Characterization and Optimization of a Fully Passive Flapping Foil in an Unsteady Environment for Power Production and Propulsion

Thesis by

Morgan Louise Hooper

In Partial Fulfillment of the Requirements for the
Degree of
Doctor of Philosophy in Aeronautics

Caltech

CALIFORNIA INSTITUTE OF TECHNOLOGY
Pasadena, California

2022
Defended May 25, 2022

ACKNOWLEDGEMENTS

For my parents, for my sisters, for found family near and far, and for me.

It has been my privilege to complete this work with the support of my adviser, Dr. Beverley McKeon. I have learned so much from her, not only about science (of course!), but also about mentorship, community building, and leadership within academia and beyond. Thank you for helping me discover my role in our community, and for helping me to grow as a scientist and individual, authentically as me.

Thank you as well to my thesis and candidacy committee members, Joanna Austin, Dan Meiron, Steve Brunton, and Tim Colonius. Our conversations have provided many fresh perspectives, as well as valuable insights for the work in this thesis. In addition, the many conversations I've had with Professors Meiron and Austin in their role(s) as Option Representative, regarding pedagogy and the graduate student experience, were formative to my identity as an educator. Thank you as well to the Center for Teaching, Learning and Outreach (CTLO), and in particular to Jennifer Weaver, for helping me develop my interest in pedagogy. To the wonderful administrative professionals I have had the privilege to interact with, Jamie Meighen-Sei, Christine Ramirez, and Liza Bradulina, thank you for everything you have done to support me, our community, and to make GALCIT a welcoming place.

This experience would not have been the same without the wonderful folks in the McKeon Research Group. Although all of you have impacted both me and this work in more ways than I can list, a few members of the group have had particular influence on my growth as a scientist and individual. Thank you to Jacqueline Tawney for teaching me to be brave, and to seek balance both scientifically and personally. Thank you to Angeliki Laskari for showing me how to be a more resilient scientist, and for encouraging me to chase my dreams. And thank you to Benedikt Barthel for believing in me since day one, especially when I didn't believe in myself.

Finally, to my family: my parents, my sisters, my partner, and my found families in both Pasadena and Toronto, I would never have made it this far without your love and support. Thank you for talking science with me, for keeping me grounded, for believing in me, and for reminding me that I am enough.

Support for this work from the US Army Research Office under grant number W911NF-17-1-0306, as well as that from the Natural Sciences and Engineering Research Council of Canada (NSERC) is gratefully acknowledged.

ABSTRACT

This thesis presents an experimental window into the duality between thrust production and energy harvesting by a flapping foil subject to unsteadiness in an oncoming flow. In particular, an airfoil is placed downstream of a circular cylinder, and allowed to interact with the vorticity shed in its wake to produce motions in both the transverse and streamwise directions. It is shown that under the right conditions, passive fluid-structure interactions arising from such a configuration can permit simultaneous extraction of energy from the flow, coupled with net thrust larger than net drag experienced by the airfoil. This observation was made previously by Beal et al. (2006), where in addition they showed that a dead fish under similar conditions appeared to swim upstream.

The contributions of the present work are threefold. Firstly, we provide measurements of the forces acting on a flapping foil in the wake of a circular cylinder and the airfoil motion that arises, for cases where the flapping motion is both active (the foil is driven through a pre-planned trajectory) and fully passive (the foil is allowed to react to the fluid forcing it experiences). These are coupled with simultaneous Particle Image Velocimetry (PIV) measurements of the flow field in the region of the airfoil. These measurements allow us to directly observe fluid-structure interactions which give rise to both thrust production and power extraction potential for the airfoil, illuminating the mechanisms driving each. It is determined that for the sinusoidal trajectories considered here, the dynamics of a fully passive flapping foil can be tuned to mirror the behaviour of a similar driven one, and the measured forces as well as fluid-structure interactions taking place are similar between the two cases.

The second focus of this work is on the optimization of the behaviour of a compliantly mounted, fully passive flapping foil for energy harvesting. A framework based on 2^{nd} order linear systems theory is proposed to guide the optimization of a simplified flapping foil energy harvester, where the dynamics are determined based on spring-mass-damper characteristics of the mounting system. Tuning efforts are shown to yield significant improvements to power extraction performance relative to a naïve choice of mounting parameters, however nonlinear feedback between airfoil motion and the aerodynamic forces it experiences acts to temper the improvements seen in experiments relative to predicted power extraction performance. In addition, the effects of parasitic dynamics due to friction in the mounting mechanism are investigated, and the resulting changes to power production performance are quan-

tified. The action of friction induces emergent behaviours for the foil not seen in the ideal case; thus, understanding these effects is key to predicting and optimizing the performance of a real engineering system.

Finally, in addition to transverse flapping, we explore the behaviour of a fully passive airfoil when it is allowed to react to the oncoming cylinder wake in the streamwise direction as well. Since the airfoil produces net thrust larger than its net drag, we observe it translating upstream, while simultaneously extracting energy from the flow. We confirm through PIV imaging that the airfoil begins translating upstream well outside of the suction region induced by the presence of the upstream cylinder, such that all thrust generated is due to its interactions with vorticity in the cylinder wake. These observations are enabled through Cyber-Physical Fluid Dynamics (CPFD), where both the transverse and streamwise behaviour of the airfoil is determined through feedback control based on measured forces acting on the foil. Such a system allows access to simulated fully passive dynamics over a range of parameter space challenging to reach using a conventional experimental setup.

PUBLISHED CONTENT AND CONTRIBUTIONS

Hooper, M. L., & McKeon, B. J. (2021). A representative driven system to interrogate passive dynamics of an airfoil in the wake of a cylinder [Published Online. Available from https://ispiv21.library.iit.edu/]. *Proceedings of the 14th International Symposium on Particle Image Velocimetry*, *1*(1)

Morgan Hooper designed the experiments, collected and analyzed the data, created the figures, and was the primary author/presenter of the contribution(s).

Hooper, M. L., Shamai, M., & McKeon, B. J. (2019). V0017: A Blacklight Ballet: Flow visualization reveals intricacies in the wake of a streamwise-oscillating cylinder. *Proceedings of the 72th Annual Meeting of the APS Division of Fluid Dynamics*. https://doi.org/https://doi.org/10.1103/APS.DFD.2019.GFM.V0017

Morgan Hooper collaborated on completion of experiments and subsequent analysis, and created the video submission.

TABLE OF CONTENTS

LIST OF ILLUSTRATIONS

LIST OF TABLES

Number *Page*

Chapter 1

INTRODUCTION

1.1 Motivation

It all started with a dead fish. In 2006, Beal et al. showed that such a body placed downstream from a circular cylinder in an oncoming flow could passively interact with the wake of the upstream object to produce enough thrust to overcome its net drag, and give the appearance of swimming upstream. This macabre tableau is more than just a curiosity. In the same study, Beal et al. were able to show that a compliantly mounted airfoil placed in a similar wake was not only able to passively generate net thrust larger than its net drag, but it was also able to simultaneously extract net power from the flow through passive flow-induced motion in both heave and pitch (Beal et al., 2006). This dual power and thrust production phenomenon, and its potential utility for passive flow-based energy harvesting and propulsion in unsteady environments, is the topic of this thesis.

The use of flapping foils for energy harvesting is not novel. In 1981, McKinney and DeLaurier proposed an 'Oscillating-Wing Windmill' as an alternative to the conventional rotary windmill, which has been in use for hundreds of years. Through the application of unsteady aerodynamic theory as well as wind tunnel testing, they found that such a configuration had the potential to achieve power production efficiencies comparable to rotary windmills (McKinney & DeLaurier, 1981). In the years since, technology for conventional rotary windmills (the modern incarnation of which are referred to as Horizontal-Axis Wind Turbines or HAWTs) has advanced considerably, with gains in real operating efficiency significantly outpacing alternative designs for flow-based energy harvesters (Sivaram et al., 2018). Recently however, interest in mechanically simple, small-scale energy harvesting devices for low-power, urban, and environment-agnostic energy production has grown (Zhu, 2011): a cross-section of such alternative flow-based energy harvesters was presented by Wang et al. (2020). The configuration discussed in this thesis represents an archetype for an energy harvesting device poised to take advantage of naturally occurring unsteadiness in the environment, both natural and man-made. Although the development of an engineering device for use in application is outside the scope of this work, detailed investigation of the fluid-structure interactions which give rise to favourable power production characteristics for such a configuration provide the backbone upon which future devices may be designed and built.

In addition to utility as an energy harvester, the configuration considered here also holds interest to the field of efficient navigation in unsteady fluid environments. Investigations of the performance of flying and swimming animals and vehicles across a range of scales and complexities interacting with unsteadiness have revealed that despite the challenges inherent in navigating such environments, significant gains to lift performance and propulsive efficiency can be realized by taking advantage of structure in the oncoming flow. This topic has a long history (see for example Wu and Chwang (1975); Streitlien et al. (1996); Liao et al. (2003); Lefebvre and Jones (2019); SureshBabu et al. (2021)), and although significant progress has been made in understanding the complex interactions between vorticity in the free-stream and unsteady manoeuvring, difficulty in performing experiments or computations under sufficiently realistic and repeatable unsteady conditions as well as sensitivity to instantaneous behaviour of the flow continues to render generalizable results elusive (Jones et al., 2022).

The experiments and analysis presented in this thesis advance understanding of the interactions between oncoming vortical flow disturbances and an aerodynamic body undergoing unsteady manoeuvres in order to achieve performance enhancement, in a setting where unsteady effects such as flow separation are dominant contributors to the observed dynamics. We consider a nominally simple experimental configuration: an airfoil is placed downstream from a circular cylinder and experiences a cyclic, vortical disturbance to the mean oncoming flow due to classical von Kármán vortex shedding. This forcing is both realistic, as such flows are ubiquitous in nature as well as in the built environment, and relatively repeatable; it thus offers an appropriate compromise between complexity and tractability for the present experiments. Since this configuration can lead to the simultaneous potential for thrust production and energy extraction, as described by Beal et al. (2006), these experiments bridge the gap between studies primarily of thrust performance (where active flapping is presumed to require net energy expenditure), and power production performance (where net power extraction is possible despite the required energy input). Cyber-Physical Fluid Dynamics (CPFD) capabilities are leveraged to perform unsteady manoeuvres on top of unsteadiness in the oncoming flow, based on force-feedback from the airfoil itself. Such experiments allow detailed observations of fluid-structure interactions that give rise to the observed power and/or thrust performance for both an actively and passively flapping foil. In addition, the use of CPFD provides easy access to a wide range of unsteady airfoil dynamics, to facilitate the optimization of energy extraction performance for the system subject to realistic engineering constraints.

The remainder of this chapter provides a brief summary of the relevant background for the work presented in this thesis, including discussion of active, semi-passive and fully passive flow energy harvesters, flapping foil propulsion, and airfoils encountering disturbance(s) to the oncoming flow. In addition, a brief history of the use of CPFD to facilitate such studies in unsteady aerodynamics, and their impact on the present work, is provided.

1.2 Flapping Foils in a Uniform Free Stream for Power and Propulsion

Flapping foils represent a simple engineering model for a wide variety of systems, from bird and insect wings and fish fins (Wu, 2011) to Vertical-Axis Wind Turbine (VAWT) blades (Dunne & McKeon, 2015). Unlike conventional fixed-wing architectures common in modern aerospace vehicles (aircraft, for example), which often suffer performance penalties when encountering unsteady flow conditions (Jones et al., 2022), flapping-wing systems often rely directly on unsteadiness and related fluid-structure interactions for the production of both lift and thrust (Corke and Thomas, 2015; Liu et al., 2016; Kinsey and Dumas, 2008).

Even when limited to the simplified case where the oncoming free-stream is uniform, the combined fields of flapping-foil propulsion and energy extraction represent a significant body of work, far too broad to be exhaustively summarized in the handful of pages available here. Therefore in this section, those studies which formed the basis for the contributions presented in this thesis are discussed, to provide context for further analysis presented in the following chapters. In addition, key concepts, metrics, and terminology commonly used throughout the thesis are presented.

1.2.1 Anatomy of a Flapping Foil

For airfoils (or any other aerodynamic body) operating in a uniform free-stream, characterized by an oncoming flow speed U_∞, a dualism exists between the goals of thrust production and power extraction. For propulsive systems, flapping motions of a foil are largely *active*, that is, external energy is put into the system to achieve some specified trajectory (often sinusoidal) for both heave and pitch motions. Proper selection of these trajectories leads to propulsion, or the condition that the net thrust generated by the flapping motion is larger than the experienced drag.

By contrast, flapping foils used for power extraction from a uniform oncoming free-stream must generate constructive unsteadiness that enhances their ability to extract energy from the flow. This is commonly achieved in one of two ways: active or semi-passive flapping, or exploitation of aeroelastic effects to passively generate oscillatory motion. The latter method relies on fluid-structure coupling at a frequency intrinsic to the foil itself; by contrast, the behaviours considered in the present study rely on the application of an external forcing, either due to active flapping motion(s) of the foil, or in later discussions due to external cyclic flow forcing. Thus, self-excited airfoil motions due to aeroelastic flutter or other related effects are outside the scope of the this discussion; instead, we consider flapping-foil energy harvesters where the frequenc(ies) of interest are externally enforced.

The general configuration for both energy-harvesting and thrust-producing flapping foils is shown in Figure 1.1. It is important to note that in contrast to propulsive foils, which are usually considered to be free to translate (at least in the streamwise direction) the dynamics of an energy harvesting device depend explicitly on the mounting system used to physically extract energy from the flow. Thus, the dynamics of such a flapping foil are considered to be the combined dynamics of both the foil itself, and its (potentially active) mounting hardware.

Several metrics are commonly used to describe the performance of flapping-foil

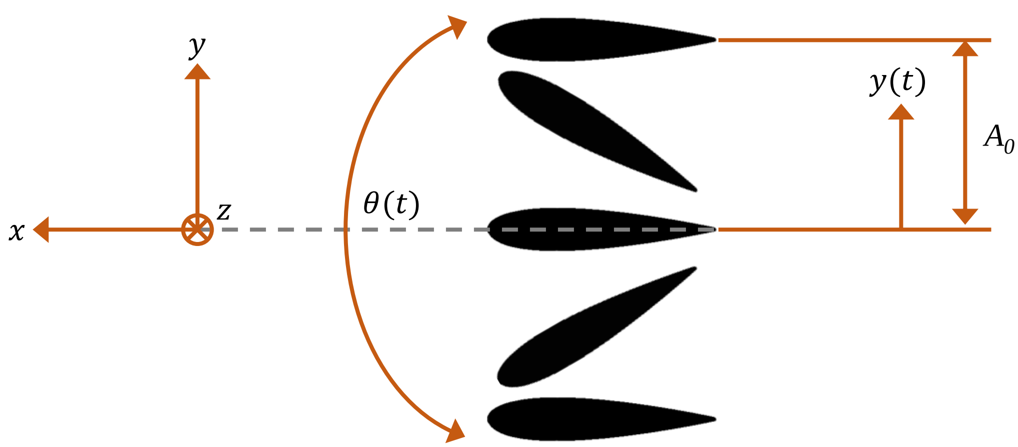

Figure 1.1: A basic schematic of a flapping foil for propulsion or energy harvesting. $\theta(t)$, $y(t)$ represent arbitrary motions in pitch and heave as a function of time, though for flapping foil motions it is conventional for these trajectories to be periodic with the same frequency f. A_0 represents the maximum excursion of the foil from its neutral position during a cycle, commonly determined based on either the quarter-chord location, or the location of the trailing edge.

systems. The Coefficient of Lift (C_L) and Thrust (C_T) are given as

$$C_L = \frac{L}{\frac{1}{2}\rho U_\infty^2 sc} \qquad\qquad C_T = \frac{T}{\frac{1}{2}\rho U_\infty^2 sc}, \tag{1.1}$$

where ρ is the fluid density [kg/m^3], U_∞ is the oncoming free-stream flow speed [m/s], s is the span of the foil [m], and c is its chord length [m]. The quantity $\frac{1}{2}\rho U_\infty^2$ is commonly referred to as the *dynamic pressure*, and denoted p_∞ throughout this thesis. These definitions coincide with those conventionally used in steady aerodynamics, a thorough discussion of which is omitted here; the interested reader can find a detailed discussion of that field in Anderson (2011), for example. One notable difference in the present case is the common use of C_T in place of the Coefficient of Drag (C_D) to describe the dynamics of the foil in the streamwise direction: since particularly in the case of propulsive systems the streamwise force is oriented upstream, it is more common to use $C_T > 0$ for such forces, rather than an equivalent $C_D < 0$ (though this is seen in studies of power extraction, for example Beal et al. (2006)). Throughout this thesis, $C_T > 0$ indicates the force acting on the foil in the streamwise direction is oriented upstream against the oncoming flow.

To characterize the power extraction of a flapping foil energy harvester, we consider the work done on the foil by the flow as a function of time. For a foil oscillating in both heave and pitch, the power and resulting Power Coefficient (C_P) are given by

$$P(t) = F_y(t)\dot{y}(t) + \tau_z(t)\dot{\theta}(t), \tag{1.2}$$

$$C_P = \frac{P}{\frac{1}{2}\rho U_\infty^3 sc} = \frac{P}{U_\infty p_\infty sc}, \tag{1.3}$$

where $F_y(t)$ and $\tau_z(t)$ are the transverse force [N] and pitching torque [Nm] acting on the foil respectively, and $\dot{y}(t)$ and $\dot{\theta}(t)$ are the corresponding translational and angular velocities [m/s, 1/s]. For propulsive systems, the definition above is also commonly used; however, in that case the quantity P generally refers to the power supplied to the foil by motors and actuators to do work on the flow. Thus, there is often an implicit negative sign included such that power consumed by the airfoil's motion corresponds to $P > 0$, in contrast to the energy harvesting case where energy extracted from the flow corresponds to $P > 0$. This generates some notational confusion, especially for studies similar to the present case where the sign of P as written in Equation 1.2 is not known a priori, and may cross zero during one cycle (for example, in the case of Beal et al. (2006)). In this thesis, power is calculated strictly according to Equation 1.2, where $F_y(t)$ and $\tau_z(t)$ represent the force/torque exerted on the airfoil by the flow, and $P > 0$ therefore corresponds to power extraction by the foil. For periodic foil motions of interest in this thesis, it is practical to consider a cycle-averaged power input/output to the foil:

$$\overline{P} = \int_0^T P(t)\ \mathrm{dt}, \tag{1.4}$$

where T is the period of oscillation. The form of $\overline{P}$ above suggests that for periodic airfoil dynamics, $\overline{P}$ is maximized when $F_y(t)$ is in-phase with $\dot{y}(t)$ and $\tau_z(t)$ is in-phase with $\dot{\theta}(t)$, while a phase shift of $\pm 90°$ between these quantities would result in zero net power extraction.

A final metric used to characterize the performance of both propulsive and energy-extracting flapping foils is the efficiency, η. This metric is defined differently in these different fields however, reflecting the often competing goals of energy harvesting and thrust production. The propulsive efficiency, η_P, is defined as

$$\eta_P = \frac{U\overline{F_x}}{-\overline{P}}, \tag{1.5}$$

where U is the oncoming flow speed (accounting for forward motion of the foil), $\overline{F_x}$ is the cycle-average thrust it generates, and $\overline{P} < 0$ is the power input required to sustain motion (such that when thrust is positive and input power is required, $\eta > 0$). By contrast, for studies of energy extraction a more appropriate definition is the power extraction efficiency, η_E, given by

$$\eta_E = \frac{\overline{P}}{U_\infty p_\infty A_s}, \tag{1.6}$$

where A_s is the 'swept area' of the foil's motion, or the vertical extent of its transverse trajectory multiplied by the span s. For the foil motion presented in Figure 1.1, $A_s = 2A_0 s$. This represents the ratio of power extracted by the foil relative to the kinetic energy flux through the region the foil 'sweeps out' during one operating cycle, and is positive for net energy extraction ($\overline{P} > 0$).

1.2.2 Efficient Flapping Motions for Thrust and Power Production

Though in general both η_P and η_E depend strongly on the specifics of a particular system under test, several general trends regarding system performance have been observed. For systems that oscillate in both heave and pitch, parameters of interest include the heaving amplitude A_0, the pitching amplitude θ_0, the phase angle between pitch and heave motions, and the frequency of these motions parameterized by the

Strouhal Number,

$$St = \frac{2A_0 f}{U_\infty}, \tag{1.7}$$

where f [Hz] describes the frequency of oscillation in both pitch and heave. Also of interest is the effective angle of attack for the foil, α_{eff}, given by

$$\alpha_{\text{eff}}(t) = \text{atan}\left[\frac{-\dot{y}(t)}{U_\infty}\right] + \theta(t) \tag{1.8}$$

for the systems in question (although modifications to Equation 1.8 are required in the case that the oncoming flow itself exhibits unsteadiness or the airfoil translates in the streamwise direction, as discussed later in this thesis). Many studies have determined regions of this parameter space where either thrust or power production efficiency is enhanced; the following discussion presents several studies of particular relevance to the experimental configuration considered in this thesis.

Read et al. (2003) conducted an extensive experimental study of flapping foil propulsion in a uniform oncoming free stream, at a Reynolds number very close to that considered in the present experiments. They found that sinusoidal trajectories with pitch leading heave motions by approximately 90° led to both high thrust and high propulsive efficiency η_P. They also found that increasing St, either through increasing f or A_0 generally led to larger values of C_T over the range of St included in the study ($St \approx [0.1, 0.6]$); however this did not always lead to larger η_P. Large maximum values of α_{eff} in the vicinity of 35° were found to enhance thrust performance at the higher St values tested; however such large $\alpha_{\text{eff, max}}$ also often led to reduced efficiencies. The highest propulsive efficiency the authors recorded ($\eta_P = 0.715$) occurred at $St = 0.16$, $\alpha_{\text{eff, max}} = 15°$, and $A_0/c = 0.75$, though they make no claims as to the optimality of this result. In fact, in a prior study of propulsion by a similar flapping foil, Anderson et al. (1998) demonstrated efficiencies up to $\eta_P = 0.87$ by tuning the heaving and pitching motions of the foil for optimal formation of leading-edge vorticity and a Reverse von Kármán-type wake. The connection between flapping foil performance and wake formation is discussed further in Section 1.2.4.

Energy extraction performance for flapping foils, described by the extraction efficiency η_E, is governed by a similar set of parameters to the propulsive efficiency; however η_E is maximized for cases when $\overline{P} > 0$ is large. Energy harvesters based on flapping foils may undergo driven motions in both heave and pitch (as in the thrust producing case), in only one of the two axes, or neither. Xiao and Zhu (2014) present an overview of the performance of such active, semi-passive, and self-sustained system archetypes respectively. One general result, discussed for active flapping by Zhu (2011) and for semi-passive flapping by Su and Breuer (2019) is that optimal power extraction performance is achieved when the phase difference between heave and pitch motions is -90°(pitch leading heave), similar to the propulsion case. Both Zhu (2011) and Su and Breuer (2019) also found that the contribution of pitching motions to power extraction ($\tau_z(t)\dot{\theta}(t)$) was in general quite small compared to the contribution of transverse airfoil motions at conditions for optimal energy extraction. However, as discussed in Section 1.2.3, topological changes to the flow induced by airfoil pitching may still strongly affect power extraction due to transverse motion.

One study of flapping foil energy harvesting in particular has significantly informed the analysis presented in this thesis, and is referred to throughout the following chapters. Su and Breuer (2019) considered a semi-passive flapping-foil energy harvester, where a foil was driven in the pitching direction and allowed to react passively in the direction transverse to the (uniform) oncoming flow. The transverse dynamics of the foil were determined based on parameters of the mounting system, which provided linear 2$^{\text{nd}}$ order dynamics corresponding to a spring-mass-damper. The power extraction performance for this system was optimized when the mounting system was tuned such that external forcing (in this case, the frequency of pitching oscillations) coincided with the natural frequency of the mounting system, and transverse force $F_y(t)$ and velocity $\dot{y}(t)$ experienced by the airfoil were in-phase. The ability to achieve these conditions, related to the conditions for optimal power extraction efficiency discussed by Zhu (2011) and Kinsey and Dumas (2008) for example, by tuning mounting system parameters provides motivation and inspiration for the current work.

For systems that have potential to produce both net thrust and extract net power, the composite nature of the propulsive efficiency η_P makes it somewhat challenging to interpret. For example, in their numerical study of a foil interacting with oncoming vorticity, Streitlien et al. (1996) found that in general, high propulsive efficiencies were associated with the production of large thrust values, which simultaneously required high input power to sustain motion. These high-efficiency interactions corresponded to what those authors deemed the 'Close Interaction Mode', where the foil's motion brought it into close contact with oncoming vorticity. By contrast, based on their interpretation of that study and several others (including those by Anderson et al. (1998) and Gopalkrishnan et al. (1994), which are discussed in the following sections), Beal et al. (2006) identified a high-thrust interaction mode corresponding to close interactions with oncoming vorticity, but identified an alternative type of interaction which they deemed the 'Slaloming Mode' as that corresponding to high propulsive efficiency. This could reflect an underlying difference in the parameter spaces considered in these studies: Streitlien et al. (1996) did not consider choices for the flapping parameters where the input power approached zero, which represents the condition $\eta_P \to \infty$ if even a small amount of thrust is produced. Exactly such systems form the basis for the study by Beal et al. (2006). By considering such systems, the meaning of η_P as a metric for performance is distorted due to the presence of a pole at $P = 0$. This discrepancy illustrates the difficulty with η_P for systems that operate near this condition, as in the present study.

Based on this discussion, throughout this thesis the performance of the system studied is discussed without reference to efficiency, either η_P or η_E. This decision was made to ensure that the results are interpretable in the context of either power extraction or thrust production, and difficulties such as that illustrated above are avoided. Instead, performance is characterized by considering the values of C_P and C_T directly.

1.2.3 The Role of Unsteadiness in Flapping Foil Performance

Flapping foil motions such as those discussed above can produce aerodynamic unsteadiness beyond that simply due to foil motion. The origin of such unsteadiness is the occurrence of stall, a phenomenon by which the boundary layer over the 'top' (relative to the angle of attack) or suction side of an airfoil at high angle of attack separates, and the flow itself becomes unsteady often with intrinsic time scale(s) different from those imposed by airfoil motion. Studies of airfoil stall, and predictions regarding the angle of stall onset and subsequent evolution of the resulting values for C_L and C_D are numerous; Le Fouest et al. (2021) present a modern study of the process of static stall for a NACA 0018 airfoil under similar experimental conditions to those in this thesis. They find that the onset of static stall for this profile occurs at an angle of attack of 14.2°. Once stall has been initiated, the lift coefficient for the section drops precipitously from $C_L \approx 1$ to $C_L \approx 0.6$ (Le Fouest et al., 2021). These values are strongly effected by particular experimental conditions, including free-stream turbulence intensity, surface roughness, Reynolds Number (defined as $Re = U_\infty c/\nu$ for a flapping foil, where ν is the kinematic viscosity [m^2/s]), blockage of the experimental facility used for testing, and more. Therefore, Melani et al. (2019) provide a recent review of experimental and numerical studies for the same

NACA 0018 section, with many such operating conditions considered.

Although static stall generally leads to the deterioration of both lift and drag performance for a steady airfoil, unsteady motion can lead to an enhancement of lift beyond the static stall angle. The process giving rise to this unsteady lift enhancement is termed Dynamic Stall, and has been clearly visualized in a recent review by Corke and Thomas (2015), for example. Dynamic stall represents one instance in which the unsteady motion of an airfoil causes the development of a Leading-Edge Vortex (LEV); the formation and subsequent shedding of this LEV as the airfoil exceeds its static stall angle is responsible for the transient aerodynamic loads it experiences, including a significant enhancement of lift (Eldredge and Jones, 2019; Corke and Thomas, 2015; McCroskey, 1982). LEVs which arise cyclically due to oscillatory motion of a lifting surface with an amplitude large enough to exceed the static stall angle are often termed Dynamic Stall Vortices (DSVs).

Of particular interest to the present work is the effect of dynamic stall and related formation and shedding of LEVs on the thrust and power production performance for a flapping foil. It is well known that flying and swimming creatures across a wide range of scales take advantage of LEVs to enhance lift production (Eldredge & Jones, 2019). Similarly, Kinsey and Dumas (2008) showed that for a flapping foil at Reynolds numbers relevant to the experiments in this thesis, if the Strouhal number St and pitching amplitude θ for the flapping motion are chosen such that an LEV is formed on the suction side of the airfoil as it approaches a position extremum (and therefore a velocity of zero), the foil is able to maintain higher lift magnitudes in-phase with airfoil velocity. The formation of an LEV and its presence near the airfoil surface generates a region of low pressure, which temporarily enhances lift with the same sign as the foil velocity. Then, as the sign of velocity is reversed the LEV is shed, leading to a large drop in the experienced lift which facilitates a similar sign change in F_y. This leads to significant improvement in cycle-averaged power extraction from the system (Kinsey & Dumas, 2008). By contrast, formation and shedding of an LEV during dynamic stall dramatically increases the drag experienced by a flapping foil (Corke and Thomas, 2015; McCroskey, 1982).

The onset of dynamic stall is sensitive to the time history of airfoil motion up to the initiation of the growth of a DSV, including the rate of change of the angle of attack as the foil passes its static stall angle. This, coupled with strong hysteresis and variability in aerodynamic loading once a DSV has been shed leads to significant cycle-to-cycle variation in the time history of aerodynamic forcing for a foil undergoing dynamic stall (Corke & Thomas, 2015).

Of course, no discussion of the unsteady loading of aerodynamic bodies is complete without a brief mention of added mass. When a body is accelerated through a fluid, added mass represents the additional force required to achieve such acceleration due to the necessity to accelerate some quantity of fluid in the region surrounding the body. A particularly elegant development of the topic was recently put forward by Corkery et al. (2019), where the authors showed an experimental correspondence between the added-mass derived vorticity generated near the surface of an accelerating body and that arising from the motion of that body in the framework of potential flow theory, even in cases where significant unsteadiness led to flow separation in the real system. Similar to the discussion by Su and Breuer (2019), in the present study the primary role of added mass is to alter the apparent mass of the airfoil for the determination of structural properties, such as the natural frequency of the system ω_n, as well as to contribute to experienced aerodynamic forcing at the frequency of the airfoil's acceleration. However, as noted by those authors, the effect of added mass can be reduced by considering a relatively massive airfoil, as is the case in the present study.

1.2.4 Flapping Foil Performance and the Formation of Wakes

The formation of wakes is intimately and intrinsically linked to aerodynamic forces experienced by a flapping foil. Many authors have discussed the correspondence between the production of such forces and resulting wake structures: several examples relevant to the configuration in this thesis are provided by Cros et al. (2018), Andersen et al. (2017), Schnipper et al. (2009), and Anderson et al. (1998). In general, based on a momentum conservation argument we expect that when a fluid creates drag on an object, the wake created by that object will exhibit a mean momentum deficit (that is, the mean value for the streamwise velocity profile U will be slower than the nominal flow speed U_∞), whereas for a body producing thrust the mean wake is energized, exhibiting jet-like behaviour. Kinsey and Dumas (2008) provide a particularly clear formulation of this idea for the case of a flapping foil which may exhibit either thrust-producing or drag-producing behaviour.

Varying the frequency and amplitude of oscillation for a flapping foil can lead to a spectacular array of vortical wakes, with some particularly beautiful experiments for a pitching-only configuration performed by Schnipper et al. (2009) revealing wake configurations with 8 vortices shed per cycle or more. Classification of such vortex wakes is carried out by identifying both single vortices (denoted S) and counter-rotating pairs of vortices (denoted P) shed per cycle, and ascribing a number to each. For example, the classic von Kármán Vortex Street, typically formed in the wake of a body experiencing net drag, corresponds to a 2S wake indicating that two single vortices are shed per cycle. Cros et al. (2018) provides very clear flow visualizations of several common wake types arising from flapping foils, including this 2S-type wake.

For oscillating foils in a uniform free-stream, the formation of different wake patterns arises from the interaction between the unsteady processes of vortex shedding from the leading and trailing edges (and interactions between such vortices), and the frequency of forcing provided by the unsteady motion. As discussed briefly in Section 1.2.2, experiments performed by Anderson et al. (1998) showed that for a foil operating at optimal conditions for propulsive efficiency, these interactions gave rise to a Reverse von Kármán type wake, classified as type 2S*. The star indicates that the vortices in this wake rotate in opposite directions to those observed for classical von Kármán shedding. This inversion of the sense of rotation gives rise to a mean U profile with jet-like structure, consistent with the observed thrust production of the foil. Optimal conditions for thrust production were previously suggested by Triantafyllou et al. (1993) based on the excitation of a wake instability which leads to the formation of jet-like wake structures for a range of St limited to 0.25 to 0.35. A later study by Zhu (2011) described a 'foil-wake resonance' mechanism similar to that of Triantafyllou et al. (1993) by which pitch-heave oscillations at a high pitching amplitude and $St \approx 0.15$ excite a wake instability which gives rise to favourable conditions for power extraction. These studies provide evidence that the formation and shedding of wake structures near a foil's surface is tightly coupled both to the production of thrust and/or drag, as well as the resulting wake far downstream.

Of particular interest in the present study is the formation of 2P wakes by an oscillating foil. 2P wakes consist of two pairs of counter-rotating vortices shed per cycle. Schnipper et al. (2009) provides a clear visualization of one formation mechanism giving rise to such a wake for pitch only motion in a uniform free-stream. In addition, Andersen et al. (2017) provides a 'Wake Map' for heave-only motion, which indicates several regions in amplitude-frequency space which give rise to 2P-type wakes. Interestingly, these wakes can correspond to cases where both mean thrust and mean drag is produced. For 2P wakes giving rise to mean thrust, jet-like features are observed above and below the mean location of the airfoil, in contrast to 2S* wakes where a jet is located near the centerline (Andersen et al., 2017).

1.3 Airfoils in a Wavy Stream

The previous section focused on the case of an unsteady airfoil which oscillates in a uniform oncoming flow. This provides context for the experiments conducted in this thesis; however the conditions considered in the following chapters are significantly more complicated. Instead of unsteadiness induced solely by the motion of the foil itself through either active or passive flapping, we additionally consider an oncoming flow generated by vortex shedding from a circular cylinder, which has been well characterized (for example, see Williamson, 1996). Interactions between the flapping foil and this vortical, periodic inflow give rise to a host of dynamics which could not be realized in a uniform oncoming flow, including the dual thrust-production and power extraction phenomenon described by Beal et al. (2006) of interest in this thesis.

This section provides a brief summary of relevant work relating to both modelling and observation of airfoils encountering unsteady oncoming flow conditions. Such conditions take two different but related forms in the following discussion: we consider either an isolated 'gust,' or a periodic, vortical wake. Emphasis is placed on studies that discuss conditions similar to those in this thesis. In particular, the experiments of Gopalkrishnan et al. (1994), and related simulations by Streitlien et al. (1996) provide a basis from which we interpret the present results. Though the response of an airfoil to an isolated gust (as opposed to a periodic inflow) is not considered explicitly in later chapters of this thesis, modelling and experimental efforts on this topic have informed the interpretation of our results, and are thus included in the discussion below.

1.3.1 Modelling Interactions with Oncoming Unsteadiness

The development of analytical models to describe the interactions between an airfoil and an oncoming unsteady flow has a long history, which includes contributions from several prominent aerodynamicists of the early 20th century (von Kármán and Sears, 1938; Küssner, 1936). The collection of such models has been recently reviewed by Hufstedler (2017), and Jones et al. (2022). Despite their origins in linear thin-airfoil theory, these models remain useful in predicting loads induced by an oncoming unsteady flow when the degree of unsteadiness is moderate (Andreu-Angulo et al., 2020; Jones et al., 2022).

Unfortunately, these linear theories are not able to accurately predict the behaviour resulting from unsteady interactions which deviate too far from idealized conditions. In particular, they rely on the assumption of inviscid, incompressible oncoming flow at a small angle of attack. Flow is assumed to remain attached at the leading edge, and any shed wake structures must remain confined to a horizontal plane containing the (stationary) airfoil and consist only of vortices shed from the trailing edge (Hufstedler and McKeon, 2019; Andreu-Angulo et al., 2020). Andreu-Angulo et al. (2020) showed that for gust encounters that produced significant leading edge separation and wake deflection in experiments, the Küssner Function prediction of C_L diverged from the measured one as the strength of the gust became large. In addition, these functions do not explicitly account for vorticity in the oncoming free-stream nor do they account for the motion of the airfoil relative to the oncoming disturbances, both of which are relevant to the current study.

To address the latter concern, a similar linear, analytic model was developed by Wu and Chwang (1975) that explicitly considers the phase of interaction between an oscillating foil and a sinusoidal oncoming free stream. This model provides an analytical basis for the observation of simultaneous power and thrust production for a flapping foil based on its interactions with an oncoming unsteady flow. However, as this model is still rooted in linear theory, it does not include the effects of free stream vorticity, nor does it necessarily provide an accurate prediction for the performance of a system subject to strong viscous effects.

A final modelling effort of interest to the current study is that of Streitlien et al. (1996), which presents potential flow-based simulations of a flapping foil encountering an idealized oncoming von Kármán vortex street. In contrast to the models mentioned above, these simulations explicitly consider vorticity in the oncoming unsteady flow. The propulsive efficiency of the simulated airfoil compares qualitatively favourably to experimental results from a similar system investigated by Gopalkrishnan et al. (1994), which is discussed in Section 1.3.2. A principal result of interest from this study is the emergence of two modes of interaction between the oncoming vortex street and the motion of the foil, separated in phase by 180°. As discussed briefly in Section 1.2.2, the first is a close interaction mode, where the motion of the foil takes it to within close proximity of each oncoming vortex. The second is an avoidance mode in which the foil's motion causes it to weave between vortices, reaching its position extrema when each oncoming vortex passes its streamwise location. These modes of interaction give rise to very different power and thrust production behaviours, which inform the analysis of the present experiments in the following chapters. Although there is some difficulty in making a direct comparison, as discussed in Section 1.2.1, these modes appear to share salient characteristics with the Vortex Interception and Slaloming modes identified by Beal et al. (2006), discussed previously.

Modelling efforts regarding the effects of unsteady oncoming flows on aerodynamic performance have continued into the 21$^{\text{st}}$ century, with a particular focus on the application to gust response and mitigation for vehicles operating within a fluid. A recent review of that topic by Jones et al. (2022) presents many modern extensions to the classical analytical methods discussed above, but indicates that additional work is necessary to accurately model the response to oncoming unsteadiness which either includes or induces effects which are not accounted for in linear theory.

1.3.2 Airfoil Response to Realistic Unsteady Oncoming Flows

In light of the limited applicability of current models to the characterization of the behaviour of a flapping foil encountering oncoming unsteadiness under more realistic inflow conditions, experiments and simulations elucidating the behaviour of such systems are essential. In contrast to the modelling efforts discussed in the previous section, which tend to present an overly simplified picture of oncoming unsteadiness, a challenge in experimental studies of similar systems remains the creation of repeatable, stable, and easily characterized inflow conditions, particularly in the case of impulsive and vortical gusts (Jones et al., 2022). To address this, in the present study unsteadiness is introduced through relatively repeatable and well-characterized vortex shedding from an upstream circular cylinder. A brief review of several similar studies which provide context for the present work follows.

Lefebvre and Jones (2019) recently performed force and flow velocity measurements for a stationary airfoil placed downstream of a circular cylinder. For an airfoil at 3 diameters downstream and with the same chord as the cylinder diameter, they observed the formation of a Leading Edge Vortex (LEV) for static, geometric angles of attack ranging from 0-20$^{\text{o}}$. Consequences of LEV formation and shedding are discussed in Section 1.2.3. They also demonstrated mean thrust production, even for a stationary airfoil, for most of the cases considered in the study when the geometric angle of attack remained moderate. Finally, they described a topological change to the flow when the airfoil was moved to within $2D$ of the cylinder, where the airfoil was engulfed in an extended suction region and periodic vortex shedding was disrupted. Building from that study, Jarman et al. (2019) performed simulations of a similar configuration, but where the airfoil was translated upstream towards the circular cylinder. They found that in contrast to Lefebvre and Jones (2019), the onset of the topological changes described above could be altered based on the approach speed of the airfoil.

A foundational study combining both driven oscillatory motion in the transverse direction and interactions with oncoming vorticity is that by Gopalkrishnan et al. (1994). They found that a foil oscillating in both heave and pitch with the same frequency as vortex shedding was able to alter the structure of an oncoming cylinder wake (a 2S drag-type wake) to produce a reverse von Kármán wake (a 2S*-type wake associated with thrust production, as discussed in Section 1.2.4), depending on the phase between the airfoil motion and oncoming vortices. They also observed the formation of an 'Expanding Wake' of type 2P when vorticity derived from the airfoil's surface interacted with oncoming cylinder vorticity to form counter-rotating pairs. They then showed that the propulsive efficiency of the foil exhibited strong peaks as a function of phase, indicating that interactions with oncoming vorticity, and the resulting wake structure due to vortex repositioning, have the potential to either enhance or reduce propulsive efficiency. This is reminiscent of the observation from (for example) Kinsey and Dumas (2008) that the switch from propulsion to energy extraction for an oscillating foil manifests as a switch from a thrust wake to a drag wake; however a key distinction between these studies is that for an airfoil interacting with oncoming vorticity rather than a uniform free-stream, mean power extraction and thrust production may be occurring simultaneously.

1.3.3 Vortex Interactions for Passive Upstream Swimming

To provide context for experiments conducted in this thesis which demonstrate the ability for for a flapping foil in an unsteady oncoming flow to passively move upstream, a brief review of past work on this topic is warranted. As discussed previously, the pioneering experimental study in this area was conducted by Beal et al. (2006), though the theoretical basis for such passive thrust production was established previously (for example by Wu and Chwang, 1975). Beal et al. (2006) found that a dead fish was able to produce enough thrust to overcome its own drag and 'swim' upstream through interactions with a von Kármán Vortex Street shed by an upstream cylinder. They then demonstrated that a compliantly mounted airfoil placed in a similar wake produced net thrust larger than its net drag, while simultaneously passively extracting energy from the oncoming flow. In a prior study of live fish swimming in a similar cylinder wake, Liao et al. (2003) found that in addition to 'drafting' (taking advantage of the region of slower flow in the wake of the cylinder), the fish also seemed to synchronize their motion to the wake. They provided a hypothesized mechanism by which such a specific body posture, which they termed the Kármán Gait, could lead to passive thrust production in the cylinder wake.

Several studies have also investigated passive upstream swimming computationally. For example, Eldredge and Pisani (2008) performed simulations at Re = 100 of an articulated 'fish-like' body consisting of three rigid plates joined by hinges placed in the wake of a circular cylinder. They observed passive thrust production for both the articulated fish model and a similar rigid one where the hinges were locked; however they attribute this largely to the envelopment of the fish model in the suction zone produced behind the cylinder, which was shown to be significantly extended by the presence of the fish. In the present study, as well as the recent experimental study by Lefebvre and Jones (2019) discussed in Section 1.3.2, such a dramatic extension of the suction region due to the presence of a downstream body is not observed. In the study by Lefebvre and Jones (2019), the downstream airfoil remains outside of the suction zone at downstream distances from the cylinder larger than $2D$. This discrepancy is likely due to the strong dependence of the size and character of the suction zone on Reynolds number (which is more than 100 times larger in the present study as well as those by Lefebvre and Jones (2019) and Beal et al. (2006)), reported for example by Bloor (1964).

Potential for such interactions with the upstream circular cylinder highlight a critical difference between one-dimensional swimming, where the foil produces net thrust but is held stationary some distance downstream from the cylinder, and two-dimensional swimming, where such thrust induces upstream motion of the foil. This upstream motion changes both the frequency with which the foil encounters oncoming vorticity, as well as the strength and coherence of oncoming vortices; the latter effect was explored by Gopalkrishnan et al. (1994), who investigated 1D flapping at a range of distances downstream from a circular cylinder. In addition, interactions between the foil and the formation and shedding of cylinder vorticity can fundamentally change fluid-structure interactions taking place in the flow (Jarman et al., 2019, Lefebvre and Jones, 2019).

Considering instead an idealized vortex wake, Oskouei and Kanso (2013) conducted simulations of cylindrical and elliptic passive swimmers in inviscid wakes composed of point vortices, and found that although periodic trajectories leading to passive upstream translation exist, they are unstable to perturbations in their initial conditions. The authors propose that the effects of viscosity, body geometry, and elasticity, which were unaccounted for in those simulations, may help to stabilize such motions.

More recent work by Hang et al. (2022) has recently explored the effect of passive vs active flexion on the posterior portion of a fish body, and its contribution to drag/thrust production, swimming speed, and efficiency. They found that propulsive efficiency could be enhanced by allowing the posterior portion of the fish body to interact passively with the vortical structures shed by the upstream portion of the body. This provides further evidence that passive interactions between a body and upstream shed vorticity can contribute to the dual goals of thrust production and energy extraction.

1.4 Cyber-Physical Fluid Mechanics for Unsteady Aerodynamics Testing

The experiments presented in this thesis leverage a Cyber-Physical Fluid Mechanics (CPFD) facility in the NOAH Water Channel at the California Institute of Technology, described in more detail in Chapter 2 as well as by Shamai et al. (2020). This Captive Trajectory System (CTS) allows for the realization of both idealized and realistic passive airfoil motion, and facilitates the optimization of a virtual, passive mounting system with respect to power extraction performance.

CPFD can be briefly summarized as the use of real-time, active feedback control to create 'virtual dynamics' for a test object during an otherwise conventional set of fluid mechanics experiments. That is, a test article is mounted to some set of motors and actuators, which based on time-resolved measurements made from the surrounding flow or the object itself are able to realize a desired trajectory for the object through real-time feedback control (Mackowski & Williamson, 2011). This approach allows the researcher access to a wider array of fluid-structure interaction possibilities than those available through conventional (electromechanical) means; however specialized facilities are required to realize such experiments.

Pioneering CPFD facilities created to study both Vortex-Induced Vibration (VIV) and vortex interactions relevant to fish swimming are described by Hover and Triantafyllou (2000). These facilities represented significant innovation in experimental research, as they allowed researchers to easily change the behaviour of an object subject to fluid forcing in software, rather than having to physically re-build the experimental setup to swap out springs, dampers or other mechanical components.

In addition, the software implementation of a mounting system is guaranteed to be free from nonidealities that are common to physical experimental setups, though computer control introduces the new potential for phase delays and other control-related deviations from the behaviour of a theoretically equivalent physical system (Hover & Triantafyllou, 2000).

In the early 2000's, CPFD had not yet been widely adopted. The only study mentioned by Williamson and Govardhan (2004) in their review of progress in VIV research to incorporate CPFD was the original study by Hover and Triantafyllou. Since then however, adoption of the technology has increased and similar CPFD setups have proliferated. For example, Beal et al. (2006) used CPFD to carry out the experiments regarding interactions between a cylinder wake and a downstream airfoil that inspired the work presented in this thesis. Mackowski and Williamson (2011) present the details of a more recent CPFD setup that expanded the original concept from one-dimensional motion to multiple dimensions, similar to the setup used in the present experiments. Their discussion of sensor noise mitigation, computation of desired trajectories and phase errors in achieved dynamics have informed the CPFD implementation used in this thesis. Su and Breuer (2019) demonstrated the use of their CPFD facility to study a semi-passive flapping foil energy harvesting device. Lambert et al. (2016) present a completely different CPFD architecture, integrating vision-based motion tracking of a test object, as well as active flow control elements in addition to actuation through 6 degrees of freedom. Development of a CPFD facility by Jones et al. (2021) which supports multiple test objects for studies of tandem oscillating foils is also underway.

1.5 Current Approach and Summary of Thesis Contents

Experiments investigating, characterizing and optimizing the simultaneous potential for power and thrust production by an airfoil placed downstream of a circular cylinder form the bulk of the contributions in this thesis. Chapter 2 provides a detailed description of the experimental setup and data analysis upon which the following chapters are based.

Within the general experimental architecture, three distinct types of experiments were performed:

1) A mechanical apparatus was constructed to allow the airfoil to move in the transverse direction (perpendicular to the oncoming free stream) only, purely due to impinging vorticity shed from the upstream cylinder. This apparatus is referred

to as the Mechanical Free-Response System (MFRS) throughout the thesis.

2) The Captive Trajectory System (CTS) was used to drive the airfoil through pre-planned sinusoidal trajectories in the transverse direction with a specific frequency and phase with respect to oncoming vorticity. The frequency and phase of the motion were determined based on preliminary experiments performed with the MFRS. Experiments where the airfoil was driven by the CTS through a pre-planned trajectory are referred to as 'Driven Airfoil Experiments' throughout the thesis.

3) The CTS was programmed to perform real-time feedback control of the airfoil, allowing it to realize specified dynamics based on measured forces that it experienced. The dynamics chosen for the experiments in this thesis represent passive behaviour of an airfoil mounted to a linear spring-mass-damper system, or translating freely in space governed by $F = ma$. The use of CPFD to realize the mounting system dynamics allowed for a large degree of flexibility in tuning the passive behaviour of the airfoil: this was principally used to realize airfoil behaviour that was similar to the Driven Airfoil Experiments in the transverse direction. In addition, simultaneous motion of the airfoil in 2 dimensions, both transverse and streamwise, was made possible through the CTS mounting system. Collectively, these experiments are referred to as 'Passive Captive Airfoil Experiments' throughout the thesis.

Chapter 3 contains discussion of the Driven Airfoil Experiments, which lay the groundwork for understanding airfoil performance and related interactions with oncoming vorticity for this system, as well as providing connections to several previous studies discussed in this chapter.

Chapter 4 presents results from Passive Captive Airfoil Experiments in the transverse direction only. The mounting system is tuned to achieve motion that is similar to the driven motion in the previous chapter, and the behaviour of the Passive Captive Airfoil is compared to the representative driven case. Then, mounting system dynamics are optimized to improve power production potential.

Chapter 5 focuses on the addition of passive motion in the streamwise direction, under the action of the mean thrust generated due to interactions with oncoming vorticity.

Chapter 6 explores the behaviour of the MFRS, and the effects of friction at play in the observed airfoil motion. The CTS is used to implement a simplified model of the experienced friction, and the effect of such a realistic engineering response on the airfoil behaviour is explored.

Chapter 7 provides a summary of the contributions in this thesis, as well as concluding remarks and future directions for the work.

Chapter 2

EXPERIMENTAL SETUP AND DATA PROCESSING METHODOLOGY

2.1 Introduction

Experiments performed in the NOAH Water Channel at the California Institute of Technology form the basis for the contributions presented in this thesis. The configuration tested consists of an airfoil placed downstream of a circular cylinder, such that interactions between the von Kármán wake shed by the cylinder and its impact on the downstream airfoil can be studied. Experiments can be described as conforming to one of three distinct types, as enumerated in the previous chapter: Mechanical Free-Response (MFRS) Experiments, Driven Airfoil Experiments, or Passive Captive Airfoil Experiments.

Information regarding experiment development as well details of the resulting experimental setup(s) are presented chronologically in this chapter, to facilitate discussion of how data from earlier sets of experiments informed choices made for airfoil behaviour in later sets. In the following chapters however, results from the experiments described here are presented in an order which proceeds more naturally from foundational observations to those which are more specified to particular engineering goals. In addition, this chapter includes discussion of the data processing pipeline used to analyze the data collected from experiments.

2.2 Basic Experimental Configuration for All Experiments

Several aspects of the experimental setup were common for all tests presented in this thesis. All tests were conducted in the NOAH Water Channel, the performance of which has been described in general terms in several recent works detailing previous experimental studies (see Shamai, 2021; Huynh, 2019; Hufstedler, 2017). Figure 2.1 provides an overview of the experimental configuration used for the present experiments.

A circular cylinder with a diameter of $D = 11.5$ cm was mounted upstream from a NACA 0018 airfoil with a chord length of $c = 10.0$ cm $= 0.9D$. The foil was mounted vertically, secured above the tunnel to allow translation in the y-direction (transverse/heave), while its z-direction (vertical/spanwise) motion was constrained.

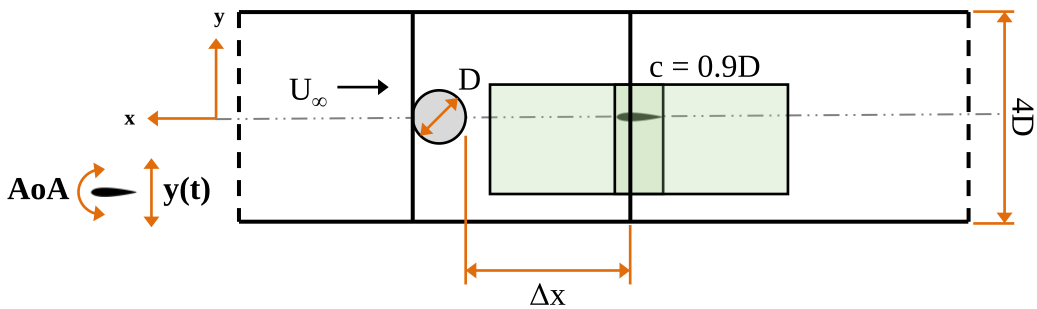

Figure 2.1: Basic schematic showing the test section of the NOAH Water Tunnel facility where all experiments were conducted. Schematic is not to scale. Test section measures 150x46 cm (13x4D), and the water depth is approximately 46 cm. Flow travels from left to right, first encountering an upstream circular cylinder, which sheds vorticity to encounter the airfoil located some distance Δx downstream. Green regions indicate the Field of View (FOV) for PIV – the position and extent of these regions varied slightly between experiments.

For the MFRS, Driven Airfoil Experiments, 1D Passive Captive Airfoil Experiments, and experiments with realistic friction, the streamwise or x-direction motion of the airfoil was also constrained; however, this was relaxed for 2-dimensional Passive Captive Airfoil Experiments. The submerged span of the airfoil is denoted as s = 0.45 m, which was measured when the tunnel was off: this represents a slight overestimation of the true submerged span, as when the tunnel is operating the water level drops slightly due to pump suction. Due to airfoil motion and the desire to use measured forces for real-time force feedback end plates were not used, though for all tests the submerged edge of the airfoil was within approximately 1 cm of the tunnel bottom.

For all experiments in this thesis, unless specifically mentioned otherwise, the speed of the oncoming free stream flow (without the presence of the cylinder) was approximately U_∞ = 0.32 m/s verified by capturing Particle Image Velocimetry (PIV) data with the tunnel operating at the same speed as testing, but with neither the cylinder nor airfoil installed. Due to slight variations in tunnel fill level and testing duration for each set of experiments, the speed of the oncoming free stream varies slightly between sets of experiments; however the Reynolds Number based on cylinder diameter for all tests remains close to 40,000. In addition, the presence of the cylinder in the tunnel creates a large change in the mean downstream conditions due to a high blockage ratio. This is a desired outcome for the experiments within this study, as the large shed vortices which create regions of energized flow as well as a large region of mean velocity deficit all form the backdrop for the behaviours of interest; in addition, frequency content from vortex shedding is directly determined for each trial. Thus, blockage ratio effects are not explicitly considered separately from other analyses of flow behaviour.

Several sources of potential asymmetry are apparent in this experimental setup. Firstly, though efforts were made to ensure that the circular cylinder was centered in the NOAH water channel for all experiments, due to the high blockage ratio (and therefore high gap-to-diameter ratios on either side of the cylinder), a relatively small error in centering could lead to asymmetry in the flow over the cylinder (Bearman & Zdravkovich, 1978). Similarly, an offset in the neutral position of the foil relative to the cylinder centerline could also lead to asymmetry in the aerodynamic forces experienced by the airfoil. The estimated transverse offset between the cylinder and airfoil centerlines is approximately 5-10% of D. In addition, the shape of the airfoil was affected by a slight warp along the span, in addition to small uncertainty in the mounting point through which aerodynamic forces are recorded. Since linear forces rather than moments are principally of interest in this study, the latter concern has only a very minor effect if any; however the existence of these asymmetries in the setup could contribute to behavioural asymmetries, discussed in later chapters.

2.2.1 Force Sensor Characterization

For all experiments discussed in this thesis, an ATI Mini40 IP68 6-axis force and torque sensor was used to measure forces in the direction parallel to the airfoil's chord, as well as the direction perpendicular to it. For experiments with a geometric angle of attack of zero ($\alpha_0 = 0$), these axes correspond to the directions of Thrust and Lift acting on the airfoil, as well as x and y-direction forces F_x and F_y, used interchangeably throughout this thesis. For cases where α_0 was non-zero, force sensor data collected in the resulting rotated reference frame was 'de-rotated' prior to any further processing, such that the reported values of Lift, C_L, or F_y always correspond to the force in the transverse tunnel direction (perpendicular to the free-stream), and Thrust, C_T, or F_x always correspond to force in the upstream direction, against the oncoming flow.

In addition, as the force sensor is mounted in an accelerating reference frame, an additional pre-processing step was carried out to remove the fictitious force(s) recorded by the force sensor due to acceleration according to

$$F_{\text{aero}} = F_{\text{meas}} + m_{\text{eff}} a, \tag{2.1}$$

where m_{eff} was determined through measuring the mass of the NACA 0018 airfoil

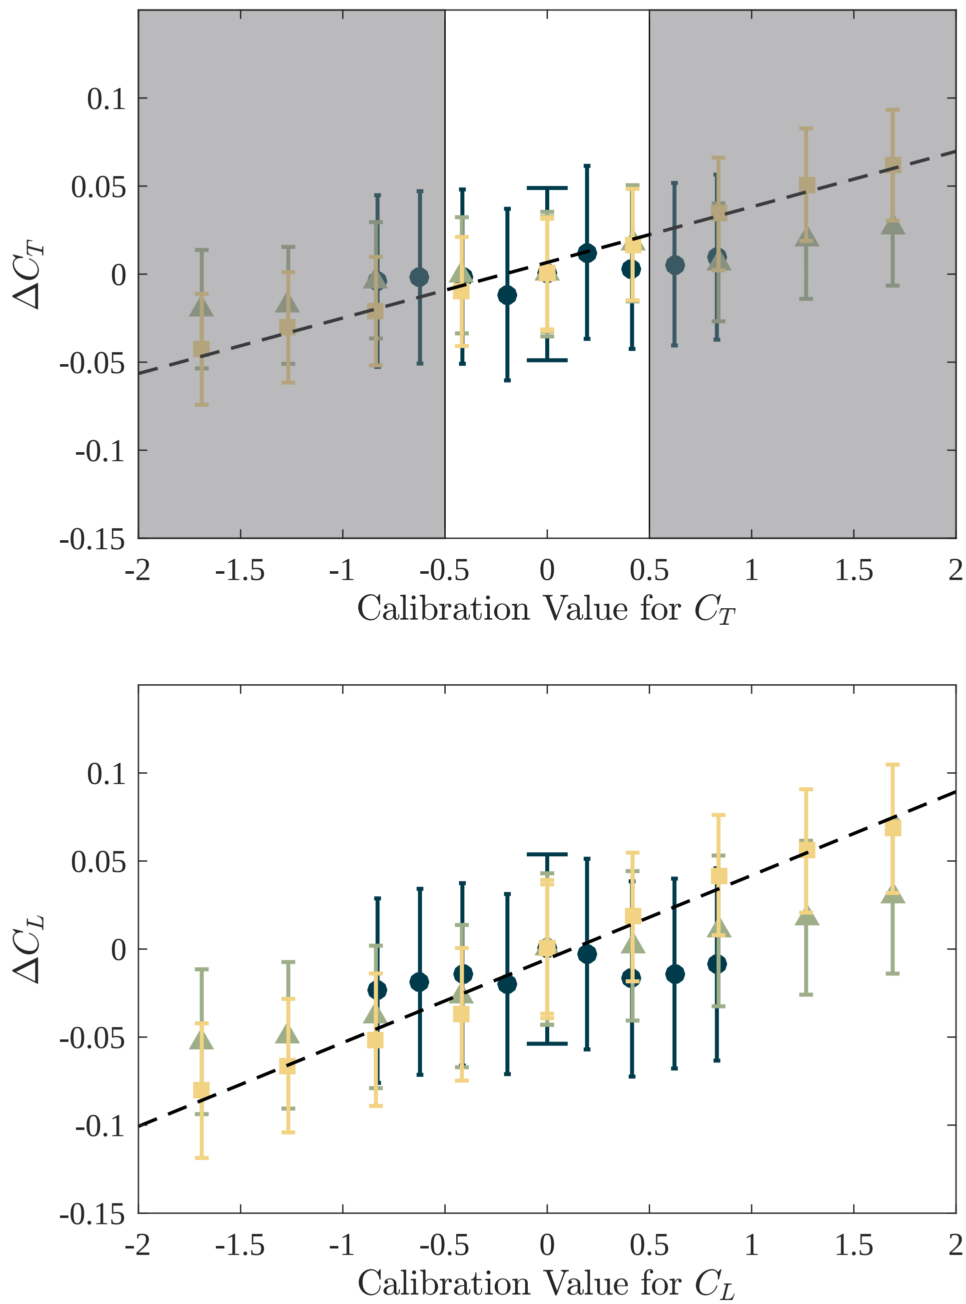

Figure 2.2: Deviation of measured force values from known calibration weights as a function of measured force in the x (top) and y (bottom) directions. Performance data were collected on three different days corresponding to the colour/symbol combinations in the plot: ● Data set 1; △ Data set 2; and ■ Data set 3. Black dashed lines indicate worst-case linear fit to $\Delta F / p_\infty sc$ as a function of the known calibration force for any set of data collected. In the top panel, the greyed out regions indicate thrust values outside the expected measurement range in experiments.

used, as well as extraction of the appropriate self-weight for the sensor and mounting system by oscillating the system with no airfoil attached and with a known frequency and amplitude, while measuring the resulting (fictitious) forces.

To characterize the performance of the force sensor in the axes of interest, calibration/performance data was collected using an external bench top setup. To collect calibration data, the force sensor was uninstalled from the experiment(s) and clamped through its CTS attachment hardware to the bench top, such that gravity was aligned with one of the positive or negative x or y axes of the sensor. Then, known masses were placed such that the force measured by the sensor should be equal to $F_{\text{meas}} = gm_{\text{cal}}$ in the axis of interest for a particular test.

Figure 2.2 shows the results of three such bench top tests. The difference between the known calibration weight and the mean reading of the sensor with the weight placed on it for 1 minute in each axis is shown, such that

$$\Delta C_X = \frac{1}{p_\infty sc}\left[\overline{F_{\text{meas},\,X}} - gm_{\text{cal}}\right]. \tag{2.2}$$

Error bars indicate the standard deviation of the sensor measurement over the 1 minute period. Forces and standard deviations are normalized by $p_\infty sc$, such that x-direction force measurements correspond to C_T, and y-direction force measurements to C_L for the sensor as installed in the experimental setup. Each colour/symbol combination indicates a different set of performance tests, taken throughout the experimental campaigns conducted for this thesis. The earliest data is shown in dark blue while the most recent, collected approximately one year later, is shown in yellow. In addition, although both x and y-direction forces were characterized over the same range, the maximum measured value of C_T in any experiment is much smaller than that of C_L: regions of the calibration that are outside of the bounds of measured F_x values are greyed out in the Figure. We note that the resolution of the sensor in the axes of interest is 0.01 N, provided by the manufacturer.

