

MURI: Closed-loop Control of Vortex Formation in Separated Flows With Application to Micro Air Vehicles

Final Report

AFOSR Grant FA9550-05-1-0369

Tim Colonius (PI) and Morteza Gharib
California Institute of Technology

Clarence W. Rowley
Princeton University

Gilead Tadmor
Northeastern University

David R. Williams
Illinois Institute of Technology

Pasadena, CA
October 25, 2010

Contents

1	Executive summary	
2	Introduction	
3	Control of vortex shedding in wind tunnel experiments	1
4	Simulation, modeling, and control tools	45
5	Control of vortex shedding in numerical simulations	104
6	Three-dimensional numerical simulations	139
7	Oil Tunnel studies at low Reynolds number	169
8	Publications resulting from this award	213
9	Acknowledgements	219
	References	220

1 Executive summary

This research is concerned with developing theory, algorithms, and applications of model-based, closed-loop flow control in order to enable robust and agile flight of micro air vehicles. The efforts are centered around the application to integrated closed-loop flow and flight control for stabilization and regulation of separated flows occurring on unmanned and micro air vehicles (UAV/MAV). Increased lift associated with controlled flows will lead to dramatic improvements in maneuverability, gust resistance, and a wider flight envelope.

In demonstrations in an unsteady wind, a computer-controlled captive trajectory system has been implemented to examine pitch, roll, and plunge maneuvers of the wing in response to the unsteady freestream flow. The ability to dynamically cancel lift fluctuations associated with gusting freestream flow was investigated in detail. Transient force measurements in response to pulse-like disturbances from leading-edge actuators were used to obtain linear models of the separated flow, and it was shown that the transient lift response can be used as a filter kernel to predict the response of the wing to more complex actuator input signals. A series of controllers of increasing complexity were designed to suppress lift fluctuations in gusting conditions. The most robust controllers were able to suppress lift fluctuations associated with a broad-band spectra of freestream velocity, up to a limit of about a factor of about 10 times slower than the intrinsic vortex shedding frequency of the natural separated flow. The bandwidth limitations were found to be associated with, for this control architecture, the timescale associated with the time for the actuator-produced leading-edge vortex to be generated and swept downstream along the wing. Circumventing this limitation in future will require controllers to interact directly, and utilize information from, the vortex formation process, a topic which was investigated directly using numerical simulations and model-based control, as discussed below.

A suite of computational fluid dynamics (CFD), reduced-order modeling (ROM), and control design tools were also developed in order to study advanced concepts for model-based control of aerodynamic forces. The algorithms provide for a fast computational environment where advanced analytical techniques for model reduction and control design are directly integrated to provide a closed-loop flow control toolkit, with general applicability to feedback control in external flows. The complex physics associated with small-scale, highly three-dimensional wing shapes and separated, unsteady flows have been mapped, and studied in companion experiments in a novel oil-tunnel facility that provides state-of-the-art, real-time and holographic flow reconstructions to complement and validate the computational models.

Major breakthroughs are reported in the mathematics and algorithms for reducing complex CFD models to low degree-of-freedom systems suitable for application in practical, real-time controllers. In the fluid dynamics community, the predominant technique for reduced-order modeling is Proper Orthogonal Decomposition (POD), in which one gathers data from simulations or experiments, and extracts an orthogonal set of modes that are optimal in the sense that they capture the most energy in the given data set. In practice, however, POD models tend to be fragile. This fragility is understood by the control theory community, and alternative techniques such as balanced truncation and optimal Hankel norm reduction offer more robust performance and theoretical bounds on errors. However, these techniques have previously been far too computationally expensive to perform on a full CFD model. This reports on development of approximate balanced truncation, called balanced POD, that is computationally tractable for large systems, and produces models almost identical to those from exact balanced truncation. The technique is extended to unstable and periodic systems in order to use them to model vortex shedding on two and three-dimensional wings.

A hierarchy of two and three-dimensional model problems are posed and solved in order to investigate reduced-order modeling and theoretical control of vortex shedding at low Reynolds number. In particular, BPOD models are used to design observer-based control that is able to completely suppress vortex shedding

in nonlinear, two-dimensional simulations. In addition, we examine the possibility of using actuation to produce high-lift limit cycles, where the synchronization of vortex shedding leads to pushing the vortices closer to the suction surface, resulting in higher lift. These limit cycles are not robust to disturbances, however, and feedback control, based on real-time tracking of the phase of the lift fluctuations, is successfully implemented in order to stabilize these high-lift limit cycles. In addition, optimal control theory is used to find actuator waveforms that maximize the period-averaged lift, and these waveforms are then used together with the phase-lock loop feedback control to achieve optimal-lift limit-cycle behavior. Three-dimensional simulations were performed for natural and actuated flows over low-aspect ratio flat plates with rectangular, semi-circular, and delta-wing planforms at a variety of angles of attack. For very low aspect ratio plates, the tip vortices have a stabilizing influence on the vortex shedding. Open-loop actuation is able to attenuate and, in some cases, completely suppress vortex shedding. Finally, closed-loop control is applied in the three-dimensional simulations using an extremum-seeking approach, which is able to find optimal frequencies of actuation.

2 Introduction

In this section, we review the recent literature on natural and actuated flows over two- and three-dimensional flat plates and airfoils, and effort aimed at open- and closed-loop control of separation for lift enhancement, drag reduction, and other objectives. We include in the discussion some of our own results that are amplified in later sections in this report.

2.1 Separation and vortex shedding

We first consider flow field and forces associated with a flat plate or airfoil at an angle of attack to an otherwise uniform, steady, stream of speed U , under the action of continuous, but unsteady, forcing from an actuator. The chord length is c , and, unless otherwise mentioned, the planform is rectangular with breadth $2b$ and an aspect ratio $AR = 2b/c$. The relevant Reynolds number is $Re = \frac{Uc}{\nu}$. We first discuss the natural (unforced) flow, and, in the next section, actuated (forced) flows.

As the angle of attack, α , is increased, different regions of separated flow appear as a function of the airfoil shape and Reynolds number. For laminar airfoils, onset of separation would typically occur near the trailing edge, and progress upstream as α is increased. For flat plates and other thin airfoils, the sharp leading edge promotes separation and a separation bubble (with typically turbulent reattachment on the suction surface) may precede the fully-stalled condition. We are concerned here with completely separated (stalled) flow over the entire suction surface, when the airfoil behaves as a bluff body with vortex shedding, oscillatory forces, and the formation of a Kármán vortex street in the wake. The frequency of vortex shedding, at least for the high AR case, follows a Strouhal scaling (Fage & Johnsen 1927) with $St = f c / U = 0.15$ to 0.2 , where $c \sin \alpha$ is the projected area in the direction of the stream, and the Strouhal number is nearly constant at high Re .

As is discussed more fully in section 5.1.3, for a strictly 2D flat plate, the onset of vortex shedding occurs, much as it does for a bluff body, as a Hopf bifurcation at critical value of Re or α . Ahuja & Rowley (2010a) found $Re_{crit} = 23^*$ at $Re = 100$ and Chen et al. (2010) found an $Re_{crit} = 80$ at a fixed $\alpha = 30^\circ$. For $Re > Re_{crit}$, the same variation in shedding frequency with Re is observed as that on a circular cylinder (Roshko 1955). Based on the similarity with the flow over a circular cylinder (e.g. Barkley & Henderson 1996), it could be expected that at a higher Re the 2D vortex shedding would undergo a further bifurcation to 3D flow (even for an infinite plate), but these instabilities are only recently beginning to be studied (Rodriguez & Theofilis 2011). For low aspect ratio, 3D flat plates at low Re , vortex shedding still occurs, but data is very limited. Taira & Colonius (2009a) investigated low AR flat plates for $Re = 300$ and 500 over a range of α and for rectangular, elliptical, semi-circular, and delta-shaped planforms. The onset of vortex shedding is delayed to higher Re and α as the aspect ratio is decreased, due to a stabilizing influence of the tip vortices. As AR is increased beyond about 3, the initial bifurcation to vortex shedding coincides with the value for strictly 2D flow.

As in bluff bodies, vortex shedding persists at high α when Re is increased (Williams-Stuber & Garib 1990). Nominally 2D airfoils, and in particular, the symmetric NACA series, have been studied the most. Huang et al. (2001) measured the frequency of vortex shedding in the wake of a NACA 0012 over a wide range of post-stall values of α up to $Re = O(10^4)$. At sufficiently large Re , the thin shear layer bounding the separation displays a Kelvin-Helmholtz instability (with St about an order of magnitude higher than vortex shedding), and ultimately becomes turbulent (Brundet & Mueller 1988). The dominant shear layer instability frequency shows a power law dependence on Reynolds number, $f \sim Re^*$, similar to circular cylinders Yarusso et al. (2009). Like the circular cylinder, separated airfoil flows typically show a broad band of frequencies centered around the nominal shedding value in the wake, potentially due to the interaction of

the differing timescales and the effects of increasingly complicated and turbulent wakes (Yarusso et al. 2009). The situation is more complicated at values of α near the onset of fully separated flow, where further increases in Re can lead to reattachment of the separated region prior to the trailing edge. Finally, we note that for flexible, membrane airfoils, Rojstrostrik et al. (2009) found that the natural vortex shedding at post-stall angles was coupled to the membrane oscillations.

2.2 Steady-state actuated flows

Following the pioneering work of Prandtl (1904), classic separation control techniques (Lachmann 1961) such as steady blowing and suction attempts to energize a boundary layer thereby delaying or preventing separation. Lately, unsteady blowing, zero-net-mass actuators (synthetic jets), piezoelectric flaps, plasma actuators, and other unsteady actuators have been shown to achieve similar performance but with far lower mass, momentum, and/or energy fluxes than steady blowing or suction (Greenblatt & Wygnanski 2000a, Seifert et al. 2004). For airfoils and flaps, lift and lift-to-drag ratio can be substantially increased, but the underlying mechanisms associated with lift enhancement or drag reduction are still debated. The dominant idea, discussed in detail by Greenblatt & Wygnanski (2000a), is that excitation of vortical structures in the separated shear layer leads to enhanced entrainment and the attendant suction of the shear layer to the nearby surface, eliminating or reducing the extent of the separated region, and leading to a time-averaged flow field closer to the ideal, potential flow. Entrainment likely plays a role regardless of whether vortical structures are generated as part of a Kelvin-Helmholtz instability in the shear layer, or via a global instability of the wake or separation bubble, but, as discussed in the previous section, the frequencies at which these instabilities occur are distinct. Especially for low frequencies and high angles of attack, though, lift enhancement has also been explained in terms of vortex lift (e.g. Wu et al. (1998)), through concentration of vortical structures closer to the surface of the plate.

For a specified periodic actuator, airfoil geometry, and angle of attack, any nondimensional performance metric χ is at most be a function of the Reynolds number, the actuator waveform and nondimensional parameters expressing its frequency and amplitude. For the frequency, the most common choice is a reduced frequency $f^* = \frac{f c}{U}$. For amplitude, the momentum coefficient is often used, $C_p = \frac{\rho U^2 A_c}{2 q c}$, where U_c and A_c characterize the mean and/or fluctuating (rms) velocity and area, respectively, of injection, and A is the planform area. For both these quantities, the chord length is sometimes replaced by the length of the (natural, unforced) separation bubble when the flow is not fully separated. Mean and fluctuating values of C_p as low as about 10^{-4} can be effective (Greenblatt & Wygnanski 2000a).

Greenblatt & Wygnanski (2000a) showed that C_p collapses data obtained with several different actuators, but, in general, it is difficult to compare data from different actuators and waveforms, or with characterizing the actuator performance in terms of a velocity (or mass flux) that can depend on the plumbing for the actuator and whether the performance is measured with or without flow. Most studies observe a lower threshold and upper saturation limit of actuation that define the range of proportional control that can be achieved. In the controls community this is known as the *static map* and it is an important step in the design of a control system. Reynolds number effects (e.g. Seifert et al. 2004) have also been studied. Here the most pressing issue is whether some reports of lift enhancement or drag reduction could be explained by the mechanism of tripping the boundary layer to delay or prevent separation. It is clear, however, that there remains an effect of forcing at Re lower than those for which tripping can lead to a turbulent boundary layer, as well as at high Re when the boundary layer was turbulent even in the absence of tripping (Seifert et al.

*To be fully general, we should also account for the Mach number, and the possibility that the density of injected fluid is not equal to the ambient density, resulting two additional parameters.

2004).

The effect of actuation frequency on performance has also been widely studied, and gains (lift enhancement, drag reduction, and other goals) have been realized across a wide range of frequencies. Here we employ the term “low frequency” when the actuation frequency is below and up to the vortex shedding frequency (discussed above), and “high frequency” to refer to everything significantly above it, and in particular to excitation of shear-layer instabilities. As discussed by Raju et al. (2008) (hereafter referred to as RMC), a third distinct timescale exists when the (mean) flow forms a closed recirculation bubble on the airfoil surface. While these three distinct timescales—shear layer, wake (vortex shedding), and (in some cases) separation bubble—can be identified by examining velocity spectra measured at different locations (RMC), the lack of such data makes it difficult to make definitive statements about their values in past experiments. For example, for a turbulent separation, Greenblatt & Wygnanski (2000a) report an optimal value of $f^* = 1$ for a deflected flap (with the flap length as the length scale), where “optimal” refers to that frequency at which a minimal actuation amplitude was required for reattachment (in the mean), and they associate this timescale with shear-layer instabilities. Using the simple scaling for the vortex shedding timescale discussed in the last section, $St = 0.15 = \frac{f c}{U}$, where St is the flap deflection angle, we conclude that a typical vortex shedding frequency would have $f^* > 1$ when $St > 8^*$; one can infer that wake and/or separation bubble instabilities may equally have played a role. Seifert et al. (1996a) used oscillatory blowing at the leading edge of a NACA0015 airfoil at $Re = 10^5$, and found lift enhancement and drag reduction over range $0 < f^* < 2$, with a broad maximum around $f^* = 0.75$, which, for the range of $16^\circ < \alpha < 22^\circ$ considered, gives $0.2 < St < 0.24$.

For lower Reynolds numbers, Hsiao et al. (1994) acoustically forced the flow near the leading edge of a fully separated NACA 633-018 airfoil and observed a strong enhancement of vortex shedding, and the mean lift, but only when excitation was close to the natural vortex shedding frequency. The computational study of RMC showed that for a NACA 4418 at $Re = 40,000$ and $\alpha = 18^\circ$, actuation with $f^* > 6$ was increasingly ineffective; in their case, $f^* = 12$ was identified as the maximally amplified frequency in the separating shear layer. These results are corroborated by the experimental findings of Cierpka et al. (2008) who subjected a NACA 0015 ($\alpha = 20^\circ$) and an inclined flat plate ($\alpha = 15^\circ$) subjected to an electromagnetic actuation near the leading edge. For the plate, with $0.5 < f^* < 3$, the flow was reattached (the lift enhancement was best at $f^* = 0.7$ ($St = 0.24$), whereas $f^* > 6$ had little effect on the separation. For their inclined flat plate, they employed a wavelet algorithm to detect large-scale vortices in time-resolved PIV data, which revealed an interesting coalescence of smaller vortices produced near the actuator into one large coherent vortex advecting down the plate per cycle of actuation, in the case of $f^* = 1$ ($St = 0.23$), and to two coherent vortices per cycle of actuation when $f^* = 0.5$, both of which imply vortex shedding at $St = 0.23$. A similar configuration (with a dielectric barrier discharge actuator) was studied by Greenblatt et al. (2008) (flat plate) and Bernard et al. (2008) (NACA0015). For the flat plate at $\alpha = 20^\circ$, $0.3 < f^* < 0.6$ provided the best lift enhancement, whereas $f^* > 3$ was ineffective, and smoke visualization at $f^* = 0.4$ showed a strong vortex advecting downstream along the chord, $f^* = 1.5$ was optimal for a NACA 0015 at $\alpha = 16^\circ$.

Recent studies have also documented the effect of the waveform on performance. It appears that periodic but pulsed actuation or modulated high frequency sinusoidal oscillations can produce performance equal or greater to sinusoidal actuation at the same frequency (Amayri & Glezer 2002a, Woo et al. 2008, Cierpka et al. 2008, Greenblatt et al. 2008, Joe & Colonius 2010). Indeed, it appears that pulses with as low a duty cycle as 5% can be effective. (Greenblatt et al. 2008). In our work presented below (section 5.2.4), we have employed an aperiodic-based approach in a low Reynolds number DNS to find the actuator signal (a body force in this case) that gave the highest lift. Indeed, the optimal signal was a nearly periodic pulsed forcing at a frequency close to that of the vortex shedding.

All of the above studies have focused on 2D geometries, but our research, described in section 5 has shown that leading-edge actuation offer similar performance benefits, with similar values of momentum coefficient and frequency of actuation, on low aspect-ratio three-dimensional wings.

However, we note that relatively few studies have reported unsteady flow metrics associated with fluctuating forces as function of the actuation frequency. As pointed out by Amayri & Glezer (2002a), enhancement of vortex shedding by forcing near its natural frequency also increases the fluctuating lift and drag compared to the baseline. Aside from being potentially detrimental in application, such unsteady effects are essential to understand if closed-loop control approaches are to be successful. This is discussed in greater detail in the next sections.

2.3 Leading-edge vortex

When a flat plate at high angle of attack is impulsively started, the separating vortices sheet at the leading and trailing edges roll up into a coherent leading and trailing edge vortices (LEV/TEV). The TEV (start-up vortex) is rapidly shed into the wake, while the LEV continues to grow for about a chord-length of travel, where the lift and drag reach their maximal values (without regard to the additional added mass force during the acceleration). Once the LEV begins to shed, the lift decreases to a minimum that occurs after 4–5 chord-lengths of travel, which appears to be only weakly dependent on the Reynolds number or acceleration rate (Chen et al. 2010). This timescale is consistent with the so-called universal time-scale of vortex formation that is observed in a variety of flows (e.g. Dabiri 2009). Rotational forces on flapping wings, on the other hand, appear to prolong (stabilize) the LEV structure (Lentink & Dickinson 2009). For a translating wing, after the initial LEV is shed, alternating TEV and LEV of diminishing amplitude are shed until periodic or quasi-periodic vortex shedding is attained. The “vortex” lift force of the initial LEV has been measured to be as much as 80% above the steady-state (time-averaged) value (Dickinson & Gotz 1993a). The LEV, and its associated lift increment, are similar to the dynamic stall vortex that is produced and shed during rapid pitch up to high angles of attack (e.g. Carr 1988a).

2.4 Transient response to actuation

As discussed above, understanding the transient response to initiation, termination, or other changes to actuator parameters is essential to the development of closed-loop flow and/or flight control strategies. Relatively few experimental and computational studies have addressed this issue. Amayri & Glezer (Amayri & Glezer 2002a, 2006) examined the response of a symmetric NACA airfoil at $Re = 3 \times 10^5$ and $\alpha = 17.5^\circ$ to a burst of high-frequency synthetic jet actuation. The response of the actuator consists of several ejection/suction phases of the actuator ($f^* = 10$) with rapidly diminishing amplitude. The phase-average circulation flux in the wake was considerable despite the low actuation amplitude $C_p \approx 10^{-3}$, and consisted of an initial negative contribution to the circulation, followed by a large positive vortex being shed into the wake, and additional oscillations of diminishing amplitude. The total duration of the transient response was long compared to the actuation, lasting several hundred actuation cycles, or about 10 convective time units ($t^* = \frac{U}{c} t = 10$). The response to step changes in actuation has also been studied (Amayri & Glezer 2002a, Dabiri & Wygnanski 2004a,b) and shows similarly long transient times when toggling between fully separated and controlled flow states. Similar to the burst of actuation, the forced reattachment first results in a negative vortex followed by the shedding of large, positive vortices. Dabiri & Wygnanski (2004a) studied flow over a deflected flap, and showed that the response scales well with convective time; at a minimum, when the actuation amplitude is sufficiently large, the controlled state is reached after about $t^* = 20$. At lower Reynolds number, the response is more oscillatory, and seems to consist of consecutive shedding of

large-scale vortices from the flap. Toggling off control results in transients of similar duration, but the early response includes the formation of a large (positive) vortex which Darabi & Wygnanski (2004b) liken to the DSV.

Similar experiments for $Re = 10^4$ to 10^5 have been conducted in the semicircular airfoil in the ITT wind tunnel with leading-edge pulsed-jet or synthetic jet actuators. These are discussed in detail in section 3, but we note here that they reveal strikingly similar transients to the 2D airfoils and flaps previously studied.

2.5 Biological benchmarks for MAV performance and control

Birds, insects, and bats offer enticing benchmarks for MAV and UAV performance including low cruising speed, agility, and propulsive efficiency, and sensor-based control. These are important benchmarks for MAVs and small UAVs where loiter/observance capability and navigation in urban confined spaces are critical to mission capabilities. Some unique properties of bio-flyers stem from the role of a stable LEV that allows high lift at extremely high angles of attack (AOA). The stability is partly due to the low Reynolds number (Re), but more strongly affected by three-dimensionality due to the low aspect ratio (wing span to chord length, AR) typical of biological wings. Studies indicate that stability of the LEV is related to the spanwise transport of vorticity through an axial flow toward the tip vortex (TV). Even during purely translational motion (as opposed to flapping), strong three-dimensionality leads to a stabilizing interplay between flow separation LEV and TV at low Re (102 to 103) (Birch et al. 2004a). Moreover, lift maximization (Wang & Kretz 2000) and rapid maneuvering in bio-flyers are achieved through careful synchronization of vortex shedding as regulated by asymmetry and timing of wing-strokes. High-speed video of insect maneuvers during predator-flight show drastic changes in flight direction that occur on the same timescale as the wing stroke (Fry et al. 2003).

Unfortunately, similar mechanisms cannot presently be applied to fixed wing aircraft, except through control surfaces that cannot respond (due to mechanical bandwidth constraints, stress limitations, and aircraft inertia) to the fast timescales required for agility. Fixed wing designs typically use high AR to obtain better efficiency (induced drag scales with the inverse of AR), higher minimum flight speed, and higher payload capacity. However, high AR severely limits the maximum AOA, increasing minimum cruising speed and decreasing agility. If flow control can be used to broaden the envelope (in both AR and Re) of high-lift aerodynamics, then decreased stall speed and enhanced agility can be obtained without sacrificing efficiency. For example, at low to moderate Re (102 to 105) the maximum lift coefficient for a very low AR wing (even a flat plate) can approach 1.5 to 2 (Birch et al. 2004a, Torres & Mueller 2004a). If this lift coefficient were available to an $AR=2$ wing, stall speed could be reduced by 30 to 50%.

The objective of closed-loop actuation near the leading edge and tip is to enable benefits associated with low AR and low Re aerodynamics to be achieved with fixed wings at higher AR and Re . In this research, we pursue the hypothesis that closed-loop control can stabilize the LEV/TV system (preventing or delaying shedding/stall at high AOA), and, when desirable for maneuver, synchronize vortex shedding to produce controlled roll, yaw, and pitching moments. In addition to higher performance, flow control actuators may also render conventional control surfaces redundant. Note that in what follows we use the terms stabilization and synchronization to denote the differing goals of control during both steady lift generation (cruise) and unsteady force generation (maneuver), rather than in any strict mathematical sense.

2.6 Closing the loop

While open-loop actuation is capable of enhancing post-stall lift under steady flight conditions, there can distinguish two goals of controlling actuation with sensor-based feedback. The first, and more ambitious

goal is to alter the dynamics of the flow in ways inaccessible to open-loop actuation. For example, is it possible to eliminate vortex shedding? The second, more modest goal is to improve flight performance in unsteady flight, especially in regimes where conventional control surfaces may not be effective. Obviously the distinction between these becomes blurred as the timescales of imposed unsteadiness approach the intrinsic fluid dynamic timescales associated with either vortex shedding or shear layer instabilities.

Most previous work on flow control has been open-loop, and a majority of previous studies of the potential benefits of feedback have been theoretical (e.g. Bewley et al. 2000, Kim & Bewley 2007). Aside from the work developed as part of this MURI, closed-loop flow control has been demonstrated in the laboratory for controlling vortex shedding on bluff bodies for drag reduction (e.g. Passaro et al. 2008, Siegel et al. 2006), combustion instabilities (e.g. Dowling & Morgans 2005), and cavity oscillations (e.g. Rowley & Williams 2006).

In the context of airfoils and MAV, for relatively slow changes in operating conditions, one may schedule operating parameters in a way that is not fundamentally different from open-loop actuation, though there remain challenges such as the hysteresis associated with separating and attached flow (e.g. Darabi & Wygnanski 2004a). For example, Magill et al. (2003) used pressure fed back to a dynamic stall model to detect imminent separation during cyclic pitching of an airfoil and apply pulsed vortex generator jets only over a portion of the pitch cycle and achieve similar lift increases as continual operation. Pivier et al. (2007) controlled incipient separation with a proportional feedback signal based on the first global mode obtained using POD. Benard, Bonnet, Moreau, Griffin & Catafesta (2010) exploited hysteresis in the separation/reattachment process by detecting the signature of incipient separation with a pressure sensor, and then lowering the actuator voltage when the flow is already attached.

A next level of complication is to optimize performance as a function of input parameters. Benard, Moreau, Griffin & Catafesta (2010) used the measured lift to adjust the voltage to a dielectric barrier discharge actuator on a NACA 0015 airfoil to autonomously reattach the flow at different speeds, and for step changes in the operating conditions. Becker et al. (2007) implemented multiple-input multiple-output extremum seeking control for spanwise distributed actuation and (pressure) sensing for pulsed-jet actuators on a flap, and were able to achieve higher lift than open-loop actuation, including lower angles of attack where the flow was not fully separated and where open-loop control showed little effect. Taira et al. (2010) also used extremum seeking to optimize actuation frequency in numerical simulations of 3D airfoils at low Re . Muto et al. (2008) used a neural network adaptive controller to control the pitch-phase motion of an airfoil in a wind tunnel.

For sufficiently fast changes in operating conditions there are bandwidth limitations for a particular architecture, i.e. the controller and the dynamic response of actuators, sensors, and the inherent response of flow fluctuations to actuation and to changes in operating conditions. As is documented in section 3 the lift response to actuation in the separated regime is governed by the convective time scale and reaches a peak transient lift in about 3 convective units. For MAV, this implies a full-scale frequency on the order of 1 to 100 Hz, and this is likely to be considerably lower than bandwidth limitations associated with the actuation and sensing. A major aspect of the present research has been to document this limitation, in the context of unsteady wind tunnel experiments described in section 3, and to show how, at least in theory and computation, it may be overcome in order to achieve stabilization and synchronization of vortex shedding on its own intrinsic timescale. In actual application, as we document, even achieving control at timescales slower than vortex shedding involves major technical advances.

2.7 Reduced-order models

Most available methods for designing closed-loop controllers require knowledge of a mathematical model of the system to be controlled, incorporating the effects of inputs and outputs. For problems in fluid mechanics, the governing equations are known (the Navier-Stokes equations), and in principle, these may be used for control design. For instance, model-based optimal controllers have been applied to simulations of channel flow (Bewley et al. 2001), Håberg et al. (2003) and jet noise (Wei & Freund 2006), but these strategies are far too computationally expensive to run in real time. In order to produce model-based controllers that may be implemented in practice, it is necessary to use some form of model reduction, in which the high-fidelity, high-dimensional model is replaced with a simpler low-dimensional approximation.

Most previous work in this area has focused on Proper Orthogonal Decomposition (POD) (Holmes et al. 1996), a method which extracts the energetically dominant features in a flow. However, in many fluids problems, low-energy features have been shown to have an important effect on the dynamics, and as a result, models based on POD often perform poorly (Ik & Rowley 2008). For instance, even models of stable phenomena can be unstable, and can be quite fragile when parameters such as the Reynolds number are changed.

One of the goals of this project is to improve the state of the art of low-order modeling, and develop systematic techniques for designing models suitable for control design. Our efforts here focus on an approximate version of balanced truncation Moore (1981) called Balanced Proper Orthogonal Decomposition (Balanced POD) (Rowley 2005). Balanced POD systematically incorporates the effects of sensors and actuators, and typically produces models that are much more accurate and robust than corresponding POD models. In previous work, the balanced POD procedure has been developed for stable linear systems, and here we investigate its application to unstable equilibrium points, as well as unstable periodic orbits (such as vortex shedding).

One legitimate criticism of the Balanced POD method is that it cannot be applied to experimental data, as it requires information from *adjoint simulations* of the flow. These adjoint simulations in essence describe the sensitivity of the flow to perturbations in different regions, and this information is critical to the effectiveness of the balanced POD method. Here, we also pursue methods of *adjoint-free* balanced model reduction, so that systematic reduced-order models may be produced directly from experimental data, without the need for adjoint simulations.

2.8 MURI Research Objectives

The overall goal of the MURI research was to develop integrated closed-loop flow and flight control for MAV applications. Specific objectives were to

- Develop closed-loop flow control to extend the parameter space for which steady lift can be maintained at high angle of attack.
- Use control to synchronize vortex shedding and improve maneuverability and gust response, which will ultimately eliminate conventional control surfaces made redundant by flow control actuators.
- Use experiments with numerical simulations to obtain insight into the flow physics of separated flows of low aspect-ratio wings.
- Develop a reliable, systematic approach to reduced-order modeling for feedback flow control, using data from simulations and experiments.

- Create a common integrated computational framework for direct numerical simulations (DNS), linearized and adjoint Navier-Stokes, global stability analysis, and reduced-order models.
- Design and evaluate closed-loop control laws based on the reduced-order models developed.
- Explore closed-loop control strategies such as phase-lock loop and extremum seeking, for which models are not required.

2.9 Summary of Accomplishments

This report provides detailed results from experiments, simulations, and modeling efforts aimed at the objectives described within the previous section. We provide here a brief summary of the main accomplishments of the overall effort.

- (Section 3) Closed-loop flow and flight control were demonstrated in the laboratory in order to enhance lift and to maintain steady lift in an unsteady, gusting flow. Reduced-order models of the lift response to actuation were obtained using novel system identification techniques and were used to design several generations of controllers that successively demonstrated the ability to suppress lift fluctuations with increasing bandwidth. The transient response to actuation and its scaling with pulse duration, amplitude, and freestream speed was studied in detail to provide data for the models, and to understand the physics associated with the generation of vorticity at the leading edge during actuation. The ultimate bandwidth of lift fluctuation suppression achieved in the laboratory was found to be associated with the inherent dynamics associated with the convection of an actuator-generated vortex along the suction surface.
- (Section 4) A new formulation of the immersed boundary (IB) method was developed for computation of flows around low-aspect ratio flat plates and wings at low to moderate Reynolds number. Algorithmic advances achieved an order of magnitude improvement in efficiency compared to the traditional IB method. A common framework was developed to implement linearized and adjoint linearizations of the governing equations that are discretely consistent with flow solvers. Simulations were used to investigate the flow physics, and to provide data for reduced-order modeling efforts.
- (Section 4) A systematic procedure for reduced-order modeling was developed, and extended to unstable systems, and systems with periodic orbits (such as vortex shedding). The resulting models incorporate the effects of actuators and sensors, and are much more effective for this problem than existing techniques such as traditional Proper Orthogonal Decomposition. A second approach of using mean-field models and statistical turbulence closures was explored to develop low and least-order Galerkin models for control design.
- (Section 5) Reduced-order models were obtained for the separating flow past a flat plate at high angles of attack, forced by actuators at the leading and trailing edge. Observer-based feedback controllers were designed, and were able to stabilize leading-edge vortices at $Re = 100$, in direct numerical simulations. In addition, gradient-based optimization was used to find lift-maximizing actuator inputs for higher angles of attack where vortex shedding is inevitable. Optimized waveforms were also implemented with a simple phase-lock-loop controller in order to yield optimal performance robustly (in the presence of disturbances) in a way suitable for application.
- (Section 6) The natural and actuated (open-loop) flows on low aspect ratio (3D) flat plates were simulated for a wide parameter regime including Reynolds number, angle of attack, actuation strength and

frequency, and planform shape. Simulations revealed the stabilizing influence of tip vortices on vortex shedding, an effect which can be enhanced with steady and unsteady actuation leading to major lift enhancement at low Reynolds number. A model-free extremum-seeking controller was also designed, and was able to rapidly find optimal frequency of sinusoidal forcing applied at the trailing edge in direct numerical simulations.

6. (Section 7.A novel, recirculating oil tunnel was designed and constructed for flow control studies. Advanced laser-based diagnostics including real-time and three-dimensional particle-image velocimetry were used to examine the three-dimensional flows on low aspect ratio wings.

13

3 m/s and over a bandwidth from 0.1 Hz to 30 Hz. The turbulence level decreased as the flow speed increased.



Fig. 3.2.1: Drawing of unsteady flow wind tunnel test section showing PIV system, model and shutters used to create an unsteady freestream.

3.3. First generation wing – rectangular planform

The ability of active flow control actuators to modify the leading-edge and tip vortex (LEV/TV) system was explored in the first phase of wind tunnel experiments using a rectangular planform wing. Three different orientations of the actuators were explored – upstream, downstream and cross-flow actuation (45° to the mean chord). Control of the leading edge vortex was expected to enhance maneuverability by stabilizing or synchronizing vortex shedding during pitch, yaw, and roll motions, and in response to gusts. The longer term goal was to extend the range of aspect ratios and Reynolds numbers for which steady lift could be maintained at very high angles of attack (AOA), while realizing benefits associated with higher aspect ratio aerodynamics during cruise.

The concept of lift augmentation by cross-flow steady blowing was first studied by Dixon (1969), who showed that the leading edge vortex could be prevented from shedding when a jet of air was blown laterally over the suction surface of the airfoil. He speculated that spanwise blowing created leading-edge sweep effects, similar to delta-wings, and devised a single jet positioned near the root of the wing and at the $c/4$ location on a rectangular flat-plate wing with aspect ratio 4.7. Steady blowing with very large momentum coefficients from $C_h = 0.29$ to 0.82 increased the maximum lift coefficients from $C_L = 1.3$ to 2.5, which was attributed to stabilization of a leading edge vortex.

We explored the possibility of achieving a similar stabilization of the leading edge vortex, but with orders of magnitude less blowing, by adding a spanwise component to the pulsed-blowing jets from actuators and by distributing them along the leading edge and tip regions of airfoils with rectangular planforms. The response of the airfoil flow to individually controlled actuators with open loop forcing was documented to provide the baseline information needed to develop a closed loop control model. A particularly important measurement was to determine the degree of influence the pulsed-blowing actuators had on the spanwise circulation distribution over the wing.

15

3 Control of vortex shedding in wind tunnel experiments

3.1. Introduction

Wind tunnel experiments aimed at controlling forces on finite span wings by modulating the strength of the leading edge vortex were performed at Illinois Institute of Technology in the Andrew Fejer Unsteady Flow Wind Tunnel. The experiments were coordinated with the numerical simulations to gain better understanding of the flow physics associated with control by LEV modulation. Although the Reynolds number of the wind tunnel experiments was larger than the numerical simulations, useful connections were made between the two approaches, which provided insight into the physics of separated flows and closed-loop control. A description of the experimental facility is provided in Section 3.2.

The first generation wing had a rectangular planform with a 2:1 aspect ratio. The influence of the direction (upstream, downstream and skewed) of actuation on the leading edge vortex and separated flow was explored. Key results from this investigation are described in Section 3.3. The second generation wing had a semi-circular planform wing with an aspect ratio of 2.54. The semi-circular leading edge enhanced the ability of the pulsed-blowing actuators to control the leading edge vortex and the associated lift increment. Details of the design and performance of the semi-circular wing with and without active flow control are provided in Section 3.4.

After establishing the ability of open-loop actuation to stabilize the leading edge vortex, the investigations expanded to dynamic control, i.e., the use of active flow control during changing flight conditions. First, the transient response of the lift to pulse-type actuation was examined, which led to scaling relations for the response of the separated flow region to actuation. Section 3.5 discusses the lift response to single and multiple-pulse actuation. Next the unsteady aerodynamic response of the wing to an oscillating freestream flow was investigated. Modern system identification techniques were useful in obtaining linear transfer functions for the transient response to actuation and the unsteady aerodynamics. Closed-loop control algorithms based on these linear models were used for randomized gust suppression demonstrations. These results are discussed in Section 3.6. Finally, to investigate the influence of actuator bandwidth on the performance of the semi-circular wing in dynamic conditions, a wing with piezo-electric (zero net mass flux) actuators was constructed. Pulse-response measurements for the piezo-electric wing are presented in Section 3.7.

3.2. Experimental set up

The leading edge and tip vortex interaction studies were conducted under steady flow and dynamic conditions in the Andrew Fejer Unsteady Flow Wind Tunnel shown in figure 3.2.1. The test section cross-section is 0.61m by 0.61m. Flow speeds up to 30m/s could be achieved, although the majority of the measurements were done in the 3m/s to 5m/s range. Chord Reynolds numbers were varied from $Re_c = 30,000$ to 100,000. A computer controlled shutter system at the downstream end of the test section allowed the freestream speed to be modulated at frequencies up to 3 Hz, and velocity fluctuation amplitudes up to 10 percent of the mean flow speed. The highest level of freestream turbulence level was measured to be 0.6 percent at an average speed of

14

A summary of the results follows in the next two subsections, and additional details of the investigation can be found in Williams, et al. (2007).

3.3.1. Rectangular Planform Wing design

The rectangular airfoil models used in the experiments were mounted on a sting connected to a force balance and pitching mechanism. The NACA 0012 and flat plate airfoils had a chord $c = 203\text{mm}$, span $b = 406\text{mm}$, and aspect ratio $AR = 2$. The flat plate airfoil had a thickness of 14 mm (56 in.) giving a thickness ratio $t/c = 0.07$. A 5:1 elliptic leading and trailing edges were used similar to the design of Torres and Mueller (2004). Actuation of the NACA 0012 consisted of four pulsed-blowing slots positioned along each wing tip, and eight slots located along the leading edge. Each actuator was isolated from its neighbor, and could be individually activated. For the flat plate (thin airfoil), 16 actuators were located along the leading edge. The leading edge was modular, and could be exchanged with other leading edges to change the direction of actuation. A drawing of the leading edge actuator for the upstream blowing configuration is shown in figure 3.3.1. The diameter of each actuator jet was 0.79 mm (0.031 in.), and angled 10° upward from the chord line of the airfoil. Straight-blowing actuation used jets aligned with the flow, whereas, "outward blowing" actuation used the same configuration as shown in figure 3.3.1, with the addition of a 45° spanwise component to the jets.




Fig. 3.3.1: Nose section for upstream actuation (dimensions are in inches.)

When pulsed-blowing actuation was used, there were two forcing frequencies of interest; namely, a low frequency at $F^+ = \omega c/U = 0.67$ (~ 5 Hz) corresponding to the global instability of the separated flow region, and a high frequency ($F^+ \approx 2$) associated with the Kelvin-Helmholtz instability of the separated shear layer.

3.3.2. PIV Measurements of Shear Stress

A qualitative measure of the effect of upstream actuation on the separated region comes from the PIV measurements of the Reynolds stress as shown in figures 3.3.2a-d. The size of the

16

recirculation region is decreased with steady blowing but has not been completely, which can be seen by comparing the streamlines. This is consistent with the flow visualization images at the mid-span. The center of the recirculation appears as a focus in the streamline patterns, which we believe is the result of a strong spanwise component of flow in the recirculation region.

The maximum values of negative Reynolds stress ($u'v'$) were found near the dividing streamline over the separated flow region, and indicate a transfer of energy from the mean flow into the turbulent flow. In the actuated case (figure 3.3.2b) the Reynolds stress shows a strong positive value near the leading edge, possibly associated with the strong favorable pressure gradient in the leading edge region. Comparing figures 3.3.2c and 3.3.2d we see that the upstream actuation accelerates the formation of the negative Reynolds stress region, resulting in a reduction of the recirculation region.




