stokes equations by selective frequency damping". *Physics of Fluids* 18.6, pp. 1–5.
- ALAUZET, F. and LOSEILLE, A. (2016). "A decade of progress on anisotropic mesh adaptation for computational fluid dynamics". *CAD Computer Aided Design* 72, pp. 13–39.
- ALBANO, E. and RODDEN, W. P. (1969). "A Doublet-Lattice Method for Calculating Lift Distributions on Planar and Non-Planar Configurations in Subsonic Flows". *AIAA Journal* 7.2, pp. 279–285.
- AMANDOLESE, X., MICHELIN, S., and CHOQUEL, M. (2013). "Low speed flutter and limit cycle oscillations of a two-degree-of-freedom flat plate in a wind tunnel". *Journal of Fluids and Structures* 43, pp. 244–255.
- AMANDOLÈSE, X. (2016). "Low Speed Flutter and Post-Critical Behaviour of Flat Plate and Naca0018 Section Models in a Wind Tunnel". In: *First International Symposium on Flutter and its Application*, pp. 447–455.
- AMESTOY, P., DUFF, I., L'EXCELLENT, J.-Y., and KOSTER, J. (2001). "A fully asynchronous multifrontal solver using distributed dynamic scheduling". *SIAM Journal on Matrix Analysis and Applications* 23.1, pp. 15–41.
- AMIRALAEI, M. R., ALIGHANBARI, H., and HASHEMI, S. M. (2010). "An investigation into the effects of unsteady parameters on the aerodynamics of a low Reynolds number pitching airfoil". *Journal of Fluids and Structures* 26.6, pp. 979–993.
- ANTHEAUME, S. and CORRE, C. (2011). "Implicit Time Spectral Method for Periodic Incompressible Flows". PhD thesis.
- ASHBY, M. F. (1999). *Material Selection in Mechanical Design*. Butterworth. Oxford. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).
- ASSEMAT, P., FABRE, D., and MAGNAUDET, J. (2012). "The onset of unsteadiness of two-dimensional bodies falling or rising freely in a viscous fluid: A linear study". *Journal of Fluid Mechanics* 690, pp. 173–202.
- BADCOCK, K. J., TIMME, S., MARQUES, S., KHODAPARAST, H., PRANDINA, M., MOTTERSHEAD, J. E., SWIFT, A., DA RONCH, A., and WOODGATE, M. A. (2011). "Transonic aeroelastic simulation for instability searches and uncertainty analysis". *Progress in Aerospace Sciences* 47.5, pp. 392–423.

- BADCOCK, K. J., WOODGATE, M. A., ALLAN, M. R., and BERAN, P. S. (2008). "Wing-rock limit cycle oscillation prediction based on computational fluid dynamics". *Journal of Aircraft* 45.3, pp. 954–961.
- BADCOCK, K. J., WOODGATE, M. A., and RICHARDS, B. E. (2005). "Direct Aeroelastic Bifurcation Analysis of a Symmetric Wing Based on the Euler Equations". *Journal of Aircraft* 42.3, pp. 731–737.
- BADIA, S. and CODINA, R. (2007). "On some fluid–structure iterative algorithms using pressure segregation methods. Application to aeroelasticity". *International Journal for Numerical Methods in Engineering* 72.1, pp. 46–71.
- BAGHERI, S., ÅKERVIK, E., BRANDT, L., and HENNINGSON, D. S. (2009a). "Matrix-Free Methods for the Stability and Control of Boundary Layers". *AIAA Journal* 47.5, pp. 1057–1068.
- BAGHERI, S., SCHLATTER, P., SCHMID, P. J., and HENNINGSON, D. S. (2009b). "Global stability of a jet in crossflow". *Journal of Fluid Mechanics* 624, pp. 33–44. arXiv: [1010.3766](https://arxiv.org/abs/1010.3766).
- BALAY, S., ABHYANKAR, S., ADAMS, M. F., BROWN, J., BRUNE, P., BUSCHELMAN, K., DALCIN, L., DENER, A., EIJKHOUT, V., GROPP, W., KARPEYEV, D., KAUSHIK, D., KNEPLEY, M., MAY, D., MCINNES, L. C., MILLS, R. T., MUNSON, T., RUPP, K., SANAN, P., SMITH, B., ZAMPINI, S., ZHANG, H., and ZHANG, H. (2019). *PETSc Web page*.
- BARKLEY, D. and TUCKERMAN, L. S. (1997). "Stokes preconditioning for the inverse power method". In: *Lecture Notes in Physics : Proc. of the Fifteenth Int'l. Conf. on Numerical Methods in Fluid Dynamics*. Ed. by P. KUTLER, J. FLORES, and J.-J. CHATOT. New York: Springer, pp. 75–76.
- BARKLEY, D. and HENDERSON, R. D. (1996). "Three-dimensional Floquet stability analysis of the wake of a circular cylinder". *Journal of Fluid Mechanics* 322, pp. 215–241.
- BARNES, C. J. and VISBAL, M. R. (2018). "On the role of flow transition in laminar separation flutter". *Journal of Fluids and Structures* 77, pp. 213–230.
- BAZILEVS, Y. and HUGHES, T. J. (2007). "Weak imposition of Dirichlet boundary conditions in fluid mechanics". *Computers and Fluids* 36.1, pp. 12–26.
- BEAUME, C. (2017). "Adaptive Stokes Preconditioning for Steady Incompressible Flows". *Communications in Computational Physics* 22.02, pp. 494–516.
- BENDIKSEN, O. O. (2011). "Review of unsteady transonic aerodynamics: Theory and applications". *Progress in Aerospace Sciences* 47.2, pp. 135–167.
- BENZI, M. (2008). "Block preconditioning of real-valued iterative algorithms for complex linear systems". *IMA Journal of Numerical Analysis* 28, pp. 598–618.
- BENZI, M., GOLUB, G. H., and LIESEN, J. (2005). "Numerical solution of saddle point problems". *Acta Numerica* 14, pp. 1–137.
- BENZI, M. and OLSHANSKII, M. A. (2006). "An Augmented Lagrangian-Based Approach to the Oseen Problem". *SIAM Journal on Scientific Computing* 28.6, pp. 2095–2113.

- BENZI, M. and OLSHANSKII, M. A. (2011a). "Field-of-Values Convergence Analysis of Augmented Lagrangian Preconditioners for the Linearized Navier–Stokes Problem". *SIAM Journal on Numerical Analysis* 49.2, pp. 770–788.
- BENZI, M., OLSHANSKII, M. A., and WANG, Z. (2011b). "Modified augmented Lagrangian preconditioners for the incompressible Navier-Stokes equations". *International Journal for Numerical Methods in Fluids* 66.4, pp. 486–508.
- BENZI, M. and WANG, Z. (2011c). "Analysis of Augmented Lagrangian-Based Preconditioners for the Steady Incompressible Navier–Stokes Equations". *SIAM Journal on Scientific Computing* 33.5, pp. 2761–2784.
- BENZI, M. and WANG, Z. (2013). "A parallel implementation of the modified augmented Lagrangian preconditioner for the incompressible Navier-Stokes equations". *Numerical Algorithms* 64.1, pp. 73–84.
- BERGEON, A., HENRY, D., BENHADID, H., and TUCKERMAN, L. S. (1998). "Marangoni convection in binary mixtures with Soret effect". *Journal of Fluid Mechanics* 375, pp. 143–177.
- BERNARDOS, L. F., RICHEZ, F., and GLEIZE, V. (2019). "RANS modeling of Laminar Separation Bubbles around Airfoils at Low Reynolds conditions".
- BESCH, P. K. and LIU, Y.-N. (1971). *Flutter and divergence characteristics of four low mass ratio hydrofoils*. Tech. rep. Naval Ship Research and Development Center.
- BHAT, S. S. and GOVARDHAN, R. N. (2013). "Stall flutter of NACA 0012 airfoil at low Reynolds numbers". *Journal of Fluids and Structures* 41, pp. 166–174.
- BISPLINGHOFF, R. L., ASHLEY, H., and HALFMAN, R. L. (1955). *Aeroelasticity*.
- BLONDEAU, C. and LIAUZUN, C. (2019). "A modular implementation of the time spectral method for aeroelastic analysis and optimization on structured meshes". In: *International Forum on Aeroelasticity and Structural Dynamics*. Savannah, Georgia.
- BOUDREAU, M., DUMAS, G., RAHIMPOUR, M., and OSHKAI, P. (2018). "Experimental investigation of the energy extraction by a fully-passive flapping-foil hydrokinetic turbine prototype". *Journal of Fluids and Structures* 82, pp. 446–472.
- BROOKS, A. N. and HUGHES, T. J. R. (1982). "Streamline upwind/Petrov-Galerkin formulations for convection dominated flows with particular emphasis on the incompressible Navier-Stokes equations". *Computer Methods in Applied Mechanics and Engineering* 32.1-3, pp. 199–259.
- BRUNO, L. and FRANSOS, D. (2008). "Evaluation of Reynolds number effects on flutter derivatives of a flat plate by means of a computational approach". *Journal of Fluids and Structures* 24.7, pp. 1058–1076.
- BRUNTON, S. L. and ROWLEY, C. W. (2013). "Empirical state-space representations for Theodorsen's lift model". *Journal of Fluids and Structures* 38, pp. 174–186.
- BRYNJELL-RAHKOLA, M., TUCKERMAN, L. S., SCHLATTER, P., and HENNINGSON, D. S. (2017). "Computing Optimal Forcing Using Laplace Preconditioning". *Communications in Computational Physics* 22.05, pp. 1508–1532.