From Figure 2.2, we can extract several informative metrics to characterize the performance of the sensor. Firstly, the standard deviation of all measurements appears roughly constant as a function of measured force, though it appears slightly larger in the earliest trial than in the later ones. The mean value of the standard deviation over all measurements in the Figure provides a precision estimate for measured forces of $\delta F_x = \pm 0.09$ N, or $\delta C_T = \pm 0.04$ for streamwise forces, and $\delta F_y = \pm 0.1$ N, or $\delta C_L = \pm 0.04$ for transverse ones. In addition, we can consider the linearity of our sensor over the range of interest by fitting a line to the mean measured force value as a function of calibration weight for each day. Fitting was accomplished using a least-squares method, and the worst-case (largest slope) line of the three calibration trials is shown for both ΔC_T and ΔC_L in the Figure. Based on these fits, we see that the force measured is very linear, and for F_x the sensor appears accurate to within approximately 3% of the measured value, while for F_y the force measured is accurate to within approximately 5% of the measured value. This corresponds to an accuracy error at the largest measured values approximately on the order of the standard deviation.

It is interesting to note that the accuracy of the sensor appears to be becoming worse with time: for future measurement campaigns, re-calibration by the manufacturer is recommended.

2.3 Experimental Setup for the Mechanical Free-Response System

For the Mechanical Free-Response System (MFRS) experiments, performed chronologically first during the author's residence at the Institute[1], the airfoil was secured to a linear motion cart and allowed to move transversely in a purely passive way in response to oncoming free stream vorticity. The airfoil was mounted to the cart through a rotary bearing on a shaft passing through its quarter-chord location; however, very little change in angle of attack was observed during any individual experimental run, so the airfoil is assumed to be stationary in pitch throughout the analysis in this thesis (although for each experimental run, the stationary pitch angle was different). The airfoil was fixed at a distance $\Delta x = 2.7D$ downstream from the circular cylinder in the streamwise direction. An experimental schematic for the MFRS is given in Figure 2.3.

The position of the cart was measured using a Keyence LK-G502 laser distance sensor, and the signal was numerically differentiated using a Savitzky-Golay filter to obtain the cart velocity and acceleration as described by Schafer (2011). This method was validated through comparison with the measured acceleration of the cart, obtained using an ADXL 337 accelerometer. The Angle of Attack (AoA) of the airfoil was measured using a Vishay 351 Hall-Effect (HE) rotary encoder. Finally, six-axis forces and torques acting on the airfoil were measured using the ATI Mini40 IP68 force/torque sensor described in Section 2.2.1. The laser distance sensor and force/torque sensor included appropriate signal conditioning through dedicated analog-to-digital conversion provided by the manufacturer. The acceleration and

[1]The author gratefully acknowledges contributions to the original concept and prototype of the MFRS Experiments by Benedikt Barthel, as part of the course Ae104c in Spring, 2017.

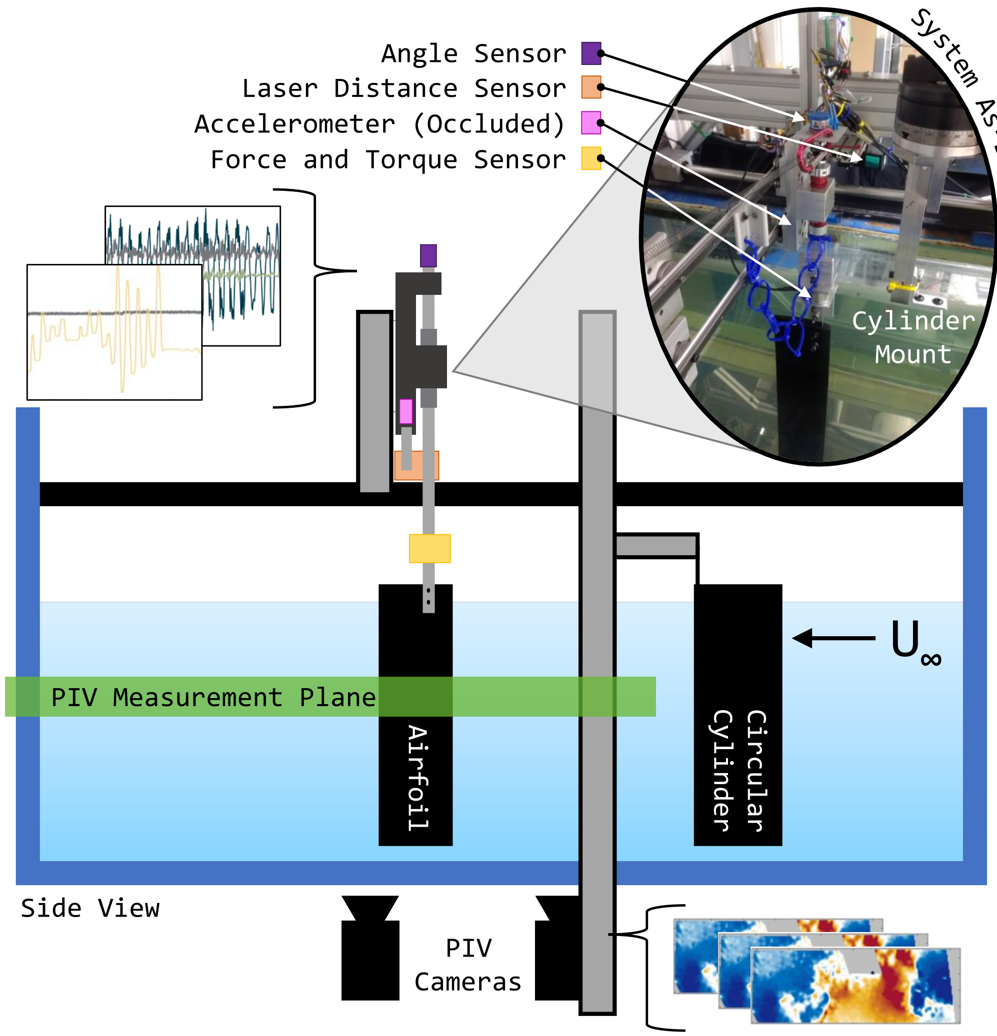

Figure 2.3: Schematic showing the mounting system and sensor architecture used for MFRS Experiments. Not to scale.

angle of attack signals were low-pass filtered prior to data acquisition to ensure compliance with the Nyquist Criterion, and eliminate high-frequency noise. Data were sampled at 25 kHz simultaneously for all sensors.

2.4 The Captive Trajectory System and Captive Airfoil Experiments

For the remainder of the experiments performed in support of this thesis, the Captive Trajectory (CTS) was used to provide the mounting system for the airfoil. A photograph of the airfoil mounted to the CTS downstream of the circular cylinder is provided in Figure 2.4.

The CTS is able to precisely drive a test object through a prescribed trajectory in the three translational axes, as well as pitch. It can also be integrated with

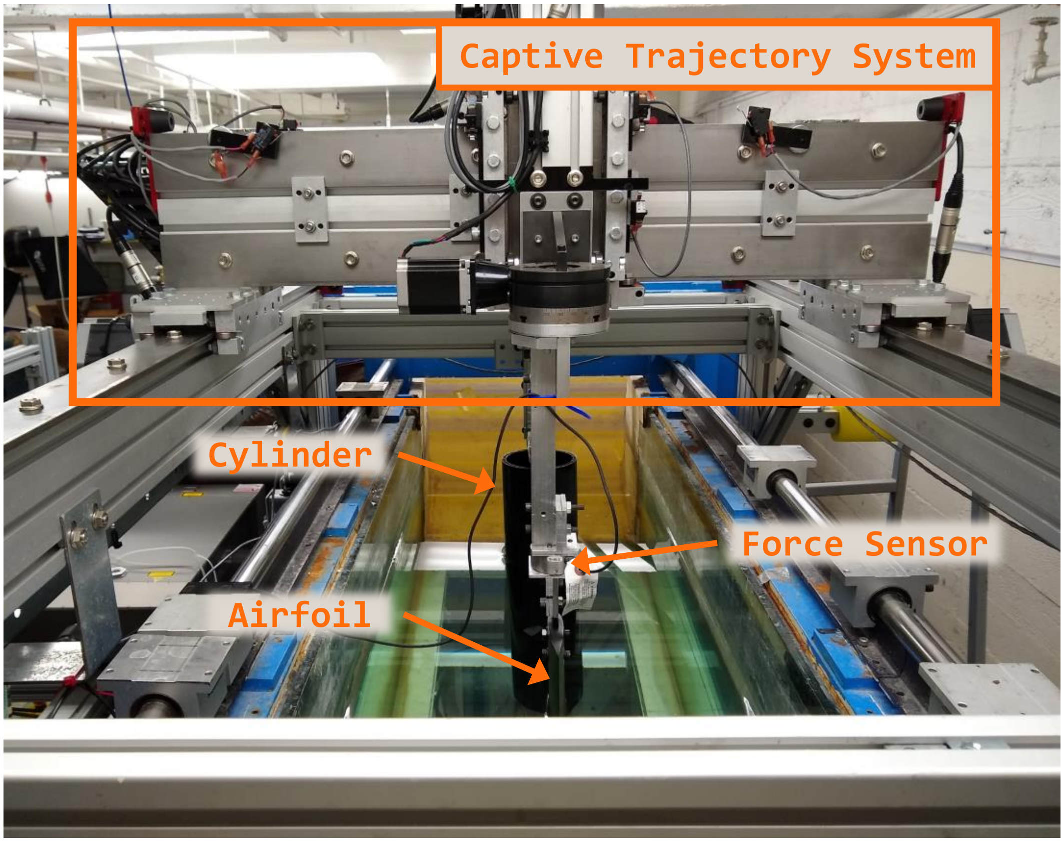

Figure 2.4: Photograph of the airfoil (foreground) mounted to the Captive Trajectory System (CTS) downstream of a circular cylinder (background). Motion of the CTS is actuated by stepping motors, which individually allow for translation of the system in three translational axes and pitch. The airfoil is mounted to the CTS through its quarter-chord location, and a force sensor is used to measure the forces acting on it.

additional hardware to provide feedback for real-time control of a test object's motion, referred to as a *captive trajectory*; more details regarding CTS performance were reported by Shamai (2021). Both prescribed and captive trajectories were created for experiments in support of this thesis, as described in the following sections. Forces and torques on the airfoil were measured using the same force sensor as in the MFRS Experiments, the performance of which is discussed in Section 2.2.1; however airfoil position, velocity, and force data were available only as digital outputs from the CTS at a rate limited to 200 Hz, a significant reduction from the MFRS Experiments. This reduction in sample rate, a constraint imposed by the CTS itself, limits the frequency range and resolution possible for analysis of measured signals; however it is still several orders of magnitude larger than the expected dynamics in our flow.

For all experiments conducted using the CTS, the airfoil's angle of attack was fixed such that it remained stationary throughout the duration of each run, regardless of the pitching torque acting on the foil. This behaviour is different from that in the MFRS Experiments described earlier, as well as the pioneering experiments by Beal et al. (2006) in which they allowed for changes in pitch angle due to flow forcing. As the observed pitching motions during the MFRS Experiments (which permitted free pitching, subject to the action of friction) were small, the geometric angle of attack was fixed to simplify the experimental setup in further Driven and Passive Captive Airfoil Experiments.

2.4.1 Motion through a Prescribed Trajectory

For all Driven Airfoil Experiments, discussed in detail in Chapter 3, the airfoil was mounted to the CTS and actively driven through a transverse sinusoidal trajectory with a known frequency and phase relative to oncoming vorticity, of the form

$$y(t) = A_p \sin(2\pi f t + \psi), \tag{2.3}$$

where $y(t)$ is the transverse position as a function of time. The parameters A_p, f and ψ were fixed based on data taken during previous MFRS experiments, which represent a truly passive mounting system (with no motors or other actuators to drive airfoil behaviour). Though the motion of the MFRS is unideal, in that it is not itself purely sinusoidal, choosing the parameters which determine the driven airfoil trajectories studied in this thesis based on characteristics of the MFRS allows us to directly connect these driven motions with the behaviour of a truly passive system.

To determine the appropriate driving frequency f, the transverse force and position of the airfoil from previous MFRS experiments were analyzed. For vortex shedding at a Reynolds Number of 40,000, the value of the Strouhal number,

$$St = \frac{fD}{U_\infty}, \tag{2.4}$$

has been shown to remain fixed around a value of 0.2 (for example Lefebvre and Jones, 2019, and references therein). This gives an expected dimensional shedding frequency of approximately 0.6 Hz. To confirm that this frequency was dominant in the airfoil dynamics, measured transverse force and position for all MFRS Experiments were windowed into temporal bins varying in length from approximately 5.5 s to 20.0 s (3.4T to 12.5T), and linear de-trending was applied to the position signal. Then, the frequency content of each bin was computed using an FFT. In all cases, the windowed signals (with bins of any length) exhibited a strong primary peak in the region of 0.6 Hz; however the position signal sometimes exhibited additional strong peaks at lower frequencies, perhaps due to long-term meander or dynamics induced by the mounting system. To remove this as a source of error in determining dominant oscillatory behaviour due to vortex shedding, only the transverse force signals were used to compute the oscillation frequency for the driven experiments. The peak locations in all transverse force signal bins were averaged resulting in the selection of 0.6226 Hz, presented with fewer significant digits throughout the remainder of this thesis as 0.62 Hz, which gives a Strouhal number of 0.22.

To determine the appropriate phase offset between the measured force and the airfoil motion (a proxy for the phase between the vortex shedding and the airfoil motion), the difference in phase angle computed from the FFT at the dominant forcing frequency in each bin between the force and motion was averaged. Although this phase difference varied widely especially for shorter window lengths, the mean value was determined to be quite close to $\pi/2$. This value confirms a visual trend identified in the passive airfoil data, as well as enforcing the condition that the airfoil's velocity is in-phase with the forcing. This is intuitively satisfying as it enforces that the rate of work done on the airfoil by the surrounding flow,

$$P(t) = F_y(t)\dot{y}(t), \tag{2.5}$$

is always positive. Since there is no power input available for the airfoil to do mean work on the flow, this is a physically meaningful constraint; thus, for the Driven Airfoil Experiments presented in this thesis, the phase ψ for the motion is set such that velocity and transverse force are (approximately) in-phase. This choice however represents a significant simplification (and idealization, in terms of power extraction potential) relative to the true behaviour of the passive foil, for which friction-induced dynamics obscure the phase shift corresponding to an ideal free-airfoil system.

Finally, the appropriate amplitude for the driven sinusoidal motion was determined through analogy with the time-mean kinetic energy $\bar{E}$ of the airfoil in the MFRS Experiments. Considering a pure sinusoid,

$$\bar{E} = \frac{\bar{E}_k}{\frac{1}{2}\rho_\infty} = \lim_{T\to\infty} \frac{1}{T}\int_0^T [A_v \sin(\omega t)]^2 dt = \frac{A_v^2}{2}. \tag{2.6}$$

Then, the mean kinetic energy of a signal that is not perfectly periodic can be matched to an equivalent value for a pure sinusoid. We consider

$$\bar{E} = \frac{A_v^2}{2}, \tag{2.7}$$

$$A_v = \sqrt{2\bar{E}} = \omega A_p, \tag{2.8}$$

$$A_p = \frac{\sqrt{2\bar{E}}}{\omega}, \tag{2.9}$$

where A_v is the amplitude of the velocity, and A_p is the amplitude of the position for a pure sinusoid. By calculating $\bar{E}$ as the mean kinetic energy over all MFRS trials (i.e., by summing square velocity at each recorded time step for all experimental runs then dividing by total time steps), then applying Equation 2.9, we can estimate the appropriate amplitude for our driven position signal as $A_p = 5.1$ mm, or approximately 5% of the airfoil chord. This is a relatively small amplitude of oscillation compared to previous studies (discussed further in Chapter 1), which usually consider flapping foil motions on the order of 0.5-1.0D; however, the amplitude of oscillation observed by Beal et al. (2006) for a passively flapping foil exhibiting both thrust production and power extraction was much smaller, approximately $0.2D$. In addition, such a small amplitude of oscillation is still seen to give rise to a host of dynamics of interest in this study.

In summary, for the Driven Airfoil Experiments presented in the thesis, the airfoil was driven through a prescribed sinusoidal trajectory given by the following equation:

$$y(t) = 0.0051\sin(2\pi(0.6226)t + \psi), \tag{2.10}$$

where ψ was fixed for each individual trial to ensure that the velocity was in-phase with the force acting on the airfoil due to vortex shedding, and significant digits for f and A_p are preserved in accordance with CTS precision. This trajectory is accurately realized by the CTS for all Driven Airfoil Experiments.

2.4.2 Motion computed from Real-Time Force Feedback

For the Passive Captive Airfoil Experiments discussed in Chapters 4 through 6, the CTS is used to provide real-time control of the airfoil's motion, based on programmed dynamics. CPFD apparatuses of this type allow the experimenter access to much more detailed control of the behaviour of a test object, as discussed further in Chapter 1. In the present case, moving from an all-mechanical mounting system (the MFRS, discussed in detail in Section 2.3) to a cyber-physical realization through the CTS allowed for the specification of an idealized, friction-free passive mounting system, and enabled much more precise control over the 'physical' details of such a system than would be possible using real mechanical components.

A flow chart showing the high-level operation of the CTS is given in Figure 2.5. Low-level behaviour of the CTS (trajectory tracking, motor control, data acquisition etc) corresponding to the `Measure`, `Output`, and `Actuate` steps in Figure 2.5 is controlled internally, and such behaviour was not changed from basic settings tuned in a recent study by Shamai (2021). The high-level behaviour, or the determination of a desired trajectory based on measured sensor inputs (the `Model` step in Figure 2.5), was changed for each set of Passive Captive Airfoil Experiments as necessary to produce the desired airfoil dynamics.

The desired trajectory for a test object, in this case the airfoil, is specified at each CTS time step (a rate of 200 Hz) based on an update programmed by the user. For all Passive Captive Airfoil Experiments in this thesis, the trajectory update at each time step was calculated based on an impulse method, following Mackowski and Williamson (2011). The first step of this method is to calculate the real Impulse I applied to the airfoil over the previous time step. The real impulse can be approximated as the measured force multiplied by the time step, including an additional term necessary to counteract the fictitious force in the measurement due to the sensor's real mass m_{real} and accelerating reference frame (discussed in Section 2.2.1):

$$I_n = \int_{t_n}^{t_{n+1}} F_{\text{meas}}(t) + m_{\text{real}}\ddot{y}(t)\ \mathrm{dt} \approx F_n \Delta t + m_{\text{real}} \Delta \dot{y}_{n+1}. \qquad (2.11)$$

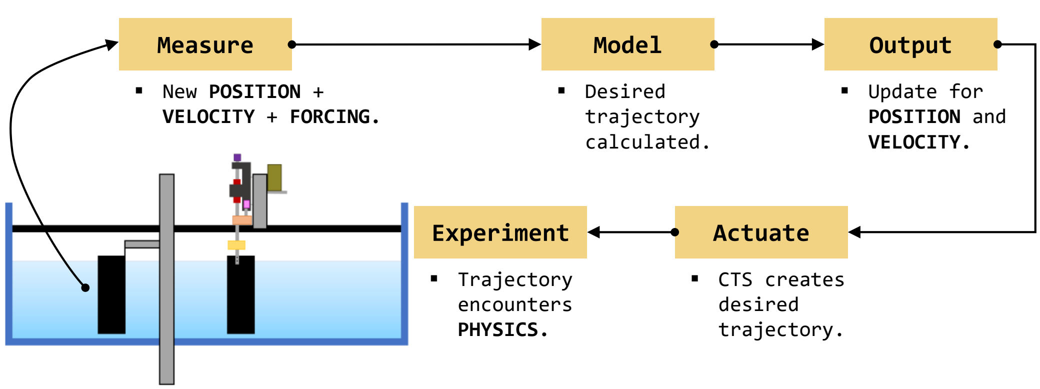

Figure 2.5: Flow chart showing the high-level behaviour of the Captive Trajectory System (CTS). The CTS allows the user to specify a model which governs a physical system's response to measured inputs, in this case, measured forces acting on an airfoil. It then actuates the physical system in accordance with the desired behaviour.

Then, impulses applied by the virtual mounting system (for example, the action of a virtual linear spring or damper, as considered in the transverse direction for all Passive Captive Airfoil Experiments in this thesis) are computed in a similar fashion,

$$I_{\text{virt},\,n} = \int_{t_n}^{t_{n+1}} F_{\text{virt}}(t)\ \text{dt} \approx F_{\text{virt},\,n}(\Theta, \Delta\dot{y}_{n+1}, y_n)\Delta t, \tag{2.12}$$

where the specific parameters Θ giving rise to the virtual forces correspond to parameters in the dynamics under test, for example $[m, b, k]$ for the spring-mass-damper system. Using the above expressions, an implicit update for the quantity of interest, $\Delta\dot{y}_{n+1}$ is made possible according to

$$m_{\text{virt}}\Delta\dot{y}_{n+1} = I_n + I_{\text{virt},\,n}. \tag{2.13}$$

Position is then updated using an Implicit Euler method. Trigger conditions appropriate to each set of experiments are also included, such that trajectory updates are initiated only after trigger conditions are met. Before implementation on the CTS, trajectory update code was tested on an emulated CTS system with no hardware in the loop and simulated forcing, and results were validated against a simulated dynamical system with the desired dynamics created in MATLAB's Simulink environment. The following code skeleton provides a general outline for the format of trajectory updates used for Passive Captive Airfoil Experiments in this thesis:

`IF` (Trigger Condition Met)

  Compute implicit update:

    Solve Equation 2.13 for $\Delta\dot{y}_{n+1}$ based on measured and virtual forces.

  Update trajectory:

    $\dot{y}_{n+1} = \dot{y}_n + \Delta\dot{y}_{n+1}$

    $y_{n+1} = y_n + \Delta t\dot{y}_{n+1}$

`ELSE`

  Wait for Trigger Condition

In general, performance of the CTS allowed for the desired trajectories computed at each time step to be realized with a high degree of accuracy. This is particularly true for transverse-only motion of the airfoil discussed in Chapter 4. For the 2D experiments discussed in Chapter 5, slight deviations from the desired trajectory in the x-direction motion only were apparent. The strict cause of these deviations is unknown, but they appeared to be more prevalent at higher forward speeds. To address the issue, a detailed review of controller gains and other motor-control level parameters set internally in the CTS is necessary. As these deviations did not seem to critically affect the observations made regarding streamwise airfoil motion, such a detailed investigation is left for future experimental campaigns.

Another source of error in the computed trajectories stems from uncertainty regarding the real mass of the airfoil (and associated mounting hardware) necessary for the computation of the real impulse, I_n. The value of m_{real} for all Passive Captive Airfoil Experiments in this thesis was set to a value of 0.8005 kg based on a preliminary investigation of fictitious forces induced in the sensor, which was later found to be in error. Correcting the error and using an independently measured mass for the foil itself, the correct mass was found to be 1.2 kg, constituting a mass error $\delta m = 0.4$ kg. Though this is a relatively large error in terms of the real mass value $m_{\text{real}} = 1.2$ kg, it is a much smaller error relative to the *virtual* mass of the system, for which the smallest value considered in this study is 6.1 kg. Propagating this error in mass to determine the effect on the force used to compute a trajectory, we find based on the expected acceleration of the system that the error should not be larger than 0.06 N, which is smaller than the uncertainty in the measured force itself (from Section 2.2.1, $\delta F \approx 0.1$ N). Thus, although every effort should be made to correct the erroneous mass in future tests, the effect of this error on the observed behaviour of the system is limited. In addition, the real mass of the system is a known source of error for CPFD systems of this type: Mackowski and Williamson (2011) acknowledge the real mass of their similar setup as a source of error in realized trajectories (despite its accurate specification), and recommend ensuring that $m_{\text{real}} \ll m_{\text{virt}}$ for optimal performance. For the majority of Passive Captive Experiments in this thesis, m_v was 18.4 kg or larger.

2.5 Particle Image Velocimetry

Particle Image Velocimetry (PIV) is used to collect quantitative information about the flow field around the airfoil in 2 dimensions. For all experiments described in this thesis, 2D2C PIV was performed simultaneously with measurements of airfoil dynamics. PIV images were captured at a rate of 800 Hz using two Phantom Miro Lab 320 cameras with a pixel resolution of 1920x1200 px. The two fields of view partially overlapped to form a single, continuous region of interest (see Figure 2.1). Both cameras were placed directly beneath the tunnel, perpendicular to the transparent tunnel bottom, and equipped with 35mm Nikkor lenses. Neutrally buoyant tracer particles were illuminated by a Photonics DM20-527(nm) YLF laser in single-pulse mode, which was expanded through a cylindrical lens to form a sheet. The laser sheet entered the tunnel through a side wall parallel to the tunnel bottom.

Prior to data collection on each day of testing, calibration images were taken of a standard LaVision Type 11 calibration target placed near the center of the combined field of view. Since a new calibration was collected for each set of experiments conducted, the underlying spatial field(s) varied slightly from day to day. This has implications for data processing, discussed further in Section 2.6.

PIV image acquisition and processing was completed using commercial software (LaVision DaVis, version 10.1.2, with the exception of the earlier MFRS Experiments, which were recorded using DaVis 8). Raw images from each independent trial were first averaged to compute a background image, which was subtracted from each frame prior to processing. In addition, the outline and shadow created by the airfoil was masked out based on its measured position in each frame. Prepared images were then processed sequentially to produce vector fields. A multi-pass correlation algorithm was used, where an initial pass using a 64x64 pixel square window was subsequently reduced to yield a final circular window size of 16x16 pixels, through three additional passes. Each pass included a 50% overlap between windows, and outliers were removed between each pass based on a minimum correlation value threshold of 0.4, as well as median filtering. Obtained vector fields were then post-processed using an additional median filter to remove outliers, as well as a 3x3 smoothing filter to reduce noise in the final fields.

The resulting size of the field of view varied slightly between sets of experiments as mentioned above, but was close to 0.65 x 0.24 m (5.7 x 2.1 D) for all experiments. The resulting pixel resolution was approximately 5 px/mm, giving rise to a final interrogation window size of approximately 3x3 mm. Thus, the resolution of the resulting vector fields, considering 50% overlap used in the calculation of velocity vectors is approximately 1.5 x 1.5 mm.

2.6 Data Processing Methods

Several methods of processing the data obtained from the experiments described previously are common to many further analyses performed, as presented in the following four chapters. This section provides details and justification for the use of these common data processing steps, which are referred to throughout the thesis.

2.6.1 Visualizing Vorticity

A common challenge in fluid mechanics is the visualization of vorticity based on measured velocity fields. In addition to the ill-defined nature of the boundaries of regions of high vorticity, and the impact of this nebulousness on the detection of 'coherent structures' within a flow, the calculation of vorticity involves numerical differentiation of flow fields which may be challenging especially when flow data arises from noisy experimental fields. Vortex identification algorithms, several of which were recently summarized with respect to their utility in an unsteady aerodynamic setting by Huang and Green (2015), aim to mitigate one or both of these concerns.

The vortex identification algorithm selected for use throughout this thesis is the Γ_2 Criterion, developed by Graftieaux et al. (2001). This particular algorithm was chosen for the present analyses because in addition to avoiding numerical differentiation, this method actually considers an integral of the velocity field, allowing the user to further suppress experimental noise as well as selecting an interrogation window size which attenuates the effect of vortical motions smaller than a certain size on the resulting visualization. This has particular utility for our relatively high Reynolds number setting, where in addition to the large-scale vortex shedding from the cylinder and airfoil (the coherent structures of interest) there is significant free-stream turbulence as well as turbulence induced by vortex shedding. Suppression of these small-scale structures by selecting a relatively large interrogation window helps to highlight large-scale coherent motions by suppressing background turbulence, while also mitigating experimental noise contributions.

The definition for the Γ_2 Criterion, approximated for computation from spatially discretized PIV data at some point in the flow $\mathbf{P} = [x_P, y_P]$ is given by (adapted from Graftieaux et al., 2001):

$$\Gamma_2(\mathbf{P}) = \frac{1}{N}\sum_{i=1}^{N} \frac{[\mathbf{PN_i} \times (\mathbf{U_{N_i}} - \overline{\mathbf{U_P}})] \cdot \mathbf{e_z}}{||\mathbf{PN_i}|| \cdot ||\mathbf{U_{N_i}} - \overline{\mathbf{U_P}}||}. \tag{2.14}$$

In Equation 2.14, N represents the number of locations adjacent to the point of interest $\mathbf{P}$ considered in a (square) interrogation window of some size, and $\mathbf{N_i}$ represents the i^{th} location. Thus, $\mathbf{PN_i}$ is the vector distance from the point $\mathbf{P}$ to the point $\mathbf{N_i}$, and $\mathbf{U_{N_i}}$ is the local flow velocity at $\mathbf{N_i}$. The velocity $\overline{\mathbf{U_P}}$ represents the mean velocity vector over the whole window at point $\mathbf{P}$, which is included to make the criterion Galilean invariant. The Γ_2 Criterion thus provides a metric for the tendency of the flow to swirl about the point $\mathbf{P}$, with an intensity bounded by $|\Gamma_2| \leq 1$. This provides a signed proxy for vorticity in the flow beyond some minimum scale of interest given by the choice of interrogation window. For all contours of the Γ_2 Criterion presented in this thesis, the interrogation window considered was 9x9 velocity vectors in size, corresponding to approximately 14mm to a side. This was found to be an acceptable compromise between mitigating experimental noise and the effect of background turbulence, while still allowing the visualization of structures of interest.

2.6.2 Calculating the Effective Angle of Attack

The effective angle of attack of the airfoil, α_{eff} is the angle of attack experienced by the airfoil both due to its geometry relative to the oncoming free stream, plus the effect of any inclination of the flow in the region of the airfoil relative to its nominal direction. This can be expressed as

$$\alpha_{\text{eff}} = \alpha_0 + \alpha_v = \alpha_0 + \text{atan}\left[\frac{V - \dot{y}}{U_\infty + \dot{x}}\right], \tag{2.15}$$

where V is the transverse fluid velocity in the region of the airfoil and $[\dot{x}, \dot{y}]$ are the airfoil's forward and transverse velocities. Figure 2.6 shows the above quantities on a diagram of the airfoil, where the nominal oncoming flow direction is aligned with $-x$, leading to the difference in signs between the numerator and denominator in Equation 2.15. In the experimental configuration(s) discussed throughout this thesis, significant flow velocities in the transverse direction are experienced by the airfoil induced by vortex shedding from the upstream circular cylinder, leading to a significant contribution to the effective angle of attack due to α_v.

To calculate α_{eff}, an appropriate value of V, the transverse flow velocity in the region of the airfoil must be determined. Since the flow is highly unsteady and measured V fields from PIV are noisy, an averaging method was used to estimate the quantity V for each frame. In particular, the flow velocity within a vertical slice of each V-field centered on the airfoil's location such that the slice just contains the full length of the airfoil was averaged for each PIV frame. This value serves as an estimate for V

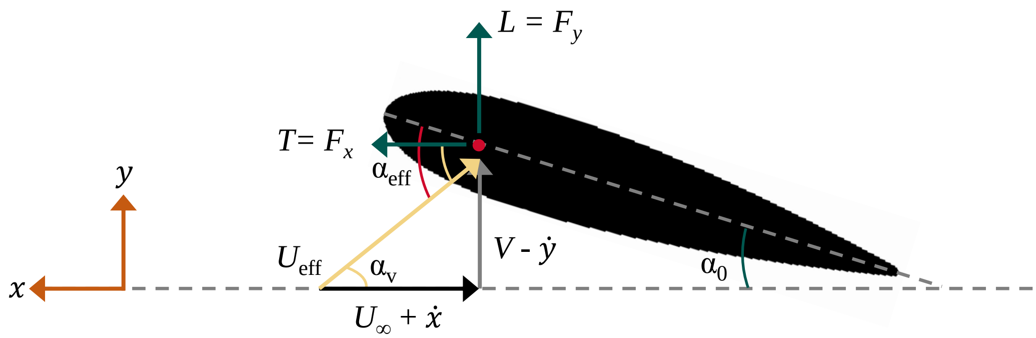

Figure 2.6: Diagram showing α_0, α_v and the resulting α_{eff} for the airfoil.

in the computation of α_{eff} above. It is important to note that in the experimentally obtained V fields, a large portion of this slice is occluded by the airfoil shadow. Thus, the mean value computed for V only represents velocities on the positive y-direction side of the airfoil, as the other side is occluded in all frames.

This calculation method results in a relatively smooth, quasi-sinusoidal variation in V for each trial; however, it likely provides a significant underestimate of the true value of α_{eff}. The region averaged not only includes occlusion from the laser shadow, but it also includes unsteady effects due to flow separation induced by the airfoil's presence, which include contributions in the opposite direction (or with a much reduced magnitude) to that of the bulk flow. In addition, the use of U_∞ to describe the nominal magnitude of the local flow speed in the x-direction is a clear over-estimate, since the airfoil sits in the velocity deficit region downstream of the cylinder. These factors together act to reduce the calculated value of α_{eff} relative to the true experience of the airfoil in the flow. Despite these challenges, the presented method allows for the visualization of trends in α_{eff}, and provides a conservative estimate for the true value of the effective angle of attack achieved by the airfoil.

2.6.3 Phase Averaging and Phase Mismatch Removal

Frequently, the data presented throughout this thesis represent phase-averaged quantities over one vortex shedding cycle. In cases where such data are presented, the following procedure is used to determine these phase-averaged quantities.

Step 0: For Driven Airfoil Experiments only, the phase match between the measured force and velocity is assessed to determine whether at each time instant it is within a specified tolerance. The procedure for this is given in Section 2.6.3.1. For Passive Captive Airfoil Experiments and MFRS Experiments, the phase match is considered acceptable at all time instants.

Step 1: Dynamic data (if relevant) measured either by MFRS sensors or through the CTS (position, velocity, x and y-direction forces) are interpolated onto simultaneous Particle Image Velocimetry (PIV) data, using measured triggering signals to ensure data are properly temporally aligned.

Step 2: A flow-based phase reference is extracted from each PIV snapshot, to determine the phase of vortex shedding at each time instant. This phase reference is determined based on analysis of a Proper Orthogonal Decomposition (POD) of the flow downstream of the circular cylinder, discussed in Section 2.6.3.2.

Step 3: PIV and dynamic data are binned based on the flow-based phase reference for each snapshot into phase bins with a width of $1°$, or 0.02 radians.

Step 4: Data points in each phase bin are averaged together.

This procedure produces phase-averaged data locked to the intrinsic behaviour of the oncoming flow. It is important to note that a bin width of 0.02 radians corresponds to an approximate time bin of 0.004 seconds, based on the expected period of vortex shedding. As this is wider than the time sampling of the flow through PIV ($\Delta t = 1/800 \approx 0.001s$), data points averaged in each bin are not necessarily statistically independent observations.

The following sections provide more details regarding phase matching for the Driven Airfoil Experiments, as well as a description of the POD-based method used to determine vortex shedding phase.

2.6.3.1 Determining Phase Match for Driven Airfoil Experiments

For the Driven Airfoil Experiments discussed in this thesis, the airfoil's position conforms to a pre-planned trajectory of the form given by Equation 2.10. The phase of its motion relative to oncoming vortex shedding, ψ, is therefore fixed at the moment of triggering based on conditions at this initial time in the trajectory. For this set of experiments, we are only interested in times when the airfoil's velocity is in-phase with the transverse force acting on it, as discussed in Section 2.4.1. Although triggering is programmed based on measurement of flow conditions to attempt to ensure this condition is met (discussed further in Chapter 3), meander in the frequency and phase of forcing due to vortex shedding is present in the obtained signals. Therefore, before computing phase-averaged quantities, pre-processing of the Driven Airfoil Experiment data is necessary to ensure that only moments in time with good phase match are included in phase averages.

Though there are many algorithmic options to accomplish this phase match validation, any potential method should satisfy the following criteria:

- Signals should be analyzed on a per-period rather than full-record basis, to help ensure the history of the airfoil behaviour is preserved without rejecting full records (which each contain several potentially well-aligned periods).
- The analysis method should provide access to information about the frequency and amplitude of the forcing for each period, to facilitate the collection of statistics regarding these quantities.

In view of these criteria, the following method was selected.

Step 1: Synchronized force and velocity signals post-triggering are windowed into segments with length T, the estimated vortex shedding period. The mean transverse force over the total record length is calculated prior to windowing.

Step 2: For each segment, a sine wave is fit to both the velocity and force data. In the case of velocity, the frequency and amplitude of the signal are known (as these are directly controlled by the CTS), so the only unknown parameter in the sine fit is the phase. For the forcing, a sine wave with an unknown frequency, amplitude and phase and an offset equal to the mean forcing for the whole record is fit. Sine fitting for both segments (force and velocity) is accomplished using the MATLAB built-in function `fminsearch()`, with the function to be minimized the mean square error of the fit relative to the observed data. Fits for all data considered were validated by eye before use in the thesis.

Step 3: Segments where the frequency of transverse forcing recovered through sine fitting differs from the expected forcing frequency of 0.6 Hz by more than 10% are rejected.

Step 4: Segments where the mean difference between the times the sinusoids fit to force and velocity achieve an extremum is greater than 5% of T are rejected. If the force and velocity have different numbers of extrema within the segment, the segment is rejected.

Step 5: Remaining segments (those not flagged for rejection in Steps 3 and 4) are included in phase averaging. The amplitude and frequency recorded for all transverse force segments, including those flagged for rejection, are saved to contribute to statistics regarding forcing to the airfoil.

2.6.3.2 Determining a Flow-Based Phase Reference from POD

Flow behind a circular cylinder exhibits von Kármán vortex shedding, which has a relatively compact representation in a basis formed by modes computed through the Proper Orthogonal Decomposition (POD); see for example Brunton et al. (2016). This can be exploited to create a phase reference for vortex shedding based on direct observations of the flow field, as developed by van Oudheusden et al. (2005), and later used by Lefebvre and Jones (2019) in a similar study to the present case. This method is particularly useful for determining a quantitative phase reference for the experiments presented in this thesis, since the airfoil's motion is not always perfectly periodic.

Both mathematical aspects as well as a variety of practical implementations of POD relevant to the analysis of fluid flows have been developed in recent decades (for example, see Berkooz et al. (1993)). As succinctly summarized for experimental data by Manohar et al. (2018), POD modes corresponding to an experimental data set may be computed based on the Singular Value Decomposition (SVD) of a snapshot matrix $\mathbf{S}$, where each column $\mathbf{S}_i$ corresponds to an individual observation of the field of interest (here, the transverse flow velocity V) at some time t_i. Then, the POD modes describing the data set $\mathbf{S}$ correspond to its left singular vectors, denoted $\mathbf{\Psi}$, computed such that

$$\mathbf{S} = \Psi\Sigma\mathbf{V}^T. \tag{2.16}$$

Columns of the matrix $\mathbf{\Psi}$, each corresponding to an individual POD mode $\mathbf{\Psi}_j$, are ranked by energetic importance through the associated singular value $\Sigma_{jj} = \sigma_j$. A low-rank representation of the flow field at each time step can then be determined through projection of energetically important (large-σ) modes onto each snapshot to determine a mode coefficient $A_j(t_i)$, where

$$A_j(t_i) = A_{ji} = \Psi_j^T\mathbf{S}_i, \tag{2.17}$$

and $\mathbf{S}_i$ is the i^{th} column of the matrix $\mathbf{S}$, corresponding to the i^{th} flow snapshot.

For the analysis presented in the following chapters, the global phase of the cylinder-airfoil system is computed based on underlying POD modes of the flow field in the region upstream of the location of the airfoil, in the following manner.

Step 1: For PIV data collected on a particular day of testing, the leading two POD modes for the flow upstream of the airfoil's location are computed. In the particular implementation used in this thesis, the modes of interest Ψ_j are recovered from the SVD of $\mathbf{S}^T\mathbf{S}$, by considering

$$\mathsf{svd}(\mathbf{S}^T\mathbf{S}) = \mathbf{V}\Sigma^2\mathbf{V}^T, \tag{2.18}$$

$$\Psi = \mathbf{S}\mathbf{V}\Sigma^{-1}, \tag{2.19}$$

$$\Psi_j = \frac{1}{\sigma_j}\mathbf{S}\mathbf{V}_i. \tag{2.20}$$

This method is used for legacy reasons and to conform with discussion in van Oudheusden et al. (2005), and does not represent the most efficient means of calculating

Ψ_j, σ_j; however as this is a pre-computation step, no updates were made to enhance efficiency.

The POD modes above are determined based on snapshots of V (y-direction velocity) from up to 10 trials collected on each day. The number of trials (each containing 3675 snapshots) was limited to 10 or fewer, even if additional trials were available on a particular day, since the time and computer memory required to compute POD modes depends strongly on the number of snapshots considered. Due to recalibration of the optical equipment for PIV on each day when data were collected, the underlying spatial field for data collected on each day varies slightly, and corresponding mode shapes are therefore required. Only the region of the field of view upstream of the location of the airfoil was considered for building the POD modes.

Step 2: Once appropriate leading-order modes have been computed from multiple trials, the modes are projected onto the corresponding region in each flow snapshot (each PIV time step) in a particular trial, to obtain projection coefficients for the leading order modes as a function of time corresponding to that trial (according to Equation 2.17).

Step 3: The phase angle representing the phase of vortex shedding for each snapshot is calculated as in van Oudheusden et al. (2005):

$$\phi_i = \operatorname{atan}\left(\frac{A_{2,i}\sigma_1}{A_{1,i}\sigma_2}\right). \tag{2.21}$$

In the above equation, $A_{1,i}$ and $A_{2,i}$ are the first and second POD mode coefficients computed by projecting the determined mode shape on a particular snapshot i according to Equation 2.17, and σ_1, σ_2 are the singular values associated with the matrix $\mathbf{S}$, which contains snapshots (with units of velocity) corresponding to the many trials used to determine the underlying modes for each day of testing.

This method allows for the recovery of a 'global' phase reference based on conditions in the flow, which is particularly important for the creation of phase-averaged flow fields. It also allows for the association of airfoil behaviour strongly with a particular portion of a vortex shedding cycle, which helps to explicitly link airfoil and fluid behaviour.

In general, the POD modes computed for each day of testing in order to determine this global phase reference are very similar; however, since a single pair of modes cannot be used for all experimental data across days (due to variations in the underlying spatial fields), there is no guarantee that the zero-phase location ($\phi = 0$) for each day will be the same. That is, data sets with global phase determined based on different sets of POD modes may exhibit small phase offsets relative to one another not because the airfoil's interaction with the oncoming flow is different, but because the location of $\phi = 0$ may vary slightly between data sets. Thus, comparisons of the global phase of events taking place in the flow between data sets should be made with caution. In the discussion of Driven Airfoil Experiments in Chapter 3, when there is a clear phase reference based on airfoil motion (enforced by our regulation of the phase difference between force and velocity as discussed in Section 2.6.3.1), phase-averaged data are presented such that the $\phi = 0$ location aligns with the phase where the airfoil achieves its minimum position. For later chapters discussing Passive Captive Airfoil motion, no adjustment is made to the $\phi = 0$ location for each set of data (since there is no expectation that airfoil interactions with oncoming vorticity have a standard phase relationship); however the Γ_2 Criterion field associated with $\phi = 0$ for each set of experiments is provided. Importantly, this complication does not affect the interpretation of phase-averaged signals relative to one another within the same set of experimental trials.

2.6.4 Filtering of Instantaneous Γ_2 Criterion Fields

In cases where phase averaging is not possible due to the aperiodic nature of the airfoil's behaviour in time, space, or both, temporal and spatial filtering (in addition to that provided inherently by the Γ_2 Criterion) is often necessary to clarify instantaneous vorticity behaviour.

In these cases, a standard filtering procedure is applied to Γ_2 fields. First, a Savitzky-Golay filter with a width of 10 grid points (approximately 15mm) is applied to each snapshot in the x and y directions, and these horizontally and vertically filtered fields are averaged together at each spatial location. Then, a non-causal moving average is used to additionally smooth each spatial location in time. The average is computed over 11 snapshots (5 previous and 5 future snapshots), representing an averaging period of approximately 0.01s.

2.6.5 Wake Visualization using Taylor's Hypothesis of Frozen Flow

It is illustrative to make use of Taylor's Hypothesis of Frozen Flow to visualize large wake structures arising from interactions between airfoil and cylinder-derived vorticity. The key assumption in this hypothesis is that convection by the mean flow is the only driving force in the evolution of structures, and other effects such as vortex-vortex interactions and turbulent fluctuations, are insignificant. Limitations of these assumptions are numerous, especially over long distances, as discussed for example by Dennis and Nickels (2008).

In the present case, vortex interactions are at least as important as the mean flow, the structures formed behind the airfoil dissipate with time and distance downstream, and have a convection velocity with a component in the direction orthogonal to the mean flow. Thus, visualizations using Taylor's Hypothesis certainly do not provide an accurate *spatial* picture of the flow. However, if we interpret the resulting visualization not as a representation of what wake structures look like in space, but rather their evolution in time in a limited spatial region near the airfoil trailing edge for example, we are able to extract useful and interpretable information about wake formation that is not easily seen using other visualization methods.

To create a frozen flow visualization from the time-resolved PIV data collected as described in Section 2.5, the following steps are taken:

Step 1: Γ_2 criterion data is phase averaged over one vortex shedding period, as described in Section 2.6.3. This creates a smoother and more representative field from which to create the wake visualization. A particular x-location in the PIV field of view, x_{ROI}, is selected to determine the region of interest for our visualization.

Step 2: A moving average filter in time is applied to the Γ_2 Criterion fields, and then filtered fields are downsampled in time by the filter width. This further smooths the Γ_2 fields, and creates representative fields that are father separated in time (and therefore in space in our visualization, applying Taylor's Hypothesis).

Step 3: The free-stream velocity U_∞ is used to determine the displacement of frozen flow from one frame to the next. Although this is a serious simplification, the same free stream velocity value is used throughout the frame (we assume that all regions of the flow translate with the same convection velocity, U_∞). This allows us to consider rectangular 'slices' of flow from each frame taken starting at the spatial location x_{ROI}, with a (dimensional) width of $U_\infty \Delta t$ where Δt is the time between downsampled frames. The height of each slice is the full height of the PIV frame.

Step 4: For each frame of our downsampled, phase-averaged data, a slice as determined above is added to our overall frozen flow image. The slice from the first frame is added at the farthest downstream position, and each subsequent slice from a frame later in time is placed upstream of that. This procedure is repeated until all frames are included.

Step 5: The created field over one cylinder shedding period is repeated (copied) downstream, to visualize several (identical) periods of vortex shedding.

Step 6: The region upstream of our slice location in the last frame of phase averaged data is added to the visualization to show the phase of vortex shedding upstream, as well as the airfoil location.

Step 7: A Gaussian smoothing filter is applied to the reconstructed field to minimize dislocation artifacts between slices.

Using the procedure above, a simplified representation of the time-based evolution of the flow at some particular spatial location x_{ROI} is made available. This visualization is not quantitatively accurate in a spatial sense, as effects of dissipation, vortex-vortex repulsion/interaction and differing convection velocities due to the wake deficit region are entirely ignored; in addition, wake structures leaving the PIV field of view are not considered. However, it allows us to visualize simplified fields corresponding to the vorticity shed from the cylinder, the airfoil, and/or their interaction(s), which proves useful in the analysis presented in the following chapters.

Chapter 3

CHARACTERIZATION OF THE BEHAVIOUR OF AN AIRFOIL DRIVEN IN THE WAKE OF A CYLINDER

3.1 Introduction

The behaviour of an airfoil placed in the wake of an upstream circular cylinder forms the common thread throughout all chapters of this thesis. The purpose of this, the opening chapter regarding the set of experiments and related investigations conducted in support of this thesis, is to orient the reader in the dynamics of such a system in its simplest possible implementation, and extend previous studies describing its myriad behaviours.

In this chapter, all experiments presented are Driven Airfoil Experiments as discussed in Chapter 2. The airfoil is mounted a fixed distance downstream of the cylinder trailing edge, and actively oscillated in the transverse direction with a fixed amplitude and frequency in-phase with the transverse force it experiences. These experiments present only a single point within the large parameter space governing the interaction of such a driven foil with oncoming vorticity, which may be influenced by dozens of additional factors not explored here. The particular combination of parameters chosen for these foundational experiments provides a backdrop for further experiments regarding the behaviour of a similar *passive* system, which we discuss as an archetype for a fully passive flapping foil energy harvesting device. By closely interrogating both the power and thrust production for this repeatable and relatively well characterized system, and connecting those outputs from the system to interactions with oncoming vorticity shed from the cylinder, the stage is set for further studies of the more complicated fully passive case.

In this chapter, both the oncoming flow from the cylinder and the resulting transverse force acting on the airfoil are characterized. Then, the phase-averaged dynamics of the airfoil and its resulting thrust and power extraction potential are presented. Detailed discussion of interactions with the oncoming flow connects the observed thrust and power production to the combined airfoil-cylinder wake. A mechanism is proposed for the formation of wake structures arising from these interactions. Finally, the effect(s) of changing the airfoil's static angle of attack on these interactions (and the resulting changes to power and thrust production) are explored.

3.2 Characterization of the Cylinder Wake

Though there exists a wealth of literature on the topic of vortex shedding from a cylinder (see for example Williamson (1996) for a thorough review of the topic), here we provide salient characteristics of the particular vortex wakes which form the forcing to our airfoil in the remainder of this thesis. To characterize the oncoming flow to the airfoil, approximately 29 periods of vortex shedding behind the cylinder were captured with Particle Image Velocimetry (PIV) over 10 individual trials, without the airfoil present. The oncoming flow speed was the same as for all experiments presented in this thesis. Figure 3.1 shows the mean flow in the x-direction for this configuration, denoted $\bar{U}$ throughout the thesis. The figure shows contours of $\bar{U}$, as well as three profiles of $\bar{U}$ in the y-direction, each corresponding to flow past a station at the x-position indicated. The profiles show the expected behaviour in the cylinder wake: that is, there is a momentum deficit region generated by the cylinder in the region downstream, whose effect is attenuated at subsequent downstream locations. The high blockage ratio of 0.25 causes flow between the cylinder and the wall to be accelerated, and may play a role in determining the Strouhal number of the cylinder shedding experienced in these experiments (discussed for example by Beal et al., 2006; Liao et al., 2003; Bearman and Zdravkovich, 1978); therefore, the frequency of vortex shedding was directly measured by placing virtual probes in the PIV fields for all obtained data. The location of the two virtual probes is shown as the grey and yellow triangles in Figure 3.1.

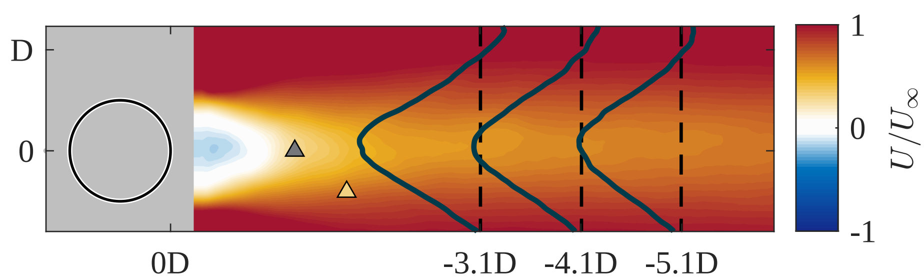

Figure 3.1: Contours of $\bar{U}$, the mean flow velocity in the x-direction. Flow velocity here is shown positive downstream, though this corresponds to negative values of x. Overlaid on contours are three y-direction profiles of $\bar{U}$ (▬), located at the x-positions indicated. An x-position of $-3.1D$ corresponds to the location of the airfoil's quarter-chord location in the following discussions of Driven Airfoil Experiments. $\bar{U}$ profiles passing to the right of their corresponding dashed line indicate y-locations where $\bar{U} > U_\infty$. Triangle icons correspond to virtual velocity probe locations giving rise to the spectra pictured in Figure 3.2.

The 10 individual PIV trials considered here each have a length of approximately $2.9T$, where T is the estimated period of vortex shedding based on preliminary investigations of airfoil behaviour described in detail in Chapter 2. To confirm that a period of $T = 1.6$ s is representative of vortex shedding from the cylinder, FFTs of the y-velocity of the flow at the virtual probe locations in Figure 3.1 were computed for each trial, and the resulting frequency spectrum was ensemble averaged. The mean spectrum for the centerline location (shown in grey) as well as for an off-center location (shown in yellow) for the virtual probes are given in Figure 3.2. A strong peak is visible in both spectra at approximately 0.65 Hz. This confirms that preliminary observations of the vortex shedding frequency, which arose from measuring the y-direction forcing acting on the airfoil, are largely accurate. Since those observations were made with a much higher frequency resolution (the limitations of which can be seen for the present case in Figure 3.2), the value of the shedding frequency determined as described in Chapter 2, 0.62 Hz, is used both for the frequency of the driven airfoil motion, and to describe the frequency of vortex shedding throughout this thesis. This value gives rise to a Strouhal number for vortex shedding of approximately $St = 0.22$.

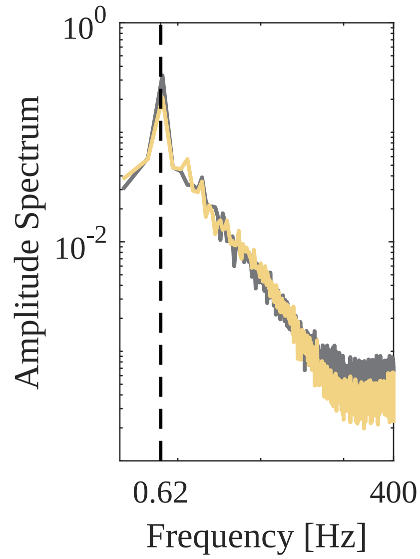

Figure 3.2: Ensemble-averaged frequency spectra for y-direction flow velocity with only the cylinder placed in the flow. ▬ Spectrum from measurements at the grey probe location in Figure 3.1; ▬ Spectrum from yellow probe.

To visualize the cycle-averaged behaviour of the cylinder wake, PIV data were phase-averaged based on the phase of vortex shedding in each frame, using the methods described in Chapter 2. Each phase bin had a width of 1°, and no fewer than 80 PIV frames were averaged per bin. Phase-averaged contours of the Γ_2 Criterion (discussed in detail in Chapter 2), a proxy for vorticity, were then used to create a frozen-wake visualization to reveal the pattern of vortices shed from the circular cylinder, shown in Figure 3.3. A 2S-type pattern is visible, corresponding to a classic von Kármán wake. During each shedding period, one Clockwise (CW) or negative vortex and one Counterclockwise (CCW) or positive vortex is shed from the cylinder. Although wake vorticity shed from either side of the cylinder appears

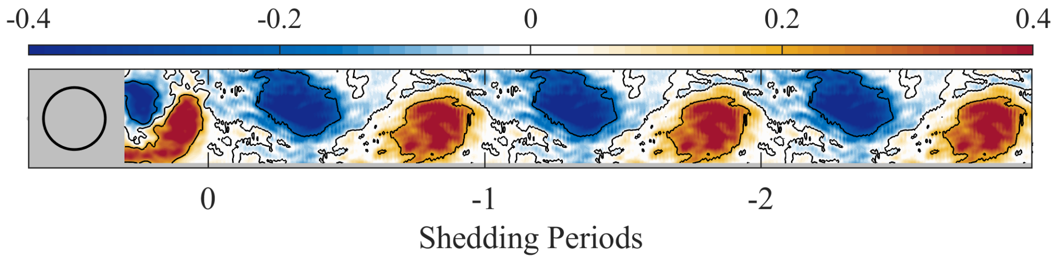

Figure 3.3: Frozen flow wake visualization for the Cylinder-Only experiments discussed in this section. Phase averaged contours of the Γ_2 Criterion are used to create the visualization, according to the method presented in Chapter 2. Contour lines indicate levels of ±0.2 and 0.0 in the figure, as indicated in the colour bar.

to be largely similar, slight asymmetry in the extent of the regions of positive and negative vorticity is observed. This could be due to an offset in the placement of the cylinder in the tunnel, as discussed in Chapter 2. This asymmetry is very slight however, and is unlikely to lead to large differences in forcing experienced by the downstream airfoil. Wake visualizations of this type presented in the following sections highlight the effect that the presence of the driven airfoil has on the wake, as airfoil and cylinder-derived vorticity interact.

3.3 Experimental Setup for Driven Airfoil Experiments

For the Driven Airfoil Experiments discussed in this chapter, testing conditions conformed to those described in detail in Chapter 2. The NACA 0018 airfoil was mounted a distance of $\Delta x = 3.1D$ downstream of the circular cylinder, and the CTS was used to drive it through a pre-planned sinusoidal trajectory with a specified mean position, an amplitude of 0.0051 m ($0.04D$), and a frequency of 0.6226 Hz (the estimated frequency of vortex shedding). Time-resolved Particle Image Velocimetry (PIV) was used to record the flow velocity in the region of the airfoil at a rate of 800 Hz, as well as velocities in the cylinder wake upstream and the combined airfoil-cylinder wake downstream. The CTS simultaneously recorded the airfoil's position and velocity as well as forces acting on the airfoil at a rate of 200 Hz. For all of the experiments presented, the airfoil's mean position was directly downstream of the cylinder, a position denoted $y = 0$. Unless specified otherwise, the airfoil's geometric angle of attack was held at 0° ; Section 3.11 presents results of changing the mean geometric angle of attack on airfoil performance. Cases tested are summarized in Table 3.1.

For each set of experiments corresponding to a particular case given in the Table, data

Case Name	Abbreviation	Δx	Periods Recorded	Mean y-Position	AoA
Basic	BA	$3.1D$	$157T$	$0.0D$	0°
High Positive AoA	HAoA+	$3.1D$	$107T$	$0.0D$	10°
High Negative AoA	HAoA-	$3.1D$	$128T$	$0.0D$	-10°

Table 3.1: A summary of configurations tested during the Driven Airfoil Experiments described in this chapter. Δx indicates the distance from the cylinder trailing edge to the airfoil's quarter-chord location.

were recorded as multiple independent trials, each with a duration of $2.9T$. Before each trial, the airfoil was moved to its maximum position in the tunnel. As vortices are shed and convect past the airfoil, they generate approximately sinusoidal forcing that is measured by the CTS. Since the goal of the Driven Airfoil Experiments was to drive the airfoil in-phase with measured forcing, the airfoil's motion was triggered at a moment when the measured force was passing through zero, with a negative gradient. In this way, at the moment of triggering, force and velocity are aligned. A causal averaging filter (where at each time step, a mean is computed over the current and previous $n - 1$ measurements) with a width of $n = 10$ measurements was used to smooth the measured force for triggering only. This was necessary to improve the robustness of the triggering algorithm to noise in the measured force signal and prevent false triggers. To compensate for the phase lag induced at the moment of triggering due to this averaging process, the desired airfoil position and velocity programmed for the moment of triggering were advanced through the phase lag. Then, as motion is initiated, the airfoil's motion is already 'caught up' to the true forcing phase.

Because no attempt was made to correct phase errors arising from slight mismatches in frequency and/or triggering phase between the forcing and the velocity, during each trial there is noticeable phase drift between these two quantities. As the motion is triggered with the appropriate phase relationship, this effect generally becomes more noticeable further from the moment of triggering. To ensure that we restrict our analysis only to moments in time when forcing and velocity are in-phase, pre-processing of the recorded force and velocity data was carried out as described in detail in Chapter 2 to remove periods of airfoil motion when the phase difference between force and velocity exceeded 5% of T. Thus, all phase-averaged dynamic and flow data discussed in this section are constrained to exhibit the desired phase relationship. For more information and justification for this constraint, please see discussion of the Driven Airfoil Experiments in Chapter 2.

3.4 Frequency and Amplitude Variation in Oncoming Forcing

Statistics regarding the nature of the oncoming forcing, $F_y = L$ or lift relative to U_∞, were collected for the Driven Airfoil Experiments. Figure 3.4 shows a histogram of forcing amplitudes fit to each measured period of the y-force, as described in Chapter 2. Overlaid on the histogram is a Normal Distribution fit to the amplitude data using MATLAB's built-in function `fitdist()`. The black dashed line shows the mean of the fit distribution. The total number of periods considered is 157. In contrast to discussion of phase match validation prior to phase averaging this data for further processing, these 157 periods include those where the force and velocity are misaligned, often significantly. This may have impact on the experienced forcing, as airfoil dynamics are closely mediated by the phase of its interactions with oncoming vorticity; however, we see that the mean forcing amplitude recovered, 2.0 N, corresponds to a lift coefficient of $C_{L,\max} = 0.88$. This is consistent with the maxima in the lift coefficient for the phase averaged data presented the following section, where these mismatched times are not included.

The standard deviation of the fit distribution is 0.5 N or equivalently $\Delta C_{L,\max} = 0.19$. Even accounting for measurement noise (which should be largely attenuated by the sine fitting procedure used to obtain the amplitudes), which from Chapter 2 causes

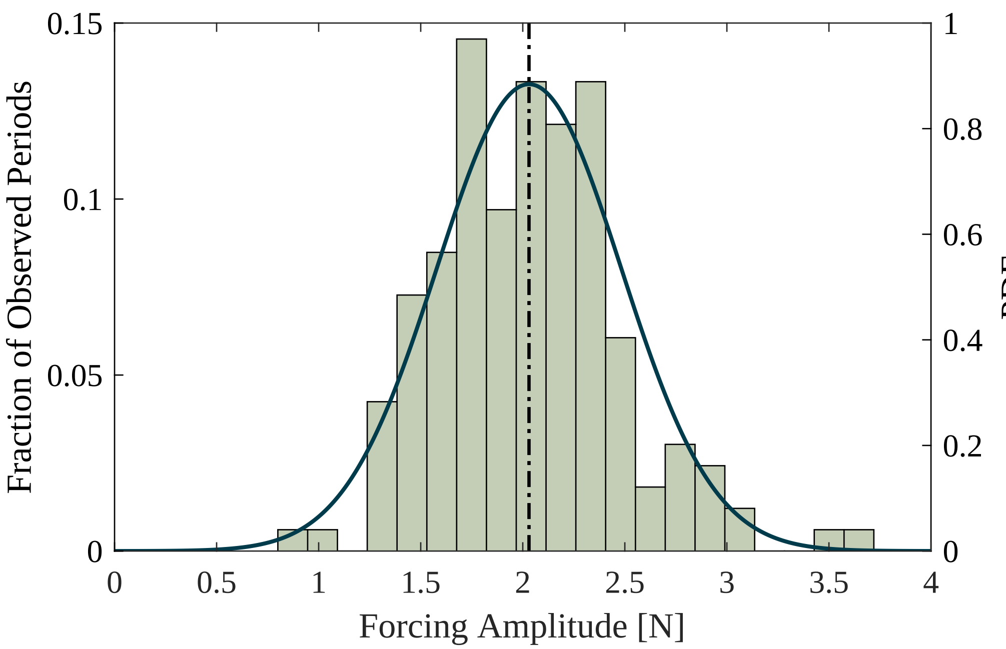

Figure 3.4: Histogram (left axis) and corresponding Normal Distribution fit (right axis) to forcing amplitudes recovered from 157 periods of the Driven Airfoil Experiments. ▪▬ Mean of Fit distribution.

