FLUID-STRUCTURE INTERACTIONS OF WALL-MOUNTED FLEXIBLE SLENDER STRUCTURES

A thesis submitted to the University of Manchester
for the degree of Doctor of Philosophy
in the Faculty of Science and Engineering

2018

By
Joseph M. O'Connor
School of Mechanical, Aerospace and Civil Engineering

Contents

Word Count: 42566

List of Tables

List of Figures

Nomenclature

Recurring abbreviations and symbols are summarised here. Note that due to the cross-disciplinary nature of this work, some symbols have multiple definitions. In this case, the intended definition will be made explicitly clear in the text.

Abbreviations

CFD	Computational fluid dynamics
CSS	Conventional serial staggered
DdQq	d-dimensional q-velocity lattice model
IBM	Immersed boundary method
FEM	Finite element method
FSI	Fluid-structure interaction
HPC	High-performance computing
LGCA	Lattice gas cellular automata
LBM	Lattice boltzmann method
MLUPS	Millions of lattice updates per second

PEL	Poroelastic
RMS	Root mean square
VIV	Vortex induced vibration
WSS	Wall shear stress

Dimensionless Quantities

α	Womersley number
Ca	Cauchy number
Fr	Froude number
$\mathcal{K}$	Dimensionless bending stiffness
$\mathcal{K}^*$	Modified dimensionless bending stiffness
$\mathcal{M}$	Mass ratio
$\mathcal{M}^*$	Modified mass ratio
Ma	Mach number
Re	Reynolds number
St	Strouhal number
$\mathcal{T}$	Dimensionless tensile stiffness
$\mathcal{U}$	Reduced velocity

$\mathcal{V}$	Vogel exponent

Greek Symbols

α, δ	Newmark integration parameters
α_f, α_s	Ratio of spectral amplitudes for fluid and structure
β	Angle of corotated frame with respect to global frame
Γ	Lagrangian domain
δ_x, δ_y	Tip deflection
δ'_x, δ'_y	Tip velocity
$\tilde{\delta}$	Discrete Dirac Delta function
ϵ	Scaling factor for Lagrangian force
θ	Angle, momentum thickness
ν	Kinematic viscosity
ξ	Elemental local coordinate
$\boldsymbol{\xi}$	Molecular velocity
Π	Stress tensor
ρ_f, ρ_s	Fluid and structural mass densities
$\Delta\rho$	Extra mass due to presence of structure

τ	Relaxation time-scale, wall shear stress
Ω	Collision operator, Eulerian domain
ω	Relaxation factor

Roman Symbols

A	Cross sectional area
$\mathbf{B}$	Strain-displacement matrix
$\mathbf{C}$	Stress-strain matrix
C_D, C_L	Drag and lift coefficients
$\mathbf{c}_i$	Vector of discrete lattice velocities
c_s	Lattice speed of sound
D_s, D_{ssss}	Differential operators (first and fourth-order)
d	Dimensions of system (e.g. 2D/3D)
E	Young's modulus
$\mathbf{F}$	Lagrangian force density, vector of internal forces
F_i	Discrete lattice force
$\Delta\mathbf{F}$	Vector of incremental internal forces
f	Probability distribution function, frequency

f^{eq}	Equilibrium probability distribution function
f^{neq}	Non-equilibrium probability distribution function
$\mathbf{f}$	Eulerian force density
$\mathbf{g}$	Vector of gravitational acceleration
H	Domain height
h	Structure thickness
I	Second moment of area
$\mathcal{I}$	Interpolation operator
$\mathbf{K}$	Stiffness matrix
K_B	Bending rigidity (EI)
$k, k+1$	Iteration counter
L	Characteristic length (usually structure length)
$\mathbf{M}$	Mass matrix
$\mathbf{N}$	Vector of elemental shape functions
$n, n+1$	Time step counter
$\mathbf{n}$	Normal vector
p	Pressure

q	Number of discrete lattice velocities
q, r, s	Lagrangian coordinates
$\mathbf{R}$	Vector of externally applied loads
$\mathbf{r}$	Residual vector
$\mathcal{S}$	Spreading operator
s	Parametric coordinate along structure
Δs	Lagrangian marker spacing
T	Tension, time period
t	Time
Δt	Time step
$\mathbf{U}, \dot{\mathbf{U}}, \ddot{\mathbf{U}}$	Vector of nodal displacements, velocities, and accelerations
$\Delta\mathbf{U}$	Vector of incremental nodal displacements
u, v	Axial and transverse displacements
$\mathbf{u}$	Vector of fluid velocity
u_x, u_y	Components of fluid velocity vector
$\mathbf{u}^*$	Vector of predicted fluid velocity
u_∞	Freestream velocity

V	Volume
W	Domain length
w_i	Lattice-specific weighting factors
X, Y, Z	Boundary coordinates in Eulerian frame
$\mathbf{X}$	Vector of boundary coordinates in Eulerian frame
$\mathbf{X}^*$	Vector of extrapolated boundary coordinates in Eulerian frame
x, y, z	Eulerian coordinates
$\mathbf{x}$	Vector of Eulerian coordinates
Δx	Lattice spacing

Abstract

The fluid-structure interactions of wall-mounted slender structures, such as cilia, filaments, flaps, and flags, play an important role in a broad range of physical processes: from the coherent waving motion of vegetation, to the passive flow control capability of hair-like surface coatings. While these systems are ubiquitous, their coupled nonlinear response exhibits a wide variety of behaviours that is yet to be fully understood, especially when multiple structures are considered.

The purpose of this work is to investigate, via numerical simulation, the fluid-structure interactions of arrays of slender structures over a range of input conditions. A direct modelling approach, whereby the individual structures and their dynamics are fully resolved, is realised via a lattice Boltzmann-immersed boundary model, which is coupled to two different structural solvers: an Euler-Bernoulli beam model, and a finite element model. Results are presented for three selected test cases – which build in scale from a single flap in a periodic array, to a small finite array of flaps, and finally to a large finite array – and the key behaviour modes are characterised and quantified.

Results show a broad range of behaviours, which depend on the flow conditions and structural properties. In particular, the emergence of coherent waving motions are shown to be closely related to the natural frequency of the array. Furthermore, this behaviour is associated with a lock-in between the natural frequency of the array and the predicted frequency of the fluid instabilities.

The original contributions of this work are: the development and application of a numerical tool for direct modelling of large arrays of slender structures; the characterisation of the behaviour of slender structures over a range of input conditions; and the exposition of key behaviour modes of slender structures and their relation to input conditions.

Declaration

No portion of the work referred to in the thesis has been submitted in support of an application for another degree or qualification of this or any other university or other institute of learning.

Copyright

Acknowledgements

I would first like to express my sincere gratitude to Alistair. Your guidance and support over the past few years cannot be understated. It is impossible to count the number of times I came into our Wednesday morning meeting completely lost and deflated, to then 30 minutes later having ten new ideas and a reignited enthusiasm. Aside from the technical support, the mini-workshops on careers and other extracurricular skills are one of the many other reasons I feel privileged to have had you as a supervisor.

Thanks also to Philip and Partha, who always made themselves available and ensured I was (semi) regular with organising our supervisory meetings.

To Ben, Josh, and Sam – AKA the CFDemons Crew. While I probably would have got more work done without the distractions, it certainly wouldn't have been as fun. My highlights of the week were definitely Football Thursdays, Changos Fridays, and the occasional trip to Garret/5th Ave. I'm sure you can guess what order.

To Fee, your support over the past few years, especially over recent weeks, has been very much needed and I'm looking forward to a nice long holiday together. Also, thanks for at least trying to be interested in my work, but please stop calling it 'lettuce bowman'.

Finally, to my family – Carol, Paddy, Dan and Mat – thanks for supporting me for the best part of a decade while I led the student life and refused to join the 'real world'. Maybe it's about time now.

Chapter 1

Introduction

1.1 Background

The coupled interactions between fluids and structures is a general topic encountered in a broad range of physical processes. In nature, the pumping and transport of blood through compliant blood vessels is crucial for supplying oxygen to cells and tissues. Furthermore, animal locomotion modes, such as swimming or flying, usually involve some form of morphing or shape-changing. In industry, the issue of fluid-structure interaction (FSI) is a key consideration in the design of buildings and bridges, where resonant interactions with wind loads can have catastrophic consequences. Although FSI is often a critical concern in terms of design safety, the dynamic interactions of fluids and structures can also be exploited for beneficial purposes, such as in the energy industry, where FSI-based power generation techniques (e.g. wind turbines, hydropower) are becoming increasingly popular.

One particular class of problems that is especially susceptible to FSI is the coupled behaviour of slender structures. These structures play important roles in a range of applications in both nature and industry, from cell propulsion [1] to flow control [2]. The distinguishing feature of slender structures is that their length is significantly larger than at least one of their other dimensions. This classification therefore includes hair-like structures, such as filaments and cilia, but also includes plate-like structures, where only the thickness is significantly smaller than the length (e.g. flaps and flags), as shown in Figure 1.1. Large aspect ratios (length-to-thickness ratio) are the defining feature of slender structures; however, there is no standard

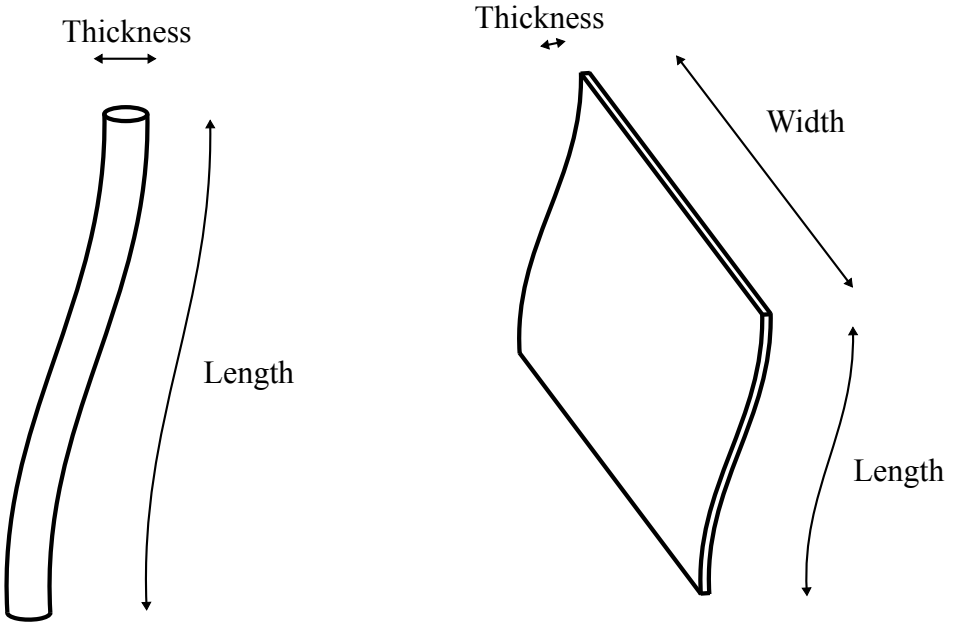

Figure 1.1: Representative slender structures. Hair-like structure (left) and plate-like structure (right).

value in the literature that can be used to classify a structure as slender. Therefore, for the purposes of the present work, a slender structure is defined as having an aspect ratio of 20 or greater. This classification is made as this is also the aspect ratio above which Euler-Bernoulli beam theory starts to become applicable [3]. Due to their large aspect ratios, slender structures are especially susceptible to bending deformations. Furthermore, the motion of slender structures in response to a fluid flow can vary greatly depending on the particular conditions, from static deflection, to dynamic waving or flapping, or even chaotic motion.

The nonlinear FSI of single slender structures has been a topic of interest for centuries [4]. Examples of such applications include flags flapping in the wind [5], or cell locomotion via propulsive flagella [6]. While such problems are certainly important, in many applications it is common for multiple slender structures to be packed together in close proximity to form dense arrays of structures. Such configurations are common in cilia and vegetation, for example. While arrays of multiple slender structures share some behavioural similarities to single structures, they also exhibit certain distinct differences. In particular, coherent interactions through the array, as well as with the fluid, are known to emerge.

One example of these coherent interactions is the waving of crops in a field. This phenomenon occurs due to Kelvin-Helmholtz instabilities, which arise from the

drag associated with the canopy [7]. These instabilities roll up to form canopy-scale vortices which propagate the canopy and induce a travelling wave of deflection on the individual plants [8]. Similar waving motions have been observed in flow control applications, where arrays of slender structures are attached to bluff bodies to improve aerodynamic performance [2]. In particular, a lock-in effect, where the natural frequencies of the fluid and structure synchronize, has been linked to special case conditions, such as minimal drag or lift fluctuations [9]. While these interactions and the mechanisms that initiate them are clearly important, they are yet to be fully understood. This is partly due to the large parameter space and complex behaviours that govern these processes, but also due to the difficulty in measuring and modelling these systems.

Experimental investigations have been, and still are, a popular choice for studying slender structures [10, 11]. However, due to the drawbacks associated with some experimental methods (e.g. limited number and type of measurements) [12] and the advances in compute power and modelling techniques, numerical modelling of slender structures is becoming an increasingly attractive approach. Of the various modelling techniques, direct modelling approaches, which resolve the individual structures and the flow around them, provide the most detail. However, while such approaches are popular for studying single structures on their own [13, 14], their use is limited in studies regarding arrays of slender structures. Instead, homogenised models, which use a bulk drag coefficient to account for the array, are common [2, 15]. This is significantly cheaper than directly modelling the individual structures, and for dense arrays where the global behaviour is of interest, this is a fair approximation. However, due to the spatial averaging implicit in these approaches, these techniques can have difficulties incorporating heterogeneity and do not provide details regarding the local effects of individual structures. Furthermore, these models usually need to be calibrated to specific conditions; and in cases where experimental data is difficult to obtain, direct modelling can provide a reference for this calibration. Considering this, there is scope for developing an approach for simulating large arrays of slender structures where each individual structure is directly modelled.

1.2 Aims & Objectives

In light of the above discussion, the purpose of the present work is to numerically study the FSI of arrays of wall-mounted slender structures via a direct modelling approach, where the individual structures are fully resolved, and quantify their behaviour over a range of input conditions. To satisfy this, three primary research aims are identified:

- to develop a numerical model that can simulate large arrays of slender structures by fully resolving the individual structures and their dynamics
- to apply this model to multiple configurations involving arrays of slender structures and characterise their behaviour over a range of input conditions
- to highlight key behavioural mechanisms (e.g. waving or lock-in) and relate them to the input parameters (such as flow conditions and structural properties)

To address these aims, a lattice Boltzmann-immersed boundary model is developed. To incorporate the structural dynamics, this model is coupled to two different structural solvers: one based on the Euler-Bernoulli beam equation, and the other based on the Finite Element Method (FEM). Simulations are performed for three test cases and the key behaviour modes are characterised and quantified.

Figure 1.2 shows schematics of the three main test cases that form the core results chapters. These cases were selected as they build in scale, from a single flap in a periodic array, to a small finite array of flaps, and finally to large finite array of flaps. The array length is small in the first two cases; therefore, to enhance the appearance of dynamic instabilities, these cases are tested under oscillating flow conditions. In the final case, the array is large enough to allow these instabilities to develop naturally, and so this is only tested under steady flow conditions.

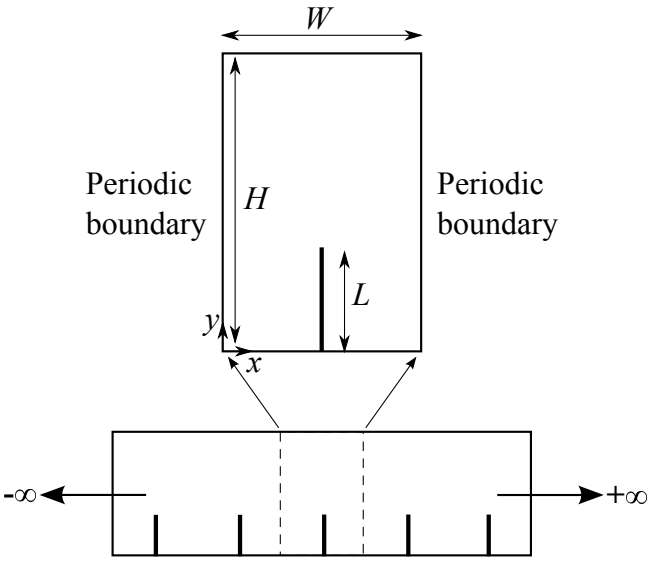

(a) Case I – Chapter 7

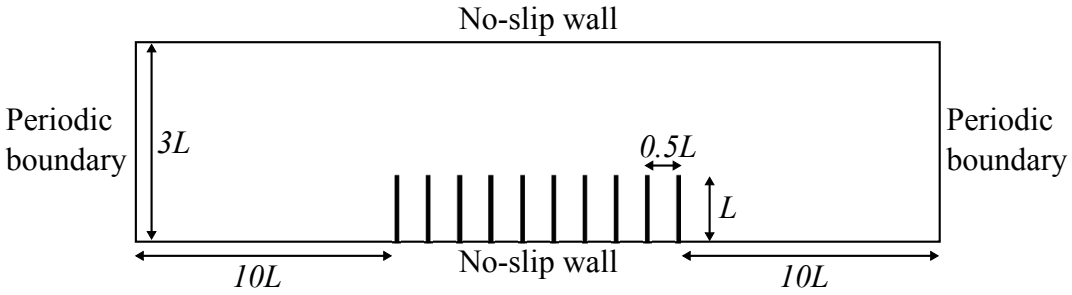

(b) Case II – Chapter 8

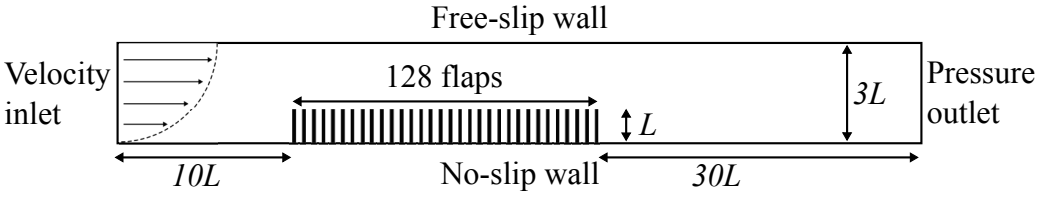

(c) Case III – Chapter 9

Figure 1.2: Schematics of all three main test cases.

1.3 Scope

The scope of this work is restricted to wall-mounted slender structures in a crossflow. Since much of the work on modelling slender structures has focussed on either direct modelling of axially aligned structures (e.g. flapping flag) or homogenised modelling of wall-mounted structures (e.g. vegetation), this focusses the present work into a relatively novel area. Furthermore, the flows of interest

are restricted to low-moderate Reynolds numbers, where $Re = \mathcal{O}(100)$. This regime was chosen so as to allow inertia-driven instabilities to develop, while at the same time avoiding instabilities arising from transition and turbulence effects. Finally, since the main behaviour of interest is the in-plane bending, this work is restricted to 2D. The main motivation for this is simplicity, as it allows a better understanding of the basic behaviours before moving to more complex 3D problems. Since the reduced blockage associated with hair-like structures leads to strong 3D effects, this limits the findings of the present work to plate-like structures, where the out-of-plane depth is sufficiently large so that edge effects become negligible. In general, this is not a drastic simplification and there are many examples of fundamental studies which focus solely on 2D behaviour [13, 16, 17].

1.4 Research Outputs & Thesis Format

The outcomes of this work led to a series of research outputs, in the form of journal publications. To accommodate this, this thesis is presented in the *Journal Format* style, as permitted by the University of Manchester, whereby each of the core results chapters is written in the format of a self-contained journal paper. Due to the inherent space limitations placed on journal publications, the traditional literature review and methodology chapters precede the results and provide details that would be otherwise unnecessary in a journal publication. While effort has been made to limit the repetition between chapters, the requirement of self-containment between the core results chapters means that this is unavoidable in some cases.

The full list of outputs generated from this work, and their significance, are given below (in chronological order):

1) **J. O'Connor, A. Revell, P. Mandal, and P. Day. Application of a lattice Boltzmann-immersed boundary method for fluid-filament dynamics and flow sensing. *J. Biomech.*, 49(11), 2016 [18]**

   This work examined the FSI of an idealised array of wall-mounted flaps under both steady and oscillating flow conditions. The configuration is an idealised one in the sense that only a single flap is directly modelled, whereas periodic boundary conditions are used to simulate the effect of the

array. Tests showed that at large Reynolds numbers the recirculation region between the flaps fully bridged the gap. This led to a reduced lag between the flow and structural response, as well as decreased flap deflection. This output forms Chapter 7 and is the first of the core results chapters.

2) **J. O'Connor, P. Day, P. Mandal, and A. Revell. Computational fluid dynamics in the microcirculation and microfluidics: what role can the lattice Boltzmann method play? *Integr. Biol.*, 8(5): 589–602, 2016** [19]

This review was a product of the initial literature review conducted in the early stages of the present work. The review discusses fluid modelling techniques and their application to the microcirculation and microfluidic devices, with a particular focus on the lattice Boltzmann method. After this review was completed, the main focus of the present research altered somewhat. As such, the review does not directly fall within the scope of this thesis, and is therefore omitted.

3) **J. Favier, C. Li, L. Kamps, A. Revell, J. O'Connor, and C. Brücker. The PELskin project - part I: fluid-structure interaction for a row of flexible flaps: a reference study in oscillating channel flow. *Meccanica*, 52(8), 2017** [20]

This paper forms part of a series of outputs generated as part of the EU-funded PELskin project, of which our research group was a member. The main aim of this project was to investigate the amelioration of aerodynamic performance of bluff bodies via a poroelastic (PEL) coating. Part I in the series introduces a simplified experimental case, which was used to isolate the key flow control mechanisms and consisted of a row of 10 flexible flaps embedded in an oscillating channel flow. This test case was also used to validate the present numerical model. Since my contribution to this paper consisted of only the model validation, this section of the paper is combined with Part II in the series to form Chapter 8, which is the second of the core results chapters.

4) **A. Revell, J. O'Connor, A. Sarkar, C. Li, J. Favier, L. Kamps, and C. Brücker. The PELskin project: part II - investigating the**

physical coupling between flexible filaments in an oscillating flow. *Meccanica*, 52(8), 2017 [21]

Part II in the PELskin series was led by our research group, and consisted of a detailed numerical study involving the baseline case described in Part I. A range of mass ratios were tested and its effect was characterised by examining the flap motion and near-field fluid velocity. Tests showed that at large mass ratios the coherence between the fluid and flap motion breaks down. Since my contribution to this paper is significant, it is combined with the validation study in Part I to form Chapter 8, which is the second of the core results chapters.

5) **A.R.G. Harwood, J. O'Connor, J.S. Muñoz, M.C. Santasmasas, and A.J. Revell. LUMA: A many-core, Fluid-Structure Interaction solver based on the Lattice-Boltzmann Method. *SoftwareX*, 7:88–94, 2018 [22]**

Over the last few years our research group has developed the *Lattice Boltzmann @ University of Manchester (LUMA)* code, which was recently made open-source and published in a software-based journal. As one of the main developers of LUMA, I was specifically responsible for the development of the FSI module and its validation. While this work is relevant to the present thesis, the actual publication, which is focussed on code design and structure, is not relevant and is therefore omitted.

6) **J. O'Connor and A. Revell. Dynamic Interactions of Multiple Flexible Flaps. *In Preparation*, 2018**

This final article is still in preparation and is yet to be submitted. It completes the present work by examining the FSI of a large array of wall-mounted flexible flaps under steady flow conditions. A range of mass ratios and bending rigidities are tested, and the global and local behavioural modes are mapped out. In particular, the waving motion and associated lock-in behaviour is characterised based on the input conditions. This paper forms Chapter 9 and is the final of the core results chapters.

To focus the narrative of this thesis to the FSI of slender structures, only Papers

1, 3, 4, and 6 are included in the results section. Furthermore, since Papers 3 and 4 formed part of a collaborative multipart series, these two have been combined into one chapter so that only work which I conducted is included. To maintain consistency throughout the thesis, each results chapter has been reformatted with respect to the original publication. Furthermore, all of the references from each separate publication have been combined into one bibliography at the end of the thesis. Finally, my contribution to each output is clearly stated in the preface to each results chapter.

1.5 Thesis Outline

Figure 1.3 provides an outline of the thesis. A description of each chapter is also given below:

Chapter 1 introduces the background, aims, and approach of the current research.

Chapter 2 provides a review of the FSI of slender structures, including applications, relevant physics, and modelling techniques.

Chapter 3 describes the fluid solver component of the present numerical model.

Chapter 4 describes the two structural models adopted in the present work.

Chapter 5 describes the coupling of the field solvers, including the transfer of quantities between the two solvers and the implicit coupling scheme.

Chapter 6 discusses developments that were made to the model as part of the present work and provides a validation of the individual field solvers and the fully coupled model.

Chapter 7 is the first of the three results chapters and presents results for an idealised case of a single flexible flap in a periodic array, under both steady and oscillating flow conditions.

Chapter 8 extends the previous chapter to incorporate a finite array of flexible flaps in an oscillating channel flow.

Chapter 9 is the final results section and presents results for a large array of flexible flaps under steady flow conditions.

Chapter 10 concludes by synthesising the main findings and provides suggestions for future work.

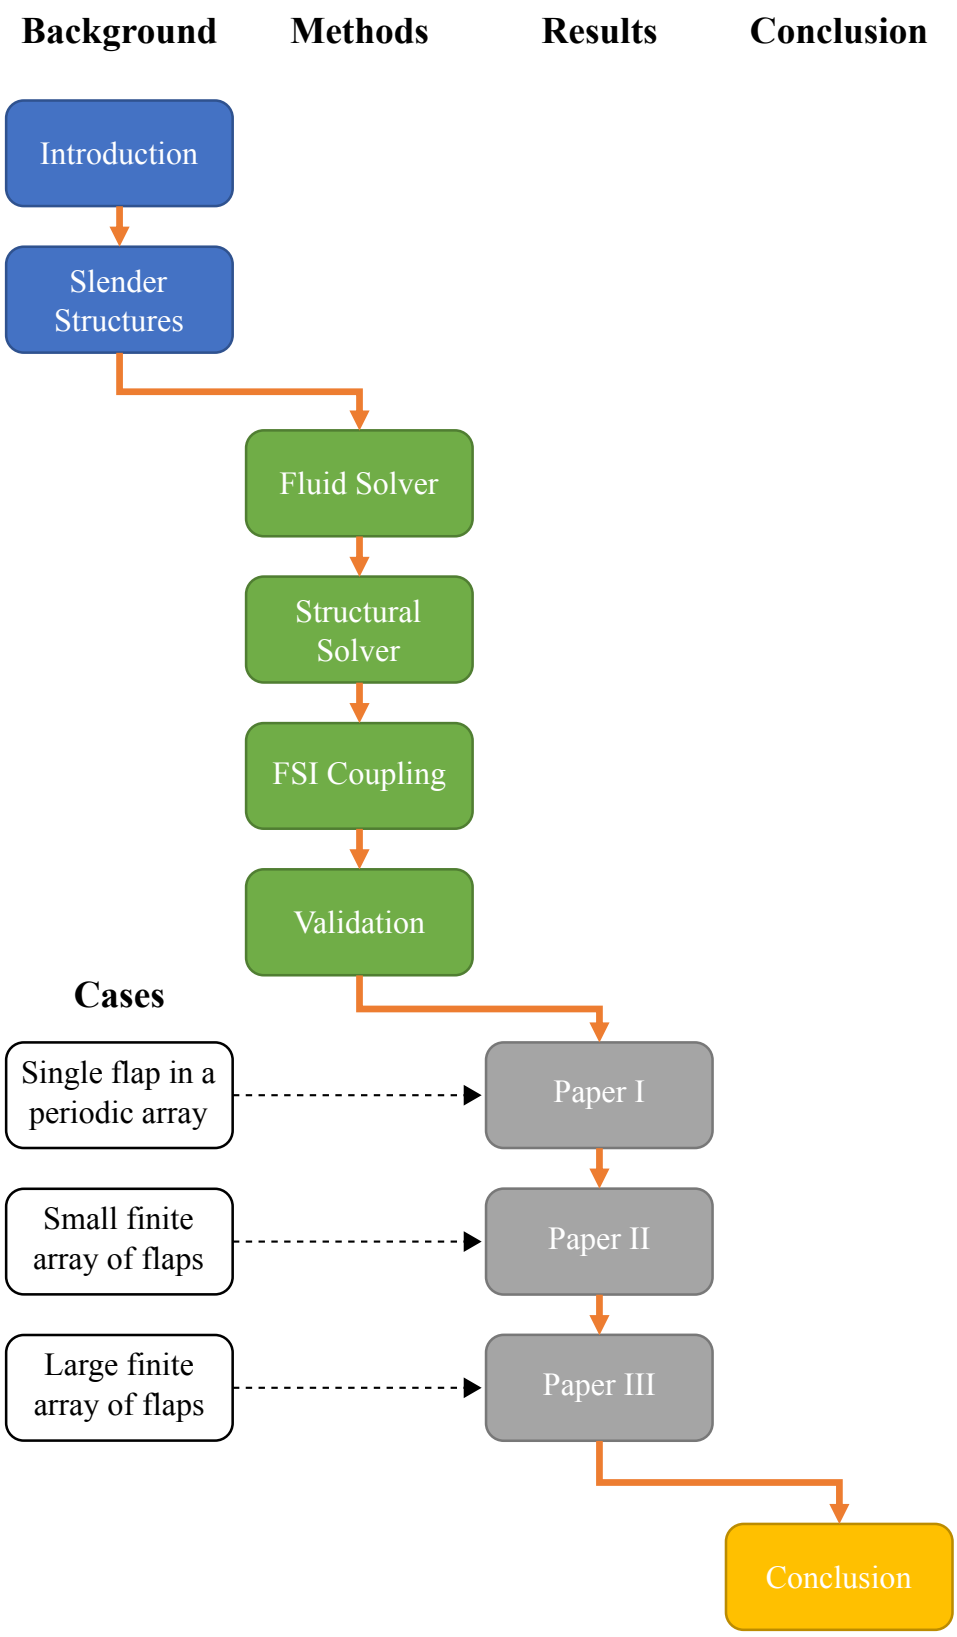

Figure 1.3: Thesis outline.

Chapter 2

Slender Structures

Following from the aims described in Chapter 1, this chapter introduces the topic of slender structures. To highlight their relevance, common examples of slender structures in nature and industry are discussed. By reviewing previous studies related to single and multiple structures, the relevant physics, typical behaviours, and common modelling approaches are described. Finally, the key points from this review, and how they relate to the present work, are summarised.

2.1 Introduction to Slender Structures

The distinguishing feature of slender structures is that their length is significantly larger than their thickness. Examples of such types of structures include hair, cilia, filaments, fibres, flaps, and flags, to name a few. Due to the large variety of roles they play, slender structures operate under a wide range of configurations (e.g. single/array, free/wall-mounted), scales (e.g. micro to macro-scales), and flow conditions (e.g. low/high Reynolds numbers, steady/oscillating flow). Furthermore, depending on their geometry, the resulting flow topology can vary significantly, from mostly 2D (in the case of plate-like structures) to highly 3D (in the case of hair-like structures).

Due to their large aspect ratios (length-to-thickness ratio), the main deformation mode in slender structures is usually bending, as opposed to axial extension. Furthermore, in certain classes of problems it is not uncommon for the translational

deformation to be of the order of the structural length scale. Such large deformations introduce complex nonlinear behaviour into the structural dynamics. When immersed in a fluid, the complex interactions between hydrodynamic, inertial, and elastic forces can lead to a broad range of physical behaviours, including static reconfiguration, dynamic waving and flapping, or even chaotic motion. While such problems are certainly interesting in their own right, they also govern many crucial processes in both nature and industry.

2.2 Applications of Slender Structures

2.2.1 Slender Structures in Nature

One of the most common examples of fluid-immersed slender structures is cilia. They are ubiquitous in nature and bear many important roles, from keeping our airways clean to trapping food particles in suspension feeders [23, 24]. Along with flagella (see Figure 2.1), cilia also play a role in cell propulsion and can generate substantial velocities, given their size [1]. Many common and lethal diseases (e.g. cystic fibrosis) are directly linked to ciliary dysfunction and so understanding how

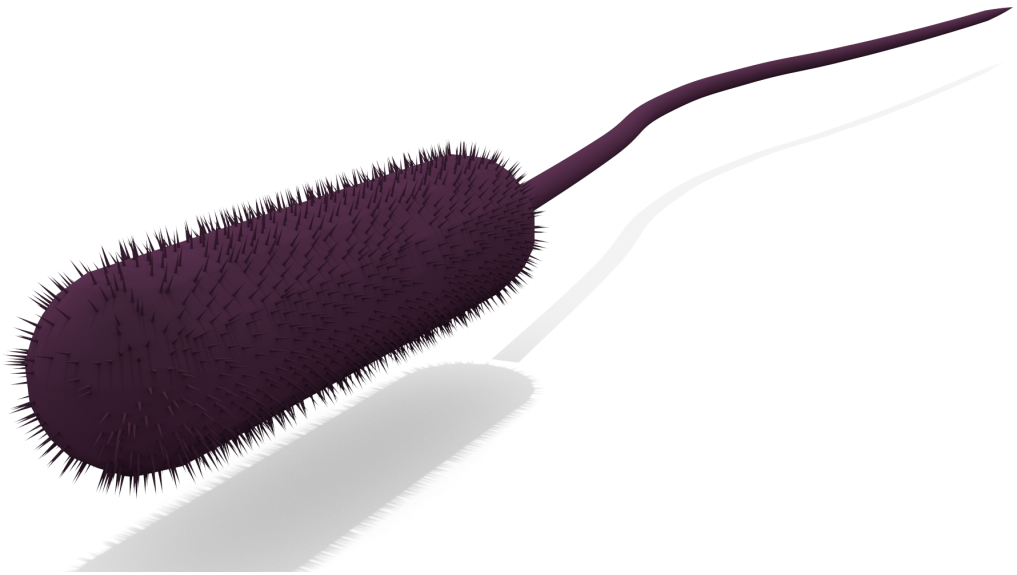

Figure 2.1: The bacterial flagellum is a long whip-like structure used for propelling bacteria (e.g. Salmonella or E. coli).

these systems behave is of particular interest.

Another example is the glycocalyx layer – a nanoscopic brush-like structure that lines the inner walls of our blood vessels. One of the primary roles of the glycocalyx layer is to act as a mechanotransduction sensor [25]. During periods of increased flow rate (e.g. during exercise) the glycocalyx will deform. Receptors in the vessel wall then sense this deformation and transmit signals to initiate an appropriate physiological response (e.g. vasodilation). This process is critical for maintaining normal vascular function and impairment of the glycocalyx can lead to serious medical conditions, such as hypertension, atherosclerosis, and thrombosis [26].

Many organisms use specialised hairs to sense the flow around them. These types of flow sensing systems are usually very efficient and extremely sensitive. Crickets, for example, can perceive air movements as low as 0.03 mm/s, which helps them escape lunging predators, such as spiders (who also have their own hair-based flow sensors) [27]. Bat wings are covered in directionally-sensitive microscopic hairs which can detect reverse flow over the wing [28]. This helps to prevent the onset of stall at low-speed flight. Furthermore, fish use arrays of neuromasts in their lateral line system to detect minute water movements [29]. In low-light conditions certain species can build 3D flow and object maps through the information provided via this sensory system [30].

Bird feathers have several important functions, one of which is to aid flight manoeuvrability and improve aerodynamic performance. During low-speed flight, such as on approach to landing, covert feathers pop up as the flow starts to separate near the trailing edge of the wing [31], as shown in Figure 2.2. The feathers prevent the separated region from moving further forward towards the leading edge [32]. This delays the onset of stall and is crucial for performing precise and controlled manoeuvres. Furthermore, the mechanism by which the feathers are activated is passive and occurs entirely as a response to the reverse flow. Such systems serve as excellent sources of inspiration for biomimicry.

The interactions between vegetation and fluid flows occur almost everywhere and have far reaching consequences. In terrestrial flows, strong winds can damage crops and uproot forests, resulting in reduced crop yield and habitat destruction [8]. Aquatic vegetation plays a key role in stabilising the river bed and regulating oxygen and nutrient levels [34]. Furthermore, coastal vegetation helps to protect

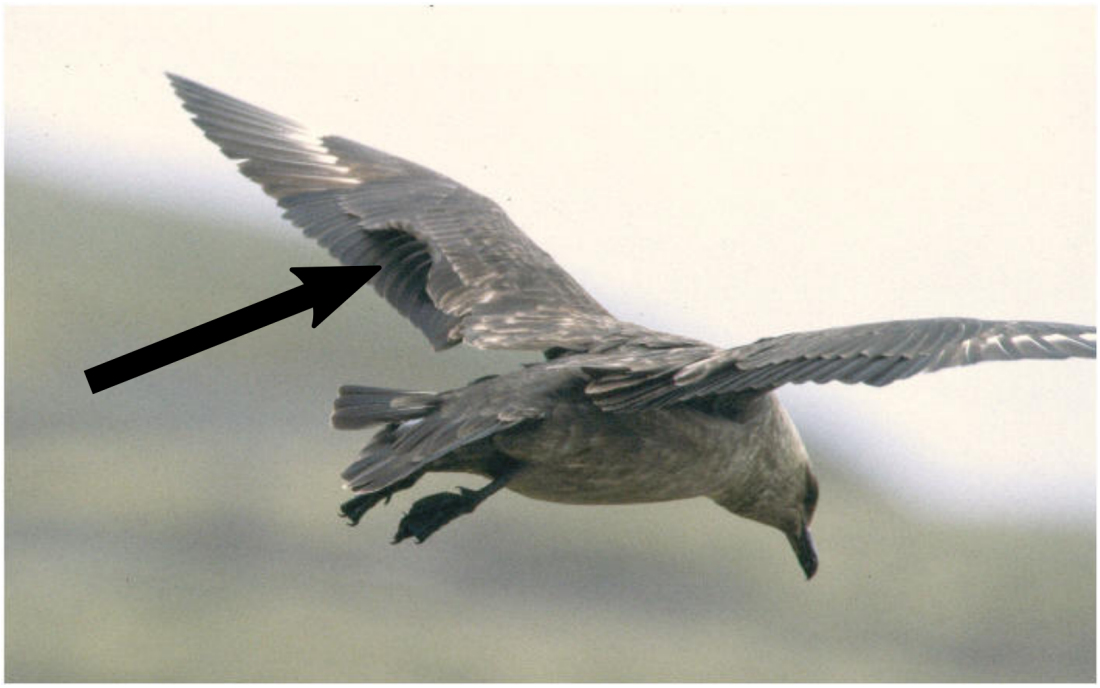

Figure 2.2: Skua bird during landing phase. The passive covert feathers pop up to delay flow separation. Photograph by Rechenberg et al. [33].

against coastal inundation (waves and storm surges) and shoreline erosion [35].

2.2.2 Slender Structures in Biomimetics/Industry

Over millions of years evolution has had the opportunity to adapt, refine, and improve upon nature's own designs. As a result, many organisms have evolved to become highly adept at certain tasks and functions. It is natural then for engineers and designers to emulate these designs and apply them to real-world industrial problems. This is the field of biomimetics and designs based on slender structures have found many uses in a wide range of applications.

Advances in fabrication techniques at the microscale have led to a surge of interest in using artificial cilia to control processes in microfluidic devices (e.g. lab-on-a-chip). Based on the ciliary systems found in nature, numerous studies have examined how artificial cilia can be used to pump fluids [1], promote mixing [36], or sort suspended particles [37] (see Figure 2.3). Depending on the specific application, these structures can be passive, where their motion is entirely driven by the fluid, or actuated through a range of external stimuli, for example via an

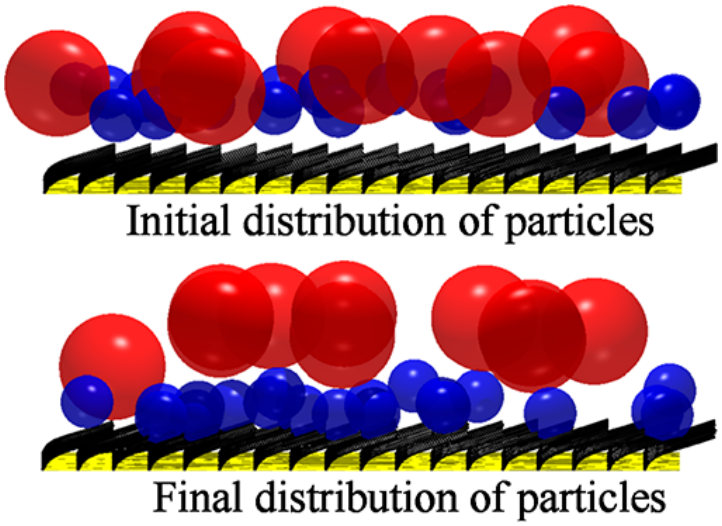

Figure 2.3: Initial and final positions of particles with varying size suspended within a fluid. Due to interplay between the attractive and lift forces generated by adhesion effects and cilia actuation, particles of a smaller size are attracted to the wall whereas larger particles are repelled. Reprinted from Tripathi et al. [37], Copyright (2013), with permission from American Chemical Society.

electrostatic or magnetic field [1].

As discussed in the previous section, biological flow sensors have evolved to become extremely sensitive to even the slightest of disturbances. Other advantages include the ability to focus on specific frequency bands or filter noisy signals [38]. With advances in micromachining and microelectromechanical systems-based design, a lot of effort has been dedicated to mimicking these systems for industrial flow sensing applications. These miniaturised devices have the potential to offer highly sensitive, efficient, and reliable flow sensing capability [30].

Flow control for aerodynamic surfaces is another area that has drawn inspiration from nature. Many studies have examined how self-adaptive hairy coatings can be exploited to improve aerodynamic performance [2, 9, 39]. By modifying the wake pattern and vortex shedding behind a bluff body (see Figure 2.4), these systems can significantly reduce drag. Furthermore, some studies have also reported large reductions in lift fluctuations. This is particularly relevant for structures such as underwater cables, pipelines, or offshore platforms, where vortex induced vibrations (VIV) can eventually lead to fatigue failure [40–42].

Due to the increasing demand for sustainable energy production, alternative forms of power generation are attracting a lot of attention. Energy harvesting,

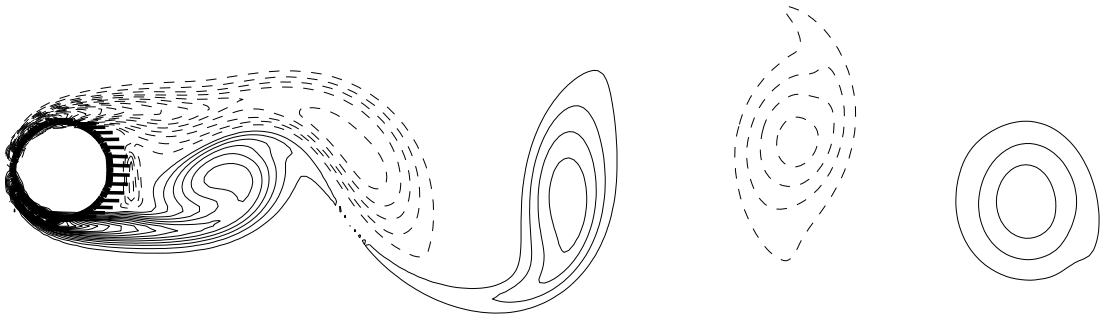

Figure 2.4: Vorticity contours aft of cylinder with a row of passive flaps. Reprinted from Favier et al. [2], Copyright (2009), with permission from Cambridge University Press.

whereby energy is extracted from ambient sources (e.g. wind, tidal, or river flows), is one such area. Using fluid flows to generate electricity is not a new concept and has found relative success in the form of wind turbines or hydropower. However, recently there has been a considerable amount of work examining how the dynamic interactions of slender structures can be used to extract energy from fluid flows. Devices such as the energy harvesting eel [43, 44] exploit the flapping flag instability to convert fluid kinetic energy into strain energy within the structure. The strain energy is then converted into electrical energy via piezoelectric patches attached to the structure. The compact design of these systems makes them particularly attractive and allows multiple devices to be arranged in close proximity. These designs are expected to be especially useful in applications such as portable electronics or remote sensing networks [45, 46], where self-sustaining energy production can dramatically reduce running and maintenance costs.

2.3 Physics of Slender Structures

2.3.1 Nondimensional Parameters

It is common in fluid-structure interaction (FSI) problems, and wider engineering problems in general, to group relevant parameters into dimensionless groups. This provides physical insight by highlighting important terms and their relative influence. Furthermore, since these parameters are dimensionless, they can be used to relate different problems across different physical scales. For the problems

of interest in this work, the relevant nondimensional parameters are the Reynolds number, mass ratio, and bending stiffness:

$$Re = \frac{u_\infty L}{\nu}, \qquad \mathcal{M} = \frac{\rho_s h}{\rho_f L}, \qquad \mathcal{K} = \frac{EI}{\rho_f u_\infty^2 L^3} \tag{2.1}$$

where u_∞ is the freestream velocity, ν is the kinematic viscosity, L and h are the length and thickness of the structure, ρ_f and ρ_s are the fluid and structural densities, E is the Young's modulus, and I is the second moment of area. Note that for 2D problems the out-of-plane depth is set to unity.

The first term is the Reynolds number, which describes the ratio between inertial and viscous forces in the flow. At low Reynolds numbers, viscous forces dominate and the flow is typically characterised as laminar. In contrast, at high Reynolds numbers inertial forces become dominant and the flow typically exhibits signs of turbulence or other inertia-driven instabilities (e.g. von Kármán vortex shedding). For problems involving slender structures, the Reynolds number can vary greatly depending on the specific application, from $\mathcal{O}(10^{-5})$ in cell propulsion [47] to $\mathcal{O}(10^7)$ for trailing-edge flaps on commercial aircraft.

The other two dimensionless parameters describe the structural material properties (density and Young's modulus). The first term is the nondimensional mass of the structure, which indicates whether the structural motion is primarily inertia-driven or driven by the flow. The second term is the nondimensional bending stiffness and relates the internal elastic forces to the hydrodynamic forces exerted by the fluid (dynamic pressure). When this ratio is $\mathcal{O}(1)$ or lower, one may expect fluid-induced deformation of the structure. Sometimes the inverse of this ratio, the Cauchy number, is adopted instead:

$$Ca = \frac{\rho_f u_\infty^2 L^3}{Eh^3} \tag{2.2}$$

Note that this definition follows de Langre [8], which accounts for the structural slenderness by modifying the classical form of the Cauchy number by the aspect ratio, L/h. For problems involving dynamic interactions, it may be more appropriate

to describe the system in terms of the reduced velocity:

$$\mathcal{U} = u_\infty L\sqrt{\frac{\rho_s h}{EI}} \tag{2.3}$$

The reduced velocity is the ratio between the period of the undamped natural frequency of the structure and the characteristic time-scale of the fluid. When this ratio is $\mathcal{O}(1)$, dynamic effects, such as resonance or lock-in, can be expected [8]. Since the reduced velocity depends on both material properties, it can be related to the dimensionless mass ratio and bending stiffness via:

$$\mathcal{U} = \sqrt{\frac{\mathcal{M}}{\mathcal{K}}} \tag{2.4}$$

Some studies suggest modifying the mass ratio and bending stiffness to account for the extra mass due to the presence of the structure [14]:

$$\mathcal{M}^* = \frac{\Delta\rho}{\rho_f L}, \qquad \mathcal{K}^* = \frac{EI}{\Delta\rho u_\infty^2 L^2} \tag{2.5}$$

where $\Delta\rho = (\rho_s - \rho_f)h$ is the extra mass due to the presence of the structure. These values can be related back to the previously defined values via:

$$\mathcal{M}^* = \mathcal{M} - \frac{h}{L}, \qquad \mathcal{K}^* = \frac{\mathcal{K}}{\mathcal{M}} = \frac{1}{\mathcal{U}^2} \tag{2.6}$$

Clearly, the difference between the two definitions for the mass ratio becomes negligible as the aspect ratio increases. However, things are less clear for the stiffness term since the modified definition has a dependence on the structural density. As will be shown later, this can make it difficult to separate the effects of different material properties, if adopting these modified definitions.

2.3.2 Single Structures

For sufficient flow speed and structural compliance, a slender structure will undergo drag-induced deformation in response to the fluid loading. This type of behaviour is common in vegetation, where it is known as reconfiguration. Reconfiguration is a drag-reducing response and plays an important role in protecting plants against lodging and uprooting [48]. In classical fluid mechanics, rigid structures in a high-speed flow will experience a drag that scales with the square of the velocity, such that:

$$F_D \propto u_\infty^2 \tag{2.7}$$

However, for deformable structures the shape of the geometry is also a function of the velocity. The resulting deformation reduces the drag by decreasing the frontal area and streamlining the structure (thereby reducing the wake) [49]. This modifies the drag scaling such that:

$$F_D \propto u_\infty^{2+\mathcal{V}} \tag{2.8}$$

where $\mathcal{V}$ is the Vogel exponent [50, 51], which is negative for drag-reducing reconfiguration. Typical values for the Vogel exponent in different plant species are -0.2 to -1.2 [52].

Alben et al. [53, 54] studied the drag-reducing reconfiguration of a flexible fibre in a flowing soap film, both experimentally and numerically. Soap film flows are used often in experimental fluid mechanics to approximate 2D flows [4]. They reported a self-similar bending which resulted in a Vogel exponent of $-2/3$, leading to a drag scaling:

$$F_D \propto u_\infty^{4/3} \tag{2.9}$$

Using dimensional analysis and an empirical drag formulation, Gosselin et al.

[55] predicted the same value for the Vogel exponent for a plate of finite width. However, their experiment could not reach this asymptotic value, due to the onset of flutter. Extending the work of Alben et al. [53, 54] to include viscosity, Zhu and Peskin [56] and Zhu [57] performed numerical simulations and reported a dependence of the Vogel exponent on Reynolds number, with the Vogel exponent increasing in magnitude towards $-2/3$ as the Reynolds number increased.

Much of the work concerning the dynamic interactions of single slender structures has revolved around the flapping flag problem. Not only is it an interesting problem that exhibits a rich set of nonlinear physics, but it also has direct applications in areas such as structural resilience, energy harvesting, and flow control, to name a few. Zhang et al. [58] seem to have initiated the recent resurgence of interest in this problem with their soap film experiment. By attaching a silk thread within the soap film and increasing its length, Zhang et al. [58] found three behavioural states: a static stretched state, a regular flapping state, and an irregular flapping state. Figure 2.5 shows the flap behaviour and flow topology for each of these three states. Zhang et al. [58] also found a region of bistability, where the thread could be tripped into a particular state by either knocking it (stretched state to flapping state) or by holding it still for a short period of time (flapping state to stretched state). They also tested two side-by-side threads and found a correlation between the level of interaction and separation distance, with in-phase flapping observed for small separation, out-of-phase flapping for intermediate separation, and uncorrelated flapping for large separation. Subsequent numerical studies [59–61] supported these findings and reported similar dynamic states and dependencies. Interestingly, Zhu and Peskin [59] reported similar behaviours at a Reynolds number two orders of magnitude lower than that given in the experiment, indicating a limited dependence on Reynolds number. Furthermore, Zhu and Peskin [59] observed that a massless flag cannot exhibit flapping motion, and attribute this to the inability of the flag to extract work from the flow, since it has no momentum.

While the previous studies examined the role of the flag length in destabilising the behaviour, the observations made by Zhu and Peskin [59] suggest that the structural inertia also plays a part. Shelley et al. [62] examined the response of heavy flags and attributed the onset of flapping motion to the destabilising structural inertia overcoming the stabilising effects of stiffness and tension. Connell and Yue [13] tested a range of mass ratios and found the same three dynamic

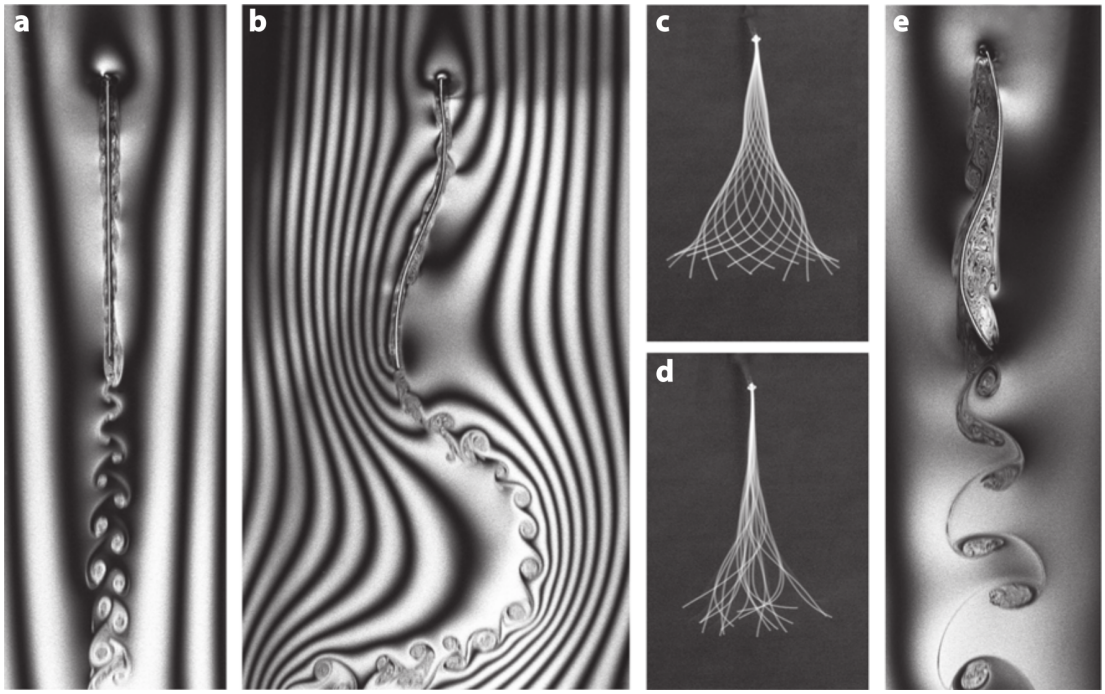

Figure 2.5: Flow visualisation of a flexible filament immersed in a soap film. (a) Stretched state. (b) Flapping state. (c) Snapshots of flapping profile. (d) Snapshots of irregular flapping profile. (e) Flow around a stiff wire which has been permanently bent into the shape of a flapping filament. Reprinted from Shelley and Zhang [4], Copyright (2011), with permission from Annual Reviews.

states for increasing structural mass (static stretched, regular flapping, and chaotic flapping). By mapping out critical mass ratios for the transition between these states over a range of Reynolds number and flag flexibility, they also highlighted the destabilising effects of these terms. Relating these findings back to the previous studies, which examined length as the controlling parameter, Equation 2.1 shows that the length of the flag affects all of the relevant dimensionless quantities in a destabilising manner. Also noteworthy is the effect of flow speed, which is also destabilising, and many studies have reported critical flow velocities for describing the onset of flutter [17, 63, 64].

In most engineering applications flutter is an undesired phenomenon. However, many studies have looked to exploit the flapping flag instability for energy harvesting applications. While some studies have found success with this approach [65, 66], the regular flag configuration is limited by the large critical flow velocities required to realise flapping and the relatively low strain energies associated with the flapping motion [67]. Recently, the inverted flag configuration, whereby the trailing edge is clamped and the leading edge is free to deflect, has become a

promising alternative to the regular configuration. Initially proposed by Kim et al. [68], the inverted flag offers several advantages over its classical counterpart, including lower critical velocities for flapping motion and significantly larger strain levels. Furthermore, the main deformation mode is the first mode. This means the curvature of the deformation has the same sign along the length of the flag, and thus the piezoelectric strips can interact constructively [45].

Although conceptually similar, the stability and dynamical properties of the inverted flag exhibit some distinct differences compared to the regular configuration. In contrast to the regular configuration, where flapping occurs above a critical velocity, the experiments of Kim et al. [68] showed that the flapping of an inverted flag occurs over a window of flow velocity (or dimensionless bending stiffness). They classified the dynamics of the flag into three states (see Figure 2.6): a straight state with negligible deflection, a flapping state with large amplitude deflection, and a deflected state where the flag is fully deflected to one side. Also in contrast to the regular configuration, they noted that the mass ratio has little effect on the stability boundaries between these states, although it does affect the actual dynamics of the flapping. Subsequent numerical studies reported similar findings [70–72] and suggested that the transition to flapping motion is a static divergence instability. Sader et al. [69] showed through stability analysis that this is indeed the case, which explains why the structural mass has little effect on the instability boundary. Sader et al. [69] also examined the mechanism of the flapping motion, and determined it to be a VIV. As the flag deflects due to

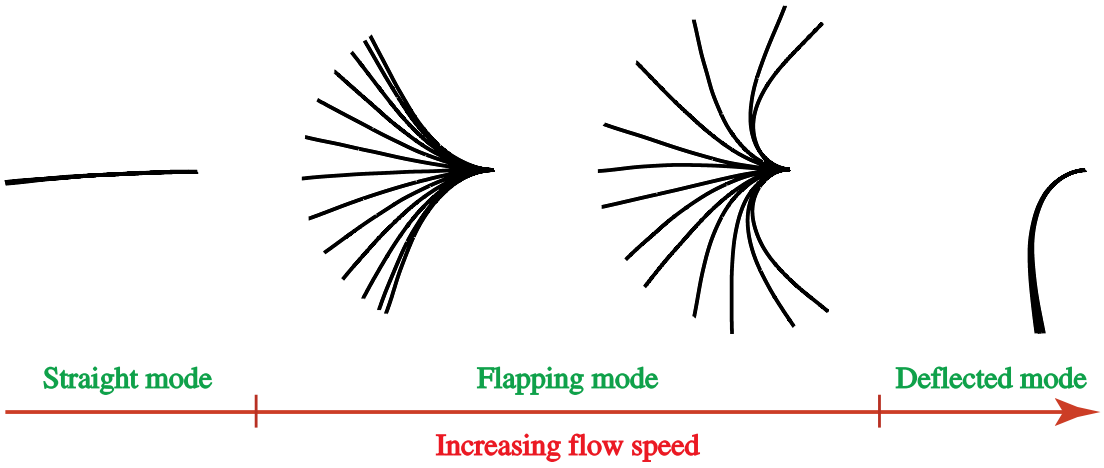

Figure 2.6: Flapping states of inverted flag for increasing flow speed. Reprinted from Sader et al. [69], Copyright (2016), with permission from Cambridge University Press.

the divergence instability, the flow becomes separated at the leading edge. This leads to unsteady vortex shedding, which synchronises with the flapping motion of the flag. The VIV mechanism for the flapping motion had also been suggested in previous studies, and explains why flapping motion cannot be realised at low Reynolds numbers ($Re < 50$) [71].

2.3.3 Multiple Structures

Like single structures, arrays of multiple structures can also undergo a static reconfiguration [12, 73]; although, there are few studies which have attempted to characterise the Vogel exponent for such configurations. Gosselin and de Langre [74] considered the flow through a poroelastic ball covered in flexible fibres. For increasing surface density, they showed a transition from the classic reconfiguration regime, where $\mathcal{V} = -2/3$, to the homogeneous porous body regime, where $\mathcal{V} = -1$. Barsu et al. [75] performed experiments on a submerged canopy model to examine the sheltering effect of adjacent plants. Their results showed that all plants within the canopy undergo approximately the same amount of deflection. Furthermore, below a critical plant spacing the Vogel exponent is independent of canopy density.

