Department of Mechanical Engineering

Novel Computational Methods for the Study of Compliant-Wall Fluid-Structure Interaction

Jarrad Stephen Kapor

This thesis is presented for the Degree of
Doctor of Philosophy
of
Curtin University of Technology

September 2012

Declaration

To the best of my knowledge and belief this thesis contains no material previously published by any other person except where due acknowledgement has been made. This thesis contains no material which has been accepted for the award of any other degree or diploma in any university. All investigations presented in this thesis are my own except where specific reference has been made to the work of others. Some of the work presented in this thesis has been published in the following:

J. S. Kapor, A. D. Lucey, and M. W. Pitman. Boundary-Layer Hydrodynamics using Mesh-Free Modelling. In *Proceedings 18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation*, July 2009

J. S. Kapor, A. D. Lucey, and M. W. Pitman. Fluid-structure interaction using mesh-free modelling. In *Proceedings 6th Australasian Congress on Applied Mechanics*, Perth, Australia, December 2010

J. S. Kapor, A. D. Lucey, and M. W. Pitman. A mesh-free compliant-wall fluid-structure interaction model. In *Proceedings MODSIM2011, 19th International Congress on Modelling and Simulation*, December 2011

J. S. Kapor and A. D. Lucey. Flow-induced vibrations of a flexible panel in a boundary-layer flow. In *Flow-Induced Vibration, 10th International Conference on Flow-Induced Vibration (& Flow-Induced Noise) FIV2012*, pages 647–654, Trinity College Dublin, Ireland, July 2012

Signature: _____

Date: _____

Acknowledgements

While there are many people that I could thank, the following have played a significant role throughout the course of my PhD:

My supervisor, Professor Anthony (Tony) Lucey from Curtin University of Technology, for sparking my interest in Fluid Mechanics at an undergraduate level and providing five years of valuable guidance and advice throughout my PhD.

Dr Mark Pitman, for being a good friend and getting me started on the pathway of DVM modelling and for trying to warn me of the perils and many sleepless nights that would lay ahead as a result.

Dr Andrew King, Dr James Jewkes, Dr Novak Elliott and Dr Richard Howell for being good friends and colleagues, particularly those who shared an intensely cluttered and crowded office.

Administrative staff, Sucy Leong, Kim Yap and Margaret Brown who have assisted in managing my teaching positions, conference trips, supervisor meetings and the copious amount of university red tape.

My family and in-laws for supporting my wife and I throughout the hard times, particularly through the period in which we were a single-income family. Our lives would have been very unpleasant without their continued support.

Most importantly I would like to thank my wife, Zaneta, who has stuck by me throughout all the highs and lows of the last ten years. Her many sacrifices have provided me with the continued emotional and financial stability needed to push through to the end. I have no doubt that, without her by my side, I would never have completed this work.

Abstract

A numerical tool for studying the 2D Fluid-Structure Interactions (FSI) of a finite-length compliant wall in a boundary-layer flow is developed. Understanding the complex behaviour that results from simultaneously coupling two wave-bearing media is of great importance to Engineering in areas such as aero/hydro-elasticity, biomechanical systems, energy harvesting and, of particular interest herein, drag-reducing technologies. Applying compliant-wall coatings to otherwise rigid surfaces can delay the onset of laminar-turbulent transition, offering marked reductions in skin-friction drag and therefore energy loss, particularly in marine applications.

A novel Discrete-Vortex Method is used to accommodate readily the nonlinearity in the Navier-Stokes equations and represent a generic Computational Fluid Dynamics tool. Efficiency is maintained through the use of a Fast-Multipole Method while viscosity is modelled using a vortex core-spreading method and a Generalised Minimum Residual Method rediscretisation scheme. This is coupled to an arbitrarily deforming surface using the Boundary-Element Method for the outer potential flow while the structural mechanics are modelled with a Kramer-type one-dimensional nonlinear Euler-Bernoulli beam model solved numerically using the Finite-Difference Method. To obtain the system's behaviour, flow and structural dynamics are coupled semi-implicitly and solved using the Newton-Krylov method, thereby complementing the matrix-free nature of the flow-solution methodology.

Using a standard desktop computer, the investigations presented address harmonic oscillations of a compliant wall, inviscid flow over both rigid bodies and compliant walls, rotational flows, and full boundary-layer flow over rigid and compliant walls, demonstrating the three fundamental FSI instabilities of Tollmien-Schlichting Waves, Travelling-Wave Flutter (TWF) and divergence. The model allows the study of the effects of the boundary-layer on the divergence-onset speed, nonlinear-divergence oscillations and TWF in a finite-wall system.

The divergence-onset flow speed is shown to be dependent upon the rela-

tive boundary-layer thickness where the highest speeds are observed for thick boundary-layers, while for thin boundary-layers the critical speed tends towards that of the potential-flow case. At post-divergence flow-speeds, the nonlinear oscillatory behaviour of the boundary-layer FSI system is seen to remained largely unchanged from that predicted by potential-flow studies. The main effect of the boundary-layer flow is to both reduce the total energy of the oscillations and to provide a dissipative effect similar to weak structural damping. However, at pre-divergence flow speeds, the system can establish temporally-unstable downstream standing waves adjacent to the fixed trailing end of the wall that are fuelled by TWF and its self-excitation at the upstream wall end.

Contents

List of Figures

Chapter 1

Introduction

Fluid-structure interaction (FSI) has been an increased focus of research in recent years because it concerns the multi-physics dynamics of an immersed structure interacting with its surrounding fluid. One of the simplest forms of FSI is that of a rigid body immersed in a fluid, either in motion or at rest, modifying the fluid-flow dynamics. A classic engineering example of FSI interest is the aero-/hydro-dynamics of transportation and wind-loaded structures where the specific geometry of the immersed body/structure influences the magnitude of the drag forces and wind loads. However, further complications arise when looking at the FSI of flexible structures where the combined system dynamics are governed by two wave-bearing media. Thus, the fluid-flow behaviour is influenced by the dynamic deformation of the structure and the structure's geometry is dependent upon the dynamic loads of the fluid.

Understanding this type of complex FSI has become of interest to the Engineering community with various biomedical applications (Heil and Hazel, 2011) such as snoring and sleep apnoea (e.g. Elliott et al., 2011), the blood-flow through the flexible-walled arteries/veins of the human cardiovascular system (e.g. Davies and Carpenter, 1997b; Gerbeau et al., 2005; Grotberg and Jensen, 2004; Heil and Jensen, 2003; Pedley, 1992) and the study of diseases such as Syringomyelia (e.g. Berkouk et al., 2003; Bertram et al., 2005; Carpenter et al., 2003; Elliott, 2012). Classical engineering interest is in studying flexible body/wall FSI to understand

failure mechanisms in engineering systems. This includes studying the resonant effects of wind-/ocean-loaded flexible structures, with a well-known example being the Tacoma Narrows Bridge seen in Figure 1.1. These failures occur when

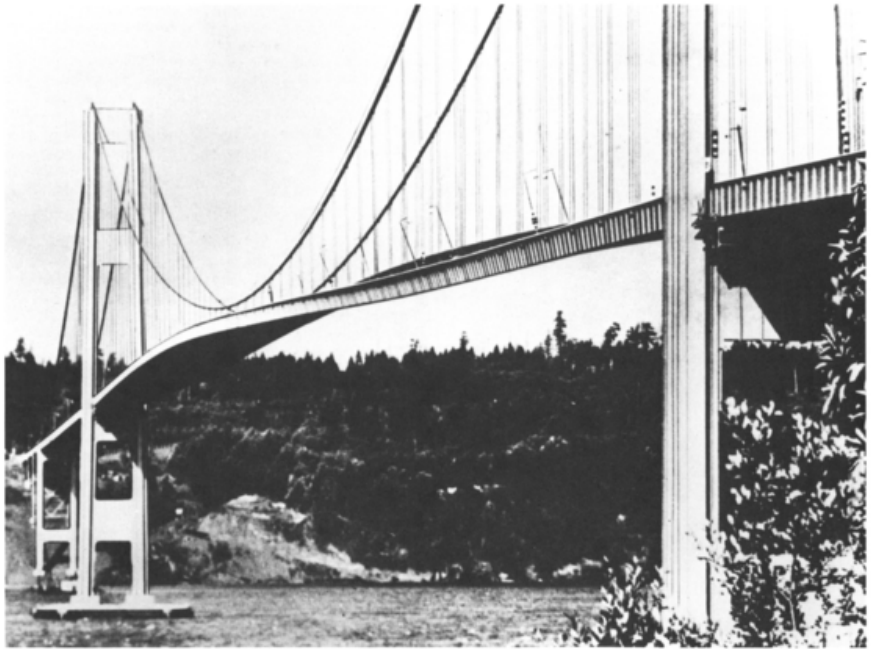

Figure 1.1: Photograph reproduced from Gunns (1981) showing the flow-induced vibrations of the Tacoma Narrows Bridge: "The writhing motion of the bridge deck was caused in part by wind action on the solid sideplates. (James Bashford, Photography Coll., University of Washington Libraries)" (Gunns, 1981).

the natural vibration modes of the coupled systems are excited, leading to irreversible and potentially dangerous energy transfer from the fluid into the structure. Without sufficient damping to dissipate this transferred energy, the system can undergo catastrophic failure. Under similar mechanisms for energy transfer there is also an emerging interest in technologies for ('green') energy harvesting (Howell and Lucey, 2012; Tang et al., 2009). However, with the current political focus on improving energy efficiency, there is a renewed technological interest in the application of compliant coatings to otherwise rigid moving objects to exploit their potential for drag-reduction. This technology was inspired by observations of dolphins (e.g. Kramer, 1957, see Figure 1.2), and operates by using the compliant wall as a passive boundary-layer control device, where the boundary-layer

is the interface of viscous fluid between the principally inviscid outer mean flow and the outer skin of an object/vehicle. A compliant wall with tuned properties can interfere with the mechanisms that cause laminar-to-turbulent transition of the boundary-layer and encourage transition postponement (e.g. Gaster, 1987). Through this, a marked reduction in the skin-friction drag can be achieved, reducing a significant component of parasitic energy loss in transportation, particularly in marine applications.

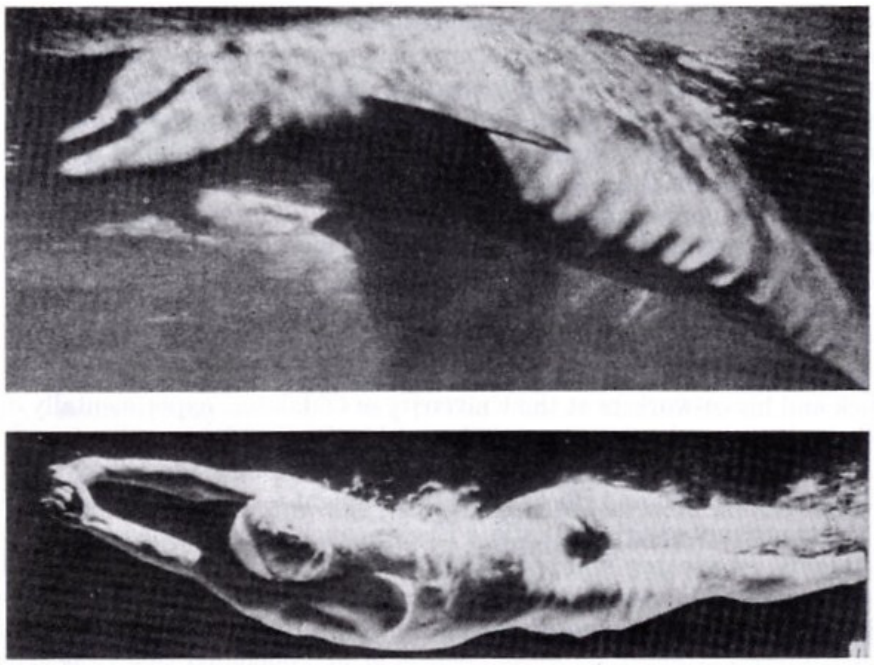

Figure 1.2: Photographs reproduced from Aleyev (1977) showing compliance in the skin of a dolphin and a human during high-speed swimming.

1.1 Objectives

In this work we develop a novel numerical tool for studying the FSI of a boundary-layer flow over a finite-length compliant wall deforming at nonlinear amplitudes. When compared to experimental methods, the numerical approach has been chosen because it overcomes the inherent time and monetary costs of conducting physical (often 'trial and error') studies. While analytical methods require the least time and monetary cost to elucidate underlying system dynamics, there

is no unifying analytical solution to the finite-wall, nonlinear boundary-layer FSI problem that can adequately deal with all fundamental (drag-reduction relevant) system instabilities. Therefore, the objective of this work is to develop a numerical tool that can effectively combine the capabilities of current analytical approaches to allow the evaluation, development and optimisation of drag-reducing compliant-wall technologies. However, drag-reducing dynamics and true boundary-layer FSI instabilities only occur for a small subset of system properties. To maintain engineering relevance the model must also capture dynamics encountered outside of the drag-reducing realm to yield insight into simpler FSI systems. Also to have potential as an engineering design tool, it must remain scalable and efficient to allow the computation on a standard (e.g. desktop) computer.

We identify the elementary sub-objectives of the present research that each have engineering interest and are consequently demonstrated in their own right as;

1. Compliant-wall vibrations *in vacuo*;
2. High Reynolds-number (inviscid) flows over an arbitrarily shaped rigid wall;
3. FSI of high Reynolds-number (inviscid) flows over a compliant wall;
4. High Reynolds-number flows that exhibit viscous/rotational effects;
5. Viscous (transitional) boundary-layer flows over a rigid wall; and
6. FSI of viscous (transitional) boundary-layer flows over a compliant surface.

An added benefit of developing individually working sub-models is that a greater wealth of pre-existing data/models can be used to validate the present numerical model. With the model suitably validated investigations are performed to gain new knowledge that is currently unobtainable when using existing analytical methods.

1.2 Methodology

We present methods and models for boundary-layer flows that are two-dimensional (2D) in nature. This is sufficient to capture the transition mechanisms that begin as 2D Tollmien-Schlichting waves and fundamental 2D wall instabilities such as divergence (observed in Figure 1.3). It is noted that the numerical methods

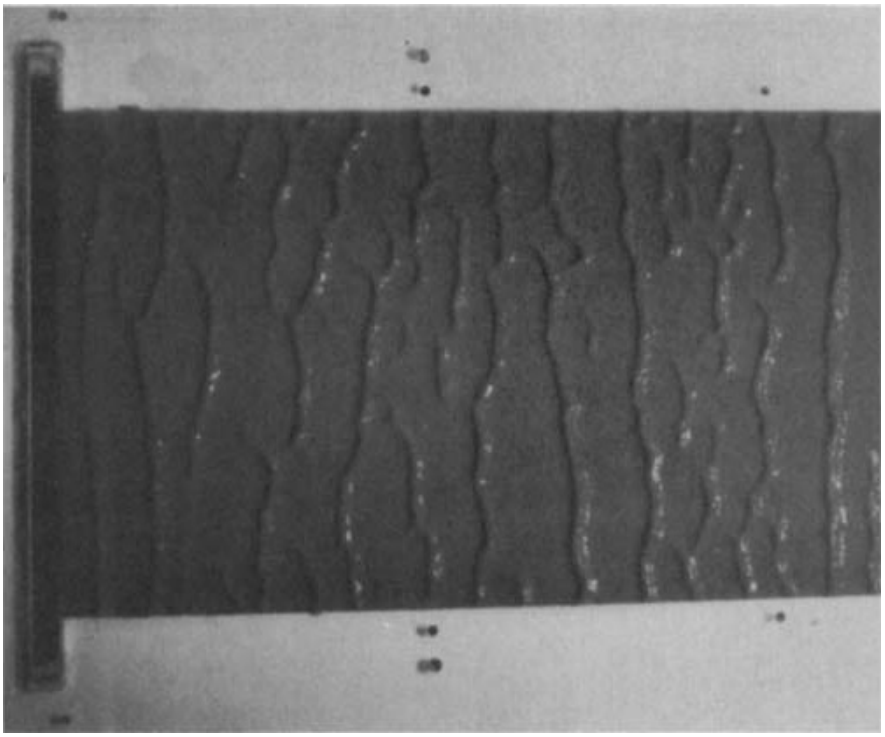

Figure 1.3: Figure 5 reproduced from Gad-el-Hak et al. (1984) showing experimental observation of (2D) static-divergence waves on a compliant wall under a turbulent boundary layer. The fluid flow is from the left to right.

herein are chosen for their scalability and extendibility to more complicated 3D systems for future works.

We develop a replacement to the frequently used infinite and flat-wall based Orr-Sommerfeld equation, by developing a novel Discrete-Vortex Method (DVM) boundary-layer flow model that can be viewed as a generic Computational Fluid Dynamics (CFD) tool. The DVM is a Lagrangian solution to the vorticity formulation of the governing Navier-Stokes equations and offers an alternative approach to Eulerian-grid based (e.g. finite-difference based DNS) methods. A Lagrangian

(and grid-free) approach is advantageous for the strong coupling of the nonlinear structural dynamics that can occur in drag-reducing compliant-wall systems. The DVM boundary-layer model is used as a viscous interface between the structural dynamics and the outer-region of potential flow as seen in Figure 1.4. Emphasis

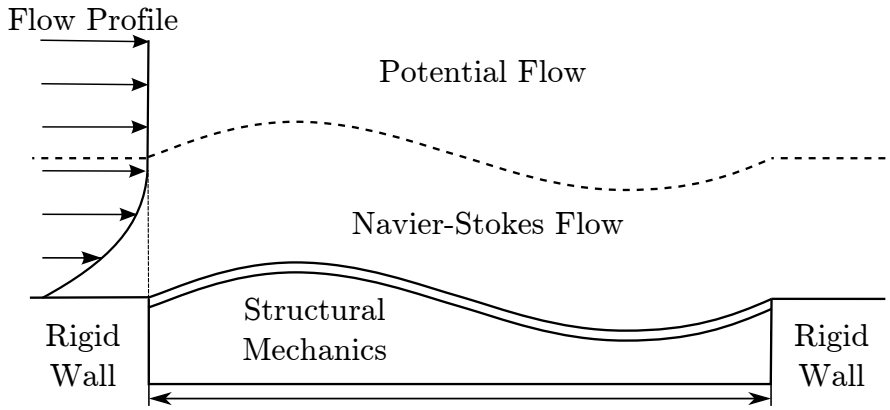

Figure 1.4: Simplified schematic of the boundary-layer flow compliant-wall FSI system.

is placed on developing efficient methods that overcome the severe difficulties in obtaining a stable numerical coupling between the nonlinear structural dynamics and the nonlinear boundary-layer flow. In this work, structural dynamics are limited to a nonlinear extension of the canonical 'Kramer-type' compliant-wall (Carpenter and Garrad, 1985, 1986). While this is based on the one-dimensional (1D) Euler-Bernoulli beam model, it is solved using methods that can be easily extended to different and/or more complex wall models.

1.3 Thesis Layout

The thesis is structured into chapters that broadly represent the increasing level of complexity for the system dynamics and numerical methods so as to reach the fully-coupled nonlinear boundary-layer FSI model. Chapter 1 provides an introduction to the field of compliant-wall FSI and a brief literature review of the major types of instabilities that are relevant to drag-reducing technologies and the established methods of past works. We also discuss the literature upon

which we draw techniques for DVM boundary-layer modelling. Chapter 2 introduces the modelling of potential flow over rigid surfaces of arbitrary shape. This includes the use of the Boundary Element Method (BEM) along with supporting numerical methods such as the Fast Multipole Method (FMM), the Generalised Minimum Residual Method (GMRES) and associated pre-conditioning techniques. Chapter 3 introduces the structural dynamics of a compliant wall based on a nonlinear 1D Euler-Bernoulli beam model. The dynamic response of the system is solved by an implicit time-stepping scheme using the Newton-Krylov (NK) and Finite Difference Method (FDM). Chapter 4 presents the FSI of an inviscid flow coupled to a compliant wall and investigates different coupling methods, their scaling efficiency and the nonlinear behaviour of the system at post-divergence flow speeds. Chapter 5 introduces the Discrete Vortex Method (DVM) and FMM for the modelling of lightly viscous flows where rotational flow dynamics can no longer be neglected. Chapter 6 presents the modelling of a viscous boundary-layer flow over a flat rigid wall. The effects of viscous diffusion are introduced through a vortex core-spreading method that requires a radial-basis rediscretisation scheme (based on the GMRES) to ensure numerical accuracy for long-time simulations. To validate the model flow-based instabilities in the form of Tollmien-Schlichting waves are investigated. Chapter 7 presents the full nonlinear boundary-layer compliant-wall FSI model. An efficient and robust method for coupling the nonlinear flow to the structural-dynamics is presented using a y-momentum integration method to obtain the pressure at the wall-fluid interface. The effects of a boundary-layer on the system's divergence-onset flow speed and nonlinear divergence oscillations are investigated. The mechanisms behind travelling-wave flutter and its behaviour in a finite-wall system are also investigated. Chapter 8 provides a summary on the achievements and observations of the entire work, including comments on the future potential and shortcomings of the developed model.

1.4 Background Literature

Due to the availability of in-depth review articles on the last 50 years of its history, progress throughout the decades and the present state of the art, we only provide an overview of the literature that inspired the present research. For a greater wealth of information, the reader is directed to the works of Bushnell et al. (1977); Carpenter (1990, 1998); Carpenter et al. (2000); Dowell (1975); Gad-el-Hak (1986, 1987, 1996, 1998, 2002); Matsuzaki (1986) and Riley et al. (1988).

In 1957, it was observations of the dolphin's paradoxically high swimming speeds that led to the pioneering experiments of Kramer (1957) into compliant-wall technology. By varying properties of a compliant coating that replicated the dolphin's epidermis, Kramer measured up to 60% reductions in drag forces on a body. This result inspired many to study Kramer's findings such as the experimental works by Nisewanger (1964); Puryear (1962); Ritter and Messum (1964) and Ritter and Porteous (1965). Most works disputed Kramer's compliant-wall finding by concluding that the transition from laminar to turbulent flow occurred much earlier than for the case of a rigid wall and therefore caused an increase in skin-friction drag. Meanwhile, analytical works such as Benjamin (1960, 1963); Betchov (1959); Kaplan (1964) and Landahl (1962) concluded that compliant-wall drag-reduction was possible (through transition delay) but the wall properties required would be extremely difficult to achieve in practice. As a consequence the focus changed from compliant coatings as a transition-delay strategy, to a drag-reduction mechanism for turbulent flows shown by Smith and Blick (1969) and Fisher et al. (1975).

During the early 1980's there was an injection of funding into compliant-wall research by the Office of Naval Research in the United States and the Procurement Executive of the Ministry of Defence in Great Britain (Gad-el-Hak, 2002). This renewed interest sparked the experimental work of Gaster (1987) which revealed an excellent agreement with analytical works based on linear hydrodynamic theory, such as Carpenter and Garrad (1985, 1986); Lucey and Carpenter

(1995) and Willis (1986). These linear methods became the accepted standard of the time and were based on reducing the Navier-Stokes equations to the Orr-Sommerfeld equation by imposing a continuous and linearised disturbance to a mean flow profile. An important outcome from these linear works (and various others such as Dixon et al., 1994; Yeo, 1988) was that compliant coatings demonstrated transition-delaying potential based on known physical principles with correctly-tuned wall parameters at specific Reynolds numbers. Carpenter and Garrad (1985) also demonstrated that Kramer's original compliant-wall coatings would offer marginal delays in boundary-layer transition and offered explanations as to the cause of many of the early experimental failures. In a much later work, Carpenter (1998) concludes that indefinite transition delay can be achieved (at-least theoretically) for optimised multiple panel arrangements, with relatively short panels showing the most promise to resist transition-promoting instabilities (see Carpenter et al., 2001; Carpenter, 1988, 1990, 1993; Davies and Carpenter, 1997a,b; Lucey and Carpenter, 1995). However, it is worth noting that transition-delay by compliant-wall technologies was established to be practically unobtainable[1] for aerodynamic applications due to the large density ratio of the structure to fluid (see Carpenter et al., 2001; Carpenter, 1990).

1.4.1 Compliant-Wall System Instabilities

With linear stability theories reaching an accepted level of maturity, a consensus was made on the three fundamental compliant-wall FSI instabilities. As shown in Figure 1.5, for a boundary-layer flow over a Kramer-type compliant wall, the fundamental instabilities are those of Tollmien-Schlichting Instabilities (TSI), static divergence and Travelling-Wave Flutter (TWF). We neglect the fourth identified instability as being fundamental because it is based on the coalescence of TSI and TWF.

The flow-based TSI is the main mechanism for instigating boundary-layer transition from laminar to turbulent flow in low disturbance environments (Car-

[1]However, other turbulent flow mechanisms may offer compliant-wall drag-reduction mechanisms (e.g. Choi et al., 1997; Lee et al., 1995).

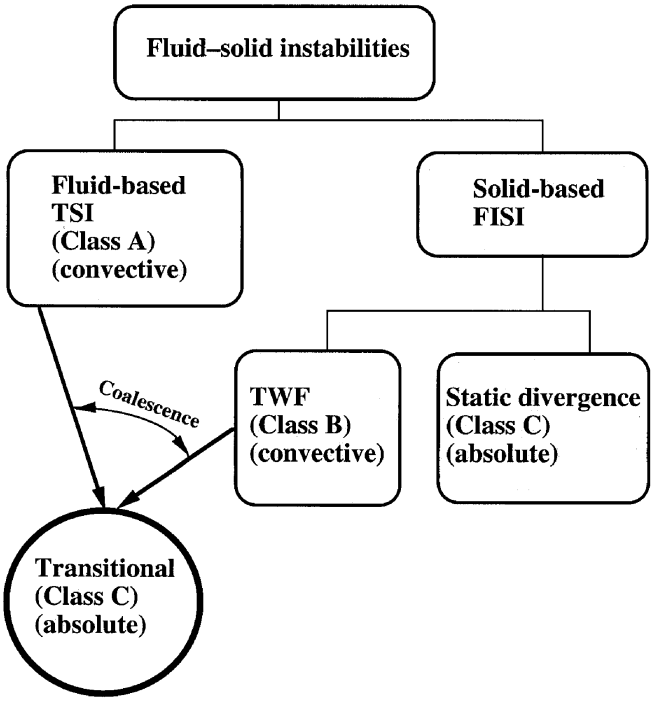

Figure 1.5: Figure 4 reproduced from Gad-el-Hak (2002) showing a summary of historical compliant-wall instability classification schemes.

penter and Garrad, 1985; Schlichting, 1979). In flows over a rigid flat surface, they manifest themselves through stream-wise propagating waves that begin as a small disturbance in the laminar regime (e.g. from a surface irregularity or a localised disturbance in the mean flow) and grow as a convective instability due to the action of viscosity. These waves are known to be stabilised by compliance in the surface and destabilised by factors such as structural-damping and adverse pressure gradients in the flow.

TWF is a flow-induced structural instability (FISI) that involves surface waves that are convectively unstable and grow as they travel downstream (Benjamin, 1963; Carpenter and Garrad, 1986). They can occur at pre-divergence flow speeds and are caused by the irreversible energy transfer that occurs due to shifts in the phase relationship of the boundary-layer flow pressure and wall velocity. As a mode of instability they are less destructive than divergence because they can be controlled by structural damping and the convective nature of the waves ensures that they are not self-exciting in (infinitely) long walls.

Divergence is a FISI that is characterised by quasi-static waves that grow when the dynamic fluid force exceeds the restorative forces in the wall (as seen in Figure 1.3, Gad-el-Hak et al., 1984). As an absolute instability they are the most destructive because they are self-exciting and their critical flow speed is unaffected by structural damping[2].

1.4.2 Methodology

In response to claims that indefinite transition postponement is obtainable for finite-length compliant-wall systems, we develop a model that can (partially) bridge the gap between theory and practical technologies by alleviating (some of) the restricting assumption in the linear methods. The previous agreement of theoretical and experimental works (see Lucey and Carpenter, 1995) was based mainly on the observations of TSI and TWF. The disadvantage of these methods was they lacked the ability to model walls of finite length and extract important information such as the nonlinear effects of a viscous boundary-layer. Carpenter (1998) also notes that with respect to predicting divergence instabilities, the common methods based on potential-flow theory (e.g. Carpenter and Garrad, 1986; Duncan et al., 1985; Garrad and Carpenter, 1982) are inadequate and offer a conservative prediction of the divergence-onset speed for laminar flows. This excludes many suitable compliant-wall material configurations when using optimisation methods to achieve transition delay (such as Dixon et al., 1994).

Lucey and Carpenter (1992) pioneered a numerical investigation into divergence instabilities by coupling a Boundary-Element Method (BEM) and a Finite-Difference Method (FDM) to solve for the FSI of a finite-length compliant-wall immersed in a uniform-flow. Pitman and Lucey (2009) achieved success in subjecting the linear BEM system to a state-space eigen-analysis approach while Lucey et al. (1997b) investigated the hydro-elastic response of compliant panels and walls to a uniform flow at nonlinear amplitudes. To model the effects of a viscous shear-flow, Davies and Carpenter (1997a,b) used a linearised velocity-

[2]This is because they are static waves exactly at the critical flow speed.

vorticity method to successfully model the transition-delaying potential of a finite-length compliant-wall immersed in channel flow. Meanwhile, to model viscous effects that were lacking in Lucey et al. (1997b), Cafolla (1997) and Lucey et al. (1997a) developed an extension to include the effects of a boundary-layer flow. This was based on modelling inviscid perturbations to a boundary-layer using the Lagrangian Discrete Vortex Method (DVM); however, this was limited to linear disturbances only. The boundary-layer was shown to reduce the magnitude of the pressure for flows with low-velocity (and therefore shear) gradients near the wall; a result further substantiated by Yeo et al. (1996, 1999). Pitman (2007) and Pitman and Lucey (2004) extended this DVM-based model beyond inviscid flow and linear perturbations by using a deterministic vortex method that modelled viscosity and allowed nonlinear wall amplitudes.

It is through the combined works of Cafolla (1997); Lucey et al. (1997a,b); Pitman (2007) and Pitman and Lucey (2004) that the present approach is founded. The grid-free BEM and the FDM are used to model the outer potential-flow and compliant-wall dynamics respectively while the DVM is used to model the dynamics of the inner boundary-layer flow (laminar or turbulent).

The Discrete Vortex Method

As the development of the DVM model is a significant component of the present work, we provide a brief overview of DVM-based turbulent mixing-layer and full boundary-layer modelling.

The DVM is a subset of broader vortex methods that are based upon the velocity-vorticity formulation of the Navier-Stokes equations. It is a numerical modelling scheme that is based on using the interactions of discrete vortex particles to approximate solutions to the two-dimensional (2D) Euler equations in a Lagrangian framework. The main advantage of the DVM is its grid-free nature and efficiency when compared to Eulerian grid-based methods. This is because the Lagrangian nature of the scheme allows the particle 'grid' to continuously deform and adapt to areas of interest in the flow such as those with high vortic-

ity/velocity gradients. These adaptive characteristics are particularly important when modelling fluid-structure interactions which may involve large amplitude boundary (wall) deformations.

Since its inception in 1931, much progress has been made using the DVM with excellent review articles being published by Aref (1983); Clements and Maull (1975); Leonard (1980, 1985); Perlman (1985); Saffman and Baker (1979) and Sarpkaya (1989). What is evidenced by these reviews is that it lends itself well to modelling applications of incompressible fluid flows such as aerodynamics (Chorin and Bernard, 1972; Leonard and Spalart, 1981), turbulent mixing (Aref and Siggia, 1980; Ashurst, 1977; Inoue, 1985; Inoue and Leonard, 1987; Leonard, 1980; Wang et al., 1999), combustion (Ashurst, 1981) and of particular interest in this study, for boundary-layers (Cafolla, 1997; Chorin, 1973, 1978, 1980; Lucey et al., 1997a; Pitman, 2007; Pitman and Lucey, 2004).

The foundations of the DVM were laid in 1858 by Helmholtz who was the first to show that in an inviscid fluid, vortices advect along streamlines of the flow and that flows with vorticity could be modelled with an approximate circulation and infinitely small cross section. However, the first serious attempt at vortex modelling is often cited as being the work by Rosenhead (1931) who studied the Kelvin-Helmholtz instability of vortex sheets. Rosenhead used a distribution of finite elemental vortices along the length of a vortex sheet and allowed the movements of the vortices to be determined by numerical time stepping. It was not long after that Westwater (1935) also reported the same success in showing smooth vortex sheet roll-up when modelling vortex shedding on a wing. Approximately thirty years later these works were re-examined by Birkoff and Fisher (1959) with the use of many more elemental vortices and smaller time steps. These results revealed an impossible motion for the vortex sheet with discrete vortices rolling up in an irregular fashion and sometimes crossing each others' pathways. They consequently concluded that while the main characteristics of Rosenhead (1931) were correct, the smooth roll-up was due to Rosenhead's large time steps. Similarly Takami (1964) and Moore (1971) reinvestigated Westwater's work and drew

similar conclusions; the paths taken by the vortices were irregular and unrealistic, while Moore also noted that increasing the number of discrete vortices would worsen the situation. However this discrepancy was finally resolved by Chorin and Bernard (1972) who revealed that a discrete approximation to the vortex sheet can only be accurate if the vorticity gradient along the sheet is small compared to the inter-vortex spacing.

One of the first successful attempts at using the DVM for boundary-layer modelling was that of Chorin (1973) who successfully modelled boundary-layer growth over a flat plate using the Random-Walk Method (RWM) to stochastically model viscous diffusion. The later works of Cafolla (1997); Lucey et al. (1997a); Pitman (2007) and Pitman and Lucey (2004) differed in approach by using the DVM to model perturbations to an established boundary-layer by stacking layers of shear[3] upon one-another, allowing the discretisation of the velocity profile of any boundary-layer flow. However, a key difference was that the effects of viscous diffusion were neglected in Cafolla (1997) while Pitman (2007) used the Corrected-Core Spreading Vortex Method (CCSVM) (Rossi, 1996). This method was based on modelling viscous diffusion by spreading Gaussian-cored elements at a rate that solves the viscous component of the vorticity transport equations (Leonard, 1980). Its deterministic nature offered an advantage over the RWM by not requiring the use of random motions[4].

Initial core-spreading methods were criticised by Greengard (1985) due to incorrect modelling of convection when an element's core-size grew too (relatively) large. Rossi (1996) mitigated this error using the CCSVM by forcing the Gaussian vortex elements to split regularly in a conservative manner. While this scheme is attractive from a simplicity point of view, its major weakness is that it causes the number of particles to grow exponentially. It therefore requires the use of particle merging schemes (such as Rossi, 1997) to avoid exacerbating the already high cost of the DVM's 'N-Body problem' (Takeda et al., 1997). While Huang (2005) later

[3]Where a single layer is similar to the turbulent mixing layer models that began with Rosenhead (1931).

[4]For a thorough comparison and review of the various schemes for modelling viscous diffusion in the DVM see Takeda et al. (1997)

revised the CCSVM to include vortex splitting (and merging) schemes to offer a greater level of accuracy, the localised nature of the split/merge event causes difficulties in enforcing global discretisation standards (such as particle overlap and resolution).

Barba et al. (2003, 2005) notes that when the vortex 'grid' becomes distorted due to simulation events such as splitting/merging or particle convection within the CCSVM, it can lead to a loss of initial accuracy of up to five orders of magnitude. This can lead to invalid chaotic behaviour (such as the observations in Birkoff and Fisher, 1959; Moore, 1971; Takami, 1964) that causes the accuracy of any long-time simulation to become questionable. The solution used herein was designed by Barba and Rossi (2010); Barba et al. (2003, 2005) and Yokota et al. (2010) to replace the local split/merge events of the CCSVM with a global Lagrangian-grid rediscretisation scheme. This is based on mapping the perturbed vorticity field onto a separate Lagrangian grid using radial-basis function interpolation. The main strength in the method is that it uses the Generalised Minimum Residual Method (GMRES) (Saad and Schultz, 1986) which is an iterative technique for solving linear systems. However, as the GMRES only plays a secondary role in the DVM boundary-layer model, we refer the reader to Saad and Van Der Vorst (2000) (or alternatively to Section 2.2) for a general history and overview on using iterative methods to solve linear systems.

Fast-Particle Algorithms

Despite the DVM's advantage of being an essentially grid-free method, its large-scale use (and also the BEM's to a smaller degree) was hampered by the scalability issues associated with the N-Body problem. This is where the fundamental field calculations of N elements requires N^2 unique calculations[5]. This property of the DVM (and all N-Body schemes, including the BEM) causes increases to resolution in a system to quickly become prohibitively expensive. Consequent demand for a fast and efficient method for solving N-Body problems resulted in the de-

[5]Computational complexity order $O(N^2)$.

velopment of various tree-base algorithms and solvers. Early solutions to this problem for the DVM came in the form of vortex-in-cell methods first presented by Christiansen (1973), with later work from others such as Baker (1979). As the DVM is similar in its governing equations to other areas in classical physics such as celestial mechanics involving point masses and plasma physics involving point charges, it is possible to use early monopole calculations and divide-and-conquer strategies developed by Barnes and Hut (1986) and Appel (1985). Shortly after, Greengard and Rokhlin (1987a) developed a similar algorithm titled the Fast Multipole Method (FMM). This algorithm has a computational complexity of $O(N \log N)$ and has since been cited as one of the top ten algorithms of the 20^{th} century (Dongarra and Sullivan, 2000).

While the FMM was very popular with widespread use in particle physics, its published adoption to real applications in fluid mechanics has thus far been limited. Work by Pringle (1994) involved applying the two-dimensional FMM to the DVM, introducing optimisations such as the 'Dynamic-P' principles and investigating its use for parallel computation. Hamilton and Majda (1995) (and others such as Cruz and Barba, 2009) also applied the FMM to the DVM using Gaussian vortex particles and analysed its accuracy and efficiency, concluding that care must be taken to ensure errors are not of significant magnitude. Shortly after publishing their 2D FMM algorithm, Greengard and Rokhlin (1987b) published a FMM algorithm for 3D problems. However, its complexity and the lack of computing power at the time resulted in its practical adoption being limited. Cheng et al. (1999) later published optimisations to the algorithm using new mathematical compression techniques that achieved considerable speed-ups for no cost in accuracy. An example of its adaptation for the purpose of the DVM was provided by Lindsay and Krasny (2001) who used the 3D FMM to simulate the roll-up of circular-disk vortex sheets into vortex rings.

1.5 Summary

The experimental and theoretical investigations that show compliant-wall systems can offer significant drag (and therefore energy) reductions in various modes of transport has sparked a renewed interest in compliant-wall technologies. This outcome has been established through the discovery of three fundamental (from a linear stability perspective) system instability mechanisms; the fluid-based Tollmien-Schlicting Instabilities (TSI) and the two flow-induced structural instabilities of Travelling-Wave Flutter (TWF) and static divergence. However, it has thus far been unknown as to why drag-reduction can only be achieved in tightly controlled laboratory experiments.

Herein we present a numerical tool that can help provide such answers by overcoming the restrictive assumptions of the established analytical methods. We therefore develop a single model that can capture the three main instabilities of a boundary-layer flow over a finite-length compliant surface at nonlinear amplitudes. The Lagrangian Discrete Vortex Method (DVM) is used in conjunction with the Fast Multipole Method (FMM) to capture the (2D) boundary-layer dynamics, yielding results similar to that of Direct Numerical Simulation (DNS) at a fraction of the computational cost. This flow model is coupled to the nonlinear mechanics of a Kramer-type compliant wall using the Boundary Element Method (BEM) and finite-difference method (FDM). The flexibility of the methods chosen allows either the simplified or full systems to be analysed, with potential use being in the optimisation of compliant walls for transition delay.

Chapter 2

Potential-Flow Modelling

We look at the class of fluid flows that can be characterised as incompressible, inviscid and irrotational, known as potential flows. Under these assumptions the governing Navier-Stokes equation reduces to the Laplace equation,

$$\frac{\partial^2 \phi}{\partial x^2} + \frac{\partial^2 \phi}{\partial y^2} + \frac{\partial^2 \phi}{\partial z^2} = \nabla^2 \phi = 0. \tag{2.1}$$

Where ϕ is the scalar velocity potential function that satisfies,

$$U = \frac{\partial \phi}{\partial x}, \quad V = \frac{\partial \phi}{\partial y}, \quad W = \frac{\partial \phi}{\partial z}, \tag{2.2}$$

where U, V and W are the x, y and z components of the velocity in the Cartesian plane. As the Laplace equation is a linear partial differential equation, if $\phi_1, \phi_2 ... \phi_n$ are known solutions, their sum will also be a valid solution,

$$\nabla^2(\phi_1 + \phi_2 ... + \phi_n) = 0. \tag{2.3}$$

One elementary solution to Equation (2.1) is that of a two-dimensional (2D) point source/sink, an element whose flow is solely in the radial direction with a potential function of

$$\phi(\boldsymbol{r}) = \frac{\sigma}{2\pi} \ln \boldsymbol{r}, \tag{2.4}$$

where σ is the element strength and $\boldsymbol{r}$ is the relative position vector to the evaluation point. If the element has a central position of (x_0, y_0), the induced velocity at the evaluation point (x, y) is,

$$U = \frac{\sigma}{2\pi} \frac{x - x_0}{(x - x_0)^2 + (y - y_0)^2} = I_u \sigma, \tag{2.5}$$

$$V = \frac{\sigma}{2\pi} \frac{y - y_0}{(x - x_0)^2 + (y - y_0)^2} = I_v \sigma, \tag{2.6}$$

where I_u and I_v have been introduced to represent the influence of an element based upon its position relative to the evaluation point. As singular point elements are not particularly useful for modelling complex flows of immersed surfaces/bodies, we define the velocity potential of a finite length source/sink line element, referred to as a panel from here on. The panel consists of a constant strength (zero-order) distribution that is formed through linear superposition of a series of singular point elements coincident with the panel. The velocity potential function is therefore obtained through integration as (Katz and Plotkin, 1991),

$$\phi(x, y) = \frac{\sigma}{4\pi}(\Delta x_l \ln {R_l}^2 - \Delta x_r \ln {R_r}^2 + 2\Delta y(\theta_r - \theta_l)), \tag{2.7}$$

where,

$$\Delta y = (y - y_0), \tag{2.8}$$

$$\Delta x_k = (x - x_k), \tag{2.9}$$

$$R_k = \sqrt{({\Delta x_k}^2 + \Delta y^2)}, \tag{2.10}$$

$$\theta_k = \tan^{-1} \frac{\Delta y}{\Delta x_k}, \tag{2.11}$$

for panel ends, $k = l, r$, with a corresponding velocity,

$$U = \frac{\sigma}{4\pi} \ln \frac{{R_l}^2}{{R_r}^2} = I_u \sigma, \tag{2.12}$$

$$V = \frac{\sigma}{2\pi}(\theta_r - \theta_l) = I_v \sigma. \tag{2.13}$$

Panel elements such as these form the core of what is known as the Boundary Element Method (BEM). For potential fluid flows, the BEM that arose from the seminal work of Hess and Smith (1967) has been in use in academia and the aeronautics industry for many years. While many reference textbooks have been published since its inception, we refer to Katz and Plotkin (1991) for a wealth of background theory and potential-flow modelling discussion. In its essence the BEM involves placing potential flow elements about a body/surface immersed in an infinite potential, or mean, flow to solve for the entire flow field that is perturbed by the body/surface. Its key advantage is in eliminating the need for a fluid-based grid, such as that required by finite-volume techniques, because the entire flow-field is condensed down onto a series of surface elements. While the elements collectively describe the flow field at any point in space, their long range effects cause it to be classified as an 'N-Body' method, a difficulty which will be further illustrated.

If the physical body/surface that we wish to model is discretised into N elements, the velocity induced at any point in space, using Equations (2.12) and (2.13), is

$$U = I_{u0}\sigma_0 + I_{u1}\sigma_1 + \cdots + I_{uN}\sigma_N, \tag{2.14}$$

$$V = I_{v0}\sigma_0 + I_{v1}\sigma_1 + \cdots + I_{vN}\sigma_N. \tag{2.15}$$

Evaluating this at M locations and looking at the normal (V) velocity component only,

$$V_1 = I_{v0,1}\sigma_0 + I_{v1,1}\sigma_1 + \cdots + I_{vn,1}\sigma_n, \tag{2.16}$$

$$V_2 = I_{v0,2}\sigma_0 + I_{v1,2}\sigma_1 + \cdots + I_{vn,2}\sigma_n, \tag{2.17}$$

$$\vdots \tag{2.18}$$

$$V_m = I_{v0,m}\sigma_0 + I_{v1,m}\sigma_1 + \cdots + I_{vn,m}\sigma_n. \tag{2.19}$$

In matrix notation we have,

$$\{V\} = [I_v]\{\sigma\}, \tag{2.20}$$

where $\{V\}$ and $\{\sigma\}$ are each column vectors of length M and N respectively and $[I_v]$ represents a rectangular matrix of shape, $M \times N$. This reveals that any velocity field calculation will be of computational complexity $O(MN)$.

2.1 Direct Solutions

At the heart of the BEM, and the most computationally expensive process, is solving for element strengths using suitable boundary conditions. This is required as the panel strengths (σ) that form a given perturbed flow field is rarely known beforehand. All immersed surfaces in this work are modelled using constant strength (zero-order) 2D source/sink panels. The strengths of these are determined by setting the induced normal velocity-flux at the panel centre to zero to ensure no mass can flow across the fluid/solid interface. This is achieved numerically by enforcing,

$$V_p + V_{p\infty} = 0, \tag{2.21}$$

$$I_v\sigma = -V_{p\infty}, \tag{2.22}$$

where V_p represents the normal velocity induced by all the boundary elements and $V_{p\infty}$ represents the apparent normal velocity of the surrounding flow field. By knowing beforehand that the panels will only be moving in the y-direction, the apparent flow velocity is given by,

$$V_{p\infty} = (V_\infty - \dot{\eta}_p)\cos\theta_p - U_\infty \sin\theta_p. \tag{2.23}$$

This is where U_∞ and V_∞ are the x and y velocity components of the mean flow, θ_p is the angle of the panel to the $x+$ axis and $\dot{\eta}_p$ is the velocity of the panel

centre in the $y+$ direction as shown in Figure 2.1.

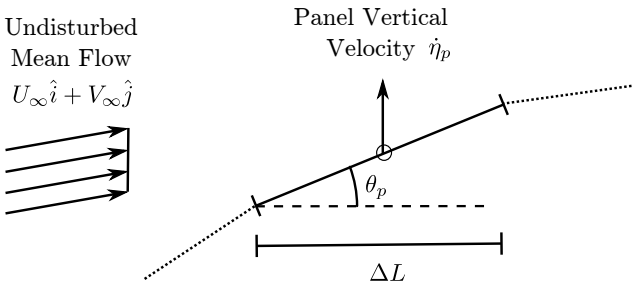

Figure 2.1: Schematic of the BEM notational and arrangement.

Equation (2.22) must be solved across all boundary elements, resulting in a $[I_v]$ matrix that is dense, non-symmetric and for most cases herein, diagonally dominant. One solution to this problem is through linear algebra,

$$\{\sigma\} = [I_v]^{-1}\{-V_{p\infty}\}, \tag{2.24}$$

where we form the square matrix of influence coefficients, $[I_v]$, then determine its inverse. From a computational point of view the inverse is the most expensive operation, with common algorithms such as Gaussian elimination (GE) or LU Factorisation (LU, Bartels and Golub, 1969) being of complexity $O(N^3)$[1]. This is compounded by the fact that in order to perform the direct inverse, we require an explicit form of the $[I_v]$ matrix, which for the standard BEM requires $O(N^2)$ time and storage to produce a dense $N \times N$ matrix.

The scaling of the BEM algorithm remains of critical importance as the accuracy of the flow solution is directly dependent on the level of discretisation of panel elements. Fortunately when dealing with bodies/surfaces that are static or undergoing small amplitude motion, the $[I_v]$ and $[I_v]^{-1}$ only need to be computed once at the start of each simulation. As the aim of this work was to develop a generalised model that could study relatively large, nonlinear surface motions, caching of large dense matrices was not an adequate solution. This is because high amplitude nonlinear motions cause $[I_v]$ to change significantly in time (ie.

[1]However LU offers greater efficiency than GE in practice.

$\partial[I_v]/\partial t \not\approx 0$).

2.2 Iterative Solutions

For a general history and current state of the art of iterative methods for solving linear systems see Saad and Van Der Vorst (2000). The Generalised Minimal Residual method (GMRES) (Saad and Schultz, 1986) is an iterative technique for solving linear systems such as,

$$[I_v]\{\sigma\} = \{-V_{p\infty}\}. \tag{2.25}$$

The method iteratively obtains a solution to $\{\sigma\}$ by minimizing the norm of the residual vector over a Krylov subspace. By requiring only matrix-vector products, the expensive inverse computations are eliminated. However, the iterative nature of the algorithm requires suitable pre-conditioning to improve the rate of convergence. The strength of the pre-conditioned GMRES (or other iterative methods) is that the explicit generation and storage of the dense $(N \times N)$ $[I_v]$ matrix is never required. This allows it to be coupled to the wide range of fast-particle algorithms that are also matrix-free[2] methods.

2.2.1 Pre-Conditioning

It is noted that Benzi (2002) provides a good reference on the history of pre-conditioning and the current direction of research efforts. Pre-conditioning refers to the multiplication of a suitable pre-conditioner $[P]$ that is an approximate inverse to the $[I_v]$ matrix (i.e. $P \approx [I_v]^{-1}$),

$$[P][I_v]\{\sigma\} = [P]\{-V_{p\infty}\}. \tag{2.26}$$

[2]For this work we define a matrix-free method as a computational algorithm that does not require explicit formation of the full system coefficient matrix.

This causes the system equations to be better conditioned and therefore increases the iterative convergence rate. A perfect pre-conditioner would be an exact inverse, $P = [I_v]^{-1}$, obtained through GE or LU, allowing the GMRES to converge in a single iteration. However as this negates the benefits of using an iterative solver, we must find a compromise that offers a good approximation, $P \approx [I_v]^{-1}$, that is computationally cheap to generate and apply, and which offers a net improvement in solution time.

Iterative methods of pre-conditioning (Saitoh and Kamitani, 2004) are an attractive option because the pre-conditioner is nested inside the GMRES loop and applied using a second iterative solver to perform the system pre-conditioning. While this may seem counter-intuitive, the pre-conditioner iterations require only loose tolerances to apply the effects of an approximated inverse. The main strength in this technique is that it is purely matrix-free, allowing solution of extremely large problem sizes. However, for this work moderate discretisations (with respect to computational storage) are to be used, resulting in iterative pre-conditioning being an inefficient (with respect to computational time) choice.

Part of the problem for large systems, particularly when using matrix-free algorithms (such as those discussed in Section 2.3), is that the system matrix ($[I_v]$) is never explicitly constructed and to do so would be cumbersome. An alternative is to form a sparsified approximation to $[I_v]$ that contains only the strongest elements for the purpose of pre-conditioning. The diagonally dominant nature of the BEM problem in this work allows prescribing a sparsifying template beforehand that preserves the K entries nearest the main diagonal, thus controlling the balance between pre-conditioner complexity and quality. As a similar technique, the Sparse Approximate Inverse (SPAI) (Alléon et al., 1997) method is one where an approximate inverse is computed via a Frobenius norm approach using a prescribed non-zero pattern. This procedure suits the utilisation of such a sparse approximation to $[I_v]$ (or the use of matrix-free fast-particle algorithms) to generate a sparse approximate inverse $[P]$, in a column-wise fashion. It has been shown (Rui and Chen, 2007) to benefit from coupling with other matrix-

free BEM algorithms (such as the Fast Multipole Method), and to be suitable for parallel computing (Chow, 2001; Huckle and Grote, 1997). However, preliminary investigations revealed that for the moderate problems sizes encountered in this work it was approximately equal to, or slightly slower than, the popular Incomplete LU Factorisation (ILU) method.