- CAHOUE T A N, J. and CHABARD, D. J.-P. (1988). "Some fast 3D finite element solvers for the generalized Stokes problem". *International Journal for Numerical Methods in Fluids* 8, pp. 869–895.
- CALIARI, M., KANDOLF, P., OSTERMANN, A., and RAINER, S. (2014). "Comparison of software for computing the action of the matrix exponential". *BIT Numerical Mathematics* 54.1, pp. 113–128.
- CANUTO, C., HUSSAINI, M. Y., QUARTERONI, A., and ZANG, T. A. (1988). *Spectral Methods in Fluid Dynamics*. Springer-Verlag.
- CARINI, M., AUTERI, F., and GIANNETTI, F. (2015). "Centre-manifold reduction of bifurcating flows". *Journal of Fluid Mechanics* 767, pp. 109–145.
- CARTE, G., DUŠEK, J., and FRAUNIÉ, P. (1995). "A spectral time discretization for flows with dominant periodicity". *Journal of Computational Physics* 120.2, pp. 171–183.
- CHAE, E. J., AKCABAY, D. T., LELONG, A., ASTOLFI, J. A., and YOUNG, Y. L. (2016). "Numerical and experimental investigation of natural flow-induced vibrations of flexible hydrofoils". *Physics of Fluids* 28.7, p. 075102.
- CHAE, E. J., AKCABAY, D. T., and YOUNG, Y. L. (2013). "Dynamic response and stability of a flapping foil in a dense and viscous fluid". *Physics of Fluids* 25.
- CHAN, T. F. and VAN DER VORST, H. A. (1997). "Approximate and Incomplete Factorizations". In: *Parallel Numerical Algorithms*. Springer, pp. 167–202.
- CHARRU, F. (2011). *Hydrodynamic Instabilities*. Cambridge University Press.
- CHOI, S., LEE, K., POTSDAM, M. M., and ALONSO, J. J. (2014). "Helicopter rotor design using a time-spectral and adjoint-based method". *Journal of Aircraft* 51.2, pp. 412–423.
- CHOMAZ, J.-M. (2005). "Global Instabilities in Spatially Developing Flows: Non-Normality and Nonlinearity". *Annual Review of Fluid Mechanics* 37.1, pp. 357–392.
- CHRISTODOULOU, K. N. and SCRIVEN, L. E. (1988). "Finding Leading Modes of a Viscous Free Surface Flow: An Asymmetric Generalized Eigenproblem". *J. Sci. Comput.* 3.4, pp. 355–406.
- CISONNI, J., LUCEY, A. D., ELLIOTT, N. S., and HEIL, M. (2017). "The stability of a flexible cantilever in viscous channel flow". *Journal of Sound and Vibration* 396, pp. 186–202.
- CITRO, V., LUCHINI, P., GIANNETTI, F., and AUTERI, F. (2017). "Efficient stabilization and acceleration of numerical simulation of fluid flows by residual recombination". *Journal of Computational Physics* 344. April, pp. 234–246.
- CITRO, V., GIANNETTI, F., LUCHINI, P., and AUTERI, F. (2015). "Global stability and sensitivity analysis of boundary-layer flows past a hemispherical roughness element". *Physics of Fluids* 27.8, p. 084110.
- CITRO, V., TCHOUFAG, J., FABRE, D., GIANNETTI, F., and LUCHINI, P. (2016). "Linear stability and weakly nonlinear analysis of the flow past rotating spheres". *Journal of Fluid Mechanics* 807.2016, pp. 62–86.

- CLIFFE, K. A., GARRATT, T. J., and SPENCE, A. (1993). "Eigenvalues of the discretized Navier-Stokes equation with application to the detection of Hopf bifurcations". *Advances in Computational Mathematics* 1.3, pp. 337–356.
- COLLER, B. and CHAMARA, P. (2004). "Structural non-linearities and the nature of the classic flutter instability". *Journal of Sound and Vibration* 277.4-5, pp. 711–739.
- COLLIS, S. S. and HEINKENSCHLOSS, M. (2002). *Analysis of SUPG Method Applied to the Solution of Optimal Control Problems*. Tech. rep. March. Houston: Rice University.
- COSSU, C. and MORINO, L. (2000). "On the Instability of a Spring-Mounted Circular Cylinder in a Viscous Flow at Low Reynolds Numbers". *Journal of Fluids and Structures* 14.2, pp. 183–196.
- CRIVELLINI, A. and BASSI, F. (2011). "An implicit matrix-free Discontinuous Galerkin solver for viscous and turbulent aerodynamic simulations". *Computers and Fluids* 50.1, pp. 81–93.
- CROON, G. de, PERÇIN, M., REMES, B., RUIJSINK, R., and DE WAGTER, C. (2016). *The DelFly*.
- CURET, O. M., SWARTZ, S. M., and BREUER, K. S. (2013). "An aeroelastic instability provides a possible basis for the transition from gliding to flapping flight". *Journal of the Royal Society Interface* 10.80.
- CYR, E. C., SHADID, J., and WILDEY, T. (2014). "Approaches for Adjoint-Based A Posteriori Analysis of Stabilized Finite Element Methods". *SIAM Journal on Scientific Computing* 36.2, A766–A791.
- DE LANGRE, E. (2002). *Fluides et solides*. Editions Ecole Polytechnique.
- DECONINCK, B. and NATHAN KUTZ, J. (2006). "Computing spectra of linear operators using the Floquet-Fourier-Hill method". *Journal of Computational Physics* 219.1, pp. 296–321.
- DENG, J. and CAULFIELD, C. P. (2016). "Dependence on aspect ratio of symmetry breaking for oscillating foils: Implications for flapping flight". *Journal of Fluid Mechanics* 787, pp. 16–49.
- DEPARIS, S., FORTI, D., GRANDPERRIN, G., and QUARTERONI, A. (2016). "FaCSI: A block parallel preconditioner for fluid–structure interaction in hemodynamics". *Journal of Computational Physics* 327, pp. 700–718.
- DIJKSTRA, H. A., WUBS, F. W., CLIFFE, A. K., DOEDEL, E., DRAGOMIRESCU, I. F., ECKHARDT, B., GELFGAT, A. Y., HAZEL, A. L., LUCARINI, V., SALINGER, A. G., PHIPPS, E. T., JUAN, S. U., SCHUTTELAARS, H., TUCKERMAN, L. S., and THIELE, U. (2014). "Numerical bifurcation methods and their application to fluid dynamics: Analysis beyond simulation". *Communications in Computational Physics* 15.1, pp. 1–45.
- DIMITRIADIS, G. (2017). *Introduction To Nonlinear Aeroelasticity*. John Wiley & Sons, Ltd.
- DIMITRIADIS, G. and LI, J. (2009). "Bifurcation Behavior of Airfoil Undergoing Stall Flutter Oscillations in Low-Speed Wind Tunnel". *AIAA Journal* 47.11, pp. 2577–2596.