Fig. 3.3.2: Arbitrary streamlines and Reynolds stress contours for the baseline and upstream forcing cases at the mid-span and $y/(b/2) = 0.33$ span of the flat plate airfoil. Reynolds stress contour levels range from $-20 \text{ m}^2/\text{s}^2$ (blue) to $+15 \text{ m}^2/\text{s}^2$ (red). The amplitude of the steady blowing was $C_p = 8.3 \times 10^{-3}$ percent. The x and y axis labels are in millimeters. a) baseline, $y/(b/2) = 0$; b) upstream actuation, $y/(b/2) = 0$; c) baseline, $y/(b/2) = 0.33$; d) $y/(b/2) = 0.33$.

Measurements of the overall lift and drag were obtained with a six-component force balance. The effects of upstream and downstream actuation on lift with steady and pulsed blowing are shown as a lift increment $\Delta C_L = C_{L, \text{actuated}} - C_{L, \text{baseline}}$ in figure 3.3.3. The error bars shown are based on the maximum variations observed after repeating the experiments several times. The

17

downstream oriented actuation is a common configuration used in active flow control, and the values of the momentum coefficient ($C_p \sim 0.1$ percent) required to enhance the lift are typical of those observed by other investigators. Lift is gradually and continuously increased as the momentum coefficient is increased. The lift does not appear to saturate with downstream directed actuation, and presumably larger forcing amplitudes would result in even higher C_L values.

Significantly different response of the lift coefficient is seen with the upstream actuation, shown in figure 3.3.3b. A close-up of the rapid change in C_L is shown in figure 3.3.3c. The maximum increase in lift occurred at very low amplitude forcing and saturated almost immediately. The lowest resolvable supply pressure with our control system pressure regulator was 860 Pa (0.125 psig). The maximum lift increment occurred at 1.72 kPa (0.25 psig) corresponding to $C_p = 8.3 \times 10^{-3}$ percent, which is two orders of magnitude lower than that achieved with downstream actuation. While this is a very encouraging result, suggesting that forces on an airfoil may be controlled with extremely low forcing amplitudes, we caution that the reasons for the upstream actuation efficiency are not fully understood yet. Furthermore, the PIV and smoke visualizations show that the saturated state with upstream actuation does not completely eliminate the flow separation in the mid-span region of the airfoil.

Four different types of actuation are shown in figure 3.3.3: namely, (1) steady-straight, (2) steady-crossflow, (3) pulsed-straight and (4) pulsed-crossflow, where straight refers to being along the x -axis and outward indicates a 45° angle toward the tips of the airfoil. The effects of the four types of downstream-oriented actuation are shown in figure 3.3.3a, where it can be seen that the outward-span pulsed blowing at $\phi = 51.1^\circ$ ($F = 0.67$) was the most effective at increasing the lift coefficient. However, when upstream actuation is used, then straight-steady actuation is the most effective at the low amplitudes but steady-crossflow becomes more effective at the higher forcing amplitudes.

3.4. Second Generation Wing - Semi-Circular Planform

The semi-circular planform provides a continuously varying sweep angle from 0° at the center span to 90° at the tip. In comparison to the rectangular planform wing it was expected that the leading edge vortex would be more receptive to the pulsed-blowing actuation and less likely to shed the vortex. This assumption turned out to be correct, and the semi-circular planform was chosen as the test article for the investigations dealing with the use active flow control in dynamic flow situations.

The experimental effort in the wind tunnel experiments was aimed at constructing closed-loop control systems for modulating the strength of the leading edge vortex in unsteady flow conditions. Pressure sensors on the suction surface of the airfoil were used to detect the early stages of stall, which were coincident with the formation of the leading edge vortex, although the overall lift force was ultimately used as the feedback signal. For the first attempt at feedback control for this problem, we used a quasi-static approach to closed-loop control to adjust the strength of the leading edge vortex in response to an oscillating free stream, which was described in Williams, et al. (2008a).

18




Fig. 3.3.3: Effect of actuator configuration on lift increment at $\alpha = 16^\circ$: (a) downstream actuation; (b) upstream actuation; (c) close-up view of upstream blowing, $C_p < 5 \times 10^{-4}$ percent.

3.4.1. Semi-Circular Planform Wing Components

A photograph of the disassembled airfoil model is shown in figure 3.4.1. The plenum covers are removed to expose the 16 micro-valves that control the pulsed-blowing to the leading edge. The planform is a semi-circle with a centerline chord $c = 203 \text{ mm}$, and span $b = 406 \text{ mm}$ and aspect ratio $= 2.54$. Although steady spanwise directed blowing is known to stabilize the leading edge vortex on rectangular wings, the mass flow rate requirements were quite large, and in the case of the semi-circular wing the lift increments were negative with steady blowing. With pulsed-blowing, on the other hand, lift coefficient increases up to 40 percent (depending on angle of attack) could be achieved. Each micro-valve actuator was isolated from its neighbor, and could be individually activated to produce traveling wave patterns, however, all actuators were driven

19

in-phase for these measurements. To document the open-loop forcing effects on performance, the actuators were operated at a 25 Hz pulse rate, and a $C_p = 0.0074$.



Fig. 3.4.1 - View of the disassembled wing model with the plenum cover plate removed. The 16 micro-valve actuators can be seen positioned radially along the circular leading edge.

The transient response of the leading edge vortex and the tip vortex system to open-loop forcing by the actuators, such as, pulse and step inputs, was obtained for modeling and validation purposes. The convection of the leading edge vortex over the airfoil was identified from surface pressure measurements at $x/c = 0.42$ and $x/c = 0.72$. Force and moment measurements were done with an ATI nano-25 or -18, six-component force balance system.

3.4.2. Smoke wire flow visualization

Flow visualization of the flow at two spanwise locations over the semi-circular airfoil at 19° angle of attack is shown in figures 3.4.2 a-d. The smoke sheet is positioned at a center span of the wing in figures 3.4.2a and 3.4.2c, and aligned with the quarter span in figures 3.4.2b and 3.4.2d. As expected, without flow control the flow is fully separated at this angle of attack, figure 3.4.2a and figure 3.4.2b. Activation of the 16 pulsed-blowing actuators (all in phase with each other) along the leading edge partially reattaches the flow as shown in figure 3.4.2c and 3.4.2d. An intensified leading edge vortex can be seen in figure 3.4.2d, indicating that the pulsed blowing actuation has captured and intensified the LEV. PIV data confirmed that the vorticity along the leading edge was concentrated by the actuation. Additional details are provided in Williams, et al. (2008b).

20




Fig. 3.4.2: Smoke wire visualization of the flow over the wing, $\alpha = 19^\circ$, $Re_c = 68,000$ a) no forcing, center span; b) no forcing, quarter span; c) actuation on, center span; d) actuation on, quarter span.

3.4.3. Response to pitching maneuver

It is well known that airfoils and wings pitched upward at high pitch rates form dynamic stall vortices. The extra circulation associated with these vortices leads to higher lift coefficients than can be achieved under steady conditions. The lift coefficient results in figure 3.4.3 indicate that our semi-circular wing exhibits similar behavior. An approximate steady state lift curve is obtained by measuring lift during a slow pitch rate of $\dot{\alpha} = 0.9$ deg/sec. Static stall occurs at $\alpha = 16^\circ$, while at the higher pitch rates of $\dot{\alpha} = 40$ and 80 deg/sec stall is delayed to $\alpha = 24^\circ$ and 32° , respectively. The presence of a dynamic stall vortex increases the lift coefficients to $C_L = 1.4$ and $C_L = 1.8$ for the two pitch rates. Pitch-down maneuvers at the same constant rates are also shown in the figures to demonstrate the symmetry of the result. Since the wing has no camber, the anti-symmetry in C_L was expected about $\alpha = 0^\circ$, when the flow was attached. A hysteresis effect in the lift curve response was seen when the flow is separated.

21




Fig. 3.4.3: Lift coefficients at pitch rates $\dot{\alpha}/dt = 0.9, 40$, and 80 deg/sec.




Fig. 3.4.4: Lift dependence on angle of attack. Forcing on at 25 Hz, $C_p = .0074$ and pitch rates $\dot{\alpha}/dt = 0.9$ and 40 deg/sec. Forcing in the "static" case produces same lift coefficient as the dynamic case.

22

3.4.4. Open-loop forcing

The effect of open-loop forcing on the lift coefficient is shown in figure 3.4.4. For reference purposes, a pitch-up baseline case $\dot{\alpha}/dt = 40$ deg/sec with actuators turned off is shown by the green curve. The red data curve corresponds to actuators continuously pulsed at 25 Hz, $C_p = .0074$, with a slow pitch rate $\dot{\alpha}/dt = 0.9$ deg/sec. The baseline steady lift case with no actuation is shown by the blue line. The effect of actuation is similar to the dynamic stall effect. In both cases stall is delayed until $\alpha = 23^\circ$ where a maximum lift coefficient of $C_L = 1.4$ is reached. This is indirect evidence supporting the earlier observation that actuation has the effect of stabilizing the leading edge vortex.

Since closed-loop control was to be used to obtain the same high lift coefficient values during wing maneuvers and flight in an unsteady freestream, knowledge of the flow state on the wing was an essential element in the development of the controller. In Williams et al. (2008b) we explored feedback signals based on the lift force and on pressure taps located at $x/c = 0.42$ and 0.72 . Figures 3.4.5a and 3.4.5b show the mean pressures measured at $x/c = 0.42$, corresponding to $\dot{\alpha}/dt = 0.9$ and 40 deg/sec pitch-up and pitch-down maneuvers. The quasi-steady data at $\dot{\alpha}/dt = 0.9$ is used as input to a second order differential equation model, which predicts the effects of the pitch rate on the pressure. Figure 3.4.5a shows that the pressure decreases linearly prior to flow separation, irrespective of the pitch rate. During a pitch-up maneuver, the separation is delayed (red dashed line), and during pitch down the reattachment of the flow is delayed (dash-dot line).

The effect of actuator forcing on the surface pressure during slow and rapid pitching maneuvers is shown in figure 3.4.5b. The pressure response to actuation is somewhat different from the response to pitch rate shown in figure 3.4.5a. Consequently, it is necessary to include the effect of actuation in the plant model in order to build an effective closed-loop control system. This topic will be revisited in Section 3.6.




Figure 3.4.5: Effect of pitch rate on the surface pressure measured at $x/c = .42$. a) No forcing at wing pitch rates $\dot{\alpha}/dt = 0.9$ deg/sec (blue and cyan) and 40 deg/sec (red and brown). b) Forcing on at 25 Hz, $C_p = .0074$ and pitch rates $\dot{\alpha}/dt = 0.9$ (blue and cyan) and 40 deg/sec (red and brown).

23

3.4.5. Lift spectra dependence on angle of attack

Lift fluctuations are driven by vortex shedding from the wing at frequencies that are dependent on the angle of attack. The dominant frequency obtained from the spectrum of the lift signal is plotted in figure 3.4.6 for angles of attack varying from $\alpha = 9^\circ$ to 31° . From figure 3.4.4, we saw that stall begins at approximately $\alpha = 15^\circ$, and the onset of stall corresponds to a decrease in the frequency of the lift fluctuation. The red dashed line in the figure 3.4.6 corresponds to a Strouhal number, $St = f^*c^*/U = 0.2$, which was originally proposed by Fage & Johansen (1927) for two-dimensional flows. In this version of the Strouhal number the length scale of the frequency is based on the wake width (or equivalently the projected mid-span chord) $c^*\sin(\alpha)$. It is interesting that a low-aspect ratio, three-dimensional wing shows good agreement with this scaling after the wing is completely stalled.




Fig. 3.4.6: Dominant lift fluctuation frequency dependence on angle of attack and comparison with $St = f^*c^*\sin(\alpha)/U = 0.2$.

3.5. Scaling of transient lift response to pulsatile actuation

For the most part, active flow control studies have focused on the steady-state behavior of the flow, where only time-continuous (e.g., sinusoidal) actuation is needed for producing changes. Early studies (Aluja, K. & Burris, R. 1984; Neuberger, D. & Wygnanski, I. 1987) determined that effective actuation frequencies should be scaled with convective time, $t^* = c/U$, which is the time for disturbances to advect over a certain characteristic length of the wing. The amplitude of the steady-state lift response is usually correlated with the momentum coefficient, C_p . The frequency and amplitude scaling parameters have physical meaning through their connection with flow instabilities. When the airfoil is in a fully stalled state, then the convective time scale is comparable to the period of vortex shedding in the wake. There is a coupling between vortical structures in the wake and the separating flow from the airfoil, so that steady-state actuation at $St = O(1)$ affects the coupling between the wake and the airfoil (Wu, et al. 1998; Brzozowski & Cleary 2006). More detailed studies using numerical simulations of flow over a two-dimensional airfoil (Raju et al. 2008) identified three naturally occurring flow instabilities, which exist during

24

steady-state conditions and are important to the dynamics of the actuator-to-flow interaction mechanism that must be modeled for control. The instabilities are connected with the shear layer, the separation bubble and the wake, and each has its own length scale and specific frequency scaling parameters. Adding unsteady aerodynamic effects on top of this already complex mix of instabilities suggests that new approaches in flow control may be necessary.

In attempting to better understand the problem, we studied the response of the separated flow system to individual pulses from the actuator. The transient behavior of the forces acting on wings in response to pulse-type or step-input disturbances can be significantly different from the steady-state response. An extensive study of 'step-input' transient flow associated with reattachment and separation was conducted by Darabi and Wygnanski (2004a, 2004b) on a two-dimensional deflected flat plate. Using step inputs from a zero-net-mass, voicecoil driven actuator they showed that the total time it takes for the flow to reattach on a deflected flat plate was long ($0.20-0.50 \tau^*$) in terms of convective time units, and was to a large extent dependent on the frequency and amplitude of excitation. For fixed values of the forcing parameters C_a and F^+ , the transient lift response to a step input scaled with the dynamic pressure (C_a) and the convective time-scale (τ^*).

The transient response of the separated flow on 2D airfoils to pulse-type actuation input was investigated by Amitay and Glezer (2002, 2006). They documented the effect of flow transients occurring at the onset and termination of actuation. Glezer and co-workers (Benjamin et al. 2006, 2006, 2006, 2006) also studied the lift response to short duration pulses produced by combustion actuators. They showed that energetic pulses from the actuators with time scales as short as $O(0.05 \tau^*)$ were effective in producing a momentary increase in circulation around the airfoil. The relaxation time of the circulation back to the undisturbed state was long, on the order of $O(5-10 \tau^*)$. The effect of different pulse sequences on the circulation was also explored.

Here and in the paper by Williams et al. (2009a) we use short-duration pulsatile disturbances for the purpose of system identification on the separated flow around a low aspect ratio, three-dimensional wing. The output from the separated flow "system" response to a single pulse gives a sort of "impulse response" model that can, within certain limits, be used to predict the large-scale features of the lift response to various time varying actuator input signals. The model can also be used as part of a control algorithm to simulate actuator response.

By studying the lift force transient to pulse-like disturbances from the actuator, it is possible to obtain input-output models that provide insight into the flow physics of the separated flow, in addition to obtaining scaling relations for the flow's response to actuator input and simple models that are useful for designing closed-loop controllers. In this section the response of a wing fixed at $\alpha = 20^\circ$ to pulse-type disturbances with varying amplitude and different freestream speeds is examined.

The peak lift coefficient increment above the unforced steady state value is shown in figure 3.5.1 for the response to a single pulse from the actuator. A wide range of actuator pressures, flow speeds and pulse time intervals are shown. A nearly linear increase with the square root of the actuator pressure coefficient can be seen until saturation occurs. The square root of the pressure coefficient correlates with the actuator jet exit velocity. Using hot-wire measurements of the actuator jet exit velocity, we found that $C_a^{0.5} = 556 U/U_1$.




Figure 3.5.1: The peak in the transient lift coefficient dependence on actuator pressure coefficient. The actuator supplied only a single pulse with different supply pressures, flow speeds and valve open times.

A variety of lift coefficient transients are plotted against the convective time scale $\tau^* = U/c$ in figure 3.5.2. Each curve corresponds to a different freestream speed and actuator pulse duration time, Δt_{act} , which are chosen so that the non-dimensional pulse duration time is kept constant, $\Delta t_{act}^* = \Delta t_{act} U/c = 0.5$. At flow speeds below 4 m/s the peak lift increment saturates as $C_a^{0.5}$ exceeds 50.

The lift response to a single pulse can be treated as a filter kernel, and can be used to predict the lift time history for arbitrary actuator input signals, at least up to $St = 0.2$. Even though the detailed interactions between the actuator input and the response of the lift force are almost certainly non-linear, the results in figure 3.5.1 indicate some degree of linear behavior over the range of operating conditions. The phase-averaged single-pulse response is used as an approximation of an impulse response kernel, $K(t)$, in the convolution to obtain a predicted output signal, $w(k) = C \sum_j K(t_j) u(k-j)$, where $u(k)$ is the arbitrary input signal. The magnitude input signal $u(k)$ is arbitrarily given an amplitude of 1.0. To find the calibration constant, C , the total impulse of the 1-, 3-, 5- and 10-pulse experiments was compared with the corresponding predicted total impulse using the model.

The predicted and measured lift transients are shown in figure 3.5.3 for the 3-, 5- and 10-pulse cases. The agreement between the model and the measurements are satisfactory, although the model over-predicts transient overshoot and undershoot in the lift response for the 10-pulse signal.




Fig. 3.5.2: transient lift coefficient increment dependence on convective time. Different curves correspond to different flow speeds and pulse duration times so that the non-dimensional pulse duration time is constant, $\Delta t_{act}^* = \Delta t_{act} U/c = 0.5$.




Fig. 3.5.3: Comparison of experimentally measured lift response and the lift predicted by convolving a 10-pulse square wave signal with the impulse response kernel. Freestream speed was $U = 5$ m/s, and 34.5 kPa supply pressure to valves.




Fig. 3.5.4: Comparison of measured lift coefficient (dashed line) and the predicted lift (solid line) is shown for a forcing frequency of 0.4 Hz. The baseline lift coefficient without actuation is 0.8.

The low-pass filter character of the system becomes apparent when the forcing frequency is increased to 1.4 Hz. The data in figure 3.5.5 show that the square wave input to the valves is rounded at the corners of the lift signal, because of the attenuation of higher frequencies. The phase delay between actuator input and lift response becomes significant at this forcing frequency. The phase shift between the actuator input and the lift at 1.4 Hz is now increased to $\phi = 79.8^\circ$. The linear model predicts a phase shift of $\phi = 82.5^\circ$, so it does a good job of reproducing the phase.

When the forcing frequency is increased to 5 Hz as shown in figure 3.5.6, then the amplitude of the lift fluctuation is significantly reduced by the filtering effect of the kernel. The amplitude of the lift fluctuations is significantly attenuated, and the phase shift is $\phi = 182^\circ$. At this frequency the model under-predicts the fluctuating lift force, and over predicts the phase ($\phi = 243^\circ$) between actuator input and lift fluctuation.

The actuation can only produce positive lift perturbations, which has the effect of increasing the average lift above the baseline state. In the case of figure 3.5.5, the mean lift coefficient shifts from $C_L = 0.8$ to $C_L = 1.1$. The lift fluctuations oscillate about the new mean lift value rather than the baseline lift. Ordinarily one might interpret the shift in the mean lift as a nonlinear effect, but the results show that the linear model gives a good prediction of the new mean lift.




Fig. 3.5.5: Comparison of measured lift (dashed line) and predicted lift (solid line) is shown for a forcing frequency of 1.4 Hz. The baseline lift coefficient without actuation is 0.8.




Fig. 3.5.6: Comparison of measured lift (dashed line) and predicted lift (solid line) is shown for a forcing frequency of 5.0 Hz. The baseline lift coefficient without actuation is 0.8.

29

In the following sections, the use of system identification techniques to obtain two linear models which approximate the separated flow response to actuator input disturbances is obtained. The first model is a higher order which is compared to the pulse response results. A second lower order model is then obtained using similar system identification techniques, which acts as a plant model for the design of a standard proportional-integral (PI) controller. The ability of the feed forward, PI-controller to maintain a constant lift force is tested with randomized "step changes" in the freestream speed of the wind tunnel. In Section V the identified models are used to speculate about the possible improvements in system response that can be achieved with closed loop flow control.

3.6.1. Model the lift response to actuation

The range of possible control is first determined by creating a static map, which is the steady state lift response to pressure actuation. The angle of attack is fixed at 20° , the micro-valves are pulsed continuously at 29Hz, and the plenum pressure is varied in 1kPa increments. Data is acquired for 60 seconds at each pressure magnitude and the mean value of the lift is found. The process is repeated for varying flow speeds. Figure 3.6.1 shows the mean lift response at flow speeds 4m/s through 7m/s with the measurements at 5m/s and 7m/s repeated as a repeatability test. The data collapses to a single "linear" curve when plotted as the change in lift coefficient versus the square root of the jet pressure coefficient, see Williams et al. (2010b and 2010c).




Fig. 3.6.1: Pulsed blowing wing's static map of lift response to pressure actuation

30

The understanding of the transient response to actuator input and the ability to predict the lift response given an arbitrary actuator input set the stage for using active flow control in dynamic flow situations. The next Section 3.6 explores the use of active flow control in an unsteady freestream flow. Much of the controller design and system identification resulted from work done in collaboration with Professor Rüdiger King and his students at the Technische Universität Berlin, Berlin Germany.

3.6. Dynamic active flow control – flow control under changing flight conditions

The conventional analysis of flight control in gusting flow conditions, Hoblit (1988), considers only the time-averaged statistical properties of the flow unsteadiness. If the control system is somehow to be able to react to instantaneous changes in flow speed and direction associated with gusts, then it may be possible to fly the vehicle in a way that allows it to extract energy from the gusts. With proper application of active flow control techniques, it may be possible to realize significant range and endurance enhancements in flight vehicles that use real-time control of the vehicle in response to flow unsteadiness. Real-time control of flight through unsteady and gusting flows requires control systems that account for both the unsteady aerodynamics and the flow system response to the actuator. Quasi-steady models are usually not sufficient for closed loop control when the flight vehicle response is expected to be faster than $k \cdot \text{ref}/U > 0.05$, because of the time lags and amplitude changes associated with both the unsteady aerodynamics and the actuator response.

Previous measurements of the transient flow response to actuation showed relatively long time delays in the flow response to the onset of actuation. Normalized time delays $t^* = 5.3$ were measured for response to sinusoidal inputs from the actuator. Phase lags between the freestream speed oscillation and the lift force were measured to be $\phi = 30^\circ$ at $k = 0.2$. These are large delays, and an effective controller must be able to quickly compensate for these time delays and for changes in amplitude associated with the flow disturbances.

The ability of a simple feed-forward controller to suppress the lift force fluctuations produced by a sinusoidally oscillating freestream flow was demonstrated by Williams, et al. (2009a and 2010a). However, that control approach was limited to a specific frequency of oscillation. The controller required manual changes in amplitude and phase delay in order to account for any frequency or amplitude changes. A more general approach is to use a system model that contains the amplitude and phase information of the lift coefficient response to actuation over a wide band of operating conditions.

Because the wing is in a fully stalled state with a fixed $\alpha = 20^\circ$, one might expect nonlinear behavior to dominate the response to actuation. The use of neural networks or look-up tables as nonlinear models is an option for the control approach. However, the pulse response experiments described in Section 3.5 indicated that linear models can be used effectively within certain limitations. A linear system model can almost always be obtained, but the question is over what range of conditions will it be valid? An equally important question is given a linear model of the system, can a useful controller be designed? We investigated these questions using experimental data to obtain black-box system models, and using conventional linear controller design techniques.

To identify dynamic models of the wing's response to actuation, black-box system identification methods are employed. Black-box modeling requires measuring the system response to pseudo-random binary signal (PRBS) step inputs in plenum pressure. The input signal is the recorded "desired pressure" and the output signal is the change in lift signal measured by the force balance. Horan the change in lift is defined as the deviation from the steady-state lift that corresponds to actuation at the lower pressure level. One example of the input and output measurements is shown in figure 3.6.2.




Fig. 3.6.2: Example of input-output data used to obtain black-box dynamic models of response to pressure actuation. Dark line shows one of the identified models.

The experiments were repeated at varying magnitudes of input pressures and pressures for two flow speeds of approximately 5 m/s and 7 m/s. A family of 21 linear black-box dynamic models was identified using the Prediction-Error-Method. First order models with a time delay of the form,

$$G_p(s) = \frac{k}{Ts + 1} e^{-\tau_d s} \quad (3.6.1)$$

fit the data well, as can be seen in figure 3.6.2 from the comparison of the measured and simulated response for one of the identified models. The consideration of the time delay results in a significantly better fit of the experimental data when compared to models identified in earlier investigations, Williams, et al. (2010a).

The frequency response of all identified models is shown in figure 3.6.3. Each model family identified for a fixed flow speed is characterized by a variation of parameters corresponding to the actual nonlinear response of the flow to the actuation. This relates also to the nonlinearity seen in the static maps for the different flow speeds shown in figure 3.6.1.

32

Furthermore the majority of the models identified at the lower flow speed of $U = 5$ m/s show a smaller gain than the ones identified at $U = 7$ m/s, which is mostly due to the fact that the plant models are defined with respect to dimensional variables. This in turn allows for an easier controller design and implementation, since the control objective is mainly to reject disturbances while maintaining a constant lift force.




Fig. 3.6.3: Frequency response diagram of identified models at flow speeds 5 m/s and 7 m/s. The dashed line is the nominal model found from the mean of the family model parameters for 7 m/s. To obtain a rational transfer function the time delay of the nominal model is approximated by an all-pass transfer function.

The control design presented in this paper focuses on achieving good disturbance suppression at a nominal flow speed of 7 m/s. Therefore, a nominal model was found by taking the mean of only the transfer functions identified at this flow speed. In order to obtain a rational transfer function, the dead-time element corresponding to the mean time delay of $\theta = 0.157$ s is approximated by a third order all-pass transfer function. Its coefficients are determined based on a least squares method proposed by Frank (1996), which minimizes the difference between the step responses of the original and the approximated transfer function. However, the approximation leads to a deviation of the phase for frequencies larger than about 6 Hz as can be seen from figure 3.6.3. This is acceptable, because the deviation lies well above the frequency range of interest for the controlled plant.

33

loop part of the control architecture provides set-point tracking at zero steady-state error, which accounts for model uncertainties and compensates the disturbances acting on the plant at low frequencies.




Fig. 3.6.5: Controller architecture used for closed loop experiments.

The output disturbance y_d corresponding to a deviation of the lift force is caused by fluctuations in the flow speed $d = U$. This represents the unsteady aerodynamics of the wing and can be modeled by the black-box disturbance model G_d . As the flow speed is measured online, the input d to the disturbance model is known. This information can be exploited by using a feedforward controller $K_d(s)$, which acts on the plant input u to enhance the disturbance compensation.

To account for actuator saturation due to the limited actuation pressure the control loop is augmented with a dynamic anti-windup compensator based on a method suggested by Park (1993). It is not shown in Figure 3.6.5 for the sake of conciseness.

A robust H_∞ feedback controller $K(s)$ is synthesized using the mixed-sensitivity loop-shaping approach Skogestad and Postlethwaite (1996). This closed-loop control strategy has been successfully applied in several active flow-control experiments (Willms, et al. 2010b, Hennings, et al. 1997, Hein, et al. 2010) and is augmented here by a feedforward controller $K_d(s)$ for improved disturbance rejection. By choosing appropriate loop-shaping weights, the mixed-sensitivity control synthesis guarantees robust stability and performance of the closed loop for a given model family. To do so, the maximum deviation of all models $G_p(s)$ within the model family Π from the nominal model $G_n(s)$ is described by a multiplicative uncertainty

$$l_i(\omega) = \max_{G_p(s) \in \Pi} \left| \frac{G_p(j\omega) - G_n(j\omega)}{G_n(j\omega)} \right| \quad (3.6.2)$$

Hence, the model family can be described by

$$\Pi: G_p(s) = G_n(s) [1 + w_1(s) \Delta_1(s)] \quad |\Delta_1(j\omega)| \leq l_1(\omega), \quad \forall \omega. \quad (3.6.3)$$

35

3.6.2. Modeling the unsteady aerodynamics

The wing's response to a time varying longitudinal "gusting" velocity is itself dynamic. The separated flow and the low aspect ratio of the wing do not lend themselves to any available theory so a separate black-box model is identified from experimental data. The lift response of the wing to sinusoidal velocity inputs at several frequencies is measured. The amplitude ratio of the lift force to the velocity amplitude is determined from the energy at the fundamental frequency in the power spectrum. The phase between the resulting lift force and the velocity, from hotwire measurements, is determined from the cross power spectrum. Figure 3.6.4 shows the ratio of the lift force amplitude to the velocity amplitude, and the phase between these signals, plotted against the frequency of velocity fluctuations. The frequency response of the wing to gusting conditions is then used to identify a black-box dynamic model. The resulting model becomes the disturbance model in the control architecture; see below for a description of the control architecture. The final form of the disturbance model is also shown in figure 3.6.4.




Fig. 3.6.4: Unsteady aerodynamic model; frequency response of lift force to longitudinal gusting flow with a mean flow speed of 7 m/s. Individual points represent measurements from sinusoidal velocity forcing and the solid line is the model generated from the frequency response measurements.

3.6.3. Controller Architecture and Synthesis

The main control objective is to maintain a constant lift by suppressing disturbances caused by sudden flow speed variations. This is achieved by employing a two degrees-of-freedom controller shown in figure 3.6.5. The output y of the plant model G_p is perturbed by a disturbance y_d . Therefore, the actual lift force y is measured and compared against the reference value r . A robust feedback controller $K(s)$ regulates the lift force by adjusting the actuation pressure p_1 . The closed-

34

wherein $\Delta_1(s)$ denotes a normalized uncertainty with a frequency dependent weight $w_1(s)$ comprising all identified transfer functions. Figure 3.6.6a shows the multiplicative uncertainty $l_1(\omega)$ and the magnitude of the corresponding weight $|w_1(\omega)|$ for the family of models identified for the wing. The uncertainty could be reduced by inverting the static map shown in Figure 3.6.1 for one fixed flow speed and using it as a pre-compensator to account for the steady-state part of the nonlinearities. This was examined in earlier experiments by the authors but turned out not to be necessary in the current control design, since the closed-loop performance is limited by the time delay of the plant transfer function. This limitation will be discussed further towards the end of this section.




Fig. 3.6.6: Multiplicative uncertainty for the identified model family (a) and loop-shaping weights with corresponding transfer functions for the synthesized H_∞ controller (b)

To tune the controller $K(s)$ one considers the closed-loop response of the nominal plant, which is given by

$$y = \frac{G_n K}{1 + G_n K} r + \frac{1}{1 + G_n K} \left(\frac{G_n K}{s} \Delta_1 \right) \frac{r}{s}. \quad (3.6.4)$$

wherein T represents the complementary sensitivity function relating to tracking performance and measurement noise. S denotes the sensitivity function relating to suppression of disturbances acting on the output of the closed loop. Finally, S_0 can be interpreted as a feedforward sensitivity function (Skogestad & Postlethwaite 1996). The sensitivity function S and the complementary sensitivity function T are shaped by the weights $w_1(s)$ and $w_2(s)$, respectively. A third weight $w_3(s)$ is used to put a bound on the control effort Ks . In order to obtain the controller a cost functional

36

$$\min \|N(K(s))\|, \text{ with } N = \begin{bmatrix} w_1 S & w_2 T & w_3 K S \end{bmatrix}, \quad (3.6.5)$$

has to be minimized, wherein $K(s)$ denotes the optimal controller. The frequency response of the closed-loop transfer functions with the corresponding loop-shaping weights is shown in figure 3.6.6b. Note that plant model was scaled to an input and output variable range during the loop-shaping process to allow for easier choice of weights. Adjusting the weight $w_1(s)$ such that

$$|T(j\omega)| \leq 1/\mu_s(j\omega), \quad \forall \omega, \quad (3.6.6)$$

ensures robust stability of the closed-loop for all models identified for the flow speeds 5 m/s and 7 m/s. Note that the magnitude of the uncertainty $|w_1(s)|$ exceeds unity for frequencies larger than approximately 2 Hz. This puts an upper limit on the achievable bandwidth ω_{BT} with respect to the set-point tracking performance. However, figure 3.6.6b reveals that the magnitude of the complementary sensitivity T is still well below this limit. Other limitations arise on the one hand from constraints on the physically possible control effort, and on the other hand from the right-half-plane zeros corresponding to the approximation of the time delay θ in the nominal model. It can be shown that for systems with time delays the closed-loop bandwidth is limited to be less than $1/\theta$ (Skogestad & Postlethwaite 1996). Due to these limitations, a bandwidth of about $\omega_B \approx 2.7$ rad/s or 0.43 Hz is achieved when just considering the feedback part of the controller, as can be seen from Figure 3.6.7. Here, the bandwidth ω_B is defined as the frequency where the Sensitivity S crosses the -3dB line for the first time from below. Note that the feedback controller shows a worse performance than the uncontrolled case in a frequency band above approximately 0.7 Hz. This can be explained by the so called second waterbed formula, which is based on a weighted sensitivity integral. It states that reducing the sensitivity of a plant with right half-plane (RHP) zeros at low frequencies will cause a large peak in the sensitivity over a limited frequency range.

Since the input d to the disturbance model G_d can be measured online, the bandwidth can be improved by using a feedforward controller K_d , which is calculated by

$$K_d = G_d \hat{G}_c^{-1} G_p, \quad (3.6.7)$$

Herein $\hat{G}_c$ denotes the allpass-free part of the nominal plant model to yield a stable inverse, and G_p represents a fast first order filter to render the transfer function K_d causal. Figure 3.6.7 shows that the feedforward controller increases the bandwidth of the controlled plant with respect to the measured disturbances to about 0.7 Hz. However, this comes at the price of increasing the sensitivity even further at a frequency band above approximately 0.8 Hz.

3.6.4. Cancellation of lift fluctuations in gusting flow

The performance of the controller is evaluated by subjecting it to a pseudo-random velocity signal (PRS). The voltage signal to the shutters is constructed of pseudo-random amplitude steps summed with pseudo-random amplitude sinusoidal signals of frequencies less than 1 Hz. The velocity input has a bandwidth of approximately 1 Hz. Figure 3.6.8 a) shows the magnitude of the velocity plotted against time, and b) shows the power spectrum magnitude of the velocity plotted

against frequency. The velocity ranges from a minimum of 6.25 m/s to a maximum of 7.25 m/s with the mean flow speed of 6.9 m/s. The same signal is repeated 15 times to reduce the uncertainty in the amplitude of the power spectrum to below 26% and the resulting time series measurements are averaged to reduce uncorrelated measurement noise.




Fig. 3.6.7: Frequency response of the plant output to sinusoidal disturbances in the flow speed for the uncontrolled plant (blue line), the feedback controlled plant (green line) and the feedback controlled plant augmented by a feedforward disturbance compensation (red line).




Fig. 3.6.8: a) Measured pseudo random velocity time series and b) power spectrum used to test the controller

37

38

The controller is commanded to maintain a constant reference lift of 1.8N during the PRS experiments. The reference lift is above the maximum value of the uncontrolled lift, this was done to reduce the effect of measurement noise. Figure 3.6.9 shows the averaged experimental controlled and uncontrolled lift force time series along with simulation results and the desired lift force. The simulation results are obtained using the averaged, experimentally measured velocity profile as an input to the disturbance model and the same reference lift from experiment. The resulting signals are passed through the closed loop and feed forward disturbance controllers and the plant model response to actuation. Without control the lift reaches a minimum of 1.1N and a maximum of 1.5N. With control the lift ranges from 1.7N to 1.8N, where the minimum is from a point where the actuator input is saturated and the required change in lift exceeds the maximum possible value.




Fig. 3.6.9: Phase averaged controlled and uncontrolled lift time series and comparison with model(s) (disturbance model is used in calculating the power spectrum giving an uncertainty in the peaks in the spectrum of less than 26% and a frequency resolution of approximately 0.01 Hz).

Figure 3.6.10 shows the power spectrum magnitude of the controlled and uncontrolled lift fluctuations plotted against frequency. Fifteen records of the length of an entire period of the pseudo-random velocity signal are used in calculating the power spectrum giving an uncertainty in the peaks in the spectrum of less than 26% and a frequency resolution of approximately 0.01 Hz.

The simulation results, shown in figure 3.6.9, agree well with the experimental results suggesting that the unsteady aerodynamics and dynamics of the response to pressure actuation are captured well by the linear, black-box, models even though the underlying process is highly complex and nonlinear. The black-box models reduce the infinite dimensional system from the solution of the Navier-Stokes equations to a single input-single output system (SISO). This also suggests that the linear superposition of the response to actuation and the response to the time




Fig. 3.6.10: Power spectrum comparison of fluctuating lift force during controlled and uncontrolled pseudo random velocity input at design conditions.

varying velocity has validity for practical controller synthesis. The use of linear models enables the use of a wide range of relatively simple controller synthesis techniques and analysis tools.

The controller is effective at reducing lift fluctuations at low frequencies, less than ~0.8 Hz ($k=0.08$), but begins to amplify disturbances above this frequency. The Bode integral formula shows that noise attenuation over some frequency band is accompanied by noise amplification over some other frequency band for systems with a pole excess of at least two. Ideally this range of frequency amplification would occur at frequencies too fast for the plant to respond at, but for systems with RHP zeros more severe limitations apply. A weighted sensitivity integral shows that the amplification of disturbances must occur over a limited frequency range. Because of the time delay present in the plant modeling which can be approximated by an all-pass transfer function containing RHP zeros and the disturbance model this amplification occurs from ~1 Hz ($k=0.09$) to ~5.5 Hz ($k=0.5$). The sensitivity of the two-degree of freedom controller used in experiment is given by SS_p (Skogestad & Postlethwaite 1996). The controller is capable of suppressing disturbances where the overall sensitivity is below 1 (0dB) and disturbances are amplified when the overall sensitivity is greater than 1. Figure 3.6.7 shows the overall sensitivity of the modeled plant and disturbance and the point where the overall sensitivity crosses 1 from below at near 1 Hz and crosses 1 from above at 5.5 Hz. This agrees with the experimentally obtained power spectrum of lift fluctuations (shown in figure 3.6.10) where the controlled disturbances are amplified over the uncontrolled fluctuations, again suggesting that the linear models capture the dynamics well.

The nominal time delay from desired pressure input to jet velocity is 0.023 seconds ($\tau=0.8$) while the nominal time delay from desired pressure to lift increase is 0.157 seconds ($\tau=5$), suggesting a time delay from jet velocity to the initiation of lift increase of $\tau=4$. This is believed to be due to the time for a disturbance issued from the actuators to roll up and convect over the wing. This method of actuation does not show the initial decrease in normal force, circulation or lift as seen by Darabi & Wygnanski (2004), Brzezinski & Glesner (2006), or Williams, et al. (2009a). This is believed to be due to the first order model behavior of the jet velocity obtained

39

40

by a desired step increase in pressure as opposed to the step increase in jet velocity observed while maintaining a constant pressure within the wing's plenum, which does show an initial decrease in lift. This fluid dynamic time delay limits the bandwidth of possible control, as discussed above.

The relatively low bandwidth of the pressure regulator and the time delay between step inputs of desired pressure to jet velocity raises the question, if a faster actuator is used would the bandwidth of control be increased? Comparisons with a zero-net-mass-flux (ZNMF) wing (described in the next Section 3.7) with piezoelectric actuators shows negligible time delay between a desired input signal and measured output of jet velocity. Consequently, the jet velocity bandwidth of the piezoelectric actuators is an order of magnitude larger than the pulsed-blowing wing's jet velocity. The ZNMF wing does show the initial decrease in lift (non-minimum phase behavior) as observed by other investigators. The non-minimum phase behavior implies a right half plane (RHP) zero in the transfer function. A RHP-zero imposes control limitations at either low or high frequencies. One can achieve tight control at frequencies below approximately $\pi/2$, where π is the magnitude of the RHP-zero or at frequencies above 2π by reversing the sign of the controller gain. Black-box models of PRBS voltage inputs to the piezoelectric actuators, which agree well with measured data, show a peak undershoot at $t^* = 1.2$ and have a RHP-zero located at 16.5 (Quach et al. 2010). This zero implies the ability to achieve control below $f = 3.1\text{ Hz}$ ($\lambda = 0.12$) or control above $f = 5.3\text{ Hz}$ ($\lambda = 0.48$), which is comparable to the region where disturbances are amplified with the pulsed-blowing wing modeled with a pure time delay. As a result, even with faster actuators, the range of frequencies of possible control is limited by the fluid dynamic response to actuation, not the bandwidth of the actuators.