The ILU method (Chan and Van Der Vorst, 1997) is a direct GE type method that effectively computes a matrix inverse which, unlike GE, preserves the sparsity of the original matrix, allowing it to be more time and memory efficient. A particular appeal of the ILU approach is also due to the maturity and availability of computational libraries such as SuperLU (Demmel et al., 1999), which are easily accessible through the SciPy (Jones et al., 2001–) Python libraries/interfaces. As this study involves a surface that undergoes motion about an undisturbed flat position, the form of the pre-conditioner also remains consistent in time, i.e. $\partial[P]/\partial t \approx 0$. It is therefore acceptable to form this pre-conditioner only once and cache it for use throughout the whole simulation.

2.3 The Fast Multipole Algorithm

It is well known that while schemes such as the BEM have the advantage of being essentially grid-free when compared to finite volume methods, they suffer the N-Body problem. This is a significant problem because the accuracy of the BEM relies on the discretisation of the surface/body. Fortunately as the potential-flow source/sink elements in the BEM are a solution to Laplace's equation, the Fast Multipole Method (FMM) (Greengard and Rokhlin, 1987a) can be used. For an introduction to the FMM and its background literature, see Section 1.4.2. While the full FMM algorithm is a process which is best understood graphically or through visualisations such as those produced by Greengard and Rokhlin (1987a), Pringle (1994) or Wang et al. (2005), a brief summary of the process is included here. However, additional texts (Greengard and Rokhlin, 1987a) should be consulted for its theoretical formulation, proof and error characterisation (Cruz and Barba, 2009).

The basis of the FMM is the conversion of the log $\boldsymbol{z}$ potential function, where $\boldsymbol{z}$ is a position vector in the complex plane, into an equivalent infinite Laurent series. For well separated particles and within reasonable error bounds, the Laurent series of relatively close particles can be summed and shifted to a collocation point. As the truncation and manipulation of an infinite Laurent series can be cumbersome, Greengard and Rokhlin (1987a) proposed using the natural hierarchy of a quad tree to speed-up calculations. Under this scheme the FMM broadly involves three sequential steps - the tree-building, the upward pass and the downward pass. In the tree-building phase a suitably dimensioned quad tree is constructed such that it encapsulates the parent-child and neighbouring relations of all boxes at each level. The tree's lowest levels (the leaves) encapsulate the individual source and target particles where the multi-pole expansions are evaluated. The upward-pass begins at the lowest level of the tree. The influence of all the particles in a box are converted into a multi-pole about its centre (see Greengard and Rokhlin, 1987a). These lowest level multi-pole expansions are then translated up through the parents to the highest level of the tree. The downward pass then begins from the highest level where the upward expansions are translated in the far field, onto well separated nearest neighbours. For this purpose a neighbouring box is classified as well separated if its centre is located a distance $3r$ away from the current box encapsulated in radius r. This condition ensures that the error on the translation process is bound to 2^{-p} where p is a chosen integer value to which the infinite-series is truncated. Once all neighbour interactions are completed, the multipole expansions are translated down the tree through children boxes until the lowest level is reached. The accumulated downward expansion contains the influence of all boxes (particles) that are well separated, allowing their evaluation back onto the real targets. Any particle interactions that could not be refined down to a resolution such that it could be classified as far field, the conventional $O(N^2)$ method is applied. This results in an algorithm approximately of the order $O(N \log N)$.

While the basic FMM algorithm offers good efficiency for uniform particle

distributions, the algorithm in this work also uses the Dynamic-P optimisations presented by Pringle (1994) and is multi-levelled to ensure computational memory and time efficiency. The code is fully implemented in C++ (although interfaced mainly in Python) and is built using the C++ Standard Template Libraries (STLs) to store efficiently the adaptive-tree data structures. Due to modern desktop computers advancing along the multi-cored, shared memory path, the code is also multi-threaded using C-Pthreads.

The implications of the FMM algorithm is that for any N-Body problem, the computational time for field evaluations is drastically reduced and, often more importantly, conducted in a memory-efficient matrix-free manner. A side effect of the FMM's matrix-free structure is that it couples extremely well with the GMRES discussed in Section 2.2. This is where a call to the FMM algorithm, returns the evaluation of the matrix-product $[I_v]\{\sigma\}$. When used in this context further optimisations can be made by caching the FMM tree between GMRES iterations, as particle positions which form the basis of the quad tree do not change.

Unless otherwise noted, the FMM algorithm will be used throughout the work for the majority of velocity/potential field and GMRES matrix-product evaluations. Performance testing of the specific implementation will be presented in Chapter 5.

2.4 Pressure Formulation

To determine the pressure in a potential flow field at any point, including on a surface or body, one can apply the unsteady Bernoulli equation along a streamline. By comparing the fluid state at a point in the unperturbed flow infinitely upstream to the point of interest, the pressure difference is,

$$\Delta p = \rho_f \left(-\frac{U_\infty^2}{2} + \frac{\partial \phi}{\partial t} + \frac{U^2 + V^2}{2} \right). \tag{2.27}$$

This is where ρ_f is the fluid density, U_∞ is the undisturbed upstream velocity and ϕ, U and V are the velocity potential, horizontal and vertical velocities respectively at the evaluation point in the fluid.

2.5 Illustrative Results – Rankine Oval

To illustrate the use of the mesh-free framework available for potential flow modelling using the BEM, we apply it to the case of potential flow around a Rankine Oval. The theoretical Rankine Oval is formed through the superposition of three potential-flow elements; a uniform flow (U_∞) and a point-source and point-sink element separated on an axis parallel to the uniform flow by a distance of $2a$, with equal, but opposing strengths of m. Due to symmetry of the flow, only half of the oval is modelled with the inclusion of upstream and downstream splitter plates. Figure 2.2 illustrates this arrangement and the shape of the streamline formed about the body's surface. The expression used to define this streamline

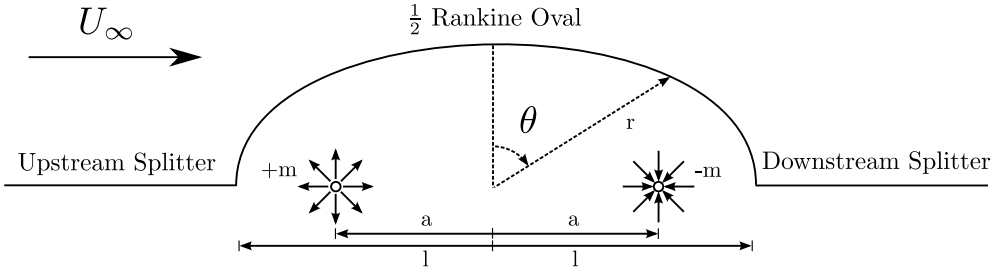

Figure 2.2: Schematic of a Rankine oval with splitter plate immersed in a potential flow.

is (Munson et al., 2002),

$$0 = r' \sin\theta - \frac{m'}{2\pi} \tan^{-1}\left(\frac{2r' \sin\theta}{r'^2 - 1}\right), \tag{2.28}$$

where

$$r' = \frac{r}{a}, \quad m' = \frac{m}{U_\infty a}, \quad l' = \frac{l}{a}. \tag{2.29}$$

By applying the steady-flow assumption, we use the Bernoulli equation (Equation (2.27)) to obtain the surface pressure,

$$\Delta p = \frac{1}{2}\rho_f U_\infty^2 - \frac{1}{2}\rho_f \left((U_\infty + U_p)^2 + {V_p}^2 \right). \tag{2.30}$$

For the theoretical case, U_p and V_p are the sum of the velocity induced by the source and sink point element. We also define the dimensionless coefficient of pressure,

$$C_p = \frac{\Delta p}{\frac{1}{2}\rho_f U_\infty^2}. \tag{2.31}$$

To test the suitability of the BEM, GMRES and FMM combination, we superimpose discretised source panels elements over the top of the splitter plate and Rankine Oval surface. The discretisation of the panels is based on a uniform angular separation ($\Delta\theta = \text{const}$) due to its simplicity. The perturbation caused by the body in the flow is enforced through the no-normal flow boundary condition at the panel centres. To obtain the pressure on the surface we use Equation (2.30) and substitute the sum of the velocities induced by all N panels for U_p and V_p. The GMRES (Section 2.2) is used to solve for the element strengths ($\{\sigma\}$) and the FMM algorithm (Section 2.3) is used for all $[I_*]\{\sigma\}$ evaluations, including that of the GMRES. To distinguish the errors induced by BEM discretisation and the FMM/GMRES combination, we also obtain a numerical solution without the use of the FMM/GMRES using linear algebra techniques (GE inverse). We will refer to this linear algebra solution as the 'standard BEM'.

Figure 2.3 illustrates the solution obtained for the pressure at the surface for the two different BEM methods with panel discretisations about the Rankine Oval set to $N = 100$. The simulation parameters are set to $m' = 2.06$, $l' = 1.29$ with splitter plates of length $l' = 1.29$. For the FMM/GMRES case, a maximum error tolerance of 1×10^{-9} (equivalent to a series trucation limit of $p = 30$) and 1×10^{-8} are used for the FMM and GMRES respectively.

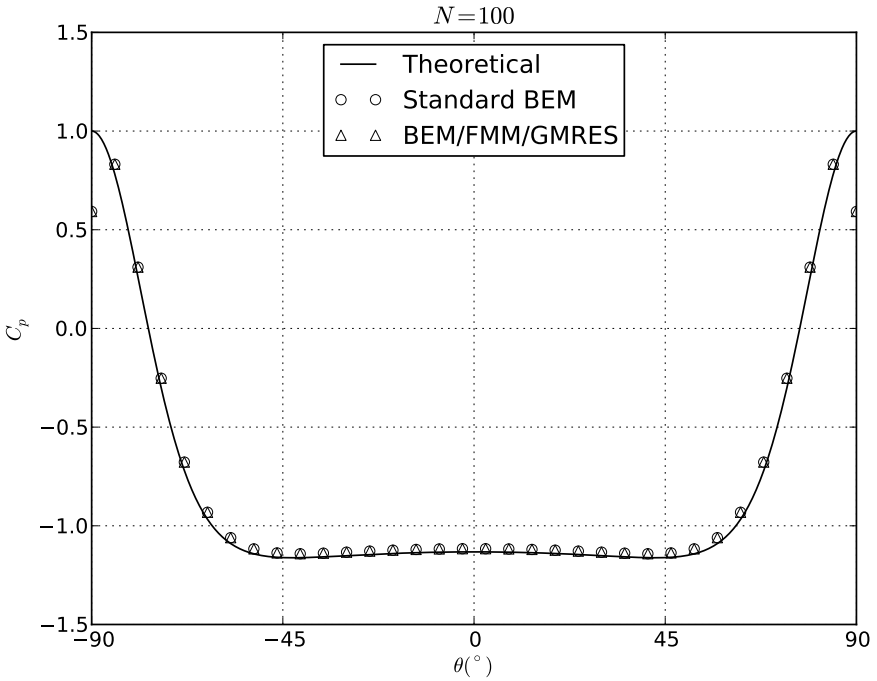

Figure 2.3: Coefficient of pressure comparison at various locations on the Rankine Oval surface using 100 panels.

It is evident that there is good agreement for both methods, with the exception of results near $\theta = \pm\pi/2$. This is as a result of the large relative panel angle (90°) where the splitter plate and the Rankine Oval join. Resolution of this issue was not sought as all uses of the BEM in this work are based on relatively smooth surface geometries where pressure about splitter plate joints and other surface discontinuities are not required. Figure 2.4 shows the convergence of the BEM solutions to the theoretical solution with increasing panel discretisation based on a reference number of nodes of $N_0 = 25$. For calculations of the RMS error, the first and last nodes were left out of the calculation to prevent skewing of the results. The results show that for the tested FMM and GMRES tolerances the only error induced in the pressure coefficient at the surface is that due to the BEM discretisation itself. Figure 2.5 shows the relative time taken for the standard BEM compared to the FMM/GMRES combination as the wall discretisation increases. The reference number of nodes remains the same as for

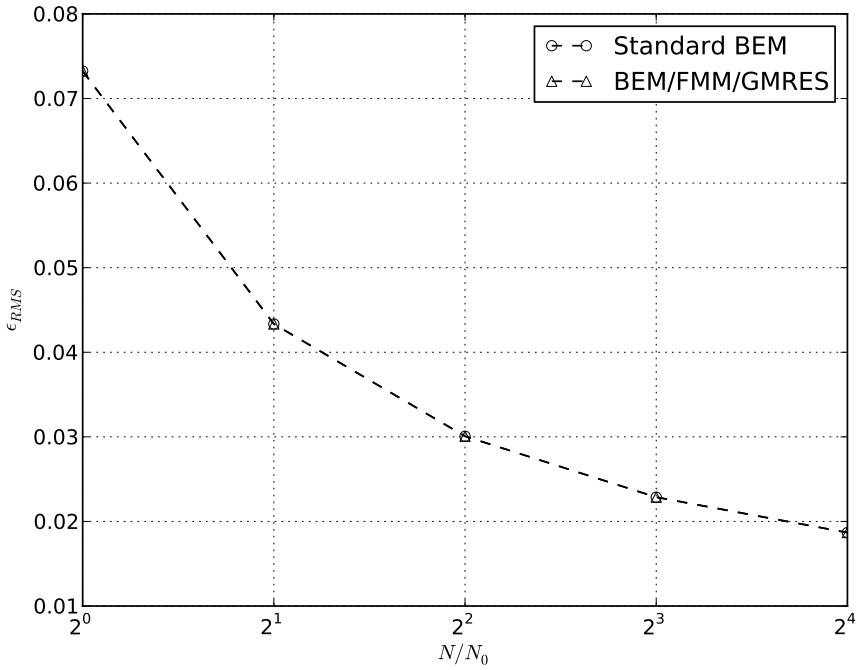

Figure 2.4: RMS error in C_p for numerical methods using various panel discretisations.

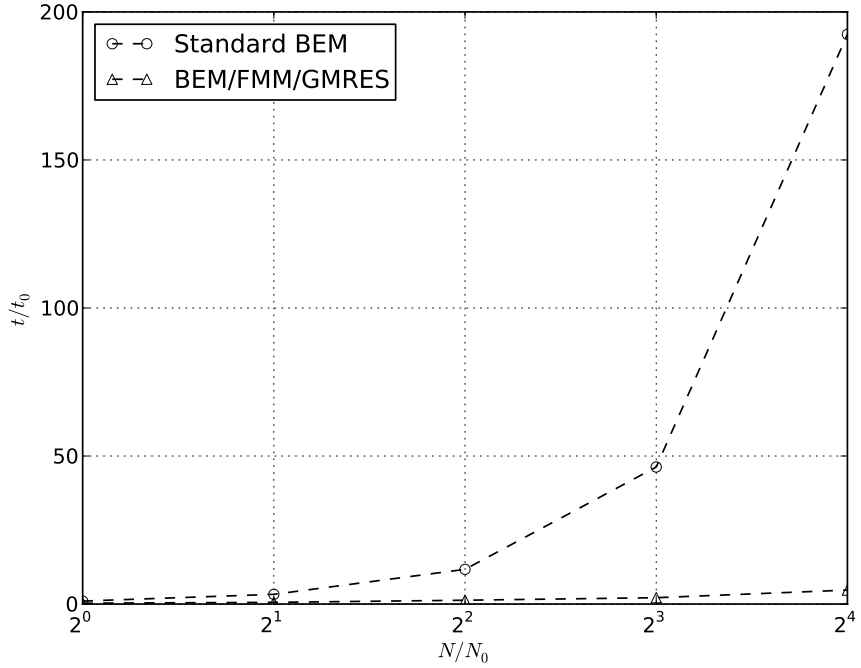

Figure 2.5: Relative time taken for different BEM numerical methods.

Figure 2.4 while the reference time is based on the time taken for the standard BEM with a discretisation of N_0. The FMM/GMRES algorithm demonstrates approximately $O(N \log N)$ scaling, whereas the standard method yields $O(N^2)$ scaling. For this case the FMM/GMRES combination is faster for all discretisation levels with the FMM/GMRES at $N = 400$ taking the same computational time as the standard algorithm at $N = 50$. Typically low element numbers favour the standard methods due to the inherent overhead in running the FMM, GMRES and its necessary pre-conditioning algorithms. However, that was not reflected in these results because the standard BEM algorithm was called through the FMM/GMRES framework for the purpose of validation. While calculation times for the standard method could have been improved for low element numbers (through overhead reductions), even when a coarse discretisation is required, the FMM/GMRES combination is favoured due to its flexibility and potential for vastly reduced computational time and computational memory savings at even mildly increased discretisation levels. This benefit of the FMM/GMRES combination is also amplified when used within an iterative framework that requires many flow-field evaluations throughout a single simulation.

2.6 Summary

Based on the agreement and inference from these results, unless otherwise stated, for the remainder of this study the BEM will be used to represent physical boundaries. Where the BEM is involved, the GMRES with an error tolerance of 1×10^{-8} will be used in conjunction with the FMM with an error tolerance of 1×10^{-9} (equivalent to a series truncation limit of $p = 30$). From these results, an ideal BEM discretisation cannot be inferred because the problem is too specific to yield general discretisation guidlines. It does, however, illustrate that even with high relative angles between panels and coarse discretisations, the RMS error in pressure remains acceptably small.

Chapter 3

Compliant-Wall Modelling

With surface modelling of the potential fluid flow achieved through the BEM, we move to developing a numerical method for modelling the structural dynamics of an immersed surface/wall. We use an Euler-Bernoulli approach that consists of a one-dimensional (1D), thin beam which undergoes small (linear) amplitudes of motion. However, to extend this model to nonlinear ranges we utilise the same method as Lucey et al. (1997b) by including a nonlinear tension term. If the vertical wall motion (η) is a function of horizontal position (x) and time (t), i.e. $\eta = \eta(x,t)$, the governing differential equation for the modified Euler-Bernoulli beam is,

$$\rho h\ddot{\eta} + B\frac{\partial^4 \eta}{\partial x^4} - T_I\left(\eta\right)\frac{\partial^2 \eta}{\partial x^2} = F(x,t), \tag{3.1}$$

where ρ is the wall density, h the thickness and F is the forcing function applied at the wall's surface. The time differentials are represented by,

$$\dot{\eta} = \frac{\partial \eta}{\partial t}, \quad \ddot{\eta} = \frac{\partial \dot{\eta}}{\partial t} = \frac{\partial^2 \eta}{\partial t^2}. \tag{3.2}$$

In addition B denotes the flexural rigidity of the plate and T_I is the non-local induced tension coefficient given by,

$$B = \frac{Eh^3}{12\left(1-{v_p}^2\right)} \quad , \quad T_I\left(\eta\right) = \frac{Eh}{L\left(1-{v_p}^2\right)} \int_0^L \left(\sqrt{1+\left(\frac{\partial \eta}{\partial x}\right)^2} - 1 \right) dx, \quad (3.3)$$

where v_p refers to Poisson's ratio, E the elastic modulus and L the total length of the wall. The inclusion of T_I captures the nonlinearity that arises from high amplitudes of wall displacement, an element that tends to zero for small amplitudes. For the purpose of simulating a Kramer-type compliant surface, a spring-damper foundation has been included (Carpenter and Garrad, 1985, 1986; Lucey et al., 1997b) as shown in Figure 3.1, changing the system equation to,

$$\rho h \ddot{\eta} + B\frac{\partial^4 \eta}{\partial x^4} - T_I\left(\eta\right)\frac{\partial^2 \eta}{\partial x^2} + D\dot{\eta} + K\eta = F(x,t), \quad (3.4)$$

where K is the stiffness coefficient of the foundation springs and D the damping coefficient.

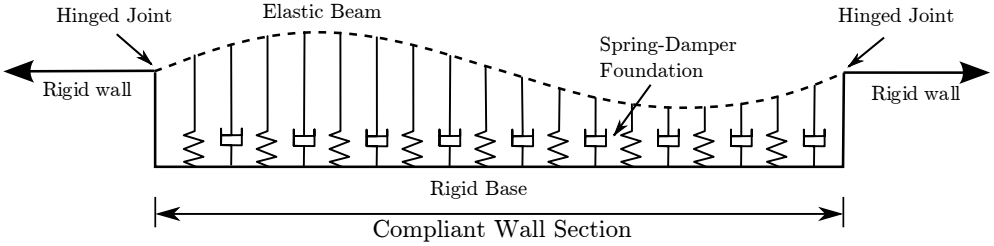

Figure 3.1: Schematic of Kramer-type compliant wall.

3.1 Numerical Solution

All spatial differentials are solved using the Finite Difference Method (FDM). The FDM is the most suitable method due to the wall being slender, 1D and with uniform wall properties to allow for regular discretisation. Finite Element Analysis (FEA) would be the natural progression for future works where structural models of higher complexity are used. While the current model does not

warrant the additional complexity and computational cost, all numerical techniques utilised in this work are to be directly amenable to such models and their associated solution methods.

The FDM solution is achieved numerically by discretising the continuous wall into N nodes of mass, uniformly spaced by Δx in the horizontal direction,

$$[\eta] = \begin{bmatrix} \eta_{x_1} \\ \vdots \\ \eta_{x_N} \end{bmatrix}. \tag{3.5}$$

By maintaining uniform separation of the nodes we can apply second-order central-difference approximations (Trefethen, 1996) to all spatial differentials. Looking at node i we have,

$$\frac{\partial \eta_i}{\partial x} \approx \frac{\eta_{i+1} - \eta_{i-1}}{2\Delta x}, \tag{3.6a}$$

$$\frac{\partial^2 \eta_i}{\partial x^2} \approx \frac{\eta_{i-1} - 2\eta_i + \eta_{i+1}}{\Delta x^2}, \tag{3.6b}$$

$$\frac{\partial^4 \eta_i}{\partial x^4} \approx \frac{\eta_{i-2} - 4\eta_{i-1} + 6\eta_i - 4\eta_{i+1} + \eta_{i+2}}{\Delta x^4}. \tag{3.6c}$$

Use of higher-order accuracy approximations for the first and second-order differentials (Equations (3.6a) and (3.6b)) is unnecessary because the fourth-order differential is the component with the greatest magnitude in most cases herein.

As suitable boundary conditions to the FDM, we implement hinged end-joints through the addition of phantom nodes at the wall ends. The resulting discretised system after applying the FDM is

$$\begin{aligned} \rho h \ddot{\eta}_i + B\frac{\eta_{i-2} - 4\eta_{i-1} + 6\eta_i - 4\eta_{i+1} + \eta_{i+2}}{\Delta x^4} \dots \\ -T\frac{\eta_{i-1} - 2\eta_i + \eta_{i+1}}{\Delta x^2} - d\dot{\eta}_i + K\eta_i - F_i = 0, \end{aligned} \tag{3.7}$$

where

$$T = \frac{Eh}{L\left(1 - {v_p}^2\right)} \sum_{i=1}^{n} \left(\sqrt{1 + \left(\frac{\eta_{i+1} - \eta_{i-1}}{2\Delta x} \right)^2} - 1 \right) \Delta x. \tag{3.8}$$

The full system forms a series of coupled nonlinear ordinary differential equations which can be solved as an initial-value problem using suitable time integration schemes. For this work we use a second-order, implicit, trapezoidal time-stepping method. The use of explicit or semi-implicit time-stepping methods such as that used by Lucey et al. (1997b) and Pitman (2007), has been avoided as the numerical stiffness of the equations of motion results in the convergence rate and stability of the problem being highly dependent on wall discretisation and time-step sizes. Thus, an increase in spatial discretisation requires a much larger increase (of order $\Delta t \propto \Delta x^2$) in time discretisation to ensure numerical stability (Richtmyer and Morton, 1967). By using an implicit method, time-step sizes only impact the solution accuracy and no longer impact the numerical stability of the system. Using the implicit trapezoidal method, all time differentials take the form,

$$\eta_t = \eta_{t-1} + \Delta t \frac{\dot{\eta}_{t-1} + \dot{\eta}_t}{2}, \tag{3.9}$$

$$\dot{\eta}_t = \dot{\eta}_{t-1} + \Delta t \frac{\ddot{\eta}_{t-1} + \ddot{\eta}_t}{2}, \tag{3.10}$$

illustrating the implicit coupling of $\ddot{\eta}_t$ in Equation (3.4). To arrive at the semi-implicit (predictor-corrector) technique of Lucey et al. (1997b) and Pitman (2007) we apply a predictor ($\ddot{\eta}_{t^*}$) to Equation (3.10) using $\ddot{\eta}_t = \ddot{\eta}_{t^*}$ to effectively decouple the acceleration (inertial) based terms in Equation (3.4) from that of the velocity (damping) and displacement (stiffness). Therefore,

$$\eta_t = \eta_{t-1} + \Delta t \frac{\dot{\eta}_{t-1} + \dot{\eta}_t}{2}, \tag{3.11}$$

$$\dot{\eta}_t = \dot{\eta}_{t-1} + \Delta t \frac{\ddot{\eta}_{t-1} + \ddot{\eta}_{t^*}}{2}. \tag{3.12}$$

This allows the system to be solved explicitly using a Jacobi or Gauss-Seidel method, updating the predictor ($\ddot{\eta}_{t^*}$) until it converges on the corrector ($\ddot{\eta}_t$). While the method offers a simple technique for solving a subset of specific (typically small) systems it does not maintain sufficient flexibility or robustness that will be required in future works; however, this will be discussed further in Chapter 4.

3.1.1 Implicit-Time Solution

The trade-off with using an implicit time-stepping method is that to obtain a solution we must find the roots of the system equation in the form $f(\ddot{\eta}) = 0$. Where the system equations are analytically defined beforehand, methods such as the Newton-Raphson (NR) approach can be used. However, the inclusion of nonlinear tension makes obtaining analytical forms of the derivative (or Jacobian, $[J]$, for systems of equations) impractical. While numerical approximation and inversion of the Jacobian matrix is possible through finite-difference and Gaussian-elimination, with increasing levels of wall discretisation and the Jacobian's square shape ($N \times N$), these operations becomes prohibitively expensive on the order of $O(N^2)$ and $O(N^3)$ respectively.

A better alternative for large systems of nonlinear equations is the Newton-Krylov method (NK, Knoll and Keyes, 2004). This is based on the use of a finite-difference approximation (typically first-order) to the Jacobian-vector product used in the NR. The advantage of this formulation is that one can use a Krylov subspace method, such as the Generalised Mean Residual method (GMRES, see Section 2.2), to iteratively determine the vector product. This method is suited to large systems of nonlinear equations because only system function (f) evaluations are required, allowing the whole scheme to be conducted in a matrix-free (refer Section 2.2) manner.

Pre-Conditioning

For good performance of any iterative Krylov method, the system of equations must be well conditioned. To achieve this we require a pre-conditioner that can be used in the Krylov calculations to provide a good approximation for the inverse of the Jacobian matrix. Here the problem lies in the fact that: i) the Jacobian does not explicitly exist, and obtaining it would be numerically expensive ($O(N^2)$), ii) even with a fully defined Jacobian matrix, its direct inversion is numerically expensive ($O(N^3)$), and iii) the pre-conditioner must be able to reduce the total number of Kyrlov and Newton iterations. Fortunately the Jacobian matrix exhibits properties similar to the BEM influence coefficient matrix discussed in Section 2.2.1, allowing similar pre-conditioning methods to be utilised.

A method for obtaining a numerical approximation to $[J]$ can be obtained using a finite-difference approach by applying small perturbations to the system input and measuring changes in the system output,

$$[J] \approx \begin{bmatrix} \Delta f_1/\Delta\ddot{\eta}_1 & \cdots & \Delta f_1/\Delta\ddot{\eta}_N \\ \vdots & \Delta f_i/\Delta\ddot{\eta}_j & \vdots \\ \Delta f_N/\Delta\ddot{\eta}_1 & \cdots & \Delta f_N/\Delta\ddot{\eta}_N \end{bmatrix}, \tag{3.13}$$

where,

$$\Delta f_i/\Delta\ddot{\eta}_j = \frac{f_i\left(\ddot{\eta}_j + \epsilon\right) - f_i\left(\ddot{\eta}_j\right)}{\epsilon}, \tag{3.14}$$

and ϵ is a small (eg. $\epsilon = 1 \times 10^{-7}$) perturbation applied to individual node displacements. The resulting matrix is diagonally dominant due to the FDM stencils using a 3 and 5 point non-zero pattern on the leading diagonal. This is suited for sparsification based on a non-zero template that includes only terms nearest the leading diagonal. The Incomplete-LU Factorisation (ILU) method is used to provide an inverse to $[J]$, however, due to its sparse nature, the Sparse Approximate Inverse (SPAI) method would also be a natural selection for much larger systems, or those undergoing significant change. For all scenarios in this work, the NK

pre-conditioner will be generated only once and cached for re-use throughout the whole simulation. This is because the underlying linear dynamics remain static in $[J]$ and dominate the majority of the system behaviour. While this remains true for the mildly nonlinear systems treated herein, extremely nonlinear cases where the Jacobian is dense and changes rapidly in time ($\partial[J]^{-1}/\partial t \not\approx 0$) would warrant the pre-conditioner being recalculated frequently using a more efficient method such as the SPAI.

3.2 Illustrative Results – Wall In Vacuo

To test the suitability of the NK coupled with the FDM we apply it to the case of a vibrating beam in a vacuum with spring backing. While there are known analytical solutions to the motion of an Euler-Bernoulli beam in a vacuum, the inclusion of nonlinear tension renders these solutions invalid. However, if we constrain our system to motions of small amplitudes, we can obtain the dynamic response using numerical methods and use linear analytical methods for validation. With linearisation, no external loading or damping; $T_I = 0$, $F(x,t) = 0$ and $D = 0$. Therefore the system equation (Equation (3.4)) simplifies to,

$$\rho h \ddot{\eta} + B\frac{\partial^4 \eta}{\partial x^4} + K\eta = 0. \tag{3.15}$$

By looking at undamped harmonic oscillations of the wall in the linear region we obtain the theoretical natural frequency (ω_n) from,

$$\omega_n = \sqrt{\frac{K + \left(\frac{n\pi}{L}\right)^4 B}{\rho h}}, \tag{3.16}$$

where n is equal to the oscillation mode number. For this case backing spring stiffness is set to that of a springless-beam in deflection mode one,

$$K = \left(\frac{\pi}{L}\right)^4 B. \tag{3.17}$$

This gives an equivalent weighting to the bending stiffness and the backing spring stiffness.

Figures 3.2 to 3.4 show a time-lapsed history of the wall undergoing a single oscillation of free vibration given initial displacements in the form of vibration modes one, two and three respectively. To determine the validity of the nodal

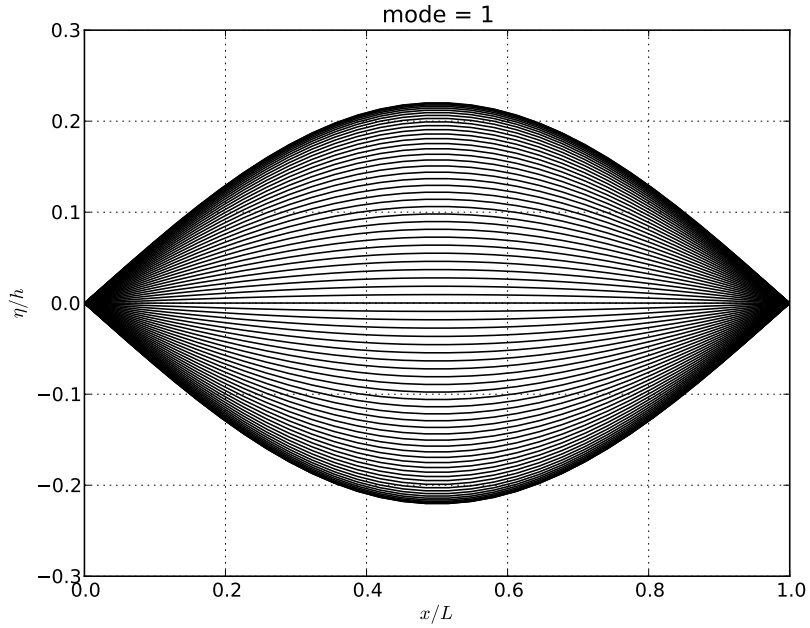

Figure 3.2: Wall positions throughout a single harmonic oscillation of the first mode shape.

discretisation and its influence on accuracy, the error in the numerically measured natural frequency and the theoretically predicted frequency is compared using,

$$\epsilon_\omega = \left| 1 - \frac{\omega_{\text{numerical}}}{\omega_n} \right|, \tag{3.18}$$

where $\omega_{\text{numerical}}$ is the numerical frequency measured over 25 full oscillations. Only the first mode of vibration was investigated. The minimum discretisations determined for this mode can be extrapolated to other frequencies and modes

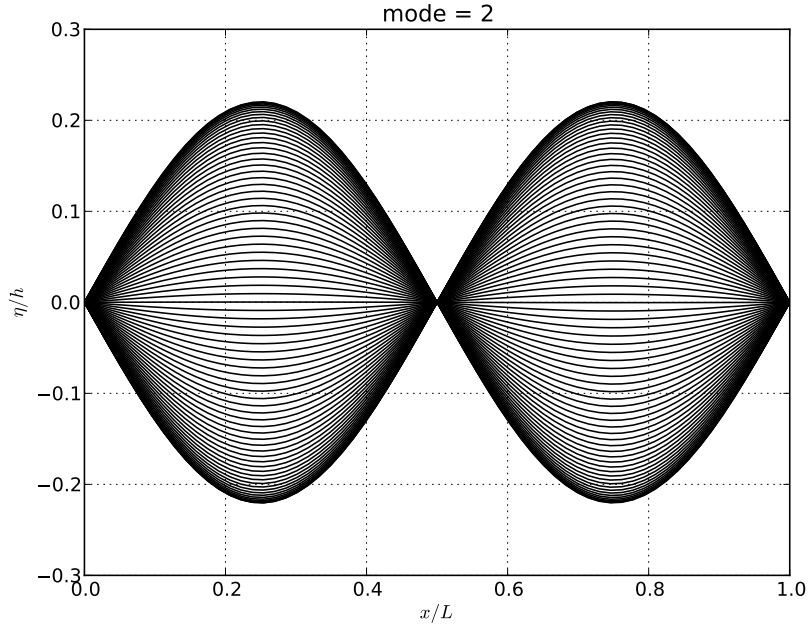

Figure 3.3: Wall positions throughout a single harmonic oscillation of the second mode shape.

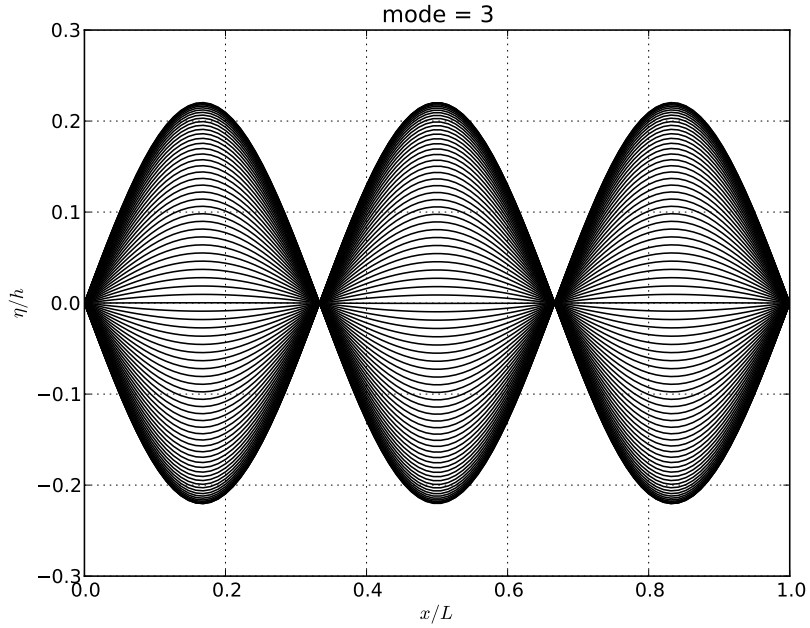

Figure 3.4: Wall positions throughout a single harmonic oscillation of the third mode shape.

using the relationship,

$$N_n = n(NPW_{n=1}), \tag{3.19}$$

$$\Delta t_n = T_n(TPO_{n=1}), \tag{3.20}$$

where N_n is the number of spatial nodes required to resolve mode number n, NPW is the following determined minimum nodes-per-wavelength, Δt_n is the time-step size required to resolve the time period of oscillation T_n and TPO is the following determined minimum time-step-per-oscillation.

Figure 3.5 illustrates the effect of time discretisation on the numerical solution by comparing different TPOs. This was performed using a nodal discretisation of NPW = 64. The ability of the NK to yield stable results as low as $TPO = 4$

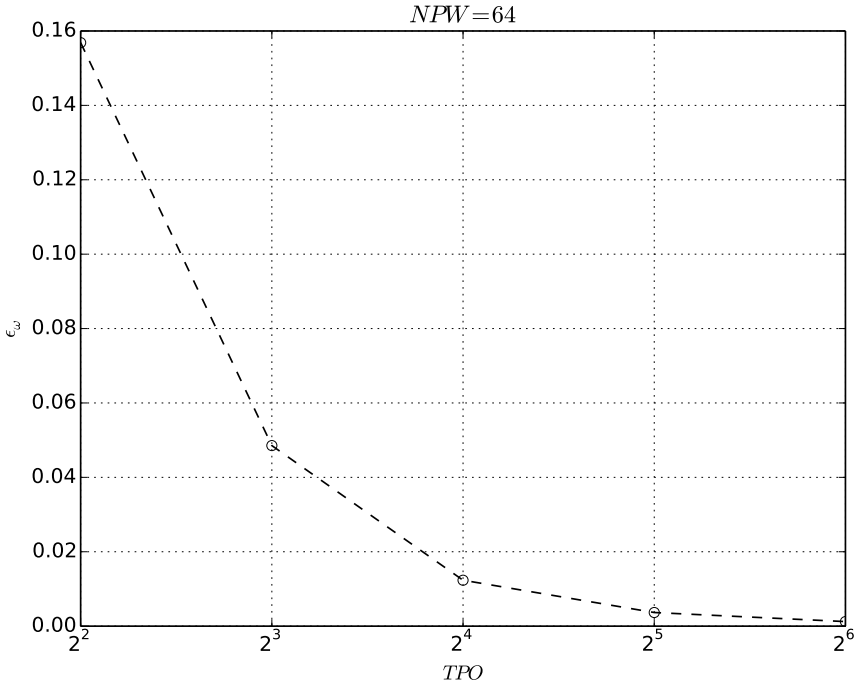

Figure 3.5: Relative error in oscillation frequency for different time discretisations.

demonstrates the stability advantage of using an implicit time-stepping over a semi-implicit (or explicit) scheme. Use of a semi-implicit scheme would necessitate time discretisations to be much smaller and coupled to nodal discretisations

to ensure solution stability, a behaviour demonstrated in Section 4.3.5. It is observed that the frequency error reduces with increasing time discretisation, with $< 1\%$ error achievable using $TPO = 16$. By fixing $TPO = 16$ we investigate the dependency of the frequency error on FDM nodal discretisation, as shown in Figure 3.6. When compared to Figure 3.5 we see that wall discretisation of-

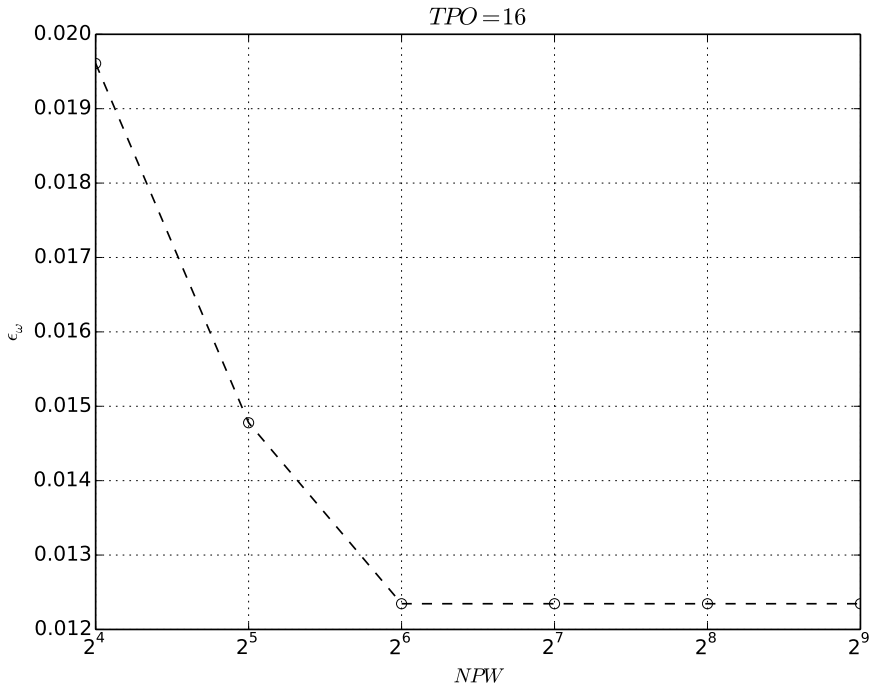

Figure 3.6: Relative error in oscillation frequency for different spatial discretisations.

fers a relatively smaller influence on the frequency error, evidenced by the lowest nodal discretisation of $NPW = 16$ exhibiting relatively low frequency errors at $< 2\%$. Increasing spatial discretisation to $NPW > 32$ offers no improvement in frequency error because the accuracy is limited by the fixed time discretisation.

To demonstrate the scalability of the FDM and NK combination, an element of prime importance in this work, we show the computational cost in run-time[1] (T_C) arising from various spatial discretisations in Figure 3.7. This is where the base nodal discretisation for comparison was set as $NPW_0 = 16$, corresponding

[1]Run-time is defined as the elapsed real time required to complete a computation.

to a base run-time of $T_{C0} = 0.02$s per time-step on an Intel Q9650 desktop PC. From these results, the NK method coupled with FDM demonstrates its excellent

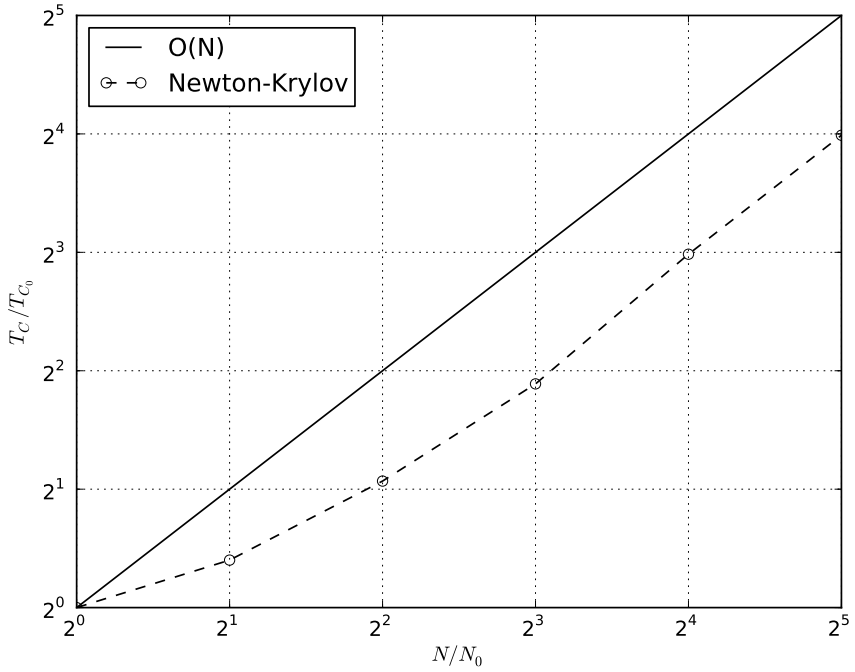

Figure 3.7: Relative computational run-time for various nodal discretisations.

scalability to the order of $\approx O(N)$. The appearance of improved scaling for lower levels of NPW in Figure 3.7 is acknowledged as an artefact of the fixed overhead in the NK implementation.

Although this test case is limited to the linear range, the results obtained and numerical methods are not. This is because the equivalent of a system matrix inverse has been constructed in the NK at every time-step without linearising simplifications. While the power of the NK may not be necessary for linear cases such as that tested, its stability, accuracy, flexibility and computational efficiency, in addition to its matrix-free nature, make this the ideal choice for any time-dependent finite-difference/finite-element method. As a consequence, for the remainder of this work the NK will be used in conjunction with the FDM for all time dependent compliant-wall simulations.

3.3 Summary

The inherent stability benefits of using an implicit time-stepping scheme when coupled with the Finite Difference Method (FDM) were demonstrated. The Newton-Krylov (NK) method offers a suitable technique for overcoming the inherent disadvantages of implicit solution methodologies. The combination of the FDM and NK results in an accurate and efficient, matrix-free tool for the solution of time dependent, nonlinear compliant-wall systems, even those based on wall mechanics more complicated than that presented here. While the simplistic nature of these results do not allow the selection of an ideal wall and time discretisation, they allow confidence in the accuracy of capturing underlying (linear) dynamics when investigating the natural frequency and mode shapes of harmonic oscillations. The benefit of using the NK is that it could easily be extended to finite-element methods that permit the inclusion of more complicated structural systems.

Chapter 4

Inviscid Fluid-Structure Interaction

Thus far, structural modelling has been demonstrated for the case of a beam undergoing free vibration *in-vacuo* using the Finite-Difference (FDM) and implicit-Newton-Krylov (NK) methods. While coupling of the NK and FDM for these systems represents an improvement on prior works (using semi-implicit methods), we are more interested in systems of higher complexity involving coupled fluid-structure interactions (FSI). These systems are of interest to the Engineering community as such interactions can be the cause of destructive flow-induced vibrations.

Real flows involve dynamics caused by boundary layers, the region of fluid flow largely dominated by viscous forces and rotational behaviour. This behaviour occurs in regions that are near a wall/structure that, in many cases, represents only a small portion of the total flow. While these boundary-layer dynamics can be important to FSI, much of the system behaviour of interest to the Engineering community can be captured by modelling the effects of the dominant irrotational, outer region of the flow beyond the thin boundary layer. The particular advantage of modelling the FSI system while assuming irrotational potential-flow behaviour is that it offers fluid-side simplicity and a foundation that is easily extendible to other flow regimes. Where relevant, the effects of the viscous and rotational

near-wall region can sometimes be provided in the form of corrections, which is left as a topic for Chapter 7.

Herein we study the features of the potential-flow FSI model in Figure 4.1.

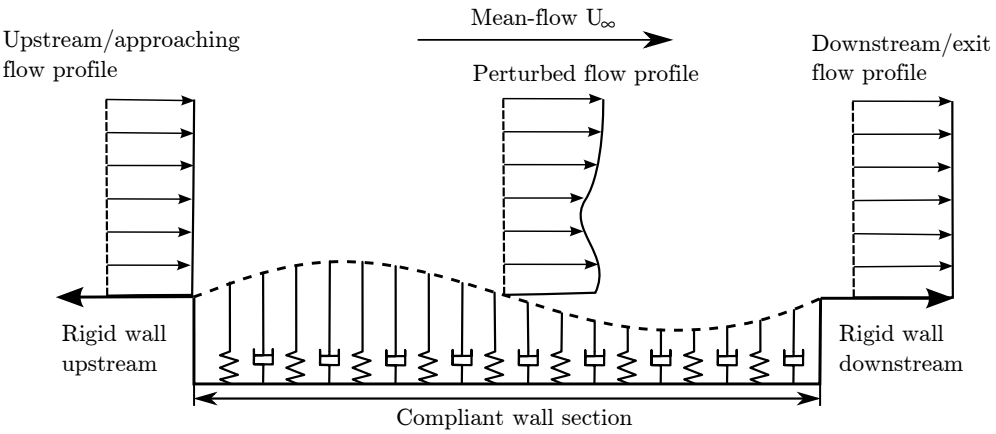

Figure 4.1: A model schematic for the compliant-wall immersed in a potential-flow.

4.1 Fluid Modelling

We have previously developed and presented (in Chapter 2) a suite of tools for modelling potential flow over surfaces using the Boundary Element Method (BEM). The BEM is a technique that has been applied to the nonlinear compliant-wall potential-flow FSI problem in Lucey et al. (1997b) and Pitman and Lucey (2009) with good success. The key features that make the BEM ideal for the present problem are; i) its ability to model the entire flow field using a surface integral method, ii) its mesh-free nature allows easy deformation of the fluid-wall boundary, and iii) pressure forces at the fluid-wall boundary can be calculated with relative ease.

Under the BEM the fluid-wall coupling is achieved using source/sink panels that follow the wall motion and are located between the discretised mass nodes of the FDM, as shown in Figure 4.2. By recalling the equations of the fluid-field

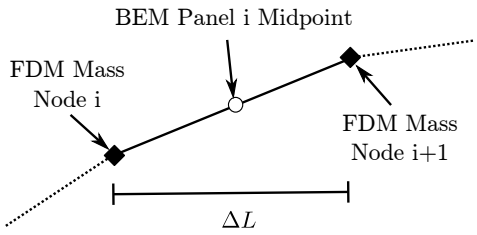

Figure 4.2: Illustration of the FDM mass node and BEM panel layout.

calculations we have,

$$\{\phi\} = [I_\phi]\{\sigma\} \quad , \quad \{U\} = [I_U]\{\sigma\} \quad , \quad \{V\} = [I_V]\{\sigma\}, \tag{4.1}$$

where $\{\sigma\}$ is the $N \times 1$ sized vector of source/sink panel strengths and $[I_\phi]$, $[I_U]$ and $[I_V]$ are corresponding influence matrices of size $N \times N$. As the BEM panels are fixed to the wall, their spatial location and rotation are governed by their local wall displacement (η). As influence matrices are a function of panel angles and positions, they will change in time according to the structural behaviour. Recalling that to determine the element strengths ($\{\sigma\}$) we solve the set of equations that enforce the zero normal velocity at the centre of each panel,

$$0 = [I_v]\{\sigma\} + \{V_{\infty p}\}, \tag{4.2}$$

$$\{\sigma\} = [I_v]^{-1}\left\{-U_\infty \sin\theta_p - \frac{\partial \eta_p}{\partial t}\cos\theta_p\right\}, \tag{4.3}$$

where I_v, θ_p and $\dot{\eta}_p$ also change in time according to the structural behaviour. The changing wall motion impacts the flow-field in three ways via;

1. the position affecting the relative angle of the panel to the mean-flow direction,

2. the curvature in the panel-normal matrix, and

3. directly through the coupled (nonzero) wall velocity.

Despite the changes to the flow-system, the Fast Multipole Method (FMM) and Generalised Mean Residual method (GMRES) maintain their matrix-free nature,

avoiding the formation and manipulation of large square matrices, a complement to the existing structural-modelling framework.

4.2 Structural Modelling

So far the interactions of the fluid with the bounding structure represented FSI systems where the *in-vacuo* structural behaviour dominated. Such systems are of interest to the Engineering community, for example, in the design of Micro-Electronic-Membrane systems that can be used to dynamically impact fluid behaviour. However, we are more interested in two-way FSI systems where the fluid behaviour can also modify the structural dynamics.