- DOLEAN, V., JOLIVET, P., and NATAF, F. (2015). *An Introduction to Domain Decomposition Methods: Algorithms, Theory and Parallel Implementation*. SIAM.
- DONEA, J., HUERTA, A., PONTHOT, J.-P., and RODRÍGUEZ-FERRAN, A. (2004). "Arbitrary Lagrangian–Eulerian Methods". In: *Encyclopedia of Computational Mechanics*. Chap. 14.
- DOWELL, E., EDWARDS, J., and STRGANAC, T. (2003). "Nonlinear aeroelasticity". *Journal of Aircraft* 40.5, pp. 857–874.
- DOWELL, E. H., CURTISS, H. C., SCANLAN, R. H., and SISTO, F. (1989). *A modern Course in Aeroelasticity*. Springer.
- DRAZIN, P. G. and REID, W. H. (2004). *Hydrodynamic stability*. Cambridge university press.
- EHRENSTEIN, U. and GALLAIRE, F. (2005). "On two-dimensional temporal modes in spatially evolving open flows: The flat-plate boundary layer". *Journal of Fluid Mechanics* 536, pp. 209–218.
- EKICI, K. and HALL, K. C. (2011). "Harmonic Balance Analysis of Limit Cycle Oscillations in Turbomachinery". *AIAA Journal* 49.7, pp. 1478–1487.
- ELMAN, H. C., HOWLE, V. E., SHAHID, J., SHUTTLEWORTH, R., and TUMINARO, R. S. (2008). "A taxonomy and comparison of parallel block multi-level preconditioners for the incompressible Navier–Stokes equations". *Journal of Computational Physics* 227.3, pp. 1790–1808.
- ELSTON, J. R., SHERIDAN, J., and BLACKBURN, H. M. (2004). "Two-dimensional floquet stability analysis of the flow produced by an oscillating circular cylinder in quiescent fluid". *European Journal of Mechanics, B/Fluids* 23.1, pp. 99–106.
- FABRE, D., CITRO, V., FERREIRA SABINO, D., BONNEFIS, P., SIERRA, J., GIANNETTI, F., and PIGOU, M. (2018). "A Practical Review on Linear and Nonlinear Global Approaches to Flow Instabilities". *Applied Mechanics Reviews* 70.6, p. 060802.
- FALGOUT, R. D., FRIEDHOFF, S., KOLEV, T. V., MACLACHLAN, S. P., and SCHRODER, J. B. (2014). "Parallel time integration with multigrid". *SIAM Journal on Scientific Computing* 36.6, pp. C635–C661.
- FARRELL, P. E., MITCHELL, L., and WECHSUNG, F. (2018). "An augmented lagrangian preconditioner for the 3d stationary incompressible navier–stokes equations at high reynolds number". arXiv: [arXiv:1810.03315v1](https://arxiv.org/abs/1810.03315v1).
- FLOQUET, G. (1883). "Sur les équations différentielles linéaires à coefficients périodiques". In: *Annales scientifiques de l'École normale supérieure*. Vol. 12, pp. 47–88.
- FORTIN, M. and BREZZI, F. (1991). *Mixed and Hybrid Finite Element Methods*. Springer-Verlag.
- FRAZZA, L. (2018). "3D anisotropic mesh adaptation for Reynolds Averaged simulations". PhD thesis. Université Pierre et Marie Curie.
- GAI, G. and TIMME, S. (2016). "Nonlinear reduced-order modelling for limit-cycle oscillation analysis". *Nonlinear Dynamics* 84.2, pp. 991–1009.
- GALLAIRE, F., BOUJO, E., MANTIC-LUGO, V., ARRATIA, C., THIRIA, B., and MELIGA, P. (2016). "Pushing amplitude equations far from threshold: Application to the

- supercritical Hopf bifurcation in the cylinder wake". *Fluid Dynamics Research* 48.6, pp. 1–12.
- GAO, C. and ZHANG, W. (2020). "Transonic aeroelasticity: A new perspective from the fluid mode". *Progress in Aerospace Sciences* August, p. 100596.
- GAO, C., ZHANG, W., LI, X., LIU, Y., QUAN, J., YE, Z., and JIANG, Y. (2017). "Mechanism of frequency lock-in in transonic buffeting flow". *Journal of Fluid Mechanics* 818, pp. 528–561.
- GARRIGUES, E. (2018). "A Review of Industrial Aeroelasticity Practices at Dassault Aviation for Military Aircraft and Business Jets". *AerospaceLab Journal* 14, pp. 1–34.
- GEUZAIN, C. and REMACLE, J.-F. (2009). "Gmsh: A 3-D finite element mesh generator with built-in pre- and post-processing facilities". *International Journal for Numerical Methods in Engineering* 79.11, pp. 1309–1331.
- GIANETTI, F. and LUCHINI, P. (2007). "Structural sensitivity of the first instability of the cylinder wake". *Journal of Fluid Mechanics* 581, pp. 167–197.
- GILMORE, R. J. and STEER, M. B. (1991). "Nonlinear circuit analysis using the method of harmonic balance—a review of the art. Part I. Introductory concepts". *International Journal of Microwave and Millimeter-Wave Computer-Aided Engineering* 1.1, pp. 22–37.
- GLAUERT, H. (1930). "The force and moment on an oscillating aerofoil". *Vorträge aus dem Gebiete der Aerodynamik und verwandter Gebiete* 1, pp. 88–95.
- GOPINATH, A. K. and JAMESON, A. (2005). "Time Spectral Method for Periodic Unsteady Computations over Two-and Three-Dimensional Bodies". In: *43rd AIAA Aerospace Sciences Meeting and Exhibit*. Reno.
- GOPINATH, A. K. and JAMESON, A. (2006). "Application of the Time Spectral Method to Periodic Unsteady Vortex Shedding". In: *44th AIAA Aerospace Sciences Meeting and Exhibit*. Reno, Nevada.
- GOZA, A., COLONIUS, T., and SADER, J. E. (2018). "Global modes and nonlinear analysis of inverted-flag flapping". *Journal of Fluid Mechanics* 857, pp. 312–344. arXiv: [1709.09745](https://arxiv.org/abs/1709.09745).
- GUERMOND, J. L., MINEV, P., and SHEN, J. (2006). "An overview of projection methods for incompressible flows". *Computer Methods in Applied Mechanics and Engineering* 195.44-47, pp. 6011–6045.
- GÜNER, H., DIMITRIADIS, G., and TERRAPON, V. E. (2018). "Inviscid and viscous flow modeling for FAS transonic flutter calculations". In: *31st Congress of the International Council of the Aeronautical Sciences, ICAS 2018*, pp. 1–10.
- HALL, K. C., THOMAS, J. P., and CLARK, W. S. (2002). "Computation of unsteady nonlinear flows in cascades using a harmonic balance technique". *AIAA Journal* 40.5, pp. 879–886.
- HE, S., JONSSON, E., MADER, C. A., and MARTINS, J. (2018). "A Coupled Newton–Krylov Time Spectral Solver for Flutter Prediction". *2018 AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference* January, pp. 1–17.