The first order models with a time delay fit the measured data much better than the previous modeling with first order models, Williams, et al. (2010a). The improved modeling leads to a better agreement between experiment and theory. The range of frequencies of the current controller is increased over the range in Williams, et al. (2010a). This is on the one hand due to the incorporation of an unsteady aerodynamic model and the feedforward disturbance compensation. On the other hand the better modeling also improves the performance.

3.7. Piezoelectric actuated Wing-II

To achieve higher bandwidth from the actuators a second semi-circular wing was constructed that used zero-net-mass (piezo-electric) actuation. The wing-II model containing 8 piezo-electric actuators with 16 exit ports is shown in Figure 3.7.1. It has the same planform and dimensions as the semi-circular wing that used pulsed-blowing actuation. Since the actuation effect is based on the volume displacement of the piezo-electric oscillation, it does not need an external regulated pressure supply to operate, which reduces actuator time delays, weight, and mechanical complexity. It contains eight piezoelectric devices that operate at a resonance frequency $f = 320\text{ Hz}$. At peak velocity, the actuators can output jets of air at velocity higher than 20 m/s . Wing-II has a higher bandwidth in operating frequency than the pulsed blowing wing, referred to as wing-I. The wing-I bandwidth is limited by the response of the pressure regulator and the tubing leading compressed air into the wing, and there do not exist in the wing-II.

The lift response of the wing-I and wing-II to a single pulse input from the actuator are described in this section. Unless otherwise noted, the single-pulse time of the wing-I $\Delta t_{\text{pulse}} = 0.017$

s and of the wing-II $\Delta t_{\text{pulse}} = 0.025\text{ s}$. These are small compared to the response time of the flow.




Fig. 3.7.1. (A, left) View of the disassembled wing-I with the plenum cover plate removed. The 16 micro valve actuators can be seen positioned radially along the circular leading edge. (B, right) Bottom view of the wing-II with the 8 synthetic jet actuators.

The data in figure 3.7.2a shows the lift response of the wing-I to a single pulse input disturbance at different angles of attack varying from 12 deg to 20 deg . The free stream speed in the tunnel test section and the supplied pressure inside the wing were set at 5 m/s and 34.5 kPa (5 psi) respectively. At $\alpha = 12\text{ deg}$ and $\alpha = 14\text{ deg}$, still in the attached flow region, the single pulse input disturbance does not generate any lift increase. On the other hand, at $\alpha = 16\text{ deg}$ and above, the flow around the semi-circular wing is fully stalled, and the same disturbance creates a positive change in lift.

In figure 3.7.2b, the lift response of the wing-II to a single pulse input disturbance at different angles of attack is shown. Similar to the lift response of the wing-I, when the flow over the wing is not separated, the momentum from the ZNMF jets does not generate any gain in lift from the steady state. In addition, at the pre-stalled angles of attack, $\alpha = 14\text{ deg}$ vs. $\alpha = 12\text{ deg}$ for the wing-I and wing-II respectively, the flow above the wings is very sensitive to the actuator inputs. The single pulse disturbance in fact causes separation for a short moment, and that leads to the lift decrease.

Possibly due to small differences in the leading edge design, the ZNMF wing stalls at $\alpha = 14\text{ deg}$ and above, whereas the previous wing model stalls at $\alpha = 16\text{ deg}$ and above. The lift response at each angle of attack is different, but in general the two wings behave in a very similar manner to the single pulse input disturbance. The lift response first increases then decreases as the angle of attack moves further into the deeply stalled regime.

The change in lift coefficient of the wing-I and wing-II at $\alpha = 20\text{ deg}$ has been normalized by its maximum value and shown in figure 3.7.3. The two curves have similar trend although the response of the wing-II at $\alpha = 20\text{ deg}$ is quicker than the response of the wing-I. The non-minimum phase behavior is observed for both curves. It occurs at $t^* = 0.64$ for wing-II and $t^* =$

1.33 for wing-I. The response of wing-II also achieves its maximum value, at $t^* = 2.66$, faster than the wing-I, at $t^* = 3.01$.




Fig. 3.7.2: Lift responses of the wings-I and -II to a single pulse input disturbance at different angles of attack with 34.5 kPa (5 psi) supplied pressure and 5 m/s free stream speed. (a) Wing-I with pulsed blowing actuation. (b) Wing-II with constant input voltage to the piezo-electric actuators.




Fig. 3.7.3 Comparison of the normalized AC of the wing-I and wing-II at $\alpha = 20\text{ deg}$.

3.8. Summary

The flow physics of controlling leading edge vortices for time-varying lift enhancement on low aspect ratio wings was explored experimentally in an unsteady flow wind tunnel at HIT. The effect of actuator configuration on leading edge vortex formation was studied with rectangular and semi-circular planform wings. The ability to stabilize the LEV was determined to be most effective with the semi-circular wing, and subsequent studies used that planform.

Transient force measurements in response to pulse-like disturbances from the actuator were used to obtain linear models of the separated flow or 'plant'. Surface pressure and phase-averaged PIV measurements identified a strong vortex that convects at a relatively slow speed over the surface of the wing was formed by the actuator pulse, and is responsible for the time delay to actuation. A number of important observations resulted from these studies. 1) The fundamental fluctuating lift force frequency, $St = 0.2$, scales with the projected chord of the wing when the wing is stalled. 2) An initial lift reversal and corresponding time delay occur when the actuator issues a single pulse disturbance. 3) The lift coefficient increment scales with the square root of the actuator pressure coefficient, or equivalently the maximum actuator jet velocity normalized by the freestream speed. 4) The transient lift response to a pulse input can be used as a filter kernel to predict the response of the wing to more complex actuator input signals. This understanding allowed the development of closed-loop control architectures for enhanced maneuverability and gust suppression studies.

Closed-loop control of the leading edge vortex was used to modulate the lift force in an unsteady flow freestream that simulated a 'gusting' flow environment created by the unsteady flow wind tunnel. The limitations of a 'quasi-steady' approach to flow control were quickly realized, and efforts concentrated on understanding how flow control should be done in an unsteady flow situation. In particular, different approaches were used to increase the gust suppression bandwidth. It was determined that models for the plant and the unsteady aerodynamic response of the wing were necessary for effective control. However, control architectures based on lift as a feedback signal were bandwidth limited by the time-delay in the response to actuation as described in the preceding paragraph. The limitation on bandwidth is independent of the type of actuator used, because the time delay originates from the LEV formation and convection time. To circumvent the limitation will require different control architectures that act on a fluid dynamic time scale, possibly at the level of instabilities in the separated shear layer.

4 Simulation, modeling, and control tools

4.1 Immersed boundary method

A new formulation of the immersed boundary method with a structure algebraically identical to the traditional fractional step method is presented for incompressible flow over bodies with prescribed surface motion. Like previous methods, a boundary force is applied at the immersed surface to satisfy the no-slip constraint. This extra constraint can be added to the incompressible Navier-Stokes equations by introducing regularization and interpolation operators. The current method gives prominence to the role of the boundary force acting as a Lagrange multiplier to satisfy the no-slip condition. This role is analogous to the effect of pressure on the momentum equation to satisfy the divergence-free constraint. The current immersed boundary method removes slip and non-divergence-free components of the velocity field through a projection. The boundary force is determined implicitly without any constitutive relations allowing the present formulation to use larger CFL numbers compared to some past methods. Symmetry and positive-definiteness of the system are preserved such that the conjugate gradient method can be used to solve for the flow field. Examples show that the current formulation achieves second-order temporal accuracy and better than first-order spatial accuracy in L_2 -norms for one- and two-dimensional test problems. Results from two-dimensional simulations of flows over stationary and moving cylinders are in good agreement with those from previous experimental and numerical studies.

4.1.1 Background

Immersed boundary methods (IBMs) have gained popularity for their ability to handle moving or deforming bodies with complex surface geometry (Peskin 2002, Mittal & Iaccarino 2005). Peskin (1972) first introduced the method by describing the flow field with an Eulerian discretization and representing the immersed surface with a set of Lagrangian points. The Eulerian grid is not required to conform to the body geometry as the no-slip boundary condition is enforced at the Lagrangian points by adding appropriate boundary forces. The boundary forces that exist as singular functions along the surface in the continuous equations are described by discrete delta functions that smear (regularize) the forcing effect over the neighboring Eulerian cells.

Peskin originally used the IBM to simulate blood flow inside a heart with flexible valves, where the forcing function was computed by Hooke's law (Peskin 1972). This technique was later extended to rigid bodies by taking the spring constant to be a large value (Beyer & LeVeque 1992, Lai & Peskin 2000). Goldstein et al. (1993) applied the concept of feedback control to compute the force on the rigid immersed surface. The difference between the velocity solution and the boundary velocity is used in a proportional-integral controller. For the aforementioned techniques to model flow over rigid bodies, the choice of gain (stiffness) remains *ad hoc* and large gain results in stiff equations. Our intention is to remove all tuning parameters and formulate the IBM in a general framework for rigid bodies (as well as bodies with prescribed surface motion).

In our formulation, we treat the boundary forces in a manner analogous to the discretized pressure. For the incompressible Navier-Stokes equations, pressure may be viewed as a Lagrange multiplier required to satisfy the divergence-free constraint. Similarly, boundary forces can be regarded as Lagrange multipliers that satisfy the no-slip constraint (Glowinski et al. 1998). By introducing regularization and interpolation operators and grouping the pressure and force unknowns together, the discretized incompressible Navier-Stokes equations can in fact be formulated with a structure algebraically identical to the traditional fractional step method. Although previous research has implemented immersed boundary techniques with the tradi-

tional fractional step algorithm, the entire IBM itself has not been regarded as a fractional step (projection) method, as reported here. We follow the algebraic approach of Perot (1993), where the fractional step method is written as a block-LU decomposition.

In the next section, we review the traditional fractional step method as it is the fundamental basis for our IBM. In Section 3, we introduce the immersed boundary projection method. This formulation is compared to previous methods in Section 4; namely the original IBM (Peskin 1972), the direct forcing method (Mohd-Yousaf 1997), the immersed interface method (IIM) (Lee & LeVeque 2003), and the distributed Lagrange multiplier (DLM) method (Glowinski et al. 1998). In Section 5, numerical examples are considered to assess the temporal and spatial accuracy of the current method. Flows over stationary and moving cylinders are simulated and results are compared to previous experimental and numerical studies. Section 6 summarizes the current formulation.

4.1.2 The fractional Step Method

We consider the incompressible Navier-Stokes equations

$$\frac{\partial \mathbf{u}}{\partial t} + \mathbf{u} \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u}, \quad (4.1.1)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (4.1.2)$$

where $\mathbf{u}$, p , and Re are the suitably non-dimensionalized velocity vector, pressure, and the Reynolds number, respectively. Following references (Chorin 1968, Temam 1969, Kim & Moin 1985, Perot 1993, Chang et al. 2002), the equations are discretized with a staggered-mesh finite volume formulation using the implicit Crank-Nicolson (CN) integration for the viscous terms and the explicit second-order Adams-Bashforth (AB2) scheme for the convective terms. This produces an algebraic system of equations,

$$\begin{bmatrix} A & G \\ D & 0 \end{bmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ \phi \end{pmatrix} = \begin{pmatrix} \mathbf{r} \\ \mathbf{b} \end{pmatrix}, \quad (4.1.3)$$

where $\mathbf{u}^{n+1}$ and ϕ are the discretized velocity flux and pressure vectors. The discrete velocity can be recovered by $\mathbf{u}^{n+1} = R^{-1} \mathbf{u}^{n+1}$, where R is a diagonal matrix that transforms the discrete velocity into the velocity flux. Sub-matrices G and D correspond to the discrete gradients and divergence operators, respectively. The operator resulting from the implicit velocity term is $A = \frac{1}{2}M - \frac{1}{2}L$, where M is the (diagonal) mass matrix and L is the discrete (vector) Laplacian. We construct the Laplacian to be symmetric, hence making A symmetric as well. The right-hand side of Eq. (4.1.3) consists of the explicit terms from the momentum equation, $\mathbf{r}$, and inhomogeneous terms from the boundary condition, $\mathbf{b}_1$ and $\mathbf{b}_2$. Details on the discretization of Eqs. (4.1.1) and (4.1.2) can be found in Perot (1993) and Chang et al. (2002). It is interesting to note that $G = -D^T$ for the staggered grid formulation.

The traditional fractional step method by Chorin (1968) and Temam (1969) was introduced to solve Eq. (4.1.3) in an efficient manner by using an approximation for A^{-1} . In the present analysis, we adopt the observation made by Perot (1993) that the fractional step method can be regarded as an LU decomposition of Eq. (4.1.3):

$$\begin{bmatrix} A & 0 \\ -G^T & G^T B^N G \end{bmatrix} \begin{bmatrix} \mathbf{r} & B^N G \\ 0 & I \end{bmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ \phi \end{pmatrix} = \begin{pmatrix} \mathbf{r} \\ \mathbf{b} \end{pmatrix} + \begin{pmatrix} \mathbf{b}_1 \\ \mathbf{b}_2 \end{pmatrix} + \begin{pmatrix} -\frac{1}{2}G^T(LM^{-1})^N G \phi \\ 0 \end{pmatrix}, \quad (4.1.4)$$

where B^N is the N -th order Taylor series expansion of A^{-1} :

$$A^{-1} \approx B^N = \Delta t M^{-1} + \frac{\Delta t^2}{2} (M^{-1} L) M^{-1} + \dots - \frac{\Delta t^N}{2^{N-1}} (M^{-1} L)^{N-1} M^{-1} \\ = \sum_{j=1}^N \frac{\Delta t^j}{2^{j-1}} (M^{-1} L)^{j-1} M^{-1}. \quad (4.1.5)$$

The last term in Eq. (4.1.4) is the leading order error resulting from the truncation in B^N . Let us note that B^N is symmetric and can be made positive-definite with appropriate choices of Δt and N (Perot 1993). In the current situation, there also exists a second-order temporal discretization error from the AB2 and CN methods. As discussed in Perot (1993), the fractional step error can be absorbed by the discrete pressure if LM^{-1} and G are commutative (for example, in the case of periodic domains); otherwise there remains an N -th order error.

Equation (4.1.4) is more conveniently written in three steps:

$$A \mathbf{q}^* = \mathbf{r} + \mathbf{b}_1, \quad (\text{Solve for intermediate velocity}) \quad (4.1.6)$$

$$G^T B^N G \phi = G^T \mathbf{q}^* + \mathbf{b}_2, \quad (\text{Solve the Poisson equation}) \quad (4.1.7)$$

$$\mathbf{q}^{n+1} = \mathbf{q}^* - B^N G \phi. \quad (\text{Projection step}) \quad (4.1.8)$$

Since both A and $G^T B^N G$ are symmetric positive-definite matrices, the conjugate gradient method can be utilized to solve the above momentum and Poisson equations in an efficient manner. In general, for non-symmetric matrices, various other Krylov solvers can be employed.

Here the discrete pressure is denoted by ϕ without any superscript for its time level, as we regard pressure as a Lagrange multiplier (Chang et al. 2002). There has been extensive discussion on the exact time level of the discrete pressure variable for various treatments of pressure in fractional step methods (Strikwerda & Lee 1999, Brown et al. 2001). For the present method, ϕ is a first-order accurate approximation of pressure in time, viz. $\phi^{n+1/2}$. Since the first-order accuracy of ϕ does not affect the temporal accuracy of the velocity variable (Perot 1993), we use ϕ as a simple representation of the pressure variable. If a second-order accurate pressure is desired, Brown et al. (2001) should be referred to for further modifications to the fractional step method.

Although a detailed stability analysis is not offered in this paper, we demonstrate that the present method described in the next section can stably solve for the flow field for CFL numbers up to 1, as shown in Section 5. We mention that fractional step methods for incompressible flow can suffer numerical instability if Δt is decreased arbitrarily (Guermund & Quartapelle 1998). The time step is limited by a lower bound of $\Delta t \geq c \Delta x^{2/l+1}$ if equal orders of interpolation are used for velocity and pressure, as in the present case (c is a constant and l is the interpolation order of velocity, here $l = 2$). While remedies are offered in Guermund & Quartapelle (1998) and Codina (2001), we have not utilized them here since a much larger Δt is usually selected based on the CFL constraint.

We note in passing that the form of Eq. (4.1.3) is known as the Karush-Kuhn-Tucker (KKT) system that appears in constrained optimization problems (Nocedal & Wright 1999). This system minimizes a term similar to the kinetic energy:

$$\min_{\mathbf{u}, \phi} \left[\frac{1}{2} (\mathbf{u}^{n+1})^T A \mathbf{u}^{n+1} - (\mathbf{u}^{n+1})^T (\mathbf{r} + \mathbf{b}_1) \right] \quad \text{subject to} \quad D \mathbf{u}^{n+1} = 0 + \mathbf{b}_2. \quad (4.1.9)$$

It is interesting that the discrete pressure ϕ does not play a direct role in time advancement, but acts as a set of Lagrange multipliers to minimize the system energy and satisfy the kinematic constraint of divergence-free velocity field.

4.1.3 Immersed Boundary Projection Method

The Discretized Navier-Stokes equations with boundary force Since the discretized Navier-Stokes equations, Eq. (4.1.3), are observed to be a KKT system with pressure acting as a set of Lagrange multipliers to satisfy the continuity constraint, one can imagine appending additional algebraic constraints by increasing the number of Lagrange multipliers. Hence we incorporate the no-slip constraint from the IBM into the fractional step framework.

The IBM introduces a set of Lagrangian points, $\mathbf{x}_s$, that represent the surface, $\partial \mathcal{B}$, of an immersed object, $\mathcal{B}$, within a computational domain, $\mathcal{D}$, whose geometry need not conform to the underlying spatial grid. At the Lagrangian points, appropriate surface forces, $\mathbf{f}_s$, are applied to enforce the no-slip condition along $\partial \mathcal{B}$. Figure 4.1.3 illustrates the setup of the spatial discretization. Since the location of the Lagrangian boundary points does not necessarily coincide with the underlying spatial discretization, two operators are required: one that passes information from the boundary points to the neighboring staggered grid points and another one that conveys information in the opposite direction.




Figure 4.1.1: Staggered grid discretization of a two-dimensional computational domain $\mathcal{D}$ and immersed boundary formulation for a body $\mathcal{B}$ depicted by a shaded object. The horizontal and vertical arrows ($\rightarrow$, $\uparrow$) represent the discrete u and v velocities locations, respectively. Pressure p is positioned at the center of each cell ($\times$). Lagrangian points, $\xi_i = (x_i, y_i)$, along $\partial\mathcal{B}$ are shown with filled squares ($\blacksquare$) where boundary forces $\mathbf{f}_b = (f_{bx}, f_{by})$ are applied ($\circ$, ∇).

We consider the continuous version of the incompressible Navier-Stokes equations and explain how the IBM can be discretized into a KKT system and solved with a fractional step/projection algorithm. The incompressible Navier-Stokes equations with a boundary force, $\mathbf{f}$, and the no-slip condition can be considered as the continuous analog of the IBM

$$\frac{\partial \mathbf{u}}{\partial t} - \nabla \cdot \nabla \mathbf{u} = -\nabla p + \frac{1}{Re} \nabla^2 \mathbf{u} + \int_{\mathcal{B}} \mathbf{f}(\xi, t) \delta(\xi - \mathbf{x}) d\xi, \quad (4.1.10)$$

$$\nabla \cdot \mathbf{u} = 0, \quad (4.1.11)$$

$$\mathbf{u}(\xi, t) = \mathbf{u}_b(\xi, t) \quad \text{on } \partial\mathcal{B}, \quad (4.1.12)$$

where $\mathbf{x} \in \mathcal{D}$ and $\xi \in \partial\mathcal{B}$. The boundary $\partial\mathcal{B}$, parametrized by s , is allowed to move at a velocity $\mathbf{u}_b(\xi, t)$. Convolutions with the Dirac delta function δ are used to allow the exchange of information from $\partial\mathcal{B}$ to $\mathcal{D}$ and vice versa in Eqs. (4.1.10) and (4.1.12), respectively.

The discretization of the above system results in

$$\begin{bmatrix} A & G & -H \\ D & 0 & 0 \\ E & 0 & 0 \end{bmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ p \\ \mathbf{f} \end{pmatrix} = \begin{pmatrix} \mathbf{r}^n \\ 0 \\ \mathbf{bc}_1 \end{pmatrix} + \begin{pmatrix} \mathbf{bc}_2 \\ 0 \\ 0 \end{pmatrix}, \quad (4.1.13)$$

where $H\mathbf{f}$ corresponds to the last term in Eq. (4.1.10) with $\mathbf{f} = (f_x, f_y)^T$. Similar to the discrete pressure, we do not place a superscript for time level on $\mathbf{f}$ to emphasize its role as a Lagrange multiplier. The no-slip condition, Eq. (4.1.12), is enforced using the constraint, $E\mathbf{u}^{n+1} = \mathbf{u}_b^{n+1}$. Here A , G , and D are the implicit operator for the velocity flux, the discrete gradient operator, and the discrete divergence operator, respectively, and $\mathbf{r}^n$, $\mathbf{bc}_1$, and $\mathbf{bc}_2$ are the explicit terms in the momentum equation, the boundary condition vector

49

Similarly, the regularization operator is a discrete version of the convolution operator in Eq. (4.1.10) that passes information from the Lagrangian points, ξ_i , to the neighboring staggered grid points, $\mathbf{x}$. Defining H in a manner similar to E , we obtain

$$H\mathbf{f}_b = \beta \hat{M} d(\xi_i - \mathbf{x}_i) d(\eta_i - \mathbf{y}_i) = \frac{\beta}{\Delta x \Delta y} \hat{E}^T \mathbf{f}_b, \quad (4.1.19)$$

where β is the numerical integration factor proportional to $d\mathbf{x}$. Note that a diagonal matrix $\hat{M}$ is included for consistency with the fractional step formulation. It should be observed that E and H are symmetric up to a constant if the diagonal matrices R^{-1} and $\hat{M}$ are absent.

Next, let us achieve symmetry between the (3,1) and (1,3) block entries in the presence of R^{-1} and $\hat{M}$ in Eq. (4.1.13). We absorb the effect in scaling into the unknown boundary force by introducing a transformed forcing function $\tilde{\mathbf{f}}$ that satisfies

$$H\tilde{\mathbf{f}} = -E^T \mathbf{f}. \quad (4.1.20)$$

The original boundary force can be retrieved by $\mathbf{f} = -\text{inv}(EH)E^T \tilde{\mathbf{f}}$. In the case of using a uniform Cartesian grid with $\Delta x = \Delta y$, the relation simplifies to $\mathbf{f} = -\frac{1}{\beta} R \tilde{\mathbf{f}}$.

The discrete delta function of Eq. (4.1.14) currently requires the use of a uniform grid in the vicinity of $\partial\mathcal{B}$ to satisfy a set of properties moment conditions (Roma et al. 1999). Since the range and domain of E and H are only limited to the neighborhood of $\partial\mathcal{B}$, non-uniform discretization can still be applied away from the body. Although it is not pursued here, it would be interesting to generate discrete delta functions that are suitable for a non-uniform spatial discretization around the immersed body.

Note that symmetry between E and H is not necessary for discretization, but it allows us to solve the overall system in an efficient manner. There are unexplored possibilities using different discrete delta functions for interpolation and regularization operators. Beyer & LeVeque (1992) consider such cases in a one-dimensional model problem.

Immersed boundary method via projection Now that we have formulated the sub-matrices G and D such that $D = -G^T$ and introduced a transformed force, $\tilde{\mathbf{f}}$, the overall system of equations, Eq. (4.1.13), becomes

$$\begin{bmatrix} A & G & E^T \\ G^T & 0 & 0 \\ E & 0 & 0 \end{bmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ p \\ \tilde{\mathbf{f}} \end{pmatrix} = \begin{pmatrix} \mathbf{r}^n \\ 0 \\ \mathbf{bc}_1 \end{pmatrix} + \begin{pmatrix} \mathbf{bc}_2 \\ 0 \\ 0 \end{pmatrix}. \quad (4.1.21)$$

As previously discussed, both the discrete pressure and boundary forcing functions are Lagrange multipliers and, algebraically speaking, it is no longer necessary to make a distinction between the two. Thus organizing the sub-matrices and vectors in Eq. (4.1.21) in the following fashion:

$$Q = [G, E^T], \quad \lambda = \begin{pmatrix} p \\ \tilde{\mathbf{f}} \end{pmatrix}, \quad \mathbf{r}_1 = \mathbf{r}^n + \mathbf{bc}_1, \quad \mathbf{r}_2 = \begin{pmatrix} \mathbf{bc}_2 \\ 0 \end{pmatrix}, \quad (4.1.22)$$

Eq. (4.1.21) can be simplified to a KKT system

$$\begin{bmatrix} A & Q \\ Q^T & 0 \end{bmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{pmatrix}, \quad (4.1.23)$$

which is now in a form identical to Eq. (4.1.3), providing motivation to apply the same fractional step technique in solving the overall system as in Section 2. Performing an LU decomposition of Eq. (4.1.23),

$$\begin{bmatrix} A & 0 \\ Q^T & -Q^T B^{-1} Q \end{bmatrix} \begin{pmatrix} \mathbf{u}^{n+1} \\ \lambda \end{pmatrix} = \begin{pmatrix} \mathbf{r}_1 \\ \mathbf{r}_2 \end{pmatrix} + \begin{pmatrix} -\frac{1}{B} (LM^T)^T Q \lambda \\ 0 \end{pmatrix}, \quad (4.1.24)$$

resulting from the Laplacian operator, and the boundary condition vector generated from the divergence operator, respectively. Note that these sub-matrices and vectors (A , G , D , $\mathbf{r}^n$, $\mathbf{bc}_1$, and $\mathbf{bc}_2$) are identical to those that appear in the traditional fractional step method, Eq. (4.1.3).

The two additional sub-matrices H and E are introduced to regularize (smear) the singular boundary force over a few cells and interpolate velocity values defined on the staggered grid onto the Lagrangian points, respectively. We will refer to these sub-matrices as regularization (H) and interpolation (E) operators.

Interpolation and regularization operators The operators H and E are constructed from the regularized discrete delta function. Among the variety of discrete delta function available, we choose to use the one by Roma et al. (1999) which is specifically designed for use on staggered grids (where even-odd decoupling does not occur). This function has the form:

$$d(r) = \begin{cases} \frac{1}{\Delta x} \left[5 - 3\frac{r}{\Delta x} - \sqrt{-3\left(1 - \frac{r}{\Delta x}\right)^2 + 1} \right] & \text{for } 0.5\Delta x \leq r \leq 1.5\Delta x, \\ \frac{1}{\Delta x} \left[1 + \sqrt{-3\left(\frac{r}{\Delta x}\right)^2 + 1} \right] & \text{for } r \leq 0.5\Delta x, \\ 0 & \text{otherwise,} \end{cases} \quad (4.1.14)$$

where Δx is the cell width of the staggered grid in the x -direction. This discrete delta function is supported over only three cells, which is an advantage for computational efficiency. We have not found significant differences in the results for the current formulation with alternative discrete delta functions. References Peskin (2002), Beyer & LeVeque (1992) may be consulted for a variety of delta functions.

As observed by Peskin (1972) and Beyer & LeVeque (1992), discrete delta functions can be used both for regularization and interpolation. The interpolation operator can be derived from discretizing the convolution of $\mathbf{u}$ and δ ,

$$\mathbf{u}(\xi) = \int_{\mathcal{D}} \mathbf{u}(\mathbf{x}) \delta(\mathbf{x} - \xi) d\mathbf{x} \quad (4.1.15)$$

yielding

$$\mathbf{u}_b = \Delta x \Delta y \sum_{\mathbf{x}_i} \mathbf{u}_i d(\mathbf{x}_i - \xi_i) d(\eta_i - \eta_b) \quad (4.1.16)$$

for the two-dimensional case, where $\mathbf{u}_i$ is the discrete velocity vector defined on the staggered grid (x_i, y_i) and $\mathbf{u}_b$ is the discrete boundary velocity at the k -th Lagrangian point (ξ_k, η_k) . For the three-dimensional case an extra factor of Δz ($\eta_k - \xi_k$) is needed. Letting λ denote the factor preceding the summation, the interpolation operator for Eq. (4.1.16) can be written as:

$$\mathbf{E}_{k,i} = \lambda d(\mathbf{x}_i - \xi_k) d(\eta_i - \eta_k), \quad (4.1.17)$$

so that the no-slip condition is represented by

$$\mathbf{E}_{k,i} \mathbf{u}_i^{n+1} = \mathbf{E}_{k,i} \mathbf{u}_b^{n+1} = \mathbf{u}_{b,i}^{n+1} \quad (4.1.18)$$

where $\mathbf{E} \equiv \mathbf{E}^T R^{-1}$ to allow the use of the flux, $\mathbf{q}^{n+1} = R\mathbf{u}^{n+1}$, from the fractional-step formulation. The last is used to represent the original operator and is removed once a transformation (e.g. R^{-1}) is applied.

50

As in the original fractional step method, there is an N -th order splitting error. Note that this error cannot be absorbed by the Lagrange multiplier, λ , because LM^{-1} and Q do not commute (even for periodic domains). Hence, a third-order expansion for B^2 is recommended, as discussed in Perot (1993) and Section 5.

Thus, the immersed boundary projection method consists of the same three steps as Eqs. (4.1.6-4.1.8) but with λ replacing ϕ and Q replacing G :

$$\mathbf{q}^n = \mathbf{r}_1, \quad \text{(Solve for intermediate velocity)} \quad (4.1.25)$$

$$Q^T B^2 Q \lambda = Q^T \mathbf{q}^n - \mathbf{r}_2, \quad \text{(Solve the modified Poisson equation)} \quad (4.1.26)$$

$$\mathbf{q}^{n+1} = \mathbf{q}^n - B^2 Q \lambda, \quad \text{(Projection step)} \quad (4.1.27)$$

The main differences between the present and the traditional fractional step methods are in the Poisson equation and the projection step. Here, the pressure and boundary force are determined implicitly from the modified Poisson equation. The projection step removes the non-divergence-free and slip components of the velocity from the intermediate velocity field in one step. The numerical constraint of no-slip exists only at the Lagrangian points, hence making the dimensions of H and $\tilde{\mathbf{f}}$ considerably smaller than those of G and ϕ . Thus it is encouraging that there is no significant increase in size of $Q^T B^2 Q$ in the modified Poisson equation from $Q^T B^2 G$ in the classical fractional step method.

We can still solve Eqs. (4.1.25) and (4.1.26) with the conjugate gradient method as both left-hand side operators are symmetric and positive-definite. Some care must be taken to make $Q^T B^2 Q$ positive-definite and well-conditioned. First, as in the traditional fractional step method, one of the discrete pressure values must be pinned to a certain value to remove the zero eigenvalue.² Second, no repeating Lagrangian points are allowed to avoid $Q^T B^2 Q$ from becoming singular. Also, to achieve a reasonable condition number and to prevent penetration of streamlines caused by a lack of Lagrangian points, the distance between adjacent Lagrangian points, Δs , is set approximately to the Cartesian grid spacing.

In the case of moving immersed bodies, the location of the Lagrangian points must be updated at each time and so must E , i.e.,

$$\mathbf{E}_{k,i} = \mathbf{E}_{k,i}^{(n+1)} = E(\xi_k^{(n+1)}, \mathbf{x}_i) \quad (4.1.28)$$

and similarly for H . These operators can be pre-computed at each time step by knowing the location of the Lagrangian points *a priori*. The current technique is not limited to rigid bodies and can model flexible moving bodies if we are provided with the location of $\partial\mathcal{B}$ at time level $n+1$. For deforming bodies, the volume of the body must be isochoric to satisfy the incompressibility constraint. The current formulation treats the density of the body and the outer fluid to be equal to each other.

4.1.4 Comparison with other immersed boundary methods

Let us compare our current formulation with a few other IBMs, in particular the original IBM (Peskin 1972), the direct forcing approach (Mohd-Yousuf 1997, Fadun et al. 2000), the IIM (Lee & LeVeque 2003), and the DLM method (Glowinski et al. 1998) to clarify the fundamental differences. Since we only select a few IBMs that are most similar to the current formulation, Peskin (2002) and Mittal & Iaccarino (2005) should be consulted for additional IBMs. The same notation introduced earlier is used in this section. Because the comparison of fundamental mechanisms for satisfying the no-slip condition along the immersed boundary is

²There are alternatives to pinning the solution of the modified Poisson equation. Bouchet & Lohmeier (2000) discuss such techniques in detail for the Poisson equation with a Neumann boundary condition. Although the current staggered grid formulation does not require any explicit pressure boundary conditions, their analysis provides insight into the algebraic properties of the discretized Poisson equation.

of interest here, we consider methods for simulating both rigid and elastic bodies. Some details such as the time integration schemes, the updating algorithms for the Lagrangian points, and the constitutive relations for the boundary forces are omitted for clarity of discussion. The discrete spatial operators and the temporal treatment of the discrete pressure variable may not be identical to our version but remain conceptually similar.

The original immersed boundary method (IBM) The original IBM (Peskin 1973) is a modification to the traditional fractional step method, Eq. (4.1.5-4.1.8), to simulate flow over a flexible body. An explicit boundary force term Hf^b computed with Hooke's law is added to the right-hand side of the momentum equation.

$$Aq^* = r^* + bc_1 + Hf^b, \quad (4.1.29)$$

$$G^T B^T G\phi = G^T q^* + bc_2, \quad (4.1.30)$$

$$q^{n+1} = q^* - B^T G\phi. \quad (4.1.31)$$

At every time step, the location of the Lagrangian points on the elastic surface is updated. Although it is not considered here, a source/sink can be added to the pressure Poisson equation to apply a correction to the continuity equation (Kim et al. 2001).

Let us discuss how the original IBM may conceptually be related to our method. Hooke's law can be written as: $f = \kappa(\xi^n - \xi)$, where κ is the spring constant and ξ^n is the equilibrium position for the boundary surface. If we are to differentiate and discretize this relation, we obtain:

$$\frac{f^{n+1} - f^n}{\Delta t} = \kappa(\xi^{n+1} - \xi q^{n+1}), \quad (4.1.32)$$

using the implicit Euler time discretization. Adding the boundary force to the momentum equation, we observe that the overall system has the form:

$$\begin{bmatrix} A & G & -H \\ D & 0 & 0 \\ E & 0 & \frac{1}{\Delta t} \kappa \end{bmatrix} \begin{bmatrix} q^{n+1} \\ \phi \\ \xi^{n+1} \end{bmatrix} = \begin{bmatrix} r^* + bc_1 \\ bc_2 \\ \xi^{n+1} - \xi q^{n+1} \end{bmatrix}. \quad (4.1.33)$$

For rigid body simulations, $\kappa \gg 1$ is chosen to reduce the effect from the (3,3) sub-matrix (Beyer & LeVeque 1992, Lai & Peskin 2000). In the limit of $\kappa \rightarrow \infty$, we recover our current formulation, Eq. (4.1.13). The above formulation, Eq. (4.1.33), has a structure identical to the artificial compressibility method (Chorin 1967) that approximately satisfies the continuity equation with: $\frac{1}{\Delta t} \frac{Dp}{Dt} + \nabla \cdot u = 0$, where u is an artificial speed of sound. This artificial parameter is typically set to a large value similar to the spring constant, κ , in Eq. (4.1.33). Instead of Hooke's law, a feedback controller ($f = -\kappa_1 \int_0^t (u(\xi, \tau) - u_0(\xi, \tau)) d\tau$) with large gains ($\kappa_1 \gg 1$ and $\kappa_2 \gg 1$) has also been used to compute the boundary force (Goldstein et al. 1993), which results in an identical structure to Eq. (4.1.33).

However, large gains used in such constitutive relations add stiffness to the governing system, thus prohibiting the use of high CFL numbers. For instance, CFL numbers used in Lai & Peskin (2000) and Goldstein et al. (1993) are $O(10^{-1})$ to $O(10^{-2})$ for simulations of flow over a rigid circular cylinder. It is possible to use higher CFL numbers by lowering the gains at the expense of relaxing the no-slip condition. In contrast, the current projection method solves for the boundary force implicitly with no constitutive relations and behaves similarly to the traditional fractional step method in terms of temporal stability. Hence simulations can be performed with CFL numbers as high as 1, which is reported later in Section 5.

53

We note in passing that Linnick & Fasel (2005) recently developed a high order IBM that employs one-sided finite differences to obtain jump conditions for higher-order derivatives. Their results along with other numerical and experimental studies for flow over a stationary cylinder are compared to our results in Section 5.

The distributed Lagrange multiplier (DLM) method The most similar method to our formulation is the DLM method by Glowinski et al. (1998), who used in a variational principle (finite element) framework. Their work is closely related to ours as they introduce Lagrange multipliers (i.e. body force) on the immersed rigid body to satisfy the no-slip condition, essentially through projection. The main difference between our formulation and the DLM method lies in how the projection is applied to the velocity field.

Conceptually speaking, we consider the DLM method as a different operator splitting applied to Eq. (4.1.13). Their overall system is solved with the Marchuk-Yanenko fractional step scheme (Yanenko 1971, Marchuk 1975) that decomposes the overall operations into three operators related to: (i) the divergence-free condition and pressure, (ii) the convective and diffusive operators, and (iii) the no-slip condition and boundary force. Because the projection operators that remove the non-divergence-free and no-slip conditions are applied separately at different sub-time levels, these two constraints cannot be simultaneously satisfied by the velocity field.

In our formulation, there is only one projection step that simultaneously removes both the non-divergence-free and slip component from the velocity field. We also note that our formulation achieves second-order accuracy in time by choosing a suitable approximation for A^{-1} .

summary on the comparisons In the first three approaches, the presence of an immersed object is treated as a corrective term to account for the no-slip condition. The fundamental difference between the aforementioned methods and our formulation is the implicit treatment of both the pressure and boundary force as a single set of Lagrange multipliers in the modified Poisson equation. Once the pressure and the force are determined, the continuity equation and the no-slip condition are satisfied through a projection at the same time level in our formulation. The DLM method is found to be the most similar method but differs in how the projections are applied. Our overall IBM is viewed as a projection method to allow further generalization and numerical investigation from an algebraic point of view.

4.1.5 Validation and Verification

We numerically investigate the temporal and spatial convergence of the current method in one- and two-dimensional model problems; namely the Stokes' problem and flow inside two concentric cylinders, respectively. Also, flow over a circular cylinder is considered to validate the current method in steady-state and transient flow. At last, a moving body example of an impulsively started circular cylinder is considered.

Since the present method is a combination of the immersed boundary and the fractional step methods, we expect convergence analyses from both methods to carry over to the current formulation. The temporal accuracy of the immersed boundary projection method should follow the analysis from the fractional step algorithm as shown in Eq. (4.1.24). In all of the problems below, second-order finite volume discretization (except for H and E) is applied. For the problems of flow over a cylinder, a non-uniform grid is employed, making the scheme formally first-order accurate. However, we suppress the first-order spatial error by using a very smooth grid stretching, effectively keeping the overall error to second-order. In the vicinity of the body, the spatial grid is kept uniform with its finest resolution and $\Delta x_{\text{min}} = \Delta x_{\text{max}} = \Delta x$. Unless stated otherwise, $N = 3$ is chosen for approximating A^{-1} .

55

previous methods, it is not clear how the gains or the magnitude of the forcing function relate to how well the no-slip condition is satisfied. On the other hand, our method satisfies the continuity equation and the no-slip condition exactly to machine precision or, if desired, to a prescribed tolerance.