When the fluid-flow perturbs under the wall motion, it provides a reactionary response onto the wall as a result of its change in momentum and thus surface pressure. The local change in pressure (Δp) in an unsteady potential-flow field can be determined by applying the unsteady-Bernoulli equation along the surface streamline,

$$\Delta p\left(\ddot{\eta}, \dot{\eta}, \eta\right) = \rho_f \left(-\frac{U_\infty^2}{2} + \frac{\partial \phi}{\partial t} + \frac{U^2 + V^2}{2} \right). \tag{4.4}$$

As previously, (see Equation (2.30)) the velocity components U and V represent the velocity induced by all potential-flow elements, i.e. the boundary elements and the mean-flow. Pressure terms involving velocity (U, V) and its potential (ϕ), change dynamically according to the structural dynamics, while the unperturbed velocity term U_∞ remains constant. In contrast to the static BEM model, the differential of the velocity potential in time is nonzero and can be obtained by a standard, second-order, backwards-difference approximation across time-steps (Trefethen, 1996),

$$\dot{\phi}_t \approx \frac{3\phi_t - 4\phi_{t-1} - \phi_{t-2}}{2\Delta t}. \tag{4.5}$$

The response of the fluid to its forced coupling can be included in the structural

model through the surface forcing term ($F(x,t)$),

$$F(x,t) = -\Delta p\left(\ddot{\eta}, \dot{\eta}, \eta\right). \tag{4.6}$$

Rewriting the full potential-flow (over an elastic panel) FSI system of equations, substituting Equations (4.4) and (4.6) into Equation (3.1) we have,

$$\rho h \frac{\partial^2 \eta}{\partial t^2} + B\frac{\partial^4 \eta}{\partial x^4} - T_I\left(\eta\right)\frac{\partial^2 \eta}{\partial x^2} = -\rho_f\left(-\frac{U_\infty^2}{2} + \frac{\partial \phi}{\partial t} + \frac{U^2+V^2}{2}\right). \tag{4.7}$$

The matrix-free NK and FDM have been formulated previously to allow the inclusion of any forcing function. The inclusion of the fluid force does not adversely impact the efficiency of this solution framework in exploiting the matrix-free nature of the FMM and GMRES.

4.3 System Solution Method

In Chapter 3 we implicitly solved the *in-vacuo* structural system in the form $f(\ddot{\eta}) = 0$. We could therefore use linear algebra to shift the additional fluid based forcing term ($F(x,t)$) onto the left-hand side (LHS) of Equation (4.7) to allow a fully implicit solution via the NK. However, with the desire to keep the solution method as flexible and generic as possible (to allow later extensions) we elect to utilise a hybrid method. Under the hybrid method, the LHS of the system equation is to be solved implicitly for $\ddot{\eta}$ using the NK, while the right-hand side (RHS) forcing function is to be coupled semi-implicitly and solved iteratively. This hybrid method is achieved numerically by decoupling the LHS and RHS equations into a semi-implicit (SI) predictor-corrector scheme. The forcing function (RHS) is solved as the explicit predictor, Δp_{t*}, and system acceleration (LHS) is solved

implicitly (using the NK) according to[1]

$$\underbrace{\left[\rho h\frac{\partial^2\eta}{\partial t^2}+B\frac{\partial^4\eta}{\partial x^4}-T_I\left(\eta\right)\frac{\partial^2\eta}{\partial x^2}+\rho_f\frac{U_\infty^2}{2}-P^*\right]_t}_{\text{Implicit NK}}=\underbrace{-\left[\rho_f\left(\frac{\partial\phi}{\partial t}+\frac{U^2+V^2}{2}\right)-P^*\right]_{t^*}}_{\text{Explicit}}. \tag{4.8}$$

Convergence of the SI iterations is dependent on the strength of the fluid response to changes in the structural motion. An additional P^* term has been included on both sides of the equation to act as a 'conditioner' on the SI iterations. This term can be used to shift the fluid pressure components that are sensitive to the structural changes inside the implicit NK scheme - the selection of such components will be discussed further below. This leaves the relatively minor and steady pressure terms, that are indifferent to rapid structural changes, to be handled by the SI scheme. The role of this term is therefore to determine where the burden of numerical convergence is distributed between the NK and the SI schemes. Its flexible nature is particularly advantageous where the FMM process is used to determine the flow pressure. This is because the FMM-calculated pressure components cannot be algebraically reduced (without major reformulation) to its corresponding $\ddot{\eta}, \dot{\eta}$ and η components. However, without the FMM one could hand pick the components critical to SI convergence, such as any terms involving beam acceleration (e.g. the fluid inertia).

Finding the optimum choice for P^* is not a trivial process due to its system dependence. Selection must be conducted in a manner similar to that of pre-conditioner selection for the BEM and NK; to find a balance between formation/application-cost and net improvement in solution time and stability. While there are many options available, we focus on presenting three distinct cases, each with pros and cons tailored to suit different FSI systems.

[1]The notation $[...]_{t^*}$ and $[...]_t$ represent the predictor and corrector components respectively that are solved using a SI scheme

4.3.1 Explicit Coupling

When looking at the computational complexity of the FSI system we note that most of the fluid pressure terms involve dense square ($N \times N$) matrices, ones which are never explicitly formed when using the GMRES/FMM combination. Dynamic evaluation (a necessity as they depend on $\ddot{\eta}$, $\dot{\eta}$ and η) of these terms inside the implicit NK solution process for many cases would be expensive and therefore undesirable. This is because every internal call to the system equation (f) inside the NK would require an additional $O(N^2)$ component, whereas, all other structural components involve static, banded matrices (excluding the induced tension) and as a result are of computational complexity $O(KN)$. Including the full nonlinear fluid pressure would also cause NK convergence rate to decrease over the previous *in-vacuo* cases. With this consideration we present a scheme that involves no implicit fluid coupling,

$$P^* = P_1^* = 0. \tag{4.9}$$

Substituting this back into the system equation we have,

$$\underbrace{\left[\rho h\frac{\partial^2\eta}{\partial t^2} + B\frac{\partial^4\eta}{\partial x^4} - T_I\left(\eta\right)\frac{\partial^2\eta}{\partial x^2} + \rho_f\frac{U_\infty^2}{2}\right]_t}_{\text{Implicit NK}} = \underbrace{-\left[\rho_f\left(\frac{\partial\phi}{\partial t} + \frac{U^2+V^2}{2}\right)\right]_{t^*}}_{\text{Explicit}}. \tag{4.10}$$

By allowing no implicit coupling of the fluid pressure it offers the advantage of being computationally simple from an implementation point of view. As the NK is essentially solving a statically-loaded *in-vacuo* beam[2] the NK pre-conditioner also remains unchanged. The fundamental difference with the purely *in-vacuo* case is that the structure now must be re-solved many times using different (effectively static) surface loads.

[2]As full $O(N^2)$ fluid calculations are restricted to each SI iteration.

4.3.2 Linearised Inertia Coupling

While no implicit fluid coupling allows the NK scheme to converge at the same rate as the *in-vacuo* case, it does so at the cost of shifting the burden of system convergence onto the sensitive SI scheme. Although this may offer faster overall convergence due to the simplicity of the coupling, this only works for a subset of problems where reducing time-step sizes can assist in ensuring SI convergence. Unlike the velocity components of the fluid pressure, the inertia component ($\partial\phi/\partial t$) does not get scaled by time-step sizes. Stability of the SI scheme for these systems can therefore be dependent on the system-specific density ratio of the fluid and wall, ρ_f/ρ_w. As this ratio grows larger in magnitude, the SI scheme will have more difficulty in converging on a solution because the motion-sensitive fluid inertia term also becomes large.

In an attempt to overcome inertial (time-step independent) instabilities in the SI scheme, we look to developing an inertial based implicit coupling which can maintain a low evaluation complexity for inclusion into the NK. Direct inclusion of the fluid inertia ($\partial\phi/\partial t$) term into the NK would drastically slow the speed of convergence as it represents the equivalent of a time-dependent dense square ($N\times N$) matrix operation. In addition to this, algebraically reducing the nonlinear inertial pressure into its strong and weak components, such as the approach by Lucey et al. (1997b), is not possible as the calculation is conducted in a matrix-free manner using the FMM. An alternative is to develop a linearised version of the inertia that can be coupled implicitly and allow the remaining nonlinear components of the inertial term to be coupled semi-implicitly.

We begin with the exact differential for $\partial\phi/\partial t$ using the chain rule,

$$\frac{\mathrm{d}\{\phi\}}{\mathrm{d}t} = \frac{\mathrm{d}\left([I_\phi]\{\sigma\}\right)}{\mathrm{d}t}, \tag{4.11}$$

$$= [\dot{I_\phi}]\{\sigma\} + [I_\phi]\{\dot{\sigma}\}. \tag{4.12}$$

Assuming linearised motion of our system, with $\sin\theta_p \approx \theta_p$ and $\cos\theta_p \approx 1$ and recalling that σ is given by the no-flux boundary condition applied at the BEM

panel centres, gives

$$\{\dot{\sigma}\} = \frac{\mathrm{d}\left([I_v]^{-1}(\{\dot{\eta_p}\cos\theta_p\} + \{U_\infty \sin\theta_p\})\right)}{\mathrm{d}t}, \tag{4.13}$$

$$\approx \frac{\mathrm{d}\left([I_v]^{-1}\{\dot{\eta_p}\}\right)}{\mathrm{d}t} + \{U_\infty \dot{\theta_p}\}, \tag{4.14}$$

$$\approx \dot{[I_v]^{-1}}\{\dot{\eta_p}\} + [I_v]^{-1}\{\ddot{\eta_p}\} + \{U_\infty \dot{\theta_p}\}. \tag{4.15}$$

Linearisation also implies,

$$[I_v]^{-1} \approx 2[I], \quad [\dot{I_\phi}] \approx [\dot{I_v}] \approx 0, \tag{4.16}$$

where $[I]$ represents the Identity matrix. Substituting Equations (4.15) and (4.16) back into Equation (4.12) we have,

$$\dot{\phi} \approx [I_\phi]\left(\{2\ddot{\eta_p}\} + \{U_\infty \dot{\theta_p}\}\right). \tag{4.17}$$

Thus, to implicitly couple only the linear component of the fluid inertia pressure we use,

$$P^* = P_2^* = \rho_f [I_\phi]\left\{2\ddot{\eta_p} + U_\infty \dot{\theta_p}\right\}, \tag{4.18}$$

and by applying this to the full system equation we have,

$$\underbrace{\left[\rho h \frac{\partial^2 \eta}{\partial t^2} + B\frac{\partial^4 \eta}{\partial x^4} - T_I(\eta)\frac{\partial^2 \eta}{\partial x^2} + \rho_f\left(\frac{U_\infty^2}{2} - [I_\phi]\left\{2\ddot{\eta_p} + U_\infty\dot{\theta_p}\right\}\right)\right]_t}_{\text{Implicit NK}} = \tag{4.19}$$

$$\underbrace{-\left[\rho_f\left(\frac{\partial \phi}{\partial t} + \frac{U^2+V^2}{2} - [I_\phi]\left\{2\ddot{\eta_p} + U_\infty\dot{\theta_p}\right\}\right)\right]_{t^*}}_{\text{Explicit}}. \tag{4.20}$$

The addition of this coupling allows the remaining pressure terms in the SI scheme to become indifferent to changes in the wall acceleration. While this results in a faster SI convergence and higher numerical stability, it does so at the cost of a single FMM call ($[I_\phi]$ evaluation) for every call of f in the NK. However,

this single FMM call is still at an advantage to a full nonlinear flow coupling due to the linear approximation for panel strengths. Where full nonlinear coupling is required, the panel strengths must be solved using the GMRES which requires multiple calls to the FMM.

The linearising assumption ensures $[I_\phi]$ is time-independent and is therefore suitable for caching. Explicit caching is not possible in this work because, when using the FMM, this matrix is never explicitly formed. However, the FMM quad tree could be cached with the panels in the unperturbed position, adjusting only element strengths with every call.

As the use of this implicit coupling results in a change in the implicit system f, the NK pre-conditioner must also be modified. The highly diagonal nature of this new linearised term complements the existing methods for obtaining the sparse approximation to the Jacobian inverse. The additional benefit of this scheme is that the linear nature of the coupling ensures the Jacobian pre-conditioner does not need to change with time and can continue to be cached as previously.

4.3.3 Fully-Implicit Coupling

The final option is to solve the fully coupled nonlinear potential-flow system in essentially one sweep of the NK using,

$$P^* = P_3^* = \rho_f \left(\frac{\partial \phi}{\partial t} + \frac{U^2 + V^2}{2} \right). \tag{4.21}$$

Rewriting the system equation we have,

$$\underbrace{\rho h \frac{\partial^2 \eta}{\partial t^2} + B \frac{\partial^4 \eta}{\partial x^4} - T_I(\eta) \frac{\partial^2 \eta}{\partial x^2} + \frac{\rho_f}{2} \left(U_\infty^2 - U^2 - V^2 - 2 \frac{\partial \phi}{\partial t} \right)}_{\text{Implicit NK}} = 0. \tag{4.22}$$

Here the time differential in ϕ is solved numerically and the BEM panel strength (σ) is solved via the full GMRES/FMM combination to satisfy the no-flux boundary condition. The advantage of this method is that the RHS of the system equation reduces to zero with the full nonlinear fluid pressure being solved on the

LHS. With the burden of solution shifted entirely to the implicit NK method, the need for SI iterations and repeat NK evaluations are eliminated. The disadvantage of the scheme is that system function calls inside the NK are even more complicated over the linearised scheme, requiring many calls to the FMM to solve the strengths using the GMRES. The modification to the inverse Jacobian pre-conditioner for the NK remains similar to that for the linear version, with fluid pressure terms remaining diagonally dominant despite now representing a dense matrix. Consequently this coupling scheme responds well to the existing pre-conditioning methods. Although the nonlinear fluid terms are now time dependent, the practice of pre-conditioner caching still offers significant efficiency improvements.

4.3.4 Coupling Complexity Analysis

We now summarise all three coupling schemes and illustrate the different computational complexity resulting from each by first recalling,

$$P_1^* = 0, \tag{4.23}$$

$$P_2^* = \rho_f [I_\phi] \left\{ 2\ddot{\eta}_p + U_\infty \dot{\theta}_p \right\}, \tag{4.24}$$

$$P_3^* = \rho_f \left(\frac{\partial \phi}{\partial t} + \frac{U^2 + V^2}{2} \right). \tag{4.25}$$

For comparison and later validation we also investigate using the P_1^* coupling with a purely SI method to solve the LHS and RHS of the system equation. By taking into account semi-implicit convergence iterations (S), Newton iterations (N), Krylov iterations (K), GMRES calls (G) and FMM calls (F) we arrive at Table 4.1.

Table 4.1: Number of operations for different coupling P^*

Method	Left eval	Right eval	Total
$SI - P_1^*$	S_{SI}	S_{SI}(G(F) + 2(F))	S_{SI}(1 + F(G+2))
P_1^*	$S_1(N_1(K_1))$	S_1(G(F) + 2(F))	$S_1(N_1K_1$ + F(G+2))
P_2^*	$S_2(N_2(K_2(F)))$	S_2(G(F) + 3(F))	S_2(F(N_2K_2 + G + 3))
P_3^*	$N_3(K_3($ G(F) + 2F))	0	$N_3(K_3$(F(G + 2)))

This table shows that for each case the complexity of G and F remain constant as they are functions of the numbers of structural nodes (n). However S varies vastly ($S_{SI} >> S_1 > S_2$) depending on the conditioning of the RHS equation, with K increasing ($K_1 < K_2 < K_3$) for greater functional (f) matrix density and N increasing ($N_1 < N_2 < N_3$) for greater nonlinearity and reducing accuracy of K.

When using P_1^* we offer no implicit coupling between the fluid and the wall. While this method under many circumstances will be the most efficient, it cannot be used effectively for all FSI systems. As the magnitude of the unsteady pressure term is governed mainly by the fluid density, not time-step size, the convergence behaviour of the P_1^* coupled system is reliant on the fluid-wall density ratio, ρ_f/ρ_w. When this is relatively high, such as flows involving water or where the ratio is greater than approximately 10^{-2}, this method cannot converge on a solution due to the strong acceleration sensitive inertia forces. For low ratios, however, this inertial effect is weak and allows an increase in S to be traded for simple NK calculations without the loss of stability.

When higher density ratios are required, the linearised fluid inertia coupling (P_2^*) must be used at a minimum for numerical stability and offers a good trade-off between the cost of the NK and the number of S calculations. This method excels in cases where wall motion remains dominantly linear and where inertia forces drive the system dynamics. For cases where S becomes too large, such as for highly nonlinear behaviour or where fluid elements (other than potential-flow) are used that are vastly more expensive to calculate (such as with the DVM), fully-implicit nonlinear fluid coupling (P_3^*) would be preferred. This would be principally aimed at reducing the number of S iteration loops but would only apply where a full potential-flow coupling offered a reasonable approximation to the real flow.

4.3.5 Illustrative Results – Optimum Coupling Method

We now demonstrate the difference between the various fluid-coupling methods in conjunction with the structural solution methodologies. To ensure numerical stability of all coupling methods, the density ratio of the fluid and wall is set at $\rho_f/\rho_w = 1 \times 10^{-3}$, akin to that of air flow over an elastic rubber-type wall. This ensures that inertia effects of the fluid are relatively small to allow the use of the SI coupling scheme, P_1^*. For our FSI model we define the dimensionless flow speed as,

$$\Lambda = \frac{\rho_f {U_\infty}^2 L^3}{B}, \tag{4.26}$$

where all variables retain their previous definitions. For this case, flow speed is set at $\Lambda = 200$ to ensure the onset of a nonlinearly saturated divergence instability. In addition we select a wall thickness ratio of $h/L = 1 \times 10^{-2}$ and no spring/damper foundation. Figures 4.3 and 4.4 illustrate historical snap-shots of the growth and decay respectively of a single nonlinear oscillation using the hybrid NK-SI scheme with a fully-implicit nonlinear fluid pressure coupling (P_3^*) and a FDM discretisation of $N = 128$. While it is clear that the first vibration mode is excited in this case, the flow speed is sufficient to also introduce excitations of vibration mode two. For comparison purposes we define the non-dimensional time as,

$$t' = \frac{\sqrt{\frac{E^*}{\rho_w}}}{L} t. \tag{4.27}$$

A trace of the vertical position of the central wall node in non-dimensional time for the first three nonlinear oscillations is shown in Figure 4.5. This illustrates the difference between the two structural solution methodologies; NK-SI-hybrid and purely-SI, and the three fluid-coupling schemes; P_1^*, P_2^* and P_3^*. The system indicates energy-conservative nonlinear oscillatory behaviour through the peaks and troughs in Figure 4.5 remaining consistent as time increases. Confidence is gained in the implementation of each numerical scheme due to the absolute

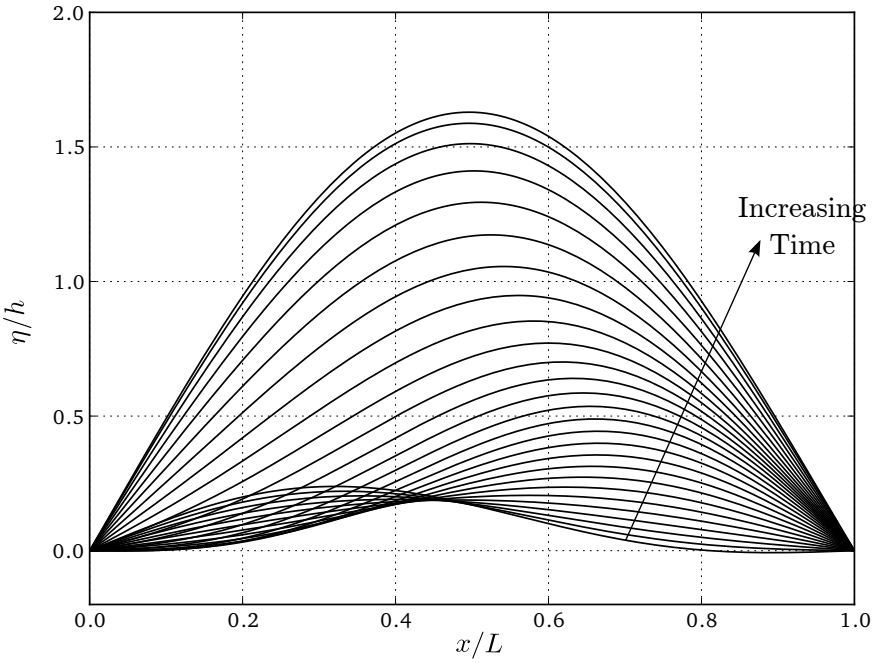

Figure 4.3: Historical plot of wall position for a single growth cycle of a nonlinear divergence instability of air flow ($\Lambda = 200$) over rubber ($\rho_f/\rho_w = 1 \times 10^{-3}$).

agreement in results for all methods using identical system discretisation and properties.

Figure 4.6 shows the total computational time required for a full three-oscillation simulation using all numerical methods and varying only spatial (FDM) discretisations. Computational time has been non-dimensaionalised to the run-time for the $SI - P_1^*$ method using 8 nodes, at $T_{C,N=8} = 18.7$s, measured on a standard desk-top computer (Intel Q9650 3.0GHz processor). Data for $N = 128$ using the $SI - P_1^*$ scheme was estimated based on previous scaling rates because running this experiment would have been infeasible due to its poor computational scaling. Figure 4.7 shows a zoomed scale of Figure 4.6 to illustrate the difference between the fluid-load coupling methods under the NK scheme. While it is seen that the $SI - P_1^*$ method offers the fastest computation time for the coarsest spatial discretiations, the NK demonstrates vastly superior scaling by outperforming from even mildly increased discretiations. The brief advantage of the $SI - P_1^*$ can be

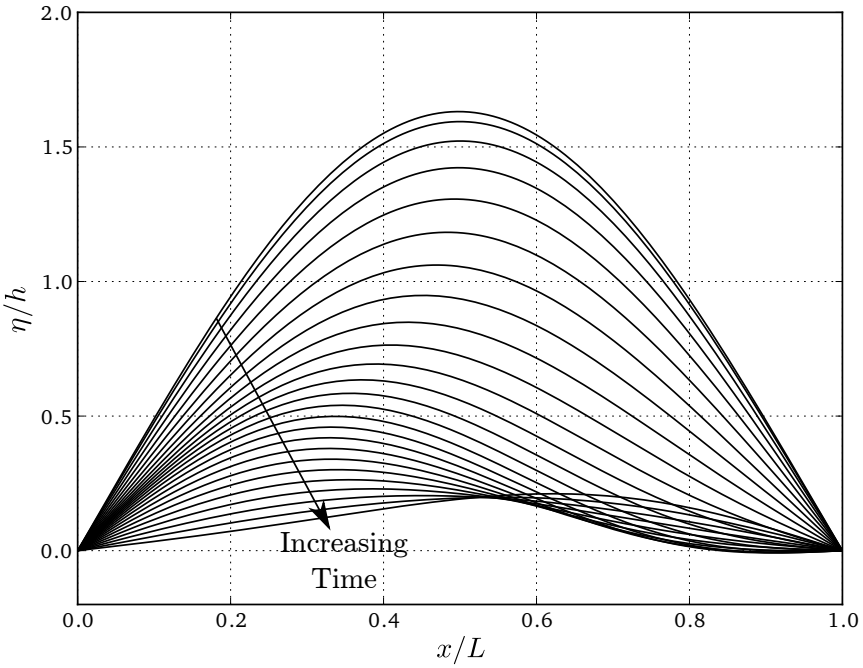

Figure 4.4: Historical plot of wall position for a single decay cycle of a nonlinear divergence instability of air flow ($\Lambda = 200$) over rubber ($\rho_f/\rho_w = 1 \times 10^{-3}$).

attributed to the inherent overhead that must be overcome in the NK method. However, it is noted that time-step sizes for the NK have been matched to suit the stability requirements of the sensitive $SI - P_1^*$ method at $N = 8$. In Figure 4.8 we compare the relative time-step sizes that were required to maintain numerical stability to investigate the cause of the exceptionally poor scaling of the $SI - P_1^*$ method (as seen in Figure 4.6). It is seen that the poor scaling of the pure SI method is due to the need to decrease time discretisation at the rate of the square of the spatial discretisation, i.e.,

$$\Delta t_1 \approx \left(\frac{\Delta x_1}{\Delta x_0}\right)^2 \Delta t_0. \tag{4.28}$$

This refinement is required to ensure stability of the numerical structural solution (Richtmyer and Morton, 1967). While each iteration of the SI method may be completed in a similar time to the NK-SI hybrid methods, the higher total number of iterations for a given simulation is what causes poor scaling. Under these SI

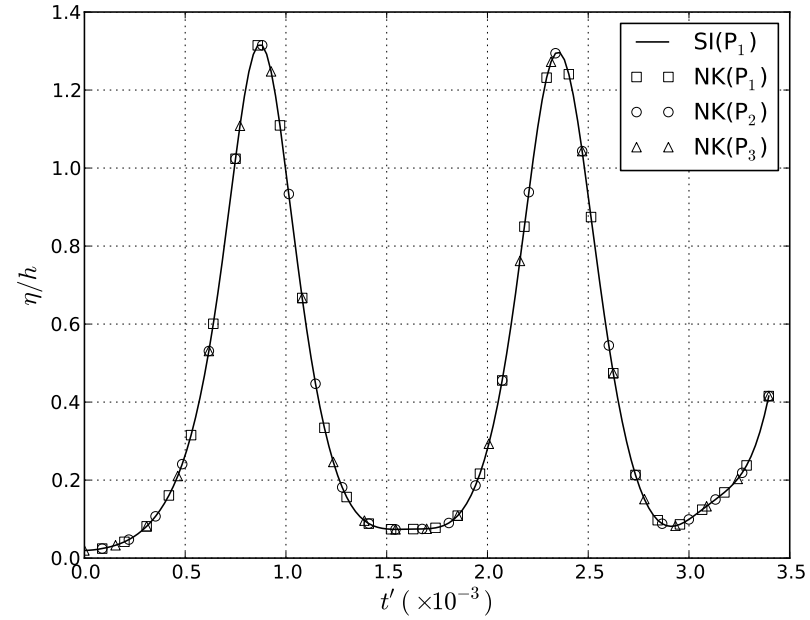

Figure 4.5: Simulated vertical displacement in time of the midpoint of a wall engaged in a nonlinear divergence instability using various numerical schemes and coupling methods.

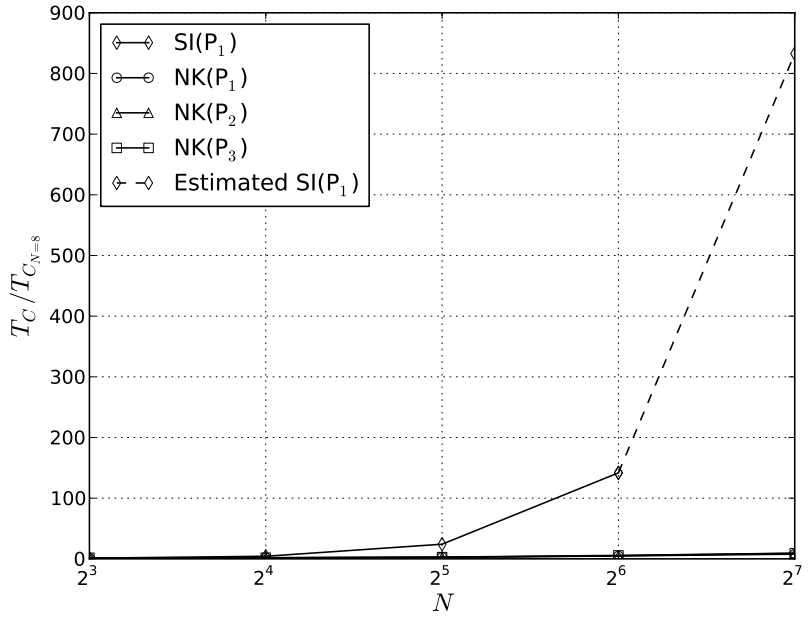

Figure 4.6: Computational run-time required to complete three nonlinear oscillations of a divergence instability for varying spatial discretisations and numerical schemes. Note a detailed view of this plot follows in Figure 4.7.

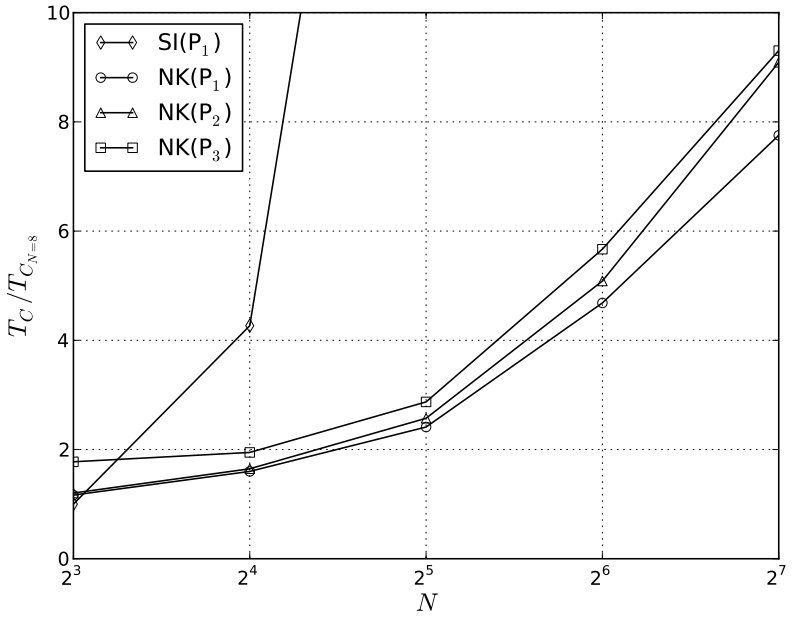

Figure 4.7: Zoomed data of Figure 4.6 emphasising the computational run-time required for three oscillations of a nonlinear divergence instability when varying spatial discretisations and fluid-coupling methods under the NK scheme.

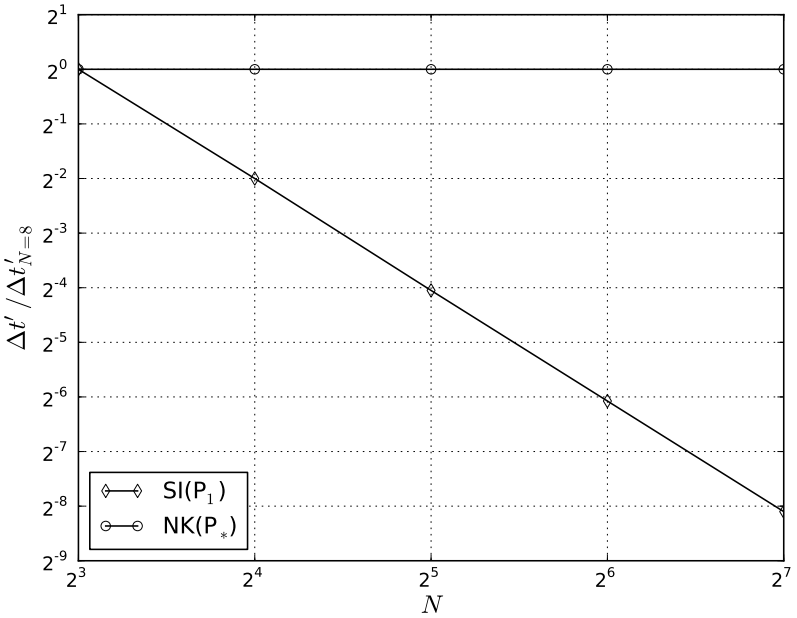

Figure 4.8: Minimum time-step size required to maintain numerical stability when obtaining data for Figure 4.6.

vs NK comparisons the time-steps for the NK methods were based on achieving the same minimum level of discretisation as the SI method. With the NK scheme being an implicit method, it is unconditionally stable and allows the use of any time-step size that achieves the desired solution accuracy. For the present system, however, numerical dissipation occurs when extremely course time discretisations are used as a result of the backwards differencing of the unsteady pressure ($\partial\phi/\partial t$) term.

When comparing fluid-coupling methods in Figure 4.7, P_1^* is the most time efficient for all discretisations with P_2^* and P_3^* trailing respectively. While this may indicate that P_1^* should always be the preferred method, this is only for a select set of density ratios of approximately $\rho_f/\rho_w < 10^{-2}$. For ratios higher than this the P_1^* scheme fails to converge. Based on the presented results the NK offers the most flexibility and superior performance to SI for nonlinear FSI. The ideal choice of coupling method for all scenarios cannot be deduced from this test case as the chosen parameters advantaged the P_1^* and P_2^* schemes in a manner not replicated in following sections. For the remainder of this work the P_3^* will be used due to the flexibility and stability offered by its fully-implicit nonlinear flow coupling.

4.4 Illustrative Results – Divergence Onset

One measure of a numerical method's ability to model FSI correctly is through capturing the divergence onset flow speed, the point where the fluid force exceeds the restorative force in the wall. Previous works such as Pitman and Lucey (2009) used a linear state-space, eigen-analysis approach to determine divergence-onset speeds. The nonlinear deformation amplitudes of the model used in this work prevents the use of the eigen-analysis approach, particularly where the FSI is in the presence of more complicated and noisy fluid flows such as those involving a boundary-layer. As discussed in Balint and Lucey (2005), divergence onset is characterised largely by two main features; i) the system gains energy due to the action of the fluid flow and ii) the energy growth occurs in a quasi-static manner.

We can observe the rate of change in wall kinetic and strain energy using the equations,

$$0 < \frac{\partial E}{\partial t}, \tag{4.29}$$

$$< \frac{\partial}{\partial t}\left(\frac{1}{2}B\int_0^L \left(\frac{\partial^2 \eta}{\partial x^2}\right)^2 \mathrm{d}x + \frac{1}{2}\rho h \int_0^L \left(\frac{\partial \eta}{\partial t}\right)^2 \mathrm{d}x\right). \tag{4.30}$$

However, energy growth alone does not fully describe divergence onset due to possible energy transfers from phase shifts in the surface motion and the interfacial pressure. To confirm divergence onset we also check that energy growth is due to the total fluid force at the surface being greater than the restorative force in the wall by,

$$F_f > F_w, \tag{4.31}$$

$$\int_0^L (-\Delta p)\mathrm{d}x > B\int_0^L \frac{\partial^4 \eta}{\partial x^4}\mathrm{d}x. \tag{4.32}$$

With the knowledge that any FSI system will consist of a flow speed that is either above, equal or below the critical divergence-onset flow speed, we define a suitable algorithm for iteratively determining the critical speed for any time-dependent system, linear or nonlinear.

1. Start with any initial flow speed with a nonzero wall displacement, eg. small-amplitude vibration mode one.

2. Commence the time-dependent simulation, averaging (to account for noise and initial disturbance waves) over many cycles the energy growth ($\partial E/\partial t$) and relative force amplitude (F_f/F_w). If both energy rate and force ratio are below zero and one respectively, then flow speed must be below the critical speed. If both are positive and/or growing then the flow speed is post-critical. This is recorded as either a lower bounding (pre-critical) or upper-bounding (post-critical) speed.

3. If either the upper or lower bounds are yet to be found, a new guesses is

attempted until the opposing bound is found.

4. With divergence onset now bounded by an upper and lower flow speed, we apply a binary search method to find the future trial speeds. By recording the final energy-growth rate and force ratio for the previous trials, we can use a Newton-Raphson (NR) approach (to find $\partial E/\partial t \approx 0$ and $F_f/F_w \approx 1$) to augment and improve the binary-search solution.

5. With the new trial speed, Step 2 is repeated and followed by Step 4 until the trial speeds converge to the desired precision.

Where the FSI systems are highly unstable, modifications can be made such as over-damping the wall and removing the effects of nonlinear tension. These modifications will not effect the converged critical speed because divergence is a quasi-static behaviour that exists in linear ranges of motion.

We now apply this divergence finding algorithm to the potential-flow FSI system using the NK and fully-implicit flow coupling (P_3^*). Figure 4.9 represents the convergence process of the trial solution for each iteration of the algorithm. It is seen that the onset speed is found approximately after the first few trial iterations, a result of using the NR guess augmentation process. However, as the trial speed approaches that of the divergence onset, the NR guess must be discarded because the averaged nature of the energy growth ($\partial E/\partial t$) and relative force amplitudes do not carry enough accuracy to allow meaningful gradient approximations. From this point the time-average interval is adaptively adjusted to gain greater confidence in the pre/post divergence decision (Step 2 of algorithm), with the binary search method solely used to reach the desired precision. For this test case, the algorithm converges on a flow speed of $\Lambda_{\text{div}} = 39.6$ (with a tolerance of %1), which offers strong agreement to Lucey et al. (1997b) and Pitman and Lucey (2009). This reveals that the developed algorithm can accurately predict the divergence-onset flow speed in an efficient and reliable manner than can be used in studying FSI systems involving (noisy) viscous boundary-layer flows.

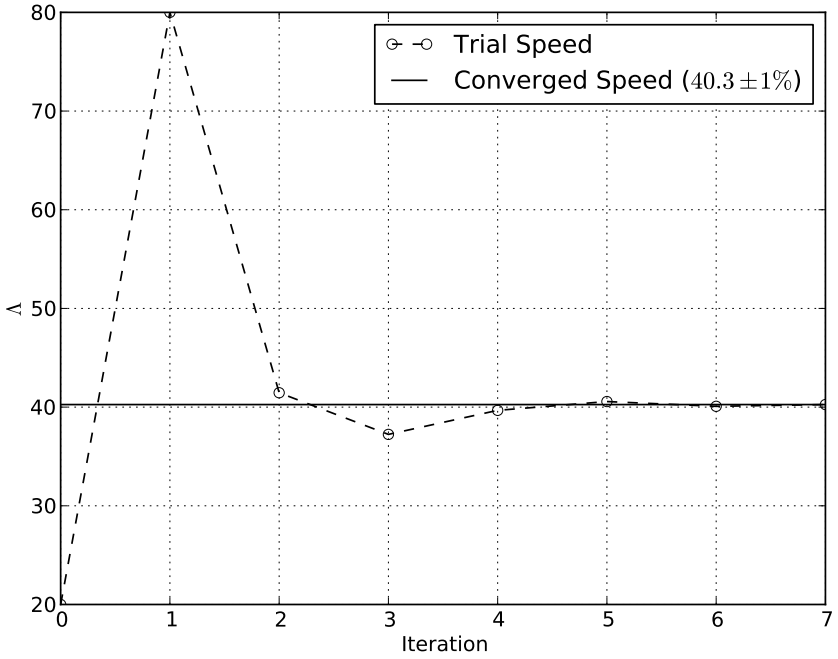

Figure 4.9: Intermediate trial and final converged solutions of the divergence onset finding algorithm for a potential-flow FSI system.

4.5 Illustrative Results – Nonlinear Divergence Oscillations

With the divergence-onset flow speed being captured correctly we now attempt to capture energy-stable nonlinear divergence oscillations in the FSI model. This behaviour arises from the wall amplitude growth due to the post-divergence flow speed and consequent energy transfer from the fluid. The wall amplitude continues to grow until the nonlinear tension provides sufficient force to overcome the energising action of the fluid, at which point it causes a 'snap-back' towards the wall's neutral position. By simulating the case of $\rho_f/\rho_w = 3.85\times10^{-1}$, akin to water flow over a thin aluminium wall, with a wall thickness ratio of $h/L = 1\times10^{-2}$, a flow speed of $\Lambda = 61$ and no wall backing, we can compare to the work of Lucey et al. (1997b). A system with these properties requires the fully-implicit, nonlinear potential-flow coupling (P_3^*) to be used. A relatively high density ratio such

as this would not allow the use of P_1^* method due to numerical instability caused by the significant fluid inertia forces. While the P_2^* scheme could have been used from a stability perspective, the nonlinear dynamics of the system would result in inferior performance when compared to P_3^*.

Figures 4.10 and 4.11 show the historical snapshots of the system undergoing a single nonlinear oscillation. Figure 4.12 shows the trace of the vertical posi-

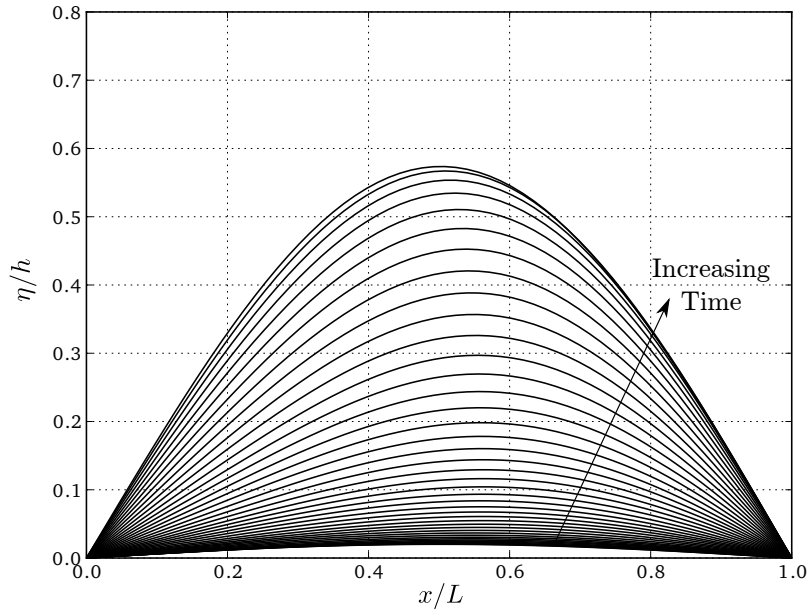

Figure 4.10: Historical plot of wall position for a single growth cycle of a nonlinear divergence instability of water flow ($\Lambda = 61$) over Aluminium ($\rho_f/\rho_w = 3.85 \times 10^{-1}$).

tion in time of the central wall node, illustrating good amplitude and frequency agreement with Figure 4.13 that is obtained from Lucey et al. (1997b). A small shift in frequency and amplitude can be attributed to the coarse discretisation, of $N = 10$ wall nodes, used in Lucey et al. (1997b). Such a low level of FDM nodes does not produce discretisation-independent results, necessitating the use of the finer nodal resolution in this work of $N = 64$.

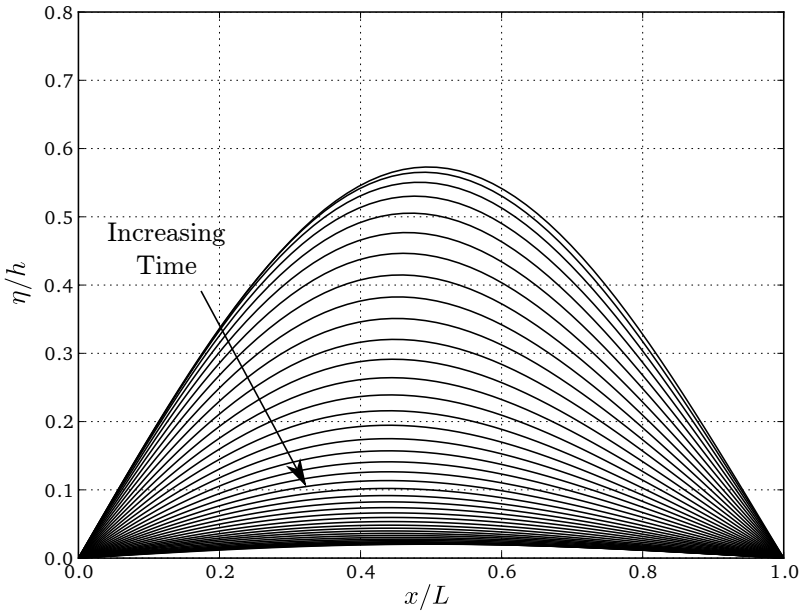

Figure 4.11: Historical plot of wall position for a single decay cycle of a nonlinear divergence instability of water flow ($\Lambda = 61$) over Aluminium ($\rho_f/\rho_w = 3.85 \times 10^{-1}$).

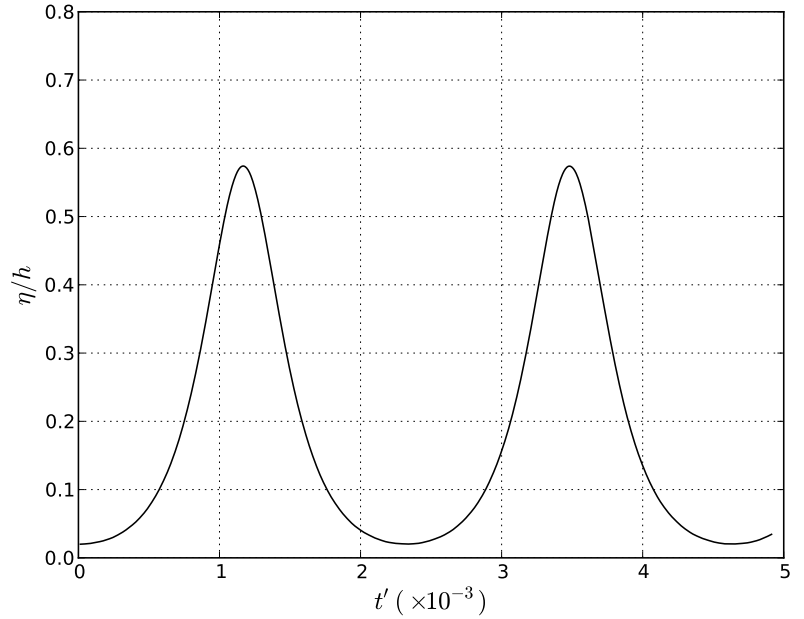

Figure 4.12: Simulated vertical displacement in time of the midpoint of a wall engaged in a nonlinear divergence instability using the fully-implicit NK method for $\Lambda = 61$, $\rho_f/\rho_w = 3.85 \times 10^{-1}$ and $h/L = 1 \times 10^{-2}$.

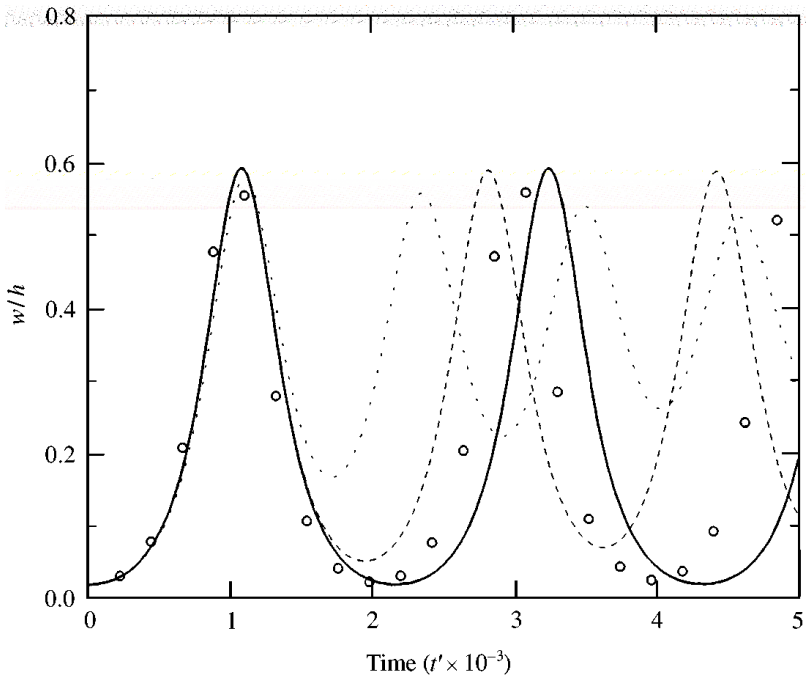

Figure 4.13: Figure 2 reproduced from Lucey et al. (1997b) showing (solid line) the vertical displacement in time of the midpoint of a wall engaged in a nonlinear divergence instability for $\Lambda = 61$, $\rho_f/\rho_w = 3.85 \times 10^{-1}$ and $h/L = 1 \times 10^{-2}$.

4.6 Summary

An efficient and numerically stable method for the Fluid-Structure Interaction (FSI) of infinite Reynolds-number (potential) flows over a flexible surface was developed. The Boundary Element Method (BEM) was used to model the effects of the deforming boundary while the fluid response was coupled to the structure using the unsteady-Bernoulli equation. The Finite Difference Method (FDM) was used to solve the structural dynamics based on the nonlinear Euler-Bernoulli beam model.

Developing an efficient and stable solution method for the FSI system required a flexible iterative scheme that involved a mixture of implicit solution (using the Newton-Krylov (NK) method) and a semi-implicit (SI) coupling of the selected fluid pressure terms. A demonstration case of three distinct fluid-coupling methods showed that the fully-implicit pressure coupling was the most computationally intensive. However, it was chosen as the method for the remain-

der of this work due to its unconditional stability and suitability to later models that include boundary-layer effects. Overall, the hybrid NK-SI method demonstrated vastly improved scaling over the previously established SI method. This property permitted the finer simulation resolution that was required to overcome the discretisation dependence observed in prior works.

A suitable algorithm for finding the divergence-onset flow speed for unsteady nonlinear FSI systems was also demonstrated. This algorithm was designed with sufficient flexibility to allow its use with later FSI systems that include the complex effects of a viscous boundary layer.

Chapter 5

Rotational-Flow Modelling

The fluid-flow models discussed thus far have been based on inviscid and irrotational (potential-flow) assumptions that correspond to an infinite Reynolds-number flow. Modelling with these assumptions captures the behaviour of the bulk flow but neglects viscous effects that cause boundary layers to form over surfaces. For laminar and transitional Reynolds-number flows the boundary-layer effect cannot be neglected and must be captured.

We introduce the Discrete Vortex Method (DVM) that can be used to model the effects of viscosity and rotationality in fluid flows at moderate (i.e. transitional) Reynolds numbers. The initial models and methods developed will focus on systems characterised largely by inviscid-flow behaviour. This allows the dissipative effects of viscous diffusion to be neglected, leaving its inclusion a topic to be discussed in Chapter 6.

For a general introduction to DVM modelling and its founding literature, see Section 1.4.2. However, for an excellent background text to the mathematical foundations of the DVM see Lewis (1991). For inviscid, incompressible, unforced two-dimensional flow, the formulation of the DVM begins with Euler's equation,

$$\dot{\vec{u}} + (\vec{u} \cdot \nabla)\, \vec{u} = -\frac{1}{\rho_f} \nabla p, \tag{5.1}$$

where $\vec{u} = U\mathbf{i} + V\mathbf{j}$ is the velocity vector and p is the scalar pressure. This is

coupled with the two-dimensional mass conservation equation,

$$\nabla \cdot \vec{u} = 0. \tag{5.2}$$

By taking the curl ($\nabla\times$) of Equations (5.1) and (5.2) we are left with,

$$\frac{\partial \omega}{\partial t} + (\vec{u} \cdot \nabla)\, \omega = 0, \tag{5.3}$$

where $\omega = \dfrac{\partial U}{\partial y} - \dfrac{\partial V}{\partial x}$ is the vorticity (clock-wise positive). The advantage of this formulation is that the system is decoupled from the flow pressure and mass conservations is conserved implicitly. As we wish to avoid the downside of using grid-based Eulerian schemes, we can track the vorticity (ω) in a Lagrangian reference frame resulting in,

$$\frac{D\omega}{Dt} = 0. \tag{5.4}$$

This shows that vorticity is always conserved along particle trajectories, allowing the system to be defined numerically by a series of discrete packets of constant vorticity that follow the streamlines of the flow field. To model the particle trajectories we use the first-order accurate Euler time-stepping method,

$$\vec{r}_{t+1} = \vec{r}_t + \vec{u}_t \Delta t, \tag{5.5}$$

where $\vec{r}_t$ represents the position vector to the particle's centre at the current time-step, Δt is the time-step size and $\vec{u}_t$ is the velocity evaluated at the particle's centre. Higher order time-stepping schemes such as the second-order accurate Adams-Bashforth scheme or second-order accurate Runge-Kutta schemes were investigated but deemed unsuitable. This is because their multi-step nature is not easily compatible with the iterative framework used for the fluid-structure interaction system.