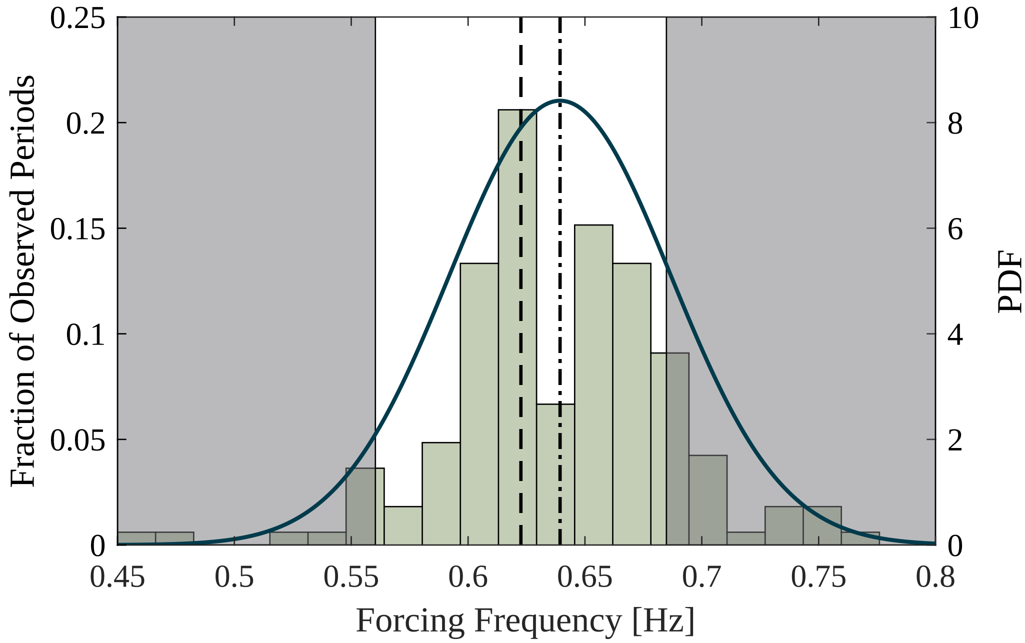

Figure 3.5: Histogram (left axis) and corresponding Normal Distribution fit (right axis) to forcing frequencies recovered from 157 periods of the Driven Airfoil Experiments. ▪▬ Mean of Fit distribution; and ▬▬ Expected frequency of vortex shedding from previous analysis. Grey regions show recovered frequencies beyond which data were rejected from phase averages presented earlier.

an uncertainty of $\delta F_y = \pm 0.1$ N or $\delta C_L = \pm 0.04$ in the measured forces, there is relatively large variation in the experienced lift from cycle to cycle for the airfoil. Though it is unclear to what extent mismatched phase between forcing and velocity plays a role in widening the observed force distribution, this still provides a useful benchmark to understand the extent to which cycle-to-cycle forcing varies in time. Although in the Driven Airfoil Experiments presented in this section, the airfoil's motion is determined by the programmed trajectory and therefore not sensitive to such variations in forcing amplitude, we expect that for the Passive Captive Airfoil experiments discussed in Chapters 4 – 6, this variation will play a significant role.

In addition to forcing amplitude, statistics regarding the frequency of the oncoming forcing were also collected. Figure 3.5 shows a histogram and corresponding Normal Distribution describing the recorded frequencies recovered from the same sinusoid fit to each of the 157 observed periods as for the forcing amplitude. As in the amplitude case, the fact that the airfoil is in motion relative to the flow when F_y is measured has the potential to influence the observed forcing frequencies. Although the inertial effects due to the accelerating reference frame of the force sensor are removed prior to determining frequency content in the force signal (as discussed in Chapter 2), added mass effects from the surrounding fluid are not. These effects could provide a bias in the recorded frequencies towards the driving frequency. As the driving frequency corresponds to the the expected forcing frequency, such effects are impossible to disentangle from the desired measurement of frequency content due to vortex shedding. However, for the purposes of determining the *actual* experienced forces acting on the airfoil, frequency content due to added mass (or other effects) is relevant to airfoil behaviour. Therefore this analysis again provides a useful baseline from which to understand the variation in forcing frequency that a similar passive system may experience.

From Figure 3.5, we see that the recovered mean forcing frequency of 0.64 Hz (given by the dash-dot line) corresponds closely to our expected forcing frequency of 0.62 Hz (shown by the even dashed line). The standard deviation recovered from the normal distribution fit is 0.05 Hz, or approximately 7% of the mean forcing. The recovered frequency from sine fitting was also used to reject periods for phase averaging where the estimated frequency differed from the driving frequency for the airfoil (0.6226 Hz) by more than 10%: this represents approximately one standard deviation. The bounds for frequency outside of which periods were rejected for phase averaging are given by the greyed out regions in the Figure.

3.5 Lift, Thrust and Power Produced by a Driven Airfoil

To describe the performance of the airfoil as it executes its pre-planned trajectory through the flow, we consider phase-averaged dynamical quantities plotted over one vortex shedding cycle. To compute these phase averages, each PIV frame is tagged with a phase angle, as described in detail in Chapter 2. Then, all frames corresponding to a particular phase angle, as well as the corresponding dynamic quantities measured at the same time as each PIV field are averaged together. As described in the previous section, only frames where the measured y-direction force is in-phase with the velocity are included in phase averages. Figure 3.6 shows a histogram of the number of PIV frames averaged in each phase bin. From the Figure, all bins contain at least 40 observations. The bin size is 1°, or 0.02 radians.

Figure 3.7 shows phase-averaged dynamic quantities as a function of vortex shedding phase angle ϕ. From the figure, the motion of the airfoil seems to have a consistent phase relative to oncoming shedding, indicating our phase-match discrimination has adequately rejected times where the force and velocity are misaligned. The shape and amplitude of the velocity and position of the airfoil are consistent with the driven trajectory input, and have relatively low spread within each phase bin (indicated by the limited interquartile range at each phase, shaded in green in the figure). In addition, the Lift Coefficient C_L experienced by the airfoil as a function of phase also appears relatively consistent, with a roughly sinusoidal shape which mirrors that of the velocity.

Higher variability in the recorded values for apparent angle of attack, thrust, and power coefficient are apparent in Figure 3.7. In the case of C_P, this is a result of the composite nature of the plotted signal; C_P is not directly measured, but is computed as a product of C_L and y/U_∞. Therefore, the moderate variability in each of those quantities is amplified in C_P. Similarly, α_{eff} is also a computed quantity, based on the y-direction velocity of both the airfoil and the flow. Though averaging and other efforts are made to extract a relatively smooth and accurate value for α_{eff}, it again suffers from composite variability contributed from its primary constituents.

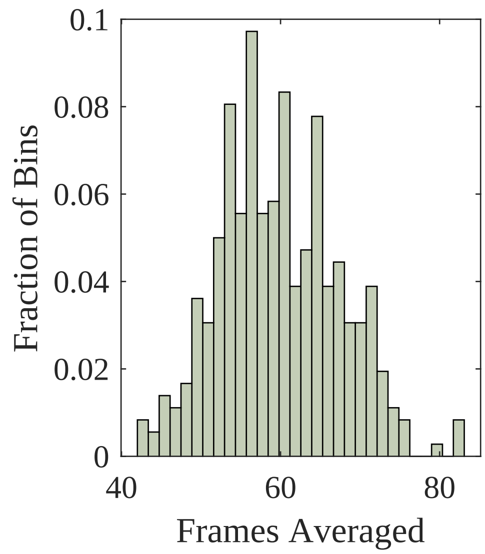

Figure 3.6: Histogram showing the fraction of phase bins with each number of observations included in phase averaging for the Driven Airfoil Experiments.

In contrast to these two signals, the thrust coefficient C_T is directly measured, but seems to exhibit the highest variability of all. Measuring C_T constitutes the most significant measurement challenge in this experimental setup, since the same sensor is used to measure both x and y-direction forces; the latter has maxima an order of magnitude larger than the former, which necessitates a large dynamic range for the sensor used. In addition, the sensor itself is mounted to the CTS and actuated during a run. Although contributions to the measured forces due to the accelerating reference frame are removed before interpreting the force values, vibration and other mechanical contributions are still present in the measured signals, and affect the smaller measured thrust values more than the larger values for lift. These issues are particularly important for experiments where the airfoil is actuated in the x-direction (all of the experiments in this section discuss y-direction or transverse

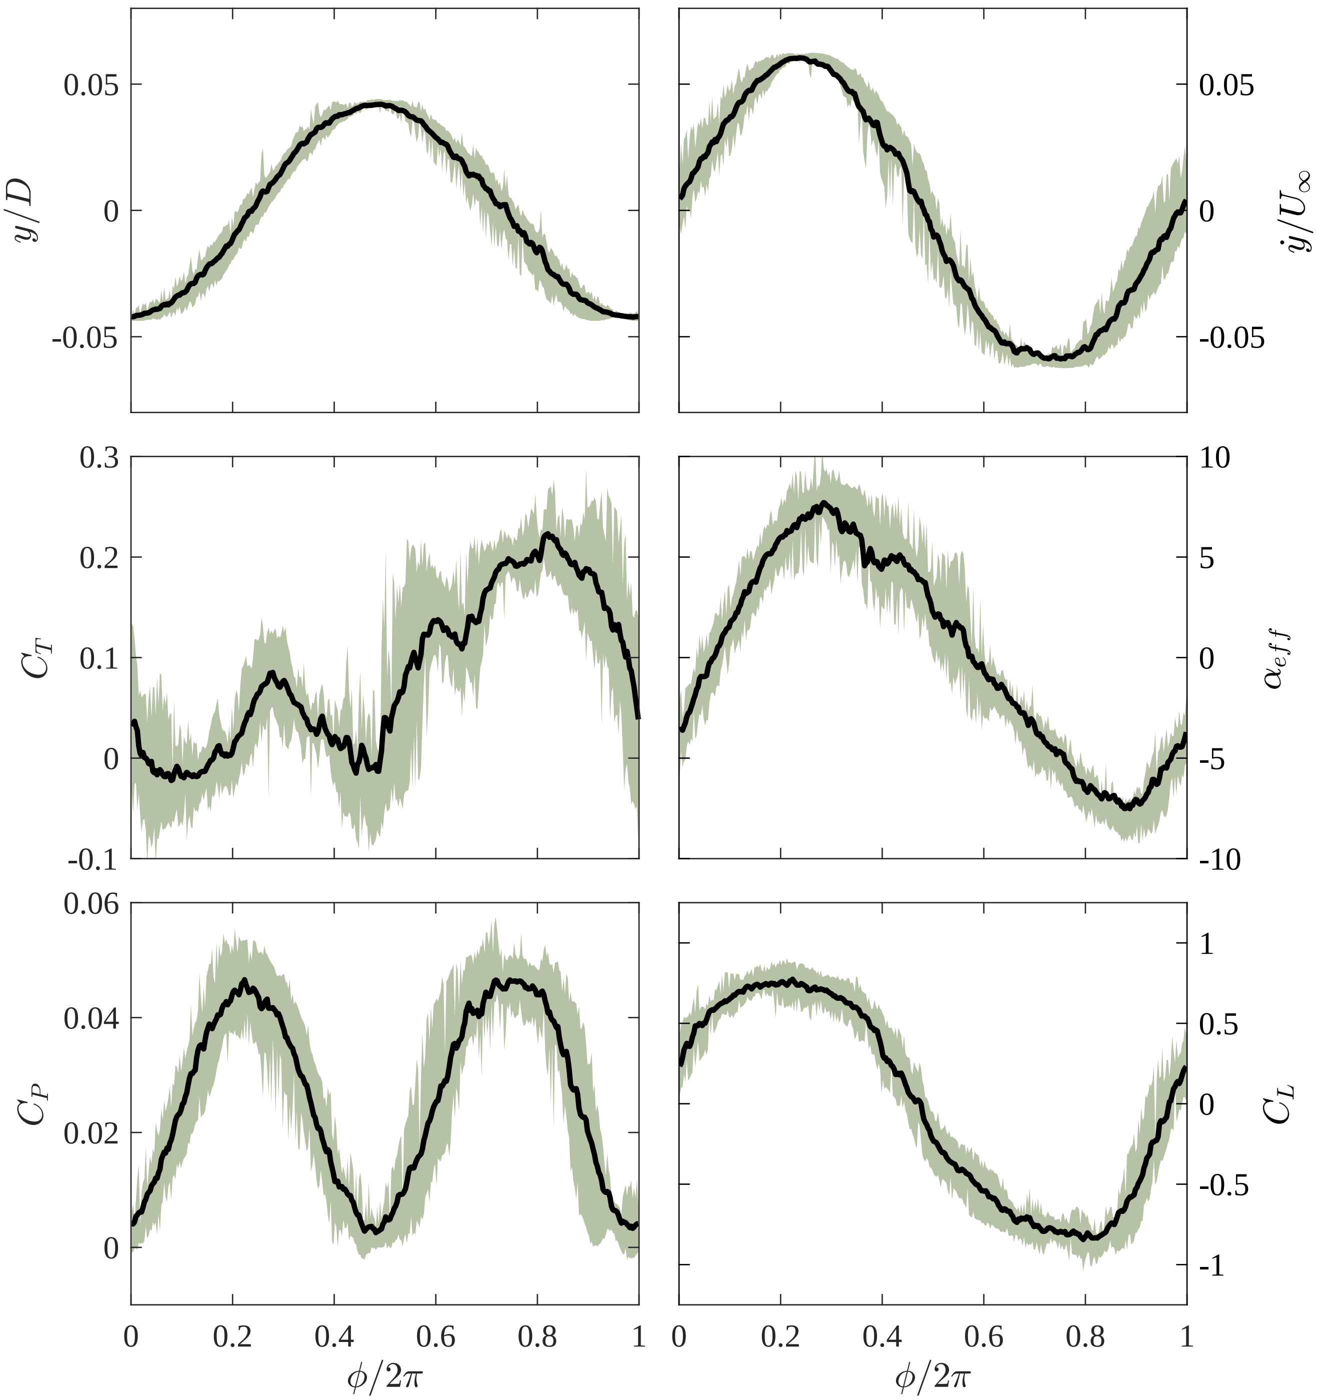

Figure 3.7: Phase-averaged dynamic quantities for the Driven Airfoil Experiments. Top Left: Airfoil Position. Top Right: Airfoil Velocity. Middle Left: Thrust Coefficient. Middle Right: Effective Angle of Attack. Bottom Left: Power Coefficient. Bottom Right: Lift Coefficient. Shaded regions indicate the interquartile range for data averaged in each phase bin. Bold black lines indicate the mean values recorded in each bin, with a 3rd order Savitzky-Golay filter applied to smooth the curves. The filter width was set to 5°.

motion only), and a more thorough discussion of sensor performance is therefore provided in Chapter 5. In the present case, we clearly extract general trends in the thrust signal, despite the higher variability exhibited in this quantity.

Overall, we see that the airfoil is performing the desired sinusoidal trajectory, such that lift and velocity are in-phase. Though in this case the airfoil is driven and therefore it consumes power and does not produce it, the work done on the airfoil by the flow per unit time, given by C_P, is positive over the whole cycle. This indicates that the system has the potential to extract net power while undergoing this trajectory. In addition, the mean thrust coefficient $\overline{C_T}$ is also positive: the airfoil is therefore simultaneously producing net thrust. One interesting feature in the phase-averaged signals recovered is the asymmetry in the magnitude of the thrust experienced by the airfoil when its velocity is positive, vs when it is negative. This is not reflected in any of the other measured quantities, which appear to have symmetric behaviour whether the airfoil is moving in the positive or negative y-direction. Though the cause of this asymmetry is unknown, it is likely the result of the interplay between geometric thrust producing effects (the Katzmayr Effect, discussed in detail in Section 3.6) and potentially intermittent fluid-structure interactions which either enhance or reduce lift and drag acting on the airfoil. Such interactions and their potential effects on thrust production by the airfoil are discussed at more length in the following sections.

3.6 Quasi-steady Thrust Production and the Katzmayr Effect

The Katzmayr Effect, the name of which refers to an early set of laboratory experiments performed in 1922 by R. Katzmayr, is a geometric effect which allows lifting bodies to experience a net thrust when there is a mismatch between the oncoming flow direction and the desired direction of travel. Early experiments quantifying the effect remarked that subjecting an airfoil to a ‘wavy stream’ or a sinusoidal variation in oncoming flow direction gave rise to a non-zero net thrust, or a net force pointing upstream in the mean flow direction. The same effect was not observed when the airfoil’s angle of attack was varied through a similar range of angles as the effective angle of attack induced by the wavy stream; in fact, such oscillations of the airfoil itself in either heaving or pitching reduced the lift to drag ratio for all oscillation frequencies tested (Katzmayr, 1922). This effect is equivalent to that which allows sailors to travel upstream against the mean wind direction by ‘tacking’, or travelling through a zig-zag trajectory across the mean wind direction, such that the time-mean lift generated by the sail points upwind.

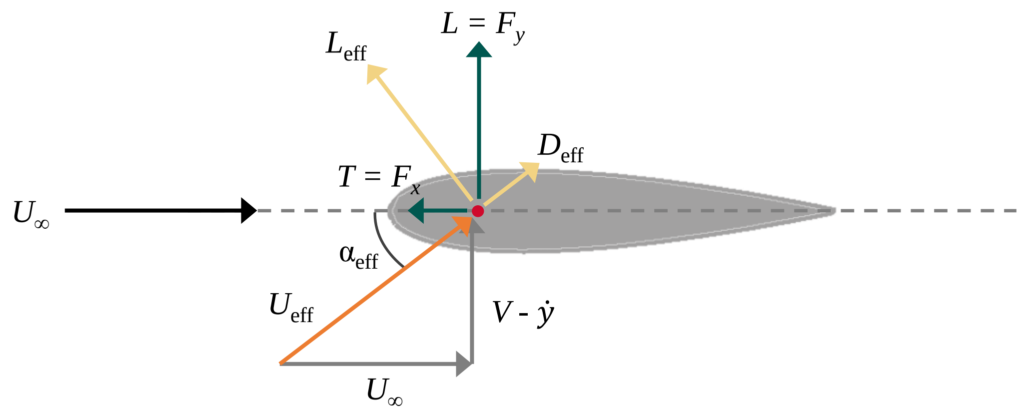

Figure 3.8: Free-body diagram showing effective flow direction, and resulting thrust vector in the direction of travel (the Katzmayr Effect).

Figure 3.8 shows a force diagram illustrating the Katzmayr Effect, as relevant to the present study. The mean flow direction, U_∞, is indicated by the black arrow upstream from the airfoil. In the illustration, the airfoil has a geometric angle of attack of 0°(it is pointing directly upstream). At the illustrated moment in time, the airfoil is experiencing an effective oncoming flow vector given by U_{eff}, shown as the orange arrow in the figure. The effective flow velocity is composed of the free-stream velocity U_∞ in the x-direction, and the expression $V - \dot{y}$ in the y-direction, or the vector subtraction of the airfoil's velocity from the y-velocity of the flow (denoted V). The effective angle of attack specific to the configuration shown in Figure 3.8 is given by

$$\alpha_{\text{eff}} = \text{atan}\left[\frac{V - \dot{y}}{U_\infty}\right], \tag{3.1}$$

since the airfoil is not moving in the x-direction. Considering a reference frame aligned with the effective flow direction (or equivalently, rotating the foil in Figure 3.8 clockwise through an angle α_{eff}), the situation is identical to that of an airfoil encountering an oncoming flow with some angle of attack α_{eff}. Thus, the airfoil generates lift perpendicular to the oncoming flow direction, as well as some amount of drag in the direction of the oncoming flow: these values L_{eff} and D_{eff} are shown in yellow in Figure 3.8. However, in this case, there is a mismatch between the oncoming flow direction and the desired direction of travel, here in the direction of U_∞. Thus, to recover the 'true' lift and drag vectors in the original reference frame, we must project the force vector generated by the foil relative to the effective wind direction onto the geometric coordinates, x, y. This results in the green force vectors in the Figure, which represent the relevant lift and thrust experienced by the airfoil.

From the Figure, we see that the thrust vector (net positive x-direction force) generated by U_{eff} is given by

$$T = F_x = L_{\text{eff}} \sin(\alpha_{\text{eff}}) - D_{\text{eff}} \cos(\alpha_{\text{eff}}). \tag{3.2}$$

Non-dimensionalizing, we can equivalently write

$$C_{T,\,\text{Katz}} = C_{L,\text{eff}} \sin(\alpha_{\text{eff}}) - C_{D,\text{eff}} \cos(\alpha_{\text{eff}}). \tag{3.3}$$

We can use this expression to make highly simplified predictions of the Katzmayr Thrust $C_{T,\,\text{Katz}}$ that the airfoil is experiencing during the Driven Airfoil Experiments in this chapter by approximating the lift and drag coefficients for the NACA 0018 profile used in this study. There is much available literature regarding such coefficients; one recent review by Melani et al. (2019) includes data for the Reynolds number range of interest in the present study. Based on their findings, we can reasonably approximate the lift coefficient of the airfoil using the classical thin-airfoil theory result (which is developed in detail by Anderson (2011) for example), $C_{L,\text{eff}} = 2\pi\alpha_{\text{eff}}$. Also from that study, we see that the drag coefficient C_D is relatively constant over the range of angles of attack relevant here ($\alpha_{\text{eff}} \in [-10, 10]$), so we choose a fixed value $C_{D,\text{eff}} = 0.025$. Then, starting from Equation 3.3, we make a small angle of attack approximation to write

$$C_{T,\,\text{Katz}} \approx 2\pi\alpha_{\text{eff}}^2 - C_{D,\text{eff}}. \tag{3.4}$$

Figure 3.9 shows the approximation for the Katzmayr thrust coefficient given by Equation 3.4 overlaid on the measured phase-averaged value of C_T from the present experiments. As expected, there is a two-peak structure in the Katzmayr thrust, and the mean thrust value is positive, $\overline{C_{T,\,\text{Katz}}} = 0.024$. In addition, for the first peak near the phase location $\phi/2\pi = 0.25$, there is very good agreement between the measured value of C_T and the estimated Katzmayr thrust coefficient. This implies that the thrust experienced by the airfoil at this point in the cycle can be largely attributed to the Katzmayr Effect, developed using results from steady thin airfoil theory. Moreover, unsteady effects from airfoil motion, dynamic stall, and impinging vortices do not appear to be playing a major role in the thrust production for this phase location, despite their dominant presence in the surrounding flow field. This simple, quasi-steady mechanism was hypothesized by Liao et al. (2003) to be

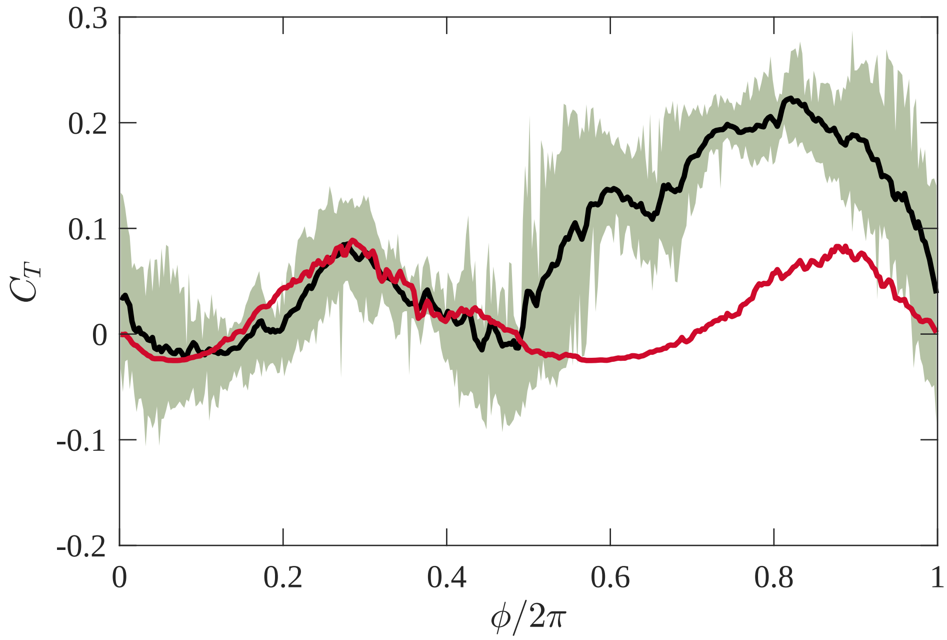

Figure 3.9: Comparison between estimated Katzmayr thrust coefficient and measured thrust coefficient over one vortex shedding cycle. ▬ Measured phase-averaged thrust coefficient C_T; and ▬ Estimated Katzmayr thrust coefficient $C_{T,\,\mathrm{Katz}}$, based on Equation 3.4. Shaded regions show interquartile range of phase averaged measurement data.

responsible for the behaviour of fish entraining in the wake of a circular cylinder, and later by Beal et al. (2006) to explain the passive thrust phenomenon they observed for a system similar to the present study.

Interestingly, agreement between Katzmayr thrust and measured thrust is much less satisfactory over the latter half of the cycle. The cause of the large degree of asymmetry in the system is unknown; however the disparity between Katzmayr and measured thrust implies that unsteady effects may be playing a more significant positive role in the thrust production for these moments in phase. For example, enhanced lift due to unsteady effects beyond the simple thin airfoil theory prediction would lead to enhanced thrust. An alternative interpretation, based on the hypothesis presented in Chapter 2 that α_{eff} may be underestimated in the present study due to the method of calculation of V from PIV fields, is that the true contribution of Katzmayr thrust is actually larger at all phases than that shown in Figure 3.9. Then at $\phi/2\pi \approx 0.25$, unsteady effects are acting to increase drag on the airfoil, which leads to a reduction in the observed thrust relative to the Katzmayr case. The reality is likely some combination of the two, with intermittent negative effects reducing the height of the first observed peak and intermittent positive ones enhancing the second. Asymmetry in the occurrence and strength of such unsteady effects as a function of phase could be attributed to asymmetry in the shape of the foil itself due to a slight spanwise warp, or an offset in its mean position downstream of the cylinder, as discussed in Chapter 2. The mean position offset for these tests is estimated to be approximately +6 mm (5% D), estimated based on analysis of PIV images and the shape of the momentum deficit region for the cylinder. This is significant relative to the small oscillation amplitude of the airfoil, and could influence interactions with oncoming vorticity despite remaining small relative to the size of oncoming vortices. Unfortunately, without access to PIV data on both sides of the airfoil (due to the airfoil shadow occluding one side in all tests), it is difficult to confirm the extent to which this induces asymmetry in the airfoil behaviour.

3.7 Phase-Averaged Interactions with Upstream Vorticity

To directly connect the observed flow field to the phase-averaged thrust and power production discussed in Sections 3.5 and 3.6, it is illuminating to consider the evolution of the flow field as a function of upstream vortex shedding phase. Figure 3.10 provides phase-average snapshots of the Γ_2 Criterion, a proxy for vorticity as discussed in Chapter 2, at five different vortex shedding phases as indicated by the dashed lines in the left-hand panels. Vortex shedding phase advances from top to bottom in the images on the right in the figure, with the corresponding moment in phase from the left-hand column indicated by the line type of the panel border.

The top-right panel in Figure 3.10 shows the flow at the moment corresponding to the first of two power production peaks observed in the cycle. The airfoil is passing through its neutral position behind the cylinder, and experiencing its maximum upwards velocity. It is centered in a region of upwards flow generated by upstream vortex shedding, which can be more easily seen in Figure 3.11 showing V, the y-direction flow velocity for the same moments in phase. Flow over the top of the airfoil is relatively well-ordered, though a region of CW rotating (negative signed) vorticity can be seen near the leading edge, partially occluded by masking. In addition, a large CW rotating cylinder vortex, which generates a low-pressure region (discussed for example by Liao et al. (2003)), has just passed over the airfoil.

As discussed in Chapter 1, in a study of energy extraction by a flapping foil in a uniform oncoming flow Kinsey and Dumas (2008) noted that properly timed formation and shedding of Dynamic Stall Vortices (DSVs) due to the airfoil's

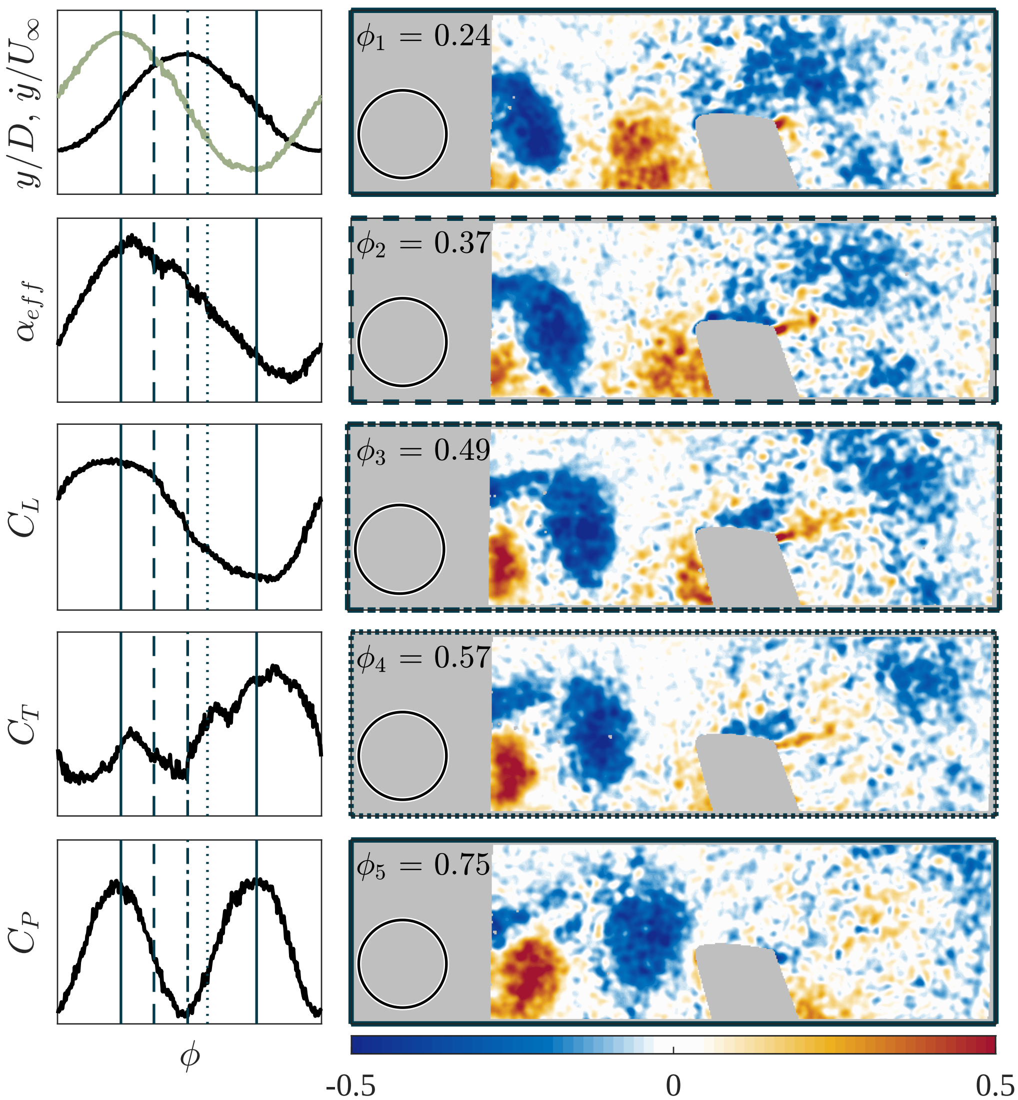

Figure 3.10: Snapshots of phase-averaged Γ_2 Criterion values at indicated points in the vortex shedding cycle. Colour bar indicates contour level for the criterion, which is constrained so that $|\Gamma_2| \leq 1$. Left column reproduces data for phase-averaged quantities from Figure 3.7. In the top left panel: ▬ y/D; and ▬ $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angle ϕ is presented as a fraction of one cycle (2π). ▬ $\phi/2\pi = \phi_1 = 0.24$; ╍ $\phi/2\pi = \phi_2 = 0.37$; ▬ · $\phi/2\pi = \phi_3 = 0.49$; ··· $\phi/2\pi = \phi_4 = 0.57$; ▬ $\phi/2\pi = \phi_1 = 0.75$.

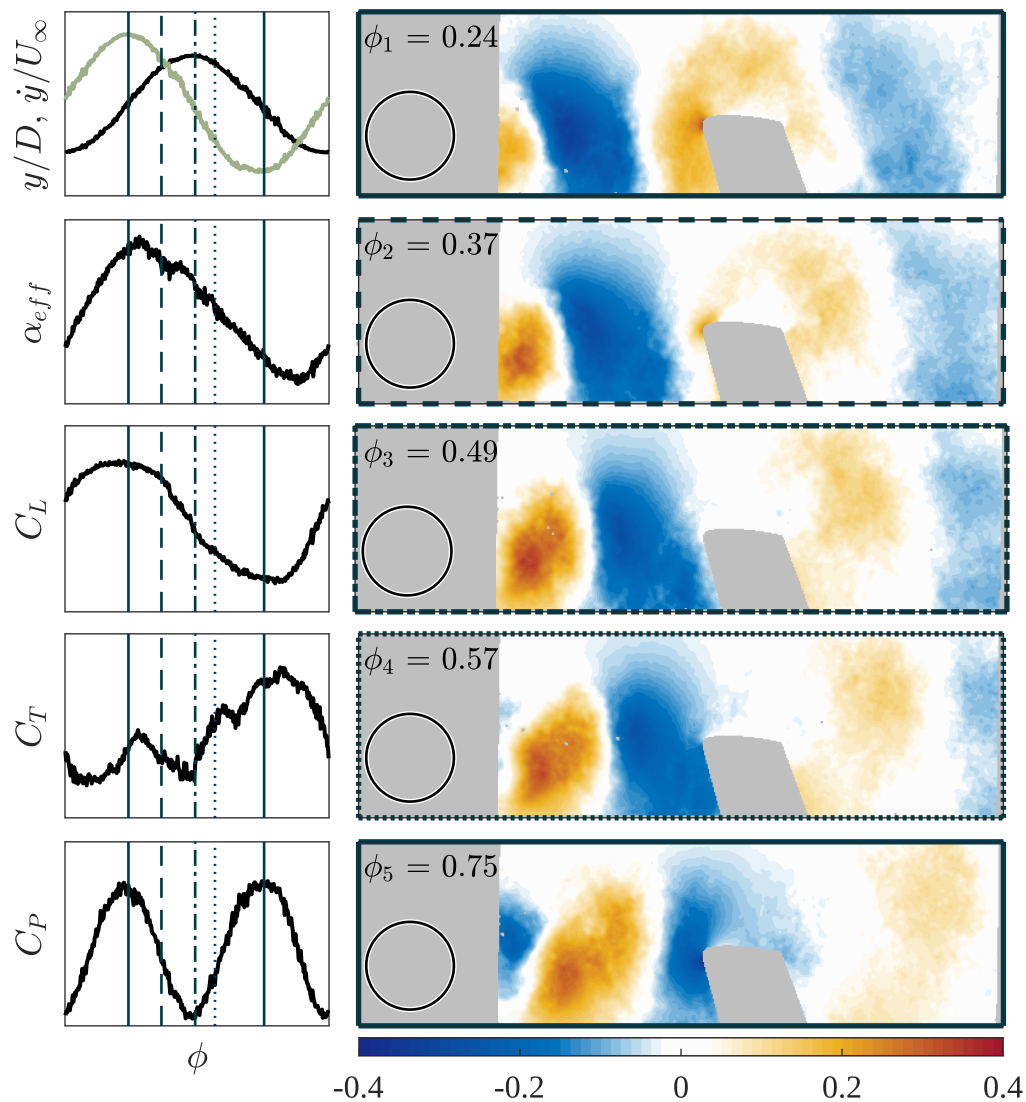

Figure 3.11: Snapshots of phase-averaged y-direction flow velocity (V) at indicated points in the vortex shedding cycle. Colour bar indicates flow speed in m/s. Left column reproduces data for phase-averaged quantities from Figure 3.7. In the top left panel: ▬ y/D; and ▬ $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angles ϕ_x are presented as a fraction of one cycle (2π). ▬ $\phi/2\pi = \phi_1 = 0.24$; ‒ ‒ $\phi/2\pi = \phi_2 = 0.37$; ‒ · $\phi/2\pi = \phi_3 = 0.49$; ··· $\phi/2\pi = \phi_4 = 0.57$; ▬ $\phi/2\pi = \phi_1 = 0.75$.

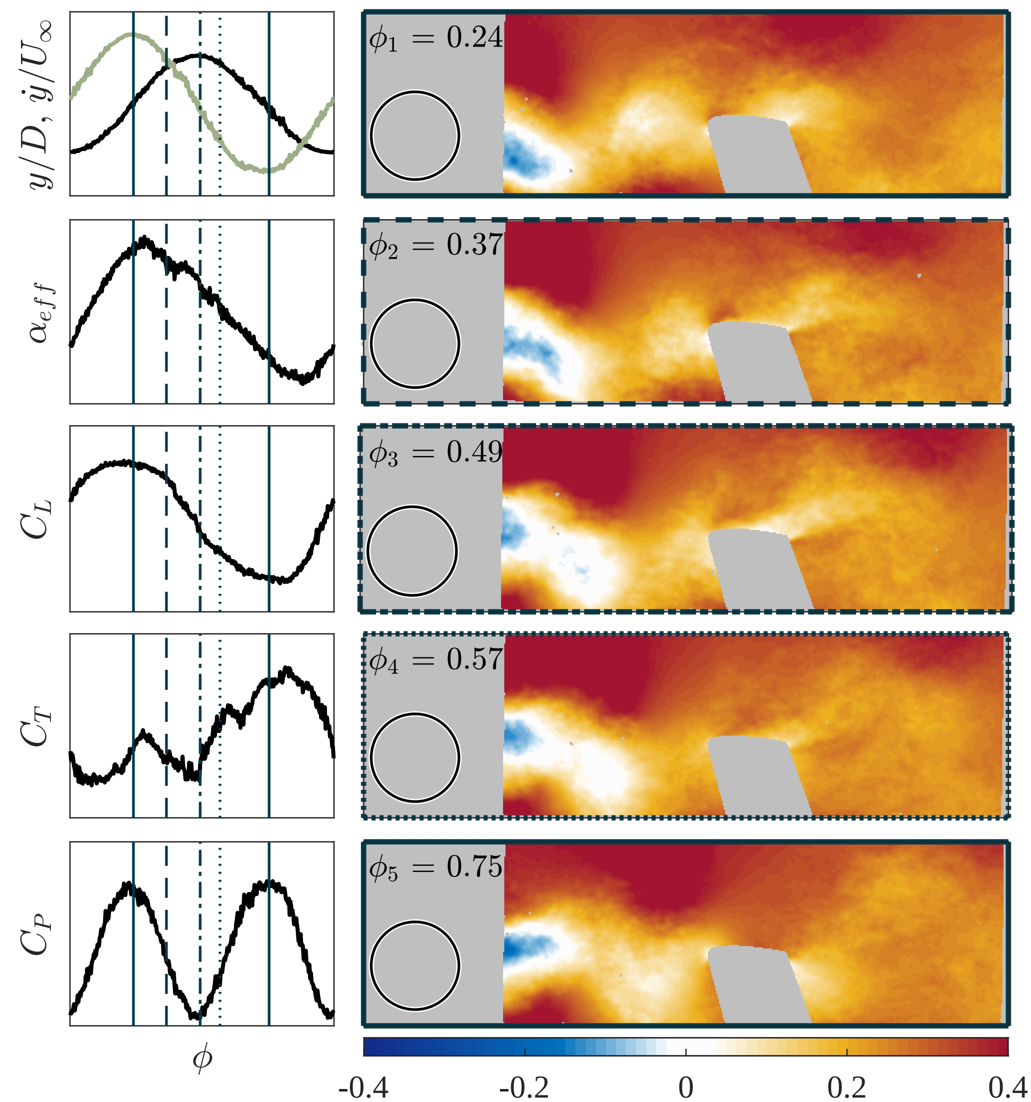

Figure 3.12: Snapshots of phase-averaged x-direction flow velocity (U) at indicated points in the vortex shedding cycle. Colour bar indicates flow speed in m/s. Left column reproduces data for phase-averaged quantities from Figure 3.7. In the top left panel: y/D; and $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angles ϕ_x are presented as a fraction of one cycle (2π). $\phi/2\pi = \phi_1 = 0.24$; $\phi/2\pi = \phi_2 = 0.37$; $\phi/2\pi = \phi_3 = 0.49$; $\phi/2\pi = \phi_4 = 0.57$; $\phi/2\pi = \phi_1 = 0.75$.

motion was critical for efficient power extraction. They observed that the low-pressure region created above the airfoil as a Leading-Edge Vortex (LEV) was formed and shed kept the lift experienced by the airfoil high over a larger portion of each cycle, in alignment with its flapping velocity. In the present case, where free-stream vorticity is dominated by large cylinder vortices and motion amplitudes and effective angles of attack for the airfoil are relatively small, this role appears to be played by oncoming cylinder vorticity. The passage of the cylinder vortex over top of the airfoil similarly creates a low pressure zone that enhances the lift value, as well as the portion of the cycle where the airfoil experiences high lift, contributing to overall power production.

In general thrust production for this system stems from the Katzmayr effect, or the exploitation of the misalignment between the direction of travel and the direction of the oncoming flow through the effective angle of attack α_{eff} (discussed at length in Section 3.6). We therefore note that the peak in the thrust lags slightly behind the peaks in lift and power in Figure 3.10, corresponding to a similar lag observed in the value of α_{eff}. This indicates a strong dependence on the oncoming flow direction, which depends on the location of the cylinder vortices (*global* flow conditions), rather than the local conditions surrounding the airfoil itself. In the top right panel of Figure 3.10, we see that the airfoil is located downstream and above a counter-clockwise rotating vortex, but upstream and below a clockwise rotating vortex and is therefore experiencing a large effective angle of attack due to the large positive y-direction flow velocity V (pictured in Figure 3.11). In addition, both of these structures contribute to upstream forcing on the airfoil (though it is emphasized that they do *not* create a region of reverse flow near the airfoil).

The second right-hand panel in Figure 3.10 shows a moment midway between the peak and trough in the observed power production, and at the beginning of a trough in the thrust. There is evidence of the growth of an LEV attached to the airfoil as well as deepening flow separation over the airfoil's surface, which has become easier to distinguish from cylinder-derived vorticity as the large CW rotating cylinder vortex has largely passed over the airfoil. The C_L value at this moment remains high, with a downwards slope beginning at phases just beyond: the growth of leading-edge vorticity may be helping to keep the lift high for the high-velocity portion of the cycle (in addition to the continuing contribution from the CW rotating vortex passing over) resulting in augmented power production. It is interesting to note that the phase-averaged effective angle of attack at this moment has not passed the static stall angle for a NACA 0018 section (approximately 10°–14°, based on recent work by Melani et al. (2019) and Le Fouest et al. (2021)), despite the formation of a leading edge vortex and the apparent initiation of dynamic stall. This could arise from sensitivity to instantaneous flow conditions for the occurrence of these phenomena; it could be that some periods included in phase averaging have sufficient α_{eff} to initiate and sustain growth of a LEV, while others do not, Therefore, some mean version of two disparate flow situations is represented in the phase average. In addition, as noted in Chapter 2, the method used to determine α_{eff} could be leading to an under-prediction of its true value.

The center-right panel (or third instant in phase) in Figure 3.10 shows a moment of minimum power production, which coincides with a maximum in airfoil position and therefore a zero-point in the velocity. Since $P = F_y(t)\dot{y}(t)$, this velocity zero crossing is largely responsible for the dip in power output. At this moment, we see that the airfoil is perched between regions of upwards and downwards flow (engulfed in a CCW rotating vortex), and the flow over the airfoil's surface is deeply separated including near the leading edge. The thrust coefficient is near a local minimum, and has started to diverge from the estimated thrust based on the Katzmayr Effect (see Figure 3.9). As discussed by Corke and Thomas (2015), the formation of an LEV leads to increased drag over the foil, which could be contributing negatively to the observed thrust at this location in the cycle.

In the fourth panel of Figure 3.10, the airfoil is experiencing zero effective angle of attack, as the CCW rotating (positive signed) vortex passes underneath it. At this point in the cycle, the lift is already slightly negative, as the low-pressure region induced by the CCW vortex (largely occluded by the airfoil shadow) starts to exert influence on the foil in excess of any quasi-steady lift. The airfoil has also passed through its maximum position extent and has started to move downwards in the frame, leading to increasing positive power experienced at this instant. There is evidence that the leading edge vortex and separation-induced vorticity that formed on the suction side of the airfoil as it was moving upward has been shed, and flow reattachment is in progress. This is mediated by the advancement of a region of strong downwards-flowing fluid towards and over top of the leading edge (seen in the corresponding panel in Figure 3.11), which promotes the development of a favourable pressure gradient on the new pressure (top) side of the airfoil for the latter half of the cycle. The shedding of the attached vorticity from the leading edge could be contributing to the recovery of the observed thrust coefficient, which is now rising in the latter half of the cycle ahead of the Katzmayr thrust, which largely follows the effective angle of attack.

Finally, in the bottom (fifth) panel of Figure 3.10, the airfoil is again producing maximum lift and power as it passes through its neutral position heading down; this allows us to view the pressure side of the airfoil at a moment which is equivalent to that shown in the first panel ($\phi/2\pi = 0.24$), but with the directions of lift, velocity, α_{eff}, and vorticity reversed. There is no remaining evidence of flow separation on the top, now pressure side of the airfoil, as at this instant there is fast downwards-moving fluid passing over the foil. The airfoil is near its maximum thrust production over the whole cycle; however the occlusion of the suction side of the airfoil makes it challenging to identify fluid dynamic factors which could be contributing to the strong asymmetry in thrust peak heights between the first and second halves of each cycle.

The net effect on thrust production due to dynamic stall or other unsteady effects leading to separation of the flow near the foil's surface represents an interesting balance between thrust-producing and drag-producing interactions. We expect large lift values and large effective angles of attack to contribute positively to experienced Katzmayr thrust (based on Equation 3.3, in the previous section); however, these unsteady effects precipitously increase the drag, which acts to pull the net direction of the experienced aerodynamic force downstream. Since the effect of drag depends on the cosine of α_{eff} while lift contributions depend on the sine, we expect that unsteady events such as dynamic stall have a net negative effect on the thrust production for the airfoil, though further investigation of this topic is warranted.

Another interesting feature visible in Figure 3.10 is the formation and shedding of Trailing-Edge (TE) vorticity, having a positive sign in the first four panels, and a negative one in the final panel. TE vorticity is primarily shed from the pressure side of the airfoil as a cylinder vortex approaches the leading edge. The sign of the TE vorticity is the same as the new impinging vortex, and opposes the sign of the previous cylinder vortex; this is seen for example in the first panel, where a positive signed (CCW) Trailing Edge Vortex (TEV) is being shed, and interacting with the large CW cylinder vortex which has advanced into that region. As the power coefficient, velocity, and lift drop, the strength of the TE vorticity increases, fed by the advancing upstream cylinder vortex. The maximum strength of the TE vorticity occurs in the vicinity of minimum power production, as the airfoil's surface(s) are transitioning between suction and pressure (or equivalently, as the lift vector changes direction).

In the latter half of the cycle, shown in the $\phi/2\pi = 0.57$ and $\phi/2\pi = 0.75$ panels, positive-signed trailing-edge vorticity continues to form as the airfoil begins to move downward with the surrounding flow; however the strength of the shed vorticity is quickly attenuated. In the $\phi/2\pi = 0.57$ panel, we see that airfoil-derived CW vorticity formed during the upward portion of the airfoil's motion (the first three panels) has detached from the leading edge and convected downstream. As a region of downwards velocity approaches, there is evidence that this separation-derived vorticity is pushed along the airfoil surface towards the trailing edge, where it merges with negative-signed TE vorticity beginning to be shed in the latter half of the cycle due to the oncoming CW cylinder vortex.

In addition to Figures 3.10 and 3.11 referenced throughout this discussion, Figure 3.12 shows phase averaged U (x-direction velocity) fields at the same moments in phase.

3.8 Time Evolution of the Combined Airfoil-Cylinder Wake

Using the Frozen Flow Visualization method described in Chapter 2, the combined airfoil-cylinder wake can be qualitatively visualized as a function of time. Figure 3.13 shows the phase-averaged wake behind the circular cylinder with no airfoil, as well as the phase-averaged combined airfoil-cylinder wake, both repeated over several periods. Using this visualization technique, interactions between the cylinder wake and airfoil-derived vorticity are illuminated.

Classical von Kármán shedding from behind the cylinder has a characteristic 2S pattern. This is clearly illustrated by the alternating pattern of CW and CCW vortices shed in the cylinder-only case in the top panel of Figure 3.13, and in the region upstream of the airfoil in the combined case (the bottom panel of the Figure). These vortices are considered two single (S) vortices rather than one pair (P) of vortices per cycle, since each one could be considered 'paired' with another either upstream or downstream in the wake. By contrast for the combined wake downstream of the airfoil, pictured in the bottom panel of the Figure, there are four distinct, dynamically relevant vortices illustrated per cycle: the oncoming Cylinder Vortices, CV+ (with a CCW orientation) and CV- (with a CW orientation), as well as Trailing Edge Vorticity shed from the airfoil, TEV+ and TEV-.

Starting from the left hand side in the bottom panel of Figure 3.13, oncoming cylinder vortices from the 2S wake are repositioned by the existence of the foil.

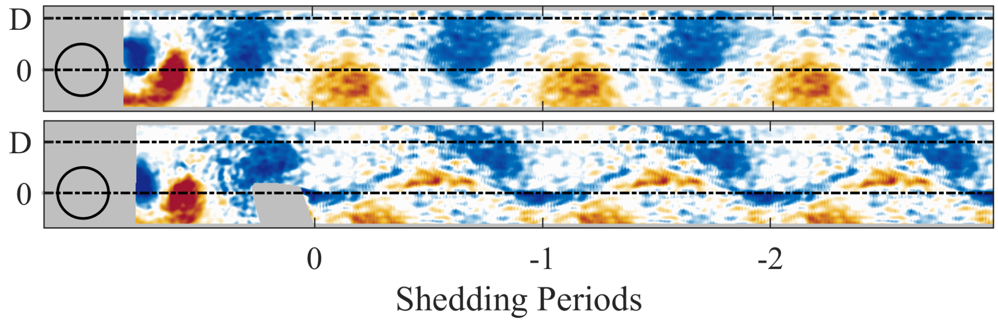

Figure 3.13: Frozen flow visualizations of both the cylinder-only wake (top) and the combined airfoil-cylinder wake (bottom), for comparison. Visualizations were created from phase-averaged Γ_2 Criterion fields using the method described in Chapter 2, and have the same contour levels. The top panel shows cylinder-only frozen flow from the same spatial location behind the cylinder as in the bottom panel, showing the combined airfoil-cylinder wake. Dot-dashed lines indicate the locations $y = 0$ and $y = D$ in each panel, as there is a vertical shift in the frame position between these data sets. Slight differences in the underlying spatial fields for these data sets lead to the very small misalignment of the x-axes apparent in the figure.

A CV- vortex has been pulled upwards away from its usual path to pass over top of the airfoil, as evidenced by the bulk of CW vorticity at approximately $0.25T$ that has passed outside of the line at $y = D$ compared to the cylinder only case. There is also evidence that the cylinder vortex may have been split, with some CW vorticity passing underneath the foil. Such vortex repositioning due to airfoil suction is discussed by Gopalkrishnan et al. (1994) for an analogous system operating at a variety of phases relative to vortex shedding. The CV- vortex over top of the airfoil then generates a Trailing Edge Vortex (TEV) of the same orientation that is shed into the flow behind the airfoil, as discussed with reference to flow snapshots in the previous section. This vortex, TEV-, then pairs up with a cylinder vortex of the opposite sign that has passed under the airfoil (CV+), and is now located slightly downstream, to form a counter-rotating pair. Moving away from the airfoil, backwards in time, similar pairs of counter-rotating CV and TEV-type vortices are visible, with alternating orientations due to cyclic vortex shedding. Thus, per cycle two pairs of counter-rotating vortices are formed giving rise to a 2P-type wake as recently visualized by Cros et al. (2018), for example. The existence of the airfoil and its motion through the space have repositioned and modified vortices from the oncoming 2S wake, as well as adding newly generated vorticity from the airfoil's surface to form a wake of type 2P instead. In addition to this topology change, the wake width also seems to have been expanded, as evidenced by the fraction of each large CW cylinder vortex located further than a distance D from the cylinder centerline in each panel of Figure 3.13. This observation is confirmed in the following section, by considering the mean flow velocity in the combined wake region.

It is illustrative to compare the formation mechanism(s) for the observed combined wake to a much simpler uniform flow case, as discussed by Schnipper et al. (2009). In that study, a detailed account of the formation of a 2P wake from a pitching foil in a uniform free stream was presented, showing that the critical parameter in the formation of a 2P wake is the timing between boundary layer vorticity roll-up and shedding vs shedding of a TEV. For the pitching airfoil, this timing is mediated by the forced pitching motion, and the dynamics in the boundary layer(s) themselves. When the frequency of pitching motion is outside of some particular values, the 2P pattern is not observed (Schnipper et al., 2009). In the present case, the timing of vortex shedding appears to be strongly linked to the passage of cylinder vortices. Instead of an intrinsic natural shedding frequency based on airfoil motion or boundary layer development, vortex shedding both from the surface of the airfoil (i.e., LEVs) as well as at the trailing edge (TEVs) appears to be strongly locked into the frequency of oncoming vortex shedding. It seems as though this frequency locking acts to enforce an interaction between oncoming vorticity and vorticity formed near the airfoil's surface that consistently leads to the formation of 2P structures.

3.9 Thrust and Power Production Effects on the Mean Airfoil-Cylinder Wake

For the driven airfoil experiments discussed throughout this chapter, both net power and net thrust are generated over one vortex shedding cycle. Therefore, we expect that the presence of the airfoil will induce changes to the oncoming cylinder wake, both by repositioning cylinder-derived vorticity, and by qualitatively changing the structure of the combined wake region. This was confirmed in the previous sections, which describe time-dependent interactions between the airfoil and oncoming vorticity, as well as the resulting vortical structures formed in the near-airfoil region. To make a quantitative link between the thrust production and power extraction of the airfoil and the resulting mean combined airfoil-cylinder wake is more challenging however, due to several complicating factors present in this system.

Firstly, due to the relatively high Reynolds number, wake vorticity quickly becomes disorganized downstream of the airfoil, and the finite size of the field of view means that vorticity is convected out of the frame. This makes it challenging to qualitatively view the wake structures without the aid of frozen-flow visualizations (discussed in detail in Chapter 2), which provide a temporal rather than spatial picture of wake evolution, and therefore illuminate wake behaviour only in a very limited spatial region. Secondly, the upstream circular cylinder induces its own drag wake in the region near the airfoil, which forms a backdrop upon which the action of the airfoil is added; since the presence of the airfoil has the potential to alter the drag characteristics of the cylinder itself, determining wake effects due to the airfoil's presence may not solely reflect the action of thrust and/or drag on the foil only. Lastly, the foil extracts net power from the flow while simultaneously producing thrust. Since energy is extracted from the flow by the foil, we expect that the combined wake region will be de-energized relative to the cylinder case, in direct opposition to our expectation of jet formation (energized flow) in the near wake due to the production of thrust.

To begin to untangle these effects, Figure 3.14 shows mean profiles of the x-direction velocity, $\bar{U}$ for one station upstream of the airfoil, and three stations downstream. For comparison, profiles at the same locations in the flow with no airfoil present are also included. For the first station located at $-1.1D$ from the cylinder trailing edge (ahead of the airfoil location), the profiles correspond quite closely, showing that the airfoil's presence has a limited effect on the mean flow at this location. The slightly larger velocity deficit for the airfoil case could be the result of a small deviation between sets of data in the free-stream velocity, U_∞. This value was assumed to be 0.32 m/s for all data collected, but in reality varied slightly from day to day in experiments due to small variations in tunnel fill level.

By contrast, we see that the airfoil's presence has a strong impact on the observed downstream profiles. Close to the airfoil's trailing edge at approximately $-4.1D$, we see that instead of the parabolic-type profile observed for the cylinder case, the mean velocity has two troughs surrounding a central peak (jet) region, which occurs near the mean y-position of the airfoil. This energized jet region is a result of the shedding of trailing edge vorticity, which appears in alternating 'stripes' behind the airfoil in the wake visualization for the Driven Airfoil Experiments in Figure 3.13. As large cylinder vortices encounter the airfoil, they generate high-velocity flow from the airfoil's trailing edge, some component of which points in the downstream direction. This can be seen in Figure 3.12, where regions of energized U are visible near the airfoil's trailing edge particularly in panels 2 and 3. The two troughs in U

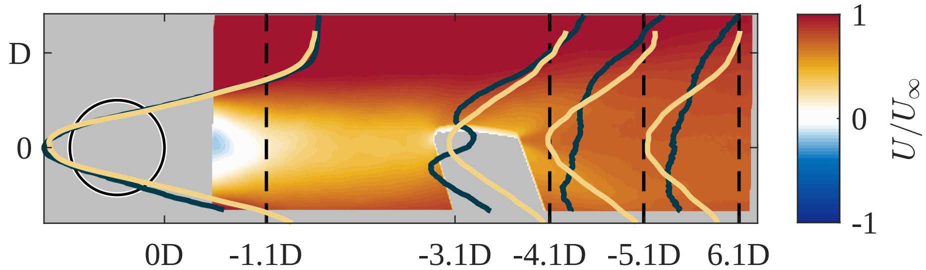

Figure 3.14: Mean x-direction velocity ($\bar{U}$) field for the Driven Airfoil Experiments, with y-direction profiles shown at 3 stations. Blue lines (—) show $\bar{U}$ for the Driven Airfoil Experiments, while yellow lines (—) show $\bar{U}$ for Cylinder-Only experiments (discussed in Section 3.2) at the same locations in x, y space for reference. Velocity profiles are shown with U_∞ subtracted, so that portions of profiles to the right of their corresponding black dashed lines indicate velocities larger than U_∞.

pictured above and below this jet region (such that the velocity deficit at the trough locations is larger than the corresponding cylinder-only profile) could be a result of the formation of LEVs or other flow separation near the airfoil's surface, which is then shed into the flow above and below the airfoil, and convected downstream. It is interesting to note the asymmetry in the deficit magnitude above and below the airfoil - this could be linked to the asymmetry in thrust behaviour observed over one vortex shedding cycle. We also note that the velocity deficit region appears to have been expanded in the y-direction, lending support to the observation that cylinder vortices are repositioned farther away from the centerline by the presence of the airfoil.

At the farther downstream stations, located at $-5.1D$ and $-6.1D$, we see that although the effect of airfoil-derived vorticity on the U-velocity has been attenuated, the overall profiles of $\bar{U}$ remain wider in the y-direction than in the cylinder-only case. This implies that the expansion of the cylinder wake by the airfoil persists even far downstream of the foil, consistent with the idea of an expanding wake discussed in the previous section.

Comparing again to the simpler uniform free-stream case, for a symmetric foil undergoing heave-only motion Andersen et al. (2017) found that 2P-type wakes were observed in both thrust and drag producing regimes. However, the unifying feature that they observed in cases where thrust was produced (in agreement with many previous studies) was that at 4 chord lengths downstream of the trailing edge, the mean flow in the streamwise direction was positive, or in the mean the airfoil's wake formed a jet. Thrust-producing 2P wakes were shown to have energetic regions away from the centerline that created jet-like flow strong enough to overcome a region of net drag directly behind the airfoil, which leads to a positive mean momentum flux. It is interesting to note the emergence of an opposite pattern here, linked to the upstream orientation of the generated pairs of counter-rotating vortices, most clearly seen in Figure 3.15 in the following section. Computing mean momentum flux at each station in Figure 3.14, we also find that when the airfoil is present, the mean momentum flux is slightly reduced compared to that for the cylinder-only case, although this calculation is highly sensitive to the mean flow velocity on a particular day of testing. This is consistent with the idea that the foil extracts net energy from the flow, or that not all of the available momentum harvested from oncoming vorticity is used to generate net thrust.

3.10 Summary of Airfoil Interactions with Oncoming Vorticity

To summarize results from the Driven Airfoil Experiments presented in the previous sections, Figure 3.15 shows an idealized picture of interactions between cylinder-derived and airfoil-derived vorticity, as well as idealized airfoil behaviour. For the Driven Airfoil Experiments, the position and velocity of the airfoil are fixed a priori (controlled by the CTS): these can be considered 'inputs' to the system, in addition to the experienced y-force which is determined by vortex shedding from the upstream cylinder. In the figure, the y-force is idealized as sinusoidal with the same frequency and phase as the airfoil's velocity, an assumption that has proven to largely represent reality for the experiments described throughout the previous sections. This assumption directly leads to an idealized value of C_P with sine-squared character, as P is a direct product of force and velocity. For the Driven Airfoil Experiments, the power and thrust produced can be conceptualized as an 'outputs' from our system. This input-output framework is useful for describing the behaviour of the airfoil in the Passive Captive Airfoil Experiments discussed in the following chapters.