The direct forcing method The direct forcing method (Mohd-Yousif 1997) approximates the boundary force for rigid bodies with an intermediate velocity field q^* . The force is not actually computed but implemented directly into the momentum equation by substituting the regularized no-slip condition near the immersed boundary. Conceptually speaking, the momentum equation, Eq. (4.1.25), is modified to yield

$$(\tilde{M} - HE)Aq^* - \frac{1}{\Delta t}HEq^* = (\tilde{M} - HE)(r^* + bc_1) + \frac{1}{\Delta t}HEq^{n+1}, \quad (4.1.34)$$

$$G^T B^T G\phi = G^T q^* + bc_2, \quad (4.1.35)$$

$$q^{n+1} = q^* - B^T G\phi. \quad (4.1.36)$$

Here HE interpolates and then regularizes a vector, which acts as a filtering operator to extract the velocity field near $\partial\mathcal{B}$. A diagonal mass matrix $\tilde{M}$ is placed for scaling such that $\tilde{M} - HE = 0$ near $\partial\mathcal{B}$. Factors of $1/\Delta t$ are inserted in Eq. (4.1.34) to keep the order with respect to Δt consistent (note that $A = O(1/\Delta t)$). Conceptually, the above equation becomes $Eq^* = q^{n+1}$ near the immersed boundary and reduces to $Aq^* = (r^* + bc_1)$ away from the body. The difference between the modified momentum equation, Eq. (4.1.34), and the momentum equation from the traditional fractional step method, Eq. (4.1.6), can be expressed as the boundary force for the direct forcing method:

$$f^{n+1} = \frac{q^{n+1} - Eq^*}{\Delta t} = EAq^* - E(r^* + bc_1). \quad (4.1.37)$$

Note that this method enforces the no-slip condition on q^* but not on q^{n+1} . A projection step is applied later to project the intermediate velocity, q^* , onto the solenoidal solution space. In order to satisfy the no-slip condition exactly, iterations over the entire fractional step algorithm is required for each time level. Although slip in q^{n+1} is reported to be small (Fadus et al. 2000), the magnitude of the error cannot be estimated in a deductive manner.

The immersed interface method (IIM) Next, we consider representing the IIM (Lee & LeVeque 2003) for elastic membranes in an algebraic form. In the IIM, the boundary force is decomposed into tangential and normal components (f_t and f_n , respectively). A regularized tangential component of the force, Hf_t , is included in the momentum equation as an explicit term and the explicit normal boundary force is implemented into the pressure Poisson equation in terms of a pressure jump condition across the interface. The overall method can be described as:

$$Aq^* = r^* + bc_1 + Hf_t, \quad (4.1.38)$$

$$G^T B^T G\phi = G^T q^* + bc_2 + G^T B^T Hf_n, \quad (4.1.39)$$

$$q^{n+1} = q^* - B^T G(\phi - b). \quad (4.1.40)$$

where $b = H(\int_0^t [u] d\tau)$ is a corrective term to calculate the pressure gradient ($G\phi - b$) taking the jump condition, $[u]$, into consideration. Since the normal component of the boundary force is implemented directly into the pressure Poisson equation rather than in the momentum equation, a sharp velocity solution in the vicinity of the interface can be achieved resulting in second-order spatial convergence for some test problems. However, the construction of the correction term b requires explicit knowledge of the boundary force, and is not easily made implicit as desired in our formulation.

54




Figure 4.1.2: Setup for the one-dimensional Stokes' problem.

One-dimensional Stokes' problem We first assess the accuracy of the current method using a one-dimensional Stokes' problem where an infinitely long flat plate is impulsively set into motion with $u_{\text{plate}} = 1$ in an initially quiescent viscous fluid with $\nu = 1$ (Figure 4.1.5). The initial condition for the simulation is set to the exact solution to the Stokes' problem after a finite time of $t_0 = 0.1$ has elapsed in order to avoid the temporal discontinuity due to the impulsive start from interfering with the convergence study. Simulations are performed in a periodic computational domain in both x - and y -directions with uniform grid discretization. The top and bottom boundaries are placed far enough to avoid periodicity from interfering with the velocity profile near the translating plate. Spatial and temporal convergence is analyzed in terms of the L_∞ and L_2 norms of the horizontal velocity error, $e_1 = u(y_j) - u_j$, over the domain $y_j \in [0, 1]$ (in non-dimensional length: $y_j/\sqrt{\nu t_0} \in [0, 3.162]$).

Figure 4.1.5(a) assesses the temporal L_∞ error for various sizes of non-dimensional time steps, $\nu\Delta t/\Delta y^2$. The error was computed by comparing the solution to a temporally refined reference solution at fixed grid resolution to isolate the spatial discretization error. We calculate the error at $t = 0.11$ with $\Delta y = 10^{-3}$. The three convergence curves on the plot result from the use of different orders of expansion N for B^T (or A^{-1}). Note that the splitting error from Eq. (4.1.24) is larger in magnitude than the underlying second-order error resulting from the time integration schemes. Hence this splitting error directly influences the temporal accuracy for the range of Δt considered. As discussed in Perot (1993), the splitting error cannot be absorbed by 2 because LM^{-1} and Q are not commutative even for a periodic domain.

Next we perform simulations with a very fine time step ($\Delta t = 10^{-6}$) and compare the results to the exact solution at $t = 0.101$ for varying Δy . The velocity profile in the vicinity of the plate is influenced by the regularization of the Dirac delta function. This alters the velocity derivative at the immersed boundary causing the first-order accuracy of the L_∞ norm as shown in Figure 4.1.5(b). Fortunately, this smearing effect is dominant only in close proximity of the plate and the underlying second-order convergence is achieved in the L_2 sense.

Flow inside two concentric cylinders For a two-dimensional test problem, we simulate flow between two concentric hollow cylinders with radii $r_1 = 1/2$ and $r_2 = 1$ as well as the flow inside the smaller cylinder as shown in Figure 4.1.5. The outer cylinder is held stationary while the inner cylinder is rotated with angular velocity Ω .

$$\Omega = \frac{u_\theta(r_1)}{r_1} = 1 + \tanh\left(\frac{1-0.2}{0.05}\right), \quad (4.1.41)$$

56




Figure 4.1.3: Error norms from the one-dimensional Stokes' problem. (a) Temporal L_2 norm errors with different orders of expansion, N , for A^{-1} ; $N=1$: $\circ$, $N=2$: $\square$, and $N=3$: $\triangle$. (b) The L_2 : $\circ$ and L_2 : $\square$ spatial velocity error norms.

moving the initially quiescent fluid at $t=0$. We take a periodic computational domain of size $[-1.05, 1.05] \times [-1.05, 1.05]$ with uniform spatial resolution and compute the azimuthal velocity error, $e_\theta = u_\theta(r_j) - u_{\theta,j}$, over $r_j \in [0, r_2]$ (including flow inside the inner cylinder) reporting the L_2 and L_2 norms.

We study the impact of the splitting error from Eq. (4.1.24) on the temporal convergence by comparing our results to a reference solution obtained with a very fine time step, $\Delta t = 5 \times 10^{-6}$, and spatial resolution, $\Delta x = \Delta y = 2.1 \times 10^{-3}$. The spatial resolution is kept constant and viscosity is set to $\nu = 1$. Figure 4.1.5(a) shows that the order of expansion N for A^{-1} again influences the behavior of convergence in a fashion similar to the one-dimensional case. As it can be seen from the $N=3$ case, the second-order time integration error starts to affect the total error at the smallest shown time step. Based on both the one- and two-dimensional test problems, we recommend the use of third-order expansion N for practical problems. There also is an advantage in choosing $N=3$ for achieving positive-definiteness of the modified Poisson equation with larger choice of Δt (Pérol 1993).

Next we consider the spatial accuracy of our method at steady-state by comparing our results to the exact solution. The viscosity is reduced to $\nu = 0.01$ in order to use a fine Δx while satisfying $\nu\Delta t/\Delta x^2 \leq 1$ to keep B^* positive-definite. Figure 4.1.5(b) shows the rate of decay for the spatial errors to be 1 and about 1.5 in the L_2 and the L_2 norms. Although the first-order convergence is expected from the use of discrete delta functions, further investigation is required to explain why second-order accuracy from the underlying spatial discretization cannot be achieved in an L_2 measure.

The spatial accuracy of the pressure is also studied by comparing the current solution to the exact solution at steady-state. Because the pressure based on the current scheme only solves up to a constant (since we pin the pressure to remove the zero eigenvalue), we compare the solutions by matching the pressure at $r = 0$ for all cases and compute the error norms along the x -axis from 0 to r_1 . The infinity and L_2 error norms are plotted against the grid size in Figure 4.1.5(c) for the same problem considered in assessing the spatial accuracy of velocity. As expected, the spatial accuracy follows the same trend as the velocity shown in Figure 4.1.5(b). Due to the presence of the discrete Delta function along the immersed boundary, the pressure distribution is affected limiting the spatial accuracy to orders of one and about 1.5 for the infinity and L_2 norms, respectively.

57




Figure 4.1.4: Setup for the problem of two concentric cylinders (inner cylinder rotates with angular velocity Ω).




Figure 4.1.5: Error norms from the problem of two concentric cylinders. (a) Temporal L_2 norm errors with different orders of expansion, N , for A^{-1} ; $N=1$: $\circ$, $N=2$: $\square$, and $N=3$: $\triangle$. (b) The L_2 : $\circ$ and L_2 : $\square$ spatial velocity error norms. (c) The L_2 : $\circ$ and L_2 : $\square$ spatial pressure error norms.

58

Flow over a stationary cylinder We consider flow over a circular cylinder as another test problem because the dimensions of the recirculation zone and the force on the cylinder at various Reynolds numbers are readily available from previous experimental and numerical studies. For the numerical studies, we list results from the IBM of Lai & Pook (2000) and the IBM of Linnick & Fasel (2005) among others when the data are available. Our two-dimensional simulations are performed by introducing a cylinder of diameter $d = 1$ in a large computational domain $\mathcal{D}$ with initially uniform flow, $u = u_\infty = 1$. Reynolds numbers of $Re = u_\infty d/\nu = 20, 40$, and 200 are chosen for validating the current method at steady-state and periodic vortex shedding conditions (ν is the kinematic viscosity).

The computational domain is discretized non-uniformly in both x - and y -directions, while the grid spacing is kept uniform with its finest size ($\Delta x_{\min}$) in the vicinity of the cylinder. Table 4.1.5 summarizes the parameters used in the simulations, where n_x and n_y are the number of cells in the x - and y -directions and n_θ is the number of Lagrangian points on the surface of the cylinder with $\Delta\theta = \Delta x_{\min} = \Delta y_{\min}$. Computations are performed with different sizes of $\mathcal{D}$ to ensure that the boundary conditions along $\partial\mathcal{D}$ do not influence our solution. Left (inflow) and lateral boundary conditions along $\partial\mathcal{D}$ are set to uniform flow of $(u, v) = (u_\infty, 0)$ and are placed far away from the cylinder. At the outlet, the convective boundary condition $(\partial u/\partial t + u_\infty \partial u/\partial x = 0)$ is applied to allow vorticity to exit the domain freely. Various spatial and temporal resolutions are chosen to ensure that reliable solutions are obtained. We record the maximum CFL number ($CFL_{\max} = u_{\max} \Delta t/\Delta x_{\min}$) in Table 4.1.5 from cases of $Re = 40$ and $Re = 200$. Note that the current method yields a stable solution even with $CFL_{\max} = 0.81$.

Table 1: Parameters for spatial and temporal discretization used in the simulations. The maximum CFL numbers are reported from $Re = 40$ (*) and $Re = 200$ (†) cases.

	$n_x \times n_y$	$\mathcal{D}$	$\Delta x_{\min}$	Δt	$CFL_{\max}$	n_θ
Case A	150×150	$[-30, 30] \times [-30, 30]$	0.04	0.005	0.22*	78
Case B	300×300	$[-30, 30] \times [-30, 30]$	0.02	0.005	0.46*	157
Case C	300×300	$[-15, 45] \times [-30, 30]$	0.0333	0.0125	0.81†	94
Case D	300×300	$[-10, 10] \times [-30, 30]$	0.0333	0.0125	0.75†	94

For comparison, we compute the force on the body applied by the flow in terms of the drag and lift coefficients: $C_D = F_x/4\rho u_\infty^2 d$ and $C_L = F_y/4\rho u_\infty^2 d$, respectively, where $\rho u_\infty^2 d = 1$. The force on the cylinder, $\mathbf{F}$, can be obtained simply by

$$\mathbf{F}(t) = \begin{pmatrix} F_x(t) \\ F_y(t) \end{pmatrix} = - \int_{\mathcal{S}} \int_0^1 \eta \xi(x, r) \hat{\mathbf{b}}(\xi, x, r) - \mathbf{x} \eta dx dr \quad (4.1.42)$$

$$= - \sum_i H_{i,j} f_i \Delta x \Delta y$$

using the regularization operator and the boundary forcing function. Summation over i is implied to take place separately for each direction of the force vector.

First, simulations are performed for $Re = 20$ and 40 to validate the steady-state characteristics. The resulting wake dimensions and drag coefficients are compared to values reported in the literature. The size of the wake is characterized by l , a , b , and θ (appropriately non-dimensionalized by the diameter) defined in Figure 4.1.5 following the notation used in Cousteau & Bouard (1977a). The parameters, l , a , and b represent the length of the recirculation zone, distance from the cylinder to the center of the wake vortex, and the gap between the centers of the wake vortices, respectively. The separation angle is denoted by θ measured from the x -axis. The steady-state vorticity contours and streamlines from Case B are shown in Figure 4.1.5 for $Re = 20$ and 40. The flow profiles are in close agreement with those reported in the literature. The wake properties from Cases A and B are compared against previous experimental and numerical studies in Table 4.1.5 and are also found to be in accord.

Next, we consider flow over a cylinder at a Reynolds number of 200 to reproduce periodic vortex shedding. A short time after simulations are initiated from uniform flow, a perturbation in a form of an asymmetric body force is added to trigger the shedding instability. Numerical results replicate the periodic shedding of vortices to form the Kármán vortex street as shown in the vorticity contour of Figure 4.1.5. The resulting lift and drag coefficients and the Strouhal number, $St = f_s d/u_\infty$, where f_s is the shedding frequency, are compared to previous studies in Table 4.1.5. Results obtained from Cases B, C, and D are found to be in good agreement with previous findings.

Results from Case D compared to Cases B and C suggest that the placement of the outflow boundary is not too critical. As a pair of positive and negative vortices convect downstream, their effect on the cylinder become less important since their far-field induced velocity would appear to cancel. On the other hand, we have observed pronounced interference from the lateral boundary conditions when the height of the computational domain is shortened.

Flow around a moving cylinder As our last test problem, we simulate flow around a circular cylinder in impulsive translation to validate the present method for moving bodies. The simulation is performed

59

60

Table 2: Comparison of experimental and numerical studies of steady state wake dimensions and drag coefficient from flow over a cylinder for $Re = 20$ and 40. Experimental studies are listed with (*).

		l/d	a/d	b/d	θ	C_D
$Re = 20$	Coutanceau & Bouard (1977a)*	0.93	0.33	0.46	45.0°	-
	Tritton (1959)*	-	-	-	-	2.09
	CiteDennisJFM70	0.94	-	-	43.7°	2.05
	Linnick & Fasel (2005)	0.93	0.36	0.43	45.3°	2.06
	Present (Case A)	0.97	0.39	0.43	44.1°	2.07
	Present (Case B)	0.94	0.37	0.43	43.3°	2.06
$Re = 40$	Coutanceau & Bouard (1977a)*	2.13	0.76	0.59	53.8°	-
	Tritton (1959)*	-	-	-	-	1.59
	Dennis & Chang (1970)	2.35	-	-	53.8°	1.52
	Linnick & Fasel (2005)	2.28	0.72	0.60	53.6°	1.54
	Present (Case A)	2.33	0.75	0.60	54.1°	1.55
	Present (Case B)	2.30	0.73	0.60	53.7°	1.54

Table 3: Comparison of Strouhal number and coefficients of drag and lift for flow over cylinder from experimental and numerical studies at $Re = 200$. Experimental studies are listed with (*).

	St	C_D	C_L
$Re = 200$	Belov et al. (1995)	0.193	1.19±0.042 ±0.64
	Liu & Kawachi (1998)	0.192	1.31±0.049 ±0.69
	Lai & Perkins (2000)	0.190	-
	Roshko (1954)*	0.19	-
	Linnick & Fasel (2005)	0.197	1.34±0.044 ±0.69
	Present (Case B)	0.196	1.35±0.048 ±0.68
	Present (Case C)	0.195	1.34±0.047 ±0.68
	Present (Case D)	0.197	1.36±0.043 ±0.69




Figure 4.1.5: Definition of the characteristic dimensions of the wake structure.




Figure 4.1.8: Snapshot of the vorticity field with contour levels from -3 to 3 in increments of 0.4 for $Re = 200$.

by moving the Lagrangian body points at each time step. As these points shift their positions in time, the regularization and interpolation operators are updated according to Eq. (4.1.28). We initially position the cylinder with unit diameter ($d = 1$) at the origin and impulsively set it into motion to the left with a constant velocity of $u_0 = -1$. Results are presented for Reynolds numbers of $Re = |u_0|d/\nu = 40$ and 200.

The computational domain $\mathcal{D}$ is taken to be $[-16.5, 13.5] \times [-15, 15]$ with no-slip boundary condition applied along $\partial\mathcal{D}$. Non-uniform grid is used with uniform grid in the near field having a resolution of $\Delta_{min} = 0.02$, resulting in a grid size of 425×250 . A constant time step of $\Delta t = \Delta x_{min}/2$ is chosen such that the maximum CFL numbers are limited to 0.98 and 0.81, respectively for $Re = 40$ and 200 during the simulation from a non-dimensional time of $t^* \approx |u_0|t/d = 0$ to 3.5. Quiescent flow is used for the initial condition.

We present snapshots of the flow field at non-dimensional time of $t^* = 1, 2.5$ and 3.5 in Figure 4.1.5. Left and right figures illustrate the vorticity field for $Re = 40$ and 200, respectively. The flow fields are in agreement with those in Coutanceau & Bouard (1977b) and Koumoutsakos & Leonard (1995) for $Re = 40$. For $Re = 200$, the flow exhibits a generation of stronger vortex pair in the wake of the cylinder. In the two cases, the solutions are resolved well even near the boundary and the difference in the effect of viscous diffusion is nicely captured.

The drag coefficients for the two cases are also computed by Eq. (4.1.42) during the simulation and are plotted in Figure 4.1.5. Computational results based on vortex methods from Koumoutsakos & Leonard (1995) and Cottet et al. (2001) along with the analytical series solution (Bar-Lev & Yang 1975) valid for early time are superposed on the current results. The current scheme reveals the singular behavior of the drag at the start up time ($\mathcal{O}(1/\sqrt{t^*})$) experienced by the cylinder due to the impulsive motion (Bar-Lev & Yang 1975). Our drag coefficients are about 4 to 5% larger than those from the vortex method. Additional simulations were performed with smaller grid spacings and larger computational domains. However, there were no noticeable changes in our solutions to account for the differences.

We also measure the length of the reevaluation zone, previously defined as l/d in Figure 4.1.5, in the frame of reference of the cylinder ($\mu - u_0, \nu$) for validation over time. In Figure 4.1.5, these lengths are compared with the reported curves from a numerical study of Collins & Dennis (1973) and experimental




Figure 4.1.7: Vorticity contours (top) for steady state flow over a cylinder, where contour levels are set from -3 to 3 in increments of 0.4, and corresponding streamlines (bottom). For left and right plots, $Re = 20$ and 40, respectively.

findings of Coutanceau & Bouard (1977b) and are found to be in excellent agreement shown by the overlaps for both Reynolds numbers.

Three-dimensional examples To demonstrate that the IBPM can be implemented in three dimensions, we briefly describe results for three-dimensional flow over a low-aspect-ratio flat plate at angle of attack. As an example, a rectangular flat plate of aspect ratio, $AR = 2$, at an angle of attack of $\alpha = 30^\circ$ is instantaneously generated in a uniform flow field at $t = 0$. The Reynolds number is set to $Re = 100$ and the computational domain is taken to be $[-4, 6.1] \times [-5, 5] \times [-5, 5]$ (normalized by the chord) with a grid size of $125 \times 55 \times 80$ (streamwise, vertical, and spanwise directions, respectively). Here, grid stretching is applied to regions away from the plate, while keeping uniform resolution in the close proximity of the immersed body. The time step and the minimum grid size are set to 0.01 and 0.04, respectively, to limit the maximum Courant number to 0.5 during the simulation.

In Figure 4.1.12, the spanwise vorticity contours at the midspan are compared to digital particle image velocimetry (DPIV) measurements acquired from a companion experiment performed in an oil tow tank. Simulation results and the DPIV data are found to be in agreement along with force measurements on the plate validating the three-dimensional immersed boundary projection method. The corresponding three-dimensional wake structures are presented in Figure 4.1.13 to illustrate the formation of leading-edge, trailing-edge, and tip vortices. The isosurface here are generated for unit Q -value (second invariant of the velocity gradient tensor) to show flow regions with significant rotation³. Streamlines are also depicted to illustrate the tip-effects. Initially a strong trailing-edge vortex is formed convecting downstream while the leading-edge and tip vortices stay stably attached to the plate ($t = 1.5$). Later at steady-state ($t = 13$), the diffused leading-edge vortical structure is still stably attached to the plate. In the case of three-dimensional flow, the viscous diffusion of vorticity in the spanwise direction and the tip-effect stabilizes the wake structure at this low Re . Results with various aspect ratio, angles of attack, and planform geometries are examined in further detail in Taira et al. (2007).

4.2. Fast, multidomain algorithm

In order to accelerate our IB method, we further implement a multispace (discrete streamfunction) method that allows the divergence-free constraint to be automatically satisfied to machine roundoff. By employing a fast sine transform technique, the linear system to determine the forces can be solved efficiently with direct or iterative techniques. A multi-domain technique is developed in order to improve far-field boundary conditions that are compatible with the fast sine transform and account for the extensive potential flow induced by the body as well as vorticity that advects/diffuses to large distance from the body. The multi-domain and fast techniques are validated by comparing to the exact solutions for the potential flow induced by stationary and propagating Oseen vortices and by an impulsively started circular cylinder. Speedups of more than an order-of-magnitude are achieved with the new method.

In the next section, implement a multispace (discrete streamfunction) method (Hall 1985, Chang et al. 2002) that allows the divergence-free constraint to be automatically satisfied to machine roundoff. We show that if the grid is kept uniform throughout space (with equal spacing in all directions), the Poisson-like equation for the forces can be efficiently solved either directly for stationary bodies or iteratively for moving bodies through the use of a fast sine transform. While uniform grid spacing is in fact required

³The Q -value (the second invariant of $\nabla\mathbf{u}$) is defined as $Q = \frac{1}{2}(\|\boldsymbol{\omega}\|^2 - \|\mathbf{S}\|^2)$, for incompressible flow where $\boldsymbol{\omega}$ and $\mathbf{S}$ are the asymmetric and symmetric components of $\nabla\mathbf{u}$, respectively (Hass et al. 1988). Compared to the vorticity norm, positive Q -values can highlight vortical structures by removing regions of high shear.




Figure 4.1.9: Snapshots of the vorticity field around an impulsively moving circular cylinder for $Re = 40$ and 200 at non-dimensional time of $t^* = 1.0, 2.5$, and 3.5 . Contour levels from -3 to 3 in increments of 0.4 are chosen.




Figure 4.1.10: History of the drag coefficient of the body for $Re = 40$ and 200 (—) compared with numerical solutions from Koumoutsakos & Leonard (1995) ($Re = 40$, ---) and Conet et al. (2001) ($Re = 200$, - - -) and analytical solution by Bar-Lev & Yang (1975) (—) valid for early time.




Figure 4.1.11: Length of the recirculation zone, l/d , in the frame of reference of the moving cylinder as a function time, t^* , for (a) $Re = 40$ and (b) $Re = 200$ compared with previous studies. Present results: —; experimental measurements of Cousteau & Bouard (1977b) ($Re = 40$, ---); and numerical study of Collin & Dennis (1973) ($Re = 200$, - - -).




Figure 4.1.12: Snapshots of spanwise vorticity (ω_z) profiles along the midspan ($x = 0$) at $Re = 100$ for a rectangular flat plate of $AR = 2$ and $\alpha = 30^\circ$ based on simulations and the DPIV measurements.




Figure 4.1.13: Top view of vortical structure behind a rectangular plate of $AR = 2$ and $\alpha = 30^\circ$ represented by an iso-surface of $Q = 1$ for $Re = 100$ at different times. Streamlines are overlaid with color contour indicating the local velocity norm from blue to red in increasing magnitude. Flow direction from top left to bottom right.

in the vicinity of the body by the discrete delta function that is used to regularize the surface force, it is relatively inefficient for external flows where the domain needs to extend to large distances from the body. In the original IBPM, this difficulty is overcome by stretching the mesh away from the body, but this is incompatible with the multispace/fast sine transform formulation introduced here. To overcome this restriction, we derive in section 4.2.2 improved far-field boundary conditions that are compatible with the fast method and allow the domain to be more snug around the body. The new boundary conditions account for the extensive potential flow induced by the body as well as vorticity that advects/diffuses to large distances from the body. The boundary conditions rely on a multi-domain approach whereby the Poisson equation is solved (with the fast sine transform) on a series of increasingly larger, but coarser, computational domains. Validation examples presented in sections 4.2.3 and 4.2.4 demonstrate the efficacy and improved efficiency, respectively, of the revised formulation.

4.2.1 Nullspace method for the immersed boundary method

Nullspace approach The *nullspace or discrete streamfunction* approach (Hall 1985; Chang et al. 2002) is a method for solving the system (4.1.21) without the immersed boundary formulation. In this case, the flow only needs to satisfy the incompressibility constraint, which leads us to the use of discrete streamfunction, ψ , such that

$$\mathbf{q} = \mathbf{C}\psi, \quad (4.2.1)$$

where $\mathbf{C}$ represents the discrete curl operator. This operator is constructed with column vectors corresponding to the basis of the nullspace of $\mathbf{D}$. Chang et al. (2002) should be consulted for details. Hence, these operators enjoy the following relation

$$\mathbf{D}\mathbf{C} \equiv \mathbf{0}, \quad (4.2.2)$$

which automatically enforces incompressibility at all time; $\mathbf{D}\mathbf{q}^{n+1} = \mathbf{D}\mathbf{C}\psi^{n+1} = \mathbf{0}$. This discrete relation is consistent with the continuous version of the vector identity: $\nabla \cdot \nabla \times \mathbf{u} = 0$.

Pre-multiplying the momentum equation with $\mathbf{C}^T$, the pressure gradient term can also be removed from the formulation since $\mathbf{C}^T \mathbf{G}\mathbf{p} = -(\mathbf{D}\mathbf{C})^T \mathbf{p} = 0$, resulting in only a single equation to be solved for each time step:

$$\mathbf{C}^T \mathbf{A} \mathbf{C} \psi^{n+1} = \mathbf{C}^T (\mathbf{f}_i + \mathbf{b} \mathbf{c}_i) \quad (4.2.3)$$

In this method, the most computationally expensive component of the fractional step method, namely the pressure Poisson solver, is eliminated while the continuity equation is exactly satisfied. Moreover the fractional step error arising from using an approximate $\mathbf{A}^{-1}$ is not present since an approximate LU decomposition is not required. This feature led Chang et al. (2002) to call this technique the *exact fractional step method*.

We note that the operator $\mathbf{C}^T$ is another discrete curl operation, and that:

$$\boldsymbol{\gamma} = \mathbf{C}^T \mathbf{q}, \quad (4.2.4)$$

is a second-order-accurate approximation to the circulation in each dual cell (vorticity multiplied by the cell area normal to the vorticity component).

This method may in general be used on unstructured meshes in two and three dimensions (Chang et al. 2002), including, as a special case, the simple Cartesian mesh used in IB methods. In two dimensions, the discrete streamfunction and circulation have a single component (in the direction normal to the plane).

[†]Note that we have set $\mathbf{b} \mathbf{c}_i = 0$ which is the case for the boundary conditions we consider here. More general situations that require $\mathbf{b} \mathbf{c}_i \neq 0$ can be handled by finding a particular solution for the inhomogeneous vector and adding the solution to Eq. (4.2.1).

which is naturally defined at the cell vertices (see Figure 4.2.1), Chang et al. (2002). In three dimensions there are three components of the streamfunction and circulation that are defined at the centers of the edges of the Voronoi (dual) cell, analogously to the velocity components on the primal mesh.




Figure 4.2.1. Location of variables on staggered 3D mesh. Velocity components are defined at the center of each edge. Streamfunction and circulation are defined similarly for the Voronoi cell—in this case a cell that is offset by half a cell length in each direction.

Nullspace approach with an immersed boundary In order to satisfy both the incompressibility and the no-slip conditions with the nullspace technique, it would be necessary to derive a basis for the nullspace of Q^T . Although, a singular value decomposition of Q^T can be performed to numerically determine the nullspace, the result is not in general a sparse representation which is desirable for computational feasibility. An analytical derivation of the nullspace operator does not seem to be an easy task either. Moreover, in the general case where the body is moving, the nullspace representation would need to be recomputed at least once per timestep.

To circumvent this difficulty, we once again rely on a projection approach. Consider the system that is obtained by incorporating C^T and q^{n+1} to Eq. (4.1.21). The incompressibility constraint and the pressure variable are eliminated and we arrive at another KKT system:

$$\begin{bmatrix} C^T AC & C^T E^T \\ EC & 0 \end{bmatrix} \begin{pmatrix} q^{n+1} \\ f \end{pmatrix} = \begin{pmatrix} C^T q^n \\ 0 \end{pmatrix}. \quad (4.2.5)$$

The left-hand-side matrix is symmetric but in general indefinite, making a direct solution less efficient. The projection (fractional step) approach yields

$$C^T AC q^n = C^T q^n, \quad (4.2.6)$$

$$EC(C^T AC)^{-1}(C^T E^T)^T f = EC q^n - u_B^{n+1}, \quad (4.2.7)$$

$$q^{n+1} = q^n - (C^T AC)^{-1}(C^T E^T)^T f, \quad (4.2.8)$$

where we have as not yet inserted an approximation for the inverse of $C^T AC$. Direct solution of this system in the general case requires a nested iteration to solve the modified Poisson equation. This may be feasible

69

in general. In the case where the body is not moving, it is moreover possible to perform a Cholesky decomposition of $EC(C^T AC)^{-1}(C^T E^T)^T$ once and for all, since the dimension of the system scales with the number of degrees of freedom for the immersed boundary. In this case a system of equations of the form $C^T ACx = b$ need be solved once for each Lagrangian force at the beginning of the computation.

Fast method for uniform grid and simple boundary conditions In this section we revert to the semi-discrete momentum equation,

$$M \frac{dq}{dt} + Gp + E^T f = A^T(q) + Lq + bc, \quad (4.2.9)$$

where symbols are as defined previously. The divergence free and no-slip constraints are unchanged.

We now show that with simplification, the system of equations may be solved using fast sine transforms, resulting in a significant reduction in computational work. When the grid is uniform (with equal grid spacing in all coordinate directions), the mass matrix M is the identity matrix. We assume for the moment that the values of the velocity are known in the region outside the computational domain. We apply simple Dirichlet boundary conditions to the velocity tangent to the sides/edges of the computational domain, and a Neumann boundary condition to the velocities normal to the sides. Lacking further information, one could specify, for example, a no-penetration BC for the normal component of velocity and a zero vorticity (or no-stress) condition for the remaining tangent components. These are natural boundary conditions for an external flow around the body, provided the domain is large. In the next section we will show how improved estimates for the velocities outside the computational domain can be obtained via a multi-domain approach.

With these simplifications, we operate on Eq. (4.2.9) with C^T (which eliminates the pressure) and we obtain

$$\frac{dq}{dt} - C^T E^T f = -\beta C^T C q + C^T A^T(q) + bc_p, \quad (4.2.10)$$

In deriving this equation we have used that $Lq = -\beta C^T C q = -\beta C y$ provided that $Dq = 0$. Here β is a constant equal to $1/(R\Delta^2)$, where Δ is the uniform grid spacing. This identity mimics the continuous identity $\nabla^2 u = \nabla(\nabla \cdot u) - \nabla \times \nabla \times u = -\nabla \times \nabla \times u$.

With uniform grid and the aforementioned boundary conditions, the matrix $-\beta C^T C$ is the standard discrete Laplacian operator on a 5- or 7-point stencil in two and three spatial dimensions, respectively. The boundary conditions discussed above result in simple Dirichlet boundary conditions for y . This discrete Laplacian is diagonalized by a sine transform that can be computed in $\mathcal{O}(N \log N)$ operations (where N is the dimension of y). We denote here the sine transform pair:

$$y = S \gamma \quad \gamma = S^T y, \quad (4.2.11)$$

where the circumscribed denotes the Fourier coefficients. In writing the transform pair, we have used the fact that the sine transform can be normalized so that it is identical to its inverse. Further, we may write symbolically $A = S^T \Lambda S$ where Λ is a diagonal matrix with the eigenvalues of $C^T C$. These are positive and known analytically and we note that there is no zero eigenvalue (since the boundary conditions are Dirichlet).

Applying the same time-marching schemes used previously we obtain the transformed system:

$$\begin{aligned} S \left(I + \frac{\beta \Delta t}{2} \Lambda \right) S^T \gamma &= \left(I - \frac{\beta \Delta t}{2} \Lambda \right) \gamma \\ &+ \frac{\Delta t}{2} (3C^T A^T(q^n) - C^T A^T(q^{n+1})) \\ &+ \Delta t bc_p, \end{aligned} \quad (4.2.12)$$

70

$$EC \left(S \Lambda^{-1} \left(I + \frac{\beta \Delta t}{2} \Lambda \right)^{-1} S^T \right) (EC)^T f = EC S \Lambda^{-1} S^T q^n - u_B^{n+1}, \quad (4.2.13)$$

$$q^{n+1} = q^n - S \left(I + \frac{\beta \Delta t}{2} \Lambda \right)^{-1} S^T (EC)^T f, \quad (4.2.14)$$

The velocity, needed for the next time step, may be found by introducing the discrete streamfunction:

$$q^n = C u^n + bc_q, \quad u^n = S \Lambda^{-1} S^T q^n + bc_u. \quad (4.2.15)$$

Each of the vectors $bc_{u,q}$ involves the assumed known values of velocity at the edge of the computational domain. Their values are discussed in detail in the next section.

In the new system of equations, only one linear system need be solved, Eq. (4.2.13), with a positive definite left-hand-side operator. The matrix is positive definite can be seen by inspection. The dimensions of the matrix are now $N_f \times N_f$, and thus many fewer iterations are required than the original modified Poisson equation. To be more precise, each iteration on Eq. (4.2.13) requires $\mathcal{O}(N(2 \log_2 N + N_{bw} + 4d))$ operations, where N is the number of vorticity unknowns and N_{bw} is the bandwidth of the body-force regularization/interpolation operators³, and d is the dimensionality of the flow (2 or 3 for 2D or 3D, respectively). For the discrete Delta function with a support of 3Δ , we have $N_{bw} = 3\Delta$. For the original Poisson equation, the cost per iteration is $\mathcal{O}(N(N_{bw} + (2d+1)j) + 4d)$, where j is the order of the approximate Taylor-series inverse of A and the factor $2d+1$ is the stencil of the discrete Laplacian. Furthermore, using standard estimates for the number of iterations required for convergence of the conjugate gradient method (Shewchuk 1994) along with the known eigenvalues of $C^T C$, we can estimate that the operation count per time step for the Poisson solution has been reduced from⁴

$$\mathcal{O}(N^{1/2} N(7d + (2d+1)j)) \quad \text{operation count for original method,}$$

to

$$\mathcal{O}(N^{1/2} N(2 \log_2 N + 7d)) \quad \text{operation count for Eq. (4.2.13).}$$

For example, in a three-dimensional case with $N = 128^3$, $N_f = 10^3$, $d = 3$, and $j = 3$ the estimated speedup is about 30. For a two-dimensional case with $N = 128^2$, $N_f = 200$, $d = 2$ and $j = 3$, the speedup is about 10. This is for the Poisson solve alone. Additional speedup occurs because it is no longer necessary to solve a system $Ax = b$ for the momentum equation. Numerical experiments in Section 4.2.4 for the two-dimensional case confirm at least the order-of-magnitude of the speedup (the actual speedup is faster than predicted). Finally, we recall that the new system of equations results in no iterative error in satisfying the divergence-free constraint (it is automatically zero to round-off).

If the body is stationary, then the Poisson-like equation for the forces can be efficiently solved using a triangular Cholesky decomposition. This results in a vastly lower work per time-step, since the operation count for the Poisson solve is simply $\mathcal{O}(N^3)$. In this case the computational speed is limited only by the solution of Eq. (4.2.12).

To summarize, if the grid is uniform and simple boundary conditions are used, it is vastly preferable to solve Eqs. (4.2.12) to (4.2.14). We refer to this in what follows as the *fast method*. Unfortunately, for

³We have used the fact that $N_f \ll N$ in arriving at the estimate.

⁴Here the factors $N^{1/2}$ or $N^{2/3}$ are the estimated number of iterations of the conjugate gradient solver, the $20 \log_2 N$ factor comes from two (fast) sine transforms, the $(2d+1)j$ factor from the Laplacian, and $7d$ from the interpolation, regularization, C^T , and C operations together.

external flows, the simplified boundary conditions are not effective unless the computational domain is quite large. Since the grid is also required to be uniform, even far away from the body, the larger domain would quickly negate the benefit of fast method. However, in the next section we discuss an alternative strategy for implementing boundary conditions in the fast method that has a more modest cost penalty.

4.2.2 Far-field boundary conditions: a multi-domain approach

The fast method relies on simplified far-field boundary conditions, namely known velocity normal to the boundary and known vorticity. These can be set to zero if the computational domain is sufficiently large. For smaller domains, this will lead to significant errors and, in particular, the forces computed on the body will suffer a significant *blockage* error. The error arises from two sources. The first is the extensive, algebraically decaying potential flow induced by the body (or equivalently, the system of forces). The second is that vorticity may advect or diffuse through the boundary. In our original method discussed in Section 4.1, these errors are minimized by using a large domain with a highly stretched Cartesian mesh near the far-field boundaries (but retaining uniform grid spacing near the body), as well as by using an approximate convective outflow boundary condition. Unfortunately, stretched meshes are incompatible⁵ with direct Fourier methods for solution of the Poisson equation. In this section, we show how to pose an accurate far-field boundary condition that is also compatible with the fast method described in the last section.

We start by briefly reviewing relevant boundary conditions designed to reduce one or both of the aforementioned errors. For errors associated with the slowly decaying potential flow, a few techniques have been posed in the past to *patch* in the potential flow extending from the truncated computational boundary to infinity. Renshaw & Loh (1997) propose a technique for two unbounded directions and one periodic direction. Their method is based on matching the numerical solution to analytical representation of the solution to Laplace equation outside a cylindrical volume. They report a 50% increase per time step for a typical large-scale computation, but this cost is more than offset by the ability to use much more compact domains. Wang (1999) presents a similar approach for two-dimensional flow in the form of a correction to a trial solution that satisfies an incorrect Dirichlet boundary condition. Vortex particle methods in principle automatically account for the extensive potential flow generated by the vorticity. However, in practice it is often necessary to remove particles that advect to large distance from the region of interest. An interesting technique to reduce errors associated with removal of particles is called *merging*, whereby the circulations of several vortex particles are combined into a single element when they are sufficiently far from the body (Shields 1998, Rousi 1997).

The second type of error associated with vorticity advecting or diffusing through the boundary is typically handled by posing *outflow* boundary conditions. For incompressible flow these are usually called *convective boundary conditions*, whereas in compressible flow the term *non-reflecting boundary condition* is often used. Another technique is to selectively apply damping in a region near the computational boundary. Methods that employ this technique vary from *ad hoc* specification of layer width, damping strength, etc., to techniques that theoretically specify the damping parameters according to a model. An example is the *perfectly matched layer* (Berenger 1994) for linear wave equations (including linearized compressible Euler equations (Hu 1996)) that uses analytical solutions to the governing equations to derive damping terms that prevent reflection of waves from the interface. Another technique called *super-grid* (Colomias & Rm 2002) is based on an analogy with turbulence modeling: that the effect of the turbulence model is to model scales too fine to be resolved in the computational mesh, whereas the effect of the boundary condition is to

⁵In certain special circumstances stretched meshes can be combined with Fourier-transform methods for elliptic equations, e.g. Cao et al. (1986).

model scales too large to be resolved in the computational domain. A full discussion of these techniques is beyond the scope of this paper; we refer the reader to some recent references for further details (Sani & Ghosh 1994, Jia & Brune 1993, Olshanski & Staroverov 2000, Colonius 2004). These techniques are designed to remove vorticity from the domain as smoothly as possible thereby preventing undesirable reflections or aliasing. Most do not account for the velocity induced by vorticity that has already exited the domain (a non-local effect).