- HE, X. and VUIK, C. (2016). "Comparison of Some Preconditioners for the Incompressible Navier-Stokes Equations". *Numerical Mathematics* 9.2, pp. 239–261.
- HECHT, F. (2012). "New development in FreeFem++". *J. Numer. Math.* 20.3-4, pp. 251–265.
- HECHT, F. (2006). *BAMG : Bidimensional Anisotropic Mesh Generator*. Tech. rep. Université Pierre et Marie Curie.
- HEISTER, T. and RAPIN, G. (2013). "Efficient augmented Lagrangian-type preconditioning for the Oseen problem using Grad-Div stabilization". *International Journal for Numerical Methods in Fluids* 71.1, pp. 118–134.
- HERNANDEZ, V., ROMAN, J. E., and VIDAL, V. (2005). "SLEPc: A Scalable and Flexible Toolkit for the Solution of Eigenvalue Problems". *ACM Transactions on Mathematical Software* 31.3, pp. 351–362.
- HICKEN, J. E. and ZINGG, D. W. (2011). "The role of dual consistency in functional accuracy: Error estimation and superconvergence". *20th AIAA Computational Fluid Dynamics Conference 2011* June, pp. 1–18.
- HILL, G. W. (1886). "On the part of the motion of the lunar perigee which is a function of the mean motions of the sun and moon". *Acta Mathematica* 8.1, pp. 1–36.
- HODGES, D. H. and PIERCE, G. A. (2011). *Introduction to structural dynamics and aeroelasticity, second edition*. Vol. 9780521195, pp. 1–247.
- HOURIGAN, K., THOMPSON, M., and TAN, B. (2001). "Self-sustained oscillations in flows around long blunt plates". *Journal of Fluids and Structures* 15.3-4, pp. 387–398.
- HUANG, H. and EKICI, K. (2014). "Stabilization of High-Dimensional Harmonic Balance Solvers Using Time Spectral Viscosity". *AIAA Journal* 52.8, pp. 1784–1794.
- HUERRE, P. and MONKEWITZ, P. A. (1990). "Local and global instabilities in spatially developing flows". *Annual review of fluid mechanics* 22.1, pp. 473–537.
- IORIO, M. C., GONZÁLEZ, L. M., and FERRER, E. (2014). "Direct and adjoint global stability analysis of turbulent transonic flows over a NACA0012 profile". *International Journal for Numerical Methods in Fluids* 76.3, pp. 147–168.
- JACKSON, C. P. (1987). "A finite-element study of the onset of vortex shedding in flow past variously shaped bodies". *Journal of Fluid Mechanics* 182, pp. 23–45.
- JALLAS, D., MARQUET, O., and FABRE, D. (2017). "Linear and nonlinear perturbation analysis of the symmetry breaking in time-periodic propulsive wakes". *Physical Review E* 95.6, pp. 1–15.
- JIAN, Z. and JINWU, X. (2009). "Nonlinear Aeroelastic Response of High-aspect-ratio Flexible Wings". *Chinese Journal of Aeronautics* 22.4, pp. 355–363.
- JOLIVET, P., DOLEAN, V., HECHT, F., NATAF, F., PRUD'HOMME, C., and SPILLANE, N. (2012). "High performance domain decomposition methods on massively parallel architectures with freefem++". *Journal of Numerical Mathematics* 20.3-4, pp. 287–302.

- JOLIVET, P. and TOURNIER, P.-H. (2016). "Block iterative methods and recycling for improved scalability of linear solvers". In: *Proceedings of the International Conference for High Performance Computing, Networking, Storage and Analysis*. IEEE Press, p. 17.
- KARNIADAKIS, G. E., ISRAELI, M., and ORSZAG, S. A. (1991). "High-order splitting methods for the incompressible Navier-Stokes equations". *Journal of Computational Physics* 97.2, pp. 414–443.
- KARYPIS, G. and KUMAR, V. (1998). "A Parallel Algorithm for Multilevel Graph Partitioning and Sparse Matrix Ordering". *Journal of Parallel and Distributed Computing* 48.1, pp. 71–95.
- KAY, D., LOGHIN, D., and WATHEN, A. (2002). "A Preconditioner for the Steady-State Navier–Stokes Equations". *SIAM Journal on Scientific Computing* 24.1, pp. 237–256.
- KHOLODAR, D. B., DOWELL, E. H., THOMAS, J. P., and HALL, K. C. (2004). "Limit Cycle Oscillation of a Typical Airfoil in Transonic Flow". *Journal of Aircraft* 41.5, pp. 1067–1072.
- KIM, J. and MOIN, P. (1985). "Application of a fractional-step method to incompressible Navier-Stokes equations". *Journal of Computational Physics* 59.2, pp. 308–323.
- KRACK, M. and GROSS, J. (2019). *Harmonic Balance for Nonlinear Vibration Problems*, p. 159.
- KRYLOV, N. M. and BOGOLIUBOV, N. N. (1949). *Introduction to non-linear mechanics*. Princeton University Press.
- LABRYER, A. and ATTAR, P. J. (2009). "High dimensional harmonic balance dealiasing techniques for a Duffing oscillator". *Journal of Sound and Vibration* 324.3-5, pp. 1016–1038.
- LAURENSEN, R. M. and TRN, R. M. (1980). "Flutter analysis of missile control surfaces containing structural nonlinearities". *AIAA Journal* 18.10, pp. 1245–1251.
- LAZARUS, A. and THOMAS, O. (2010). "A harmonic-based method for computing the stability of periodic solutions of dynamical systems". *Comptes Rendus Mécanique* 338.9, pp. 510–517.
- LE MAÎTRE, O. P., SCANLAN, R. H., and KNIO, O. M. (2003). "Estimation of the flutter derivatives of an NACA airfoil by means of Navier-Stokes simulation". *Journal of Fluids and Structures* 17.1, pp. 1–28.
- LEE, B. H. K., PRICE, S. J., and WONG, Y. S. (1999a). "Nonlinear aeroelastic analysis of airfoils: bifurcation and chaos". *Progress in Aerospace Sciences* 35, pp. 205–334.
- LEE, B. H. K. and JIANG, L. (1999b). "Flutter of an Airfoil With a Cubic Restoring Force". *Journal of Fluids and Structures* 13, pp. 75–101.
- LEE, B. H., LIU, L., and CHUNG, K. W. (2005). "Airfoil motion in subsonic flow with strong cubic nonlinear restoring forces". *Journal of Sound and Vibration* 281.3-5, pp. 699–717.
- LEFFELL, J. I., MURMAN, S. M., and PULLIAM, T. H. (2014). "Time-Spectral Rotorcraft Simulations on Overset Grids". June, pp. 16–20.

- LEFFELL, J. I., SITARAMAN, J., LAKSHMINARAYAN, V. K., and WISSINK, A. M. (2016). "Towards Efficient Parallel-in-Time Simulation of Periodic Flows". January, pp. 4–8.
- LI, H. and EKICI, K. (2018). "A novel approach for flutter prediction of pitch–plunge airfoils using an efficient one-shot method". *Journal of Fluids and Structures* 82, pp. 651–671.
- LI, X. and FLEETER, S. (1997). "Dynamic stall generated airfoil oscillations including chaotic responses". *Collection of Technical Papers - AIAA/ASME/ASCE/AHS/ASC Structures, Structural Dynamics and Materials Conference 1*, pp. 11–23.
- LI, X. S. (2005). "An Overview of SuperLU: Algorithms, Implementation, and User Interface". *ACM Transactions on Mathematical Software* 31.3, pp. 302–325.
- LIESEN, J. and TICHÝ, P. (2004). "Convergence analysis of Krylov subspace methods". *GAMM-Mitteilungen* 27.2, pp. 153–173.
- LINKE, A., REBHOLZ, L. G., and WILSON, N. E. (2011). "On the convergence rate of grad-div stabilized Taylor-Hood to Scott-Vogelius solutions for incompressible flow problems". *Journal of Mathematical Analysis and Applications* 381.2, pp. 612–626.
- LISSAMAN, P. B. S. (1983). "Low-reynolds-number airfoils", pp. 223–239.
- LIU, G., LV, Z. R., LIU, J. K., and CHEN, Y. M. (2018). "Quasi-periodic aeroelastic response analysis of an airfoil with external store by incremental harmonic balance method". *International Journal of Non-Linear Mechanics* 100.December 2017, pp. 10–19.
- LIU, L., THOMAS, J. P., DOWELL, E. H., ATTAR, P., and HALL, K. C. (2006). "A comparison of classical and high dimensional harmonic balance approaches for a Duffing oscillator". *Journal of Computational Physics* 215.1, pp. 298–320.
- LIU, L. and DOWELL, E. H. (2004). "The Secondary Bifurcation of an Aeroelastic Airfoil Motion: Effect of High Harmonics". *Nonlinear Dynamics* 37, pp. 31–49.
- LIU, L. and DOWELL, E. H. (2005). "Harmonic Balance Approach for an Airfoil with a Freeplay Control Surface". *AIAA Journal* 43.4, pp. 802–815.
- LIVNE, E. (2018). "Aircraft active flutter suppression: State of the art and technology maturation needs". *Journal of Aircraft* 55.1, pp. 410–450.
- LOISEAU, J. C., BUCCI, M. A., CHERUBINI, S., and ROBINET, J. C. (2019). "Time-stepping and krylov methods for large-scale instability problems". *Computational Methods in Applied Sciences* 50, pp. 33–73. arXiv: [1804.03859](https://arxiv.org/abs/1804.03859).
- LOISEAU, J. C., ROBINET, J. C., CHERUBINI, S., and LERICHE, E. (2014). "Investigation of the roughness-induced transition: Global stability analyses and direct numerical simulations". *Journal of Fluid Mechanics* 760, pp. 175–211.
- LUCIA, D. J., BERAN, P. S., and SILVA, W. A. (2004). "Reduced-order modeling: New approaches for computational physics". *Progress in Aerospace Sciences* 40.1-2, pp. 51–117.
- MACK, C. J. and SCHMID, P. J. (2010). "A preconditioned Krylov technique for global hydrodynamic stability analysis of large-scale compressible flows". *Journal of Computational Physics* 229.3, pp. 541–560.