Based on the discussion of thrust production for this system due to the Katzmayr Effect in Section 3.6, we postulate that the idealized thrust production also exhibits sine-squared character, since it depends approximately on the square of the effective angle of attack through Equation 3.4, and α_{eff} is itself derived from cyclic flow velocities due to vortex shedding. The link between thrust production and fluid-structure interactions is less clear than in the case of power; however observed thrust coefficients do to some extent reflect this simplified underlying pattern. It is interesting to note that α_{eff} (and therefore C_T) appears to have a slightly different

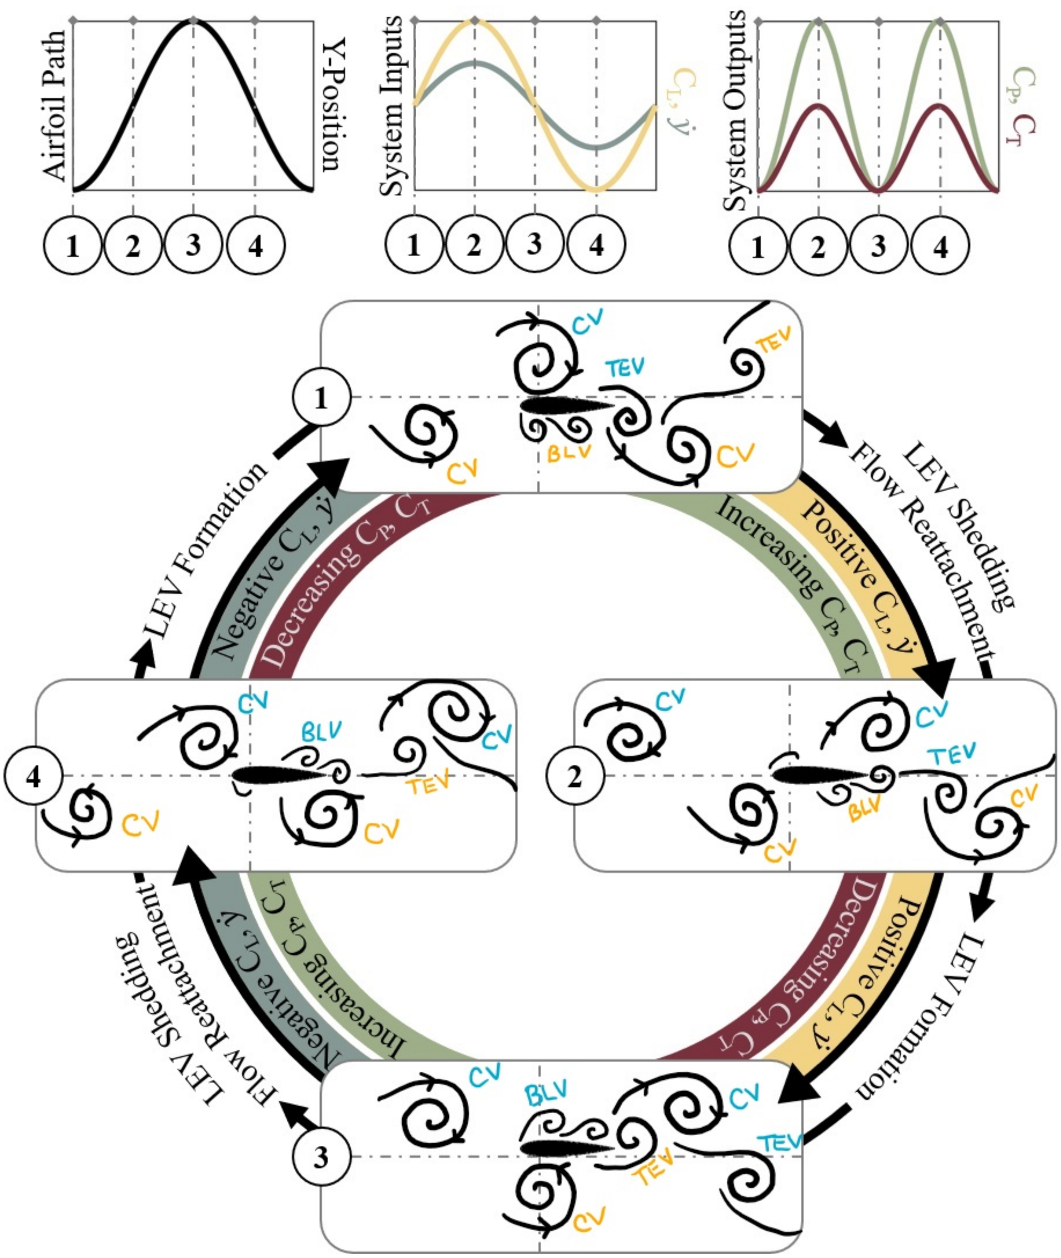

Figure 3.15: Summary of interactions between the airfoil and vorticity shed by the upstream circular cylinder over one vortex shedding cycle. Top three panels show idealized airfoil position y (top left), C_L, $\dot{y}$ (top middle), and C_P, C_T (top right). Trends in these quantities are indicated in the concentric circles in the bottom panel, between indicated phase locations 1-4. Images in the bottom panel show a simplification of both cylinder and airfoil-derived vorticity and their approximate interactions at each phase location in one vortex shedding cycle. Vortices are labelled according to their origin and sign of rotation. CV: Cylinder Vortex; TEV: Trailing-Edge Vortex; BLV: Boundary-Layer or other Near-Airfoil Vortex (including LEVs, and any vorticity due to flow separation over the airfoil's surface). ▬ Negative (CW) vorticity; ▬ Positive (CCW) vorticity. In addition, cylinder vortices have arrows indicating sign of vorticity.

phase relative to oncoming shedding than that of the observed lift and commanded velocity. This shift reflects an offset between the phase angle giving rise to the highest lift magnitude (approximately $\phi/2\pi = 0.25, 0.75$), and the phase where the airfoil is centered in a region of upwards or downwards flowing fluid. Though noticeable for example in Figure 3.11 showing the y-direction flow velocity (V), the shift is small and it is neglected in Figure 3.15.

The bottom portion of Figure 3.15 shows a simplified picture of the formation mechanism for the 2P wake generated by the airfoil, discussed in Section 3.8 and 3.9. We first consider phase instant 1 in Figure 3.15, when the airfoil is at its minimum position in the frame. A large CW cylinder vortex is located directly over top of the foil, which is is causing a TEV of the same sign to be shed into the flow downstream. In addition, there is a region of separated flow on the foil's bottom surface (labelled BLV in the Figure). In the wake, trailing edge vorticity is pairing up with a CCW cylinder vortex located downstream and below the foil.

At instant 2, the oncoming vorticity has advanced through a quarter cycle, such that the large cylinder vortex over top the airfoil in the previous instant is now located just past the trailing edge. The advancement of the upstream CCW cylinder vortex towards the leading edge has energized the flow on the airfoil's bottom surface, causing previously shed BLV to detach from the leading edge and convect along the foil's surface towards the trailing edge, initiating flow reattachment. This shed BLV begins to coalesce into a TEV to be shed later in the cycle. The airfoil is moving upwards with its maximum velocity, and a region of flow separation has started to form on the top side near the leading edge. Any separation in the trailing edge region is difficult to distinguish from vorticity due to the large CW cylinder vortex downstream. In the wake, the TEV shed at the previous instant has detached from the airfoil, paired up with a CCW cylinder vortex, and the pair are now convecting downstream and away from the centerline, forming the 2P-type wake.

At instant 3, a moment representing the dual of instant 1 but where airfoil motion, forces and vorticity are reversed, flow separation over the top surface of the airfoil which was initiated at the previous instant has deepened, with airfoil-derived vorticity remaining localized near the top surface. Cylinder vorticity has advanced such that the airfoil is directly over top of a CCW cylinder vortex and at its maximum position extent. The presence of the CCW cylinder vortex has caused BLV shed at the previous instant to coalesce, and has further fed the formation of a CCW TEV. This TEV is in the process of pairing with the downstream CW cylinder vortex to form the next 2P-type wake structure.

Finally, at instant 4, the dual of instant 2, we see that the advancement of the oncoming CW cylinder vortex above the airfoil has caused BLV from the top surface of the airfoil to be shed towards the trailing edge as the flow reattaches to the top side. This BLV begins to coalesce into the CW TEV seen in the next instant 1. The previously shed CCW TEV has paired with the downstream CW cylinder vortex to form a 2P wake structure, which convects downstream and away from the centerline. This is the second pair of vortices shed per cycle, constituting a 2P wake.

Based on the relationship between the airfoil's position and the passage of oncoming cylinder vorticity, we see that the airfoil moves to avoid oncoming vortex cores, achieving position maxima as cores are passing through the airfoil's x-location. This appears to correspond to a *Slaloming Mode* of interaction as identified by Beal et al. (2006). In agreement with the current study, the slaloming mode of interaction was found to give rise to pairs of counter-rotating vortices in the combined wake region, as well as high observed propulsive efficiency (Beal et al., 2006). Though conceptually the propulsive efficiency of our current system is very high, as it extracts thrust without requiring energy input at all and in fact extracts net power, under the conventional definition it would be less than zero ($\eta < 0$). Therefore, this is not a particularly appropriate performance metric for hybrid propulsive/energy extracting systems.

The principal difference between similar past studies and the current experiments, and the key factor in distorting efficiency as a metric for this system, is twofold. Firstly, in the present case the magnitude of the airfoil's velocity is much smaller than the maximum magnitude of flow velocity in the region surrounding it. Secondly, the flow and the foil's velocities are always aligned. These factors together lead to the net power extraction from the flow, since the flow in the region of the airfoil puts energy into the airfoil's motion, similar to the action of pushing a child on a swing set. This leads to a qualitatively different interpretation of efficiency as a metric, since in the classical framework used to study propulsion, the computed efficiency for our system is always negative. Moreover, for systems of this type, the pole in the equation for efficiency when the required input power to generate thrust in a system passes through zero distorts the interpretation further. This was discussed more completely in Chapter 1.

Despite these challenges in interpretation of similar studies with larger foil amplitudes, it is interesting to compare the present vortex formation mechanism with those discussed in the pioneering work of Gopalkrishnan et al. (1994). In doing so, we see that the mode of interaction considered in the current study represents a different mechanism than those presented previously. Though those authors discuss an ‘Expanding Wake’ mode which has qualitative similarities to the observations made in this study (for example, it also leads to a combined wake of type 2P that convects away from the centerline), in their study vortex repositioning by the airfoil is a dominant mechanism, dragging cylinder vortices across the centerline to pair with previously shed airfoil-derived vorticity. In our study, it is the trailing edge vorticity which crosses the airfoil centerline to pair with downstream cylinder-derived vorticity. The extensive vortex repositioning is made possible by the much larger foil ($c = 2D$) used in that study, as well as the much larger transverse motion ($A_0 = 0.5D - 0.833D$). In the present case, the change in pressure field induced by the airfoil is not large enough to affect the trajectories of cylinder-derived vorticity in such an extreme way, and the extent of the wake expansion is correspondingly limited.

In addition, the Expanding Wake mode identified by Gopalkrishnan et al. (1994) appears to correspond most closely to the *Interception Mode* identified by Streitlien et al. (1996), though significant differences in assumptions and experimental and/or parametric frameworks between the the studies make a true apples-to-apples comparison challenging. In the work of Streitlien et al. (1996), the interception mode corresponds to the case where the airfoil encounters vortex cores head on, instead of weaving between them as in the avoidance/slaloming mode. This mode represents a simultaneous maximum in required input power to sustain motion, but results in good efficiency since the thrust produced is also large. In addition, this mode results in the expansion of the combined wake signature by the action of the airfoil. Although Gopalkrishnan et al. (1994) report variable efficiencies for experiments associated with the Expanding Wake mode (which they found challenging to reproduce, in contrast to the current study), the phase of vortex interaction appears similar, and in both cases an expanding wake signature is the principal feature.

If this interpretation is correct, the present study represents a region of phase space not visualized by Gopalkrishnan et al. (1994). The present experiments, corresponding to the slaloming mode of interaction, operate at a phase of interaction 180° out of phase with the Expanding Wake mode discussed in that study, and hold similarity with the low-power/avoidance mode identified by Streitlien et al. (1996).

3.11 Effect of Static Angle of Attack Offset on Airfoil Behaviour

In addition to the Driven Airfoil Experiments carried out in the Basic configuration (Case BA from Table 3.1) which were discussed at length in the previous sections, additional experiments were performed with a static Angle of Attack offset applied to the airfoil (Cases HAoA+ and HAoA- in Table 3.1). Changing the static angle of attack of the airfoil, referred to as α_0 throughout this thesis, significantly alters the thrust and power production behaviour of the foil.

Similar to Case BA, phase averaging was used to extract cyclic behaviour of the airfoil for static angle of attack offsets of 10° (Case HAoA+) and -10° (Case HAoA-). Figure 3.16 shows histograms of the number of frames averaged for each phase bin in Cases HAoA±. For both HAoA±, fewer frames are averaged per bin as compared to Case BA discussed previously. This is both because fewer experimental trials were completed for these cases as compared to Case BA (see the disparity in total recorded periods in Table 3.1), and for these high angle of attack offsets it was also more challenging to obtain an acceptable phase match between force and velocity. This resulted in fewer periods of those recorded being included in phase averages. Despite these difficulties, there are still a minimum of 10 frames averaged in each phase bin, with a large majority of bins containing more than 30 frames.

Figures 3.17 and 3.18 show phase averaged dynamic quantities of interest for Case

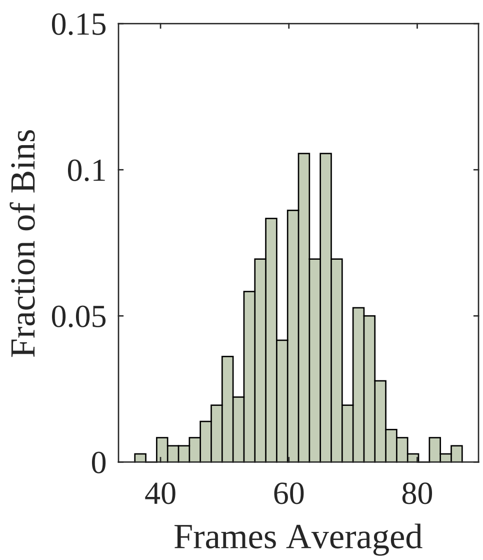

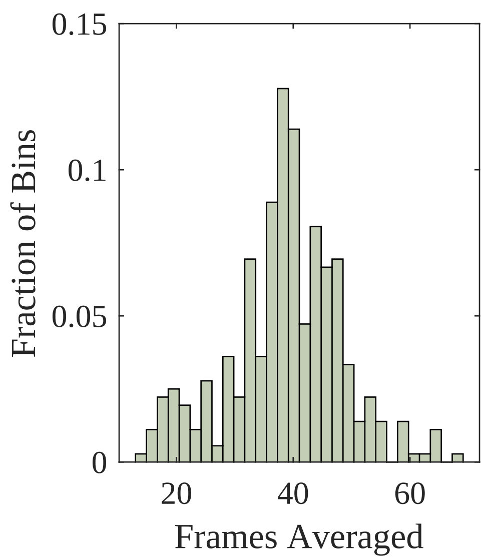

Figure 3.16: Histograms showing number of frames averaged per phase bin for the high static angle of attack (HAoA±) Driven Airfoil Experiments. Left: HAoA- ($\alpha_0 = -10°$). Right: HAoA+ ($\alpha_0 = 10°$).

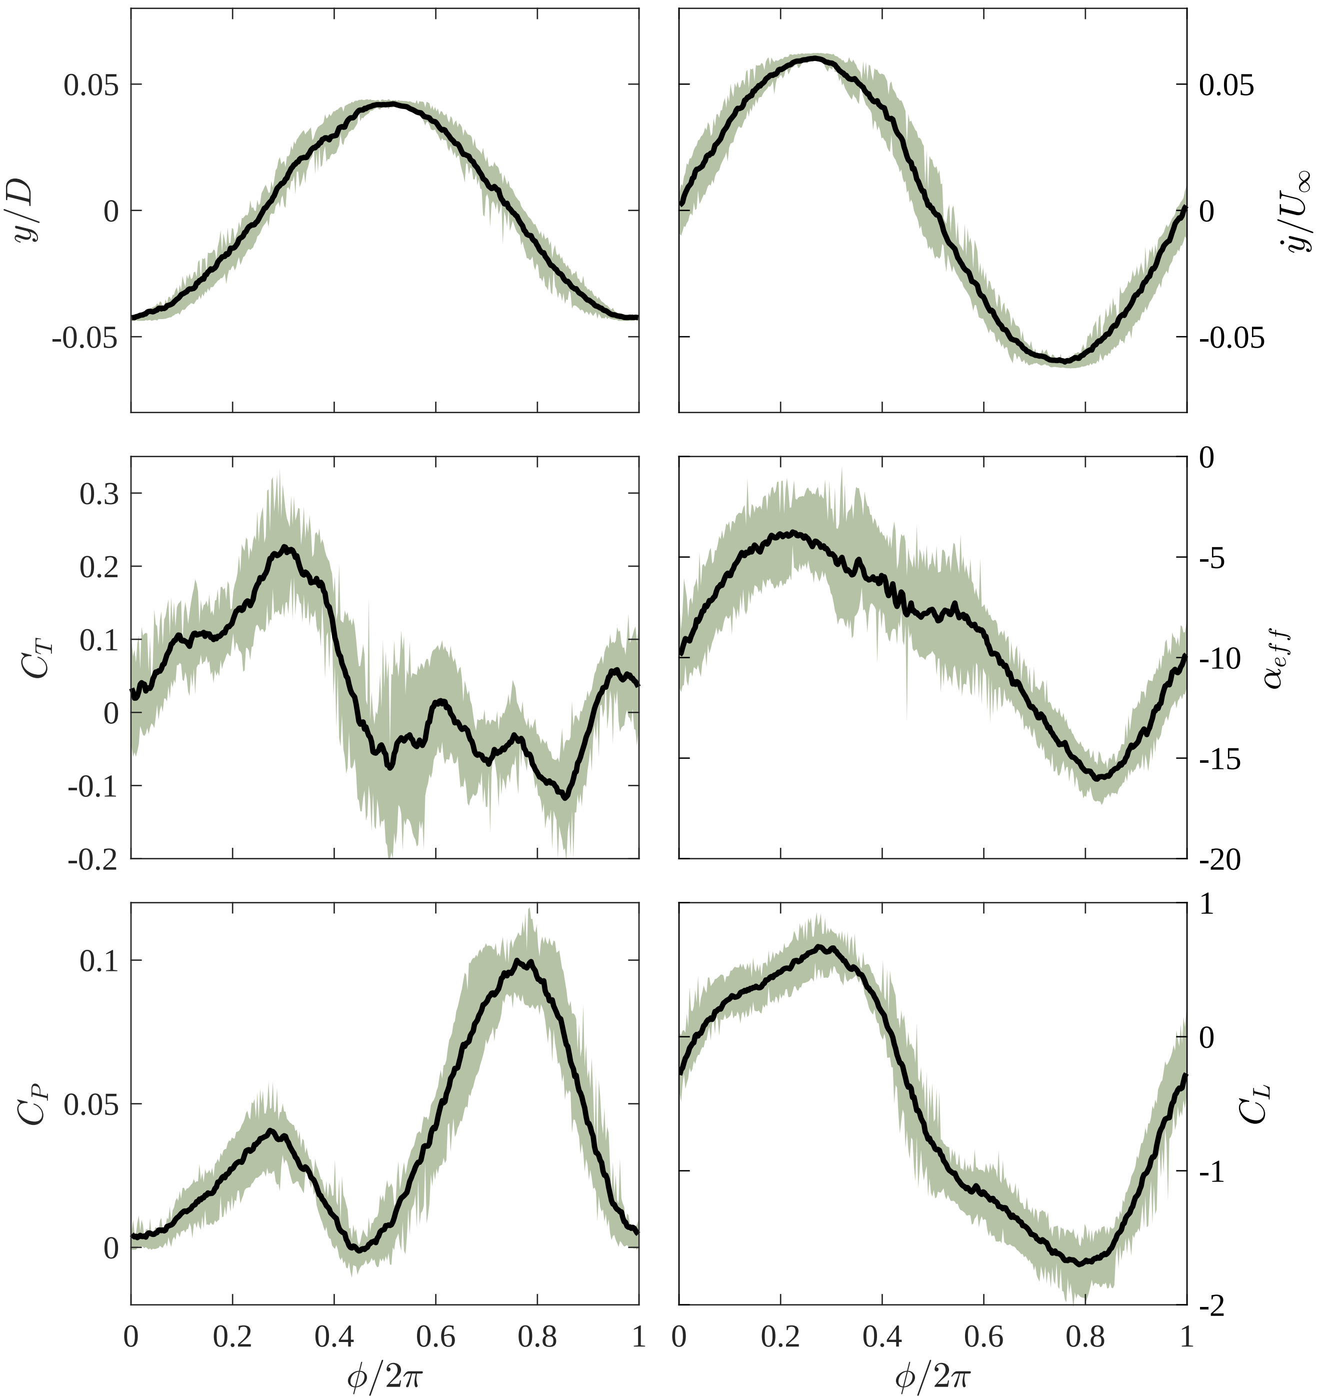

Figure 3.17: Phase-averaged dynamic quantities for Case HAoA-, or Driven Airfoil Experiments with $\alpha_0 = -10°$. Top Left: Airfoil Position. Top Right: Airfoil Velocity. Middle Left: Thrust Coefficient. Middle Right: Effective Angle of Attack. Bottom Left: Power Coefficient. Bottom Right: Lift Coefficient. Shaded regions indicate the interquartile range for data averaged in each phase bin. Bold black lines indicate the mean values recorded in each bin, with a 3rd order Savitzky-Golay filter applied to smooth the curves. The filter width was set to 5°.

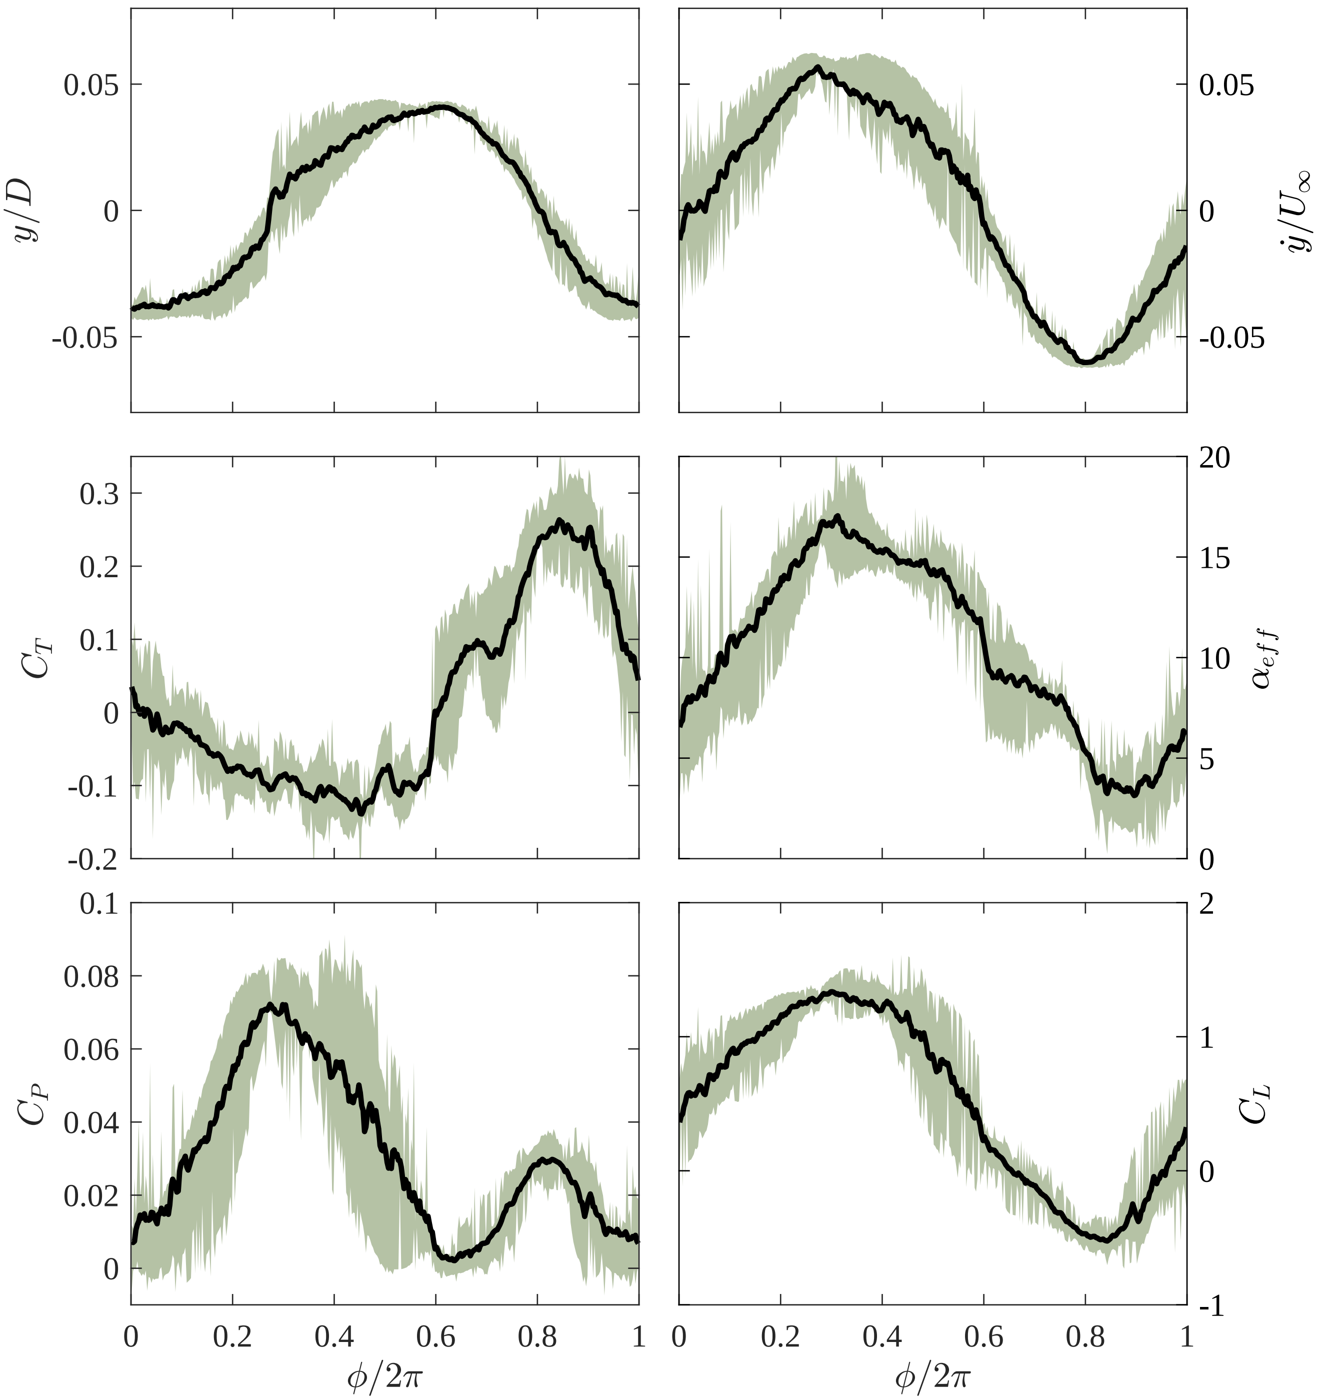

Figure 3.18: Phase-averaged dynamic quantities for Case HAoA+, or Driven Airfoil Experiments with $\alpha_0 = +10°$. Top Left: Airfoil Position. Top Right: Airfoil Velocity. Middle Left: Thrust Coefficient. Middle Right: Effective Angle of Attack. Bottom Left: Power Coefficient. Bottom Right: Lift Coefficient. Shaded regions indicate the interquartile range for data averaged in each phase bin. Bold black lines indicate the mean values recorded in each bin, with a 3rd order Savitzky-Golay filter applied to smooth the curves. The filter width was set to 5°.

HAoA- and Case HAoA+ respectively. The most notable change in the behaviour of the airfoil relative to Case BA is the increased asymmetry in the thrust and power coefficients output in Cases HAoA±, as well as the mean lift offset apparent in C_L. In addition, particularly for Case HAoA+ there appears to be increased variability in all measured quantities, especially in the first half of the cycle. This could be due to the relatively few observed phase-matched frames, coupled with greatly increased unsteadiness in the flow over that region in phase.

Of particular interest is the behaviour of our system outputs, C_P and C_T in response to the deviation of α_0 from 0. We see that for both quantities, one of two peaks observed per cycle is suppressed while the other is enhanced relative to Case BA. For C_P, suppression is apparent for the peak corresponding to the location in phase where oncoming flow acts to reduce the magnitude apparent angle of attack ($\phi/2\pi \approx 0.25$ for HAoA-, and $\phi/2\pi \approx 0.75$ for HAoA+). Interestingly, the opposite trend is true for C_T, which appears to be strongly suppressed at phases where $|\alpha_{\text{eff}}|$ is large.

The observed suppression/enhancement of the peaks in power is a direct result of the shift in C_L induced by a non-zero α_0. Momentarily neglecting unsteadiness, both in the oncoming free stream and due to the onset of stall at high values of α_{eff}, one would expect based on thin airfoil theory arguments (discussed in detail by Anderson (2011), for example) that over one cycle, the lift produced by the airfoil would simply be offset by a set value $\Delta C_L = 2\pi\alpha_0$. This corresponds to the idealized lift production for an airfoil at $\alpha = \alpha_0$, and for the values of interest here gives a lift coefficient offset of $\Delta C_L \approx 1$. Thus, instead of the range of C_L values experienced by the airfoil in Case BA, approximately $C_{L,\,\text{BA}} \in [-1, 1]$, we would expect to see $C_{L,\,\text{HAoA-}} \in [-2, 0]$, and $C_{L,\,\text{HAoA+}} \in [0, 2]$. Since for the Driven Airfoil Experiments the asymmetry in experienced lift does not affect airfoil velocity, which is fixed, these values of lift would result in a half-period where the signs of C_L and $\dot{y}/U_\infty$ were misaligned, leading to negative values of C_P and local minima with $C_P < 0$ (though the magnitude of the lift, and therefore the resulting magnitude of C_P over this portion of a cycle would be relatively small). By contrast, over the half-period where the signs of lift and velocity correspond, the increased lift magnitude would lead to higher positive values for C_P as compared to Case BA.

Of course, in the present flow case unsteadiness is a dominant feature, and the airfoil is not experiencing an ideal, uniform oncoming flow. Though strong asymmetry is observed between the peaks in C_P, power output remains positive over the majority of the cycle, and retains both local maxima with amplitudes larger than zero, as observed in Case BA. This stems from the observation that although the observed value of α_{eff} does not cross zero for either of Cases HAoA±, the lift does change sign over the cycle, in-phase with the airfoil velocity. It also fails to achieve the highest values of $|C_L|$ predicted using a simple, steady framework based on thin airfoil theory. This highlights the important role unsteadiness plays in the lift behaviour for these cases.

We first consider the contribution of stall to the observed values of C_L, starting from the simplest quasi-steady arguments. The effects of static stall on the production of lift by a very similar airfoil in a quasi-steady setting were recently characterized by Le Fouest et al. (2021), who noted an abrupt drop in C_L (through a distance $\Delta C_L \approx 0.4$) once an angle of attack of approximately $14°$ was reached and exceeded. Though in the present experiments we expect the onset of airfoil stall to conform more closely to a dynamic stall case rather than this static one, the mean lift coefficient observed over one cycle is reduced through a ΔC_L of approximately 0.5 relative to the thin airfoil theory prediction, which is reminiscent of the results of Le Fouest et al. (2021).

To explore the relevant fully unsteady fluid-structure interactions taking place in the flow, Figures 3.19 and 3.20 show contours of the phase-averaged Γ_2 Criterion for several instants in phase, similar to what was presented previously for Case BA. From the Figures, we see that the offset $\alpha_0 = \pm 10°$strongly affects flow in the region of the airfoil, as well as in the combined airfoil-cylinder wake region; however, it does not change the fundamental interaction of the airfoil with oncoming vorticity which is determined by its pre-planned trajectory. In both the present cases and Case BA, the foil slaloms between oncoming cylinder vortices. Thus, low-pressure zones created by the passage of cylinder vortices create the conditions for positive power extraction over the whole cycle, similar to observations made for Case BA.

For example in Case HAoA+, near a phase angle of $\phi/2\pi = 0.80 = \phi_5$ in Figure 3.20, the airfoil is experiencing a positive value of α_{eff} with a small magnitude close to zero. Considering the progression from top to bottom in the right-hand column of the Figure, we see that a large CW rotating Leading-Edge Vortex (LEV) has been recently shed from the foil (the size and strength of which is far beyond that observed for Case BA) and the process of reattachment on the airfoil's top surface is in progress in the final panel; the foil is moving down in the frame with its maximum velocity. Despite the small but positive value of α_{eff}, the foil is producing negative lift due to vortex suction from the large CCW vortex underneath it. Due to the recent

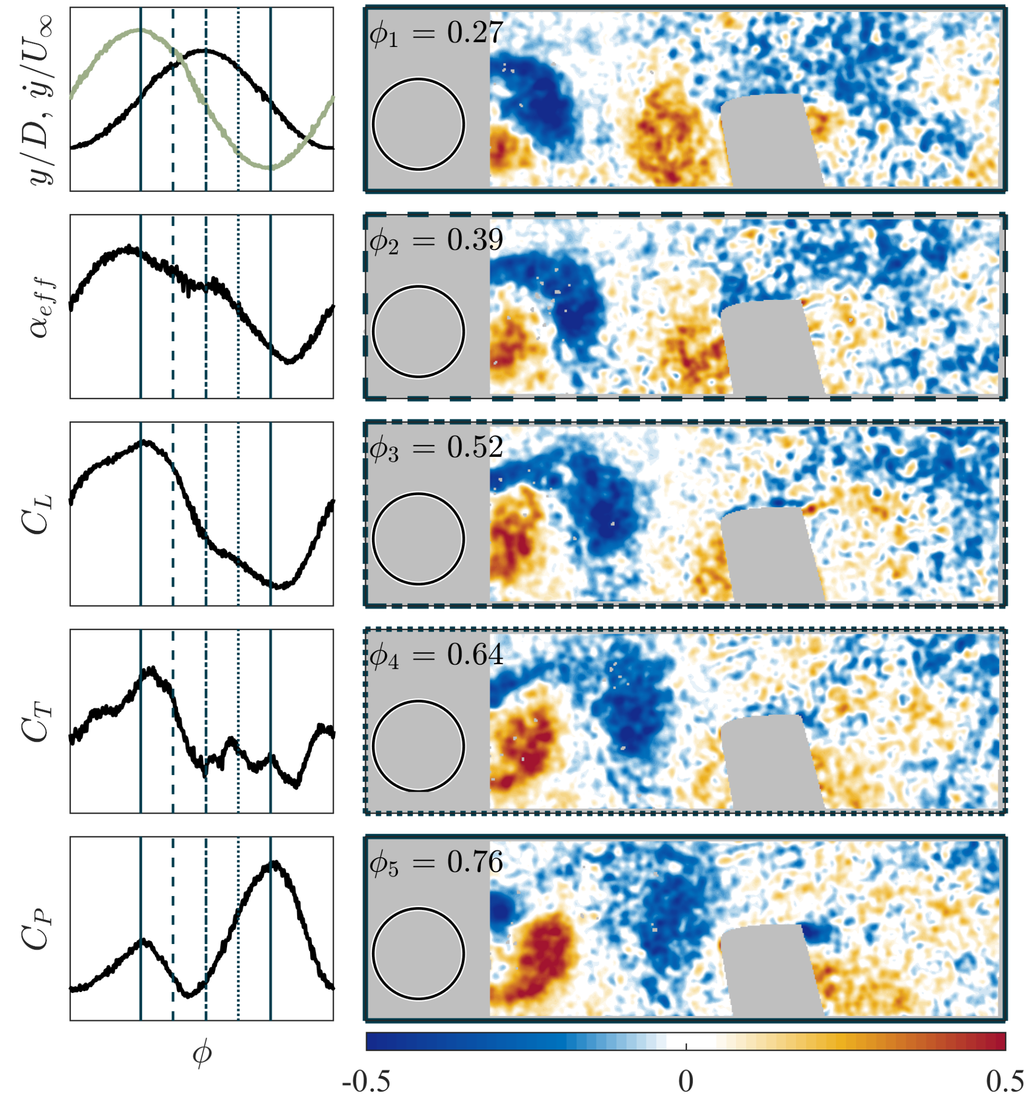

Figure 3.19: Snapshots of phase-averaged Γ_2 Criterion values at indicated points in the vortex shedding cycle, for an airfoil with a static angle of attack offset $\alpha_0 = -10°$, corresponding to Case HAoA-. Colour bar indicates contour level for the criterion, which is constrained so that $|\Gamma_2| \leq 1$. Left column reproduces data for phase-averaged quantities from Figure 3.17. In the top left panel: ▬ y/D; and ▬ $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angle ϕ is presented as a fraction of one cycle (2π). ▬ $\phi/2\pi = \phi_1 = 0.27$; - - $\phi/2\pi = \phi_2 = 0.39$; -· $\phi/2\pi = \phi_3 = 0.52$; ··· $\phi/2\pi = \phi_4 = 0.64$; ▬ $\phi/2\pi = \phi_1 = 0.76$.

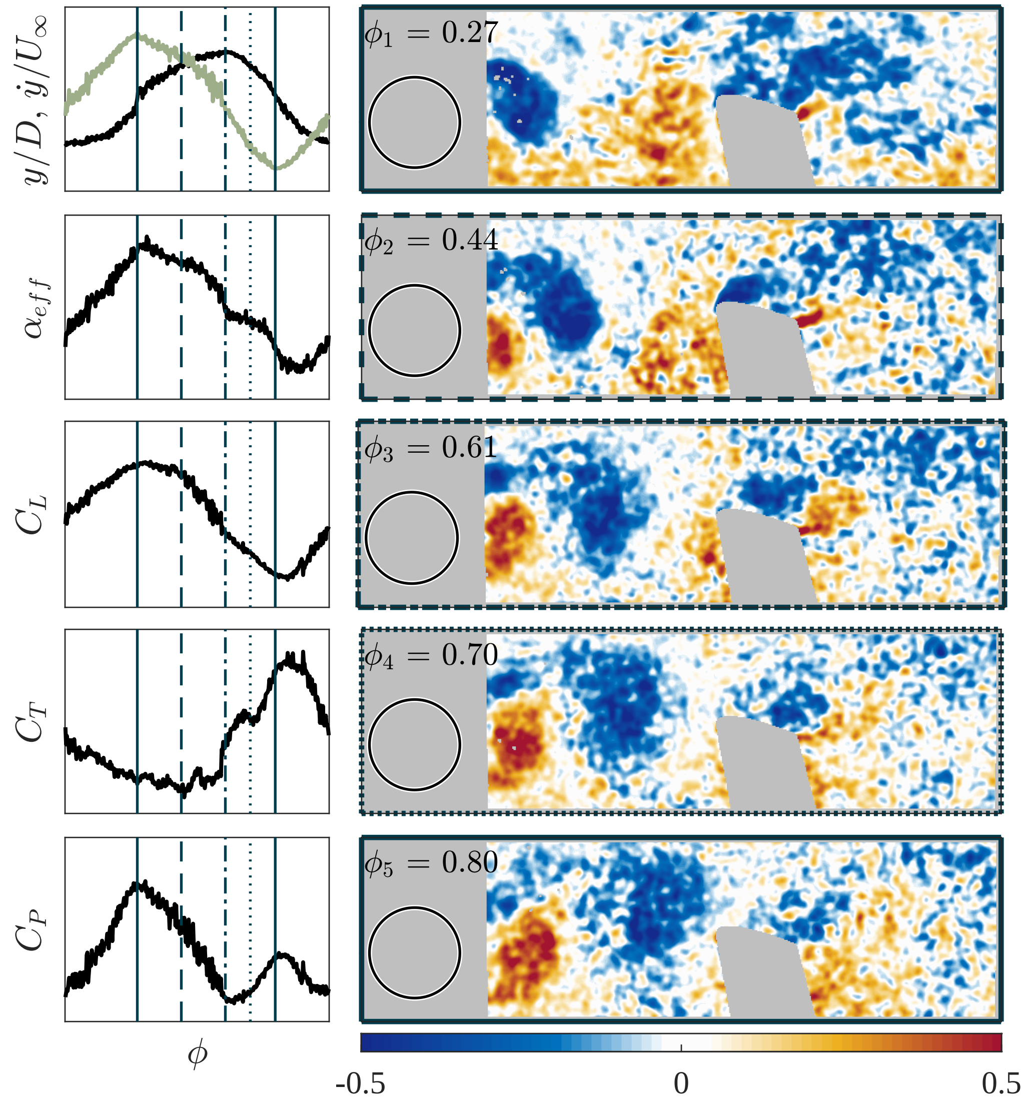

Figure 3.20: Snapshots of phase-averaged Γ_2 Criterion values at indicated points in the vortex shedding cycle, for an airfoil with a static angle of attack offset $\alpha_0 = +10°$, corresponding to Case HAoA+. Colour bar indicates contour level for the criterion, which is constrained so that $|\Gamma_2| \leq 1$. Left column reproduces data for phase-averaged quantities from Figure 3.18. In the top left panel: y/D; and $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angle ϕ is presented as a fraction of one cycle (2π). $\phi/2\pi = \phi_1 = 0.27$; $\phi/2\pi = \phi_2 = 0.44$; $\phi/2\pi = \phi_3 = 0.61$; $\phi/2\pi = \phi_4 = 0.70$; $\phi/2\pi = \phi_1 = 0.80$.

LEV growth and separation events in the time history of the airfoil, its ability to produce lift through quasi-steady means has been reduced, as discussed for example in a recent review by Corke and Thomas (2015). Therefore, the transverse force due to vortex suction overcomes any quasi-steady lift generated by the effective angle of attack, resulting in a negative lift (and a corresponding positive C_P).

The formation and shedding of this LEV contributes to both the enhanced power production at high $|\alpha_{\text{eff}}|$, and the simultaneous suppression of any peak at all in C_T at the same location in phase. As discussed for Case BA previously, the formation of an LEV during the dynamic stall process helps to keep experienced lift high as the foil passes its static stall angle. This contributes to enhancing lift in phase with the airfoil's velocity for increased $|\alpha_0|$, leading to the larger power peak experienced by the airfoil. In both Cases HAoA±, the maximum value of C_P is significantly larger than for either peak in Case BA, and an increase in the phase-averaged power per cycle is observed. For Case BA, $\overline{C_P} = 0.027$, while for Case HAoA-, $\overline{C_P} = 0.038$, and for case HAoA+, $\overline{C_P} = 0.031$. Such mean results should be interpreted with caution however, due to the relatively small number of observations included in these mean values for Cases HAoA±.

By contrast, as discussed for example by Corke and Thomas (2015), formation of a Dynamic Stall Vortex (DSV) as the foil exceeds its static stall angle leads to a precipitous increase in drag. The large LEV seen forming in the second panel of Figure 3.20 creates a low pressure zone above and downstream of the foil surface, which simultaneously acts to increase the lift and the drag. This increase in pressure drag is so large that it cancels out any thrust production the airfoil may be experiencing due to the Katzmayr effect, and leads to net drag $C_T < 0$ over the portion of the vortex shedding cycle where the LEV remains near the airfoil surface. Interestingly, considering Figure 3.19 showing the evolution of vorticity near the airfoil's top surface for $\alpha_0 = -10$, we see that the static angle of attack seems to suppress separation and LEV formation on the 'underside' of the foil relative to the angle of attack: though there is evidence of flow separation due to the passage of the large CW rotating cylinder vortex over top of the foil in the first half of the cycle, LEV formation and the deepening flow separation seen at the same phase(s) for Case BA are not present. It is unclear whether this lack of formation of an LEV has a significant impact on the thrust performance of the airfoil relative to Case BA, as asymmetry in the thrust peaks experienced in that case makes a direct comparison somewhat unfair; however there does seem to be a small benefit to thrust production

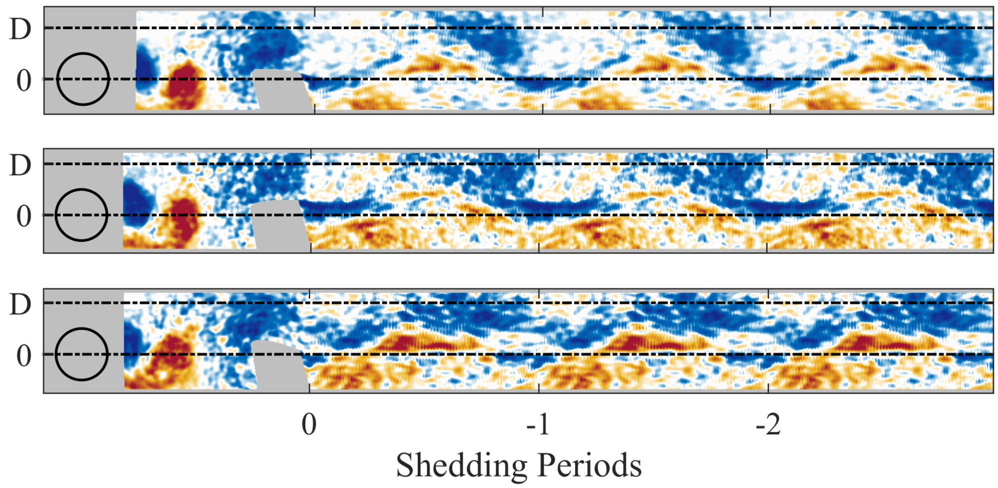

Figure 3.21: Frozen flow visualizations of Driven Airfoil Experiments performed with a high static angle of attack (Cases HAoA±), as well as with a static angle of attack of zero (Case BA) for comparison. Visualizations were created from phase-averaged Γ_2 Criterion fields using the method described in Chapter 2, and have the same contour levels in each panel. Top: Case BA. Middle: Case HAoA- with $\alpha_0 = -10°$. Bottom: Case HAoA+ with $\alpha_0 = +10°$. Dot-dashed lines indicate the locations $y = 0$ and $y = D$ in each panel. Slight differences in the underlying spatial fields for these data sets lead to the very small misalignment of the x-axes apparent in the Figure.

at phases where LEV shedding is suppressed.

A final observation regarding changes to the system behaviour in Cases HAoA± is a noticeable increase in asymmetric intensity of Trailing-Edge Vortex (TEV) shedding relative to Case BA. Figure 3.21 shows Frozen Wake Visualizations for Cases BA, HAoA- and HAoA+. For HAoA- (the middle panel in the figure), there is a noticeable increase in the intensity and extent of CW rotating trailing edge vorticity in the combined wake; in addition the intensity and extent of the CCW rotating vorticity below the airfoil appears somewhat enhanced relative to Case BA (the top panel in the Figure), though this trend is less clear. Correspondingly, for Case HAoA+ pictured in the bottom panel, these trends are reversed.

This enhancement of TEV shedding corresponds to the enhancement of lift experienced by the airfoil due to the addition of a static angle of attack α_0. For Case HAoA- shown in the middle panel, $\alpha_0 = -10°$ biases lift production towards negative values, or downward in the frame of Figure 3.21. As the lift produced by the airfoil is constantly changing, we expect that vorticity will be formed and shed at the trailing edge, similar to a starting vortex (though in this case, continuous unsteadiness takes the place of an isolated 'starting' event). Since more negative lift is generated in Case HAoA-, we expect that a correspondingly larger CW rotating TEV will be formed, fed by the large velocity gradient induced by the surrounding flow near the trailing edge. In the dual situation, in Case HAoA+ the airfoil produces a larger positive lift, and we correspondingly find an increase in the intensity of CCW TEVs in the combined cylinder-airfoil wake. Finally, LEV formation and shedding of the opposite sign to the TEVs with enhanced strength in each case could be contributing a slight overall increase in vorticity either above (Case HAoA+) or below (Case HAoA-) the airfoil.

3.12 Chapter 3 Interim Summary and Conclusions

This chapter presented results from an experimental campaign illuminating fluid-structure interactions taking place between vorticity shed by an upstream circular cylinder, and a downstream airfoil driven in the transverse direction, in-phase with the forcing it experiences. This system represents the simplest experimental setup of practical utility in investigating interactions between a cylinder wake and a downstream airfoil allowed to move passively in response to experienced forcing. Although the experiments discussed in this chapter involve actively driving the airfoil through a pre-planned trajectory, this trajectory was built to represent the behaviour of a passive airfoil. Thus these experiments shed light on the fluid-structure interactions that are likely to take place in a fully passive framework, while remaining experimentally simple. This chapter presents the following notable results:

1) For the experimental configuration discussed in this chapter, an airfoil driven transversely in the wake of a circular cylinder simultaneously maintained a positive power coefficient C_P and thrust coefficient C_T, indicating the potential for such a system to simultaneously produce net thrust and extract net power from the flow. Power extraction is facilitated through the maintenance of phase alignment between the transverse force and airfoil velocity.

2) Contributions to the observed net positive thrust due to the Katzmayr Effect, as well as unsteady aerodynamic effects induced by oncoming cylinder vortices were explored and characterized. The oscillation of the oncoming free-stream due to cyclic cylinder vortex shedding plays a significant role in the thrust production for this system.

3) Interactions of the airfoil with oncoming cylinder vortices give rise to the formation of a 2P-type wake in the combined airfoil-cylinder wake region,

through repositioning of cylinder-derived vorticity, as well as the addition of airfoil-derived vorticity to the wake. This 2P wake results in a wider wake region than that observed to arise from 2S-type shedding from the cylinder alone.

4) Increasing the magnitude of the airfoil's geometric angle of attack leads to improved power performance over one half of a vortex shedding cycle, at the expense of reduced power production over the other. Increased size and strength of leading-edge vortices formed due to this larger angle of attack suppress thrust production due to their negative impact on drag.

Chapter 4

PASSIVE CAPTIVE AIRFOIL MOTION IN 1 DIMENSION: OPTIMIZING PASSIVE BEHAVIOUR

4.1 Introduction

Building from the characterization of the driven airfoil in the wake of a circular cylinder provided in Chapter 3, this chapter presents the behaviour of a similar but distinct *passive captive* system. In the experiments presented in the following sections, the same airfoil as previous was mounted to the Captive Trajectory System (CTS), which was programmed to react passively to oncoming forces experienced by the airfoil. In contrast to the driven case, the mounting system reacts in the transverse direction as if it consists not of motors and actuators, but of a simple, linear spring and damper. Thus, the dynamics of the mounting system conform to the second-order linear dynamics of a spring-mass-damper system. Using this relatively simple setup, passive behaviour of an airfoil in the wake of an upstream cylinder is interrogated, and optimization of a fully passive flow-driven energy harvester based on these dynamics is demonstrated.

Second-order linear dynamics represent a popular choice in the literature surrounding Cyber-Physical Fluid Mechanics (CPFM) studies of flow energy harvesting. The current study was motivated by early results from Beal et al., where those authors used these simplified dynamics to simulate the behaviour of a dead fish interacting with the wake shed by an upstream D-shaped cylinder (Beal et al., 2006). As discussed in Chapter 1, this system is of considerable practical and scientific interest as it consistently demonstrates simultaneous net power extraction from the oncoming flow, coupled with net thrust larger than net drag. The mechanisms giving rise to these effects have been discussed in the previous chapter, through experiments using an airfoil driven through a pre-programmed trajectory. Here, those results are shown to conform closely to what is observed in the fully passive case.

The purpose of this chapter is twofold. Firstly, we demonstrate strong similarity between the behaviour of a Passive Captive Airfoil and a similar driven one, provided the passive mounting system is tuned to give rise to similar harmonic behaviour to that enforced in the driven case. Small differences in the behaviour of the passive airfoil due to its ability to react to changes in the oncoming flow are described.

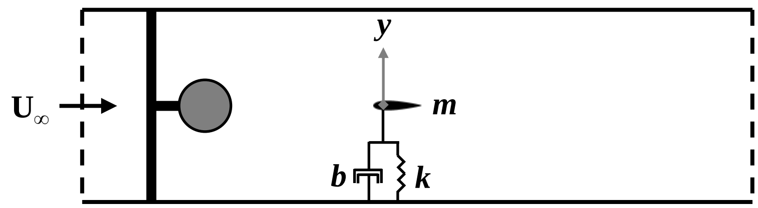

Figure 4.1: Basic Schematic of the cylinder-airfoil system showing spring-mass-damper dynamics in the transverse (y) direction.

Secondly, theory is presented to allow for the optimization of the power extraction potential of the Passive Captive Airfoil, subject to realistic engineering constraints on practical energy harvesting devices using this architecture. Experiments are performed to confirm that improved choices for mounting system dynamics in fact lead to improved power extraction performance. Potentially nonlinear feedback due to the airfoil's interaction with the oncoming flow is shown to influence this power extraction performance, outside of the predictions from simple linear theory.

4.2 Experimental Setup for the Passive Captive Airfoil

For all experiments described in this section, the Captive Trajectory System (CTS) was configured to allow the airfoil to respond to measured forces as a simulated passive system; therefore all experiments presented in this chapter are Passive Captive Airfoil Experiments, as described in detail in Chapter 2.

4.2.1 1-Dimensional Airfoil Motion

For the experiments described in this chapter, the motion of the airfoil was constrained to a single direction, the tunnel spanwise or airfoil heaving direction, here denoted as y as shown in Figure 4.1. The CTS was programmed to move the airfoil as if it were attached to a spring-mass-damper system, which has canonical 2$^{\text{nd}}$ order dynamics described later in this chapter.

Initial parameters under test were selected based on common conditions in literature, as well as practical considerations for preliminary testing and troubleshooting. To ensure good phase match between force and transverse velocity, a condition for optimal power production across a range of foil behaviours as described for example by Kinsey and Dumas (2008) or Su and Breuer (2019), a natural frequency of $\omega_n = 4$ s^{-1} was selected to roughly coincide with the observed mean forcing frequency from vortex shedding. As a conservative choice, our first experiments were conducted

Case Name	ω_n [1/s]	ζ	$M_{\omega_{\max}}$ [m/N]	m [kg]	b [Ns/m]	k [N/m]
Case 0	4.00	1.00	0.0034	18.38	147.06	294.12
Case 1	4.46	0.45	0.0034	18.38	73.81	365.94
Case 2	4.07	0.25	0.0034	36.76	74.72	607.53

Table 4.1: Parameters used to specify dynamics for each of the three Passive Captive cases tested. Case 0 corresponds to the conservative mounting parameters discussed in Section 4.3, while Cases 1 and 2 are described in Sections 4.5 and 4.6 respectively.

with $\zeta = 1$ (the critically damped case). This was chosen to ensure safe operation of the system during characterization and testing as this system experiences no resonance. In addition, for energy harvesting a large structural damping may be required, since this is the means by which fluid energy is transferred to the device ('harvested'). Finally, the amplitude response, $|H(j\omega)|$ was chosen such that the amplification at the natural frequency from force to position was $M_{\omega_n} = A_0/F_0$, so that the approximate observed amplitude of the motion would coincide with that for the driven case. This resulted in physically reasonable values for the dimensional parameters m, b and k given in Table 4.1. This preliminary, conservative case is referred to as Case 0 throughout the following discussion.

Once the basic setup had been tested and data concerning the motion as well as thrust and power production for the airfoil had been obtained, two additional cases were tested with different values for the parameters $[m, b, k]$. These cases present optimal tuning of the mounting system subject to different sets of engineering constraints, as discussed in the following sections. For completeness, these cases are also described in Table 4.1.

4.2.2 Filtering through Dynamics and Filter-based Phase Lag Mitigation

An additional benefit of using simple spring-mass-damper dynamics for our transverse motion is the reduced need for filtering of signals from the force sensor. By applying the transfer function for a tuned spring-mass-damper system (given by Equation 4.5, later) to our measured force, the contribution of noise in the sensor signal is greatly attenuated, since noise frequencies tend to be much higher than those of interest dynamically in this system. In this way, we eliminate the need to add additional causal filters to the dynamics of the airfoil, and can therefore mitigate any associated phase lag from such filters.

4.3 Passive Captive Motion in 1 Dimension with Case 0 Mounting Parameters

Experiments to characterize the behaviour of the passive captive airfoil with conservative mounting parameters were performed in the NOAH water channel as described in Chapter 2 as well as in Section 4.2. The free-stream velocity was approximately 0.32 m/s for all tests, giving rise to a Reynolds number based on cylinder diameter of approximately 40,000. The airfoil was placed 3.1D downstream of the circular cylinder, and constrained to move only in the transverse direction according to the programmed dynamics discussed in Section 4.2. For more details regarding the experimental setup, please see Chapter 2.

In this configuration, the airfoil performed quasi-harmonic motions in response to the forcing signal from the oncoming flow, similar in magnitude and frequency to those enforced in the driven case as described in Chapter 3. The parameters of the mounting system were selected purposefully and tuned to achieve this similarity between the motions. Figure 4.2 shows approximately 20 periods of typical airfoil motion, as well as the measured transverse (lift) force. The phase agreement between the lift and the velocity is excellent, showing little to no phase variation despite cycle-to-cycle fluctuations in measured lift. This demonstrates an immediate benefit of captive motions over the similar driven ones in terms of ensuring that force

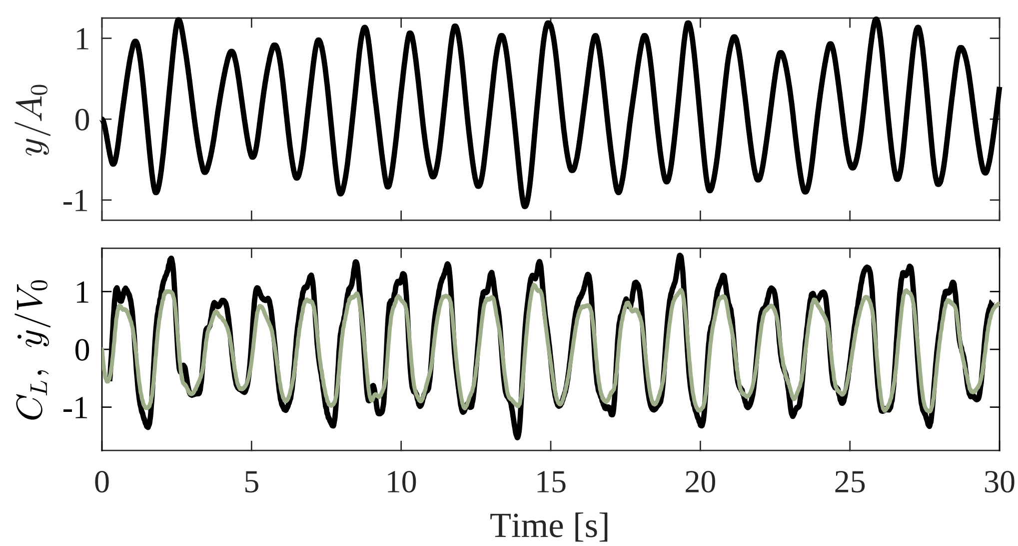

Figure 4.2: Time series behaviour of the Passive Captive Airfoil with Case 0 mounting parameters. Top: Airfoil transverse position, normalized by the standard amplitude for the driven experiments discussed in Chapter 3. Bottom: Lift Coefficient (▬), and Airfoil velocity normalized by the standard velocity amplitude for the driven experiments discussed in Chapter 3 (▬).

and velocity are well aligned over a long time horizon. This phase alignment is a direct result of tuning the natural frequency of the mounting system to coincide with the expected forcing frequency. This is an advantageous quality for a system operating in the wake of a cylinder at relatively high Reynolds numbers, as such wakes have been well established to exhibit cycle-to cycle irregularities due to vortex dislocations among other phenomena (Williamson, 1996). For systems with large damping ratios, the phase shift associated with an input forcing frequency deviation from ω_n is smaller than for systems with small damping due to the slope of the phase response in the region of ω_n, an advantage of operating at high damping ratios.

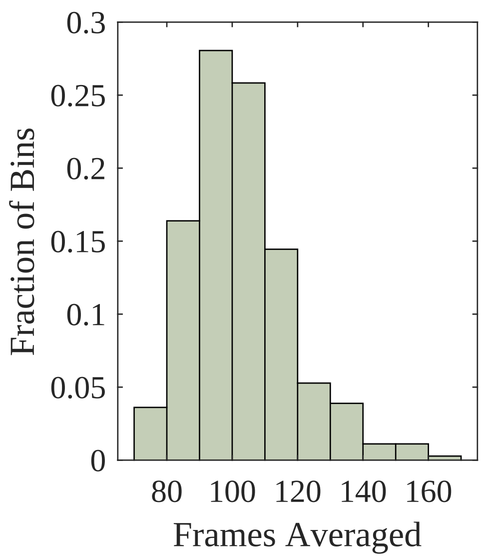

Figure 4.3: Histogram showing the fraction of phase bins with each number of observations included in phase averaging for Case 0.

To further characterize the behaviour of the airfoil, it is illustrative to consider phase-averaged dynamic quantities using the process described in Chapter 2 and demonstrated for the driven case in Chapter 3. Unlike the driven case however, here the airfoil is free to respond to oncoming forcing and all recorded snapshots are representative of 'true' airfoil behaviour regardless of phase match. This leads to a wider variation in the positions and velocities recorded as a function of phase than in the driven case, where frames with divergent behaviour were discarded prior to phase averaging. Figure 4.3 shows the distribution of number of observations recorded in each phase bin, where each bin has a width of 1°. From this figure, we see that the mean number of observations per bin is approximately 95, and there are no bins with fewer than 70 observations. Figure 4.4 shows phase-averaged dynamical quantities, along with shading which indicates the interquartile range for the observations captured in each phase bin. This means that for each phase bin, the middle 50% of observations fall within the shaded region.

From these observations, we can confirm that though there is cycle-to-cycle variability between periods in our dynamical quantities, the airfoil travels through a very closely sinusoidal path with a mean amplitude of approximately 0.04D, or

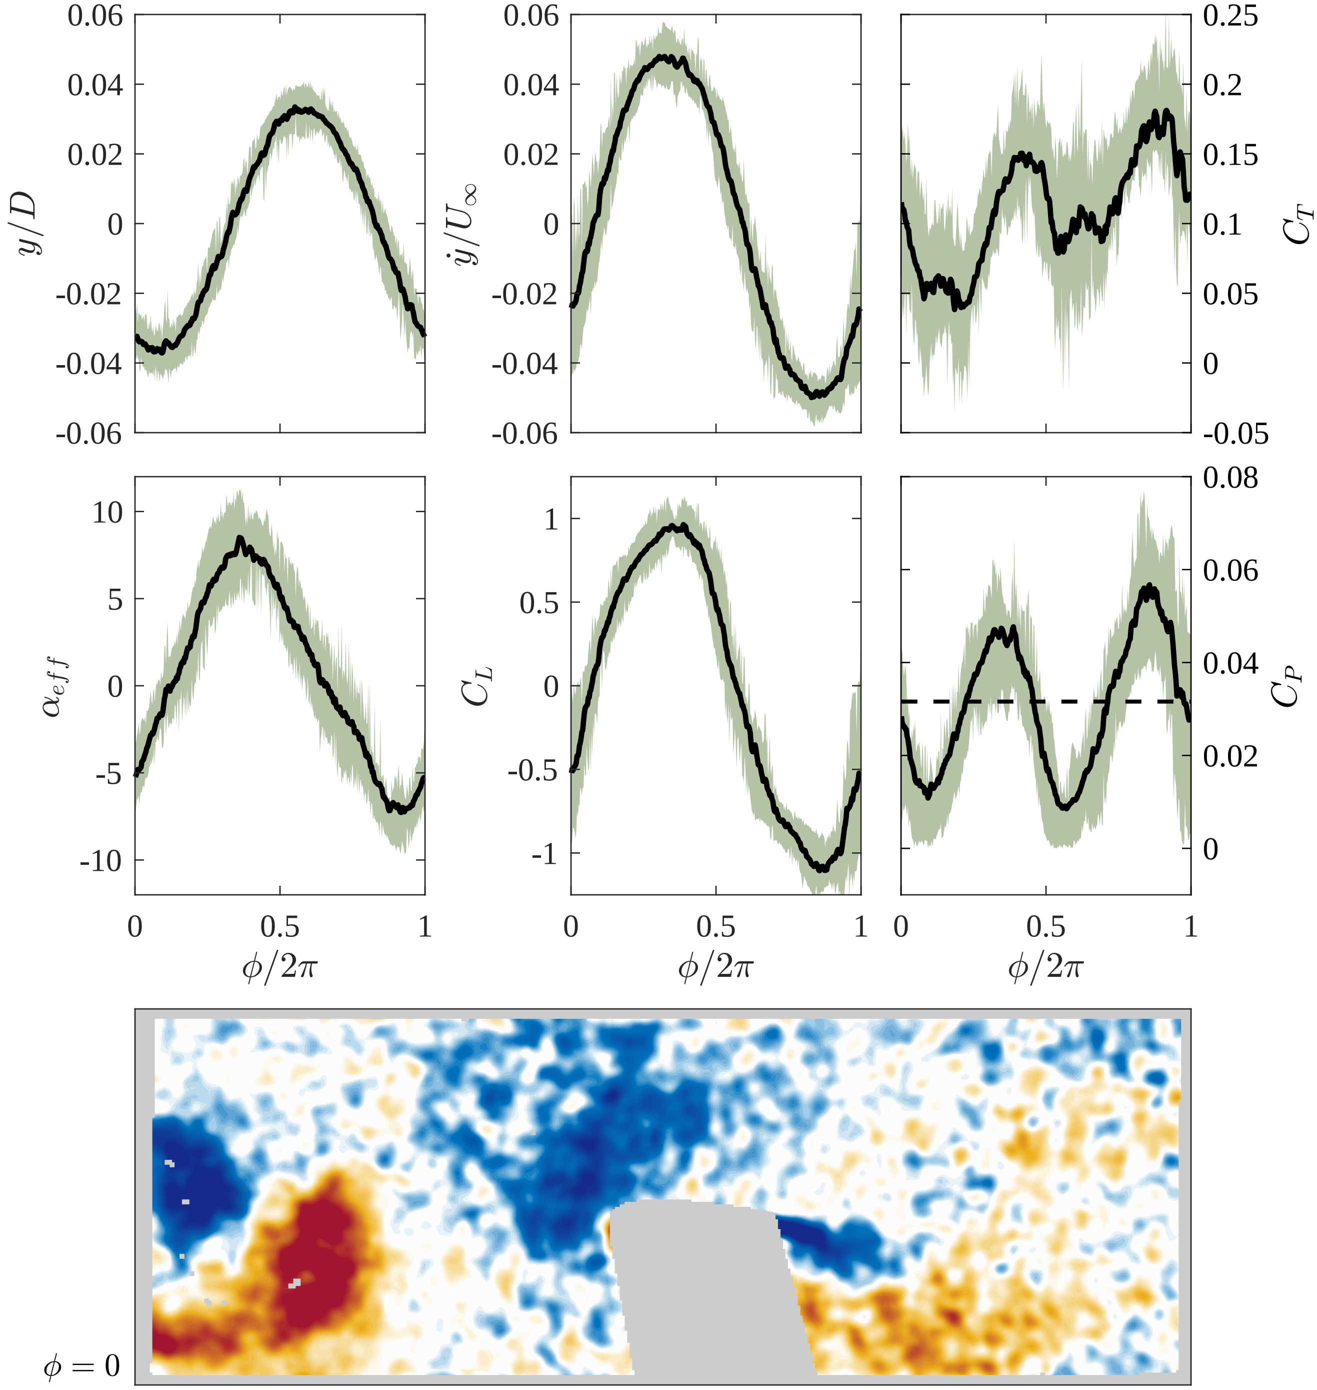

Figure 4.4: Phase-averaged dynamical quantities for the Case 0 Passive Captive Airfoil Experiments. Top Left: Airfoil Position. Top Center: Airfoil Velocity. Top Right: Thrust Coefficient. Middle Left: Effective Angle of Attack. Middle Center: Lift Coefficient. Middle Right: Power Coefficient. Bottom Panel: The flow state at $\phi = 0$, shown here as contours of the Γ_2 Criterion. Shaded regions in the top 5 figures indicate the interquartile range for data averaged in each phase bin. Bold black lines indicate the mean values recorded in each bin, with a Savitzky-Golay filter applied to smooth the curves. The filter width was set to 5°.