A sizable portion of the research concerning multiple slender structures has been geared towards studies of vegetation. The velocity profile over a patch of vegetation (canopy) is modified by the drag associated with the presence of the canopy, as shown in Figure 2.7. The discontinuity in drag at the top of the canopy separates

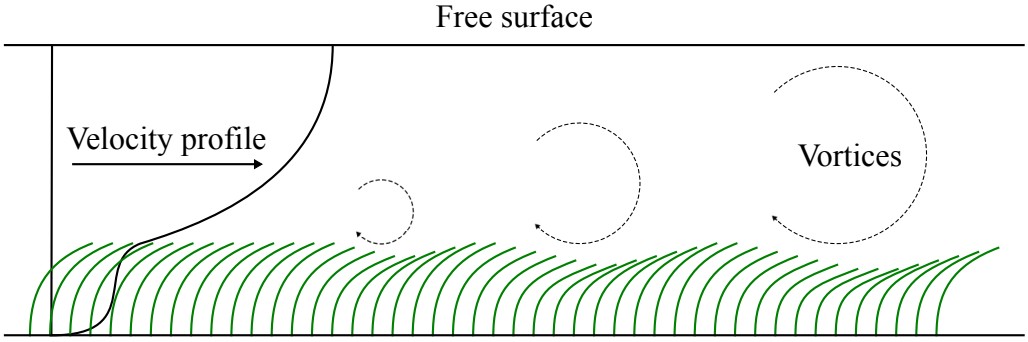

Figure 2.7: Velocity profile and monami motion over a submerged canopy. An inflection point is formed due to the drag associated with the canopy. This leads to Kelvin-Helmholtz instabilities which roll up to form coherent vortices.

the flow into two regions: the bulk flow over the top of the canopy and a region of low-velocity flow within the canopy. For sufficiently dense canopies, the shear layer at the top of the canopy will form an inflection point in the velocity profile [76]. This makes the flow susceptible to Kelvin-Helmholtz instabilities, similar to a mixing layer. The Kelvin-Helmholtz instabilities roll up to form coherent vortical structures which are advected through the canopy [7]. For terrestrial and deeply submerged canopies, these vortices are highly 3D due to interactions with the upper boundary layer turbulence [77]. However, for shallow submerged vegetation, where the presence of the free surface means there is no upper boundary layer, the Kelvin-Helmholtz vortices dominate the flow and thus remain predominantly lateral [78]. For aquatic vegetation, the classification of submergence is given by the ratio of flow depth (H) to plant height (L), with $^{H}/_{L} > 10$ for deeply submerged, $^{H}/_{L} < 5$ for shallow submergence, and $^{H}/_{L} < 1$ for emergent vegetation [77].

Although the canopy shear layer is similar in principle to a mixing layer, there are some important differences. In a normal free shear layer, the Kelvin-Helmholtz instabilities grow continually. However, for submerged vegetation the instabilities grow over a development length until they reach a fixed size and penetration depth into the canopy, which can be of the order of the canopy height [77]. In a normal mixing layer, the Kelvin-Helmholtz vortices are centred on the inflection point and thus advect with the same velocity as the velocity at the inflection point. However, in a canopy shear layer the vortices are displaced upwards from the inflection point, where the flow velocity is higher. The frequency of the instability is related to the mean flow speed, shear layer thickness, and momentum thickness, and can be estimated via [7, 79]:

$$f_{KH} = St_n \left(\frac{u_1 + u_2}{2\theta} \right) \tag{2.10}$$

where St_n is the natural Strouhal number associated with mixing layers, θ is the momentum thickness, and u_1 and u_2 are the low and high-stream velocities (see Figure 2.8). Through theory and experiment, the natural Strouhal number is known to be approximately $St_n \approx 0.032$ [79]. Equation 2.10 states that the frequency of the vortices is proportional to their advection speed and inversely proportional to the momentum thickness (which is an indicator of vortex size).

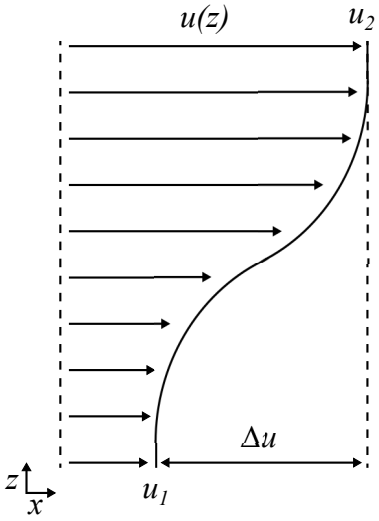

Figure 2.8: Mixing layer velocity profile with key quantities for Equation 2.10.

The momentum thickness is given by:

$$\theta = \int_{-\infty}^{\infty} \left[\frac{1}{4} - \left(\frac{u - \frac{1}{2}(u_1 + u_2)}{u_2 - u_1} \right)^2 \right] \mathrm{d}z \tag{2.11}$$

The passage of the Kelvin-Helmholtz vortices generates large variations in the flow velocity at the top of the canopy. For sufficient flow velocity and plant flexibility, the vortices deflect individual plant stems as they pass by [34]. This is observed as a travelling wave of deflection which propagates through the canopy. This phenomenon is termed honami for terrestrial flows, and monami for aquatic flows. The mechanism by which this deflection occurs is described by Ghisalberti and Nepf [11] as a strong forward sweep in front of the vortex which deflects the canopy down and forwards, followed by a short ejection which pushes the canopy up and backwards. In addition to the honami/monami motion, three other regimes of plant motion exist. In order of increasing flow speed, these four regimes are characterised as erect, gentle swaying, strong coherent swaying (honami/monami), and prone [80]. Figure 2.9 illustrates these four regimes.

Ghisalberti [7] suggests that the honami/monami effect is strictly governed by the mixing layer instability, and that the waving motion is a forced response to the Kelvin-Helmholtz vortices. To support this, Ghisalberti performed experiments on a scaled seagrass model and showed strong correlations between the velocity spectra, waving frequency, and predicted frequency of the Kelvin-Helmholtz

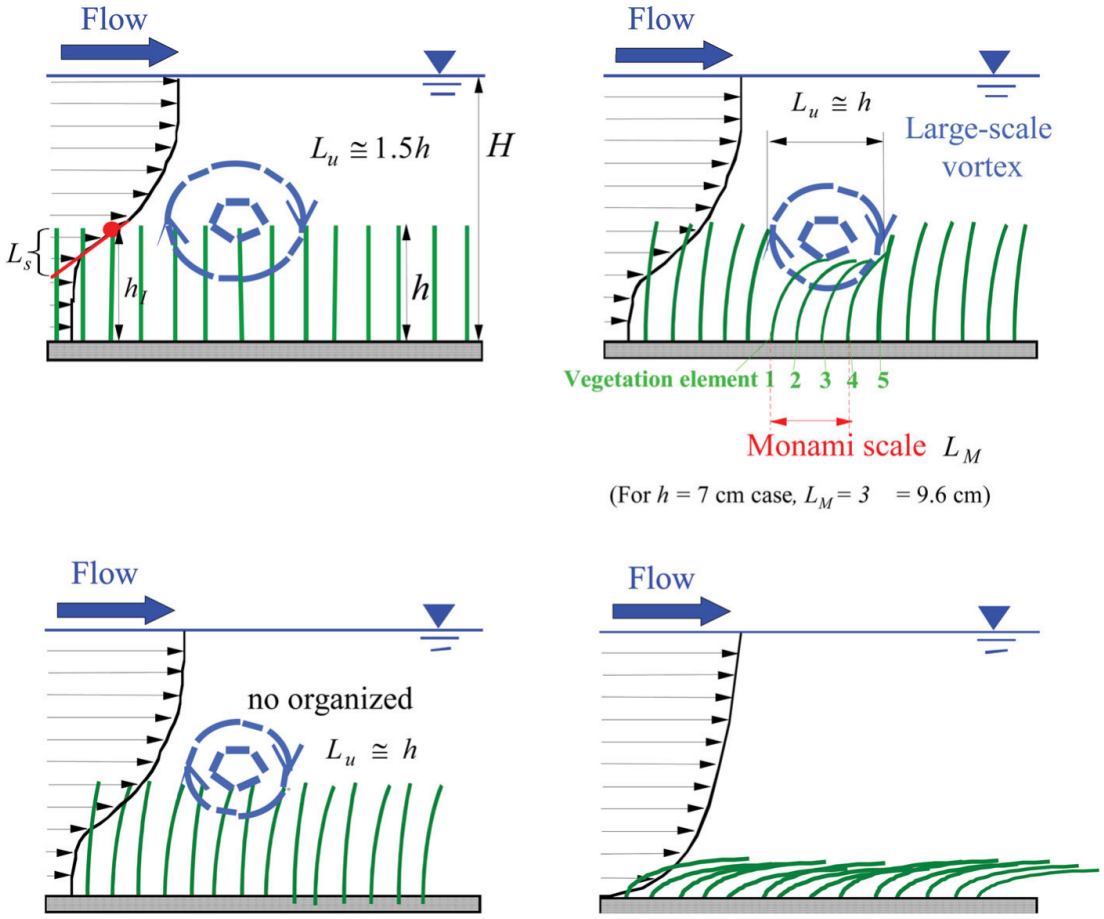

Figure 2.9: Flow patterns and deflection modes for a dense canopy of flexible vegetation. The four regimes of plant deflection are: erect/rigid (top left); swaying (bottom left); monami (top right); and prone/reconfigured (bottom right). Reprinted from Okamoto et al. [12], Copyright (2016), with permission from Taylor & Francis.

vortices (given by Equation 2.10). Furthermore, Ghisalberti noted that the observed frequencies were much lower than the natural frequencies of the model plants. However, combining simple beam theory with the model parameters given in the experiment suggests that the quoted natural frequency for the model plants is the undamped (vacuum) natural frequency. Since the plant models are submerged, the damped natural frequency, which arises from the fluid damping and is likely to be lower, should be used for comparison instead.

In contrast to Ghisalberti [7], multiple studies have reported results which suggest that the honami/monami motion is not a purely fluid-driven response, but instead has some dependence on the structural properties [81–83]. Py et al. [81] developed a simplified fluid model combined with a flexible porous layer for the canopy. While the main instability mechanism was still driven by the Kelvin-Helmholtz instability,

the model indicated that the characteristics of this instability are significantly modified when the canopy compliance is considered. In a subsequent study, Py et al. [82] used an image-correlation technique [10] combined with biorthogonal decomposition [84, 85] to extract the wavelengths and frequencies of the coherent structures over two different plant species (alfalfa and wheat). These two species were selected as they share similar geometric properties (height and spacing) but different material properties (density and stiffness), and therefore different natural frequencies. Interestingly, for both species, the waving motion was found to occur near their own (damped) natural frequency. Moreover, this was observed over the full range of tested wind speeds. Extending their previous model [81], with increasing wind speed they observed a lock-in effect as the frequency of the Kelvin-Helmholtz instability approached the natural frequency of the canopy. In contrast to Ghisalberti [7], this indicates that the honami/monami effect is not solely driven by the mixing layer instability, and in fact the material properties of the vegetation also play a role. That being said, a subsequent study using large eddy simulation combined with a flexible canopy model failed to reproduce the lock-in effect [15], which they attribute to a nonlinear saturation mechanism. Clearly, further work is required in this area to fully understand the lock-in effect and how it relates to the honami/monami instability.

A similar lock-in effect has been observed in flow control applications. Favier et al. [2] used a homogenised porous layer model to examine the effect of a row of passive flaps attached to the aft of a circular cylinder (see Figure 2.4). Under optimum conditions, they observed drag reductions of 15 % and reductions in the lift fluctuations of 40 %, which they attributed to a stabilisation of the wake. In the optimal regime, they observed a travelling wave through the array with a frequency that matched the vortex shedding frequency. Niu and Hu [39] reported similar aerodynamic improvements in their experiment, shown in Figure 2.10. However, they did not observe this travelling wave within the array, probably because of the sparseness of the flaps within the array. Nevertheless, several studies examining densely packed arrays have observed this travelling wave, and many report a lock-in at optimal conditions between the travelling wave frequency and the shedding frequency [9, 86, 87]. These similarities between the honami/monami effect and the travelling wave observed in flow control applications, particularly the lock-in effect, are striking and clearly warrant further investigation.

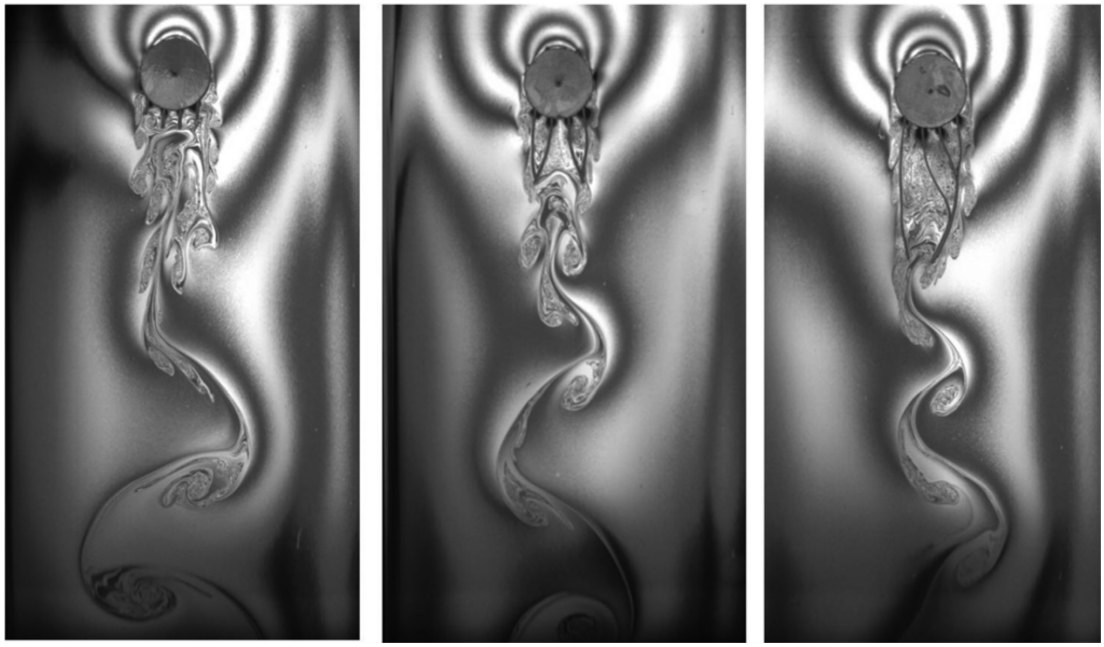

Figure 2.10: Flow visualisation of wake shedding behind a hairy disk with different hair lengths. The ratio between the hair length and disk diameter is 0.25 (left), 1 (centre), and 2 (right). Reprinted from Niu and Hu [39], Copyright (2011), with permission from AIP Publishing.

2.4 Modelling Slender Structures

Experimental and field studies are common for studying the FSI of slender structures [10, 53, 58, 82]. Compared to field measurements, experimental studies have the advantage that the input conditions are well-defined and can be carefully controlled. Furthermore, some field measurements are extremely difficult to obtain, such as the *in vivo* dynamics of cilia, or 3D flow maps of coherent vortices propagating through vegetation canopies. However, translating findings in the lab, which are usually based on idealised cases, to real-world applications is not always straightforward. Moreover, the number and type of measurements that are available is usually limited in some form, which can make analysing complex 3D flow fields difficult [12]. With increasing compute power and more sophisticated modelling techniques, theoretical and numerical modelling is becoming an increasingly attractive approach for these types of systems. Such techniques have the potential to provide detailed datasets for problems that are otherwise experimentally intractable, and a wide range of modelling approaches have already been applied to slender structures. This section serves as a broad overview of these techniques, including theoretical models, homogenised models, and state-of-the-art

high fidelity models that directly account for the individual structures.

2.4.1 Theoretical Models

Theoretical modelling relies on simplifying assumptions to derive rudimentary models that capture the minimum amount of physics required to describe the problem. The details of the assumptions depend on the specific model and problem. Therefore, theoretical modelling encompasses a broad class of different approaches. As such, it is difficult to make any generalisations regarding these models, other than they all rely on some form of simplification. However, common examples of the assumptions used in theoretical models of slender structures include potential flow [53, 88, 89] or piecewise linear velocity profiles [82]. Theoretical models have been applied to several of the aforementioned applications regarding slender structures, including predicting stability boundaries for flapping flags [62, 88, 90, 91], or calculating the frequency and growth rates of the Kelvin-Helmholtz instability in canopy flows [15, 82, 83, 92]. These models are usually inexpensive and often reduce to analytical expressions. Furthermore, they provide clear physical insight by describing the problem with a reduced set of parameters. While their simplicity is the key to their appeal, it is also their main limitation. Due to the simplifications used to derive them, theoretical models are usually only valid over a specific range of conditions. Moreover, by design, they do not include all of the physical effects, and therefore may miss certain subtle behaviours. Also, many models are linear and therefore may not be able to provide detailed information about the full behaviour (such as the amplitudes of unstable modes) [65, 82]. Nevertheless, theoretical models provide efficient and insightful solutions to FSI problems and they have found widespread use in applications involving slender structures.

2.4.2 Homogenised Models

Numerical modelling involves using computers to solve mathematical problems that are too complex to be solved analytically. The field of Computational Fluid Dynamics (CFD) falls within this classification. Due to the fact that the underlying mathematical models are more complex, numerical modelling

techniques are capable of capturing more complicated physics, such as viscous effects or nonlinear structural deformations, compared to theoretical models. In regards to modelling arrays of slender structures, many studies have coupled high fidelity CFD techniques with simplified descriptions of the array. An example of such an approach is the use of homogenised porous layers to represent the array. The porous layer acts as a momentum sink within the fluid equations and takes the form of a bulk drag coefficient for the array. This approach is particularly popular in studies of vegetation [15, 93, 94]; however, it has also found uses in flow control applications [2, 86]. In most vegetation and flow control studies, where the array is densely packed and the global behaviour is of interest, this is usually a fair approximation and leads to significant savings in compute time. However, in the case of deformable structures the porous layer must also deform, and modelling compliance in porous layers is not a straightforward task. Furthermore, these models are not capable of capturing effects at the scale of individual structures, such as stem-scale turbulence, individual structural motion, or spatial heterogeneity, which are known to have significant effects on local quantities, such as the instantaneous velocity field [95].

2.4.3 Direct Models

If local effects are to be captured then the dynamic motion of individual structures and the flow around them must be modelled directly. While there are many examples in the literature of studies that directly model the coupled behaviour of single structures [13, 14, 96], there are very few examples of fully resolved simulations for arrays of slender structures [95, 97]. Moreover, the effects of structural deformation are usually neglected. This is because of the huge computational cost required to solve these problems, since resolving the flow within the array requires a very fine mesh at the sub-structure scale (which for a densely packed 3D array can result in a very large mesh). Nevertheless, these models provide unrivalled detail about the complex nonlinear behaviours that govern the interactions of slender structures, and their popularity is rising.

2.5 Conclusions

In light of this review, a few key observations can be made about the current understanding and capability regarding the FSI of slender structures:

- single slender structures and arrays of slender structures play important roles in a broad range of applications
- the coherent waving motion in arrays of slender structures is exhibited across multiple application areas
- this waving motion and its associated lock-in seem to be linked to special case conditions in flow control applications (e.g. minimal drag and lift fluctuations)
- the mechanism by which this waving instability occurs is yet to be fully understood
- dense arrays of slender structures are typically modelled via an homogenised porous layer which neglects local effects
- the studies that do resolve the individual structures have mostly focussed on single structures aligned with the flow (e.g. flapping flag)
- there is a limited amount of research regarding the direct numerical modelling of wall-mounted arrays of slender structures

Based on these observations, it seems natural to develop a model that incorporates the direct modelling of individual structures and apply it to dense arrays of wall-mounted slender structures, with a particular focus on the coherent wave-like motion and its associated lock-in effect. This is the topic of the present work. While a portion of the examples discussed in this chapter involve hair-like structures, which are characterised by highly 3D flow fields, this work focusses on 2D applications. The main reason for this is simplicity, as it allows a better understanding of the basic behaviours before moving to more complex applications. The following chapters introduce the components of the numerical model, which

is described in terms of three modules: the fluid solver, the structural solver, and the coupling approach between the two solvers. Chapter 3 will now describe the fluid solver.

Chapter 3

Fluid Solver

The following three chapters describe the core components of the numerical model: the fluid solver, the structural solver, and the coupling approach. This chapter focusses on the fluid solver, which is based on the lattice Boltzmann method (LBM). To begin with, the basic theory of the LBM and its governing equations are introduced. Then, specific extensions (e.g. forcing scheme, boundary conditions) pertaining to the present work are described in detail. Finally, some key aspects concerning the LBM and how it compares to traditional Navier-Stokes-based solvers are discussed.

3.1 Lattice Boltzmann Method

There are two extremes to modelling fluid dynamics. One approach is to consider the fluid at its most basic level: a collection of individual molecules which are continuously moving and colliding. Atomistic techniques, such as molecular dynamics, explicitly track the motion of individual fluid molecules and calculate their trajectories based on Newton's equations of motion. This approach describes the fluid at its most fundamental level and allows highly detailed time histories of the complete system state. However, the number of molecules in a given system is typically of the order of Avogadro's constant (6×10^{23}), and with today's processing power this limits the availability of these techniques to nanoscale problems [98]. Furthermore, these methods are extremely sensitive to initial conditions and any uncertainty in the initial state of the system is likely to grow

exponentially as the simulation progresses [99].

The second extreme is to neglect the details of the individual molecules and instead model the fluid as a continuum. Such an approach assumes that the ratio between the molecular mean free path and a characteristic length scale (Knudsen number) is much smaller than one. The Navier-Stokes equations are founded on this assumption; making them well suited to large scale problems. While their efficiency is one of the reasons for their widespread popularity, continuum-based methods can have difficulties incorporating certain molecular-level physics which govern complex fluid behaviour (e.g. multiphase flows).

The lattice Boltzmann method (LBM) exists as a middle ground between these two extremes (see Figure 3.1) and shares similarities with both. The LBM derives from kinetic theory, which aims to describe macroscopic gas behaviour by considering the microscopic motion of molecules [100]. Although molecular interactions are considered, individual molecules are not explicitly tracked as in atomistic methods. Instead, kinetic theory relies on a distribution function to describe the statistical properties of a collection of molecules and their evolution. By realising that the macroscopic behaviour of the fluid is not sensitive to the finer details of individual molecular motion, the mesoscopic description provided by the LBM can correctly capture the governing macroscopic laws [101].

The LBM developed out of the need to address some of the shortcomings of lattice

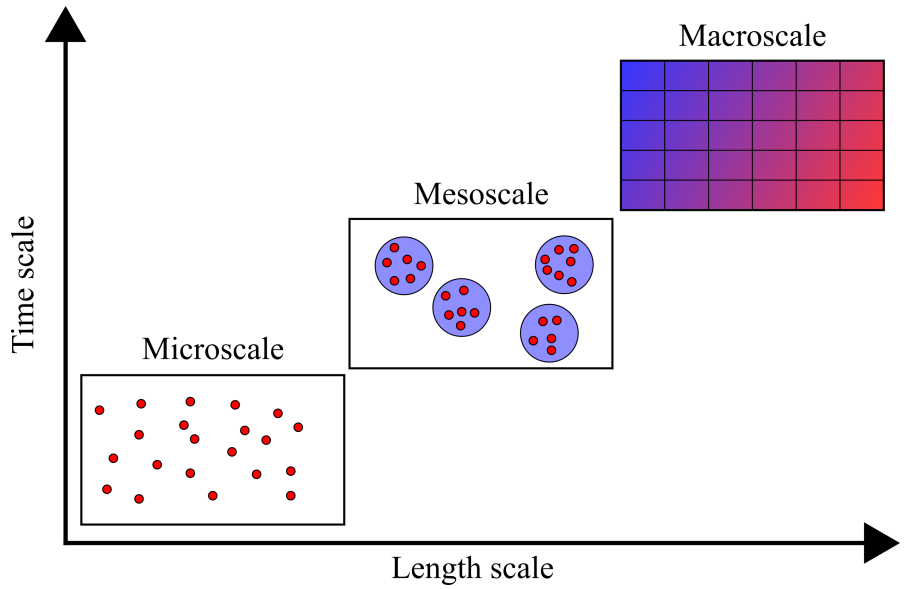

Figure 3.1: Spatial and temporal scales for the micro, meso and macro-scales.

gas cellular automata (LGCA) [102]. LGCA, which itself derives from cellular automation methods, attempts to describe fluid motion through discrete particles which propagate and collide on a discrete lattice according to simple collision rules [103]. One of the key advantages of this approach is that it only requires Boolean values to represent either an empty or occupied site. This eliminates round-off error which is always present in any calculations involving floating-point arithmetic. Another advantage is the calculations are local, making LGCA very amenable to parallel computation on high-performance computing (HPC) systems. Hardy et al. [104] were the first to introduce LGCA; however, the square lattice they used lacked the rotational invariance needed to correctly recover the Navier-Stokes equations. Frisch et al. [105] addressed this issue by introducing a hexagonal lattice model. The hexagonal lattice provided sufficient symmetry to satisfy Galilean invariance and thus, for the first time, LGCA could reproduce the Navier-Stokes equations. Although this was a landmark development, there were still several issues that prevented the widespread use of LGCA for practical applications. One of these stemmed from one of its main advantages: Boolean arithmetic. The use of Boolean values generated statistical noise which needed spatial and temporal averaging to smooth it out, making it quite inefficient. McNamara and Zanetti [106] solved this by abandoning the Boolean values and replacing them with real-valued distribution functions. While this removed the noise, the notion of a round-off free approach to solving fluid problems also had to be abandoned. This development, along with enhancements made to the collision operator [107, 108], is considered to be the birth of the LBM from LGCA.

Since the LBM derives from LGCA, it inherits many of the advantages that initially made LGCA so appealing. The LBM retains the locality of the calculation, making it very well suited to massively parallel computing. Furthermore, its foundation in kinetic theory provides a clear link to the microscopic interactions that govern certain macroscopic behaviours. That being said, by neglecting the unnecessary microscopic details, the LBM can handle fluid problems on a scale comparable to traditional Navier-Stokes-based approaches. This has led to a steady growth in popularity and there are now many examples in the literature of LBM-based applications for a wide range of flow problems; including fluid-structure interactions [42, 109–111], biological flows [112, 113], microfluidics [114, 115], and particle suspensions [116–118], to name a few.

3.2 Boltzmann Equation

The cornerstone of kinetic theory is the Boltzmann equation:

$$\frac{\partial f}{\partial t} + \boldsymbol{\xi} \cdot \nabla f = \Omega(f) \tag{3.1}$$

where $\boldsymbol{\xi}$ is the molecular velocity, and $\Omega(f)$ is a collision operator. For now, while the main ideas are developed, only the unforced form of the Boltzmann equation is considered. The distribution function, $f(\mathbf{x}, \boldsymbol{\xi}, t)$, represents the proportion of molecules at position $\mathbf{x}$, and time t, with velocity $\boldsymbol{\xi}$. Equation 3.1 shows that the transport of the distribution function is governed by a linear advection term (streaming) and a nonlinear term which describes the intermolecular interactions (collision). The Boltzmann equation is one of the most fundamental equations for describing gas dynamics [119]. Contrary to the Navier-Stokes equations, which operate on the macroscopic quantities (e.g. density, pressure, velocity), the Boltzmann equation describes the evolution of the distribution function. Even so, this fine-grained description is still able to recover the correct macroscopic behaviour.

The familiar macroscopic fluid quantities can be calculated from the weighted moments of the distribution function, which are integrated over velocity space. The first three moments are:

$$\rho(\mathbf{x}, t) = \int f(\mathbf{x}, \boldsymbol{\xi}, t) \, \mathrm{d}\boldsymbol{\xi} \tag{3.2}$$

$$\rho\mathbf{u}(\mathbf{x}, t) = \int \boldsymbol{\xi} f(\mathbf{x}, \boldsymbol{\xi}, t) \, \mathrm{d}\boldsymbol{\xi} \tag{3.3}$$

$$\rho E(\mathbf{x}, t) = \frac{1}{2} \int |\boldsymbol{\xi}|^2 f(\mathbf{x}, \boldsymbol{\xi}, t) \, \mathrm{d}\boldsymbol{\xi} \tag{3.4}$$

where $\rho(\mathbf{x}, t)$ is the fluid density, $\mathbf{u}(\mathbf{x}, t)$ is the macroscopic velocity, and $E(\mathbf{x}, t)$ is the specific energy.

Central to the idea of kinetic theory is the notion of equilibrium. In a gas left undisturbed for a sufficient amount of time, the cumulative effect of repeated collisions will tend to evenly distribute the molecular velocities around a mean macroscopic velocity [120]. This state is termed the equilibrium distribution and gives rise to the Maxwell-Boltzmann distribution:

$$f^{eq}(\rho, T, \boldsymbol{\xi}, \mathbf{u}) = \rho \left(\frac{1}{2\pi RT} \right)^{d/2} \mathrm{e}^{-|\boldsymbol{\xi}-\mathbf{u}|^2/(2RT)} \tag{3.5}$$

where f^{eq} is the equilibrium distribution function, T is the temperature, and R is the gas constant.

3.3 Discretisation

The continuous form of the distribution function has seven independent variables, all of which are continuous and therefore have infinite degrees of freedom. To make the problem computationally tractable the Boltzmann equation must be discretised. This is done in two steps. First, it is discretised in velocity space by restricting the advection of the distribution function to discrete velocity directions. This results in the semi-discrete Boltzmann equation:

$$\frac{\partial f_i}{\partial t} + \mathbf{c}_i \cdot \nabla f_i = \Omega(f_i), \qquad i = 0 \ldots q-1 \tag{3.6}$$

where $\mathbf{c}_i$ represents the ith component of the q discrete velocities. This discretisation is performed through either a small Mach number [121, 122] or Hermite series expansion [119, 123]. Depending on the order of the chosen polynomials, the Hermite expansion permits approximations of the Boltzmann equation up to any arbitrary order, even those not captured by the Navier-Stokes equations [123].

The second step in the discretisation process is to discretise in space and time.

This gives rise to the lattice Boltzmann equation:

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) - f_i(\mathbf{x}, t) = \Delta t \, \Omega_i(\mathbf{x}, t) \tag{3.7}$$

Equation 3.7 is the driving equation behind the LBM. The left-hand side shows a propagation of the distribution function through space whereas the right-hand side describes a local collision which is nonlinear.

The discrete versions of the equilibrium function and macroscopic relations are:

$$f_i^{eq}(\mathbf{x}, t) = w_i \rho \left(1 + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{c}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2} \right) \tag{3.8}$$

$$\rho(\mathbf{x}, t) = \sum_i f_i(\mathbf{x}, t) \tag{3.9}$$

$$\rho \mathbf{u}(\mathbf{x}, t) = \sum_i \mathbf{c}_i f_i(\mathbf{x}, t) \tag{3.10}$$

Note that the equilibrium function depends only on local macroscopic quantities. Here, w_i are velocity-specific weighting factors, and c_s is the lattice speed of sound. These values are all dictated by the chosen lattice model (discussed in Section 3.6). For the standard isothermal LBM, the pressure is directly related to density via the equation of state:

$$p = c_s^2 \rho \tag{3.11}$$

3.4 BGK Approximation

The original Boltzmann collision operator is a nonlinear double integral and notoriously difficult to solve. To overcome this, several alternative collision operators have been suggested. The most popular collision operator is the Bhatnagar-

Gross-Krook (BGK) approximation [124]. The BGK approximation simplifies the collision operator by reducing the relaxation rates of all modes into a single relaxation time [125, 126]:

$$\Omega_i(\mathbf{x}, t) = \frac{1}{\tau} \left[f_i^{eq}(\mathbf{x}, t) - f_i(\mathbf{x}, t) \right] \tag{3.12}$$

where τ is the relaxation time-scale. Equation 3.12 shows a relaxation towards local equilibrium and is valid for any fluid. This is because, while seeming relatively simple, Equation 3.12 contains the equilibrium function (f_i^{eq}), which is nonlinear and can take various forms depending on the type of physics to be modelled (e.g. Navier-Stokes equations) [127]. Substituting this into Equation 3.7 leads to the lattice Boltzmann-BGK (LBGK) equation:

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \frac{\Delta t}{\tau} \left[f_i^{eq}(\mathbf{x}, t) - f_i(\mathbf{x}, t) \right] \tag{3.13}$$

Equation 3.13 is the final form of the Boltzmann equation used throughout most of the literature. It is fully explicit and formally linear in terms of the distribution function. The nonlinearity is instead contained within the local equilibrium function – leading to the mantra that *nonlinearity is local and nonlocality is linear*[1]. This is one of the main advantages of the LBM and allows the recovery of nonlinear macroscopic behaviour from a seemingly simple transport equation.

The relaxation time-scale is related to the fluid viscosity via:

$$\nu = \left(\tau - \frac{1}{2} \right) c_s^2 \Delta t \tag{3.14}$$

Equation 3.14 shows that reducing the relaxation time-scale leads to a lower viscosity and faster relaxation towards equilibrium. It also shows that for a real fluid with positive viscosity the relaxation time-scale must be greater than $1/2$ [128]. In practise, the LBGK scheme exhibits stability problems as the relaxation time-scale approaches $1/2$. The multiple-relaxation-time model [129], developed around the same time as LBGK, helps to alleviate these stability issues by allowing

[1] Krüger et al. [100] attributes this description to Sauro Succi.

different modes to relax at different rates. This allows so-called 'ghost modes' not relevant to the hydrodynamics to be fully relaxed within each collision step. However, selecting appropriate relaxation times for each mode is not always straightforward and is typically problem-dependent. In any case, for low-moderate Reynolds numbers the performance of the LBGK scheme is often sufficient and it is therefore adopted in the present work.

3.5 Forcing

External forces, such as gravity, play a significant role in certain fluid problems. Furthermore, in the present work an accurate forcing scheme is even more crucial due to the approach used to couple the fluid and structural dynamics, which imposes the no-slip condition on the surface of the boundary via an appropriate force distribution applied to the fluid equations. A variety of forcing schemes have been developed over the years, some of which are compared in Guo et al. [130]. The approach adopted in this work modifies the transport and macroscopic equations such that [130]:

$$f_i(\mathbf{x} + \mathbf{c}_i\Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \frac{\Delta t}{\tau}\left[f_i^{eq}(\mathbf{x}, t) - f_i(\mathbf{x}, t)\right] + \Delta t F_i(\mathbf{x}, t) \tag{3.15}$$

$$\rho(\mathbf{x}, t) = \sum_i f_i(\mathbf{x}, t) \tag{3.16}$$

$$\rho\mathbf{u}(\mathbf{x}, t) = \sum_i \mathbf{c}_i f_i(\mathbf{x}, t) + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x}, t) \tag{3.17}$$

where $F_i(\mathbf{x}, t)$ is the discrete form of the force density $\mathbf{f}(\mathbf{x}, t)$, given by:

$$F_i(\mathbf{x}, t) = w_i\left(1 - \frac{1}{2\tau}\right)\left(\frac{\mathbf{c}_i - \mathbf{u}}{c_s^2} + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^4}\mathbf{c}_i\right) \cdot \mathbf{f}(\mathbf{x}, t) \tag{3.18}$$

3.6 Velocity Model

The discretisation step in Section 3.3 implies a discrete lattice representation of the flow domain. To ensure isotropy, which is essential for recovering the Navier-Stokes equations, the lattice structure must be symmetric and contain a sufficient number of velocity directions [101]. The most popular two-dimensional lattice structure, and the one adopted in this work, is the D2Q9 model (where DdQq implies a d-dimensional, q-velocity lattice). Figure 3.2 shows the D2Q9 lattice along with the corresponding velocities.

The corresponding velocities, weights, and lattice speed of sound for the D2Q9 model are given by:

$$\mathbf{c} = \left(\begin{array}{c|cccc|cccc} 0 & 1 & -1 & 0 & 0 & 1 & -1 & 1 & -1 \\ 0 & 0 & 0 & 1 & -1 & 1 & -1 & -1 & 1 \end{array} \right) \frac{\Delta x}{\Delta t} \tag{3.19}$$

$$w_i = \begin{cases} 4/9 & i = 0 \\ 1/9 & i = 1 \ldots 4 \\ 1/36 & i = 5 \ldots 8 \end{cases} \tag{3.20}$$

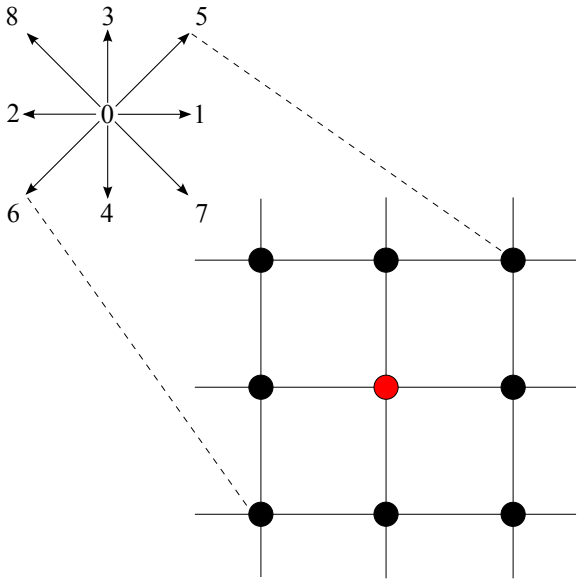

Figure 3.2: D2Q9 lattice model with numbering scheme.

$$c_s = \frac{\Delta x}{\Delta t}\frac{1}{\sqrt{3}} \tag{3.21}$$

3.7 Initial & Boundary Conditions

Like other numerical solvers for fluid dynamics, the LBM cannot be solved without specifying initial and boundary conditions. How these conditions are imposed is extremely important as often they dictate the stability and accuracy of the entire simulation. In traditional Navier-Stokes-based solvers these conditions are usually given in the form of the macroscopic quantities. Since the LBM does not directly operate on these, it is necessary to map these conditions to the mesoscale and specify them in terms of the distribution function. However, this is a significant challenge as the translation from the macroscale to the mesoscale is under-determined and therefore not unique. Closure relations must then be formed and over the years many different approaches have been proposed, each with varying degrees of accuracy and stability.

The most common approach to setting initial conditions is to assign each distribution function its equilibrium value based on the initial local density and velocity. However, for time-dependent problems with non-trivial initial conditions this approach is not strictly correct [131]. For the present work, trivial initial conditions (e.g. uniform density and zero-velocity) are used. This ensures consistency between the macroscopic and mesoscopic conditions at initialisation.

During the streaming step the distribution functions are propagated along the velocity links to their neighbouring sites. After this step, some of the velocity directions on lattice sites which form the boundary of the domain will be unpopulated – specifically, any directions with a wall-normal component pointing into the domain. This is because there are no lattice sites which stream into these locations, as shown in Figure 3.3. The purpose of specifying boundary conditions is to calculate appropriate values for these missing distribution functions, which recover the desired macroscopic conditions and satisfy the conservation laws. The remainder of this section will introduce the key boundary conditions used throughout the present work.

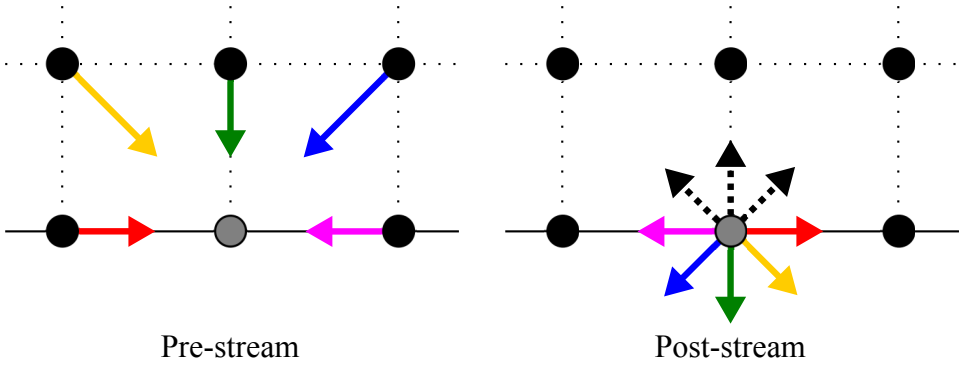

Figure 3.3: Known (solid) and unknown (dashed) components of the distribution function pre and post-stream. The solid grid line indicates the boundary.

3.7.1 Periodic

Periodic boundary conditions are the simplest to implement. They are used to model flows that have a periodically repeating pattern. In the LBM, this is realised by wrapping the outgoing distribution functions back to the boundary on the opposite side of the domain. In many implementations this is the default boundary condition and is incorporated directly into the streaming step. For the 2D problem shown in Figure 3.4, the periodic condition is:

$$f_{\hat{i}}(x, y, t) = f_i(x + L - \mathbf{c}_i, y, t), \qquad \text{at } x = 0 \tag{3.22}$$

$$f_{\hat{i}}(x, y, t) = f_i(x - L - \mathbf{c}_i, y, t), \qquad \text{at } x = L \tag{3.23}$$

where $f_{\hat{i}}(x, y, t)$ indicates the unknown components of the distribution function at $\mathbf{x} = (x, y)$, and L is the length of the domain.

3.7.2 Bounceback

The bounceback condition is the most common (and simple) way to impose the no-slip condition. It is inherited from LGCA and works by reflecting outgoing distribution functions in the opposite direction. There are two approaches to implementing the bounceback condition: full-way or midway. Contrary to what

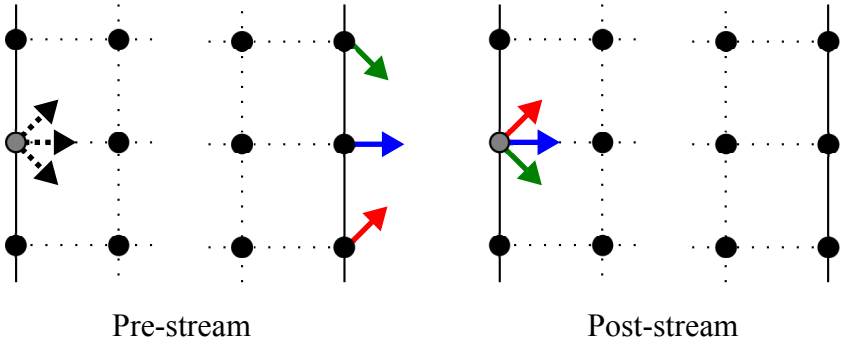

Figure 3.4: Known (solid) and unknown (dashed) components of the distribution function after applying periodic boundary condition. Here, the periodic boundary condition is incorporated directly into the streaming step. The solid grid line indicates the boundary.

their names suggest, both approaches impose the no-slip condition approximately halfway between the lattice sites. The difference lies in when they bounce back the distribution functions. The full-way approach does not bounce back the distribution functions until the next time step (after another collide-stream step). On the other hand, the midway approach bounces them back within the same time step, as shown in Figure 3.5. Although conceptually similar, the full-way approach is only first-order accurate whereas the midway approach yields second-order accuracy [132]. This is because the distribution functions technically make contact with the wall halfway through the time step and should therefore be fully bounced back within the same time step.

In the present implementation a buffer layer is placed around the bounceback sites. This eliminates the need to modify the streaming step. For a bounceback site located at $\mathbf{x}$, the boundary condition is given by:

$$f_{\hat{i}}(\mathbf{x}, t) = f_j(\mathbf{x} - \mathbf{c}_i, t) \tag{3.24}$$

where j is the component of the distribution function opposite to i.

The simplicity of the bounceback condition and the fact it conserves mass exactly make it one of the most popular approaches for imposing the no-slip condition [133]. However, the standard bounceback algorithm cannot represent curved surfaces (although there are extensions which do allow this [134]). This leads to the 'staircase' effect which severely degrades the accuracy of the boundary represen-

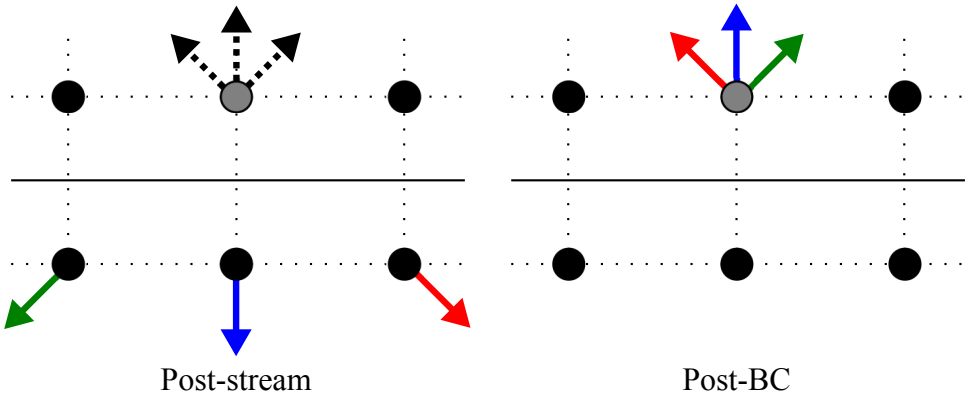

Figure 3.5: Known (solid) and unknown (dashed) components of the distribution function after applying the midway bounceback boundary condition. The solid grid line indicates the boundary.

tation. Furthermore, the bounceback condition introduces a slip velocity at the wall and the actual position of the no-slip condition depends on the relaxation time-scale [135, 136].

3.7.3 Regularised

The regularised boundary condition [137, 138] is used heavily throughout this work to impose Dirichlet conditions. Unlike the previous schemes, the regularised approach calculates new values for all of the distribution functions on a boundary site, not just the unknown ones. The main idea is to reconstruct the distribution functions from the stress tensor, $\mathbf{\Pi}$, which can be computed from the known quantities. The following describes the main steps for implementing this approach.

Given the velocity on the wall at location $\mathbf{x}$ and time t, the density (pressure) can be evaluated by recognising that there are three separate contributions to the density:

$$\rho_+ = \sum_{\{i|\mathbf{c}_{ik}=\mathbf{n}_k\}} f_i \tag{3.25}$$

$$\rho_0 = \sum_{\{i|\mathbf{c}_{ik}=0\}} f_i \tag{3.26}$$

$$\rho_- = \sum_{\{i|\mathbf{c}_{ik}=-\mathbf{n}_k\}} f_i \tag{3.27}$$

where ρ_0, ρ_+, and ρ_- are the zero, inner, and outer-normal components of the distribution function, $\mathbf{n}$ is the normal vector which points into the domain, and k is the dimension along which the normal is directed. Here, ρ_+ contains all the unknown components after the streaming step. Substituting these into the macroscopic relations (Equations 3.9–3.10) gives:

$$\rho = \rho_+ + \rho_0 + \rho_- \tag{3.28}$$

$$\rho \mathbf{u}_k = \rho_+ - \rho_- \tag{3.29}$$

which can be rearranged to give the density in terms of the known quantities:

$$\rho = \frac{1}{1 - \mathbf{u}_k}(2\rho_- + \rho_0) \tag{3.30}$$

Note that pressure boundaries can be imposed in a similar manner; where the pressure and tangential velocities are known, the normal velocity can be calculated via:

$$\mathbf{u}_k = 1 - \frac{2\rho_- + \rho_0}{\rho} \tag{3.31}$$

Once the macroscopic quantities have been evaluated, intermediate values for the unknown components of the distribution function are calculated by making use of the non-equilibrium bounceback assumption [139]:

$$f_{\hat{i}}^{neq} = f_j^{neq} \tag{3.32}$$

The intermediate values for the distribution function are then used to calculate the first-order component of the stress tensor, $\mathbf{\Pi}^{(1)}$, given by:

$$\mathbf{\Pi}^{(1)} = \sum_i \mathbf{c}_i \otimes \mathbf{c}_i f_i^{(1)} \tag{3.33}$$

Here, $f_i^{(1)}$ is the first-order term in the power series expansion of the distribution function:

$$f_i = f_i^{(0)} + \epsilon f_i^{(1)} + \mathcal{O}(\epsilon^2) \tag{3.34}$$

where $\epsilon \ll 1$ is a smallness parameter (Knudsen number) and:

$$f_i^{(0)} = f_i^{eq} \tag{3.35}$$

$$f_i^{(1)} \approx f_i^{neq} \tag{3.36}$$

Finally, all components of the distribution function are reconstructed from the stress tensor via:

$$f_i = f_i^{eq} + \frac{w_i}{2c_s^4}\mathbf{Q}_i : \mathbf{\Pi}^{(1)} \tag{3.37}$$

where:

$$\mathbf{Q}_i = \mathbf{c}_i \otimes \mathbf{c}_i - c_s^2\mathbf{I} \tag{3.38}$$

The regularised algorithm is applicable to any lattice and can be used to specify no-slip, velocity, or pressure boundaries. Furthermore, it is second-order accurate and shows very good stability properties [137, 138]. Because of these characteristics,

the regularised condition is used extensively throughout the present work.

3.8 LBM vs Navier-Stokes

There are many examples throughout the literature that show the LBM asymptotically recovers the Navier-Stokes equations [100, 102, 140]. The procedure rests on the assumption of small Knudsen number, small Mach number and small density fluctuations. Therefore, the LBM can be considered as a numerical solver for the weakly compressible Navier-Stokes equations. In this regard the LBM is similar to artificial compressibility methods used in traditional Navier-Stokes-based solvers [141, 142]. However, it is important to stress that the LBM is not a direct discretisation of the Navier-Stokes equations. Instead, it is based on the Boltzmann equation, which describes gasses at a finer level of detail.

The LBM converges to the incompressible Navier-Stokes equations with error terms $\mathcal{O}(\Delta x^2)$, $\mathcal{O}(\Delta t^2)$, and $\mathcal{O}(Ma^2)$ [137, 143], where Ma is the Mach number. This implies the LBM is second-order accurate in space and time. However, care must be taken when interpreting this result as there is also a compressibility-error term, which scales with $Ma^2 \propto {}^{\Delta t^2}/_{\Delta x^2}$. This implies that the time step should scale with $\Delta t \propto \Delta x^2$ to reduce the compressibility error at the same rate as the $\mathcal{O}(\Delta x^2)$ term. This is equivalent to keeping the lattice viscosity (relaxation time-scale) constant, and is often referred to as diffusive scaling. In practise, this means the LBM really behaves like a first-order method in time, which has serious implications in terms of efficiency. As an example, consider the effect of increasing the mesh resolution by a factor λ. The number of lattice sites scales with λ^d, and the number of time steps needed to simulate the same amount of time scales with λ^2. This means the increase in computational cost scales with λ^{d+2}. Therefore, doubling the mesh resolution in 2D will increase the compute time by a factor of 16 (32 in 3D).

Like any numerical solver the LBM has certain advantages and disadvantages which make it more amenable to some applications than others. The major characteristics of the LBM are summarised below:

Advantages

- algorithm is simple and easy to implement
- heaviest part of the calculation is local (making it well suited to parallel calculation on HPC systems)
- pressure is related through the equation of state (this negates the need to solve the Poisson equation)
- advection term is linear and exact
- foundation in kinetic theory allows molecular level physics to be incorporated more readily

Disadvantages

- velocity and space discretisations are coupled (this leads to square lattices which are inefficient for a large class of problems)
- LBGK scheme becomes expensive for very low ($Re < \mathcal{O}(10^{-3})$) and very high Reynolds numbers ($Re > \mathcal{O}(10^{4})$) [144]
- inherently time-dependent algorithm makes it inefficient for steady-state problems
- streaming step is memory-access intensive (this can be a significant performance bottleneck)
- standard algorithm is isothermal

One of the main reasons for adopting the LBM in this work is the fact it is very amenable to parallel implementation. This is essential for modelling large arrays of slender structures, which is expensive (even in 2D) and requires parallel

computation. Furthermore, the simplicity of the algorithm allows easier implementation and code-coupling. Moreover, some of the disadvantages of the LBM are less relevant for the types of problems under consideration in this work. For example, the square lattice limitation is actually an advantage in the immersed boundary module (discussed in Section 5.1), where the regular structured mesh makes locating support points significantly easier. Furthermore, the LBGK scheme is well suited to the range of Reynolds numbers examined in this work, where $Re = \mathcal{O}(100)$.

3.9 Summary

This chapter has introduced the fluid solver component of the present numerical model. Specifically, the LBM has been discussed and specific details of this technique pertaining to the present work have been highlighted. While the LBM is capable of recovering the Navier-Stokes equations, its theoretical basis is fundamentally different from traditional Navier-Stokes-based solvers. The mesoscopic description provided by the LBM offers a number of advantages relevant to the present work; most notable of which include its parallel performance and ease of implementation. The following chapter will now describe the structural solver component of the numerical model.

Chapter 4

Structural Solver

Following from the previous chapter, which described the fluid solver component of the numerical model, this chapter introduces the structural solver. Two different approaches are used throughout this work to solve the structural dynamics. The first approach is a direct discretisation of the Euler-Bernoulli beam equation, whereas the second approach adopts a finite element method (FEM) for slender structures. As a requirement, both of these approaches permit large (nonlinear) deformations of the structure. The first part of this chapter describes the Euler-Bernoulli beam solver, whereas the second part describes the FEM approach.

4.1 Euler-Bernoulli Beam Solver

The early stages of the work described in this thesis, including the results presented in Chapters 7 and 8, made use of an already existing model for slender structures. The structural solver within this existing model consisted of an approach based on a direct discretisation of the Euler-Bernoulli beam equation. While the limitations of this model meant it was replaced with an FEM-based approach in the later stages of this work, two of the results chapters presented in this work made use of this model, and therefore it is described here. It should also be noted that a significant portion of the present work involved making several other improvements to the original model. A detailed discussion of these developments is presented in Section 6.1.

4.1.1 Governing Equations

The key assumption in deriving the governing equation for an Euler-Bernoulli beam is that the cross sections of the beam do not deform and remain plane and normal to the neutral axis. This implies there is no transverse shear through the beam, which is only a valid approximation for slender structures. Timoshenko theory, a more general theory for beams, permits transverse shear through the thickness and is therefore applicable to beams with an aspect ratio (length-to-thickness ratio) approaching unity [145]. For aspect ratios larger than about 20, both Euler-Bernoulli and Timoshenko theory give similar results (within approximately 5–10 % for the first few eigenvalues) and the relative difference between the two reduces significantly as the aspect ratio is further increased [3]. Therefore, since the focus of the present work is slender (high aspect ratio) structures, the assumption of plane and normal cross sections (Euler-Bernoulli theory) is sufficient.

The governing equation for the Euler-Bernoulli beam solver is obtained by seeking a stationary point of the action integral $J = \int_0^t \mathcal{L}\,\mathrm{d}t$, where $\mathcal{L}$ is the Lagrangian [110]:

$$\Delta\rho\frac{\partial^2\mathbf{X}}{\partial t^2} = \frac{\partial}{\partial s}\left(T\frac{\partial\mathbf{X}}{\partial s}\right) - K_B\frac{\partial^4\mathbf{X}}{\partial s^4} + \Delta\rho\mathbf{g} - \mathbf{F} \tag{4.1}$$

where s is a parametric coordinate along the beam, $\mathbf{X} = (X(s,t), Y(s,t))$ are the spatial coordinates, T is the axial tension, $\mathbf{g}$ is the gravitational acceleration, and $\mathbf{F}$ is an externally applied distributed load. The other two terms describe the material properties of the beam, where $\Delta\rho = (\rho_s - \rho_f)h$ is the extra mass per unit length due to the presence of the structure, and $K_B = EI$ is the bending rigidity. Here, ρ_s and ρ_f are the structural and fluid mass densities, E is the Young's Modulus, and I is the second moment of area. The inclusion of the fluid density in the acceleration term of the structural equation helps to improve the stability of the present approach by limiting the appearance of so-called added-mass instabilities [146, 147].

The inextensibility condition is used to close the system and serves to prevent axial extension of the beam. This condition is satisfied by the tension term in

Equation 4.1, and is given by [14]:

$$\frac{\partial \mathbf{X}}{\partial s} \cdot \frac{\partial \mathbf{X}}{\partial s} = 1 \tag{4.2}$$

Equations 4.1 and 4.2 are nondimensionalised by introducing appropriate reference quantities. These reference quantities are given by L for the spatial coordinates, L/u_∞ for time, $\Delta\rho u_\infty^2$ for the tension term, $\Delta\rho u_\infty^2 L^2$ for the bending rigidity, u_∞^2/L for gravity, and $\Delta\rho u_\infty^2/L$ for the forcing term. Here, L is the length of the beam and u_∞ is the freestream velocity. The resulting nondimensional equations of motion are given by:

$$\frac{\partial^2 \mathbf{X}}{\partial t^2} = \frac{\partial}{\partial s}\left(T\frac{\partial \mathbf{X}}{\partial s}\right) - K_B \frac{\partial^4 \mathbf{X}}{\partial s^4} + Fr\frac{\mathbf{g}}{|\mathbf{g}|} - \mathbf{F} \tag{4.3}$$

$$\frac{\partial \mathbf{X}}{\partial s} \cdot \frac{\partial \mathbf{X}}{\partial s} = 1 \tag{4.4}$$

where K_B is now the dimensionless bending rigidity, $K_B = \mathcal{K}^*$, and $Fr = |\mathbf{g}|L/u_\infty^2$ is the Froude number. Note that this is not the standard definition of the Froude number, which is usually given as $Fr = u_\infty/\sqrt{|\mathbf{g}|L}$.

4.1.2 Discretisation

A staggered grid arrangement is used to discretise the beam, whereby the tension force is defined on the interfaces, while all other quantities are defined on the nodes. Figure 4.1 shows a schematic of this arrangement, where the beam is discretised into $N+1$ nodes. Equations 4.3 and 4.4 are discretised using a finite difference approximation:

$$\frac{\mathbf{X}^{n+1} - 2\mathbf{X}^n + \mathbf{X}^{n-1}}{\Delta t^2} = D_s(T^{n+1} D_s \mathbf{X}^{n+1}) - K_B D_{ssss}\mathbf{X}^n + Fr\frac{\mathbf{g}}{g} - \mathbf{F}^n \tag{4.5}$$

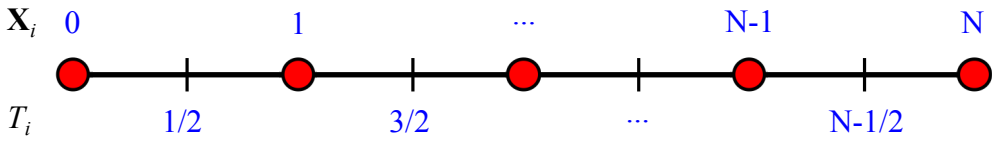

Figure 4.1: Staggered grid arrangement for discretising the Euler-Bernoulli beam equations.

$$D_s\mathbf{X}^{n+1} \cdot D_s\mathbf{X}^{n+1} = 1 \tag{4.6}$$

where n represents the solution at the current time step, and D_s and D_{ssss} represent the first and fourth-order finite difference approximations with respect to the parametric coordinate.