- MADER, C. A. and MARTINS, J. R. (2012). "Derivatives for time-spectral computational fluid dynamics using an automatic differentiation adjoint". *AIAA Journal* 50.12, pp. 2809–2819.
- MAHBUB ALAM, M., ZHOU, Y., YANG, H. X., GUO, H., and MI, J. (2009). "The ultra-low Reynolds number airfoil wake". *Experiments in Fluids* 48.1, pp. 81–103.
- MALHER, A. (2016). "Amortisseurs passifs non linéaires pour le contrôle de l'instabilité de flottement". PhD thesis. Université Paris-Saclay.
- MAMUN, C. K. and TUCKERMAN, L. S. (1994). "Asymmetry and Hopf bifurcation in spherical Couette flow". *Physics of Fluids* 7.1, pp. 80–91.
- MARQUET, O. and LARSSON, M. (2015a). "Global wake instabilities of low aspect-ratio flat-plates". *European Journal of Mechanics - B/Fluids* 49, pp. 400–412.
- MARQUET, O. and LESSHAFFT, L. (2015b). "Identifying the active flow regions that drive linear and nonlinear instabilities". arXiv: [1508.07620](https://arxiv.org/abs/1508.07620).
- MCCROSKEY, W. J. (1982). "Unsteady airfoils". *Annual Review of Fluid Mechanics* 14, pp. 285–311.
- MCMULLEN, M., JAMESON, A., and ALONSO, J. (2006). "Demonstration of Nonlinear Frequency Domain Methods". *AIAA Journal* 44.7, pp. 1428–1435.
- MEERBERGEN, K. and ROOSE, D. (1996). "Matrix transformations for computing rightmost eigenvalues of large sparse non-symmetric eigenvalue problems". *IMA Journal of Numerical Analysis* 16.3, pp. 297–346.
- MELIGA, P. and CHOMAZ, J. M. (2011). *An asymptotic expansion for the vortex-induced vibrations of a circular cylinder*. Vol. 671, pp. 137–167.
- MELIGA, P., CHOMAZ, J. M., and SIPP, D. (2009). "Global mode interaction and pattern selection in the wake of a disk: A weakly nonlinear expansion". *Journal of Fluid Mechanics* 633, pp. 159–189.
- MELIGA, P., GALLAIRE, F., and CHOMAZ, J.-M. (2012). "A weakly nonlinear mechanism for mode selection in swirling jets". *Journal of Fluid Mechanics* 699, pp. 216–262.
- MENON, K. and MITTAL, R. (2019). "Flow physics and dynamics of flow-induced pitch oscillations of an airfoil". *Journal of Fluid Mechanics* 877, pp. 582–613.
- MERCADER, I., BATISTE, O., and ALONSO, A. (2006). "Continuation of travelling-wave solutions of the Navier – Stokes equations". *International Journal for Numerical Methods in Fluids* 52.7, pp. 707–721.
- MICHELIN, S. and LLEWELLYN SMITH, S. G. (2009). "Resonance and propulsion performance of a heaving flexible wing". *Physics of Fluids* 21.7. arXiv: [0906.2804](https://arxiv.org/abs/0906.2804).
- MITTAL, S. and SINGH, S. (2005). "Vortex-induced vibrations at subcritical Re". *Journal of Fluid Mechanics* 534, pp. 185–194.
- MOUGIN, G. and MAGNAUDET, J. (2002). "The generalized Kirchhoff equations and their application to the interaction between a rigid body and an arbitrary time-dependent viscous flow". *International Journal of Multiphase Flow* 28.11, pp. 1837–1851.

- MUELLER, T. J. and DELAURIER, J. D. (2003). "Aerodynamics of small vehicles". *Annual Review of Fluid Mechanics* 35, pp. 89–111.
- MUNDIS, N. L. and MAVRIPLIS, D. J. (2014). "An Efficient Flexible GMRES Solver for the Fully-coupled Time-spectral Aeroelastic System". January, pp. 1–25.
- MUNDIS, N. L. and MAVRIPLIS, D. J. (2015). "Wave-number Independent Preconditioning for GMRES Time-spectral Solvers". *53rd AIAA Aerospace Sciences Meeting* January, pp. 1–21.
- MUNDIS, N. L. and MAVRIPLIS, D. J. (2017). "Toward an optimal solver for time-spectral fluid-dynamic and aeroelastic solutions on unstructured meshes". *Journal of Computational Physics* 345, April, pp. 132–161.
- NAVROSE and MITTAL, S. (2016). "Lock-in in vortex-induced vibration". *J. Fluid Mech* 794, pp. 565–594.
- NAYFEH, A. H., GHOMMEM, M., and HAJJ, M. R. (2012). "Normal form representation of the aeroelastic response of the Goland wing". *Nonlinear Dynamics* 67, pp. 1847–1861.
- NAYFEH, A. H. (2011). *The Method of Normal Forms*. Second, Up. Wiley-VCH.
- NAYFEH, A. H. and BALACHANDRAN, B. (1995). *Applied nonlinear dynamics : analytical, computational, and experimental methods*. Wiley, p. 685.
- NEGI, P. S., VINUESA, R., HANIFI, A., SCHLATTER, P., and HENNINGSON, D. S. (2018). "Unsteady aerodynamic effects in small-amplitude pitch oscillations of an airfoil". *International Journal of Heat and Fluid Flow* 71, November 2017, pp. 378–391.
- NEGI, P. S. (2019). "Stability and transition in pitching wings". PhD thesis. Royal Institute of Technology in Stockholm.
- NEGI, P. S., HANIFI, A., and HENNINGSON, D. S. (2019). "Global stability of rigid-body-motion fluid-structure-interaction problems". 1962, pp. 1–38. arXiv: [1910.09605](https://arxiv.org/abs/1910.09605).
- OLSHANSKII, M., LUBE, G., HEISTER, T., and LÖWE, J. (2009). "Grad-div stabilization and subgrid pressure models for the incompressible Navier–Stokes equations". *Computer Methods in Applied Mechanics and Engineering* 198.49, pp. 3975–3988.
- OLSHANSKII, M. A. and REUSKEN, A. (2004). "Grad-Div stabilization for Stokes equations". *Mathematics of Computation* 73.248, pp. 1699–1718.
- OLSHANSKII, M. A. (2002). "A low order Galerkin finite element method for the Navier–Stokes equations of steady incompressible flow: a stabilization issue and iterative methods". *Computer Methods in Applied Mechanics and Engineering* 191.47, pp. 5515–5536.
- OLSHANSKII, M. A. and BENZI, M. (2008). "An Augmented Lagrangian Approach to Linearized Problems in Hydrodynamic Stability". *SIAM Journal on Scientific Computing* 30.3, pp. 1459–1473.
- ONOUE, K., SONG, A., STROM, B., and BREUER, K. S. (2015). "Large amplitude flow-induced oscillations and energy harvesting using a cyber-physical pitching plate". *Journal of Fluids and Structures* 55, pp. 262–275.