90% of the driven amplitude. The lift coefficient for the airfoil exhibits similarly quasi-sinusoidal behaviour, with the lift force well-aligned with the velocity. The variability in these quantities is notably smaller than for the effective angle of attack, thrust or power coefficients, the latter two of which exhibit roughly cyclic variations with double the frequency of the motion. Though the interquartile ranges for C_T and C_P are large, the values are both clearly larger than zero throughout a majority of each cycle, indicating that this system is simultaneously extracting net thrust and net power per cycle by passively responding to the vorticity in the oncoming flow. This mirrors the trends we observed for the driven airfoil, validating that simpler system as a model for the behaviour of this more complex passive one.

To connect these observed dynamics with the oncoming vorticity explicitly, we consider the evolution of the flow in the region of the airfoil as a function of phase, similar to the discussion presented in Chapter 3. We will consider five moments in phase as shown in Figure 4.5.

Overall, the interaction of the airfoil with oncoming vorticity is very similar to what was observed in the driven case, presented in Chapter 3. In the top panel of the Figure, when the airfoil is located directly behind the cylinder and moving upwards with its maximum velocity (and correspondingly generating close to its maximum power), we see that the airfoil is located between a region of positive or CCW vorticity (orange in the figure) and negative or CW vorticity (blue in the figure), and moving upward to avoid the oncoming positive vortex core. This is consistent with the slaloming mode of interaction with oncoming vorticity, as discussed by Beal et al. (2006) and references therein, as well as in Chapters 1 and 3. This behaviour is expected, as the slaloming mode for similar driven airfoil studies corresponds to low required input power (Streitlien et al., 1996; Beal et al., 2006). As the amplitude of the airfoil motion is very small compared to the size of the oncoming vortices, there is some evidence of vortex splitting: this is especially evident in the second panel as the oncoming CCW vortex impacts the airfoil, and some CCW vorticity is pushed up over top of the region of separation beginning to form near the airfoil's surface. The presence of the large negative vortex above and downstream of the airfoil makes it difficult to distinguish vorticity generated at the airfoil surface vs that in the free-stream; however, as discussed for a lower Re system by Wei and Zheng (2017), especially for relatively small amplitude motions this large vortex above the airfoil is critical in establishing a low-pressure zone on the suction side of the airfoil, significantly augmenting the experienced lift in a similar manner to the

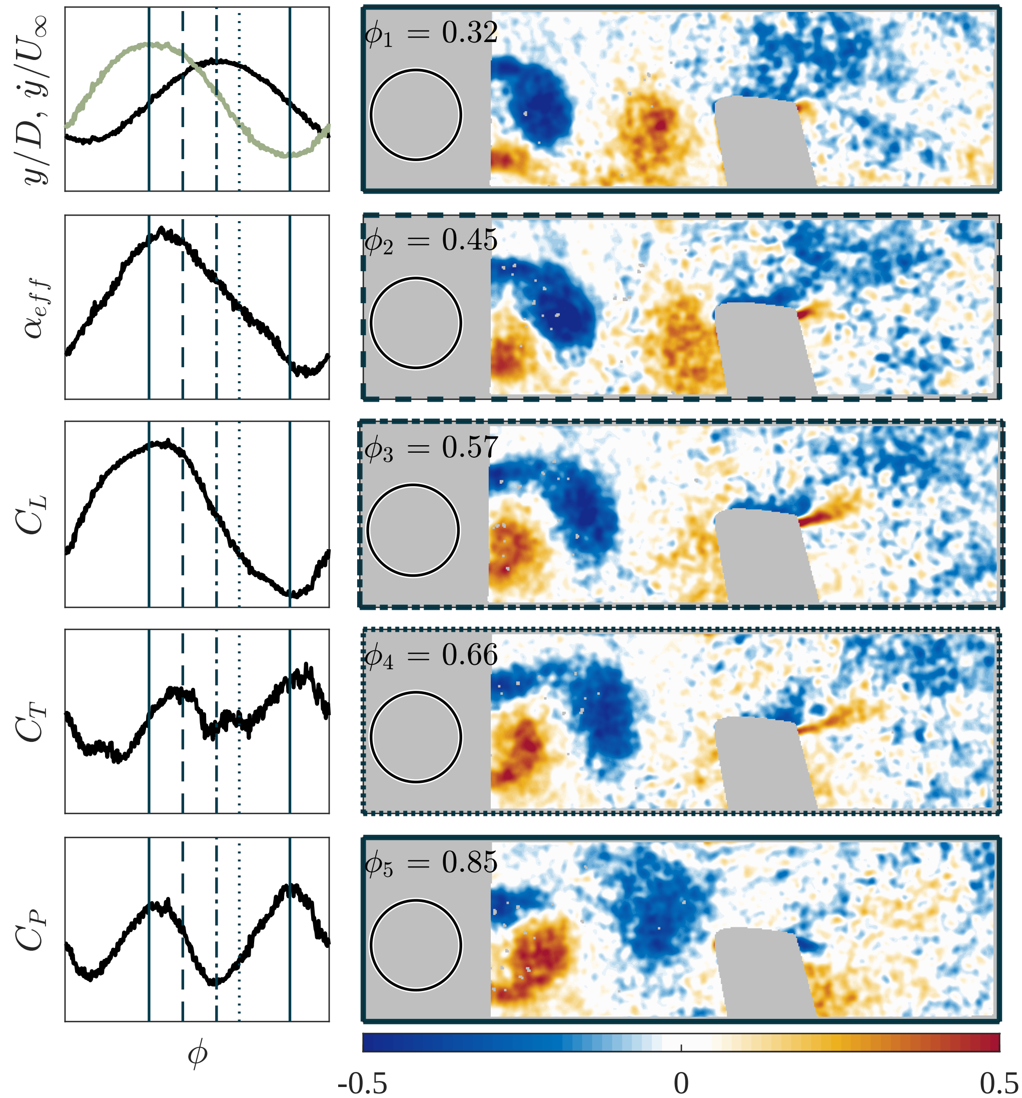

Figure 4.5: Snapshots of phase-averaged Γ_2 Criterion values at indicated points in the vortex shedding cycle, for an airfoil with Case 0 mounting parameters. Colour bar indicates contour level for the criterion, which is constrained so that $|\Gamma_2| \leq 1$. Left column reproduces data for phase-averaged quantities from Figure 4.4. In the top left panel: ▬ y/D; and ▬ $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angle ϕ is presented as a fraction of one cycle (2π). ▬ $\phi/2\pi = \phi_1 = 0.32$; ▬ ▬ $\phi/2\pi = \phi_2 = 0.45$; ▬ · $\phi/2\pi = \phi_3 = 0.57$; ··· $\phi/2\pi = \phi_4 = 0.66$; ▬ $\phi/2\pi = \phi_1 = 0.85$.

interaction of an airfoil with its own self-generated LEV.

In the second panel, the airfoil is decelerating as it continues to move upward and over top of the oncoming positive vorticity. It is fully engulfed in a region of upwards-flowing fluid, and the lift and effective angle of attack remain high, though this location in the cycle marks the beginning of a swift drop in the lift generated by the airfoil. While the lift remains positive (throughout the first three panels), the presence of the oncoming CCW rotating vortex causes the formation and subsequent shedding of a CCW-rotating trailing-edge vortex (TEV). The passing CW-rotating cylinder vortex contributes a positive effect on the output power by helping to keep the lift larger and positive while the airfoil is still moving upwards. As the lift passes its zero point in the third frame, the strength of this TEV is greatly reduced and it is detached and shed into the wake, but not before it has begun to interfere with the action of the CW-rotating vortex on the airfoil, preparing for the top side to transition from suction to pressure.

Although the flow in the region of the airfoil's surface appears to separate, one noticeably absent feature is the formation and subsequent shedding of a strongly coherent Leading-Edge Vortex (LEV), as described for the driven case in Chapter 3. One factor contributing to a reduced tendency for the airfoil to form and subsequently shed a coherent LEV during the cycle could be the ability of the passive airfoil to react individually to each oncoming vortex. From Figure 4.2, the oncoming forcing experienced by the airfoil is not exactly sinusoidal; in fact it varies not only in frequency, but each vortex induces a transverse force with a slightly different magnitude. Similar variations in the amplitude of the experienced lift were observed in the driven case, and discussed in Chapter 3. In the driven case, the airfoil executed the same pre-planned trajectory regardless of the flow conditions. By contrast, the Passive Captive Airfoil discussed here moves more quickly for a stronger impacting vortex, and more slowly for a weaker one. This acts to moderate the apparent angle of attack of the airfoil, reducing it in the case of a particularly strong vortex encounter and consequently reducing the tendency for the airfoil to form and shed coherent leading-edge vorticity. Considering the mean effective angle of attack over the cycle, the airfoil does not achieve the necessary α_{eff} to initiate dynamic stall in the classic sense (the process for which was described in Chapter 1, and has been thoroughly reviewed by McCroskey (1982) and Corke and Thomas (2015), for example), and any separation near the surface is more likely a direct result of interactions with the oncoming vorticity.

Despite the lack of coherent LEV formation, there is significant hysteresis in the behaviour of the airfoil throughout one cycle. Figure 4.6 shows C_L and C_P as a function of angle of attack, as well as C_L as a function of $\dot{y}/U_\infty$. The latter relationship, shown in the second panel in the Figure, confirms that the alignment between force and velocity is very good over the whole cycle. The first panel shows a relatively linear relationship between angle of attack and C_L, as we would expect for an airfoil operating in a quasi-steady setting (Corke & Thomas, 2015); however the lift achieved on the upstroke and downstroke of the motion are very different. As previously stated, this hysteresis is unlikely to be caused by classical dynamic stall (as this would require larger α_{eff}), and is instead induced by direct interactions between the airfoil and oncoming vorticity. Comparing the first and fourth panels in Figure 4.5, we see instants that correspond to the (near) maximal and zero locations for α_{eff} respectively. In the top panel, the airfoil is moving between regions of positive and negative vorticity while engulfed in a region of upwards-flowing fluid. It is achieving a high angle of attack despite its upward motion in-phase with the fluid velocity because the amplitude of its motion is small compared to the surrounding flow. By contrast, in the fourth panel the airfoil has begun to move downwards in the frame towards the centroid of a region of CW vorticity. The effective angle of attack is small, as the airfoil is experiencing both upwards and downwards-flowing oncoming fluid near its surface; however the presence of the large vortex below it helps to maintain a region of low pressure on the new suction side of the airfoil,

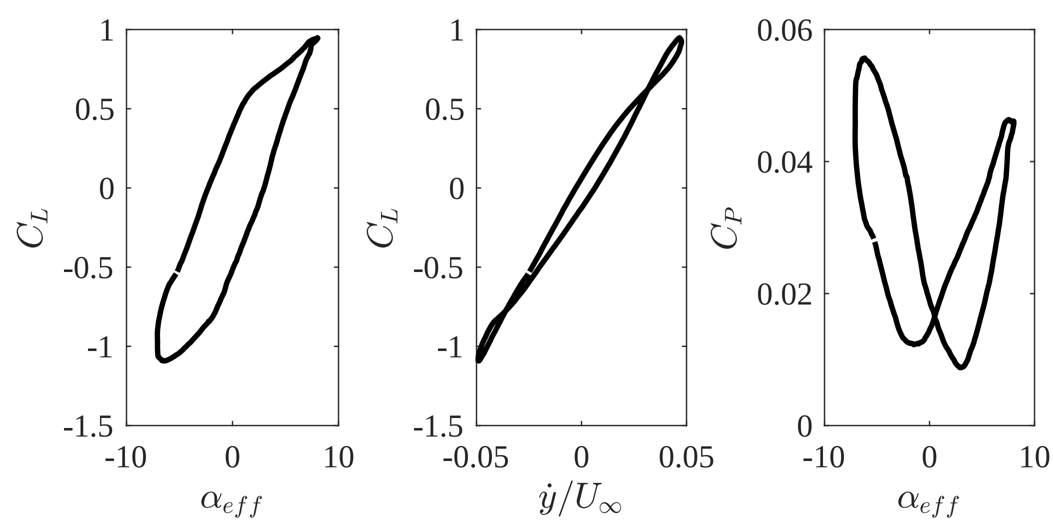

Figure 4.6: Phase portraits showing the behaviour of the Case 0 Passive Captive Airfoil. A Savitzky-Golay filter with a width of 10% of T was used to smooth phase-averaged data to improve clarity.

augmenting the negative lift (or increasing the lift in the downward direction). The presence of oncoming vortices therefore contributes to broadening the peaks of the lift curve while sharpening the transition from negative to positive relative to the uniform flow case, a similar effect to that of LEV formation and shedding in a driven airfoil setting (described by Kinsey and Dumas (2008), for example). In the present case, instead of the airfoil forming and shedding the vortex itself, a large vortex of the same sign is convected past the airfoil. Though these situations are dynamically distinct, both promote the maintenance of a high C_L when the velocity is positive, and vice versa.

The third panel of Figure 4.6 directly depicts the positive effect on power production of this hysteresis in lift. For the first quarter of the cycle as the lift initially increases, we see a steep ascent in the power augmented by the oncoming vortex suction. As the foil passes through its lift maximum, from the first panel in Figure 4.5 we see that the negative signed vortex has passed over the airfoil and lost some coherence, and a coherent positive vortex is now approaching the airfoil's leading edge. Though the effective angle of attack remains positive, there is a steep reduction in lift from panel 2 to panel 4 as the vortex passes under the airfoil, creating the hysteresis in the right-hand lobe of the 3rd plot in Figure 4.6. However, as the airfoil passes through its lift (and therefore power) minimum, the continued suction effect from the positive vortex causes the magnitude of the lift to again rise quickly, giving rise to the augmented power production portion of the left-hand lobe.

The formation and shedding of leading-edge and trailing-edge vorticity due to interactions with an oncoming vortex street was studied numerically by Wei and Zheng (2017). Though this was a low Reynolds number study compared to the present case, the authors observed very similar patterns of interaction between oncoming, leading edge, and trailing edge vorticity as in the current study. This lends credence to the theory that although the present study considers flow at a much higher Reynolds number, the unsteady effects observed are largely derived from interactions with the surrounding flow field rather than dynamic stall or other effects that are often determine the behaviour of an airfoil in a uniform oncoming flow.

As the airfoil begins to move downward in the fourth and final panels of Figure 4.5, the strength of the vorticity in the region of the top of the airfoil is greatly attenuated, and can be seen to separate from the leading edge and convect downstream along the airfoil's surface. As in the driven case, we see that as the top side of the airfoil

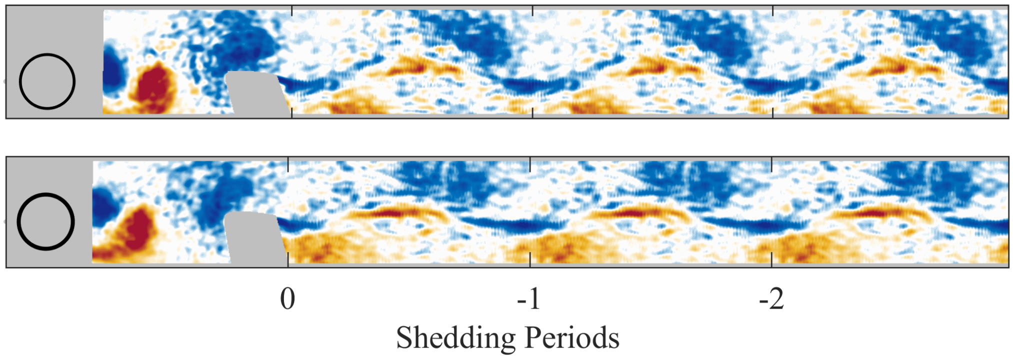

Figure 4.7: Frozen flow wake visualization for Case 0 mounting parameters, as compared to the Driven case described in Chapter 3. Top Panel: Wake visualization for the driven case, reproduced from Chapter 3. Bottom Panel: Wake visualization for the Case 0 passive captive airfoil. Both panels show contours of the Γ_2 Criterion with the same contour levels.

transitions from suction to pressure, some of the convected vorticity coalesces into trailing edge vorticity to be shed in the latter half of the cycle.

To compare the effect on the combined wake vorticity of the Passive Captive Airfoil to what was observed in the driven case, Figure 4.7 provides a frozen-flow wake visualization for both cases. Comparing the wakes generated by the driven airfoil (top panel) and the passive one (bottom), it is clear that the wake structures in both cases are quite similar, with trailing edge vorticity shed each cycle pairing up with oncoming vorticity to form a 2P-type wake as discussed in Chapter 3. The strength of the trailing edge vorticity shed from the airfoil seems to be similar in both cases, with shedding taking place over a similar duration in both cases as well. This is further evidence that the interaction of the Passive Captive Airfoil with the oncoming vorticity is similar to that in the driven case, despite the lack of consistent shedding of an LEV in each cycle.

4.4 Optimizing 1-Dimensional Passive Motions for Energy Harvesting

After successfully reproducing the behaviour of the driven airfoil and validating the behaviour of our passive captive system, we next consider methods to enhance the power production of this system through tuning the mounting system parameters. Throughout this process, consideration of practical engineering constraints on the airfoil motion is paramount; CPFD allows for the simulation of systems that would not be realizeable with standard mechanical components, so care is taken throughout the following discussion to ensure that the systems tested are practically realizeable without the use of a (powered) feedback control system.

4.4.1 Background: Second Order Dynamic Systems

Second order, linear dynamic systems are widely taught in the engineering and physics communities. Some basic discussion and definitions are included here to motivate and contextualize the work that follows. Much of this discussion is based on publicly available resources including online course material such as that from Trumper and Dubowsky (2005) and Cheever (2021), for example.

A canonical second-order linear system in mechanics, the spring-mass-damper system is described by the following differential equation:

$$F(t) = m\ddot{y}(t) + b\dot{y}(t) + ky(t), \tag{4.1}$$

where m is the mass in kg, b is the damping in Ns/m, and k is the spring constant in N/m. For such a system, it is customary to define the following quantities that describe the behaviour of the system more intuitively:

$$\omega_n = \sqrt{\frac{k}{m}}, \tag{4.2}$$

$$\zeta = \frac{b}{2\sqrt{km}}, \tag{4.3}$$

$$\omega_d = \omega_n\sqrt{1-\zeta^2}. \tag{4.4}$$

The undamped natural frequency ω_n describes the frequency of the system response to a non-harmonic input for ζ, the damping ratio, = 0. For $\zeta \in (0, 1)$, the system will oscillate at the damped natural frequency, ω_d. For $\zeta \geq 1$, oscillations in the response are damped out entirely. The system is said to be critically damped for the case $\zeta = 1$, which means that subject to an impulsive forcing, the system returns to its steady state position along the most efficient path with zero overshoot.

If we consider the response of the above system in terms of its transfer function, we can write the response to a harmonic input as

$$\frac{Y(j\omega)}{\mathcal{F}(j\omega)} = H(j\omega) = \frac{1}{(k - m\omega^2) + jb\omega}. \tag{4.5}$$

Considering the magnitude and phase functions and the relationships to our system quantities above, we can write

$$|H(j\omega)| = \frac{1}{\left((k - m\omega^2)^2 + b^2\omega^2\right)^{1/2}} = \frac{(1/k)\omega_n^2}{\left((\omega_n^2 - \omega^2)^2 + 4\zeta^2\omega_n^2\omega^2\right)^{1/2}}, \tag{4.6}$$

$$\phi = -\text{atan}\left[\frac{b\omega}{(k - m\omega^2)}\right] = -\text{atan}\left[\frac{2\zeta\omega_n\omega}{(\omega_n^2 - \omega^2)}\right]. \tag{4.7}$$

Using these relationships, we can determine the input frequency of maximum amplification, or the frequency where the gain from input force to output position is largest, by differentiating $|H(j\omega)|$. Doing so, we find

$$\omega_{\max} = \omega_n\sqrt{1 - 2\zeta^2}. \tag{4.8}$$

For cases where $\zeta \geq 1/\sqrt{2}$, the frequency of maximum amplification is $\omega = 0$. It is interesting to note that the most amplified frequency $\omega_{\max}$ in general does not coincide with either ω_n or ω_d for non-zero damping.

4.4.2 Choosing Mounting System Parameters: Mass and Gain Constraints

From the preceding discussion, three free parameters are required to specify the behaviour of the airfoil, corresponding to the choice of $[m, b, k]$. A common alternative to this is a description based on the parameters $[\omega_n, \zeta]$, but to fully specify the system with these parameters, an additional choice is required. Common choices in literature include choosing $m = 1$, a unit mass, or $k = 1$, a unit stiffness. Although these are convenient in theory, in the design of a real system realizing these choices may not be possible or practical. Instead, it may be convenient to specify the mass of the system as a known non-unitary value: for example, a particular energy harvesting device prototype under test could have a pre-determined mass. In addition, an important parameter for the safe operation of the device is the maximum force-to-position gain, or $|H(j\omega)|_{\max} = M_{\max}$. This governs the largest excursions from the neutral position that the device will experience based on the expected incoming force. With these engineering considerations in mind, we specify the system in terms of the constraints $[m, M_{\max}]$, and then choose a third parameter to maximize the energy harvesting potential for the device.

This can be accomplished by incorporating information about the oncoming flow. For the conservative case described in the previous section, the natural frequency of the system was chosen to correspond to the vortex shedding frequency ($\omega_n \approx \omega_f$) based on references such as Su and Breuer (2019) and Kinsey and Dumas (2008). Examining Equation 4.7, the condition ϕ = -90° occurs at $\omega = \omega_n$ regardless of the system damping, and we correspondingly observed that the lift force and velocity were very well aligned over a long time horizon in the conservative case (see Figure 4.2). However, from Equation 4.8, the maximum amplification from force to velocity occurs at another location, $\omega = \omega_{\text{max}}$, which is a function of ζ. This implies that in a magnitude-constrained system operating largely at a specified frequency determined by the vortex shedding from the upstream cylinder, there is a trade-off between choosing $\omega_n = \omega_f$ to achieve optimal alignment between the force and the velocity, or to choose $\omega_{\text{max}} = \omega_f$, so the system operates at the maximum velocity magnitude allowable for our M_{max} constraint, but has a slight phase shift relative to the lift force.

A mathematical description of these competing factors is developed as follows. We consider our 2$^{\text{nd}}$ order system dynamics subject to a harmonic forcing described by Equations 4.6 and 4.7, so that an applied force $F(t)$ gives rise to a sinusoidal variation in position:

$$F(t) = F_0 \cos(\omega t - \phi(\omega)), \tag{4.9}$$

$$y(t) = F_0 A(\omega) \cos(\omega t), \tag{4.10}$$

where for notational convenience, $A(\omega) = |H(j\omega)|$.

Here $\phi(\omega)$ is the phase shift from force to position (as described by Equation 4.7); it is subtracted from the argument of the cosine in the force rather than added to that for position for convenience in later developments.

By differentiating the expression for the position $y(t)$, we find that

$$\dot{y}(t) = -F_0 \omega A(\omega) \sin(\omega t), \tag{4.11}$$

$$\ddot{y}(t) = -F_0 \omega^2 A(\omega) \cos(\omega t). \tag{4.12}$$

Re-arranging the above expressions, we have that

$$\sin(\omega t) = \frac{-\dot{y}(t)}{F_0 \omega A(\omega)}, \tag{4.13}$$

$$\cos(\omega t) = \frac{-\ddot{y}(t)}{F_0 \omega^2 A(\omega)}. \tag{4.14}$$

Following discussion by Su and Breuer, 2019, we re-write the expression for the applied force in terms of contributions aligned with the airfoil's velocity, and with its acceleration. Thus, we consider:

$$F(t) = F_0 \big(\cos(\phi) \cos(\omega t) + \sin(\phi) \sin(\omega t)\big), \tag{4.15}$$

$$F(t) = \frac{\cos(\phi)}{-\omega^2 A(\omega)} \ddot{y}(t) + \frac{\sin(\phi)}{-\omega A(\omega)} \dot{y}(t). \tag{4.16}$$

To harvest energy with this system while operating at a steady state, any power extracted from the flow must be fully dissipated (commonly, either transferred to useable electrical energy or released as waste heat) through damping. Power extraction occurs when the flow does useful work on the airfoil; that is, when the applied force results in a velocity with some component in the direction of forcing. Using Equation 4.16 above, and making the assumption that all useful work is in fact dissipated (the further consequences of which are discussed at more length by Su and Breuer (2019)), power dissipation $P_D(t)$ for this system is given by

$$P_D(t) = F(t) \cdot \dot{y}(t) = \frac{-\sin(\phi)}{\omega A(\omega)}\dot{y}^2(t) = -\sin(\phi)F_0^2\omega A(\omega)\sin^2(\omega t) \tag{4.17}$$

since for sinusoidal motion, $\dot{y}(t) \perp \ddot{y}(t)$. Finally, cycle-averaged power output is given by

$$\overline{P_D(t)} = \frac{1}{T}\int_0^T P_D(t)\ \mathrm{d}t = \frac{-\sin\big(\phi(\omega)\big)F_0^2\omega A(\omega)}{2}. \tag{4.18}$$

This expression is similar to that developed by Su and Breuer (2019), but explicitly accounts for the variation in the phase and magnitude responses as a function of the forcing frequency. For a known operating frequency $\omega = \omega_f$ and forcing amplitude F_0, the above expression provides an estimate for the power dissipated (available for extraction) by an energy harvester with behaviour described by $A(\omega)$ and $\phi(\omega)$.

Figure 4.8 shows a map of the estimated cycle-averaged power available to an energy harvesting system as a function of the choice of ω_n and ζ, subject to the constraint that the maximum gain is less than some specified value, $M_{\max}$. In the figure, $M_{\max}$ is set to 0.0034 m/N, to correspond with the Case 0 experiment described previously. The operating frequency ω is set to ω_f = 3.91 rad/sec to correspond with the observed mean forcing frequency from cylinder shedding. The parameters $[m, b, k]$ vary as a function of $[\zeta, \omega_n]$ in the figure, as the system is specified based on the characteristics $[\zeta, \omega_n, M_{\max}]$ rather than using the classic parameters.

From the figure, the naïve optimal choice to achieve the highest power output per cycle is to choose $\zeta = 0$ and $\omega_n = \omega$. This corresponds to an undamped mounting system having $b = 0$, which causes the damped natural frequency given by Equation 4.4 and the frequency of maximum amplification given by Equation 4.8 to coincide with the operating frequency, while preserving the optimal phase shift between force and velocity. Although these conditions do provide an optimal compromise between exploiting a maximum amplitude gain and maintaining an optimal phase shift, the

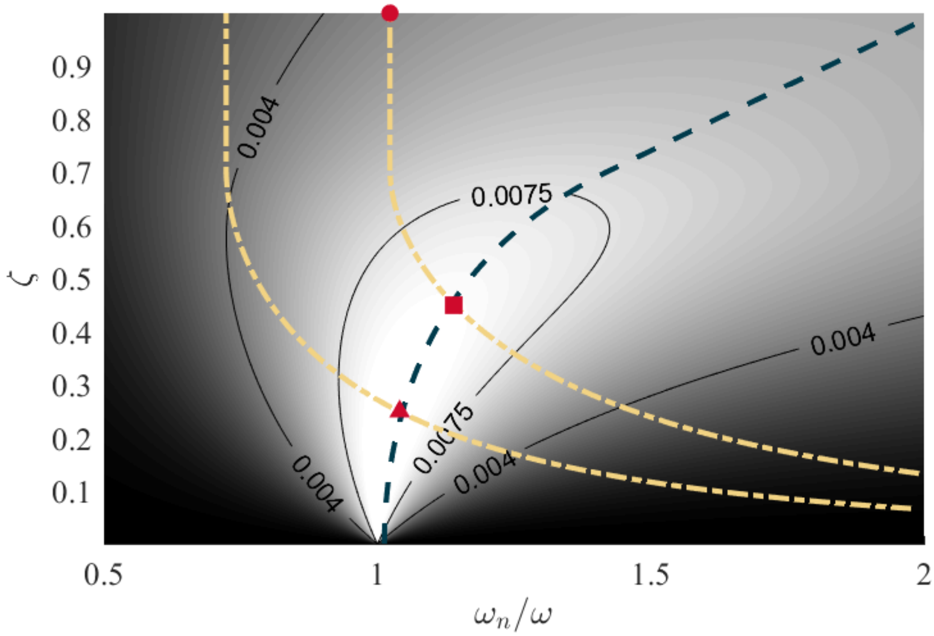

Figure 4.8: Contour plot showing the cycle-averaged power extraction potential per square newton input ($\overline{C_p}/F_0^2$) for the system, subject to the constraint $A(\omega) < M_{\text{max}}$. Three symbols indicate the test conditions described in Table 4.1: ● Case 0. ■ Case 1. △ Case 2. Symbols are joined by three lines: lines of constant mass where m = 18.38 kg (top line, intersecting ● and ■), and m = 36.76 kg (bottom line, intersecting △). a line of constant damping where b = 74.27 Ns/m.

configuration is not achievable for a real engineering system. As the damping approaches zero, the amplitude response becomes large in the vicinity of the natural frequency. Since we have placed a constraint on the maximum amplitude gain, the stiffness k becomes large to attempt to stop the motion of the system from exceeding the constraint. Since the natural frequency is also fixed at a finite number, from Equation 4.2 the mass of the system must also become large. Thus, the condition $\zeta = 0$, $\omega_n = \omega_f$ subject to a maximum gain constraint constitutes a singularity, where $b \to 0$, while $m, k \to \infty$. This singularity is also responsible for creating the unsatisfactory situation where maximum power extraction occurs with zero power dissipation ($b = 0$), further confirming that this is not a realistic operating condition.

To resolve this issue, an additional constraint on the physical parameters of our system is required. Referring back to the discussion at the beginning of this section, for a real engineering system it may be convenient to consider the mass m fixed by the system design. Thus, we consider contours of fixed mass within the parameter space, shown in Figure 4.8 as yellow dashed curves. The value of the mass along these yellow curves are 18.38 kg (top line) and 36.76 kg (bottom line), corresponding to the mass tested in Case 0 described in the previous section, and a value of mass double that. Then, to fully specify the system in terms of these constraints, the damping ratio ζ can be selected based on Figure 4.8, at the location of maximum power extraction potential along these lines of constant mass.

To determine the b and k values corresponding to this selection, we re-arrange Equations 4.2-4.8 to enforce the conditions that $|H(j\omega_{\max})| = M_{\max}$ and that the mass m is fixed for any chosen value of ζ. The former constraint results in the following expression giving k in terms of $[M_{\max}, \zeta]$:

$$k = \begin{cases} \frac{1}{M_{\max}} & \zeta \geq 1/\sqrt{2} \\ \frac{1}{2M_{\max}\zeta\sqrt{1-\zeta^2}} & \zeta < 1/\sqrt{2} \end{cases} \tag{4.19}$$

where the condition $\zeta = 1/\sqrt{2}$ corresponds to the disappearance of the peak in the amplitude response such that $|H(0)| = M_{\max}$. Once k has been determined using Equation 4.19, then Equations 4.2 and 4.3 can be used to determine the corresponding value for b (and ω_n) using the fixed value of m. Thus, we have fully specified the system based on the choice of $[M_{\max}, m, \zeta]$.

4.4.3 Choosing Mounting System Parameters: Optimal Structural Damping

Again considering Figure 4.8, it is apparent that should the mass of the system be increased, there is further potential for improved power extraction. Comparing the two yellow dashed lines in the Figure, we see that the larger mass curve passes through a region of larger power extraction potential as compared to the smaller mass curve. Choosing the optimal damping value for the larger mass, $m = 36.76$ kg, according to Figure 4.8 and then using the method described in the previous section, the appropriate choices of b and k to satisfy the same constraint on $M_{\max}$ can be determined. Interestingly, although m and k (and correspondingly ω_n) are different, the value determined for the optimal damping is strikingly similar to the smaller mass case. A line of constant b is plotted in Figure 4.8 in dark blue, which corresponds closely with the contour of maximum power extraction potential as a function of ω_n/ω.

To explain this, we begin by combining Equations 4.5, 4.7 and 4.18 to express the cycle-averaged power extraction potential for a general system as

$$\overline{P_D} = \frac{F_0^2}{k}\left[\frac{\zeta\omega_n^2\omega^2}{(\omega_n^2 - \omega^2) + 4\zeta^2\omega_n^2\omega^2}\right]. \tag{4.20}$$

To apply the constraint on the maximum gain, we substitute the definition for k given in Equation 4.19 for the case $\zeta < 1/\sqrt{2}$, to obtain

$$\overline{P_D} = F_0^2 M_{\max}\left[\frac{2\zeta^2\omega_n^2\omega^2\sqrt{1-\zeta^2}}{(\omega_n^2-\omega^2)+4\zeta^2\omega_n^2\omega^2}\right]. \tag{4.21}$$

To determine the location of the ridge of maximum power extraction potential, we differentiate Equation 4.21 with respect to ω_n, and set the result to zero:

$$\frac{\partial\overline{P_D}}{\partial\omega_n} = 2F_0^2 M_{\max}\left[\frac{\zeta^2\omega_n^2\omega^2\sqrt{1-\zeta^2}\left(3\omega^4+2\omega^2\omega_n^2(2\zeta^2-1)-\omega_n^4\right)}{\left((\omega_n^2-\omega^2)+4\zeta^2\omega_n^2\omega^2\right)^2}\right] = 0. \tag{4.22}$$

Since $\zeta < 1/\sqrt{2}$, this simplifies to finding the roots of the following polynomial:

$$\omega_n^4 - 2\omega^2(2\zeta^2-1)\omega_n^2 - 3\omega^4 = 0.$$

Since we are only interested in real, positive frequencies, we consider the positive real root and obtain the solution

$$\omega_n = \omega\left(2\sqrt{\zeta^4-\zeta^2+1}+2\zeta^2-1\right)^{1/2}. \tag{4.23}$$

Equation 4.23 describes the values of ω_n corresponding to the optimal choice for power extraction potential given a specified value of ζ, while respecting the $M_{\max}$ constraint described in the previous section. As expected, the location of this ridge is not dependent on the forcing amplitude or the value of the constraint; in addition, Equation 4.23 specifies the ratio ω_n/ω, and therefore scales with operating frequency.

Considering instead a description for lines of constant damping, b, using Equations 4.2 and 4.3 we can write for a general mass-spring-damper system

$$\zeta = \frac{b\omega_n}{2k}. \tag{4.24}$$

Again substituting Equation 4.19 for the case $\zeta < 1/\sqrt{2}$ and solving for ω_n,

$$\omega_n = \frac{1}{M_{\max}b\sqrt{1-\zeta^2}}. \tag{4.25}$$

Equation 4.25 represents ω_n corresponding to a given value of ζ along a line of constant b, with the $M_{\max}$ constraint enforced. To specify the value of b corresponding to good power extraction performance, we can consider the case where $\zeta \to 0$, since we know optimal performance in this case is for the case $\omega_n \to \omega$. Then, we can evaluate b at this point using Equation 4.25 to obtain

$$b = 1/M_{\max}\omega. \tag{4.26}$$

Substituting this value back in for b, we obtain the relationship

$$\omega_n = \frac{\omega}{\sqrt{1-\zeta^2}}. \tag{4.27}$$

Comparing Equations 4.23 and 4.27, we see that though they are not equivalent, in the region of interest $\zeta \in [0, 1/\sqrt{2}]$, they describe relationships that are extremely similar. Taking the difference between the two and expanding in a Taylor Series, the difference in $\omega_n(\zeta)$ described by Equations 4.23 and 4.27 is $O(\zeta^4)$. Thus, for values of ζ in the range of interest, Equation 4.26 provides a close approximation to the optimal choice for the parameter b for any values of $[M_{\max}, m, \zeta < 1/\sqrt{2}]$ and a given operating point $\omega = \omega_f$. This is perhaps not entirely unexpected, as this corresponds to setting the operating point at the damped natural frequency of the system, according to Equation 4.4. This operating point thus represents a near-optimal compromise between providing a good phase match, and operating at the maximum gain location $\omega_{\max} < \omega_d$ for $\zeta < 1/\sqrt{2}$. From Equation 4.26, given the choice of $M_{\max}$ discussed in previous sections and the operating frequency imposed by vortex shedding, the optimal choice for the damping (for any mass and small damping ratio) is therefore fixed at $b = 75.2$ Ns/m.

4.5 Optimized Energy Harvesting Results: Constrained Mass and Gain

To experimentally verify improvements to the power extraction of a Passive Captive system corresponding to an improved choice of mounting parameters while keeping mass constant, experiments analogous to those described for the conservative case (Case 0) in Section 4.3 were performed for a new virtual mounting system. All aspects of the experimental setup were kept the same, except that the virtual damping and stiffness were altered to correspond to those for Case 1, as described in Table 4.1.

To determine the parameters $[m, b, k]$ for Case 1, potential power extraction of the system was optimized subject to the same mass and amplitude constraints set for Case 0. The location for $\overline{P_D}_{\max}$ along the top constant-mass curve in Figure 4.8 is (in dimensional units) $\omega_n = 4.46$ rad/sec, and $\zeta = 0.45$. This is quite different

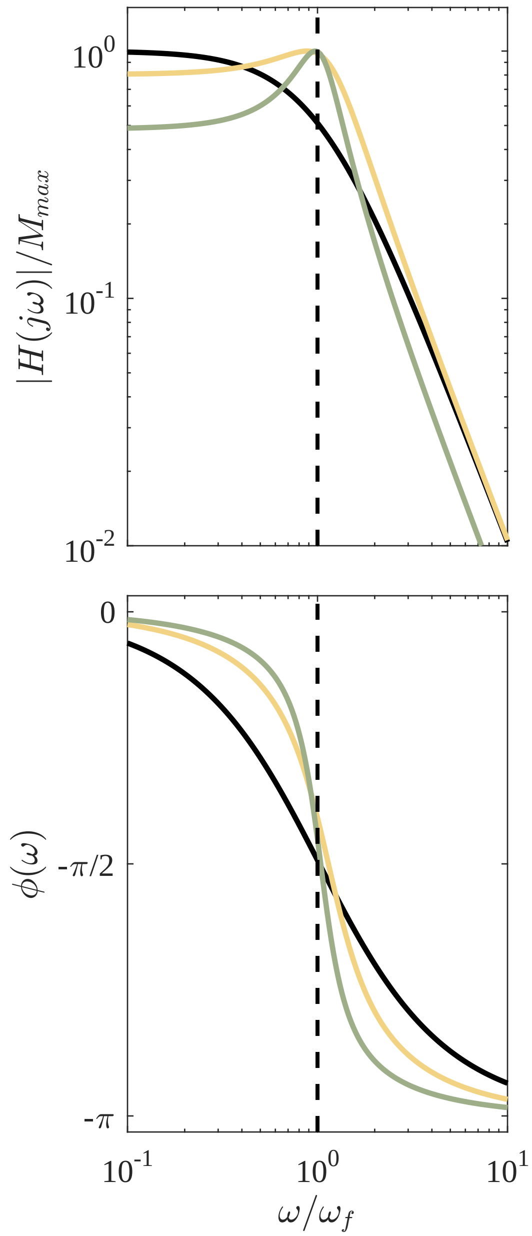

Figure 4.9: Bode plot showing frequency responses for a Passive Captive Airfoil with the three sets of mounting parameters tested. Magnitude responses are normalized by the maximum amplitude gain constraint $M_{\max}$, and frequency axes are normalized by the estimated operating frequency, ω_f. ▬ Case 0, ▬ Case 1 , and ▬ Case 2 mounting parameters as given in Table 4.1. ‒ ‒ Operating condition, $\omega = \omega_f$.

than the the initial conservative parameters - the system is under-damped, and the natural frequency does not coincide with the mean frequency of operation of 3.91 rad/sec. Figure 4.9 shows the frequency and phase responses of the system with Case 0 and Case 1 parameters (as well as those for Case 2, to be discussed in Section 4.6). Comparing the Case 0 (black) and Case 1 (yellow) frequency responses, the optimized Case 1 system exhibits a broad amplitude gain peak in the vicinity of the operating frequency; using Equation 4.8 and substituting values for ζ, ω_n, we find that for Case 1, $\omega_{\max} = 3.44$ rad/sec. The broadness of the magnitude response however leads to an elevated gain at the operating frequency ($\omega_f = 3.91$ rad/sec) as compared to Case 0: the gain at this frequency is approximately double that of the Case 0 system. The trade-off for this improved magnitude response is a deviation from the optimal phase offset value of $-\pi/2$ radians, which would lead to a velocity which is perfectly in-phase with forcing. Comparing the phase responses for Cases 0 and 1, the deviation from this optimal phase for Case 1 is approximately +0.29 radians. Considering the sine dependency on the phase difference in Equation 4.18, the reduction in power extraction potential due to this deviation is $\sin(\phi) = 0.96$, or a 4% reduction compared to a perfect phase match. Clearly, the improvement in the gain at the frequency of interest more than compensates for the deviation in phase between force and velocity. From Figure 4.8, we expect that the cycle-averaged power dissipation for our optimized Case 1 system will be approximately 180% of that for the conservative Case 0 system, in alignment with the above discussion.

The behaviour of the airfoil with Case 1 parameters is qualitatively similar to that for Case 0: the airfoil still undergoes quasi-harmonic oscillations that closely track the incoming forcing signal from upstream vortex shedding. However, the amplitude of the motion is much larger, as shown in Figure 4.10. The position achieves maxima approaching $2A_0$, with a corresponding increase in airfoil velocity. In addition, there is a phase shift apparent between velocity and force. Computing the cross-correlation, the phase shift between the signals shown in the figure is approximately +0.22 radians, with velocity leading force. This is quite close to the expected phase shift of +0.29 radians from the previous discussion.

Figure 4.11 shows experimentally determined phase-averaged quantities for the Case 1 mounting system, computed in the same fashion as described in Section 4.3 for Case 0. Figure 4.12 confirms that the amount of data averaged per phase bin is also similar to Case 0, with a mean number of frames per bin of approximately 95, and no bins with fewer than 70 frames. From Figure 4.11, we confirm that the mean power coefficient per cycle is larger, taking a value of 0.044 as compared to the coefficient of 0.032 reported for Case 0. This constitutes an increase in power of approximately 140% from the naïve choice of mounting parameters, while maintaining a physically realistic system and constraining the maximum gain as well as the mass.

Though the power extraction potential for the Case 1 system constitutes a clear

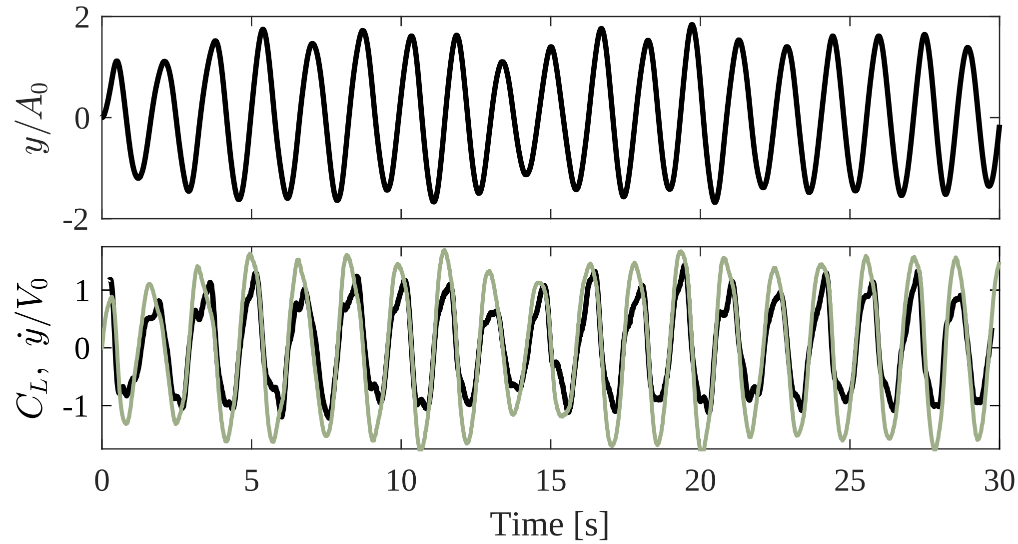

Figure 4.10: Time series behaviour of the Passive Captive Airfoil with Case 1 mounting parameters. Top: Airfoil transverse position, normalized by the standard amplitude for the driven experiments discussed in Chapter 3. Bottom: Lift Coefficient (▬), and Airfoil velocity normalized by the standard amplitude for the driven experiments discussed in Chapter 3 (▬).

improvement, it was not able to realize the expected increase in power of approximately 180% relative to the Case 0 system. This shortfall may be due to fluid dynamical effects that were not taken into account in the development of the prediction for $\overline{P_D}$. In particular, it was assumed that the force experienced by the airfoil $F(t) = F_0\cos(\omega_f t)$ is the same for all mounting parameters. In reality, increasing the airfoil's velocity in-phase with the force leads to a reduction in the apparent angle of attack, defined again for convenience here as

$$\alpha_{\text{eff}}(t) = \arctan\left[\frac{V - \dot{y}(t)}{U_\infty}\right] \tag{4.28}$$

Equation 4.28 shows that as the velocity of the airfoil ($\dot{y}$) increases in-phase with the flow velocity in the airfoil's vicinity (V), the effective angle of attack is reduced.

Making a quasi-steady assumption, which for a uniform oncoming flow and prestall effective angle of attack is justified, this reduced apparent angle of attack leads directly to a reduction in lift force generated by the airfoil. However, from the discussion of the Case 0 system in the previous section, for our highly dynamic oncoming flow this clearly does not provide a complete picture of the forces affecting the airfoil. To facilitate similar observations as made for Case 0, we again consider snapshots of the phase-averaged Γ_2 criterion, shown in Figure 4.13.

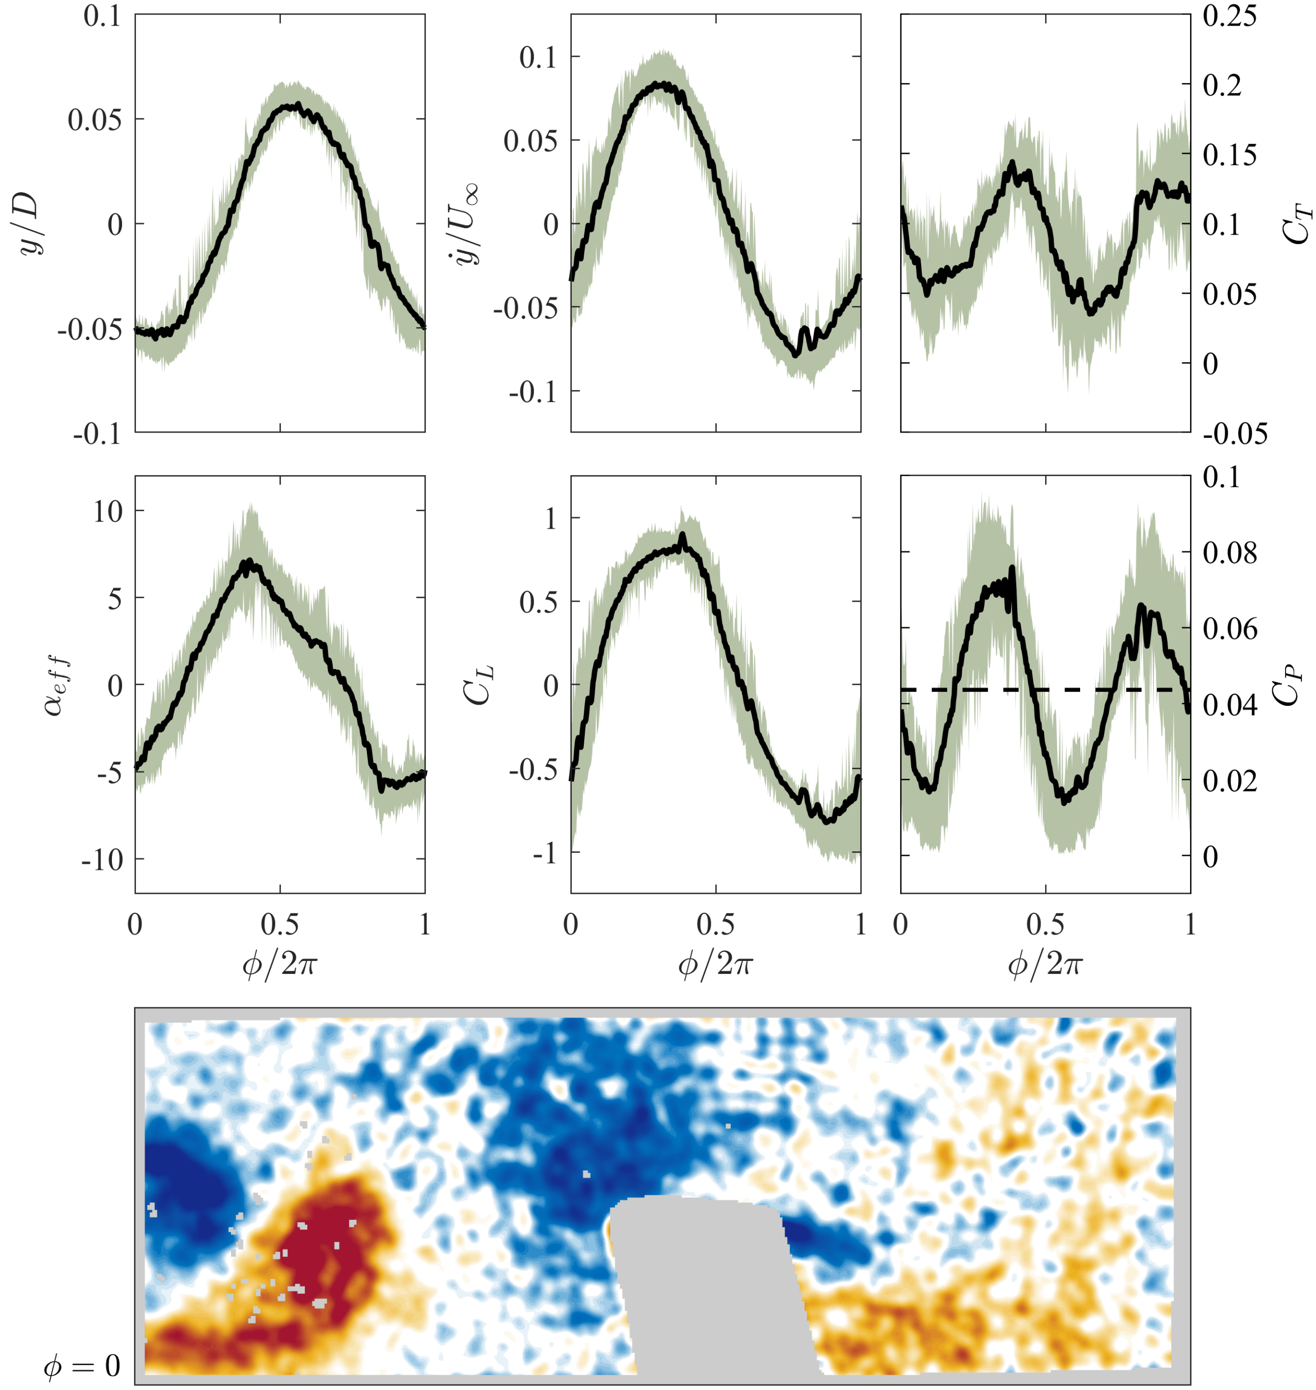

Figure 4.11: Phase-averaged dynamical quantities for the Case 1 Passive Captive Airfoil. Top Left: Airfoil Position. Top Center: Airfoil Velocity. Top Right: Thrust Coefficient. Middle Left: Effective Angle of Attack. Middle Center: Lift Coefficient. Middle Right: Power Coefficient. Bottom Panel: The flow state at $\phi = 0$, shown here as contours of the Γ_2 criterion. Shaded regions in the top 5 figures indicate the interquartile range for data averaged in each phase bin. Bold black lines indicate the mean values recorded in each bin, with a Savitzky-Golay filter applied to smooth the curves. The filter width was set to 5°.

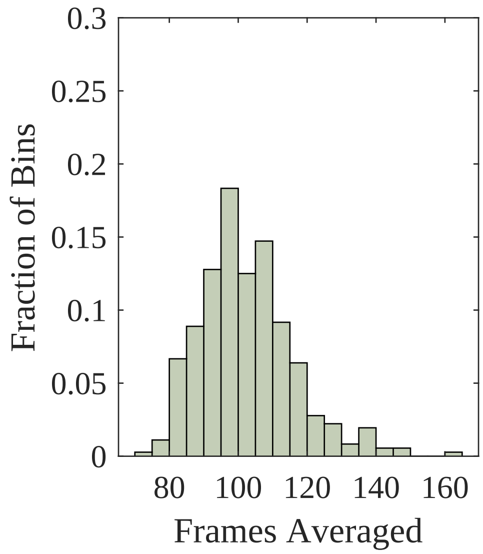

Figure 4.12: Histogram showing the fraction of phase bins with each number of observations included in phase averaging for Case 1.

The phase difference between force and velocity is noticeable in the Figure, with the extrema in lift, α_{eff} and power all occurring later than the extrema in velocity. This is as expected, as our Case 1 system purposefully operates away from $\omega = \omega_n$, which induces this phase shift.

The phase shift does not appear to interfere with the basic pattern of interaction between the airfoil and the oncoming vorticity observed in either Chapter 3 or for the Case 0 Passive Captive Airfoil, as it is very small compared to one cycle. A slaloming mode of interaction is again observed, with the airfoil moving to wind between the oncoming vortices. Although the amplitude of motion in Case 1 is large compared to that of Case 0, it is still only a small fraction (about 10%) of the size of the cylinder D and its shed vortices; therefore the motion of the airfoil relative to the wake structures is not that different from Case 0 to Case 1, though there is evidence that the vortex splitting that was observed in Case 0 is less pronounced here.

One difference is that for the faster-moving airfoil in the present case, both the CW vorticity bound to the suction side of the airfoil as well as the CCW TEV appear to have reduced intensities as compared to Case 0 when the airfoil is moving up (panels 1-3 in the Figure). Thus, compared to Case 0 the increased airfoil velocity in Case 1 appears to further suppress the formation of LEVs and the buildup and subsequent shedding of trailing edge vorticity, which in concert with the reduction in effective angle of attack suppress the lift force experienced by the airfoil. These relationships are particularly pronounced in Figure 4.14, which shows the lift and power coefficients as a function of α_{eff}, as well as the relationship between force and velocity.

The central panel of the figure confirms that there is a slight phase shift between lift and velocity, again as expected due to the difference between ω_n and ω_f for this case. Although a reduction in the lift coefficient is apparent as compared to

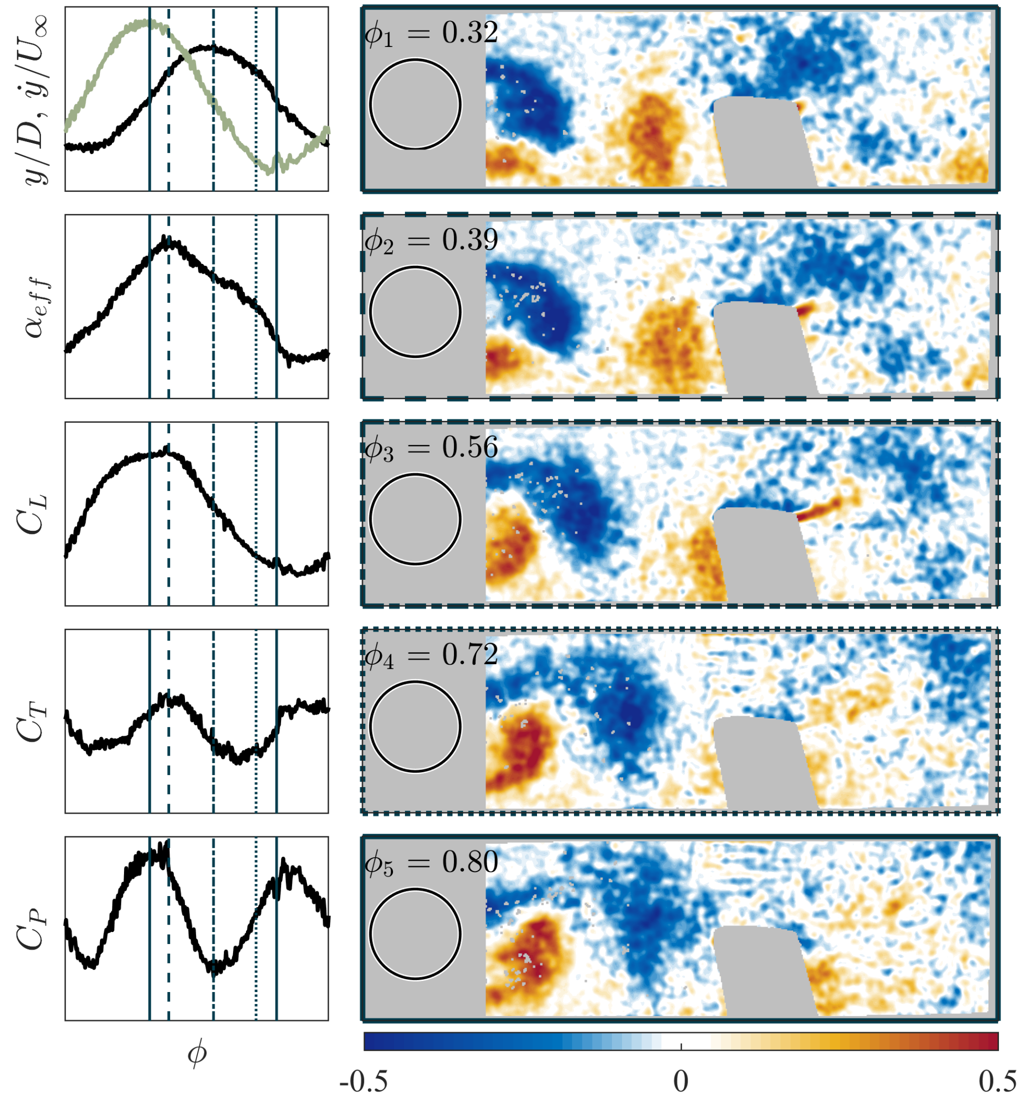

Figure 4.13: Snapshots of phase-averaged Γ_2 Criterion values at indicated points in the vortex shedding cycle, for an airfoil with Case 1 mounting parameters. Colour bar indicates contour level for the criterion, which is constrained so that $|\Gamma_2| \leq 1$. Left column reproduces data for phase-averaged quantities from Figure 4.11. In the top left panel: ▬ y/D; and ▬ $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angle ϕ is presented as a fraction of one cycle (2π). ▬ $\phi/2\pi = \phi_1 = 0.32$; ▬ ▬ $\phi/2\pi = \phi_2 = 0.39$; ▬ · $\phi/2\pi = \phi_3 = 0.56$; ··· $\phi/2\pi = \phi_4 = 0.72$; ▬ $\phi/2\pi = \phi_1 = 0.80$.

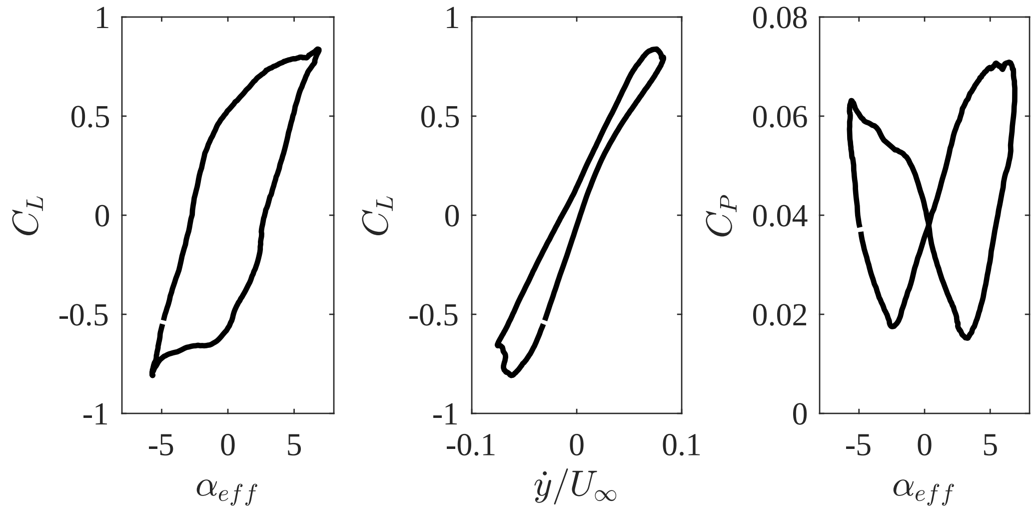

Figure 4.14: Phase portraits showing the behaviour of the Case 1 Passive Captive Airfoil. A Savitzky-Golay filter with a width of 10% of T was used to smooth phase-averaged data to improve clarity.