Introducing an extrapolated solution, $\mathbf{X}^* = 2\mathbf{X}^n - \mathbf{X}^{n-1}$, and substituting $\mathbf{X} = (x, y)$, $Fr^{\mathbf{g}}/_{|\mathbf{g}|} = (Fr_x, Fr_y)$, $\mathbf{F}^n = (-F_x^n, -F_y^n)$ and $\beta = \Delta s^2/\Delta t^2$, Equations 4.5 and 4.6 can be multiplied by Δs^2 and rearranged to give:

$$\Delta s^2 \left[D_s(T^{n+1}D_s x^{n+1}) - K_B D_{ssss} x^* + Fr_x + F_x^n\right] + \beta x^* - \beta x^{n+1} = 0 \tag{4.7}$$

$$\Delta s^2 \left[D_s(T^{n+1}D_s y^{n+1}) - K_B D_{ssss} y^* + Fr_y + F_y^n\right] + \beta y^* - \beta y^{n+1} = 0 \tag{4.8}$$

$$(\Delta s D_s x^{n+1})^2 + (\Delta s D_s y^{n+1})^2 - \Delta s^2 = 0 \tag{4.9}$$

These derivatives are then expanded and rearranged by grouping the unknown values. The remainder of this expansion is given in Appendix A. Finally, the nonlinear system is solved iteratively via Newton's multidimensional method to get the solution at the next time step.

4.1.3 Boundary Conditions

Appropriate boundary conditions are required to solve Equations 4.7–4.9. The standard beam boundary conditions are adopted in this work [14], which at the free end ($i = N$) are given by:

$$T = 0, \qquad \frac{\partial^2 \mathbf{X}}{\partial s^2} = (0, 0), \qquad \frac{\partial^3 \mathbf{X}}{\partial s^3} = (0, 0) \tag{4.10}$$

In this work, two separate boundary conditions are adopted at the fixed end: simply supported and clamped. For a simply supported beam, the conditions at the fixed end ($i = 0$) are:

$$\mathbf{X} = (X_0, Y_0), \qquad \frac{\partial^2 \mathbf{X}}{\partial s^2} = (0, 0) \tag{4.11}$$

In the case of a clamped beam, the conditions at the fixed end ($i = 0$) are:

$$\mathbf{X} = (X_0, Y_0), \qquad \frac{\partial \mathbf{X}}{\partial s} = (\cos(\theta), \sin(\theta)) \tag{4.12}$$

where θ is the clamp angle, with respect to the horizontal axis. The discretisation of these boundary conditions is also given in Appendix A.

4.1.4 Limitations

The method just described is used to calculate the structural dynamics in Chapters 7 and 8. During the course of these investigations, the Euler-Bernoulli beam solver was found to have some important limitations and nuances. Recalling the key dimensionless parameters governing the Euler-Bernoulli beam model:

$$\mathcal{M}^* = \frac{\Delta\rho}{\rho_f L}, \qquad \mathcal{K}^* = \frac{EI}{\Delta\rho u_\infty^2 L^2} \tag{2.5}$$

where $\Delta\rho = (\rho_s - \rho_f)h$. It can be seen that this approach breaks down when the fluid and structural densities are equal. In reality, tests revealed stability problems as the structural density approached the same value as the fluid density. While it should be noted that this is not a limitation of Euler-Bernoulli models in general – but is instead due to the $\Delta\rho$ term in the structural equation of motion – it does present a serious limitation of the particular approach adopted in the present work.

Another issue arose from the coupling between the dimensionless mass ratio and bending stiffness terms. In the work described in Chapter 8, the mass ratio was varied while the bending stiffness was held constant. However, examining Equation 2.5 shows that these two terms are coupled through the structural density. This means that varying the mass ratio while keeping the (dimensionless) bending stiffness fixed is equivalent to varying both the structural density and the Young's modulus at the same time, so that:

$$E \propto \rho_s - \rho_f \tag{4.13}$$

While there is not necessarily anything wrong with this, it can cause some confusion when trying to separate the effects that the different material properties have on the coupled dynamics. This helps to explain some of the findings given in Chapter 8, such as why smaller deflections were noted at larger mass ratios.

4.2 Finite Element Method

In response to the limitations just discussed regarding the Euler-Bernoulli beam solver, a finite element approach for slender structures was developed and implemented within the present model. As well as handling cases where the fluid and structural densities are the same, the FEM approach clearly separates the effects of different material properties. The present model is based on the corotational formulation, which permits large rotations of the structure [148]. Furthermore, to allow geometric nonlinearities due to large deflections, a Newton-Raphson approach is incorporated. To advance in time, the second-order implicit Newmark scheme is implemented [149]. Since the FEM is a well-established method with

decades of applications, only a brief overview of the present approach is given below.

4.2.1 Corotational Formulation

The basic idea of the FEM is to discretise the structure into an assemblage of elements. In the present work, two-noded Euler-Bernoulli beam elements are adopted. Each node contains three degrees of freedom (two translational and one rotational) which describe the displacement of the nodes. Adjacent elements are connected via the nodes they share. Similar to the Euler-Bernoulli beam solver described in Section 4.1, the assumption of rigid cross sections is made, so that the present model is only applicable to slender structures.

The present model adopts the corotational formulation, which permits large rotations of the structure. The main idea of the corotational formulation is to decompose the motion of each element within the assemblage into a rigid body motion and a purely deformational one [148]. This is achieved by adopting two

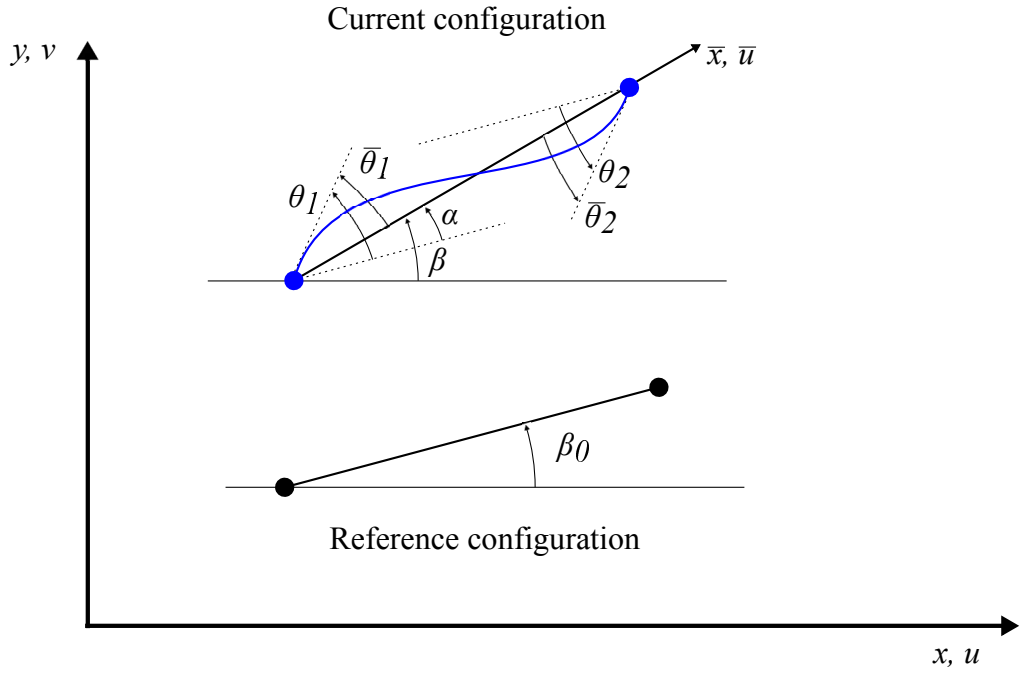

Figure 4.2: Corotational beam formulation. The motion of the beam is split into a rigid body transformation and a purely deformational motion. The rigid body motion is measured with respect to a reference (initial) configuration whereas the deformation is measured with respect to the corotated (local) frame.

reference configurations: the initial (or reference) frame, and the corotated (or local) frame. The corotated frame is element-specific and moves with the element, so that its longitudinal axis passes through the nodes of the element [150]. A schematic of these frames is given in Figure 4.2.

The rigid body motion is measured with respect to the initial configuration and can be arbitrarily large. On the other hand, the elemental deformation is measured with respect to the corotated frame and is assumed to be small within this local frame. In this regard, two sets of quantities are defined: the local quantities which are measured with respect to the corotated frame, and the global quantities which are measured with respect to the global coordinate system. In the formulation of the system matrices which appear in the governing equations, the elemental matrices are first calculated in the corotated frame before being transformed to the global coordinate frame, and then finally assembled into the global system matrices.

The kinematics of the element are described by the nodal displacements, which in global coordinates are given by:

$$\mathbf{U}_e^T = [u_1, v_1, \theta_1, u_2, v_2, \theta_2] \tag{4.14}$$

where the subscripts indicate the specific node within the element, e denotes the elemental vector, and the superscript T denotes the transpose. In local coordinates (denoted by the overbar), the elemental displacements are given by:

$$\bar{\mathbf{U}}_e^T = \left[\bar{u}_1, \bar{v}_1, \bar{\theta}_1, \bar{u}_2, \bar{v}_2, \bar{\theta}_2\right] = (\mathbf{T}\mathbf{U}_e)^T \tag{4.15}$$

Note that in the corotated (local) frame, $\bar{u}_1$, $\bar{v}_1$, and $\bar{v}_2$ are always zero. The

transformation matrix, $\mathbf{T}$, is given by:

$$\mathbf{T} = \begin{bmatrix} \cos(\beta) & \sin(\beta) & 0 & 0 & 0 & 0 \\ -\sin(\beta) & \cos(\beta) & 0 & 0 & 0 & 0 \\ 0 & 0 & 1 & 0 & 0 & 0 \\ 0 & 0 & 0 & \cos(\beta) & \sin(\beta) & 0 \\ 0 & 0 & 0 & -\sin(\beta) & \cos(\beta) & 0 \\ 0 & 0 & 0 & 0 & 0 & 1 \end{bmatrix} \tag{4.16}$$

where β is the angle between the corotated frame and the global coordinate frame.

The displacements along the beam can be obtained by the shape (interpolation) functions, which are specific for each type of element. For the standard Hermitian beam adopted in this work, the shape functions are given by:

$$\mathbf{N} = \left[\begin{array}{c|c|c|c} \frac{1}{2}(1-\xi) & \frac{1}{4}(1-\xi)^2(2+\xi) & \frac{L_e}{8}(1-\xi)^2(\xi+1) & \frac{1}{2}(1+\xi) \\ & & & \frac{1}{4}(1+\xi)^2(2-\xi) \quad \frac{L_e}{8}(1+\xi)^2(\xi-1) \end{array} \right] \tag{4.17}$$

where L_e is the length of the element and ξ is a local coordinate within the element ranging from -1 to 1 at nodes 1 and 2, respectively. Each element of the shape function vector corresponds to the local degrees of freedom (Equation 4.15). The displacement of an arbitrary point located at ξ along the beam is then given by:

$$\bar{u}(\xi) = \mathbf{N}_u \bar{\mathbf{U}}_e \tag{4.18}$$

$$\bar{v}(\xi) = \mathbf{N}_{v,\theta} \bar{\mathbf{U}}_e \tag{4.19}$$

4.2.2 Governing Equations

Following Bathe [149], the equilibrium conditions governing the nonlinear finite element solution are:

$$\mathbf{R} - \mathbf{F} = \mathbf{0} \tag{4.20}$$

where $\mathbf{R}$ are the externally applied loads and $\mathbf{F}$ are the internal forces within the structure, which correspond to the elemental stresses. In dynamic analyses, such as that described in this work, $\mathbf{R}$ also includes the inertial forces. Equation 4.20 must be satisfied throughout the whole analysis; therefore, the solution at the next time step is given by:

$$\mathbf{M}\ddot{\mathbf{U}}^{n+1} = \mathbf{R}^{n+1} - \mathbf{F}^{n+1} \tag{4.21}$$

where $\mathbf{M}$ is the global mass matrix, $\ddot{\mathbf{U}}$ are the nodal accelerations, and the superscript $n+1$ denotes the solution at the next time step. Note that here the inertial forces have been separated from the external load forces.

In the case of a nonlinear analysis, the internal forces, $\mathbf{F}$, depend on the nodal displacements (which are unknown). Therefore, Equation 4.21 is solved in an incremental fashion via Newton-Raphson iterations within the time step. Given that the solution at the current iteration, k, is known, the internal forces at the next iteration are given by:

$$\mathbf{F}^{n+1,k+1} = \mathbf{F}^{n+1,k} + \Delta\mathbf{F}^{k} \tag{4.22}$$

where $\Delta\mathbf{F}^k$ is the increment in internal forces due to the increment in nodal displacements from iteration k to $k+1$. This increment can be approximated by introducing a tangent stiffness matrix, which is the derivative of the internal forces with respect to the nodal displacements and includes the material stiffness

and the (nonlinear) geometric stiffness:

$$\Delta \mathbf{F}^k = \mathbf{K}^{n+1,k} \Delta \mathbf{U}^k \tag{4.23}$$

where $\Delta \mathbf{U}^k$ is the incremental nodal displacements. The governing equations for the solution at the next iteration are then given by:

$$\mathbf{M}\ddot{\mathbf{U}}^{n+1,k} + \mathbf{K}^{n+1,k} \Delta \mathbf{U}^k = \mathbf{R}^{n+1} - \mathbf{F}^{n+1,k} \tag{4.24}$$

$$\mathbf{U}^{n+1,k+1} = \mathbf{U}^{n+1,k} + \Delta \mathbf{U}^k \tag{4.25}$$

After each iteration, the tangent stiffness matrix and internal forces are updated according to the newly calculated displacements.

The elemental matrices used to build the system matrices are given by [149]:

$$\bar{\mathbf{M}}_e = \int_V \rho_s \mathbf{N}^T \mathbf{N} \, \mathrm{d}V \tag{4.26}$$

$$\bar{\mathbf{K}}_e = \int_V \mathbf{B}_L^T \mathbf{C} \mathbf{B}_L \, \mathrm{d}V + \int_V \mathbf{B}_{NL}^T \boldsymbol{\sigma} \mathbf{B}_{NL} \, \mathrm{d}V \tag{4.27}$$

$$\bar{\mathbf{R}}_e = \int_V \mathbf{N}^T \mathbf{f}_B \, \mathrm{d}V + \int_S \mathbf{N}^T \mathbf{f}_S \, \mathrm{d}S \tag{4.28}$$

$$\bar{\mathbf{F}}_e = \int_V \mathbf{B}_L^T \boldsymbol{\sigma} \, \mathrm{d}V \tag{4.29}$$

where $\mathbf{B}_L$ and $\mathbf{B}_{NL}$ are the linear and nonlinear strain-displacement matrices, $\mathbf{C}$ is the stress-strain material property matrix, $\boldsymbol{\sigma}$ are the stresses, and $\mathbf{f}_B$ and $\mathbf{f}_S$ are the body and surface loads. The integral is taken over the elemental volume V. Since the cross sectional area of the beam is constant, these integrations can be solved analytically. The elemental matrices are first calculated in the corotated

frame and then must be translated to the global frame before assembling into the system matrices.

4.2.3 Newmark Time Integration

The Newmark time integration scheme rests on the following assumptions:

$$\dot{\mathbf{U}}^{n+1} = \dot{\mathbf{U}}^{n} + \left[(1-\delta)\ddot{\mathbf{U}}^{n} + \delta\ddot{\mathbf{U}}^{n+1}\right]\Delta t \tag{4.30}$$

$$\mathbf{U}^{n+1} = \mathbf{U}^{n} + \dot{\mathbf{U}}^{n}\Delta t + \left[(\frac{1}{2}-\alpha)\ddot{\mathbf{U}}^{n} + \alpha\ddot{\mathbf{U}}^{n+1}\right]\Delta t^2 \tag{4.31}$$

where δ and α are parameters that can be tuned to obtain specific accuracy and stability properties. Setting $\delta = 1/2$ and $\alpha = 1/4$ gives the constant-average-acceleration (trapezoidal) method, which is second-order accurate and unconditionally stable [151–153]. This set of values is adopted in the present work. Substituting Equations 4.30 and 4.31 into Equation 4.24 and rearranging gives:

$$\hat{\mathbf{K}}^{n+1,k}\Delta\mathbf{U}^{k} = \mathbf{R}^{n+1} - \mathbf{F}^{n+1,k} - \hat{\mathbf{M}}^{n+1,k} \tag{4.32}$$

where the effective stiffness and mass matrices are given by:

$$\hat{\mathbf{K}}^{n+1,k} = \mathbf{K}^{n+1,k} + \frac{4}{\Delta t^2}\mathbf{M} \tag{4.33}$$

$$\hat{\mathbf{M}}^{n+1,k} = \mathbf{M}\left(\frac{4}{\Delta t^2}(\mathbf{U}^{n+1,k} - \mathbf{U}^{n}) - \frac{4}{\Delta t}\dot{\mathbf{U}}^{n} - \ddot{\mathbf{U}}^{n}\right) \tag{4.34}$$

Once the incremental displacements, $\Delta\mathbf{U}^{k}$, are found, the displacement at the next iteration is calculated via Equation 4.25. The iteration scheme is converged once the incremental displacements become sufficiently small, after which the

solution is advanced to the next time step.

4.3 Summary

This chapter has described the techniques adopted in the present work to solve the structural dynamics of slender structures. The first part of the chapter describes an approach based on a direct discretisation of the Euler-Bernoulli beam equation, which was provided within an already existing model. In response to certain limitations found with this approach, the second part of this chapter describes an FEM-based model for slender structures. Now that both field solvers (fluid and structure) have been introduced, the following chapter describes the coupling of these solvers.

Chapter 5

Fluid-Structure Coupling

The previous two chapters described the two separate field solvers (fluid and structure). This chapter now describes how these solvers are coupled in order to enable fluid-structure interaction simulations. The first part of this chapter details the approach used to couple the moving boundary of the structure to the fixed fluid mesh (immersed boundary method). The second part describes the data mapping approach between the non-conforming immersed boundary and finite element grids. Finally, the third part of this chapter describes the algorithm that ensures the interface conditions are properly satisfied.

5.1 Immersed Boundary Method

One of the major difficulties in CFD is grid generation. This is further compounded when complex geometries or moving boundaries are considered. The immersed boundary method (IBM) attempts to alleviate these difficulties by decoupling the solution of the fluid equations from the boundary representation. Initially developed by Peskin to model flow through the heart [154, 155], the IBM mimics the effect of the boundary by introducing a source term (force) in the governing fluid equations. Through this set of forces, the fluid responds to the imposed kinematic constraint, which can represent a range of boundary conditions (e.g. no-slip, slip, stationary, moving, flexible). In this work, it is used to impose the no-slip condition on both rigid and deforming boundaries.

The IBM uses two separate and independent grids to represent the fluid and the boundary (see Figure 5.1), each with their own reference frames. The fluid is defined on an Eulerian grid, which is Cartesian and fixed in space. The boundary, on the other hand, is represented by a Lagrangian curvilinear grid, which is made up of discrete markers and is free to move over the top of the fluid mesh. Since the grids are independent, the principal challenge in the IBM is to ensure accurate and efficient communication of data from one to the other. Furthermore, since they are non-conforming, this data transfer requires specialised interpolation/spreading operators. The required boundary conditions must then be reformulated into a distribution of forces which act to impose the desired kinematics. Because the boundary is defined in the Lagrangian frame, the force must be calculated in this frame and transferred back to the Eulerian grid. Once the force has been transferred back, the fluid equations are solved throughout the whole domain, without any specialised boundary treatment at the interface (other than the imposed force field).

One of the key advantages of the IBM is that the solution of the fluid equations is decoupled from the boundary representation. This greatly simplifies mesh generation [156] and allows arbitrarily complex geometries undergoing large deformations to be handled just as easily as simple rigid geometries, without affecting the underlying fluid mesh. This eliminates problems with poor quality cells and expensive remeshing procedures, which typically plague body-conforming approaches. That being said, the IBM does have certain drawbacks. The interpolation/spreading operators smear the boundary over a finite region in space, which results in diffused interfaces and a loss of accuracy at the boundary. Moreover, body-conforming approaches allow a greater control over mesh resolution in the region of the boundary. This has important implications for efficiency, particularly for high Reynolds number flows, where properly resolving the boundary layer is crucial [156]. Furthermore, some variants of the IBM, such as explicit approaches,

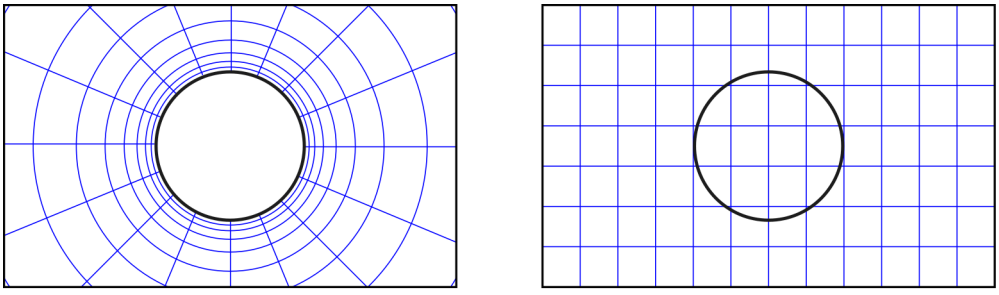

Figure 5.1: Conforming (left) and non-conforming (right) grids.

do not strictly enforce the no-slip condition [157].

Over the last few decades, many extensions to the IBM have been proposed. One of the greatest sources of variation among these extensions is the force calculation. Peskin's original method employs Lagrangian nodes linked by springs, which when displaced by the fluid generate a restoring force via Hooke's law [118]. However, this approach introduces free parameters (spring constants) which must be tuned to a particular problem. Furthermore, large stiffness values are required to model rigid boundaries. This leads to stability problems, which can only be overcome by reducing the time step size. An alternative approach is the direct forcing method [158–161]. This approach calculates the force directly from the fluid, usually through a predicted velocity field. The direct forcing approach works well for both rigid and deforming boundaries and does not require any additional free parameters. The following sections describe the approach adopted in this work: an implicit direct forcing method [110, 162, 163]. This particular approach enforces the no-slip condition by implicitly calculating a scaling factor for the force-spreading step.

5.1.1 Interpolation & Spreading Operators

The IBM uses two coordinate systems: an Eulerian system for the fluid and a Lagrangian system for the boundary. Let lower and upper-case notation define values in the Eulerian and Lagrangian frames, respectively, so that $\mathbf{x} = (x, y, z)$ defines the position in the Eulerian domain Ω, and $\mathbf{X} = (q, r, s)$ defines the position in the Lagrangian domain Γ. The interpolation and spreading operators are defined as:

$$\Phi(\mathbf{X}) = \mathcal{I}[\phi(\mathbf{x})] = \int_{\Omega} \phi(\mathbf{x})\ \delta(\mathbf{x} - \mathbf{X})\, \mathrm{d}\mathbf{x} \tag{5.1}$$

$$\phi(\mathbf{x}) = \mathcal{S}[\Phi(\mathbf{X})] = \int_{\Gamma} \Phi(\mathbf{X})\ \delta(\mathbf{x} - \mathbf{X})\, \mathrm{d}\mathbf{X} \tag{5.2}$$

where ϕ is a quantity defined in the Eulerian frame, and Φ is the same quantity defined in the Lagrangian frame. The Dirac Delta function, δ, is used to weight

the contributions and will be discussed in the next section.

5.1.2 Discretisation

Since both the fluid and boundary are defined in discrete domains, the continuous forms of the interpolation/spreading steps must be replaced with their discrete counterparts. These are given by weighted summations:

$$\Phi(\mathbf{X}) = \mathcal{I}[\phi(\mathbf{x})] = \sum_{\Omega_s} \phi(\mathbf{x})\ \tilde{\delta}(\mathbf{x} - \mathbf{X})\Delta x \Delta y \Delta z \tag{5.3}$$

$$\phi(\mathbf{x}) = \mathcal{S}[\Phi(\mathbf{X})] = \sum_{\Gamma_s} \Phi(\mathbf{X})\ \tilde{\delta}(\mathbf{x} - \mathbf{X})\epsilon \Delta q \Delta r \Delta s \tag{5.4}$$

The weighting function, $\tilde{\delta}$, is the discrete form the Dirac Delta function and defines the support region over which the summation is taken (see Figure 5.2). Note that in nondimensional lattice units $\Delta x = \Delta y = \Delta z = 1$. Furthermore, for the 1D line boundaries used in the present work $\Delta q = \Delta r = 1$. The scaling factor, ϵ, is particular to the implicit approach adopted in this work and will be discussed later in this section.

The choice of the discrete Dirac Delta function is crucial, since it governs the communication between the Eulerian and Lagrangian grids, and must be the same for both the interpolation and spreading steps [164]. Various sets of constraints have been proposed which define rules for generating new forms of this operator [164, 165]. Although there is some variation in these constraints, in general the discrete Dirac Delta function must preserve linear and angular momentum in both reference frames. Furthermore, it should be continuous and suppress the underlying lattice as much as possible, while at the same time not substantially smearing the interface. Moreover, for economical reasons it should have a compact support (act over a finite region in space). Several variants of the discrete Dirac Delta function have been proposed, some of which are shown in Figure 5.3. The form adopted in this work is the three-point stencil proposed by Roma et al.

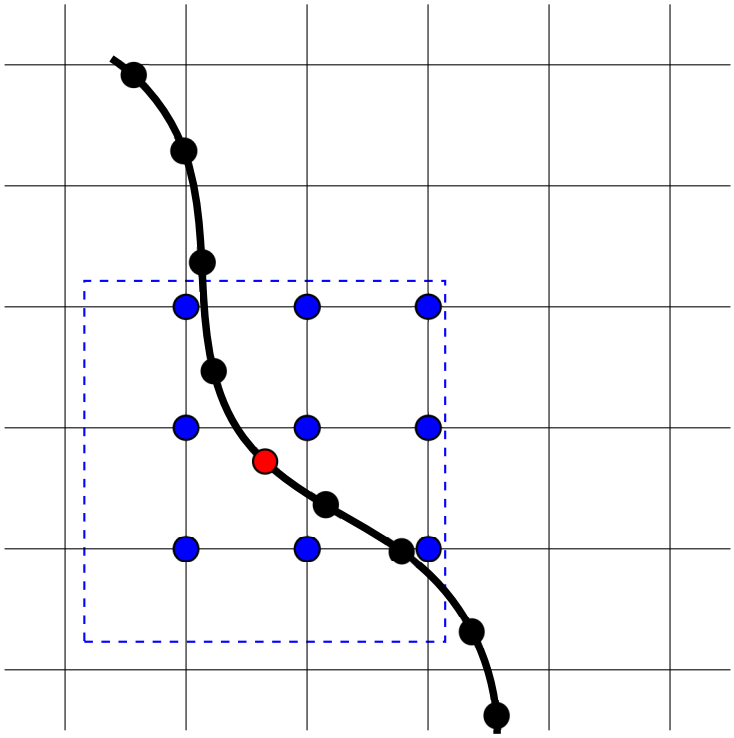

Figure 5.2: Support stencil for IBM marker (red). Dashed blue line is the support stencil over which the interpolation/spreading steps are performed. Blue lattice sites are the support points that exist within the support stencil and are directly involved in the communication.

[165]:

$$\tilde{\delta}(r) = \begin{cases} \frac{1}{3}\left(1 + \sqrt{-3r^2 + 1}\right) & |r| \leq 0.5 \\ \frac{1}{6}\left(5 - 3|r| - \sqrt{-3(1 - |r|)^2 + 1}\right) & 0.5 \leq |r| \leq 1.5 \\ 0 & \text{otherwise} \end{cases} \tag{5.5}$$

The three-point stencil was chosen for the present work, since it is more efficient than the popular four-point version (especially in 3D) and recent studies comparing the resulting hydrodynamics show little difference between the two [166, 167]. Equation 5.5 provides the one-dimensional form of the discrete Dirac Delta function; in 2D, it can be factored to give (Figure 5.3):

$$\tilde{\delta}(\mathbf{x} - \mathbf{X}) = \tilde{\delta}(x - X)\ \tilde{\delta}(y - Y) \tag{5.6}$$

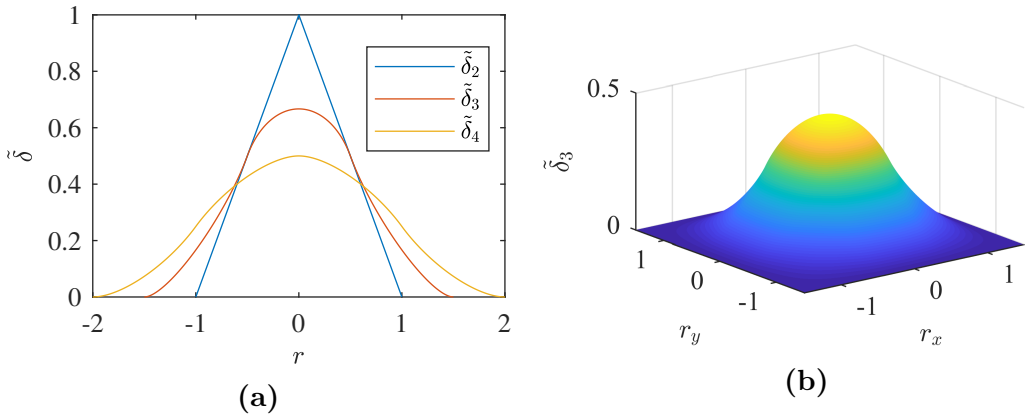

Figure 5.3: Popular one-dimensional discrete Dirac Delta functions (a) and the 2D version of the three-point stencil (b).

Due to the finite size of the support stencil, some of the support regions for different Lagrange markers will overlap each other. This leads to some lattice sites acting as support points for multiple Lagrange markers, which means during a spreading communication they will receive multiple contributions from different Lagrange markers. In practise, this means the no-slip condition will not be strictly satisfied. One approach to guarantee a consistent solution is to calculate the contributions from each Lagrange marker implicitly so that they satisfy a partition of unity [110]. Following Pinelli et al. [162], a scaling factor, ϵ, is computed for each IBM marker which ensures reciprocity of the interpolation and spreading steps. Consider any smooth function, $\phi(\mathbf{x})$, which can be interpolated onto the Lagrangian grid:

$$\Phi(\mathbf{X}) = \mathcal{I}[\phi(\mathbf{x})] \tag{5.7}$$

Replacing $\phi(\mathbf{x})$ with that obtained from the complimentary spreading step (including the scaling factor) and expanding the operators gives:

$$\Phi(\mathbf{X}) = \sum_{\Omega} \tilde{\delta}(\mathbf{x} - \mathbf{X}) \Delta x \Delta y \Delta z \sum_{\Gamma} \Phi(\mathbf{X})\ \tilde{\delta}(\mathbf{x} - \mathbf{X}) \epsilon \Delta q \Delta r \Delta s \tag{5.8}$$

After rearranging, one obtains:

$$\Phi(\mathbf{X}_n) = \sum_{m=1}^{N} A_{n,m}\epsilon_m\Phi(\mathbf{X}_m) \tag{5.9}$$

where:

$$A_{n,m} = \Delta s \sum_{\Omega} \tilde{\delta}(\mathbf{x} - \mathbf{X}_n)\ \tilde{\delta}(\mathbf{x} - \mathbf{X}_m) \tag{5.10}$$

In matrix notation, the resulting system is:

$$\mathbf{A}\boldsymbol{\epsilon} = \mathbf{1} \tag{5.11}$$

This approach is very similar to the implicit method proposed by Wu and Shu [157]. The only difference is that here the right-hand-side of Equation 5.11 is invariant across all time steps, since it is unity. For rigid problems, where the **A** matrix is constant, this means the scaling factors only need to be calculated once for the whole simulation. In Wu and Shu's approach, which solves for a corrective velocity, the right-hand-side of the equivalent system is a function of the predicted velocity field. Although the **A** matrix can be factored just once at the start of the simulation, the resulting system must be solved at every time step (e.g. via back-substitution). This minor difference is only relevant for rigid problems; for problems involving flexible bodies the **A** matrix is not constant and must be factored every time step in both approaches.

5.1.3 Coupling with LBM

The coupled LBM-IBM algorithm is given by Li et al. [163] and makes use of the fact that the LBM forcing scheme [130] decomposes the velocity into a predicted and a force-corrected term [157]. Since the IBM force enters the fluid equations as a correction to the velocity, the fluid equations only need to be solved once per

time step. Compared to the approach given by Favier et al. [110], this results in significant computational savings.

Recalling from Section 3.5, the macroscopic fluid velocity is given by:

$$\rho\mathbf{u}(\mathbf{x},t) = \sum_i \mathbf{c}_i f_i(\mathbf{x},t) + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x},t) \tag{3.17}$$

After one pass of the fluid equations, the distribution functions, f_i, and the fluid density, ρ, are known (since the fluid density is just the summation of all components of the distribution function). Therefore, the two terms on the right-hand-side of Equation 3.17 can be decomposed into a predicted velocity field and a force-correction term:

$$\rho\mathbf{u}(\mathbf{x},t) = \rho\mathbf{u}^*(\mathbf{x},t) + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x},t) \tag{5.12}$$

Since the fluid velocity at the boundary must equal the velocity of the boundary, $\mathbf{U}^b = \mathcal{I}[\mathbf{u}]$, then converting Equation 5.12 into the Lagrangian frame gives:

$$\mathcal{I}[\rho(\mathbf{x},t)]\mathbf{U}^b(\mathbf{X},t) = \mathcal{I}[\rho\mathbf{u}^*(\mathbf{x},t)] + \frac{\Delta t}{2}\mathbf{F}(\mathbf{X},t) \tag{5.13}$$

Since the velocity of the boundary, $\mathbf{U}^b$, is known, Equation 5.13 can be rearranged and solved for the corrective force density, $\mathbf{F}$. The force is then transferred back to the Eulerian frame via the spreading operator:

$$\mathbf{f}(\mathbf{x},t) = \mathcal{S}[\mathbf{F}(\mathbf{X},t)] \tag{5.14}$$

Finally, the velocity field is updated by adding the corrective force to the predicted velocity field (Equations 5.12).

The steps outlined above are sufficient to model rigid bodies using the IBM. Since some of the validation cases presented in Chapter 6 involve rigid boundaries, the

rigid body LBM-IBM algorithm is given below:

1. perform LBM step to get the density and predicted velocity field
2. interpolate the predicted momentum and density field onto the Lagrangian markers via the interpolation operator
3. calculate the corrective force in the Lagrangian frame via Equation 5.13
4. spread the corrective force back to the Eulerian grid via the spreading operator
5. correct the predicted velocity field with the force density via Equation 5.12

The algorithm for flexible bodies requires further discussion and will be given in Section 5.3.3.

5.1.4 Modelling Considerations

There is some debate about the accuracy of the IBM. Formally, it is first-order accurate [164]. This is due to the discrete form of the Dirac Delta function which smears the boundary over a finite region. However, many studies have reported close to second-order accuracy. The discrete Dirac Delta also causes problems when defining the hydrodynamic thickness of the boundary. Due to the finite support region, the boundary the fluid sees is not the true boundary prescribed by the Lagrange markers. To account for this, Feng and Michaelides [168] suggested augmenting the true radius of the boundary and proposed a scaling rule for 3D spheres. However, it is unclear how such an approach could be extended to general shapes.

When generating the Lagrangian grid, the spacing between the markers is dictated by the fluid resolution. If the markers are too far apart the boundary will become 'leaky' and fluid will be able to penetrate it. However, markers that are too close to each other adds unnecessary computational cost, without any benefit of increased accuracy [162]. Generally, a marker spacing approximately equal to the fluid mesh

spacing is suitable, so that $\Delta s \approx \Delta x$. When a Lagrange marker approaches the edge of the fluid domain some of the support points may exist outside of the domain. In the present work, the contributions to/from these support points are neglected. Since the fluid properties in the region of the boundary are prescribed by the boundary conditions, it is unlikely that such an approximation will have a significant impact. It may be possible to generate one-sided discrete Dirac Delta functions which can accommodate this scenario; however, this is not the topic of the present work.

5.2 Data Mapping

Usually the grid resolution requirements for the fluid are higher than those for the structure. Since the IBM resolution is dictated by the fluid mesh, using a one-to-one mapping between the IBM and structural meshes would be inefficient. This was the case for the Euler-Bernoulli beam solver, which required that the nodes of the structural mesh coincided with the IBM nodes (and thus scaled with the fluid resolution). Therefore, an increase in fluid resolution necessitated a (potentially redundant) increase in structural resolution. To overcome this, a data mapping approach was adopted for the FEM solver which allowed the transfer of displacements, velocities, and forces between the non-conforming IBM and structural grids. Since this data mapping approach was developed after the Euler-Bernoulli beam solver was replaced with the FEM solver, the following discussion is only applicable to the FEM solver.

5.2.1 Force Transfer

Figure 5.4 shows the non-conforming IBM and FEM grids. The force on the FEM element is made up of contributions from multiple IBM markers. Specifically, markers a, b, and c. Since the IBM force is a force density, the force from each IBM marker is distributed over a region which spans Δs, where s is a parametric coordinate along the structure. Furthermore, this distributed force is centred over each IBM marker. Recalling from Section 4.2, the elemental load vector

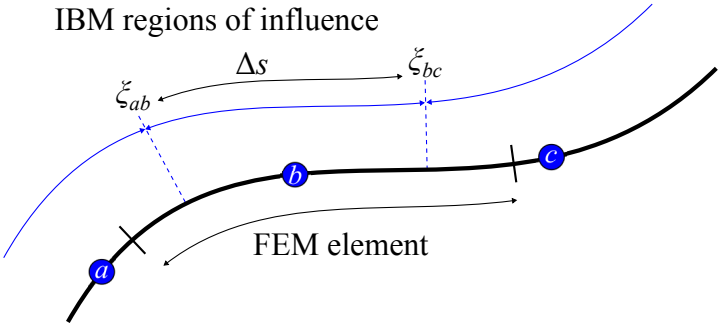

Figure 5.4: Non-conforming FEM and IBM grids. Black line indicates assemblage of FEM elements. Blue markers indicate the IBM nodes and the dashed blue line indicates the region over which the force from node b is applied to the FEM element.

(neglecting body forces) is given by:

$$\bar{\mathbf{R}}_e = \int_S \mathbf{N}^T \mathbf{f}_S \, \mathrm{d}S = \int_0^{L_e} \mathbf{N}^T \bar{\mathbf{F}} \, \mathrm{d}\bar{x} \tag{5.15}$$

where L_e is the length of the element, and $\bar{\mathbf{F}}$ is the IBM force density (after scaling with ϵ) transformed into the element-specific corotated frame.

Since the shape functions, $\mathbf{N}$, are described by the local coordinates, ξ, the integral is recast:

$$\bar{\mathbf{R}}_e = \int_{-1}^{1} \mathbf{N}^T \bar{\mathbf{F}} \frac{\mathrm{d}\bar{x}}{\mathrm{d}\xi} \, \mathrm{d}\xi \tag{5.16}$$

This integral can be decomposed into the sum of the contributions from each IBM marker that spans a portion of the FEM element. For the element shown in Figure 5.4, the contributions are:

$$\bar{\mathbf{R}}_e = \int_{-1}^{\xi_{ab}} \mathbf{N}^T \bar{\mathbf{F}}_a \frac{\mathrm{d}\bar{x}}{\mathrm{d}\xi} \, \mathrm{d}\xi + \int_{\xi_{ab}}^{\xi_{bc}} \mathbf{N}^T \bar{\mathbf{F}}_b \frac{\mathrm{d}\bar{x}}{\mathrm{d}\xi} \, \mathrm{d}\xi + \int_{\xi_{bc}}^{1} \mathbf{N}^T \bar{\mathbf{F}}_c \frac{\mathrm{d}\bar{x}}{\mathrm{d}\xi} \, \mathrm{d}\xi \tag{5.17}$$

For the Hermitian beam elements used in the present work, these integrals can be expanded analytically to obtain the load contributions from each IBM marker.

5.2.2 Position/Velocity Transfer

Mapping the FEM displacements and velocities to the IBM markers relies on the fact that the elemental shape functions describe the local kinematics within the element. Therefore, given the nodal displacements/velocities (obtained from the FEM solution), the position of an IBM marker located at an arbitrary location within the element, ξ, can be calculated via:

$$\bar{u}(\xi) = \mathbf{N}_u \bar{\mathbf{U}}_e \tag{4.18}$$

$$\bar{v}(\xi) = \mathbf{N}_{v,\theta} \bar{\mathbf{U}}_e \tag{4.19}$$

In addition to permitting non-conforming IBM and FEM grids, preliminary tests showed this mapping approach also possessed superior stability properties compared to a simple linear interpolation mapping procedure.

5.3 Coupling Scheme

In numerical modelling of fluid-structure interaction problems, as well as solving for the separate fields (fluid and structure), it is also necessary to adopt some form of coupling technique so that the fluid and structure can interact with each other. This coupling takes place at the interface and involves common boundary values, such as positions, velocities, and forces. In particular, the interface compatibility conditions must be met [169]:

$$\mathbf{u} = \frac{\mathrm{d}\mathbf{X}}{\mathrm{d}t} \tag{5.18}$$

$$\boldsymbol{\sigma}^f \cdot \mathbf{n} = \boldsymbol{\sigma}^s \cdot \mathbf{n} \tag{5.19}$$

where $\mathbf{u}$ is the fluid velocity, $\mathbf{X}$ is the position of the boundary, $\boldsymbol{\sigma}^f$ and $\boldsymbol{\sigma}^s$ are the

fluid and structural surface stresses, and **n** is the normal vector. Equation 5.18 is the kinematic condition, which states that the velocity of the fluid at the interface is equal to the velocity of the boundary (i.e. no-slip condition). Equation 5.19, on the other hand, is the dynamic interface condition and states that the tractions on the fluid and structure sides of the interface must be in equilibrium. Satisfying these interface conditions implies conservation of mass, momentum, and energy [170].

There are two main approaches to coupling FSI problems: monolithic and partitioned. The monolithic approach derives a coupled set of equations for the entire system (fluid and structure) and solves them simultaneously. The partitioned approach on the other hand treats the fluid and structure separately: both have their own separate solvers which are advanced in time independently and information between the two is exchanged only at specific points during the solution. Although the monolithic approach shows good stability and accuracy properties [171], for complex problems it can be difficult to formulate and expensive. Furthermore, since the fluid and structure are solved as one system, the choice of discretisation and numerical methods may be constrained (and therefore suboptimal). In the partitioned approach, the separate field solvers can be chosen completely independently. This means existing solvers, which have been specifically optimised for a particular problem, can be reused. Moreover, the modularity of this approach allows for better software maintainability and flexibility [172]. Since the fluid and structural dynamics are solved separately, partitioned approaches require coupling schemes which define the order of the calculations, what interface conditions are met, and how they are met. Some of the basic partitioned approaches do not strictly satisfy the interface conditions. This can lead to stability and accuracy issues, particularly for highly coupled problems with strong interactions. That being said, more advanced partitioned coupling algorithms are available which overcome these issues, at a price of increased cost and algorithmic complexity.

5.3.1 Implicit Gauss-Seidel Scheme

There are a variety of techniques available for coupling partitioned FSI schemes. The most basic of these is the explicit approach, commonly referred to as the Conventional Serial Staggered (CSS) scheme [172]. In this approach, the fluid and structural fields are solved sequentially within a single time step. Information (displacements, velocities, and forces) is then exchanged between the two before

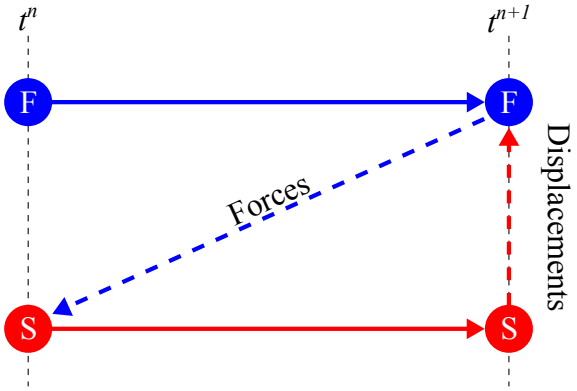

Figure 5.5: Schematic of Conventional Serial Staggered (CSS) scheme. Blue and red lines indicate fluid and structural steps. Solid lines indicate the solver step whereas the dashed lines indicate communication between the solvers.

beginning the next time step, as shown in Figure 5.5. The CSS scheme is simple, efficient, and easy to implement. However, it does not strictly satisfy the interface conditions. This is because the staggered solution of the fluid and structure introduces a lag between the two fields [171]. This leads to artificial energy generated at the interface, which in some cases can result in stability problems. This is especially true for problems involving very light and flexible structures, where this effect is termed the added-mass instability [173]. For compressible flows, reducing the time step size can help to stabilise the solution [174]. However, this can become prohibitively expensive and is therefore not always a practical solution. In addition to its stability problems, the CSS scheme is only first-order accurate in time, no matter the accuracy of the underlying field solvers [172]. Nevertheless, the CSS algorithm is simple and efficient and works well for problems involving heavy, stiff structures (e.g. aeroelastic applications).

While various techniques have been developed to improve the stability and accuracy properties of explicit partitioned schemes [172, 175–177], an alternative approach is to employ an implicit scheme. To satisfy the interface conditions, both field solvers require information about the next time step, which is unknown. Therefore, implicit approaches iterate over the field solvers within the time step until the interface conditions are met. Provided these iterations converge, implicit schemes show improved accuracy and stability over explicit approaches, even when the added mass is relatively large. A popular implicit scheme is the block Gauss-Seidel algorithm, which solves the fluid and structure sequentially, but in an iterative fashion within the time step [178]. The CSS scheme can actually be thought of as a single-iteration version of the Gauss-Seidel scheme. A schematic of the implicit

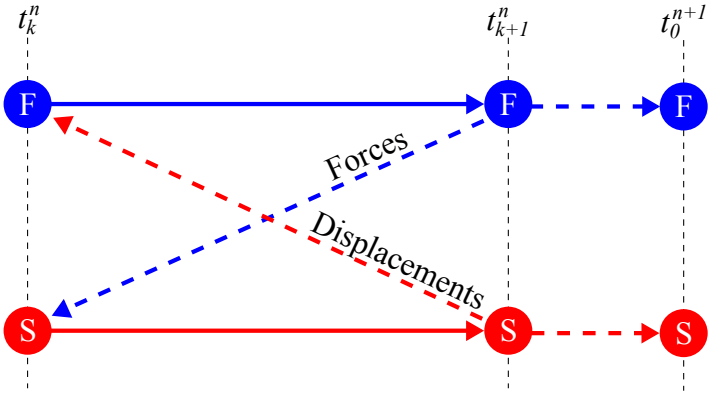

Figure 5.6: Schematic of Gauss-Seidel scheme. Blue and red lines indicate fluid and structural steps. Solid lines indicate the solver step whereas the dashed lines indicate communication between the solvers. Once the inner loop is converged both solvers are advanced to the next time step.

Gauss-Seidel algorithm is shown in Figure 5.6.

5.3.2 Aitken's Relaxation Method

Even with iterations, implicit schemes may still converge slowly, and in some cases may not converge at all. To accelerate convergence, typically the boundary displacement is relaxed by combining the current solution with that of the previous iteration:

$$\mathbf{U} = \omega\tilde{\mathbf{U}} + (1-\omega)\mathbf{U}^{k-1} \tag{5.20}$$

where $\mathbf{U}$ is the relaxed solution which is passed to the fluid solver, $\tilde{\mathbf{U}}$ is the displacement computed from the structural solver, and $\mathbf{U}^{k-1}$ is the solution from the previous iteration.

The rate of convergence and stability of the Gauss-Seidel scheme depends heavily on the relaxation factor. Too low and the iteration scheme will converge slowly; too high and it may become unstable. The optimum value for the relaxation factor is problem dependent and unknown a priori. Furthermore, the optimum value will change as the simulation progresses and the level of interaction between the fluid and structure changes. Therefore, to accelerate convergence, typically a dynamic

relaxation factor is used. A very popular approach for calculating the relaxation factor is Aitken's delta-squared method [179]. This technique dynamically adjusts the relaxation factor based on the solution at previous iterations.

$$\omega_k = -\omega^{k-1} \frac{(\mathbf{r}_{k-1})^T (\mathbf{r}_k - \mathbf{r}_{k-1})}{|\mathbf{r}_k - \mathbf{r}_{k-1}|^2} \tag{5.21}$$

where $\mathbf{r}_k$ is the residual vector, which is also used as a stopping criterion.

$$\mathbf{r}_k = \tilde{\mathbf{U}}_k - \mathbf{U}_{k-1} \tag{5.22}$$

From Equation 5.21, calculating the relaxation factor requires two previous iterations. Therefore, for the first iteration, either the value at the end of the last time step or a predetermined maximum value is assigned:

$$\omega_1^{n+1} = \min(\omega^n, \omega_{max}) \tag{5.23}$$

Aitken's method is one of the most popular approaches for accelerating implicit coupling schemes. While it is simple and easy to implement, it is also efficient, robust, and shows good convergence and stability properties.

Finally, to further accelerate the convergence of the Gauss-Seidel iterations, the present work employs a structural predictor at the start of the time step:

$$\mathbf{X}_0^{n+1} = \frac{5}{2}\mathbf{X}^n - 2\mathbf{X}^{n-1} + \frac{1}{2}\mathbf{X}^{n-2}, \qquad n \geq 3 \tag{5.24}$$

$$\mathbf{X}_0^{n+1} = 2\mathbf{X}^n - \mathbf{X}^{n-1}, \qquad n = 2 \tag{5.25}$$

$$\mathbf{X}_0^{n+1} = \mathbf{X}^n, \qquad n = 1 \tag{5.26}$$

5.3.3 Algorithm

The full implicit algorithm for flexible bodies, which was developed and implemented as part of the present work, is given below [178, 180].

1. perform the LBM step to get the predicted velocity field
2. perform the structural predictor step (Equations 5.24–5.26)
3. recompute the support points and scaling factors (ϵ) with the new boundary location
4. perform the IBM step (interpolation, force calculation, spreading)
5. perform the FEM step
6. calculate the residual via Equation 5.22
7. if converged, continue to Step 8; otherwise, relax the displacement of the boundary and go back to Step 3
8. correct the predicted velocity field with the IBM force calculated in Step 4
9. advance to the next time step and proceed to Step 1

5.4 Summary

This chapter has described the techniques used to couple the two field solvers. The first section of this chapter described the approach used to couple the fluid and structural meshes (immersed boundary method). The second section introduced the data mapping approach between the immersed boundary and finite element grids. Finally, the third section described the coupling algorithm that ensures the interface conditions are properly satisfied and a stable solution is obtained. The following chapter now discusses the original model and the developments that

were implemented as part of the present work, as well as providing a validation of the numerical model.

Chapter 6

Model Developments & Validation

Following from the previous chapters, which detailed the specific methods employed in the present model, the first part of this chapter summarises the main method developments that were implemented as part of the present work. By highlighting the limitations of various components of the original model and how they relate to the objectives of this work, the reasons for each of the developments, and their significance, are discussed. The second part of this chapter provides a validation for each component of the present approach, as well as the fully coupled model itself.

6.1 Model Developments

As mentioned in the preceding chapters, the early part of this work made use of an already existing numerical model for slender structures. This model incorporated the Euler-Bernoulli beam solver discussed in Section 4.1 and a different version of the immersed boundary algorithm to that described in Section 5.1. The first two test cases presented in this thesis, Chapters 7 and 8, made use of this model. However, several limitations of this approach were highlighted that prevented the application of this model to large arrays of slender structures with ranging material properties. Since this is the main objective of the present work, a significant amount of effort was devoted to developing and improving the original model. In

Table 6.1: Summary of model developments and their significance.

Original Model	Present Model	Benefit
Euler-Bernoulli solver	Finite element method	Robustness
CSS coupling scheme	Gauss-Seidel coupling scheme	Robustness & stability
Two-step IBM algorithm	One-step IBM algorithm	Efficiency
Serial	Parallel (OpenMP)	Speed

particular, efforts were made to improve the performance and robustness of the numerical approach. Table 6.1 highlights the limitations of the original model, the ensuing developments, and their significance. The remainder of this section now discusses these developments in detail.

6.1.1 Structural Model

As discussed in Section 4.1.4, the Euler-Bernoulli beam solver, which was part of the original model, had some important limitations that prevented its use over a wide range of structural densities. Specifically, this approach broke down when the density ratio (ρ_s/ρ_f) approached one. This limited the application of this method to problems where the structural density was larger than the fluid density. While this is true for many types of problems, there are also many cases where this is not true, such as in aquatic vegetation [181].

To overcome this problem, a nonlinear finite element approach for slender structures was developed and implemented in the present model. The FEM solver takes the actual structural density as an input, as opposed to the Euler-Bernoulli beam solver, which accounts for the extra mass due to the presence of the structure (which is zero for $\rho_s/\rho_f = 1$). Therefore, the FEM approach permits the solution of problems where the structural density is equal to (or lower than) the fluid density, as long as the added-mass instabilities are controlled via an appropriate coupling scheme (discussed next).

While not strictly a limitation, the formulation of the Euler-Bernoulli beam solver introduces a coupling of the structural material properties via the nondimensional parameters that appear in the governing equations. In the FEM, the effects of Young's modulus and structural density are clearly separated. This allows easier

insight into their physical effects. Furthermore, the original discretisation of the Euler-Bernoulli beam solver employed an explicit time-marching scheme. To ensure stability without severe time step restrictions, the current FEM approach adopts an implicit time integration scheme. While this is an advantage over the previous approach, it should be noted that the Euler-Bernoulli can also be reformulated with an implicit time integration scheme. However, in light of the other arguments, this was not explored.

6.1.2 FSI Coupling Scheme

The original model made use of the Conventional Serial Staggered (CSS) coupling scheme, which is described in Section 5.3.1. In this approach, the fluid solver is advanced in time first, then the updated velocity field is passed to the structural solver, which is then also advanced in time, before both solvers begin the next time step. With this approach, there is a lag between the fluid and structural solvers. For cases involving large density ratios, where the structural motion is dominated by its own inertia, this approach is usually acceptable. However, in cases where the density ratio approaches one, there is a strong coupling between the fluid and structural dynamics. In this case, the lag induced by the CSS scheme leads to stability problems, known as added-mass instabilities. Since the Euler-Bernoulli beam solver restricted the applicability of the original model to cases where the density ratio was larger than one, this limited the appearance of added-mass instabilities. However, in support of the FEM solver, which was developed to allow a range of density ratios, the FSI coupling scheme was also modified.

To overcome the stability problems associated with the CSS scheme, a block Gauss-Seidel implicit coupling scheme was implemented. This approach iterates over the fluid and structural solvers within the time step until the interface conditions are met, as described in Section 5.3.1. The implicit coupling scheme enhances the stability and accuracy properties of the present approach. Since efficiency is also of concern in the present work, to accelerate the convergence of these iterations, the structural displacements are relaxed via a dynamically adjusted relaxation factor. This ensures a rapid and robust convergence of the iteration scheme.

6.1.3 Immersed Boundary Algorithm

While the version of the IBM discussed in Section 5.1 only requires one pass of the fluid equations, the approach adopted in the original model required the full solution of the LBM equations twice per time step. The first step in the original implementation is to solve the LBM equations by neglecting the presence of the immersed boundary. This results in a predicted velocity field, $\mathbf{u}^*$. This predicted velocity field is then interpolated onto the Lagrangian markers via:

$$\mathbf{U}^*(\mathbf{X}, t) = \mathcal{I}[\mathbf{u}^*(\mathbf{x}, t)] \tag{6.1}$$

Using Newton's second law, the force is obtained via:

$$\mathbf{F}(\mathbf{X}, t) = \frac{\mathbf{U}^b(\mathbf{X}, t) - \mathbf{U}^*(\mathbf{X}, t)}{\Delta t} \tag{6.2}$$

where $\mathbf{U}^b$ is the velocity of the boundary. For rigid bodies with no-slip velocity $\mathbf{U}^b = 0$. This force is spread back to the Eulerian grid via the spreading operator:

$$\mathbf{f}(\mathbf{x}, t) = \mathcal{S}[\mathbf{F}(\mathbf{X}, t)] \tag{5.14}$$

Finally, the full fluid equations are solved once more, but this time with the immersed boundary force included.

Clearly, this approach is significantly more expensive than the approach described in Section 5.1, owing to the fact it computes the full fluid equations twice. However, recent studies have recognised that this approach is also unnecessary, since the force calculated from the IBM procedure appears in the fluid equations as a correction to the macroscopic velocity [157, 163]. Thus, adding the corrective velocity (obtained from the corrective force) directly to the predicted velocity field, as described in Section 5.1, negates the need for a second pass of the fluid equations. An added benefit of this approach is that, due to the finite width of the support stencil and the explicit nature of the LBM, this correction is only

required over the lattice sites which fall within the support region, leading to further savings in execution time.

6.1.4 Parallel Implementation

As mentioned in Section 3.8, one of the main benefits to adopting the LBM is its suitability to parallel implementation. However, the original model used during the early stages of the present work was only developed for serial execution and did not make use of this potential benefit. Therefore, in addition to the model developments described above, effort was made in developing the new model to integrate parallel capability via the OpenMP application programming interface. While this did lead to significant performance gains and permit the application of the new model to larger scale problems, this approach is neither novel nor the main focus of the present work. Therefore, it is only briefly discussed here. For a more detailed discussion on parallel implementations of the LBM, the reader is directed to one of the many examples in the literature [182–184].

To highlight the performance gains realised via adopting the OpenMP implementation, Table 6.2 compares the performance for both serial and parallel (24 cores) execution of a representative case from Chapter 9. The case involves a large array of 128 flaps in an open-channel flow. Both the speed – measured in millions of lattice updates per second (MLUPS) – and time of execution are tabulated. Instead of running the full serial simulation, which would take a prohibitively long time, the MLUPS value was measured in real-time until it reached a stable value (after approximately 72 hours). This value was then used to calculate a projected time of execution. As can be seen, the performance benefits of the parallel calculation are substantial and reduce the execution time from approximately one month

Table 6.2: Serial and parallel performance for a representative case involving a large array of 128 flaps in an open-channel flow. The projected serial execution time was calculated using the measured speed of execution after it reached a stable value.

Execution	**Speed (MLUPS)**	**Time (days)**
Serial	3.04	29.7 (projected)
Parallel (24 cores)	36.3	2.49

to a couple of days. It should be noted that while the performance gains are significant, some portions of the code are not fully optimised for parallel execution (e.g. output) and therefore the results do not show an ideal scaling.

6.2 Validation

Following the previous section, where the main method developments have been highlighted and discussed, the following section provides a validation of the numerical model. Each component of the model, the fluid solver and both structural solvers, as well as the fully coupled model are tested for a range of common benchmark cases. Some tests are provided as part of the results chapters, and so are not provided in this section. Table 6.3 summarises each case and where it can be found in the thesis.

6.2.1 Fluid Solver

First the fluid solver is tested on its own. Since the focus of this work is wall-mounted slender structures, a simple Poiseuille flow provides a relevant test that captures the essential features of the types of flows encountered in this work. The

Table 6.3: Summary of selected validation cases. EB corresponds to Euler-Bernoulli beam solver.

Component	Case	Section
Fluid solver	Poiseuille flow (LBM)	Section 6.2.1
	Womersley flow (LBM)	Section 7.4
Structural solver	Static deformation with tip moment (FEM)	Section 6.2.2
	Static deformation with tip load (FEM)	Section 6.2.2
	Flexible pendulum (EB & FEM)	Section 6.2.2
Fully coupled model	Rigid cylinder (LBM+IBM)	Section 6.2.3
	Flapping beam (LBM+IBM+FEM)	Section 6.2.3
	PELskin case (LBM+IBM+EB)	Section 8.5
	PELskin case (LBM+IBM+FEM)	Section 9.4

main objective of this test is to ensure the fluid solver recovers the correct velocity profile across the channel and exhibits the proper accuracy properties provided by the LBM theory, namely second-order convergence (with appropriate boundary conditions).