To solve for the velocity field of the flow we introduce the 2D point vortex element whose flow pattern is similar to that of a solid cylinder (of zero radius)

rotating in a viscous fluid. This element has a radial velocity component of zero with streamlines that form concentric circles around its core. This flow pattern is described mathematically by a velocity potential (ϕ) and stream function (ψ) of,

$$\phi(\mathbf{r}) = \frac{\Gamma}{2\pi}\theta, \tag{5.6}$$

$$\psi(\mathbf{r}) = \frac{\Gamma}{2\pi}\ln|\mathbf{r}|, \tag{5.7}$$

where Γ is the vortex circulation strength, $\mathbf{r}$ is the relative position vector to the evaluation point and θ is the relative angle to the evaluation point. These equations represent the effect of a concentration of vorticity that is distributed in an infinitely small space, an element that causes velocity to tend towards infinity at its core. From a numerical point of view, this causes problems in simulations where two particles become close together as a result of coarse time-stepping. This behaviour is also not physically representative because viscous effects would be significant at the vortex core. As an alternative we introduce a vortex particle that consists of a Gaussian vortex distribution that is an exact solution to the Navier-Stokes equations for a single vortex in an unbounded incompressible domain. The vorticity distribution for this particle is defined by,

$$\omega(\mathbf{r}) = \frac{\Gamma}{\pi\sigma^2}\exp\left(\frac{-|\mathbf{r}|^2}{\sigma^2}\right), \tag{5.8}$$

where σ is the Gaussian core-size. If the i^{th} particle centre is located at (x_i, y_i) then the velocity field induced at the evaluation point (x, y) is,

$$U = \frac{\Gamma_i}{2\pi}\frac{y-y_i}{(x-x_i)^2+(y-y_i)^2}\left(1-\exp\left(-\frac{(x-x_i)^2+(y-y_i)^2}{2\sigma_i^2}\right)\right) = I_{u,i}\Gamma_i, \tag{5.9}$$

$$V = \frac{\Gamma_i}{2\pi}\frac{x-x_i}{(x-x_i)^2+(y-y_i)^2}\left(1-\exp\left(-\frac{(x-x_i)^2+(y-y_i)^2}{2\sigma_i^2}\right)\right) = I_{v,i}\Gamma_i. \tag{5.10}$$

In the same manner as the BEM, the calculation at any point in a DVM flow field requires linear superposition of the influence of all N vortex particles,

$$U = \sum_{i=1}^{N} I_{u,i}\Gamma_i, \tag{5.11}$$

$$V = \sum_{i=1}^{N} I_{v,i}\Gamma_i, \tag{5.12}$$

and evaluating at M target positions,

$$\{U\} = [I_u]\{\Gamma\}, \tag{5.13}$$

$$\{V\} = [I_v]\{\Gamma\}, \tag{5.14}$$

results in a velocity field calculation that is of complexity $O(MN)$. Under the Lagrangian framework the velocity field must be evaluated at the centre of each particle, for every time-step, to allow particle convection. This results in a $[I_*]$ matrix that is dense and square $(N \times N)$, requiring a computational effort of the order $O(N^2)$ and classifying the DVM as an N-body method. This is a significant downside of the DVM because increasing the model resolution requires an increase in the number of discrete vortices that represent the flow. When compared to the BEM, the N-body problem of the DVM is exacerbated because a typical simulation will require more DVM particles than BEM panels. This is attributed to the BEM panels being applied on lines/curves whereas the vortex particles are scattered across a 2D plane.

5.1 The Fast-Vortex Algorithm

We investigate the use of the Fast Multipole Method (FMM) (introduced in Section 1.4.2 and developed in Chapter 2) to perform the full DVM field calculations in an efficient $O(N \log N)$ manner. The principle behind using the FMM algorithm that was used for the BEM is based on the observation that in regions 'far' outside a DVM particle's Gaussian core, the velocity influence equations tend

towards that of a BEM source/sink particle (with a $\log|\mathbf{z}|$ potential function). The main difference between the two elements is that the U and V velocity components are swapped. As a result, with appropriate modifications, we can use the existing BEM FMM framework to perform the Gaussian element's velocity-field calculations. In an approach similar to Hamilton and Majda (1995), the remaining FMM changes are relatively straight forward; using the size of the Gaussian core to define the FMM limit on what constitutes 'far' and 'near' field. Any evaluation outside the Gaussian core will treat the particle as a point vortex, allowing the full FMM approximations to apply. When inside the Gaussian core the multipole identities and truncations that the FMM were based upon are no longer valid. However, these evaluations can be treated as the 'near-field' calculations under the FMM and calculated in the standard $O(N^2)$ manner. Efficient use of the FMM requires setting a suitable tolerance for the allowable error in approximating a Gaussian particle with a standard point vortex. For this work we use the same error tolerance as the FMM truncation operations for the DVM, $\epsilon = 1 \times 10^{-6}$. This effectively defines the smallest Gaussian 'near-field' and thus the smallest allowable FMM box-size. The result is that a Gaussian core will never overflow a box's boundary and be allowed to participate in 'far-field' interactions.

5.2 Surface Modelling

A key disadvantage of the DVM is its inability to model an impermeable surface as a boundary condition. This is overcome by using source/sink panels from the BEM to enforce no-flux at the surface, an approach that is more versatile than using a method such as image vortices (Katz and Plotkin, 1991; Lewis, 1991). The calculation of the apparent normal velocity used to enforce the boundary conditions in the BEM must be adjusted to allow the influence of the vortex

particles,

$$V_{\text{norm}} = (V_\infty + V_{p,\omega} - \dot{\eta}_p)\cos\theta_p - (U_\infty + U_{p,\omega})\sin\theta_p. \tag{5.15}$$

The addition of the time dependent $U_{p,\omega}$ and $V_{p,\omega}$ terms represent the deviation of the flow-field velocity at the panel centre in the x and y directions, as a result of the introduced vortex particles. This coupling of the DVM particles to the BEM ensures that the existing formulation of the BEM and associated (eg. GMRES) solution methods remain valid. For evaluations of the full flow field, the influence of the vortex particles must also be included in addition to those of the BEM panels,

$$U = U_\infty + \sum_{i=1}^{N} U_{\omega_i} + \sum_{i=1}^{M} U_{p_i}, \tag{5.16}$$

$$V = V_\infty + \sum_{i=1}^{N} V_{\omega_i} + \sum_{i=1}^{M} V_{p_i}, \tag{5.17}$$

where N represents the number of vortex particles and M represents the number of BEM panels. These summations are conducted by the FMM algorithm, however in its current implementation they correspond to a separate call for each element type.

It is beneficial to define two additional BEM elements that can be used for solving vorticity-based boundary conditions. These boundary conditions will be of benefit to flow situations that are characterised mainly by the transport of vorticity. We define a zero-order vortex panel as a constant strength, distributed vortex panel with a finite length. It is obtained by integrating the effect of a series of constant strength point vortices with a midpoint of (x_0, y_0) and an induced flow velocity at (x, y) of (Katz and Plotkin, 1991),

$$U = \frac{\gamma}{2\pi}(\theta_r - \theta_l) = I_u\gamma, \tag{5.18}$$

$$V = \frac{\gamma}{4\pi}\ln\frac{{R_r}^2}{{R_l}^2} = I_v\gamma, \tag{5.19}$$

where γ is the vortex sheet strength and

$$\Delta y = (y - y_0), \tag{5.20}$$

$$\Delta x_k = (x - x_k), \tag{5.21}$$

$$R_k = \sqrt{(\Delta {x_k}^2 + \Delta y^2)}, \tag{5.22}$$

$$\theta_k = \tan^{-1} \frac{\Delta y}{\Delta x_k}, \tag{5.23}$$

for panel end points $k = l, r$. Unlike source/sink panels, it is difficult to model an impermeable boundary (no-flux condition) using these elements because their normal self-influence is zero. The strength of an element such as this is in enforcing no-slip/slip boundary conditions. To use vortex panels to enforce a no-flux condition we must obtain a higher-order panel element.

5.3 Illustrative Results – von-Kàrmàn Street

We apply the BEM/DVM combination to the model of an infinitely thin plate heaving in the presence of a high Reynolds-number (inviscid) uniform flow. A schematic of this model is shown in Figure 5.1. To model the fluid wake be-

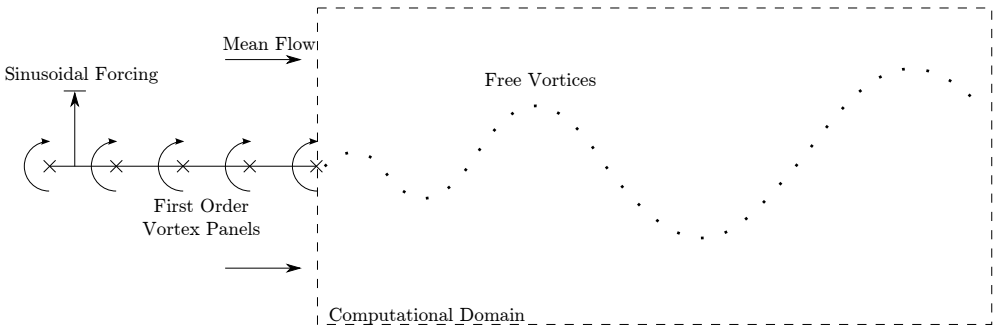

Figure 5.1: Schematic of the heaving plate model.

haviour, discrete vortices with Gaussian cores are shed from the trailing edge of the plate and are allowed to convect freely in the domain. A series of first-order vortex panels are used to model the sinusoidally heaving plate and enforce the no-flux boundary condition in the same manner as Chapter 2. The use of vortex

panels, as opposed to source/sink panels, allows easier enforcement of the Kutta condition. This is achieved by setting the trailing node of the heaving plate to zero vorticity, physically representing the no-slip where upper and lower streams meet, an inherent characteristic of a real flow.

The first-order vortex panel represents a line of integrated point vortices whose strength distribution varies linearly by $\gamma(x) = \gamma_g x + \gamma_c$ with an induced flow velocity of (Katz and Plotkin, 1991),

$$U = \frac{\gamma_c}{2\pi}(\theta_r - \theta_l) + \frac{\gamma_g}{4\pi}\left(\Delta y \ln \frac{{R_r}^2}{{R_l}^2} + 2\Delta x(\theta_r - \theta_l)\right), \tag{5.24}$$

$$V = \frac{\gamma_c}{4\pi} \ln \frac{{R_r}^2}{{R_l}^2} + \frac{\gamma_g}{4\pi}\left(\Delta x \ln \frac{{R_r}^2}{{R_l}^2} + 2\Delta y(\theta_l - \theta_r) + L\right). \tag{5.25}$$

The normal self-influence of this element at its centre is nonzero, indicating its enhanced suitability (as opposed to zero-order vortex panels) for no-flux boundary conditions. With an increasing order of complexity, we require two equations per panel to solve for the strength distribution. This is achieved by enforcing smooth transitions in wall vorticity between neighbouring panels by explicitly setting the strength of panel ends to match that of their neighbour. Determination of the strength of the released vortex is achieved implicitly by coupling its influence directly as an equation in the BEM system.

Figure 5.2 shows a snapshot after sufficient development of the wake behaviour for three different forcing frequencies of the heaving plate. Each marker in the figure represents the centre of a discrete packet of free vorticity that is convecting under the action of the resulting time-dependent flow field. The asymmetry observed in the figures for higher frequencies is attributed to the specific initial conditions of the nonlinear flow-field with the fluid and plate being at rest before the simulation is started. The numerically obtained wake pattern qualitatively matches the experimental results obtained by Lau et al. (2004) for the wake pattern generated by a heaving and pitching aerofoil in a wind tunnel. This match provides evidence to support the BEM[1] and DVM as tools for modelling

[1]While we demonstrated a vortex-panel BEM, its implementation is based on the source/sink panel BEM in Chapter 2.

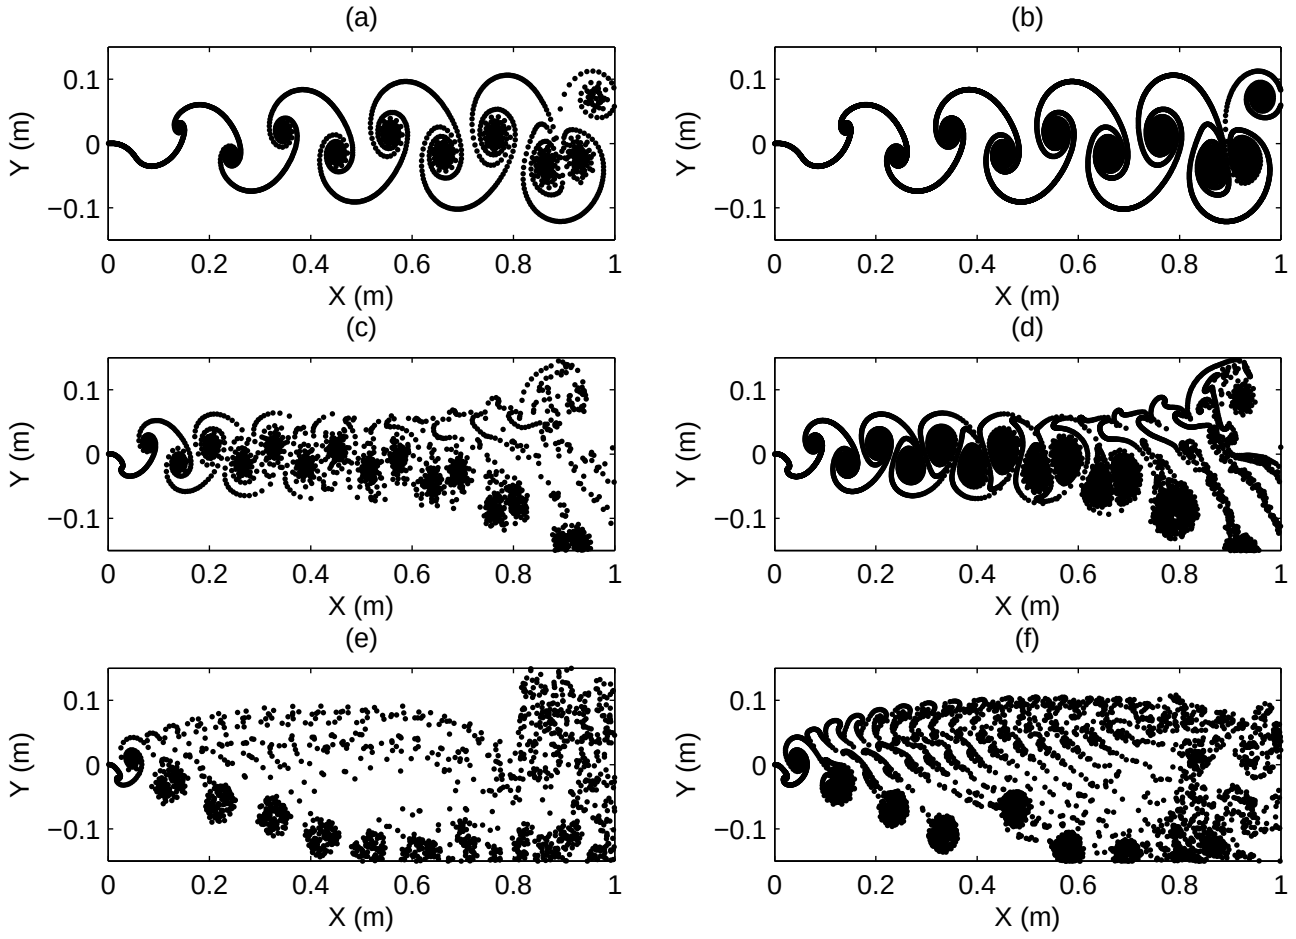

Figure 5.2: Generated wake pattern for the heaving plate model using: (a) Standard DVM, $f_f = 1$Hz, (b) FMM, $f_f = 1$Hz, (c) Standard DVM, $f_f = 2$Hz, (d) FMM, $f_f = 2$Hz, (e) Standard DVM, $f_f = 4$Hz, (f) FMM, $f_f = 4$Hz.

real flows involving surfaces and rotational flow behaviour.

While the FMM has been shown in Chapter 2 to work suitably with the BEM, the present case provides an opportunity to demonstrate its suitability and benefit to the DVM/BEM combination. As such, two separate DVM wake discretisations have been utilised in Figure 5.2 to compare the difference between the FMM and standard algorithm for velocity-field calculations. For all cases the forcing amplitude is $A_f = 0.01$m and Gaussian core-size is $\sigma = 0.005$. All FMM simulations use a time-step size of $\Delta t = 2 \times 10^{-4}$ s resulting in a total of $N \approx 12,500$ free vortices, with the standard DVM cases using $\Delta t = 1 \times 10^{-3}$ s and $N \approx 2,500$. The total particle numbers (bound by time-step size) were chosen to represent similar computational effort per time-step at $\approx$ 0.3s on an Intel Q9650 desktop computer. The increase in maximum discretisation for constant computational effort is a consequence of the superior scaling of the FMM over the standard DVM velocity field calculations.

In the lower forcing frequency cases, the FMM yields identical results to the standard DVM algorithm. This indicates that at the resolution of $N = 2,500$ particles (and $\Delta t = 1 \times 10^{-3}$ s), the predicted wake behaviour is discretisation independent, an outcome which is not true for higher forcing frequencies. The use of the FMM for DVM field calculations demonstrates that improvements in simulation resolution are available at no cost to solution accuracy or computational time (per time-step). The additional flow-field resolution also revealed a different wake pattern for high forcing frequencies. This is a result that would have been missed when using the low resolutions of the standard method.

5.4 Mixing-Layer Model

Another demonstration case for the DVM and its supporting infrastructure is that of a turbulent mixing layer. Physically this structure occurs when two separated streams running parallel at constant, yet different speeds are allowed to mix and undergo subsequent convection. Numerically this can be represented

by the model of a single convecting shear layer superimposed onto a uniform inviscid (viscous diffusion is neglected) flow field as shown in Figure 5.3. This

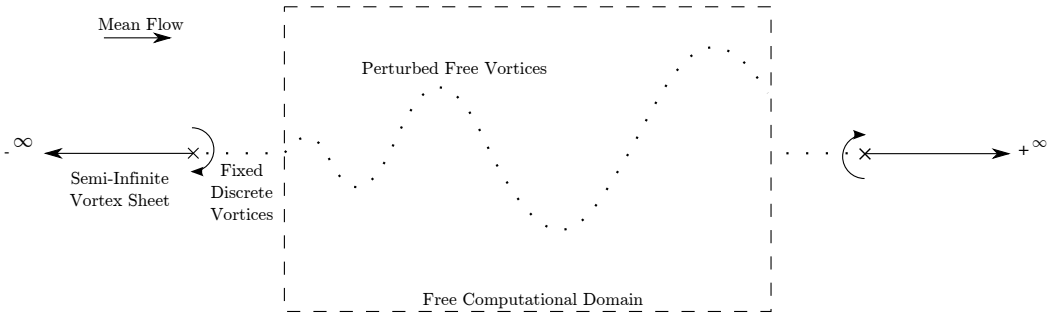

Figure 5.3: Schematic of the mixing layer model.

model resembles the founding numerical experiments of Rosenhead (1931) with the exception that our vortex sheet is subjected to a mean flow. The numerical foundations of the present model can be seen in similar works such as Basu et al. (1992, 1995); Cafolla (1997) and Inoue (1985). The main component of the model is a series of discrete vortices with Gaussian cores that represent the shear-layer. The vortices are free to convect according to the resulting time-dependent flow field to simulate the flow's mixing behaviour. Vorticity injection occurs with every time-step at the leading edge of the domain to conserve shear-layer inlet vorticity and vorticity removal occurs at the trailing edge of the domain as a result of the flow convection. While a layer of free vortices represents the shear layer in the computational domain, semi-infinite vortex sheets (discussed in the following paragraphs) are used to approximate the far-field influence of the shear layer outside of the domain. Approximation of the far-field influence using semi-infinite vortex sheets for this case is essential, without them the free vortices would undergo solid body rotation as soon as time-stepping begins.

A semi-infinite vortex sheet mathematically represents the inverse of a vortex panel; the influence of a line of distributed point vortices of constant strength, integrated from/to $\pm\infty$, with a finite section at its midpoint removed. In the present model this removed section represents the free computational domain that is modelled by the free vortices. The velocity field induced at $(x, \pm y)$ by the

semi-infinite sheet with a midpoint of (x_0, y_0) is,

$$U = \frac{\gamma}{2\pi}(\theta_r - \theta_l \pm \pi), \tag{5.26}$$

$$V = \frac{\gamma}{4\pi} \ln \frac{{R_l}^2}{{R_r}^2}, \tag{5.27}$$

where γ is the strength per unit length of the vortex sheet. If the two modelled parallel stream speeds are U_1 and U_2 then the resulting vortex-sheet strength is,

$$\gamma = U_1 - U_2, \tag{5.28}$$

and effective mean-flow speed (U_∞) is,

$$U_\infty = \frac{U_1 + U_2}{2}. \tag{5.29}$$

An additional feature of the model is a small section of discrete vortices that are fixed in the region of the domain's leading and trailing edge to aid in the transition from the semi-infinite vortex sheets to free Gaussian vortices.

5.4.1 Illustrative Results – Kelvin-Helmholtz Instability

By setting the stream speeds to $U_1/U_\infty = 1.5$ and $U_2/U_\infty = 0.5$, Figure 5.4 illustrates the profile obtained when sampling the velocity field about the unperturbed shear-layer model. The smooth transition from above and below the layer illustrates the smoothing effect of the Gaussian vortices.

When conducting time-dependent numerical experiments the free shear-layer is initialised with a sinusoidal disturbance to speed-up the inevitable destabilisation process. However, starting the simulation with free vortices in a straight line will still result in mixing after sufficient time due to its unstable nature and the noise induced by time-stepping. Three consecutive snapshots of the unsteady mixing layer simulation are shown in Figure 5.5 with a spacing of 150 time-steps. This is where time-steps are $\Delta t U_\infty / L = 1.33 \times 10^{-3}$ and the domain is length L. These figures are obtained after the initial disturbance has convected beyond the

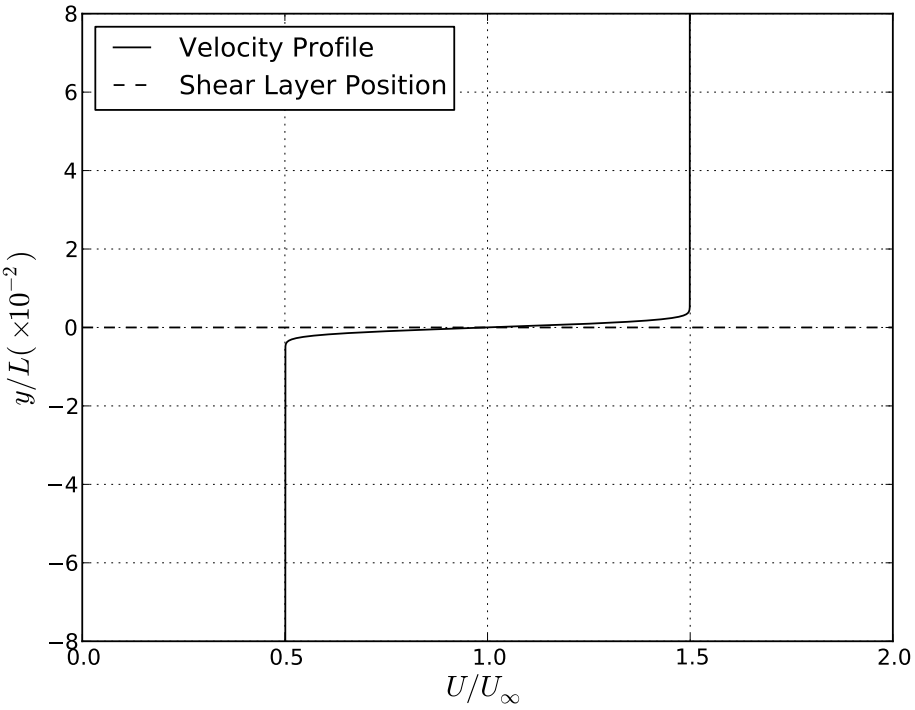

Figure 5.4: Horizontal velocity profile obtained numerically for the unperturbed mixing layer model with $U_1/U_\infty = 1.5$ and $U_2/U_\infty = 0.5$.

free domain. Each marker in the figure represents the spatial (x, y) position of the centre of a free Gaussian vortex that convects under the action of the surrounding time-dependent flow field. This result demonstrates the evolution of the mixing layer instability and qualitatively matches the results of Basu et al. (1992, 1995), Cafolla (1997, see Figure 5.6) and Inoue (1985). However, further investigation using this model is not within the scope of this work. It has been introduced with the sole purpose of forming the basis of the DVM based boundary-layer model in subsequent chapters.

5.5 Summary

The Discrete Vortex Method (DVM) is a grid-free, Lagrangian method that is based on the velocity-vorticity form of the Navier-Stokes equations. It suffers from the N-body problem that is mitigated by the use of the Fast Multipole Method (FMM). Re-use of the previously developed FMM algorithm required special

attention when dealing with Gaussian vortices, achieved by approximating point vortices in the far-field. This was demonstrated successfully in conjunction with the use of the existing Boundary Element Method (BEM) for surface modelling under the heaving plate model.

A numerical model for the evolution of a 2D mixing-layer was demonstrated. This model will form an integral part in the development of a robust wall-bounded boundary-layer model.

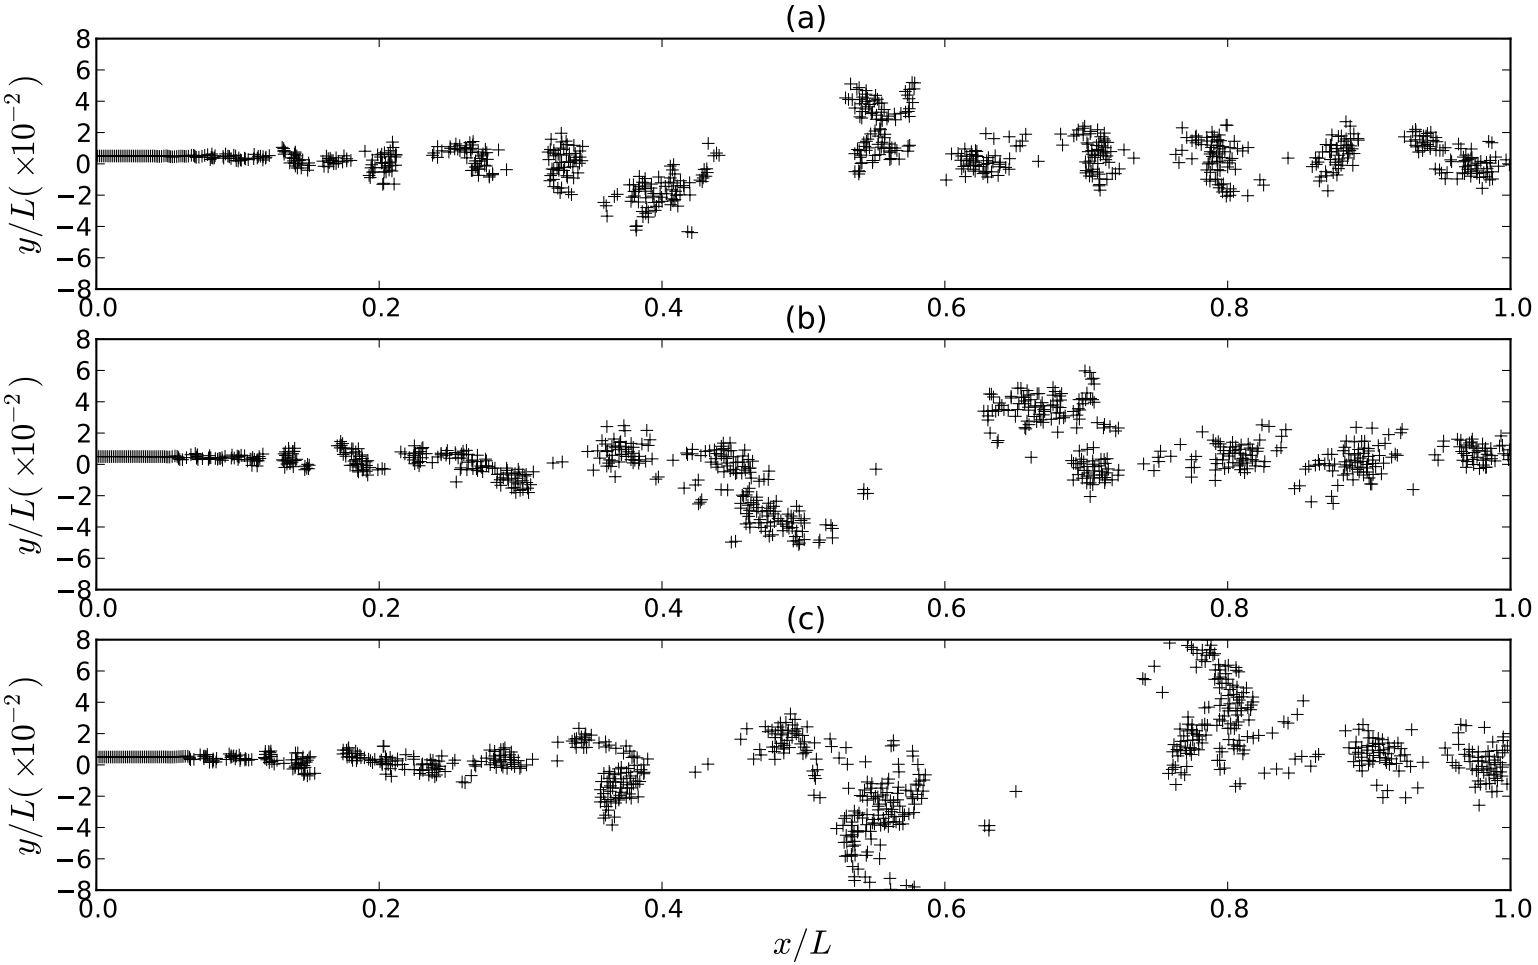

Figure 5.5: Consecutive snapshots from top to bottom, at 150 time-step ($\Delta t U_\infty / L = 1.33 \times 10^{-3}$) spacings, for the mixing layer model. See following Figure 5.6 for qualitative comparison to the various flow regions described in Cafolla (1997).

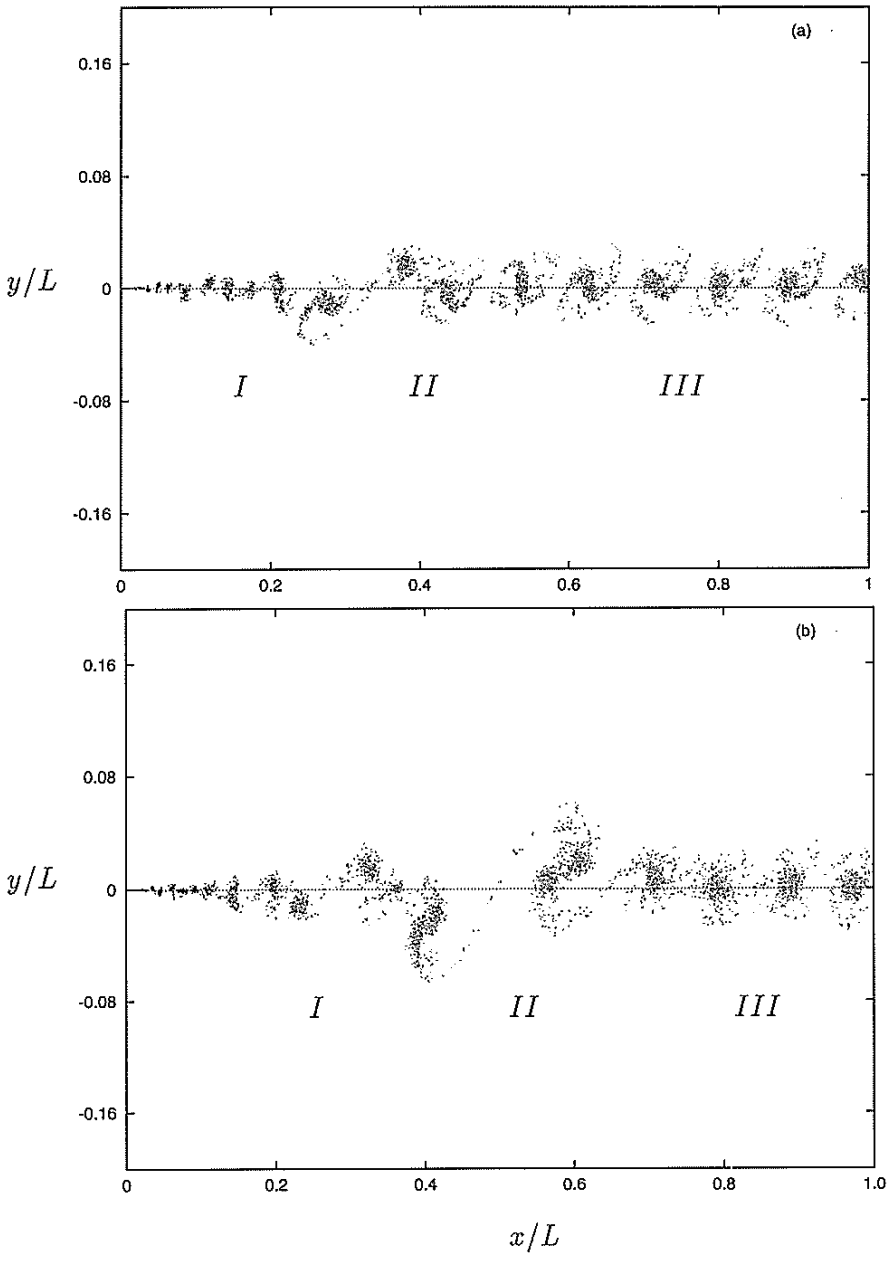

Figure 5.6: Reproduced Figure 3.4 from Cafolla (1997), illustrating vortex roll-up (time increasing from top to bottom) of a 2D mixing layer using the DVM. Region I is the developing mixing layer and Region III is the region of vortex roll-up as a result of the non-physical initial conditions. Region III illustrates large vortex clustering at the transiton between the two regions.

Chapter 6

Boundary-Layer Flow Modelling

The Discrete Vortex Method (DVM) models presented thus far have consisted of high Reynolds-number flows which exhibit rotationality but are largely inviscid in nature. For wall-bounded flows at moderate (i.e. transitional) Reynolds numbers, the effects of viscosity such as the no-slip condition, the attendant formation of a boundary-layer velocity profile, and viscous diffusion must be included in the fluid model.

In this work we use a parallel boundary-layer model that is similar in its foundations to that of Cafolla (1997) and Pitman (2007) and focuses on the perturbations to a boundary-layer that has already established a given thickness (δ). Thus, we assume that the boundary-layer growth is negligible for the typical stream-wise length scales and Reynolds-number flows of interest.

In its unperturbed state, a boundary-layer velocity profile can be modelled using a series of flat shear-layers stacked upon one another to represent the different levels of shear that occur due to the action of viscosity and the no-slip condition. We use the previously developed mixing-layer model (in Section 5.4) as an elemental building block for the stacked shear-layers as shown in Figure 6.1. Fundamentally the model consists of a computational 'free-field' that is populated by overlapping Gaussian vortices while far-field effects are modelled by semi-infinite vortex sheets that satisfy the parallel boundary-layer assumption. Fixed Gaussian vortices are included in the up- and down-stream regions of the

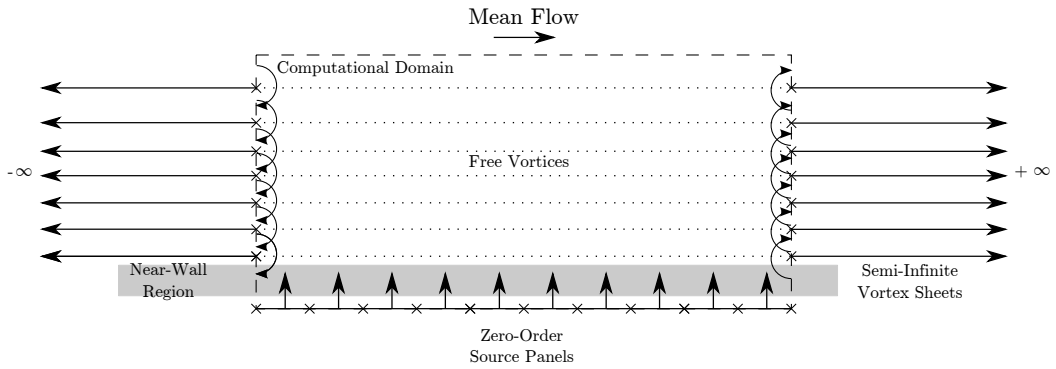

Figure 6.1: Schematic of the DVM model for a boundary-layer flow over a flat wall.

free-field to smooth the transition from the semi-infinite sheet elements. The length of these regions are equal to the stream-wise free-domain length, L.

When the Gaussian core of a free vortex overlaps a solid boundary, its strength can be mis-represented by allowing its vorticity to effectively leak through the wall. A local solution for elements that overlap a solid boundary is to replace their Gaussian core with that of a point-vortex element ($\sigma = 0$) (Ploumhans and Winckelmans, 2000). However, this approach can result in a flow-field that is not always continuous and would therefore violate the methods used herein for wall-pressure evaluations. An alternative solution is to introduce a near-wall region that eliminates local particle convection. This region replaces free particles with specialised elements that encapsulate the behaviour of the essentially wall-bounded flow (Pitman, 2007). The method used herein is a hybrid of both methods, based on modifying the behaviour of the near-wall region[1] that encapsulates the lowest free shear-layer. The Gaussian vortices of this layer are replaced by zero-order vortex panels which are fixed to the motions of the nearest BEM panel in the wall, shown in Figure 6.2. The velocity induced by the panels is interpolated[2] across the sub-layer region to improve the near-wall stream-wise velocity profile.

The model is discretised in the vertical (y) direction based on shear-layers

[1]Also known as the linear viscous sub-layer.

[2]The stream-wise velocity jump across the panel is smoothed using a linear function.

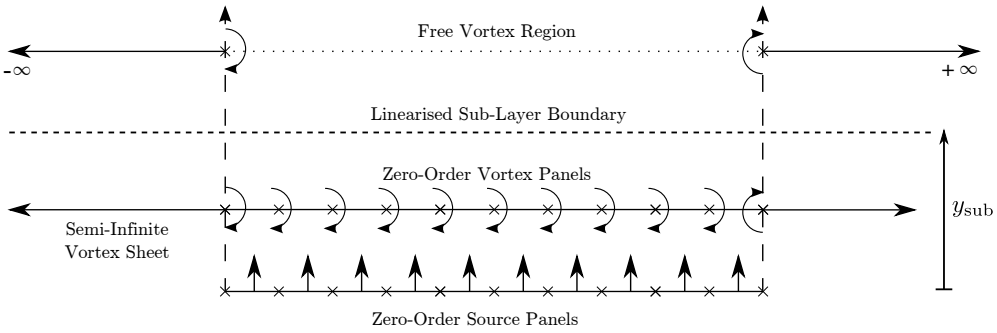

Figure 6.2: Additional schematic of the near-wall region for the boundary-layer flow model.

with fixed separations of Δy. The strengths of the shear-layers are adjusted to fit the curve of the given boundary-layer mean-flow velocity/vorticity profile. Each shear-layer is discretised in the horizontal (x) direction into a series of overlapping Gaussian vortices with a spacing of Δx and initial core-size of σ. We use a rectangular grid for our initial Lagrangian discretisation, where $\Delta x = \Delta y = h$, with a Gaussian overlap ratio of $\beta = h/\sigma < 1$.

6.1 Background Velocity Profiles

The mean-flow velocity profile of the unperturbed boundary-layer is given by a 4^{th} order Pohlhausen approximation[3] to a laminar Blasius boundary-layer,

$$\frac{U}{U_\infty} = 2\left(\frac{y}{\delta}\right) - 2\left(\frac{y}{\delta}\right)^3 + \left(\frac{y}{\delta}\right)^4, \tag{6.1}$$

where y/δ is the normalised vertical position within the boundary-layer and U/U_∞ is the stream-wise velocity normalised to the outer-flow speed.

To demonstrate the suitability of the DVM to adequately model the behaviour of the unperturbed Blasius profile, we take a sample of the velocity generated within a vertical slice above a flat wall. With a linearised sub-layer from $y_{\text{sub}} < 0.15\delta$ the simulated velocity profile is compared to the desired Blasius profile in Figure 6.3. The current arrangement of the DVM models the desired Blasius

[3]More accurate approximations are available, however, this approximation is chosen to allow comparisons to the results of previous works such as Cafolla (1997); Pitman (2007).

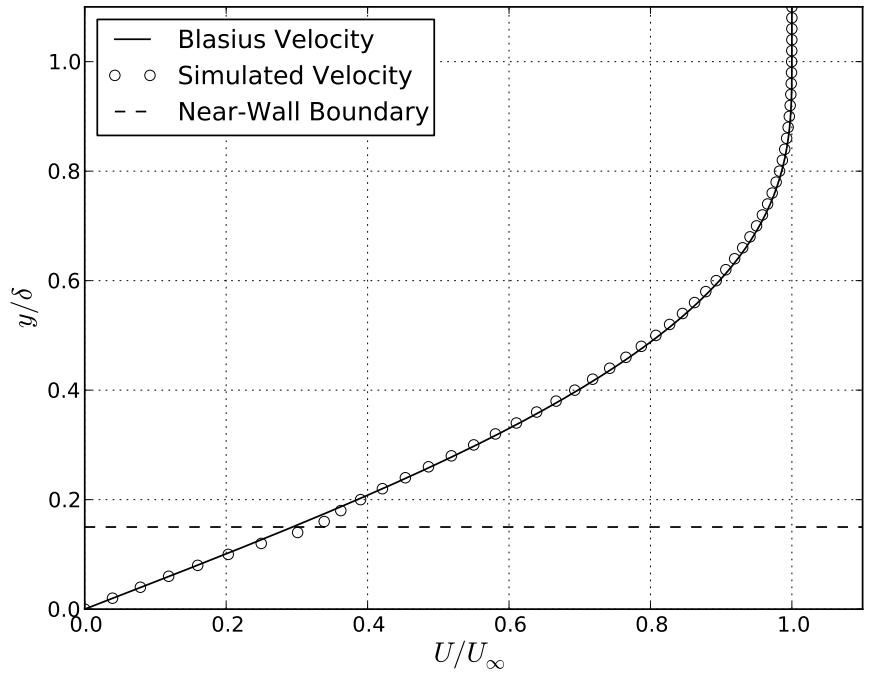

Figure 6.3: Unperturbed velocity profile for the DVM-based laminar boundary-layer model with $y_{\text{sub}} < 0.15\delta$ and $\beta = 1$.

velocity profile to a high level of accuracy. However, it is noted that this model introduces an inflection point at the boundary of the near-wall region. A method to minimise this undesirable feature will be discussed in the sections to follow.

The accuracy of this model is quantified by comparing the boundary-layer displacement thickness of both the model and the theoretical case (δ^*_{th}) using,

$$\delta^* = \int_0^\delta \left(1 - \frac{U(y)}{U_\infty}\right) \text{dy}. \tag{6.2}$$

With $\delta^*_{\text{th}} = 0.301\delta$, the resulting error in the displacement thickness of the model is 0.34%. We investigate changes to the boundary-layer model by increasing the Gaussian particle density. This is achieved by increasing the Gaussian overlap ratio to $\beta = 0.5$ or reducing the near-wall cut-off at $y_{\text{sub}} < 0.1\delta$, seen respectively in Figures 6.4 and 6.5. This yields marginal improvements in model accuracy, evidenced by the reduced displacement-thickness error of 0.16% and 0.15% respectively. To illustrate the cost of these measures we present a small segment of

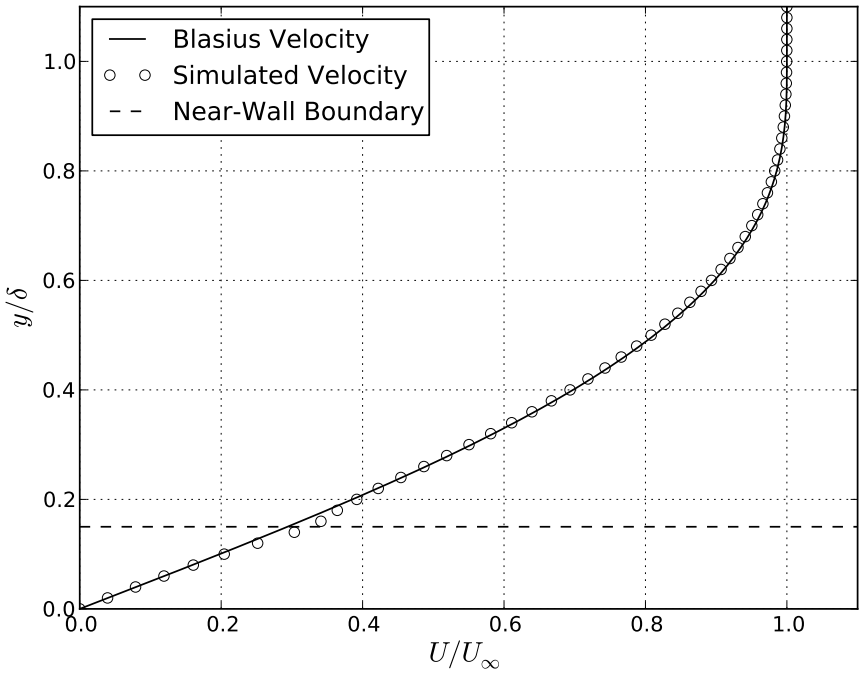

Figure 6.4: Unperturbed velocity profile for the DVM-based laminar boundary-layer model with $y_{\text{sub}} < 0.15\delta$ and $\beta = 0.5$.

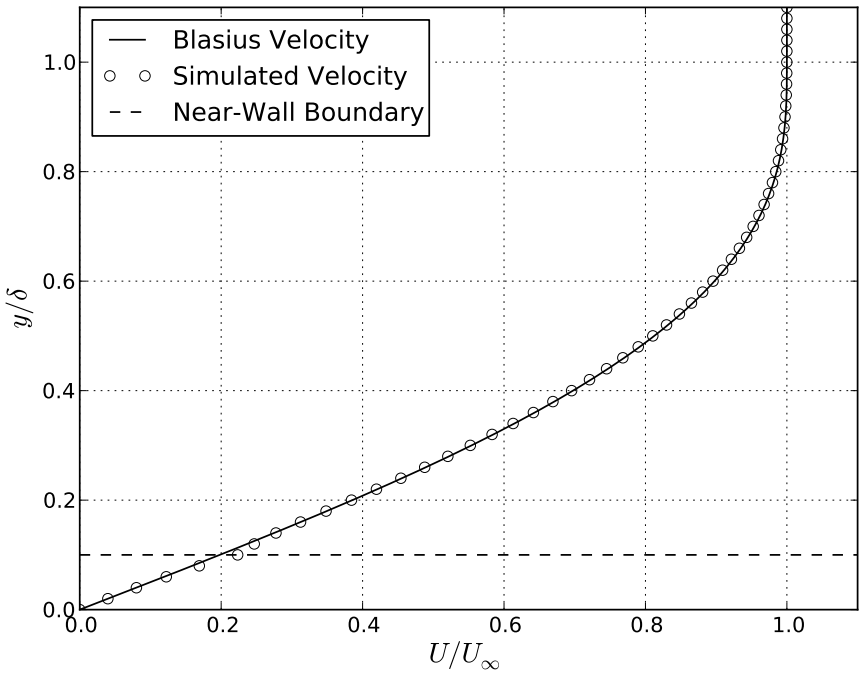

Figure 6.5: Unperturbed velocity profile for the DVM-based laminar boundary-layer model with $y_{\text{sub}} < 0.1\delta$ and $\beta = 1.0$.

each Langrangian grid respectively in Figures 6.6 to 6.8. Markers in the figures represent the centre of a free vortex while the shaded circles represent area within their respective core ($r < \sigma$) and corresponding particle overlap. The initial case

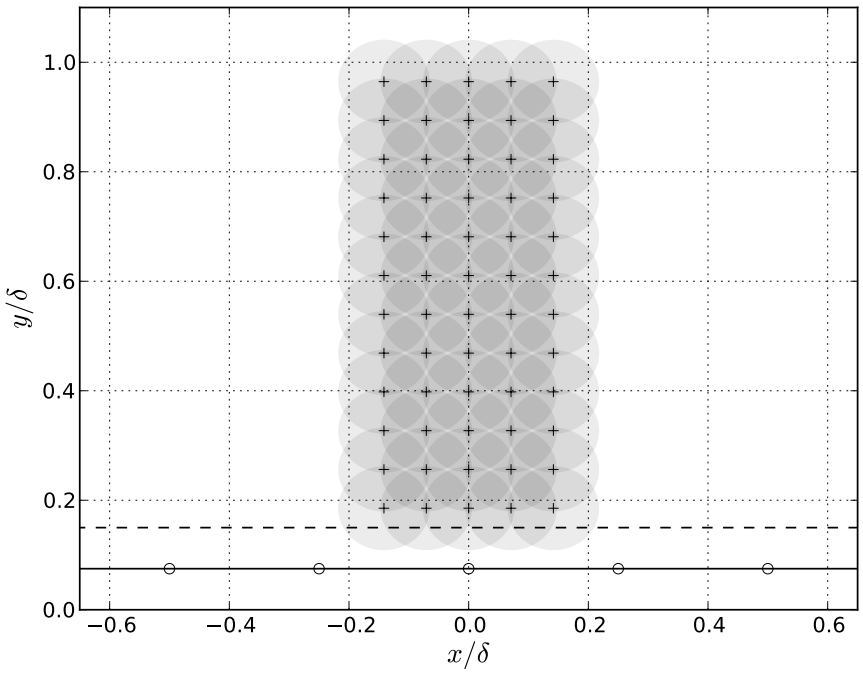

Figure 6.6: A segment of the Lagrangian grid showing Gaussian-vortex centres [+] and their corresponding core-size/overlap [shading] for the boundary-layer model with $y_{\text{sub}} < 0.15\delta$ and $\beta = 1$.

of $\beta = 1$ and $y_{\text{sub}} < 0.15\delta$ consisted of a y-discretisation of 12 shear-layers. On a rectangular grid this corresponds to a total number of particles of $169 \times L/\delta$ for an entire computational domain with horizontal length L. Reducing the overlap ratio to $\beta = 0.5$ required an increase in y-discretisation to 23 shear-layers and $624 \times L/\delta$ particles per domain, while lowering the near-wall region to $y_{\text{sub}} < 0.1\delta$ required 18 shear-layers and $362 \times L/\delta$ particles.

Increasing the accuracy of the modelled velocity-profile requires a large increase in particle density. High particle densities become particularly prohibitive when used in a time-dependent scheme and compounded even further when nested within an iterative fluid-structure interaction (FSI) framework. Consequently for the remainder of this work we restrict the Gaussian overlap ratio to $\beta = 1.0$ and the linearised sub-layer to $y_{\text{sub}} < 0.15\delta$.