Case 0, the shape of the lift cycle is similar, with hysteresis observed again largely due to the vortex suction from the oncoming wake. In the regions of peak lift in Figure 4.14 there appears to be a plateau towards the maximum value, which is in contrast to the behaviour in Case 0 which exhibited a stronger peak (with the pre- and post-maxima lift values remaining more similar). The increased hysteresis here indicates a broadening of the lift curve relative to Case 0, induced by the reduction of secondary helpful effects (TEV shedding, and bound CW vorticity on the airfoil's surface during deceleration). Despite this reduction in useful lift, the increase in airfoil velocity more than compensates, ensuring that the power production is high throughout the cycle compared to Case 0.

4.6 Optimized Energy Harvesting Results: Increased Mass

To further investigate the interaction between linear systems theory results presented in Section 4.4 and the nonlinear fluid forcing, experiments were carried out with the Case 2 mounting parameters, as described in Table 4.1. These parameters represent a system with the same maximum amplitude gain constraint, but with double the mass of the Case 0 and 1 systems. As discussed in Section 4.4, the damping coefficient b is approximately equal for both of the optimized cases. The experiments for Case 2 were performed under the same conditions as in Cases 0 and 1, with the only change being the virtual mounting parameters.

Comparing the frequency and phase responses shown in Figure 4.9, the behaviour of

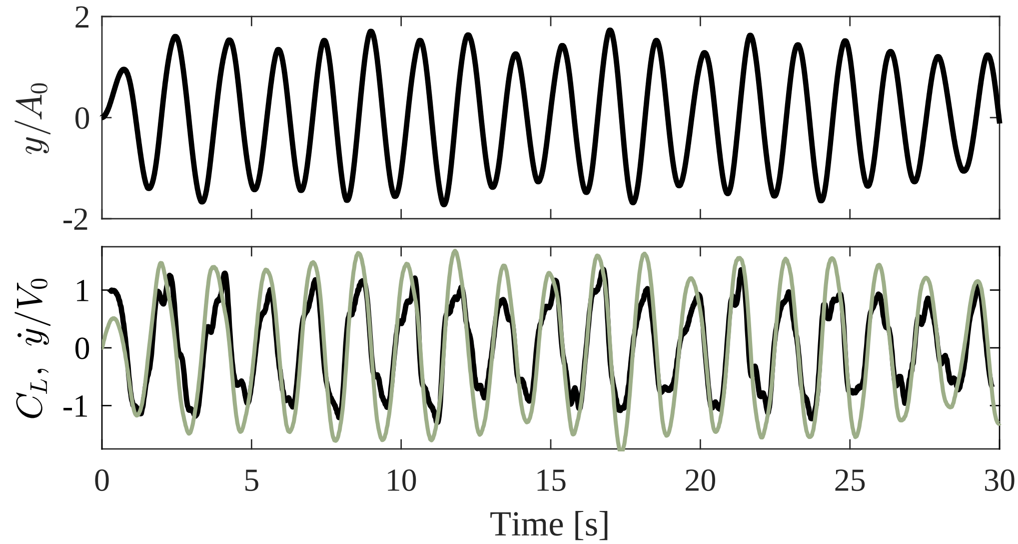

Figure 4.15: Time series behaviour of the Passive Captive Airfoil with Case 2 mounting parameters. Top: Airfoil transverse position, normalized by the standard amplitude for the driven experiments discussed in Chapter 3. Bottom: Lift Coefficient (▬), and Airfoil velocity normalized by the standard velocity amplitude for the driven experiments discussed in Chapter 3 (▬).

the Case 2 system should be a compromise between the behaviours of the Case 0 and Case 1 systems. The magnitude response in the vicinity of ω_f is nearly identical to Case 1, but the phase shift between force and velocity has been reduced to effectively zero, similar to Case 0. Despite these similarities in the response, from Figure 4.8, we expect that the power extracted from this system should be approximately 190% that of the Case 0 system, with a marginal but measurable gain in power extraction relative to Case 1. The trade-off for achieving this enhanced power extraction potential is that the range of input frequencies that produce a desirable response is much more limited.

Figure 4.15 shows typical behaviour for the airfoil with Case 2 mounting parameters, confirming that the amplitude of the motion in this case is qualitatively very similar to Case 1. Again, the amplitude of the motion is approximately double that in the Driven case discussed in Chapter 3, which is also reflected in the elevated velocity. The phase match between force and velocity is very good, as expected.

We again examine a phase-averaged picture of the airfoil behaviour, in the same manner as in previous sections. Again, data are binned by phase, and the mean number of observations per bin is approximately 95. Figure 4.16 summarizes the phase-averaged results. No bins considered have fewer than 60 observations, as

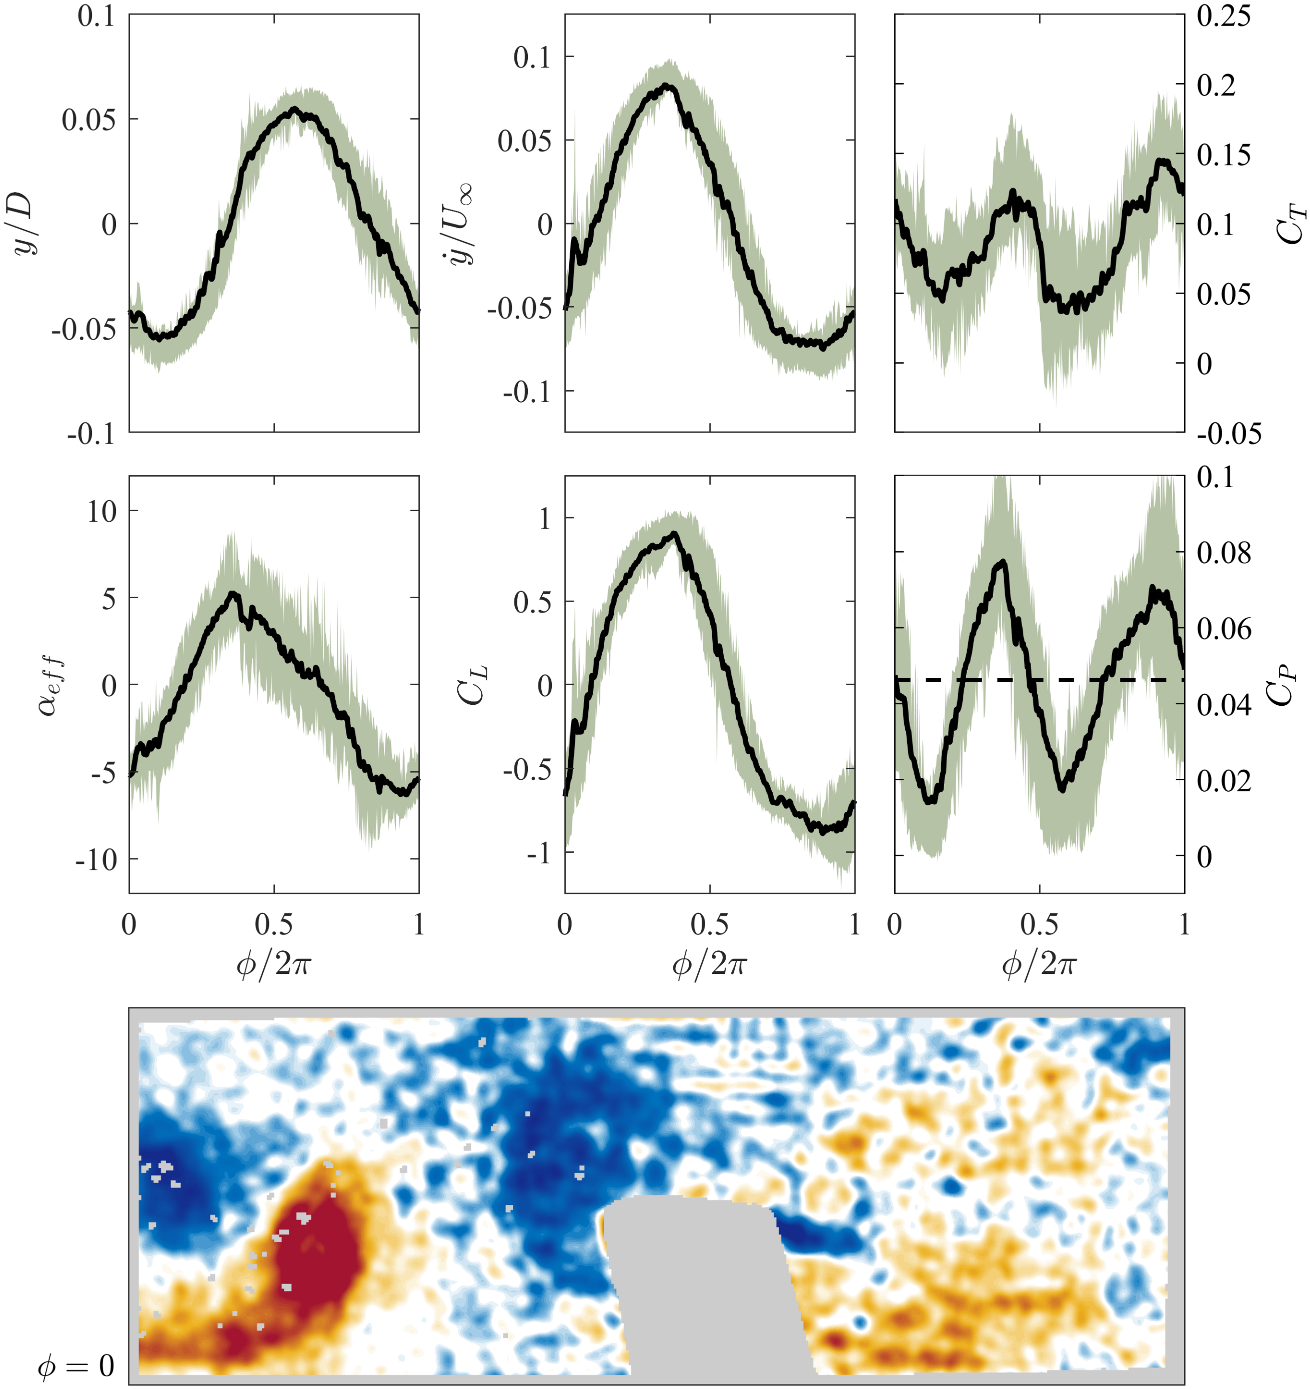

Figure 4.16: Phase-averaged dynamical quantities for the Case 2 Passive Captive Airfoil. Top Left: Airfoil position. Top Center: Airfoil Velocity. Top Right: Thrust Coefficient. Middle Left: Effective Angle of Attack. Middle Center: Lift Coefficient. Middle Right: Power Coefficient. Bottom Panel: The flow state at $\phi = 0$, shown here as contours of the Γ_2 Criterion. Shaded regions in the top 5 figures indicate the interquartile range for data averaged in each phase bin. Bold black lines indicate the mean values recorded in each bin, with a Savitzky-Golay filter applied to smooth the curves. The filter width was set to 5°.

shown in Figure 4.17.

Comparing the power extracted by the Case 2 system to that of the Case 0 system, the 190% increase expected is not realized. The mean power coefficient for Case 2 is 0.048: compared to the mean C_p in Case 0, this represents an approximately 150% increase. Though we do observe a marginal improvement in power production relative to Case 1, as expected, it is clear that fluid-dynamic factors are again acting to reduce the power extracted relative to linear theory, as described in Section 4.5.

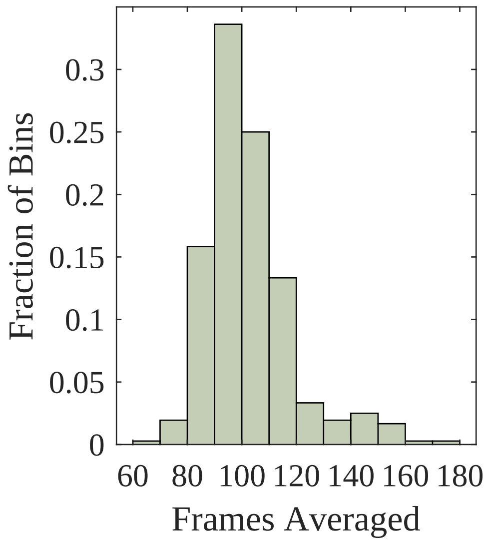

Figure 4.17: Histogram showing the fraction of phase bins with each number of observations included in phase averaging for Case 2.

To investigate this further, Figure 4.18 shows sequential frames of the Γ_2 Criterion for the Case 2 system, similar to the figures presented for Case 0 and Case 1 in the previous sections. By contrast to either of those cases, the first panel in Figure 4.18 represents a moment when the airfoil is experiencing simultaneously its maximum velocity, effective angle of attack, lift, and power production. In the previous cases, the velocity was found to lead the maxima of these other quantities by a small amount in Case 0, and a more significant one in Case 1 (due to its operating point away from ω_n). The behaviour however is otherwise qualitatively similar to the previous two cases. As the airfoil is travelling more quickly than in Case 0, we again see that it is able to more effectively slalom between the oncoming vortices.

As in Case 1, in the third and fourth panels which show the airfoil at its minimum velocity (maximum position extent) and minimum effective angle of attack respectively, show a reduced strength and extent of vorticity generated near the airfoil's surface. This is particularly apparent in the fourth panel as the airfoil's top surface transitions from suction to pressure, and the vorticity it had generated is shed. The magnitude of the vorticity shed from the trailing edge in Case 2 is also noticeably attenuated compared to both Case 0 and Case 1. Since the airfoil's motion is most closely aligned to that of the surrounding fluid in this case (due to its improved phase match relative to Case 1), a reduction in shed trailing edge vorticity (again due to a

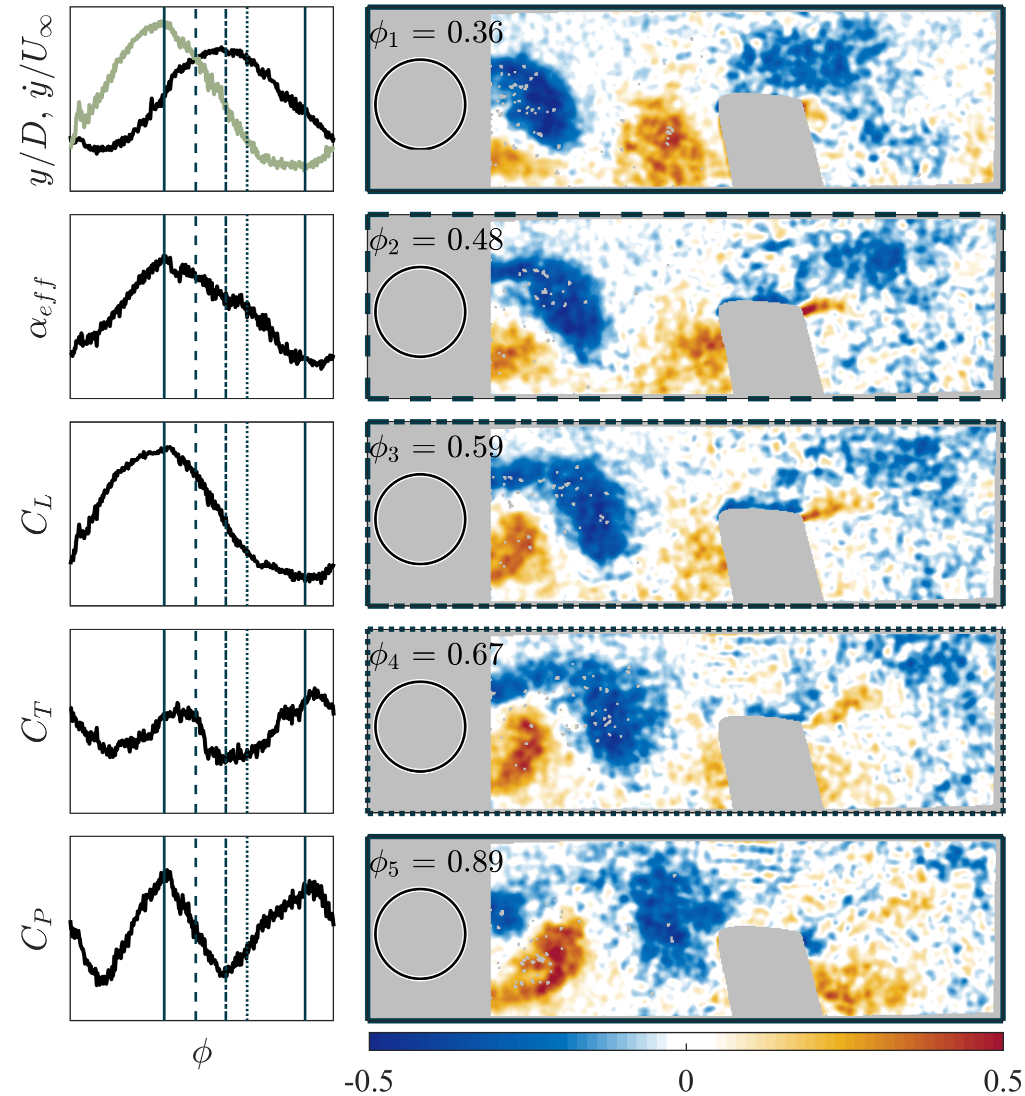

Figure 4.18: Snapshots of phase-averaged Γ_2 Criterion values at indicated points in the vortex shedding cycle, for an airfoil with Case 2 mounting parameters. Colour bar indicates contour level for the criterion, which is constrained so that $|\Gamma_2| \leq 1$. Left column reproduces data for phase-averaged quantities from Figure 4.16. In the top left panel: y/D; and $\dot{y}/U_\infty$. Left-hand figures also indicate the phase of each snapshot shown on right, where the border line type corresponds to the indicated moments in phase. Phase increases from top to bottom in the right-hand column. Here, the phase angle ϕ is presented as a fraction of one cycle (2π). $\phi/2\pi = \phi_1 = 0.36$; $\phi/2\pi = \phi_2 = 0.48$; $\phi/2\pi = \phi_3 = 0.59$; $\phi/2\pi = \phi_4 = 0.67$; $\phi/2\pi = \phi_1 = 0.89$.

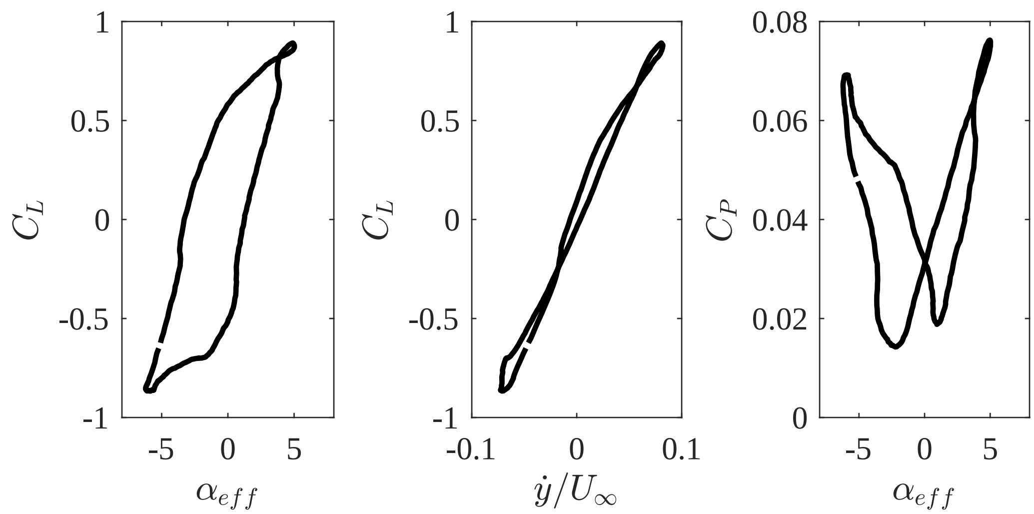

Figure 4.19: Phase portraits showing the behaviour of the Case 2 Passive Captive Airfoil. A Savitzky-Golay filter with a width of 10% of T was used to smooth phase-averaged data to improve clarity.

corresponding reduction in α_{eff}) is expected.

Figure 4.19 shows phase portraits of the behaviour of the Case 2 system. The effect of the improved phase match is clearly visible relative to that displayed in the corresponding figures in Section 4.5: there is little to no misalignment between the lift and velocity seen in the middle panel. In addition, the lift hysteresis is noticeably reduced as compared to Case 1, especially in the regions of maximum lift amplitude. This directly leads to higher power outputs over a larger portion of the cycle, as shown in the third panel of that figure. This reduction in hysteresis appears responsible for the marginal power production gains from Case 1 to Case 2.

To examine the interactions between shed vorticity from the airfoil surface and that from the upstream circular cylinder, Figure 4.20 presents a frozen-flow visualization of the wake in the combined airfoil-cylinder region for all three discussed cases. The most notable change in the wake structures from top to bottom is the intensity of shed trailing-edge vorticity. As the velocity of the airfoil conforms more closely to the motion of the surrounding fluid, the apparent acceleration of the airfoil is reduced, and there is a corresponding reduction in TEV shedding. This also has implications for the shape of the lift curves experienced by the airfoil in all three cases, as discussed earlier in this section and the previous ones. It is particularly interesting to compare the wake structures shed by the Case 1 and Case 2 airfoils, which are undergoing very similar quasi-harmonic motions separated only by a small phase shift. In the Case 1 wake field, the large CW vortices that pass over

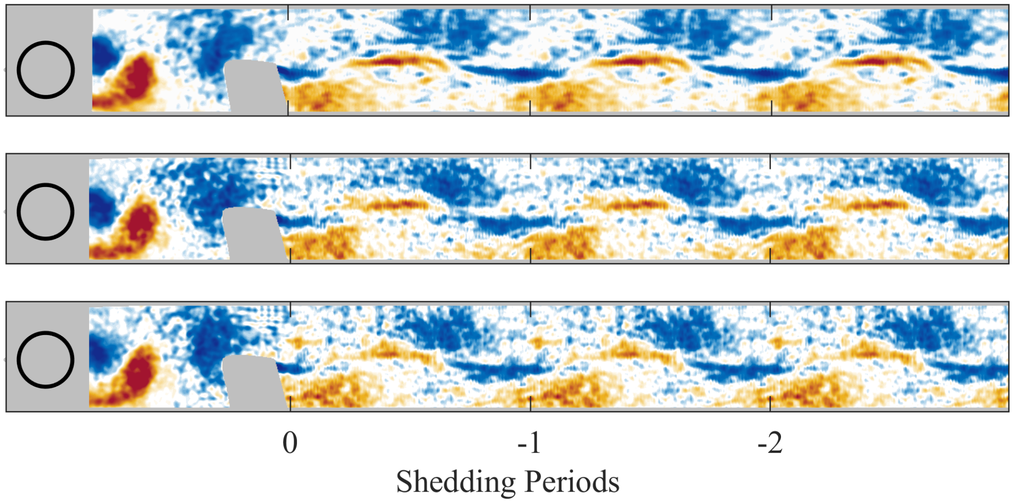

Figure 4.20: Frozen flow wake visualization for Case 2 mounting parameters, as compared to Case 0 and Case 1. Top Panel: Wake visualization for Case 0, reproduced from Section 4.3. Middle Panel: Wake visualization for the Case 1 Passive Captive Airfoil. Bottom Panel: Wake visualization for the Case 2 Passive Captive Airfoil. All three panels show contours of the Γ_2 Criterion at the same level.

the airfoil seem to stretch around the trailing-edge vorticity that they pair with as shedding occurs, where as this effect is not as pronounced for either Case 0 or 2. This implies that trailing-edge vorticity is shed slightly sooner in the cycle for Case 1: this is consistent with the slight lead in the airfoil's velocity relative to the oncoming flow. The slightly earlier shedding could be reducing the effect of vortex suction from the cylinder vortices as they pass by the airfoil slightly earlier than the other cases, which causes a reduction of lift slightly earlier than is optimal. However, this effect, if present, is very small.

4.7 Thrust Production with the Spring-Mass-Damper Mounting System

The previous sections describe three configurations for a passive captive airfoil that all exhibit simultaneous net thrust and power production, purely through interactions with oncoming vorticity. The development in this chapter has focused primarily on the power extraction potential for this system, rather than the production of thrust: this section presents a justification for this imbalance.

The principal generator of thrust for the systems considered in the previous sections is the Katzmayr effect, discussed at length in Chapter 3. This is a strictly geometric effect, that allows an airfoil to take advantage of an incoming flow at an angle to its forward direction. The Katzmayr Effect has been shown to provide net thrust larger than net drag for low-drag, lifting bodies in wavy streams as early as its first observation in the 1920s (Katzmayr, 1922). The effect is dependent on the effective angle of attack of the the airfoil relative to the free-stream, and thus in our discussions here is sensitive to the airfoil motion. The optimal case for the airfoil to produce thrust through the Katzmayr effect is the case where it is not moving, or when α_{eff} is maximized. Thus, a more interesting avenue of inquiry is to consider the thrust production *beyond* the Katzmayr effect, specifically due to the airfoil's motion and its interaction with oncoming vorticity.

As discussed in Chapter 1, previous studies have investigated the link between the phase of driven airfoil motion relative to oncoming vorticity shed by a von Kármán Vortex Street, and the resulting propulsive efficiency. For example, Streitlien et al. (1996) used an analytical model to demonstrate that thrust and efficiency for an airfoil driven in such a wake are simultaneously maximized, while input power required to sustain the airfoil's motion is minimized at a location in phase 180°away. For those authors, a phase angle $\phi = 0$ corresponded to an 'interception mode' where the airfoil's motion is always oriented towards the oncoming vortex cores. This causes significant changes to the topology of the wake, and the authors show that for this mode of interaction, those changes tend to result in an increase in thrust experienced by the airfoil. That increase comes at the cost of increased input power required to maintain the motion; however the gain in thrust is significant enough that this mode of interaction results in optimal efficiency for the airfoil, where the efficiency is given by

$$\eta = \frac{\overline{F_x} U_\infty}{-\overline{P}}, \tag{4.29}$$

reproduced from Chapter 1. The negative sign here accounts for the difference between *input* power, of interest in driven systems, and the *output* power, or work per time done by the fluid on the airfoil, which has been described as $\overline{P_D}$ throughout this work.

By contrast, the authors identified another regime at $\phi = 0$ where the airfoil moved to avoid the oncoming vortex cores; this is the slaloming mode identified by Beal et al. (2006), in which the airfoil moves with the surrounding fluid and leaves relatively little impact on the downstream wake. This slaloming mode corresponds to the condition for maximum power extraction through alignment of the transverse force and velocity, as demonstrated in the preceding sections as well as for the driven case in Chapter 3. As the airfoil moves through the vortex street, it experiences lift in-phase with its velocity due to interactions with alternating low-pressure vortex core regions, which augment suction due to leading-edge or trailing-edge vortex formation from airfoil movement (if any). Since in our study, the maximum y-velocity of the airfoil is much smaller than the velocity of the surrounding fluid, this results in $P > 0$ over the whole cycle, and the airfoil extracts the net excess energy through damping ($\overline{P_D} > 0$).

Since this is a purely passive system, assuming a steady-state is reached it must be the case that cycle averaged power extracted from the flow is greater than or equal to zero ($\overline{P} \geq 0$) with the excess dissipated through mechanical damping, since there is no mechanism by which the airfoil can do net work on the flow. It is not the case however that there can be no *instantaneous* work done on the flow: energy storage in the form of spring energy allows for moments throughout a cycle where $P < 0$. This is encoded in the phase response for the mounting system, given for the various cases considered here in Figure 4.9. The range of achievable phase shifts between the airfoil motion and the measured forcing is $[0, -\pi]$, with the operating condition for maximum power extraction, or the slaloming mode of interaction located close to $\phi = -\frac{\pi}{2}$ (this corresponds to a phase location of $\phi = \pi$ for Streitlien et al. (1996)). According to that study, to transition from the location of minimum power input to that of maximum thrust, the phase of interaction between the airfoil and the surrounding flow needs to be offset by approximately 180°. This is not possible with our simple mounting parameters, as this offset is outside the range of the phase response. This makes sense, as such a shift would result in the force and velocity being perfectly out of phase, or $\overline{P_D} < 0$, which is a physical impossibility for a purely passive system with second-order linear dynamics.

Therefore, the best we could hope to achieve with the current mounting setup is a misalignment between force and velocity of $\pm\pi/2$, which would ideally lead to $\overline{P_D} = 0$, with all of the energy extracted from the flow reinvested as work done by the airfoil. Though there is evidence that this operating condition leads to enhanced thrust compared to the maximum power extraction case (see for example Kinsey and Dumas, 2008; Streitlien et al., 1996; Gopalkrishnan et al., 1994) for a driven system, its implementation for the passive system has several practical challenges. Considering the magnitude function for the mounting systems shown in 4.9, we see that by selecting an operating frequency to achieve a phase offset approaching $-\pi$ results in a quickly diminishing amplitude of the response, potentially negating any benefits to thrust relative to the Katzmayr case by reducing the amplitude of airfoil motion (and therefore, ability to reposition vorticity, as is necessary to achieve enhanced thrust). On the other hand, the magnitude of the response remains large as the operating frequency is reduced; however in order for the fixed operating frequency to become small relative to ω_n, the natural frequency must be increased. This corresponds either to increasing k or reducing m (which is undesirable in a practical setting), according to Equation 4.2. Unfortunately, increasing k again directly reduces the amplitude of the response for small operating frequencies, according to Equation 4.5. Thus, it is clear that 2nd order dynamics, while very useful for resonant energy harvesting, are significantly more resistant to optimization for thrust production over and above that which is achieved via the Katzmayr effect.

4.8 Chapter 4 Interim Summary and Conclusions

In this chapter, the behaviour of a simulated fully passive flapping foil was interrogated using a Cyber-Physical Fluid Dynamic (CPFD) approach. First, the behaviour of a similar driven system discussed in the previous chapter was reproduced through tuning of the virtual mounting parameters responsible for the airfoil's behaviour. Then, mounting parameters were optimized to improve the energy extraction performance based on linear systems theory, while respecting realistic engineering constraints on system behaviour. Further experiments were then performed to confirm that such improved performance was realized. The following conclusions can be drawn from this work:

1) The behaviour of a fully passive flapping foil and its interactions with the oncoming vorticity shed by an upstream circular cylinder can be made both qualitatively and quantitatively similar to that for a similar foil actively driven through a sinusoidal trajectory. Such a driven system then provides a reference system with a simpler experimental implementation with which to study the behaviour of a passive foil.

2) A fully passive flapping foil energy harvester is capable of both extracting net energy from the flow while producing net thrust larger than its net drag through interactions with oncoming vorticity shed from an upstream circular cylinder.

3) Tuning the behaviour of the system through adjusting the properties of the mounting system can lead to improved airfoil performance in terms of power

extraction potential, while still respecting relevant engineering constraints on system behaviour.

4) The frequency response of a linear mounting system can be used to predict the power extraction performance of a passive flapping foil energy harvester; however changes to airfoil behaviour lead to changes to the fluid-structure interactions taking place in the flow, which feed back (potentially non-linearly) to reduce the achievable performance of such systems.

Chapter 5

PASSIVE CAPTIVE MOTION IN 2 DIMENSIONS: AN AIRFOIL SWIMS UPSTREAM

5.1 Introduction

A natural extension of the one-dimensional Passive Captive Airfoil motion described in Chapter 4 is to consider the effect of allowing the airfoil to translate in the streamwise direction as it generates thrust, that is, allowing the airfoil to swim upstream. This undertaking was inspired by the dead fish experiments by Beal et al. (2006), which have been discussed throughout the thesis. Those authors observed a dead fish interacting with oncoming vorticity to produce upstream motion, and showed that an airfoil under similar conditions produced net thrust larger than net drag. They could not however demonstrate the airfoil actually translating upstream, due to experimental constraints (Beal et al., 2006). The present study seeks to rectify this, by demonstrating passive captive motion of an airfoil in 2 dimensions. In addition to the satisfaction of the author's curiosity, this configuration is also of some practical interest. As discussed further in Chapter 1 and by authors such as Streitlien et al. (1996) or Jarman et al. (2019), the ability for a vehicle to not only safely navigate the unsteady environment downstream of an obstacle but also to extract useful thrust from it is highly practical when operating in such environments on an ongoing or repeated basis.

This chapter presents experiments completed using the Captive Trajectory System (CTS) to realize 2-dimensional captive trajectories for an airfoil downstream of a circular cylinder in the same configuration as discussed in previous chapters. Motion in the transverse direction was governed by the same dynamics as described in Chapter 4, while the streamwise dynamics were programmed to follow $F_x = ma_x$. This posed some practical challenges, as these dynamics are unstable to the positive mean thrust experienced by the airfoil. In addition, a single sensor was used for feedback in both the transverse and streamwise directions. Since the magnitude of the forces experienced in the transverse direction are in general more than an order of magnitude larger than those in the streamwise direction, selecting a sensor with appropriate dynamic range and noise characteristics poses an obstacle. Despite these challenges, repeatable behaviour is observed as the airfoil translates upstream, while oscillating transversely under the action of oncoming vorticity.

The purpose of this chapter is to describe the 2-dimensional Passive Captive Airfoil Experiments conducted, and to characterize the airfoil's behaviour during these tests. Simultaneous Particle Image Velocimetry (PIV) data allows for direct visualization of interactions with oncoming vorticity during upstream swimming; such interactions are compared to similar interactions in the 1-dimensional case. Similarities and differences between the two sets of experiments are explored.

5.2 Experimental Setup for 2-Dimensional Airfoil Motion

The setup for Passive Captive Airfoil Experiments in 2 dimensions was very similar to that for the 1-dimensional case, but without the restriction that the airfoil only move in the transverse direction. Full details regarding the basic experimental setup are provided in Chapter 2. The below sections describe specifications unique to the 2D trajectory experiments.

5.2.1 Equations Governing 2-Dimensional Airfoil Motion

For the 2-dimensional experiments described in this chapter, the airfoil exhibited passive captive motion in both the streamwise and transverse directions, here denoted x and y and pictured in Figure 5.1. The motion in y was governed by the same set of equations as described for the 1-dimensional case in Chapter 4, with the conservative mounting parameters (Case 0) as described therein. In the streamwise or x direction, the airfoil obeyed the equations of motion

$$F_x(t) = m\ddot{x} + b_x\dot{x}, \tag{5.1}$$

where a very small simulated damping $b_x \ll b_y$ was added to make the system more stable, and to limit the magnitude of airfoil motion caused by low-frequency

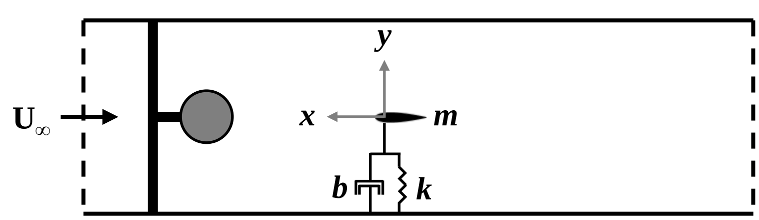

Figure 5.1: Basic Schematic showing the airfoil-cylinder system with spring-mass-damper dynamics in the transverse (y) direction, and free motion in the streamwise (x) direction.

Case Name	ω_n [1/s]	ζ	$M_{\omega_{\max}}$ [m/N]	$m_{x,y}$ [kg]	b_x [Ns/m]	b_y [Ns/m]	k_x [N/m]	k_y [N/m]
2D	4.00	1.00	0.0034	18.38	1.00	147.06	0.00	294.12

Table 5.1: Parameters used to specify dynamics for the 2D airfoil motion discussed in this chapter. Note that the virtual mass $m_{x,y}$ is the same for motion in the x and y directions.

sensor drift and zero-point set errors. In addition, the force signal was filtered using a 10-point causal moving average filter (a filter length of 0.05 s) to limit the effect of sensor and mechanical noise as well as feedback due to structural bending on airfoil motion when the measured forces are close to zero. The linear phase lag induced by this filter is approximately 5° relative to the period of vortex shedding, and should therefore play a relatively minor role in our dynamics (as discussed by Mackowski and Williamson (2011) and Su and Breuer (2019)). The mounting parameters implemented on the CTS for the 2D motion are summarized in Table 5.1. Further discussion of the selection of the virtual damping b_x is included in the following section.

5.2.2 System Characterization: Mounting Parameters

To better understand the dynamics under test, it is illustrative to consider the frequency responses in both the x and y dimensions. For the transverse motion, the airfoil's behaviour is the same as that discussed at length in Chapter 4; however, due to the lack of a virtual spring force in the streamwise direction, the behaviour in response to x-direction forcing is quite different.

Paralleling the discussion of the second-order system presented in the previous chapter, we consider the response of the system in the x direction to a harmonic forcing, writing the transfer function associated with Equation 5.1 as

$$H(j\omega) = \frac{1}{-m\omega^2 + jb_x\omega}. \tag{5.2}$$

Extracting from this the magnitude and phase responses, we have

$$|H(j\omega)| = \frac{1}{\omega\sqrt{m^2\omega^2 + b_x^2}}, \tag{5.3}$$

$$\phi(\omega) = \text{atan}\left[\frac{-b_x}{-m\omega}\right], \tag{5.4}$$

where signs are preserved in Equation 5.4 to ensure that the correct quadrant is selected for computation of the phase angle. The biggest difference between the responses in the x and y directions is the presence of a pole for non-zero damping in Equation 5.2 at $\omega = 0$, which corresponds to the theoretically infinite position that would be achieved by the airfoil subject to a constant force for all time.

To simulate the behaviour of a body free to move subject to forces acting on it, the ideal behaviour is that given by $F = m\ddot{x}$. This corresponds to the case in Equations 5.1-5.4 where $b_x = 0$, a system without artificial damping. However, for the reasons mentioned in the previous section, this damping plays a practical role in ensuring good behaviour of the system under test. Therefore, we characterize the effect of this additional virtual damping by comparing the frequency responses of systems with and without this addition. Figure 5.2 shows the frequency response for the model described by Equations 5.1-5.4 for three different values of the virtual damping b_x, along with the ideal $b_x = 0$ case. From the figure, we see that the principal effect of adding the virtual damping is to reduce the magnitude of the response to low-frequency signals, as well as to change the phase of the response relative to the input forcing at those frequencies.

Importantly, the addition of the virtual damping b_x does not change the airfoil's response to high-frequency inputs such as sensor noise. As discussed in the previous section, for the transverse direction implementing second-order dynamics is akin to placing a practically zero-phase-lag low-pass filter on the input force, since the large phase lag generated by the dynamics is actually the desired transverse behaviour in the region of our known operating frequency. For the streamwise direction, the same reduction in magnitude of sensor noise is achieved even without the addition of b_x as the dynamics are still of second order; however the phase shift corresponding to the addition of b_x creates a potentially undesirable phase lead in the region of our operating frequency. The bottom panel of Figure 5.2 shows the phase as a function of frequency for values of b_x ranging from 0.1 to 10 Ns/m, all less than 5% of b_y. The grey horizontal dashed line in the figure indicates the location of a 5° phase lead relative to the desired behaviour, which according to $F = ma$ should be a phase lag of 180° between position and force.

At the two frequencies of interest indicated in the figure, the phase error is very small for both $b_x = 0.1$ and $b_x = 1$. For $b_x = 10$, the phase response seems to pass through the 5° location at the frequency of interest to the thrust measurements, but is rather large for frequencies smaller than this. This could present an interesting opportunity to compensate for the lag induced by the use of the causal filter discussed in the previous section; however the highly non-linear phase nature of the action of b_x

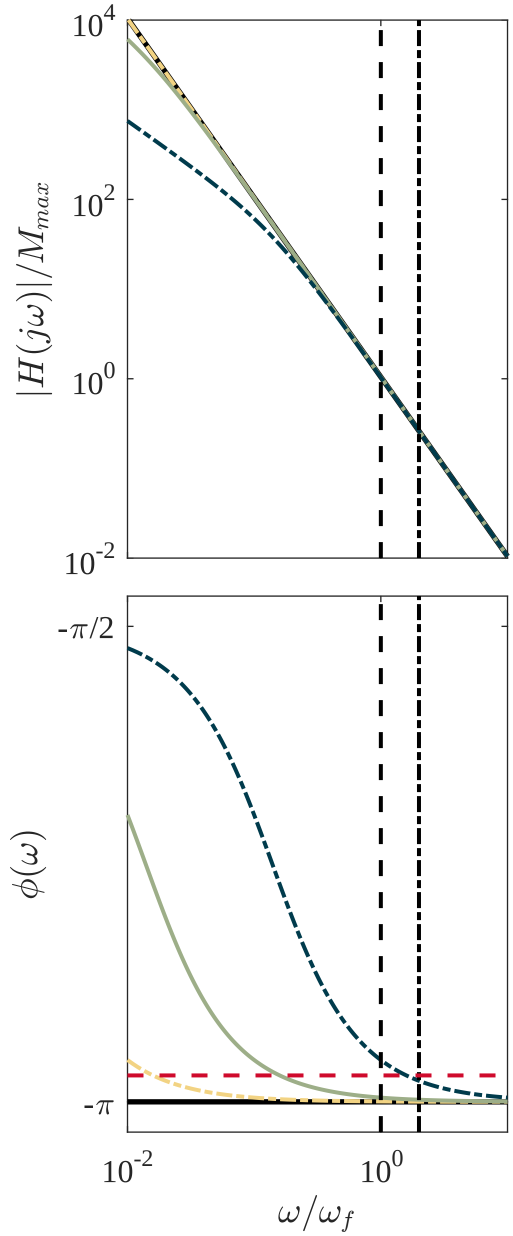

Figure 5.2: Bode plot showing the frequency response in the streamwise direction for a system described by Equations 5.1-5.4 with four different values of b_x, the virtual damping. Magnitude Responses shown normalized by the maximum amplitude gain constraint for the y-direction M_{max}, and frequency axis shown normalized by the estimated operating frequency, ω_f. ▬ $b_x = 0$ Ns/m; ▬ ▪ $b_x = 0.1$ Ns/m; ▬ $b_x = 1$ Ns/m; and ▬ ▪ $b_x = 10$ Ns/m. Bold dashed and dash-dot lines indicate points of interest in the figure: ▬ ▬ Operating condition, $\omega = \omega_f$; ▬ ▪ Dominant frequency in thrust signal, $\omega \approx 2\omega_f$; and ▬ ▬ 5° phase difference from desired phase response.

makes this a risky proposition when there is variation in the operating frequency (in this case, due to irregularities in the vortex shedding, and the advance of the airfoil through the vortex street). To ensure that the behaviour of the filter(s) applied to our signal is well understood and uniform for a variety of incoming frequency content, the value $b_x = 1$ was selected as a compromise between limiting gain at low frequencies (a desirable behaviour), and mitigating changes to the phase response at the operating frequency. The potential usefulness of these dynamics for lag compensation is left for future exploration.

5.2.3 System Characterization: Uniform Free-Stream Behaviour

To ensure that the 2D motion of the CTS with the dynamics described in the previous section produces trajectories that are representative of the behaviour of a passive system, the dynamics were tested in the case that the upstream circular cylinder was not present. These tests were conducted both at the nominal free-stream velocity U_∞, as well as at a reduced speed of approximately $U_\infty/2$, to account for the reduction in oncoming flow experienced by the airfoil when the cylinder is present. The top panel of Figure 5.3 shows the x-position as a function of time for the airfoil at both free-stream velocities. Two trials at each speed are shown. In all cases, the airfoil is blown backwards after being released at time $t = 10$s, more quickly in the high-speed oncoming flow cases than for low flow speeds. The airfoil's transverse position remains constant at its neutral position throughout all four tests. This behaviour is largely as expected in response to a uniform incoming free-stream due to the action of drag on the airfoil.

Looking at Figure 5.3 however, several unideal aspects of the airfoil behaviour are apparent. For the higher tunnel speed, the action of drag appears similar in both trials; however, for the slower speed, there is significant variation between the streamwise trajectories in the two essentially identical trials. This behaviour arises when measured drag values encounter the noise floor of the system, that is, when instantaneous measured drag is so small that noise on the reading becomes larger than the signal. The bottom panel of Figure 5.3 shows power spectra computed from the data from the trials shown in the top panel. Comparing the yellow (high-speed) and green (low-speed) spectra, it is clear that the mean drag experienced by the airfoil is quite near or even below the noise floor for the low speed tests. In addition, there is a large spike in noise power at approximately 12 Hz, which is larger than the mean value signal of interest for both speeds. The causal filter implemented on the raw signal helps to remove some of this higher frequency noise before the signal is

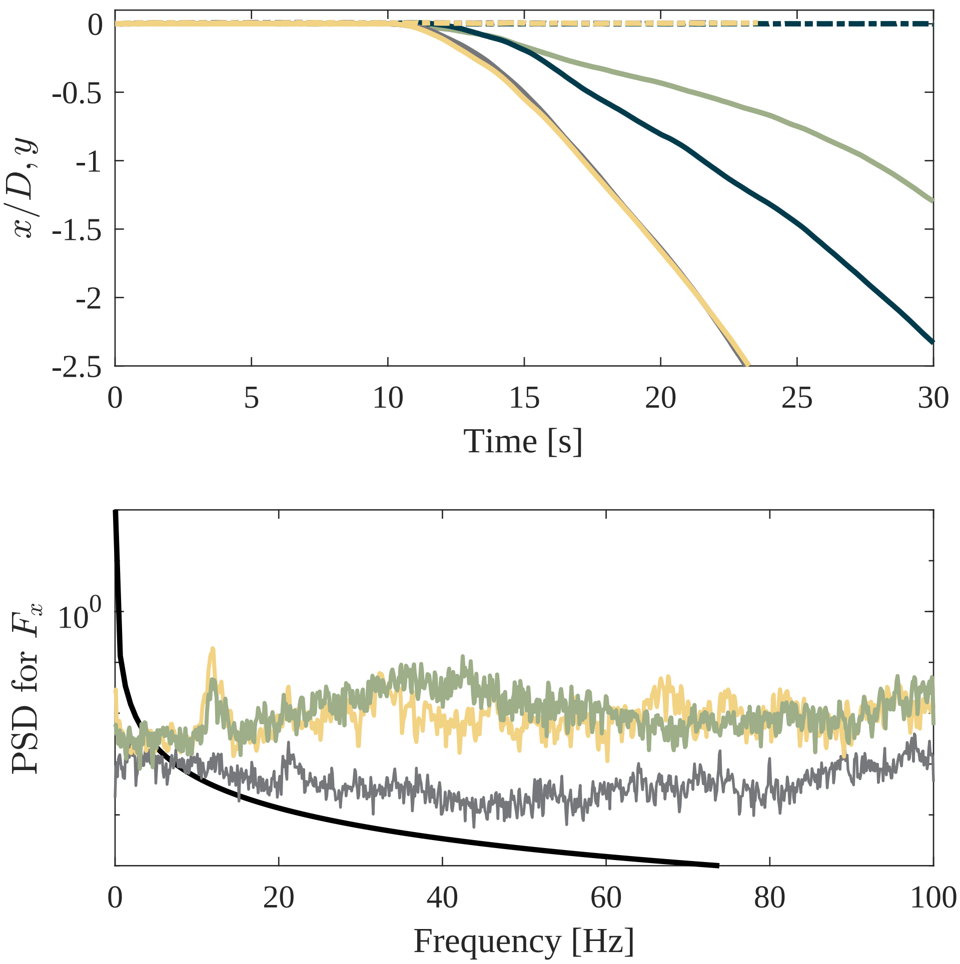

Figure 5.3: Top Panel: Position of airfoil as a function of time during uniform free-stream validation testing. Airfoil motion in y direction is initiated at $t = 0$ s, while motion in the x-direction is initiated at $t = 10$ s. Four trials are recorded at two different free-stream velocities: — — $U_\infty/2$, and — — U_∞. Solid lines indicate x-position, and dot-dashed lines indicate y-position. Bottom Panel: Power spectral densities of F_x: — U_∞, and — $U_\infty/2$. These PSDs are computed based on both trials at each speed for times after motion in the x-direction has started at $t = 10$ s using Welch's Method. Bold black line shows frequency response of the airfoil in the x-direction. Thin grey line shows the PSD for F_x before triggering occurs for all trials, with mean value subtracted. This represents the pre-trigger noise floor.

used for feedback control; however since the measured value of drag is very small the signal still drifts with a large amplitude relative to the true value, leading to the divergent behaviour displayed for the two low-speed trials in Figure 5.3.

In addition to behavioural errors caused by sensor noise, the system also exhibited sensitivity to the choice of zero set point for the force in the x-direction. During testing and validation of the system, it was noted that even a very small change to the zero-force set point sometimes resulted in non-physical behaviour such as the airfoil drifting upstream during some part of a test with uniform oncoming flow. This was only observed for flow speeds slower than the nominal U_∞ used for the experiments described in this chapter, and was not observed for either flow speed discussed here with the set-point used for the present experiments (though the extent to which this was due to the limited number of trials for the low speed case is questionable).

The bold black line in the bottom panel of Figure 5.3 shows the action of the dynamics as a filter on the incoming force signal. The dynamics strongly filter out noise in regions away from $f = 0$, which leads to the relatively smooth changes in position observed in the top panel. Unfortunately, this also explains the system's sensitivity to set-point errors. Even a very small difference between a measured value and the zero set-point is strongly amplified by the dynamics; this includes noise power in the vicinity of $f = 0$, and drift in the sensor output as a function of time. Thus, even a small change in this set point can lead to divergent behaviour for small measured forces, especially if it further reduces the apparent mean value of the drag.

These undesirable characteristics arise from a classic measurement challenge, as the force sensor itself was selected to withstand the forces and torques experienced in the transverse direction during a test, which are more than an order of magnitude larger than the thrust and drag. To measure the small drag values, a sensor with a smaller dynamic range and lower noise floor would be ideal, but this is challenging due to simultaneous large forces in the y-direction. In addition, the mechanical noise environment in which these experiments are performed compounds the challenge of measuring small thrust/drag values accurately, regardless of the noise floor inherent to the sensor. In the second panel of Figure 5.3, the thin grey line shows the PSD of a composite of the four signals in the top panel before the moment of triggering, with mean values removed. This represents the noise floor of the system when the CTS is not in motion in the x-direction (and though y-direction motion is permitted, it is very small). The noise power is clearly smaller without x-direction motion, though there do appear to be small spikes, likely due structural resonance of the airfoil itself, and shaking of the entire tunnel due to the action of the recirculating pump. However, once x-direction motion is initiated, there is a significant increase in measurement noise, both broadband and localized at specific frequencies (in particular, near 12 Hz). It is logical to assume this noise frequency arises from mechanical vibration due to the motion of the CTS carriage, which is actuated by stepping motors on a rack and pinion system. This mechanical noise is measured above the inherent noise floor of the current sensor, which makes the use of more sensitive equipment ineffective, if interfering mechanical noise cannot first be dampened.

Though these are serious concerns regarding the sensitivity to noise in the simulated motion of the airfoil in the x-direction, the behaviour that was observed in response to oncoming forcing from the cylinder (where mean thrust values are approximately an order of magnitude larger than the $U_\infty/2$ case discussed here) was quite consistent over many trials, including those conducted on different days and with different zero set points for the x-direction force. All data presented in this chapter are from a single testing day, and are calibrated in the same way as the above trials, which showed the desired response to uniform oncoming forcing. For further testing, a more precise sensor coupled with further efforts to damp mechanical noise should be used to mitigate these concerns.

5.3 Passive Captive Airfoil Motion in 2 Dimensions

Using the experimental setup described in Section 5.2, 10 trials were conducted to demonstrate the passive airfoil swimming upstream against the oncoming flow. The airfoil started at position $x = -3.7D$, where it was constrained to move only in the transverse direction for the first 10 seconds of each trial. This was done to ensure that any transient behaviour in the transverse direction (which is governed by spring-mass-damper dynamics) had died away before motion was initiated in the streamwise direction. Triggering for the initiation of 2D motion took place when the airfoil was at a local maximum position in the frame, determined based on a measured velocity $v_y = 0$. At $t = 10$s, the trigger to release the constraint on the streamwise motion was armed, such that starting from the next local maximum in transverse motion, the airfoil was allowed to translate in both the x and y directions according to the dynamics discussed in the previous section. Particle Image Velocimetry (PIV) was also started simultaneously with streamwise motion. For more details on PIV, including the experimental setup and processing, please see the complete discussion in Chapter 2. For each trial, the airfoil was stopped when it reached a location 30 cm or $2.6D$ upstream of its starting location ($1.1D$ downstream of the cylinder), since unlike the study of the dead fish by Beal et al., here the airfoil could not be allowed to physically contact the cylinder (Beal et al., 2006).

Figure 5.4 shows the 2D trajectories achieved by the airfoil after triggering occurred

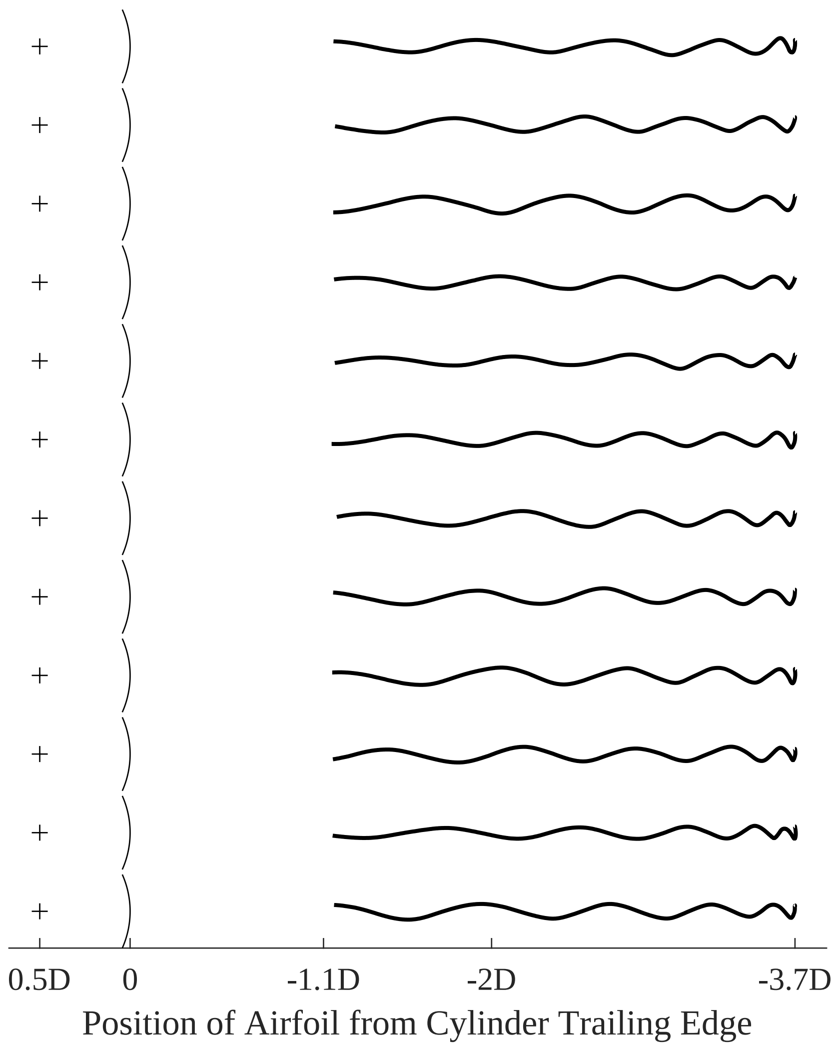

Figure 5.4: Airfoil trajectories for the 10 trials conducted for 2D Passive Captive Airfoil Experiments. Bold lines indicate individual airfoil trajectories, heading towards the upstream cylinder from right to left. Flow is from left to right. With the exception of very small times (when the airfoil is located in the vicinity of -3.7D), the motion of the airfoil is always upstream. Crosses indicate the cylinder center, and fine lines indicate the trailing edge of the cylinder for each trial. Y-axis scale is the same as that for the X-axis, with axis labels omitted for clarity.

for each trial. The airfoil oscillates in the transverse direction while travelling upstream, leading to a spatial frequency which decreases with proximity to the upstream circular cylinder as the airfoil accelerates in the streamwise direction. Though the foil begins its trajectory at a position maximum for all trials, the evolution of the trajectories does not appear to occur with consistent spatial frequency or phase. Figure 5.5 shows the populations of both x-velocity and thrust for the airfoil as a function of time over the 10 trials considered, with a single representative run shown bolded. From the figure, we see there is some spread in the mean acceleration (the slope of the line in the first panel), as well as variation in the minimal and maximal thrust coefficients achieved throughout a run. These small differences in thrust, amplified by the sensitivity of the x-direction dynamics to low-frequency inputs, translate to a spread in the spatial frequency in the airfoil motion shown in Figure 5.4.

In addition to the grey and bold, black lines shown in the bottom panel of Figure 5.5 which indicate measured thrust coefficients, the yellow line shown indicates the value of ma_x corresponding to the bolded black run. The correspondence

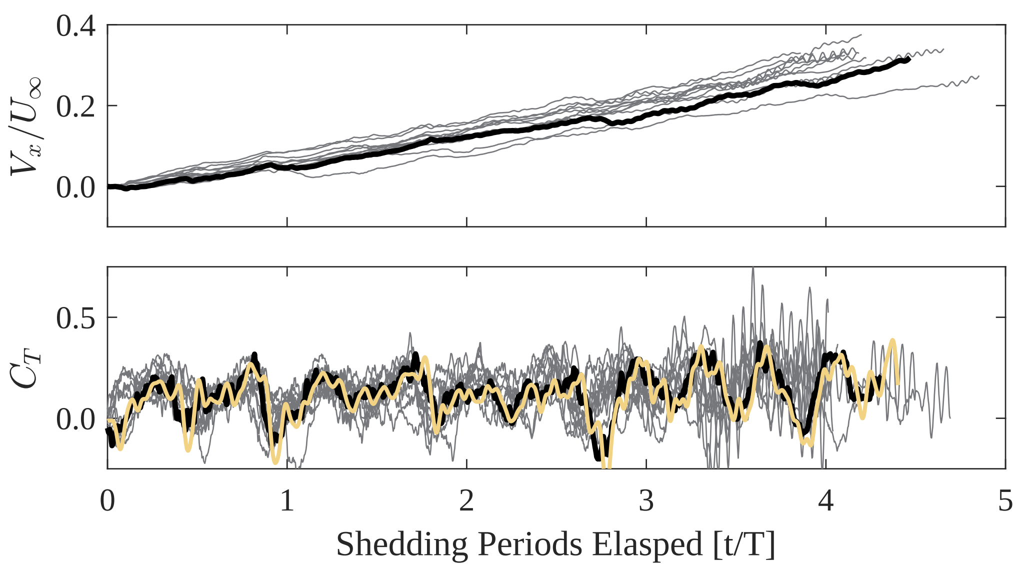

Figure 5.5: x-Direction velocity and measured thrust coefficient as a function of time. Thin grey lines show airfoil behaviour during each independent trial. Bold black lines highlight a particular trial with representative behaviour. In the bottom panel, the yellow line shows the value of the virtual mass of the airfoil times its acceleration (normalized by $p_\infty sc$): if the system was truly behaving according to $F_x = ma_x$, the bold black line and the yellow lines would be perfectly coincident for all time. Thrust coefficients are filtered using a Savitzky-Golay filter for clarity.

between the bold black and yellow lines indicates that the behaviour of the system is reasonably closely conforming to the expected behaviour given by $F_x = ma_x$ (with a small but noticeable phase lag due to filtering) despite the difficulties in realizing this cyber-physical representation discussed in Section 5.2.3. This was largely true for all recorded trials.

Throughout each trial, there is a strong correlation between moments of high thrust and transverse position of the airfoil. Considering a subset of the runs shown in Figure 5.4 where the pattern is most clear, Figure 5.6 shows the same trajectories but with moments of high thrust remaining dark in colour while moments of small thrust (or even of instantaneous drag) are lightened. From the figure, it is clear that the airfoil generates most of its useful thrust in the region of the cylinder centerline, while travelling at its maximum y-velocity. This mirrors the trend we observed in the case of the Driven and 1D Passive Captive Airfoils discussed in the previous chapters, and implies that the slaloming mode of interaction is again active for the 2D airfoil motion, as enforced by the passive mounting system dynamics. This is confirmed in the following section, where PIV fields show the interaction of the

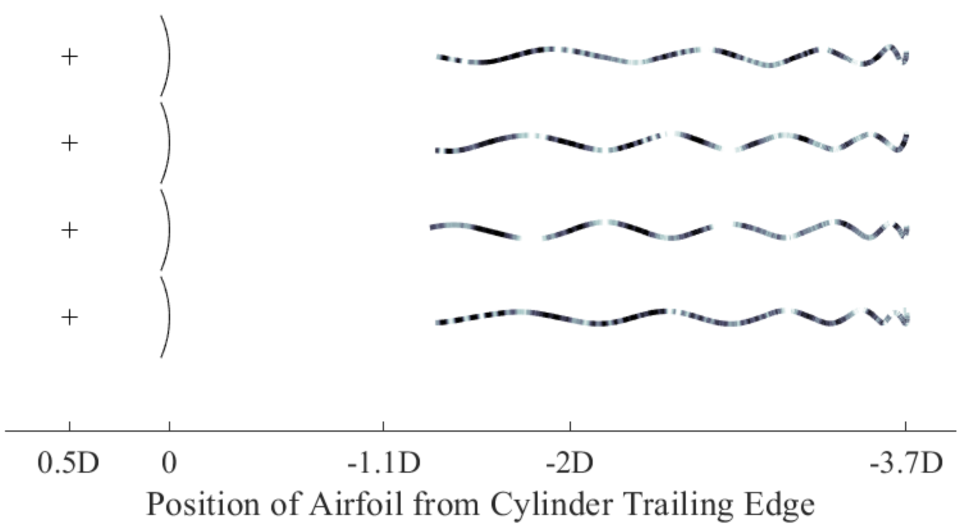

Figure 5.6: Airfoil trajectories for a subset of 2D Passive Captive Airfoil experimental trials. Individual airfoil trajectories are coloured according to experienced thrust: darker coloured portions indicate large positive thrust, while lighter coloured portions indicate smaller positive thrust. White moments in the trajectory are those for which instantaneous drag is experienced. Crosses indicate the cylinder center, and fine lines indicate the trailing edge of the cylinder for each trial. Y-axis scaling is the same as that for the X-axis, with the axis label omitted for clarity.

airfoil with oncoming flow.

One noticeable feature of the measured thrust coefficients in both Figure 5.5 and Figure 5.6 is the presence of large amplitude oscillations with a frequency of approximately 12 Hz. The origin of these high-frequency oscillations appears to be mechanical vibration due to the x-direction motion of the CTS, based on the presence of a similar noise power peak measured for the free-stream tests as discussed in Section 5.2.3. This indicates that the frequency content likely does not arise from a fluid-structure interaction due to the presence of the cylinder. The oscillations are most pronounced near the end of each trial when the airfoil is travelling with its maximum velocity. This observation was also confirmed by eye during testing where the CTS was seen to noticeably ‘shudder’ in this region during most tests, even when properly lubricated and set to operate far from any velocity or acceleration hardware limits. These oscillations make it challenging to accurately determine the effect on thrust of the airfoil’s entry into the suction region, where its presence is expected to influence shedding from the cylinder itself. This phenomenon is akin to vortex shedding suppression as observed by Strykowski and Sreenivasan (1990), and as discussed for a more similar configuration to the present case in references such as Liao et al. (2003), Beal et al. (2006), Lefebvre and Jones (2019), and Jarman et al. (2019), for example.