The idealised case considers the pressure-driven flow between two infinitely long parallel plates. This results in a 2D flow given by the classic parabolic profile:

$$u_x(y) = -\frac{\mathrm{d}p}{\mathrm{d}x}\frac{1}{2\rho\nu}y(H-y) \tag{6.3}$$

where x and y are the streamwise and transverse coordinates respectively, u_x is the streamwise velocity, $^{\mathrm{d}p}/_{\mathrm{d}x}$ is the pressure gradient along the channel, ρ is the fluid density, ν is the kinematic viscosity, and H is the height of the channel. The computed velocity profile is compared against the analytical solution and the relative error is measured as the absolute difference between the centreline velocities, which in the analytical case is given by:

$$u_x\left(\frac{H}{2}\right) = -\frac{\mathrm{d}p}{\mathrm{d}x}\frac{H^2}{8\rho\nu} \tag{6.4}$$

The flow is initially at rest. Once the simulation is started, a constant pressure gradient is applied which accelerates the flow until the viscous forces balance the pressure force and the flow reaches a steady state. Periodic boundary conditions are applied at the left and right boundaries while the regularised technique (second-order accurate) is used to enforce the no-slip condition at the top and bottom boundaries. The pressure gradient is imposed indirectly via a body force, using the forcing scheme discussed in Section 3.5. While this approach may not seem entirely physical, in an incompressible flow this is equivalent to setting a pressure gradient directly [100] and many studies have reported success with this approach for cases involving periodic domains [113, 185].

To identify the order of the numerical scheme, a range of grid resolutions are tested. As discussed in Section 3.8, the LBM is second-order accurate in space. However, this can only be realised if the compressible error term, $\mathcal{O}(Ma^2 \propto {}^{\Delta t^2}/_{\Delta x^2})$, also scales with the grid resolution. This leads to the so-called diffusive scaling,

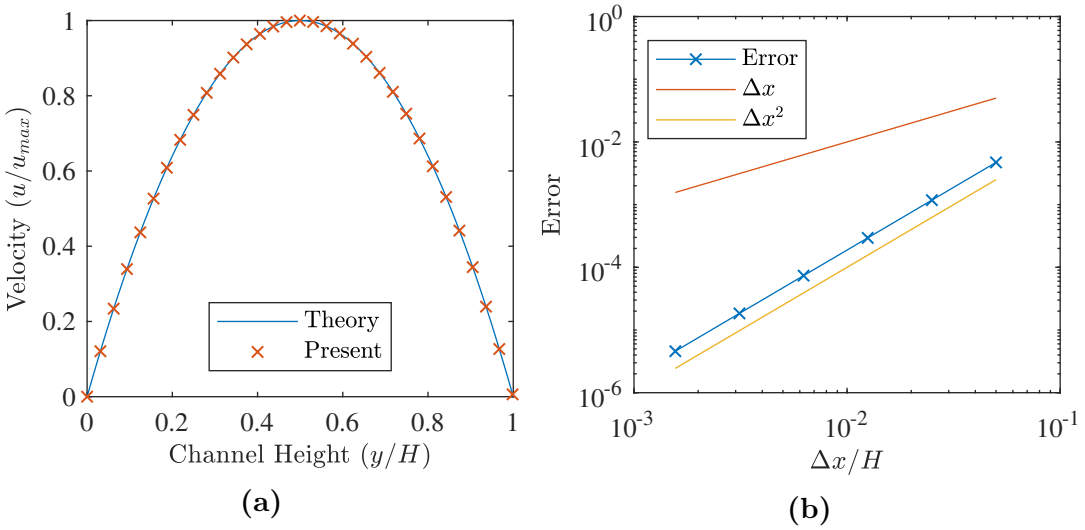

Figure 6.1: Velocity profile across the channel and error against grid resolution for Poiseuille flow test case. Order of convergence is 2.00.

and implies that the relaxation time-scale is held constant across different grid resolutions. This leads to the time step scaling such that $\Delta t \propto \Delta x^2$.

Figure 6.1 compares the analytical and computed velocity profiles at the finest grid resolution. Also shown is the relative error in the centreline velocity over a range of grid resolutions. The LBM solver correctly captures the velocity profile across the channel. Furthermore, comparing the relative error against the grid resolution shows second-order convergence, as expected.

The fluid solver has also been validated for an oscillating channel flow (Womersley flow). This validation is provided as part of Chapter 7, and so to avoid repetition it is not described here.

6.2.2 Structural Solver

Both structural solvers are tested in this section. First, the nonlinear static deflection of a cantilever beam under two different loading conditions is tested via the FEM solver. The time-dependence in the Euler-Bernoulli beam solver is built directly into the discretisation and is therefore not considered in the static case. The dynamic motion computed by both solvers is then tested in terms of a freely hanging flexible pendulum.

Static Case

To test the static deflection of the FEM solver, a cantilever beam is subjected to two types of loading conditions. The first loading is an applied tip moment. This results in a constant-curvature deflection of the beam. The relationship between the applied moment and the radius of curvature is given by:

$$\frac{1}{\kappa} = \frac{EI}{M} \tag{6.5}$$

where E is the Young's modulus, I is the second moment of area, and L is the length of the beam. Under a critical applied moment the beam rolls up to form a circle with circumference L and radius $L/2\pi$. This critical moment is given by:

$$M_{cr} = \frac{2\pi EI}{L} \tag{6.6}$$

The beam is discretised into 10 elements and the solution is obtained via Newton-Raphson iterations with 100 load increments. Figure 6.2 shows the computed beam shape under various loads, as well as the analytical solution. Excellent agreement is found between the two.

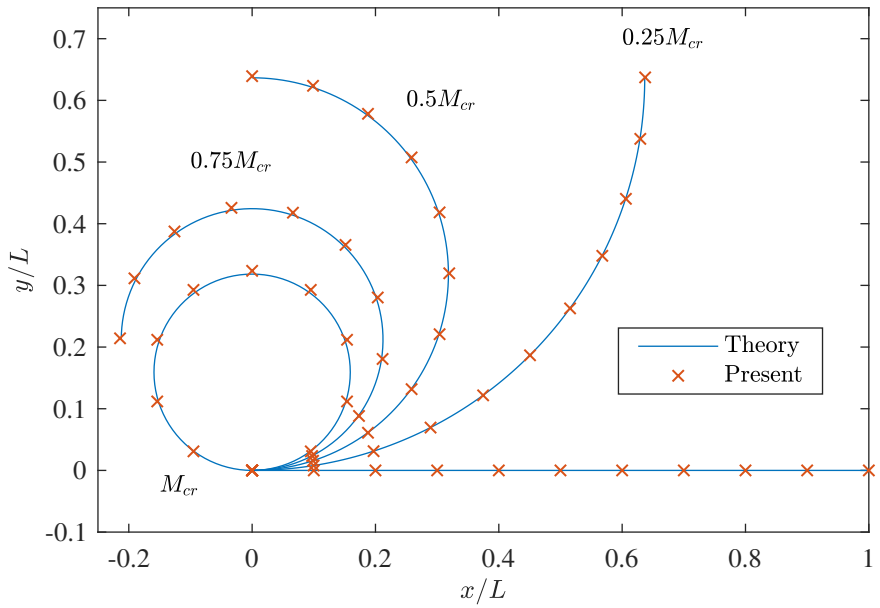

Figure 6.2: Profile shapes of cantilever beam under applied tip moment.

The second type of loading is a transverse tip load. In the case of small deformations, the tip deflection is linear and can be calculated in a straightforward manner:

$$\delta_y = \frac{F_y H^3}{3EI} \tag{6.7}$$

where F_y is the transverse tip force. However, in the case of large deflections, such as those expected in the present work, this linear behaviour breaks down due to a dependence of the structural stiffness on the deflection. Capturing this effect requires a nonlinear structural model, such as those described in Chapter 4 and adopted in the present work. The analytical solution for the nonlinear tip deflection is given by [186]:

$$\delta_y = L - 2[E(\mu) - E(\phi, \mu)]\sqrt{\frac{EI}{F_y}} \tag{6.8}$$

where $E(\mu)$ is the complete elliptic integral of the second kind, $E(\phi, \mu)$ is the incomplete elliptic integral of the second kind, and:

$$\mu = \frac{1 + \sin(\theta_L)}{2}, \qquad \phi = \arcsin\left(\frac{1}{\sqrt{2\mu}}\right) \tag{6.9}$$

The tip rotation at the end of the beam, θ_L, is obtained via:

$$\sqrt{\frac{F_y L^2}{EI}} = K(\mu) - F(\phi, \mu) \tag{6.10}$$

where $K(\mu)$ is the complete elliptic integral of the first kind, and $F(\phi, \mu)$ is the incomplete elliptic integral of the first kind. For a given load, F_y, Equation 6.10 is solved iteratively to get the tip rotation, θ_L. The transverse tip deflection is then obtained via Equation 6.8.

As before, the beam is discretised with 10 elements and the solution is obtained

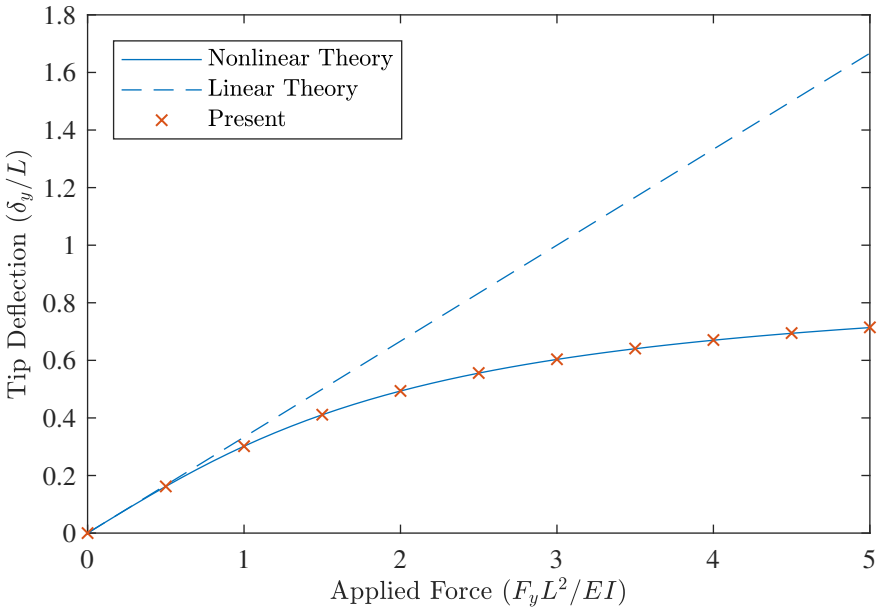

Figure 6.3: Tip deflection against transverse tip load. The linear deflection is given by Equation 6.7 whereas the nonlinear solution is given by Equation 6.8.

with 100 load increments. Figure 6.3 shows the analytical and computed tip deflections with applied load. Also shown is the linear solution obtained via Equation 6.7. Excellent agreement is found between the computed and nonlinear analytical solution. Furthermore, it can be seen that linear theory performs well for deflections below about 20 %. However, for deflections larger than this, the importance of the nonlinear structural model becomes clear.

Dynamic Case

The dynamic case, initially presented in Huang et al. [14], consists of a freely hanging filament under gravity. This is analogous to a flexible pendulum. The filament is initially held at an angle of 18° and is pinned at its fixed end. The Froude number, $Fr = |\mathbf{g}|L/u_\infty^2$, is 10. Two different dimensionless bending rigidities are tested, $\mathcal{K}^* = 0$ and $\mathcal{K}^* = 0.01$. Note that in this case, where there is no fluid, the fluid quantities involved in the nondimensionalisation are free but must be chosen consistently. Here, they are set to unity.

The filament is inextensible. In the Euler-Bernoulli beam solver this is enforced as part of the solution procedure. However, this is not the case in the FEM solver, which permits extension of the filament. To minimise the axial extension, another

nondimensional parameter is introduced; namely, the tensile stiffness:

$$\mathcal{T}^* = \frac{EA}{\Delta\rho u_\infty^2} \tag{6.11}$$

where A is the cross-sectional area of the filament ($A = h$ in 2D). Noting that $\mathcal{K}^* \propto Eh^3$ and $\mathcal{T}^* \propto Eh$, the Young's modulus and thickness can be adjusted to ensure a very high tensile stiffness at the correct value of bending stiffness. This minimises the axial extension in the FEM solution.

Figure 6.4 shows snapshots of the filament motion over one half-cycle for both bending rigidities, obtained with the Euler-Bernoulli beam solver. In the case of zero stiffness, there is a clear 'kick' motion where the tip of the filament rolls up. This is not observed in the finite stiffness case, where the filament tip remains straight throughout the cycle. Figure 6.5 shows the tip histories compared against the reference data [14] for both structural solvers, where $\mathcal{K}^* = 0.01$. Good

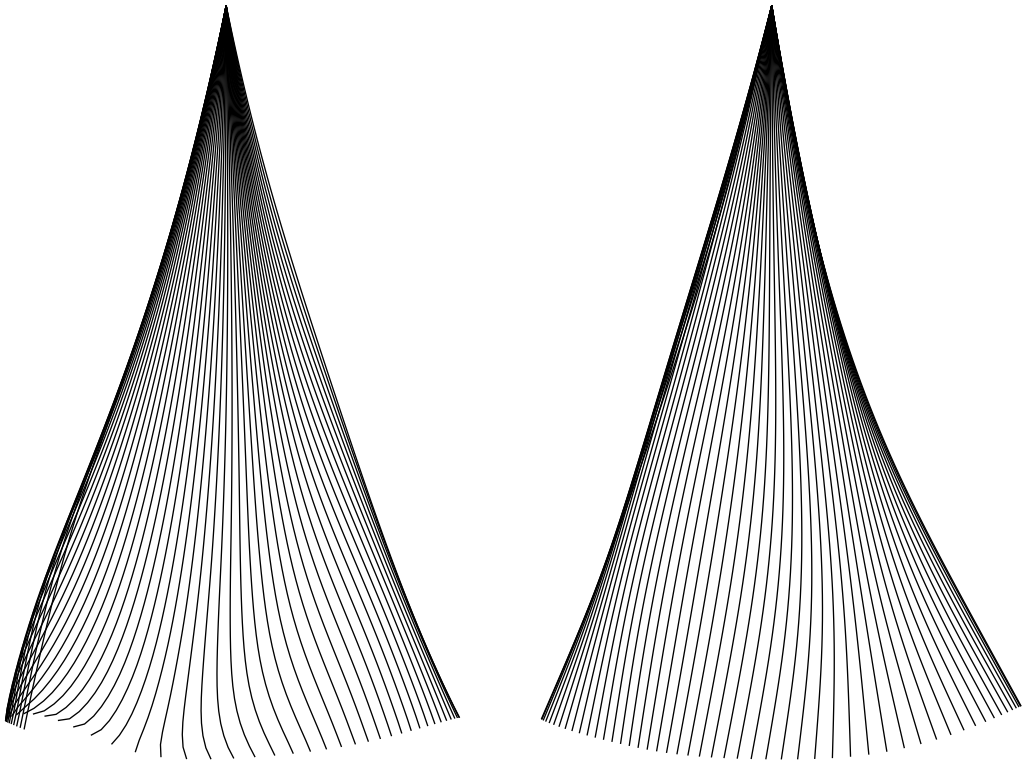

(a) $\mathcal{K}^* = 0$ **(b)** $\mathcal{K}^* = 0.01$

Figure 6.4: Snapshots of filament profile over one half-cycle obtained with the Euler-Bernoulli beam solver.

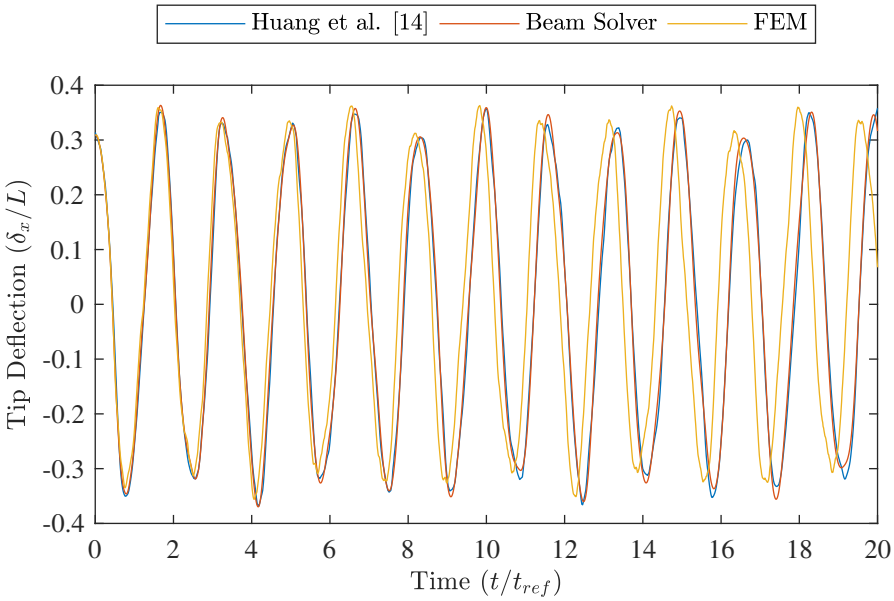

Figure 6.5: Tip histories obtained with the Euler-Bernoulli and FEM solvers compared against the benchmark data ($\mathcal{K}^* = 0.01$) from Huang et al. [14].

agreement is shown for both solvers, including the FEM, which permits (small) tensile extension. While the compute time for a single time step between both solvers is fairly similar, the implicit time integration scheme adopted in the FEM solver permits larger time steps, and therefore a shorter execution time.

6.2.3 Coupled Solver

In this final section, each of the components of the fully coupled model are tested together. First, the lattice Boltzmann-immersed boundary implementation is tested on its own for the flow around a rigid cylinder. To test the fluid-structure coupling, this case is then extended to include a flexible beam attached to the aft of the cylinder.

Rigid Case

To test the lattice Boltzmann-immersed boundary implementation, the flow around a rigid cylinder is computed and compared against benchmark data. A schematic of the case, initially presented in Schäfer et al. [187], is shown in Figure 6.6. A parabolic velocity profile is set at the inlet, no-slip on the top and bottom

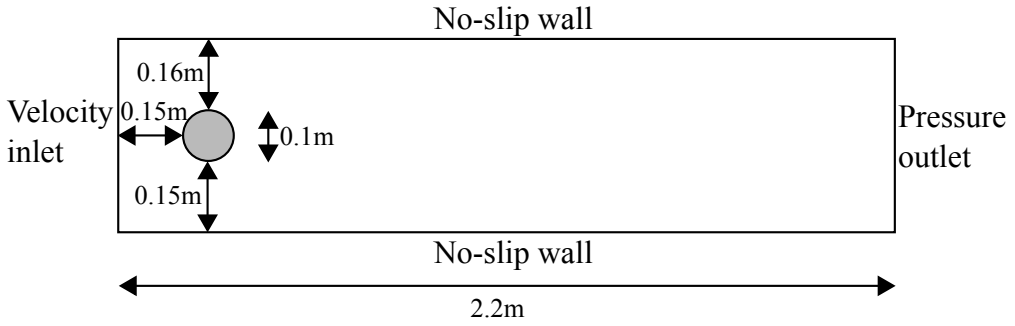

Figure 6.6: Schematic of rigid cylinder case.

walls, and a fixed pressure/density (equal to the initial value of the uniform pressure/density field) is set at the outlet. The Reynolds number, based on the mean flow velocity at the inlet (u_∞) and the cylinder diameter (D), is 100. At this Reynolds number, an unsteady von Kármán vortex street is observed behind the cylinder. Zero-velocity initial conditions are set everywhere and the inlet velocity is ramped up over the first two seconds, according to:

$$u_x(0,y,t) = \begin{cases} u_x^{in}(y)\frac{1-\cos(\pi t/2)}{2} & t < 2 \\ u_x^{in}(y) & \text{otherwise} \end{cases} \tag{6.12}$$

where $u_x^{in}(y)$ is the classic parabolic velocity profile. Figure 6.7 shows a snapshot of the instantaneous vorticity field.

The benchmark provides data for the Strouhal number and maximum lift and drag coefficients, which are measured once the initial transient has disappeared and the shedding behaviour becomes periodic. This data is provided by the benchmark in

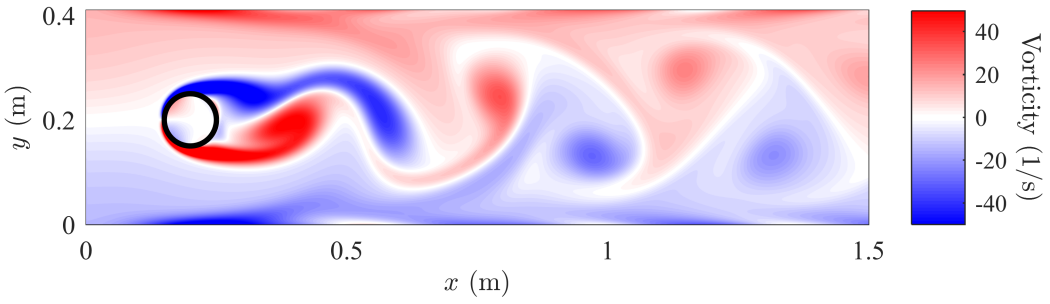

Figure 6.7: Instantaneous out-of-plane vorticity for rigid cylinder case.

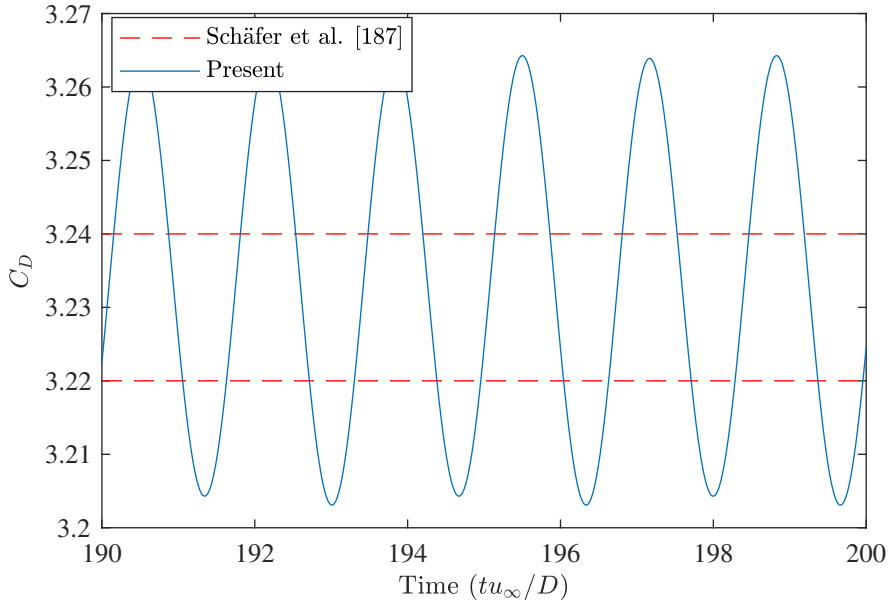

(a)

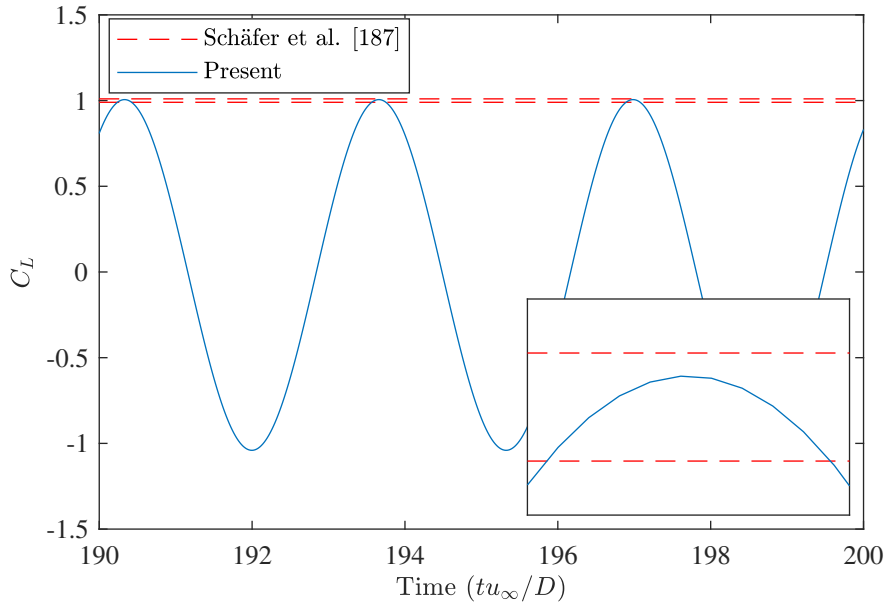

(b)

Figure 6.8: Lift and drag coefficients for rigid cylinder case compared against the benchmark data of Schäfer et al. [187]. The dashed lines show the range of maximum values reported by ten different research groups in the benchmark.

the form of ranges, where results were collected from ten different research groups, each using their own modelling approach. Figure 6.8 shows the time histories of the lift and drag coefficients, and the reported ranges of the maximum values provided by the benchmark [187]. The maximum lift coefficient of the numerical results falls within the range of reported values. However, there is a clear overprediction of drag. This can possibly be explained by the discussion given in Section 5.1.4. Due

to the finite support region required for the interpolation/spreading step in the IBM, the boundary seen by the fluid is not the true boundary prescribed by the Lagrange markers. This will affect the drag imparted on the cylinder. However, it is not thought to have a significant effect on the lift, since both the top and bottom side of the cylinder experience the same boundary augmentation.

To check that the overprediction of drag is in fact an artefact of the IBM, an in-depth comparison has been made with other studies which use a similar modelling approach. Figure 6.9 plots the time histories of the lift and drag coefficients and compares them against another study utilising a similar IBM scheme [188], showing very good agreement. Table 6.4 reinforces this point by comparing the maximum lift and drag coefficients and Strouhal numbers against two other studies using a similar IBM scheme [188, 189]. For reference, the mean values of the original benchmark results are also shown. Clearly evident is the fact that the present work and the two cited studies all show an overprediction of drag when compared to the original benchmark.

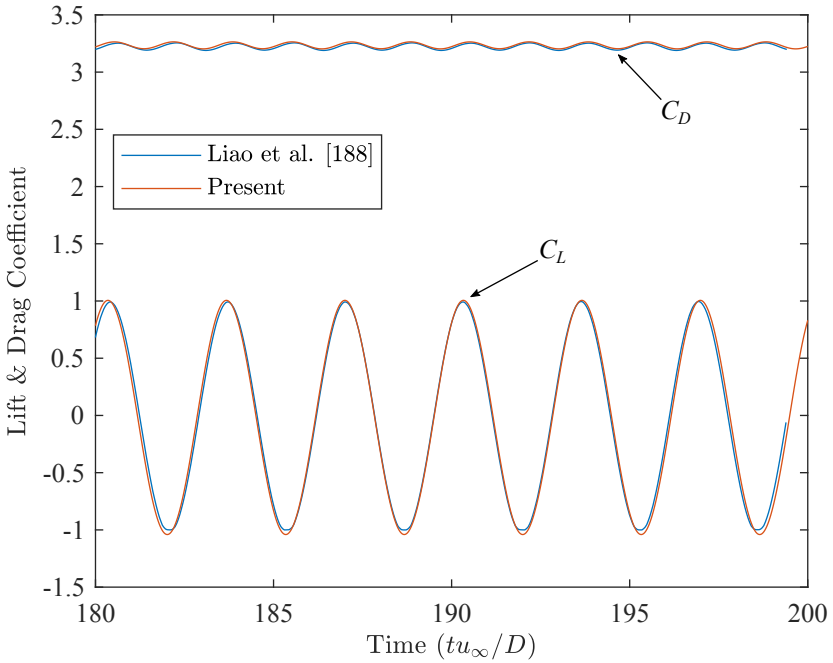

Figure 6.9: Time histories of lift and drag coefficients for rigid cylinder case compared against Liao et al. [188].

Table 6.4: Lift, drag and Strouhal numbers for selected studies which use a similar IBM scheme to the present method. For reference, the mean values of the original benchmark results are also shown.

Source	Max C_D	Max C_L	St
Present	3.26	1.00	0.30
Original Benchmark [187]	3.23	1.00	0.30
Liao et al. [188]	3.26	0.99	0.30
Lee and Lee [189]	3.29	1.04	0.30

Flexible Case

The purpose of the final case is to test all of the components of the fully coupled model together. The case considered here is very similar to the rigid cylinder case; however, this time a flexible beam ($\rho_s = 10 \times 10^3\,\mathrm{kg/m^3}$, $E = 1.4\,\mathrm{MPa}$) is attached to the aft of the cylinder [190]. A schematic of the case is shown in Figure 6.10. Like the cylinder case, the Reynolds number is 100. This results in a self-sustaining flapping motion as the wake of the cylinder interacts with the flexible beam, as shown in Figure 6.11.

Figure 6.12 shows the tip deflection in both the streamwise and transverse directions. The structural dynamics were computed using the FEM solver. The agreement between the present model and the benchmark data is good, although there are some slight differences in the tip motion. This may be because the boundary of the beam in the present model is represented by a line of IBM markers, whereas in the original benchmark case it has a finite thickness. Nevertheless, the agreement is good, considering this simplification. Table 6.5 compares the Strouhal numbers (St) of the axial and transverse tip motion against the benchmark data

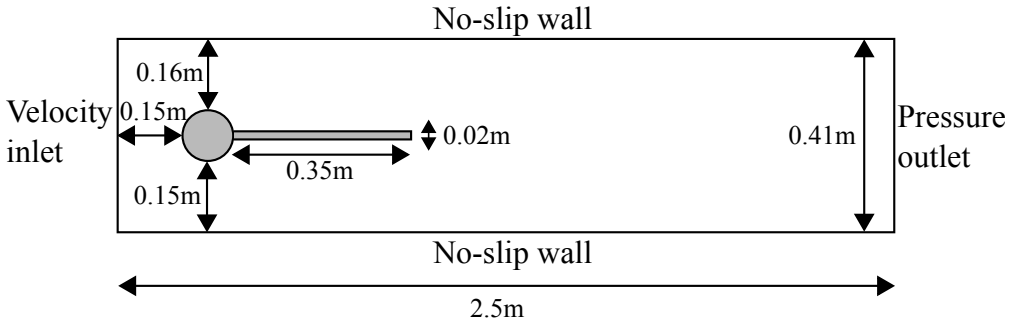

Figure 6.10: Schematic of rigid cylinder with attached flexible beam.

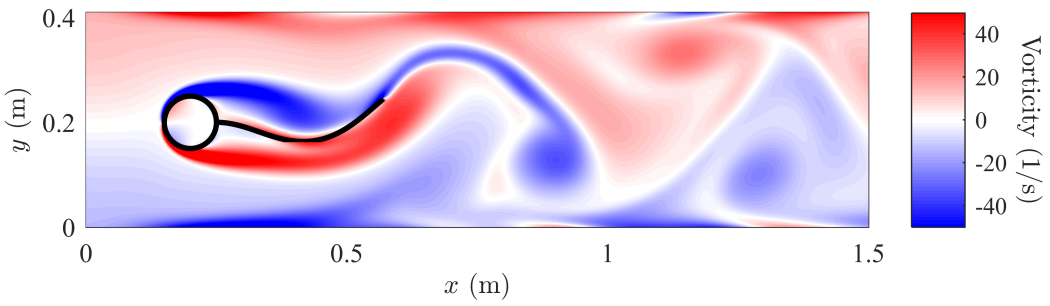

Figure 6.11: Instantaneous out-of-plane vorticity for rigid cylinder with attached flexible beam.

and again reinforces the agreement, with the values within 4 % of each other.

The fully coupled model with the Euler-Bernoulli beam solver has been validated for a separate experimental case. This validation is provided as part of Chapter 8, and so to avoid repetition it is not described here. Furthermore, this experimental case is also used as another test for the FEM solver, the details of which are provided in Chapter 9.

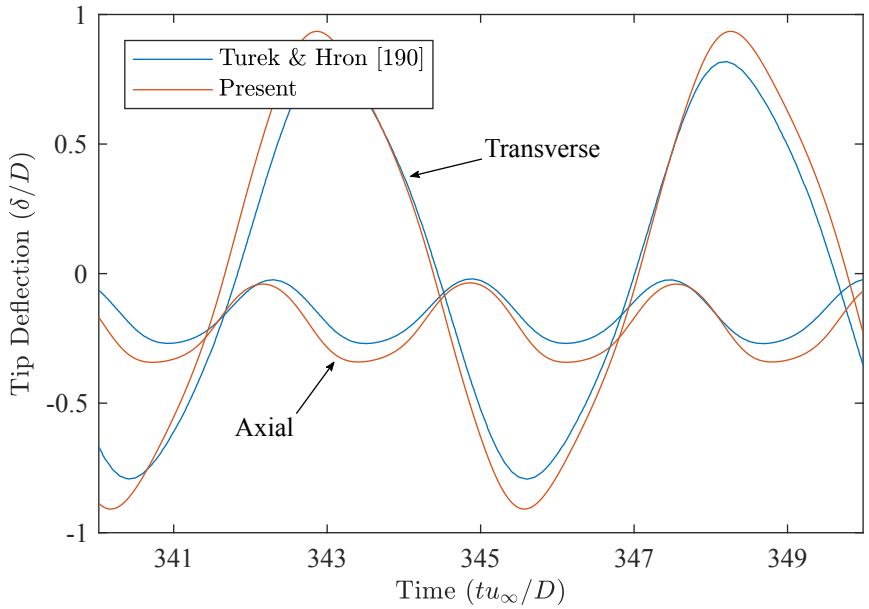

Figure 6.12: Transverse and axial tip motion histories compared against the benchmark data from Turek and Hron [190].

Table 6.5: Strouhal numbers of tip motion compared against the benchmark data of Turek and Hron [190].

Source	***St* (axial)**	***St* (transverse)**
Present	0.370	0.185
Turek and Hron [190]	0.384	0.192

6.3 Summary

This chapter has described the main method developments undertaken as part of the present work. The purpose of these developments was to enable efficient simulations of large arrays of slender structures over a range of input conditions. Furthermore, each separate component of the model, as well as the fully coupled model itself, has been validated against common benchmark cases. The following three chapters now present the application of this model to three different test cases, which build in scale from a single flap in a periodic array, to a small finite array in an oscillating flow, and finally to a large array of flaps in an open-channel flow.

Chapter 7

Single Flap in a Periodic Array

This chapter is the first of the core results chapters and is taken from the following publication:

J. O'Connor, A. Revell, P. Mandal, and P. Day. Application of a lattice Boltzmann-immersed boundary method for fluid-filament dynamics and flow sensing. *J. Biomech.*, 49(11), 2016 [18]

The paper describes the investigation into a periodic array of flexible flaps under both steady and oscillating flow conditions. Only one flap is directly modelled whereas the array is represented by appropriate periodic boundary conditions. A range of Reynolds numbers are tested and the resulting dynamic response is measured. The results show that at large Reynolds numbers the recirculation region between the flaps fully bridges the gap. This leads to a reduced lag between the flow and structural response, as well as decreased flap deflection.

My contributions to this paper are as follows: validating the numerical model, setting-up and running all of the simulations, post-processing, analysis, and writing the paper.

7.1 Abstract

Complex fluid-structure interactions between elastic filaments, or cilia, immersed in viscous flows are commonplace in nature and bear important roles. Some biological systems have evolved to interpret flow-induced motion into signals for feedback response. Given the challenges associated with extracting meaningful experimental data at this scale, there has been a particular focus on the numerical study of these effects. Porous models have proven useful where cilia arrangements are relatively dense, but for more sparse configurations the dynamic interactions of individual structures play a greater role and direct modelling becomes increasingly necessary. The present study reports efforts towards explicit modelling of regularly spaced wall-mounted cilia using a lattice Boltzmann-immersed boundary method. Both steady and forced-unsteady 2D channel flows at different Reynolds numbers are investigated, with and without the presence of a periodic array of elastic inextensible filaments. It is demonstrated that the structural response depends significantly on Reynolds number. For low Reynolds flow, the recirculation vortex aft of successive filaments is small relative to the cilia spacing and does not fully bridge the gap; in which case the structure lags the flow. At higher Reynolds number, when this gap is fully bridged, the structure and flow move in phase. The trapping of vortices between cilia is associated with relatively lower wall shear stress. At low to intermediate Reynolds, vortex bridging is incomplete and large deflection is still possible, which is reflected in the tip dynamics and wall shear stress profiles.

7.2 Introduction

Fluid-structure interactions (FSI) between slender filaments immersed in viscous flows occur often in nature and play many important biological and physiological roles, from propulsion, self-cleaning and particle transport in motile cilia, to sensing and mechanotransduction in non-motile cilia. In the latter case, evolution has found a number of creative and effective solutions which utilise flexible filaments for flow sensing, providing an excellent source of inspiration for designers [38]. These phenomena have been the subject of interest across a range of disciplines, partly due to increasing focus on miniaturisation and the need for flow sensing at microscopic scales. Elucidating the ability of a simple ciliary system to efficiently

sense changes in a flow field for initiating a control response could inspire important innovations in aerospace, microfluidics, and biotechnology alike.

Many animal species rely on external filament-like structures to sense small-scale changes in their surrounding flow field and adapt their behaviour accordingly. Various arthropods use filiform hairs to sense perturbations in the air, such as those created by a beating wing, to avoid flying predators [191]. Similarly, bats use hairs on their wings for feedback flow sensing to aid in flight control [192]. Another example in nature is the glycocalyx layer (GL) – a layer of nanoscopic hairs that cover the endothelial cells in the vascular system. The GL is thought to play a primary role in sensing changes in flow conditions within blood vessels and transmitting these to the endothelial surface [193, 194]. It is likely that this transmission of forces through the GL from the bulk blood flow to the endothelial wall regulates important vascular processes, such as vascular morphogenesis. As such, the GL plays a major physiological role in the human body.

Due to the scale and complexity of the fluid-induced motion, detailed experimental measurements of the aforementioned systems are particularly challenging to obtain. As such, numerical modelling can provide a significant advantage in understanding these interactions. Most methods however are limited to simulating bulk effects, and incorporate cilia dynamics via a porous layer model rather than resolving individual structures [195, 196]. Where cilia density is high, this is often a justified assumption and allows for a general recovery of correct flow response. Nevertheless, these approaches are usually calibrated for a specific range of conditions, beyond which individual filament-fluid dynamics and the viscous coupling between structures can lead to poor approximations. Porous layer models demarcate a pervious region of homogeneous properties, calibrated to model flow through a dense layer of filaments. These models are useful as they allow a significant reduction in complexity while still recovering the overall behaviour of the fluid. Recently, Pontrelli et al. [197] coupled the lattice Boltzmann method (LBM) with a porous layer model to demonstrate that the GL significantly reduces the endothelial wall shear stress (WSS), and thus highlighted the importance of representing the GL in any WSS studies of the endothelium. Other studies have also demonstrated the use of porous layer models for the GL [198, 199], and have theorised the further role of the GL in initiating response to more complex processes. However, their primary focus was on analysing mass transport rather than the mechanotransduction capability of the GL. At larger scales of

motion, Favier et al. [2] successfully developed a homogenised poroelastic approach to model a layer of feathers aft of a circular cylinder; highlighting a potential mechanism for drag reduction which has since attracted a lot of attention. As such, homogenised approaches have proven their potential to respond realistically to a range of needs. Reference data is then necessary for their calibration, and in the absence of experimental data this is most practically obtained from the direct modelling of the flow-structure interactions.

There have been multiple studies in recent years that have focussed on the explicit modelling of individual filaments specifically to examine the flow sensing capability of cilia, for example Bathellier et al. [200] and Cummins et al. [201]. However, both studies were somewhat limited by a simplified fluid model, restricting the range of application to Stokes flow conditions. More recently, Heys et al. [202] and Lewin and Hallam [203] developed general methods to solve the complete Navier-Stokes equations, and thus explicitly modelled both the fluid and structural dynamics of the coupled system. Both of these studies examined the viscous coupling between multiple filaments at low Reynolds numbers. It should be noted that there are many more examples in the literature of explicit fluid-filament modelling [23, 37, 204]. However, at present it seems that none have investigated the mechanotransduction and flow sensing capability of elastic filaments. As such, the aim of the present work is to provide the first steps towards developing a tool for the numerical simulation of elastic filaments immersed in a viscous flow, so that important fluid properties and the associated filament response can be assessed in detail. For this initial study the focus is on a sparse array of 2D filaments which repeats periodically. A recently developed lattice Boltzmann-immersed boundary method is employed to allow for explicit simulation of the individual fluid and solid dynamics [110]. In the following section the numerical method is briefly presented along with the simulation setup and flow parameters for both a steady and unsteady (oscillating) crossflow at moderate Reynolds numbers. Initial results are subsequently presented and key flow physics are discussed, focussing on the intra-filament flow dynamics which would otherwise be neglected in a homogeneous arrangement.

7.3 Methods

7.3.1 The Lattice Boltzmann Method

The LBM has evolved over the recent years into a powerful method, capable of outperforming more commonly used Navier-Stokes-based approaches in certain complex flow problems. For the present work, the ease with which it can be coupled with the immersed boundary method (IBM) and explicitly incorporate complicated FSI phenomena makes it an ideal choice. This section gives a brief introduction to the LBM. However, for a more in-depth overview, with examples of common applications, the reader is directed to one of the many reviews in the literature [205–207].

The driving equation behind the LBM is the Boltzmann transport equation. Discretising this equation along the q velocity directions of the lattice (see Figure 7.1) results in the discrete lattice Boltzmann equation (LBE):

$$f_i(\mathbf{x} + \mathbf{c}_i \Delta t, t + \Delta t) - f_i(\mathbf{x}, t) = \Omega_i \tag{7.1}$$

where $f_i(\mathbf{x}, t)$ is a probability function that represents the proportion of particles

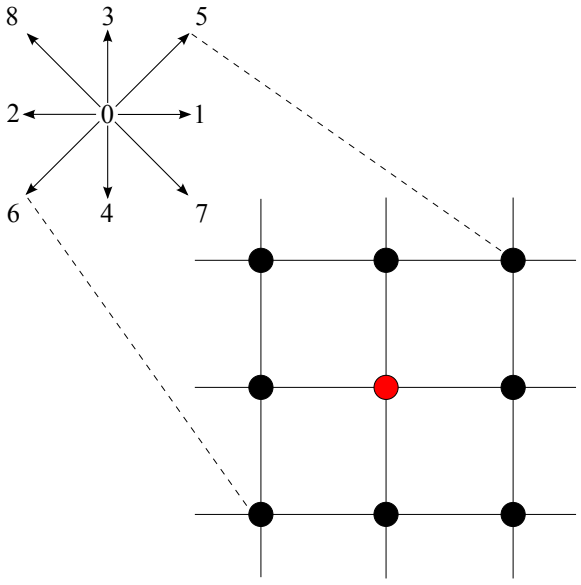

Figure 7.1: D2Q9 lattice model with numbering scheme.

at lattice site $\mathbf{x}$ and time t moving with a lattice velocity i. Using a probability function to represent clusters of fluid particles places the LBM somewhere between a true particle and true continuum method, and thus it exhibits properties of both. These 'packets' of fluid molecules evolve according to the LBE and the resulting values of the probability functions are used to calculate the averaged macroscopic quantities (velocity and density).

The collision operator, Ω_i, is notoriously difficult to solve in the LBM and typically the Bhatnagar-Gross-Krook (BGK) approximation [124] is used to simplify the collision operator to:

$$\Omega_i = \frac{1}{\tau}(f_i^{eq} - f_i) \tag{7.2}$$

where τ is the relaxation time-scale and f_i^{eq} is the local equilibrium distribution function, which depends only on local macroscopic flow properties. The resulting evolution equation for the LBGK model is:

$$f_i(\mathbf{x} + \mathbf{c}_i\Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \frac{1}{\tau}[f_i^{eq}(\mathbf{x}, t) - f_i(\mathbf{x}, t)] \tag{7.3}$$

The relaxation time-scale, τ, is directly related to the kinematic viscosity of the fluid, ν, via the relationship:

$$\nu = \left(\tau - \frac{1}{2}\right) c_s^2 \tag{7.4}$$

The equilibrium function that leads to the correct form of the Navier-Stokes equations is:

$$f_i^{eq} = \rho w_i \left[1 + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^2} + \frac{1}{2}\frac{(\mathbf{c}_i \cdot \mathbf{u})^2}{c_s^4} - \frac{1}{2}\frac{\mathbf{u}^2}{c_s^2}\right] \tag{7.5}$$

where, for the two-dimensional nine-speed (D2Q9) lattice model adopted in the

present work and illustrated in Figure 7.1:

$$c_s = \frac{1}{\sqrt{3}} \tag{7.6}$$

$$\mathbf{c} = \begin{pmatrix} 0 & 1 & -1 & 0 & 0 & 1 & -1 & 1 & -1 \\ 0 & 0 & 0 & 1 & -1 & 1 & -1 & -1 & 1 \end{pmatrix} \tag{7.7}$$

$$w_i = \begin{cases} 4/9 & i = 0 \\ 1/9 & i = 1...4 \\ 1/36 & i = 5...8 \end{cases} \tag{7.8}$$

The macroscopic quantities at each lattice point are calculated by taking the sum of the moments of the corresponding probability density function at that particular lattice site.

$$\rho = \sum_{i=0}^{8} f_i, \qquad \rho\mathbf{u} = \sum_{i=0}^{8} f_i \mathbf{c}_i \tag{7.9}$$

In the present work, a body force, $\mathbf{f}$, is introduced via the imposed pressure gradient and through the IBM (discussed in Section 7.3.2). To account for this, the evolution equation is modified such that:

$$f_i(\mathbf{x} + \mathbf{c}_i\Delta t, t + \Delta t) = f_i(\mathbf{x}, t) + \omega[f_i^{eq}(\mathbf{x}, t) - f_i(\mathbf{x}, t)] + \Delta t F_i \tag{7.10}$$

where F_i is the discretised form of the total body force, $\mathbf{f}$, given by Guo et al. [130]:

$$F_i = w_i \left(1 - \frac{\omega}{2}\right) \left[\frac{\mathbf{c}_i - \mathbf{u}}{c_s^2} + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^4}\mathbf{c}_i\right] \cdot \mathbf{f} \tag{7.11}$$

and the modified relationships for the macroscopic quantities are:

$$\rho = \sum_{i=0}^{8} f_i, \qquad \rho\mathbf{u} = \sum_{i=0}^{8} f_i \mathbf{c}_i + \frac{1}{2}\mathbf{f} \tag{7.12}$$

The total body force is the sum of the immersed boundary force and the pressure gradient force (acceleration due to gravity is zero):

$$\mathbf{f} = \mathbf{f}_{ib} + \mathbf{f}_{pressure} \tag{7.13}$$

The pressure force acts equally on all fluid lattice sites and is calculated by converting the physical pressure gradient to nondimensional lattice units. After adding it to the immersed boundary force the total body force is discretised and inserted into the LBM procedure via Equations 7.10–7.12. While treating the pressure gradient as a body force and allowing it to act equally throughout the domain may not be an entirely physical representation, this method is common in the literature and there are examples of studies involving similar geometries that use this approach [113, 208]. Furthermore, the primary focus of this study was to examine the trends in the filament dynamics under a driven flow at various flow rates.

In the context of the present application, the primary advantage to using the LBM over a continuum method is when incorporating the FSI. The underlying regular lattice in the LBM helps with this procedure – compared to a finite volume method, for example, which may have been designed to handle irregular meshes. Also, many similar biological applications involving fluid-filament dynamics exist that also include additional complex physics; such as particle-particle or particle-filament forces in particle/cell regulation and transport. Another reason for selecting the LBM is, due to its particulate nature, it is better suited to handling these additional complex phenomena and thus allows easier extension of the current work to more elaborate applications.

7.3.2 The Immersed Boundary Method

In the present work, the IBM is used to couple the fluid and structural solvers. By discretising the fluid and solid on separate, independent grids, the IBM dispenses with the need for time-consuming re-meshing procedures between each time step. However, as the solution evolves on two separate grids, the transfer of information between these grids is crucial.

Initially, the LBM step is performed on the Eulerian (fluid) grid without any immersed object in the flow, resulting in a predicted velocity field, $\mathbf{u}^*$. This velocity field is interpolated from the Eulerian grid onto the N Lagrangian (structure) markers of the immersed boundary, located at $\mathbf{X}$ in Lagrange space:

$$\mathbf{U}^*(\mathbf{X}) = \int_\Omega \mathbf{u}^*(\mathbf{x})\ \delta(\mathbf{x} - \mathbf{X})\, \mathrm{d}\mathbf{x} \tag{7.14}$$

where Ω is the computational domain on the Eulerian grid and δ is the Dirac delta function. In the present notation, lower case is used to represent variables defined in the Eulerian space whereas upper case is used for the Lagrange space.

The interpolated velocities can then be used to calculate the corrective force, $\mathbf{F}_{ib}$, required to achieve the desired velocity on the immersed boundary, $\mathbf{U}^b$:

$$\mathbf{F}_{ib}(\mathbf{X}) = \frac{\mathbf{U}^b(\mathbf{X}) - \mathbf{U}^*(\mathbf{X})}{\Delta t} \tag{7.15}$$

The penultimate step is the spreading of the force from Lagrange space back onto the Eulerian grid. This is done through a convolution with the Dirac delta function over the immersed surface, Γ:

$$\mathbf{f}_{ib}(\mathbf{x}) = \int_\Gamma \mathbf{F}_{ib}(\mathbf{X})\ \delta(\mathbf{x} - \mathbf{X})\, \mathrm{d}s \tag{7.16}$$

where s is the Lagrangian coordinate along Γ.

In the numerical procedure, a discrete approximation (mollifier) to the Dirac delta is needed. The form of the mollifier, $\tilde{\delta}$, used in the present work was proposed by Roma et al. [165]:

$$\tilde{\delta}(r) = \begin{cases} \frac{1}{3}\left(1 + \sqrt{-3r^2 + 1}\right) & |r| \leq 0.5 \\ \frac{1}{6}\left(5 - 3|r| - \sqrt{-3(1 - |r|)^2 + 1}\right) & 0.5 \leq |r| \leq 1.5 \\ 0 & \text{otherwise} \end{cases} \tag{7.17}$$

which is easily extended to two dimensions:

$$\tilde{\delta}(\mathbf{x} - \mathbf{X}) = \tilde{\delta}(x - X)\, \tilde{\delta}(y - Y) \tag{7.18}$$

It can be seen from Equation 7.17 that the mollifier has a compact support of three lattice spacings. In 2D this equates to a support kernel containing nine lattice nodes for each Lagrange marker, which Equations 7.14 and 7.16 act over. When a Lagrange point is located very close to a boundary (e.g. the marker that represents the base of the filament), there may be nodes within the support stencil for that marker that exist outside the domain. In this case, during the velocity interpolation and force spreading steps, the contributions to/from these support nodes are neglected and set to zero. This treatment is justified as the values at the solid and fluid nodes on this boundary are set later in the algorithm to the prescribed boundary conditions.

The final step in the procedure is to perform the LBM step once more on the Eulerian grid, but this time including the body force distribution obtained from the IBM step, via Equations 7.10–7.12. The result is, at the end of the time step, the fluid feels the presence of the immersed boundary through the imposed body force. The reader is directed to Favier et al. [110] and Pinelli et al. [162] for further details and validation of this approach.

7.3.3 Filament Structural Dynamics

The governing equations of motion for a flexible filament are given in nondimensional Lagrangian form as follows:

$$\frac{\partial^2 \mathbf{X}}{\partial t^2} = \frac{\partial}{\partial s}\left(T\frac{\partial \mathbf{X}}{\partial s}\right) - \frac{\partial^2}{\partial s^2}\left(K_B\frac{\partial^2 \mathbf{X}}{\partial s^2}\right) + Fr\frac{\mathbf{g}}{|\mathbf{g}|} - \mathbf{F}_{ib} \tag{7.19}$$

where T is the tension in the filament, K_B is the bending rigidity, $\mathbf{g}$ is the gravitational acceleration, and $Fr = {}^{|\mathbf{g}|L}/_{u_\infty^2}$ is the Froude number, where L is the length of the filament. For the full derivation of Equation 7.19 and the relationships for the nondimensional reference values, see Huang et al. [14] and Favier et al. [110]. This system is closed via the inextensibility condition to ensure non-physical stretching at high deflection rates. The resulting system is discretised and then solved via the Newton method and the new positions of the filament Lagrange markers are updated before beginning the next LBM time step.

7.4 Simulation Setup

Figure 7.2 shows a schematic of the 2D computational domain, which is periodic in the x-direction to represent an infinite array of filaments. The flexible filament is located equidistant from the left and right boundaries with a height L. The height and length of the channel are $H = 3L$ and $W = 2L$ respectively, and so the filament spacing is W. As the domain is two-dimensional, the filament represents a flap of infinite extent in the z-direction. This is a significant simplification versus three-dimensional structures, where the spanwise blockage ratio can be expected to be much lower. However, the development and testing of three-dimensional structures is beyond the scope of the current work. Note that while the term cilium refers to a 3D structure, the 2D flexible filaments used in the present work can be considered representative of such structures and thus the terms are used interchangeably throughout the chapter.

The simulation parameters, in both physical units and nondimensional lattice units, are shown in Table 7.1. A preliminary study indicated that both the spatial and

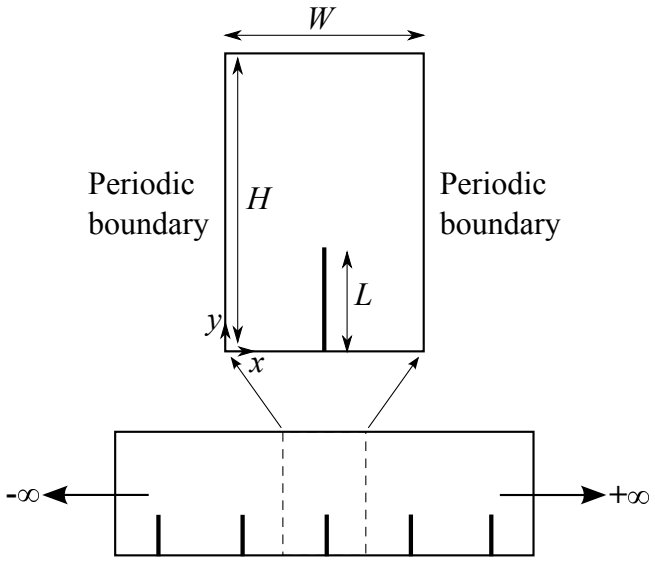

Figure 7.2: Schematic of computational domain.

temporal resolution were adequate for the flows studied. The midway bounceback boundary condition was used to enforce a no-slip velocity on the upper and lower channel walls. In all cases, the flow is driven by a pressure gradient of magnitude (or amplitude) A, which was implemented in the LBM via the forcing scheme discussed in Section 7.3.1. Two forms of the driving pressure gradient were used to model both constant (Poiseuille) and oscillating (Womersley) flow between two infinite-span parallel plates. For the Womersley flow cases, the applied pressure

Table 7.1: Simulation parameters in physical and lattice units.

Parameter	**Symbol**	**Physical Value**	**LBM Value**
Channel height	H	$0.2\,\mathrm{m}$	105
Channel length	W	$0.13\,\mathrm{m}$	70
Filament height	L	$0.067\,\mathrm{m}$	35
Fluid density	ρ	$1\,\mathrm{kg/m^3}$	1
Kinematic viscosity	ν	$5 \times 10^{-4}\,\mathrm{m^2/s}$	1.684×10^{-3}
Lattice spacing	Δx	$1.905 \times 10^{-3}\,\mathrm{m}$	1
Time step	Δt	$1.2 \times 10^{-5}\,\mathrm{s}$	1
Time period	T	$5.027\,\mathrm{s}$	411476
Womersley number	α	5	5
Stokes length	δ^*	$0.028\,\mathrm{m}$	14.85

gradient was varied sinusoidally so that:

$$\frac{\mathrm{d}p}{\mathrm{d}x} = A\cos(2\pi f t) \tag{7.20}$$

where the frequency of oscillation, f, was set to achieve the desired Womersley number, α, given by:

$$\alpha = \frac{H}{2}\sqrt{\frac{2\pi f}{\nu}} \tag{7.21}$$

The Womersley number was set to a value of $\alpha = 5$ so that the unsteady (pulsatile) effects dominated over viscous effects. This value was selected following Boyd et al. [209] as a midrange of larger arterial flow in humans ($\alpha_{\text{brachial}} = 3$, $\alpha_{\text{carotid}} = 5$, $\alpha_{\text{aorta}} = 15$). At this level of Womersley number 'plug-flow' velocity profiles are evident, whereby during the cycle the flow does not fully develop into the classic parabolic profile shape. Also, there is a phase lag between the pressure signal and velocity, as well as a change in velocity amplitude compared to the steady state [210]. The period, T, of each oscillatory cycle is:

$$T = \frac{\pi}{2\nu}\left(\frac{H}{\alpha}\right)^2 \tag{7.22}$$

indicating that there are approximately 4×10^5 time steps per cycle, as shown in Table 7.1. The Stokes layer thickness, δ^*, is given by:

$$\delta^* = \sqrt{\frac{2\nu}{\omega}} \tag{7.23}$$

where in this instance $\omega = 2\pi f$ refers to the angular frequency. The filament length, L, was selected to be roughly of the order of δ^*. The spacing used in the present study, $W = 2L$, represents a sparse configuration which allows for a range of flow recirculation for the selected values.

Table 7.2 displays a summary of the cases studied and the corresponding pressure

Table 7.2: Summary of simulated cases. Each of the eight cases shown here were run with and without the filament present (sixteen cases total).

Case	Pressure Signal	A [Pa/m]	τ (τ_{max}) [mPa]	Re_τ ($Re_{\tau_{max}}$)	u_{max} [m/s]	$Re_{u_{max}}$	Re_{δ^*}
P1	Poiseuille	0.025	2.5	10.00	0.25	100	–
P2	Poiseuille	0.100	10.0	20.00	1.00	400	–
P3	Poiseuille	0.175	17.5	26.46	1.75	700	–
P4	Poiseuille	0.250	25.0	31.62	2.50	1000	–
W1	Womersley	0.180	2.5 (3.54)	10.00 (11.89)	0.15	61.6	8.71
W2	Womersley	0.721	10.0 (14.14)	20.00 (23.78)	0.62	246.5	34.85
W3	Womersley	1.261	17.5 (24.74)	26.46 (31.46)	1.08	431.3	61.00
W4	Womersley	1.802	25.0 (35.36)	31.62 (37.61)	1.54	616.1	87.14

gradients applied in each; cases P1–4 represent Poiseuille flow, whereas cases W1–4 represent the Womersley flow cases. All simulations were run with and without the filament present. For all cases, the amplitude of the pressure gradient (A) was selected so as to achieve a specific Reynolds number. In P1–4 the Reynolds number based on maximum velocity, $Re_{u_{max}}$, was varied from 100 to 1000. These values correspond to a range of friction Reynolds number (based on WSS), Re_τ, of 10 to 31.6; where τ refers to the WSS. In cases W1–4 the RMS value of WSS, $\langle\tau\rangle$, is taken for comparison and values remain within the range of commonly quoted values for laminar flow [211]. Even considering Reynolds number based on the maximum WSS reached throughout the cycle, values are observed to remain below this threshold and as such, turbulence effects can safely be ignored. Reynolds numbers based on δ^* are also included to show that a similar range is examined to that of Boyd et al. [209].