- PADMANABHAN, M. A., DOWELL, E. H., and PASILIAO, C. L. (2018). "Computational Study of Aeroelastic Limit Cycles due to Localized Structural Nonlinearities". *Journal of Aircraft* 55.4, pp. 1531–1541.
- PAÏDOUSSIS, M. P., PRICE, S. J., and DE LANGRE, E. (2011). *Fluid Structure Interactions : Cross-Flow-Induced Instabilities*. Cambridge. New York.
- PALADINI, E., MARQUET, O., SIPP, D., ROBINET, J. C., and DANDOIS, J. (2019). "Various approaches to determine active regions in an unstable global mode: Application to transonic buffet". *Journal of Fluid Mechanics* 881.M, pp. 617–647.
- PATANKAR, S. V. and SPALDING, D. B. (1983). "A calculation procedure for heat, mass and momentum transfer in three-dimensional parabolic flows". *Numerical Prediction of Flow, Heat Transfer, Turbulence and Combustion*, pp. 54–73.
- PATIL, M. J. and HODGES, D. H. (2004). "On the importance of aerodynamic and structural geometrical nonlinearities in aeroelastic behavior of high-aspect-ratio wings". *Journal of Fluids and Structures* 19.7, pp. 905–915.
- PATIL, M. and HODGES, D. (2001). "Limit-cycle oscillations in high-aspect-ratio wings". *Journal of Fluids and Structures* 15, pp. 107–132.
- PELETAN, L., BAGUET, S., TORKHANI, M., and JACQUET-RICHARDET, G. (2013). "A comparison of stability computational methods for periodic solution of nonlinear problems with application to rotordynamics". *Nonlinear Dynamics* 72.3, pp. 671–682.
- PENG, Z. and ZHU, Q. (2009). "Energy harvesting through flow-induced oscillations of a foil". *Physics of Fluids* 21.12, pp. 1–9.
- PERSSON, P.-O., BONET, J., and PERAIRE, J. (2009). "Discontinuous Galerkin solution of the Navier–Stokes equations on deformable domains". *Computer Methods in Applied Mechanics and Engineering* 198.17-20, pp. 1585–1595.
- PFISTER, J. L. and MARQUET, O. (2020). "Fluid-structure stability analyses and nonlinear dynamics of flexible splitter plates interacting with a circular cylinder flow". *Journal of Fluid Mechanics*, pp. 1–38.
- PFISTER, J.-L. (2019). "Instabilities and optimization of elastic structures interacting with laminar flows". PhD thesis. Université Paris-Saclay.
- PFISTER, J.-L., MARQUET, O., and CARINI, M. (2019). "Linear stability analysis of strongly coupled fluid–structure problems with the Arbitrary-Lagrangian–Eulerian method". *Computer Methods in Applied Mechanics and Engineering* 355, pp. 663–689.
- PIGOLOTTI, L., MANNINI, C., BARTOLI, G., and THIELE, K. (2017). "Critical and post-critical behaviour of two-degree-of-freedom flutter-based generators". *Journal of Sound and Vibration* 404, pp. 116–140.
- POIREL, D., HARRIS, Y., and BENAÏSSA, A. (2008). "Self-sustained aeroelastic oscillations of a NACA0012 airfoil at low-to-moderate Reynolds numbers". *Journal of Fluids and Structures* 24.5, pp. 700–719.
- POIREL, D. and YUAN, W. (2010). "Aerodynamics of laminar separation flutter at a transitional Reynolds number". *Journal of Fluids and Structures* 26.7-8, pp. 1174–1194.

- POIREL, D. and MENDES, F. (2012). "Experimental Investigation of Small Amplitude Self-Sustained Pitch-Heave Oscillations of a NACA0012 Airfoil at Transitional Reynolds Numbers". In: *50th AIAA Aerospace Sciences Meeting including the New Horizons Forum and Aerospace Exposition*. Nashville, pp. 2012–40.
- PRASAD, R., KIM, H., CHOI, S., and YI, S. (2018). "High fidelity aeroelastic analysis based flutter prediction". *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2018* 210049, pp. 1–9.
- PRICE, S. J., ALIGHANBARI, H., and LEE, B. H. (1995). *The aeroelastic response of a two-dimensional airfoil with bilinear and cubic structural nonlinearities*.
- RAMEZANIAN, D., REZA AHRABI, B., and MAVRIPLIS, D. (2017). "An Order NlogN Parallel Newton_Krylov Solver for Time Spectral Problems". June, pp. 1–20.
- RAZAK, N. A., ANDRIANNE, T., and DIMITRIADIS, G. (2011). "Flutter and Stall Flutter of a Rectangular Wing in a Wind Tunnel". *AIAA Journal* 49.10, pp. 2258–2271.
- RICHTER, T. (2015). "A monolithic geometric multigrid solver for fluid-structure interactions in ALE formulation". *International Journal for Numerical Methods in Engineering* 104, pp. 372–390.
- RIGAS, G., SIPP, D., and COLONIUS, T. (2020). "Non-linear input/output analysis: application to boundary layer transition". arXiv: [2001.09440](https://arxiv.org/abs/2001.09440).
- ROMAN, J. E., CAMPOS, C., ROMERO, E., and TOMÁS, A. (2018). "SLEPc Users Manual Scalable Library for Eigenvalue Problem Computations".
- ROSTAMI, M. W. and XUE, F. (2018). "Robust Linear Stability Analysis and a New Method for Computing the Action of the Matrix Exponential". *SIAM Journal on Scientific Computing* 40.5, A3344–A3370.
- SAAD, Y. (1980). "Variations on Arnoldi's method for computing eigenvalues of large unsymmetric matrices". *Linear Algebra and its Applications* 34, pp. 269–295.
- SAAD, Y. (1993). "A Flexible Inner-Outer Preconditioned GMRES Algorithm". *SIAM Journal on Scientific Computing* 14.2, pp. 461–469.
- SAAD, Y. and SCHULTZ, M. H. (1986). "GMRES: A Generalized Minimal Residual Algorithm for Solving Nonsymmetric Linear Systems". *SIAM Journal on Scientific and Statistical Computing* 7.3, pp. 856–869.
- SAAD, Y. (2003). *Iterative Methods for Sparse Linear Systems*. Second edi. Society for Industrial and Applied Mathematics. arXiv: [0806.3802](https://arxiv.org/abs/0806.3802).
- SARTOR, F., METTOT, C., and SIPP, D. (2015). "Stability, Receptivity, and Sensitivity Analyses of Buffeting Transonic Flow over a Profile". *AIAA Journal* 53.7, pp. 1980–1993.
- SCHEWE, G., MAI, H., and DIETZ, G. (2003). "Nonlinear effects in transonic flutter with emphasis on manifestations of limit cycle oscillations". *Journal of Fluids and Structures* 18.1, pp. 3–22.
- SCHMID, P. J. and HENNINGSON, D. S. (2001). *Stability and Transition in Shear Flows*. New York: Springer.
- SEDAGHAT, A., COOPER, J. E., WRIGHT, J. R., and LEUNG, A. Y. (2000). "Limit cycle oscillation prediction for aeroelastic systems with continuous non-linearities". *41st Structures, Structural Dynamics, and Materials Conference and Exhibit* May 2014.