We present here an alternative technique that shares some features with these previous methods, especially those of Rammich & Loh (1997), Shiels (1998), and Colonius & Rani (2002). It is based on a multi-domain approach that also shares some operations with the subgrid method for solving elliptic equations. We first describe the method in words. The basic idea is to consider the domain as embedded in a larger domain but with a coarser mesh. The circulation on the inner (smaller, finer) mesh is then interpolated or counterfitted onto the outer (larger, coarser) mesh. The Poisson equation is solved (with zero boundary conditions) on the outer domain. This solution is then interpolated along the boundary of the inner mesh and the Poisson equation is solved, with the "corrected" boundary value specified, on the inner mesh.

Similar to the vortex merging method discussed above, any existing circulation in the outer portion of the larger domain is retained from the previous time level. In this way, we approximately account for circulation that has advected or diffused out of the inner domain. Clearly, the solution on the coarser mesh contains a larger truncation error for the evolution of this vorticity. However, inversion of the Laplacian is a smoothing operation. High frequency components of the solution induced by circulation in the outer mesh decay more rapidly than low-frequency components. Being interested in the flow in the vicinity of the body (and its wake), we discard the solutions in the outer region only retaining the velocity it induces on the inner domain.

We apply this technique recursively a number of times, enlarging (and coarsening) the domain in each grid level. We choose to keep the total number of grid points in each direction fixed on each mesh; we magnify the domain and coarsen the grid by a factor of 2 at each grid level. The procedure is shown schematically in Figure 4.2.2. The vorticity is repeatedly counterfitted on each progressive grid. The Poisson equation is then solved on the largest domain, in turn providing a boundary condition for the next smaller domain. The process is then repeated until we return to the original domain.

The velocity field decays algebraically in the far-field and we thus expect errors associated with the boundary condition on the largest domain to decrease geometrically as the size of the largest domain is increased. In the worst case of a two-dimensional flow with non-zero total circulation, the velocity decays with the inverse of the distance to the vorticity region. Analytical estimates given in Appendix B show that we obtain a factor of 4 reduction in the boundary error with each progressively larger grid. This, of course, is what would be obtained by simply extending the original grid to a distance equal to the extent of the largest grid, but due to the coarsening operation, the cost increases linearly with increasing extent, rather than quadratically (in two dimensions) or cubically (in three dimensions).

The method can thus be written as follows. We define the domain of each grid as $\mathcal{D}^{(k)}$, $k = 1, 2, \dots, N_g$, where $k = 1$ refers to the original (smallest) grid and $k = N_g$ refers to the largest one. We then define the multi-domain inverse Laplacian

$$\mathcal{L} = \mathcal{S}\mathcal{A}^{-1}\mathcal{S}^T, \quad (4.2.16)$$

where $\mathcal{S}$ is an arbitrary input vector (with length equal to the number of discrete circulation values on the grid), $\mathcal{L}$ is the solution (with length equal to the number of discrete streamfunction values), and the operator




Figure 4.2.2: Schematic of 3-level multi-domain solution of the Poisson equation.

$\mathcal{S}\mathcal{A}^{-1}\mathcal{S}^T$ implies the following operations:

$$q^{(k)} = \tilde{q}^{(k)}, \quad (4.2.17)$$

$$q^{(k)} = \begin{cases} \tilde{q}^{(k)} & \text{where } x \in \mathcal{D}^{(k)} \setminus \mathcal{D}^{(k-1)}, \\ p^{(k-1)} - (1 - \beta) (q^{(k-1)}) & \text{where } x \in \mathcal{D}^{(k-1)}, \end{cases} \quad (4.2.18)$$

$$q^{(N_g+1)} = 0, \quad k = 2, 3, \dots, N_g, \quad (4.2.19)$$

$$s^{(k)} = \mathcal{S}\mathcal{A}^{-1}\mathcal{S}^T q^{(k)} + \beta (p^{(k-1)} - q^{(k-1)}) (q^{(k-1)}), \quad (4.2.20)$$

$$s = \mathcal{S}\mathcal{A}^{-1}\mathcal{S}^T s^{(1)}, \quad k = N_g, N_g - 1, \dots, 1, \quad (4.2.21)$$

Here $p^{(k-1)} - (1 - \beta) (q^{(k-1)})$ is a fine-to-coarse interpolation operator and $p^{(k-1)} - (1 - \beta) (q^{(k-1)})$ is its coarse-to-fine counterpart restricted to $\mathcal{D}^{(k-1)}$ by β .

In constructing $\mathcal{L}$, it would be desirable to preserve (to machine roundoff) certain moments of the circulation distribution so that the velocity decay rate far from the body is correct. In the present implementation, we attempt to preserve only the total circulation. Switching from matrix/vector to point-operator notation, we write, for the two-dimensional case,

$$p^{(k-1)} - (1 - \beta) (q^{(k-1)})_{2 \times 2} = \frac{q^{(k-1)}}{4} + \frac{1}{2} \frac{q^{(k-1)}}{4} + \frac{1}{2} \frac{q^{(k-1)}}{4} + \frac{1}{2} \frac{q^{(k-1)}}{4} + \frac{1}{4} \frac{q^{(k-1)}}{4} + \frac{1}{4} \frac{q^{(k-1)}}{4} + \frac{1}{4} \frac{q^{(k-1)}}{4} + \frac{1}{4} \frac{q^{(k-1)}}{4}, \quad (4.2.22)$$

The 9-point stencil leads to a conservation of the total circulation and is second-order accurate based on a Taylor-series expansion. We note that the coefficients in Eq. (4.2.22) sum to 4 since the circulation in the (stair) cell is the vorticity multiplied by the area, and coarsening the grid by a factor of 2 results in a factor of 4 increase in cell area. The three-dimensional version of Eq. (4.2.22) consists of averaging Eq. (4.2.22) over two adjacent (i, j) planes of data normal to the vorticity component, for each of the three components.

For the coarse-to-fine interpolation at the boundary of the next-finer mesh, we use the value from the coarser mesh for those grid points that coincide, and a mid-point linear interpolation (again second-order-accurate) for those points in between.

We note that circulation is only strictly preserved if there is no vorticity advecting or diffusing out of the original domain. During vorticity transfer from fine to coarse mesh, circulation is only preserved to the level of discretization error, since the discretization error is different on each mesh and advection and diffusion rates are therefore slightly different. Tests below confirm that changes in circulation as structures pass between the different domains are appropriately small.

Utilizing the multi-domain description of the circulation and solution of the Poisson equation, we now

write the overall system of equations to be solved at each time-step.

$$\mathcal{S} \left(I + \frac{\beta \Delta t}{2} \mathcal{A} \right) \mathcal{S}^T q^{(k)} = \left(I + \frac{\beta \Delta t}{2} \mathcal{C}^T \mathcal{C} \right) q^{(k)} + \frac{\Delta t}{2} \left(\mathcal{C}^T \mathcal{A} (q^{(k-1)}) - \mathcal{C}^T \mathcal{A} (q^{(k-1)}) \right) + \frac{\Delta t}{2} \mathcal{C}^T \left(\left(p^{(k-1)} - (1 - \beta) (q^{(k-1)}) \right) - \left(p^{(k-1)} - (1 - \beta) (q^{(k-1)}) \right) \right), \quad (4.2.23)$$

$$\mathcal{E} \mathcal{C} \left(\mathcal{S}\mathcal{A}^{-1} \left(I + \frac{\beta \Delta t}{2} \mathcal{A} \right) \mathcal{S}^T \right) (\mathcal{E} \mathcal{C})^T q^{(k)} = \mathcal{E} \mathcal{C} \mathcal{S} \mathcal{A}^{-1} \mathcal{S}^T q^{(k-1)} + w^{(k-1)}, \quad (4.2.24)$$

$$q^{(k+1)} = q^{(k)} - \mathcal{S} \left(I + \frac{\beta \Delta t}{2} \mathcal{A} \right)^{-1} \mathcal{S}^T (\mathcal{E} \mathcal{C})^T q^{(k)}, \quad (4.2.25)$$

$$p^{(k+1)} = \mathcal{S}\mathcal{A}^{-1}\mathcal{S}^T q^{(k+1)}, \quad (4.2.26)$$

Note that in solving for the streamfunction at the next time step, Eq. (4.2.26), we save the counterfitted circulation fields and streamfunctions to use on the right-hand-side of Eq. (4.2.23) at the next time step.

When vorticity crosses the boundary of a given grid level, the $q^{(k)}$ fields are not necessarily smooth across the interfaces, especially at the coarsest levels. The propagation of a vortex through mesh levels is examined in the next section and it is possible to see some slight internal reflections of the local circulation near the boundary. However, the errors remain confined to a small region near the boundary and diffused over time by the physical viscosity.

The multi-domain technique comes with a significant increase in computational expense. Since we now solve the intermediate vorticity equation each Poisson equation N_g times, the operation count goes up by a factor of N_g . Nevertheless, it enables us to utilize the fast algorithm described in the previous section. Moreover, we find that the multi-domain is sufficiently accurate that computational domain can be made snug around the body. Run times for particular examples are discussed below.

We note that in many situations, it is desirable to specify a uniform flow about a body. This is simple to accomplish in the multipole formulation, as there is no circulation associated with it. One need only add the uniform flow to $q^{(k)}$ resulting from Eq. (4.2.15) and to $w^{(k-1)}$ in Eq. (4.2.24). In principle one could add any potential flow in this way, at least provided it satisfies the discrete Poisson equation with zero right-hand-side.

4.2.3 Validation examples

Velocity field for an Oseen vortex The two-dimensional velocity field associated with a Gaussian distribution of vorticity (Oseen vortex) is computed with the multi-domain boundary conditions. This test is used to validate the methodology since it is possible to derive analytically the expected improvement in multi-domain solution with increasing N_g for this case. As discussed above, the largest domain uses no penetration/no stress boundary conditions. An analytical solution for the velocity field with these boundary conditions may be constructed by the method of images such that the expected error for the multi-domain boundary conditions can be evaluated. The procedure is straightforward and is described in Appendix B. The results show that the error should decrease as 4^{-N_g} in general, and for the special case of a square domain, the rate improves to 16^{-N_g} .

The vorticity field is initialized with:

$$\omega(x, y) = \frac{\Gamma}{4\pi R^2} e^{-\frac{r^2}{4R^2}}, \quad (4.2.27)$$

where $r = \sqrt{x^2 + y^2}$ is the distance from the origin. The analytical solution for the azimuthal velocity is

$$u_\theta(x, y) = \frac{\Gamma}{2\pi R} \left(1 - e^{-\frac{r^2}{4R^2}}\right). \quad (4.2.28)$$

We start the computation at time $t = t_0$ and choose Γ and t_0 such that the maximum speed is U at $r = R$. In what follows, all lengths and velocities are normalized by R and U , respectively. The vorticity is evaluated at the vertices of a rectangular domain with uniform (and equal) grid spacing in both directions and the Poisson equation is solved using the multi-domain method discussed above. In Figure 4.2.3, contours of the velocity in the x -direction are plotted for a case with $N_\theta = 5$; the velocity computed on each of the 5 domains are overlaid to show that the velocity field remains smooth through the domain transitions. In Figure 4.2.4, the L_1 error of u (the entire discrete velocity field) is plotted in N_θ is varied from 1 to 5, for two different computational domains. For the rectangular domain extending to ± 4 and ± 8 in the x and y directions, respectively, the decay follows the 4^{-N_θ} theoretical estimate through $N_\theta = 5$. For the square domain extending to ± 5 in each direction, we observe the 16^{-N_θ} decay down to errors around 10^{-3} which can be shown to be roughly the level of the truncation error for the second-order finite-volume method at this grid density. For the non-square domain, we require about $N_\theta = 5$ to reduce the boundary condition error to a similar level.




Figure 4.2.3: Multidomain solution of the Poisson equation with $N_\theta = 5$ for an Oseen vortex. Contours of the velocity component in the x -direction are plotted for each of the 5 grids. The smallest grid extends to $\pm 5R$, with grid spacing $\Delta = 0.05R$. Contour levels: $\min = -0.2$, $\max = 0.2$, $\text{increment} = 0.02$.

Propagation of an Oseen vortex In order to evaluate errors associated with vorticity advection/diffusing through the computational boundary, we again use the analytical solution associated with an Oseen vortex.

77




Figure 4.2.4: L_1 error in u with multidomain solution of Poisson equation with increasing N_θ for the Oseen vortex. The horizontal black line shows the approximate level of the truncation error.

The vortex is initialized at $(x, y) = (0, 0)$ and advected by an otherwise uniform flow with speed equal to the maximum velocity of the vortex. The vorticity and azimuthal velocity are still given by Eqs. (4.2.27) and (4.2.28), respectively, but with the radius, r redefined with $r = \sqrt{(x - Ut)^2 + y^2}$. Again, Γ and the initial time, t_0 are set so that at $t = t_0$, the maximum speed associated with the vortex alone is U and occurs at $r = R$. Again we set $Re = 300$.

Figure 4.2.5 shows the error in the velocity at the origin for a domain that nominally extends to $\pm 5R$ with $\Delta = 0.05$. Since the velocity decays like $1/r$, it has a long-range effect. To achieve less than 1% error without corrected boundary conditions, the domain would need to extend to $\pm 100R$. The error is initially zero (even with the uncorrected boundary conditions) due to symmetry. As time progresses, the error increases and reaches 25% for $N_\theta = 1$. This occurs as the vortex propagates through the right boundary of the domain. With $N_\theta > 1$, the vortex is progressively transferred to the next largest mesh at intervals of time $5 \times 2^{N_\theta-1}$, $n = 1, \dots, N_\theta$. With $N_\theta = 5$, the error stays below 1% up to nondimensional time 80, when it leaves the coarsest, largest mesh. There are small oscillations in the error evident during grid-to-grid transfer times. The associated total circulation changes by at most 5% during these transfers. With $N_\theta = 10$, error from the boundary condition is undetectable up to time 100 and the error is controlled by the 2nd-order discretization error and stays below about 0.2%. The solution at time 100 is shown in Figure 4.2.6 on the largest mesh. The magnified region is shown as an inset and shows contours of the vorticity and normal velocity. By time 100, the vortex would have physically diffused to a core size of about $1.6R$, whereas the grid spacing on the largest domain is $12.8R$! The velocity field near the core is completely wrong, but the circulation is nearly conserved and this induces the correct potential flow far from the core. The physical (smallest) domain is also depicted on the plot and, as is shown in Figure 4.2.5, the error at the origin is still less than about one percent of the correct value at that time.

Potential flow over a cylinder As a final example, we consider the potential flow induced at $t = 0^+$ by an impulsively started cylinder of diameter D . The immersed boundary uses 571 equally spaced Lagrangian points and the domain is defined snugly around the body, extending to $\pm 0.55D$ in each direction with grid

78




Figure 4.2.5: Error in normal velocity at the origin for propagating Oseen vortex with $N_\theta = 1, 5$, and 10 .




Figure 4.2.6: Propagating Oseen vortex at $t = 100R/U$ on the largest domain with $N_\theta = 10$. Color contours represent the normal velocity. The inset shows a zoomed region near the vortex. The black lines are contours of the circulation which is represented on only a few grid points of the largest mesh at this time.

79

spacing $\Delta = 0.0055D$. We initiate a uniform flow with speed U and let the body "materialize" at $t = 0$. The solution is obtained by performing 1 time-step of the Navier-Stokes solution using the fast method with multi-domain boundary conditions. A flow field obtained with $N_\theta = 4$ is presented with the exact potential flow solution in Figure 4.2.7. The streamlines are found to be in agreement with a slight difference near the immersed boundary due to the regularized nature of the discrete delta function. In Figure 4.2.8, we compare the exact potential flow solution to the numerical solution along the top boundary of the innermost domain for different N_θ . We observe the estimated $\mathcal{O}(4^{-N_\theta})$ convergence (see Appendix B) down to a level of about 10^{-3} after which the leading-order error is dominated by the truncation error arising from the discrete delta functions at the immersed boundary and the discretization of the Poisson equation.




Figure 4.2.7: Streamlines around a circular cylinder for potential flow for $N_\theta = 4$ with solutions from the first two inner multi domains shown. Present result (---) and the exact solution (—) are presented.

4.2.4 Performance of the fast method

We conclude by measuring the performance of the fast multispace/multi-domain immersed boundary method compared to the original performance by the IBPM. First, we simulate flows over a stationary circular cylinder of diameter D and compare to previously published results (Linnick & Fasel 2005; Taira & Colonius 2007). Computations are performed on the domain $\mathcal{D}^{(1)} = [-1, 3] \times [-2, 2]$ with $\Delta = 0.02D$ where N_θ is varied between 1 and 5. The cylinder is centered at the origin. The flow is impulsively started at $t = 0$, and the body is stationary. Thus the Cholesky decomposition is used to solve Eq. (4.2.24).

After transient effects associated with the impulsively-started flow have died away, we examine wake structures and forces on the cylinders from for different values of N_θ . These are compared with previous results for $Re = 40$ and 200 in Tables 4 and 5, respectively. For the steady flow at $Re = 40$ we report characteristic dimensions of the recirculation bubble in the wake, and for the unsteady flow at $Re = 200$, we report shedding frequency and fluctuating lift and drag coefficients. Characteristic dimensions of the wake are illustrated in Figure 4.1.5. It is evident that as N_θ is increased, the fast method gives nearly identical results to the previously published data. It appears that $N_\theta = 4$ is sufficient to recover the previous results. Note that for the original IBPM, computations are performed over a domain of $[-30, 30] \times [-30, 30]$ by

80




Figure 4.2.8: Velocity error along the top boundary of the smallest domain for different N_x . A guide of $4^{-1.5}$ is also shown (---).

300 × 300 stretched grid points with the finest resolution of $\Delta x = \Delta y = 0.02$. The time step for all cases are chosen to be $\Delta t = 0.01$ to limit the maximum Courant number to 1.

In the tables, speed-up is defined as the time required to compute the last 50 time steps in the simulations normalized by the time elapsed for the original BPFM. By measuring the last 50 time steps, we give a conservative estimate for speed-up since the original method is iterative and typically requires many more iterations for earlier times. Thus with $N_x = 4$ the fast method reduces the computational time by a factor of about 15 for the steady flow and 65 for the unsteady flow. We have found similar speed-ups in a variety of problems on which we have tested the code. We note that we have thus far only implemented the fast method in two dimensions (the original algorithm has been validated in both two and three dimensions). Speed-ups for three-dimensional problems are likely to be more dramatic as discussed in Section 4.2.1.

Next, we compare the speed-up from a translating circular cylinder simulated by moving the Lagrangian boundary points. Now Eq. (4.2.24) is solved iteratively with the conjugate-gradient method. A cylinder originally at the origin at $t = 0$ is impulsively translated to the left with unit velocity with $Re = 200$. The innermost domain is selected as $\mathcal{D}^{(1)} = [-5, 5] \times [-1, 1]$ with $\Delta = 0.02D$ and we use $N_x = 4$ multi-domains. Inside this highly confined $\mathcal{D}^{(1)}$, the translating cylinder generates two counter rotating vortices in the wake as shown in Figure 4.1.5 for $t = 3.5$. The vorticity profile is in accord with previous results reported in Taira & Colonius (2007). Compared to a computation performed with the original approach, the present computation is found to be 43.4 times faster. Recall that a speed-up of 53.8 is observed for a stationary cylinder (Table 5), which suggests that the overall algorithm is still solved efficiently even with a moving immersed boundary.

4.3 Linearization and adjoint formulation

For deriving reduced-order models useful for control design using the approximate balanced truncation method outlined in section 4.4, we first linearize the Navier-Stokes equations about a pre-computed steady

		l/d	a/d	b/d	θ	C_D	Speed-up
$Re = 40$	Present ($N_x = 2$)	1.69	0.60	0.55	53.4°	1.92	25.8
	Present ($N_x = 3$)	2.01	0.67	0.58	54.0°	1.68	18.5
	Present ($N_x = 4$)	2.17	0.70	0.59	53.8°	1.55	14.2
	Present ($N_x = 5$)	2.20	0.70	0.59	53.5°	1.55	11.3
	Linnick & Fasel (2005)	2.28	0.72	0.60	53.6°	1.54	-
	Taira & Colonius (2007)	2.30	0.73	0.60	53.7°	1.54	-

Table 4: Comparison of results from the fast-method with previously reported values for steady-state flow around a cylinder at $Re = 40$.

		Δt	C_D	C_L	Speed-up
$Re = 200$	Present ($N_x = 2$)	0.206	1.47 ± 0.049	± 0.66	121.1
	Present ($N_x = 3$)	0.200	1.40 ± 0.052	± 0.70	84.7
	Present ($N_x = 4$)	0.197	1.36 ± 0.046	± 0.70	65.4
	Present ($N_x = 5$)	0.195	1.34 ± 0.045	± 0.68	53.0
	Linnick & Fasel (2005)	0.197	1.34 ± 0.044	± 0.69	-
	Taira & Colonius (2007)	0.196	1.35 ± 0.048	± 0.68	1

Table 5: Comparison of results from the fast-method with previously reported values for unsteady flow around a cylinder at $Re = 200$.

state. As described in more detail in Ahuja & Rowley (2010b), the linearized equations are:

$$\frac{d\gamma}{dt} + C^T E^T \gamma = -\beta C^T C \gamma + C^T \mathcal{A}_L(\gamma), \quad (4.3.1)$$

$$E C \gamma = 0, \quad (4.3.2)$$

where $\mathcal{A}_L(\gamma)$ is the spatial discretization of

$$\eta \times \nabla \gamma - \nabla \gamma \times \eta, \quad (4.3.3)$$

and the flux q is related to γ by $q = C(C^T C)^{-1} \gamma$. The boundary conditions for the linearized equations are $\delta \gamma = 0$ on the outer boundary of the largest computational domain.

The modeling technique of section 4.4 also requires certain adjoint simulations to approximate the observability Gramians. In order to derive the adjoint formulation of (4.3.1, 4.3.2), we define the following inner-product on the state-space:

$$(\gamma, \eta)_x = \int_{\mathcal{D}} \gamma(C^T C)^{-1} \eta dx. \quad (4.3.4)$$

That is, the inner-product defined on the state-space is the standard L^2 -inner product weighted with the inverse-Laplacian operator. It can be shown that the inner-product (4.3.4) induces the usual energy-norm, that is, $(\gamma, \gamma)_x = \int_{\mathcal{D}} \eta^2 dx$, which is the energy of the fluid integrated over the entire domain. This choice is convenient as it results in the adjoint equations which differ from the linearized equations only in the

nonlinear term. A derivation is outlined in Appendix B of Ahuja & Rowley (2010b), and the resulting equations are:

$$\frac{d\zeta}{dt} + C^T E^T \zeta = -\beta C^T C \zeta + (C^T C) \mathcal{A}_L(\zeta) \gamma, \quad (4.3.5)$$

$$E C \zeta = 0, \quad (4.3.6)$$

where the variables ζ , ξ and η are dual to the discrete circulation γ , stream function z and body force f , respectively, and $q_x = C \zeta$ is dual to the flux q . The adjoint of the linearized nonlinear term is $(C^T C) \mathcal{A}_L(\zeta) \gamma$, which can be shown to be a spatial discretization of

$$\nabla \times (\eta \times \zeta_x) - \nabla \zeta \times (\eta \times \zeta_x). \quad (4.3.7)$$

Since equation (4.3.5) differs from (4.3.1) only in the last term on the right hand side, the numerical integrator for the adjoint equations can be obtained by a small modification to the solver for the linearized equations.

The nature of this multi-domain scheme used to approximate the boundary conditions of the smallest computational domain results in a multi-domain discrete Laplacian that is not exactly self-adjoint to numerical precision. As a result, the adjoint formulation given by (4.3.5, 4.3.6) which also uses the same multi-domain approach, is not precise and results in small, rather insignificant, errors in the computation of the reduced-order models.

4.4 Balanced Proper Orthogonal Decomposition (BPOD)

Balanced Proper Orthogonal Decomposition (BPOD) is a model reduction method based on the idea of balanced truncation, which arose in the control theory community. Balanced truncation was developed by Moore (1981) and is valid for stable systems of the form

$$\begin{aligned} \dot{x} &= Ax + Bu \\ y &= Cx \end{aligned} \quad (4.4.1)$$

This method uses the concepts of controllability and observability of a system, and starts with defining the controllability and observability Gramians of the system (4.4.1) as follows:

$$W_c = \int_0^\infty e^{At} B B^T e^{A^T t} dt \quad (4.4.2)$$

$$\text{and } W_o = \int_0^\infty e^{A^T t} C^T C e^{A t} dt, \quad (4.4.3)$$

where asterisks are used to denote adjoint operators, defined by

$$(Bu, x)_x = (u, B^T x)_u, \quad \forall u \in \mathcal{U} \text{ and } \forall x \in \mathcal{X}, \quad (4.4.4)$$

$$(Cx, y)_y = (x, C^T y)_x, \quad \forall y \in \mathcal{Y} \text{ and } \forall x \in \mathcal{X}. \quad (4.4.5)$$

The Gramians (4.4.2, 4.4.3) have a nice physical interpretation. For controllability, the minimum amount of input energy required to drive the system from the origin at time $t = -\infty$ to a state x_0 at $t = 0$ is given by $\|x_0\|_{W_c}^2 = x_0^T W_c^{-1} x_0$. The states that can be reached using the least input energy are the most controllable states, and geometrically, they can be represented by the major axes of the ellipsoid $x^T W_c x = 1$, while the




Figure 4.4.1: A two-dimensional caricature of the balanced transformation of the controllability and observability Gramians.

minor axes represent the least controllable states; see Fig. 4.4.1. On the other hand, the output energy excited by the system starting at state x_0 is given by $\|x_0\|_{W_o}^2 = x_0^T W_o^{-1} x_0$. The states that excite the largest output energy are most observable and are given by the major axes of $x^T W_o x = 1$.

Balancing is referred to as a transformation of the system (4.4.1) to different coordinates in which the controllability and observability Gramians (4.4.2, 4.4.3) are equal and diagonal. It is always possible to find such a transformation if the system is both controllable and observable. Thus, in the balanced coordinates, the most controllable and the most observable states coincide and a reduced model can be obtained by simply truncating the least controllable and observable modes; see Fig. 4.4.1. If the balanced transformation is given by $x = \Phi \bar{x}$, the Gramians in the new coordinates are given by

$$\bar{W}_c = \Phi^{-1} W_c (\Phi^{-1})^T, \quad \bar{W}_o = \Phi^T W_o \Phi, \quad (4.4.6)$$

$$\text{and } \bar{W}_c = \bar{W}_o = \Sigma, \quad (4.4.7)$$

where $\Sigma \in \mathbb{R}^{n \times n}$ is a real, diagonal matrix whose entries $\sigma_i \geq 0$ decrease monotonically; they are called the Hankel singular values (HSVs) and are directly related to the controllability and observability of the corresponding states. A reduced-order model is obtained by truncating the states with relatively small HSVs, that is, the states which are almost uncontrollable and unobservable. Further, the HSVs are independent of the choice of the coordinates and are given by the eigenvalues of the product of the Gramians $W_c W_o$, while the (appropriately scaled) eigenvectors give the balancing transformation.

The main advantage of balanced truncation over the previously described methods is that it captures both the actuation and sensing. The resulting reduced model is guaranteed to be stable, provided that the truncation does not take place at an order between two equal HSVs. Also, there exist rigorous error bounds for the accuracy of the reduced model. In particular, if $G(t)$ is the input-output impulse response of (4.4.1) and $G_r(t)$ is the impulse response of the balanced system truncated at an order r , the error is given by

$$\|G(t) - G_r(t)\|_\infty \leq 2 \sum_{i=r+1}^n \sigma_i. \quad (4.4.8)$$

A disadvantage of the exact balanced truncation method is that it is not tractable for large-dimensional systems as it involves solution of large matrix Lyapunov equations to compute the Gramians; we now describe an approximate technique developed by Rowley (2005).

For systems of large dimension such as those encountered here, the Gramians (4.4.2, 4.4.3) are huge matrices which cannot be easily computed or stored. A computationally tractable procedure was introduced by Rowley (2005) for obtaining an approximate balancing transformation. The procedure relies on an approximate expression of the Gramians which was introduced by Moore (1981) and can be derived as by observing that the impulse response of (4.4.1) is given by $x(t) = \exp(A)tB$, where columns of $x(t) \in \mathbb{R}^{n \times p}$ are states obtained from the response of (4.4.1) to an impulse to the corresponding element of input u . The controllability Gramian (4.4.2) can be written in terms of this impulse response as

$$W_c = \int_0^\infty x(t)x^T(t)dt. \quad (4.4.9)$$

If the snapshots from this impulse response are sampled at equal time intervals δ and stacked in a matrix $X \in \mathbb{R}^{n \times pm}$ (after scaling),

$$X = \sqrt{\delta} \begin{bmatrix} x_1 & x_2 & \dots & x_m \\ e^{A\delta}x_1 & e^{A\delta}x_2 & \dots & e^{A\delta}x_m \end{bmatrix} \quad (4.4.10)$$

$$= \sqrt{\delta} \begin{bmatrix} e^{A\delta}B & e^{A\delta}B & \dots & e^{A\delta}B \end{bmatrix}, \quad (4.4.11)$$

where $t_i = (i-1)\delta$, the integral in (4.4.9) can be approximated by the quadrature sum

$$W_c \approx XX^T. \quad (4.4.12)$$

In general, the snapshots need not be sampled at equal time intervals, in which case each snapshot needs to be scaled differently by the appropriate quadrature factor. As pointed out by Rowley (2005), if the dataset (4.4.10) is used for computing POD modes, the resulting modes are the leading controllable modes; in that regard, POD captures the effect of actuation but not sensing.

The observability Gramian can also be approximated in a similar way. We first define the adjoint state-space system of (4.4.1):

$$\dot{z} = A^T z + C^T v, \quad (4.4.13)$$

$$w = B^T z, \quad (4.4.14)$$

where the adjoint matrices are given by (4.4.4, 4.4.5). The observability Gramian can be written in terms of the impulse response $z(t) = \exp(A^T)tC^T$ of (4.4.13), where columns of $z(t) \in \mathbb{R}^{n \times m}$ are states obtained from the response of (4.4.13) to an impulse to the corresponding element of input v :

$$W_o = \int_0^\infty z(t)z^T(t)dt. \quad (4.4.15)$$

If the snapshots of the impulse response are again sampled at equal time intervals δ and stacked in a matrix $Z \in \mathbb{R}^{n \times pm}$,

$$Z = \sqrt{\delta} \begin{bmatrix} z_1 & z_2 & \dots & z_m \\ e^{A^T\delta}z_1 & e^{A^T\delta}z_2 & \dots & e^{A^T\delta}z_m \end{bmatrix} \quad (4.4.16)$$

$$= \sqrt{\delta} \begin{bmatrix} e^{A^T\delta}C^T & e^{A^T\delta}C^T & \dots & e^{A^T\delta}C^T \end{bmatrix} \quad (4.4.17)$$

the integral in (4.4.15) can be approximated by

$$W_o \approx ZZ^T. \quad (4.4.18)$$

The approximate Gramians (4.4.12, 4.4.18) are huge dimensional and not actually computed due to the large storage cost, but the leading columns (or modes) of the transformation that balances these Gramians

85

4.5 Extensions of BPOD to unstable systems

The technique described in the previous section is valid only for stable systems. This is easy to see, as the technique involves computing the response to an impulsive input from an actuator (or, for the adjoint system, from the sensor), and for an unstable system, this response blows up. In this section, we describe how to extend the method for unstable systems. Further details may be found in Ahuja (2009) and Ahuja & Rowley (2010a).

4.5.1 Exact method

We briefly describe a model reduction procedure using the balanced truncation method for unstable systems developed by Zhou et al. (1999). The eigenvalues of A are assumed to be anywhere on the complex plane, except on the imaginary axis. For unstable systems, the integrals in (4.4.2, 4.4.3) are unbounded and hence the Gramians are ill-defined. A modified technique was proposed by Zhou et al. (1999) based on the following frequency-domain definitions of the Gramians:

$$W_c = \frac{1}{2\pi} \int_{-\infty}^{\infty} j\omega I - A \quad B B^T (-j\omega I - A^T)^{-1} d\omega, \quad (4.5.1)$$

$$W_o = \frac{1}{2\pi} \int_{-\infty}^{\infty} (-j\omega I - A^T)^{-1} C^T C(j\omega I - A)^{-1} d\omega. \quad (4.5.2)$$

By using Parseval's theorem, it can be shown that for stable systems, the frequency-domain definitions (4.5.1, 4.5.2) are equivalent to the time-domain definitions (4.4.2, 4.4.3). The model-reduction procedure of Zhou et al. (1999) begins by first transforming the system (4.4.1) to coordinates in which the stable and unstable dynamics are decoupled. That is, let $T \in \mathbb{R}^n$ be a transformation such that if $x = Tx$, the system (4.4.1) transforms to

$$\dot{x} = \frac{d}{dt} \begin{pmatrix} x_s \\ x_u \end{pmatrix} = \begin{pmatrix} A_s & 0 \\ 0 & A_u \end{pmatrix} \begin{pmatrix} x_s \\ x_u \end{pmatrix} + \begin{pmatrix} B_s \\ B_u \end{pmatrix} u \quad (4.5.3)$$

$$y = \begin{pmatrix} C_s & C_u \end{pmatrix} \begin{pmatrix} x_s \\ x_u \end{pmatrix}. \quad (4.5.4)$$

Here, A_s and A_u are such that all their eigenvalues are in the right- and left-half complex planes respectively, while B_s and B_u are the corresponding inputs. In the applications that we consider, A_u is low-dimensional with $O(10)$ states, while A_s is still very large. Next, denote the controllability and observability Gramians corresponding to the set (A_s, B_s, C_s) describing the stable dynamics by W_s^c and W_s^o respectively. Similarly, denote the Gramians corresponding to the set $(-A_u, B_u, C_u)$ by W_u^c and W_u^o . The Gramians of the original system (4.4.1) are then related to those corresponding to the two subsystems by:

$$W_c = T \begin{pmatrix} W_s^c & 0 \\ 0 & W_u^c \end{pmatrix} T^T, \quad (4.5.5)$$

$$\text{and } W_o = (T^{-1})^T \begin{pmatrix} W_s^o & 0 \\ 0 & W_u^o \end{pmatrix} T^{-1}. \quad (4.5.6)$$

A system is said to be balanced if its Gramians defined by (4.5.5, 4.5.6) are equal and diagonal, in which case the diagonal entries are called the *generalized Hankel singular values*. A reduced-order model is obtained by truncating the states with small generalized HSVs.

A physical interpretation of the Gramians (4.5.5, 4.5.6) was also given by Zhou et al. (1999) and is as follows. The sum of the minimum input energies required to drive the system from the origin at time

87

are computed using a cost-efficient algorithm similar to the method of snapshots using POD. It involves computing the singular value decomposition of

$$Z^T X = U \Sigma V^T = \begin{pmatrix} U_1 & U_2 \end{pmatrix} \begin{pmatrix} \Sigma_1 & 0 \\ 0 & \Sigma_2 \end{pmatrix} \begin{pmatrix} V_1^T \\ V_2^T \end{pmatrix}, \quad (4.4.19)$$

where $\Sigma_1 \in \mathbb{R}^{r \times r}$ is a diagonal matrix of the most significant HSVs greater than a cut-off which is a modeling parameter, while $\Sigma_2 \in \mathbb{R}^{(m-r) \times (m-r)}$ is a diagonal matrix of smaller and zero HSVs. Note that $Z^T X \in \mathbb{R}^{pm \times pm}$ is a relatively small matrix, where m_s and m_o are the number of snapshots of the impulse responses of (4.4.1) and (4.4.13), to each input, respectively. For fluid systems that we are interested in, the typical number of snapshots is $O(10^2-4)$, thus resulting in a reasonable computational cost, and typically $r \leq 100$. The leading columns and rows of the balancing transformation and its inverse are then obtained using:

$$\Phi = X V_1 \Sigma_1^{-1/2} \quad \text{and} \quad \Psi = Z U_1 \Sigma_1^{1/2}, \quad (4.4.20)$$

where $\Phi, \Psi \in \mathbb{R}^{n \times r}$, and the two sets of modes are bi-orthogonal; that is, $\Psi^T \Phi = I_r$. The columns of Φ and Ψ are called the balancing and adjoint modes respectively. The reduced-order model is then written as

$$\begin{aligned} \dot{z} &= \Psi^T A \Phi z + \Psi^T B u \\ y &= C \Phi z. \end{aligned} \quad (4.4.21)$$

Another comparison with POD was obtained by Rowley (2005): the models are the same as those obtained using the standard POD/Galerkin method if the inner product used in computing POD modes is weighted by the observability Gramian (4.4.3).

4.4.1 Output projection

When the number of outputs of the system (rows of C) is large, the algorithm described in section 4.4 can become intractable. The reason for this is that it involves one simulation of the adjoint system (4.4.13) for each component of v , the dimension of which is the same as the number of outputs. This number is often large in fluid systems where a good model needs to capture the response of the entire system to a given input ($C = I$). To overcome this problem, Rowley (2005) proposed a technique called *output projection*, in which the output y of (4.4.1) is projected onto a small number of energetically important modes obtained using POD. Let the columns of $\Theta \in \mathbb{R}^{m \times m}$ consist of the leading m POD modes of the dataset consisting of the outputs obtained from an impulse response of (4.4.1). Then, for the purpose of obtaining a reduced-order model, the output is approximated by

$$y = \Theta \Theta^T C x, \quad (4.4.22)$$

where $\Theta \Theta^T$ is an orthogonal projection of the output onto the first m POD modes. The resulting output-projected system is optimally close (in the L^2 -sense) to the original system, for an output of fixed rank m . With this approximation, only m adjoint simulations are required to approximate the observability Gramian; refer to Rowley (2005) for details. The number of POD modes m for output projection is a design parameter: for instance, one might choose this so that the first m modes capture at least 90% of the output energy. In the rest of this dissertation, the models resulting from this approximation of the output are referred to as *m-mode output projected models*.

The reduced-order model of the output-projected system is then given by

$$\dot{z} = \Psi^T A \Phi z + \Psi^T B u, \quad (4.4.23)$$

$$y = \Theta \Theta^T C \Phi z. \quad (4.4.24)$$

86

at $t = -\infty$ to a state x_0 at $t = 0$ and back to the origin at $t = \infty$ is given by $\|u\|_2^2 = x_0^T W_c^{-1} x_0$. For observability, if the system (4.4.1) is started with an initial condition x_0 and with no control input, the sum of the output energies: (a) excited on the stable subspace of A in forward time $t = (0, \infty)$, and (b) excited on the unstable subspace of A in the time interval $t = (-\infty, 0)$, is given by $\|y\|_2^2 = x_0^T W_o x_0$.

The properties of balanced truncation for stable systems described in section 4.4 extend to unstable systems as well: the reduced system is guaranteed to have no eigenvalues on the imaginary axis provided that the balanced model is not truncated between two equal generalized HSVs. Also, the error bound (4.4.8) holds for unstable systems, but with the time-domain impulse response $G(t)$ replaced by its frequency-domain counterpart $G(j\omega)$ (which is the transfer function from u to y). This is because, for unstable systems, $G(t)$ grows without bound, however, the ∞ -norm is well-defined if the transfer function $G(j\omega)$ is used.

A disadvantage of reducing an unstable system based on generalized HSVs is that an unstable mode, if it is almost uncontrollable or unobservable, might get truncated. Thus the reduced model will not capture all the instabilities, which might be undesirable for control. In the next section, we develop an approximate algorithm which differs from the approach of Zhou et al. (1999) in this respect; the proposed method treats the unstable dynamics exactly and obtains a reduced model of the stable dynamics.