To facilitate comparison between constant and sinusoidal cases, the pressure gradients in W1–4 were adjusted so that the RMS value of WSS was equal to the constant value of WSS in cases P1–4. This equality holds only in the absence of filaments but as such, comparisons on the relative effect of the filaments on the WSS for both constant and sinusoidal flow can be made. In the current study,

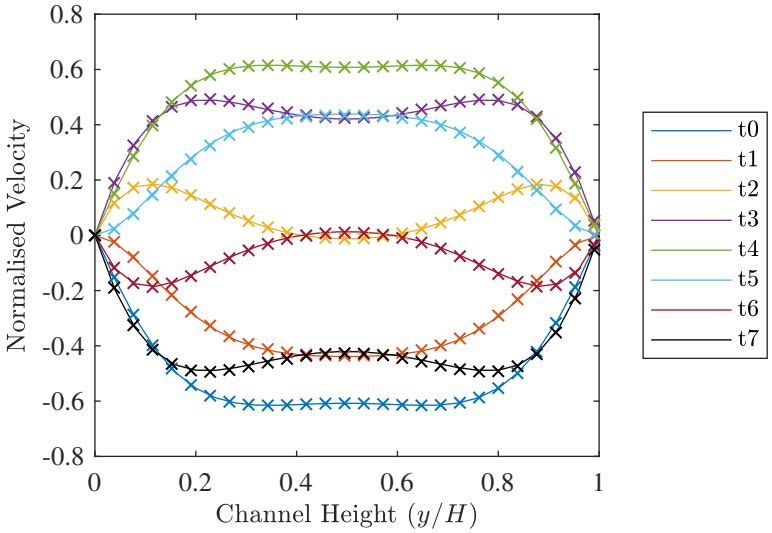

Figure 7.3: Womersley velocity profiles at eight equally spaced points throughout cycle. Present (x) and analytical solution (–). Note that the velocity has been normalised by the maximum velocity obtained from the steady case.

with $\alpha = 5$, the pressure gradient scaling factor between corresponding P and W cases was approximately 7.21, resulting in the unsteady velocity profiles presented in Figure 7.3. Despite the increase in pressure gradient, the maximum velocity was seen to reduce by around 38.4 % when compared to the corresponding Poiseuille cases, as illustrated in the figure and the corresponding $Re_{u_{max}}$ values in Table 7.2. For a detailed discussion of the effect of Womersley number on the velocity profile, the reader is referred to Loudon and Tordesillas [210].

7.5 Results & Discussion

7.5.1 Validation of WSS

The computational approach has been previously validated for fluid-filament dynamics in Favier et al. [110]. In addition to a mesh convergence study, the flow dynamics and WSS calculation procedure were also validated for the range of flow parameters under examination. The flow was initially set at rest in a 2D channel, as shown in Figure 7.2 (without the filament). On applying a constant pressure gradient, the flow begins to move and eventually develops into the classic parabolic velocity profile. The WSS at the midpoint of the channel wall was

recorded during this development stage and the final steady state value compared against the theoretical value, obtained from the relationship:

$$\tau = \frac{\mathrm{d}p}{\mathrm{d}x}\frac{H}{2} \tag{7.24}$$

This procedure was followed for a range of pressure gradients, corresponding to a Re_τ of 10 to 31.6. Figure 7.4 shows the measured WSS values over a simulated time of 60 s for both P1–4 and W1–4, in the absence of the structure. A good agreement is found between the theoretical values and the simulated steady state values, even at high Reynolds numbers where the velocity gradients are relatively large. We recall that in the unsteady cases the pressure gradient is adjusted so that the RMS value, $\langle\tau\rangle$, corresponds to the WSS from the steady cases. As such, the time variation of WSS for cases W1–4 are shown for comparison, where negative values are clipped for clarity.

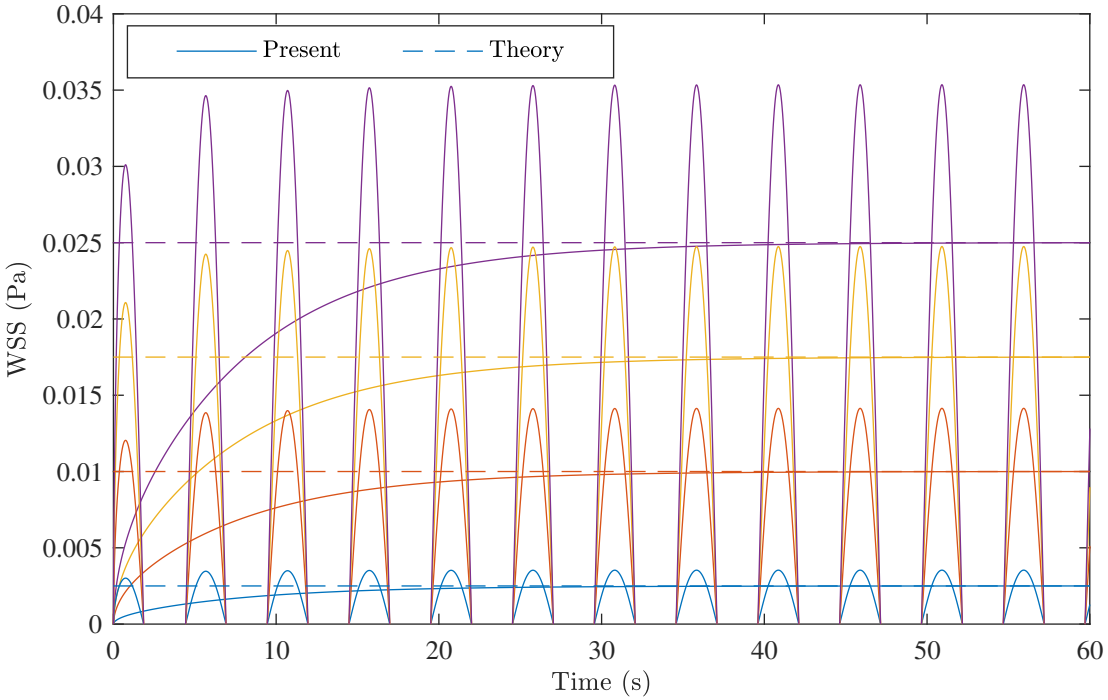

Figure 7.4: Simulated and theoretical WSS values for cases without filament. Positive Womersley values are shown for comparison.

7.5.2 Effect of Reynolds number

Figure 7.5 presents the effect that the filament has on the variation of WSS, τ, for a range of friction Reynolds numbers, Re_τ. In each case τ^* is the WSS normalised by the value from the no-filament case (i.e. the values given in Table 7.2). Since τ^* is reported to be below unity in all cases, it can be inferred that the presence of the filament acts to reduce WSS over the entire length of the channel wall, in some cases by an order of magnitude. A reduction in WSS is expected due to the blockage imposed by the filaments near the channel wall. The flow deflects the tip and forms a recirculation region between successive filaments, while the bulk of the flow passes over the filament tip itself. Figure 7.6 displays vorticity contours

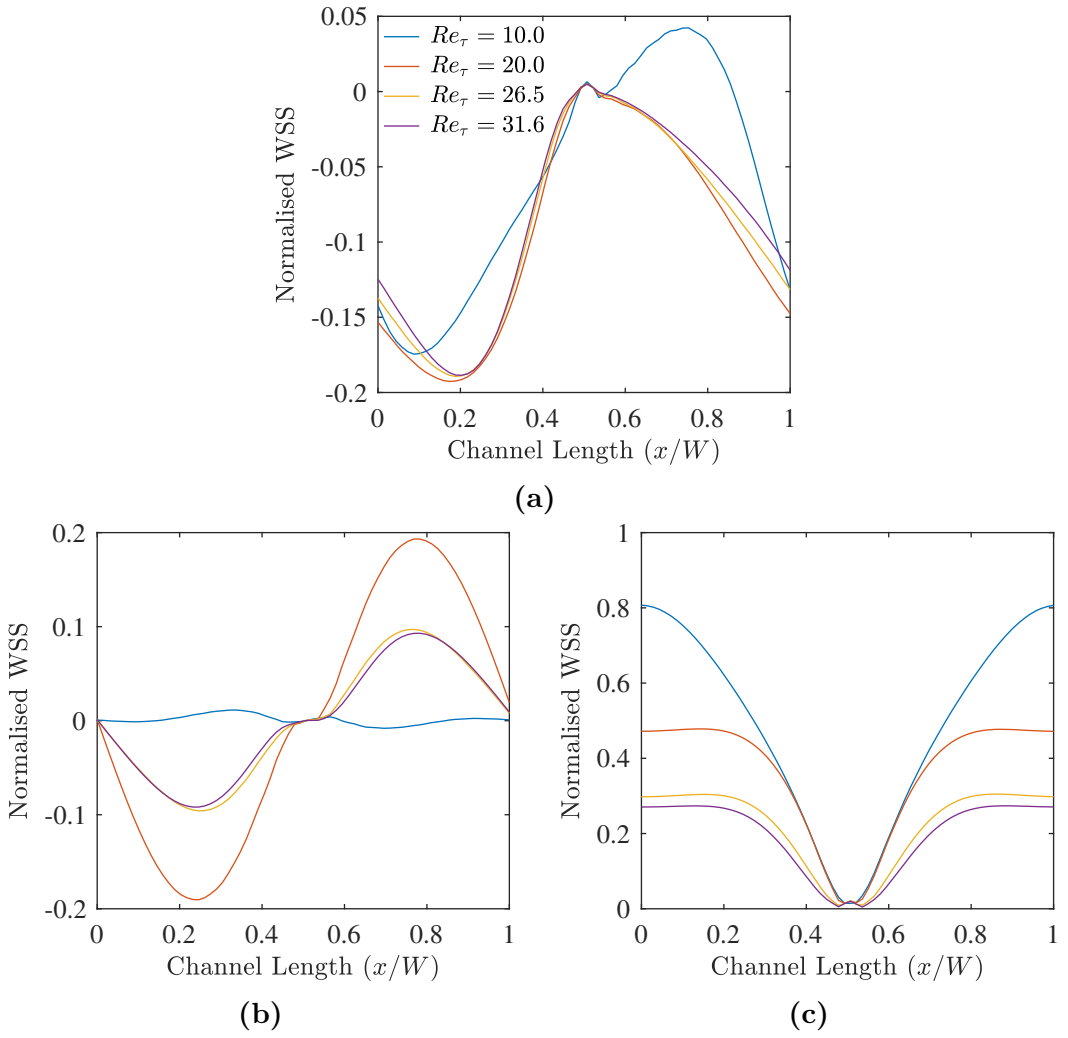

Figure 7.5: WSS along channel wall with filament present. (a) τ^* for cases P1–4, (b) $\overline{\tau}^*$ for cases W1–4 and (c) $\langle\overline{\tau}^*\rangle$ for cases W1–4. In each case the WSS has been normalised by the value obtained from the no-filament case.

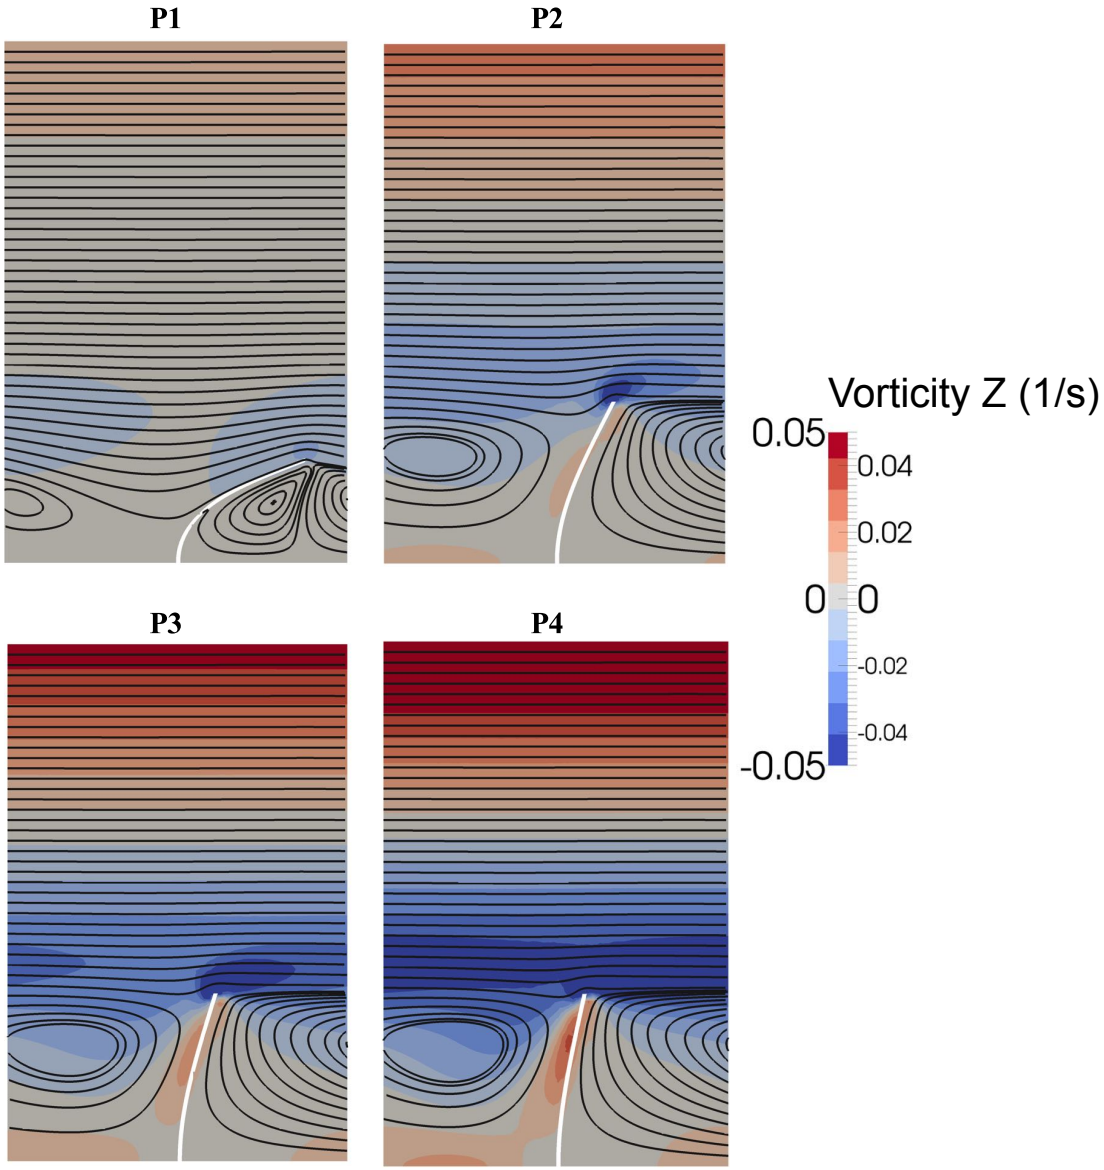

Figure 7.6: Streamlines and vorticity contours for steady flow cases P1–4 with filament present.

for cases P1–4, with the streamlines overlaid. Figure 7.7 shows tip deflection and tension at the filament base for all cases.

7.5.3 Poiseuille Flow with Filaments

For the steady flow cases, all results show a negative value of WSS in the region $0 < x/W < 0.5$, i.e. across the second half of the interval between filaments in the direction of the crossflow. This corresponds to a flow recirculation region and

it is apparent that the WSS profile is almost independent of Reynolds number for cases P2–3. However, for the lowest Reynolds number (P1) there is a change in sign of WSS in the region $0.5 < x/W < 0.85$, corresponding to a change in flow direction at the wall. This corresponds to the first section of the interval, into which the filament is deflected. Indeed, considering Figure 7.6, one can see that for P1 there is a stagnation point at the wall, directly below the deflected tip, about which the flow is rotating in opposite directions. Cases P1–4 all show a peak in the WSS close to the midpoint between the filaments, in the region $0.1 < x/W < 0.2$. Examining Figure 7.6, there is a correspondence between the spike in WSS and a peak in the levels of vorticity.

At the lower Reynolds number there is a greater tip deflection, which results in a lower vertical intrusion of the filament into the bulk flow. Connected to this, the flow velocity at the point of contact with the filament is lower, and thus the recirculation length is reduced. In terms of reducing the overall levels of WSS, case P1 is the most effective. From Figure 7.7 the maximum tip deflection is proportional to the tension computed at the base of the filament, as can be expected. It can also be seen that the difference in tip deflection between cases P1 and P2–4 is around a factor of four. Since a decrease in tip deflection for increasing flow velocity (Reynolds number) is a somewhat counter-intuitive response, further tests are needed in the lower Reynolds number regime to identify more precisely the Reynolds number at which this behaviour manifests.

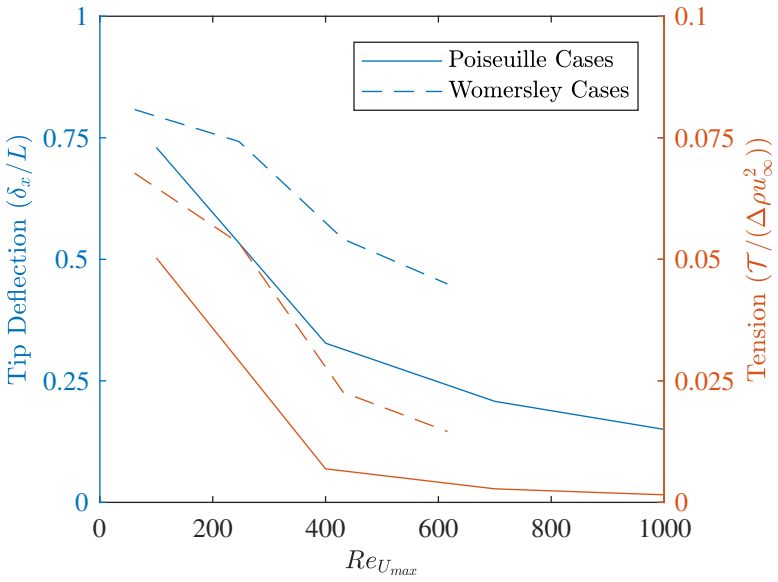

Figure 7.7: Tension at filament base and maximum filament tip deflection for all cases. Both values are given in nondimensional form.

7.5.4 Womersley Flow with Filaments

Figures 7.5b and 7.5c display two different plots of the streamwise WSS variation for cases W1–4; the time-averaged values, $\overline{\tau}^*$, and the time-averaged RMS values, $\langle\overline{\tau}^*\rangle$. Figure 7.8 provides a series of snapshots from each of the cases W1–4, taken at eight regular intervals with $t/T = n/8$, where $n = 0$ coincides with the start of the pressure gradient cycle. In each case the streamlines are provided with an attempt made to capture the main flow features in each snapshot. In addition to Figure 7.7 plotting the maximum deflection and tension for each case, Figure 7.9a plots the evolution of the tip deflection against the cyclic pressure gradient, and Figure 7.9b plots the derived lag against the normalised recirculation length.

Considering first the values of $\langle\overline{\tau}^*\rangle$ from Figure 7.5c, there is a continuous reduction with increasing Reynolds number, tending towards a maximum value of 0.26. Interestingly however, the trend is not monotonic when considering the time-averaged variation, $\overline{\tau}^*$, in Figure 7.5b. For case W1 the values of $\overline{\tau}^*$ are close to zero, after which the values reach a maximum for W2 before decreasing towards very similar variations for both W3 and W4. For a Womersley flow with no obstruction the time-averaged WSS should be zero, and so a value close to zero with the inclusion of a filament implies that its presence has a low impact over the duration of the cycle. The plot of $\langle\overline{\tau}^*\rangle$ in Figure 7.5c confirms this, by indicating that the WSS is only reduced by 20 % in the centre of the interval; though this is reduced to zero near the filament itself, by virtue of the blockage effect. For cases W2–4, the values of $\langle\overline{\tau}^*\rangle$ reach plateaus between the filaments in the regions $0 < x/W < 0.3$ and $0.7 < x/W < 1$, which are continuous in the periodic domain.

Figure 7.8 shows that, like the steady cases, vortical flow structures appear for the cases with large pressure gradients – leading to decreased tip deflection and thus lower velocities in the near wall region. Furthermore, the maximum of $\overline{\tau}^*$ occurs at both $x/W \approx 0.25$ and $x/W \approx 0.75$, corresponding to the location of maximum vorticity associated with the interspersed recirculation zones. For a larger recirculation region, a smaller range of tip motion is permitted, as is confirmed in Figure 7.7. Interestingly, it appears that the tip deflection for W1 is greater than for P1, indicating that comparison based on Re_τ alone is insufficient. Indeed, if one compares $Re_{u_{max}}$ instead, the value of W1 is less than P1, which might better describe this effect.

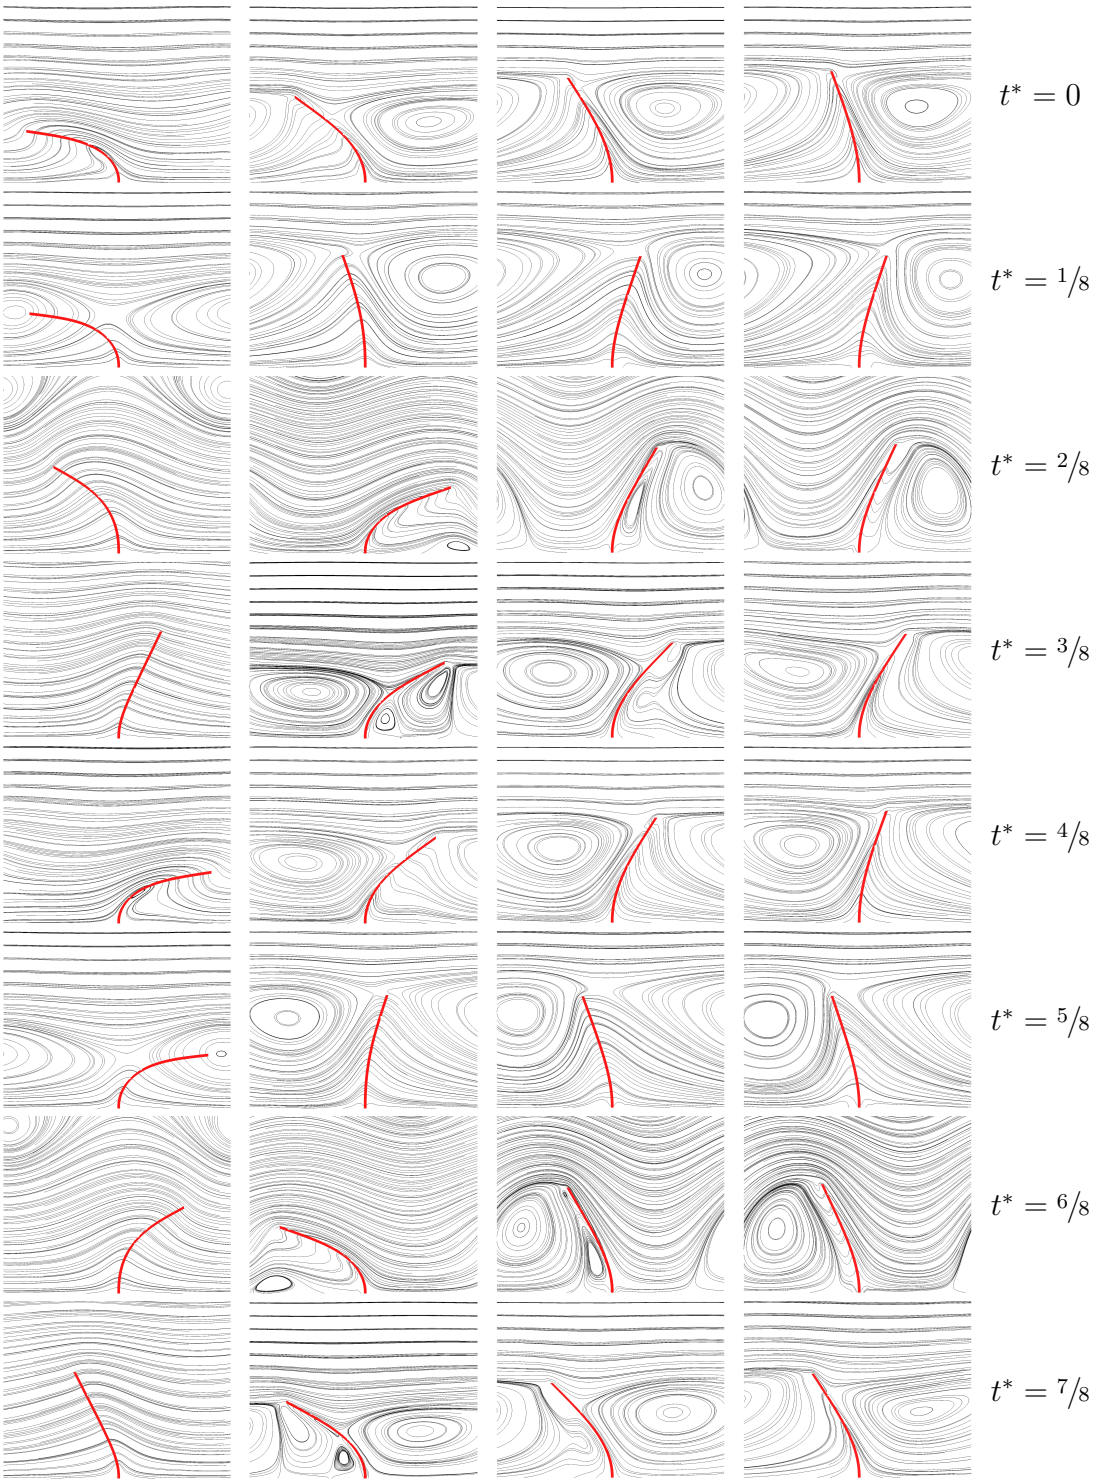

Figure 7.8: Snapshots of streamlines at eight equally spaced time steps during one full cycle for cases W1–4. $Re_{u_{max}} = 61.6$ (column 1), $Re_{u_{max}} = 247$ (column 2), $Re_{u_{max}} = 431$ (column 3) and $Re_{u_{max}} = 616$ (column 4). Note that $t^* = t/T$.

Overall there is a clear correlation between the tension at the base of the filament and the maximum tip deflection experienced, both of which are inversely proportional to Reynolds number. However, as viscous effects and filament rigidity start to dominate over the inertial effects (at lower Reynolds numbers), this trend will likely reverse itself. This behaviour was observed in Heys et al. [202] where, at the very low Reynolds numbers that were examined, increased flow velocity led to increased tip deflection, which is a physically intuitive response.

Figure 7.9a provides insight into the dynamic response of the filament over one cycle, indicating a phase lag versus the imposed pressure signal, which has been normalised to unity for clarity. The lag is greatest for the W1 case and reduces with increasing Reynolds number. Figure 7.9b displays this lag, measured as the time between $T = 0$ and the point where the tip deflection is zero, against the recirculation region normalised by the filament spacing. Larger recirculation regions, apparent at higher Reynolds numbers, act as buffers between filaments which couple the motion of consecutive filaments, thereby increasing the response to the applied pressure gradient. More measurements are required to ascertain the proportionality of this relationship.

7.6 Conclusions

The present work has demonstrated the capability of a recently developed lattice Boltzmann-immersed boundary method to investigate the flow around an infinite series of 2D flexible filaments. The approach provides a stable and reliable means of investigating complex interactions involving high structural deformation, via explicit modelling of the various physics involved.

For the selected cases, a range of nonlinear effects are observed in which fluid and structural dynamics are strongly coupled; most of which would pose a challenge for homogeneous porous layer modelling. Calibration of the pressure gradient for Womersley flows to match equivalent levels for Poiseuille flows has enabled the effects of embedded flexible filaments to be compared. In the steady flow cases, the filaments were observed to reduce WSS by generating a region of lower momentum flow reversal between subsequent filaments, with a recirculation region that spans the interval at higher Reynolds numbers. As the Reynolds number is

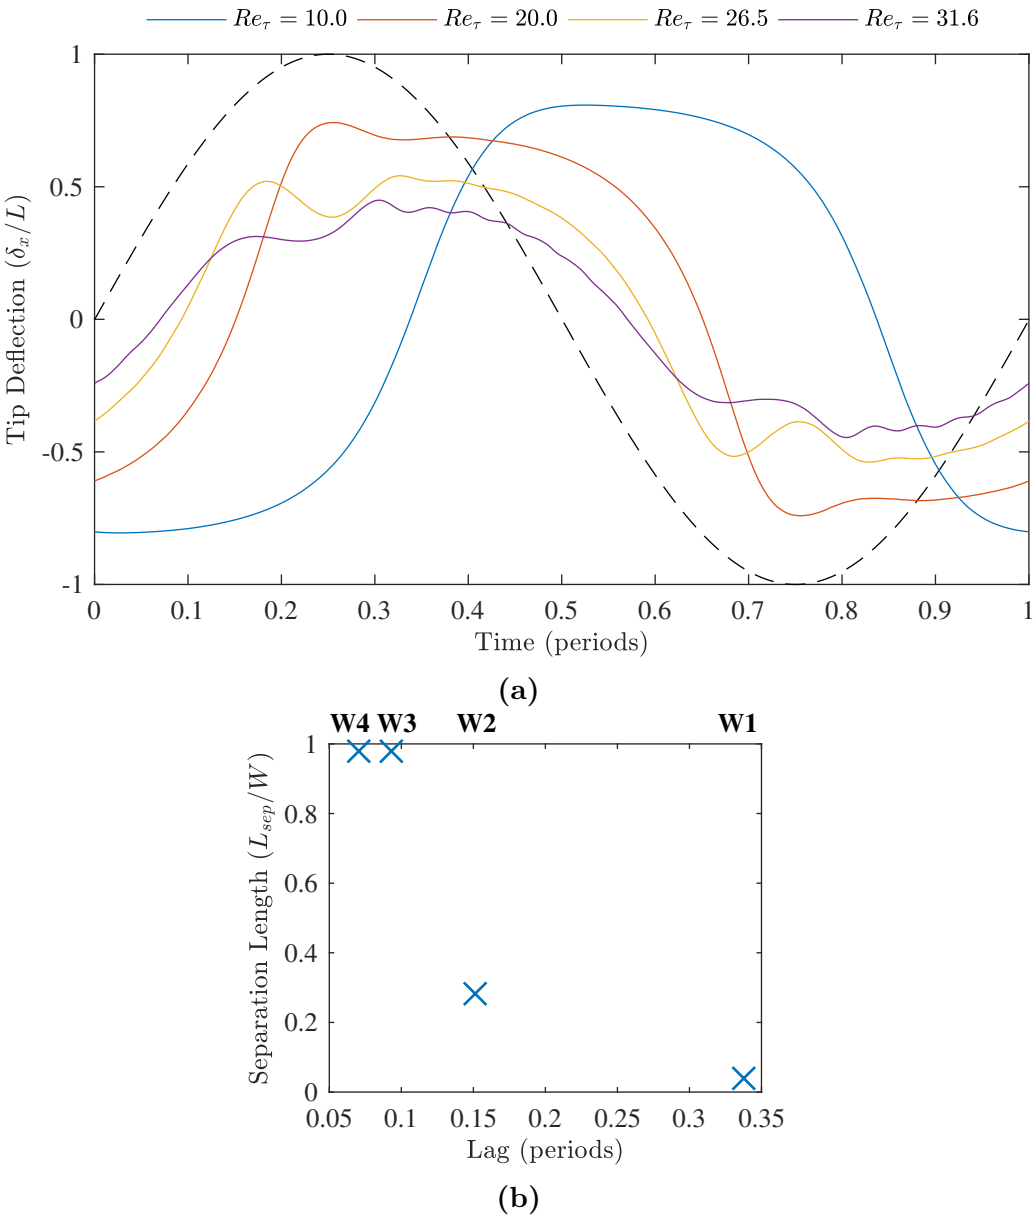

Figure 7.9: (a) Tip deflection in the x-direction for cases W1–4. Overlaid is the pressure signal (scaled to unity). (b) Separation length of the primary recirculation region and the lag of the tip response (against the pressure signal) for cases W1–4. Tip deflection is normalised by L whereas separation length is normalised by W.

lowered there is a negligible effect on the WSS profile until a critical threshold of Reynolds number, below which the recirculation length becomes shorter than the filament spacing. At this point the flow gives rise to a secondary recirculation region of opposite sign, and the shape of the WSS profile changes significantly over a small portion of the interval.

In the Womersley cases, a dynamic lag between maximum tip deflection and

peak pressure gradient was observed, which reduced with increasing Reynolds number. Analysis of the flows indicated that the bridging of the interval between consecutive filaments enabled a faster response of the filaments to the imposed flow frequency. These observations are expected to be of potential interest both to observers of nature and designers of flow control systems alike. They demonstrate the potential for tuning a system to a particular range of flow parameters, as well as highlighting the need for explicit modelling of the filaments to capture these dynamics.

One of the major limitations of the present study is the simplification to 2D – effectively simulating infinitely wide plates rather than filaments. While there are no doubt applications for which 2D plates may be preferable to individual cilia, the extension of these findings to 3D cilia is not possible. Where the structures have a finite depth the overall blockage ratio will be significantly reduced; and a portion of the flow will traverse around the structures, rather than over the top of them. This will likely lead to increased WSS, due to increased flow closer to the channel wall, as well as other complex 3D flow effects. The effect of other parameters, such as bending stiffness, filament height, and filament spacing, should also be considered to fully appreciate the nature of these configurations.

The present study is a first step towards the investigation of 2D fluid-filament interactions and provides some insight into the effect of Reynolds number and Womersley number. Future studies may focus on flow conditions relevant to specific structures such as the glycocalyx layer or flow sensing devices, although there is expected to be a wide range of parameters to investigate. As an example, the ratio of channel height to filament length in the current study is three, while in cases involving the GL it can be two orders of magnitude higher. In the present model the resolution needed to capture the filament dynamics at these sorts of ratios would increase the computation time dramatically. Future work may investigate the potential of using a refined LBM, or perhaps a multiscale coupling method, to overcome this obstacle.

Chapter 8

Multiple Flaps in an Oscillating Channel Flow

This chapter presents work undertaken as part of the EU-funded PELskin project[1]. This project consisted of several partners across Europe with the aim of investigating the amelioration of aerodynamic performance of bluff bodies via a poroelastic (PEL) coating. The projected resulted in a series of publications, two of which I co-authored:

J. Favier, C. Li, L. Kamps, A. Revell, J. O'Connor, and C. Brücker. The PELskin project - part I: fluid-structure interaction for a row of flexible flaps: a reference study in oscillating channel flow. *Meccanica*, 52(8), 2017 [20]

A. Revell, J. O'Connor, A. Sarkar, C. Li, J. Favier, L. Kamps, and C. Brücker. The PELskin project: part II - investigating the physical coupling between flexible filaments in an oscillating flow. *Meccanica*, 52(8), 2017 [21]

Part I of the series introduced a simplified experimental case, which was used to isolate the key flow control mechanisms involved in the coupled system. As part of this work, my role was to validate the present numerical model against the experiment for this simplified case. Part II of the series focussed on using the numerical model to perform a detailed parametric study which analysed the

[1]https://trimis.ec.europa.eu/project/pel-skin-novel-kind-surface-coatings-aeronautics.

effect of material properties on the coupled system. This parametric study was my main contribution to the project. However, since the validation from Part I forms an important step in this work, the validation section from Part I of the series has been extracted and combined with Part II to form this chapter.

My contributions to these papers are as follows: validating the numerical model (Part I), setting-up and running the simulations (Part II), post-processing (Part II), and jointly working on the analysis (Part II). Although I did not write the original papers, the extract given in this chapter has been re-written in my own words.

8.1 Abstract

Several studies have reported the beneficial effects of passive flexible flaps attached to the aft of a bluff body in the form of improved aerodynamic performance. The mechanism by which this occurs is thought to be linked to the emergence of travelling waves through the array, which modify the wake and shedding cycle. In the context of the EU-funded PELskin project, a simplified configuration consisting of an oscillating channel flow with an array of wall-mounted flexible flaps is investigated. The oscillating nature of the flow is emblematic of the von Kármán vortex shedding associated with bluff body wakes. A lattice Boltzmann-immersed boundary model is first validated against experimental data for this selected case. A detailed analysis is then undertaken to examine how the structural natural frequency affects this wave-like coupling mechanism. The travelling waves are characterised by examining the flow topology and spectra of the tip motion and near-field fluid velocity. The analysis indicates that the wave-like motion is associated with the formation of vortices, which arise due to the resistance the array imparts on the crossflow. Primary vortices upstream and downstream of the array are observed, which interact with the shear layer above the array. This results in a phase shift along the array, resembling the observations made in previous studies regarding bluff body flow control.

8.2 Introduction

The dynamic interactions between fluids and flexible slender structures is a general topic encountered in a broad range of applications. Examples of such interactions include the waving of plants in fields or rivers [8, 77, 82], or the dynamic motion of cilia, which have many roles throughout nature in processes such as self-cleaning and flow sensing [18, 38]. Several studies have reported the emergence of coherent waving motions when arrays of slender structures are subjected to a crossflow [10, 77, 82]. This behaviour is common in vegetative flows, where it is called honami (terrestrial vegetation) or monami (aquatic vegetation). This waving motion is driven by Kelvin-Helmholtz instabilities, which arise due to the inflection within the shear layer that separates the slow-moving flow within the array from the fast-moving flow above the array [8].

A similar wave-like motion has been observed in flow control applications, where recent studies have shown that poroelastic coatings on bluff bodies can control the flow separation and wake pattern [2]. These coatings, initially inspired by bird feathers, consist of an array of densely-packed slender structures, and are therefore similar to vegetative flows. The compliance of these coatings manipulates the flow topology within the wake. Under the right conditions this can lead to improved aerodynamic performance, such as reduced drag and lift fluctuations [2, 32]. Recent experimental studies have shown that these devices interact with the wake and modify the shedding cycle [9]. In particular, it has been noted that at a critical Reynolds number the amplitude of the waving motion increases markedly and a travelling wave-like motion through the array is observed, which results in a narrowing of the wake and a reduction in drag [9].

These observations provided the motivation for the EU-funded PELskin project. The project involved a small consortium of partners focussed on investigating the amelioration of aerodynamic performance via a Porous ELastic (PEL) coating. The aim of this project was to characterise the potential for passive slender structures to interact with the separated flow in the wake region of bluff bodies, and ultimately reduce drag. This chapter forms part of a series of outputs summarising the work of this project. Specifically, this chapter focusses on the numerical modelling of the aforementioned systems. To highlight the important mechanisms involved, a simplified setup is sought so as to examine the flap behaviour under well-defined conditions. Such a case is studied herein, whereby an oscillating channel flow with an array of flexible wall-mounted flaps is considered. This configuration captures the essential features of a bluff body wake, while significantly reducing the problem complexity.

Numerical models that resolve down to the individual structures can be prohibitively expensive, particularly in 3D. For that reason, many studies model arrays of densely packed structures via a homogenised porous layer. However, reproducing the correct kinematics of a compliant porous layer is a challenge. Therefore, the approach adopted here resolves the fluid and structural dynamics down to the individual structures. This is achievable since the problem considered is 2D and only a modest number of flaps are tested. First, the numerical approach is validated against experimental data obtained for the selected configuration. Then, to elucidate the physical mechanisms involved, the influence of the mass ratio and its effect on the coupled behaviour is analysed in detail. In particu-

lar, the previously identified wave-like motion is scrutinised and a quantitative characterisation of this mechanism is sought. The oscillating channel flow case is well suited to identifying phase lags between adjacent flaps and delineating their time-dependent behaviour. The numerical approach adopted in this work is based on a lattice Boltzmann-immersed boundary method coupled to an Euler-Bernoulli beam solver for the flap dynamics. After introducing and validating the numerical model, results are presented for a range of mass ratios, with a particular focus on the honami/monami wave-like motion.

8.3 Methods

8.3.1 Lattice Boltzmann Method

The lattice Boltzmann method (LBM) is used to model the fluid dynamics. Unlike traditional Navier-Stokes-based solvers, which operate directly on the macroscopic quantities, the LBM relies on a mesoscopic description of the flow. The governing equation in the LBM is the Boltzmann equation:

$$\frac{\partial f}{\partial t} + \boldsymbol{\xi} \cdot \nabla_{\mathbf{x}} f + \mathbf{f} \cdot \nabla_{\boldsymbol{\xi}} f = \Omega(f) \tag{8.1}$$

where f is the probability density function, $\mathbf{x}$ are the spatial coordinates, $\boldsymbol{\xi}$ is the molecular velocity, $\mathbf{f}$ is an external force vector, and Ω is the collision operator. The force term, $\mathbf{f}$, is crucial in the present work as it provides the coupling between the fluid and structural dynamics, through the immersed boundary method (to be discussed later).

The BGK approximation is used to simplify the collision operator. This approximation relies on the assumption that the probability density function relaxes towards a local equilibrium:

$$\Omega = \frac{1}{\tau}(f^{eq} - f) \tag{8.2}$$

where τ is the relaxation time-scale and f^{eq} is the equilibrium function. The equilibrium function depends only on local macroscopic quantities and can be obtained through a Taylor series expansion of the Maxwell-Boltzmann distribution [126].

Equation 8.2 is substituted into Equation 8.1 before discretising in velocity, space, and time. The discrete lattice form of the forcing term is given by [130]:

$$F_i = w_i \left(1 - \frac{1}{2\tau}\right) \left(\frac{\mathbf{c}_i - \mathbf{u}}{c_s^2} + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^4} \mathbf{c}_i\right) \cdot \mathbf{f} \tag{8.3}$$

where $\mathbf{c}_i$ represents the ith component of the discrete lattice velocities, w_i is a velocity-specific weighting factor, and c_s is the lattice speed of sound, which all take standard values.

8.3.2 Immersed Boundary Method

The immersed boundary method (IBM) is used to incorporate the moving boundaries of the flaps. The main idea of the IBM is to solve the fluid equations on a fixed Cartesian grid, which is non-conforming to the boundary. The presence of the structure is then incorporated through a set of forces, which have the effect of imposing the no-slip condition on the surface of the boundary [154, 164]. The approach adopted in the present work is the direct forcing approach [160], whereby a predicted velocity field, $\mathbf{u}^*$, is first computed by neglecting the presence of the structure. Since the boundary is non-conforming, this predicted velocity field must be interpolated from the fluid mesh onto the boundary markers, located at $\mathbf{X}$:

$$\mathbf{U}^*(\mathbf{X}, t) = \mathcal{I}[\mathbf{u}^*(\mathbf{x}, t)] \tag{8.4}$$

Using this predicted velocity field, the restoring force required to satisfy the no-slip

condition can be calculated via:

$$\mathbf{F}(\mathbf{X},t) = \frac{\mathbf{U}^b(\mathbf{X},t) - \mathbf{U}^*(\mathbf{X},t)}{\Delta t} \tag{8.5}$$

where $\mathbf{U}^b$ is the velocity of the boundary. This force is then transferred back to the fluid grid via a spreading operator:

$$\mathbf{f}(\mathbf{x},t) = \mathcal{S}[\mathbf{F}(\mathbf{X},t)] \tag{8.6}$$

Finally, the fluid equations are solved once more (Equation 8.1), but this time including the IBM force term.

8.3.3 Euler-Bernoulli Beam Model

The dynamics of the flexible flaps are calculated via the Euler-Bernoulli beam equation. In dimensionless form this is given as:

$$\frac{\partial^2 \mathbf{X}}{\partial t^2} = \frac{\partial}{\partial s}\left(T\frac{\partial \mathbf{X}}{\partial s}\right) - K_B\frac{\partial^4 \mathbf{X}}{\partial s^4} + Fr\frac{\mathbf{g}}{g} - \mathbf{F} \tag{8.7}$$

where T is the tension and K_B is the flexural rigidity. The reference quantities used in the nondimensionalisation are: $T_{ref} = \Delta\rho u_\infty^2$ is a reference tension, $K_{Bref} = \Delta\rho u_\infty^2 L^2$ is a reference bending rigidity, and $F_{ref} = {}^{\Delta\rho u_\infty^2}/_{L\epsilon\rho_f}$ is a reference force. The density difference per unit length between the structure and fluid, $\Delta\rho$, is the controlling parameter in this investigation and its effect is examined over a range of values.

Equation 8.7 is closed via the inextensibility condition:

$$\frac{\partial \mathbf{X}}{\partial s} \cdot \frac{\partial \mathbf{X}}{\partial s} = 1 \tag{8.8}$$

This is satisfied through the tension term in Equation 8.7 and ensures no axial extension of the flaps. The boundary conditions at the free end are:

$$T = 0, \qquad \frac{\partial^2 \mathbf{X}}{\partial s^2} = (0,0), \qquad \frac{\partial^3 \mathbf{X}}{\partial s^3} = (0,0) \tag{8.9}$$

and at the fixed end the clamped condition is set via:

$$\mathbf{X} = (X_0, Y_0), \qquad \frac{\partial \mathbf{X}}{\partial s} = (\cos(\theta), \sin(\theta)) \tag{8.10}$$

The resulting equations are discretised using a staggered finite difference scheme and the nonlinear solution is obtained via the Newton method. A detailed explanation of the fully coupled model has already been provided elsewhere [110].

The Euler-Bernoulli model described above breaks down when the fluid and structural densities are equal (i.e. when $\Delta\rho = 0$). This meant that during the validation stage in Section 8.5, the structural density had to be artificially increased to achieve a stable solution. However, it will be shown in Section 8.5 that even with this modification, the present model can still reproduce the system dynamics to a reasonable degree. Furthermore, since the purpose of this work is to examine how the coupled behaviour changes with mass ratio, its absolute value is not crucial. For this reason, it is not expected that this issue has a significant effect on the conclusions of the present work.

8.4 Case Description

Figure 8.1 displays a schematic of the experimental setup, which consists of an oscillating channel flow with a square cross section ($6\,\text{cm} \times 6\,\text{cm}$) over a row of 10 flexible flaps of length $L = 2\,\text{cm}$, where $H = 3L$. The test section is filled with a water/glycerine solution, which has a kinematic viscosity of $\nu = 1 \times 10^{-4}\,\text{m}^2/\text{s}$ and a mass density of $\rho_f = 1200\,\text{kg/m}^3$. The flaps are made of silicone rubber ($E = 1.2\,\text{MPa}$, $\rho_s = 1200\,\text{kg/m}^3$) so that they are easily deflected by the flow. The flaps have a thickness of $h = 1\,\text{mm}$ and a span of $b = 5\,\text{cm}$. As such, they

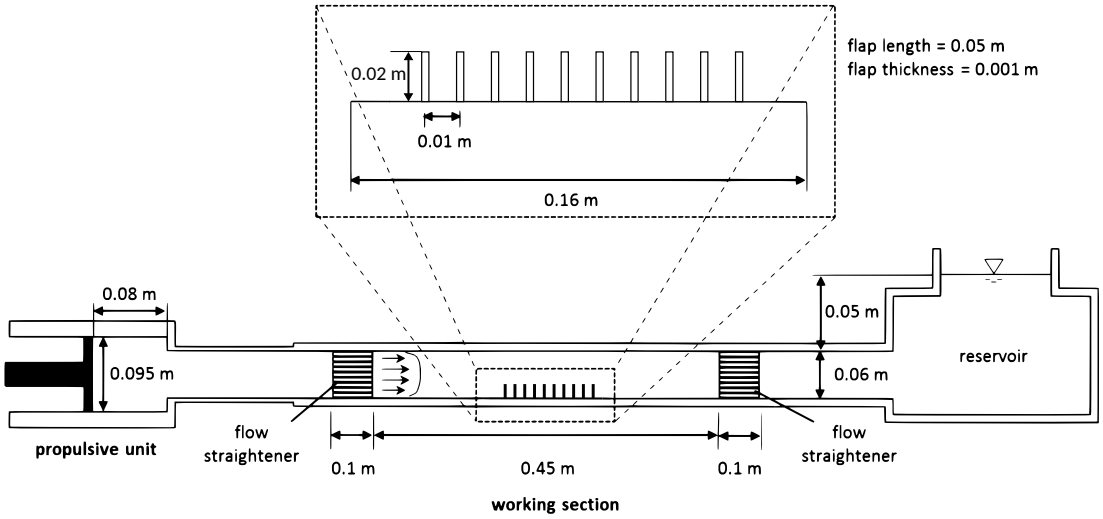

Figure 8.1: Schematic of the experimental setup.

extend across most of the channel width; this means the resulting flow field can be considered quasi-2D at the channel centreline. The fluid and structural densities are set equal so that no buoyancy effects are present. The flow is driven by a piston pump which generates an oscillating flow with a Womersley profile. The Reynolds number is approximately 120, based on the maximum velocity and flap length.

8.5 Validation

A comprehensive validation for each component of the numerical scheme, as well as the fully coupled model, has already been provided elsewhere [110]. In the present work, the availability of experimental data provides an opportunity for further validation. The simulation was set up to match the baseline case described in Section 8.4, except for the structural density, which was artificially increased to achieve a stable solution. Figure 8.2 shows the tip motion histories of the first three flaps for increasing grid resolution, where N_{IBM} is the number of IBM markers used to discretise each flap. There is little variation in the tip motion over the range of tested grid resolutions. Also shown is the relative error at each grid resolution, measured with respect to the finest grid level ($N_{IBM} = 40$). The results show approximately second-order convergence with grid resolution. All results in the following sections were obtained with $N_{IBM} = 35$.

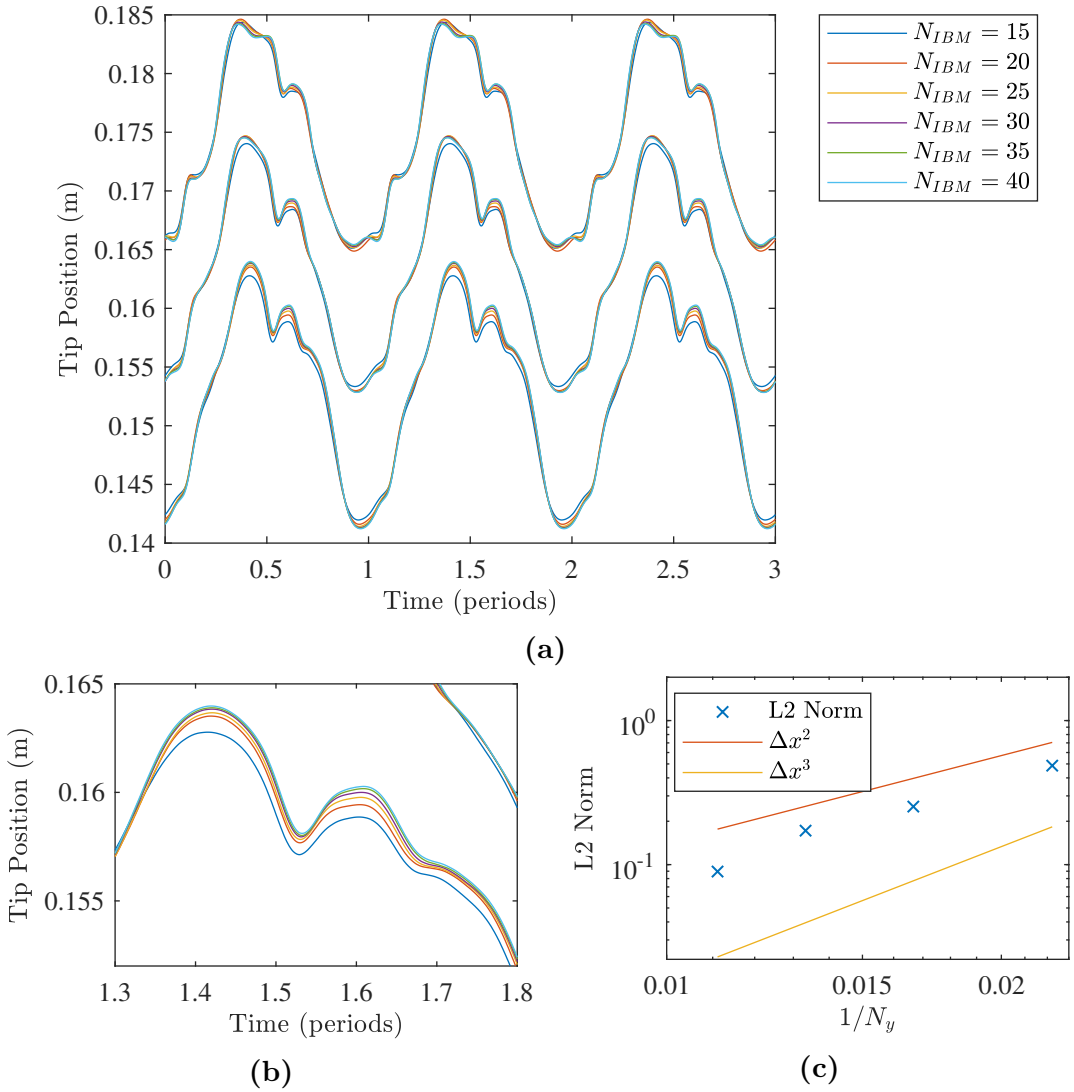

Figure 8.2: Refinement study for coupled FSI solver. (a) Streamwise tip motion histories of the centre flaps for increasing grid resolution. (b) Zoomed view. (c) L2 norm of error measured with respect to the finest resolution ($N_{IBM} = 40$).

Figure 8.3 shows the streamwise (in the direction of the flow) tip deflection for each of the 10 flaps over three periods. These values represent the final three periods after the motion reaches a steady periodic behaviour (approximately 10 cycles). The agreement between the experimental and numerical results is good, although some differences are noticeable. Specifically, small 'kicks' in the tip motion observed in the experiment are not as pronounced in the numerical solution, especially for the outer flaps. This may be attributed to small differences in the approximations of the structural parameters between the experiment and simulation. Another possible explanation is 3D effects due to the finite span of the flaps, which are not considered in the numerical solution. Nevertheless, it is

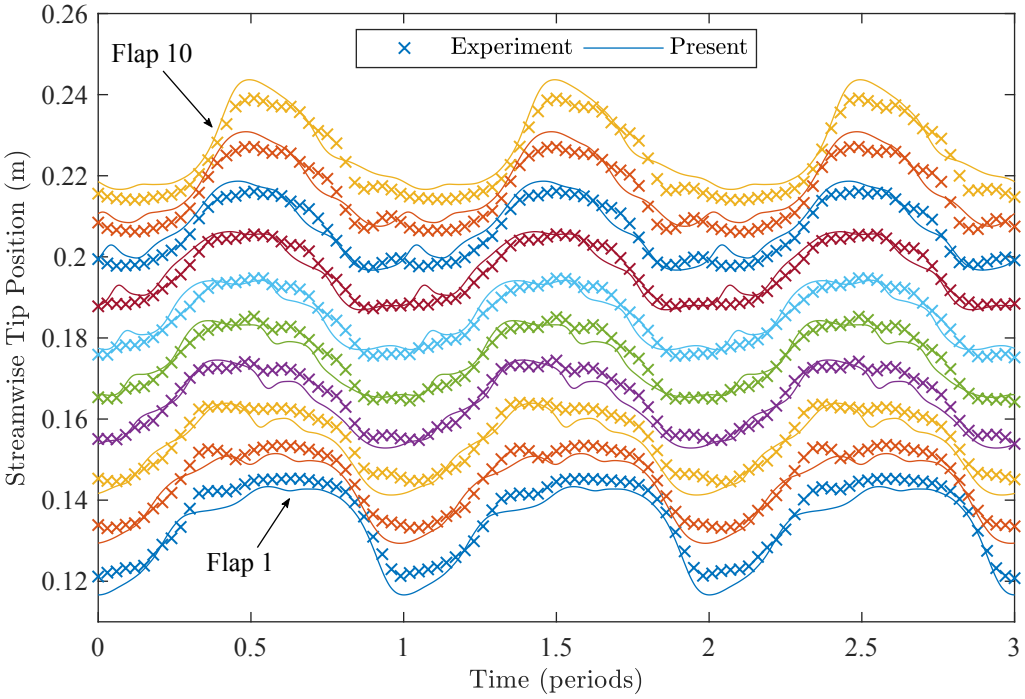

Figure 8.3: Experimental and numerical tip deflection for each flap over three periods.

clear to see the numerical model can reproduce the main dynamics very well, and these 'kicks' in the tip motion are better resolved in the flaps towards the centre of the array.

8.6 Variation of Mass Ratio

This section investigates the effect of mass ratio, μ, on the coupled dynamics of the system, where $\mu = {}^{\Delta\rho}/_{\rho_f L}$. As well as the baseline case, four additional mass ratios are tested. After normalising with respect to the baseline case, the values of the mass ratio are: $\mu^* = {}^{\mu}/_{\mu_{base}} = [0.66, 0.8, 1.0, 1.6, 4.0]$, where μ_{base} is the mass ratio from the baseline case. Only the mass ratio is varied while everything else is held constant, including the dimensionless bending stiffness $K_B = {}^{EI}/_{\Delta\rho u_\infty^2 L^2}$. The modification of the mass ratio alters the natural frequency of the flaps, which in turn will modify the structural response to the imposed flow frequency. During preliminary tests a range of bending stiffness values were also studied. However, the results for increasing bending stiffness were qualitatively similar to the results for increasing mass ratio, and so only the mass ratio tests are examined here.

This preliminary phase was also used to select the range of tested mass ratio. Starting from the baseline value, a number of cases around this value were tested and, based on the results of this parameter sweep, the most pertinent cases were selected for detailed analysis.

Figure 8.4 shows snapshots of the flap profiles during one full cycle. Clearly, variations in the dynamics exist not only across different mass ratios, but also through the array itself.