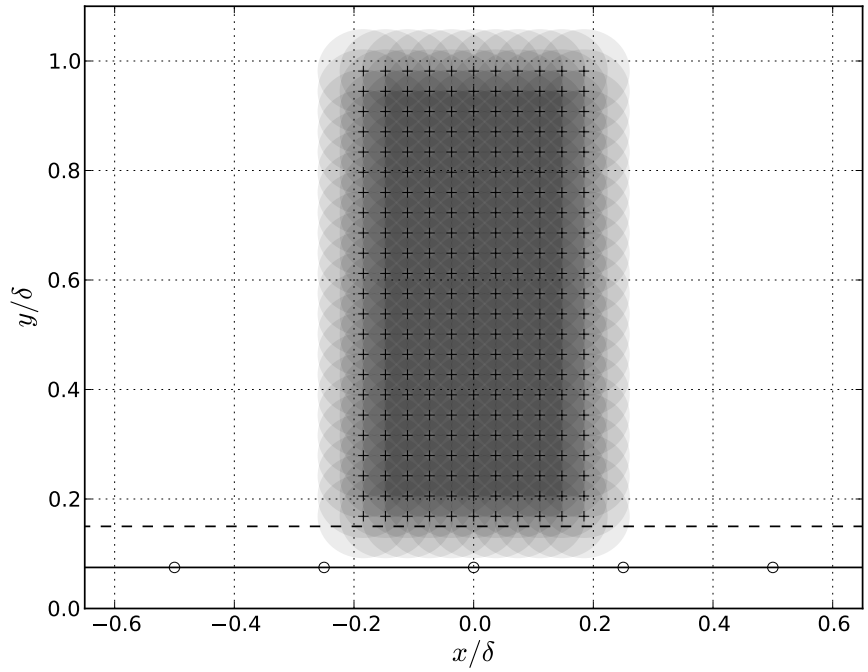

Figure 6.7: A segment of the Lagrangian grid showing Gaussian-vortex centres [+] and their corresponding core-size/overlap [shading] for the boundary-layer model with $y_{\text{sub}} < 0.15\delta$ and $\beta = 0.5$.

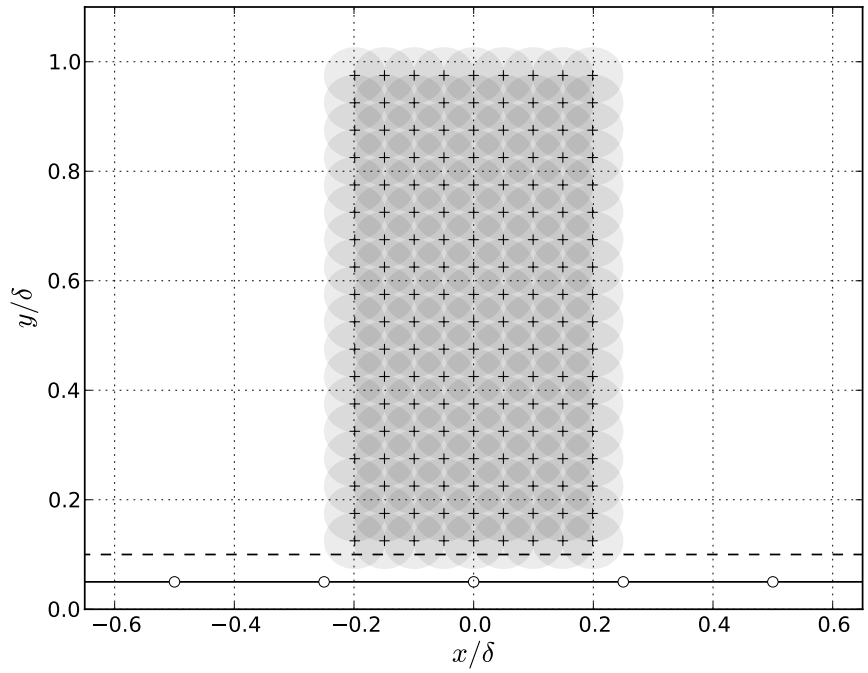

Figure 6.8: A segment of the Lagrangian grid showing Gaussian-vortex centres [+] and their corresponding core-size/overlap [shading] for the boundary-layer model with $y_{\text{sub}} < 0.1\delta$ and $\beta = 1$.

6.1.1 Variable Core-size Discretisations

While the rectangular grids previously discussed offered acceptable resolution of the unperturbed velocity profile, the constant nature of its discretisation and core-size results in a model that is over-resolved in areas of low velocity gradient, while the flow is comparatively under-resolved in regions of high gradient.

We introduce an alternative method to initialising the Lagrangian grid based on consistent resolution of stream-wise velocity-gradients. This is achieved using shear-layers that have Gaussian core-sizes that vary throughout the boundary-layer to restrict the maximum induced convection error of any particle. Under time-dependent simulations of the DVM, each Gaussian vortex represents a 'patch' of flow vorticity. This entire patch convects according to the streamline followed by its centre, however, this relies on the flow-field across it being constant. We define the convection error of a particle (ϵ_{conv}) based on its greatest local change in velocity,

$$\epsilon_{\text{conv}} = 1 - \left| \frac{U\left(y - \sigma(y)\right)}{U(y)} \right|, \tag{6.3}$$

where $U(y)$ is the velocity at a particle's centre and $U(y - \sigma(y))$ represents the velocity at the patch edge closest to the wall. The U velocity component is the only one of interest because it can be expected to exhibit the largest gradient, i.e.

$$\frac{\partial U}{\partial y} \gg \frac{\partial U}{\partial x}, \quad \frac{\partial U}{\partial y} \gg \frac{\partial V}{\partial y}, \quad \frac{\partial U}{\partial y} \gg \frac{\partial V}{\partial x}. \tag{6.4}$$

If the acceptable convection error ϵ_{conv} is fixed *a-priori* then $\sigma(y)$ represents the function for largest acceptable core-size at any position and can be solved as an optimisation problem. The result is a core-size profile that begins larger at the boundary-layer edge and then adaptively reduces towards zero at the wall.

The modelled velocity profile obtained when setting a maximum convection error of $\epsilon_{\text{conv}} = 0.125$, a linearised-sublayer at $y_{\text{sub}} < 0.15\delta$ and a minimum core-overlap ratio of $\beta = 1$, is shown in Figure 6.9. The variable core-size scheme

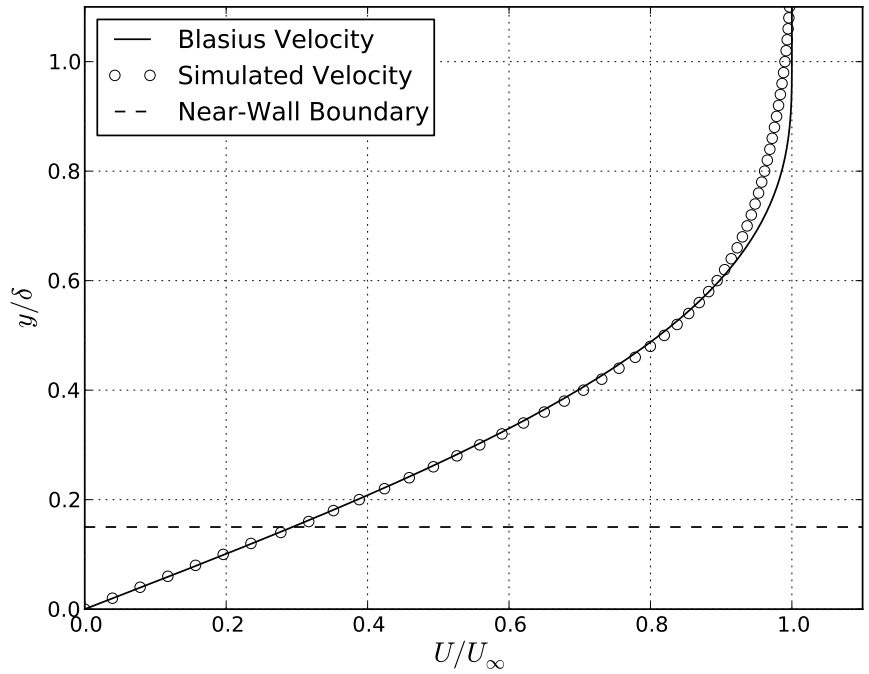

Figure 6.9: Unperturbed velocity profile for the DVM-based laminar boundary-layer model with a variable-core discretisation based on $\epsilon_{\text{conv}} = 0.125$, $y_{\text{sub}} < 0.15\delta$ and $\beta = 1.0$.

reproduces the Blasius velocity profile to an acceptable level with a displacement-thickness error of 1.9%. It is seen that the discretisation of the upper regions of the boundary-layer has been sacrificed to improve the resolution of the near-wall region, eliminating the near-wall inflection point noted earlier and seen in Figure 6.3. The significant advantage of this scheme is that the linear sub-layer can be easily reduced without impacting the resolution of the upper boundary-layer. By setting the linearised sub-layer to $y_{\text{sub}} < 0.0125\delta$, the modelled velocity profile is shown in Figure 6.10. The reduction in the size of the sub-layer has reduced the error in displacement thickness to 0.9%. However, this reduction has come at the cost of a very large increase in particle numbers. A small section of the variable core-size grid has been shown for $y_{\text{sub}} < 0.15\delta$ and $y_{\text{sub}} < 0.0125\delta$ respectively in Figures 6.11 and 6.12. With $y_{\text{sub}} < 0.15\delta$, the model consisted of a y-discretisation of 13 shear-layers and resulted in $325 \times L/\delta$ particles per domain. Reducing the linearised sub-layer to $y_{\text{sub}} < 0.0125\delta$ increased total shear-layers

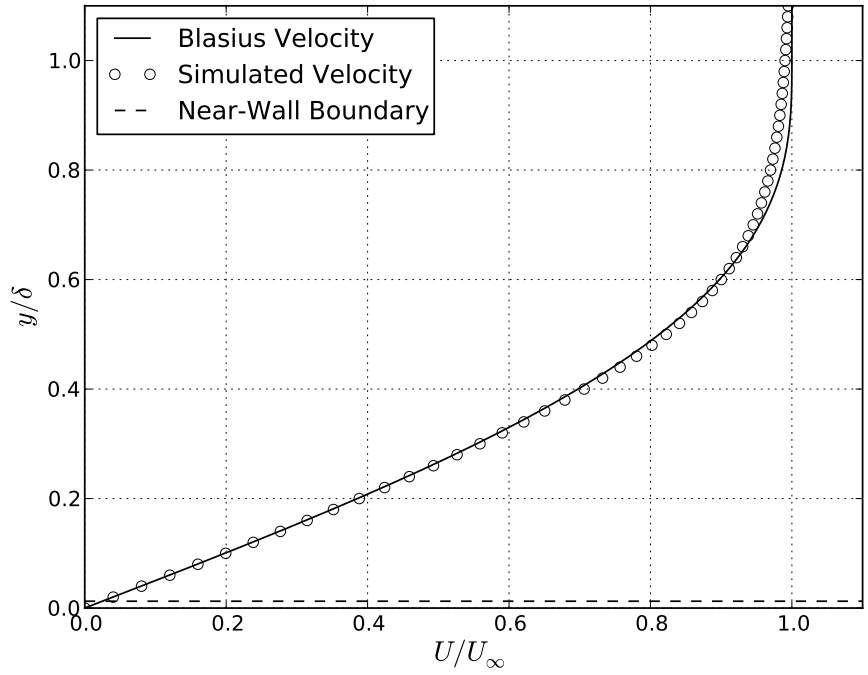

Figure 6.10: Unperturbed velocity profile for the DVM-based laminar boundary-layer model with a variable-core discretisation based on $\epsilon_{\text{conv}} = 0.125$, $y_{\text{sub}} < 0.0125\delta$ and $\beta = 1.0$.

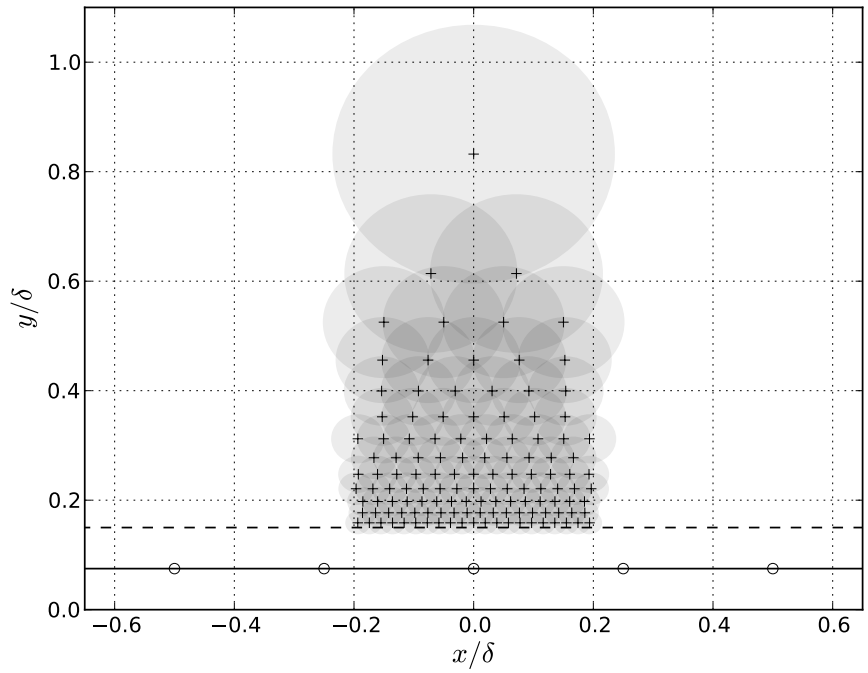

Figure 6.11: A segment of the Lagrangian grid showing Gaussian-vortex centres [+] and their corresponding core-size/overlap [shading] for the variable-core boundary-layer model with $\epsilon_{\text{conv}} = 0.125$, $y_{\text{sub}} < 0.15\delta$ and $\beta = 1.0$.

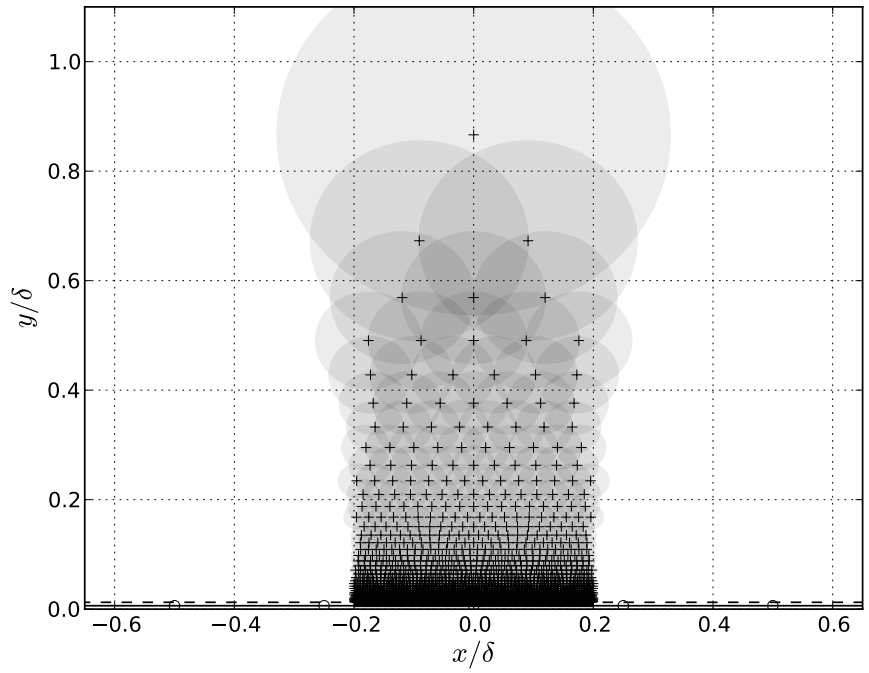

Figure 6.12: A segment of the Lagrangian grid showing Gaussian-vortex centres [+] and their corresponding core-size/overlap [shading] for the variable-core boundary-layer model with $\epsilon_{\text{conv}} = 0.125$, $y_{\text{sub}} < 0.0125\delta$ and $\beta = 1.0$.

to 37 and total particles to $6250 \times L/\delta$.

We note that the concept of using a variable core-size Lagrangian grid has mainly been provided as a stimulus for future work. It will not be used any further due to its incompatibility with the vorticity rediscretisation scheme that is introduced in Section 6.2.1. It shows great promise in being able to obtain higher quality discretisations over rectangular grids by efficiently resolving (as opposed to approximating) the entire near-wall region without using unnecessary resolution in the outer-region. However, as this still comes at the cost of increased (19×) particle numbers, the advantages gained in modelling such a low-convection and high-diffusion region of the flow would need to be evaluated on a case-by-case basis.

6.2 Time-Dependent Flows

To resolve the time-dependent behaviour of the boundary-layer model the vortex particles in the computational free-field are time-stepped using a first-order accurate Euler time-stepping scheme. Vorticity injection and deletion occurs in a conservative manner at the domain inlet (left-hand side) and exit (right-hand side) respectively, while vortices are deleted if they convect below the wall. To model the impact of the movable and impermeable flexible-plate surface, we use the BEM to enforce the zero normal-flow condition with zero-order source/sink panels. The no-slip condition at the wall is enforced using zero-order vortex panels. These elements are superimposed on the source/sink panels and are solved simultaneously with the zero normal-flow condition. To maintain the no-slip condition at each time-step, a non-zero slip velocity at the wall will require a change in the strength of the vortex panels. The perturbation in panel strength is then discretised into Gaussian vortices and released back into the free-flow domain at the lowest shear-layer. This models the creation/destruction of vorticity (positive/negative) by the wall in response to flow perturbations.

6.2.1 Viscous Diffusion and Rediscretisation

The effects of viscous diffusion are modelled using operator splitting and a core-spreading method (CSM) for the Gaussian DVM (Leonard, 1980). This is based upon solving the viscous component (right hand side) of the vorticity transport equation for viscous 2D incompressible flow,

$$\frac{\partial \omega}{\partial t} + (\vec{u} \cdot \nabla)\, \omega = \nu \nabla^2 \omega, \tag{6.5}$$

by expanding the cores of the Gaussian vortices linearly in time according to,

$$\frac{\mathrm{d}\sigma^2}{\mathrm{d}t} = 4\nu, \tag{6.6}$$

where ν is the kinematic viscosity. Thus, once a particle's motion has been calculated in a given time-step, the size of its core can be obtained using a first-order integration scheme,

$$\sigma_t = \sqrt{\sigma_{t-1}^2 + 4\nu\Delta t}, \tag{6.7}$$

where Δt is the time period over which the particle has convected/diffused. However, to coincide with the parallel boundary-layer assumption, we only apply core-spreading to the perturbation components of the flow. This prevents the diffusion and subsequent growth of the mean-flow velocity profile.

To maintain convection accuracy throughout long-time simulations, we use a Lagrangian-grid 're-mapping' (rediscretisation) scheme similar to Barba et al. (2005). This is based on mapping the perturbed vorticity field onto a separate Lagrangian grid using radial-basis function interpolation and the Generalised Minimum Residual Method (GMRES). The advantage of this scheme is that, unlike circulation re-distribution methods, there are no strict requirements on the number, position (x, y) or core-size (σ) of the particles represented by the new grid. However, use of non-rectangular and variable core-size grids (such as those presented in Section 6.1.1) can cause the system matrices to have a greater density and/or become ill-conditioned. This greatly increases the computational effort required to obtain a suitable pre-conditioner and the GMRES solution (if one can be found at all).

We apply the re-mapping process at every time-step to ensure that the quality of the Lagrangian grid remains consistent throughout the simulation. The grid template used is the same as that of the initialised unperturbed flow-field. The predefined nature of this rectangular grid allows efficient formation of a sparse and static pre-conditioner that reduces the intensity of the entire re-mapping process. In systems where the solid boundary is perturbed away from the initial flat position the re-mapping template is adjusted to follow the deformation of the wall. However, the sparse pre-conditioner used for a perturbed system remains unchanged from the flat-wall case due to the prohibitive cost of its re-formation.

From a numerical point of view, we evaluate the vorticity field $\{\omega_*\}$ about the points of the re-mapping template $\{x_*, y_*\}$ to,

$$\{\omega_*\} = [I_{\omega_*,\Gamma}]\{\Gamma\}, \tag{6.8}$$

where $[I_{\omega_*,\Gamma}]$ represents the influence matrix of the perturbed vortex particles, with strengths $\{\Gamma\}$, on the template's grid-points. Likewise, we can form another system of linear equations based on the vorticity field resulting from the Gaussian particles represented by the re-mapping template,

$$\{\omega_*\} = [I_{\omega_*,\Gamma_*}]\{\Gamma_*\}, \tag{6.9}$$

where $[I_{\omega_*,\Gamma_*}]$ represents the influence matrix of the rediscretised particles, with strength $\{\Gamma_*\}$, onto themselves ($\{x_*, y_*\}$). By substituting Equation (6.9) into Equation (6.8) we are left with,

$$[I_{\omega_*,\Gamma_*}]\{\Gamma_*\} = [I_{\omega_*,\Gamma}]\{\Gamma\}. \tag{6.10}$$

While $\{\Gamma_*\}$ (the only unknown) can be solved using linear algebra, the relatively large number of vortex particles used in this work ($N > 10^3$) necessitates the use of a pre-conditioned GMRES (see Section 2.2).

6.3 Illustrative Results – Tollmien-Schlichting Waves

Thus far, we have demonstrated the capabilities of the model to capture static boundary-layer features. We now attempt to validate the dynamic properties of the model by investigating flow-based instabilities that arise from the growth of Tollmien-Schlichting waves (TSW) over a flat plate.

One method for analytically modelling TSW dynamics is through the use of the Orr-Sommerfeld equation. This is based on the assumption of a paral-

lel boundary-layer of infinite length, with a disturbance applied throughout the domain that takes the form of the stream function,

$$\psi(x, y, t) = A(y)e^{i(\alpha x + \beta t)}. \tag{6.11}$$

The results obtained from conducting an eigen-analysis are illustrated in the stability map in Figure 6.13 from Schlichting (1979). This illustrates the range of disturbance wavelengths that are unstable for given Reynolds-number flows. From the markers in the figure labelled 'I' and 'II', Schlichting also produces Figure 6.14 to show the relative amplitude of the disturbance across a slice in the y-direction.

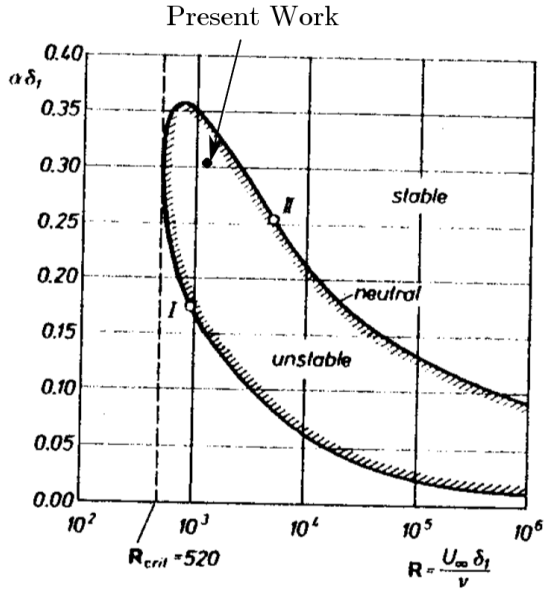

Figure 6.13: Figure 16.11 reproduced from Schlichting (1979, p. 470) showing the curve of neutral stability for disturbance wavelength ($\alpha\delta_1$) as a function of Reynolds number for the boundary layer on a flat plate. The 'Present Work' marker indicates the parameters used herein.

In an attempt to replicate Schlichting's analytical results, we initiate disturbance waves by continuously applying a small sinusoidal perturbation, in the y-direction, to the free-vortex injection point (y_{inj}) of a shear-layer located at

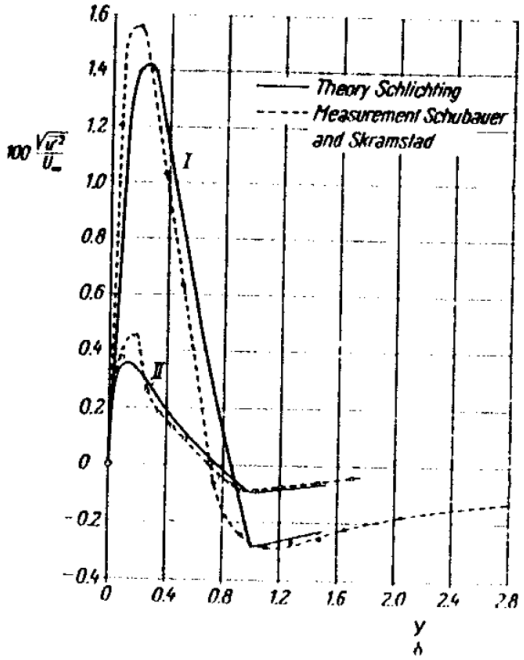

Figure 6.14: Figure 16.20 reproduced from Schlichting (1979, p. 479) showing the variation of amplitude of the velocity perturbation u' throughout the boundary layer.

$y = y_{\text{lay}}$,

$$y_{\text{inj}}(t) = y_{\text{lay}} + A \sin\left(\frac{2\pi U_{\text{lay}} t}{\lambda_{inj}}\right). \tag{6.12}$$

This is where λ_{inj} is the injection wavelength, A is the disturbance amplitude and U_{lay} is the unperturbed velocity of the flow at $y = y_{\text{lay}}$. It is noted that the chosen injection wavelength does not have to be set to that of any specific TSW as the time-stepping process introduces broad-band noise into the system. We also define the dimensionless wave-number $\bar{\alpha}$ by,

$$\bar{\alpha} = \frac{2\pi\delta}{\lambda}, \tag{6.13}$$

where λ is the wavelength and δ is the boundary-layer thickness[4]. Reynolds

[4]This relates to the boundary-layer displacement thickness by ($\delta_1 = \delta^* = .301\delta$)

number based on boundary-layer thickness is given by,

$$\mathrm{Re} = \frac{\rho_f U_\infty \delta}{\mu}. \tag{6.14}$$

Figure 6.15a shows a contour plot of the normalised velocity component in the stream-wise (U/U_∞) direction for the entire domain. This is for $A = \delta/400$, $y_{\mathrm{lay}} \approx 0.35\delta$, $\bar{\alpha}_{inj} = 1$ and $\mathrm{Re}_\delta = 4153$ at an instant where sufficient time has passed to allow the initial disturbances to convect downstream (to the right-hand side). By marking the present flow and disturbance input[5] parameters in Figure 6.13 (labelled 'Present Work') we would expect the simulation to show the growth of relatively short wavelength TSW. The disturbance injection wavelength is chosen in an attempt to excite TSW that are shorter than the most unstable for these Reynolds-number flows (that being $\bar{\alpha} \approx 0.86$) because it allows a relatively longer domain to observe the disturbance evolution with the given computational power. It is seen that the boundary-layer model consistently maintains the overall mean-flow velocity profile despite the introduction of the perturbations. However, as the small amplitude nature of the perturbation prevents any observations being made on Figure 6.15a alone, we define the stream-wise velocity perturbation u' by,

$$u'(x, y, t) = U(x, y, t) - \overline{U}(x, y), \tag{6.15}$$

where $\overline{U}(x, y)$ represents the mean-flow velocity field that is approximated numerically by time-averaging $U(x, y, t)$. Figure 6.15b shows a contour plot of the normalised velocity perturbation in the stream-wise direction (u'/U_∞). It is generally seen that a perturbation wave, of $\bar{\alpha} \approx 1$, evolves into a natural disturbance eigen-mode and then decays as it traverses the domain. However, for the given flow parameters, this dissipative behaviour does not reflect that of Schlichting (1979). Possible causes of this behaviour will be discussed after further investigation.

[5]Relating to λ_{inj}, not the broadband noise introduced through the time-stepping process.

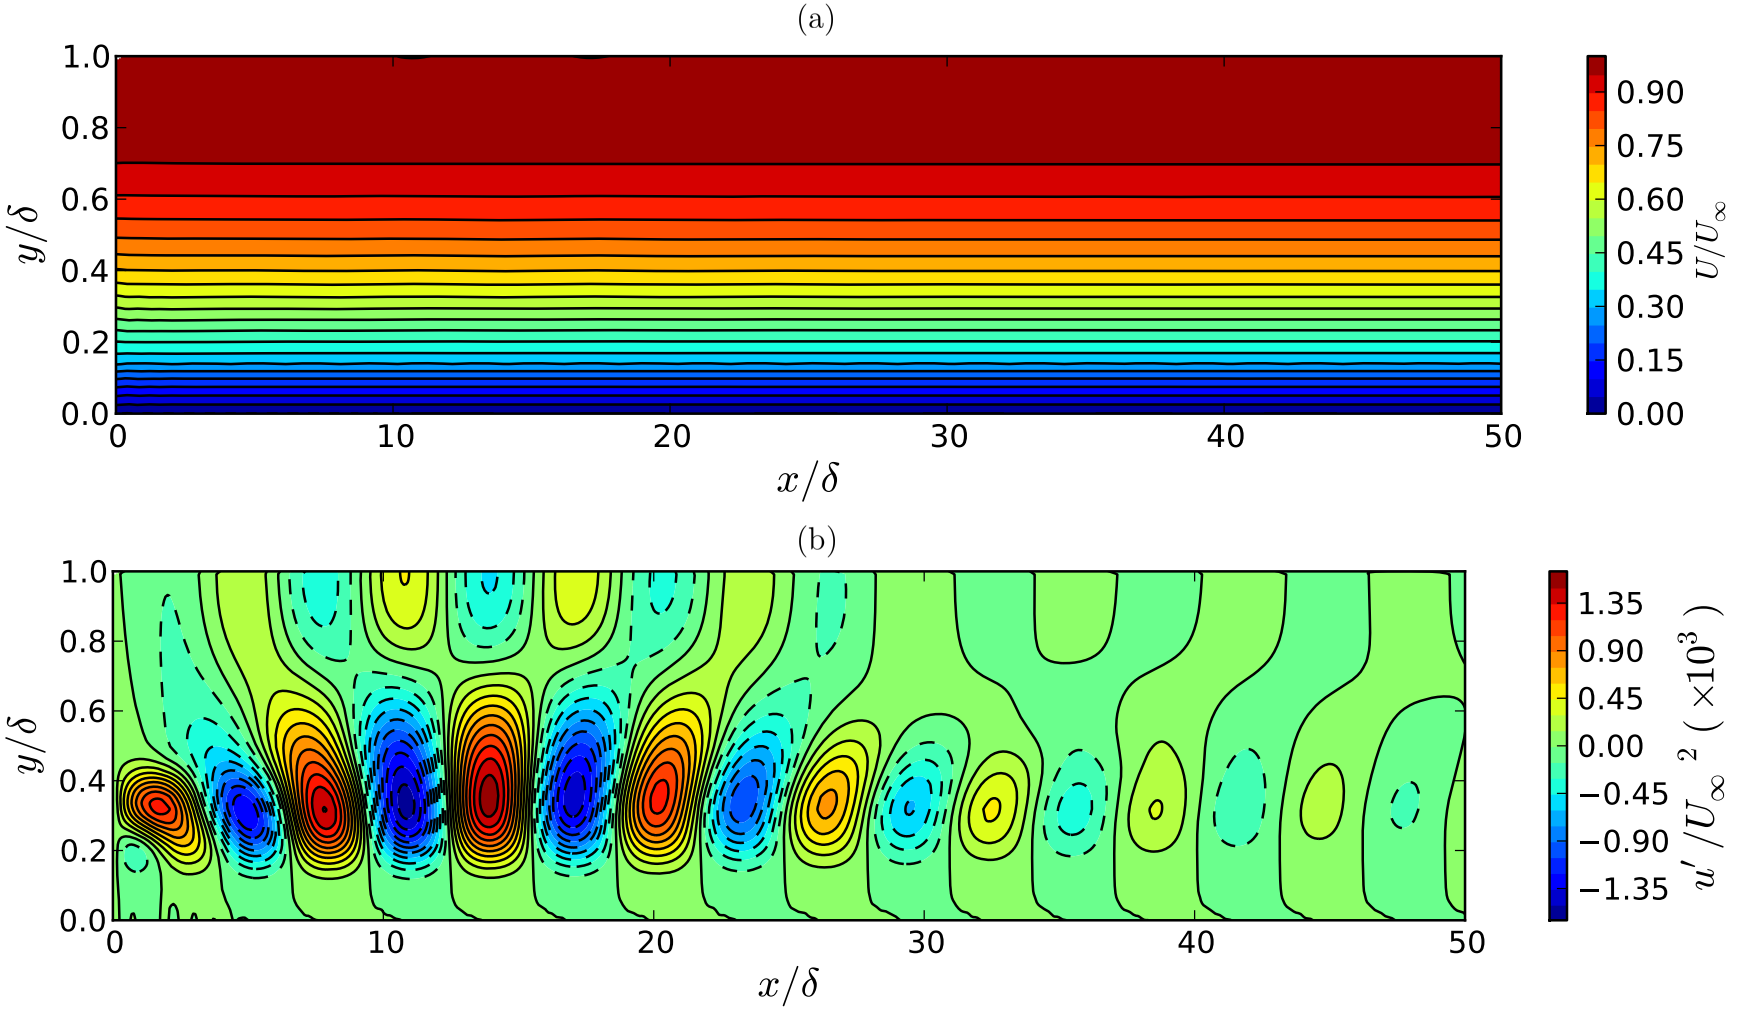

Figure 6.15: Instantaneous contour plots at $t' = 32.7$s for the time-dependent boundary-layer flow over a flat plate with $\mathrm{Re}_\delta = 4153$ for a) the normalised U velocity field and b) the normalised perturbation velocity (u'/U_∞) field.

Figure 6.16 shows both perturbation u' and v' at a constant level of $y = 0.35\delta$ throughout the domain. It is noted that the time-instant shown had sufficient time to allow the transient effects of the simulation initial conditions to convect outside of the computational domain (evidenced by Figure 6.17). This shows

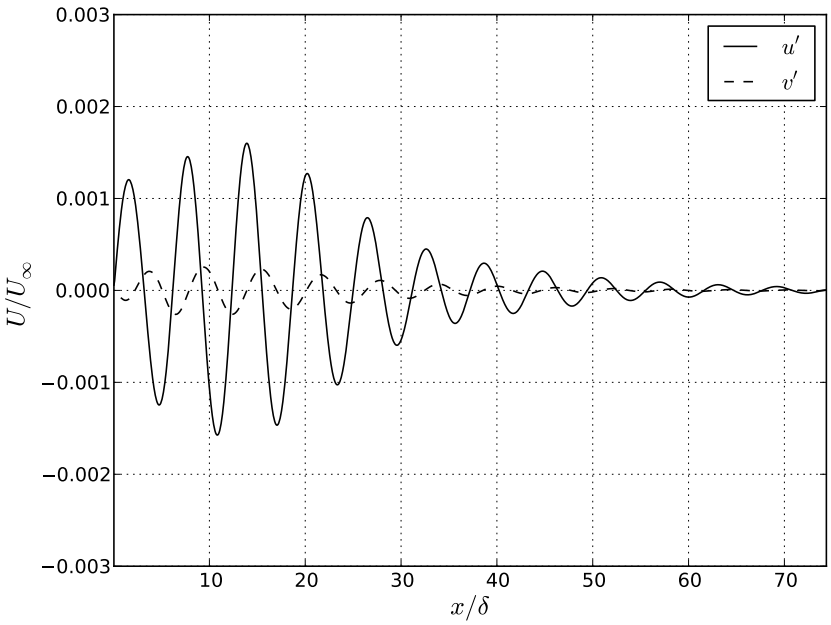

Figure 6.16: The normalised U and V perturbation velocities at a height of $y = 0.35\delta$ for a single time-instant ($t' = 32.7s$) of the time-dependent boundary-layer flow over a flat plate with $\mathrm{Re}_\delta = 4153$.

that the perturbation wave undergoes spatial growth for $x \lesssim 14\delta$. However, once the injected disturbance settles into the system eigen-mode, only decay is evident. Figure 6.17 shows the dependence of perturbation u' and v' at position $(x, y) = (14\delta, 0.35\delta)$ in non-dimensional time. This shows that the disturbance wave is temporally stable as predicted by the classical hydrodynamic stability analysis (Orr-Sommerfeld).

The energy production of a TSW within the boundary-layer is linked to the term $\overline{u'v'}\dfrac{\partial U}{\partial y}$ (Domaradzki and Metcalfe, 1987). Thus, for a disturbance wave to become unstable requires a non-zero product of the perturbation-velocities, $u'v'$, that can exceed the viscous dissipation. This can only occur where the action of viscosity causes a relative shift in the phase difference (from the neutral 90° case)

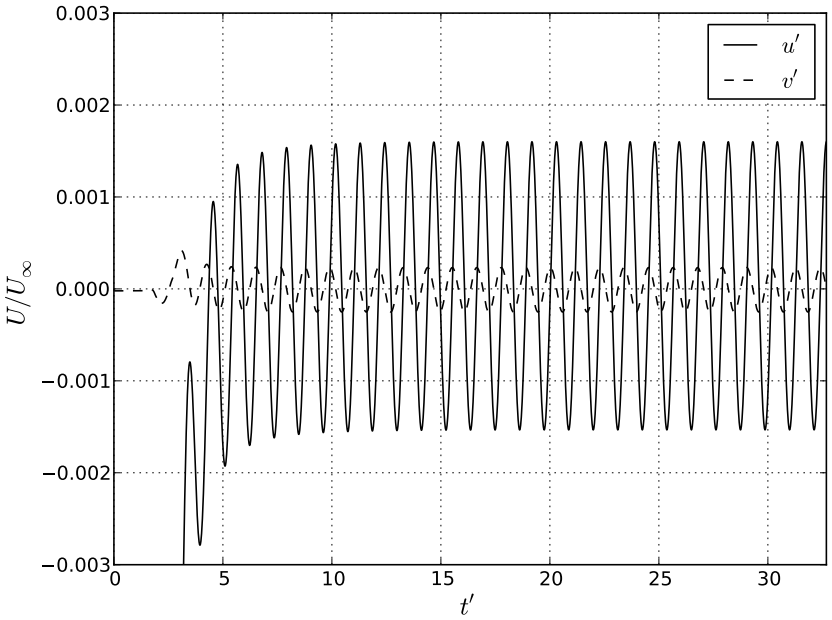

Figure 6.17: The normalised U and V perturbation velocities in time at $(x, y) = (14\delta, 0.35\delta)$ for the time-dependent boundary-layer flow over a flat plate with $\mathrm{Re}_\delta = 4153$.

of the two velocity components.

The spatial dependence of the perturbation-velocity product $(u'v')$ is shown in Figure 6.18. This shows that for the disturbance growth region, the approximated average (over $x = 3\delta$) is negative, representing the perturbation u' signal leading that of v'. Once the injected disturbance has settled, the phase relationship transitions to lagging and asymptotes to the neutral phase for $x \to \infty_+$. Figure 6.19 shows the dependence of the perturbation-velocity product at position $(x, y) = (14\delta, 0.35\delta)$ on non-dimensional time. At this inspection point the wave is temporally stable. However, the non-zero time-averaged signal illustrates the convective energy-growth of the boundary-layer at this location.

We investigate the disturbance mode-shape in the y direction by showing the stream-wise perturbation for a slice at $x = 14\delta$ in Figure 6.20. Qualitative agreement to Figure 6.14 is demonstrated with the largest peak in disturbance occurring near the wall at $y = 0.35\delta$, while a smaller and opposing peak occurs higher in the boundary-layer at $y = \delta$ and tends towards zero for $y \to \infty_+$.

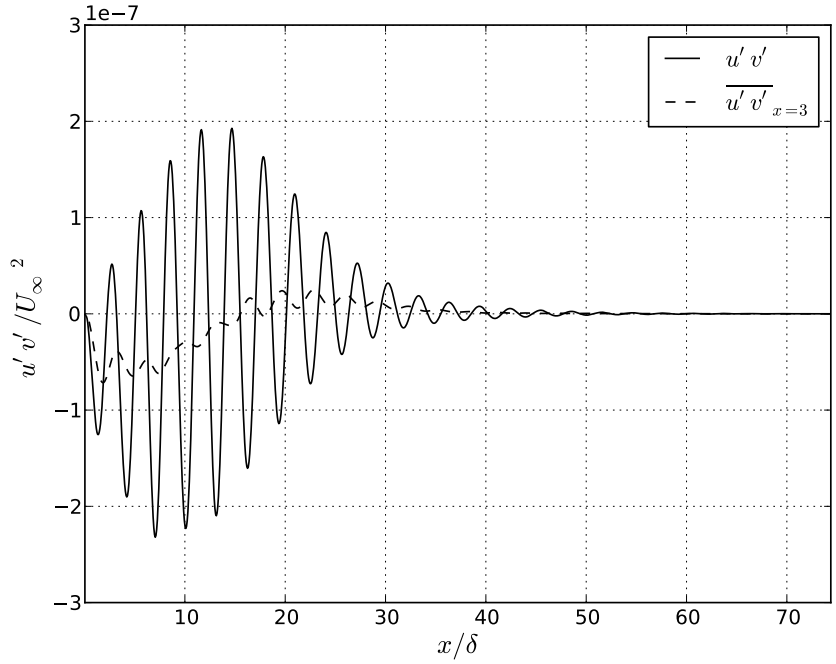

Figure 6.18: The normalised perturbation-velocity product $u'v'/U_\infty$ at a height of $y = 0.35\delta$ for a single time-step ($t' = 32.7$) of the time-dependent boundary-layer flow over a flat plate with $\mathrm{Re}_\delta = 4153$.

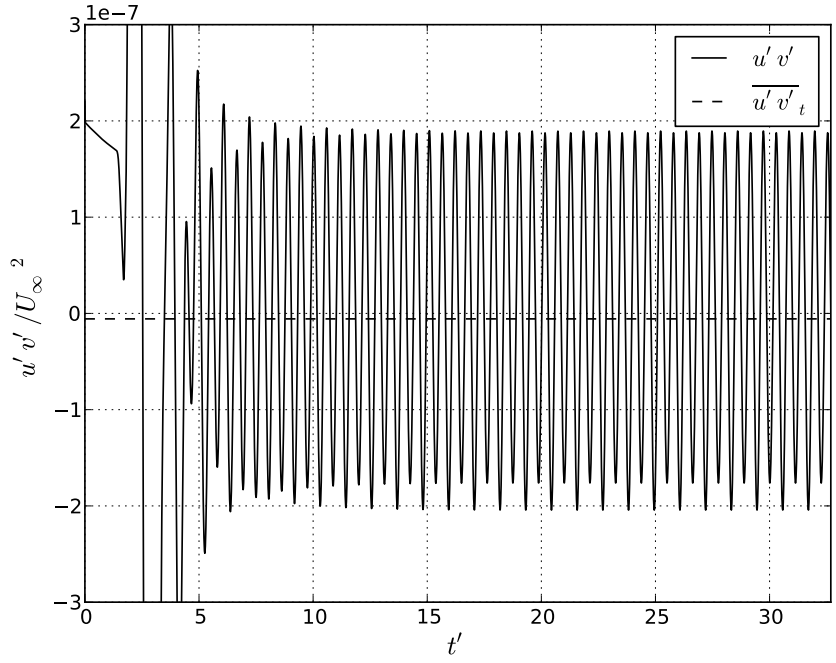

Figure 6.19: The normalised perturbation-velocity product $u'v'/U_\infty$ at $(x, y) = (14\delta, 0.35\delta)$ in time for the time-dependent boundary-layer flow over a flat plate with $\mathrm{Re}_\delta = 4153$.

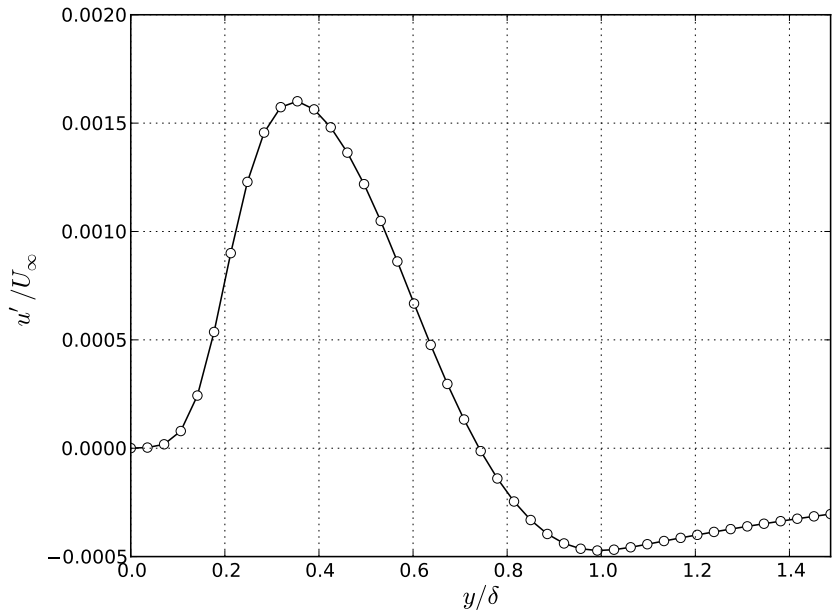

Figure 6.20: The normalised U perturbation velocity at a slice of $x = 14\delta$ for a single time-step of the time-dependent boundary-layer flow over a flat plate with $\mathrm{Re}_\delta = 4153$.

Quantitative (as opposed to qualitative) agreement with Schlichting (1979) for the evolution of TSW has not been demonstrated. Conducting investigations into the drag-reducing properties of compliant-panel technologies (via boundary-layer transition delay, similar to the analytical work of Carpenter and Garrad, 1985) will therefore remain the focus of future works. However, the ability of the model to consistently maintain the boundary-layer mean-flow velocity profile under perturbed conditions will allow the capture FSI dynamics of interest to the Engineering community, such as those involving flow-induced buckling (divergence) or travelling-wave flutter. Furthermore, the lack of agreement for elements such as the precise location of eigen-mode maxima and the general unstable wave behaviour was unsurprising considering the different approach between the eigen-analysis of the Orr-Sommerfeld equations and our time-dependent numerical method[6].

[6]We solve the DVM as a time-dependent numerical problem in a finite-domain where a discrete disturbance exists only in the computational domain after being started from a fixed (x,y) position.

6.4 Summary

We have presented a model for studying perturbations to a boundary-layer flow over a finite-wall that can be of arbitrary shape. The previous wall-bounded potential-flow model was extended to include the effects of viscosity using the Discrete Vortex Method (DVM) with stacked layers of overlapping Gaussian vortices that captured the dynamics of a boundary-layer flow with an arbitrary mean-velocity profile. The no-slip condition was strictly enforced using the Boundary Element Method (BEM) and vorticity injection while effects of viscous diffusion were modelled using a core-spreading method. Time-dependent simulations were conducted using an explicit time-stepping method while a vorticity rediscretisation scheme, based on radial-basis function interpolation and implemented using the Generalised Minimum Residual Method (GMRES), ensured the consistency of the Lagrangian grid.

An accurate reproduction of an unperturbed laminar Blasius boundary-layer velocity profile was demonstrated. Attempts to validate the stability behaviour of a Tollmien-Schlichting (TS) wave against prior analytical works (based upon the Orr-Sommerfeld equation) were met with partial success. Excellent qualitative agreement was demonstrated for the disturbance eigen-mode while localised growth/decay of the disturbance wave was exhibited as a result of shifts in the phase relationship of perturbation velocity components. Attempts to achieve complete quantitative agreement were not demonstrated (or sought) due to the differences in the present model and the analytical works.

The inherent ability of the model to maintain accurate resolution of the boundary-layer mean-flow profile, under perturbed conditions and over the entire domain, will allow the reliable study of various modes of fluid-structure interaction (FSI). This will include cases such as those involving flow-induced buckling (divergence) and travelling-wave flutter. However, investigations into the effect of wall flexibility upon TSW behaviour (for the purpose of boundary-layer transition delay) will not be conducted herein and will remain the objective of future works.

Chapter 7

Viscous Fluid-Structure Interaction

We extend the model of Chapter 6 to include the structural dynamics of a Kramer-type compliant wall as shown in Figure 7.1. The compliant-wall is modelled as a

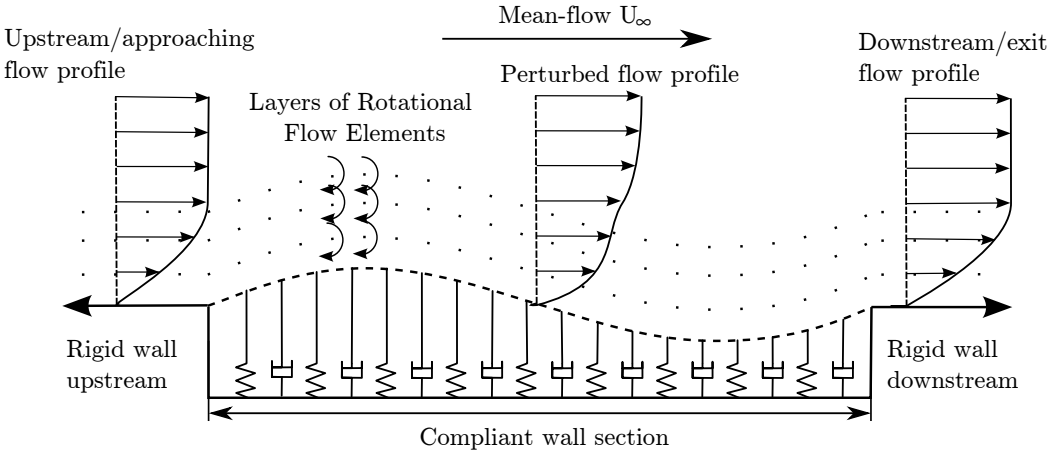

Figure 7.1: Schematic of the compliant wall interaction with a boundary-layer flow.

nonlinear Euler-Bernoulli beam with a spring-damper foundation and is governed by,

$$\rho h \frac{\partial^2 \eta}{\partial t^2} + B \frac{\partial^4 \eta}{\partial x^4} - T_I(\eta) \frac{\partial^2 \eta}{\partial x^2} + D \frac{\partial \eta}{\partial t} + K\eta = F(x,t). \tag{7.1}$$

While much of the fluid-wall coupling methodology remains the same as in Chapter 4, we present the necessary modifications to the potential-flow fluid-structure interaction (FSI) scheme before investigating the instabilities of the complete boundary-layer flow FSI system.

7.1 Flow Pressure

The pressure induced at the surface by the viscous boundary-layer flow can no longer be evaluated using the unsteady Bernoulli equation because of the inviscid-flow assumption in its formulation. We instead obtain the pressure at the wall by direct numerical integration of the fluid-momentum equations (for 2D incompressible flow) in the y-direction where,

$$\rho\left(\frac{\partial V}{\partial t}+U\frac{\partial V}{\partial x}+V\frac{\partial V}{\partial y}\right)=-\frac{\partial p}{\partial y}+\mu\left(\frac{\partial^2 V}{\partial x^2}+\frac{\partial^2 V}{\partial y^2}\right). \tag{7.2}$$

We integrate $\partial p/\partial y$ from $y=\infty_+$ down to the midpoint of a BEM-panel at $y=y_p$. However, as integrating from $y=\infty_+$ is difficult, we separate the integration into two parts based on the inviscid outer-flow region located at $y\geq\delta+y_p$, where δ is the boundary-layer thickness. By only looking at the change in pressure (Δp_μ) from the upstream flow we have,

$$\Delta p_\mu=\underbrace{-\frac{\rho_f {U_\infty}^2}{2}}_{\text{upstream}}+\underbrace{\int_{\infty_+}^{\delta+y_p}\frac{\partial p}{\partial y}dy}_{\text{potential-flow}}+\underbrace{\int_{\delta+y_p}^{y_p}\frac{\partial p}{\partial y}dy}_{\text{viscous-flow}} \tag{7.3}$$

$$=\int_{\delta+y_p}^{y_p}\left(-\rho_f\left(\frac{\partial V}{\partial t}+U\frac{\partial V}{\partial x}+V\frac{\partial V}{\partial y}\right)+\mu\left(\frac{\partial^2 V}{\partial x^2}+\frac{\partial^2 V}{\partial y^2}\right)\right)dy + \Delta p_{\phi,\delta}. \tag{7.4}$$

This is where $\Delta p_{\phi,\delta}$ is the relative pressure at the inviscid outer-flow boundary and is obtained using the unsteady Bernoulli equation in Equation (4.4). The pressure change arising from the viscous region of flow is found by numerically integrating down through the boundary layer to the wall. Sufficient samples of

the velocity-field, $U(x, y)$ and $V(x, y)$, are taken to enable the approximation of all spatial differentials using a 2nd-order accurate central-difference method. The advantage of using direct integration of the y-momentum equation is that it can be readily applied to the potential-flow case for theoretical/numerical validation. This is achieved using the viscous-flow model by turning off the influence of all rotational (vortex) elements.