4.5.2 Approximate method

The approximate balancing procedure described in section 4.4, which is essentially a snapshot-based method, does not extend to unstable systems since the impulse responses of (4.4.1) and (4.4.13) are unbounded. We could consider applying the algorithm to the two sub-systems given in (4.5.3), but the transformation T that decouples (4.4.1) itself is not available. However, when the dimension of the unstable sub-system is small, we show that it is not necessary to compute the entire transformation T and it is still possible to obtain an approximate balancing transformation. Here, we present an algorithm for computing such a transformation and also show that it essentially results in a method that is a variant of the technique of Zhou et al. (1999) presented in section 4.5.1. The idea behind the algorithm is to first project the original system (4.4.1) onto the still high-dimensional stable subspace of A . Then, one obtains a reduced-order model of the projected system using the snapshot-based procedure described in section 4.4. The dynamics projected onto the unstable subspace can be treated exactly on account of its low dimensionality.

We assume that the number of unstable eigenvalues n_u is $O(10)$ and can be computed numerically, using the computational package ARPACK developed by Lehoucq et al. (1998). We further assume that the bases for the right and the left unstable eigenspaces $\Phi_u, \Psi_u \in \mathbb{R}^{n \times n_u}$ can be computed. For the algorithm, we need the following projection operator onto the stable subspace of A :

$$P_s = I - \Phi_u \Psi_u^T, \quad (4.5.7)$$

where Φ_u and Ψ_u have been scaled such that $\Psi_u^T \Phi_u = I_{n_u}$. We use the operator P_s to obtain the dynamics of (4.4.1) restricted to the stable subspace of A as follows:

$$\dot{z}_s = P_s A z_s + P_s B u, \quad (4.5.8)$$

$$y_s = C P_s z_s, \quad (4.5.9)$$

where $z_s = P_s x \in \mathbb{R}^n$. The adjoint of (4.5.8, 4.5.9) is the same as the dynamics of (4.4.13, 4.4.14) restricted to the stable subspace of A^T using P_s^T , and is given by

$$\dot{z}_s = P_s^T A^T z_s + P_s^T C^T v, \quad (4.5.10)$$

$$w_s = B^T P_s^T z_s, \quad (4.5.11)$$

88

where $z_0 = \bar{P}_s^* z \in \mathbb{R}^n$. Then, as shown in Appendix A of Ahuja (2009), balancing the stable part of the Gramians W_s and W_u defined in (4.5.5, 4.5.6) (balancing W_s^* and W_u^*) is the same as balancing the Gramians of the stable subsystem (4.5.8, 4.5.9).

We use the procedure of section 4.4 to obtain a transformation that balances the Gramian of the stable subsystem (4.5.8, 4.5.9). We first compute the state-impulse responses of (4.5.8) and (4.5.10) and stack the resulting snapshots x_s and z_s in matrices X_s and Z_s , respectively. As in (4.4.19), we compute the singular value decomposition of $Z_s^* X_s$ and use the expressions (4.4.20) to obtain the balancing modes Φ_s and the adjoint modes Ψ_s , where again $\Psi_s^* \Phi_s = I_s$. The reduced-order models are obtained by expressing the state x as

$$x = \Phi_s a_s + \Phi_u a_u, \quad (4.5.12)$$

where $a_s \in \mathbb{R}^{n_s}$ and $a_u \in \mathbb{R}^{n_u}$. Substituting (4.5.12) in (4.4.1) and pre-multiplying by Ψ_s^* and Ψ_u^* , we obtain

$$\frac{da}{dt} = \frac{d}{dt} \begin{pmatrix} a_s \\ a_u \end{pmatrix} = \begin{pmatrix} \Psi_s^* A \Phi_s & \Psi_s^* A \Phi_u \\ \Psi_u^* A \Phi_s & \Psi_u^* A \Phi_u \end{pmatrix} \begin{pmatrix} a_s \\ a_u \end{pmatrix} + \begin{pmatrix} \Psi_s^* \\ \Psi_u^* \end{pmatrix} B u \quad (4.5.13)$$

$$y = C \begin{pmatrix} \Phi_s a_s + \Phi_u a_u \end{pmatrix} = \begin{pmatrix} C \Phi_s & C \Phi_u \end{pmatrix} \begin{pmatrix} a_s \\ a_u \end{pmatrix}. \quad (4.5.14)$$

Now, since the unstable subspace is invariant ($\text{range}(A \Phi_u) \subseteq \text{span}(\Phi_u)$), we can write $A \Phi_u = \Phi_u \Lambda$ for some $\Lambda \in \mathbb{R}^{n_u \times n_u}$, and using the properties of eigenvectors, we have $\Psi_u^* A \Phi_u = \Psi_u^* \Phi_u \Lambda = 0$. Similarly, it can be shown that $\Psi_s^* A \Phi_s = 0$. Thus, the cross terms in (4.5.13) are zero and the reduced-order model is

$$\frac{da}{dt} = \begin{pmatrix} \Psi_s^* A \Phi_s & 0 \\ 0 & \Psi_u^* A \Phi_u \end{pmatrix} \begin{pmatrix} a_s \\ a_u \end{pmatrix} + \begin{pmatrix} \Psi_s^* \\ \Psi_u^* \end{pmatrix} B u \quad (4.5.15)$$

$$y = C \begin{pmatrix} \Phi_s a_s + \Phi_u a_u \end{pmatrix} = \begin{pmatrix} \bar{C}_s & \bar{C}_u \end{pmatrix} \begin{pmatrix} a_s \\ a_u \end{pmatrix}. \quad (4.5.16)$$

The procedure described so far to obtain the reduced-order model (4.5.15, 4.5.16) is related to the procedure of Zhou et al. (1999) described in section 4.5.1. It can be shown that the transformation that balances the Gramians defined by (4.5.5, 4.5.6) results in a system in which the unstable and stable dynamics are decoupled. Furthermore, the resulting stable dynamics are the same as those given by the equations describing the a_s -dynamics of (4.5.15). The difference is that, in our algorithm, the unstable dynamics are not balanced, while they are in Zhou et al. (1999). Further, our approach does not explicitly compute the stable subsystem A_s , since it is not unstable for large systems. A disadvantage of Zhou's approach is that an unstable mode might be truncated resulting in a model which does not capture all the unstable modes, which is undesirable for control purposes.

4.5.3 Output projection for the stable subspace

For systems with a large number of outputs, the number of adjoint simulations (4.5.10) can become intractable; however, the output projection of section 4.4.1 can be readily extended to unstable systems. Instead of projecting the entire output y onto POD modes, we first express the state $x = x_u + x_s$, where $x_u = (I - \bar{P}_s)x$ and $x_s = \bar{P}_s x$ are projections on the unstable and stable subspaces of A respectively. We similarly express the output as $y = y_u + y_s = C(x_u + x_s)$. We then project the component y_s onto a small number of

When the output is the entire state or $C = I$, the entire field can be reconstructed by measuring the coefficients of the unstable modes a_u and the POD modes Θ_s of the stable subspace. That is, the output (4.5.21) can be represented as

$$y = \begin{pmatrix} \bar{C}_s & 0 \\ 0 & \bar{C}_u \end{pmatrix} \begin{pmatrix} a_s \\ a_u \end{pmatrix} \stackrel{\text{def}}{=} \bar{C} a, \quad \text{where,} \quad (4.5.25)$$

$$\bar{C}_s = I_{n_s}, \quad \bar{C}_u = \Theta_s^* \Phi_u. \quad (4.5.26)$$

Finally, if the initial state z_0 is known, the initial condition of (4.5.20) can be obtained using

$$a_0 = (\Psi_s^* \quad \Psi_u^*)^* z_0. \quad (4.5.27)$$

4.6 Extensions of BPOD to unstable limit cycles

In order to control vortex shedding, the techniques of the previous section need to be extended to periodic systems. This extension is straightforward, but somewhat complex, and is the subject of the current section. In this context, it is most convenient to consider discrete-time systems, which may be viewed as a temporal discretization of the Navier-Stokes equations.

In particular, we consider linear discrete-time periodic systems of the form

$$x(k+1) = A(k)x(k) + B(k)u(k), \quad y(k) = C(k)x(k), \quad (4.6.1)$$

with state $x \in \mathbb{C}^n$, input $u \in \mathbb{C}^p$, output $y \in \mathbb{C}^q$, and T -periodic matrix coefficients $A(\cdot), B(\cdot), C(\cdot)$. The transition matrix in (4.6.1) is $F_{j,j+1} := A(j-1)A(j-2) \cdots A(j)$ for $j > i$, where $F_{j,j} = I_{n \times n}$. Periodicity implies that the eigenvalues of $F_{j,j+1}$ are independent of j . The neutrally stable case where the spectral radius $\rho(F_{j,j+1}) = 1$ will be discussed later. For now, assume the system is exponentially stable, i.e. $\rho(F_{j,j+1}) < 1$. The controllability and observability Gramians of (4.6.1) are then well defined and are T -periodic in j (Varga 2000):

$$W_c(j) := \sum_{k=0}^{T-1} F_{j,j+1} B(k) B(k)^* F_{j,j+1}^*, \quad (4.6.2)$$

$$W_o(j) := \sum_{k=0}^{T-1} F_{j,j+1}^* C(k)^* C(k) F_{j,j+1}. \quad (4.6.3)$$

where $*$ denotes the adjoint operator.

A standard lifting procedure (Meyer & Burrus 1975) recasts (4.6.1) in T input-output (I/O) equivalent LTI form:

$$\begin{aligned} \bar{A}(j+1) &= \bar{A}_j \bar{A}_j(j) + \bar{B}_j \bar{B}_j(j); \\ \bar{y}_j(j) &= \bar{C}_j \bar{x}_j(j) + \bar{D}_j \bar{u}_j(j), \end{aligned} \quad (4.6.3)$$

with $j = 1, \dots, T$, where t is the time variable, j parameterizes the lifted systems, the state $\bar{x}_j(t) = x(j+1)T$ is periodically sampled from (4.6.1), the original inputs and outputs over each period are arranged as $\mathbb{C}^{pT}$ and $\mathbb{C}^{qT}$ column vectors $\bar{u}_j(t) = [u(j+1)T, \dots, u(j+T)T]^T$ and $\bar{y}_j(t) = [y(j+1)T, \dots, y(j+T)T]^T$, and the definitions of the constant matrices $\bar{A}_j, \bar{B}_j, \bar{C}_j$ and $\bar{D}_j$ readily follow from the variations of parameters formula in (4.6.1), e.g. $\bar{A}_j = F_{j,j+1}$. Assuming exponential stability, the controllability and observability Gramians of the j -th lifted LTI system are

$$\bar{W}_{\bar{c}} := \sum_{k=0}^{T-1} \bar{A}_j^k \bar{B}_j \bar{B}_j^* (\bar{A}_j^*)^k; \quad \bar{W}_{\bar{o}} := \sum_{k=0}^{T-1} (\bar{A}_j^*)^k \bar{C}_j^* \bar{C}_j \bar{A}_j^k. \quad (4.6.4)$$

The following statement follows from the periodicity of (4.6.1).

POD modes, of the data set consisting of the outputs from an impulse response of (4.5.8, 4.5.9). If the POD modes are represented as columns of the matrix $\Theta_s \in \mathbb{R}^{n_s \times m}$, the output of (4.5.1) is approximated by

$$y = [C(I - \bar{P}_s) + \Theta_s \Theta_s^* C \bar{P}_s] x = C_s x + \Theta_s \Theta_s^* C_s x, \quad (4.5.17)$$

Finally, with the state x expressed by the modal expansion (4.5.12), the output of the reduced-order model (4.5.15) is given by

$$y = [C \Phi_s \quad \Theta_s \Theta_s^* C \Phi_u] \begin{pmatrix} a_s \\ a_u \end{pmatrix}. \quad (4.5.18)$$

4.5.4 Algorithm

The steps involved in obtaining reduced-order models of (4.4.1), for the case with a large number of outputs, can now be summarized as follows:

1. Compute the unstable eigenvectors Φ_u and Ψ_u of the linearized and adjoint systems.
2. Project the original system (4.4.1) onto the subspace spanned by the stable eigenvectors of A in the direction of the unstable eigenvectors of A to obtain (4.5.8, 4.5.9). Compute the state and output responses from an impulse on each input of (4.5.8) and stack the state snapshots $\{x_s(i)\}$ in a matrix X_s .
3. Assemble the resulting output snapshots $\{y_s(i)\}$, and compute the POD modes Θ_s of the resulting data-set. These POD modes are stacked as columns of Θ_s .
4. Choose the number of POD modes one wants to use to describe the output (4.5.9). For instance, if 10% error is acceptable, and the first m POD modes capture 90% of the energy, then the output is the velocity field projected onto the first m modes. Thus, the output is represented as $y_s = \Theta_s^* C_s x_s$.
5. Project the adjoint system (4.4.13, 4.4.14) onto the subspace spanned by the stable eigenvectors of A^* in the direction of the unstable eigenvectors of A^* to obtain (4.5.10, 4.5.11). Compute the (state) response of (4.5.10), starting with each POD mode Θ_s as the initial condition (one simulation for each of the first m modes). Stack the snapshots $\{z_s(i)\}$ in a matrix Z_s .
6. Compute the singular value decomposition $H = Z_s X_s^* = U_s \Sigma_s V_s^*$; let U_s and V_s be the leading r columns of U_s and V_s and let $\Sigma_s \in \mathbb{R}^{r \times r}$ contain the leading rows and columns of Σ_s .
7. Define balancing modes Φ_s^j and the corresponding adjoint modes Ψ_s^j as columns of the matrices Φ_s and Ψ_s , where

$$\Phi_s^j = X_s V_s \Sigma_s^{-1/2}, \quad \Psi_s^j = U_s \Sigma_s^{-1/2}. \quad (4.5.19)$$

8. Obtain the reduced-order model using (4.5.15), which can be written as

$$\frac{da}{dt} = \begin{pmatrix} \bar{A}_s & 0 \\ 0 & \bar{A}_u \end{pmatrix} a + \begin{pmatrix} \bar{B}_s \\ \bar{B}_u \end{pmatrix} u \stackrel{\text{def}}{=} \bar{A} a + \bar{B} u, \quad (4.5.20)$$

$$y = \begin{pmatrix} \bar{C}_s & \bar{C}_u \end{pmatrix} a \stackrel{\text{def}}{=} \bar{C} a \quad \text{where,} \quad (4.5.21)$$

$$a = \begin{pmatrix} a_s \\ a_u \end{pmatrix}, \quad (4.5.22)$$

$$\bar{A}_s = \Psi_s^* A \Phi_s, \quad \bar{B}_s = \Psi_s^* B, \quad \bar{C}_s = C \Phi_s, \quad (4.5.23)$$

$$\bar{A}_u = \Psi_u^* A \Phi_u, \quad \bar{B}_u = \Psi_u^* B, \quad \bar{C}_u = \Theta_s^* C \Phi_u. \quad (4.5.24)$$

Proposition 4.1. $\bar{W}_{\bar{c}} = W_c(j)$ and $\bar{W}_{\bar{o}} = W_o(j)$ for all $j = 1, \dots, T$.

Proposition 4.1 enables us to enjoy the best of both worlds: Whereas lifting enables an appeal to LTI balanced truncation in the lifted domain, as discussed through the remainder of the paper, Gramian computations can be carried in the original periodic setting, where the dimensions of the input and output spaces are much lower: p and q instead of Tp and Tq .

4.6.1 Factorization of empirical Gramians using snapshot-based matrices

In snapshot-based methods (Lall et al. 2002, Rowley 2005), the exact Gramians are substituted by approximate empirical Gramians where the infinite series in (4.6.2) are truncated (Chablain & Van Dooren 2006, Verrier & Kailath 1983, Shokrollahi et al. 1983) at a finite $m < \infty$:

$$\begin{aligned} \bar{W}_{\bar{c}}(j; m) &:= \sum_{k=0}^{m-1} F_{j,j+1} B(k) B(k)^* F_{j,j+1}^*, \\ \bar{W}_{\bar{o}}(j; m) &:= \sum_{k=0}^{m-1} F_{j,j+1}^* C(k)^* C(k) F_{j,j+1}. \end{aligned} \quad (4.6.5)$$

When the system is exponentially stable, truncation is justified by an induced norm bound on the truncation error, obtained by a geometric series argument and an appeal to Proposition 4.1:

Lemma 4.2. Assume that the linear periodic system (4.6.1) is exponentially stable and let m be an integer multiple of the period, $m = IT$. Then the following induced norm error bounds hold:

$$\begin{aligned} \frac{\|\bar{W}_{\bar{c}}(j) - \bar{W}_{\bar{c}}(j; m)\|}{\|\bar{W}_{\bar{c}}(j)\|} &\leq \|\bar{F}_{j,j+1}^m\|_2^2, \\ \frac{\|\bar{W}_{\bar{o}}(j) - \bar{W}_{\bar{o}}(j; m)\|}{\|\bar{W}_{\bar{o}}(j)\|} &\leq \|\bar{F}_{j,j+1}^m\|_2^2. \end{aligned} \quad (4.6.6)$$

Empirical Gramians can be factorized using snapshot-based matrices.

Proposition 4.3. Let $B^{(j)}, j = 1, \dots, p$, denote the j -th column of B , and let $X^{(j)} \in \mathbb{C}^{n \times m}$ be defined as

$$X^{(j)}(j; m) := [F_{j,j+1} B^{(j)}(j), \dots, F_{j,j+1} B^{(j)}(j-m)],$$

for each $j = 1, \dots, T$ and the horizon length m . Finally, define the matrix of snapshots

$$X(j; m) := [X^{(1)}(j; m), \dots, X^{(p)}(j; m)] \in \mathbb{C}^{n \times mp}.$$

Then $\bar{W}_{\bar{c}}(j; m) = X(j; m) X(j; m)^*$.

As illustrated in Figure 4.6.1(a), the columns of $X(j; m)$ are snapshots of impulse-response simulations of the system (4.6.1), justifying the term empirical Gramian. Invoking the T -periodicity of $B(\cdot)$ and $F(\cdot, \cdot)$ (e.g. $F_{j,j+1-m+T+1} = F_{j,j+1-m+1}$), one observes that the m columns of $X^{(j)}(m; j)$ are samples at times $j - mT, j - (m-1)T, \dots, j - T$ of trajectories of simulations initiated at $x(j-m+1)T = B^{(j)}(j-m+1)$, $j = 1, \dots, T$, assuming $m = IT$. In total, Tp simulations and mp snapshots are needed to construct $X(j; m)$.

An analogous observation applies to the empirical observability Gramian.




Figure 4.6.1: (a) The T impulse-response simulations corresponding to the i -th control input. (b) The T adjoint impulse-response simulations corresponding to the i -th adjoint control input.

Proposition 4.4. Let $C^{(i)} = 1, \dots, q$ denote the i -th row of C , and let $Y^{(i)} \in \mathbb{C}^{n \times m}$ be defined as

$$Y^{(i)}(j, m) = \begin{bmatrix} F_{j, m-1}^{(i)} C^{(i)}(j-m-1) \\ \vdots \\ F_{j, m-2j}^{(i)} C^{(i)}(j-m-2) \end{bmatrix}$$

for each $j = 1, \dots, T$ and the horizon length m finally, let

$$Y(j, m) = [Y^{(1)}(j, m), \dots, Y^{(q)}(j, m)] \in \mathbb{C}^{n \times mq}.$$

Then $W_{\text{adj}}(j, m) = Y(j, m)Y(j, m)^*$.

As illustrated in Figure 4.6.1(b), $T(j, m)$ can be obtained from simulations of the adjoint system

$$x(k+1) = A(k)x(k) + \hat{C}(k)u(k) \quad (4.6.7)$$

where $k = j, \dots, j+m-1$, $z \in \mathbb{C}^n$, $v \in \mathbb{C}^q$, $\hat{A}(k) := A(2j+m-k-1)^*$ and $\hat{C}(k) := C(2j+m-k-1)^*$. By periodicity, T adjoint simulations, and in total mq snapshots taken at time $j = kT$, $k = 1, \dots, l$ are needed to construct $Y(j, m)$.

93

The I/O map of the j -th lifted LTI system (4.6.3) is determined by the $Tq \times Tp$ dimensional impulse-response matrices $\{G_i(t)\}$. The output-projected lifted system

$$\begin{aligned} \hat{z}_j(t+1) &= \hat{A}_j \hat{z}_j(t) + \hat{B}_j \hat{u}_j(t); \\ \hat{y}_j(t) &= \hat{C}_j \hat{z}_j(t) + \hat{D}_j \hat{u}_j(t); \end{aligned} \quad (4.6.10)$$

is designed to best approximate the exact impulse responses of the original lifted system. Ideally, the low-rank orthogonal projection matrix $\hat{P}_j$ should thus satisfy

$$\hat{P}_j = \underset{\{P \in \mathbb{R}^{n \times n}\}}{\text{argmin}} \left(\sum_{i=0}^{\infty} \|G_i(t) - \hat{P}_j G_i(t)\|^2 \right), \quad (4.6.11)$$

where $\mathbb{R}^{n \times n}$ is the space of orthogonal projections of rank $r_{\text{app}} \ll Tq$. When the Frobenius norm $\|\cdot\|_F$ is used in (4.6.11), it becomes a standard projection problem. Its solution is $\hat{P}_j = \hat{\Theta}_j \hat{\Theta}_j^*$, where the columns of $\hat{\Theta}_j$ are the leading r_{app} POD modes of the dataset $\{G_i(t)\}_{i=0}^{\infty}$.

As described above, the optimal $\hat{P}_j$ is generically a full matrix. Thus, $\hat{P}_j(t) = \hat{P}_j \hat{P}_j(t)$ is no longer the lifted representation of the output of a periodic system, and the projected system cannot be "un-lifted". Rather, for each t , the value of $\hat{P}_j(t)$ is determined by the original response along an entire period. In particular, we lose the ability to compute the Gramian in the original periodic setting. To avoid this problem we impose on (4.6.11) the additional condition that the projection has a block diagonal form

$$\hat{P}_j = \text{diag}\{\hat{P}_j(1), \dots, \hat{P}_j(T)\} \quad (4.6.12)$$

where each $q \times q$ diagonal block is a rank- r_{app} orthogonal projection with $r_{\text{app}} = r_{\text{app}} T$. This enables to un-lift the projected lifted system (4.6.10) to an output-projected time-periodic system

$$\begin{aligned} x(k+1) &= A(k)x(k) + B(k)u(k); \\ y(k) &= P(k)C(k)x(k), \end{aligned} \quad (4.6.13)$$

where the T -periodic, rank- r_{app} orthogonal projection P is defined by $P(i+T-i) = P(i) = \hat{P}_j(i+1)$, $i = 0, \dots, T-1$. The constrained optimization problem (4.6.11)-(4.6.12) is solved as an equivalent set of unconstrained problems in the periodic setting. Invoking the correspondence of the $T, q \times pT$ dimensional blocks of $G_i(t)$, $G(i+T-i, j)$, $i = 0, \dots, T-1$, to the impulse response of (4.6.1), as detailed in Barnish & Pearson (1992):

Proposition 4.5. Using the Frobenius norm, the solution of the constrained optimization problem (4.6.11) and (4.6.12) is equivalent to the combined solution of the problems

$$\hat{P}_j(i+1) = \underset{\{P \in \mathbb{R}^{n \times n}\}}{\text{argmin}} \left(\sum_{j=0}^{\infty} \|G(j+T-i, j) - P(i+1)G(j+T-i, j)\|_F^2 \right),$$

for $i = 0, \dots, T-1$.

Proof. By a reduction to a standard projection problem. $\square$

95

4.6.2 Balanced truncation using the method of snapshots

Fix a time $1 \leq j \leq T$. Justified by Propositions 4.3 and 4.4 and Lemma 4.2, let $X(j, m_e)$ and $Y(j, m_e)$ be computed, allowing $m_e \neq m_o$ as factors of the empirical Gramians $W_{xx}(j, m_e)$, $W_{yy}(j, m_e)$. By Proposition 4.1, they can be also used as factors of the empirical Gramians of the j -th lifted system (4.6.3). The method of snapshots presented in Rowley (2005) then leads to approximate balanced truncations in the lifted LTI setting, as follows: Compute the SVD $Y(j, m_e)^* X(j, m_e) = U \Sigma V^*$, and the transformations Φ, Ψ that exactly balance the empirical Gramians of the lifted system

$$\Phi = X(j, m_e) V \Sigma^{-1/2}, \quad \Psi = Y(j, m_e) U \Sigma^{-1/2} \quad (4.6.8)$$

Let Φ_r, Ψ_r be the first r columns of Φ and Ψ , comprising the leading bi-orthogonal balancing and adjoint modes of the j -th lifted system. (Note that to simplify notation, the dependence of $\Phi, \Psi, \Phi_r, \Psi_r$ on j is suppressed.) The reduced state $\hat{z}_j(t) \in \mathbb{C}^r$ is defined by the projection $\hat{z}_j(t) = \Psi_r^* z_j(t) = \Psi_r^* x(j+T-i)$ and the estimated full state $\hat{x}_j(t) = \Phi_r \hat{z}_j(t)$. The reduced model, of order r , in the lifted setting, reads

$$\begin{aligned} \hat{z}_j(t+1) &= \hat{A}_r \hat{z}_j(t) + \hat{B}_r \hat{u}_j(t); \\ \hat{y}_j(t) &= \hat{C}_r \hat{z}_j(t) + \hat{D}_r \hat{u}_j(t). \end{aligned} \quad (4.6.9)$$

I/O equivalence of (4.6.1) to the lifted (4.6.3) means that the reduced-order system provides the sought I/O approximation of (4.6.1). Note that improved numerical stability of the computations above can be achieved by first representing each of the factors $X(j, m_e)$ and $Y(j, m_e)$ in terms of leading orthogonal bases, obtained, e.g., by SVD or by Krylov methods.

We comment in closing on the possibility to "un-lift" the reduced-order lifted system. As discussed in Varga (2000), the exact Gramians solve an allied periodic Lyapunov equation, thus providing an exact periodic balancing and an "un-lifted" balanced truncation in the periodic setting. Using the method of snapshots, there are two computational shortcomings to that approach in the current problem. First, the computational burden is high when $T \gg 1$. Second, the truncated empirical Gramians used here do not form an exact solution of the periodic Lyapunov equation. Un-lifting is nonetheless a simple task if the balancing requirement is limited to the periodically sampled system (i.e., to a lifted system for one, fixed j). The following inductive procedure is one possible solution: Fix $\Phi(j) := \Phi_r$ and $\Psi(j+T-1) := \Psi_r$. Let $P(j) \in \mathbb{R}^{n \times n}$ be the rank- r orthogonal projection on $\mathcal{R}(\Phi(j+T-i))$ and let $\Phi(j+1) := \Psi_r^* x(j+1) \in \mathbb{C}^{n \times r}$, $i = 0, \dots, T-2$, satisfy $P(j+1) = \Phi(j+1) \Phi(j+1)^*$. Then a periodic realization of the reduced order system is defined with $A_r(k) := \Psi_r^* A(k) \Phi_r(k)$, $B_r(k) := \Psi_r^* B(k)$ and $C_r(k) := C(k) \Phi_r(k)$.

4.6.3 Output projection method

The computations delineated above require a vast number of adjoint simulations when very high dimensional outputs are considered, e.g., when the output is not identical to the state, such that one can use state response data in design of an optimal controller (e.g. linear-quadratic regulator) or to analyze system dynamics in detail. In the LTI case Rowley (2005) proposed to project the output on the (few) leading POD modes of the dataset formed by the impulse response simulations. This one involves the kinematic significance of POD modes, to reduce the dimension of the output space, but avoids the weakness of standard POD models that use them as dynamic states. Here we extend the output projection method to periodic systems.

The computation of the structurally constrained optimal $\hat{P}_j$ of the form (4.6.12) is thus reduced to T unconstrained optimization problems for each $P(k)$, $k = j, \dots, j+T-1$, in the periodic setting. Following standard POD rationale, the solutions are $P(k) = \Theta(k) \Theta(k)^*$, where the r_{app} columns of $\Theta(k)$ are the leading POD modes of the dataset $\{G(i+T-i, j)\}_{i=0}^{\infty}$, and the approximation error between the output-projected system and the original system is

$$\begin{aligned} \sum_{i=0}^{\infty} \|G_i(t) - \hat{P}_j G_i(t)\|_F^2 \\ = \sum_{i=0}^{\infty} \sum_{j=0}^{j+T-1} \|G(j+T-i, j) - \hat{P}_j(i+1)G(j+T-i, j)\|_F^2 \\ = \sum_{i=0}^{\infty} \sum_{j=0}^{j+T-1} \lambda(i)_j \end{aligned}$$

where for each i , $\lambda(i)_j, \dots, \lambda(i)_n$ are the descending-ordered eigenvalues $\sum_{j=0}^{j+T-1} G(i+T-i, j)G(i+T-i, j)^*$. The POD modes can be computed by the method of snapshots (Sirovich 1987), applied to datasets comprising the columns of the impulse-response matrices $\{G(i+T-i, j)\}_{i=0}^{\infty}$. Conveniently, provided that $m_e \geq (j+1)T$, periodicity implies that data required to compute these snapshots have already been obtained during the computation of $X(j, m_e)$, as described in § 4.6.1. For instance, the matrix $C(j)X(j, m_e)$ includes the columns of matrices $\{G(j+T-i, j)\}_{i=0}^{m_e/T}$.

The empirical factor $Y(j, m_e)$ of the corresponding observability Gramian

$$W_{yy}(j) = \sum_{i=0}^{\infty} F_{j, j+i}^* C(i)^* \Theta(i) \Theta(i)^* C(i) F_{j, j+i}$$

is needed in order to realize the snapshot-based approximate balanced truncation for the output-projected system (4.6.13). This is accomplished with only $T r_{\text{app}}$ ($r_{\text{app}} \ll q$) impulse-response simulations of the adjoint time-periodic system corresponding to the output-projected system (4.6.13), whose control input is r_{app} -dimensional.

In closing we note that, for additional simplicity and a requirement of a single SVD computation, one can also use a single, time-invariant output projection. Under this constraint, the optimal selection is $P = \Theta \Theta^*$, where the columns of Θ are the leading POD modes of the entire impulse-response $\{(G(i+T-i, j))_{i=0}^{j+T-1}\}_{j=0}^{T-1}$ of (4.6.1). This stronger constraint implies further reduction in matching, when compared with the optimal solution in the lifted domain.

4.6.4 Summary: procedures of balanced POD for periodic systems

Following the terminology in Rowley (2005), the approximate balanced truncation method for linear, time-periodic systems is termed a *lifted balanced POD*. Its main steps include:

- Step 0: Fix a time j , $1 \leq j \leq T$, as the time point for lifting.
- Step 1: Run $T p$ impulse-response simulations to obtain $m_e p$ snapshots and form the $n \times m_e p$ dimensional $X(j, m_e)$ as described in § 4.6.1.
- Step 2: Compute $y = Cx$ from stored states in simulations carried to compute $X(j, m_e)$. Solve for the POD problems for the periodically sampled $y(j+T-i)$, to obtain the output-projection matrices $\Theta(j+i)$, $i = 0, \dots, T-1$.

96

- Step 3: Run T_{exp} impulse-response simulations of the adjoint output-projected system, to form the $n \times m_{\text{exp}} r_{\text{exp}}$ dimensional matrix $\tilde{Y}(j; m_{\text{exp}})$ as described in § 4.6.1.
- Step 4: Compute the SVD of $\tilde{Y}(j; m_{\text{exp}})^* X(j; m_{\text{exp}})$ and the balancing modes for the lifted system given by (4.6.8).
- Step 5: Compute the reduced lifted system (4.6.9).

Various include skipping Step 2, when the output dimension q is small, and using a single, time-invariant output projection, as discussed in § 4.6.3. The reduced system can be un-lifted to a periodic system, e.g., as described in closing § 4.6.2. As in Rowley (2005), an obvious dual version of the algorithm addresses the case of a high-dimensional input space, with only few outputs. This case is motivated by systems susceptible to distributed disturbances, simultaneously affecting the entire state (e.g., $B = I$).

4.6.5 The neutrally stable case

Consider a linear periodic system (4.6.1) that arises from linearization of a system around an asymptotically stable periodic orbit. By Floquet theory (Hartman 1964), in this case $\tilde{A}_j = F_{j+T, j}$ is only neutrally stable, due to one unity eigenvalue that corresponds to persisting perturbations along the periodic orbit in the linearization. Balanced truncation cannot be directly applied to a neutrally stable system, as the infinite series used to define Gramians may diverge.

Ahija & Rowley (2009) presented an extended version of balanced POD for unstable LTI systems that have small unstable dimensions. Following the idea presented in Zhou et al. (1999), it decomposes the system dynamics into stable and unstable parts. Then it applies approximate balanced truncation to the stable dynamics while keeping the unstable dynamics exactly. This method is conceptually applied here to periodic systems through the lifted setting, with all computations executed in the periodic setting. First, for a given lifting time j , define a projection onto the stable subspace $E^s(\tilde{A}_j)$ by $\tilde{P}_j = I_{n, n} - \sum_{k=0}^{\infty} \tilde{A}_j^k \tilde{v}_j \tilde{v}_j^*$ where $\tilde{v}_j, \tilde{v}_j^* \in \mathbb{C}^n$ are the left/right eigenvectors of $\tilde{A}_j$ corresponding to the unity eigenvalue. Dynamics of the neutrally stable lifted system (4.6.3) is thus restricted to the stable subspace of $\tilde{A}_j$:

$$\begin{aligned} \tilde{x}_j(t+1) &= \tilde{A}_j \tilde{x}_j(t) + \tilde{P}_j \tilde{B} \tilde{u}_j(t); \\ \tilde{y}_j(t) &= \tilde{C}_j \tilde{P}_j \tilde{x}_j(t) + \tilde{D} \tilde{u}_j(t), \end{aligned} \quad (4.6.14)$$

where $\tilde{x}_j(t) = \tilde{P}_j x_j(t)$. Lifted balanced POD can be realized to this projected system describing stable dynamics. Let $\Phi_{\tilde{A}_j}^r$ and $\Psi_{\tilde{A}_j}^r$ be the matrices including the leading r_s balancing and adjoint modes of the projected system (4.6.14). Then, a reduced model of order r , $r = r_s + 1$, for the neutrally stable lifted system (4.6.3) can be obtained in the form of (4.6.9), where now $\Phi_r = [\Phi_{\tilde{A}_j}^r, \tilde{v}_j]$, $\Psi_r = [\Psi_{\tilde{A}_j}^r, \tilde{v}_j^*]$. The reduced system keeps the one-dimensional neutrally stable dynamics exactly, while the exponentially stable dynamics is reduced to the order of r_s .

Numerically, the neutrally stable eigenvectors of $\tilde{A}_j$ can be calculated using a Krylov method, or even the power method. By running a control-free simulation of the periodic system (4.6.1) with an arbitrary initial condition $x_{j,0} \notin E^s(\tilde{A}_j)$, one can approximate $\tilde{v}_j$ by $x(j-T)$, with a large T . Similarly, a long-time control-free simulation of the adjoint periodic system (4.6.7) is needed to approximate $\tilde{v}_j^*$. Then, when computing the transformations $\Phi_{\tilde{A}_j}^r$ and $\Psi_{\tilde{A}_j}^r$ for the projected system (4.6.14), one follows exactly the same procedures given in § 4.6.4. The only difference is that in the T_p simulations of the periodic system (4.6.1) described in § 4.6.1, the states should be projected onto $E^s(\tilde{A}_j)$ by $\tilde{P}_j$ at time $j-m-T$. The simulations

then resume with these states as new initial conditions. Similarly, in the adjoint simulations, the adjoint states should be left-multiplied by $\tilde{P}_j^*$ at time $j-T$ before the simulations resume.

By construction, this method is applicable to other neutrally stable/unstable periodic systems, with small neutrally stable/unstable dimensions. For unstable systems, in impulse-response simulations one can repeatedly project the states once each period, using $\tilde{P}_j$, to numerically confine the dynamics to the stable invariant subspace.

4.6.6 Numerical example

To illustrate the balanced POD algorithm, consider an exponentially stable example (similar to that in Farhood et al. (2005)): a linear periodic system (4.6.1) with period $T = 5$, state dimension $n = 30$, output dimension $q = 30$, control input dimension $p = 1$, and $\{A(k)\}_{k=1}^5$ are randomly generated diagonal matrices with diagonal entries bounded in $[0.16, 0.96]$, guaranteeing asymptotic stability. The matrices $B(k)$ and $C(k)$ are also randomly generated, with entries bounded in $[0, 1]$.

Here we pick the "lifting time" $j = 1$. Choose $m_s = m_u = 37 = 15$. Figure 4.6.2(a) shows the error plots of the infinity norm, $\|\tilde{G} - \tilde{G}_r\|_{\infty}/\|\tilde{G}\|_{\infty}$ versus r , the order of the reduced lifted system. Here $\tilde{G}_r$ is the impulse-response matrix of the reduced lifted system of order r . We see that the snapshot-based balanced truncation gives a good approximation of exact balanced truncation. Further, the balanced POD, even with low orders of output projection r_{exp} , generates satisfying results. Recall that, for the lifted system, the order of output projection is $r_{\text{exp}} = r_{\text{exp}} T$.

Figure 4.6.2(b) shows comparisons between balanced POD results with the same order of output projection, one set based on T -periodic projection matrices along one period, and the other using single time-invariant projection matrix (see § 4.6.5). For the cases where r_{exp} are low, these two approaches give almost identical results, or even the latter one gives better results. However, when the order of output projection r_{exp} increases, the results based on T -periodic projection matrices are better than those by a single projection matrix, as we expect.

This algorithm has also been applied to a neutrally stable, time-periodic system obtained by linearizing the Ginzburg-Landau partial differential equation about its exponentially stable time-periodic solution; see Ma & Rowley (2008).

4.7 Low and least order Galerkin models

4.7.1 Overview

Without exception, feedback design requires a model that predicts the dynamic response of the subject system over a time horizon that is sufficiently longer than that of the controlled phenomenon. The complexity of design models range with system characteristics and design objectives, from rudimentary assumptions of input-output monotonicity, in "model free", extremum seeking schemes, to dissipative feedback stabilization that require detailed predictions of the unsteady phases and amplitudes. In these cases, modeling is often the single most significant component of a successful feedback design.

The control of aerodynamic forces over an airfoil illustrates both extremes: Slow modulation of long time averages in attached flows in the domain of traditional flight control. Yet when the low mass and rise of a MAV narrow the gap between time constants pertinent to flight and to unsteady aerodynamics, feedback control is required to tackle an altogether different dynamic range, countering rapid variations in the effective Reynolds number, pitch, yaw and roll angles. Here the modeler walks a tightrope, balancing precision, robustness and simplicity requirements. We focused predominantly on low order Galerkin models




Figure 4.6.2: Error $\|\tilde{G} - \tilde{G}_r\|_{\infty}/\|\tilde{G}\|_{\infty}$ for lifted balanced POD approach: (a) For exact balanced truncation (—), balanced truncation by the method of snapshots but without output projection (---), balanced POD with $r_{\text{exp}} = 1$ (···), balanced POD with $r_{\text{exp}} = 3$ (- · -), and the lower bound for any model reduction scheme (-). All output projections are T -periodic. (b) Time-varying T -periodic output projections versus time-invariant output projections: balanced POD with $r_{\text{exp}} = 1$ (···), balanced POD with $r_{\text{exp}} = 3$ (- · -) and balanced POD with $r_{\text{exp}} = 5$ (—). Solid lines correspond to cases using T -periodic projection matrices, and dashed lines using one single projection matrix.

(LOGMs), and on their use in the design of feedback flow control. To meet the needs of sufficiently large operational envelopes, a major thrust in our groups work has been (i) the identification of intrinsic inconsistencies between the traditional LOGM structure and practice, and flow physics, and (ii) the utilization of our observations as the foundation of modified model structure and modeling methods, that at once remove the inconsistencies at their roots, and, by design, insert the needs of feedback applications.