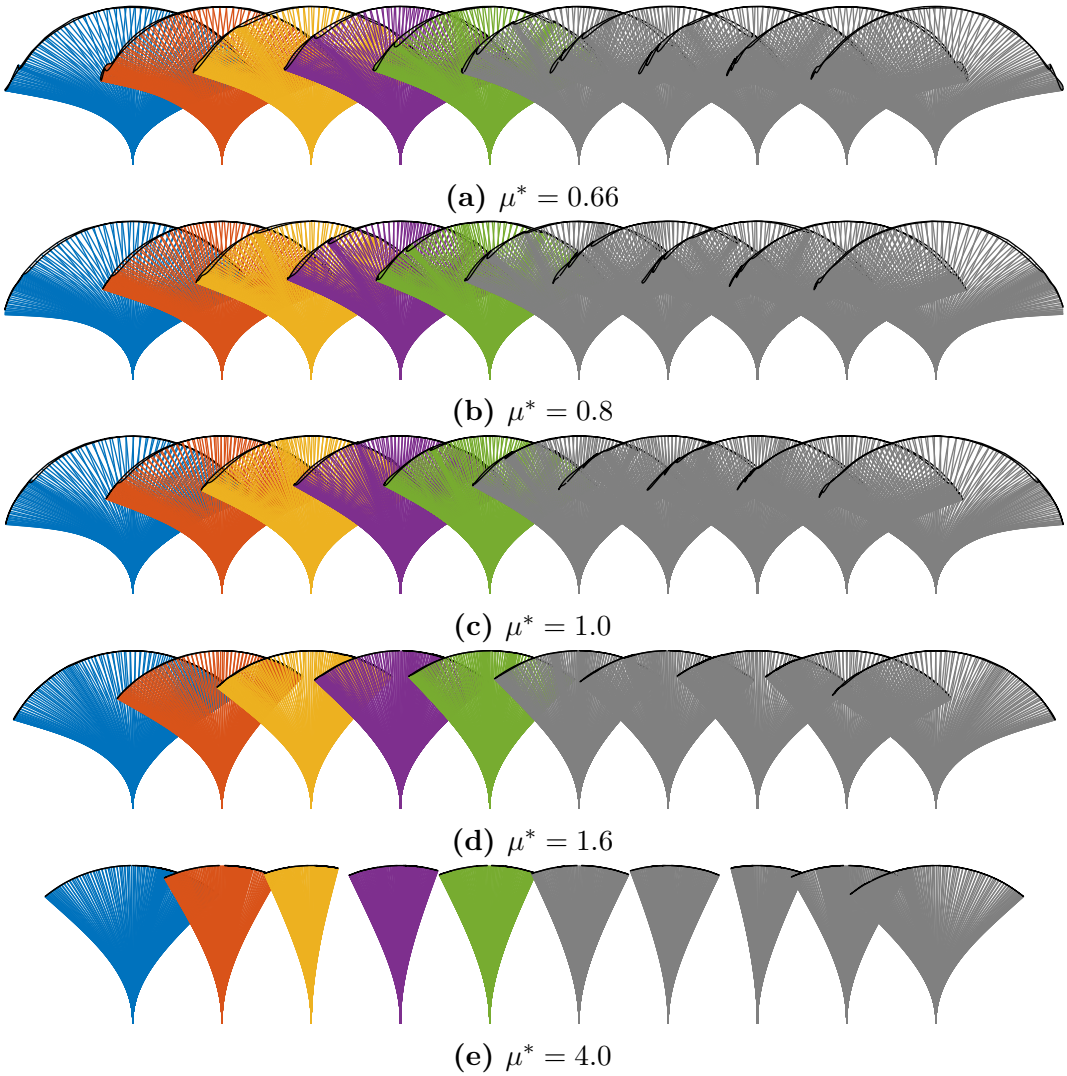

(a) $\mu^* = 0.66$

(b) $\mu^* = 0.8$

(c) $\mu^* = 1.0$

(d) $\mu^* = 1.6$

(e) $\mu^* = 4.0$

Figure 8.4: Flap shapes during one period for each mass ratio. Only the first five flaps are coloured.

8.6.1 Flow Dynamics

Figure 8.5 shows the streamwise velocity profiles at five locations along the channel. In regions sufficiently far away from the array, the velocity follows the normal Womersley profile. However, near the array the velocity profiles are heavily distorted. In the centre of the row the bulk flow above the array resembles a skewed Womersley profile, with a near-zero velocity within the array itself. At the edges the profiles are asymmetric themselves but antisymmetric to each other. Furthermore, at these locations the velocity profiles within the array seem to suggest a net flow towards the centre.

Figure 8.6 shows the streamwise velocity over three cycles at six probe locations. These locations are given by positions P1–P5 and P11 in Figure 8.7, which are located at a height of $1.1L$. This height was selected based on inspection of the flow dynamics, which showed a rich set of flow behaviour within this region. Spectral analysis shows that these time histories have a dominant frequency matching the forcing frequency (1 Hz), with a second peak at the first harmonic (2 Hz).

Figures 8.8–8.10 show snapshots of the instantaneous vorticity through one half-cycle for $\mu^* = [0.66, 1.6, 4.0]$. The results for $\mu^* = [0.8, 1.0]$ were found to be qualitatively similar to $\mu^* = 1.6$ and are therefore omitted. The flap motion is strongly affected by the primary bulk vortex, which is shed from the edges of

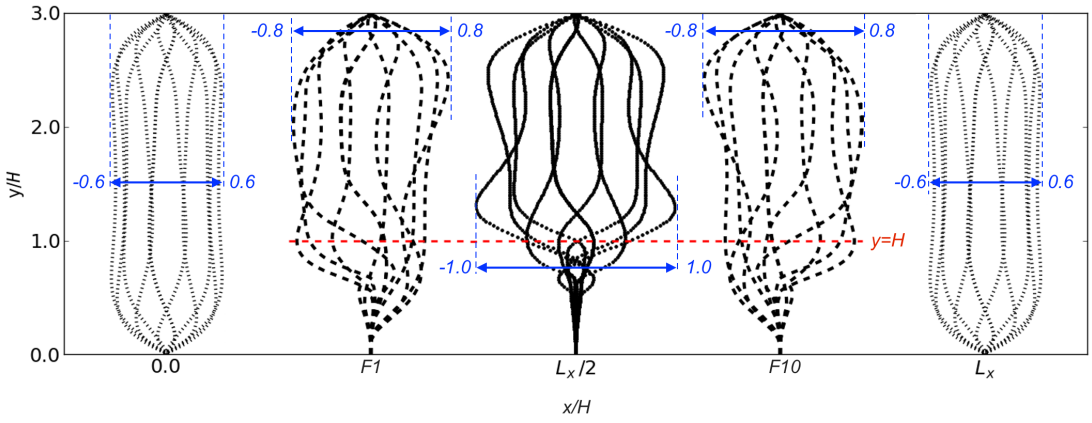

Figure 8.5: Streamwise velocity profiles at various locations along the channel over one oscillation cycle. From the left to right the positions are (i) channel entrance; (ii) flap 1; (iii) channel centre; (iv) flap 10; (v) channel exit. Red line indicates the height of the flaps and blue lines indicate nondimensional range of velocity profiles.

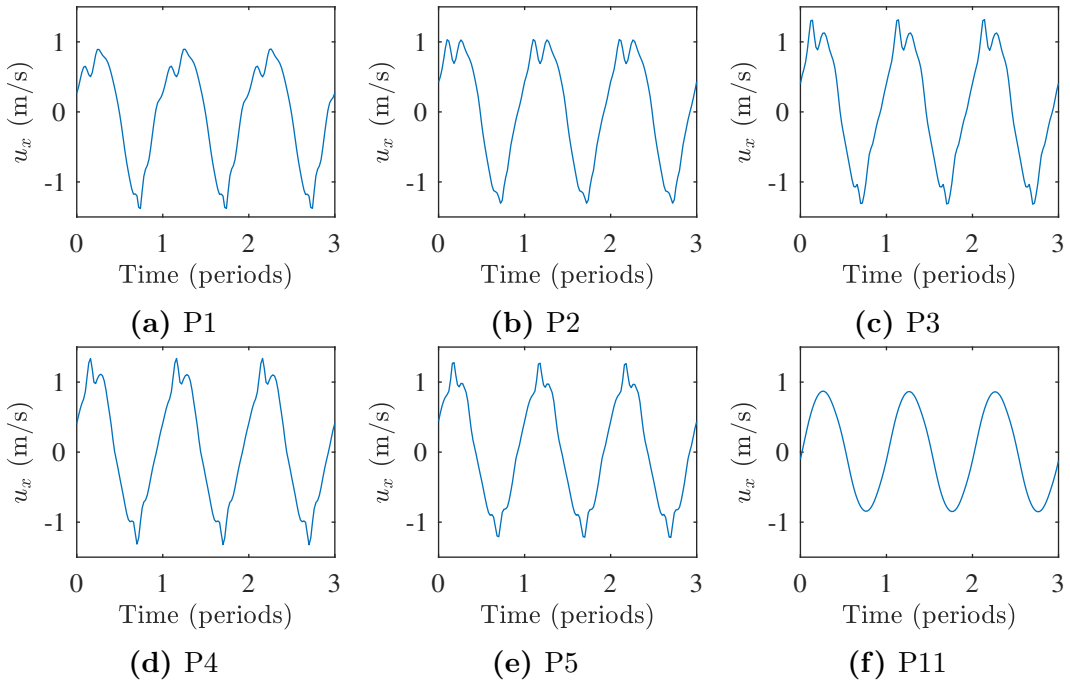

Figure 8.6: Streamwise velocity histories at six probe locations along the channel.

the array, and the mass ratio influences this interaction. Examining the velocity contours, at low mass ratio the bulk flow passes over the array without significant deflection. However, at larger mass ratios the flaps yield less to the flow. This exacerbates the blockage effect and accelerates the flow upwards away from the array, reducing the chance of a lock-in behaviour. These observations are supported by the traces of the tip motion, given in Figure 8.11. Further to this, at large mass ratios the recirculation at the trailing edge is more pronounced and the magnitude of vorticity is increased. This deflects the outermost flap back towards the array. Moreover, the low-pressure region in the vortex core appears to act to detach the flow from the top wall of the channel.

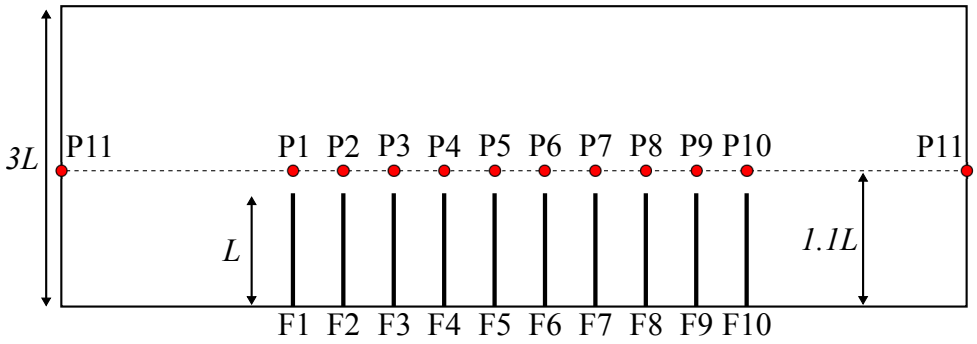

Figure 8.7: Computational domain with probe locations.

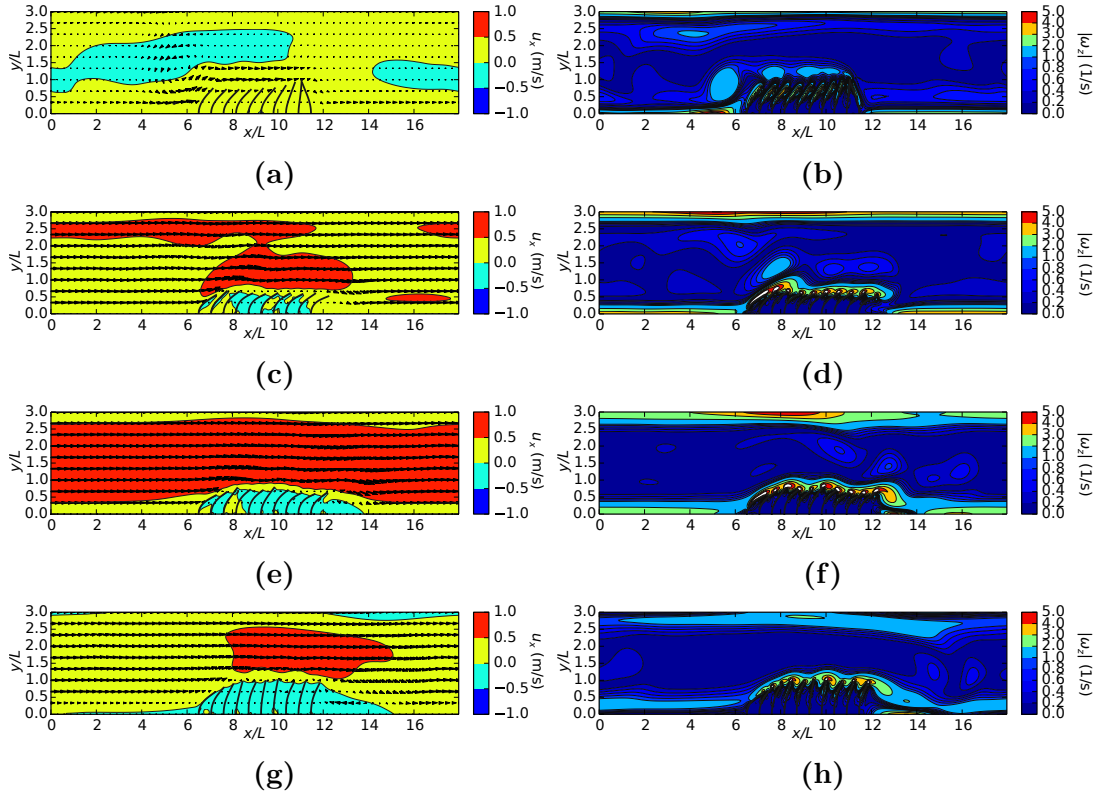

Figure 8.8: Contours of instantaneous streamwise velocity (a,c,e,g) and out-of-plane vorticity (b,d,f,h) for $\mu^* = 0.66$ over one half-cycle.

For large mass ratios the flaps stay well separated from each other and do not come into contact. However, at the lower mass ratios there are short instances where the flaps come within close proximity. A collision model is not incorporated in this work and therefore there is nothing to prevent the flaps from making contact with each other. However, there does not seem to be any adverse behaviour from these events, since the flaps are capable of separating themselves without getting stuck. For cases where the contact frequency is high and substantial, a collision model would be required. However, in the current work these events are rare and thus it is not expected they will have a significant impact on the present conclusions.

8.6.2 Flap Dynamics

Figure 8.11 shows the normalised and superimposed trace of the streamwise tip deflection for $\mu^* = [0.66, 1.6, 4.0]$. Only the first five flaps are shown since the motions of the other flaps are antisymmetric. Also shown is the power spectrum of the trace for each flap, where the forcing frequency of the fluid is 1 Hz. For

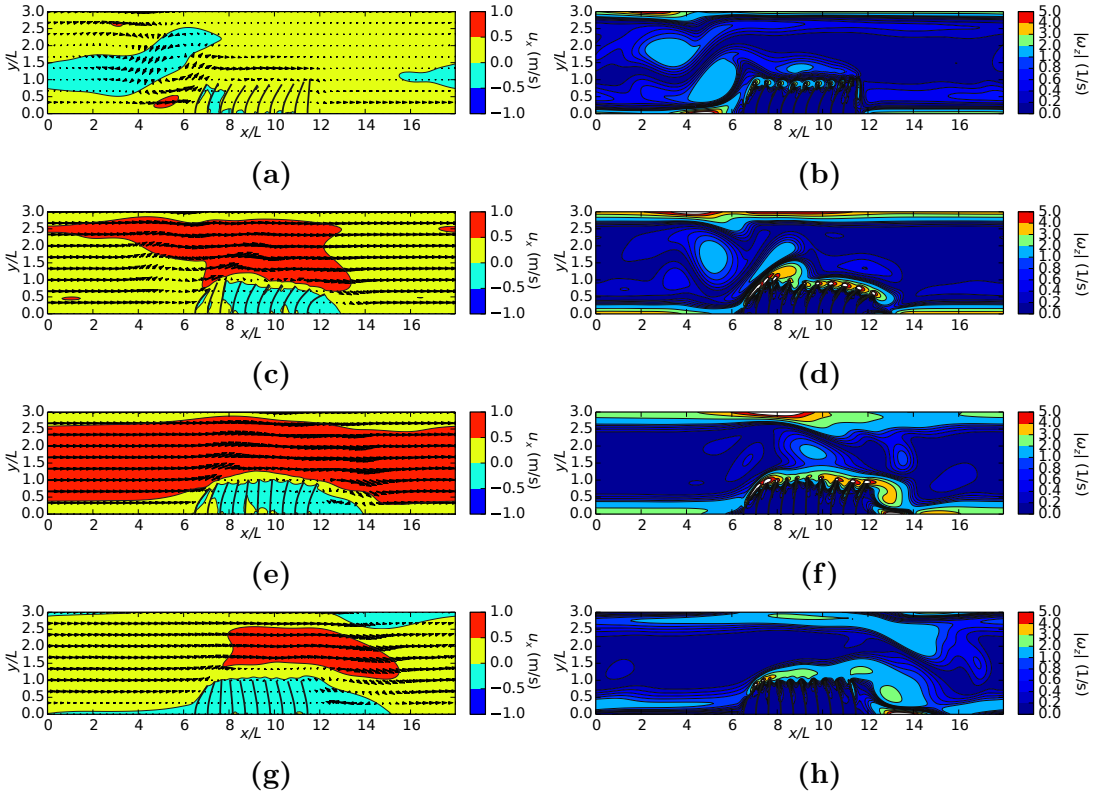

Figure 8.9: Contours of instantaneous streamwise velocity (a,c,e,g) and out-of-plane vorticity (b,d,f,h) for $\mu^* = 1.6$ over one half-cycle.

$\mu^* = 0.66$, the flap motion is relatively coherent through the array, with all flaps exhibiting a similar range of motion. Furthermore, it is clear to see the main mode of motion is the 1 Hz dynamic, which is a result of the imposed fluid motion. A phase lag between the flaps is also evident; however, this is relatively small compared to the intermediate mass ratio case ($\mu^* = 1.6$). At the intermediate mass ratio, the range of motion is reduced and there is a marked drop in coherence between the flaps, particularly in the second half of the cycle where sharp 'kicks' in the tip motion are clearly observable. This trend continues to the largest mass ratio ($\mu^* = 4.0$), where over the entire cycle there seems to be almost no correlation between the neighbouring flaps or the driving fluid. These observations are also clearly demonstrated in the power spectra, which show a significant drop in amplitude at 1 Hz as the mass ratio is increased.

Figure 8.12 plots the normalised amplitudes of the power spectral density at the driving frequency, and its first harmonic, against mass ratio. The ratio between these two amplitudes, α_s, is also shown. There is a decrease in power at the driving frequency as the mass ratio increases, as shown in Figure 8.12a. This

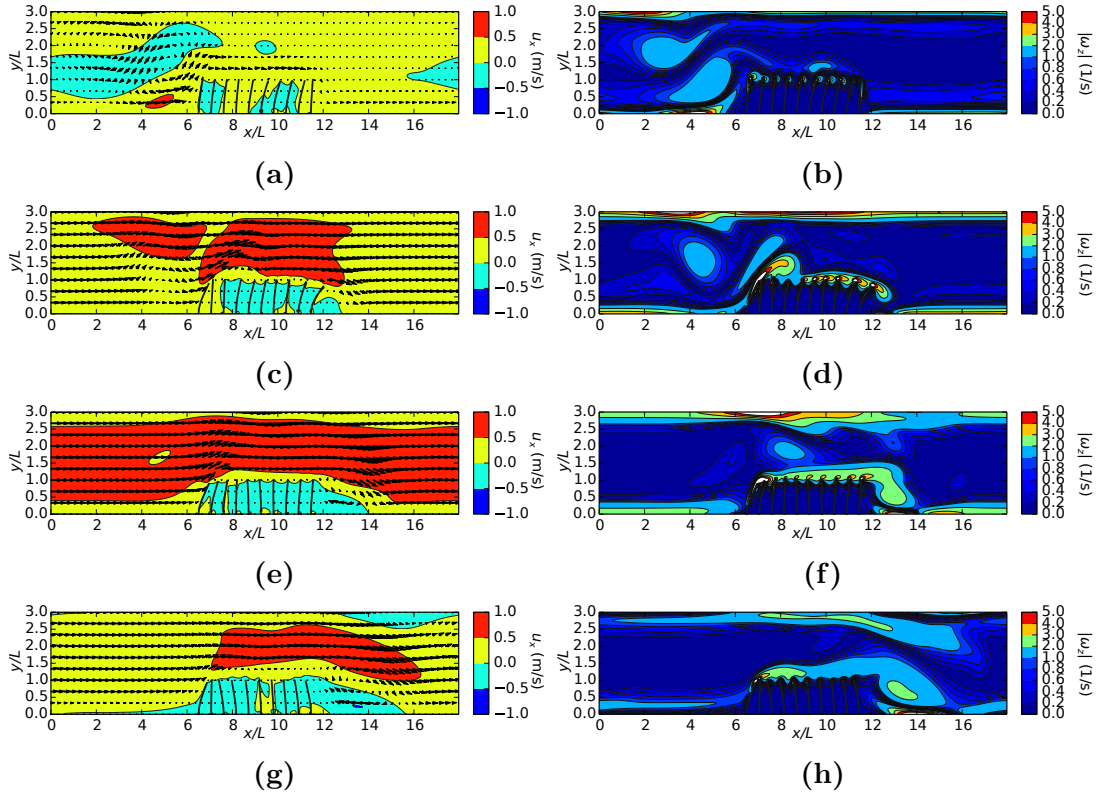

Figure 8.10: Contours of instantaneous streamwise velocity (a,c,e,g) and out-of-plane vorticity (b,d,f,h) for $\mu^* = 4.0$ over one half-cycle.

happens rapidly at first, before seemingly levelling-off. For all cases, the range of motion of the first flap is larger than the rest of the flaps, and the relative difference increases with increasing mass ratio. The rest of the flaps all show a similar range of motion to each other. In contrast, the power distribution across the array at the first harmonic shows a broad range of values for the different mass ratios (see Figure 8.12b). At the lowest mass ratio, the power varies almost linearly across the array. Furthermore, the first, fourth, and fifth flaps show almost constant values over the range of tested mass ratios. The second and third flaps show an initial increase in power until they peak at $\mu^* = 1.6$, before decreasing again at the largest mass ratio. This is related to the 'kick' motion observed in the tip trace and suggests an amplification of this effect in the region of these two flaps. Figure 8.12c plots the ratio between the two amplitudes and shows that, as mass ratio is increased, more energy is manifested in the higher frequency spectrum. This is also clearly observable in the tip trace (Figure 8.11).

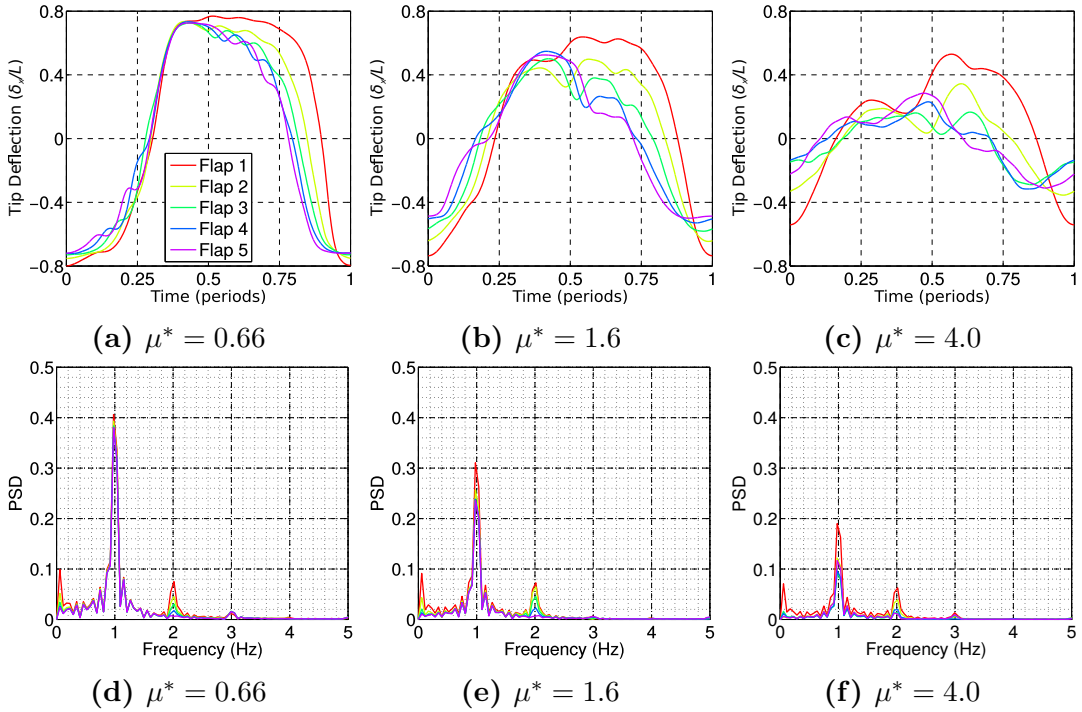

(a) $\mu^* = 0.66$ **(b)** $\mu^* = 1.6$ **(c)** $\mu^* = 4.0$

(d) $\mu^* = 0.66$ **(e)** $\mu^* = 1.6$ **(f)** $\mu^* = 4.0$

Figure 8.11: Normalised and superimposed trace of streamwise tip deflection and the resulting power spectra for the first five flaps.

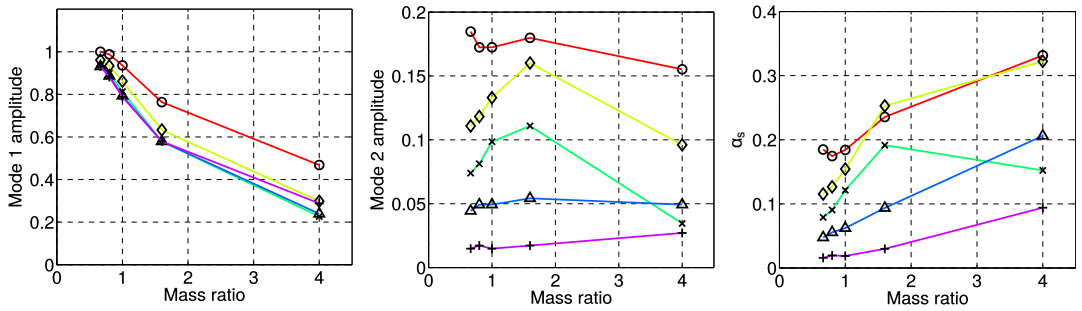

(a) **(b)** **(c)**

Figure 8.12: Normalised amplitude of the flap motion against mass ratio. Amplitude of primary frequency (a); amplitude of first harmonic (b); ratio of amplitudes (c). Legend corresponds to Figure 8.11.

8.6.3 Spectral Analysis

This section aims to correlate the dynamics of the flaps with the near-field fluid velocity. In addition to α_s, a second parameter (α_f) is introduced which describes the ratio between the dominant amplitudes of the power spectral density of the streamwise fluid velocity at the probe locations (P1–P5). Furthermore, a third

parameter is also introduced to measure the asymmetry of the flap motion:

$$\Delta t = |\Delta t_2 - \Delta t_1| \tag{8.11}$$

where Δt_1 is the amount of time the flap spends deflected to the right of its equilibrium position during one period, and Δt_2 is the time spent deflected to the left. Clearly, if $\Delta t = 0$ then the flap spends an equal amount of time deflected both left and right.

Figure 8.13 shows the variation of α_f, α_s, and Δt for each flap across the five different mass ratios. These values all show a similar trend: relatively large values at the extremities of the array which decrease towards the centre. At low-moderate mass ratios the quantities are remarkably close. This is especially true for the baseline case, which was specifically chosen to exhibit a lock-in behaviour between the fluid and the flaps. Towards the centre of the array, all cases show a decrease in Δt towards zero. This indicates a symmetric motion for the flaps in this region.

For $\mu^* = 1.6$, the flaps at the centre of the array show almost the same ratio between the primary and secondary frequencies. However, towards the edges there is a loss in correlation due to the presence of the primary vortex which is shed from the array, particularly for the second and ninth flaps. This is confirmed by examining Figure 8.9. At $\mu^* = 4.0$ this coherence becomes even weaker, leading to almost completely uncorrelated dynamics between the fluid and the flaps. In this case the primary vortex is deflected upwards away from the array, thereby reducing its effect on the flaps. This can be seen in Figure 8.10, which is in contrast to Figure 8.8, where the vorticity is transmitted through the array and therefore leads to a more coherent motion.

8.7 Conclusions

This study has investigated the coupling mechanism between an oscillating channel flow and a row of wall-mounted flexible flaps. The phase difference, which leads to the honami/monami waving motion, has been characterised by means of several analyses. This has led to several observations regarding the coupled behaviour.

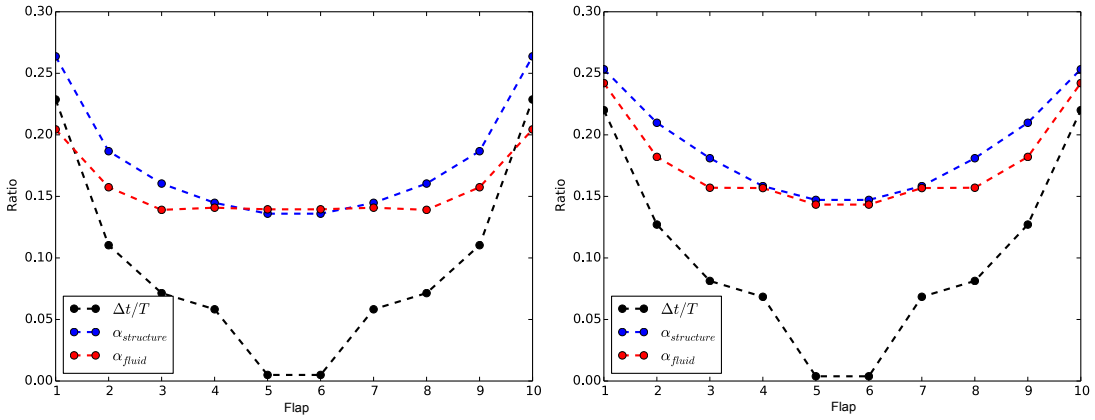

(a) $\mu^* = 0.66$

(b) $\mu^* = 0.8$

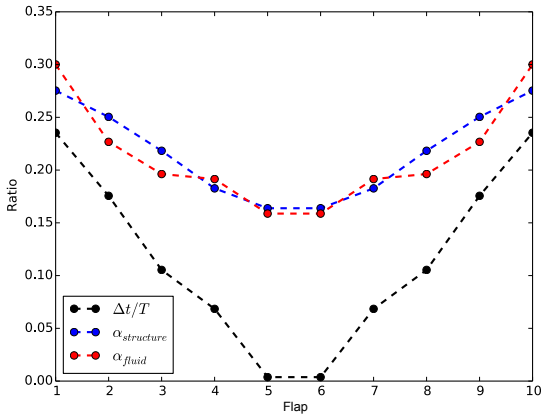

(c) $\mu^* = 1.0$

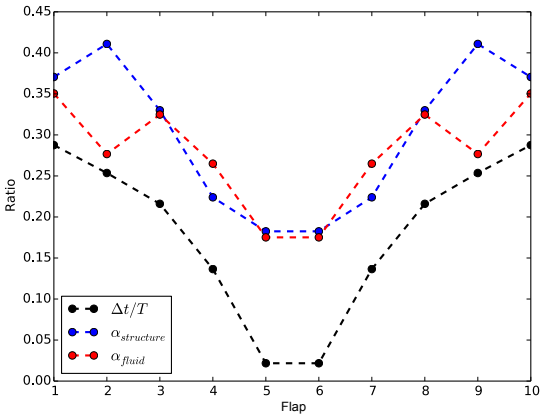

(d) $\mu^* = 1.6$

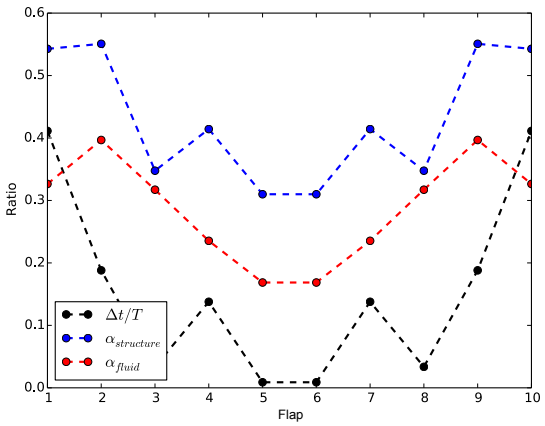

(e) $\mu^* = 4.0$

Figure 8.13: Spectral ratios and asymmetry metric for each flap across all mass ratios. Spectral ratios for the fluid are measured at probe locations P1–10 in Figure 8.7.

The primary bulk vortex, which occurs at the leading and trailing edges of the array, is the cause of the phase shift between adjacent flaps and leads to a smoothly varying phase difference through the array. By analysing the power spectra of

the tip motion and near-field fluid velocity, the coupled interactions between the fluid and the flaps have been quantified. In the region of the array, the spectral ratios and the asymmetry in the trace of the tip motion are found to be similar. This indicates the interaction between the fluid and the flaps is indeed coupled, as opposed to a purely fluid-driven response.

Over the course of the cycle there is a net flow from the outer flaps towards the inner flaps. As a result, the outer flaps spend most of the cycle deflected inwards towards the centre of the array. For small mass ratios there is a strong coupling between the motion of the fluid and the motion of the flaps. However, as the mass ratio is increased, the coupling breaks down and the coherence between the fluid and the flaps is reduced. For large mass ratios, there is almost no correlation between the flap and fluid motion, since the bulk flow is deflected upwards and away from the array, thereby reducing the opportunity for interaction. These results show there seems to be a fairly narrow region where the fluid and structural motions are coupled. This has important implications for flow control. However, due to the limited parameter space examined in this study, it is not possible to identify an optimum mass ratio based on these results and further work is required to link these findings to unsteady bluff body wakes.

Chapter 9

Large Array of Multiple Flaps in an Open-Channel Flow

This chapter is based on the following article:

J. O'Connor and A. Revell. Dynamic Interactions of Multiple Flexible Flaps. *In Preparation*, 2018

which is currently in preparation and is planned for journal publication. Following from the two previous chapters – which build in scale from a single flap in a periodic array to a small finite array – this chapter examines a large array of 128 flaps under steady flow conditions. The main goal of this chapter is to characterise the behavioural states of an array of flaps over a wide range of material properties and compare it to the behaviour of a single flap. In particular, the waving mode (honami/monami) and its associated lock-in are examined in detail. The results show that the material properties of the array have a strong effect on the dynamic response of the flaps. This implies that the waving mode is in fact a coupled response between the fluid and the array, as opposed to a purely fluid-driven instability. Furthermore, an additional mode, similar to the flapping mode exhibited by single flaps, is observed for low bending rigidities.

My contributions to this paper are as follows: developing, implementing, and validating the numerical model, setting-up and running all of the simulations, post-processing, analysis, and writing the paper.

9.1 Abstract

Coherent waving interactions between vegetation and fluid flows are known to emerge under certain conditions. This behaviour is associated with the mixing layer instability that arises due to the canopy drag. A similar waving motion has also been observed in flow control applications, where arrays of slender structures are used to augment the wake topology behind bluff bodies. While their existence is well reported, the mechanisms which govern this behaviour, and their dependence on structural properties, are not yet fully understood. This work investigates the coupled interactions of a large array of slender structures in an open-channel flow via numerical simulation. A direct modelling approach, whereby the individual structures and their dynamics are fully resolved, is realised via a lattice Boltzmann-immersed boundary-finite element model. For steady flow conditions at low-moderate Reynolds number ($Re = 80$), the dynamic response of the array is measured for a range of structural material properties. Both the mass ratio and bending rigidity are varied over a range spanning two orders of magnitude, and the ensuing response is characterised and quantified. The results show that a broad range of behaviours are exhibited, which depend on the structural properties. These behaviours are classified into distinct states: static reconfiguration, regular waving, irregular waving, flapping, and static with small deflection. Furthermore, the regular waving regime is found to occur when the natural frequency of the array approaches the estimated frequency of the mixing layer instability. These findings indicate that the coherent waving motion observed in vegetation and flow control studies is in fact a coupled instability, as opposed to a purely fluid-driven response, and that this specific regime is triggered by a lock-in between the fluid and structural natural frequencies.

9.2 Introduction

The dynamic interactions between fluid flows and arrays of slender structures play an important role in a broad range of physical processes. In nature, the metachronal motion of dense arrays of cilia serve a number of crucial physiological functions, from cell propulsion [1] to particle regulation [24]. In industrial applications, bioinspired designs based on slender structures are finding uses in areas such as flow control [2, 9] and sensing [30]. Furthermore, the interactions between

vegetation and fluid flows have important consequences in terms of agriculture [8] and coastal protection [35].

In vegetative flows, the emergence of coherent waving motions between adjacent plants is known to appear under certain conditions (see Figure 9.1). This phenomenon is known as honami for terrestrial flows, and monami for aquatic flows. This behaviour is initiated by the discontinuity in drag between the interstitial flow within the canopy and the bulk flow over the top of the canopy, which produces an inflection point within the velocity profile [76]. The inflection point makes the flow susceptible to Kelvin-Helmholtz instabilities, which roll up to form coherent vortical structures and propagate through the array [7]. The passage of the vortices over individual plants causes a local deflection, which propagates through the canopy with the flow structures. This is observed as a travelling wave through the array.

Ghisalberti [7] suggests that the honami/monami effect is strictly governed by the mixing layer instability, and that the waving motion is a forced response to the Kelvin-Helmholtz vortices which are generated due to the inflectional velocity profile. To support this, Ghisalberti performed experiments on a scaled seagrass model and showed strong correlations between the velocity spectra, waving frequency, and predicted frequency of the Kelvin-Helmholtz vortices, which were estimated via:

$$f_{KH} = St_n \left(\frac{u_1 + u_2}{2\theta} \right) \tag{9.1}$$

where St_n is the natural Strouhal number associated with mixing layers, θ is the

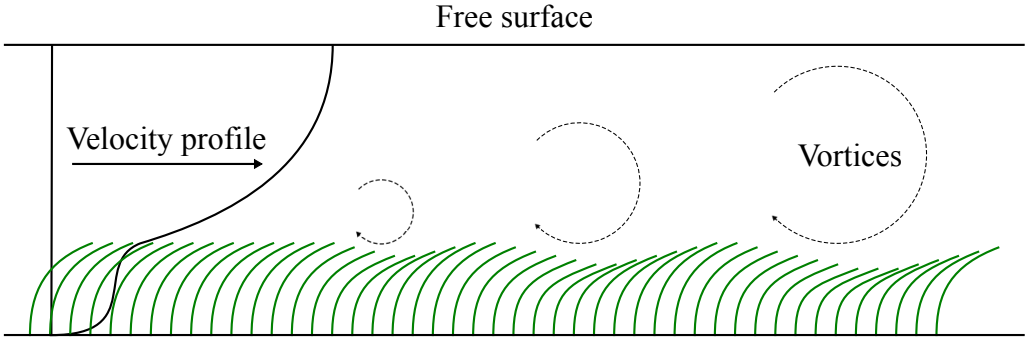

Figure 9.1: Velocity profile and monami motion over a submerged canopy.

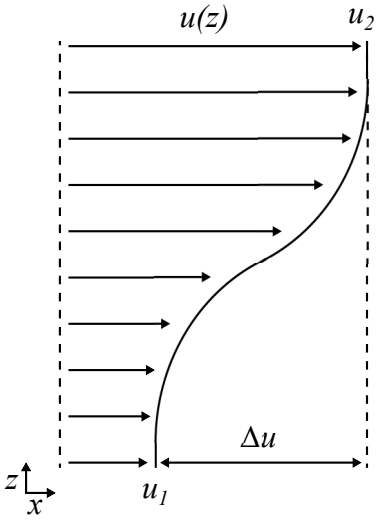

Figure 9.2: Mixing layer velocity profile with key quantities for Equation 9.1.

momentum thickness, and u_1 and u_2 are the low and high-stream velocities (see Figure 9.2). Through theory and experiment, the natural Strouhal number is known to be approximately $St_n \approx 0.032$ [79]. The momentum thickness is given by:

$$\theta = \int_{-\infty}^{\infty} \left[\frac{1}{4} - \left(\frac{u - \frac{1}{2}(u_1 + u_2)}{u_2 - u_1} \right)^2 \right] \mathrm{d}z \tag{9.2}$$

Ghisalberti also noted that the observed waving frequencies were much lower than the natural frequencies of the model plants. However, the values quoted for the natural frequency of the plant models suggest the theoretical (undamped) natural frequency was used for comparison. In the case of structures immersed in a flow, the (damped) natural frequency is likely to be lower.

More recent studies have suggested that the structural properties do in fact play a role in the waving instability. Py et al. [81] developed a simplified fluid model combined with a flexible porous layer for the canopy. While the main instability mechanism was still driven by the Kelvin-Helmholtz instability, the model indicated that the characteristics of this instability are significantly modified when the canopy compliance is considered. In a subsequent study, Py et al. [82] used an image-correlation technique [10] to extract the wavelengths and frequencies of the coherent structures over two different plant species (alfalfa and wheat). These two species were selected as they share similar geometric properties (height

and spacing) but different material properties (density and stiffness), and therefore different natural frequencies. Interestingly, for both species, the waving motion was found to occur near their own (damped) natural frequency. Moreover, this was observed over the full range of tested wind speeds. Extending their previous model [81], with increasing wind speed they observed a lock-in effect as the frequency of the Kelvin-Helmholtz instability approached the natural frequency of the canopy.

A similar lock-in effect has been observed in flow control applications, where arrays of slender structures are used to augment the wake topology behind bluff bodies. Favier et al. [2] used a homogenized porous layer model to examine the effect of a row of passive flaps attached to the aft of a circular cylinder. Under optimum conditions, they observed drag reductions of 15 % and reductions in the lift fluctuations of 40 %, which they attributed to a stabilisation of the wake. In the optimal regime, they observed a travelling wave through the array with a frequency that matched the vortex shedding frequency. Several other studies have also observed this travelling wave, and many report a lock-in at optimal conditions between the travelling wave frequency and the shedding frequency [9, 86, 87].

While several studies have observed these coherent waving interactions, the mechanisms which govern this behaviour, and their dependence on structural properties, are not yet fully understood. To address this issue, this work investigates the coupled interactions of a large array of slender structures in an open-channel flow via numerical simulation. A direct modelling approach, whereby the individual structures and their dynamics are fully resolved, is realised via a lattice Boltzmann-immersed boundary-finite element model. For steady flow conditions at low-moderate Reynolds number ($Re = 80$), the dynamic response of the array is measured for a range of structural material properties. Both the mass ratio and bending rigidity are varied over a range spanning two orders of magnitude, and the ensuing response is characterised and quantified. The following sections provide details and validation of the fully coupled numerical model. After introducing the case, results are then presented and the main findings discussed, with a particular focus on the coupled waving instability and its associated lock-in.

9.3 Methods

9.3.1 Lattice Boltzmann Method

The LBM has evolved over recent years to become an attractive alternative to the Navier-Stokes equations. Derived from kinetic theory, the LBM relies on a mesoscopic description of the fluid to compute its macroscopic behaviour. The driving equation behind the LBM is the Boltzmann equation, which in its discrete form is given by:

$$f_i(\mathbf{x} + \mathbf{c}_i\Delta t, t + \Delta t) - f_i(\mathbf{x}, t) = \frac{\Delta t}{\tau}\left[f_i^{eq}(\mathbf{x}, t) - f_i(\mathbf{x}, t)\right] + \Delta t F_i(\mathbf{x}, t) \tag{9.3}$$

where $\mathbf{x}$ are the spatial coordinates, t is time, $\mathbf{c}_i$ is the ith component of the lattice velocity vector, and $F_i(\mathbf{x}, t)$ is an external force discretised on the lattice. The probability distribution function, $f_i(\mathbf{x}, t)$, describes the proportion of fluid molecules within an elemental volume located at $\mathbf{x}$ and time t moving with velocity $\mathbf{c}_i$. Equation 9.3 can be split into two steps: a streaming step, where the components of the distribution function propagate through the lattice along discrete velocity links; and a collision step, where the distribution function relaxes towards a local equilibrium. The relaxation time-scale, τ, is related to the (lattice) viscosity via:

$$\nu = \left(\tau - \frac{1}{2}\right) c_s^2 \Delta t \tag{9.4}$$

where c_s is the lattice speed of sound. The equilibrium function, $f_i^{eq}(\mathbf{x}, t)$, is a function of local macroscopic quantities only, and can be obtained via a Taylor series [121, 122] or Hermite polynomial [119, 123] expansion of the Maxwell-Boltzmann distribution, giving:

$$f_i^{eq}(\mathbf{x}, t) = w_i \rho \left(1 + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^2} + \frac{(\mathbf{c}_i \cdot \mathbf{u})^2}{2c_s^4} - \frac{\mathbf{u} \cdot \mathbf{u}}{2c_s^2}\right) \tag{9.5}$$

where w_i is a velocity-specific weighting factor related to the discrete lattice. The macroscopic quantities, $\rho(\mathbf{x}, t)$ and $\mathbf{u}(\mathbf{x}, t)$ can be calculated by taking the leading moments of the distribution function:

$$\rho(\mathbf{x}, t) = \sum_i f_i(\mathbf{x}, t) \tag{9.6}$$

$$\rho\mathbf{u}(\mathbf{x}, t) = \sum_i \mathbf{c}_i f_i(\mathbf{x}, t) + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x}, t) \tag{9.7}$$

The force density, $\mathbf{f}(\mathbf{x}, t)$, is crucial for the present work as it provides the coupling between the fluid and structural dynamics. Its Cartesian form is discretised on the lattice to give [130]:

$$F_i(\mathbf{x}, t) = w_i \left(1 - \frac{1}{2\tau}\right) \left(\frac{\mathbf{c}_i - \mathbf{u}}{c_s^2} + \frac{\mathbf{c}_i \cdot \mathbf{u}}{c_s^4}\,\mathbf{c}_i\right) \cdot \mathbf{f}(\mathbf{x}, t) \tag{9.8}$$

9.3.2 Finite Element Method

The present model adopts the corotational formulation of the finite element method (FEM) to solve the structural dynamics. The main idea of the corotational formulation is to decompose the motion of each element within the assemblage into a rigid body motion and a purely deformational one [148]. This is achieved by adopting two reference configurations: the initial (or reference) frame, and the corotated (or local) frame which is attached to the element and moves with it.

The rigid body motion is measured with respect to the initial configuration and can be arbitrarily large. On the other hand, the elemental deformation is measured with respect to the corotated frame and is assumed to be small within the local frame. In this regard, two sets of quantities are defined: the local quantities which are measured with respect to the corotated frame, and the global quantities which are measured with respect to the global coordinate system. In the formulation of the system matrices which appear in the governing equations, the elemental matrices are first calculated in the corotated frame before being transformed to

the global coordinate frame, and then finally assembled into the global system matrices.

Following Bathe [149], the equilibrium conditions governing the nonlinear Newton-Raphson solution are:

$$\mathbf{M}\ddot{\mathbf{U}}^{n+1,k} + \mathbf{K}^{n+1,k}\Delta\mathbf{U}^{k} = \mathbf{R}^{n+1} - \mathbf{F}^{n+1,k} \tag{9.9}$$

$$\mathbf{U}^{n+1,k+1} = \mathbf{U}^{n+1,k} + \Delta\mathbf{U}^{k} \tag{9.10}$$

where n and k are the time step and iteration counters, $\mathbf{M}$ is the mass matrix, $\mathbf{K}$ is the tangent stiffness matrix, $\mathbf{R}$ is the external load vector, $\mathbf{F}$ are the internal forces within the structure, $\ddot{\mathbf{U}}$ are the nodal accelerations, and $\Delta\mathbf{U}$ are the incremental nodal displacements.

To advance in time, the constant-average-acceleration version of the Newmark integration scheme is adopted. The solution at the next time step is obtained by iterating within the time step. Once the incremental displacements, $\Delta\mathbf{U}^{k}$, are found, the displacement at the next iteration is calculated via Equation 9.10 and the tangent stiffness matrix and internal forces are updated. The iteration scheme is converged once the incremental displacements become sufficiently small, after which the solution is advanced to the next time step.

9.3.3 Immersed Boundary Method

The immersed boundary method (IBM) provides the link between the fluid and structural dynamics. It relies on two separate grids to represent the fluid and the boundary, which are independent of each other and non-conforming. The fluid is defined on an Eulerian grid (the lattice), which is Cartesian and fixed in space. The boundary, on the other hand, is represented on a Lagrangian grid, which is curvilinear and free to move across the computational domain. The main idea of the IBM is to mimic the effect of the boundary by introducing a source term (force) in the governing fluid equations. These forces are calculated in such a way

so that the fluid feels the presence of the boundary through this force, and the no-slip condition is satisfied along the surface.

Since the two grids are separate from each other, communication between them is crucial. Moreover, since they are non-conforming, this data transfer requires specialised interpolation/spreading operators. Specifically, the momentum field must be interpolated from the fluid grid (Eulerian) to the boundary grid (Lagrangian), and the resulting forces must be spread back from the boundary grid to the fluid grid. The operators defining these communication steps are given by:

$$\Phi(\mathbf{X}) = \mathcal{I}[\phi(\mathbf{x})] = \sum_{\Omega_s} \phi(\mathbf{x})\ \tilde{\delta}(\mathbf{x} - \mathbf{X})\Delta x \Delta y \Delta z \tag{9.11}$$

$$\phi(\mathbf{x}) = \mathcal{S}[\Phi(\mathbf{X})] = \sum_{\Gamma_s} \Phi(\mathbf{X})\ \tilde{\delta}(\mathbf{x} - \mathbf{X})\epsilon \Delta q \Delta r \Delta s \tag{9.12}$$

where $\mathbf{x} = (x, y, z)$, $\mathbf{X} = (q, r, s)$, ϕ is a quantity defined in the Eulerian frame, Φ is the same quantity defined in the Lagrangian frame, and ϵ is a scaling factor which ensures reciprocity between the interpolation and spreading steps [162]. Here, lower-case notation denotes values in the Eulerian frame, and upper-case notation denotes values in the Lagrangian frame. The present work adopts the three-point version of the discrete Dirac Delta function, $\tilde{\delta}$, proposed by Roma et al. [165].

The full LBM-IBM algorithm is given by Li et al. [163] and makes use of the fact that the LBM forcing scheme [130] decomposes the velocity into a predicted and a force-corrected term [157]. After one pass of the LBM equations, the macroscopic fluid velocity can be written:

$$\rho\mathbf{u}(\mathbf{x}, t) = \rho\mathbf{u}^*(\mathbf{x}, t) + \frac{\Delta t}{2}\mathbf{f}(\mathbf{x}, t) \tag{9.13}$$

where $\mathbf{u}^*$ is the predicted velocity field. The density and predicted velocity field are the only known quantities at this stage. Since the fluid velocity at the boundary must equal the velocity of the boundary, $\mathbf{U}^b = \mathcal{I}[\mathbf{u}]$, then converting Equation 9.13

into the Lagrangian frame gives:

$$\mathcal{I}[\rho(\mathbf{x},t)]\mathbf{U}^b(\mathbf{X},t) = \mathcal{I}[\rho\mathbf{u}^*(\mathbf{x},t)] + \frac{\Delta t}{2}\mathbf{F}(\mathbf{X},t) \tag{9.14}$$

Since the velocity of the boundary, $\mathbf{U}^b$, is known, Equation 9.14 can be rearranged and solved for the corrective force density, $\mathbf{F}$. The force is then transferred back to the Eulerian frame via the spreading operator, $\mathcal{S}$. Finally, the velocity field is updated by adding the corrective force to the predicted velocity field (Equations 9.13).

9.3.4 Fluid-Structure Coupling

To overcome the so-called added-mass instability at low mass ratios, a block Gauss-Seidel implicit coupling scheme is adopted. This technique iterates over the separate field solvers within the time step until the interface conditions are met. However, this does not guarantee stability and the coupling scheme may still converge slowly. Therefore, after each iteration the solution is relaxed by combining the structural displacements at the current and previous iterations, such that:

$$\mathbf{U} = \omega\tilde{\mathbf{U}} + (1-\omega)\mathbf{U}^{k-1} \tag{9.15}$$

where $\mathbf{U}$ is the relaxed solution which is passed to the fluid solver, $\tilde{\mathbf{U}}$ is the displacement computed from the structural solver, and $\mathbf{U}^{k-1}$ is the solution from the previous iteration.

To accelerate the convergence of the coupling scheme, the relaxation factor is adjusted dynamically according to Aitken's delta-squared method [179]:

$$\omega_k = -\omega^{k-1}\frac{(\mathbf{r}_{k-1})^T(\mathbf{r}_k - \mathbf{r}_{k-1})}{|\mathbf{r}_k - \mathbf{r}_{k-1}|^2} \tag{9.16}$$

where $\mathbf{r}_k$ is the residual vector, which is also used as a stopping criterion.

$$\mathbf{r}_k = \tilde{\mathbf{U}}_k - \mathbf{U}_{k-1} \tag{9.17}$$

9.4 Validation

To validate the current approach, the model is tested against experimental data obtained for an oscillating channel flow with a row of 10 flexible wall-mounted flaps [20]. A schematic of the computational setup is shown in Figure 9.3. In the original experiment, the flaps spanned most of the channel width and therefore the flow can be considered 2D at the channel centreline. Table 9.1 presents the experimental parameters; these values were also used as inputs for the numerical model. The Womersley and Reynolds numbers are calculated based on the flap length, and the reference velocity is given as the maximum centreline velocity obtained for a clean channel (no flaps).

Figure 9.4 shows the experimental and numerical results for the streamwise tip deflection of each flap. Overall the agreement is good; however, some slight discrepancies are observed. These may be attributed to differences in the approximation of the experimental parameters. In particular, the bending rigidity of the

Table 9.1: Experimental parameters for validation case.

Parameter	**Symbol**	**Value**
Channel height	H	6 cm
Flap height	L	2 cm
Flap thickness	h	1 mm
Fluid density	ρ_f	$1200\,\mathrm{kg/m^3}$
Kinematic viscosity	ν	$1 \times 10^{-4}\,\mathrm{m^2/s}$
Flap density	ρ_s	$1200\,\mathrm{kg/m^3}$
Young's Modulus	E	1.2 MPa
Frequency	f	1 Hz
Womersley number	α	5
Reynolds number	Re	120

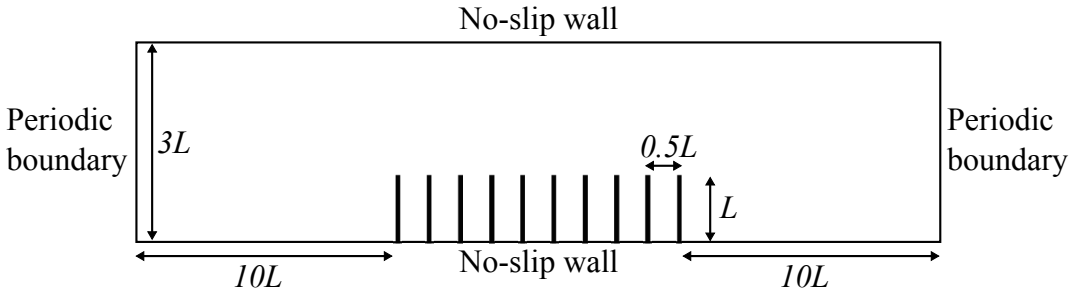

Figure 9.3: Schematic of computational domain for validation case.

flaps is extremely sensitive to the flap thickness, with $EI \propto h^3$. Another possible explanation for the noted differences is 3D effects, due to the finite span of the flaps in the experiment, which are not captured in the 2D simulation. Nevertheless, the agreement is good and shows that the present model can reproduce the dynamic motion of multiple wall-mounted flaps.

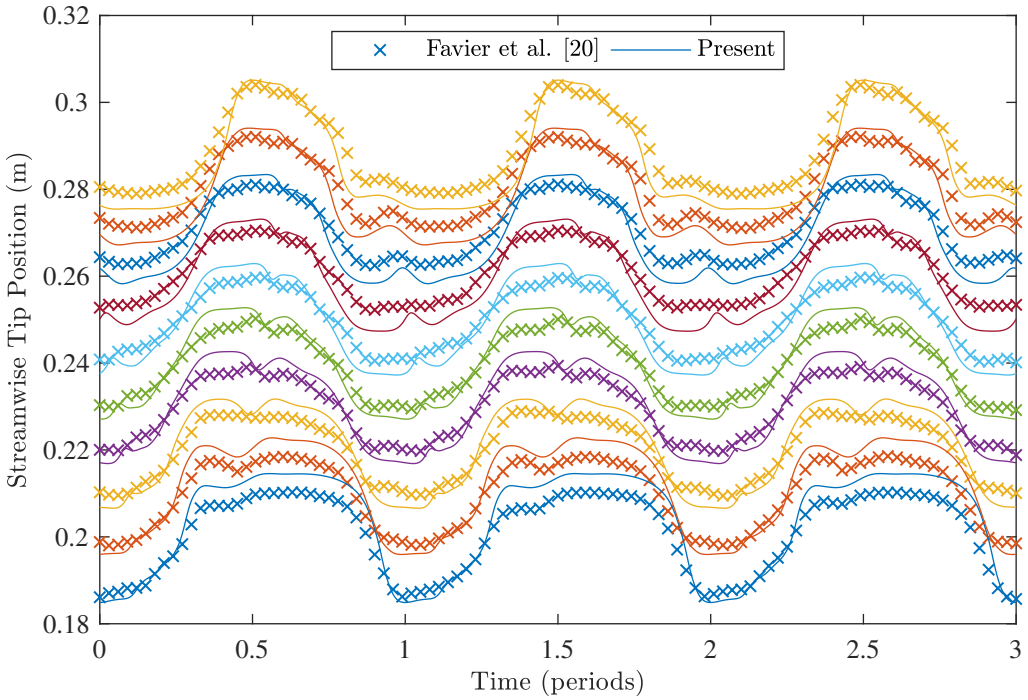

Figure 9.4: Validation of present model. Streamwise tip positions for each flap are compared against the experiment of Favier et al. [20].

9.5 Simulation Setup

To limit the edge effects and isolate the waving instability, a large array of flaps is tested. Figure 9.5 shows a schematic of the computational domain. The setup consists of an open-channel flow with an array of 128 flaps. The flaps have a length L and are equally spaced by $0.5L$. The channel height is given by $H = 3L$, so that the configuration resembles that of a shallow submerged vegetation canopy [77]. The inlet is placed $10L$ upstream from the array, whereas the outlet is placed $30L$ downstream. Preliminary tests indicated this was sufficient to limit the effect of the boundaries on the global behaviour of the array and is in-line with other studies of wall-mounted flaps at similar Reynolds numbers [212, 213].

The inlet velocity profile is set according to:

$$u_x(0, y, t) = 1.5\bar{u}\frac{y(2H - y)}{H^2} \tag{9.18}$$

where $\bar{u}$ is the mean (bulk) velocity. Zero-velocity initial conditions are set and the inlet velocity is ramped up over a short period of time. This limits the appearance of density waves, which arise due to the weakly compressible nature of the LBM. A fixed pressure is set at the outlet, no-slip on the bottom wall, and a free-slip (rigid lid) condition is used for the top boundary [15]. Based on the bulk velocity and flap length the Reynolds number is 80. This was chosen from a preliminary study, which indicated this is close to the minimum Reynolds number capable of exhibiting waving instabilities.