- SEGAL, A., UR REHMAN, M., and VUIK, C. (2010). "Preconditioners for incompressible Navier-Stokes solvers". *Numerical Mathematics* 3.3, pp. 245–275. arXiv: [f1d.1](#) [DOI: [10.1002](#)].
- SHAABANI-ARDALI, L., SIPP, D., and LESSHAFFT, L. (2019). "Vortex pairing in jets as a global Floquet instability: modal and transient dynamics". *Journal of Fluid Mechanics* 862, pp. 951–989.
- SHEN, S. F. (1959). "An Approximate Analysis of Nonlinear Flutter Problems". *Journal of the Aerospace Sciences* 26.x, pp. 25–31.
- SHROFF, G. M. and KELLER, H. B. (1993). "Stabilization of unstable procedures: the recursive projection method". *SIAM Journal on Numerical Analysis* 30.4, pp. 1099–1120.
- SHYY, W., AONO, H., CHIMAKURTHI, S. K., TRIZILA, P., KANG, C. K., CESNIK, C. E., and LIU, H. (2010). "Recent progress in flapping wing aerodynamics and aeroelasticity". *Progress in Aerospace Sciences* 46.7, pp. 284–327.
- SICOT, F., PUIGT, G., and MONTAGNAC, M. (2008). "Block-Jacobi Implicit Algorithms for the Time Spectral Method". *AIAA Journal* 46.12, pp. 3080–3089.
- ŠIDLOF, P., VLČEK, V., and ŠTĚPÁN, M. (2016). "Experimental investigation of flow-induced vibration of a pitch-plunge NACA 0015 airfoil under deep dynamic stall". *Journal of Fluids and Structures* 67, pp. 48–59.
- SILVESTER, D., ELMAN, H., KAY, D., and WATHEN, A. (2001). "Efficient preconditioning of the linearized Navier–Stokes equations for incompressible flow". *Journal of Computational and Applied Mathematics* 128.1-2, pp. 261–279.
- SINGH, S. P. and MITTAL, S. (2005). "Vortex-induced oscillations at low reynolds numbers: Hysteresis and vortex-shedding modes". *Journal of Fluids and Structures* 20.8, pp. 1085–1104.
- SIPP, D. and LEBEDEV, A. (2007). "Global stability of base and mean flows: A general approach and its applications to cylinder and open cavity flows". *Journal of Fluid Mechanics* 593, pp. 333–358.
- SIPP, D., MARQUET, O., MELIGA, P., and BARBAGALLO, A. (2010). "Dynamics and Control of Global Instabilities in Open-Flows: A Linearized Approach". *Applied Mechanics Reviews* 63.3.
- STANFORD, B. and BERAN, P. (2013). "Direct flutter and limit cycle computations of highly flexible wings for efficient analysis and optimization". *Journal of Fluids and Structures* 36, pp. 111–123.
- STEWART, G. W. (2002). "A Krylov-Schur Algorithm for Large Eigenproblems". *SIAM Journal on Matrix Analysis and Applications* 23.3, pp. 601–614.
- SU, X. and YUAN, X. (2010). "Implicit solution of time spectral method for periodic unsteady flows". *International Journal for Numerical Methods in Fluids* 63.7, pp. 860–876.
- SUN, Q. and BOYD, I. D. (2004). "Flat-plate aerodynamics at very low Reynolds number". *Journal of Fluid Mechanics* 502.March 2004, pp. 199–206.
- SUN, Q., BOYD, L. D., and CANDLER, G. V. (2001). "Numerical simulation of gas flow over micro-scale airfoils". *35th AIAA Thermophysics Conference* June.

- TANG, D. M. and DOWELL, E. H. (1992). "Flutter and Stall Responses of a Helicopter Blade with Structural Nonlinearity". *Journal of Aircraft* 29.5, pp. 953–960.
- TANG, D., DOWELL, E. H., and VIRGIN, L. N. (1998). "Limit cycle behavior of an airfoil with a control surface". *Journal of Fluids and Structures* 12.7, pp. 839–858.
- TANG, D. M. and DOWELL, E. H. (2004). "Effects of geometric structural nonlinearity of flutter and limit cycle oscillations of high-aspect-ratio wings". *Journal of Fluids and Structures*.
- TANG, D. and DOWELL, E. H. (2002). "Limit-Cycle Hysteresis Response for a High-Aspect-Ratio Wing Model". *Journal of Aircraft* 39.5, pp. 885–888.
- TANG, D., DOWELL, E. H., and HALL, K. C. (1999). "Limit Cycle Oscillations of a Cantilevered Wing in Low Subsonic Flow". *AIAA Journal* 37.3, pp. 364–372.
- TCHOUFAG, J., FABRE, D., and MAGNAUDET, J. (2014a). "Global linear stability analysis of the wake and path of buoyancy-driven disks and thin cylinders". *Journal of Fluid Mechanics* 740, pp. 278–311.
- TCHOUFAG, J., FABRE, D., and MAGNAUDET, J. (2015). "Weakly Nonlinear Model with Exact Coefficients for the Fluttering and Spiraling Motion of Buoyancy-Driven Bodies". *Physical Review Letters* 115.11, p. 114501.
- TCHOUFAG, J., MAGNAUDET, J., and FABRE, D. (2014b). "Linear instability of the path of a freely rising spheroidal bubble". *Journal of Fluid Mechanics* 751, pp. 1–12.
- THEODORSEN, T. (1935). *General Theory of Aerodynamic Instability and the Mechanism of Flutter*. Tech. rep. Langley: NACA Report 496.
- THEOFILIS, V. (2003). "Advances in global linear instability analysis of nonparallel and three-dimensional flows". *Progress in Aerospace Sciences* 39.4, pp. 249–315.
- THEOFILIS, V. (2011). "Global Linear Instability". *Annual Review of Fluid Mechanics* 43.1, pp. 319–352.
- THOMAS, J. P., DOWELL, E. H., and HALL, K. C. (2004). "Modeling viscous transonic limit-cycle oscillation behavior using a harmonic balance approach". *Journal of Aircraft* 41.6, pp. 1266–1274.
- THOMAS, J. P., DOWELL, E. H., and HALL, K. C. (2002). "Nonlinear inviscid aerodynamic effects on transonic divergence, flutter and limit cycle oscillations". *AIAA Journal* 40.4, pp. 638–646.
- THOMAS, J. P. and DOWELL, E. H. (2018). "A fixed point iteration approach for harmonic balance based aeroelastic computations". *AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference, 2018* 210049, pp. 1–10.
- THOMAS, J. P., HALL, K. C., and DOWELL, E. H. (2005). "Discrete adjoint approach for modeling unsteady aerodynamic design sensitivities". *AIAA Journal* 43.9, pp. 1931–1936.
- TREFETHEN, L. N. (1996). *Finite Difference and Spectral Methods for Ordinary and Partial Differential Equations*. unpublished text. arXiv: [arXiv:1011.1669v3](https://arxiv.org/abs/1011.1669v3).

- TUCKERMAN, L. S., BERTAGNOLIO, F., DAUBE, O., LE QUÉRÉ, P., and BARKLEY, D. (2000a). "Stokes preconditioning for the inverse Arnoldi method". In: *Continuation Methods for Fluid Dynamics*. Ed. by D. HENRY and A. BERGEON. Aussois, pp. 241–255.
- TUCKERMAN, L. S. (1989). "Steady-state solving via stokes preconditioning; Recursion relations for elliptic operators". In: *11th International Conference on Numerical Methods in Fluid Dynamics*. Ed. by D. L. DWOYER, M. Y. HUSSAINI, and R. G. VOIGT. Berlin, Heidelberg: Springer Berlin Heidelberg, pp. 573–577.
- TUCKERMAN, L. S. (2015). "Laplacian Preconditioning for the Inverse Arnoldi Method". *Communications in Computational Physics* 18.05, pp. 1336–1351.
- TUCKERMAN, L. S. and BARKLEY, D. (2000b). "Bifurcation Analysis for Timesteppers". In: *Numerical Methods for Bifurcation Problems and Large-Scale Dynamical Systems*. Springer, New York, NY, pp. 453–466.
- TUCKERMAN, L. S., LANGHAM, J., and WILLIS, A. (2019). "Order-of-magnitude speedup for steady states and traveling waves via stokes preconditioning in channelflow and openpipeflow". *Computational Methods in Applied Sciences* 50, pp. 3–31. arXiv: [1804.05087](https://arxiv.org/abs/1804.05087).
- VAN DER WEIDE, E., GOPINATH, A. K., and JAMESON, A. (2005). "Turbomachinery applications with the time spectral method". *35th AIAA Fluid Dynamics Conference and Exhibit*, pp. 1–12.
- VAN ROOIJ, A., NITZSCHE, J., and DWIGHT, R. P. (2017a). "Prediction of Aeroelastic Limit-Cycle Oscillations Based on Harmonic Forced-Motion Oscillations". *AIAA Journal* 55.10, pp. 3517–3529.
- VAN ROOIJ, A., NITZSCHE, J., and DWIGHT, R. P. (2017b). "Bifurcations of limit-cycle oscillations of a two degree-of-freedom airfoil caused by aerodynamic nonlinearities". In: *58th AIAA/ASCE/AHS/ASC Structures, Structural Dynamics, and Materials Conference*. Reston, Virginia: American Institute of Aeronautics and Astronautics.
- VIO, G. A. and COOPER, J. E. (2005). "Limit cycle oscillation prediction for aeroelastic systems with discrete bilinear stiffness". *International Journal of Applied Mechanics and Mechanics* 3, pp. 100–119.
- VIO, G. A., DIMITRIADIS, G., and COOPER, J. E. (2007). "Bifurcation analysis and limit cycle oscillation amplitude prediction methods applied to the aeroelastic galloping problem". *Journal of Fluids and Structures* 23.7, pp. 983–1011.
- VON KÁRMÁN, T. and SEARS, W. R. (1938). "JOURNAL OF THE Volume 5 Airfoil Theory for Non-Uniform Motion". *Journal of the Aeronautical Sciences* 5.10, pp. 379–390.
- WILLIAMSON, C. H. and GOVARDHAN, R. (2004). "Vortex-induced vibrations". *Annual Review of Fluid Mechanics* 36.1982, pp. 413–455.
- WOOD, R. J., AVADHANULA, S., STELTZ, E., SEEMAN, M., ENTWISTLE, J., BACHRACH, A., BARROWS, G., SANDERS, S., and FEAR, R. S. (2007). "An autonomous palm-sized gliding micro air vehicle". *IEEE Robotics and Automation Magazine* 14.2, pp. 82–91.