4.7.2 Broadband representation: Mean field models

Referring to the triple Reynolds decomposition

$$\mathbf{u} = \mathbf{u}^F + \mathbf{u}^C + \mathbf{u}^d, \quad (4.7.1)$$

generic flow control objectives concern a limited range of length and time scales, captured in the coherent flow component $\mathbf{u}^C$. This applies, in particular, to the unsteady regulation of aerodynamic forces, at the heart of our MURI project. Deeply ingrained in common practice is the view that slow variations in the base flow $\mathbf{u}^F$ may be ignored, and that suppressing the short length and time scales of the "stochastic" component $\mathbf{u}^d$ is at the essence of achieving the sought low model order. Perhaps the main contribution of the seminal study by Aubrey et al. (1988), is in refuting both, highlighting the essential roles of mean field dynamics and of subgrid representations of the energy cascade. With these observations as our starting point, we derived modified model structures and new modeling methods, as enablers for the least feasible complexity in effective accounting for the dynamic contributions of $\mathbf{u}^F$ and $\mathbf{u}^d$. Key to our approach is the focus of modeling structure on flow physics first principles at the Navier-Stokes equation (NSE) level.

The apparent success of numerous LOGMs over more than two decades, naturally puts in question our claim regarding the unequivocal and universal necessity of mean field and subgrid models. Our answer is




Figure 4.7.1: The shift mode for the cylinder wake flow (right) is the normalized difference between the unstable steady solution (left) and the mean of the attractor flow (center).




Figure 4.7.2: A subgrid model for a LOGM comprising of the leading attractor POD mode pair and the shift mode. The turbulence energy state $u_T = \sqrt{K^2}$, is well approximated by slaving it to the squared leading oscillations amplitude (left). A LOGM containing nonlinear subgrid model (right, dashed) eliminates the over-prediction of the NSE attractor (right, solid) amplitude and reduces the transient overshoot of the original LOGM (right, dotted). The remaining mismatch is the result of model deformation, as discussed below.

twofold: A generic observation is that successful LOGMs invariably do contain ingredients that represent mean field variations and subgrid corrections for the truncated energy cascade, in either front- or back-door versions. Examples include cubic terms, POD modes extracted with late transient and calibration methods. Our assertion is that model performance is guaranteed to improve once haphazard, back-door representations are substituted by systematic counterparts.

At a more rigorous level, in Tadmor et al. (2010) we have introduced *practical energy balance analysis* as a quantitative design and analysis tool. We then used that tool to irrevocably demonstrate the necessity of a dynamic mean field representation. Following are some added details.

Mean field models are the LOGM equivalents of the Reynolds equation at the NSE level. We first introduced this concept in Noack et al. (2003a), where we demonstrated the necessity for a Galerkin-Reynolds




Figure 4.7.3: A parameterized LOGM enables a near perfect prediction of the acceleration force amplitude (left, solid) and the oscillation frequency (center, solid), when compared with the local Galerkin projections (dots, both), successfully compensating for the residuals in Figure 4.7.2.

model component at the LOGM level. A detailed analysis at the NSE level, in Tadmor, Gonzalez, Lehmann, Noack, Morzyński & W. Stankiewicz (2007), Tadmor et al. (2010), covered both model structure and the aforementioned transient energy balance analysis. Extending the concept of the Reynolds equation to frequency filtered versions of the NSE, we highlighted and quantified the essential roles of the bilateral interplay between distinct frequency bands and, in particular, between u^b and u^f . These analytical tools have also been used in deriving an ideal structure for the Galerkin Reynolds equation, its state(s) and their interactions with the traditional, u^f -centered LOGM.

4.7.3 Broadband representation: Statistical closure and subgrid models

Significant dynamic prediction mismatch between LOGMs and the NSE, including instabilities, motivated the growing prevalence of *model calibration*. The tuning of model parameters by empirical data and, in particular, tuned (model) eddy viscosities, trace to Aubry et al. (1988). Calibration has the aesthetic shortcoming of forfeiting the confidence associated with the first principles foundation the Galerkin projection of the NSE. Perhaps not surprisingly, it is also plagued by all too common γ s of tuned models. Addressing this issue is yet another significant contribution to the development of effective model design and analysis tools.

This thrust had three components: Structural analysis of energy transfer mechanisms at the NSE and the (ideal, infinite) Galerkin system levels (Noack et al. 2007, 2008) was our starting point. It reveals the root cause of failures in the structural mismatch between the linear eddy viscosity dissipation term(s), and the triadic energy exchanges, created by quadratic Galerkin and NSE terms, in an exact representation. In essence, the calibration of the linear model component fails because the linear term terms are structurally precluded from mimicking the nonlinear nature of the truncated energy cascade. At low and least model orders, a successful subgrid remedy must endeavor to match the nonlinearity of the suppressed physical mechanism.

The second component of this thrust faced the call to revisit subgrid Galerkin representations at the structural level. Clearly, a least order model should be focused only on lumped, slow quantities, i.e., the stored energy. That focus inevitably reawakened the statistical closure monster, where it is least welcome, i.e., in a least complexity modeling context. Cognizant of the protracted issues of turbulence modeling, dating to Kolmogorov, Onsager, Batchelor and Kraichnan, we chose to start, *de novo*, with an axiomatic framework and derived a remarkably simple closure theory, expressly for Galerkin models (Noack et al. 2007, 2008). The new framework has been validated in systems ranging in complexity from low order Burgers equation, to a ≈ 1500 modes spectral compression of the 3D homogeneous shear turbulence.

In a third step we revisited the issue of broadband LOGMs (Tadmor et al. 2009, Noack et al. 2010), cre-

dynamics resolved by a single mode in commonplace. This approach, described in Tadmor, Centuri, Noack, Luchtenburg, Lehmann & Morzyński (2007), Tadmor et al. (2008, 2010) has the additional advantage of a simple derivation of a nearly gradient mode interpolation scheme.

4.8 Summary

In this section, we have presented a suite of tools for developing reduced-order models of complex fluid flows, suitable for designing controllers for closed-loop flow control. The version of the immersed boundary method presented in Section 4.1 is fast, and well-suited to the vortex shedding problems considered in the next section. The Balanced POD method, a close approximation to balanced truncation, is a model-reduction method based on linearization about a particular flow state. In its most basic form, it is valid only for equations linearized about a stable equilibrium (e.g., a stable solution of Navier-Stokes equations). Here, the method has been extended to unstable equilibria (such as an unstable steady solution of Navier-Stokes, with no vortex shedding), as well as periodic orbits (such as a particular vortex-shedding cycle). When nonlinearities are important, minimal-order Galerkin models allow one to include these effects, by modeling both the effect on the mean flow and on truncated modes.

ating a hybrid framework that integrates mean field and subgrid representations into the LOGM. In all, the new framework introduces altogether new components into the realm of LOGMs, including mean field *shift* modes that resolve mean field variations, and *nonbalance energy states*, that represent the energy content of suppressed flow components, but for the first time, are not the time coefficients of spatial, expansion modes. First principles model derivation from the NSE thus combine the Galerkin projection of the NSE and the NSE-based energy balance equation.

4.7.4 Mode deformation, modal interpolation and manifold embedding

The very objective of flow control in the MURI project, the shaping and orchestrating of periodic cycles of LEV shedding and regrowth over an airfoil at a high angle of attack (α_{OA}) can be stated in terms of the deformation of leading flow structures. This example is generic to flow control, in general, as well as in reference to the changes in the flow as it traverses natural transients and responds to unsteady ambient conditions.

Motivated by earlier results (Gerhard et al. 2003), we demonstrated in Lehmann et al. (2005), Luchtenburg et al. (2006) both the severe deleterious effects on closed loop performance of ignoring mode deformation, and the feasibility of very simple remedies, based on a parametrized model that utilizes a continuously interpolated mode set. That is, the starting point of the LOGM is a parametrized Galerkin expansion:

$$u(x, t) \approx u^b(x, \alpha) + \sum_{n=1}^N \phi_n(t) u_n(x, \alpha), \quad (4.7.2)$$

where $\alpha \in \mathcal{M}$ is an exogenous parametrization of the operating point. This approach can be formulated in terms of efficient embedding of the approximated flow state in a low dimensional nonlinear manifold, as opposed to much higher dimensional embedding in the host linear subspace. It is similar in this respect to the framework of approximate inertial manifolds. The significant difference, and the innovation in the proposed approach, is in the transition from a single global model, in inertial manifolds, to the utilization of local bases, which the generic differential geometric approach.

Once again, this approach stretches previous boundaries of the Galerkin paradigm. Its advantage is in enabling to maintain a least model order, affecting neither the predictive power nor the robustness of the model. This approach is contrasted, e.g., with approaches that cover mode deformations by enlarging the expansion set, hence both complexity and numerical sensitivity. Benchmark studies included flow separation over natural and actuated transients in flow separation over static and pitching 2D airfoils, and a counterpart range in 2D wake flows (Morzyński et al. 2006, Morzyński, Stankiewicz, Noack, King, Thiele & Tadmor 2007, Morzyński, Noack & Tadmor 2007, Stankiewicz, Morzyński, Roszak, Noack & Tadmor 2008, Stankiewicz, Morzyński, Noack & Tadmor 2008). In these articles we have demonstrated, in particular, the ability to effectively interpolate expansion modes between few, explicitly computed expansions. The new approach enables near perfect state resolution with a least order representation, where higher order traditional models fail.

4.7.5 Harmonically specific modal expansions

Design models typically target few distinct, albeit possibly time varying frequencies. Effective modeling will therefore employ frequency-specific states, hence frequency specific modes. We have formulated and tested a conceptual and computational framework, based on frequency specific modal expansion, as an advantageous alternative to the use of POD modes, for whom the mixing of multiple frequencies in the

5 Control of vortex shedding in numerical simulations

5.1 BPOT models for suppression of vortex shedding

We apply the model reduction techniques developed in Section 4 to the uniform flow past a flat plate in two spatial dimensions, at a low Reynolds number, $Re = 100$. We obtain reduced-order models of a system actuated by means of a localized body force near the trailing edge of the flat plate; the vorticity and velocity contours of the flow field obtained on an impulsive input to the actuator are shown in Fig. 5.1.1. Using these reduced-order models, we develop feedback controllers that stabilize the unstable steady state at high angles of attack. We first assume full-state feedback, but use output projection described in section 4.4.1 to considerably decrease the number of outputs in order to make the model computation tractable. Later, we relax the full-state feedback assumption, and develop a more practical observer-based controller which uses a few velocity measurements in the near-wake of the flat plate (shown in Fig. 5.1.1) to reconstruct the entire flow.




Figure 5.1.1: Actuator modeled as a localized body force near the trailing edge of the flat plate, with the angle of attack fixed at $\alpha = 35^\circ$. Vorticity contours are plotted, with negative contours shown by dashed lines. The velocity-sensor locations are marked by solid circles.

5.1.1 Numerical parameters

The grid size used is 250×250 , with the smallest computational domain given by $[-2, 3] \times [-2.5, 2.5]$, where lengths are non-dimensionalized by the chord of the flat plate, with its center located at the origin. We use 5 domains in the multiple-grid scheme, resulting in an effective computational domain 2^5 times larger the size of the smallest domain; thus the largest domain is given by $[-32, 48] \times [-40, 40]$. The timestep used for all simulations is $\Delta t = 0.01$.




Figure 5.1.2: Forces on a flat plate at a fixed angle of attack α and at $Re = 100$, showing a transition from a stable equilibrium to periodic vortex shedding at $\alpha = 26$. Shown are the force coefficients corresponding to the stable (—) and unstable (---) steady states, and the maximum and minimum (---), and the mean (---) values during periodic vortex shedding. Also shown are the vorticity contours (negative values in dashed lines) of steady states at $\alpha = 15^\circ, 55^\circ$ and the flow fields corresponding to the maximum and minimum force coefficients at $\alpha = 55^\circ$.

105

2. Velocity measurements at two near-wake sensor locations, shown in Fig. 5.1.1, which are used to develop observer-based feedback controllers.

The control goal is to stabilize unstable steady states using the above actuator and sensors, for which we first develop reduced-order models using the method presented in section 4.5.2. We also test the robustness of these controllers in the presence of certain random disturbances.

5.1.3 Steady-state analysis

Since our approach is to obtain reduced-order models of the flow linearized about a given steady state, we first need to compute these steady states. The model-reduction of unstable systems involves projecting the dynamics onto a stable subspace, for which we also need to compute the right and left eigenvectors of the linearized dynamics. This section concerns this steady-state analysis, using a "timestepper-based" approach as outlined in Tuckerman & Barkley (2000) and Kelley et al. (2004).

A simple way of computing stable steady states is by simply evolving the time-accurate simulation to stationarity. However, unstable steady states cannot be found in this manner, and stable steady states near a bifurcation point could take very long to converge. Instead, we use a timestepper-based approach which involves writing a computational wrapper around the original computational routine to compute the steady states using a Newton iteration. If the numerical timestepper advances a circulation field γ^k at a timestep k to a circulation field $\gamma^{k+1} = \Phi_T(\gamma^k)$ after T timesteps, the steady state is given by the field γ_0 that satisfies

$$g(\gamma_0) = \gamma_0 - \Phi_T(\gamma_0) = 0. \quad (5.1.4)$$

The steady states are given by zeros of $g(\gamma_0)$, which could, in principle, be solved for using Newton's method. However, the standard Newton's method involves computing and inverting Jacobian matrices at each iteration, which is computationally infeasible due to the large dimension of fluid systems. Instead of computing the Jacobian, we use a Krylov-space based iterative solver called Generalized Minimal Residual Method (GMRES) developed by Saad & Schultz (1986) to compute the Newton update (see Kelley (1995) and Trefethen & Bau (1997) for a description of the method). This method requires computation of only Jacobian-vector products $Dg(\gamma) v$, which can be approximated using finite differences as $[g(\gamma + \epsilon v) - g(\gamma)]/\epsilon$, for $0 < \epsilon \ll 1$. So, the Jacobian-vector products can also be computed by invoking the appropriately-initialized timestepper. A nice feature of GMRES is relatively fast convergence to the steady state when the eigenvalues of the Jacobian $Dg(\gamma_0)$ occur in clusters; see Kelley (1995) and Kelley et al. (2004) for details. For systems with multiple time-scales, such as Navier-Stokes, most of the eigenvalues of the continuous Jacobian lie in the far-left-half of the complex plane. Thus, the corresponding eigenvalues of the discrete Jacobian $D\Phi_T$, for a sufficiently large value of T , cluster near the origin.

The procedure described above is used to compute the branch of steady states for the angles of attack $0 < \alpha < 90^\circ$; the parameter T in (5.1.4) is fixed to 50 timesteps. The lift and drag coefficients, C_L and C_D , and their ratio C_L/C_D with changing α are plotted in Fig. 5.1.2. As with flow past bluff bodies with increasing Reynolds number (for example, see Provansal et al. (1987)), the flow undergoes a Hopf bifurcation from a steady flow to periodic vortex shedding as the angle of attack α is increased beyond a critical value α_c , which in our computations is $\alpha_c \approx 27^\circ$. Also plotted in the figure are the maximum, minimum, and mean values of the forces during shedding for $\alpha > \alpha_c$. We see that the (unstable) steady state values of the lift coefficient are smaller than the minimum for the periodic shedding till $\alpha \approx 75^\circ$, after which they are slightly higher, but still smaller than the mean lift for the periodic shedding. The (unstable) steady state drag is much lower than the minimum value for periodic shedding. The ratio C_L/C_D of the (unstable) steady state is close to the mean value for shedding. Thus, if the large fluctuations in the forces are undesirable at high angles

107

5.1.2 Input and output

The actuation is modeled as a localized body force near the trailing edge of the flat plate. The flow-field obtained from an impulsive input ($u(t) = \delta(t)$) consists of two counter-rotating vortices, where the circulation of each vortex is given by

$$B_i^*(r) = \pm c(1 - e^{-r^2})e^{-r^2}, \quad i = 1, 2 \quad (5.1.1)$$

$$\text{where } r^2 = (x - x_0)^2 + (y - y_0)^2. \quad (5.1.2)$$

The constants a and c determine the radius and strength of the vortices, while (x_0, y_0) determine the location of the centers of these vortices. The velocity fields corresponding to the functions B_i^* do not satisfy the no-slip boundary conditions at the plate surface; a projection step is used to enforce these conditions and the resulting fields are used to model actuation, that is, $B = B_1^* - B_2^*$, where the field B is plotted in Fig. 5.1.1. The control is implemented in the numerical solver by simply adding a term of the form Bu to the right hand side of (4.5.1).

In previous research, Taira & Colonius (2009a) considered actuators modeled as body forces smeared over a few grid points in studying the effect of open loop constant forcing on three-dimensional flows past a low aspect ratio flat plate, while Williams, Collins, Jankov, Colonius & Tadmor (2008) performed experiments on semi-circular planforms using periodic blowing through slots on the leading edge. The actuation above is a simplistic model of blowing and suction, although our aim here is not to have an accurate representation of blowing or suction, but rather to demonstrate the effectiveness of the algorithm presented in section 4.5.2 by developing simple controllers. Several other actuators were also considered by varying the constants a, c in (5.1.1), and one of the examples that resulted in successful control is reported here.

The energy input from the actuation, in studies using open-loop control by steady or periodic forcing, is often quantified in terms of the momentum coefficient C_μ (Greenblatt & Wygnanski 2000a, Taira & Colonius 2009a) which is defined as:

$$C_\mu = \frac{\rho U_\infty^2 \sigma_{act}}{\frac{1}{2} \rho U_\infty^2} \quad (5.1.3)$$

where U_∞ is the constant actuator velocity in case of steady forcing, σ_{act} is the actuator width, and c is the flat plate chord length. With feedback control, the input u is a function of time and so is U_∞ , and thus the momentum coefficient is time-dependent. However, for the sake of quantifying the control input, we assume that the input u has unit amplitude and is a constant. Later, we will see that the maximum amplitude of u is $O(1)$, so this assumption is reasonable. Here, the maximum velocity of actuation is $U_{act}/U_\infty = 0.07$, while the actuation width is $\sigma_{act}/c = 0.3$, which gives $C_\mu = 0.15\%$. This value is within the standard range $C_\mu = 0.01\%$ to 10% used in studies with steady actuation (Greenblatt & Wygnanski 2000a, Taira & Colonius 2009a).

We consider two different outputs of the system, and they are:

1. The velocity field over the entire fluid domain, which is used for developing full-state feedback controllers. As discussed in section 4.4.1, for large dimensional flows, the model reduction procedure using approximate balanced truncation becomes intractable as the number of adjoint simulations needed is the same as the number of outputs. Hence, output projection is used and the observables are considered to be the velocity field projected onto (i) unstable eigenvectors and (ii) leading POD modes of the stable subspace dynamics (impulse-response).




Figure 5.1.3: (a) Streamlines and vorticity contours (in color) of the unstable steady state at $\alpha = 35^\circ$. (b) C_L vs. time, with the steady state as an initial condition.




Figure 5.1.4: The real (μ_r) and imaginary (μ_i) part of the eigenvalues, representing the growth rate and frequency of the corresponding eigenmodes, of the flow linearized about the steady states in the range $20 \leq \alpha \leq 90^\circ$. Also shown is the frequency of the periodic vortex shedding for $\alpha \geq 27^\circ$ (---).

of attack, it would be useful to stabilize the unstable state. If higher lift is required, it would be desirable to stabilize the state with maximum lift during vortex shedding, but since that state is not a steady state of the governing equations, our method cannot be used to achieve that control goal with the present flow configuration.

The steady state at $\alpha = 35^\circ$ is shown in Fig. 5.1.3(a), and a time history of the lift coefficient C_L with this steady state as an initial condition is shown in Fig. 5.1.3(b). Since the steady state is unstable, the numerical perturbations excite the instability, and the flow eventually transitions to periodic vortex shedding.

We also compute a basis spanning the right and left unstable eigenspaces (Φ_u and Ψ_u) of the flow linearized about the unstable steady states, which are required in our model reduction procedure, for restricting dynamics onto the stable subspace. As the flow undergoes a Hopf bifurcation, a complex pair of eigenvalues crosses the imaginary axis from the left half of the complex plane; thus the dimension of the unstable subspace is two. For solving the linearized eigenvalue problem we use the implicitly restarted Arnoldi method which was implemented by Lehoucq et al. (1998) in the form of a freely available Fortran-77 library called ARPACK. This library can be used to compute a small number of eigenvalues (and eigenvectors) with user-specified properties such as the largest or smallest magnitude, largest or smallest real part, etc. to a desired accuracy. We use ARPACK to compute the leading eigenvectors of the linearized and adjoint equations, that is, those corresponding to the eigenvalues with the largest magnitude.

108




Figure 5.1.5: Basis vectors of the unstable eigenspace of the linearized (left) and the adjoint (right) equations. Vorticity contours are plotted (negative contours are dashed).

The eigenvalues μ of the continuous operator are related to the eigenvalues λ of the discrete operator by $\mu = \log \lambda / (T \Delta t)$, where we fix $T = 30$ timesteps. We computed two eigenvalues with the largest magnitude for the range of angle of attack $20^\circ \leq \alpha \leq 90^\circ$, and found that they form a complex pair, implying an oscillating eigenmode. The real and imaginary parts of these eigenvalues, which correspond to the growth rate and frequency of the instability, are plotted in Fig. 5.1.4. The real part of the eigenvalue becomes positive (or the eigenvalues cross the imaginary axis into the right-half complex plane with a non-zero speed) at $\alpha_c \approx 27^\circ$, confirming Hopf bifurcation. For the post-bifurcation values of α , we also plot the frequency of vortex shedding, which departs considerably from the frequency of the linear instability growth, consistent with the finding of Barkley (2006) for the flow past a cylinder. The real and imaginary parts of the right and left (linear and adjoint) unstable eigenvectors of the flow linearized about the steady state at $\alpha = 35^\circ$ are plotted in Fig. 5.1.5. These modes are qualitatively similar to the structures during periodic vortex shedding, but have different spatial wavelengths, as reported in earlier studies by Noack et al. (2003b) and Barkley (2006).

5.1.4 Reduced-order models

We now describe the process involved in deriving reduced-order models of the input-output responses of (4.4.1), which in this example are the linearized incompressible Navier-Stokes equations (4.3.1, 4.3.2). The actuator used is a localized body force close to the trailing edge of the first plate, plotted in Fig. 5.1.1. The models are derived using the procedure outlined in section 4.5.2. As soon in equation (4.5.15), the output of the system is considered to be the entire velocity field, observed as a projection onto (a) the unstable eigenspace, and (b) the span of the leading POD modes of the impulse response restricted to the stable subspace.

The first step in computing the reduced-order models is to project the flow field B onto the stable subspace of (4.3.1, 4.3.2) using the projection operator P_s defined in equation (4.5.7); the unstable eigenvectors computed in section 5.1.3 are used to define P_s numerically. The vorticity contours of the corresponding flow field $P_s B$ are plotted in Fig. 5.1.6a. The next step is to compute the impulse response of (4.5.8). Instead, for practical reasons, we compute the impulse response of

$$\dot{x}_s = P_s A x_s + P_s B u; \quad (5.1.5)$$

109




Figure 5.1.6: Vorticity contours of (a) the flow field shown in Fig. 5.1.1, projected onto the stable subspace, and (b,c) the first- and fifth-most energetic POD modes of the impulse response, restricted to the stable subspace.

that is, at each timestep of integration, we project the state x_s onto the stable subspace of A . Because the stable subspace is an invariant subspace for the linearized dynamics (4.3.1), theoretically, the impulse responses of equations (4.5.8) and (5.1.5) are exactly the same, and they are the same as that obtained by restricting the impulse response of (4.4.1) to its stable subspace. However, due to the (small) numerical inaccuracy of the projection P_s (which is a result of the numerical inaccuracy of the unstable eigenspaces Φ_u and Ψ_u), the dynamics of (4.5.8) is not strictly restricted to the stable subspace and, in the long term, grows without bound in the unstable direction. Thus, the state is projected at each timestep to ensure that it remains constrained to the stable subspace. Next, we compute the POD modes θ_j^s of the impulse response of (5.1.5), and consider the output of (5.1.5) to be the state x_s projected onto a certain number of these POD modes. Here, 200 snapshots spaced every 50 timesteps were used to compute the POD modes. The leading 4 and 10 POD modes contain 85.00% and 99.06% of the energy respectively and, as it has been observed in previous studies (see Deane et al. 1991, Bak & Rowley 2008), these modes come in pairs in terms of their energy content, a characteristic of traveling structures; the leading first and third POD modes are shown in Fig. 5.1.6.

The next step is to compute the adjoint snapshots, with the POD modes of the impulse response (projected onto the stable subspace of the adjoint) as the initial conditions. As the linearized impulse response, these simulations are also restricted to the stable subspace. Again, instead of computing the response of (4.5.10), we compute that of the following system:

$$\dot{z}_s = P_s^* A^* z_s + P_s^* C^* u. \quad (5.1.6)$$

The snapshots of the impulse responses of systems (5.1.5) and (5.1.6) are stacked as columns of X and Z , and using the expressions (4.4.19) and (4.4.20), we obtain the balancing modes ϕ_j^s and the adjoint modes ψ_j^s .

110




Figure 5.1.7: Vorticity contours of the leading (in the order of Hankel singular values of the stable subspace dynamics) first and third balancing (left) and adjoint (right) modes.




Figure 5.1.8: The empirical Hankel singular values (—) and the diagonal elements of the controllability (---) and observability (---, x) Gramians of a 25-mode model with a 4, 10, and 20-mode output projection, for the unstable steady state at $\alpha = 35^\circ$.

111




Figure 5.1.9: Outputs (projection of the flow field onto POD modes) from a reduced-order model obtained using a 20-mode output projection. The first (top figure) and eleventh (bottom figure) outputs of the DNS (—) are compared with predictions of models with 4 (---), 10 (---, $\circ$), and 20 (---, $\times$) modes.

112

We used 200 snapshots of the linearized simulation and 200 snapshots of each adjoint simulation, with the spacing between snapshots fixed to 50 timesteps, to compute the balancing transformation. These number of snapshots and the spacing were sufficient to accurately compute the modes; further reduction in the spacing did not significantly change the singular values from the SVD computation (A.4.19). We considered the output to be a projection onto 4, 10 and 20 POD modes (corresponding to 4, 10 and 20 mode output-projections), as introduced in section 4.4.1. Using these modes, we use the expressions in equation (4.5.20, 4.5.25) to obtain the matrices A_r, B_r, C_r defining the reduced-order model of the stable-subspace dynamics. The vorticity contours of the balancing and the adjoint modes, for a 10-mode output projection system, are plotted in Fig. 5.1.7. The adjoint modes provide a direction for projecting the linearized equations onto the subspace spanned by the balancing modes. Since these modes are quite different from the POD and the balancing modes, the resulting models are also quite different from those obtained using the standard POD-Galerkin technique wherein an orthogonal projection is used. Since the models obtained using balanced truncation are known to perform better than the POD-Galerkin models, as reported by Ilak & Rowley (2008), the better performance could be attributed to a better choice of projection using the adjoint modes.

Since the reduced-order models of the stable-subspace dynamics are approximately balanced, the controllability and observability Gramians of the a_r -dynamics of (4.5.20) are approximately equal and diagonal. Further, their diagonal values are approximately the same as the Hankel singular values σ_i obtained by the SVD (A.4.19). The diagonal values of the Gramians and the singular values for different output projections are plotted in Fig. 5.1.8 for a 30-mode reduced-order model. With increasing order of output projection, the HSVs converge to the case with full-state output, and the number of converged HSVs is roughly equal to the order of output projection, as was observed by Ilak & Rowley (2008). We see that the diagonal elements of both the Gramians are very close to the HSVs for the first 20 modes. For higher modes, the diagonal elements of the observability Gramians are inaccurate, which is due to a small inaccuracy of the adjoint formulation mentioned in section 4.3. For controller design, we use models of order ≤ 20 , for which these Gramians are sufficiently accurate.

Finally, to test the accuracy of the reduced-order models, we compare the impulse responses of system (5.1.5) (that is, restricted to the stable subspace) with that of the model (4.5.20), restricting $\phi_a = 0$. In particular, we compare the outputs of the two systems, which are the projection onto the POD modes; a representative case in Fig. 5.1.9 shows the results of 4, 10 and 20 mode models of a system approximated using a 20-mode output projection (the outputs are projection onto the leading 10 POD modes). The first output, which is a projection onto the first POD mode, is well captured by all the models until $t = 60$, while the 20-mode model performs well for all time. Also shown is the eleventh output, which is well captured only by the 20-mode model.

5.1.5 Full-state feedback control

The resulting models can now be used along with standard linear control techniques to obtain stabilizing controllers. We use Linear Quadratic Regulator (LQR) to compute the gain K so that the eigenvalues of $(A + BK)$ (where the matrices were defined in (4.5.20)) are in the left-half of the complex plane, and the input $u = Kx$ minimizes the cost

$$J(u, x) = \int_0^\infty (x^T Q x + u^T R u) dt, \quad (5.1.7)$$

where Q and R are positive weights computed as follows. We choose Q such that the first term in the integrand of (5.1.7) represents the energy, that is, we use $Q = C^T C$, with C defined in (4.5.21). The weight R

113




Figure 5.1.12: Lift-coefficient C_L vs. time t , for full-state feedback control, with control turned on at different times in the base uncontrolled simulation. The base case with no control (—) has the unstable steady state as the initial condition, and transitions to periodic vortex shedding. The control is tested for different initial conditions, corresponding to $t = 170, 180, 210$ of the base case, and stabilizes the steady state in all the cases (—).

is chosen to be a multiple of the identity cI , and typically c is chosen to be a large number $\sim 10^4$, to avoid excessively aggressive controllers. The control implementation steps are sketched in Fig. 5.1.10; first compute the reduced-order state x , using the expression (4.5.27), then the control input is given by $u = Kx$. Here, we derive the gain K based on a 22-mode reduced-order model (with 2 unstable and 20 stable modes), using $R = 10^4$, and include the same in the original linearized and nonlinear simulations. The output is approximated using a 4-mode output projection. The difference between the linear and nonlinear simulations is that, in the latter, the steady state field x_0 is subtracted from the state x , before projecting onto the modes to compute the reduced-order state a .

Fig. 5.1.11 compares the model predictions with the projection of data from the simulations of the linearized system (4.3.1, 4.3.2), with a control input. The initial condition used is the flow field obtained from an impulsive input to the actuator. Both the states shown in the figure eventually decay to zero, which implies that the perturbations decay to zero, thus stabilizing the unstable steady state. More importantly, the model predicts the outputs accurately for the time horizon shown in the plots.

We now use the same controller in the full nonlinear simulations and test the performance of the model for various perturbations of the steady state. A plot of the lift coefficient C_L vs. time t , with the control turned on at different times of the base simulation, is shown in Fig. 5.1.12. The initial condition for the base case (no control) is the unstable steady state; eventually, small numerical errors excite the unstable modes and the flow transitions to periodic vortex shedding. In separate simulations, control is turned on at times $t = 170, 180, 210$ corresponding to the base case. As the figure shows, the control is effective and is able to stabilize the steady state in each case, even when the flow exhibits strong vortex shedding. We remark that the latter two of these perturbations are large enough to be outside the range of validity of the linearized system, but the control is still effective, implying a large basin of attraction of the stabilized steady state. We also compare the output of the reduced-order model with the outputs of the nonlinear simulation; the plots are shown in Fig. 5.1.13. The models perform well for the initial transients, but for longer times

115




Figure 5.1.10: Schematic of the implementation of full-state feedback control in the nonlinear simulations. The entire velocity is first projected onto the unstable eigenvectors and the stable subspace POD modes to compute the reduced-order state a . The state is then multiplied by the gain K , computed based on the reduced-order model using LQR, to obtain the control input u .




Figure 5.1.11: Comparison of the outputs $y_{a,1}$ and $y_{a,2}$ of a 12-mode reduced-order model (—) with the projection of data from the linearized simulation (---). The control gain is obtained using LQR, and the initial condition is that obtained by an impulsive input to the system. Control is turned on at $t = 0$.

114




Figure 5.1.13: Outputs of a system with full-state feedback control. The control gain is obtained using LQR, and the initial condition is that corresponding to $t = 180$ of the uncontrolled case plotted in Fig. 5.1.12. Comparison of the outputs $y_{a,1}$ and $y_{a,2}$ of a 12-mode (2 unstable and 10 stable modes) reduced-order model (—) with the projection of data from the full nonlinear simulation (---).

fail to capture the actual dynamics. This is not surprising as these perturbations are outside the range of validity of the linear models. For control purposes, it appears to be sufficient to capture the initial transients (approximately one period), during which the instability is suppressed to a great extent. We remark that one could possibly compute nonlinear models by projecting the full nonlinear equations onto the balancing modes, or enhance the model subspace by adding POD modes of vortex shedding and the shift modes as proposed by Noack et al. (2005) to account for the nonlinear terms.

Finally, we note that the reduced-order model (4.5.20) decouples the dynamics on the stable and unstable subspaces, and also, the dynamics on the unstable subspace can be computed only using the unstable eigenvectors Φ_u and Ψ_u . Thus, we could derive a control gain $K \in \mathbb{R}^{2 \times n}$, based only on the two-dimensional unstable part of the model, such that the eigenvalues of $(\tilde{A}_u - \tilde{B}_u K)$ are in the left half complex plane. That is, we can obtain a stabilizing controller without modeling the stable subspace dynamics. We have performed simulations to test such a controller and found that it also is capable of suppressing the periodic vortex shedding and thus results in a large basin of attraction for the stabilized steady state. The choice of weight matrices Q and R in the LQR cost (5.1.7) needs to be different to obtain a comparable performance. However, as shown in the next section, it is essential to model the stable subspace dynamics to design a practical controller based on an observer that reconstructs the entire flow field using a few sensor measurements.

5.1.6 Observer-based feedback control

The full-state feedback control of section 5.1.5 is not directly useful in practice, since it is not possible to measure the entire flow field. Here, we consider a more practical approach of measuring certain flow quantities at a small number of sensor locations. We assume that we can measure the velocities at the sensors shown in Fig. 5.1.1, in the near-wake of the plate. We remark that, even though these sensors may not be realizable in applications, they serve as a good testing ground for our models.

5.1.7 Reduced-order models

The method described in detail in section 5.1.4 to obtain models of a system with the full-state output is first used to obtain models of a system with the output represented by the two sensor measurements. For this

116




Figure 5.1.14: Schematic of the implementation of observer-based feedback control in the nonlinear simulation. The control input u and the sensor measurements y are used as inputs to the observer, which reconstructs the reduced-order state δ . This state is then multiplied by the gain K , to obtain the control input u . Both, the controller and observer gains K and L are computed based on the reduced-order model using LQR and LQG respectively.

case, the output matrix C in equation (4.4.1) has two rows and is sparse with each row filled with 0s except for the entry corresponding to a sensor measurement, which is 1. Since the dimension of the output is small, the output projection step of the algorithm outlined in section 4.5.4 is not required. Two adjoint simulations for each sensor location are performed, with the initial condition obtained in two projection steps: first, the velocity field with a unit v -velocity at the sensor location is projected onto the space of flow-fields satisfying the incompressibility constraint and the no-slip boundary condition at the flat plate surface and second, the resulting field is then projected onto the stable subspace of the adjoint using the global eigenmodes (that is, using the projection operator P_s^*). The resulting initial condition for the adjoint simulations is plotted in Fig. 5.1.15. The snapshots from the two adjoint simulations are stacked as columns of Z , and the expressions (4.4.19) and (4.4.20) are used to compute the balancing modes ϕ_i^* and the adjoint modes ψ_i^* . We again used 200 snapshots of the linearized simulation and 200 snapshots of each adjoint simulation, with the spacing between snapshots fixed to 50 timesteps, to compute the balancing transformation. Using these modes, we use the expressions in equation (4.5.15) to obtain the matrices $\tilde{A}$, $\tilde{B}$, $\tilde{C}$, defining the reduced-order model of the stable-subspace dynamics. The resulting balancing modes are qualitatively similar to those plotted in Fig. 5.1.7, however the adjoint modes which are plotted in Fig. 5.1.16 are different from those for the full-state output and the leading modes have support near the sensor locations. The resulting models are again balanced; a 22-mode model (with 2 unstable and 20 stable modes) exhibits good performance and is used to compute the feedback gain K and to design reduced-order observers to estimate the reduced-order states.

5.1.8 Observer design

Using the models derived in section 5.1.7, we design an observer using a Linear Quadratic (LQ) estimator, or Kalman filter. This method assumes that the errors in representing the state δ and the measurement y (due to the inaccuracies of the model) are stochastic Gaussian processes, and results in an estimate $\hat{\delta}$ of the state δ that is optimal in the sense that it minimizes the mean of the squared error; refer to Skogestad & Postlethwaite (2005) for details. We now discuss briefly our procedure for modeling these noises; consider the reduced-order model (4.5.20), but with process noise w and sensor noise v which enter the dynamics as

$$\dot{\delta} = \tilde{A}\delta + \tilde{B}u + w \quad (5.1.8)$$

$$y = \tilde{C}\delta + v \quad (5.1.9)$$

A key source of the process (state) noise w arises from model truncation, and second, from ignoring the nonlinear terms in the reduced-order model. The nonlinearity of the dynamics is important, for instance, when the model is used to suppress vortex shedding. A source of the sensor noise arises from two sources; first, the state x is approximated as a sum of a finite number of modes (4.5.12), and second, in the output projection step, the output is considered as a projection of the (approximated) state x onto a finite number of POD modes (4.5.18). Here, we approximate these two noises as Gaussian processes whose variances are

$$Q = E(w w^*), \quad w = f(\hat{a}_{\text{trunc}}) - \hat{a}_{\text{trunc}}, \quad (5.1.10)$$

$$R = E(v v^*), \quad v = y - \tilde{C}\hat{a}_{\text{trunc}}, \quad (5.1.11)$$

and $E(\cdot)$ gives the expected value. Here, $f(\cdot)$ is the operator obtained by projecting the nonlinear Navier-Stokes equations onto the balancing modes Φ , using the adjoint modes Ψ . The state $\hat{a}_{\text{trunc}}$ is obtained by projecting the snapshots, obtained from a representative simulation of the full nonlinear system, onto the balancing modes. While w is not actually a Gaussian white-noise process, for the purposes of observer design, all we require is an approximate measure of the size of the modeling errors (here modeled as external disturbances), and for this purpose, the Gaussian approximation suffices. The representative simulation we used here is the base case, with no control, shown in Fig. 5.1.12, which includes the transient evolution from the steady state to periodic vortex shedding. The resulting estimator is of the form

$$\dot{\hat{\delta}} = \tilde{A}\hat{\delta} + \tilde{B}u + L(y - \tilde{C}\hat{\delta}), \quad (5.1.12)$$

$$\hat{y} = \tilde{C}\hat{\delta}, \quad (5.1.13)$$

where $\hat{\delta}$ is the estimate of state δ , $\hat{y}$ is the estimated output, and L is the observer gain. The estimator is then used along with the full-state feedback controller designed in section 5.1.5 to determine the control input; a schematic is shown in Fig. 5.1.14.

5.1.9 Observer-based control

The models obtained in section 5.1.7 are used to design dynamic observers based on the vertical (v -) velocity measurements at the sensor locations. A 22-mode reduced-order model, with 2 and 20 modes describing the dynamics on the unstable and stable subspaces respectively, is used to design a Kalman filter for producing an optimal estimate of the velocity field based on Gaussian approximations of error terms (5.1.10, 5.1.11). This estimate is then used along with reduced-order model controller to determine the control input, as shown in Fig. 5.1.14. The results of this observer-based controller (or compensator) are shown in Figs. 5.1.17, 5.1.18. The compensator again stabilizes the unstable operating point, and furthermore, the observer reconstructs the reduced-order model states accurately. Initially, the observer has no information about the states (the initial condition is zero), but it quickly converges to and follows the actual states.

Finally, to test the robustness of the resulting controller, an external disturbance is added to the flow upstream of the flat plate. The disturbance is modeled using the same functional form (5.1.1) used to model the actuation but with the parameters $a = 4$ and $c = 0.05$. The vorticity contours of the resulting field are plotted in Fig. 5.1.19 and the disturbance has support over a much wider region as compared to the




Figure 5.1.15: Contours of the initial condition for the adjoint simulation corresponding to the left sensor location shown in Fig. 5.1.1.