The expected behaviour in the steady case is that the airfoil’s presence beginning from approximately $2D$ downstream of the cylinder should begin to interrupt the formation and shedding of vortices from the upstream cylinder, instead becoming engulfed in an extended suction zone (Lefebvre & Jones, 2019). However, for the unsteady case this behaviour is modified when the speed of the airfoil is large relative to the formation and shedding time for the upstream vortices (Jarman et al., 2019). Since thrust measurements may be unreliable for this region in time, we instead investigate the lift experienced by the airfoil. The magnitude of measured lift is large compared to the interfering noise, and thus a much more reliable signal is recovered.

Following the example presented by Jarman et al. (2019), in Figure 5.7 we plot the lift coefficient for the airfoil for the same representative trials shown in Figure 5.6. Plots of C_L both as a function of downstream distance from the airfoil as well as a function of time normalized by the estimated shedding period T are presented. Similar to the findings of those authors, we see a slight increase in the magnitude of the experienced lift as the airfoil approaches the cylinder. This is due to the

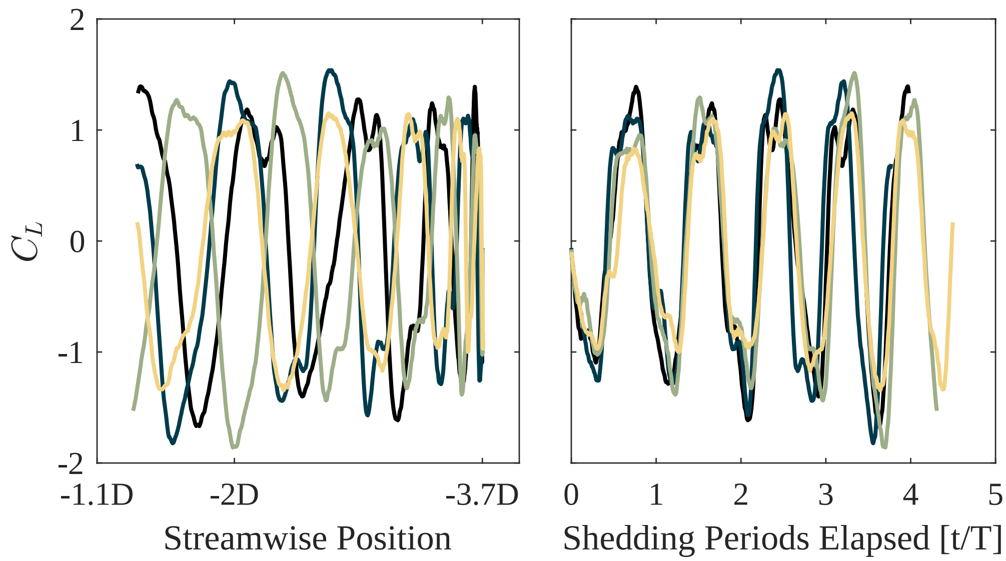

Figure 5.7: Lift Coefficient for the airfoil during the four tests presented previously in Figure 5.6. The left-hand panel shows C_L as a function of distance downstream from the cylinder (with the airfoil travelling from right to left). The right-hand panel shows C_L as a function of time, normalized by the expected shedding period.

presence of stronger and more coherent vorticity closer to the cylinder, which acts to increase the apparent angle of attack experienced by the airfoil. Since we do not see a subsequent reduction in the lift once the airfoil has approached closer than $2D$ to the cylinder, we can infer that the airfoil is approaching the cylinder at a high enough velocity that there is insufficient time for vortex shedding to be disrupted. This was the conclusion reached by Jarman et al. (2019) as well, for trials with similar V_x/U_∞ at the end of each trial.

The biggest difference between the current work and that study is that the airfoil in the present case is moving passively under its own power, and thus does not maintain a constant speed throughout each trial. This leads to the large variation in spatial frequency of the experienced lift force over an individual trial, unlike what was observed in the constant-speed tests by Jarman et al. (2019). It is interesting to note however that despite the quick divergence of the airfoil's behaviour as a function of space, the lift experienced as a function of time remains relatively similar between the presented trials. This hints at some similarity between airfoil interactions with oncoming vorticity, despite the fact that these give rise to quite different trajectories through the space. To explore the interactions between the oncoming flow and the airfoil directly, the next section presents results from PIV that accompanied these tests.

5.4 Interactions with Oncoming Flow

To further interrogate the airfoil behaviour, we consider contours of the Γ_2 Criterion both upstream and in the region of the airfoil, as in previous chapters; however, here each individual trial has unique spatial behaviour rendering a phase-averaging approach to extracting dynamically relevant behaviour inappropriate. Therefore, fields corresponding to a single representative trial are presented. Temporal and spatial filtering of the fields is applied to clarify flow behaviour, as described in Chapter 2.

Figure 5.8 presents six flow fields from the same trial, with relevant data regarding the airfoil behaviour shown in the bottom panels. As expected, there is a close phase-match between the transverse velocity and the observed lift coefficient. As in the 1D Passive Captive Airfoil case described in detail in Chapter 4, this behaviour is enforced by the spring-mass-damper mounting used in the transverse direction. By contrast, clearly no such relationship exists in the x-direction. The bottom panel of Figure 5.8 shows that in response to the non-zero mean thrust produced, the airfoil accelerates throughout the trial, with the acceleration modulated by periodic variations in thrust about its mean value.

The left hand column of snapshots in the top half of the Figure shows a series of stills corresponding to times indicated by the set of three dark blue dashed lines in the bottom two panels. The stills are outlined with the dash type corresponding to the moment in time that they represent, and are ordered from first to last in time as indicated by the arrows. In the top left panel, the airfoil is located at its bottom-most transverse position in the frame, and has zero velocity. A CW-rotating (negative signed) vortex is located just upstream and above the airfoil. A large CCW-rotating (positive signed) vortex is located farther upstream and below the airfoil. Although the presence of the large negative vortex is causing separation over the top surface of the foil, the suction effect from its passage has begun to increase the lift experienced by the airfoil, which rises more rapidly than the airfoil's transverse velocity. The thrust experienced by the airfoil is near a local minimum in this panel, as the small value of lift and small induced angle of attack produced by the oncoming flow leads to a limited contribution to thrust from the Katzmayr effect; however, the passage of vorticity acts to reduce the flow speed in the region of the airfoil, leading to favourable conditions for reduced drag.

In the second panel on the left hand side, the airfoil is located directly downstream of the cylinder and moving upwards with its maximum velocity. The region of negative

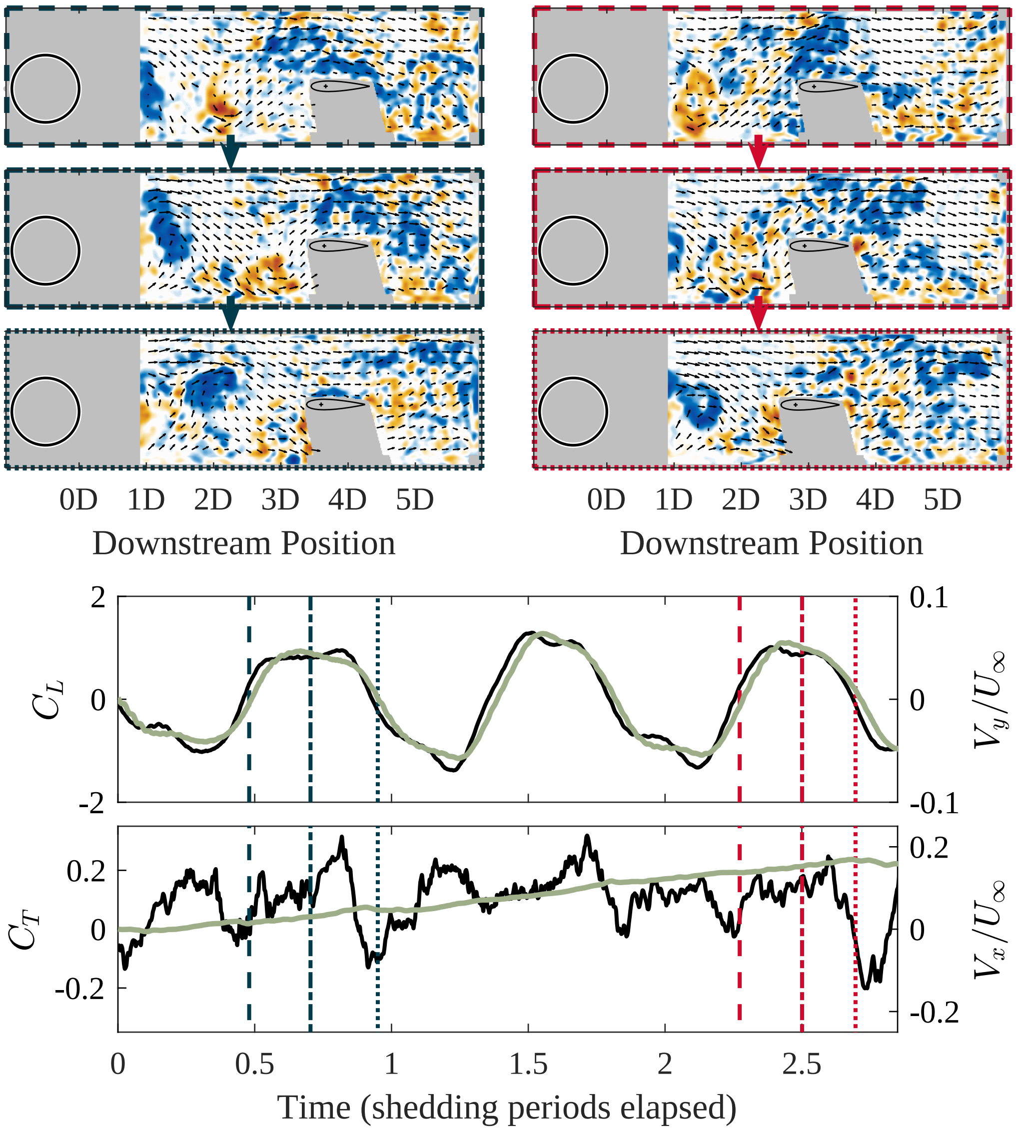

Figure 5.8: Representative behaviour of the 2D Passive Captive Airfoil, captured with simultaneous measurement of the surrounding flow field. Bottom two figures show behaviour of the airfoil as a function of time throughout the run: ▬ C_L and C_T; and ▬ x and y-direction velocities normalized by the oncoming flow speed U_∞. Two columns at top of figure show snapshots of the Γ_2 Criterion taken at time instants indicated in bottom two plots: dark blue snapshots correspond to earlier times, and red snapshots to later times as indicated. Line types on figure borders indicate corresponding time in bottom figures; time proceeds from top to bottom in each column as indicated by large arrows. Arrows overlaid on contour plots show flow velocity magnitude and direction.

vorticity is now over top and slightly behind the airfoil, and creating a low-pressure region contributing to high lift. The upstream positive vortex has moved closer to the leading edge, and is beginning to pass underneath the airfoil. The effect of these advancing vortices is to engulf the airfoil in a region of upward-flowing fluid, which increases the apparent angle of attack and correspondingly the Katzmayr thrust and lift of the airfoil. Considering the flow in the region of the airfoil, at this point in the cycle both the CW vortex above and the CCW vortex below the airfoil are contributing to reducing the magnitude of the oncoming flow near the airfoil's surface. This reduction, along with the large effective angle of attack lead to the high thrust values experienced at this point in the cycle, consistent with the trend noted in Figure 5.6.

Finally, considering the bottom left-hand panel in Figure 5.8 where the airfoil has reached its transverse position maximum and again has zero velocity, the oncoming CCW vortex is passing under the airfoil leading to the reversal in lift direction at this time. The flow in the region of the airfoil appears energized relative to its state in the previous two panels, and it therefore makes sense that the airfoil is achieving a local minimum of thrust. The precipitous drop in the instantaneous thrust at this location seems to indicate a particular interaction taking place in the flow; a similar drop is observed as the airfoil reverses direction again later, near $2.75T$. Though drops in thrust are seen for all locations where lift is reduced (due to the Katzmayr Effect), the steepness of the reduction in this case hints at additional effects at play. By analogy with the similar 1D case, there is evidence of a large positive signed trailing-edge vortex being shed in the bottom left-hand panel. Perhaps instantaneous conditions in the flow could lead to cycle-to-cycle variation in trailing-edge vortex shedding, which would provide a variable contribution to drag on the airfoil; however, the measured thrust here is relatively small, and such subtleties are challenging to detect with certainty both in the computed Γ_2 fields, and in the measured thrust under these experimental conditions.

Overall, the process of interaction between the airfoil and the oncoming flow shown in the dark blue left-hand series in Figure 5.8 is extremely similar to that presented for the 1D Case 0 Passive Captive Airfoil in Chapter 4. We again see evidence that the airfoil is operating in the slaloming mode of interaction with oncoming vorticity, which in general favours power production over thrust production and arises naturally from the tuning of the mounting system in this study. This similarity in behaviour is expected, since the x-direction velocity of the airfoil is small in this portion of the trial, rendering the 2D dynamics qualitatively similar to their 1D counterparts. By contrast, the right-hand column in Figure 5.8, outlined with red and corresponding to a later time in the same trial presents a series of snapshots for the same conditions in the transverse direction but where the airfoil is travelling with a considerable forward velocity. It has also begun to approach the cylinder more closely, though it has not reached the $2D$ threshold identified to correspond to the onset of coupling between the foil and the vortex shedding process itself (Lefebvre and Jones, 2019; Jarman et al., 2019).

It is interesting to note that although the airfoil is translating in the x-direction with an appreciable velocity, the phase relationship between transverse velocity and lift remains very similar even near the end of the run. The forward speed of the airfoil acts to increase the frequency of encounters with oncoming vorticity, in effect, increasing the frequency of the oncoming forcing (as observed for the constant velocity case by Jarman et al., 2019). Such a frequency increase is small however, making it challenging to visually observe over the duration of the trial. To estimate it, we consider the vortex spacing and convection velocity in the system (assuming the airfoil does not change the vortex shedding process):

$$2\lambda = TU_{\text{conv}} = U_{\text{conv}}/f, \tag{5.5}$$

where λ is the vortex spacing, and U_{conv} is the convection velocity of the vortices, here roughly approximated as 0.2 m/s based on the velocity deficit from the presence of the cylinder at the airfoil location. By allowing the airfoil to translate upstream, we are effectively increasing U_{conv}, such that our perturbed frequency is given by

$$f' = \frac{U_{\text{conv}} + V_x}{2\lambda}. \tag{5.6}$$

Finally, considering the nominal frequency and convection velocity,

$$\frac{f'}{f} = \frac{U_{\text{conv}} + V_x}{U_{\text{conv}}}. \tag{5.7}$$

Considering the maximum x-velocity of the airfoil during the PIV period shown of approximately $0.25U_{\text{conv}}$, this leads to a maximum expected frequency increase of 25%, from approximately 0.6 Hz to 0.76 Hz. Not only would this change be challenging to detect by eye, but even performing a frequency analysis of the 1D and 2D airfoils yields little result, as the change is small relative to the frequency resolution of the signals, and we expect a broadening of the frequency content in the region of the shedding frequency due to the airfoil's acceleration.

Such a small increase in experienced frequency should cause a very small phase lag to occur between force and velocity due to the spring-mass-damper dynamics in this direction. This does appear to be the case in the latter half of the dynamics shown in Figure 5.8; however cycle-to cycle variability in the instantaneous data considered here makes it challenging to confirm the cause is the slight frequency increase expected. In any case, as discussed in Chapter 4, such a small phase shift should have little impact on the transverse dynamics. This explains the apparent insensitivity of the time history of C_L to the airfoil's streamwise velocity despite divergent behaviour of the airfoil in space, as shown both here and for several additional trials in Figure 5.7 in the previous section. Small differences in the frequency of vortex encounters result in only small differences in transverse oscillation, and a very small temporal phase shift between velocity and lift. These similar oscillations are then stretched into very different trajectories in space by differences in the airfoil's forward velocity.

Connecting this to interactions with oncoming vorticity, in the right-hand column of Figure 5.8 outlined in red, we see that the advance of the foil has altered the global phase of the interactions with oncoming vorticity in a noticeable way. The airfoil is closer to the cylinder, and in the first panel showing a moment where the airfoil is at a local minimum in transverse position (the same phase of motion as in the top left panel), it is clear that vortex shedding from the foil is at an earlier phase than what is shown in the corresponding left-hand panel. There is a CCW vortex just downstream of the beginning of the field of view on the right; the same vortex has advanced a considerable distance on the left. However, the interactions of the foil itself with the oncoming vorticity are largely similar between the earlier and later snapshots (left and right hand columns in the figure). We again see that the airfoil generates its maximum lift and thrust as it passes directly behind the cylinder, underneath a large negative signed oncoming vortex, and that the subsequent passage of a positive-signed vortex beneath the airfoil causes a large drop in thrust, and hastens the transition between positive and negative lift generation. These interactions just happen with a higher frequency as the airfoil advances. Similar to the 1D case, this demonstrates the utility of passive transverse dynamics in allowing the behaviour of the airfoil to adapt to moderate changes in the frequency content of oncoming forcing.

5.5 Chapter 5 Interim Summary and Conclusions

This chapter presented experimental results from the investigation of a Passive Captive Airfoil placed downstream of a circular cylinder allowed to move in two dimensions, both in the transverse and streamwise directions of the flow. The behaviour is quite similar to that described in the 1-dimensional case presented previously, with the fundamental difference that the airfoil translates upstream towards the circular cylinder under the action of mean thrust. The following conclusions can be drawn from this work:

1) An airfoil allowed to move passively in the wake of a circular cylinder can be made to generate net thrust larger than its net drag, which allows it to swim upstream towards the cylinder. This behaviour is qualitatively similar to that exhibited by the dead fish studied by Beal et al. (2006).

2) For mounting system dynamics tuned as discussed in Chapter 4, the interaction of the airfoil with oncoming vorticity is similar between 2-dimensional motion and motion only in the transverse direction: the airfoil's transverse velocity increases in frequency tracking similar changes to the experienced lift as the foil approaches the cylinder. This allows for similar interactions between the airfoil and oncoming vorticity as seen in the 1-dimensional case, while the airfoil remains far enough downstream that it does not interfere with the vortex formation and shedding process.

Chapter 6

REALISTIC ENGINEERING BEHAVIOUR AND THE EFFECTS OF FRICTION

6.1 Introduction

In previous chapters, we examined the behaviour of an idealized system: an airfoil was either driven through a pre-planned trajectory, as in Chapter 3, or its motion was governed by very simple, linear second-order dynamics as in Chapters 4. These dynamics were implemented cyber-physically, which allowed for the elimination of nonidealities associated with an engineering system needed to realize such behaviour in application. The dynamics of the mounting system in the previous chapters, as well as in past studies of active or semi-passive flapping foils (for example Gopalkrishnan et al., 1994; Su and Breuer, 2019), are determined by constraints due to a particular mechanical setup or by programming a desired behaviour, which is then realized by active components of the mounting system (motors and actuators). In this chapter, we explore a ubiquitous effect that we have observed to significantly affect the performance of a truly passive mounting system: friction.

Descriptions of the action of friction are plentiful. Though simple models for sliding friction based on a linear relationship between normal force at a contact and the friction experienced may be adequate for systems operating at steady-state, accurately describing the dynamics of friction for systems that start and stop as well as undergoing other unsteady manoeuvres is much more complicated. Even more complicated again is considering how the action of friction may affect a system in the context of control, reviewed for example by Armstrong-Hélouvry et al. (1994). A further consideration for the experiments at hand is how friction in the mounting system used to facilitate passive airfoil motion may affect the nature of the airfoil's interactions with oncoming vorticity, and its ability to generate thrust while extracting energy from the flow.

To explore this topic, experiments akin to those discussed in previous chapters were conducted using an all-mechanical mounting system for the airfoil, the Mechanical Free-Response System (MFRS). In brief, this system was created to ensure that truly passive motion of the airfoil was observed (since the MFRS contains no motors or other actuators to do net work on the flow, unlike the CTS), but it suffers from very high friction in the mechanism that allows for 'free' transverse motion of the airfoil. This friction changes the character of the observed motion of the airfoil from smooth and approximately sinusoidal to impulsive, with frequent starts and stops. Accurately modelling the behaviour of the airfoil subject to these strongly nonlinear and highly setup-specific effects presents a challenge that is not solved in this chapter: instead, a highly simplified model for the action of friction is presented. Although it does not model the underlying physics associated with friction in the system, this model does adequately reproduce the qualitative behaviour of the airfoil, as well as statistics describing its motion. Through this simple realization of friction in the mounting system, access to reproducible and tuneable behaviour is made available to facilitate interrogation of the effect of friction on interactions with vorticity as well as system performance.

The purpose of this chapter is to describe the behaviour of the truly passive MFRS, and to present a simplified model for friction that adequately reproduces its behaviour. This simplified model is then used to interrogate the effects of stick-slip friction on the performance of a passive energy harvesting system.

6.2 Characterization of Behaviour for the Mechanical Free-Response System

Engineering effects of friction on system performance were first investigated through the creation of a fully passive, all-mechanical mounting system referred to throughout the thesis as the Mechanical Free-Response System (MFRS). The experimental setup and details for this system are provided in Chapter 2. To summarize, the same NACA 0018 airfoil used in the experiments for all previous chapters was mounted a distance of $2.7D$ downstream of the same circular cylinder, and its all-mechanical mounting allowed for the airfoil to undergo fully passive motion in the transverse (y) direction only. The free-stream velocity U_∞ was kept the same as in previous experiments, leading to similar forces experienced by the airfoil. The biggest difference between the behaviour of this system and those discussed up to this point is the presence of significant friction in the mechanism that allowed for the 'free' translational motion: the friction observed in the system causes a clear and noticeable difference in behaviour relative to idealized sinusoidal motion.

Qualitatively, this difference is observed in Figure 6.1. The figure shows three realizations of the motion of the MFRS over approximately $18T$, each at a different initial transverse position relative to the circular cylinder. Experimental conditions for each of these three trials are given in Table 6.1. From the Table, we see that

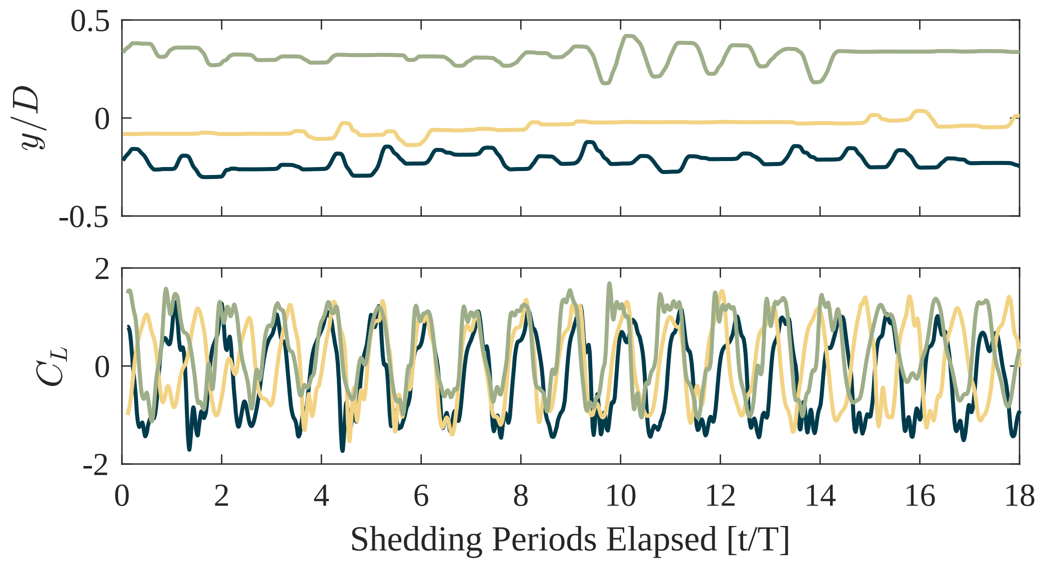

Figure 6.1: Transverse position (top) and lift coefficient (bottom) for three independent trials of MFRS motion. Table 6.1 shows experimental conditions for each trial. — Case MFRS10. — Case MFRS11. — Case MFRS12.

each realization has both positive mean power and positive mean thrust: the airfoil is both producing net thrust and extracting net power, despite the action of friction. Unlike previous experiments however, the airfoil's motion is far from sinusoidal. It seems to go through periods of inactivity where the action of friction is so great that the airfoil remains stationary despite oncoming fluid forcing, punctuated by periods where the airfoil undergoes roughly sinusoidal motion. The onset of motion appears to have significant hysteresis, where once the airfoil has started to move a smaller lift force is able to maintain its motion. This creates a highly irregular trajectory despite similarities in lift amplitude over a whole test, shown in the bottom panel of Figure 6.1. The characteristic shape of the y-direction motion once the airfoil has started to move appears to be a truncated sinusoid, where at position extrema the force applied to the airfoil is insufficient to re-initiate airfoil motion as the direction of forcing is reversed.

Case #	**Mean y-Position**	**AoA [º]**	**Mean C_T**	**Mean C_P**	**Colour**
MFRS10	-0.2D	-1	0.12	0.03	—
MFRS11	-0.05D	2	0.12	0.015	—
MFRS12	0.3D	6	0.068	0.009	—

Table 6.1: Case number and experimental conditions for representative MFRS trials discussed throughout this chapter. The final column shows the colour used to represent each trial in figures in this section.

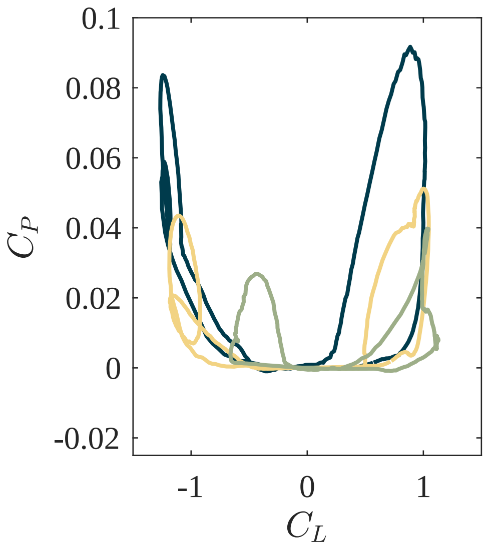

Figure 6.2: Phase portrait of C_p for three representative MFRS trials. Data are phase averaged over each trial individually, and resulting data is smoothed using a Savitzky-Golay filter with a width of 10% of T for clarity. ▬ Case MFRS10. ▬ Case MFRS11. ▬ Case MFRS12.

Small differences in maximum and minimum lift coefficient for the different trials in Figure 6.1 stem from different angles of attack, as given in Table 6.1. Correspondingly, the lift extrema seem to be slightly more positive for MFRS12 with an angle of attack of 6°, and slightly more negative for MFRS10 with an angle of attack of -1°. In addition, the MFRS does not enforce a mean position of zero (or another mean position set-point), and the mean position of the airfoil was observed to drift under the action of fluid forcing over long time horizons. The figure also highlights the drift in the frequency of vortex shedding. Although the shedding frequency is similar throughout the trials pictured in Figure 6.1, there are periods where the forcing seems to align for all trials, and later periods where they appear completely out of phase. This is due to slight shifts in frequency both between runs, and over the course of a single trial. Finally, in contrast to the Driven and Passive Captive Airfoil Experiments discussed in previous chapters, the phase shift between transverse force and velocity for the MFRS Experiments at any particular moment in a trial varied widely, though in the mean, the behaviour resembles that enforced in the driven case. This is discussed at more length in Chapter 2.

Though the motion of the foil is highly irregular, the power coefficient C_p remains either positive or zero as seen in Figure 6.2. This fits with our intuition since the MFRS does not have any means with which to store energy (other than some very soft springs used for station keeping) for later doing work on the flow. Phase portraits of the phase-averaged power coefficient shown in the Figure were computed from a segment $2.9T$ in duration contained within the signals shown in Figure 6.1, corresponding to times where PIV images are available as required for phase averaging (see Chapter 2). Though the portraits exhibit a butterfly-type structure similar to what was observed in the friction-free cases, the action of friction is apparent. The airfoil spends much of its time stationary, at which time it is generating zero power. This leads to the region of zero C_P for non-zero C_L between the two lobes visible in the figure – all of the power production is localized in regions of very high lift only. Contributions from low-speed instants are attenuated, since the action of friction often causes the airfoil to come to a stop when lift is small or moderate; in addition, due to phase averaging, the mean power for phases linked to small lift is attenuated by many instances where the airfoil is not moving.

This highlights a conceptual issue with phase averaging for data sets of this nature. The airfoil does not necessarily have a mean behaviour as a function of phase; the dynamic quantities of interest are more accurately described by a bimodal or even multimodal distribution, based on effects that rely on the time history of the airfoil's motion. Despite these challenges, Figure 6.2 still illuminates the hysteresis in the power produced as a function of the measured lift. Due to the nonlinear nature of friction, once motion has been initiated smaller speeds can be maintained without stopping than before the initiation of motion. This leads to higher power production recorded for the same values of lift when the lift is decreasing vs when it is increasing, which results in the vertical offset between instants in phase with the same value of C_L seen in the figure. One must use caution when interpreting Figure 6.2 however, as the varying height of the lobes (instances of large power production) may be misleading. This height is influenced by the fraction of recorded times averaged where friction has prevented motion, which may not be constant across the three trials. Since each trial in the figure represents a very short time (only 2.9 vortex shedding cycles), these height differences do not necessarily represent a statistically significant difference in behaviour between trials.

To address these issues, we present a statistical description of the behaviour of the MFRS in the following section.

6.2.1 Statistical Picture of MFRS Motion

To further interrogate the characteristics of the MFRS, we consider a statistical description of the dynamic quantities of interest. To compute statistics for MFRS motion discussed in this section, 9 trials each 30s in length are considered. The trials have different mean positions in the tunnel, as well as varying angles of attack. These quantities are summarized for each individual trial in Table 6.2. The total number of shedding periods considered is approximately 168. Dynamic quantities

MFRS Case #	4	5	6	7	8	9	10	11	12
AoA [°]	6	7	7	2	3	4	-1	2	6
Mean y/D	0.26	0.28	0.38	-0.03	0.12	0.22	-0.22	-0.05	0.32

Table 6.2: Case number and experimental conditions for all MFRS trials included in the statistical analysis of airfoil behaviour in this section.

are sampled at a rate of 25,000 Hz, as discussed in Chapter 2.

We first consider the lift acting on the airfoil, as this represents the forcing or input to our system. We expect this to be largely driven by oncoming vorticity and thus similar to all other cases presented in previous chapters; however, as discussed previously, airfoil motion in-phase with the oncoming flow acts to reduce both the experienced effective angle of attack and lift on the airfoil. Since the airfoil is much more stationary in these trials than in previously discussed cases with sinusoidal or quasi-sinusoidal motion, we expect slightly larger experienced forces. In addition, as summarized in Table 6.2, we see that for most of the considered trials, the airfoil has a positive angle of attack. Thus, we also expect that overall, there will be a larger positive lift experienced by the airfoil than the corresponding negative one. We stress however that since the airfoil experiences cyclic forcing which induces an effective angle of attack that is on the order of (or larger than) the static angles of attack in Table 6.1, the sign of the static angle of attack offset should not pose any further complications to the statistical analysis of this system.

Figure 6.3 shows a histogram of the lift coefficients experienced by the airfoil over all trials summarized in Table 6.2. As expected, the lift distribution appears to be bimodal, with two mean values of lift corresponding to the mean positive and mean negative lift values experienced by the airfoil. Also as expected, the mean positive lift experienced is slightly larger than the mean negative lift – the whole distribution appears to be shifted towards slightly positive values of C_L due to the the positive angle of attack of the airfoil for many of the trials.

To make these observations more quantitative, a Gaussian Mixture (GM) distribution model was fit to the C_L data. This is a probability distribution composed of a sum of N normal distributions each with a different mean and standard deviation, giving rise to a Probability Density Function (PDF) described by

$$P(x) = \sum_{i=1}^{N} \frac{A_i}{\sigma_i\sqrt{2\pi}} \exp\left[\frac{-(x-\mu_i)^2}{2\sigma_i^2}\right], \tag{6.1}$$

where μ_i and σ_i are the means and standard deviations of the underlying distribution(s), and A_i are the proportionality coefficients for each constituent distribution, such that $\sum_{i=1}^{N} A_i = 1$. MATLAB's built-in fitting function `fitgmdist()` was used for fitting. While this fit provides useful and interpretable information about the behaviour of the airfoil, we emphasize that the fitting procedure is not particularly robust. For example, different choices for initial guesses for μ and σ for the underlying distributions can result in shifts in the optimal locations of these values, with resulting distributions that exhibit similar fitness. This is particularly prevalent in fits to airfoil velocity, discussed in the following paragraphs. This implies that the distributions selected are non-unique, in that a variety of choices of probability distribution describe the underlying populations equally well. Thus, caution should be exercised in interpreting these results. In all further discussions, quantitative information from GM fits to data are used only to support qualitative information from histograms. Distributions fit starting from the same initial guesses are used to qualitatively compare the behaviour of the airfoil for different mounting parameter models for friction.

With these caveats in mind, the result of fitting a GM model with $N = 2$ to the observed C_L values is shown overlaid on the histogram in Figure 6.3. From this

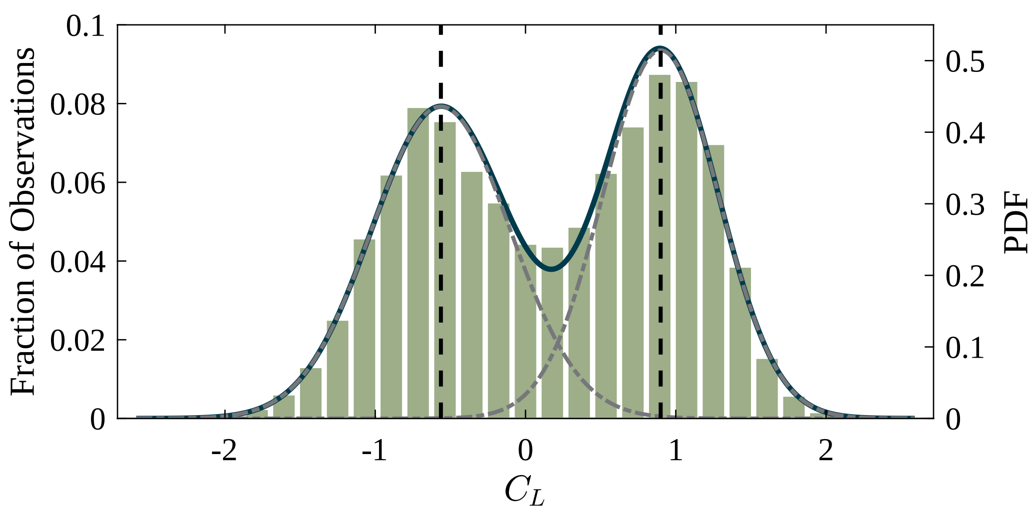

Figure 6.3: Histogram and Gaussian Mixture Distribution fit to C_L data for all recorded MFRS trials. Green bars show histogram counts, normalized by total data points considered (left axis). Line plots show Gaussian Mixture distribution fit to the data, and two underlying Normal distributions with different mean and standard deviation values (right axis). Vertical dashed lines indicate the two mean values for C_L.

fit, the mean positive value of C_L is 0.90 while the mean negative value is -0.56, which supports our observation that there is a shift towards positive values of lift due to a positive mean angle of attack. The standard deviation values for each of the two Normal distributions which form the GM distribution are similar, though the standard deviation for the negative peak is slightly larger. As expected, the the fit recovers a proportionality coefficient A_i for each of the underlying normal distributions of approximately 0.5, indicating that both peaks have similar dynamic importance.

To characterize the behaviour of the MFRS experiment (the 'output' of our system), we consider the velocity of the airfoil. Though the output dynamics are strictly specified by the position, and the mean offset position likely plays a role in forcing experienced by the airfoil, velocity is considered instead for two reasons. Firstly, the airfoil velocity is the key factor in determining the frictional behaviour of the system, which dominates the observed dynamics. Secondly, by considering velocity observations it is more likely that all recorded trials will have a similar mean value, which is clearly untrue for position.

If the airfoil dynamics were linear, we would expect that the two normally distributed populations of lift shown in Figure 6.3 would create a similar set of populations in velocity; however, the airfoil's velocity is strongly influenced by the nonlinear action of friction such that a significant proportion of observations have a velocity of 0, despite varying input lift values. This fraction, computed by dividing the number of observations below a threshold velocity of ± 1 mm/s by total observations, is 52%. To account for this in our analysis, we again consider fitting a GM distribution, but with three underlying normal distributions allowed ($N = 3$).

Figure 6.4 shows a histogram of the observed velocities, along with the best-fit GM distribution to the data. Unlike the lift coefficient data, a GM-type distribution does not seem to describe the observations particularly well; however, insight can still be gained from considering this imperfect model. The fit to the data recovers three underlying Normal distributions that describe the velocity. The first of these is a Gaussian pulse with $\mu = 0$, and a very small standard deviation. The proportion of the data attributed to this pulse is 52%, which matches the proportion of observations of zero velocity computed earlier. The remaining observations are then fit based on a similar bimodal model to that used for the lift coefficient. The recovered Normal distributions are centered around [-0.026, 0.027] and each represent approximately 24% of the observations; however these center locations were found to be quite

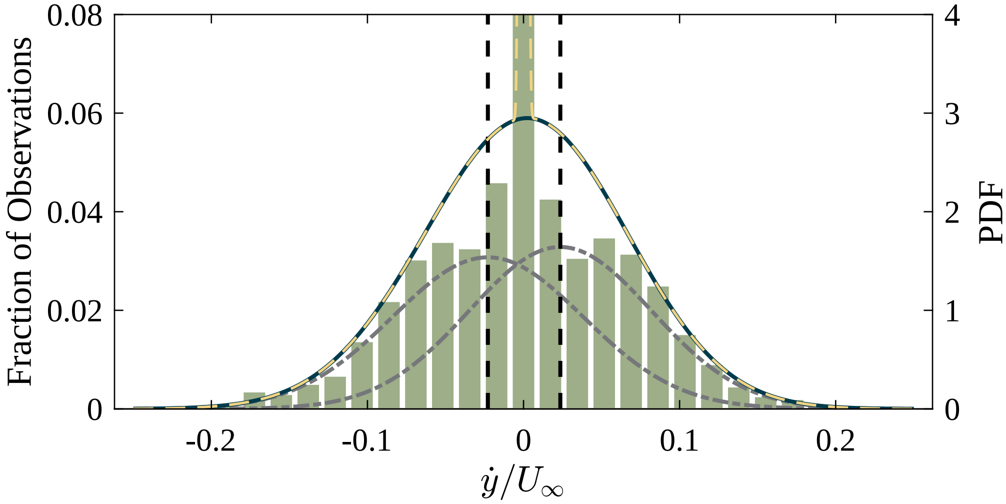

Figure 6.4: Histogram and Gaussian Mixture distribution fit to y-velocity data for all recorded MFRS trials. Green bars show histogram counts, normalized by total data points considered (left axis). Note that the data in the bin centered at zero have a proportion of approximately 0.52 - the plot is truncated for clarity. Line plots show two underlying Normal distributions with different mean and standard deviation values; sum of the two Normal distributions with non-zero mean; and Gaussian Mixture distribution fit to the data (right axis). Vertical dashed lines indicate the two non-zero mean values for $\dot{y}/U_\infty$.

sensitive to initial conditions in the fitting and must be interpreted with caution. A significant contributor to the poorness of this fit is the action of friction. In addition to increasing the proportionality of observations with zero velocity, a system with high friction also exhibits a larger proportion of observations with small velocities relative to the proportion of small lift values, as the airfoil's motion is impeded by the action of friction. This indicates that the Gaussian Mixture distribution presented here may not accurately represent the underlying behaviour of the airfoil, and another choice for probability distribution may reveal further insights regarding the behaviour of this system (in particular, one that allows skew towards small velocities). However, for simplicity and to facilitate comparisons to systems with (largely) linear dynamics in later sections, we continue with the present analysis.

6.3 Modelling Frictional Behaviour using the Captive Trajectory System

Although analyzing the behaviour of the MFRS provides many insights into the action of friction on the dynamics of the airfoil, these dynamics are extremely setup-specific. During testing, the behaviour of the airfoil was shown to change based on lubrication of the rails which allowed for translation, as well as more subtle adjustments such as the amount of play in the connection of the rails to the tunnel sides, for example. This makes the results from the above analysis challenging to generalize. Since the dynamics are not known a priori, it is also challenging to systematically vary the behaviour of the airfoil to isolate the effects of friction on the performance of the system. Thus, a cyber-physical representation (model) of the system is desired, such that parameters controlling the frictional behaviour(s) exhibited by the airfoil may be adjusted in a more systematic way.

6.3.1 Simplified Frictional Model and Predicted Behaviours

In the literature regarding frictional modelling, three parameters are included across a wide range of models: Coulomb Friction (F_C), Viscous Friction (F_V), and the Stiction Force (F_S). F_S, also called the breakaway force, represents the force that is required to overcome friction when an object is stationary and initiate motion. $F_C < F_S$ is the friction force that acts on an object after the initiation of motion, equivalent to the simple $F_f = \mu F_N$ picture of friction for steady-state motion. Viscous friction represents viscous damping acting to slow the object as a function of its velocity, $F_V = \sigma v$. Though many continuous-time models incorporate these parameters (see for example Huang et al. (2019)), we are able to take advantage of the digital nature of the CTS' operation to build a simplified representation of the nonlinear action of friction.

Building from the spring-mass-damper mounting system discussed in Chapters 4 and 5, nonlinearity is included through the addition of a condition on whether motion should be initiated at a given time step, based on both the present force acting on the airfoil and its motion history for some period preceding the current time. In particular, we consider the following logical statements evaluated at each time step in our experiment:

$\texttt{IF}\ (|F_y(t_{i+1})| > F_s \,||\, v(t_i) > 0)$

$\quad F_{\text{applied}} = \texttt{sgn}(F_y(t_{i+1}))\texttt{max}\big((|F_y(t_{i+1})| - F_C),\ 0\big)$

$\quad F_{\text{applied}} = m\ddot{y} + b\dot{y} + ky$

$\texttt{ELSE}$

$\quad v(t_{i+1}) = 0$

where $F_y(t_{i+1})$ is the current measured force, $v(t_i)$ is the velocity at the previous time step, and the parameters $[m, b, k]$ are selected based on matching the parameters ω_n and ζ from previous Case 0 Passive Captive Airfoil Experiments discussed in Chapter 4. More details regarding the tuning of these parameters are provided in

Section 6.3.2.

It is important to note that regardless of such tuning, this passive captive implementation is fundamentally different than the MFRS system it is meant to model. In particular, due to the non-zero value of k in the dynamics above, energy storage by this system is possible (whereas for the MFRS, it is not). However, based on the discussion in the following sections, this simple model gives rise to sufficiently similar behaviour to that of the MFRS, while retaining a linear character (away from moments of starting and stopping) that simplifies analysis of the system.

Stepping through the simplified model of friction above, at each time step the velocity and force acting on the airfoil are evaluated. For moments where *either* the force exceeds the stiction level F_S *or* the airfoil is already moving, the velocity update is calculated according to Case 0 spring-mass-damper dynamics. The applied force used to calculate the motion update is first reduced by a static value, F_C. In addition, the force $F_y(t_{i+1})$ is determined using a causal moving-average filter with a length of 10 samples to reduce the effect of sensor noise on airfoil behaviour. If neither of the above conditions are met, the airfoil remains motionless. This simple model does not explicitly account for viscous friction (in addition to viscous damping b, an explicit property of the mounting system) once the airfoil has started moving, as the airfoil velocity is very small and such an effect would be challenging to discern from the static offset value F_C. The `max` statement in the behaviour of the airfoil once motion has been initiated means that the force experienced by the airfoil remains continuous as the airfoil decelerates – this ensures that at times with small measured forcing the applied forces on the airfoil are continuous, which limits the effect of structural oscillations on the more sensitive low-speed (and therefore low-force) behaviour.

Figure 6.5 shows simulated behaviour of the CTS subject to idealized sinusoidal forcing, with this simplified model for friction implemented. The simulation used the same implicit time-stepping method as implemented on the CTS, which is described in detail in Chapter 2. From the simulation, sinusoidal forcing gives rise to stick-slip motion, as the airfoil comes to a stop near its position extremum when the experienced force is small. The shape of the position curve with flattened amplitude peaks, as well as the periods of zero velocity between impulsive airfoil motions is reminiscent of the behaviour of the MFRS, as desired. Since we have considered an idealized, perfectly sinusoidal forcing, the behaviour of the airfoil is the same during each pictured shedding period; however, we show in the next section

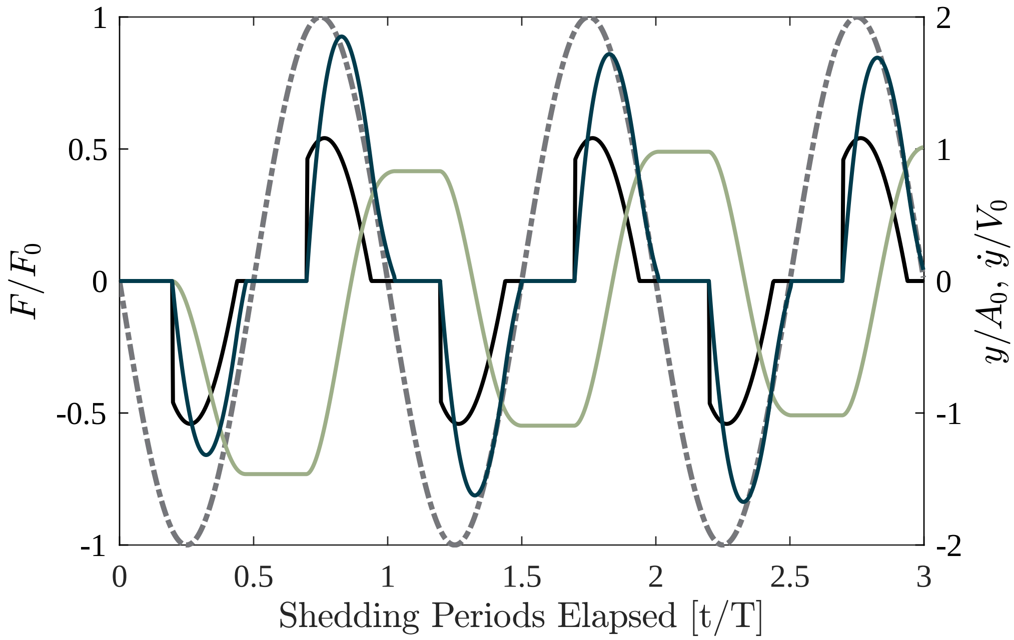

Figure 6.5: Idealized behaviour of a Passive Captive Airfoil governed by the simplified model for friction discussed in this section. Measured force acting on airfoil. Forcing experienced by airfoil due to action of friction. Resulting airfoil velocity. Resulting airfoil position. In the simulation, model parameters are set according to Case RF in Table 6.3, and $F_0 = 3.5$ N.

that subject to appropriate tuning of the dynamics, the behaviour of this simplified model qualitatively reproduces the behaviour of the MFRS.

6.3.2 Tuning the Simplified Frictional Model

Tuning of the (dimensional) values of the two parameters used in the simplified friction model (F_S and F_C) in tandem with the parameters $[m, b, k]$ for the mounting system proceeded as follows. First, ω_n and ζ were selected to coincide with those for the basic spring-mass-damper configuration considered in this thesis, the Case 0 dynamics discussed at length in Chapter 4. This choice meant $\omega_n = 4$ s^{-1} and $\zeta = 1$. Then, a preliminary choice for F_S was determined based on an analysis of available MFRS data, to determine an approximate force threshold corresponding to the initiation of airfoil motion. Finally, F_C, F_S, and $M_{\max}$, the maximum amplitude gain for the system, were manually tuned in experiment considering the following criteria:

- The percentage of time spent stationary for the Passive Captive airfoil should be similar to the MFRS Experiment, or approximately 50% of the time.

- During periods of oscillatory behaviour, the amplitude of oscillation should be close to the nominal oscillation amplitude $A_0 = 0.0051$ m.
- Airfoil motion should be qualitatively similar to the MFRS Experiment.

During the tuning process, it was observed that in general, increasing the stiction level, F_S caused the percentage of time that the airfoil was stationary to increase, with the obvious extreme values $F_S = 0$ meaning the airfoil is never stationary (no stiction), and $F_S > F_{\max}$ such that no amount of forcing is enough to overcome friction, and the airfoil is stationary 100% of the time. By contrast, increasing F_C seemed to reduce the maximum velocity achieved by the airfoil, or equivalently the width of the distribution(s) corresponding to non-zero velocity observations. This also makes sense, since the Coulomb friction acts to directly curtail the acceleration of the airfoil under the action of fluid forcing.

Using these observations, appropriate values of F_C, F_S, and $M_{\max}$ were selected to correspond to a Realistic Friction (RF) case based on the criteria listed above. In addition, based on our observations of the approximate action of F_S and F_C in our system, an additional High Friction (HF) case was developed and tested, corresponding to a system with the same stiction level, but a larger Coulomb friction F_C. Finally, as a point of comparison, a No Friction (NF) case with the same choice for $[\omega_n, \zeta, M_{\max}]$ was also tested. These conditions are summarized in Table 6.3.

To generate the simulation data presented in Figure 6.5, RF parameters given in Table 6.3 were used, along with a dimensional forcing amplitude of $F_0 = 3.5$ N. Though this forcing amplitude exceeds the estimated mean forcing amplitude of $F_y = 2.0 \pm 0.5$ N, as established in Chapter 3, it gives expected position extrema near a value of $\pm A_0$ which is consistent with the observed amplitude in our tuning efforts. It is interesting to note that though the position extrema are similar in amplitude to the cases discussed in previous chapters, the velocity extrema appear to exceed those observed for smooth airfoil motion with a similar amplitude. This arises from the larger maximum gain permitted in this system as compared to the

Case Code	m [kg]	b [Ns/m]	k [N/m]	F_S [N]	F_C [N]
HF				3.2	2.4
RF	6.13	49.02	98.04	3.2	1.6
NF				0.0	0.0

Table 6.3: Table showing model parameters used for each of three Passive Captive airfoil cases discussed in later sections in this chapter.

well-studied Case 0 dynamics, but also from the essentially transient behaviour of the airfoil as it attempts to 'catch up' with the phase of the forcing once motion is initiated. This has implications for interactions between the airfoil and oncoming vorticity as well as lift and thrust behaviour, all of which are strongly mediated by the amplitude and phase of the airfoil velocity relative to the forcing.

Validation of our choice of parameters for each of the three cases presented in Table 6.3, as well as consequences of this particular choice of friction model are discussed at length in Section 6.4. Overall, this simple model for friction allows access to two parameters governing the behaviour of the airfoil which are highly interpretable, and seem to give rise to realistic behaviour relative to the MFRS Experiments. The simplicity of the model has the added benefit of localizing the nonlinearity in the friction behaviour to the moment that the airfoil is starting or stopping – at other times in a cycle, the behaviour is governed by 2^{nd} order linear dynamics, which have been investigated at length in previous chapters.

6.3.3 Continuous-Time Model for Friction and Implementation Challenges

The very simple, discrete-time model used to simulate the action of friction presented previously is used to realize frictional behaviour for a Passive Captive airfoil in the remaining sections of this chapter; however, a brief mention of a potentially appropriate continuous time model for friction is warranted. Based on a recent review of popular friction models (Huang et al., 2019), the LuGre model of Canudas de Wit et al. (1995) was investigated for future use as a higher-fidelity model for friction in a passive mounting mechanism.

This model incorporates multiple forms of nonlinearity, and explicitly considers the deflection of the asperities (bristles) between two sliding surfaces, parameterized by the variable z. It includes 6 additional parameters that must be assumed or tuned based on the performance of the system, including F_S and F_C as discussed in the previous section:

- Bristle stiffness σ_0, and damping parameter σ_1
- Viscous damping, σ_2
- Stribeck Velocity, v_s
- Coulomb Friction level, F_C
- Stiction force level, F_S

Using these parameters, the friction model is given as (Canudas de Wit et al., 1995):

$$F_f = \sigma_0 z + \sigma_1 \frac{dz}{dt} + \sigma_2 \dot{y}, \tag{6.2}$$

$$\frac{dz}{dt} = \dot{y} - \frac{|\dot{y}|}{g(\dot{y})} z, \tag{6.3}$$

$$g(\dot{y}) = \frac{1}{\sigma_0}\left(F_C + (F_S - F_C)e^{-(\dot{y}/v_s)^2}\right). \tag{6.4}$$

Though the explicit consideration of surface asperity deflection adds considerable complexity to the model, it captures stick-slip motion and the complicated behaviour of an object during reversals in velocity direction, both of which are important for the current study. The velocity-dependent nature of the nonlinearity in this model is not captured by the simpler description in the previous section.

In a preliminary attempt to tune the parameters in the above equation to capture the behaviour of the MFRS Experiments, a model incorporating Equations 6.2-6.4 was created in MATLAB's Simulink environment. Though tuning attempts did generate several models that exhibited stick-slip motion and produced frictional forces that qualitatively reproduced the observed behaviour of the MFRS (including settling on values of F_S and F_C similar to those found by manual tuning in the previous section), the resulting additional fit parameters did not provide significantly more insight into the behaviour of the system. In addition, the ODEs in Equations 6.2-6.4 appear to be quite stiff, which presented a significant challenge for implementation on the CTS rather than in simulation. Thus, further work to implement this model in experiment is left for future investigations.

6.4 Validation of Simplified Friction Model: Real vs Simulated Friction

Using the simplified model for friction presented in Section 6.3.1, data regarding the input-output relationship between lift and the motion of the airfoil were obtained. For both the Realistic Friction (RF) and High Friction (HF) conditions described, data were collected over 16 trials each with a duration of 60 s, for a total number of vortex shedding periods of approximately 600. Data were sampled through the CTS at a rate limited to 200 Hz, as discussed further in Chapter 2.

Figure 6.6 shows histograms of the observed lift coefficient for both the RF and HF cases tested. Similar to the MFRS experiments, C_L appears to conform to a bimodal distribution, with one positive and one negative lift peak. Using the same fitting procedure as described for the MFRS Experiments in Section 6.2.1 with the same initial guesses for $[\mu, \sigma]$, a Gaussian Mixture (GM) model with two

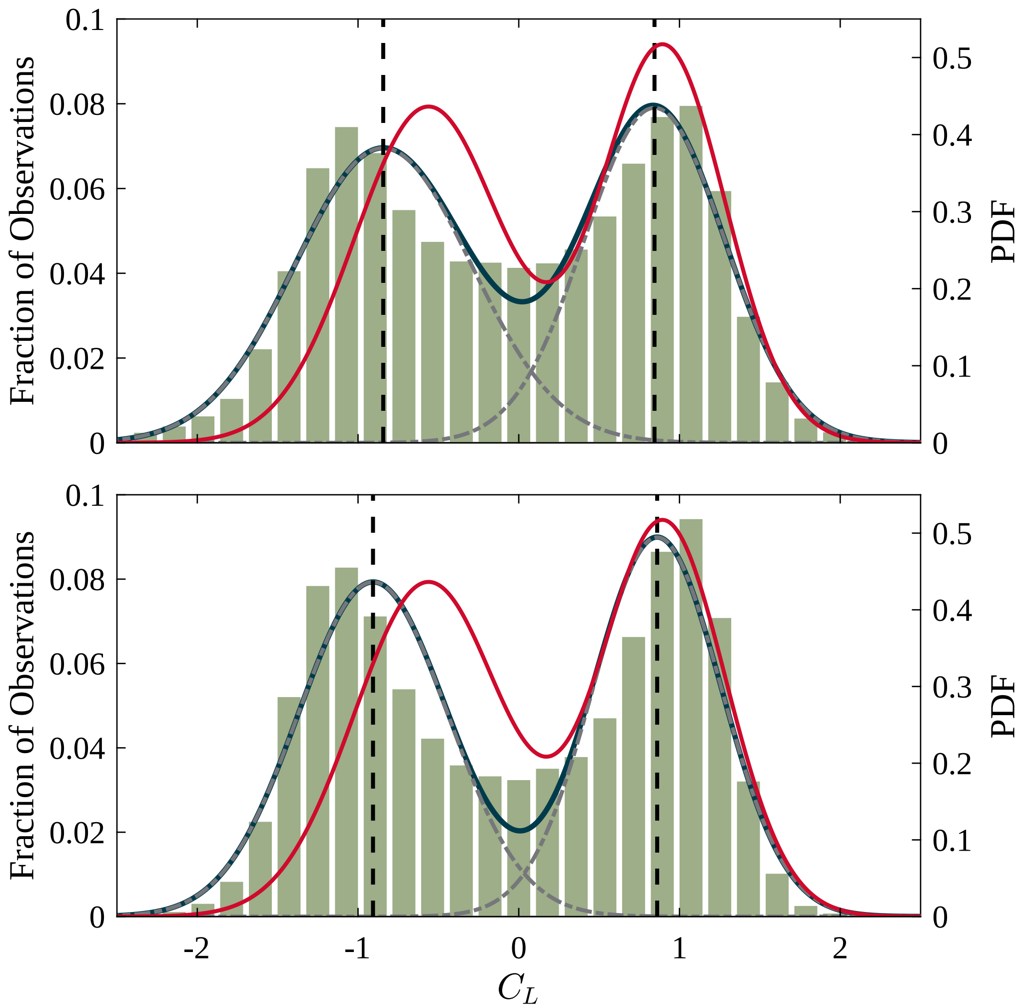

Figure 6.6: Histogram and Gaussian Mixture (GM) distribution fit to C_L data for all recorded Realistic Friction (RF) trials (Top Panel), and High Friction (HF) trials (Bottom Panel). Green bars show histogram counts, normalized by total data points considered (left axis). Line plots show Gaussian Mixture distribution fit to the data, two underlying Normal distributions with different mean and standard deviation values, and GM distribution fit to data from MFRS Experiments (right axis). Vertical dashed lines indicate the two mean values for C_L.

underlying Normal distributions was again fit to the data. These models are shown overlaid on the histograms in Figure 6.6. Similar to the MFRS case, the GM fitting procedure recovers two Normal distributions with similar standard deviations and a proportionality constant of approximately 0.5 for both the RF and HF cases. In contrast to the MFRS data, overlaid in red in Figure 6.6, the peak locations for the distributions in both the HF and RF cases appear to be quite symmetric about zero, with $\mu = [-0.84, 0.84]$ for the RF data, and $\mu = [-0.91, 0.86]$ for the HF data.

This symmetry is expected, since for all RF and HF trials recorded, the airfoil had an angle of attack of 0° and was initially located directly behind the cylinder. It is interesting to note that the impact of the non-zero angle of attack of the airfoil in the MFRS experiments seems to have affected the negative lift behaviour more strongly than the positive lift behaviour: in Figure 6.6, the blue and red curves in both panels seem to overlap for the most positive observed lift values, but diverge for the most negative. This could be due to the onset of separation and other unsteady phenomena that occur when the airfoil approaches its static stall angle, and act to reduce the maximum achievable C_L values for the airfoil when α is large. In addition, there is a bias towards positive mean tunnel positions in the MFRS experiments, which may contribute to the asymmetry in the experienced lift as well.

Comparing the lift distributions for the RF and HF cases pictured in Figure 6.6, the most notable difference between them appears to be the frequency with which the airfoil experiences lift values close to zero. Reduced lift acting on the airfoil is associated with an increase in the airfoil's transverse velocity, and a corresponding reduction in its apparent angle of attack, as discussed in previous chapters. Therefore, for the set of RF trials where the action of friction is moderate, we see that the depth of the valley between the two peaks in the probability distribution is shallower than that for the HF case, where the airfoil spends more time motionless (and therefore tends to experience both larger and more frequent lift extrema). It is interesting to note that the valley in the distribution for the MFRS Experiments is the shallowest of all: this corresponds to the smallest fraction of observations experiencing zero velocity.

To explore this further, we consider histograms of the corresponding observed velocity of the airfoil for the RF and HF trials. Since the dynamics giving rise to the airfoil's motion are largely linear (with very simple forms of nonlinearity added to create friction-like behaviour), we expect that the GM distribution will do a more adequate job capturing the behaviour of the airfoil than for the fully nonlinear MFRS Experiments.

This suspicion is confirmed in Figure 6.7, which shows histograms and corresponding GM distribution fits to the observed velocity for the RF (top) and HF (bottom) cases. A bimodal distribution in the velocity is recovered for both cases, along with a Gaussian pulse around zero velocity similar to that observed for the MFRS Experiments. The pulse width (or the standard deviation of the Normal distribution recovered from the GM fitting procedure with $\mu = 0$) for both the RF and HF trials

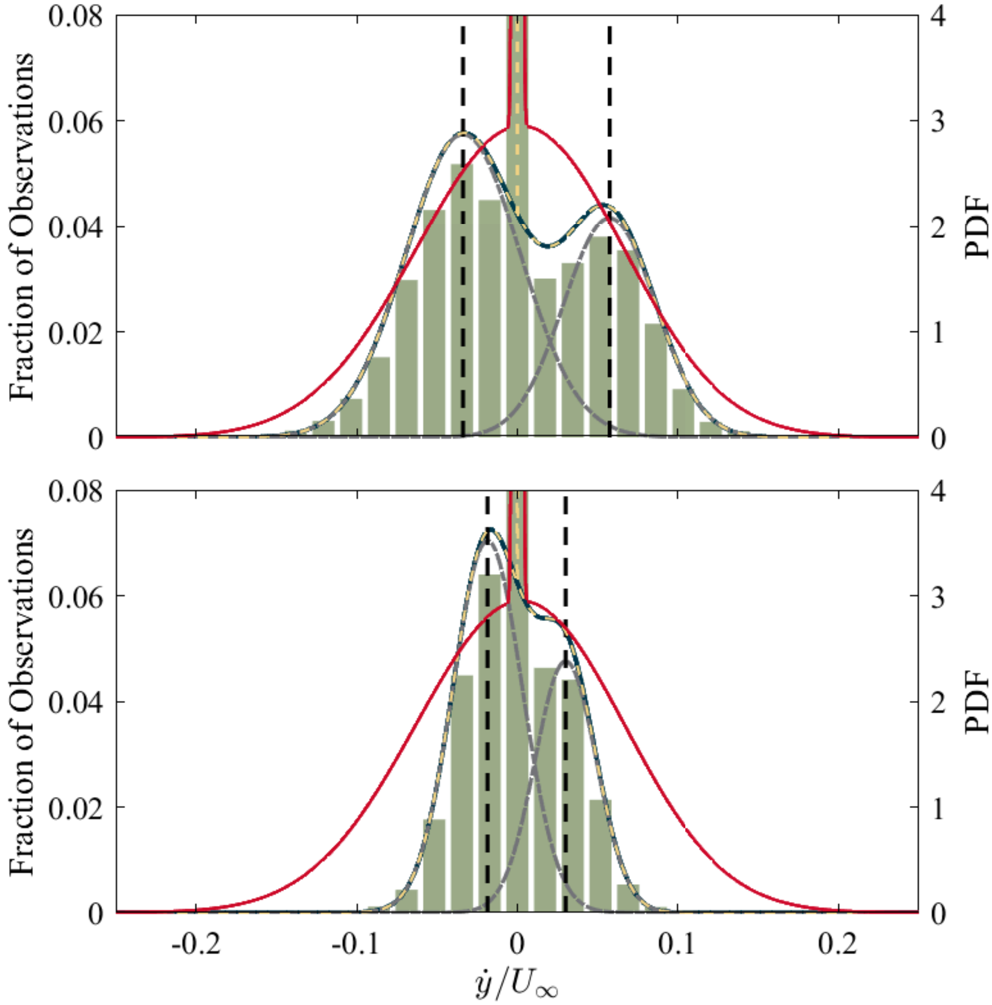

Figure 6.7: Histogram and Gaussian Mixture (GM) distribution fit to y-velocity data for all recorded Realistic Friction (RF - top panel) and High Friction (HF - bottom panel) trials. Green bars show histogram counts, normalized by total data points considered (left axis). Note that the data in the bins centered at zero extend beyond the pictured axes - the plot is truncated for clarity. Line plots show two underlying Normal distributions with different mean and standard deviation values; sum of the two Normal distributions with non-zero mean; and Gaussian Mixture distribution fit to the data (right axis). The velocity distribution from the MFRS Experiments is shown for comparison (). Vertical dashed lines indicate the two non-zero mean values for $\dot{y}/U_\infty$.

is several orders of magnitude smaller even than that recovered for the MFRS data – this reflects the synthetic nature of the simulated frictional forces, which cause a jump discontinuity between small but non-zero velocities and a velocity of zero based on thresholding.