- WOODGATE, M. A. and BADCOCK, K. J. (2007). "Fast prediction of transonic aeroelastic stability and limit cycles". *AIAA Journal* 45.6, pp. 1370–1381.
- WOODGATE, M. A. and BADCOCK, K. J. (2009). "Implicit Harmonic Balance Solver for Transonic Flow with Forced Motions". *AIAA Journal* 47.4, pp. 893–901.
- WOOLSTON, D. S., RUNYAN, H. L., and ANDREWS, R. E. (1957). "An Investigation of Effects of Certain Types of Structural Nonlinearities on Wing and Control Surface Flutter". *Journal of the Aeronautical Sciences (Institute of the Aeronautical Sciences)* 24.1, pp. 57–63.
- WOOLSTON, D. S. and CASTILE, G. E. (1951). *Some effects of variations in several parameters including fluid density on the flutter speed of light uniform cantilever wings*. Tech. rep. NACA Technical Note 2558.
- WOOLSTON, D. S., RUNYAN, H. L., and BYRDSOON, T. A. (1955). *Some Effects of System Nonlinearities in the Problem of Aircraft Flutter*. Tech. rep. NACA.
- YANG, Z. C. and ZHAO, L. C. (1988). "Analysis of limit cycle flutter of an airfoil in incompressible flow". *Journal of Sound and Vibration* 123.1, pp. 1–13.
- YAO, W. and MARQUES, S. (2015). "Prediction of Transonic Limit-Cycle Oscillations Using an Aeroelastic Harmonic Balance Method". *AIAA Journal* 53.7, pp. 2040–2051.
- YAO, W. and MARQUES, S. (2018). "A harmonic balance method for nonlinear fluid structure interaction problems". *Computers and Structures* 201, pp. 26–36.
- YOUNG, Y. L., CHAE, E. J., and AKCABAY, D. T. (2012). "Hybrid algorithm for modeling of fluid-structure interaction in incompressible, viscous flows". *Acta Mechanica Sinica/Lixue Xuebao* 28.4, pp. 1030–1041.
- YUAN, W., POIREL, D., and WANG, B. (2013). "Simulations of pitch-heave limit-cycle oscillations at a transitional reynolds number". *AIAA Journal* 51.7, pp. 1716–1732.
- ZEBIB, A. (1987). *Stability of viscous flow past a circular cylinder*. Tech. rep., pp. 155–165.
- ZHANG, L. and CHEN, F. (2017). "Stability and bifurcation for limit cycle oscillations of an airfoil with external store". *Nonlinear Dynamics* 88.1, pp. 165–187.
- ZHU, Y., SU, Y., and BREUER, K. (2020). "Nonlinear flow-induced instability of an elastically mounted pitching wing". *J. Fluid Mech* 899, pp. 35–36.

Titre : Sur la bifurcation de flottement en écoulement laminaire: méthodes modales linéaires et nonlinéaires

Mots clés : interaction fluide-structure, flottement, analyse de stabilité linéaire, analyse faiblement nonlinéaire, méthodes d'équilibrage harmonique, méthodes de Krylov préconditionnées

Résumé : L'instabilité de flottement a été le sujet de nombreuses études depuis le milieu du vingtième siècle à cause de ses applications critiques en aéronautique. Elle est classiquement décrite comme un instabilité linéaire en écoulement potentiel, mais les effets visqueux et nonlinéaires du fluide peuvent avoir un impact crucial.

La première partie de cette thèse est consacrée au développement de méthodes théoriques et numériques pour l'analyse linéaire et nonlinéaire de la dynamique d'une "section typique aéroélastique" — une plaque montée sur des ressorts de flexion et torsion — plongée dans un écoulement laminaire bi-dimensionnel modélisé par les équations de Navier–Stokes incompressibles. D'abord, on développe une analyse faiblement nonlinéaire pour étudier le régime basse amplitude, puis, une approche d'équilibrage harmonique, connue comme la Méthode Spectrale en Temps (TSM), de façon à capturer des solutions de flottement plus fortement nonlinéaires. Le défi de la résolution numérique des équations TSM est relevé grâce au développement d'une approche parallèle en temps de type Newton–Krylov, combinée à

un préconditionneur spécialement développé, dit "bloc-circulant".

La seconde partie de la thèse est dédiée à l'étude physique de la bifurcation de flottement. On commence par revisiter le problème de stabilité linéaire, en mettant en lumière, en particulier, les effets de viscosité. On poursuit avec l'étude des effets nonlinéaires fluides: les structures légères et les hauts nombres de Reynolds favorisent des bifurcations sous-critiques. On achève cette partie en étudiant l'apparition de modulations de basse fréquence sur des solutions périodiques de flottement. On explique ce comportement par une instabilité linéaire (Floquet) de cycle limite.

La dernière partie de la thèse vise à initier l'extension des différentes méthodes évoquées précédemment pour le cas de configurations tridimensionnelles à grande échelle. En guise de premier pas vers cet objectif à long terme, on développe un outil *open-source* massivement parallèle capable de réaliser l'analyse de stabilité linéaire hydrodynamique (structure figée) d'écoulements tridimensionnels possédant plusieurs dizaines de millions de degrés de liberté.

Title : On the flutter bifurcation in laminar flows: linear and nonlinear modal methods

Keywords : fluid-structure interaction, flutter, linear stability analysis, weakly nonlinear analysis, harmonic balance methods, preconditioned Krylov subspace methods

Abstract : The flutter instability has been the focus of numerous works since the middle of the twentieth century, due to its critical application in aeronautics. Flutter is classically described as a linear instability using potential flow models, but viscous and nonlinear fluid effects may both crucially impact this aeroelastic phenomenon.

The first part of this thesis is devoted to the development of theoretical and numerical methods for analyzing the linear and nonlinear dynamics of a "typical aeroelastic section" — a heaving and pitching spring-mounted plate — immersed in a two-dimensional laminar flow modeled by the incompressible Navier–Stokes equations. First, we develop a semi-analytical weakly nonlinear analysis to efficiently study the small amplitude regime. Second, we develop a harmonic balance-type method, known as the Time Spectral Method (TSM), in order to tackle highly-nonlinear periodic flutter solutions. The challenging task of solving the TSM equations is tackled *via* a time-parallel Newton–Krylov approach in combination with a new, so-called block-circulant preconditioner.

The second part of the thesis focuses on the physical investigation of the flutter bifurcation. We start by revisiting the linear stability problem using a Navier–Stokes fluid model allowing to highlight, in particular, the effect of viscosity. We continue our route on the flutter bifurcation by investigating the effect of fluid nonlinearities: low solid-to-fluid mass ratios and increasing Reynolds numbers foster subcritical bifurcations. We conclude our study by investigating the appearance of low-frequency amplitude modulations on top of a previously established periodic flutter solution. We explain this behavior by a (Floquet) linear instability of periodic solutions.

The last part of the thesis aims at initiating the extension of the different methods previously evoked to large-scale three-dimensional configurations. As a first step towards this long-term goal, we develop an open-source massively parallel tool, able to perform hydrodynamic (the structure is fixed) linear stability analysis of three-dimensional flows possessing several tens of millions of degrees of freedom.