Figure 7.2 shows a comparison of the coefficient of pressure, $C_p = \Delta p / 2\rho_f {U_\infty}^2$, for the integration scheme (o) and the Bernoulli equation (solid), induced by an axial potential-flow over the surface of a static wall with length L and a small displacement in the fundamental-mode shape ($--$). This was obtained

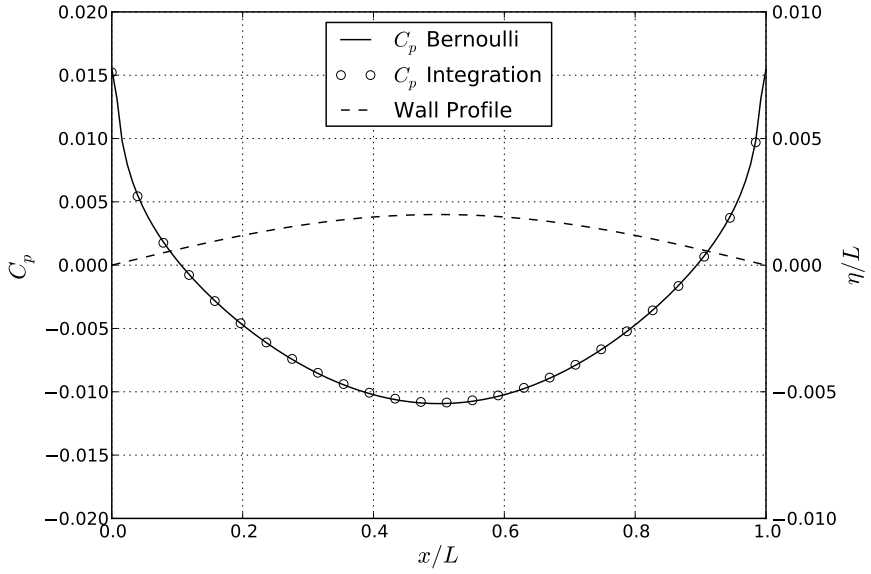

Figure 7.2: Comparison of pressure coefficient for a potential-flow over a statically deformed wall.

using $N = 128$ finite-difference mass nodes and 50 evenly spaced velocity-field samples in the y-direction. The results show excellent agreement between the two methods, evidenced by an RMS-difference of 0.9%.

For the same wall profile, Figure 7.3 shows a comparison of the pressure induced by a boundary-layer flow (using the integration method) with $\delta = L/10$ at an instant when sufficient time has passed to allow initial disturbances to convect downstream. The figure shows that when compared to the potential

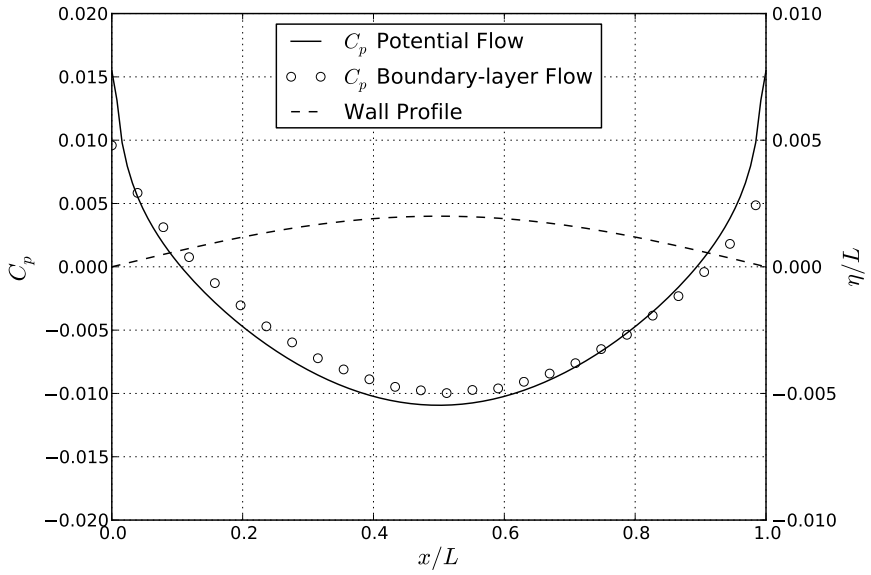

Figure 7.3: Pressure coefficient for a boundary-layer flow over a statically displaced wall.

flow the introduction of the boundary-layer causes a shift in the pressure profile towards the right-hand-side (down-stream). The boundary-layer flow also causes a reduction in the magnitude of the potential-flow pressure (Δp_Φ) that can be measured with the scaling factor β,

$$\beta = \sqrt{\frac{\sum_{i=1}^{N} \left(\Delta p_{\mu,i}\right)^2}{\sum_{i=1}^{N} \left(\Delta p_{\Phi,i}\right)^2}}. \tag{7.5}$$

It is noted that for the present flow parameters, the pressure scaling ratio is $\beta = 0.86$.

7.2 Flow Coupling and System Solution

The interfacial pressure arising from the boundary-layer flow is coupled to the compliant wall through the forcing term, $F(x,t) = -\Delta p_\mu(\ddot{\eta}, \dot{\eta}, \eta)$, in Equation (3.4) as,

$$\rho h \frac{\partial^2 \eta}{\partial t^2} + B \frac{\partial^4 \eta}{\partial x^4} - T_I\left(\eta\right) \frac{\partial^2 \eta}{\partial x^2} + D \frac{\partial \eta}{\partial t} + K\eta = -\Delta p_\mu(\ddot{\eta}, \dot{\eta}, \eta). \tag{7.6}$$

Unlike previously, the full viscous pressure cannot be implicitly coupled and solved using the Newton-Krylov (NK) framework. This is due to the Discrete Vortex Method's (DVM) extremely high (relative to other components) computational intensity. We solve the system numerically by decoupling the left and right hand sides of Equation (7.6) into a semi-implicit (SI) solution method (previously introduced in Chapter 4). To avoid the inherent instability of a pure semi-implicit method we take a hybrid approach by solving the left hand side of the equation implicitly (with the NK) while the right hand side is treated as a constant that must be iterated until convergence. In scenarios with high-density (relative to the plate density) fluid loading, the convergence behaviour of this scheme rapidly deteriorates. Thus, we apply a conditioner, $\beta\Delta p_{\Phi'}$, to both sides of the equation that represents the scaled fluid loading of the corresponding inviscid-flow case. This is evaluated using the unsteady Bernoulli equation with all rotational/viscous elements turned off. This improves the conditioning of the semi-implicit iterations, increasing the numerical stability and convergence rate. The selection of this conditioner is based upon reducing the sensitivity of the semi-implicitly coupled pressure to the wall acceleration[1] whereby,

$$\left|\frac{\partial\left(\Delta p_{\mu}-\beta\Delta p_{\Phi'}\right)}{\partial\ddot{\eta}}\right| << \left|\frac{\partial\Delta p_{\mu}}{\partial\ddot{\eta}}\right|. \tag{7.7}$$

To further improve convergence rates we adjust the scaling factor (β) dynamically based on the RMS of the inviscid and viscous flow pressures of the previous time-step,

$$\beta_t = 1.5\sqrt{\frac{\sum_{i=1}^{N}\left(\{\Delta p_{\mu,i}\}_{t-1}\right)^2}{\sum_{i=1}^{N}\left(\{\Delta p_{\Phi',i}\}_{t-1}\right)^2}}. \tag{7.8}$$

A factor of safety of 1.5 has also been included to ensure that the time-lag in β does not result in an under-prediction of the boundary-layer pressure which can lead to a numerical instability in the semi-implicit solution.

[1]We interchangeably use $\ddot{\eta} = \dfrac{\partial^2\eta}{\partial t^2}$ where necessary.

The resulting NK-SI hybrid system with improved conditioning can be expressed as,

$$\underbrace{\left[\rho h\frac{\partial^2\eta}{\partial t^2}+B\frac{\partial^4\eta}{\partial x^4}-T_I\left(\eta\right)\frac{\partial^2\eta}{\partial x^2}+D\frac{\partial\eta}{\partial t}+K\eta+\beta\Delta p_{\Phi'}\right]_t}_{\text{Implicit NK}}= \underbrace{\left[-(\Delta p_\mu-\beta\Delta p_{\Phi'})\right]_{t^*}}_{\text{Explicit}} \tag{7.9}$$

The right-hand side term now represents a viscous correction that is applied to the implicitly coupled potential flow pressure. This is updated after every application of the NK method until convergence is reached ($t^* = t$).

7.3 Illustrative Results – Divergence

7.3.1 Divergence Onset

The introduction of the boundary-layer serves to modify the potential flow results in a reduction of the overall pressure experienced by the wall. We investigate the impact of this reduction on the divergence-onset flow speed for a boundary-layer flow over a simple panel of length L. This is modelled by setting the stiffness (K) and damping (D) parameters in the compliant-wall foundation to zero. The onset speed is measured using the previously presented algorithm (in Section 4.4) that is based on iteratively finding the divergence-onset flow speed using a binary-search and Newton-Raphson method. Figure 7.4 shows the dependence of the divergence-onset (critical) flow speed on the relative boundary-layer thickness for a tolerance of $\epsilon_{\Lambda,\text{cr}} = 1\%$. The non-dimensional critical flow speed has been normalised to the potential-flow value of $\Lambda_{\Phi,cr} = 39.6$ (with a tolerance of 1%). It is seen that the introduction of the boundary-layer causes an increase in the divergence-onset speed that is dependent on the boundary-layer thickness. The highest critical speeds are seen with thicker boundary-layers due to the greater potential-flow pressure reduction (a lower β). By reducing the boundary-layer thickness towards zero, the critical flow speed approaches

that of the potential-flow case. However, it is noted that an exact asymptote towards the potential flow case cannot be demonstrated due to limitations that prevent adequate discretisation of very thin boundary layers (a topic discussed in Chapter 8).

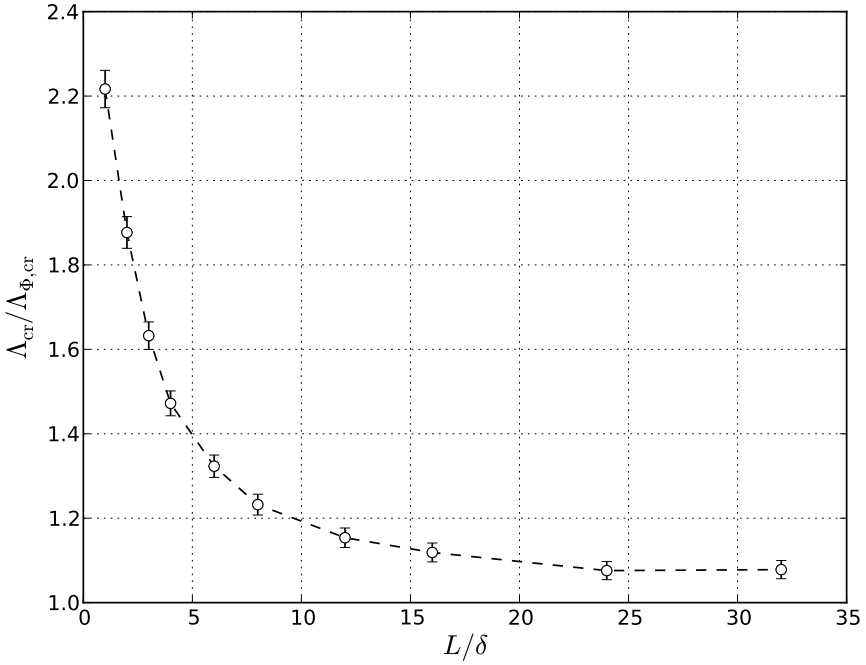

Figure 7.4: Dependence of divergence-onset flow speed on relative boundary-layer thickness.

7.3.2 Nonlinear Divergence Oscillations

Validation

Having established the validity of the boundary-layer pressure evaluation in a static framework (see Section 7.1), we reproduce the potential flow results of Section 4.5 to ensure that the evaluation of the time-dependent pressure is also valid. The system and spatial/temporal discretisation parameters are set to that of Section 4.5 whilst the integration pressure method replaces the unsteady-Bernoulli equation. Figure 7.5 shows the displacement of the panel-midpoint in time, illustrating the energy-stable nonlinear limit-cycle oscillations of the system and offers excellent agreement to the previous unsteady-Bernoulli based results. Figure 7.6

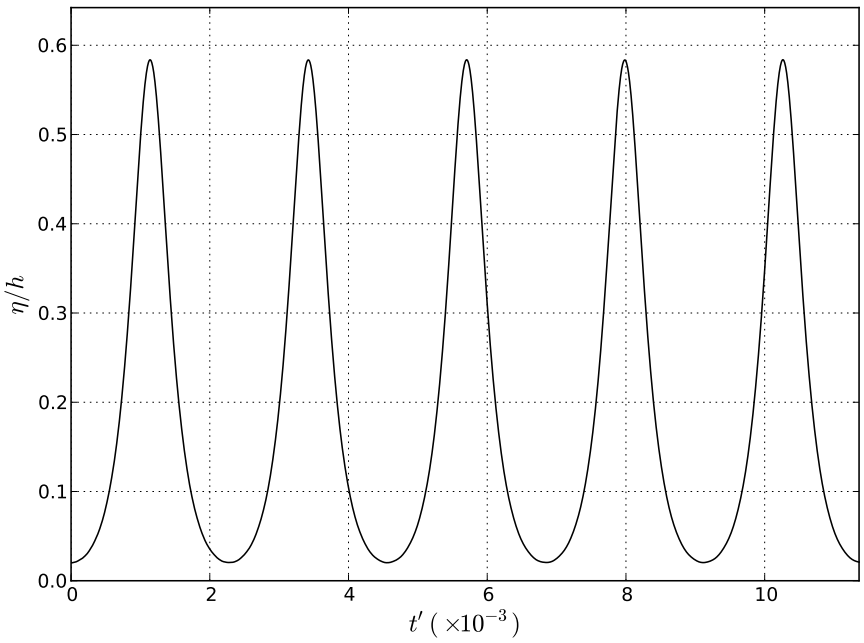

Figure 7.5: Panel-midpoint displacement in time for divergence instability in a potential-flow, using the boundary-layer flow pressure evaluation method.

shows a plot of the panel-midpoint velocity and position for the entire simulation ($t' \approx 0 \rightarrow 18 \times 10^{-3}$). This reveals that the nonlinear oscillations of the panel are neutrally stable with $\eta > 0$ throughout and the panel-midpoint orbits an attractor centred at $\eta/h \approx 0.4$. This provides excellent agreement to the results of Lucey et al. (1997b).

Low Initial Amplitude

We investigate the nonlinear limit-cycle behaviour of the FSI system for a boundary-layer flow at a post-critical speed. The system parameters are set to match those of the previous potential-flow system with a boundary layer thickness of $\delta/L = 0.1$ and flow viscosity set to achieve a Reynolds number of $Re_{\delta^*} = 2970$. Figures 7.7a and 7.7b show the nonlinear motion, through one cycle, of panel divergence occasioned by a boundary-layer flow. The overall unstable fluid-structure mode is seen to be dominated by the fundamental. This behaviour is similar to the potential-flow results albeit with a reduced maximum amplitude. Figure 7.8 shows the displacement of the panel-midpoint in time. This remains similar to the non-

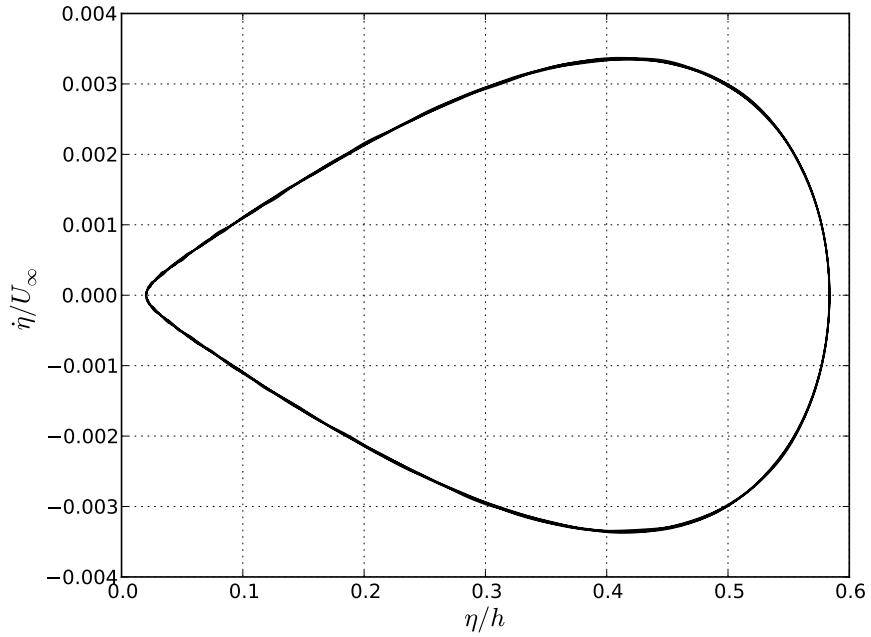

Figure 7.6: Orbit generated by the instantaneous position and velocity of the panel-midpoint in time ($t' \approx 0 \rightarrow 18 \times 10^{-3}$) for the nonlinear divergence occasioned by a potential flow and evaluated using the integration pressure method.

linear oscillations of the potential-flow case, however, the system is no-longer neutrally stable. The inclusion of the boundary-layer influences the long-time oscillations of the system in a manner similar to weak structural damping due to the dissipative effects of viscosity. For long simulation times the system would be expected to reach a steady-state equilibrium position where the restorative forces in the panel are balanced by the force of the deformed flow field (Lucey et al., 1997b). The panel-midpoint velocity and position for the entire simulation ($t' \approx 0 \rightarrow 6 \times 10^{-3}$) are shown in Figure 7.9. Nonlinear oscillations of the panel reveal an orbit that is based around an attractor centred at $\eta/h \approx 0.3$. In comparison to the potential flow case, the orbit is of a smaller radius and undergoing constant decay as a result of the dissipative effects of the boundary-layer flow. It is also seen that higher-order frequencies are present at earlier simulation times due to the impulsive release of the panel in a deformation shape that differs from the system's fundamental mode shape. The potential-flow scaling factor (β), implicitly-coupled potential-flow pressure ($\Delta p_{\Phi'}$) and the boundary-

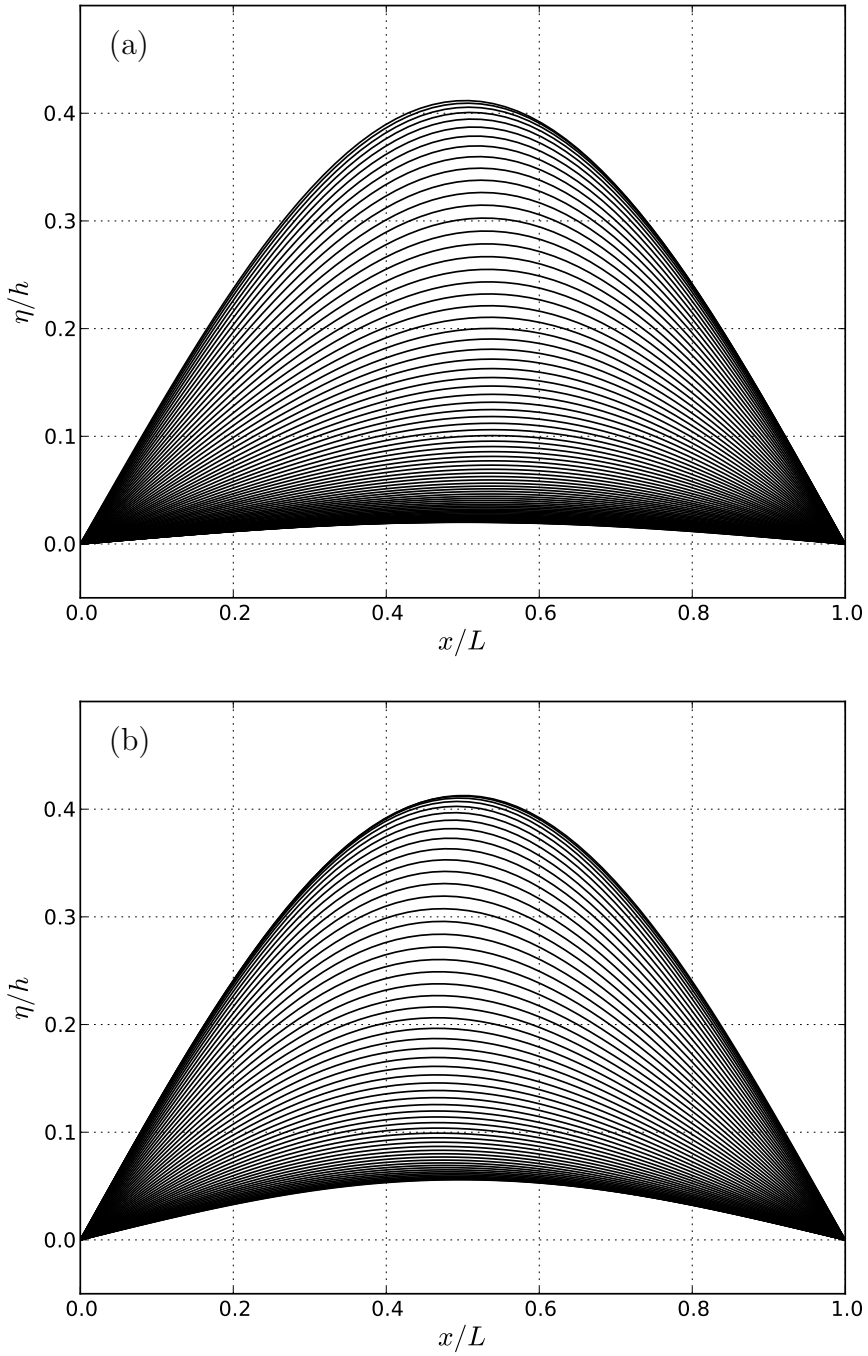

Figure 7.7: Simulation of non-linear divergence instability in a boundary-layer ($Re_{\delta^*} = 2970$, $L/\delta = 10$) water flow over an aluminium ($\rho_f/\rho_w = 0.385$) plate ($h/L = 0.01$) at $\Lambda = 61$. Series of panel deformations in time for (a) growth and (b) decay phase of a single cycle ($t' \approx 0 \rightarrow 2.5 \times 10^{-3}$).

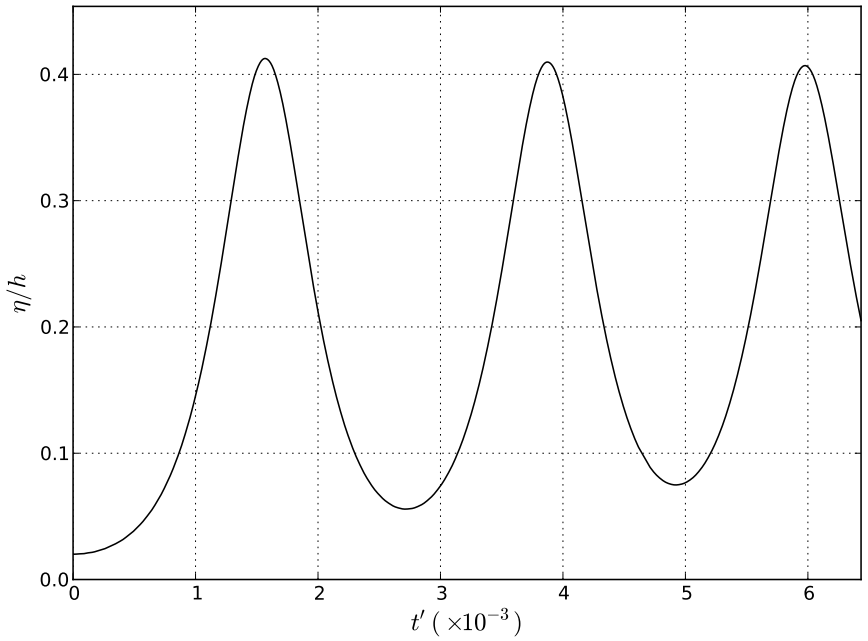

Figure 7.8: Panel-midpoint displacement in time for divergence instability in a boundary-layer flow.

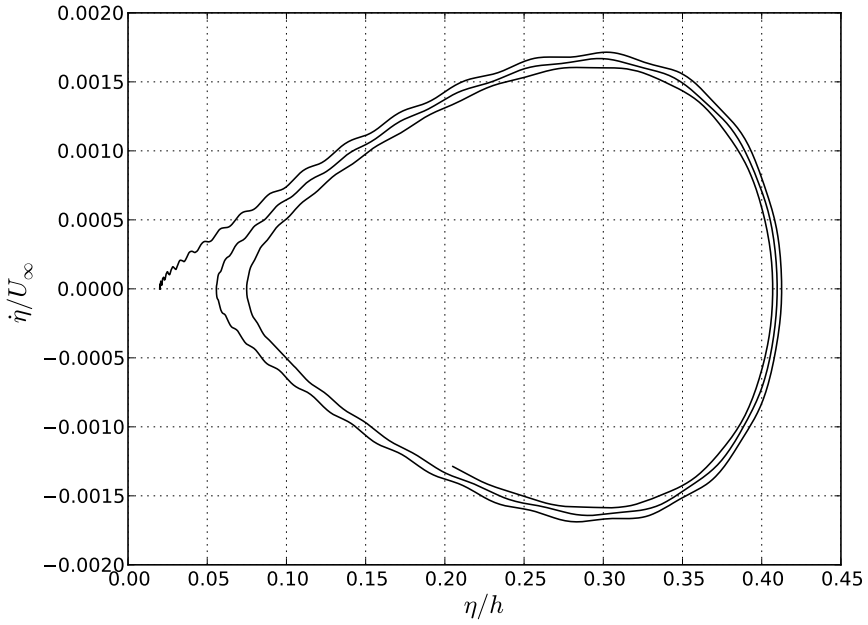

Figure 7.9: Orbit generated by the instantaneous position and velocity of the panel-midpoint in time ($t' \approx 0 \rightarrow 6 \times 10^{-3}$) for the nonlinear divergence of a panel in a boundary-layer flow.

layer pressure (Δp_μ) at the panel-midpoint in time are shown in Figure 7.10. Once the short-lived disturbances generated by the release of the panel have been

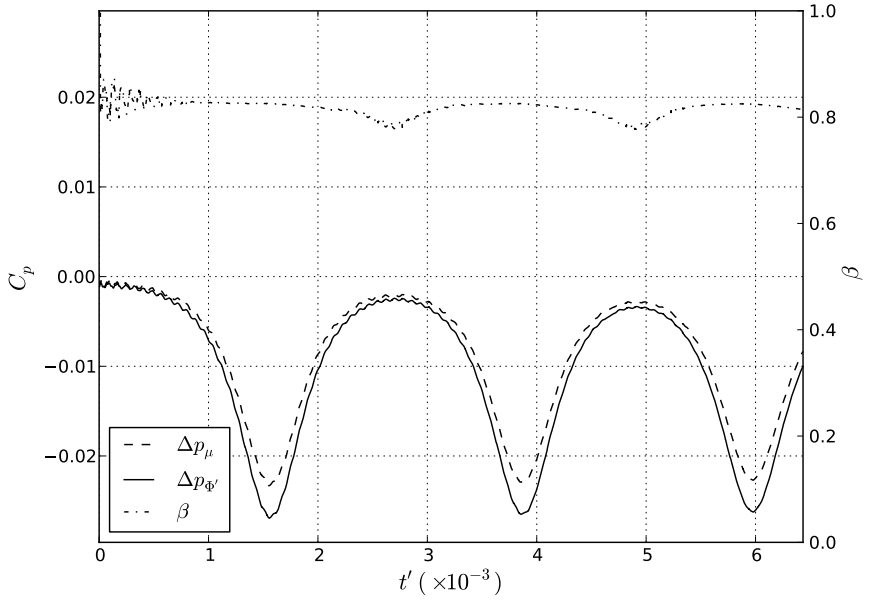

Figure 7.10: Scaling ratio {-.} (β) and coeffecient of pressure for; the boundary-layer {solid} (Δp_μ) and potential flow pressure {- -} ($\Delta p_{\Phi'}$), at the panel-midpoint in time for the nonlinear divergence of a panel in a boundary-layer flow.

dissipated, the boundary-layer pressure slowly varies in the vicinity of 80%−RMS of the potential-flow pressure. By extending upon Equation (4.30) to include the spring-foundation and effects of nonlinear tension, we define the total compliant-wall energy by,

$$E_{\text{wall}} = \frac{1}{2}\int_0^L \left(\underbrace{B\left(\frac{\partial^2 \eta}{\partial x^2}\right)^2}_{\text{Bending}} + \underbrace{\rho h \left(\frac{\partial \eta}{\partial t}\right)^2}_{\text{Kinetic}} + \underbrace{K\eta^2}_{\text{Spring}} + \underbrace{T_I(\eta)\left(\frac{\partial \eta}{\partial x}\right)^2}_{\text{Tension}} \right) \mathrm{d}x. \qquad (7.10)$$

Figure 7.11 shows the variation of the panel's total[2], bending and tension energy in time. The kinetic energy term has not been shown due it its negligible magnitude when compared to the dominant bending and tension terms. This is a result of the static-like nature of the divergence instability. The dissipative effects of the boundary-layer are seen to cause the maximum energy state of the panel to slowly

[2]As $K = 0$ for a simple panel, the foundation spring stiffness term has been neglected.

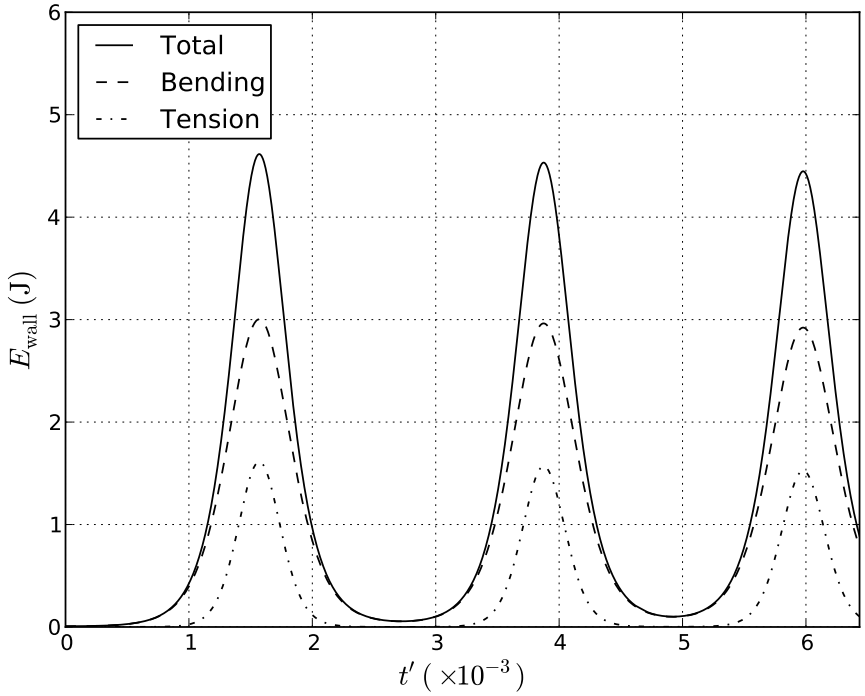

Figure 7.11: Variation of the total, bending and tension energy in time for the nonlinear divergence of a panel in a boundary-layer flow.

reduce while the minimum energy state is increasing. By allowing sufficient time to pass, the system would be expected to settle into a static buckled position. Figure 7.12 shows the total power (P_{fluid}) transferred from the boundary-layer flow to the panel in time where,

$$P_{\text{fluid}} = \int_0^L (-\Delta p_\mu \dot{\eta})\, \mathrm{d}x. \tag{7.11}$$

The fluid is seen to alternate between states of adding and removing energy from the panel. By looking at the time-average[3] (over the period $t' = 1.37 \times 10^{-3}$ to $t' = 6.26 \times 10^{-3}$) of the power we see a bias towards the negative axis which results in dissipation of panel energy with every nonlinear oscillation. As the panel approaches its long-time statically buckled state the range of power fluctuation also reduces and will be expected to converge towards a zero value.

[3]Calculated over the time interval by the sum of values at each time-step divided by the number of time-steps.

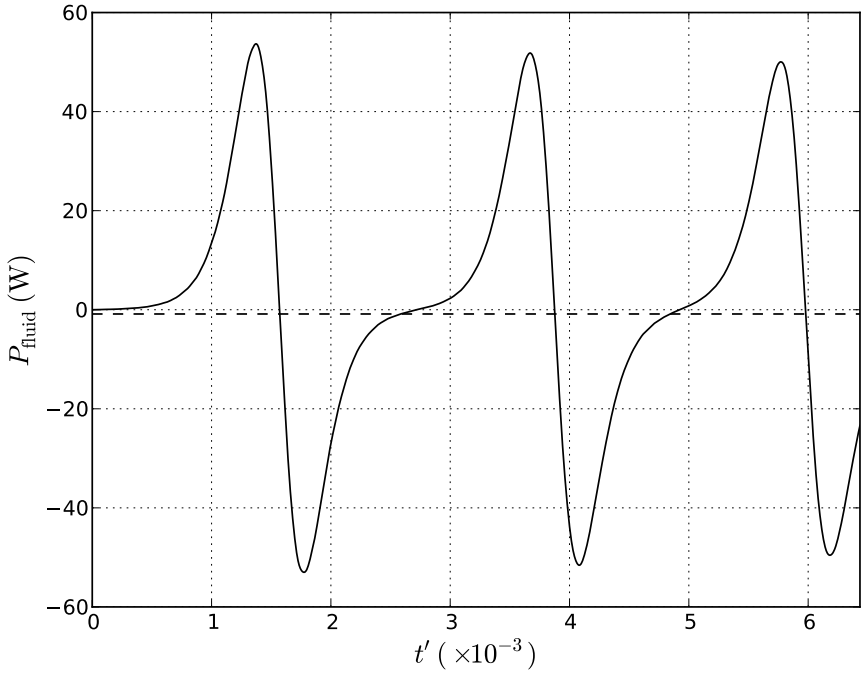

Figure 7.12: Total {solid} and average {- -} power transferred from the boundary-layer flow into the panel in time whilst undergoing nonlinear divergence oscillations.

High Initial Amplitude

We also investigate the changes induced on the present boundary-layer FSI system by initialising the panel with a larger peak amplitude of $A = 0.5h$. This amplitude exceeds the maximum amplitude observed previously and the estimated stable nonlinear (static) equilibrium position of $A \approx 0.3h$. Figures 7.13a and 7.13b show the motion of the panel through one full oscillation. The vibration mode is dominated by that of the fundamental with slight elements of higher-order mode and appears to be oscillating sinusoidally. Figure 7.14 shows the displacement of the panel-midpoint in time. The oscillatory behaviour of this system is vastly different to previously with the midpoint oscillating about a zero value in a sawtooth motion, whilst the dissipative effects of the boundary-layer are less pronounced. A plot of instantaneous values ($t' \approx 0 \rightarrow 6 \times 10^{-3}$) for the panel-midpoint velocity and position are shown in Figure 7.15. For the limited duration, the panel-midpoint appears to orbit two attractors centred at

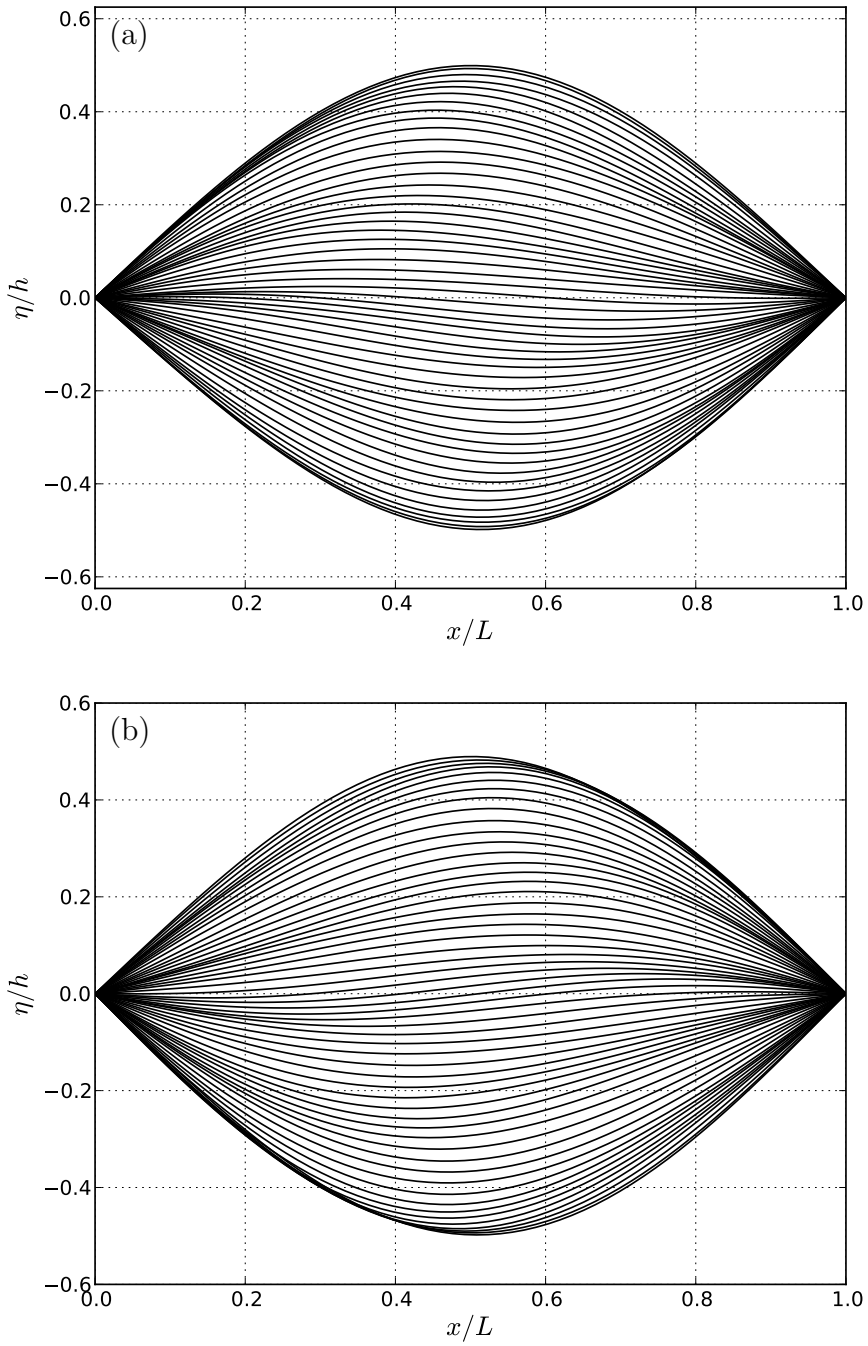

Figure 7.13: Simulation of non-linear divergence instability in a boundary-layer ($Re_{\delta^*} = 2970$, $L/\delta = 10$) water flow over an aluminium ($\rho_f/\rho_w = 0.385$) plate ($h/L = 0.01$) at $\Lambda = 61$ with a higher initial amplitude of $A = 0.5h$. Series of panel deformations in time for (a) growth and (b) decay phase of a single cycle ($t' \approx 0 \rightarrow 2$).

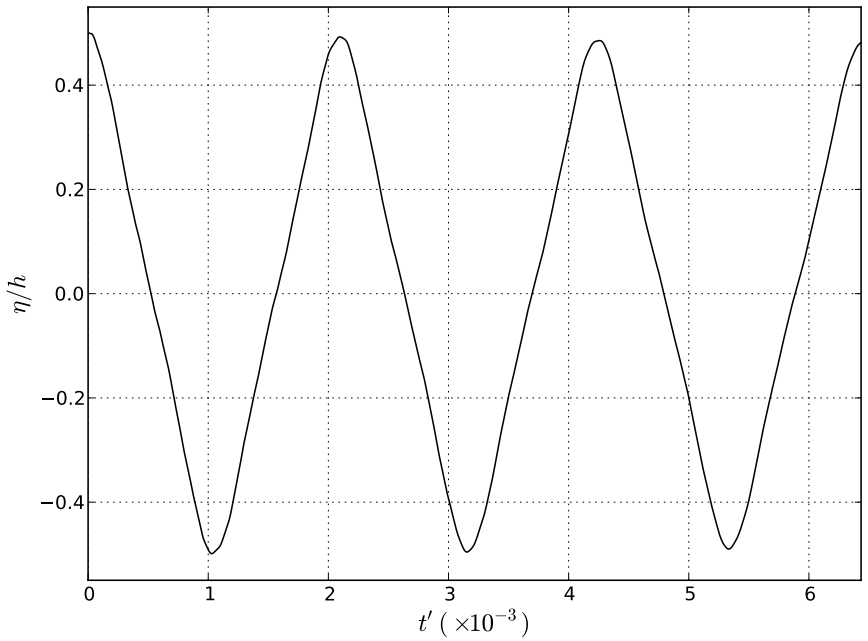

Figure 7.14: Panel-midpoint displacement in time for divergence instability in a boundary-layer flow with a large initial amplitude.

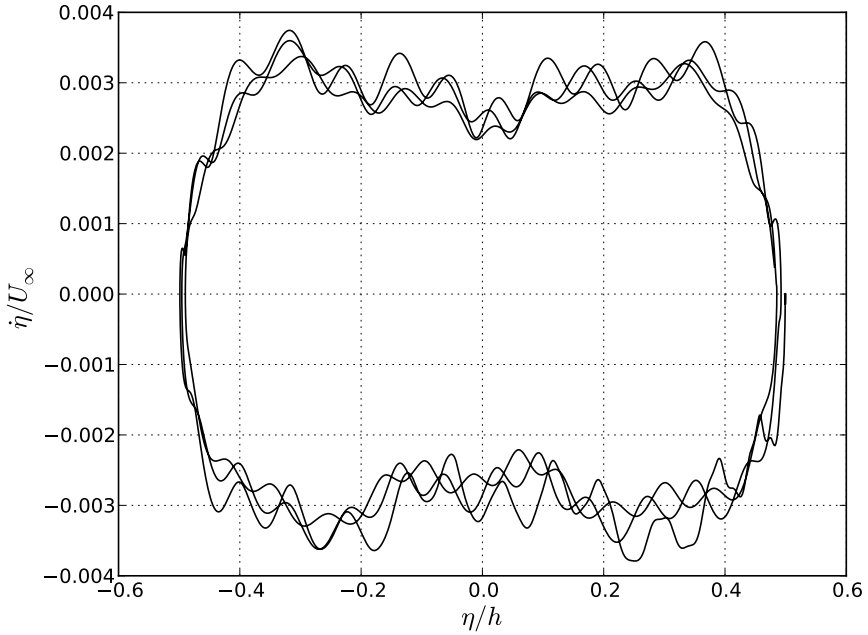

Figure 7.15: Orbit generated by the instantaneous position and velocity of the midpoint in time ($t' \approx 0 \rightarrow 6 \times 10^{-3}$) for the nonlinear divergence of a panel in a boundary-layer flow with a large initial amplitude.

$\eta/h \approx \pm 0.3$. This qualitatively agrees with the high-amplitude potential-flow results of Lucey et al. (1997b). With every cycle the overall orbit is decaying and thus reducing in radius. This would be expected to continue until it reaches the orbit radius of the low initial-amplitude case and would then switch to the motion shown in Figure 7.9. The potential-flow scaling factor (β), implicitly-coupled potential-flow pressure ($\Delta p_{\Phi'}$) and the boundary-layer pressure (Δp_μ) at the panel-midpoint in time are shown in Figure 7.16.

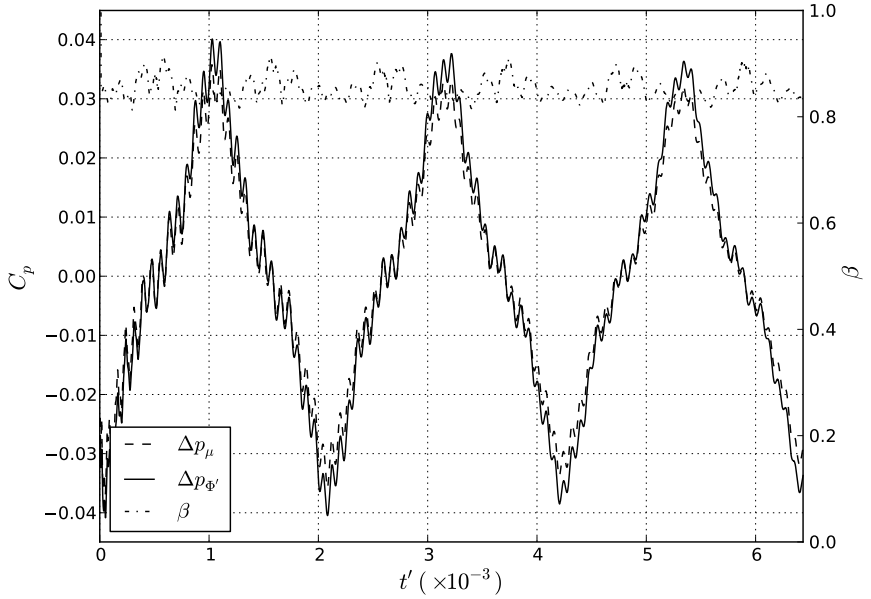

Figure 7.16: Scaling ratio {-.} (β), boundary-layer {solid} (Δp_μ) and potential flow pressure {- -} ($\Delta p_{\Phi'}$) at the panel-midpoint in time for the nonlinear divergence of a panel in a boundary-layer flow with a large initial amplitude.

Divergence Summary

Despite introducing viscous boundary-layer effects to an otherwise potential-flow compliant-wall FSI system, the underlying behaviour of nonlinear divergence oscillations in a simple panel remained mostly unchanged from the potential-flow case. It was shown that the most significant effect was through the reduced steady-state pressure which caused lower oscillation amplitudes. However, it appeared that the dissipative nature of the viscous-flow FSI system behaved in a manner similar to the inclusion of weak structural damping. Whilst this damping

would not prevent the onset of divergence instability, it appeared to offer a stabilising influence on the post-divergence behaviour of the system and is expected to encourage the panel into a statically buckled equilibrium position.

7.4 Illustrative Results – Travelling-Wave Flutter

7.4.1 Prescribed Wall Motion

The mechanism behind TWF is the irreversible energy transfer that occurs due to a favourable shift in the phase relationship of the flow-pressure and wall-velocity. It was shown by Benjamin (1963) and more recently Carpenter and Garrad (1986) that for an infinite domain, based on an inviscid shear-layer theory (analogous to very thin boundary layers), positive energy transfer between a boundary-layer flow and a travelling wall-wave can occur for wave speeds (c) in the range of $0 < c/U_\infty < 1$.

We attempt to validate the fluid component of these mechanisms in the present model by prescribing the wall motion as a travelling wave of the form,

$$\eta(x,t) = \Re\left\{Ae^{i\alpha(x-ct)}\right\}, \tag{7.12}$$

where α is the angular wave-number and c is the complex wave-speed. Interaction of the fluid-structure system is therefore restricted to one-way - from the wall to the fluid. We apply the travelling-wave motion only to the finite section of compliant-wall and set the boundary conditions to $\eta(0,t) = \eta(L,t) = 0$. The system parameters are set to match those used previously with the exception of $\delta/L = 1$, while $c/U_\infty = 2$ and $\alpha = 8\pi/L$. Figure 7.17 shows the series of wall deformations when prescribing the travelling wave form in time. The simulation time has also been normalised to the time period (T) of a single cycle of the travelling-wave. Figure 7.18 shows the relationship of the wall-velocity ($\dot{\eta}$) and coefficient of pressure (C_p) at the wall-midpoint ($x = L/2$) in time. It is

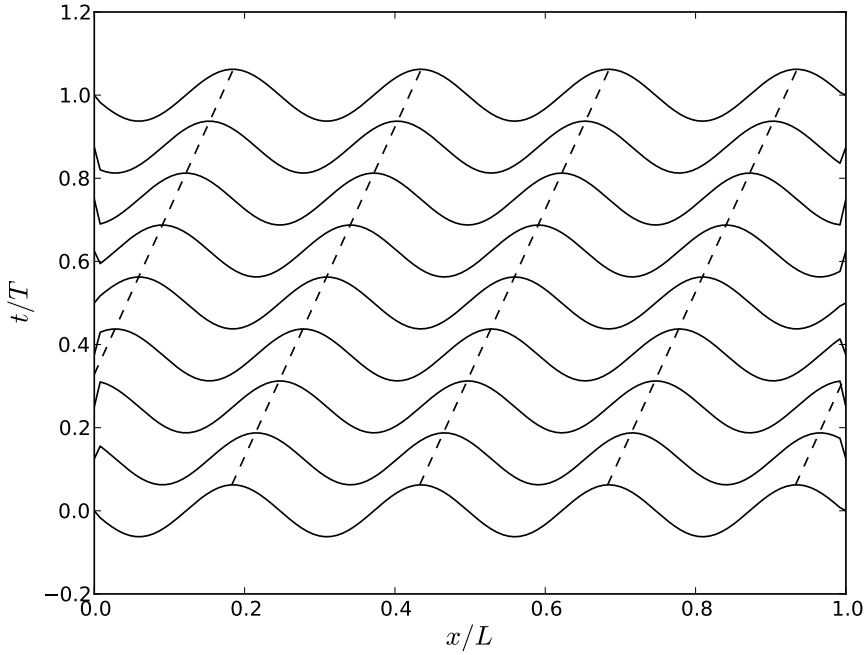

Figure 7.17: The series of wall deformations at various times {solid} for the prescribed downstream travelling-wave form (with artificially set end conditions).

noted that results are only taken after sufficient time to allow initial disturbances in the fluid (arrising from starting the fluid from a rested position) to convect downstream. By including the boundary-layer we observe a shift in the phase relationship between the pressure and velocity signal that is now lagging when compared to the potential-flow case of 90-degrees. We define the dimensionless rate of work (power) done by the fluid-flow as,

$$\bar{P} = \frac{C_p \dot{\eta}}{c} \tag{7.13}$$

Figure 7.19 shows the instantaneous and averaged power output of the fluid at the wall-midpoint in time. With an average power output that is below zero (at $\bar{P} = -408$), the fluid is absorbing energy. In a FSI system this would cause attenuation of the compliant-wall surface waves that are travelling downstream at this wave-speed (or greater), offering qualitative agreement with the conclusions of Carpenter and Garrad (1986). By using the same system parameters we also investigate the effects of the boundary-layer flow on wave-packets travelling

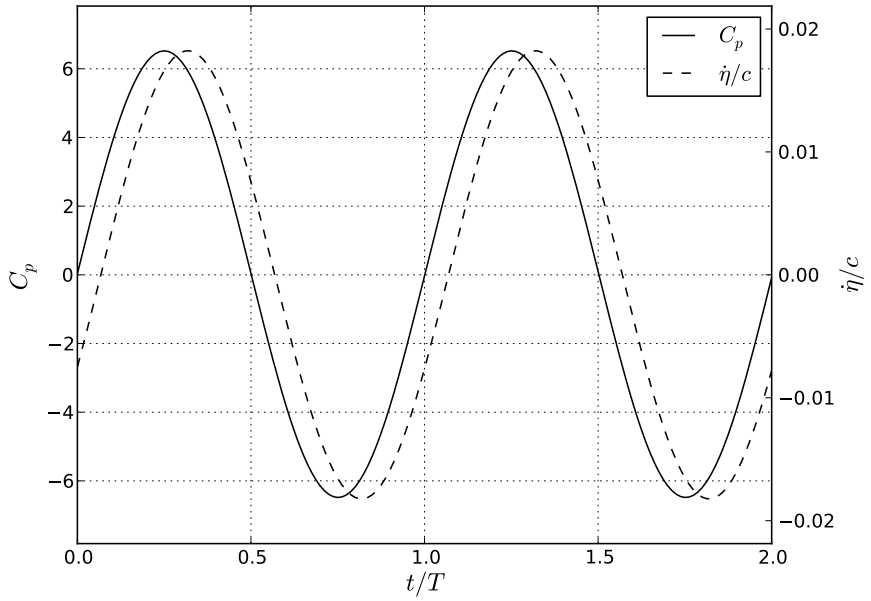

Figure 7.18: Variation of wall-midpoint velocity {- -} and flow-pressure {solid} in time for the boundary-layer flow over a prescribed travelling wave with speed $c/U_\infty = 2$.