For a fixed Reynolds number and geometry, the governing parameters are the

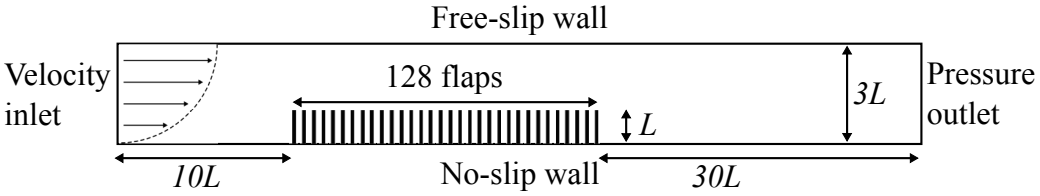

Figure 9.5: Schematic of computational domain.

dimensionless mass ratio and bending stiffness:

$$\mathcal{M} = \frac{\rho_s h}{\rho_f L}, \qquad \mathcal{K} = \frac{EI}{\rho_f \bar{u}^2 L^3} \tag{9.19}$$

In the following tests, five different mass ratios (denoted M1–5) and ten different bending rigidities (K1–10) are tested. An initial parameter search identified an appropriate range where a variety of behaviours were observed. These values are distributed exponentially, such that:

$$\mathrm{M}m = 1 \times 10^{-1+\frac{(m-1)}{2}}, \qquad m = 1 \ldots 5 \tag{9.20}$$

$$\mathrm{K}k = 5 \times 10^{-3+\frac{2(k-1)}{9}}, \qquad k = 1 \ldots 10 \tag{9.21}$$

To compare the instability mechanism of the array, a similar range of tests are performed for a single flap. Since the instability mechanism in this case is different to that observed for the array, preliminary tests found that the range of mass ratio given in Equation 9.20 did not yield significant dynamic motion for the single flap. Therefore, the range of mass ratio in the single flap tests are shifted one order of magnitude higher. To keep the numbering scheme consistent, the range of cases for the single flap test are given by:

$$\mathrm{M}m = 1 \times 10^{-1+\frac{(m-1)}{2}}, \qquad m = 3 \ldots 7 \tag{9.22}$$

$$\mathrm{K}k = 5 \times 10^{-3+\frac{2(k-1)}{9}}, \qquad k = 1 \ldots 10 \tag{9.23}$$

therefore, both the single and array cases at MmKk share the same material properties.

Each flap is discretised with 50 IBM markers/lattice sites ($\Delta x = 0.02$) and 20 FEM elements. To ensure a configuration-independent solution, tests were performed on

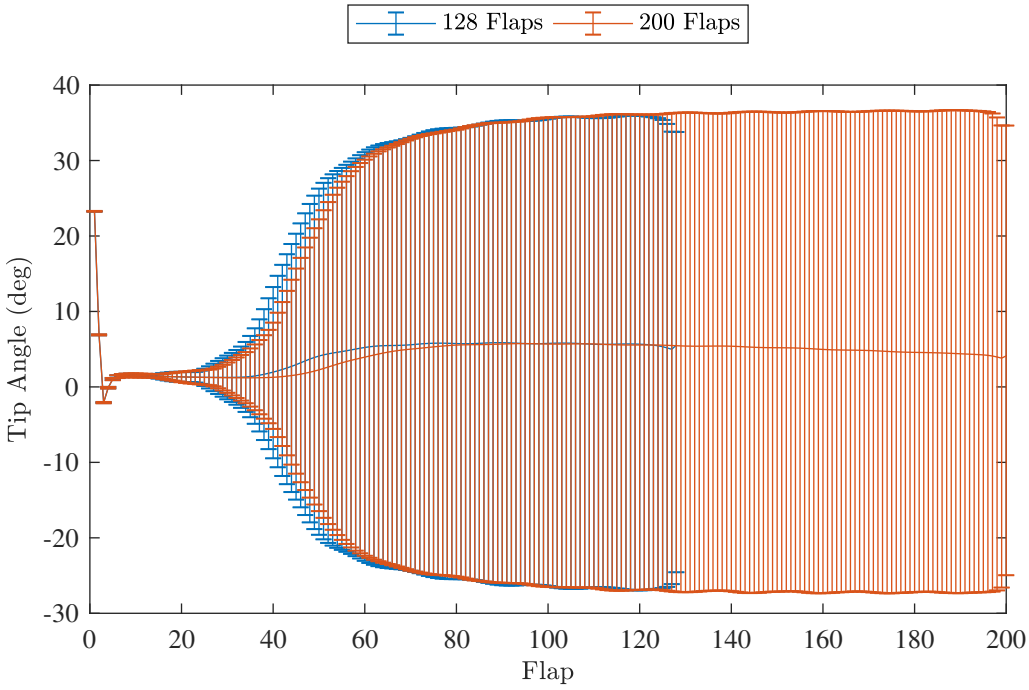

(a)

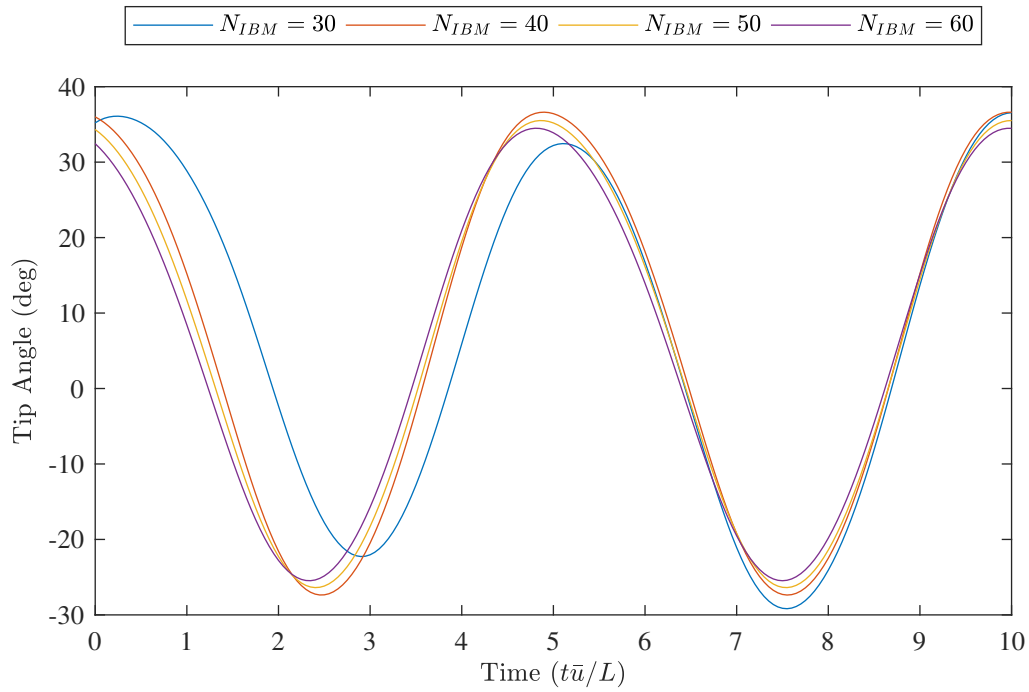

(b)

Figure 9.6: Array length and grid resolution tests. (a) Minimum, maximum and mean flap deflection across the array for two configurations with 128 and 200 flaps. The values are extracted over a period of time spanning over 100 periods of oscillation. (b) Tip angle over final 10 s, where N_{IBM} is the number of IBM markers/lattice sites used to discretise the flap.

both the array length and the grid resolution. Figure 9.6a shows the range of tip motion for each flap along the array for a representative case (M3K8) with 128 and 200 flaps. The tip angle is measured from the base of the flap as the angle between the x-axis and the flap tip, so that $\theta_{tip} = \arctan\left({}^{(L+\delta_y)}/_{\delta_x}\right)$, where δ_x and δ_y are the tip deflections. As can be seen, there is little difference between the two cases in terms of the range of tip motion. Both cases also show a significant deflection for the first flap, which is entirely exposed to the force of the fluid. Furthermore, for the 200 flap case there is little variation in the tip motion in the extended region of the array (flaps 129 to 200). Although there is a slight decreasing trend in the mean value towards the end of the array, this is relatively minor and the range of motion seems unaffected. For this reason, 128 flaps was deemed sufficient for the length of the array. Figure 9.6b shows the tip deflection for a representative flap (number 100) for a range of tested grid resolutions. After a relatively large shift in solution between the $N_{IBM} = 30$ and $N_{IBM} = 40$ cases, there is little change with further increasing grid resolution. As such, the $N_{IBM} = 50$ resolution was selected for all further simulations.

9.6 Classification of States for Single Flap

Initially, the global behaviour of the system with respect to the parameter space is examined for a single flap. Through observation of the flap dynamics, each case is classified into one of four modes: static, flapping, period-doubling, or chaotic motion. Figure 9.7 shows the parameter map for all tested cases (M3–7, K1–10). Note that cases where large mass ratios were combined with low bending rigidities (M6K1, M7K1), the large flapping motion resulted in the flap tip making contact with the bottom wall or passing straight through it. While this did not always result in stability problems (e.g. when the contact was brief), no contact model is incorporated in the present model and so, to avoid possible pollution of the data, results were not collected for these cases. This also happened for the M7K2 case; however, the period of time before contact was made was sufficient to characterise the motion.

Figure 9.7 shows dynamic states over only a small region of material properties. In particular, it is evident that bending rigidity is stabilising while mass ratio is destabilising. This is in line with the literature concerning the flapping flag

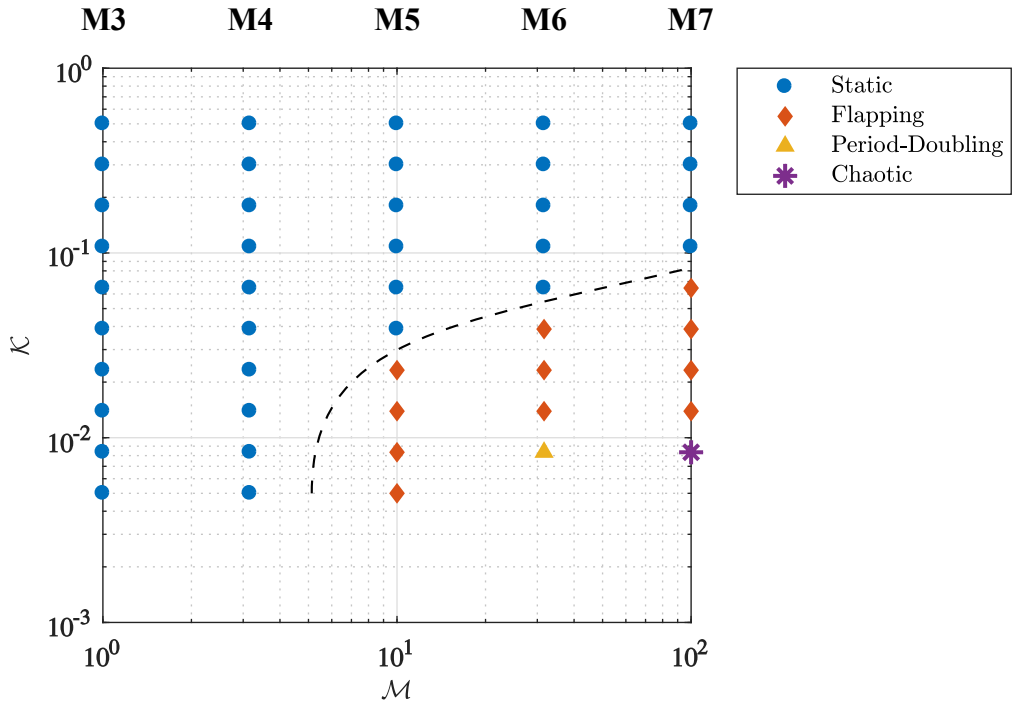

Figure 9.7: Parameter map with observed behaviour modes for single flap. The range of mass ratio and bending stiffness is given by M3–7 and K1–10 cases. Dashed line is shown to help distinguish boundaries.

problem, which report similar trends with regards to the stabilising/destabilising effects of bending stiffness/mass ratio [13, 62]. It is expected that at larger mass ratios the window of dynamic motion would continue to increase in size. However, since the main purpose here is to compare the behaviour modes against the array cases, which were conducted at mass ratios one order of magnitude lower than the single flap cases, these larger mass ratios were not tested. Furthermore, it is expected that the problems associated with the tip making contact with the bottom wall would be exacerbated at larger mass ratios.

To help characterise the motion of the flap over the three dynamical states (flapping, period-doubling, chaotic), Figure 9.8 shows the tip angle histories and power spectra for each of these three states. Also shown are the phase plots, which depict the tip trajectories based on the vertical velocity (δ_y') and displacement (δ_y) of the tip. The chaotic case (M7K2) is plotted over a different time range, since soon after this the flap tip made contact with the bottom wall. The phase plots are taken over the same period of time as that given by the tip angle histories. For the flapping case (M5K2), the flap motion is seen to have a regular motion with constant amplitude and frequency. A small 'kick' in the tip motion is observable at the extremities of the flapping cycle. This is especially noticeable in

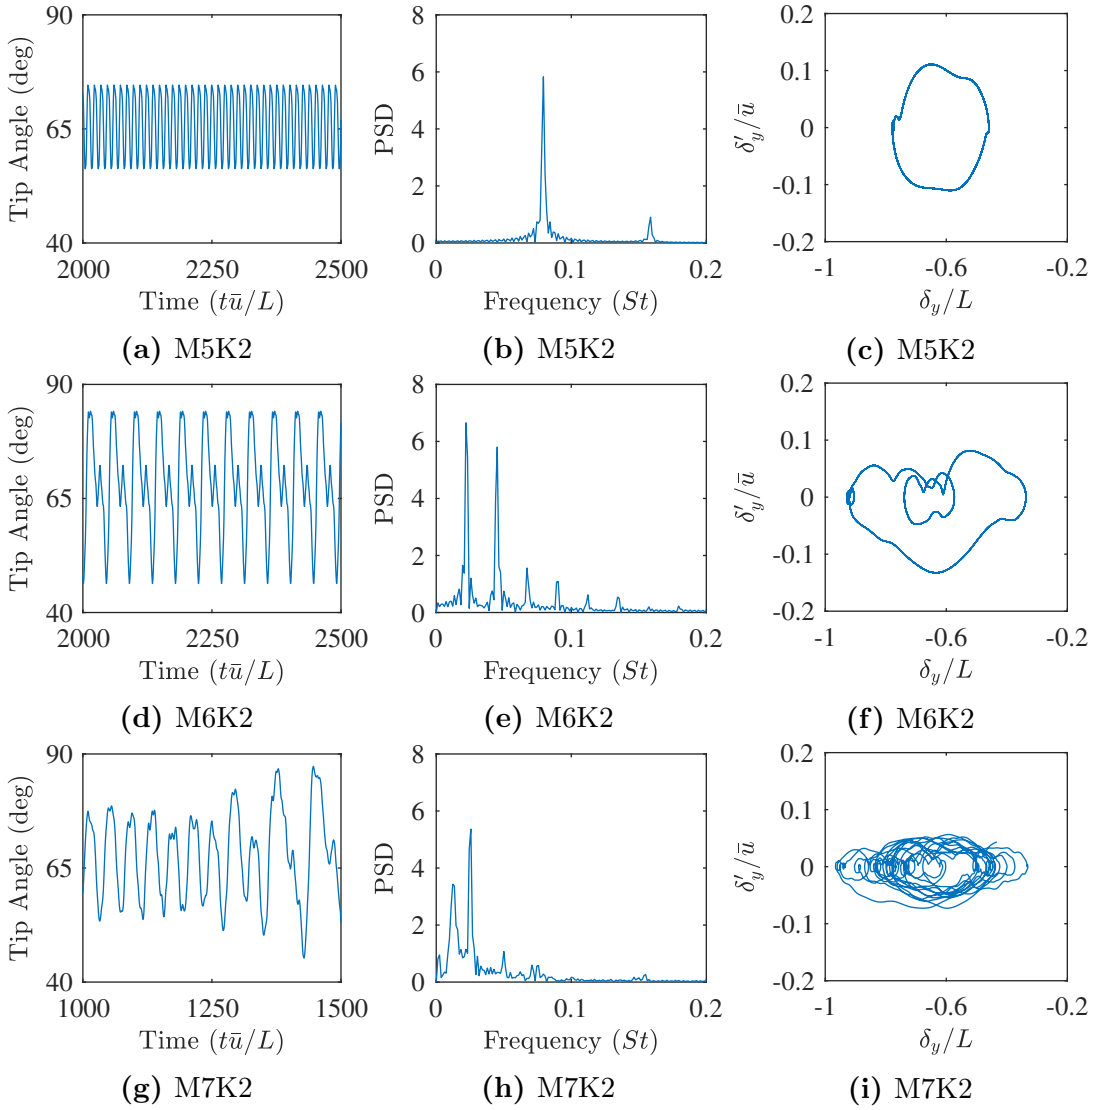

Figure 9.8: Tip angle histories, power spectra, and phase plots of vertical tip motion for flapping (M5K2), period-doubling (M6K2), and chaotic motion (M7K2) for single flap. The phase plots are taken over the same period of time as that shown in the tip angle histories.

the phase plot. At the next mass ratio (M6K2), a regular flapping motion is still observable. However, this now takes the form of a double-period cycle, with the secondary motion obtaining a significantly lower tip angle. Finally, at the largest mass ratio (M7K2) there is no clear pattern in the tip motion, with variations in amplitude and frequency. Furthermore, both the power spectrum and phase portrait are representative of chaotic motion, with the power spectrum showing a wider distribution of power around the main frequency and the phase portrait showing uncorrelated trajectories. It is interesting to note that the transition from period-one, to period-two, then to chaotic motion would suggest a period-doubling

cascade to chaos. However, a finer resolution of cases between these states would be required to determine this.

9.7 Classification of States for Array

Figure 9.9 shows the classification of the behavioural states over the full parameter space for the array. Through observation of the flap dynamics, each case is classified into one of four modes: a static behaviour, a waving behaviour, an irregular waving behaviour, and a flapping behaviour. Note that cases where large mass ratios were combined with low bending stiffness (M5K1, M5K2, M5K3), contact between adjacent flaps led to stability problems and so results were not obtained for these cases. Examining the variation in behaviour with the structural properties of the flaps shows a clear diagonally-banded region of waving motion. Immediately either side of this are regions of static behaviour, which exhibit a broad range of deflection values, depending on the bending stiffness. At large mass ratio/low bending stiffness, a region of flapping behaviour is observed. This

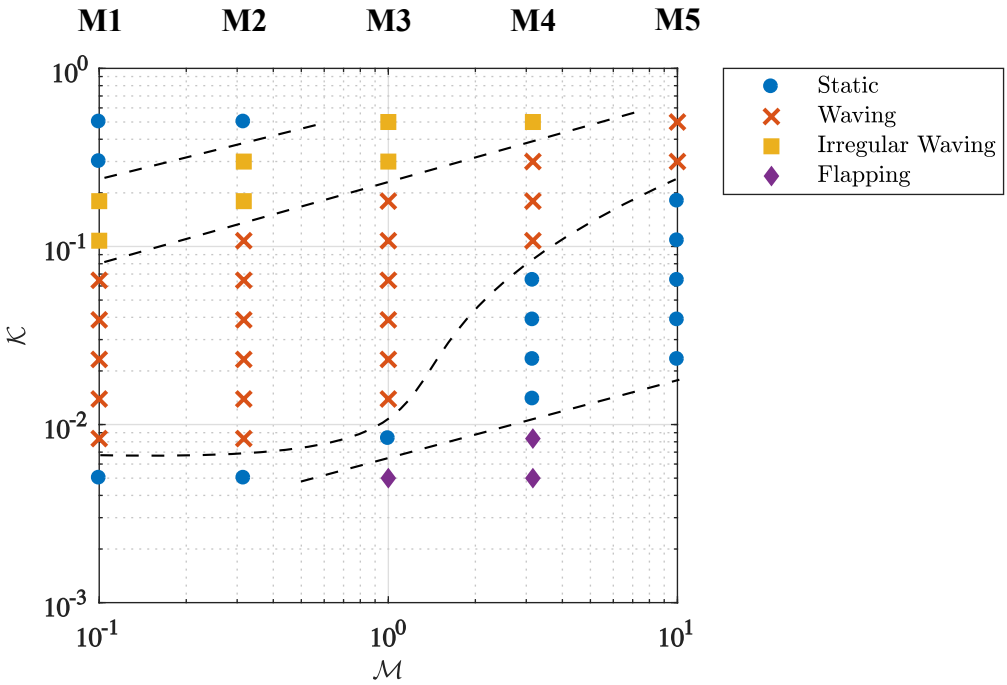

Figure 9.9: Parameter map with observed behaviour modes for array. Dashed line is shown to help distinguish boundaries. The range of mass ratio and bending stiffness is given by M1–5 and K1–10 cases. Therefore, these mass ratios are shifted one order of magnitude lower compared to the single flap case.

flapping mode will be discussed in detail later.

A distinction is made between the two types of waving motion: regular and irregular. Figure 9.10 gives an example of what is meant by this classification. Here, the tip angle history, power spectra, and phase portrait of the streamwise tip deflection are shown for a representative flap along the channel. For the case of regular waving, after the initial transient the motion of the flap settles into a regular periodic motion with constant amplitude and frequency, which corresponds to the passage of individual vortices over the flap. However, in the case of irregular flapping, as well as the main high-frequency mode, there is a secondary low-frequency dynamic. This causes large temporal and spatial variations in the waving amplitude. This is especially evident in the power spectrum; which, as well as showing a low frequency peak, has a wider distribution of power around the high frequency dynamic. The phase portrait also helps to distinguish between the regular and irregular flapping. Here, the streamwise tip deflection has been plotted over a period of 150 s, which corresponds to approximately 15 and 30 periods of oscillation for these two cases, respectively.

Figures 9.11–9.12 show snapshots of the flow field and flap deflection for four representative cases. For the first case (low bending stiffness), the behaviour is observed to be static with a drag-reducing reconfiguration. No instabilities are observed in the flow field or flap dynamics. However, as the bending stiffness is increased, instabilities start to develop approximately a third of the way along the channel, before rolling up to form large coherent vortical structures. The appearance and behaviour of this instability depends heavily on the structural properties. After an initial development length, these structures eventually establish a fixed size. As the bending stiffness is increased further, the instabilities are suppressed and the behaviour becomes static again, this time with small deflection. The four states shown here are emblematic of the four characteristic states usually used to describe vegetation behaviours. In order of decreasing flow speed, these states are usually described as prone (static reconfiguration), strong coherent swaying (regular waving), gentle swaying (irregular waving), and erect (static with small deflection) [80]. While these states are described in terms of flow speed, since the dimensionless bending stiffness is inversely dependent on flow speed (Equation 9.19), the trends observed in Figures 9.11–9.12 are in line with those in the literature.

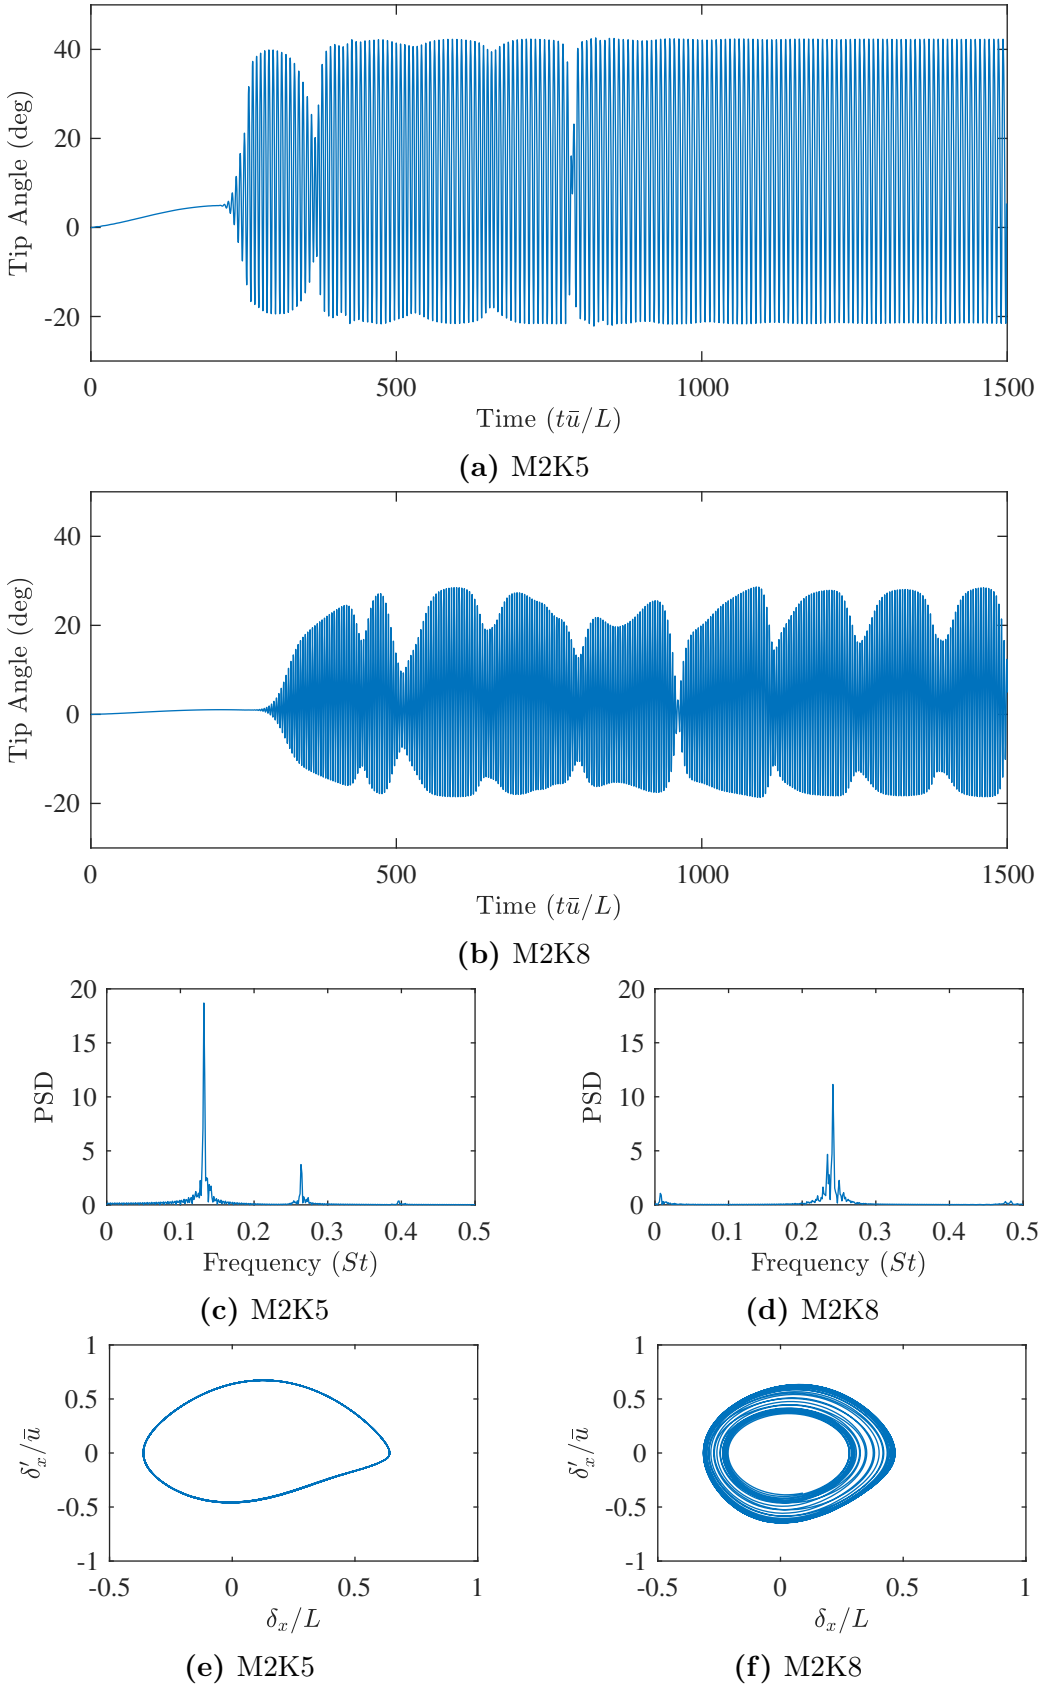

(a) M2K5

(b) M2K8

(c) M2K5

(d) M2K8

(e) M2K5

(f) M2K8

Figure 9.10: Tip angle histories, power spectra, and phase plots for regular (M2K5) and irregular waving (M2K8). These results are obtained for the 100th flap. The phase portraits are plotted over the final 150 s, corresponding to approximately 15 (M2K5) and 30 (M2K8) periods of oscillation for these two cases.

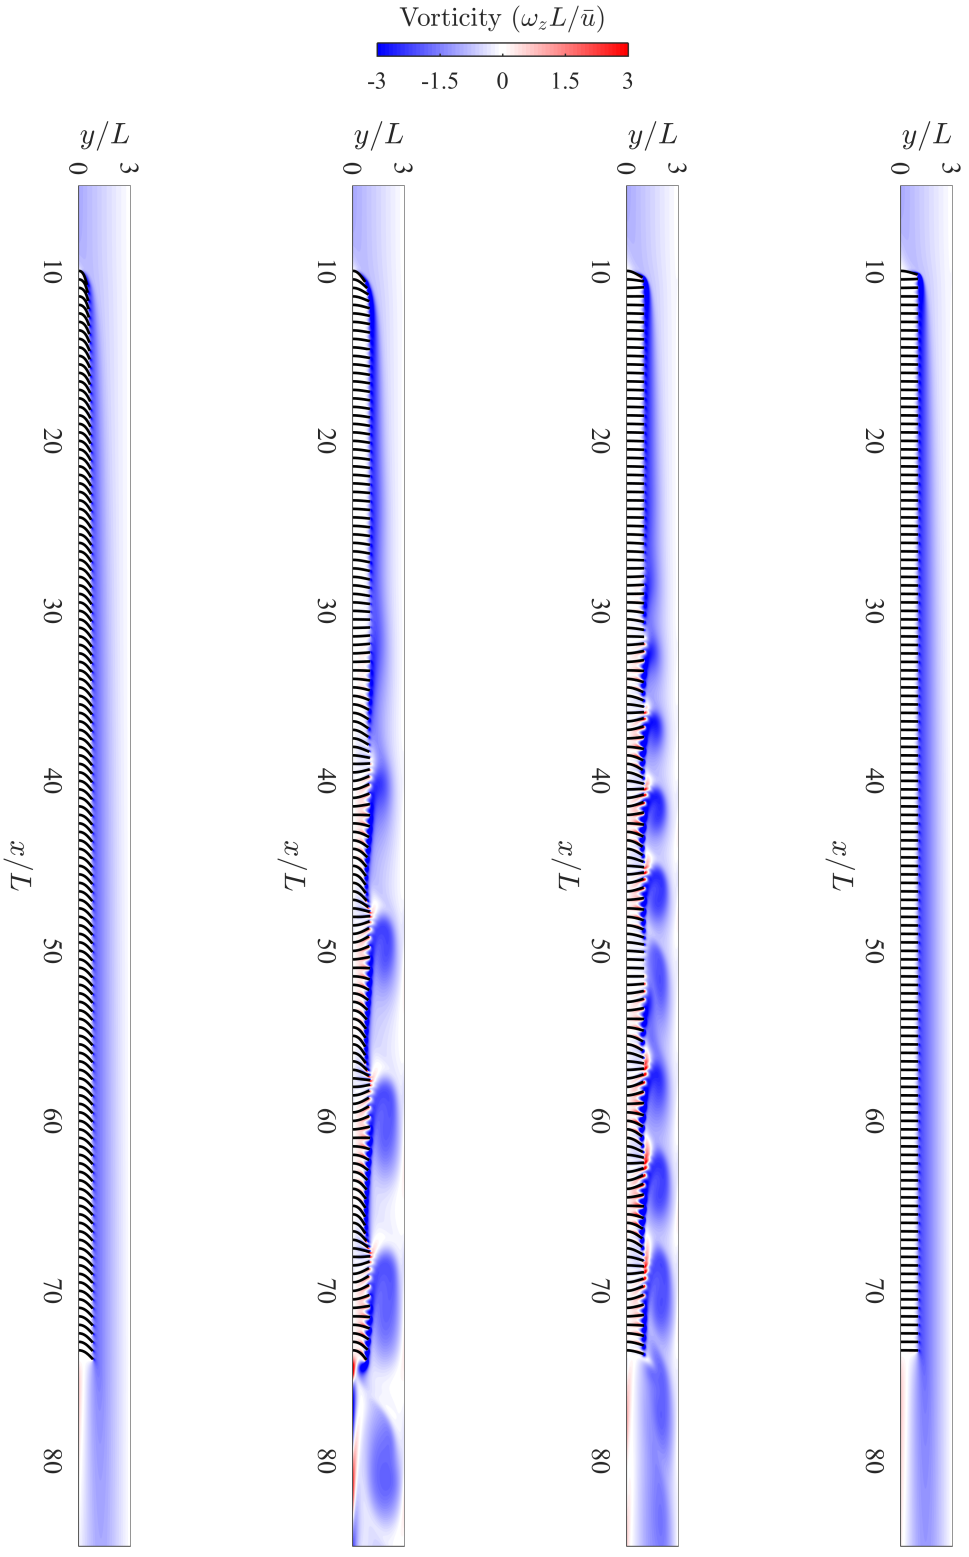

(a) M2K1 **(b)** M2K5 **(c)** M2K8 **(d)** M2K10

Figure 9.11: Snapshots of out-of-plane vorticity and flap positions across the whole array for four representative cases. (a) Static reconfiguration. (b) Waving. (c) Irregular waving. (d) Static with small deflection.

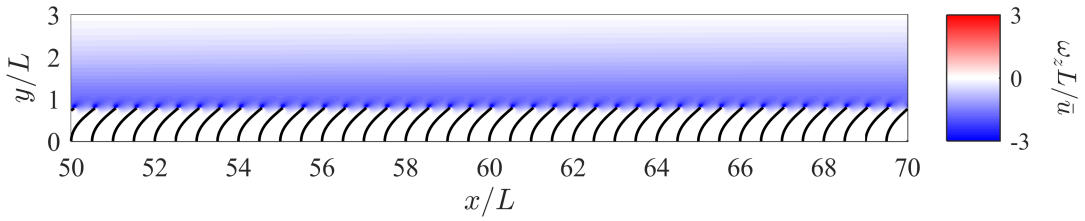

(a) M2K1

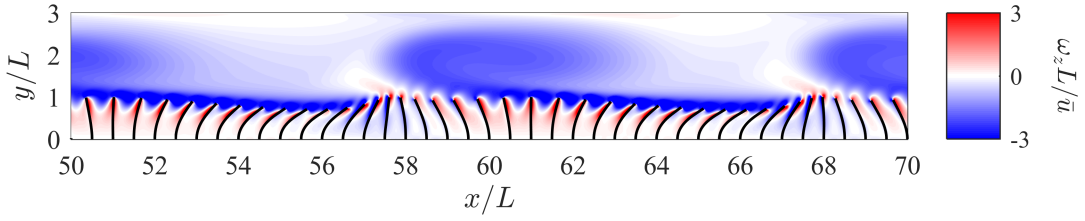

(b) M2K5

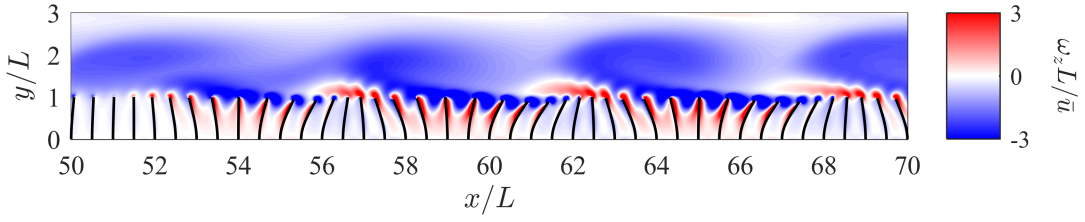

(c) M2K8

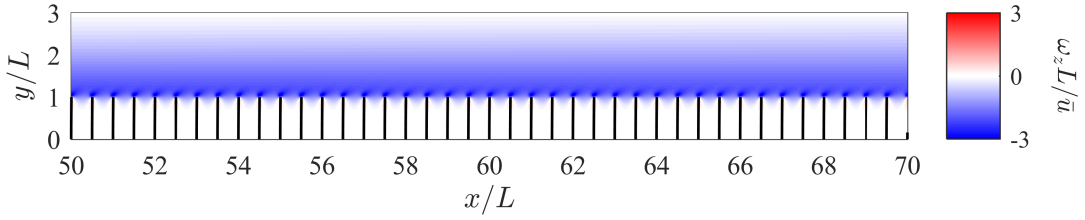

(d) M2K10

Figure 9.12: Enlarged snapshots of out-of-plane vorticity and flap positions for four representative cases. (a) Static reconfiguration. (b) Waving. (c) Irregular waving. (d) Static with small deflection.

Examining both the temporal and spatial patterns in the coupled dynamics is a challenging task. Figure 9.13 provides a convenient method for examining such patterns. For each flap along the array the instantaneous tip angle is plotted with time. This allows easy observations of a range of properties describing the system behaviour, including the development length, wavelength, frequency and wave speed of the instabilities, as well as the global behaviour mode. Clearly observable is the effect of bending stiffness and the resulting transition from static reconfiguration, through the waving mode (with increasing frequency of motion),

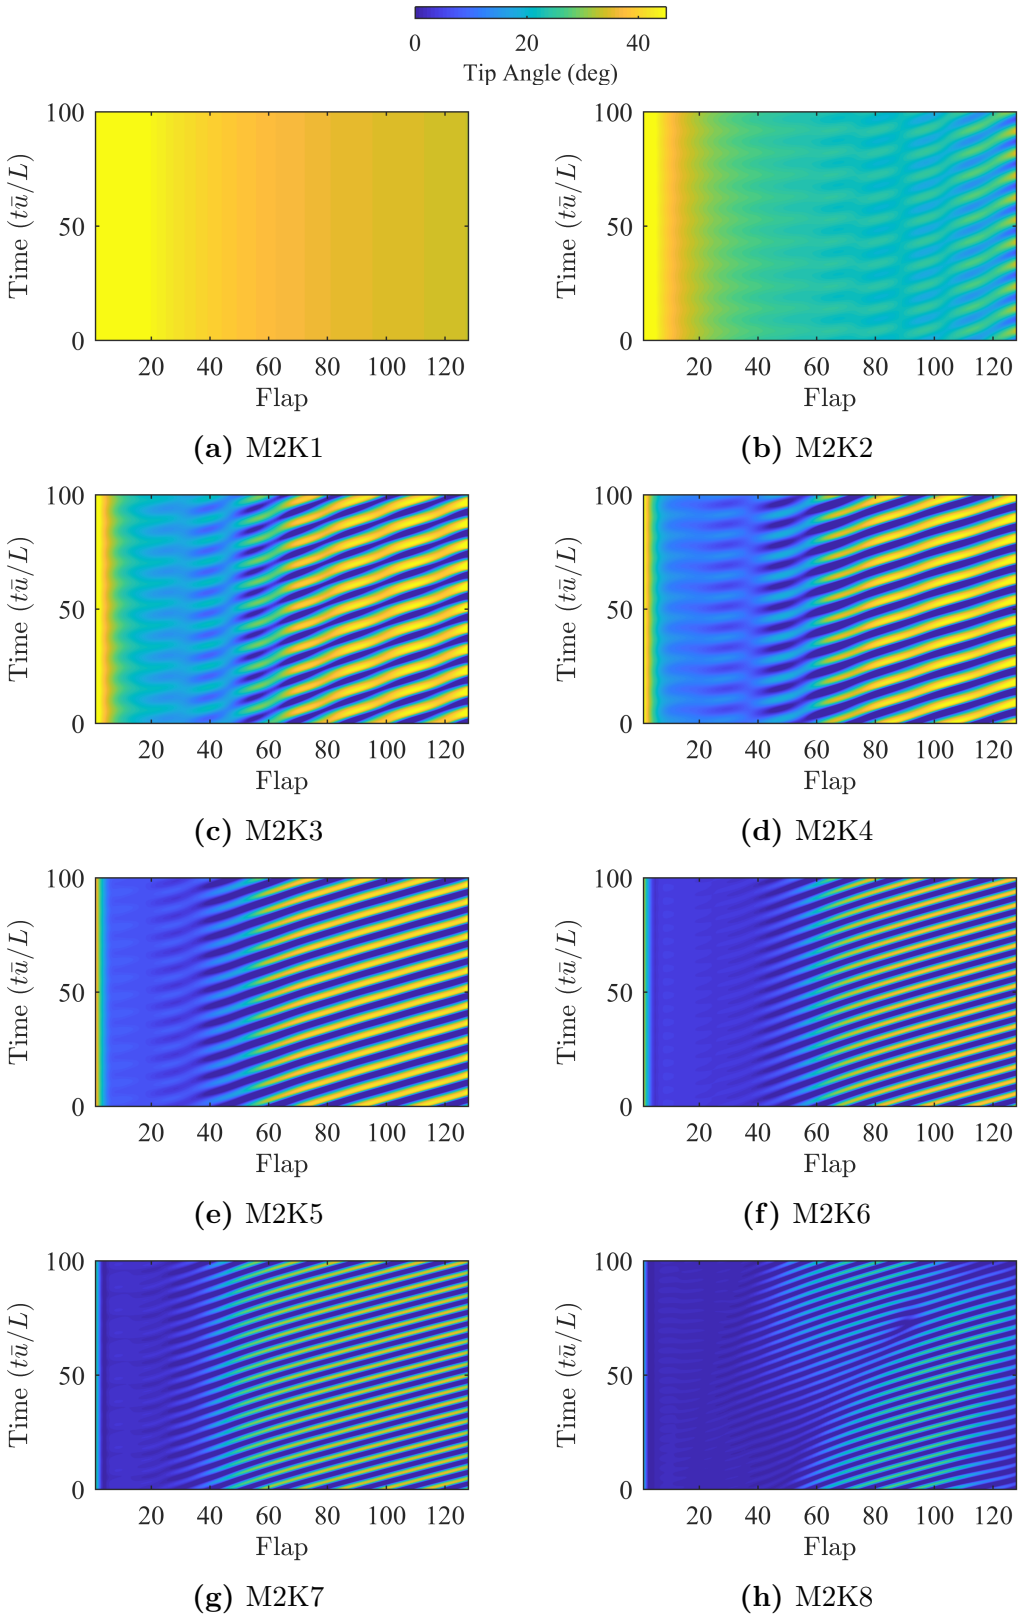

Figure 9.13: Contours of tip angle across the whole array for the final 100 s. Clearly observable is the emergence of a regular waving motion which breaks up at large bending rigidities. Furthermore, the narrowing of the waving bands indicates higher frequency motion at large bending rigidities.

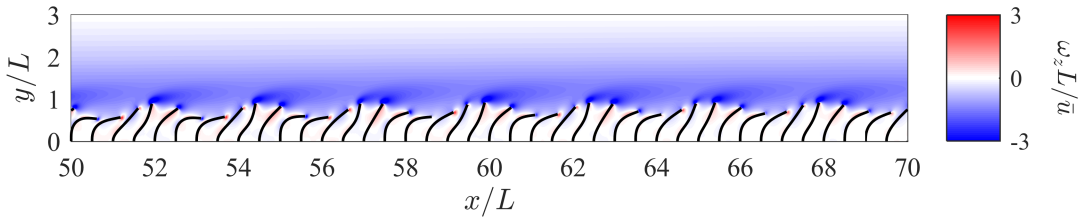

Figure 9.14: Enlarged snapshot of out-of-plane vorticity and flap positions for flapping mode (M4K1).

until a breakup of the spatial and temporal coherence through the array occurs at large bending rigidity (irregular waving mode). Also noticeable is a narrowing of the waving bands as the bending rigidity is increased. This indicates an increase in the frequency of motion.

Figure 9.14 shows the contours of the out-of-plane vorticity and flap positions for the flapping mode (M4K1). Clearly observable is the emergence of smaller scale flow structures, which span a smaller number of flaps compared to the waving mode. Furthermore, the snapshot of the flap positions indicates a different deformation mode to the one observed in the waving cases. Figure 9.15 shows snapshots of a representative flap profile over one cycle for both the waving and flapping modes. For comparison, the flapping mode observed for a single flap on its own is also shown. This figure clearly justifies the classification of the flapping mode in the array behaviour. Both the single flap and the array show a clear

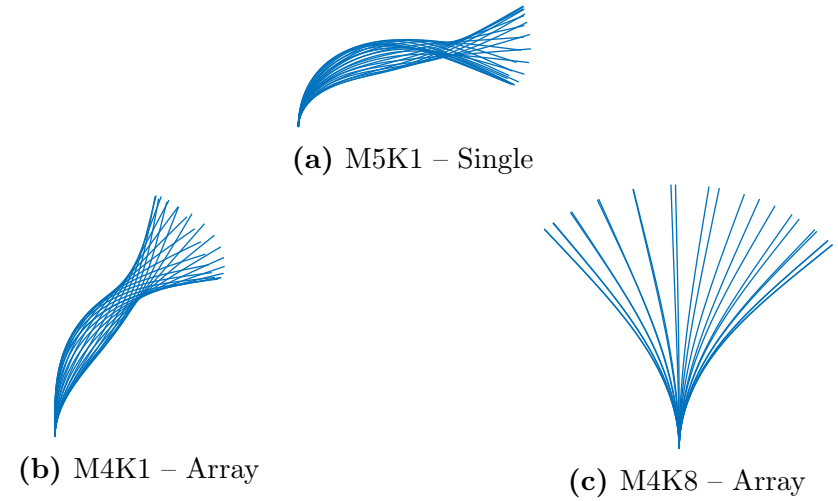

(a) M5K1 – Single

(b) M4K1 – Array

(c) M4K8 – Array

Figure 9.15: Flap profile shapes for flapping and waving modes. (a) Single flap in flapping mode (M5K1). (b) Array in flapping mode (M4K1). (c) Array in waving mode (M4K8). For the array cases the 100th flap is extracted.

necking behaviour in the profile snapshots. This is evidence of the mode-two deformation that is associated with flapping motion. Although the bending rigidity for these two cases is the same, the single flap shows a greater degree of deflection. This can be attributed to the presence of the adjacent flaps in the array, which help to impede the motion. Also noticeable is the fact that the waving motion is dominated by mode-one deformation. As will be shown in Section 9.8, the emergence of different deformation modes affects the frequency of the dynamic motion.

Figure 9.16 shows the spatial and temporal variations in tip angle for the flapping and waving modes. The first observation that can be made is the flapping mode exhibits a larger mean value for the tip angle but a smaller amplitude of motion, when compared to the waving mode. Furthermore, observations regarding the frequency and wave speed can be made by examining the size of the waving bands. The vertical distance between the bands gives the period of the motion, whereas the angle (measured with respect to the time axis) of the travelling bands gives the wave speed. While both these cases exhibit similar frequencies, the wave speed is dramatically different, with a much lower wave speed observed for the flapping case. Since the wavelength is related to the frequency and wave speed, this implies the size of the flow structures are significantly smaller for the flapping case. This observation is supported by examining the vorticity contours in Figure 9.14.

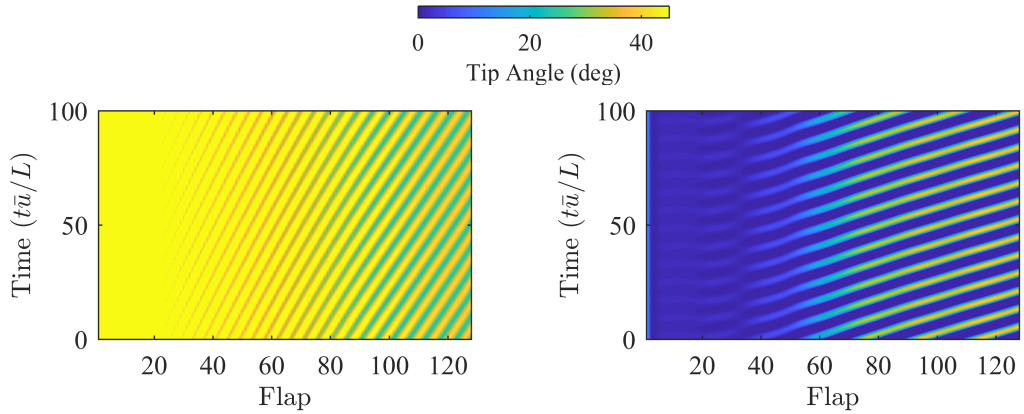

(a) M4K1 **(b)** M4K8

Figure 9.16: Contours of tip angle across the whole array for the flapping (a) and waving modes (b).

9.8 Frequency of Instability

Figure 9.17 shows the mean frequency of motion across the whole array for each dynamic case. Several key observations can be made regarding this figure. First, there is a large variation in frequency over the tested parameter space. This indicates that there is indeed a coupled dependence of the mixing layer instability on the structural properties of the array. Furthermore, the regular waving mode seems to only exist within a well-defined frequency range (St = 0.1–0.2). To estimate the frequency of the mixing layer instability, a rigid configuration was tested and the steady-state velocity profile extracted. Using Equation 9.1, the frequency of the mixing layer instability for this set of flow conditions was found to be approximately $St_{KH} = 0.13$. This falls within the range of the regular waving frequencies and suggests that the governing mechanism for the dynamic motion is a coupling of the mixing layer and structural natural frequencies. This may also explain the transition to the irregular waving state, where the imbalance between the two frequencies leads to competing behaviour between the fluid and flaps. For a sample of selected cases, the time-averaged velocity field for the dynamic cases was also used to estimate the mixing layer frequency. The range of predicted values was found to be 0.1 to 0.15, which is within the observed range of regular

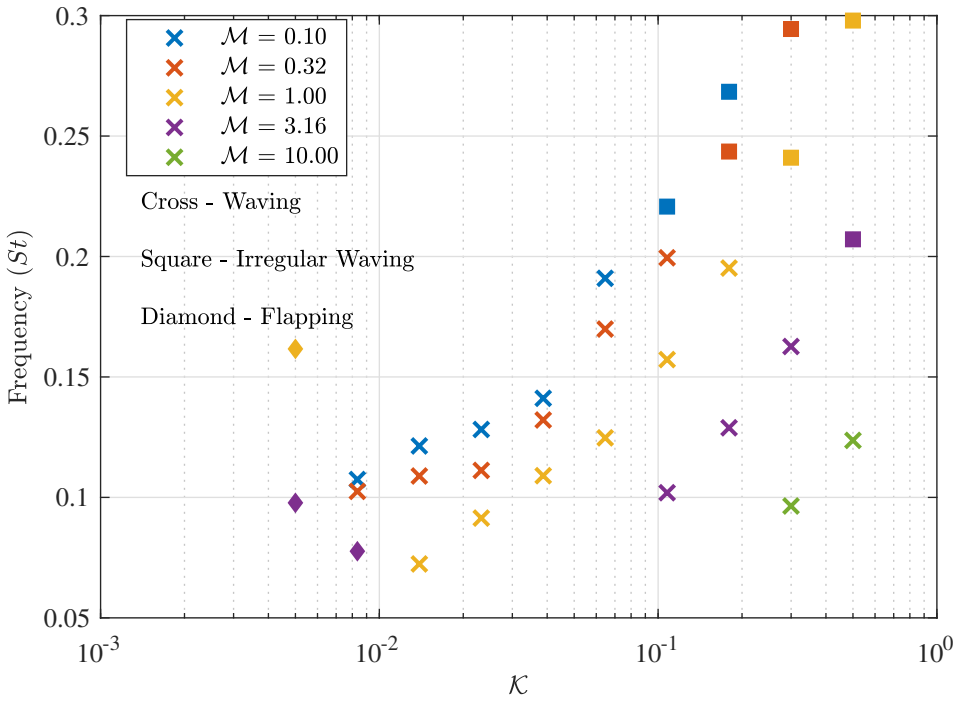

Figure 9.17: Mean frequency (Strouhal number) across the whole array for each dynamic case.

waving frequencies and thus further supports this finding.

Another observation from Figure 9.17 can be made regarding the three flapping cases. These states do not follow the same trend as the rest of the cases. Instead, they show much higher frequencies. This can be explained by the fact that the main deformation mode for the flapping states is the second mode, which is associated with higher-frequency motion. This further supports the claim that this classification is indeed a flapping mode.

Also apparent from Figure 9.17 is the dependency of the waving frequency on the structural natural frequency, with increasing frequency observed for low-mass/high-stiffness cases. To test this, the damped natural frequency of the array was measured by performing a free-oscillation test, whereby the array was placed in a stationary flow with each flap given an initial deflection. At the start of the simulation the flaps were released and left to oscillate freely. The results from this test are shown in Figure 9.18. As expected, the structural mass has the effect of decreasing the natural frequency, whereas the bending rigidity has the effect of increasing it.

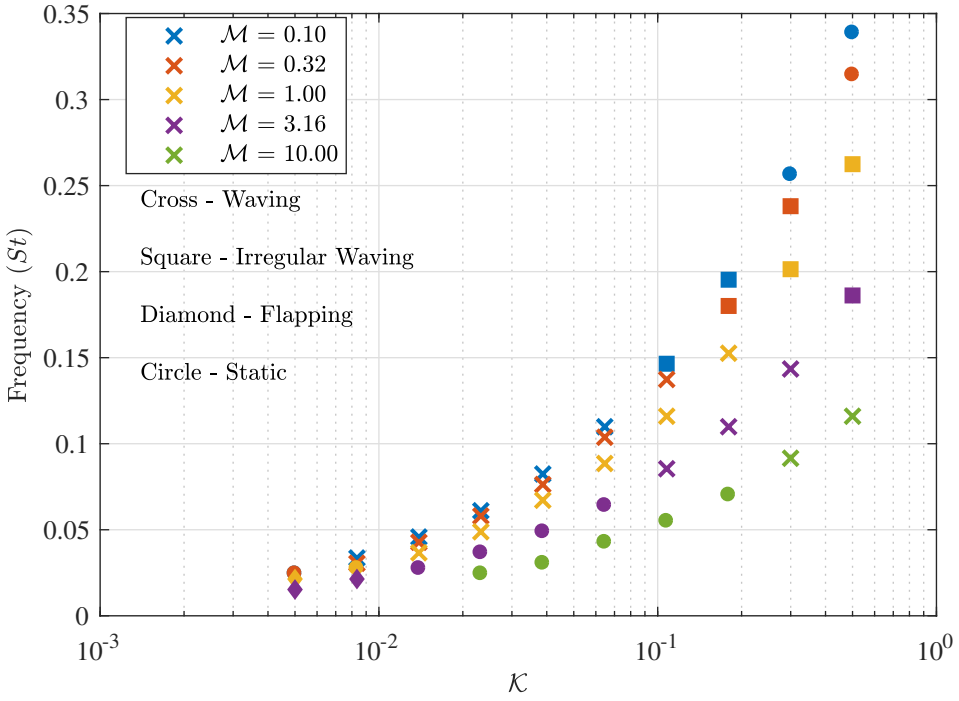

Figure 9.18: Natural frequency of array measured by free (damped) oscillation test. Note that the marker shapes are based on the states exhibited in Figure 9.17, not the free-oscillation test.

The damped natural frequencies obtained from the free-oscillation test were used to normalise the frequencies obtained in Figure 9.17. Figure 9.19 shows the results from this procedure. The normalisation collapses the data across mass ratios onto a single curve. While this plot makes separating the type of waving motion more difficult, it is clear to see that the waving motion occurs in a range of frequencies that are close to the (damped) natural frequency of the array. Furthermore, this ratio approaches one as the bending stiffness increases. This is likely due to the increased bending energy stored in the flaps, which results in the natural response of the structure dominating the motion. Also noteworthy is the large disparity between the natural frequency and the observed frequency for the flapping cases. This can be explained by the fact that the free-oscillation tests were set up so as to measure the first natural frequency of the flaps, as opposed to the second mode which dominates the flapping motion.

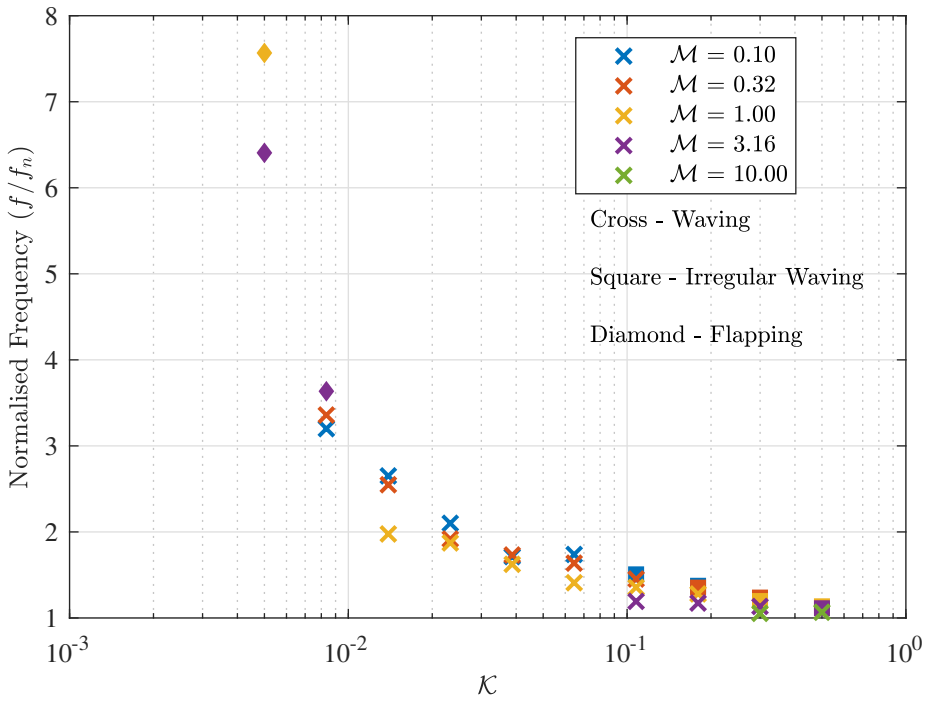

Figure 9.19: Mean normalised frequency (f/f_n) across the whole array for each dynamic case.

9.9 Transition Between States

To examine the transition over the full range of behaviours, extra cases were run for the M3 mass ratio. The resulting states and frequencies are shown in Figure 9.20. The full range of observed behaviours are shown here; starting from static behaviour at low bending rigidities, moving to the flapping mode with increasing bending rigidity, through another region of static behaviour, then through the waving and irregular waving modes, before finally becoming static again. Clearly observable is the jump in frequency between the flapping and waving modes. Furthermore, right in the middle of the waving mode there appears to be a plateau in frequency at approximately $St = 0.11$. Recalling from Section 9.8, the predicted frequency of the mixing layer was found to be $St_{KH} = 0.1$–0.15. This plateau suggests a lock-in between the natural frequency of the array and the frequency of the mixing layer instability. This further supports the claim that the regular waving mode is triggered by a lock-in effect, whereas the other modes are associated with asynchronous natural frequencies between the fluid and the array.