The fraction of observations in the RF trials considered with a velocity less than 1 mm/s is computed as 61%. From the GM distribution fit, the proportion of the PDF attributed to the Normal distribution with $\mu = 0$ is 59%, a close correspondence. This is a result of the manual tuning of the airfoil dynamics discussed in Section 6.3.1, where a goal of approximately 50% stationary observations was used to set the model parameters. It is interesting to note that the fit distribution attributes a slightly larger proportion of probability to the negative velocity distribution, 26% as opposed to 15% for the peak centered in positive velocities. This is reflected in the underlying data, as evidenced by the histogram in the background of Figure 6.7, but seems to be in contrast to the trend observed for lift distributions of the same data. In Figure 6.6, the peak corresponding to positive lift appears higher than that for negative lift; however, the standard deviation of the data for the negative lift peak is larger, and its proportion of probability is actually slightly larger than that for positive lifts. As our friction model acts to first reduce the force acting on the airfoil, then counteract its acceleration, the larger spread of experienced forces in the negative direction is squeezed into a higher peak in negative velocities, with a more limited extent. This could account for the higher peak height in the negative velocity region. The centers of the two peaks for the velocity distribution are $\mu = [-0.034, 0.057]$ m/s. These values fall outside those for the MFRS experiments ($\mu = [-0.026, 0.027]$ m/s); however, by comparing the red and blue/yellow curves in Figure 6.7 we see that in fact, the airfoil is slightly more likely to achieve the highest velocities pictured in the MFRS Experiments than in the realistic modelled friction case.

Although the behaviour of the airfoil with a modelled mounting system corresponding to the RF case is not precisely equivalent to that of the MFRS experiments, the use of the term 'Realistic Friction' to describe this set of mounting parameters is justified. The modelled dynamics have a higher fraction of observations of zero velocity (about 60%, vs 52% for the MFRS experiments), and do not achieve the highest observed velocities as often as in the MFRS case. However, the airfoil spends a relatively larger time at small, non-zero velocities which leads to larger non-zero mean velocities. This trade-off creates a probability distribution for velocity that is qualitatively similar to that extracted from the MFRS data, especially considering the simplicity of the model used to describe the highly complex action of friction.

Trends for the HF case appear qualitatively similar to those for the the RF case discussed above, with a notable reduction in the extent of the estimated PDF into regions of high velocity. For the high friction case, the model severely limits the achievable velocities for the airfoil, and leads to observed velocities that are very tightly clustered near zero. In addition, the airfoil is stationary 70% of the time, both as a fraction of observations and as a relative proportion of probability attributable to the central peak at zero velocity. These factors justify the use of the title 'High Friction' to describe this set of parameters.

As a final, qualitative validation step, Figure 6.8 shows the position as a function of time for several representative trials for both the RF and HF cases. The behaviour seems qualitatively similar to that exhibited in the representative MFRS trials pictured in Figure 6.1, with notable periods of inactivity, as well as those with roughly sinusoidal motion. The peaks in regions of sinusoidal motion are truncated, as the force acting at the position extrema is not immediately able to overcome the stiction level. In addition, though all trials start at a position of zero, there is a clear preference for the system to operate at a slightly negative mean tunnel position. This consistent asymmetry is likely due to slight asymmetry in the mechanical setup (for example, a slight negative angle of attack), but presents an interesting avenue for further study.

6.5 Effect of Increasing Friction on Power and Thrust Production

Based on the discussion in the previous section, the simplified CPFD model implemented on the CTS was used to interrogate the action of friction on the performance of the airfoil, in terms of both power and thrust production. Figure 6.9 shows histograms of the power coefficient achieved by the airfoil for the HF and RF conditions discussed in the previous section, along with a third No Friction (NF) case. This case, as described in Section 6.3.1, has the same spring-mass-damper parameters as the RF and HF cases, but with the action of friction removed. GM distributions were not fit to the C_P observations, as these observations constitute a product between two random variables (C_L and $\dot{y}/U_\infty$), which is not expected to be accurately described by a composite of Normal distributions.

From the Figure, we see that higher friction (shown in the top panel) severely limits the ability of the airfoil to extract energy from the flow. This makes sense, as much of the work done on the airfoil is repurposed into overcoming the effects of friction. Observations are clustered near zero power output, and the percentage of observations with no power output at all is approximately 70% corresponding to times with zero velocity. Considering the middle panel corresponding to the realistic friction case, there are still a large number of observations of zero power

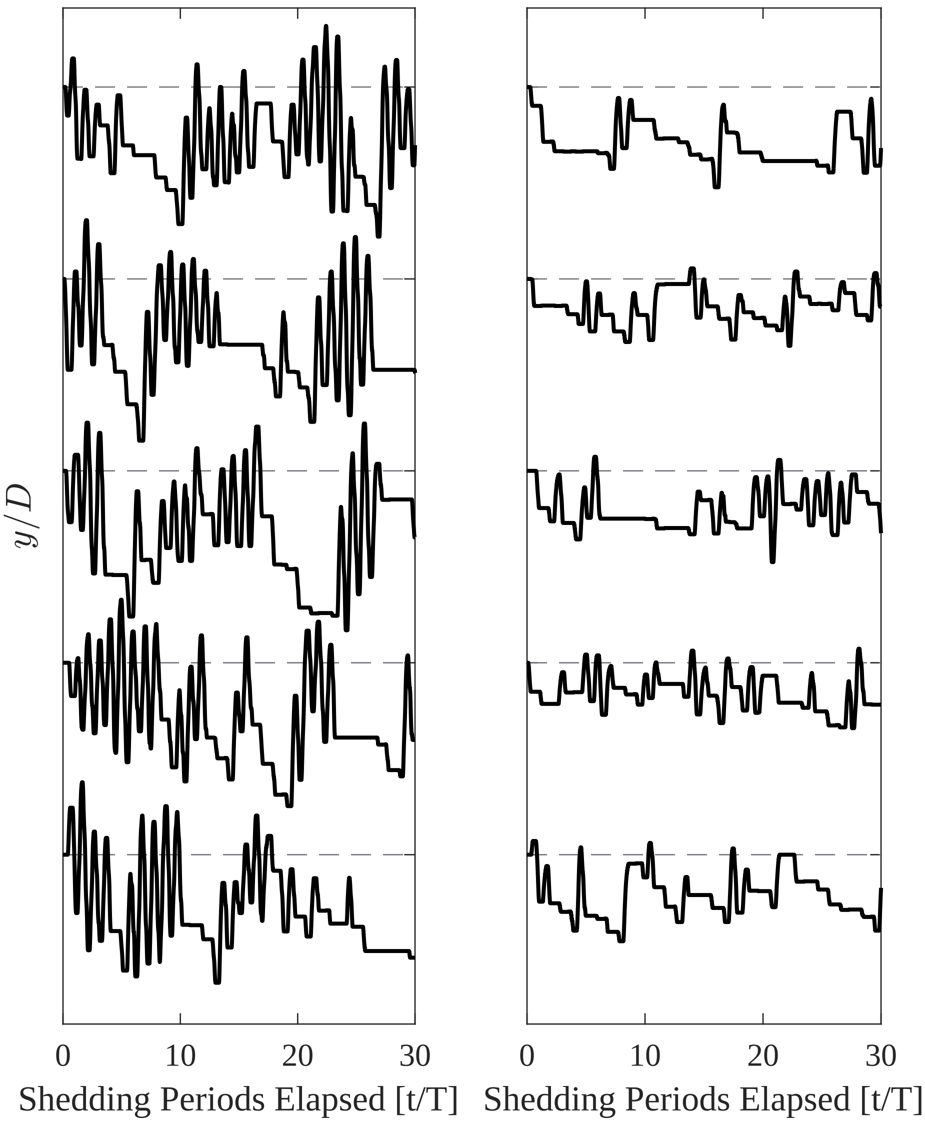

Figure 6.8: Tunnel position as a function of elapsed shedding period for five representative trials with both Realistic (RF) and High Friction (HF) mounting parameters. Black lines indicate position of the airfoil, relative to the cylinder centerline provided as a grey dotted line. Figure shows relative amplitude only. Left Panel: Realistic Friction (RF) Trials; Right Panel: High Friction (HF) Trials.

(approximately 60%); however, the airfoil achieves a limited number of moments of high power production corresponding to times when the applied forcing overcomes the effects of friction and the airfoil begins to move more freely. The frequency with which this occurs is much smaller than in the bottom panel, which presents observations of the no-friction case with the same mounting parameters otherwise.

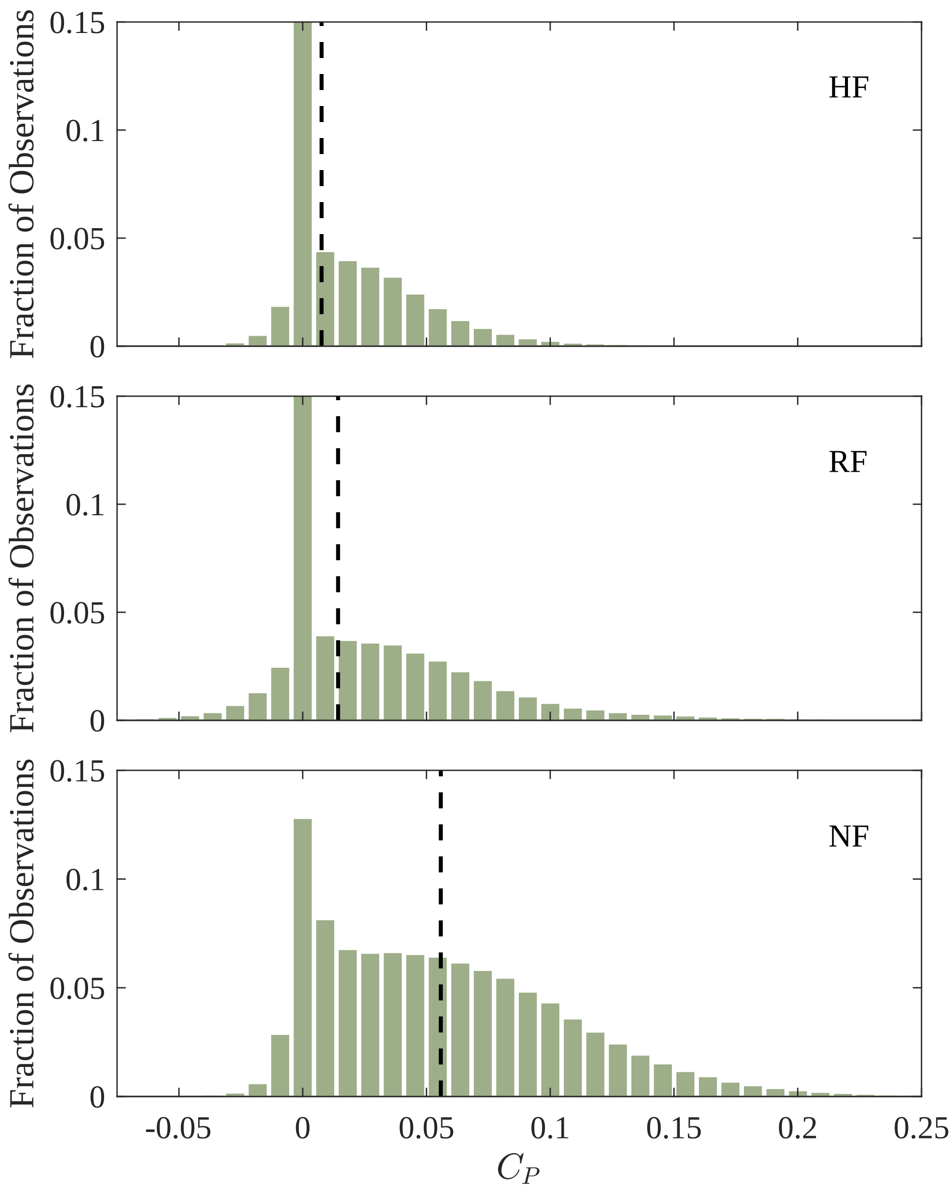

Figure 6.9: Histograms for C_P values observed during High Friction (HF) trials (top), Realistic Friction (RF) trials (middle), and No Friction (NF) trials (bottom). Dashed black lines indicate the overall mean power coefficient from all observations, including those with zero velocity. Note that the bar centered at $C_P = 0$ has a height of 0.7 in the top panel, and 0.6 in the middle panel. Axes are truncated for clarity.

It is interesting to note that there is still a significant number of observations of zero or very small power (though there are far fewer than in the cases with friction).

Comparing the mean and maximum values of C_P in each of the three panels, it seems that the addition of friction to the system drastically reduces the mean C_P

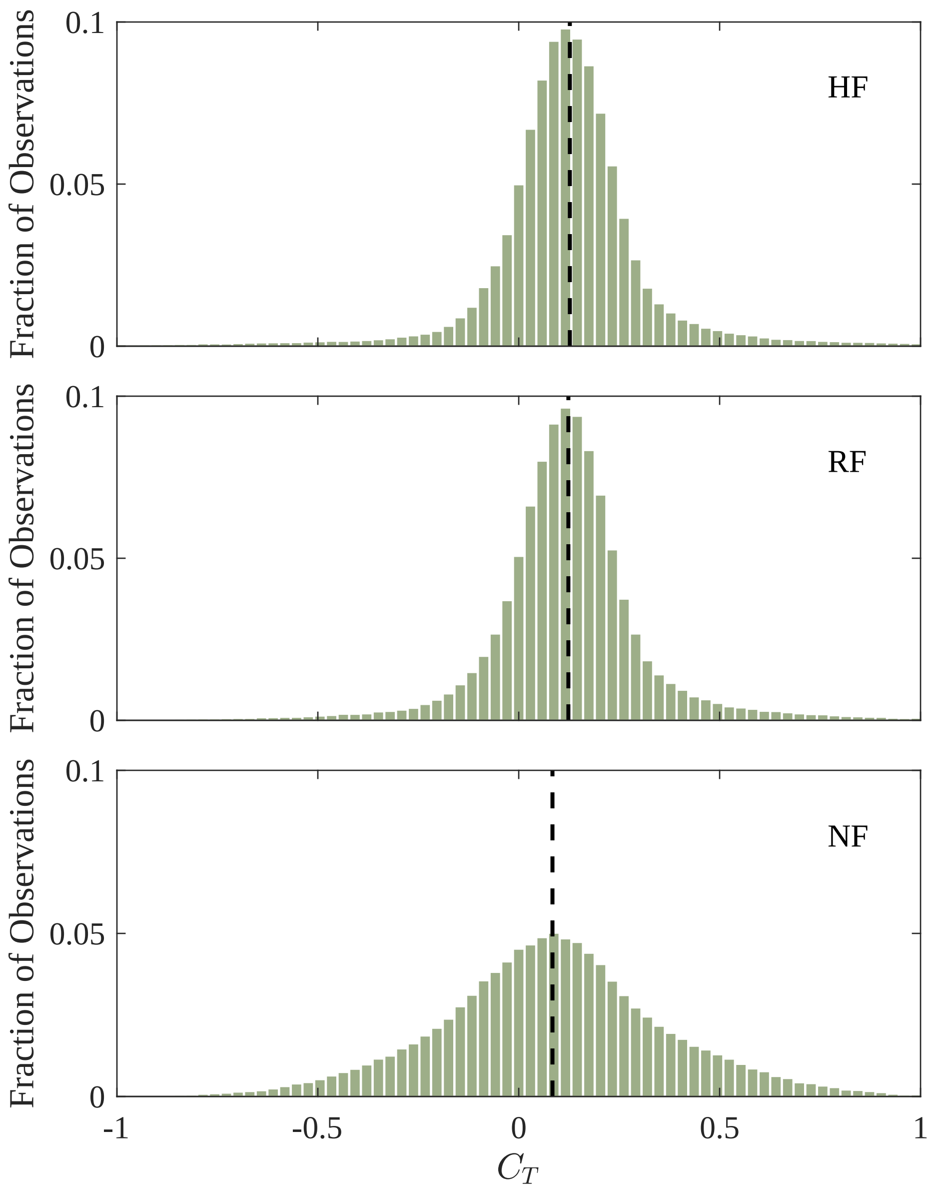

Figure 6.10: Histograms for C_T values observed during High Friction (HF) trials (top), Realistic Friction (RF) trials (middle), and No Friction (NF) trials (bottom). Dashed black lines indicate the overall mean thrust coefficient from all observations.

value experienced by the airfoil, but has a more limited impact on the maximum power achievable by the airfoil. This make sense, as times when the effects of friction are most prevalent correspond to times when the applied force is small (or equivalently, when the airfoil is located at position extrema): this is the portion of a cycle with the smallest power production even without friction, as the force and velocity are both near zero. When the airfoil should be generating its maximum power (at the neutral position behind the cylinder, with the maximum experienced velocity in the friction-free case), friction acts to reduce the value of the maximum velocity experienced. However, as discussed in previous chapters, such a reduction in velocity leads to a higher apparent angle of attack, and therefore a higher lift value. This could help to offset the loss in power production due to the reduced airfoil velocity. Finally, the transient behaviour of the airfoil after motion is initiated leads to a short time where high velocities are produced to 'catch up' with the forcing; this leads to a high velocity moment at high lift, which also contributes to positive power production.

Considering instead the thrust produced by the airfoil, a different trend is apparent. Figure 6.10 shows histograms for the thrust coefficient C_T for each of the three cases (HF, RF and NF) considered. The thrust produced by the airfoil in both cases with simulated friction appears to be very similar, exhibiting mean values for C_T of 0.29 in both cases. However, for the zero friction case, the mean thrust coefficient is noticeably reduced with a mean value of 0.19. In addition to the reduction in the mean value, there is also a much larger spread in the observed values for C_T, leading to a much larger fraction of observations where the airfoil is experiencing net drag.

This aligns with the discussion of passive thrust production in this system presented in Chapters 3 and 4. We expect that as the amplitude of the airfoil motion increases in-phase with the flow velocity (and therefore the fluid forcing), the contribution to passive thrust production from the Katzmayr Effect is reduced along with the induced angle of attack for the airfoil. Since the airfoil operates in a slaloming mode of interaction with the oncoming vortices, a condition enforced by the tuning of its passive mounting system for the benefit of power extraction, its effectiveness as a propulsor is limited. The amplitude of oscillation in the NF case presented here is the largest amplitude of any tests conducted in this study, and correspondingly exhibits a low mean thrust coefficient.

6.6 Fluid-Structure Interaction and Emergent Behaviours

Though there is large cycle-to-cycle variability in the behaviour of the airfoil subject to friction (which motivated our consideration of these systems from a statistical perspective), considering time series for forces and power production yields additional insight into flow-induced behaviour. The primary difference in behaviour induced by friction is start-stop (stiction) behaviour. Stiction acts to increase the effective angle of attack when the airfoil remains stationary as applied force increases – this could lead to an increased tendency to form Leading Edge Vortices (LEVs) relative to smooth motions of the airfoil with no friction. However, based on the idealized model for friction presented in Section 6.3.1 we also expect that the transient behaviour of our system in response to an impulsive jump in experienced force could lead to locally larger values of airfoil velocity than those observed for the basic Passive Captive cases in Chapter 4. This would lead to suppression of LEV formation and/or shedding at later points in the cycle. In addition, the action of friction is very sensitive to local flow conditions. The nonlinearity in our model is activated when the flow forcing exceeds a specific threshold – the vortex shedding phase at which this threshold is exceeded may vary from cycle to cycle. Unlike the friction-free case where the airfoil was shown to smoothly adapt to fluctuations in oncoming forcing, when the airfoil is governed by stiction such variations may strongly affect the airfoil's interactions with oncoming vorticity.

To investigate these emergent behaviours induced by friction, we consider Particle Image Velocimetry (PIV) images of the flow field in the region of the airfoil in the RF and HF cases. As discussed for the 2D Passive Captive airfoil motions in Chapter 5, in the present case there is also no consistent phase-locked behaviour, which is evidenced here by the position traces of the airfoil shown for several trials in Figure 6.8. Therefore, phase averaging of PIV fields is not appropriate, and we instead show spatially and temporally filtered instantaneous contours of the Γ_2 Criterion. The filtering process for time-series Γ_2 data is discussed in detail in Chapter 2.

We are particularly interested in moments when the effective angle of attack for the airfoil is high, and we are therefore more likely to observe the formation and potential shedding of an LEV. Unlike the friction-free airfoil discussed in Chapters 3 and 4, at large values of lift (corresponding to large effective AoA) the airfoil is not necessarily moving in the same direction as the forcing; it has often been kept stationary by the action of stiction. This leads to a phase difference between the location of maximum α_{eff} and maximum velocity. This phase difference creates an opportunity for cyclic formation and shedding of an LEV.

Such a moment is shown in the first (left-hand) column of Figure 6.11, which shows several snapshots of the Γ_2 Criterion for a representative run where the mounting system is specified using parameters for the RF case, as described in Table 6.3 earlier. The moments in time described by these snapshots are indicated by the blue dashed lines in the bottom two line plots, which show α_{eff} and the airfoil's velocity and power production as a function of time. The first (top left) snapshot corresponds

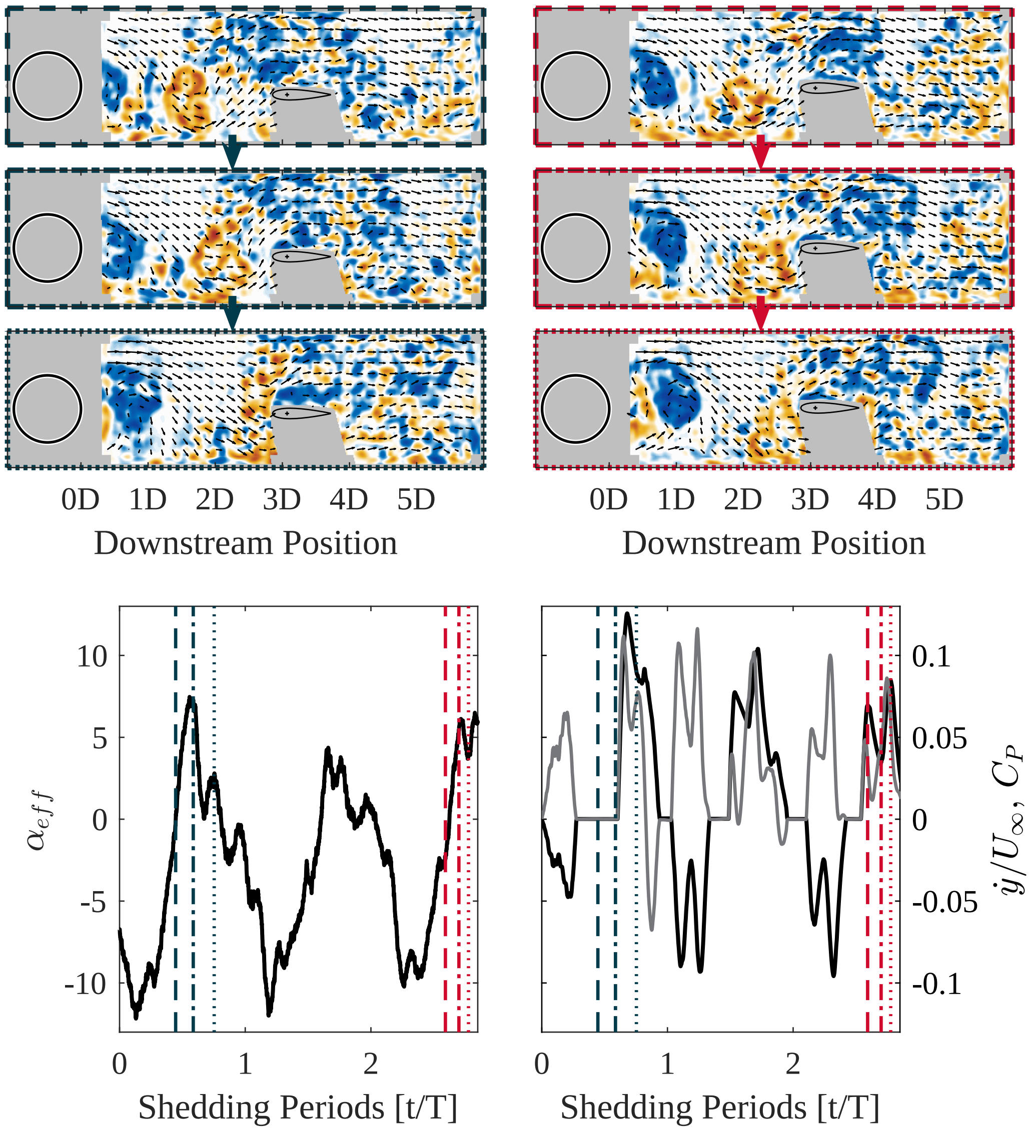

Figure 6.11: Representative behaviour of a Passive Captive airfoil with a mounting system described by Realistic Friction (RF) parameters, described in Table 6.3, captured with simultaneous measurement of the surrounding flow field. Bottom two figures show behaviour of the airfoil as a function of time throughout the run. Left Bottom: Effective angle of attack α_{eff}; Right Bottom: ▬ Airfoil velocity normalized by U_∞; ▬ Power extraction by airfoil. The two columns at the top of the figure show snapshots of the Γ_2 Criterion taken at time instants indicated in the bottom two plots: dark blue snapshots correspond to earlier times, and red snapshots to later times as indicated. Line types on figure borders indicate corresponding time in bottom figures; time proceeds from top to bottom in each column as indicated large arrows. Black arrows overlaid on contour plots show flow velocity magnitude and direction.

to a time when the airfoil is stationary due to the action of stiction. The effective angle of attack is small, and an oncoming CW rotating (negative signed) vortex is just beginning to pass over the airfoil. From our discussion of the dynamics of this system in previous chapters, the airfoil should be near a position extremum and beginning to change direction to slalom away from the oncoming CCW rotating (positive signed) vortex approaching from upstream. In this case, the airfoil is stationary as the force acting on the airfoil is not yet large enough to initiate motion in the positive transverse direction.

In the second left-hand panel, the airfoil is experiencing a local maximum in α_{eff} as it remains stationary. At this moment however, the fluid forcing has just overcome the stiction threshold and motion is about to be initiated, as seen in the right-hand bottom plot of the airfoil velocity. Though α_{eff} is still below the static stall angle, there is some evidence of coherent leading-edge vorticity forming in this snapshot. As the airfoil is engulfed in a region of CW rotating cylinder vorticity it is challenging to confidently distinguish airfoil-derived vorticity specifically; however there appears to be a localized region of relatively strong CW vorticity located very close to the leading edge, which could indicate the presence of an LEV. In addition, accurately determining α_{eff} in this setting is a challenge, since the method used in this thesis likely represents an underestimate; for more discussion of the determination of α_{eff}, please see Chapter 2.

In the third left-hand panel, the airfoil is experiencing a large positive velocity, which has acted to reduce its apparent angle of attack from the previous maximum. There is a localized region of CW rotating vorticity over the suction surface, indicating flow separation: this vorticity is likely a combination of leading-edge vorticity shed from the airfoil and cylinder-derived vorticity which has been localized on the suction surface due to the upwards velocity of the airfoil. This localized vorticity acts to improve the power production of the airfoil by maintaining a low-pressure zone above the airfoil as it moves up with high velocity – the combination of high lift and large $\dot{y}$ lead to a local maximum in power production in this region. This phenomenon is discussed at length in previous chapters, as well as in literature such as Kinsey and Dumas (2008).

Very similar trends can be seen in Figure 6.12, which shows a similar process for the airfoil with the high-friction (HF) mounting configuration. In this figure, we see the same phase lead of α_{eff} relative to the airfoil's velocity for times corresponding to both the right-hand and left-hand columns of snapshots. For the left-hand series,

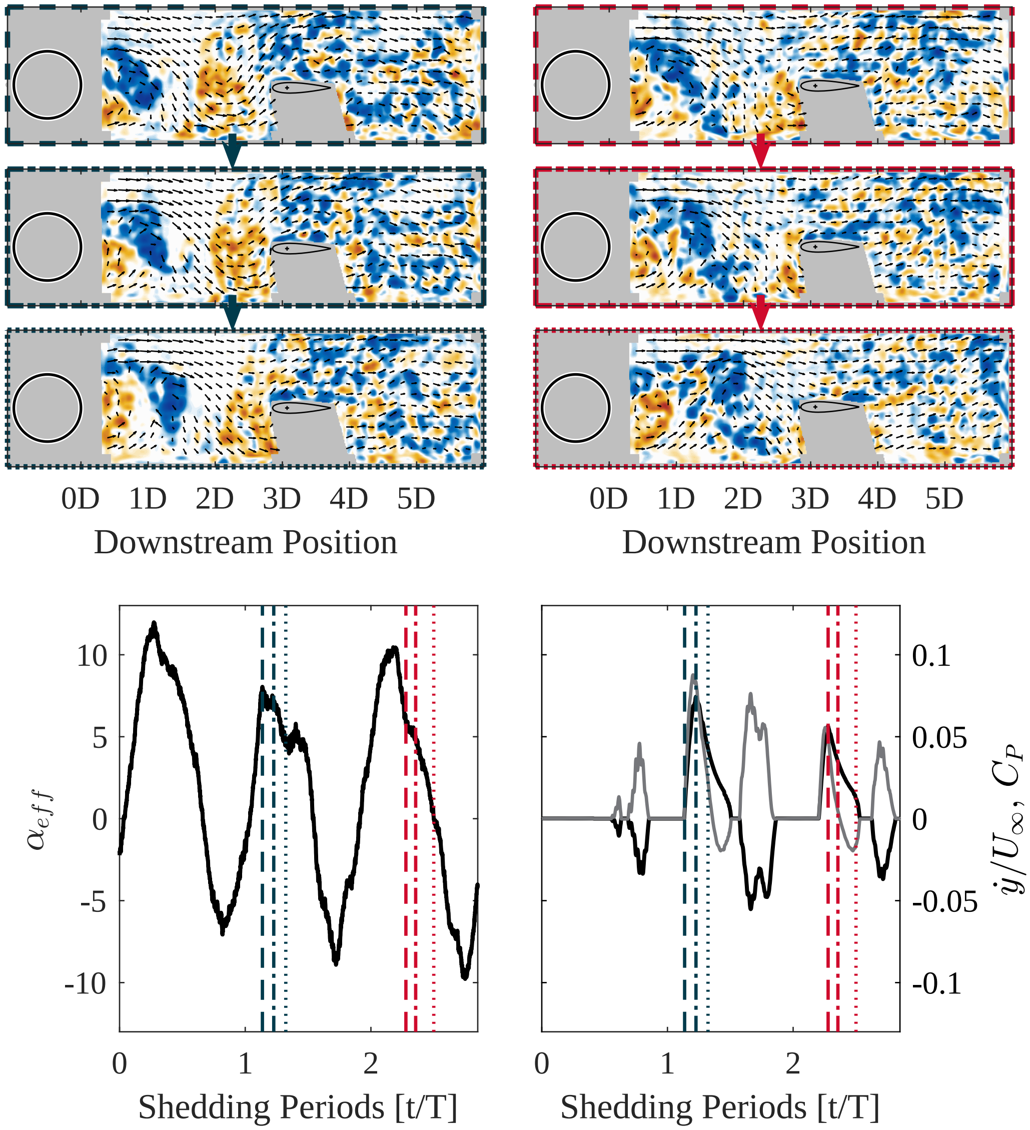

Figure 6.12: Representative behaviour of a Passive Captive airfoil with a mounting system described by High Friction (HF) parameters, described in Table 6.3, captured with simultaneous measurement of the surrounding flow field. Bottom two figures show behaviour of the airfoil as a function of time throughout the run. Left Bottom: Effective angle of attack α_{eff}; Right Bottom: ▬ Airfoil velocity normalized by U_∞; ▬ Power extraction by airfoil. The two columns at the top of the figure show snapshots of the Γ_2 Criterion taken at time instants indicated in the bottom two plots: dark blue snapshots correspond to earlier times, and red snapshots to later times as indicated. Line types on figure borders indicate corresponding time in bottom figures; time proceeds from top to bottom in each column as indicated large arrows. Black arrows overlaid on contour plots show flow velocity magnitude and direction.

the moments indicated show evidence of the formation of an LEV near the airfoil's surface while the velocity is high – especially in the final panel, there remains CW vorticity located in the region of the suction side of the airfoil. In the right-hand series, a similar trend is noted, but time relative to vortex shedding has advanced, such that as the airfoil returns to rest in the final panel, the leading-edge vorticity that was once localized near the leading edge appears to have been shed.

The primary effect of friction in these interactions is to change the phase relationship between velocity and α_{eff} relative to the friction-free case. Instead of having velocity, force, and α_{eff} in-phase over all portions of a cycle, as shown for the friction-free Passive Captive airfoil motions in Chapter 4, here the maximum value of α_{eff} occurs earlier than the maximum in velocity. This is required in order to initiate airfoil motion after the action of friction has held the airfoil at rest near a position extremum. As a result, the airfoil experiences an impulsive start and quickly achieves a velocity maximum shortly thereafter; however, by this point the lift acting on the airfoil is starting to reduce, and it cannot maintain these high velocities further. During this impulsive motion, the velocity and power production for the airfoil are larger than those experienced (in a phase-averaged sense) by the Case 0 Passive Captive airfoil, but the periods of stagnation between these impulsive airfoil motions remove the ability to harvest useful power during lower-lift moments. In addition, due to the impulsive nature of the airfoil's acceleration, there are noticeable times when the foil's velocity is still positive, but the lift has started to act in the opposite direction leading to negative power output (stored energy in the mounting system doing work on the flow). This behaviour arises due to the choice of friction model particularly in regions where the applied force is small. Since the force of friction acts to cancel out small forcing on the airfoil, it does not return to its neutral position as quickly as in the friction-free case, leading to these prominent regions of negative power.

It is interesting to note that there appears to be asymmetry in the power production characteristics between regions of positive and regions of negative lift - the negative power regions apparent for both the RF and HF trials pictured in Figures 6.11 and 6.12 seem to result only for positive values of lift. This trend is consistent beyond the trials presented here, though there are a few isolated incidents of negative power production occurring for negative lift-induced power peaks. Though the cause is unknown, and could in fact be due to a subtle modelling error or other non-physical phenomena, this observation could also be linked to asymmetry in the mean tunnel position for the airfoil. For all trials presented here, the airfoil established a mean tunnel position that was slightly negative, unlike in the friction free case where the mean position was consistently very close to zero. The explanation for these (potentially linked) behaviours provides an interesting avenue for future work.

In addition to these detrimental effects for power production, the behaviour of the airfoil subject to the action of friction exhibits very high cycle-to-cycle variability in its interactions with oncoming vorticity. Considering the right-hand column of snapshots in Figure 6.11 and the corresponding time instants indicated with red dashed lines, a different type of interaction is apparent. In this case, the first snapshot indicates a moment when the effective angle of attack is small (and even slightly negative); however motion has been initiated, and the airfoil is moving in the positive direction with a local maximum velocity. This is due to the 'slingshot' effect of releasing the hold of stiction on the airfoil. Shortly after this however, at a time when the airfoil is now achieving a local maximum in α_{eff}, the velocity of the airfoil has decreased. In the first right-hand snapshot, there is little evidence of LEV formation despite a large CW rotating vortex passing over top of the airfoil, similar to the situation in the second panel on the left-hand side (where leading-edge vorticity is seen to form). The airfoil's positive velocity in the right-hand sequence seems to suppress the formation of an LEV, and contributes to the small experienced angle of attack in the first panel. In the second panel as the airfoil slows and α_{eff} catches up, exhibiting a local maximum, there is evidence for the formation of a small amount of leading-edge vorticity near the airfoil surface. In the third panel, as the airfoil continues to move up and maintain a large value of α_{eff}, this leading edge vorticity again remains localized near the suction side similar to what was observed in the previous cycle.

Figure 6.13 provides a closer look at the lift force acting on the airfoil that leads to the observed divergence in the phase relationship between velocity and effective angle of attack during this representative trial. There is a strong frequency content peak in the lift signal (considering both RF and HF trials) at a frequency of approximately 3.6 Hz, indicated in the top panel of the Figure. The cause of this peak is unknown; however it is likely linked to the structural response of the airfoil to the impulsive starts it experiences frequently in cases where stiction is dominant. This frequency content causes large-amplitude local oscillations in the lift recorded during these trials superimposed onto the base frequency due to vortex shedding, as pictured in the bottom panel of Figure 6.13.

In the case of the second column of snapshots in Figure 6.11, the motion of the airfoil

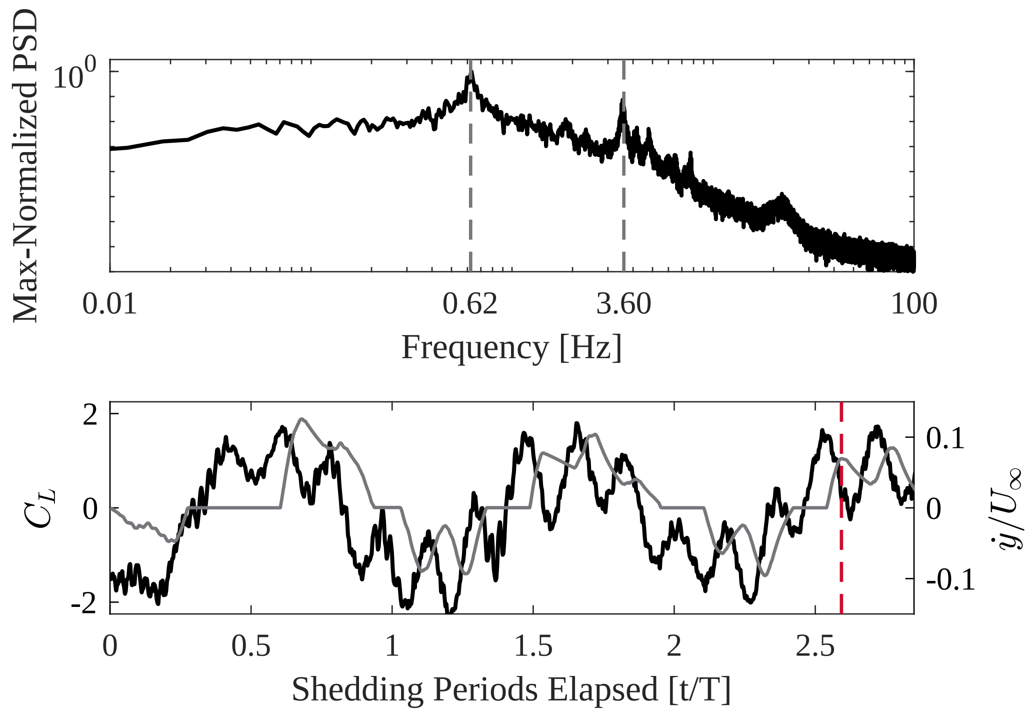

Figure 6.13: Aggregated frequency content and single realization of lift coefficient of a Passive Captive airfoil with idealized friction. Top Panel: Power spectral density of lift coefficient for representative RF and HF trials presented in Figure 6.8. PSD is normalized by its maximum value which occurs at the frequency of vortex shedding, 0.62 Hz as indicated. An additional peak at 3.60 Hz is also pictured. Bottom Panel: ▬ Coefficient of Lift for the representative RF trial shown in Figure 6.11; ▬ y-velocity of the airfoil normalized by U_∞; and ‐ ‐ Time stamp corresponding to the first snapshot in the right-hand column of Figure 6.11.

is initiated 'early' when a local lift oscillation surmounts the stiction threshold. This can be seen just prior to the location of the red dashed line in Figure 6.13, which indicates the time of the first snapshot in the right-hand column of Figure 6.11. The effective angle of attack, based on the bulk behaviour of the flow and not the measured lift, lags the initiation of motion, which attempts to react to the locally high experienced lift by impulsively starting and accelerating to a velocity maximum. The observation that this local oscillation in lift is not reflected in the angle of attack (and therefore is not induced by bulk motion of the flow in the region of the airfoil) supports the hypothesis that these 3.6 Hz oscillations are induced in the structure of the airfoil itself. As the local lift oscillation amplitude is decreased, the large airfoil velocity is maintained by the bulk increase in flow velocity, leading to a maximum region in lift experienced at high velocity.

Once airfoil motion is initiated, the behaviour of the system with friction is relatively similar to that for the friction-free cases discussed in the previous chapters; however the onset of motion is very sensitive to instantaneous forces, including those induced by self-oscillation of the airfoil, for example. Though this particular behaviour is specific to the physical setup and implementation discussed in this chapter, the idea that impulsive starts and stops induced by the action of stiction may interact with structural harmonics in a system to quantitatively alter their behaviour is a true concern for engineering systems of this type. In some sense, this is an emergent behaviour for our system due to friction, though in this case a highly setup-specific and detrimental one.

6.7 Chapter 6 Interim Summary and Conclusions

This chapter presented studies of the behaviour of both a truly passive, all-mechanical mounting system, the Mechanical Free-Response System (MFRS), and a Passive Captive realization of mounting dynamics that are subject to the effects of friction. A simplified model for such effects based on observations from the MFRS was created, and its behaviour was compared to the original all-mechanical system in both a statistical and instantaneous sense. The following conclusions can be drawn from this work:

1) A highly simplified model for the action of friction on the airfoil produces airfoil motion that is qualitatively similar to that produced by the MFRS, which is subject to ‘real’ engineering friction. By tuning parameters in the friction model, statistical descriptions of the MFRS and Passive Captive implementations with friction can be made reasonably similar, considering the complexity of the underlying dynamics of real friction.

2) The addition of friction to the mounting system dynamics changes the phase relationship between force and velocity, so although larger instantaneous velocities and power coefficients are often experienced, the overall power extracted is reduced.

3) The action of friction introduces oscillations into the measured forcing, likely due to structural excitation of the airfoil from impulsive starting and stopping motions. Structural excitation interacts with stiction to create large variability in the timing between the lift force and an impulsive start, changing the way the airfoil interacts with oncoming vorticity.

Management of the interaction of friction with the behaviour of a fully passive flapping-foil energy harvester is necessary for both improved understanding and potential optimization of such a device. Simplified frictional dynamics presented here highlight how cycle-to-cycle changes in the response of the airfoil due to instantaneous flow conditions can cause unexpected and often detrimental behaviour. Building an understanding of conditions which give rise to such behaviours may allow for the implementation of control intervention(s), to attempt to mitigate some of the consequences of mechanism friction on energy harvesting potential.

Chapter 7

CONCLUSIONS & FUTURE WORK

The present work has demonstrated that a flapping foil downstream from a circular cylinder, whether actively driven or passively reacting to structures shed in the wake, is able to simultaneously extract net power from the oncoming flow while producing net thrust larger than its net drag. This was first demonstrated for the simplest case, an airfoil driven through a sinusoidal trajectory with a frequency matched to that of vortex shedding from the cylinder. Such an airfoil, flapping with a relatively small amplitude compared to the flow that surrounds it, can be made to leverage the unsteadiness in the oncoming fluid to generate aerodynamic forces akin to those generated by active swimmers or fliers. By tuning the phase of the foil's motion such that it slaloms between periodically shed vortices from the cylinder, the foil extracts benefits similar to those due to the active formation of leading-edge vortices in flapping flight. In addition, oscillation in the direction of the oncoming flow relative to the foil's motion due to cyclic vortex shedding allows for simultaneous production of appropriate thrust to 'tack' upstream like a sail, the magnitude of which can either be enhanced or diminished by unsteady aerodynamic effects near the foil's surface.

Actively driving the flapping foil provides a useful starting point from which to study this dual passive thrust and power production phenomenon, as it involves a relatively simple experimental setup, and a results in a repeatable trajectory for airfoil motion. However, the forces that the airfoil experiences due to interactions with oncoming vorticity are themselves unsteady; thus, maintaining the desired phase relationship between the motion of the foil and the oncoming vorticity, a critical factor in the performance of the system, becomes a challenge. By contrast, a fully passive flapping foil reacts to such changes in the forcing it experiences; however to enforce the desired phase relationship, the system (mechanical or otherwise) which permits foil motion in response to oncoming forcing must be tuned.

In this work, we have shown that the behaviour of a fully passive flapping foil which is compliantly mounted through a linear spring-mass-damper system (in this case implemented with the aid of Cyber-Physical Fluid Dynamics (CPFD)) can be tuned to closely mirror the slaloming behaviour demonstrated for the actively driven foil.

Under those conditions, the fully passive airfoil also extracts net power from the flow while simultaneously producing net thrust. In addition, the mounting parameters can be tuned to improve the power extraction performance of the foil while respecting relevant engineering constraints on the system. We have presented a tuning method based on linear second-order systems theory by which the power extraction performance of the fully passive flapping foil was improved by approximately 40% in experiments, while the mass of the foil and its maximum gain from force to position at any frequency remained the same as in our uninformed choice of system parameters. These constraints represent physically reasonable engineering considerations for the design of fully passive flapping foil energy harvesters.

In addition to transverse flapping, the motion of a fully passive airfoil allowed to move in both the transverse and streamwise directions simultaneously was demonstrated. Due to the net production of thrust by the airfoil over one vortex shedding cycle, the foil was shown to passively translate upstream against the mean flow direction, in analogy to experiments performed with a dead trout by Beal et al. (2006). We show definitively that the foil achieves this passive propulsion while it is outside of the suction zone induced by the upstream circular cylinder, solely due to interactions with the unsteady oncoming flow.

Finally, the role of friction in the behaviour of an engineering system used to achieve similar passive airfoil motion, without the aid of CPFD, was briefly investigated. Based on these experiments, which considered an all-mechanical mounting system which allowed the foil to translate freely in the transverse direction but exhibited large frictional resistance, a highly simplified model for the contribution of friction to airfoil dynamics was presented. This model was used to explore both the detrimental effects of friction on passive energy harvesting performance, as well as to quantify the changes to the interactions with oncoming vorticity due to friction-mediated behavioural modifications.

This modelling effort presents the first of several avenues of further inquiry available on the topic(s) presented in this thesis. The behaviour of a fully passive flapping foil with nonlinear dynamics (such as those imposed by friction) has been shown to exhibit quantitatively different interactions with oncoming flow than observed for an ideal, linear mounting system. These nonlinear interactions, coupled with the potentially nonlinear response of the flow to airfoil motion give rise to a rich collection of potential behaviours, which may have significant impacts on both thrust production and energy harvesting performance. Therefore, to understand and control the occurrences of such behaviours in a real engineering system, more sophisticated models for the dynamics governing the airfoil's response than the one presented in this thesis are required. Such models could enable low effort control interventions to prevent negative interactions from occurring, similar to ongoing research efforts in the areas of gust response and mitigation. Alternatively, they could reveal mounting system architectures (either passive or active) that further improve power extraction performance of such small-scale passive energy harvesters.

Of course, there are several improvements and extensions to the experiments performed in support of this thesis that represent open areas of future work. In addition to the brief discussion of driven airfoil motions at a static geometric angle of attack of $\alpha_0 = \pm 10°$ presented in this thesis, additional experimental data for the case $\alpha_0 = 5°$, and where the foil is driven through the same sinusoidal trajectory but with several different mean positions offset from the cylinder centerline are available. Preliminary analyses of the latter data have indicated that a large enough mean offset leads to suppression of thrust and/or power production during one half of each vortex shedding cycle; continued analysis of this data may provide further insight into the particular fluid-structure interactions which most strongly influence thrust and power production for systems of this type.

Nonlinear interactions between the foil and the fluid surrounding it also influence the power and thrust production performance of a fully passive flapping foil. In this thesis, it was found that tuning the mounting system for such a passive foil based on results from linear systems theory did lead to improvements in power extraction performance; however, such improvements consistently under-performed relative to predicted values. One contributor to the discrepancy between linear theory and real airfoil performance is the assumption that changing the behaviour of the airfoil has no impact on the experienced forcing. This is of course false, as unsteady effects such as flow separation and dynamic stall, the importance of which was found to be strongly mediated by airfoil motion, lead to significant changes in aerodynamic forcing. Thus, explicit consideration of how changes to the behaviour of the airfoil through its mounting system may initiate such unsteady events in the flow could lead to both improved agreement between predicted and realized power extraction potential, as well as highlighting opportunities to tune mounting system parameters specifically to take advantage of beneficial unsteadiness available to the foil.

BIBLIOGRAPHY

Andersen, A., Bohr, T., Schnipper, T., & Walther, J. H. (2017). Wake structure and thrust generation for a flapping foil in two-dimensional flow. *Journal of Fluid Mechanics*, *812*(R4), R4:1–12.

Anderson, J. M., Streitlien, K., Barrett, D. S., & Triantafyllou, M. S. (1998). Oscillating Foils of High Propulsive Efficiency. *Journal of Fluid Mechanics*, *360*, 41–72.

Anderson, J. D. J. (2011). *Fundamentals of Aerodynamics* (5th ed.). McGraw-Hill.

Andreu-Angulo, I., Babinsky, H., Biler, H., Girguis, S., & Jones, A. R. (2020). Effect of Transverse Gust Velocity Profiles. *AIAA Journal*, *58*(12), 5123–5133.

Armstrong-Hélouvry, B., Dupont, P., & Canudas De Wit, C. (1994). A Survey of Models, Analysis Tools and Compensation Methods for the Control of Machines with Friction. *Automatica*, *30*(7), 1083–1138.

Beal, D. N., Hover, F. S., Triantafyllou, M. S., Liao, J. C., & Lauder, G. V. (2006). Passive propulsion in vortex wakes. *Journal of Fluid Mechanics*, *549*, 385–402.

Bearman, P., & Zdravkovich, M. (1978). Flow around a circular cylinder near a plane boundary. *Journal of Fluid Mechanics*, *89*, 33–47.

Berkooz, G., Holmes, P., & Lumley, J. L. (1993). The Proper Orthogonal Decomposition in the Analysis of Turbulent Flows. *Annual Review of Fluid Mechanics*, *25*, 539–575.

Bloor, M. S. (1964). The transition to turbulence in the wake of a circular cylinder. *Journal of Fluid Mechanics*, *19*(2), 290–304.

Brunton, S. L., Proctor, J. L., & Kutz, N. J. (2016). Discovering governing equations from data by sparse identification of nonlinear dynamical systems. *PNAS*, *113*(15), 3932–3937.

Canudas de Wit, C., Olsson, H., Åström, K. J., & Lischinsky, P. (1995). A new model for control systems with friction. *IEEE Transactions on Automatic Control*, *40*(3), 419–425.

Cheever, E. (2021). Approximation of magnitude and phase of underdamped second order systems [Copyright 2005-2021 Erik Cheever. Available from: https://lpsa.swarthmore.edu/Bode/underdamped/underdampedApprox.html. Accessed: 2022-02-21].

Corke, T. C., & Thomas, F. O. (2015). Dynamic Stall in Pitching Airfoils: Aerodynamic Damping and Compressibility Effects. *Annual Review of Fluid Mechanics*, *47*, 479–505.

Corkery, S. J., Babinsky, H., & Graham, W. R. (2019). Quantification of added-mass effects using particle image velocimetry data for a translating and rotating flat plate. *Journal of Fluid Mechanics*, *870*, 492–518.

Cros, A., Llamas, B. F., & Hernández, E. S. (2018). Vortical patterns generated by flapping foils of variable aspect ratio chord-to-thickness. *Experiments in Fluids*, *59*(152), 1–9.

Dennis, D. J. C., & Nickels, T. B. (2008). On the limitations of Taylor's hypothesis in constructing long structures in a turbulent boundary layer. *Journal of Fluid Mechanics*, *614*, 197–206.

Dunne, R., & McKeon, B. J. (2015). Dynamic stall on a pitching and surging airfoil. *Experiments in Fluids*, *56*(157).

Eldredge, J. D., & Jones, A. R. (2019). Leading-Edge Vortices: Mechanics and Modelling. *Annual Review of Fluid Mechanics*, *51*, 75–104.

Eldredge, J. D., & Pisani, D. (2008). Passive locomotion of a simple articulated fish-like system in the wake of an obstacle. *Journal of Fluid Mechanics*, *607*, 279–288.

Gopalkrishnan, R., Triantafyllou, M. S., Triantafyllou, G. S., & Barrett, D. (1994). Active vorticity control in a shear flow using a flapping foil. *Journal of Fluid Mechanics*, *274*, 1–21.

Graftieaux, L., Michard, M., & Grosjean, N. (2001). Combining PIV, POD, and vortex identification algorithms for the study of unsteady turbulent swirling flows. *Measurement Science and Technology*, *12*, 1422–1429.

Hang, H., Heydari, S., Costello, J. H., & Kanso, E. (2022). Active tail flexion in concert with passive hydrodynamic forces improves swimming speed and efficiency. *Journal of Fluid Mechanics*, *932*(A35), A35–1:25.

Hover, F. S., & Triantafyllou, M. S. (2000). *Combined simulation with real-time force feedback: A new tool for experimental fluid mechanics* (T. E. Djaferis & I. C. Schick, Eds.). Springer US. https://doi.org/10.1007/978-1-4615-5223-9_33

Huang, S., Liang, W., & Tan, K. (2019). Intelligent Friction Compensation: A Review. *IEEE/ASME Transactions on Mechatronics*, *24*(4), 1763–1774.

Huang, Y., & Green, M. A. (2015). Detection and tracking of vortex phenomena using Lagrangian coherent structures. *Experiments in Fluids*, *56*(147), 147–1:12.

Hufstedler, E. A. L. (2017). *Experimental Generation and Modelling of Vortical Gusts and Their Interactions with an Airfoil* (Doctoral dissertation). California Institute of Technology.

Hufstedler, E. A. L., & McKeon, B. J. (2019). Vortical Gusts: Experimental Generation and Interaction with Wing. *AIAA Journal*, *57*(3), 921–931.

Huynh, D. P. (2019). *Spatio-Temporal Response of a Compliant-Wall, Turbulent Boundary Layer System to Dynamic Roughness Forcing* (Doctoral dissertation). California Institute of Technology.

Jarman, L. M., Smith, M. J., Lefebvre, J. N., & Jones, A. R. (2019). Dynamics of an Airfoil Moving Through the Wake of a Circular Cylinder. *Proceedings from the AIAA SciTech 2019 Forum*.

Jones, A. R., Cetiner, O., & Smith, M. J. (2022). Physics and Modelling of Large Flow Disturbances: Discrete Gust Encounters for Modern Air Vehicles. *Annual Review of Fluid Mechanics*, *54*, 469–493.

Jones, M., Odgers, J., Kanso, E., & Luhar, M. (2021). On the Development of a Harmonic Pitching and Heaving System with Real-Time Control and Optimization Capabilities. 74th Annual Meeting of the APS Division of Fluid Mechanics, November 21-21, 2021. *Bulletin of the American Physical Society*, *66*(17).

Katzmayr, R. (1922). *Effect of Periodic Changes of Angle Attack on Behaviour of Airfoils* (Technical Manual No. 147) [Reproduced from "Zeitschrift fur Flugtechnik und Motorluftschiffahrt"]. National Advisory Committee for Aeronautics.

Kinsey, T., & Dumas, G. (2008). Parametric Study of an Oscillating Airfoil in a Power-Extraction Regime. *AIAA Journal*, *46*(6), 1318–1330.

Küssner, H. G. (1936). Zusammenfassender bericht über den instationären auftrieb von flügeln. *Luftfahrtforschung*, *13*(12), 410–424.

Lambert, T. J., Vukasinovic, B., & Glezer, A. (2016). A six degrees of freedom dynamic wire-driven traverse. *Aerospace*, *3*(11), 1–16.

Le Fouest, S., Deparday, J., & Mulleners, K. (2021). The dynamics and timescales of static stall. *Journal of Fluids and Structures*, *104*(103304), 1–11.

Lefebvre, J. N., & Jones, A. R. (2019). Experimental Investigation of Airfoil Performance in the Wake of a Circular Cylinder. *AIAA Journal*, *57*(7), 2808–2818.

Liao, J. C., Beal, D. N., Lauder, G. V., & Triantafyllou, M. S. (2003). The karman gait: Novel body kinematics of rainbow trout swimming in a vortex street. *Journal of Experimental Biology*, *206*, 1059–1073.

Liu, G., Dong, H., & Li, C. (2016). Vortex dynamics and new lift enhancement mechanism of wing–body interaction in insect forward flight. *Journal of Fluid Mechanics*, *795*, 634–651.

Mackowski, A. W., & Williamson, C. H. (2011). Developing a cyber-physical fluid mechanics facility for fluid-structure interaction studies. *Journal of Fluids and Structures*, *27*, 748–757.

Manohar, K., Brunton, B. W., Kutz, N. J., & Brunton, S. L. (2018). Data-driven sparse sensor placement for reconstruction. *IEEE Control Systems Magazine*, *38*(3), 63–86.

McCroskey, W. J. (1982). Unsteady Airfoils. *Annual Review of Fluid Mechanics*, *14*, 285–311.

McKinney, W., & DeLaurier, J. (1981). The Wingmill: An Oscillating-Wing Windmill. *Journal of Energy*, *5*(2), 109–115.

Melani, P. F., Balduzzi, F., Ferrara, G., & Bianchini, A. (2019). An Annotated Database of Low Reynolds Number Aerodynamic Coefficients for the NACA 0018 Airfoil. *AIP Conference Proceedings from the 74th ATI National Congress*, *2191*(020110), 020110-1:020110–10.

Oskouei, B. G., & Kanso, E. (2013). Stability of passive locomotion in vortex wakes. *Physics of Fluids*, *25*(021901), 021901–1:12.

Read, D. A., Hover, F. S., & Triantafyllou, M. S. (2003). Forces on oscillating foils for propulsion and maneuvering. *Journal of Fluids and Structures*, *17*, 163–183.

Schafer, R. W. (2011). What *Is* a Savitzky-Golay Filter? [Date of Publication: 15 June, 2011.]. *IEEE Signal Processing Magazine*, 111–117.

Schnipper, T., Andersen, A., & Bohr, T. (2009). Vortex wakes of a flapping foil. *Journal of Fluid Mechanics*, *633*, 411–423.

Shamai, M. (2021). *Complexity reduction of fluid-structure systems at low forcing frequencies* (Doctoral dissertation). California Institute of Technology.

Shamai, M., Dawson, S. T. M., Mezić, I., & McKeon, B. J. (2020). From unsteady to quasi-steady dynamics in the streamwise-oscillating cylinder wake [*Phys. Rev. Fluids*, to appear.]. *arXiv*, *2007.05635v1[physics.flu-dyn]*.

Sivaram, V., Dabiri, J. O., & Hart, D. M. (2018). The Need for Continued Innovation in Solar, Wind and Energy Storage. *Joule*, *2*, 1639–1642.

Streitlien, K., Triantafyllou, G. S., & Triantafyllou, M. S. (1996). Efficient Foil Propulsion Through Vortex Control. *AIAA Journal*, *34*(11), 2315–2319.

Strykowski, P. J., & Sreenivasan, K. R. (1990). On the formation and suppression of vortex ‘shedding’ at low Reynolds numbers. *Journal of Fluid Mechanics*, *218*, 71–107.

Su, Y., & Breuer, K. (2019). Resonant response and optimal energy harvesting of an elastically mounted pitching and heaving hydrofoil. *Physical Review Fluids*, *4*(064701), 064701–1:18.

SureshBabu, A., Medina, A., Rockwood, M., Bryant, M., & Gopalarathnam, A. (2021). Theoretical and experimental investigation of an unsteady airfoil in the presence of external flow disturbances. *Journal of Fluid Mechanics*, *921*(A21), A21–1:38.

Triantafyllou, G. S., Triantafyllou, G. S., & Grosenbaugh, M. A. (1993). Optimal Thrust Development in Oscillating Foils with Application to Fish Propulsion. *Journal of Fluids and Structures*, *7*, 205–224.

Trumper, D., & Dubowsky, S. (2005). 2.003 modelling dynamics and control i [Spring 2005. Massachusetts Institute of Technology: MIT OpenCourseWare, https://ocw.mit.edu. License: Creative Commons BY-NC-SA. Accessed: 2022-02-21].

van Oudheusden, B. W., Scarano, F., van Hinsberg, N. P., & Watt, D. W. (2005). Phase-resolved characterization of vortex shedding in the near wake of a square-section cylinder at incidence. *Experiments in Fluids*, *39*, 86–98.

von Kármán, T., & Sears, W. R. (1938). Airfoil Theory for Non-Uniform Motion. *Journal of the Aeronautical Sciences*, *5*(10), 379–390.

Wang, J., Geng, L., Ding, L., Zhu, H., & Yurchenko, D. (2020). The state-of-the-art review on energy harvesting from flow-induced vibrations. *Applied Energy*, *267*(114902), 114902–1:17.

Wei, Z., & Zheng, Z. (2017). Energy-Harvesting Mechanism of a Heaving Airfoil in a Vortical Wake. *AIAA Journal*, *55*(12), 4061–4073.

Williamson, C. H. K. (1996). Vortex dynamics in the cylinder wake. *Annual Review of Fluid Mechanics*, *28*, 477–539.

Williamson, C., & Govardhan, R. (2004). Vortex-Induced Vibrations. *Annual Review of Fluid Mechanics*, *36*, 413–455.

Wu, T. Y. (2011). Fish Swimming and Bird/Insect Flight. *Annual Review of Fluid Mechanics*, *43*, 25–58.

Wu, T. Y., & Chwang, A. T. (1975). Extraction of flow energy by fish and birds in a wavy stream. In T. Y.-T. Wu, C. J. Brokaw, & C. Brennen (Eds.), *Swimming and flying in nature: Volume 2* (pp. 687–702). Springer US.

Xiao, Q., & Zhu, Q. (2014). A review on flow energy harvesters based on flapping foils. *Journal of Fluids and Structures*, *46*, 174–191.

Zhu, Q. (2011). Optimal frequency for flow energy harvesting of a flapping foil. *Journal of Fluid Mechanics*, *675*, 495–517.