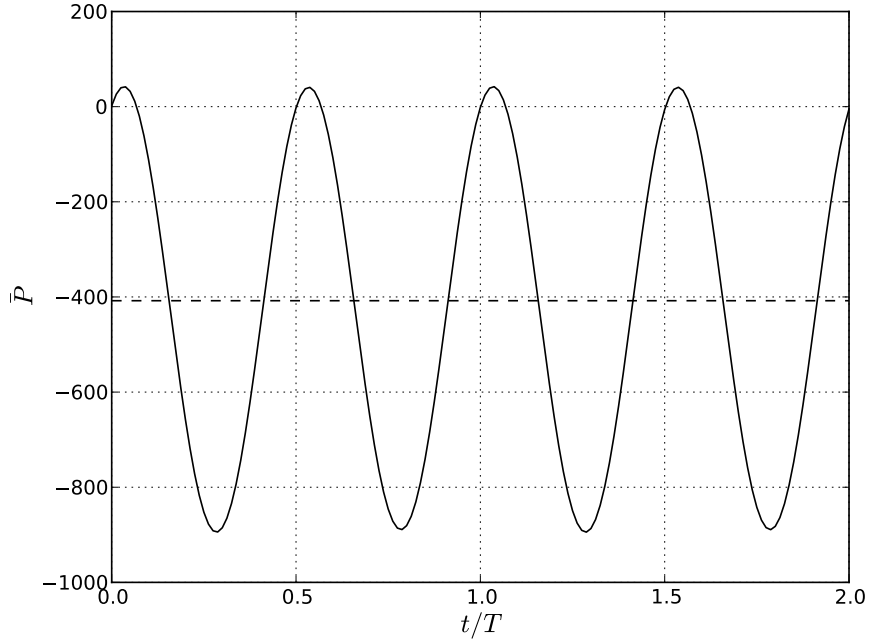

Figure 7.19: Total {solid} and average {- -} power output of the boundary-layer flow over a prescribed travelling wave with speed $c/U_\infty = 2$.

downstream at $c/U_\infty = 0.25$. Figure 7.20 shows the coefficient of pressure and the wall-velocity at the midpoint in time. It is seen that the relative phase shift of the pressure signal to the potential-flow case is now of leading type. The resulting

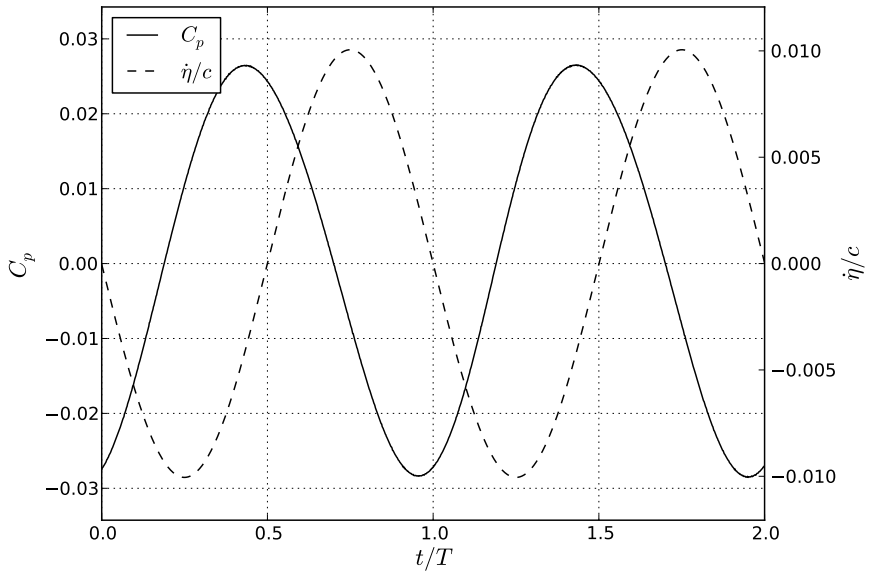

Figure 7.20: Variation of wall-midpoint velocity {- -} and flow-pressure {solid} in time for the boundary-layer flow over a prescribed travelling wave with speed $c/U_\infty = 0.25$.

instantaneous and average power production at the midpoint in time is shown in Figure 7.21. The average power output of the fluid is positive (at $\bar{P} = 0.351$) and therefore indicates that the boundary-layer has a means of injecting energy into waves travelling downstream at this speed. This behaviour further agrees with the results of Carpenter and Garrad (1986) and demonstrates the existence of the key mechanism for TWF in a compliant-wall FSI system.

7.4.2 Finite-Wall Travelling-Wave Flutter

With the underlying mechanisms that drive TWF demonstrated in our model we move to capturing the existence of a TWF instability for the full boundary-layer compliant-wall FSI system. We model the flow of water over a rubber-type compliant-wall with a spring-foundation. Whilst most system and disretisation parameters match those used in the nonlinear divergence simulations, a

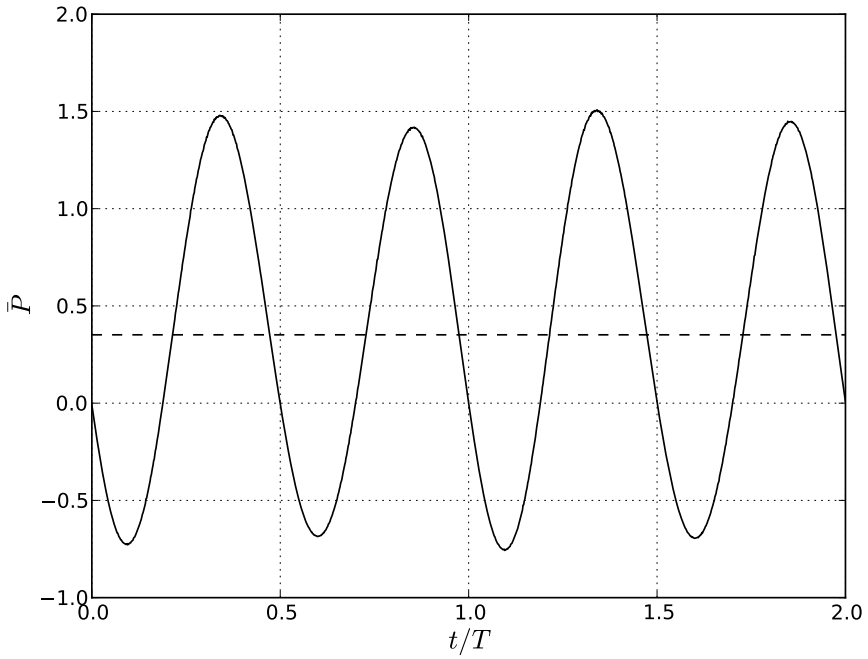

Figure 7.21: Total {solid} and average {- -} power output of the boundary-layer flow over a prescribed travelling wave with speed $c/U_\infty = 0.25$.

compliant-wall spring foundation has now been included to ensure that divergence does not dominate the system dynamics. By increasing the spring foundation stiffness the divergence on-set speed can be raised higher than the TWF on-set speed because of their respective exponents in following equations Equation (7.14) and Equation (7.15). We set $L = 1m$, $h/L = 0.04$, $\delta/L = 0.27$, $\rho_f = 1000\text{kg/m}^3$, $\rho/\rho_f = 1.2$, $B = 3.55\text{Nm}$ and $k = 5.88 \times 10^7 Pa/m$. Rigid sections of wall are also included up- and down-stream of the compliant wall at a length of $L/8$ within the computational domain. Viscosity is adjusted accordingly to set the control parameter of the flow (Reynolds-number) to $Re_{\delta^*} = 2970$. The number of finite-difference mass nodes is also increased to $N = 256$ to allow the resolution of waves of shorter length. According to the analytical work of Carpenter and Garrad (1986), the divergence-onset flow speed for a spring-backed compliant-wall (potential flow) system is,

$$U_{\text{div}} = 2\sqrt[8]{\frac{BK^3}{27{\rho_f}^4}}. \tag{7.14}$$

This predicts a divergence-onset speed for the present system with a potential flow to be $U_{\text{div}} = 40.2\text{m/s}$. However, according to the results of Section 7.3.1, the relatively thick boundary-layer will cause an increase in the divergence onset speed of approximately 25% to $U_{\text{div}} \approx 50\text{m/s}$. The onset flow speed for the TWF of a spring-backed compliant-wall system is given by (Carpenter and Garrad, 1986),

$$U_{\text{twf}} = \sqrt{\frac{2\sqrt{BK}}{h\rho}}. \tag{7.15}$$

This predicts a TWF onset speed for the present system of $U_{\text{twf}} = 24.5\text{m/s}$. We set the mean-flow speed for the present simulation at $U_\infty = 36.2\text{m/s}$ to encourage the excitement of TWF whilst maintaining a safe margin against the excitation of divergence. For flow speeds above the critical, Carpenter and Garrad (1986) predict the outer bounds of unstable wave-numbers using,

$$\alpha = \sqrt{\frac{h\rho {U_\infty}^2 \pm \sqrt{(h\rho {U_\infty}^2)^2 - 4BK}}{2B}}. \tag{7.16}$$

For the present system this is $31.5 < \alpha < 129$. We initialise the compliant-wall as a standing wave of length $\lambda = 0.2L$ ($\alpha = 31.4$) starting at $x = 0$ in the form,

$$\eta(x,0) = \frac{A}{2}\begin{cases} \cos\left(2\pi\left(\frac{x}{0.75\lambda}\right)\right) - 1 & : 0 \le x/\lambda < 0.25, \\ \cos\left(2\pi\left(\frac{x}{0.75\lambda}\right)\right) - \cos\left(2\pi\left(\frac{x-0.25\lambda}{0.75\lambda}\right)\right) & : 0.25 \le x/\lambda < 0.75, \\ -\cos\left(2\pi\left(\frac{x-0.25\lambda}{0.75\lambda}\right)\right) + 1 & : 0.75 \le x/\lambda < 1, \\ 0 & : \text{otherwise}, \end{cases} \tag{7.17}$$

to act as a smooth initial disturbance for the fluid flow (Cafolla, 1997). This is plotted in Figure 7.22a.

A key strength of the present modelling scheme is that irrespective of the initial disturbance (or even in the absence of one), the most unstable wave-forms of the system will eventually emerge.

Figures 7.22a to 7.22d and 7.23a to 7.23d show the wall-displacement for times $t/T_0 = 0 \rightarrow 3.5$ at intervals of $t/T_0 = 0.5$, where T_0 is the time taken for a full oscillation of the initial standing wave. The initial disturbance slowly convects downstream and reduces in amplitude by giving rise to dominating shorter and fast travelling waves observed from $t/T_0 = 1$ onwards. The natural emergence and spatial growth of these shorter waves gives evidence to the onset of TWF, as predicted by Carpenter and Garrad (1986) for the present system parameters. Further confirmation of TWF is achieved by measuring the dominant wave-number, which is found to be at $\alpha \approx 100$, placing it inside the unstable range of waves predicted by Equation (7.16). While TWF is often described as a purely convective instability, this is only true for infinite domains. With a finite wall and in the absence of structural damping, the energy gained by the travelling waves is accumulated/reflected once it reaches the hinged wall-ends. It is seen that the boundary-layer's ability to attenuate upstream travelling waves eventually causes all accumulated energy to be transformed effectively into a temporally unstable downstream standing wave adjacent to the fixed end of the wall (at $x/L \gtrsim 0.7$ in Figure 7.23d). The finite nature of the wall also permits the hinged end in the upstream section to act as a driver for introducing new travelling waves. Transient accumulation of energy in the downstream region is seen to encourage the upstream driving action and offers a means of self-excitation that causes large amplitude oscillations for long simulation times. Figure 7.24 shows the total, spring foundation, bending and kinetic energy of the compliant wall in nondimensional time (t/T_0). The energy in the spring-foundation term dominates until bending takes over at $t/T_0 > 2$. It is noted that the tension term has not been shown due to its negligible magnitude for this system. For time $t/T_0 < 1$ it is seen that the total energy of the system is decaying. The loss of energy of the initial waveform indicates that there is no energy transfer from the fluid to the wall to overcome the attenuating effects of the boundary-layer flow. However, for time $t/T_0 > 1$ it is seen that the average energy level is increasing, representing the emergence of the short/fast waves that characterise the onset of

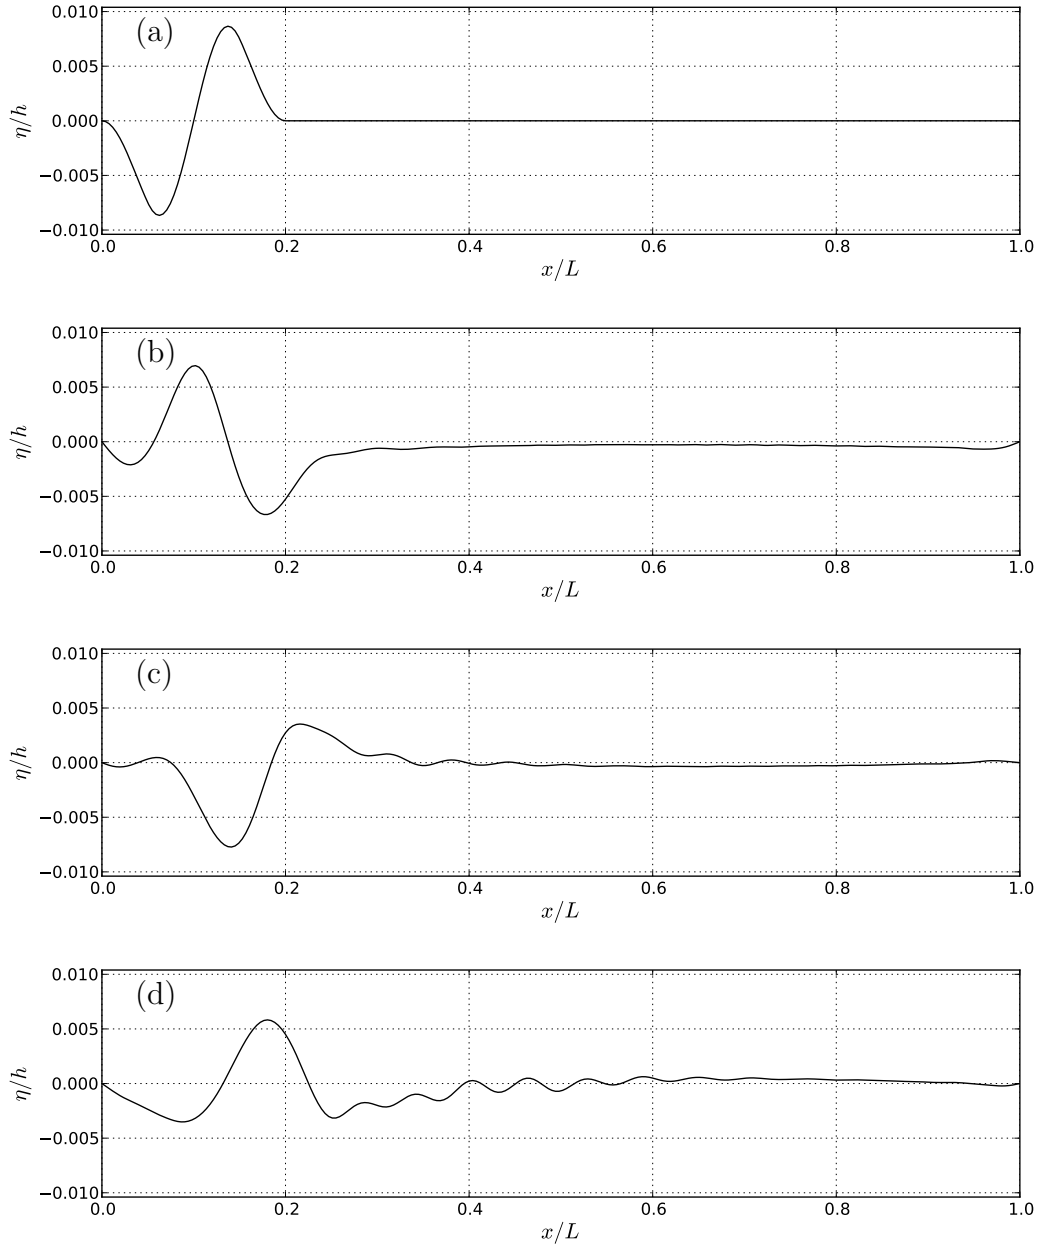

Figure 7.22: The instantaneous wall-displacement at times $t/T_0 = \{0, 0.5, 1, 1.5\}$ respectively for the TWF of a water boundary-layer flow over a rubber-type compliant wall.

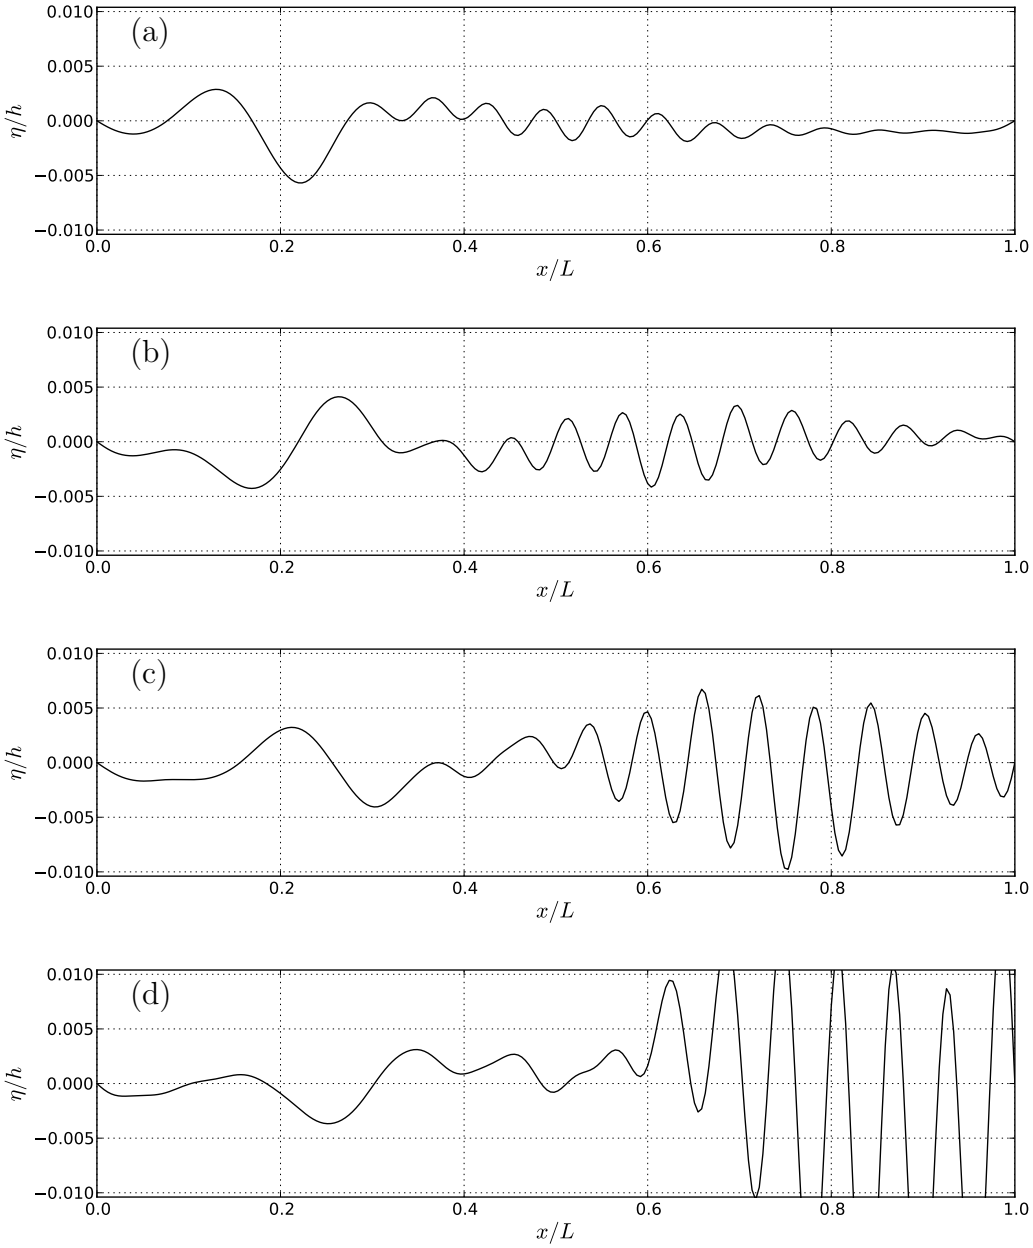

Figure 7.23: The instantaneous wall-displacement at times $t/T_0 = \{2, 2.5, 3, 3.5\}$ respectively for the TWF of a water boundary-layer flow over a rubber-type compliant wall.

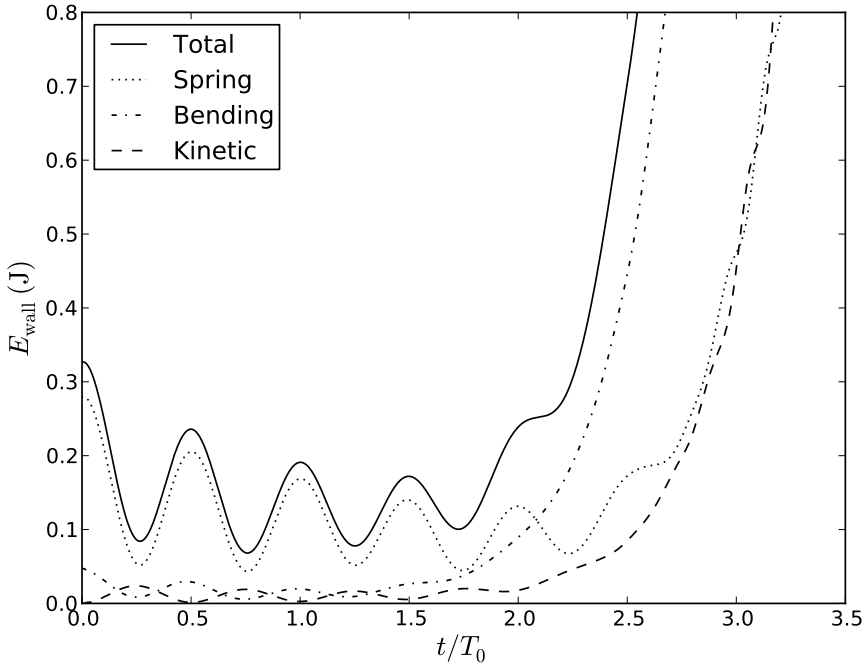

Figure 7.24: Variation of total, spring foundation, bending and kinetic energy in time for the TWF of a water boundary-layer flow over a rubber-type compliant wall.

a TWF instability. At $t/T_0 = 2$ the first set of short/fast waves are seen to reach the hinged-end of the compliant wall. The reflected/accumulated waves cause a temporal growth of energy that leads to a large-amplitude self-excited absolute type of instability for $t/T_0 > 2$ because no structural damping has been included.

Effects of Structural Damping

The absolute type of instability that is found in the finite-wall system is one that grows only due to accumulated TWF-generated energy. As we are operating below the divergence onset speed, there is no mechanism for static wave growth[4]. Consequently we investigate the use of structural damping to control the onset of the convective TWF and the resulting absolute type of instability in the finite-wall system. Figures 7.25a to 7.25d, 7.26a to 7.26d, 7.27a to 7.27d and 7.28a to 7.28d show the wall-displacement for times $t/T_0 = 0 \rightarrow 3.5$ at intervals of $t/T_0 = 0.5$

[4]The outcome of Section 7.3.2 suggests that at pre-divergence flow speeds, the boundary-layer will actually cause dissipation of static waves.

for the present system with light-damping ($D = 10$kNs/m) and heavy-damping ($D = 100$kNs/m) respectively. It is seen that TWF has not been completely eliminated, however, it is controlled by increasing the level of damping to reduce the travelling waves' peak amplitude and rate of growth. As such the use of heavy damping greatly delays the point in time at which the effects of TWF can be clearly seen. Figures 7.29 and 7.30 show the total energy in nondimensional time (t/T_0) for the compliant wall with light and heavy damping respectively. The use of light damping is seen to provide only a marginal decrease in the total energy and rate of growth when compared to the undamped case. However, the use of heavy damping is seen to rapidly dissipate the initial disturbance and have a greater effect on reducing overall energy levels. As the excitation of TWF has not been avoided we still see temporal energy growth in the finite system from $t/t_0 > 2$, however, the rate of energy growth has been greatly reduced, indicating that further increases in damping would overcome this instability entirely (see Figure 7.30).

Upstream Wall Effects

We investigate the present FSI system with no structural damping and the initial disturbance centralised about $x = L/2$ (see Figure 7.31a). Figures 7.31a to 7.31e and 7.32a to 7.32e show the wall-displacement for times $t/T_0 = 0 \rightarrow 4.5$ at intervals of $t/T_0 = 0.5$.

As previously, the initial disturbance slowly travels downstream whilst transforming into shorter/faster downward travelling waves that demonstrate the onset of TWF. It is seen that the boundary-layer attenuates all forms of upstream-travelling waves, however, TWF is seen to emerge at locations upstream of the site of the initial disturbance. These self-excited waves demonstrate the driving influence of the upstream wall-hinge in conjunction with the destabilising effect of the downstream standing waves adjacent to the fixed end of the wall. Figure 7.33 shows the total energy of the compliant wall in nondimensional time (t/T_0). For time $t/T_0 < 2$ it is seen that the system energy decays. The first

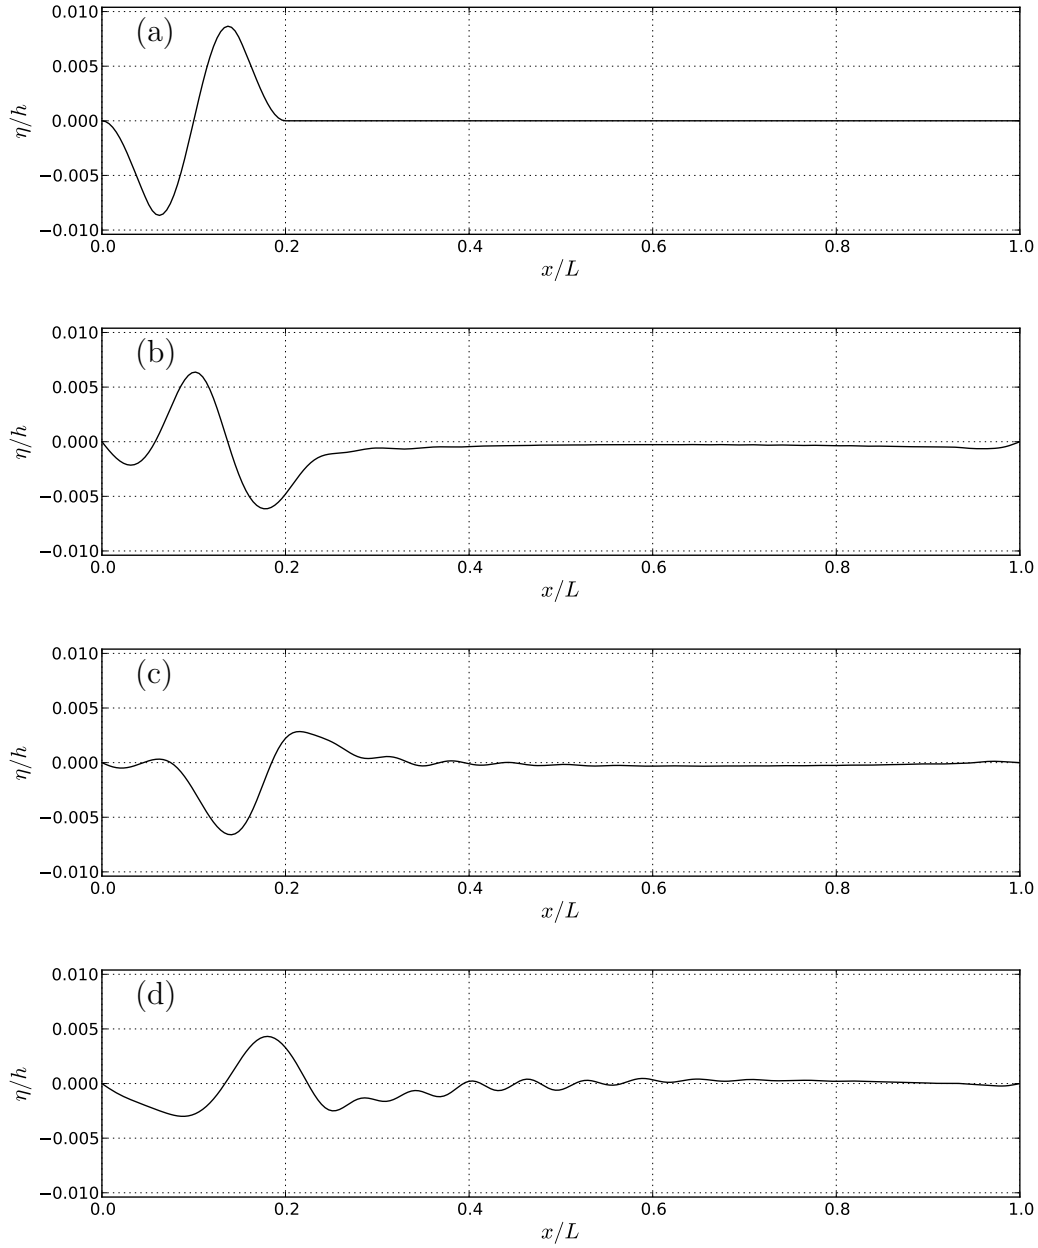

Figure 7.25: The instantaneous wall-displacement at times $t/T_0 = \{0, 0.5, 1, 1.5\}$ respectively for the TWF of a water boundary-layer flow over a rubber-type compliant wall with light ($D = 10$ kNs/m) structural damping.

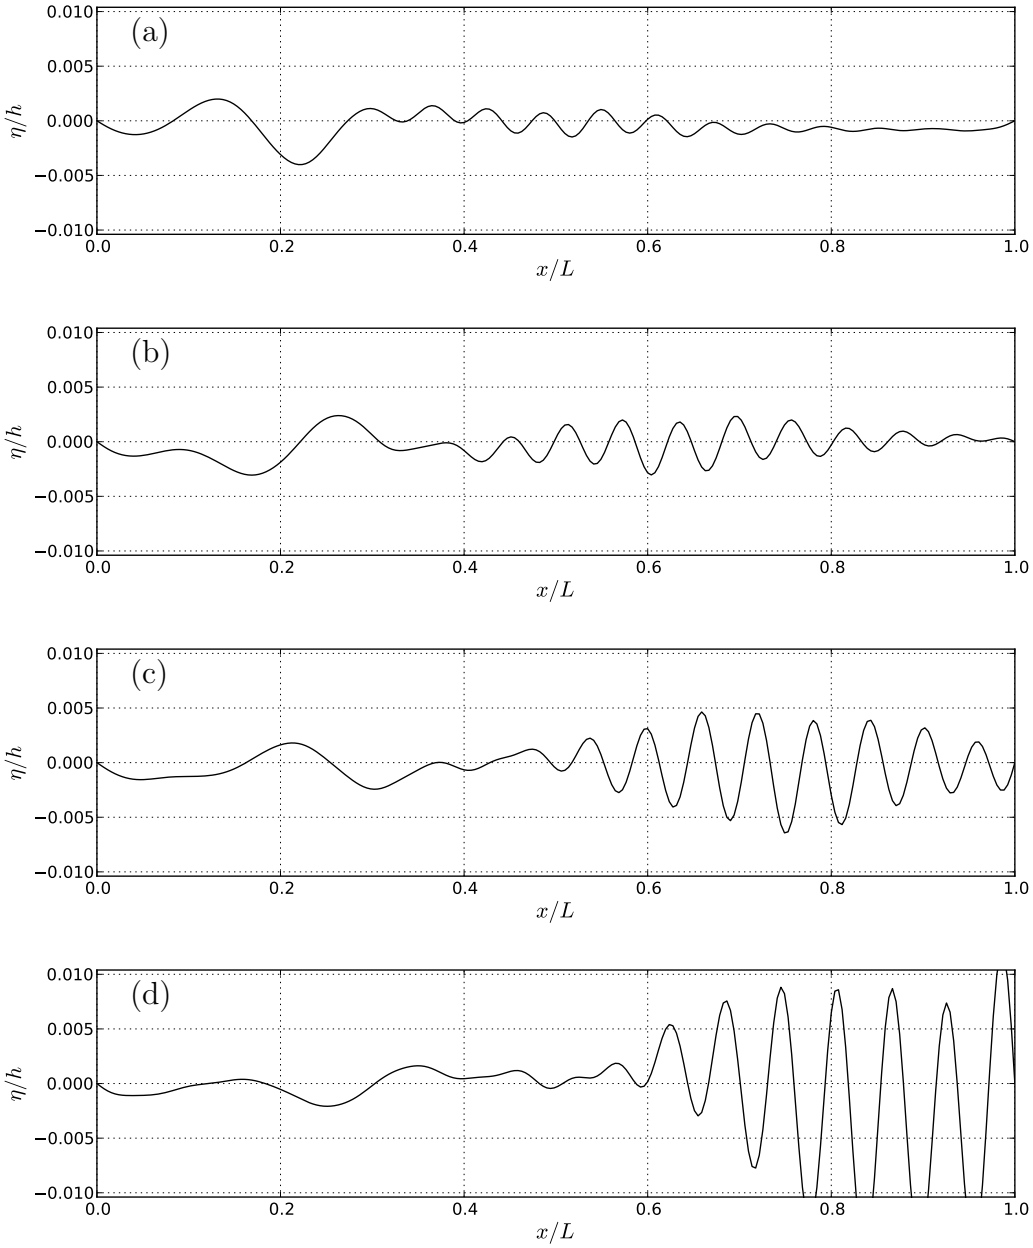

Figure 7.26: The instantaneous wall-displacement at times $t/T_0 = \{2, 2.5, 3, 3.5\}$ respectively for the TWF of a water boundary-layer flow over a rubber-type compliant wall with light ($D = 10$ kNs/m) structural damping.

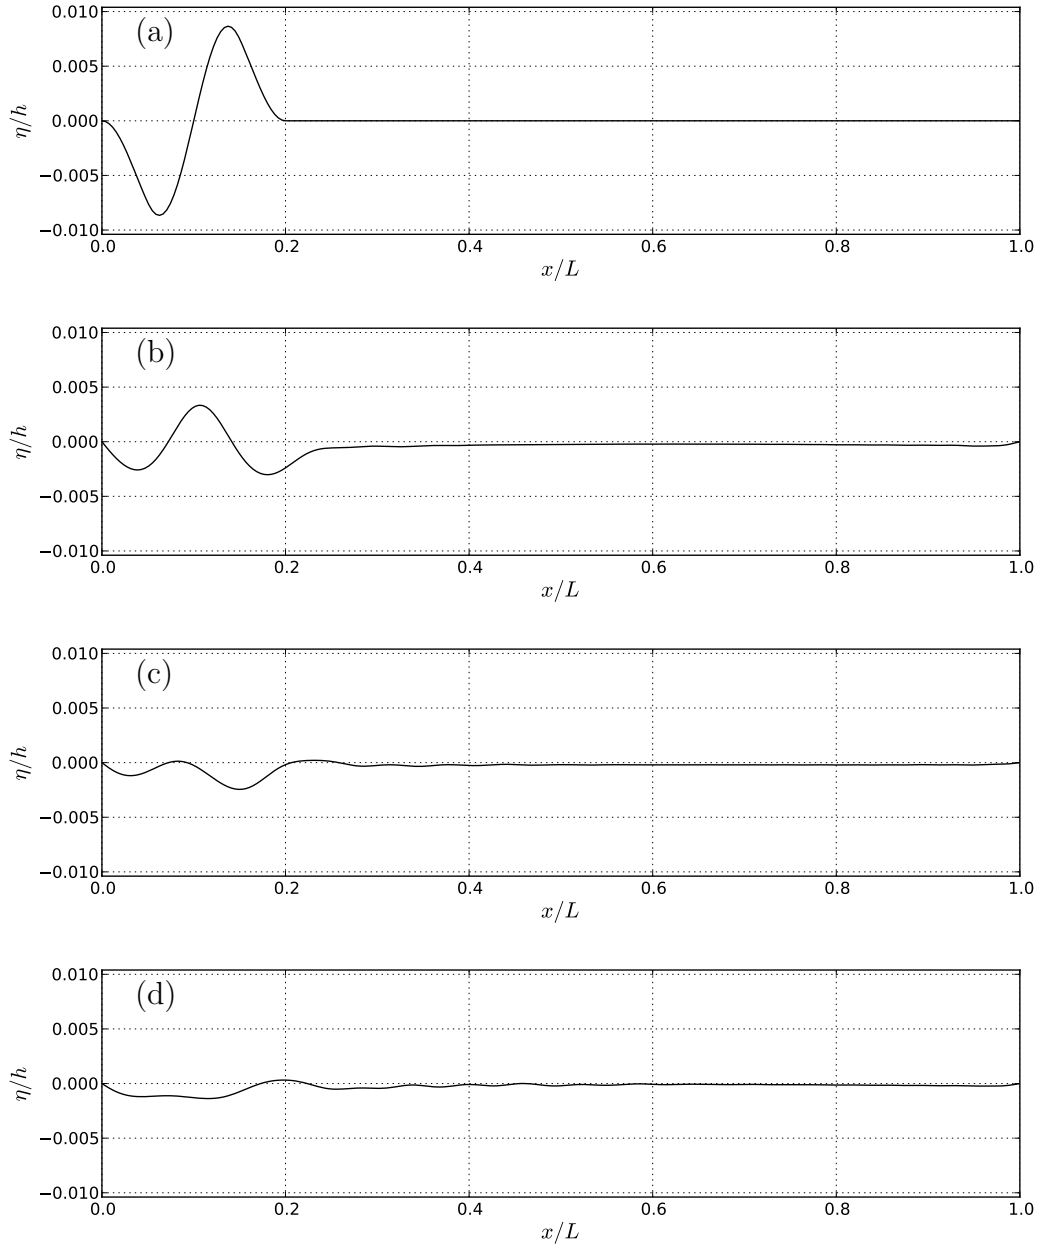

Figure 7.27: The instantaneous wall-displacement at times $t/T_0 = \{0, 0.5, 1, 1.5\}$ respectively for the TWF of a water boundary-layer flow over a rubber-type compliant wall with heavy ($D = 100$ kNs/m) structural damping.

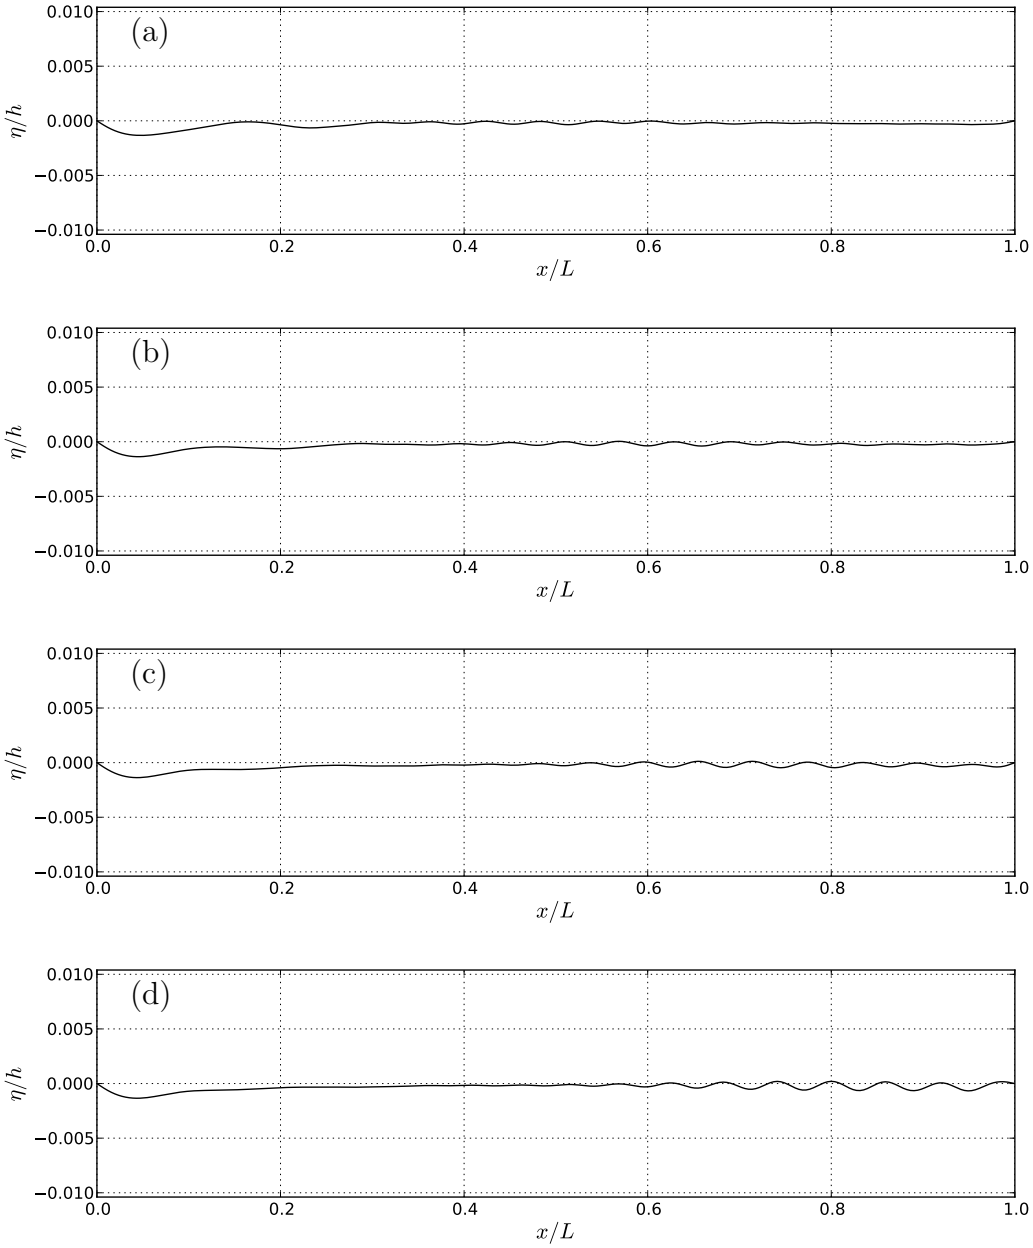

Figure 7.28: The instantaneous wall-displacement at times $t/T_0 = \{2, 2.5, 3, 3.5\}$ respectively for the TWF of a water boundary-layer flow over a rubber-type compliant wall with heavy ($D = 100$ kNs/m) structural damping.

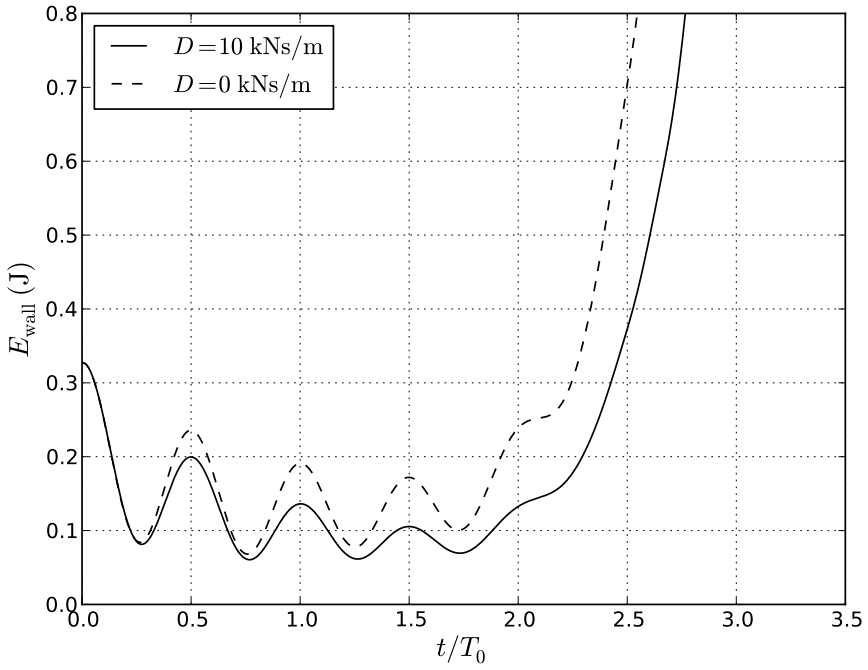

Figure 7.29: Variation of total wall energy in time for the TWF of a water boundary-layer flow over a rubber-type compliant wall with light ($D = 10$ kNs/m) structural damping.

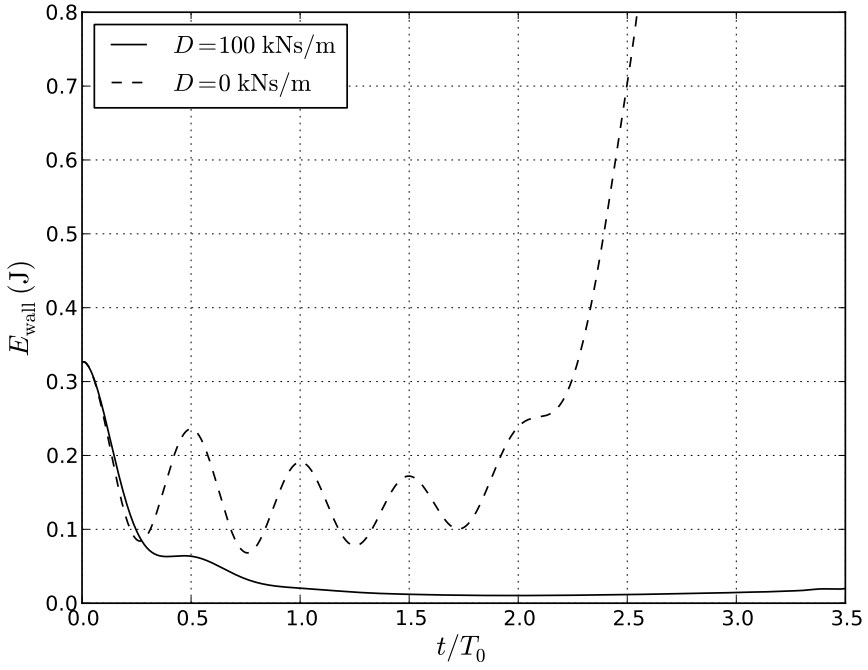

Figure 7.30: Variation of total wall energy in time for the TWF of a water boundary-layer flow over a rubber-type compliant wall with heavy ($D = 100$ kNs/m) structural damping.

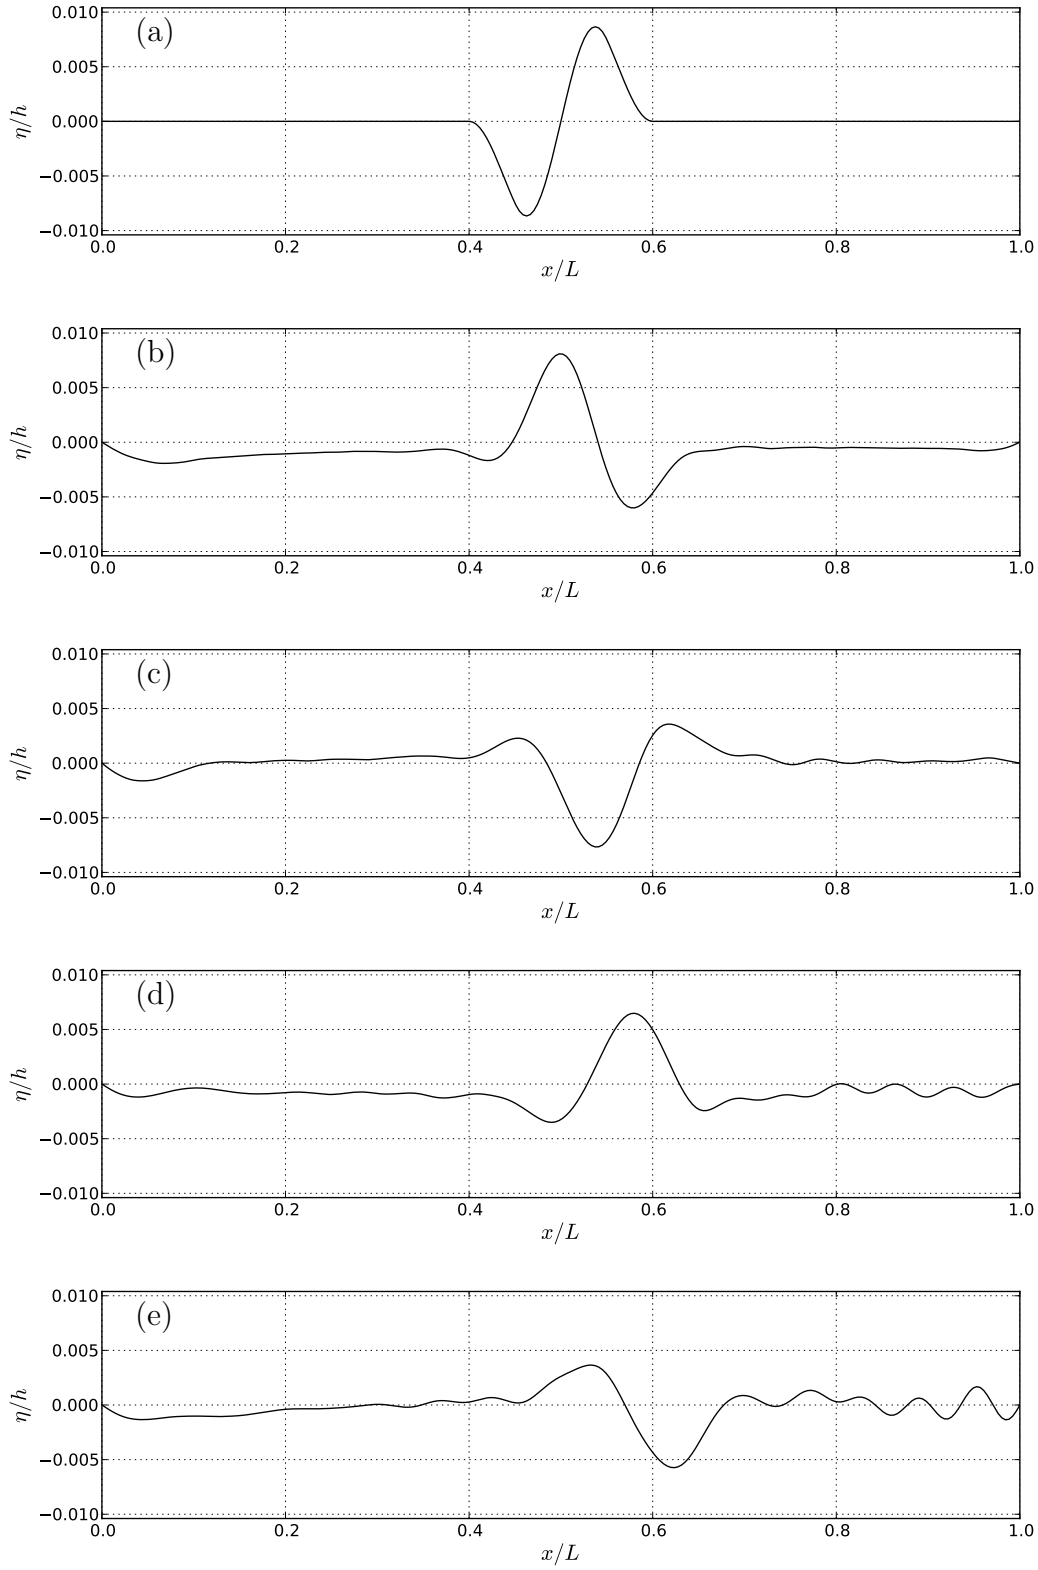

Figure 7.31: The instantaneous wall-displacement at times $t/T_0 = \{0, 0.5, 1, 1.5, 2\}$ respectively for the TWF of a water boundary-layer flow over an undamped rubber-type compliant wall with a centralised initial disturbance.