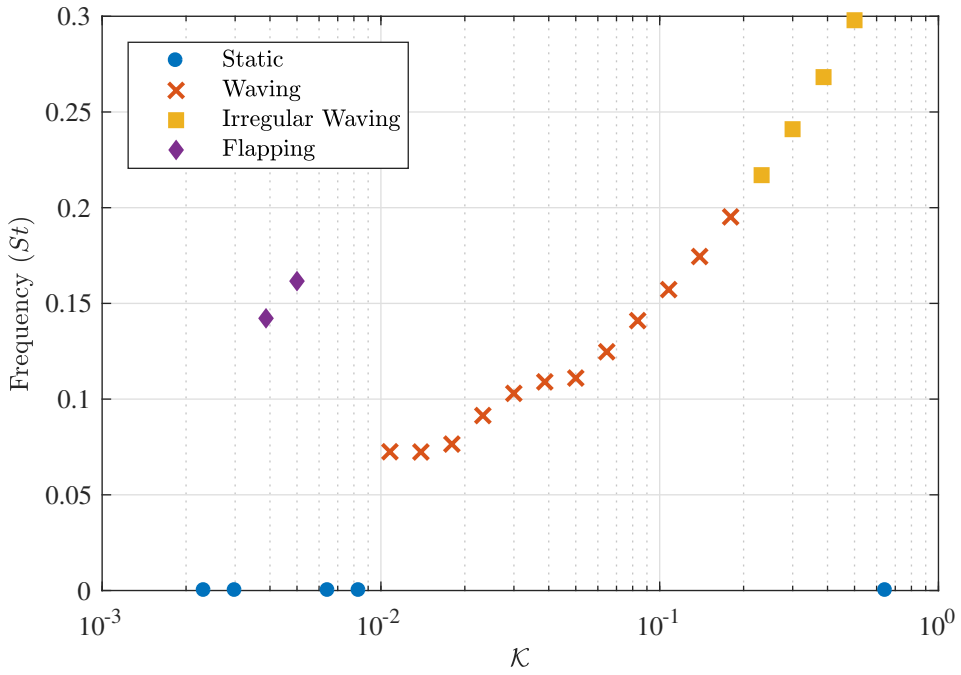

Figure 9.20: Mean frequency (Strouhal number) across the whole array for full range of M3 tests.

9.10 Conclusions

To investigate the coupled interactions between fluids and arrays of slender structures, a large array consisting of 128 wall-mounted flexible flaps in an open-channel flow has been studied via numerical simulation. The flaps were modelled directly, so that local effects and interactions were captured. In particular, the coherent waving interactions (honami/monami), and the mechanisms that drive this instability, have been characterised over a range of mass ratio and bending rigidity.

The results show a broad range of behaviours over the tested structural properties. The coherent waving interactions (regular waving mode) were found to occur when the damped natural frequency of the array approached the predicted frequency of the mixing layer instability. As the natural frequency of the array departed from the mixing layer frequency, different behavioural states emerged. For an increase in natural frequency, the regular waving mode transitioned to an irregular waving state, characterised by large spatial and temporal variations in the structural response, before eventually transitioning to a static state with small deflection. For a decrease in natural frequency, the waving mode initially transitioned into a static state with moderate deflection. Further decreasing the natural frequency led to a regime where the mixing layer frequency excited the second structural mode. The observed behaviour in this regime was found to be similar to the flapping behaviour exhibited by single flaps, and is characterised by small scale flow structures. Finally, further decreasing the natural frequency eventually led to another static regime, this time with large deflection.

These results suggest that the coherent waving motion commonly observed in vegetation is in fact a coupled response between the fluid and the array, as opposed to a purely fluid-driven instability. Furthermore, the regular waving motion is linked to a lock-in between the natural frequencies of the fluid and structure. This has important implications for flow control, where similar waving motions and lock-in effects have been observed, and suggests the potential for tuning coupled behaviour via modifying the natural frequencies.

Due to the large parameter space associated with these systems, there are several possible avenues for future research. The effects of Reynolds number, flap spacing, and blockage ratio are all worth investigating. Furthermore, incorporating het-

erogeneity effects, such as spatial variations in flap length, spacing, and material properties would provide results more applicable to real-world applications. Moreover, to test the robustness of the coupled instability and its associated lock-in, unsteady flow conditions, due to gusting or oscillating flow, should be tested over a range of frequencies.

Chapter 10

Conclusions

10.1 Summary

This thesis has described work towards direct modelling of arrays of slender structures, where the individual structures, their dynamics, and the flow around them are fully resolved. Specifically, the main aims were:

- to develop a numerical model that can simulate large arrays of slender structures by fully resolving the individual structures and their dynamics
- to apply this model to multiple configurations involving arrays of slender structures and characterise their behaviour over a range of input conditions
- to highlight key behavioural mechanisms (e.g. waving or lock-in) and relate them to the input parameters (such as flow conditions and structural properties)

To address these aims, a lattice Boltzmann-immersed boundary model was developed and coupled to two different structural models: one based on the Euler-Bernoulli beam equation, and the other based on the finite element method. This model was then applied to a series of test cases, which progressed in scale from a single flap in a periodic domain, to a small finite array, and finally to a large array of flexible flaps. Over a range of flow conditions and structural properties, the main behaviour modes were characterised and quantified.

10.2 Conclusions

The main conclusions from this work can be categorised and summarised as follows:

Single Flap in a Periodic Array

- **Recirculation bridging at high Reynolds numbers**

  At high Reynolds numbers the recirculation region between the periodic flaps fully bridged the gap. This led to the counterintuitive response of decreased deflection for increased flow rate. Such behaviour would not be captured by a homogenised model.

- **Lag between structural response and imposed pressure gradient**

  Under oscillating flow conditions there was a lag between the structural response and the oscillating pressure gradient, which reduced with increasing Reynolds number. This can be attributed to the bridging of the recirculation region, which enabled a faster response.

- **Configuration is not suitable for examining large-scale behaviour**

  While this idealised case is relatively simple and efficient to compute, it cannot capture flow structures that are larger than the periodic domain. Furthermore, the phase lags between flaps also cannot be captured. This makes the periodic configuration unsuitable for studying the typical behaviours associated with large arrays of slender structures, such as coherent waving interactions.

Multiple Flaps in an Oscillating Channel Flow

- **Waving motion is due to vortex passage**

The primary bulk vortex, which is a result of the finite array size, induces a smoothly varying phase lag between the flaps as it passes over the array.

- **Strong correlation between fluid and structure at low mass ratio**

  At low mass ratios the motion of the flaps follow the imposed fluid frequency. This is evidenced by the tip motion and its spectral behaviour.

- **Breakdown of coherence with increasing mass ratio**

  With increasing mass ratio the correlation between the fluid and structural motion breaks down. At the largest mass ratio there is almost no correlation between the fluid and structural motion.

Large Array of Multiple Flaps in an Open-Channel Flow

- **Different behaviour modes for single flap and array of flaps**

  The dynamic states and stability boundaries are markedly different for a single flap and an array of flaps. This confirms the instability mechanisms are separate.

- **Behavioural state depends strongly on material properties**

  This indicates the coherent waving motion is not a purely fluid-driven instability, but is in fact a coupled response between both the fluid and the array.

- **Coherent waving motions are associated with a lock-in effect**

  This indicates that the coherent waving motion is a coupled instability that is associated with a lock-in between the natural frequencies of the mixing layer and the array.

- **Flapping mode for high-mass/low-stiffness cases**

For high-mass/low-stiffness cases, the natural frequency of the array was shifted sufficiently low so that the mixing layer instability excited the second mode of the flaps. The deformation of the flaps in this regime is similar to the flapping mode exhibited for single flaps, and is characterised by small-scale flow structures.

10.3 Limitations

The main limitations of this work, and their significance, are described below:

- **2D simulations**

  The restriction to 2D simulations limits the applicability of the present findings to cases where the out-of-plane effects can be neglected. Therefore, while plate-like structures with negligible edge effects are considered within the scope of this work, the present conclusions cannot be extended to hair-like structures, which are characterised by highly 3D flow fields. Since many of the real-world applications of slender structures can be described as hair-like structures, this is a serious limitation of the present work. However, the restriction to 2D cases allows a basic understanding of the main behaviours before extending to more complex problems. Thus, it is a necessary step in the path to understanding the complex 3D behaviour of slender structures.

- **Low-moderate Reynolds numbers**

  For all of the cases described in this work, the Reynolds number was limited to $\mathcal{O}(100)$. While there are additional modelling challenges associated with simulating higher Reynolds numbers, this was not the primary reason for imposing this constraint. Instead, the main reason was to simplify the physics of the problem – the values used throughout this work were chosen so that they were large enough to allow inertia-driven instabilities to develop, while at the same time low enough to avoid instabilities due to transition and turbulence affects. Since many applications involving slender structures operate in regimes outside of this range, extending the present findings to these applications is not possible. Although many studies on slender

structures have reported a limited dependence on Reynolds number, it is not possible to say the same for the cases described here without further work. Furthermore, it is expected that below a critical value of Reynolds number there will be no dynamic instabilities as viscous effects begin to dominate the flow.

- **Idealised cases**

  Another limitation is the idealised cases which have been adopted in the present work. The assumptions of uniform spacing, length, and material properties do not hold in real-world applications in general. Therefore, it is difficult to translate these findings directly to real-world applications. However, one of the advantages of the present method, which directly models the individual structures, is that it can readily incorporate these effects.

- **Confined flows**

  The cases described in this work have all been restricted to either internal flows or open-channel flows with large confinement ratios. The main motivation for this was computational cost, since simulating external flows is considerably more expensive, especially in the present approach which is limited to a square lattice. Since many applications of slender structures involve either external flows or internal flows with low confinement ratios, this limits the applicability of the current findings. That being said, there are some important applications that have been discussed which do have similar configurations, such as aquatic vegetation.

- **Small parameter space**

  Finally, although a large range of parameters were tested in the present work, this is only a small portion of the parameter space. Therefore, it is important to stress that the findings of the current work are only applicable to the range of parameters that have been investigated.

10.4 Future Work

To address some of the limitations discussed above, a number of possible avenues for future work are available. The most obvious one is the extension of this work to 3D. Not only will this provide a greater range of applicability, but it is also likely to introduce a richer set of physical behaviours. However, given the extra complexity associated with modelling 3D arrays of slender structures, initial work should focus on extending current knowledge regarding the 2D behaviour. The effects of Reynolds number, flap spacing, and blockage ratio are all worth investigating. For very densely-packed arrays, where the contact between adjacent structures is substantial, some form of collision model would need to be developed. Furthermore, incorporating heterogeneity effects, such as spatial variations in flap length, spacing, and material properties, or temporal variations in the flow conditions, would provide results more applicable to real-world applications. Moreover, since the present model can incorporate such effects readily, this would provide a natural extension of the present work.

References

[1] J.M.J. den Toonder and P.R. Onck. Microfluidic manipulation with artificial/bioinspired cilia. *Trends Biotechnol.*, 31(2):85–91, 2013.

[2] J. Favier, A. Dauptain, D. Basso, and A. Bottaro. Passive separation control using a self-adaptive hairy coating. *J. Fluid Mech.*, 627:451–483, 2009.

[3] P.R. Heyliger. When beam theories fail. *J. Mech. Mater. Struct.*, 8(1):15–35, 2013.

[4] M.J. Shelley and J. Zhang. Flapping and Bending Bodies Interacting with Fluid Flows. *Annu. Rev. Fluid Mech.*, 43:449–465, 2011.

[5] W. Huang and H.J. Sung. Three-dimensional simulation of a flapping flag in a uniform flow. *J. Fluid Mech.*, 653:301–336, 2010.

[6] E. Lauga and C. Eloy. Shape of optimal active flagella. *J. Fluid Mech.*, 730: R1, 2013.

[7] M. Ghisalberti. Mixing layers and coherent structures in vegetated aquatic flows. *J. Geophys. Res.*, 107(C2):3011, 2002.

[8] Emmanuel de Langre. Effects of Wind on Plants. *Annu. Rev. Fluid Mech.*, 40(1):141–168, 2008.

[9] S. Kunze and C. Brücker. Control of vortex shedding on a circular cylinder using self-adaptive hairy-flaps. *Comptes Rendus Mécanique*, 340(1-2):41–56, 2012.

[10] C. Py, E. de Langre, B. Moulia, and P. Hémon. Measurement of wind-induced motion of crop canopies from digital video images. *Agric. For. Meteorol.*, 130(3-4):223–236, 2005.

[11] M. Ghisalberti and H. Nepf. The Structure of the Shear Layer in Flows over Rigid and Flexible Canopies. *Environ. Fluid Mech.*, 6(3):277–301, 2006.

[12] T. Okamoto, I. Nezu, and M. Sanjou. Flow-vegetation interactions: length-scale of the monami phenomenon. *J. Hydraul. Res.*, 54(3):251–262, 2016.

[13] B.S.H. Connell and D.K.P. Yue. Flapping dynamics of a flag in a uniform stream. *J. Fluid Mech.*, 581:33–67, 2007.

[14] W. Huang, S.J. Shin, and H.J. Sung. Simulation of flexible filaments in a uniform flow by the immersed boundary method. *J. Comput. Phys.*, 226(2): 2206–2228, 2007.

[15] S. Dupont, F. Gosselin, C. Py, E. de Langre, P. Hemon, and Y. Brunet. Modelling waving crops using large-eddy simulation: comparison with experiments and a linear stability analysis. *J. Fluid Mech.*, 652:5–44, 2010.

[16] S. Alben and M.J. Shelley. Flapping States of a Flag in an Inviscid Fluid: Bistability and the Transition to Chaos. *Phys. Rev. Lett.*, 100:074301, 2008.

[17] S. Michelin, S.G. Llewellyn Smith, and B.J. Glover. Vortex shedding model of a flapping flag. *J. Fluid Mech.*, 617:1–10, 2008.

[18] J. O'Connor, A. Revell, P. Mandal, and P. Day. Application of a lattice Boltzmann-immersed boundary method for fluid-filament dynamics and flow sensing. *J. Biomech.*, 49(11), 2016.

[19] J. O'Connor, P. Day, P. Mandal, and A. Revell. Computational fluid dynamics in the microcirculation and microfluidics: what role can the lattice Boltzmann method play? *Integr. Biol.*, 8(5):589–602, 2016.

[20] J. Favier, C. Li, L. Kamps, A. Revell, J. O'Connor, and C. Brücker. The PELskin project - part I: fluid-structure interaction for a row of flexible flaps: a reference study in oscillating channel flow. *Meccanica*, 52(8), 2017.

[21] A. Revell, J. O'Connor, A. Sarkar, C. Li, J. Favier, L. Kamps, and C. Brücker. The PELskin project: part II - investigating the physical coupling between flexible filaments in an oscillating flow. *Meccanica*, 52(8), 2017.

[22] A.R.G. Harwood, J. O'Connor, J.S. Muñoz, M.C. Santasmasas, and A.J. Revell. LUMA: A many-core, Fluid-Structure Interaction solver based on the Lattice-Boltzmann Method. *SoftwareX*, 7:88–94, 2018.

[23] R. Ghosh, G.A. Buxton, O.B. Usta, A.C. Balazs, and A. Alexeev. Designing Oscillating Cilia That Capture or Release Microscopic Particles. *Langmuir*, 26(4):2963–2968, 2010.

[24] H. Masoud and A. Alexeev. Harnessing synthetic cilia to regulate motion of microparticles. *Soft Matter*, 7(19):8702–8708, 2011.

[25] S. Weinbaum, J.M. Tarbell, and E.R. Damiano. The Structure and Function of the Endothelial Glycocalyx Layer. *Annu. Rev. Biomed. Eng.*, 9:121–167, 2007.

[26] B.M. Fu and J.M. Tarbell. Mechano-sensing and transduction by endothelial surface glycocalyx: composition, structure, and function. *Wiley Interdiscip. Rev. Syst. Biol. Med.*, 5(3):381–390, 2012.

[27] J. Casas and O. Dangles. Physical Ecology of Fluid Flow Sensing in Arthropods. *Annu. Rev. Entomol.*, 55:505–520, 2010.

[28] S. Sterbing-D'Angelo, M. Chadha, C. Chiu, B. Falk, W. Xian, J. Barcelo, J.M. Zook, and C.F. Moss. Bat wing sensors support flight control. *Proc. Natl. Acad. Sci.*, 108(27):11291–11296, 2011.

[29] M. Asadnia, A.G.P. Kottapalli, K. Karavitaki, M.E. Warkiani, J. Miao, D.P. Corey, and M. Triantafyllou. From Biological Cilia to Artificial Flow Sensors: Biomimetic Soft Polymer Nanosensors with High Sensing Performance. *Sci. Rep.*, 6:32955, 2016.

[30] A.G.P. Kottapalli, M. Bora, M. Asadnia, J. Miao, S.S. Venkatraman, and M. Triantafyllou. Nanofibril scaffold assisted MEMS artificial hydrogel neuromasts for enhanced sensitivity flow sensing. *Sci. Rep.*, 6:19336, 2016.

[31] C.H.J. Wang and J. Schlüter. Stall control with feathers: Self-activated flaps on finite wings at low Reynolds numbers. *Comptes Rendus Mécanique*, 340(1-2):57–66, 2012.

[32] C. Brücker and C. Weidner. Influence of self-adaptive hairy flaps on the stall delay of an airfoil in ramp-up motion. *J. Fluids Struct.*, 47:31–40, 2014.

[33] I. Rechenberg, R. Bannasch, G. Patone, and W. Müller. Aeroflexible Oberflächenklappen als 'Rückstrombremsen' nach dem Vorbild der Deckfedern des Vogelflügels, Statusbericht 1995 für das BMBF-Vorhaben 13N6536. Technical report, Inst. f. Bionik u. Evolutionstechnik, TU Berlin, 1995.

[34] H.M. Nepf and M. Ghisalberti. Flow and transport in channels with submerged vegetation. *Acta Geophys.*, 56(3):753–777, 2008.

[35] K.B. Gedan, M.L. Kirwan, E. Wolanski, E.B. Barbier, and B.R. Silliman. The present and future role of coastal wetland vegetation in protecting shorelines: answering recent challenges to the paradigm. *Clim. Change*, 106 (1):7–29, 2011.

[36] J.M.J. den Toonder, F. Bos, D. Broer, L. Filippini, M. Gillies, J. de Goede, T. Mol, M. Reijme, W. Talen, H. Wilderbeek, V.V. Khatavkar, and P.D. Anderson. Artificial cilia for active micro-fluidic mixing. *Lab Chip*, 8(4): 533–541, 2008.

[37] A. Tripathi, A. Bhattacharya, and A.C. Balazs. Size Selectivity in Artificial Cilia-Particle Interactions: Mimicking the Behavior of Suspension Feeders. *Langmuir*, 29(14):4616–4621, 2013.

[38] J. Tao and X. Yu. Hair flow sensors: from bio-inspiration to bio-mimicking - a review. *Smart Mater. Struct.*, 21(11):113001, 2012.

[39] J. Niu and D.L. Hu. Drag reduction of a hairy disk. *Phys. Fluids*, 23(10): 101701, 2011.

[40] L. Schouveiler, C. Eloy, and P. Le Gal. Flow-induced vibrations of high mass ratio flexible filaments freely hanging in a flow. *Phys. Fluids*, 17(4): 047104, 2005.

[41] L. Jia and X. Yin. Response modes of a flexible filament in the wake of a cylinder in a flowing soap film. *Phys. Fluids*, 21(10):101704, 2009.

[42] J. Wu, J. Wu, and J. Zhan. Characteristics of flow over a circular cylinder with two attached filaments. *J. Fluids Struct.*, 66:269–281, 2016.

[43] J.J. Allen and A.J. Smits. Energy Harvesting Eel. *J. Fluids Struct.*, 15(3-4): 629–640, 2001.

[44] G.W. Taylor, J.R. Burns, S.M. Kammann, W.B. Powers, and T.R. Welsh. The Energy Harvesting Eel: a small subsurface ocean/river power generator. *IEEE J. Ocean. Eng.*, 26(4):539–547, 2001.

[45] S. Orrego, K. Shoele, A. Ruas, K. Doran, B. Caggiano, R. Mittal, and S.H. Kang. Harvesting ambient wind energy with an inverted piezoelectric flag. *Appl. Energy*, 194:212–222, 2017.

[46] J.A. Dunnmon, S.C. Stanton, B.P. Mann, and E.H. Dowell. Power extraction from aeroelastic limit cycle oscillations. *J. Fluids Struct.*, 27(8):1182–1198, 2011.

[47] M.S. Kumar and P. Philominathan. The physics of flagellar motion of E. coli during chemotaxis. *Biophys. Rev.*, 2(1):13–20, 2010.

[48] T. Leclercq and E. de Langre. Drag reduction by elastic reconfiguration of non-uniform beams in non-uniform flows. *J. Fluids Struct.*, 60:114–129, 2016.

[49] R. Hua, L. Zhu, and X. Lu. Dynamics of fluid flow over a circular flexible plate. *J. Fluid Mech.*, 759:56–72, 2014.

[50] S. Vogel. Drag and Flexibility in Sessile Organisms. *Am. Zool.*, 24(1):37–44, 1984.

[51] S. Vogel. Drag and Reconfiguration of Broad Leaves in High Winds. *J. Exp. Bot.*, 40(8):941–948, 1989.

[52] E. de Langre, A. Gutierrez, and J. Cossé. On the scaling of drag reduction by reconfiguration in plants. *Comptes Rendus Mécanique*, 340(1-2):35–40, 2012.

[53] S. Alben, M. Shelley, and J. Zhang. Drag reduction through self-similar bending of a flexible body. *Nature*, 420:479–481, 2002.

[54] S. Alben, M. Shelley, and J. Zhang. How flexibility induces streamlining in a two-dimensional flow. *Phys. Fluids*, 16(5):1694–1713, 2004.

[55] F. Gosselin, E. de Langre, and B.A. Machado-Almeida. Drag reduction of flexible plates by reconfiguration. *J. Fluid Mech.*, 650:319–341, 2010.

[56] L. Zhu and C.S. Peskin. Drag of a flexible fiber in a 2D moving viscous fluid. *Comput. Fluids*, 36(2):398–406, 2007.

[57] L. Zhu. Scaling laws for drag of a compliant body in an incompressible viscous flow. *J. Fluid Mech.*, 607:387–400, 2008.

[58] J. Zhang, S. Childress, A. Libchaber, and M. Shelley. Flexible filaments in a flowing soap film as a model for one-dimensional flags in a two-dimensional wind. *Nature*, 408:835–839, 2000.

[59] L. Zhu and C.S. Peskin. Simulation of a Flapping Flexible Filament in a Flowing Soap Film by the Immersed Boundary Method. *J. Comput. Phys.*, 179(2):452–468, 2002.

[60] L. Zhu and C.S. Peskin. Interaction of two flapping filaments in a flowing soap film. *Phys. Fluids*, 15(7):1954–1960, 2003.

[61] D.J.J. Farnell, T. David, and D.C. Barton. Coupled states of flapping flags. *J. Fluids Struct.*, 19(1):29–36, 2004.

[62] M. Shelley, N. Vandenberghe, and J. Zhang. Heavy Flags Undergo Spontaneous Oscillations in Flowing Water. *Phys. Rev. Lett.*, 94:094302, 2005.

[63] C. Eloy, C. Souilliez, and L. Schouveiler. Flutter of a rectangular plate. *J. Fluids Struct.*, 23(6):904–919, 2007.

[64] L. Tang and M.P. Païdoussis. On the instability and the post-critical behaviour of two-dimensional cantilevered flexible plates in axial flow. *J. Sound Vib.*, 305(1-2):97–115, 2007.

[65] O. Doaré and S. Michelin. Piezoelectric coupling in energy-harvesting fluttering flexible plates: linear stability analysis and conversion efficiency. *J. Fluids Struct.*, 27(8):1357–1375, 2011.

[66] S. Michelin and O. Doaré. Energy harvesting efficiency of piezoelectric flags in axial flows. *J. Fluid Mech.*, 714:489–504, 2013.

[67] K. Shoele and R. Mittal. Energy harvesting by flow-induced flutter in a simple model of an inverted piezoelectric flag. *J. Fluid Mech.*, 790:582–606, 2016.

[68] D. Kim, J. Cossé, C.H. Cerdeira, and M. Gharib. Flapping dynamics of an inverted flag. *J. Fluid Mech.*, 736:R1, 2013.

[69] J.E. Sader, J. Cossé, D. Kim, B. Fan, and M. Gharib. Large-amplitude flapping of an inverted flag in a uniform steady flow - a vortex-induced vibration. *J. Fluid Mech.*, 793:524–555, 2016.

[70] P.S. Gurugubelli and R.K. Jaiman. Self-induced flapping dynamics of a flexible inverted foil in a uniform flow. *J. Fluid Mech.*, 781(2015):657–694, 2015.

[71] J. Ryu, S.G. Park, B. Kim, and H.J. Sung. Flapping dynamics of an inverted flag in a uniform flow. *J. Fluids Struct.*, 57:159–169, 2015.

[72] C. Tang, N. Liu, and X. Lu. Dynamics of an inverted flexible plate in a uniform flow. *Phys. Fluids*, 27(7):073601, 2015.

[73] T.I. Marjoribanks, R.J. Hardy, S.N. Lane, and D.R. Parsons. Does the canopy mixing layer model apply to highly flexible aquatic vegetation? Insights from numerical modelling. *Environ. Fluid Mech.*, 17(2):277–301, 2017.

[74] F.P. Gosselin and E. de Langre. Drag reduction by reconfiguration of a poroelastic system. *J. Fluids Struct.*, 27(7):1111–1123, 2011.

[75] S. Barsu, D. Doppler, J.J.S. Jerome, N. Rivière, and M. Lance. Drag measurements in laterally confined 2D canopies: Reconfiguration and sheltering effect. *Phys. Fluids*, 28(10):107101, 2016.

[76] H.M. Nepf. Hydrodynamics of vegetated channels. *J. Hydraul. Res.*, 50(3): 262–279, 2012.

[77] H.M. Nepf. Flow and Transport in Regions with Aquatic Vegetation. *Annu. Rev. Fluid Mech.*, 44:123–142, 2012.

[78] M. Ghisalberti and H.M. Nepf. Shallow Flows Over a Permeable Medium: The Hydrodynamics of Submerged Aquatic Canopies. *Transp. Porous Media*, 78:309–326, 2009.

[79] C.M. Ho and P. Huerre. Perturbed Free Shear Layers. *Annu. Rev. Fluid Mech.*, 16:365–422, 1984.

[80] H.M. Nepf and E.R. Vivoni. Flow structure in depth-limited, vegetated flow. *J. Geophys. Res.*, 105(C12):28547–28557, 2000.

[81] C. Py, E. de Langre, and B. Moulia. The mixing layer instability of wind over a flexible crop canopy. *Comptes Rendus Mécanique*, 332(8):613–618, 2004.

[82] C. Py, E. de Langre, and B. Moulia. A frequency lock-in mechanism in the interaction between wind and crop canopies. *J. Fluid Mech.*, 568:425–449, 2006.

[83] F. Gosselin and E. de Langre. Destabilising effects of plant flexibility in air and aquatic vegetation canopy flows. *Eur. J. Mech. - B/Fluids*, 28(2): 271–282, 2009.

[84] N. Aubry, R. Guyonnet, and R. Lima. Spatiotemporal analysis of complex signals: Theory and applications. *J. Stat. Phys.*, 64(3-4):683–739, 1991.

[85] P. Hémon and F. Santi. Applications of biorthogonal decompositions in fluid-structure interactions. *J. Fluids Struct.*, 17(8):1123–1143, 2003.

[86] D. Venkataraman and A. Bottaro. Numerical modeling of flow control on a symmetric aerofoil via a porous, compliant coating. *Phys. Fluids*, 24:093601, 2012.

[87] N. Mazellier, A. Feuvrier, and A. Kourta. Biomimetic bluff body drag reduction by self-adaptive porous flaps. *Comptes Rendus Mécanique*, 340 (1-2):81–94, 2012.

[88] M. Argentina and L. Mahadevan. Fluid-flow-induced flutter of a flag. *Proc. Natl. Acad. Sci.*, 102(6):1829–1834, 2005.

[89] O. Doaré, M. Sauzade, and C. Eloy. Flutter of an elastic plate in a channel flow: Confinement and finite-size effects. *J. Fluids Struct.*, 27(1):76–88, 2011.

[90] S. Alben. The flapping-flag instability as a nonlinear eigenvalue problem. *Phys. Fluids*, 20(10):104106, 2008.

[91] C. Eloy, R. Lagrange, C. Souilliez, and L. Schouveiler. Aeroelastic instability of cantilevered flexible plates in uniform flow. *J. Fluid Mech.*, 611:97–106, 2008.

[92] R. Singh, M.M. Bandi, A. Mahadevan, and S. Mandre. Linear stability analysis for monami in a submerged seagrass bed. *J. Fluid Mech.*, 786:R1, 2016.

[93] S. Ikeda, T. Yamada, and Y. Toda. Numerical study on turbulent flow and honami in and above flexible plant canopy. *Int. J. Heat Fluid Flow*, 22(3): 252–258, 2001.

[94] J. Cui and V.S. Neary. LES study of turbulent flows with submerged vegetation. *J. Hydraul. Res.*, 46(3):307–316, 2008.

[95] A. Chakrabarti, Q. Chen, H.D. Smith, and D. Liu. Large Eddy Simulation of Unidirectional and Wave Flows through Vegetation. *J. Eng. Mech.*, 142 (8):04016048, 2016.

[96] A. Goza and T. Colonius. A strongly-coupled immersed-boundary formulation for thin elastic structures. *J. Comput. Phys.*, 336:401–411, 2017.

[97] M. Maza, J.L. Lara, and I.J. Losada. Tsunami wave interaction with mangrove forests: A 3-D numerical approach. *Coast. Eng.*, 98:33–54, 2015.

[98] P. Meakin and A.M. Tartakovsky. Modeling and simulation of pore-scale multiphase fluid flow and reactive transport in fractured and porous media. *Rev. Geophys.*, 47(3):1–47, 2009.

[99] S. Succi. *The Lattice Boltzmann Equation for Fluid Dynamics and Beyond*. Oxford University Press, Oxford, 1st edition, 2001.

[100] T. Krüger, H. Kusumaatmaja, A. Kuzmin, O. Shardt, G. Silva, and E.M. M Viggen. *The Lattice Boltzmann Method: Principles and Practice*. Springer International Publishing, 1st edition, 2017.

[101] R. Benzi, S. Succi, and M. Vergassola. The Lattice Boltzmann Equation: Theory and Applications. *Phys. Rep.*, 222(3):145–197, 1992.

[102] D.A. Wolf-Gladrow. *Lattice-Gas Cellular Automata and Lattice Boltzmann Models: An Introduction*. Springer-Verlag Berlin Heidelberg, 1st edition, 2000.

[103] S. Wolfram. Cellular Automation Fluids 1: Basic Theory. *J. Stat. Phys.*, 45 (3-4):471–526, 1986.

[104] J. Hardy, Y. Pomeau, and O. de Pazzis. Time Evolution of a Two-Dimensional Classical Lattice System. *Phys. Rev. Lett.*, 31(5):276–279, 1973.

[105] U. Frisch, B. Hasslacher, and Y. Pomeau. Lattice-Gas Automata for the Navier-Stokes Equation. *Phys. Rev. Lett.*, 56(14):1505–1508, 1986.

[106] G.R. McNamara and G. Zanetti. Use of the Boltzmann Equation to Simulate Lattice-Gas Automata. *Phys. Rev. Lett.*, 61(20):2332–2335, 1988.

[107] F.J. Higuera and J. Jiménez. Boltzmann Approach to Lattice Gas Simulations. *Europhys. Lett.*, 9(7):663–668, 1989.

[108] F.J. Higuera, S. Succi, and R. Benzi. Lattice Gas Dynamics with Enhanced Collisions. *Europhys. Lett.*, 9(4):345–349, 1989.

[109] A. De Rosis, S. Ubertini, and F. Ubertini. A partitioned approach for two-dimensional fluid-structure interaction problems by a coupled lattice Boltzmann-finite element method with immersed boundary. *J. Fluids Struct.*, 45:202–215, 2014.

[110] J. Favier, A. Revell, and A. Pinelli. A Lattice Boltzmann-Immersed Boundary method to simulate the fluid interaction with moving and slender flexible objects. *J. Comput. Phys.*, 261:145–161, 2014.

[111] J. Favier, A. Revell, and A. Pinelli. Numerical study of flapping filaments in a uniform fluid flow. *J. Fluids Struct.*, 53:26–35, 2015.

[112] C. Sun and L.L. Munn. Lattice Boltzmann simulation of blood flow in digitized vessel networks. *Comput. Math. with Appl.*, 55(7):1594–1600, 2008.

[113] T. Krüger, D. Holmes, and P.V. Coveney. Deformability-based red blood cell separation in deterministic lateral displacement devices - A simulation study. *Biomicrofluidics*, 8(5):054114, 2014.

[114] J.P. Arata and A. Alexeev. Designing microfluidic channel that separates elastic particles upon stiffness. *Soft Matter*, 5(14):2721–2724, 2009.

[115] S. Kobel, A. Valero, J. Latt, P. Renaud, and M. Lutolf. Optimization of microfluidic single cell trapping for long-term on-chip culture. *Lab Chip*, 10 (7):857–863, 2010.

[116] A.J.C. Ladd. Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 1. Theoretical foundation. *J. Fluid Mech.*, 271:285–309, 1994.

[117] A.J.C. Ladd. Numerical simulations of particulate suspensions via a discretized Boltzmann equation. Part 2. Numerical results. *J. Fluid Mech.*, 271:311–339, 1994.

[118] Z. Feng and E.E. Michaelides. The immersed boundary-lattice Boltzmann method for solving fluid-particles interaction problems. *J. Comput. Phys.*, 195(2):602–628, 2004.

[119] X. Shan and X. He. Discretization of the Velocity Space in the Solution of the Boltzmann Equation. *Phys. Rev. Lett.*, 80(1):65–68, 1998.

[120] E.M Viggen. *The lattice Boltzmann method: Fundamentals and acoustics*. Phd thesis, Norwegian University of Science and Technology, 2014.

[121] X. He and L. Luo. A priori derivation of the lattice Boltzmann equation. *Phys. Rev. E*, 55(6):R6333, 1997.

[122] X. He and L. Luo. Theory of the lattice Boltzmann method: From the Boltzmann equation to the lattice Boltzmann equation. *Phys. Rev. E*, 56 (6):6811–6817, 1997.

[123] X. Shan, X. Yuan, and H. Chen. Kinetic theory representation of hydrodynamics: a way beyond the Navier-Stokes equation. *J. Fluid Mech.*, 550: 413–441, 2006.

[124] P.L. Bhatnagar, E.P. Gross, and M. Krook. A Model for Collision Processes in Gases. I. Small Amplitude Processes in Charged and Neutral One-Component Systems. *Phys. Rev.*, 94(3):511–525, 1954.

[125] H. Chen, S. Chen, and W.H. Matthaeus. Recovery of the Navier-Stokes equations using a lattice-gas Boltzmann method. *Phys. Rev. A*, 45(8): 5339–5342, 1992.

[126] Y.H. Qian, D. D'Humières, and P. Lallemand. Lattice BGK Models for Navier-Stokes Equation. *Europhys. Lett.*, 17(6):479–484, 1992.

[127] N. Delbosc. *Real-Time Simulation of Indoor Air Flow Using the Lattice Boltzmann Method on Graphics Processing Unit*. Phd thesis, University of Leeds, 2015.

[128] J.D. Sterling and S. Chen. Stability Analysis of Lattice Boltzmann Methods. *J. Comput. Phys.*, 123(1):196–206, 1996.

[129] D. D'Humières. Generalized Lattice-Boltzmann Equations. In D.P. Weaver and B.D. Shizgal, editors, *Rarefied Gas Dyn. Theory Simulations*, pages 450–458. American Institute of Aeronautics and Astronautics, 1992.

[130] Z. Guo, C. Zheng, and B. Shi. Discrete lattice effects on the forcing term in the lattice Boltzmann method. *Phys. Rev. E*, 65(4):046308, 2002.

[131] P.A. Skordos. Initial and boundary conditions for the lattice Boltzmann method. *Phys. Rev. E*, 48(6):4823–4842, 1993.

[132] X. He, Q. Zou, L. Luo, and M. Dembo. Analytic solutions of simple flows and analysis of nonslip boundary conditions for the lattice Boltzmann BGK model. *J. Stat. Phys.*, 87(1-2):115–136, 1997.

[133] M.A. Gallivan, D.R. Nobel, J.G. Georgiadis, and R.O. Buckius. An Evaluation of the Bounce-Back Boundary Condition for Lattice Boltzmann Simulations. *Int. J. Numer. Methods Fluids*, 25(3):249–263, 1997.

[134] M. Bouzidi, M. Firdaouss, and P. Lallemand. Momentum transfer of a Boltzmann-lattice fluid with boundaries. *Phys. Fluids*, 13(11):3452–3459, 2001.

[135] T. Inamuro, M. Yoshino, and F. Ogino. A nonslip boundary condition for lattice Boltzmann simulations. *Phys. Fluids*, 7(12):2928–2930, 1995.

[136] J.C.G. Verschaeve. Analysis of the lattice Boltzmann Bhatnagar-Gross-Krook no-slip boundary condition: Ways to improve accuracy and stability. *Phys. Rev. E*, 80:036703, 2009.

[137] J. Latt. *Hydrodynamic limit of lattice Boltzmann equations*. Phd thesis, Université de Genève, 2007.

[138] J. Latt, B. Chopard, O. Malaspinas, M. Deville, and A. Michler. Straight velocity boundaries in the lattice Boltzmann method. *Phys. Rev. E*, 77(5): 056703, 2008.

[139] Q. Zou and X. He. On pressure and velocity boundary conditions for the lattice Boltzmann BGK model. *Phys. Fluids*, 9(6):1591–1598, 1997.

[140] X. He and L. Luo. Lattice Boltzmann Model for the Incompressible Navier-Stokes Equation. *J. Stat. Phys.*, 88(3-4):927–944, 1997.

[141] M.B. Reider and J.D. Sterling. Accuracy of discrete-velocity BGK models for the simulation of the incompressible Navier-Stokes equations. *Comput. Fluids*, 24(4):459–467, 1995.

[142] Y.T. Feng, K. Han, and D.R.J. Owen. Coupled lattice Boltzmann method and discrete element modelling of particle transport in turbulent fluid flows: Computational issues. *Int. J. Numer. Methods Fluids*, 72(9):1111–1134, 2007.

[143] D.J. Holdych, D.R. Noble, J.G. Georgiadis, and R.O. Buckius. Truncation error analysis of lattice Boltzmann methods. *J. Comput. Phys.*, 193(2): 595–619, 2004.

[144] T. Krüger. *Computer Simulation Study of Collective Phenomena in Dense Suspensions of Red Blood Cells under Shear*. Phd thesis, Eindhoven University of Technology, 2012.

[145] A. Labuschagne, N.F.J. van Rensburg, and A.J. van der Merwe. Comparison of linear beam theories. *Math. Comput. Model.*, 49(1-2):20–30, 2009.

[146] P. Causin, J.F. Gerbeau, and F. Nobile. Added-mass effect in the design of partitioned algorithms for fluidstructure problems. *Comput. Methods Appl. Mech. Eng.*, 194(42-44):4506–4527, 2005.

[147] C. Wang and J.D. Eldredge. Strongly coupled dynamics of fluids and rigid-body systems with the immersed boundary projection method. *J. Comput. Phys.*, 295:87–113, 2015.

[148] R. de Borst, M.A. Crisfield, J.J.C. Remmers, and C.V. Verhoosel. *Non-Linear Finite Element Analysis of Solids and Structures*. John Wiley & Sons Ltd, second edition, 2012.

[149] K.J. Bathe. *Finite Element Procedures*. Prentice Hall, second edition, 2014.

[150] A. De Rosis, G. Falcucci, S. Ubertini, and F. Ubertini. A coupled lattice Boltzmann-finite element approach for two-dimensional fluid-structure interaction. *Comput. Fluids*, 86:558–568, 2013.

[151] T.C. Fung. Unconditionally stable higher-order Newmark methods by sub-stepping procedure. *Comput. Methods Appl. Mech. Eng.*, 147(1-2):61–84, 1997.

[152] O.C. Zienkiewicz, R.L. Taylor, and J.Z. Zhu. *The Finite Element Method: Its Basis and Fundamentals*. Butterworth-Heinemann, seventh edition, 2013.

[153] O.C. Zienkiewicz, R.L. Taylor, and D. Fox. *The Finite Element Method for Solid and Structural Mechanics*. Butterworth-Heinemann, seventh edition, 2014.

[154] C.S. Peskin. Flow patterns around heart valves: A numerical method. *J. Comput. Phys.*, 10(2):252–271, 1972.

[155] C.S. Peskin. Numerical analysis of blood flow in the heart. *J. Comput. Phys.*, 25(3):220–252, 1977.

[156] Rajat Mittal and Gianluca Iaccarino. Immersed Boundary Methods. *Annu. Rev. Fluid Mech.*, 37(1):239–261, 2005.

[157] J. Wu and C. Shu. Implicit velocity correction-based immersed boundary-lattice Boltzmann method and its applications. *J. Comput. Phys.*, 228(6): 1963–1979, 2009.

[158] J. Mohd-Yusof. Combined immersed-boundary/B-spline methods for simulations of flow in complex geometries. Technical report, NASA Ames/Stanford University, 1997.

[159] E.A. Fadlun, R. Verzicco, P. Orlandi, and J. Mohd-Yusof. Combined Immersed-Boundary Finite-Difference Methods for Three-Dimensional Complex Flow Simulations. *J. Comput. Phys.*, 161(1):35–60, 2000.

[160] M. Uhlmann. An immersed boundary method with direct forcing for the simulation of particulate flows. *J. Comput. Phys.*, 209(2):448–476, 2005.

[161] Z. Feng and E.E. Michaelides. Proteus: a direct forcing method in the simulations of particulate flows. *J. Comput. Phys.*, 202(1):20–51, 2005.

[162] A. Pinelli, I.Z. Naqavi, U. Piomelli, and J. Favier. Immersed-boundary methods for general finite-difference and finite-volume Navier-Stokes solvers. *J. Comput. Phys.*, 229(24):9073–9091, 2010.

[163] Z. Li, J. Favier, U. D'Ortona, and S. Poncet. An immersed boundary-lattice Boltzmann method for single- and multi-component fluid flows. *J. Comput. Phys.*, 304:424–440, 2016.

[164] C.S. Peskin. The immersed boundary method. *Acta Numer.*, 11:479–517, 2002.

[165] A.M. Roma, C.S. Peskin, and M.J. Berger. An Adaptive Version of the Immersed Boundary Method. *J. Comput. Phys.*, 153(2):509–534, 1999.

[166] B. Dünweg and A.J.C. Ladd. Lattice Boltzmann Simulations of Soft Matter Systems. In *Adv. Polym. Sci. 2008*, pages 1–78. Springer Berlin Heidelberg, 2008.

[167] Z. Li and J. Favier. A non-staggered coupling of finite element and lattice Boltzmann methods via an immersed boundary scheme for fluid-structure interaction. *Comput. Fluids*, 143:90–102, 2017.

[168] Z. Feng and E.E. Michaelides. Robust treatment of no-slip boundary condition and velocity updating for the lattice-Boltzmann simulation of particulate flows. *Comput. Fluids*, 38(2):370–381, 2009.

[169] Gene Hou, Jin Wang, and Anita Layton. Numerical Methods for Fluid-Structure Interaction A Review. *Commun. Comput. Phys.*, 12(02):337–377, 2012.

[170] S. Kollmannsberger. *ALE-type and fixed grid fluid-structure interaction involving the p-version of the Finite Element Method*. Phd thesis, Technische Universität München, 2010.

[171] F. Sotiropoulos and X. Yang. Immersed boundary methods for simulating fluidstructure interaction. *Prog. Aerosp. Sci.*, 65:1–21, 2014.

[172] C. Farhat and M. Lesoinne. Two efficient staggered algorithms for the serial and parallel solution of three-dimensional nonlinear transient aeroelastic problems. *Comput. Methods Appl. Mech. Eng.*, 182(3-4):499–515, 2000.

[173] C. Förster, W.A. Wall, and E. Ramm. Artificial added mass instabilities in sequential staggered coupling of nonlinear structures and incompressible viscous flows. *Comput. Methods Appl. Mech. Eng.*, 196(7):1278–1293, 2007.

[174] E.H. van Brummelen. Added Mass Effects of Compressible and Incompressible Flows in Fluid-Structure Interaction. *J. Appl. Mech.*, 76(2):021206, 2009.

[175] S. Piperno and C. Farhat. Partitioned procedures for the transient solution of coupled aeroelastic problems - Part II: energy transfer analysis and three-dimensional applications. *Comput. Methods Appl. Mech. Eng.*, 190(24-25): 3147–3170, 2001.

[176] C. Farhat, K.G. van der Zee, and P. Geuzaine. Provably second-order time-accurate loosely-coupled solution algorithms for transient nonlinear computational aeroelasticity. *Comput. Methods Appl. Mech. Eng.*, 195 (17-18):1973–2001, 2006.

[177] C. Farhat, A. Rallu, K. Wang, and T. Belytschko. Robust and provably second-order explicit-explicit and implicit-explicit staggered time-integrators for highly non-linear compressible fluid-structure interaction problems. *Int. J. Numer. Methods Eng.*, 84(1):73–107, 2010.

[178] J. Degroote. Partitioned Simulation of Fluid-Structure Interaction. *Arch. Comput. Methods Eng.*, 20(3):185–238, 2013.

[179] B.M. Irons and R.C. Tuck. A version of the Aitken accelerator for computer iteration. *Int. J. Numer. Methods Eng.*, 1(3):275–277, 1969.

[180] U. Küttler and W.A. Wall. Fixed-point fluid-structure interaction solvers with dynamic relaxation. *Comput. Mech.*, 43(1):61–72, 2008.

[181] M. Luhar and H.M. Nepf. Flowinduced reconfiguration of buoyant and flexible aquatic vegetation. *Limnol. Oceanogr.*, 56(6):2003–2017, 2011.

[182] V. Heuveline, M.J. Krause, and J. Latt. Towards a hybrid parallelization of lattice Boltzmann methods. *Comput. Math. with Appl.*, 58(5):1071–1080, 2009.

[183] K.R. Tubbs and F.T.C. Tsai. Multilayer shallow water flow using lattice Boltzmann method with high performance computing. *Adv. Water Resour.*, 32(12):1767–1776, 2009.

[184] M. Basha and N.A.C. Sidik. Numerical predictions of laminar and turbulent forced convection: Lattice Boltzmann simulations using parallel libraries. *Int. J. Heat Mass Transf.*, 116:715–724, 2018.

[185] S.H. Kim and H. Pitsch. A generalized periodic boundary condition for lattice Boltzmann method simulation of a pressure driven flow in a periodic geometry. *Phys. Fluids*, 19(10):108101, 2007.

[186] J. Wang, J. Chen, and S. Liao. An explicit solution of the large deformation of a cantilever beam under point load at the free tip. *J. Comput. Appl. Math.*, 212(2):320–330, 2008.

[187] M. Schäfer, S. Turek, F. Durst, E. Krause, and R. Rannacher. Benchmark Computations of Laminar Flow Around a Cylinder. In E.H. Hirschel, editor, *Flow Simul. with High-Performance Comput. II DFG Prior. Res. Program. Results 1993-1995*, volume 48, pages 547–566. Vieweg+Teubner Verlag, 1996.

[188] C. Liao, Y. Chang, C. Lin, and J. McDonough. Simulating flows with moving rigid boundary using immersed-boundary method. *Comput. Fluids*, 39(1):152–167, 2010.

[189] J. Lee and S. Lee. Fluid-structure interaction analysis on a flexible plate normal to a free stream at low Reynolds numbers. *J. Fluids Struct.*, 29: 18–34, 2012.

[190] S. Turek and J. Hron. Proposal for Numerical Benchmarking of Fluid-Structure Interaction between an Elastic Object and Laminar Incompressible Flow. In H. Bungartz and M. Schäfer, editors, *Fluid-Structure Interact. Model. Simulation, Optim.*, pages 371–385. Springer Berlin Heidelberg, 2006.

[191] T. Steinmann, J. Casas, G. Krijnen, and O. Dangles. Air-flow sensitive hairs: boundary layers in oscillatory flows around arthropod appendages. *J. Exp. Biol.*, 209(21):4398–4408, 2006.

[192] B.T. Dickinson, J.R. Singler, and B.A. Batten. Mathematical modeling and simulation of biologically inspired hair receptor arrays in laminar unsteady flow separation. *J. Fluids Struct.*, 29:1–17, 2012.

[193] A.R. Pries, T.W. Secomb, and P. Gaehtgens. The endothelial surface layer. *Pflügers Arch. - Eur. J. Physiol.*, 440(5):653–666, 2000.

[194] S. Reitsma, D.W. Slaaf, H. Vink, M.A.M.J. van Zandvoort, and M.G.A. oude Egbrink. The endothelial glycocalyx: composition, functions, and visualization. *Pflügers Arch. - Eur. J. Physiol.*, 454(3):345–359, 2007.

[195] T.W. Secomb, R. Hsu, and A.R. Pries. Effect of the endothelial surface layer on transmission of fluid shear stress to endothelial cells. *Biorheology*, 38(2-3):143–150, 2001.

[196] J.M. Tarbell and Z. Shi. Effect of the glycocalyx layer on transmission of interstitial flow shear stress to embedded cells. *Biomech. Model. Mechanobiol.*, 12(1):111–121, 2013.

[197] G. Pontrelli, I. Halliday, T.J. Spencer, C.S. König, and M.W. Collins. Modelling the glycocalyx-endothelium-erythrocyte interaction in the microcirculation: a computational study. *Comput. Methods Biomech. Biomed. Engin.*, 18(4):351–361, 2015.

[198] P.E. Vincent, S.J. Sherwin, and P.D. Weinberg. The effect of the endothelial glycocalyx layer on concentration polarisation of low density lipoprotein in arteries. *J. Theor. Biol.*, 265(1):1–17, 2010.

[199] X. Liu, Y. Fan, and X. Deng. Effect of the endothelial glycocalyx layer on arterial LDL transport under normal and high pressure. *J. Theor. Biol.*, 283(1):71–81, 2011.

[200] B. Bathellier, F.G. Barth, J.T. Albert, and J.A.C. Humphrey. Viscosity-mediated motion coupling between pairs of trichobothria on the leg of the spider Cupiennius salei. *J. Comp. Physiol. A*, 191(8):733–746, 2005.

[201] B. Cummins, T. Gedeon, I. Klapper, and R. Cortez. Interaction between arthropod filiform hairs in a fluid enviroment. *J. Theor. Biol.*, 247(2): 266–280, 2007.

[202] J.J. Heys, T. Gedeon, B.C. Knott, and Y. Kim. Modeling arthropod filiform hair motion using the penalty immersed boundary method. *J. Biomech.*, 41 (5):977–984, 2008.

[203] G.C. Lewin and J. Hallam. A computational fluid dynamics model of viscous coupling of hairs. *J. Comp. Physiol. A*, 196(6):385–395, 2010.

[204] X. Yin and J. Zhang. Modeling the dynamic flow-fiber interaction for microscopic biofluid systems. *J. Biomech.*, 46(2):314–318, 2013.

[205] S. Chen and G.D. Doolen. Lattice Boltzmann Method for Fluid Flows. *Annu. Rev. Fluid Mech.*, 30:329–364, 1998.

[206] C.K. Aidun and R.C. Clausen. Lattice-Boltzmann Method for Complex Flows. *Annu. Rev. Fluid Mech.*, 42:439–472, 2010.

[207] J. Zhang. Lattice Boltzmann method for microfluidics: models and applications. *Microfluid. Nanofluidics*, 10(1):1–28, 2011.

[208] J. Branscomb and A. Alexeev. Designing ciliated surfaces that regulate deposition of solid particles. *Soft Matter*, 6(17):4066–4069, 2010.

[209] J. Boyd, J.M. Buick, and S. Green. Analysis of the Casson and Carreau-Yasuda non-Newtonian blood models in steady and oscillatory flows using the lattice Boltzmann method. *Phys. Fluids*, 19(9):093103, 2007.

[210] C. Loudon and A. Tordesillas. The Use of the Dimensionless Womersley Number to Characterize the Unsteady Nature of Internal Flow. *J. Theor. Biol.*, 191(1):63–78, 1998.

[211] O. Lida and Y. Nagano. The Relaminarization Mechanisms of Turbulent Channel Flow at Low Reynolds Numbers. *Flow, Turbul. Combust.*, 60(2): 193–213, 1998.

[212] J.B. Lee, S.G. Park, B. Kim, J. Ryu, and H.J. Sung. Heat transfer enhancement by flexible flags clamped vertically in a Poiseuille channel flow. *Int. J. Heat Mass Transf.*, 107:391–402, 2017.

[213] J.B. Lee, S.G. Park, and H.J. Sung. Heat transfer enhancement by asymmetrically clamped flexible flags in a channel flow. *Int. J. Heat Mass Transf.*, 116:1003–1015, 2018.

Appendix A

Discretisation of Euler-Bernoulli Beam Solver

A.1 Governing Equations

From Section 4.1, the discrete equations of motion are given by:

$$\Delta s^2 \left[D_s(T^{n+1} D_s x^{n+1}) - K_B D_{ssss} x^* + Fr_x + F_x^n \right] + \beta x^* - \beta x^{n+1} = 0 \tag{4.7}$$

$$\Delta s^2 \left[D_s(T^{n+1} D_s y^{n+1}) - K_B D_{ssss} y^* + Fr_y + F_y^n \right] + \beta y^* - \beta y^{n+1} = 0 \tag{4.8}$$

$$(\Delta s D_s x^{n+1})^2 + (\Delta s D_s y^{n+1})^2 - \Delta s^2 = 0 \tag{4.9}$$

Expanding the derivatives using standard finite difference formulas results in the

following system of equations:

$$T_i^{n+1}x_{i-1}^{n+1} - (T_i^{n+1} + T_{i+1}^{n+1} + \beta)x_i^{n+1} + T_{i+1}^{n+1}x_{i+1}^{n+1} - \frac{K_B}{\Delta s^2}(x_{i-2}^* - 4x_{i-1}^* + 6x_i^* - 4x_{i+1}^* + x_{i+2}^*) + \Delta s^2(Fr_x + F_{x,i}) + \beta x_i^* = 0 \tag{A.1}$$

$$T_i^{n+1}y_{i-1}^{n+1} - (T_i^{n+1} + T_{i+1}^{n+1} + \beta)y_i^{n+1} + T_{i+1}^{n+1}y_{i+1}^{n+1} - \frac{K_B}{\Delta s^2}(y_{i-2}^* - 4y_{i-1}^* + 6y_i^* - 4y_{i+1}^* + y_{i+2}^*) + \Delta s^2(Fr_y + F_{y,i}) + \beta y_i^* = 0 \tag{A.2}$$

$$(x_i^{n+1} - x_{i-1}^{n+1})^2 + (y_i^{n+1} - y_{i-1}^{n+1})^2 - \Delta s^2 = 0 \tag{A.3}$$

for $i = 0 \dots N$.

A.2 Boundary Conditions

Free End

The boundary conditions at the free end ($i = N$) are given by:

$$T = 0, \qquad \frac{\partial^2 \mathbf{X}}{\partial s^2} = (0,0), \qquad \frac{\partial^3 \mathbf{X}}{\partial s^3} = (0,0) \tag{4.10}$$

which, after discretising, give:

$$\begin{cases} x_{N-1} - 2x_N + x_{N+1} = 0 \quad \Rightarrow \quad x_{N+1} = 2x_N - x_{N-1} \\ y_{N-1} - 2y_N + y_{N+1} = 0 \quad \Rightarrow \quad y_{N+1} = 2y_N - y_{N-1} \\ -\frac{1}{2}x_{N-2} + x_{N-1} - x_{N+1} + \frac{1}{2}x_{N+2} = 0 \quad \Rightarrow \quad x_{N+2} = 4x_N - 4x_{N-1} + x_{N-2} \\ -\frac{1}{2}y_{N-2} + y_{N-1} - y_{N+1} + \frac{1}{2}y_{N+2} = 0 \quad \Rightarrow \quad y_{N+2} = 4y_N - 4y_{N-1} + y_{N-2} \\ T_{s=1} = 0 \end{cases} \tag{A.4}$$

Fixed End

For a simply supported beam, the conditions at the fixed end ($i = 0$) are:

$$\mathbf{X} = (X_0, Y_0), \qquad \frac{\partial^2 \mathbf{X}}{\partial s^2} = (0, 0) \tag{4.11}$$

where X_0 and Y_0 are the spatial coordinates of the base. In discretised form, Equation 4.11 gives:

$$\begin{cases} x_0 = X_0 \\ y_0 = Y_0 \\ x_{-1} - 2x_0 + x_1 = 0 \quad \Rightarrow \quad x_{-1} = 2X_0 - x_1 \\ y_{-1} - 2y_0 + y_1 = 0 \quad \Rightarrow \quad y_{-1} = 2Y_0 - y_1 \end{cases} \tag{A.5}$$

In the case of a clamped beam, the conditions at the fixed end ($i = 0$) are:

$$\mathbf{X} = (X_0, Y_0), \qquad \frac{\partial \mathbf{X}}{\partial s} = [\cos(\theta), \sin(\theta)] \tag{A.6}$$

where θ is the clamp angle, with respect to the horizontal axis. In discretised form, this gives:

$$\begin{cases} x_0 = X_0 \\ \\ y_0 = Y_0 \\ \\ -\frac{1}{2}x_{-1} + \frac{1}{2}x_1 = \cos(\theta) \quad \Rightarrow \quad x_{-1} = X_0 - 2\cos(\theta) \\ \\ -\frac{1}{2}y_{-1} + \frac{1}{2}y_1 = \sin(\theta) \quad \Rightarrow \quad y_{-1} = Y_0 - 2\sin(\theta) \end{cases} \tag{A.7}$$

A.3 Full System of Equations

After expanding and rearranging Equations A.1–A.3, the $3N$ variables, $\{T_i^{n+1},\ i \in [1, N]\}$, $\{x_i^{n+1},\ i \in [1, N]\}$, $\{y_i^{n+1},\ i \in [1, N]\}$, are a solution of the $3N$ equations:

$$\begin{cases} F_1^{n+1} = A_1^{n+1} - B_1^{n+1} = 0 \\ \\ F_2^{n+1} = A_2^{n+1} - B_2^{n+1} = 0 \\ \\ F_3^{n+1} = A_3^{n+1} - B_4^{n+1} = 0 \\ \\ \ldots \\ \\ F_{3N}^{n+1} = A_{3N}^{n+1} - B_{3N}^{n+1} = 0 \end{cases} \tag{A.8}$$

where, for $i = 1 \dots N$:

$$\begin{cases} A_{3i-2}^{n+1} = (x_i^{n+1} - x_{i-1}^{n+1})^2 + (y_i^{n+1} - y_{i-1}^{n+1})^2 \\ B_{3i-2}^{n+1} = \Delta s^2 \\ A_{3i-1}^{n+1} = T_i^{n+1} x_{i-1}^{n+1} - (T_i^{n+1} + T_{i+1}^{n+1} + \beta) x_i^{n+1} + T_{i+1}^{n+1} x_{i+1}^{n+1} \\ B_{3i-1}^{n+1} = \frac{K_B}{\Delta s^2}(x_{i-2}^* - 4x_{i-1}^* + 6x_i^* - 4x_{i+1}^* + x_{i+2}^*) - \Delta s^2 (Fr_x + F_{x,i}) - \beta x_i^* \\ A_{3i}^{n+1} = T_i^{n+1} y_{i-1}^{n+1} - (T_i^{n+1} + T_{i+1}^{n+1} + \beta) y_i^{n+1} + T_{i+1}^{n+1} y_{i+1}^{n+1} \\ B_{3i}^{n+1} = \frac{K_B}{\Delta s^2}(y_{i-2}^* - 4y_{i-1}^* + 6y_i^* - 4y_{i+1}^* + y_{i+2}^*) - \Delta s^2 (Fr_y + F_{y,i}) - \beta y_i^* \end{cases} \tag{A.9}$$

The boundary conditions from Equations A.4, A.5, and A.7 are substituted into Equation A.9 as appropriate. The nonlinear system is then solved iteratively via Newton's multidimensional method to get the solution at the next time step.