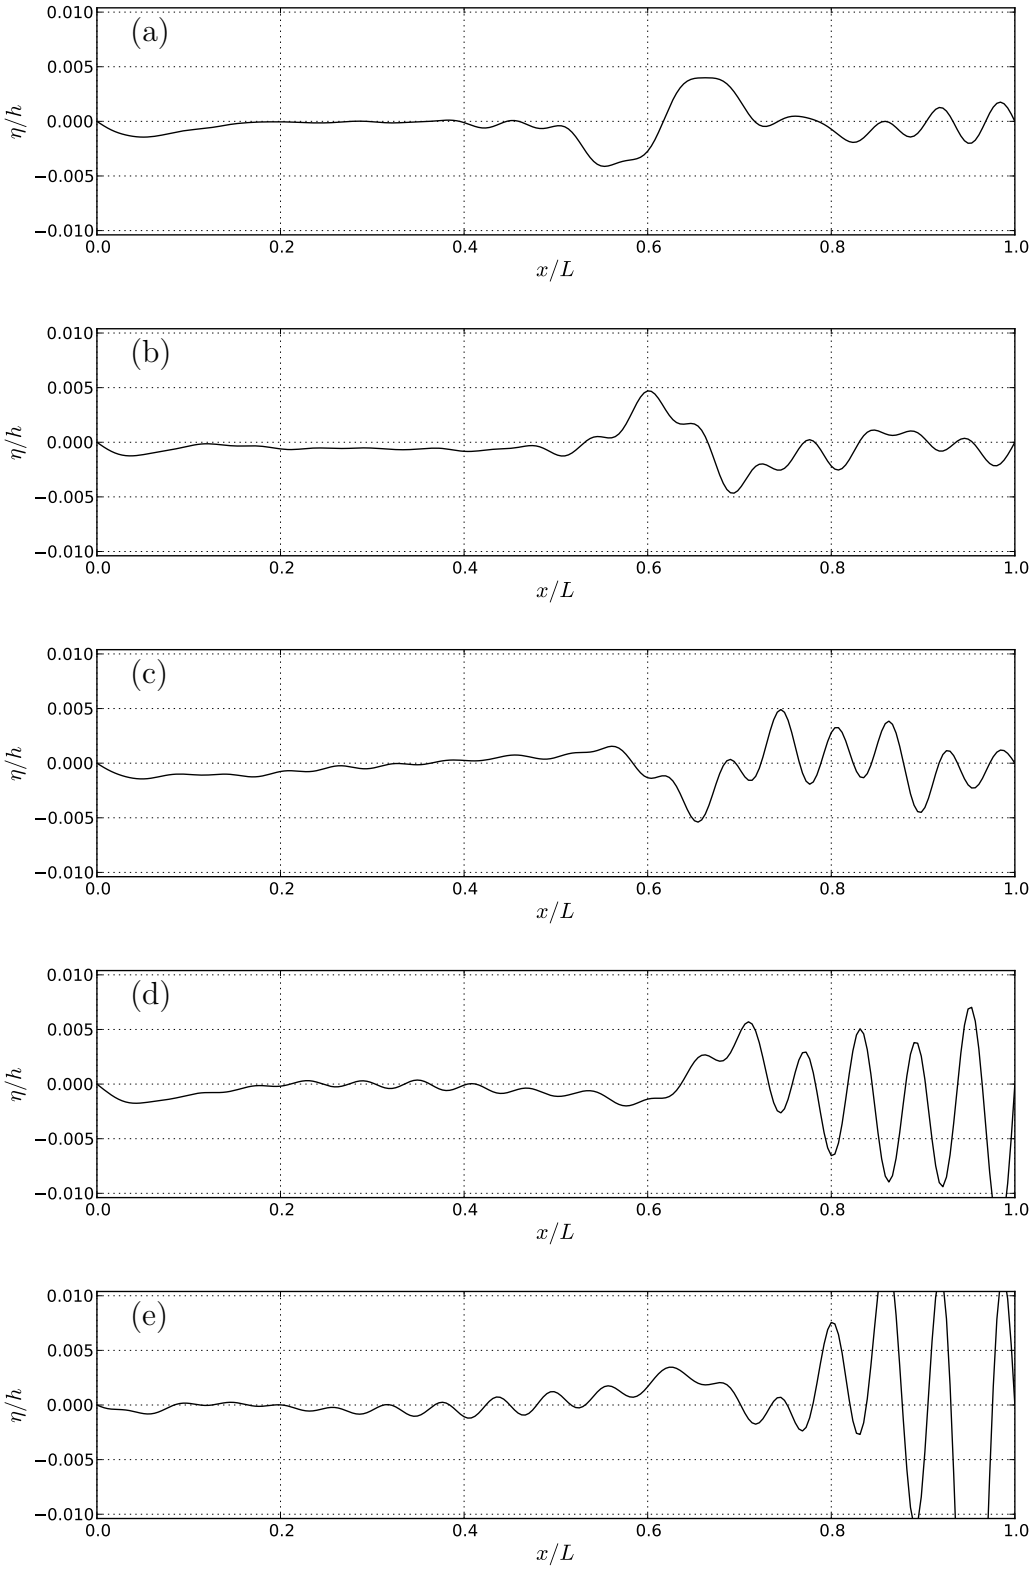

Figure 7.32: The instantaneous wall-displacement at times $t/T_0 = \{2.5, 3, 3.5, 4, 4.5\}$ respectively for the TWF of a water boundary-layer flow over an undamped rubber-type compliant wall with a centralised initial disturbance.

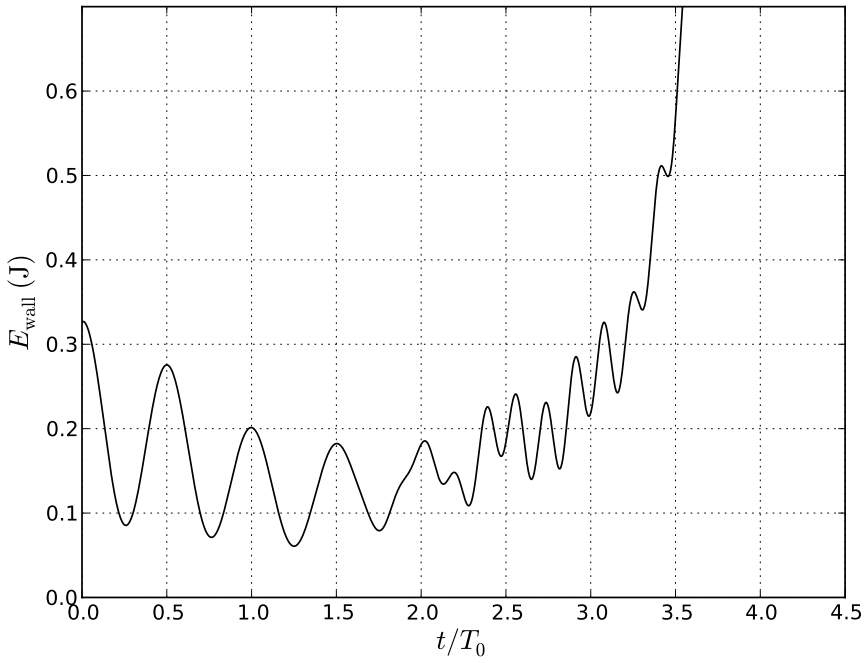

Figure 7.33: Variation of total wall energy in time for the TWF of a water boundary-layer flow over an undamped rubber-type compliant wall with a centralised initial disturbance.

set of travelling waves emerge from the initial prescribed (now centralised in the domain) waveform and reach the downstream wall at an earlier time of $t/T_0 = 1$. The shorter distance from the initial waveform to the end of the compliant wall reduces the early (simulation time) convective energy gains and delays the onset of the self-excited absolute type of instability. For time $t/T_0 > 2$ the downstream standing wave has accumulated sufficient energy to dominate the initially prescribed waveform and at the relatively (to the previous simulations) later time of $t/T_0 > 3$ the wall undergoes rapid temporal energy growth. To investigate further, Figures 7.34 and 7.35 show the variation of total energy in time for the upstream ($x < L/2$) and downstream ($x > L/2$) sections of the wall respectively. It is seen that in the upstream section, the wall immediately begins to lose energy as the initial prescribed disturbance convects downstream and upstream travelling waves are attenuated. The upstream energy level stabilises once TWF is first established ($t/T_0 > 1$) and then grows indefinitely after finite-wall effects cause self-excitation for $t/T_0 > 3$. For the downstream section, wall energy grows im-

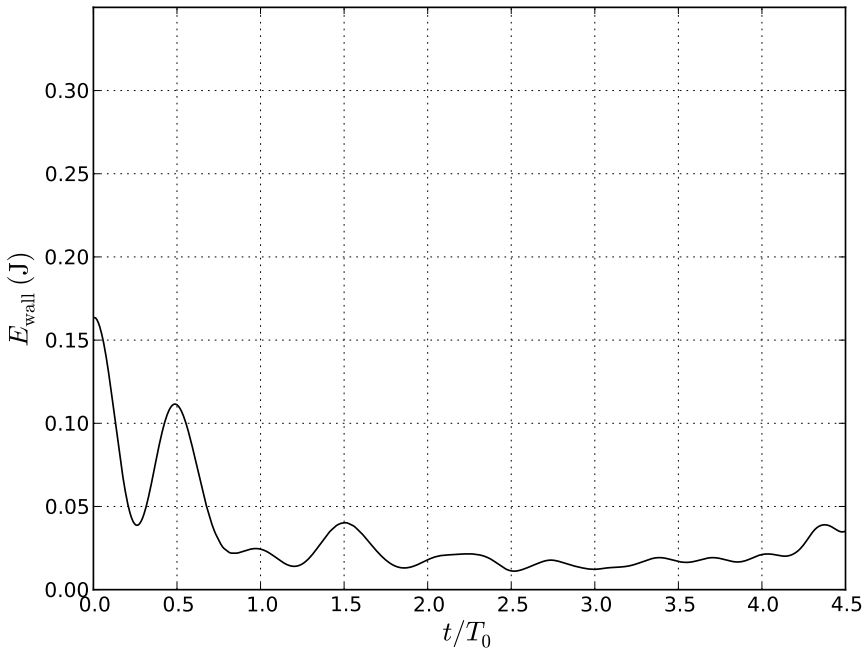

Figure 7.34: Variation of upstream ($x < L/2$) wall energy in time for the TWF of a water boundary-layer flow over an undamped rubber-type compliant wall with a centralised initial disturbance.

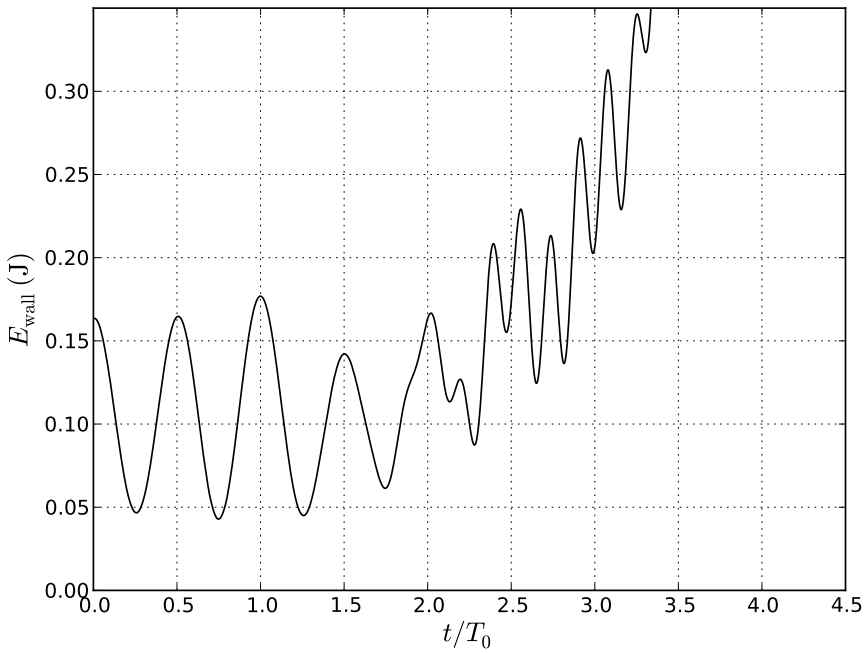

Figure 7.35: Variation of downstream ($x > L/2$) wall energy in time for the TWF of a water boundary-layer flow over an undamped rubber-type compliant wall with a centralised initial disturbance.

mediately with the initially applied deformation convecting into the downstream half. The onset of TWF causes a growth in energy while the change in oscillation frequency is attributed to the temporally unstable downstream standing wave adjacent to the fixed wall-end (at $x/L > 0.8$ in Figure 7.32e).

Travelling-Wave Flutter Summary

The inclusion of a finite-thickness boundary-layer to the previous FSI system demonstrates a mechanism for wall energy growth that revealed a mechanism for establishing potentially large-amplitude self-excited oscillations. This behaviour existed at a flow speed much lower than that for divergence because the travelling-wave flutter (TWF) and divergence onset speeds can be independently varied. However, unlike divergence, TWF can be controlled through the use of sufficient structural damping.

7.5 Summary

By combining the methods and models of a boundary-layer flow and a nonlinear potential-flow fluid-structure interaction (FSI) system, we demonstrated a single tool for the study of two-dimensional FSI of a finite-length Kramer-type compliant wall coupled with a viscous and finite-thickness boundary-layer flow at nonlinear ranges of motion.

The viscous-flow pressure was calculated by a y-momentum integration method and was coupled numerically using a semi-implicit scheme. An implicitly-coupled nonlinear potential flow solution was simultaneously used to improve efficiency and ensure numerical stability. The results indicate that the inclusion of boundary-layer effects act as a modifier to an otherwise potential-flow by causing the reduction of the steady-state flow pressure. This effect was shown to cause an increase in the divergence-onset speed that was dependent on the boundary-layer thickness. The highest critical speeds were observed with thicker boundary-layers, while reducing the boundary-layer thickness towards zero caused the critical flow

speed to approach that of the potential-flow case (as might be expected).

The inclusion of boundary-layer effects on the nonlinear limit-cycle oscillations of a post-divergent FSI system revealed that the underlying behaviour of system remained largely unchanged from the potential-flow case. While the most significant changes were through the lower steady-state pressure reducing oscillation amplitudes, the dissipative nature of the flow modified the FSI behaviour in a manner similar to the inclusion of weak structural damping. Whilst this damping would not prevent the onset of divergence instabilities, it would be expected to offer a stabilising influence on the post-divergence behaviour of the system by encouraging the wall into a statically buckled equilibrium position.

For downstream travelling waves a relative shift in the phase relationship of flow-pressure and wall-velocity was shown to exist and support irreversible energy transfer for a specific wave speed. The associated compliant-wall FSI system gave rise to TWF that was characterised by convectively growing high-frequency waves. In addition, the effects of the finite-wall ends were shown to provide a means of self-excitation that lead to an absolute type of instability; i.e. amplitude growth occurs both upstream and downstream of the position of initial excitation. This type of destabilisation can occur at flow speeds much lower than the divergence onset. However, in accordance with theoretical expectations, the growth of wall-energy induced by TWF was seen to be reduced by the use of structural damping.

Chapter 8

Conclusion

While compliant-wall technologies have shown great promise for reducing the parasitic energy loss of drag-forces for marine transportation, tools that can cheaply optimise and evaluate such technologies are not in abundance. With inherently high costs in conducting (often 'trial and error') physical experiments and the lack of general purpose analytical solutions to the full nonlinear system, the numerical modelling approach shows promise[1] as a complementary (to linear stability theory) way to aid in the design and study of future compliant-wall technologies.

We have developed a novel numerical tool that can be used to investigate how a finely-tuned compliant-wall system can beneficially modify the mechanisms that cause laminar-to-turbulent transition and offer significant drag reductions. In doing so we developed an efficient model for the problem of strongly coupled fluid-structure interaction (FSI) of a 2D boundary-layer flow over a finite-length compliant-wall at nonlinear amplitudes, capturing the three main instabilities of Tollmien-Schlichting Waves (TSW), Travelling-Wave Flutter (TWF) and divergence. While we restricted the scope of this work to 2D flow over the canonical Kramer-type compliant-wall, all methods chosen can be extended to higher-order, 3D and/or different wall models. A significant component of development was for a novel Discrete-Vortex Method (DVM) boundary-layer flow model that represented a generic Computational Fluid Dynamics (CFD) tool for solving the

[1] Particularly with available computing power increasing according to Moore's Law (Schaller, 1997).

full (incompressible) Navier-Stokes equations. This Lagrangian approach avoids the use of ad-hoc turbulence models and offers a new alternative to the (infinite and linear) Orr-Sommerfeld equation method. Currently most grid-based solvers take a segmented approach and explicitly couple independent fluid and structural models, thereby suffering associated inefficiency and numerical stability problems. We developed a new and efficient method for coupling the nonlinear boundary-layer flow directly to the nonlinear structural dynamics.

However, as true drag-reducing dynamics and boundary-layer FSI instabilities are only exhibited for a small subset of system properties, we increased the applicability of the model to engineering by providing sufficient modelling flexibility for it to be used as a general engineering tool for a broad range of problems. It also remains scalable and efficient enough to allow the computation to be executed on a standard desktop computer. The modelling advances made are as follows;

1. An inviscid (potential) flow over arbitrarily shaped (2D) bodies/walls was modelled using the Boundary-Element Method (BEM). Computational efficiency was drastically improved with the use of the Fast-Multipole Method (FMM) to solve the flow-field velocity evaluations, while a sparse Incomplete-LU pre-conditioned Generalised Minimum Residual method (GMRES) solved the BEM system strengths in an iterative matrix-free manner (using the FMM).

2. In the absence of external loading, the dynamic response of a compliant wall undergoing various modes of harmonic oscillation was modelled. The inherent stability benefits of using an implicit time-stepping scheme when coupled with the Finite-Difference Method (FDM) were proven while a pre-conditioned Newton-Krylov (NK) method showed an efficient technique for overcoming the computational difficulties of implicit solution methodologies. The combination of NK and FDM yielded an accurate and efficient matrix-free tool for the transient study of compliant wall systems undergoing nonlinear motions.

3. The unsteady FSI of an inviscid flow over a compliant surface at nonlinear amplitudes was modelled. This was achieved by implicitly coupling the pressure of the BEM to the structural mechanics to produce a numerical tool that exceeded the scalability and stability of previous works such as Lucey et al. (1997a). An efficient algorithm was also developed to iteratively find the divergence-onset flow speed of the unsteady (nonlinear) FSI system to facilitate the production of stability maps that have previously been left to linear analytical (eigen-analysis) methods.

4. The motion of essentially inviscid flows that exhibit localised rotationality effects was modelled. This included flow phenomena such as the von-Kármán street for wake-dynamics and the Kelvin-Helmholtz instability in mixing/shear layers. The BEM proved an efficient method for flow-surface modelling while the viscous free-flow was modelled with the DVM. The inherent computational difficulties associated with both N-body methods were overcome with the use of the Fast-Multipole Method (FMM) to allow particle-particle calculations to be evaluated efficiently[2].

5. A DVM-based laminar Blasius boundary-layer flow over a potentially arbitrary shaped rigid wall was modelled. This was achieved through the use of a vortex core-spreading method to model viscosity while a GMRES-based (radial-basis interpolation) rediscretisation scheme maintained Lagrangian accuracy.

6. A stable and efficient method for the study of 2D FSI of a finite-length, compliant Kramer-type wall coupled with a viscous and finite-thickness boundary-layer flow at nonlinear ranges of motion was demonstrated. This was achieved by combining the methods and models of a rigid-wall boundary-layer flow to the nonlinear potential-flow FSI system. The viscous-flow pressure was calculated by a y-momentum integration method and coupled semi-implicitly, while a nonlinear potential flow solution was implicitly

[2]A 10^6 point-vortex calculation (10^{12} operations) takes $\approx$ 10s and $\approx$ 1GB of RAM on a standard desktop computer.

coupled to improve efficiency and ensure numerical stability.

While we successfully developed a multi-purpose tool that can be used to study a broad range of phenomena encountered in compliant-wall FSI systems, we also presented new contributions to the field by demonstrating the behaviour and instabilities of strongly coupled systems that are insufficiently captured when using the existing analytical methods; these include:

- TSW: Through modelling the propagation of finite disturbances in the boundary-layer flow over a rigid wall, some agreement (e.g. areas of convective growth and decay of TSW) was demonstrated with the Orr-Sommerfeld method. This included observing elements of TSWs through the convective growth and decay of small disturbance waves.

- Divergence: Contributions to the field were made through the use of the full boundary-layer model to determine the onset of the static divergence instability. It was shown that the divergence onset speed was dependent upon the relative boundary-layer thickness. The results showed that whilst the highest divergence onset speed was for thick boundary-layers, the onset speed tended towards that of the potential-flow case (as might be expected) for thin boundary-layers.

- Nonlinear post-divergence behaviour: The unrestrictive assumptions and numerical stability of the boundary-layer FSI model contributed new knowledge that was unobtainable when using existing methods. The nonlinear limit-cycle oscillations of a post-divergent FSI system showed that the underlying behaviour remained largely unchanged from that of nonlinear potential-flow studies. The main effect of the boundary-layer flow is to both reduce the total energy of the nonlinear oscillations and to provide a dissipative effect similar to weak structural damping.

- TWF: Under the same nonlinear boundary-layer FSI model we simulated a TWF instability at pre-divergence flow speeds. This allowed the finite-wall

ends and stabilising effects of structural damping to be investigated. The FSI system was observed to gain energy through convective growth mechanisms that were ultimately self-excited by the upstream wall-end. These travelling waves then established a downstream standing wave, adjacent to the fixed trailing edge of the wall, that was temporally unstable. As the temporally unstable system was initiated by a convective energy growth mechanism, it would be difficult to use the existing infinite travelling-wave (Orr-Sommerfeld) or finite-length standing-wave analysis methods to capture this phenomenology (with or without nonlinear effects included).

8.1 Future Directions

To critique the methods presented in this thesis, the current research efforts are by no-means complete (or 'industry ready') due to various shortcomings of the present approach. We list the main problems encountered and provide some suggestions for future directions;

- With the relatively recent invention of efficient 2D and 3D fast-particle algorithms (the FMM), the computational cost of the traditional N-Body problem can be reduced. It is noted, however, that in the present work the true power of the FMM[3] was far from being demonstrated due to being held back by the computational costs of the various algorithms that supported the present FSI scheme.

- The strict need for an accurate and low-noise pressure signal for the strongly-coupled (and therefore highly sensitive) FSI system dynamics is not particularly suited to the use of the (typically noisy) Lagrangian DVM method. Consequently the DVM particle field required extremely expensive rediscretisation procedures to prevent Lagrangian distortion and ensure that the pressure signals were stable and consistently accurate. Whilst previ-

[3]For example, Yokota et al. (2011) demonstrated an 8.6×10^9 vortex simulation on a supercomputer.

ous works found success in delaying particle-field rediscretisation events, we noted significant changes in the flow behaviour and pressure signal even with the most minor of strains in the Lagrangian grid. This necessitated rediscretisation events at every time-step and negated some of the benefits of the DVM's grid-free nature, turning the grid-free method into one which relied on a pseudo-grid to maintain accuracy. It is noted, however, that in many cases the inertia of the wall can act like a low-pass filter to allow the DVM's noise/errors to appear to be 'time-averaged' away. However, this may reduce the system's (numerical) stability and also make accurate boundary-layer modelling (a requirement in modelling TSW) difficult. This impact would render the method impractical for drag-reduction studies.

- We encountered significant difficulties in obtaining sufficient (y-direction) flow resolution in the low velocity near-wall region and (x-direction) for capturing long-wave (or thin boundary-layer) phenomena due to the rediscretisation scheme requiring excessive computational effort. As a result, most of the studies in this work were based upon 'thick' boundary-layer flow phenomena. However, to overcome this challenge one 'only' needs to eliminate the prohibitively high-cost of the Lagrangian field rediscretisation.

- The advantage of solving the velocity-vorticity formulation of the Navier-Stokes equations is that the scalar pressure terms are solved implicitly in the formulation. While this reduces the degrees of freedom of the flow system, modelling FSI requires the flow pressure at the wall to be back-calculated. This is currently not a trivial problem with the DVM and often requires calculations that are of similar computational intensity (or even more in this work) to the main flow-field calculations. However it is noted that a potential solution may exist in the further development/extension of the Lagrangian 'Crocco' pressure formulation; see Lucey et al. (1997a) for 2D and Khatir and Lucey (2012) for 3D.

- The numerical stability of the semi-implicitly coupled boundary-layer FSI

solution is dependent upon the nonlinear potential-flow conditioner capturing most of the acceleration sensitive boundary-layer flow terms. It was noted that, particularly with TWF simulations, the conditioning effect was reduced and that other schemes such as implicitly coupled linearised and/or laminar boundary-layer flow solutions may offer better performance and higher stability.

- Whilst design optimisation problems can theoretically be solved with this tool, in practice the model proved fairly sensitive to specific combinations of system, discretisation and algorithm accuracy parameters. Whilst valid compliant wall FSI behaviours were demonstrated, this required a high degree of operator input in the form of simulation-specific fine-tuning and consequent testing for discretisation-independence. A solution to this may be to change the approach that allows for sub-modelling and the associated nesting of methods/algorithms by tuning the model to include only the absolutely essential system components and numerical methods of the specific problem[4]. This would therefore benefit from the input of fewer method/discretisation/accuracy parameters, but it would be at the cost of the flexibility of its current hybrid form. This would also mean that to fine-tune the model to a specific use case, one must possess *a-priori* knowledge of the system dynamics. Such knowledge may not be readily available to the wider Engineering community or even the most seasoned FSI specialists.

[4]The solution method can be significantly simplified where nonlinear effects are insignificant and/or the boundary-layer FSI dynamics are weak.

References

Y. G. Aleyev. *Nekton*. The Hague: Junk, 1977.

G. Alléon, M. Benzi, and L. Giraud. Sparse approximate inverse preconditioning for dense linear systems arising in computational electromagnetics. *Numerical Algorithms*, 16(1):1–15, 1997.

A. W. Appel. An efficient program for many-body simulation. *SIAM Journal on Scientific and Statistical Computing*, 6(1):85–103, 1985. doi: 10.1137/0906008. URL `http://link.aip.org/link/?SCE/6/85/1`.

H. Aref. Integrable, chaotic and turbulent vortex motion in two dimensional flows. *Annual Review of Fluid Mechanics*, 15:345–389, 1983.

H. Aref and E. D. Siggia. Vortex dynamics of the two-dimensional turbulent shear layer. *Journal of Fluid Mechanics*, 100(4):705–737, 1980.

W. Ashurst. Vortex simulation of a model turbulent combustor. *Combustion in reactive systems*, pages 259–273, 1981.

W. T. Ashurst. Numerical simulation of turbulent mixing layers via vortex dynamics. In *Symposium on Turbulent Shear Flows*, pages 403–413, 1977.

G. R. Baker. "Cloud in Cell" technique applied to the roll up of vortex sheets. *Journal of Computational Physics*, 31, 1979.

T. S. Balint and A. D. Lucey. Instability of a cantilevered flexible plate in viscous channel flow. *Journal of Fluids and Structures*, 20(7):893–912, 2005.

L. A. Barba and L. F. Rossi. Global field interpolation for particle methods. *Journal of Computational Physics*, 229(4):1292–1310, 2010.

L. A. Barba, A. Leonard, and C. B. Allen. Numerical investigations on the accuracy of the vortex method with and without remeshing. *AIAA paper*, 3426, 2003.

L. A. Barba, A. Leonard, and C. B. Allen. Advances in viscous vortex methods - meshless spatial adaption based on radial basis function interpolation. *International Journal for Numerical Methods in Fluids*, 47(5):387–421, 2005.

J. Barnes and P. Hut. A Hierarchical O(NlogN) Force-Calculation Algorithm. *Nature*, 324:446–449, Dec. 1986.

R. H. Bartels and G. H. Golub. The simplex method of linear programming using LU decomposition. *Communications of the ACM*, 12(5):266–268, 1969.

A. J. Basu, A. Prabhu, and R. Narasimha. Vortex sheet simulation of a plane "canonical" mixing layer. *Computers and Fluids*, 21(1):1–30, 1992.

A. J. Basu, R. Narasimha, and A. Prabhu. Modelling plane mixing layers using vortex points and sheets. *Applied Mathematical Modelling*, 19:66–75, 1995.

T. B. Benjamin. Effects of a flexible boundary on hydrodynamic stability. *Journal of Fluid Mechanics*, 9(4):513–532, 1960.

T. B. Benjamin. The threefold classification of unstable disturbances in flexible surfaces bounding inviscid flow. *Journal of Fluid Mechanics*, 16:436–450, 1963.

M. Benzi. Preconditioning techniques for large linear systems: a survey. *Journal of Computational Physics*, 182(2):418–477, 2002.

K. Berkouk, P. W. Carpenter, and A. D. Lucey. Pressure wave propagation in fluid-filled co-axial elastic tubes part 1: basic theory. *Journal of Biomechanical Engineering*, 125:852, 2003.

C. Bertram, A. Brodbelt, and M. Stoodley. The origins of syringomyelia: numerical models of fluid/structure interactions in the spinal cord. *Journal of biomechanical engineering*, 127:1099, 2005.

R. Betchov. *Simplified analysis of boundary-layer oscillations*. Douglas Aircraft Company, 1959.

G. D. Birkoff and J. Fisher. Do vortex sheets roll up? *Rendiconti del Circolo matematico di Palermo*, 2,8:77–90, 1959.

D. M. Bushnell, J. N. Hefner, and R. L. Ash. Effect of compliant wall motion on turbulent boundary layers. *Physics of Fluids*, 20:S31, 1977.

G. J. Cafolla. *Hydroelastic Instabilities of Compliant Panels*. PhD thesis, Fluid Dynamics Research Group, 1997.

P. Carpenter, A. Lucey, and C. Davies. Progress on the use of compliant walls for laminar-flow control. *Journal of Aircraft*, 38(3):504–512, 2001.

P. W. Carpenter. The optimization of compliant surfaces for transition delay. In *Turbulence Management and Relaminarisation*, volume 1, pages 305–313, 1988.

P. W. Carpenter. Status of transition delay using compliant walls. *Progress in Astronautics and Aeronautics*, 123:79–113, 1990.

P. W. Carpenter. Optimization of multiple-panel compliant walls for delay of laminar-turbulent transition. *AIAA Journal*, 31(7):1187–1188, 1993.

P. W. Carpenter. Current status on the use of wall compliance for laminar-flow control. *Experimental Thermal and Fluid Science*, 16:133–140, 1998.

P. W. Carpenter and A. D. Garrad. The hydrodynamic stability of flows over kramer-type compliant coatings. part 1. tollmien-schlichting instabilities. *Journal of Fluid Mechanics*, 155:465–510, 1985.

P. W. Carpenter and A. D. Garrad. The hydrodynamic stability of flows over kramer-type compliant surfaces. part 2. flow-induced surface instabilities. *Journal of Fluid Mechanics*, 170:199–232, 1986.

P. W. Carpenter, C. Davies, and A. D. Lucey. Hydrodynamics and compliant walls: Does the dolphin have a secret? *Current Science*, 79(6):758–765, 2000.

P. W. Carpenter, K. Berkouk, and A. D. Lucey. Pressure wave propagation in fluid-filled co-axial elastic tubes part 2: mechanisms for the pathogenesis of syringomyelia. *Journal of Biomechanical Engineering*, 125:857, 2003.

T. F. Chan and H. A. Van Der Vorst. Approximate and incomplete factorizations. *Parallel Numerical Algorithms, ICASE/LaRC Interdisciplinary Series in Science and Engeneering*, pages 167–202, 1997.

H. Cheng, V. Rokhlin, and L. Greengard. A fast adaptive multipole algorithm in three dimensions. *Journal of Computational Physics*, 155(2):468–498, 1999.

K. S. Choi, X. Yang, B. R. Clayton, E. J. Glover, M. Atlar, B. N. Semenov, and V. M. Kulik. Turbulent drag reduction using compliant surfaces. *Proceedings of the Royal Society of London. Series A: Mathematical, Physical and Engineering Sciences*, 453(1965):2229–2240, 1997.

A. J. Chorin. Numerical study of slightly viscous flow. *Journal of Fluid Mechanics*, 57(4):785–796, 1973.

A. J. Chorin. Vortex sheet approximation of boundary layers. *Journal of Computational Physics*, 27(3):428–442, 1978.

A. J. Chorin. Vortex models and boundary layer instability. *SIAM Journal on Scientific and Statistical Computing*, 1(1):1–21, 1980.

A. J. Chorin and P. S. Bernard. Discretisation of a vortex sheet, with an example of roll-up. Technical report, University of California, Berkeley, 1972.

E. Chow. Parallel implementation and practical use of sparse approximate inverse preconditioners with a priori sparsity patterns. *International Journal of High Performance Computing Applications*, 15(1):56–74, 2001.

J. P. Christiansen. Numerical simulation of Hydrodynamics by the Method of Point Vortices, J. 1 Comp. *Journal of Computational Physics*, 1973.

R. R. Clements and D. J. Maull. The representation of vortex streets by discrete vortices. *Progress in Aerospace Science*, 16,2:129–146, 1975.

F. A. Cruz and L. A. Barba. Characterization of the accuracy of the fast multipole method in particle simulations. *International Journal for Numerical Methods in Engineering*, 79(13):1577–1604, 2009. ISSN 1097-0207. doi: 10.1002/nme.2611. URL `http://dx.doi.org/10.1002/nme.2611`.

C. Davies and P. Carpenter. Numerical simulation of the evolution of tollmien-schlichting waves over finite compliant panels. *Journal of Fluid Mechanics*, 335: 361–392, 1997a.

C. Davies and P. W. Carpenter. Instabilities in a plane channel flow between compliant walls. *Journal of Fluid Mechanics*, 352:205–243, 1997b.

J. W. Demmel, S. C. Eisenstat, J. R. Gilbert, X. S. Li, and J. W. H. Liu. A supernodal approach to sparse partial pivoting. *SIAM J. Matrix Analysis and Applications*, 20(3):720–755, 1999.

A. E. Dixon, A. D. Lucey, and P. W. Carpenter. Optimization of viscoelastic compliant walls for transition delay. *AIAA Journal*, 32(2):256–267, 1994.

J. A. Domaradzki and R. W. Metcalfe. Stabilization of laminar boundary layers by compliant membranes. *Physics of Fluids*, 30:695, 1987.

J. Dongarra and F. Sullivan. The top 10 algorithms. *Computing in Science & Engineering*, 2(1):22–23, 2000.

E. H. Dowell. *Aeroelasticity of plates and shells*. Springer, 1975.

J. H. Duncan, A. M. Waxman, and M. P. Tulin. The dynamics of waves at the interface between a viscoelastic coating and a fluid flow. *Journal of Fluid Mechanics*, 158(1):177–197, 1985.

N. S. Elliott. Syrinx fluid transport: modelling pressure-wave-induced flux across the spinal pial membrane. *Journal of Biomechanical Engineering*, 134(3): 031006, 2012.

N. S. J. Elliott, A. D. Lucey, M. Heil, P. R. Eastwood, and D. R. Hillman. Modelling and simulation of fluid-structure interactions in human snoring. In *Proceedings MODSIM2011, 19th International Congress on Modelling and Simulation*, December 2011.

M. C. Fisher, L. M. Weinstein, D. M. Bushnell, and R. L. Ash. Compliant-wall turbulent skin-friction-reduction research. *AIAA Journal*, pages 75–833, 1975.

M. Gad-el-Hak. Boundary layer interactions with compliant coatings: an overview. *Applied Mechanics Reviews*, 39:511, 1986.

M. Gad-el-Hak. Compliant coatings research: A guid to the experimentalist. *Journal of Fluids and Structures*, 1:55–70, 1987.

M. Gad-el-Hak. Compliant coatings: A decade of progress. *Applied Mechanics Reviews*, 49(10):147–157, 1996.

M. Gad-el-Hak. Compliant coatings: the simpler alternative. In *Proceedings EFTS*, 1998.

M. Gad-el-Hak. Compliant coatings for drag reduction. *Progress in Aerospace Sciences*, 38:77–99, 2002.

M. Gad-el-Hak, R. F. Blackwelder, and J. J. Riley. On the interaction of compliant coatings with boundary-layer flows. *Journal of Fluid Mechanics*, 140(1):257–280, 1984.

A. D. Garrad and P. W. Carpenter. On the aerodynamic forces involved in aeroelastic instability of two-dimensional panels in uniform incompressible flow. *Journal of Sound and Vibration*, 80(3):437–439, 1982.

M. Gaster. Is the dolphin a red herring? In *Proceedings of IUTAM Symposium on Turbulent Management and Relaminarization*, pages 285–304, Bangalore, 1987. Springer.

J. F. Gerbeau, M. Vidrascu, and P. Frey. Fluid–structure interaction in blood flows on geometries based on medical imaging. *Computers and Structures*, 83 (2):155–165, 2005.

G. Greengard. The core-spreading method approximates the wrong equation. *Journal of Computational Physics*, 61:345–348, 1985.

L. Greengard and V. Rokhlin. A Fast Algorithm for Particle Simulations. *Journal of Computational Physics*, 73:325–348, 1987a. URL `citeseer.ist.psu.edu/greengard87fast.html`.

L. Greengard and V. Rokhlin. The rapid evaluation of potential fields in three dimensions. *Lecture Notes in Mathematics*, 1360:121–41, 1987b.

J. B. Grotberg and O. E. Jensen. Biofluid mechanics in flexible tubes. *Annual Review of Fluid Mechanics*, 36:121–147, 2004.

A. F. Gunns. The first tacoma narrows bridge: A brief history of galloping gertie. *The Pacific Northwest Quarterly*, 72(4):162–169, 1981.

J. T. Hamilton and G. Majda. On the Rokhlin-Greengard Method with Vortex Blobs for Problems Posed in All Space or Periodic in One Direction. *Journal of Computational Physics*, 121(1):29–50, 1995.

M. Heil and A. L. Hazel. Fluid-structure interaction in internal physiological flows. *Annual Review of Fluid Mechanics*, 43:141–162, 2011.

M. Heil and O. E. Jensen. Flows in deformable tubes and channels. *Flow past highly compliant boundaries and in collapsible tubes*, pages 15–49, 2003.

J. L. Hess and A. M. O. Smith. Calculation of potential flow about arbitrary bodies. *Progress in Aerospace Sciences*, 8(0):1 – 138, 1967. ISSN 0376-0421. doi: 10.1016/0376-0421(67)90003-6. URL `http://www.sciencedirect.com/science/article/pii/0376042167900036`.

R. M. Howell and A. D. Lucey. Energy production characteristics of a spring-mounted cantilevered-free flexible plate in a uniform flow. In *Proceedings ASME 2012 Fluids Engineering Summer Meeting: 13th Symposium on Fluid-Structure Interaction and Flow-Induced Noise in Industrial Applications*, Puerto Rico, USA, July 2012.

M. J. Huang. Diffusion via splitting and remeshing via merging in vortex methods. *International Journal for Numerical Methods in Fluids*, 48:521–539, 2005.

M. G. Y. T. Huckle and M. Grote. Parallel preconditioning with sparse approximate inverses. *SIAM Journal on Scientific Computing*, 18(3):838–853, 1997.

O. Inoue. Vortex simulation of a turbulent mixing layer. *AIAA Journal*, 23(3): 367–373, 1985.

O. Inoue and A. Leonard. Vortex simulation of forced/unforced mixing layers. *AIAA Journal*, 25(11):1417–1418, 1987.

E. Jones, T. Oliphant, P. Peterson, et al. SciPy: Open source scientific tools for Python, 2001–. URL `http://www.scipy.org/`.

R. E. Kaplan. *The stability of laminar incompressible boundary layers in the presence of compliant boundaries*. PhD thesis, Massachusetts Institute of Technology, Cambridge, MA, 1964.

J. S. Kapor and A. D. Lucey. Flow-induced vibrations of a flexible panel in a boundary-layer flow. In *Flow-Induced Vibration, 10th International Conference on Flow-Induced Vibration (& Flow-Induced Noise) FIV2012*, pages 647–654, Trinity College Dublin, Ireland, July 2012.

J. S. Kapor, A. D. Lucey, and M. W. Pitman. Boundary-Layer Hydrodynamics using Mesh-Free Modelling. In *Proceedings 18th World IMACS Congress and MODSIM09 International Congress on Modelling and Simulation*, July 2009.

J. S. Kapor, A. D. Lucey, and M. W. Pitman. Fluid-structure interaction using mesh-free modelling. In *Proceedings 6th Australasian Congress on Applied Mechanics*, Perth, Australia, December 2010.

J. S. Kapor, A. D. Lucey, and M. W. Pitman. A mesh-free compliant-wall fluid-structure interaction model. In *Proceedings MODSIM2011, 19th International Congress on Modelling and Simulation*, December 2011.

J. Katz and A. Plotkin. *Low Speed Aerodynamics: From Wing Theory to Panel Methods*. McGraw-Hill, 1991.

Z. Khatir and A. D. Lucey. A combined boundary integral and vortex method for the numerical study of three-dimensional fluid flow systems. *International Journal of Computer Mathematics*, 89(11):1504–1524, 2012.

D. A. Knoll and D. E. Keyes. Jacobian-free newton-krylov methods: a survey of approaches and applications. *Journal of Computational Physics*, 193(2): 357–397, 2004.

M. O. Kramer. Boundary-layer stabilization by distributed damping. *Journal of the Aeronautical Sciences*, 24:459, 1957.

M. T. Landahl. On the stability of a laminar incompressible boundary layer over a flexible surface. *Journal of Fluid Mechanics*, 13:609–632, 1962.

T. C. W. Lau, R. M. Kelso, and E. R. Hassan. Flow visualisation of a pitching and heaving hydrofoil. In *Proceedings of the Fifteenth Australasian Fluid Mechanics Conference*, University of Sydney, Australia, 2004.

T. Lee, M. Fisher, and W. H. Schwarz. Investigation of the effects of a compliant surface on boundary-layer stability. *Journal of Fluid Mechanics*, 288(1):37–58, 1995.

A. Leonard. Vortex methods for flow simulation. *Journal of Computational Physics*, 37(3):289–335, 1980.

A. Leonard. Computing three-dimensional incompressible flows with vortex elements. *Annual Review of Fluid Mechanics*, 17:523–559, 1985.

A. Leonard and P. R. Spalart. Computation of separated flows by a vortex tracing algorithm. In *AIAA-81-1246*, 1981.

R. I. Lewis. *Vortex Element Methods for Fluid Dynamic Analysis of Engineering Systems*. Cambridge University Press, 1991.

K. Lindsay and R. Krasny. A particle method and adaptive treecode for vortex sheet motion in three-dimensional flow. *Journal of Computational Physics*, 172 (2):879–907, 2001.

A. D. Lucey and P. W. Carpenter. A study of the hydroelastic stability of a compliant panel using numerical methods. *International Journal of Numerical Methods for Heat and Fluid Flow*, 2:537–553, 1992.

A. D. Lucey and P. W. Carpenter. Boundary layer instability over compliant walls: Comparison between theory and experiment. *Physics of Fluids*, 7:2355, 1995.

A. D. Lucey, G. J. Cafolla, and P. W. Carpenter. Numerical simulation of a boundary-layer flow interacting with a passive compliant boundary. *Lecture Notes in Physics*, pages 406–411, 1997a.

A. D. Lucey, P. W. Carpenter, G. J. Cafolla, and M. Yang. The nonlinear hydroelastic behaviour of flexible walls. *Journal of Fluids and Structures*, 11:717–744, 1997b.

Y. Matsuzaki. Stability of flat plates and cylindrical shells exposed to flows. *Encyclopedia of Fluid Mechanics*, 1:476–509, 1986.

D. W. Moore. The discrete vortex approximation of a finite vortex sheet. *Californian Institute of Technology Report*, 1971.

B. R. Munson, D. F. Young, and T. H. Okiishi. *Fundamentals of Fluid Mechanics*. Wiley, New York, 4th edition, 2002.

C. R. Nisewanger. Flow noise and drag measurements of vehicle with compliant coating. Technical report, US Naval Ordnance Test Station Report No. 8518, NOTS No. TP-3510, China Lake, California, 1964.

T. Pedley. Longitudinal tension variation in collapsible channels: a new mechanism for the breakdown of steady flow. *Journal of Biomechanical Engineering*, 114:60, 1992.

M. Perlman. On the accuracy of vortex methods. *Journal of Computational Physics*, 59:200–233, 1985.

M. W. Pitman. *An Investigation of Flow Structure Interactions on a Finite Compliant Surface Using Computational Methods*. PhD thesis, Curtin University of Technology, 2007.

M. W. Pitman and A. D. Lucey. A deterministic viscous vortex method for grid-free cfd with moving boundary conditions. *IASME Transactions*, 1(4):591–596, 2004.

M. W. Pitman and A. D. Lucey. On the direct determination of the eigenmodes of finite flow–structure systems. *Proceedings of the Royal Society A: Mathematical, Physical and Engineering Science*, 465(2101):257–281, 2009.

P. Ploumhans and G. S. Winckelmans. Vortex methods for high-resolution simulations of viscous flow past bluff bodies of general geometry. *Journal of Computational Physics*, 165(2):354–406, 2000.

G. J. Pringle. *Numerical Study of Three-Dimensional Flow using Fast Parallel Particle Algorithms*. PhD thesis, Napier University, Edinburgh, Scotland, 1994.

F. W. Puryear. Boundary-layer control drag reduction by compliant surfaces. Technical report, David W. Taylor Model Basin, 1962.

R. D. Richtmyer and K. W. Morton. *Different methods for initial value problems.* Interscience, 1967.

J. J. Riley, M. Gad-el-Hak, and R. W. Metcalfe. Compliant coatings. *Annual Review of Fluid Mechanics*, 20(1):393–420, 1988.

H. Ritter and L. T. Messum. Water tunnel measurements of turbulent skin friction on six different compliant surfaces of 1 ft length. Technical report, British Admiralty Research Laboratory Report No. ARL/N4/GHY/9/7, ARL/G/N9, London, Great Britain, 1964.

H. Ritter and J. S. Porteous. Water tunnel measurements of turbulent skin friction on a compliant coating. Technical report, British Admiralty Research Laboratory Report No. ARL/N3/G/HY/9/7, London, Great Britain, 1965.

L. Rosenhead. The formation of vorticies from a surface of discontinuity. *Proceedings of the Royal Society of London*, 134:170–192, 1931.

L. F. Rossi. Resurrecting core spreading vortex methods: A new scheme that is both deterministic and convergent. *SIAM Journal on Scientific Computing*, 17:370–397, 1996.

L. F. Rossi. Merging computational elements in vortex simulations. *SIAM Journal on Scientific Computing*, 18:1014, 1997.

P. L. Rui and R. S. Chen. An efficient sparse approximate inverse preconditioning for FMM implementation. *Microwave and Optical Technology Letters*, 49(7): 1746–1750, 2007.

Y. Saad and M. H. Schultz. Gmres: A generalized minimal residual algorithm for solving nonsymmetric linear systems. *SIAM Journal on Scientific and Statistical Computing*, 7(3):856–869, 1986.

Y. Saad and H. A. Van Der Vorst. Iterative solution of linear systems in the 20th century. *Journal of Computational and Applied Mathematics*, 123(1):1–33, 2000.

P. G. Saffman and G. R. Baker. Vortex interactions. *Annual Review of Fluid Mechanics*, 11:95–122, 1979.

A. Saitoh and A. Kamitani. GMRES with new preconditioning for solving BEM-type linear system. *IEEE Transactions on Magnetics*, 40(2):1084–1087, 2004.

T. Sarpkaya. Computational methods with vortices-the 1988 freeman scholar lecture. *Journal of Fluids Engineering*, 111:5–52, 1989.

R. R. Schaller. Moore's law: past, present and future. *IEEE Spectrum*, 34(6):52 –59, jun 1997. ISSN 0018-9235. doi: 10.1109/6.591665.

H. Schlichting. *Boundary-Layer Theory*. McGraw-Hill, New York, 1979.

R. L. Smith and E. F. Blick. Skin friction of compliant surfaces with foamed material substrate. *Journal of Hydronautics*, 3:100–103, 1969.

H. Takami. Numerical experiment with discrete vortex approximation, with reference to the rolling up of a vortex sheet. *Department of Aeronautics and Astronautics, Stanford University Report*, 202, 1964.

K. Takeda, O. R. Tutty, and A. D. Fitt. A comparison of four viscous models for the discrete vortex method. *AIAA Journal*, 1997.

L. Tang, M. P. Païdoussis, and J. Jiang. Cantilevered flexible plates in axial flow: Energy transfer and the concept of flutter mill. *Journal of Sound and Vibration*, 326(2):263–276, 2009.

L. N. Trefethen. *Finite difference and spectral methods for ordinary and partial differential equations*. Cornell University, 1996.

H. Wang, H. Zhang, X. Wang, Y. Guo, and W. Lin. Numerical simulation of mixing layer flow simulation. *Qinghua Daxue Xuebao/Journal of Tsinghua University*, 39(6):106 – 109, 1999. ISSN 1000-0054.

Z. Wang, K. S. Yeo, and B. C. Khoo. Spatial direct numerical simulation of transitional boundary layer over compliant surfaces. *Computers and Fluids*, 34:1062–1095, 2005.

F. L. Westwater. Rolling up of the surface of discontinuity behind an aerofoil of finite span. *Aeronautical Research Council Report*, 1692, 1935.

G. J. K. Willis. *Hydrodynamic stability of boundary layers over compliant surfaces.* PhD thesis, University of Exeter, 1986.

K. Yeo. The stability of boundary-layer flow over single- and multi-layer viscoelastic walls. *Journal of Fluid Mechanics*, 196:259–408, 1988.

K. S. Yeo, B. C. Khoo, and H. Z. Zhao. The absolute instability of boundary-layer flow over viscoelastic walls. *Theoretical and Computational Fluid Dynamics*, 8 (4):237–252, 1996.

K. S. Yeo, B. C. Khoo, and H. Z. Zhao. The convective and absolute instability of fluid flow over viscoelastic compliant layers. *Journal of Sound and Vibration*, 223(3):379–398, 1999.

R. Yokota, L. A. Barba, and M. G. Knepley. PetRBF - a parallel O(N) algorithm for radial basis function interpolation with gaussians. *Computer Methods in Applied Mechanics and Engineering*, 199:1793–1804, 2010.

R. Yokota, T. Narumi, L. A. Barba, and K. Yasuoka. Petascale turbulence simulation using a highly parallel fast multipole method. *Arxiv preprint arXiv:1106.5273*, 2011.