Figure 5.1.16: Contours of the leading first and third adjoint modes of the stable subspace, corresponding to the system with the outputs being sensor measurements.




Figure 5.1.17: Lift-coefficient C_L vs. time t , for estimator-based feedback control, with control turned on at different times in the base uncontrolled simulation. The base case (—) is the same as in Fig. 5.1.12, and the control is tested for different initial conditions, corresponding to $t = 170, 180, 210$ of the base case (---). In both the cases, this controller stabilizes the flow to a small neighborhood of the steady state.




Figure 5.1.18: States of the system with observer-based control; the states reconstructed (---, x) by a 22-mode observer quickly converge to the actual states (—, δ). The initial conditions used are those corresponding to $t = 180, 210$ (top and bottom) of the uncontrolled case shown in Fig. 5.1.17.




Figure 5.1.19: Estimator-based feedback control, in the presence of an external disturbance modeled using (5.1.1). Plot (b) shows the vorticity contours of the disturbance. The resulting lift coefficients are shown in plot (a), both with the control turned off (—) and on (---), with the initial condition being a velocity field from the vortex shedding regime. The controller stabilizes the flow even in the presence of these disturbances. The control input u , which remains $O(1)$, is shown in plot (c).

sensation. The disturbance amplitude is modeled as a random variable sampled from a uniform distribution in the range $[-1, 1]$. The lift coefficients, in the presence of this disturbance, for the flow with control turned off and on, are shown in Fig. 5.1.19. With the control off, the lift stays in the neighborhood of its value during vortex shedding. When the compensator is turned on, the shedding is suppressed and the steady state is again stabilized. However, the disturbance causes the lift to fluctuate around the steady state value. When the disturbance is finally turned off, the lift again converges to the steady state value.

5.2 Synchronization of vortex shedding for high-lift limit cycles

Open- and closed-loop synchronization of vortex shedding in two-dimensional flow over a flat plate at high angle of attack is now considered. In order to study the control of two basic constituents of unsteady post-stall flow (i.e. leading-edge and trailing-edge vortices) and develop a physically motivated feedback strategy, we consider a two-dimensional flow at $Re = O(10^4)$. Even though introducing camber or using Eppler airfoil shape would improve uncontrolled performance, the flat plate ensures the separation at the leading edge in the post-stall regime and allows us to avoid additional complications due to the variation of the separation point or curvature effects of a different airfoil geometry. A Reynolds number of 300 was selected to be high enough to ensure forming and shedding of large coherent structures of opposite signs from the leading and trailing edges. Even though this model is highly idealized, we hope to develop a physically motivated closed-loop control of global instability of vortex shedding without any computational burden.

As in the BPOD study, unsteady actuation is modeled as a body force near the leading or trailing edge and is directed either upstream or downstream. For moderate angles of attack, sinusoidal forcing at the natural shedding frequency results in phase locking, with a periodic variation of lift at the same frequency, leading to higher unsteady lift than the natural shedding. However, at sufficiently high angles of attack, a subharmonic of the forcing frequency is also excited and the average lift over the forcing period varies from cycle to cycle in a complex manner. It is observed that the periods with the highest averaged lift are associated with particular phase differences between the forcing and the lift, but that this highest-lift shedding cycle is not always stably maintained with open-loop forcing. We design a feedback algorithm

121

to lock the forcing with the phase shift associated with the highest period-averaged lift. It is shown that the compensator results in a stable phase-locked limit cycle for a broader range of forcing frequencies than the open-loop control, and that it is able to stabilize otherwise unstable high-lift limit cycles that cannot be obtained with open-loop control. For example at an angle of attack of 40° , the feedback controller can increase the averaged magnitude of force on the plate by 76% and increase the averaged lift coefficient from 1.33 to 2.43.

In this section, open-loop control with periodic pulsing at the natural shedding frequency is first investigated for various actuator configurations over a range of α .

In certain cases, primarily for lower angles of attack, open-loop forcing results in a phase-locked limit cycle with lift varying at the frequency of actuation. The momentum coefficient,

$$C_p = \frac{\rho U_j^2 \Delta x}{\frac{1}{2} \rho U_\infty^2 c} \quad (5.2.1)$$

is the ratio between the momentum injected by the forcing and that of the freestream. The values of C_p reported are based on the average jet velocity, U_j , fixed at 0.5, and the width of the actuator, $\Delta x = 0.02$. This corresponds to a fixed C_p of 0.01 for all of the cases considered here. For each actuation location, two cases of blowing angles are considered, one directed downstream and the other directed upstream.

For sufficiently high α , subharmonic frequencies are excited and a more complex limit cycle behavior is obtained. The period-averaged lift over one cycle of actuator forcing varies from cycle to cycle, and it is observed that higher lift is associated with a particular phase shift between the forcing and the lift. We show that feedback of the lift signal can be used to phase lock the forcing to the particular phase shift associated with the highest period-averaged lift. Similar phase-locking feedback control has been used in the aforementioned study of Pantor et al. (2008) and by Tadmor (2004).

5.2.1 Uncontrolled Flow

For the translating flat plate at $Re = 300$, steady attached flow is observed for $\alpha < 10^\circ$. At $\alpha = 10^\circ$, the flow is observed to be separated but remains steady. The flow undergoes a Hopf bifurcation between angles of attack of 12° and 15° after which vortex shedding occurs with natural shedding frequency, ω_n , which varies from 3.65 at $\alpha = 15^\circ$ to 1.39 at $\alpha = 50^\circ$. Using the vertical projection of the airfoil to the freestream, we find that ω_n can be scaled, for $\alpha \geq 30^\circ$, to a Strouhal number of $St = f_n c \sin(\alpha) / U_\infty = 0.2$, where $f_n = \omega_n / (2\pi)$. This agrees with the wake Strouhal number for vortex shedding behind two-dimensional bluff bodies (Roshko 1961; Bearman 1967; Griffin 1978). The unsteady shedding cycle consists of vortices of opposite signs alternately shed from the leading and trailing edges, creating periodic oscillations in the lift and drag. As α is increased, larger vortex structures are formed, inducing a larger amplitude of oscillation in the forces exerted on the plate. For $\alpha \geq 30^\circ$, the vortex structure on the suction side of the plate is observed to be created from the leading edge and can be viewed as a transient LEV, or, equivalently, a dynamic stall vortex (DSV) that occurs during a rapid pitch up. Maximum lift is found when the LEV is brought down to the suction side of the plate as it grows in strength. The lift decreases as the new vortex structure of the opposite sign is formed at the trailing edge. This trailing-edge vortex (TEV) pushes up the LEV sitting on the suction side of the plate, and finally halts its growth causing it to pinch-off and shed into the wake.

5.2.2 Open-loop control

In order to investigate the effect of unsteady blowing on these vortex shedding cycles, we first consider open-loop control using periodic pulsing with different blowing angles at the leading and trailing edge of

122

the plate. The nondimensional jet velocity is set as $U_j = \bar{U}_j + U'_j \sin(\omega_f t)$, where $\bar{U}_j = 0.5$ and $U'_j = 0.5$. Since this study is focused on maximizing lift from shedding of the coherent vortex structures rather than the suppression of shedding or separation, α is initially chosen to be the natural shedding frequency for each α , at which the unsteady shedding of the large coherent vortex structure will likely be amplified the most (Glezer et al. 2005; Amitay & Glezer 2002a). In the next two sections we examine leading and trailing edge actuation, respectively.

Leading-edge actuation Figure 5.2.1 shows the lift coefficient with actuation at the leading edge directed downstream (left) and upstream (right). In each figure, the uncontrolled flow (baseline) is overlaid in grey with its average in dashed grey and its maximum and minimum bounding the shaded region. Squares show the minimum and maximum of the lift signal whose overall average is shown in the circles in between. For cases where the lift is not phase locked to the forcing signal, variation in the period-averaged lift (averaged over each actuation period) is also plotted with error bar.

Blowing downstream provides extra momentum at the leading-edge in addition to that of the freestream. This amplifies the unsteady shedding of vortex structures, resulting in larger magnitudes of the lift fluctuations. The forced flow exhibits higher maximum lift, but also lower minimum lift, below that of the baseline flow. As a result, blowing downstream does not significantly benefit the overall average lift.

However, when the actuation is directed upstream, the resulting amplification of the unsteady shedding has a more positive effect on the average lift. For $\alpha < 25^\circ$, the flow locks onto the forcing 2–3 periods after the actuation is initiated. However at higher α , the flow fails to lock onto the forcing frequency and displays a more complicated limit cycle, with subharmonics of the forcing frequency also excited. An example is shown in figure 5.2.2, at $\alpha = 50^\circ$ where each subharmonic limit cycle consists of several periods with a different period-averaged lift. Figure 5.2.2 also shows the lift as a function of the jet velocity, and shows that the actuation produces the highest lift when U_j is in phase with the C_L (maximum C_L when U_j is maximum). However, the succeeding period becomes slightly out of phase and the lift decreases. Each period within the subharmonic limit cycle is observed to be associated with a particular phase shift, ϕ , between the forcing signal and the lift, yielding a particular period-averaged lift. The actuation period associated with the highest average lift is plotted in a thicker line. At each α , there is a particular ϕ , resulting in the highest average lift over an actuation period. If the feedback allows us to accordingly adjust the frequency of actuation to phase lock the flow at these ϕ , then we could repeatedly produce the highest average lift period. This feedback design will be revisited later.

It might be counter-intuitive that upstream actuation at the leading edge achieves such a lift enhancement and performs better than downstream actuation. However, experiments at Reynolds number of the order of 3×10^4 by Rullan et al. (2006) demonstrated that unsteady blowing upstream, parallel to the chord at the leading-edge of a sharp-edged, circular arc airfoil at various α beyond stall leads to averaged pressure distribution that resulted in higher lift than that of the baseline flow. They achieved lift increase as high as 70% with momentum coefficient of $C_p = C_p / \sin(\alpha) = 1\%$, scaled with the vertical projection of the airfoil and the actuation pulsing at the shedding frequency of the airfoil.

Trailing-edge actuation In figure 5.2.3, the lift performance of the open-loop actuation at the natural shedding frequency at the trailing edge is investigated in a similar manner as in figure 5.2.1. Blowing downstream exerts a negative effect on the average lift, yielding a lower minimum lift than that of the baseline flow with a similar maximum lift.

However, when the forcing is directed upstream, the forced flow displays a significant lift enhancement. The forcing excites the vortex shedding cycle even for α below the Hopf bifurcation. For $\alpha \leq 15^\circ$, the flow

123




Figure 5.2.1: Leading-edge actuation: maximum and minimum lift ($\square$) and its average over time ($\circ$) for downstream (left) and upstream (right) actuation. Average of the baseline case is plotted in dashed grey and shaded region is bounded by its maximum and minimum. Actuation is applied at the natural shedding frequency, $\omega_f = \omega_n$. For cases where the flow is not phase locked to the forcing signal, variation in period-averaged lift over each actuation period is plotted with error bar to indicate the range of values over a subharmonic limit cycle.




Figure 5.2.2: Lift as a function of time (a) and jet velocity (b) with upstream actuation at the leading edge (LE, upstream) at the natural shedding frequency ($\alpha = \alpha_n$) for $\alpha = 50^\circ$.

124




Figure 5.2.3: Trailing-edge actuation: see Figure 5.2.1 for a description.




Figure 5.2.4: Vorticity contour at the time of maximum lift for baseline (thin) and upstream actuation (thick) at the trailing edge at the natural shedding frequency ($\omega_f = \omega_b$). Dashed and solid lines represent counter-clockwise and clockwise vorticity.




Figure 5.2.5: Trailing-edge actuation: maximum and minimum lift ($\square$), average lift ($\circ$), and period-averaged lift (error bar) over a range of open-loop forcing frequency, ω_b . Maximum and minimum lift of baseline (---) case is shown as a reference.




Figure 5.2.6: Trailing-edge actuation: phase shift of the forcing signal, U_f , relative to the lift signal, C_L , for phase-locked flows, over a range of open-loop forcing frequency, ω_b ($\alpha = 30^\circ$).

locks onto the forcing after 2–3 periods. However, for $\alpha \geq 20^\circ$, the subharmonic resonance is excited. This is similar to the observation with upstream blowing at the leading edge, but the subharmonic resonance is excited at a lower α for the trailing-edge actuation than that for the leading-edge actuation.

Each period within the subharmonic limit cycle is again observed to be associated with a particular ϕ , resulting in a particular period-averaged lift. We denote the ϕ associated with the highest period-averaged lift at each α as ϕ_{opt} . Particularly at $\alpha = 30^\circ, 40^\circ$, and 50° , ϕ_{opt} was observed to be approximately $-0.25, -0.05$, and 0.0 radians, respectively. For trailing-edge actuation, the period-averaged lift at high α is, in many cases, greater than the maximum lift occurring in the baseline flow. This suggests a greater potential for the trailing-edge feedback actuation to sustain the flow with the highest period-averaged lift. Consequently, we would obtain a phase-locked flow that has an average lift as high as the maximum lift of the baseline flow (or even higher).

In general, blowing upstream at both the leading and trailing edges creates significant enhancement in the average lift. However, blowing upstream at the trailing edge provides a larger increase in lift than that of the leading-edge actuation. These findings are similar to observations made by Huang et al. (2004) who investigated the effect of blowing and suction control at various locations on the upper surface of a NACA0012 airfoil. They considered steady blowing and suction at $Re = 5 \times 10^5$ and $\alpha = 18^\circ$ and demonstrated that blowing at the leading edge directed downstream exerts a negative effect, decreasing lift and increasing drag at the same time, but suction increases lift by creating a larger and lower pressure zone on the airfoil's upper surface. They also observed that the actuation near the trailing-edge on the upper surface, $0.8c$ from the leading edge, improves lift and drag characteristics by manipulating the circulation of the TEV.

In order to understand the lift-enhancing mechanism of upstream actuation at the trailing edge, we compare the vorticity contours at the time of maximum lift for the cases of baseline and upstream actuation at the trailing edge, for 40° in figure 5.2.4. Actuation feeds extra circulation to the TEV which induces a stronger downwash near the trailing edge. As a result, the vortex structure on the suction side is pulled down closer to the plate and the backflow near the trailing edge is reduced. Particularly at $\alpha = 40^\circ$, this delays the interference of the newly forming TEV with the LEV residing on the suction side. It also lengthens the duration over which the vortex structure is formed from the leading edge. These results also agree with the observations that the period associated with the highest lift within a subharmonic cycle in figure 3(b) has a longer period than that of the baseline flow. This might indicate that there exists a forcing frequency below ω_b at which the flow becomes phase locked to the forcing at a higher lift than that of the baseline flow.

Thus, we next investigate the possibility of the existence of shedding cycles that are phase locked to the open-loop forcing signal. Figure 5.2.5 shows the lift response to the variation in open-loop forcing frequency for $20^\circ \leq \alpha \leq 50^\circ$, above which upstream actuation at $\omega_f = \omega_b$ fails to phase lock the flow.

Over a range of frequency below ω_b , the flow is phase locked to the actuation with its average lift near the maximum period-averaged lift of the flow actuated at $\omega_f = \omega_b$. As we go deeper into stall by increasing α , the domain of attraction for the phase-locked limit cycle decreases, and finally at $\alpha = 50^\circ$, actuation failed to phase lock the flow over the range of forcing frequencies considered.

Figure 5.2.6 shows the corresponding phase shift, ϕ , over this range of ω_b that achieves a phase-locked flow for $\alpha = 30^\circ$. Recall that the subharmonic cycle (excited with upstream blowing at ω_b) consists of several actuation periods with a distinct ϕ associated with a particular period-averaged lift. Also, its highest period-averaged lift is associated with $\phi = \phi_{\text{opt}} \approx -0.25$. As ω_b is decreased below ω_b , the flow phase locks onto one of those periods observed in the subharmonic limit cycle at $\alpha = 30^\circ$. As ω_b is decreased further, the flow is phase locked onto different periods with different ϕ , closer to ϕ_{opt} , with higher average lift. Finally, at $\omega_b/\omega_b0 = 0.57$, the actuation is able to lock the flow at the best period achieved with forcing at ω_b . This indicates that each phase-locked limit cycle of the vortex shedding could be characterized by its frequency

and the phase shift, yielding a particular maximum, minimum, and average lift.




Figure 5.2.7: Feedback control configuration.

If the feedback allows us to adjust the frequency of the actuation accordingly to keep the phase shift between the forcing signal and the lift constant (for example at $\phi = \phi_{\text{opt}}$), we should be able to reproduce the high-lift shedding cycles over a wide range of α . Thus in order to achieve the desired phase-locked shedding cycle, we feedback lift into the controller, whose details are described in the next section.

5.2.3 Closed-loop control

Open-loop periodic forcing can lead to limit cycles with a high average lift, but with a decreasing domain of attraction as α increases. Our goal with closed-loop control is to obtain forced limit cycles with the maximum average lift. This involves stably maintaining limit cycles that are not stable without feedback.

Since the actuated flows with the highest average lift seem to be characterized by a distinct phase shift of the forcing relative to the lift at each α , we feedback C_L in an attempt to phase lock the flow at these high-lift states. Direct feedback of C_L with appropriate gain would only allow us to force the flow to be in-phase with U_b . However, since the observed best phase shifts between C_L and U_b are negative, shifting direct feedback signal requires us to know the frequency of the forced flow *a priori*. Instead, we assume that the lift signal being fed back is approximately sinusoidal. In such cases lift can be expressed as

$$C_L(t) = -a_0 + A_L \cos(\omega_b t + \theta), \quad (5.2.2)$$

$$= a_0 + a_1 \cos(\omega_b t) + b_1 \sin(\omega_b t).$$

Assuming that A_L and θ are slowly varying in time, we can estimate a_1 and b_1 to be the Fourier mode over a moving window,

$$a_1(t) = \frac{2}{T} \int_{t-T}^t L(t') \cos(\omega_b t') dt', \quad (5.2.3)$$

$$b_1(t) = \frac{2}{T} \int_{t-T}^t L(t') \sin(\omega_b t') dt', \quad (5.2.4)$$

$$a_0 = \frac{2\pi}{T}. \quad (5.2.5)$$

Then we feedback a phase-shifted version of this demodulated lift signal as the jet velocity, U_f with appropriate gain, K_p ,

$$U_f(t) = a_0 + K_p(a_1(t) \cos(\omega_b t + \phi) + b_1(t) \sin(\omega_b t + \phi)). \quad (5.2.6)$$




Figure 5.2.8: Maximum and minimum lift ($\square$) and its average over time ($\circ$) (top) and frequency (bottom) of phase-locked limit cycles at different phase-shift, ϕ , for (a) $\alpha = 40^\circ$ and (b) 50° .

where a_0 is the average value of the output U_2 , which can be prescribed as 0.5 to fix $C_p = 0.01$. We also adjust K_p periodically, such that the rms amplitude of U_1 remains steady and similar to that of open-loop control, i.e. U_1 varies from 0 to 1.

The configuration of our feedback control is shown in figure 5.2.7. Lift is fed back to the controller which has two parameters: demodulation frequency, ω_d , and the desired phase shift ϕ_d . The controller outputs a sinusoidal, U_1 that is phase shifted relative to the dominant frequency of the lift signal. The flow system outputs C_L , which has a frequency ω_L and a phase shift ϕ_L relative to the input signal U_1 .

If C_L is phase-locked to U_1 , the frequency of U_1 will always be the same as the frequency of C_L . However, if the demodulation frequency, ω_d is not equal to the frequency of the lift signal, ω_L , then ϕ_L will be different from ϕ_d (unless $\omega_d = \omega_L$ in which case $\phi_L = \phi_d$). Thus, it is necessary to add an integral part to the algorithm to adjust ω_d such that,

$$\omega_d^{k+1} = \omega_d^k + \beta(\omega_L^k - \omega_d^k). \quad (5.2.7)$$

We can adjust ω_d until it reaches ω_L , and thus obtain the exact desired phase shift and allow the frequency constant to be determined only by the flow. Then we have a robust compensator to explore different limit cycles that are phase locked at various ϕ at different α .

Figure 5.2.8 investigates the sensitivity of the lift and the frequency of the forced phase-locked limit cycles to the changes in the phase shift, ϕ at $\alpha = 40^\circ$ and 50° . Feedback was able to phase lock the flow at any desired phase shift after $2 \sim 5$ periods over a wide range of $-0.5 \leq \phi \leq 0.5$. At $\alpha = 40^\circ$, as shown in figure 5.2.8(a), FB₁ corresponds to the limit cycle phase-locked to the actuation at ϕ_{max} . However, the phase shift that achieved the highest average lift was not ϕ_{max} . An even higher-lift limit cycle was achieved near zero phase shift, resulting in as high as 83% increase in the average lift coefficient. A broad range of ϕ ($-0.28 \leq \phi \leq 0.06$) resulted in average lift that was higher than the maximum lift of the baseline flow, that is more than 45% in the average lift enhancement. At $\alpha = 50^\circ$, the highest average lift occurred near zero phase shift, and over a range of ϕ , $-0.3 \leq \phi \leq 0.16$ the actuation achieved at least 25% enhancement over the average lift of the natural flow. At both α 's, a larger range of negative phase shift contributed more to

lift enhancement than the positive phase shift. Particularly at $\alpha = 40^\circ$, there was a sharp decrease in the lift after $\phi = 0.06$ whereas the lift decrease was more gradual at the negative phase shift. Thus, forcing seems more effective as the newly forming LEV is pulled down by the TEV (lift-increasing phase). On the other hand, forcing seems the least effective after the maximum lift occurs; when the LEV sits closest to the plate and is pushed away by the growing TEV (lift-decreasing phase). As ϕ approaches 0.5 or -0.5 (out of phase), the forced flow results in the average lift similar to that of an unforced flow, but with a slightly smaller magnitude of oscillation in lift coefficient.




Figure 5.2.9: Comparison between open-loop control case (---), that phase locked the flow at the highest average C_L (denoted as OL₁ in figure 5.2(c)), and the corresponding feedback control case (—) (denoted as FB₁ in figure 5.2.8). Forcing frequency of this open-loop control is denoted as ω_{OL_1} , and the average output frequency of C_L of the feedback control is denoted as ω_{FB_1} .

Recall in figure 5(c), we observe a very small domain of attraction near $\omega_L = 0.8$ for the phase-locked limit cycle and the resulting limit cycle has a positive phase shift, $\phi \approx -0.3$. However, the phase-locked limit cycle achieved by this feedback have a wide range of frequencies, varying from 0.8 to 0.95 with the corresponding phase shifts ranging from -0.5 to 0.5 . These limit cycles were not achieved by any of the forcing frequencies of the open-loop control in figure 5(c). The feedback algorithm results in phase-locked limit cycles that are not attainable by the open-loop forcing.

Figure 5.2.9 compares the lift signal of the two limit cycles; the best open-loop case at $\alpha = 40^\circ$, denoted as OL₁ in figure 5(c) and the corresponding feedback case, denoted as FB₁ in figure 5.2.8. With open-loop control at fixed ω_L , the flow seems to lock onto the actuation at the highest average lift cycle during earlier periods, with its phase shift closer to ϕ_{max} . But after a couple of periods, ϕ drifts away from ϕ_{max} and the flow eventually locks onto the lower average lift cycle. On the other hand, the feedback compensator prevents ϕ from drifting away and sustains the phase at ϕ_{max} , producing higher average lift than the open-loop control. Thus, we can conclude that this feedback algorithm stabilizes the limit cycle with a significant lift enhancement that cannot be obtained with the open-loop control.

To ensure that the feedback is still required to sustain the achieved phase-locked limit cycle, FB₁ is investigated further. Feedback is turned off after the phase-locked limit cycle has been achieved for a long time, and the forcing signal is continued with the open-loop forcing at a fixed frequency, ω_L , as shown in figure 5.2.10. This behavior of unstable phase relationship has also been shown with a open- and closed-loop control model of an oscillating cylinder wake by Tadmor et al. (2004). Notice that when the forcing signal is continued with the actuation of $\omega_L = \omega_{OL_1}$, the flow drifts back to the previous open-loop limit cycle. When it is continued with actuation oscillating at $\omega_L = \omega_{FB_1}$, the average frequency of the previous feedback output signal, the flow displays a loss of phase-locking to this forcing frequency and it displays a pulling-out phenomenon. These results indicate that the feedback compensator was adjusting its forcing

corresponding to the change of output frequency from the flow, and that the feedback is still required to sustain the flow at the high-lift limit cycle.

The feedback algorithm stabilizes the limit cycle with a significant lift enhancement that are not attainable by the open-loop forcing. Furthermore, even with careful tuning of the forcing frequency, open-loop forcing cannot sustain this high-lift limit cycle. Thus, the feedback achieves high-lift unsteady flow states that cannot be achieved or sustained without it.




Figure 5.2.10: Continuation of feedback control case in figure 5.2.9 with open-loop control of ω_L .

5.2.4 Waveform optimization

While feedback has been successful in locking the flow onto sinusoidal forcing, resulting in high-lift limit cycles that would be otherwise unstable, it is not clear that sinusoidal forcing is optimal. Indeed, as discussed in section 2.2, recent studies have shown the efficacy of pulsed (low duty cycle) actuation (Amiray & Glauser 2002a; Woo et al. 2008; Cierpka et al. 2008; Greenblatt et al. 2008). In this section, optimal control theory is used to determine the optimal waveform for synchronization of vortex shedding in order to maximize lift. The adjoint of the linearized perturbed equations is solved backwards in time to obtain the gradient of the lift to changes in actuation (the jet velocity), and this information is used to iteratively improve the controls.

To maximize lift, we define a cost functional to be minimized

$$J = - \int_{t_0}^{t_1} \int_{\Omega} \tilde{p}(\phi(t), x, t) dx dt + C_p \int_{t_0}^{t_1} \int_{\Omega} \phi^2(t) dx dt \quad (5.2.8)$$

where t_0 and t_1 are the start and end times of the optimization horizon and Ω is surface of the body. ϕ is the control input, in this case $\phi(t) = U_{jet}(t)$. Again, $\tilde{p}$ is y component of forces on the plate calculated in the

immersed boundary projection method. The first term is the total squared lift over the optimization horizon. The second term penalizes the actuator amplitude in order to keep C_p to a value commensurate with the open-loop control discussed previously. The control weight, C_p , is determined by trial and error and is held fixed throughout the optimization.

At each iteration of the optimization, we modify the controls according to

$$\phi^{k+1} = \phi^k + r + g(\phi^k), \quad (5.2.9)$$

where $g(\phi)$ is the gradient of the cost function with respect to the controls, and r is the generalized distance determined iteratively (using Broyden's line minimization) to minimize the cost function. $g(\phi)$ is found by solving

$$g(\phi) = F_J^T \tilde{F} - 2C_p \phi, \quad (5.2.10)$$

where $\tilde{F}$ are the force unknowns in the linearized adjoint equations (Ahuja & Rowley 2008)

$$\mathcal{A}^T(\mathbf{q}) \mathbf{q}^* = \tilde{F}^*, \quad (5.2.11)$$

Here $\mathbf{q}^*$ are the adjoint variables (discrete circulation and forces) and $\tilde{F}^*$ is given by

$$\tilde{F}^* = \begin{bmatrix} \tilde{F}_1^* \\ \tilde{F}_2^* \end{bmatrix} = \begin{bmatrix} 0 \\ 2\tilde{p}_1^* \end{bmatrix}. \quad (5.2.12)$$

The adjoint operator requires the full flow field from the (forward) Navier-Stokes simulation at every time step. However, in order to save memory, we saved the flow solution only every few time steps and used a linear interpolation in time. Several test cases were done with a different number of time steps skipped, including a case where the solution was saved at every time step, and no significant differences were noted between them.

All optimizations used zero control ($\phi = 0$) for the first iteration ($k = 1$) on each optimization horizon. At each iteration, we required roughly ten full Navier-Stokes simulation to perform the line minimization (to find r).

Optimization was done over a horizon $T = [t_0, t_1]$, where the horizon, T , is long enough to overcome transient effects, but limited by the computational effort to perform all the required iterations and to tune the control weight. We found for this problem that after about two periods the controls converged to an approximately periodic signal with each period corresponding to a vortex shedding cycle. A horizon of 6 periods gave the results presented below, and tests showed that the results were not very sensitive as the horizon was varied from about 5 to 8 periods. Once the iteration of the optimization converges, the control near the end of each optimization horizon (transient of adjoint simulation) is discarded and the optimization is begun anew. This process is depicted in figure 5.2.11.

Optimization results in a periodic control waveform after a couple of transient periods. As shown in figure 5.2.14, this periodic optimal waveform is not sinusoidal, but rather composed of two distinct pulses per shedding cycle. The larger, later pulse is roughly in phase with the maximum lift. This result will be discussed further below after feedback is designed to achieve highest-lift, phase-locked shedding cycle with a given optimal or sinusoidal control waveform. Different values of control weight, C_p , results in a periodic control waveform with similar features, but with different average control input, thus different values of C_p . For example, $C_p = 0.3$ gives the results shown in figure 5.2.14 where C_p is about two times lower than that used for the sinusoidal forcing, but comparable lift is achieved. It should be noted that, although we cannot be assured that this is a global optimal, we observed similar results with different values of control weight and different initial controls (zero, constant, or sinusoid).




Figure 5.2.15: Average lift of optimized control ($\square$) and closed-loop sinusoidal forcing ($\circ$) at different values of C_μ at $\alpha = 40^\circ$.




Figure 5.2.16: Comparison between feedback control cases with optimized waveform at $\alpha = 40^\circ$: $N_k = 10$ (dashed) and $N_k = 4$ (solid).

resulting in larger oscillations in the magnitude of force on the plate. Upstream actuation at both leading and trailing edges led to the enhancement of average lift.

We next examined closed-loop control aimed at attenuating and synchronizing vortex shedding. We applied BPOD models to the supercritical flow at $Re = 100$ to obtain reduced-order, input-output models for control design. The reduced-order models used a novel approach for applying BPOD to unstable systems where linear stability eigenfunctions are used to describe the unstable modes, and BPOD is then applied to the system projected onto a stable manifold. Feedback control was designed using a LQR approach with both full-state and observer-based feedback. Using these techniques, the control laws were able to suppress vortex shedding even in a nonlinear context when applied in the DNS.

In open-loop forcing, it was observed that the hydrodynamics were phase locked to the actuation at lower α , but at sufficiently high α , actuation near the natural shedding frequency led to the excitation of a subharmonic resonance. The subharmonic limit cycle consisted of several periods with a different period-averaged lift over each actuation period. When the forcing signal was at a particular phase shift relative to the lift signal (in-phase at $\alpha = 50^\circ$), the actuation achieved the highest lift. However, the succeeding period became slightly out of phase and the lift decreased. With the goal of obtaining forced limit cycles with the




Figure 5.2.17: Maximum and minimum lift ($\square$) and average lift ($\diamond$) of phase-locked limit cycles at different phase shift with optimized waveform ($N_k = 10$) at $\alpha = 40^\circ$. Maximum and minimum lift of baseline ($\circ$) case is shown as a reference.

maximum average lift, we designed a feedback algorithm to phase lock the flow at the desired shedding period, particularly at the phase shift that gave the highest period-averaged lift. The control utilizes the modulation and demodulation of the lift signal in order to output a sinusoidal forcing with a specified phase shift relative to the lift signal. The compensator was able to phase lock the flow at the desired vortex shedding limit cycle that was not sustainable with any of the open-loop periodic forcing.

Finally, a gradient-based (adjoint) approach was applied in a receding-horizon setting to optimize the control waveform in order to maximize the lift. The optimized control waveform is not sinusoidal, but rather is pulse-like, with each period composed of two distinct pulses (a primary, as well as a smaller earlier pulse). As a result, the optimal control achieves comparable lift with 2 times lower C_μ value as the sinusoidal forcing case. However, if applied in open loop, the waveform fails to phase lock onto the optimal waveform, degrading the lift performance. We generalized the feedback procedure that had been applied to sinusoidal waveforms to now allow the compensator to determine the phase of the output (lift) with respect to the individual Fourier components of the optimized waveform. Feedback control was then able to reproduce the high-lift limit cycle from the optimization, but starting from an arbitrary phase of the baseline limit cycle. This also allowed us to phase lock an essentially arbitrary waveform, enabling us to investigate the sensitivity of the flow to the phase shift and other features of the optimized waveform.

6 Three-dimensional numerical simulations

In this section, we report on progress towards developing model-based control of vortex shedding in three-dimensional simulations of low aspect ratio flat plates and airfoils at high angle of attack. While there are straightforward generalizations of all the techniques described in the last section, their application to three-dimensional flow is limited by the much larger computational resources and time required for their solution. Nevertheless, we have performed extensive simulations of natural and open-loop actuated flows on low aspect ratio wings at low Re , in order to examine the detailed flow physics and the lift-enhancing effects of actuation. Parametric studies have been carried out to examine the effects of aspect ratio, planform shape, actuator placement, amplitude, and frequency. Finally, we report at the end on preliminary studies implementing closed-loop extremum seeking control to three-dimensional simulations.

6.1 Natural three-dimensional flows

For purely translating low-aspect-ratio wings, Torres & Mueller (2004b) have experimentally measured the aerodynamic characteristics of low-aspect-ratio wings at Reynolds numbers around 10^4 . Aerodynamic performance (lift, drag, pitching moment, etc) of various planforms was considered over angles of attack (α) of 0° to 40° and aspect ratios (AR) of 0.5 to 2 and was observed to be quite different from that of low-aspect-ratio wings in high-Reynolds-number flows. They concluded that the most important parameter that influences the aerodynamic characteristics is the aspect ratio. Transient studies have also been conducted by Freymuth et al. (1987) by using smoke to visualize the start-up flows around low-aspect-ratio airfoils. A qualitative insight into the three-dimensional formation of wake vortices was presented. The experiments by Ringuette et al. (2007) extensively studied the wake vortices behind low-aspect-ratio plates but only at $\alpha = 90^\circ$. On the numerical side, two-dimensional simulations around translating wings were performed by Hamdani & Sun (2000). Also, studies by Mittal & Tezduyar (1995) and Corry & Vierendeels (2006) considered the three-dimensional flows around translating low-aspect-ratio planforms but focused mostly on those at low angles of attack. For wings at post-stall angles of attack, unsteady separated flows and vortex dynamics behind low-aspect-ratio wings in pure translation are still not well-documented.

To extend the previous studies, and provide an understanding of the natural flow for the control efforts reported later, we use numerical simulations to examine the aerodynamics of impulsively started low-aspect-ratio flat-plate wings under pure translation at Reynolds numbers of 300 and 500. We focus on the unsteady vortex dynamics at post-stall angles of attack, and study the influence of aspect ratio, angle of attack, and planform geometry on the wake vortices and the resulting forces on the plate. These Reynolds numbers are high enough to induce separation and unsteadiness in the wake but low enough for the three-dimensional flow field to remain laminar. The regime also includes, for a range of angles of attack, the critical Reynolds numbers at which the flow first becomes unstable to small disturbances. Depending on the aspect ratio, angle of attack, and Reynolds number, the flow at large time reaches a stable steady state, a periodic cycle, or aperiodic shedding. For cases of high angles of attack, an asymmetric wake developed in the spanwise direction at large time. The present results are compared to higher Reynolds number flows. Some non-planform planforms are also considered to examine the difference in the wakes and forces. After the impulsive start, the time at which maximum lift occurs is fairly constant for a wide range of flow conditions during the initial transient. Due to the influence of the tip vortices, the three-dimensional dynamics of the wake vortices are found to be quite different from the two-dimensional von Kármán vortex street in terms of stability and shedding frequency.



Figure 6.1.1: A typical computational domain showing the top-port side of the wake around a rectangular flat plate of $AR = 2$. The spatial discretization of this computational domain is shown for every 5 cells for the x - and y -directions and 4 cells for the z -direction.

6.1.1 Simulation setup

The numerical method and its validation were previously presented in section 4. Simulations are performed in a large rectangular box typically of size $[-4, 6.1] \times [-5, 5] \times [-5, 5]$ in the streamwise (x), vertical (y), and spanwise (z) directions. Typical grid size ranges from $125 \times 55 \times 80$ to $200 \times 88 \times 128$ with the smallest resolution of $\Delta x = 0.025$ for the case of $AR = 2$ and much larger sizes were used for simulations of flows around higher aspect ratio plates. Grid stretching is applied in all directions with finer resolution near the plate to capture the wake structure as illustrated in figure 6.1.1. Extensive studies have been performed in two and three dimensions to ensure that the present choice of grid resolution and domain size does not influence the flow field in a significant manner (previously reported in Thira et al. (2007) and Taira (2008)).

Boundary conditions along all sides of the computational boundary, $\partial\mathcal{D}$, are set to uniform flow ($U_\infty, 0, 0$) except for the outlet boundary where a convective boundary condition ($\frac{\partial}{\partial x} = U_\infty \frac{\partial}{\partial x} = 0$) is specified. Inside the computational domain, a flat plate is positioned with its center at the origin. This flat plate is instantaneously materialized at $t = 0^+$ in an initially uniform flow to model an impulsively started translating plate. Computations are advanced in time with a time step such that the Courant number based on the free-stream velocity obeys $U_\infty \Delta t / \Delta x \leq 0.5$. Both the initial transient and the large-time behavior of the flow are considered.

6.1.2 Dynamics of wake vortices behind rectangular planforms

We first consider the formation and evolution of wake vortices behind rectangular flat plates of $AR = 1, 2$, and 4 at $Re = 300$. Snapshots in figure 6.1.2 show the corresponding flow fields at a high angle of attack of $\alpha = 30^\circ$. The wake vortices are visualized with two different iso-surfaces. First, constant vorticity surface ($|\omega|_2 = 3$) is shown in light gray to capture the vortex sheets. Second, the second invariant of the velocity gradient tensor (Q), known as the Q -criterion or Q -value, is used to highlight the vortex cores depicted by the dark gray surfaces ($Q = 3$). Positive Q -values give prominence to regions of high swirl in comparison to shear to represent coherent vortices (Hunt et al. 1988). In incompressible flows, the Q -value is provided by $Q = \frac{1}{2}(|\Omega|^2 - |S|^2)$, where $\Omega \equiv \frac{1}{2}(\nabla \mathbf{u} - (\nabla \mathbf{u})^T)$ and $S \equiv \frac{1}{2}(\nabla \mathbf{u} + (\nabla \mathbf{u})^T)$.

Plates are impulsively translated in an initially quiescent flow, generating strong vortex sheets on the top and bottom surface of the plates at $t = 0^+$. Thereafter, the flow separates from the leading edge and

tips, which act as strong sources of vorticity. Vorticity is then fed into the flow as vortex sheets that roll up into the leading-edge, trailing-edge, and tip vortices. As the leading-edge vortex forms and the trailing-edge vortex advects downstream, the plate experiences a large increase in lift. This enhanced lift is generated by the low-pressure core of the initial leading-edge vortex and is an added effect on top of the lift achieved at large time (also observed in two-dimensional flows by Dickinson & Götz (1993b)). Note that the initial topology of the wake structures are the same for all finite AR considered here and the phenomena resemble dynamic stall observed behind high-aspect-ratio wings undergoing a swift pitch-up (Carr 1988b).

As features from the initial transient lose their effect on the plate through advection and diffusion, the wake behind the plate becomes strongly dependent on the aspect ratio. At $\alpha = 30^\circ$, the $AR = 1$ case slowly reaches a steady state with a pair of strong counter-rotating tip vortices that cover the entire span of the plate. The vortex sheet created from the leading edge is kept attached to the plate due to the downward induced velocity from the tip vortices.

For a plate of $AR = 2$, the vortex sheet emanating from the leading edge rolls into a leading-edge vortex that accumulates spanwise vorticity over time. The tip vortices are not strong enough to keep the leading-edge vortex attached. Around $t \approx 8$, two consecutive leading-edge vortices start to pinch off from the plate. As the detachment takes place, the disconnected vortices start to interact with the tip vortices. This interaction results in the loss of the columnar structures initially maintained by the tip vortices and reduces the downward induced velocity into the vortical structure residing above the top surface of the plate. Hence, once the initial leading-edge vortices are separated, consecutive formation of the leading-edge vortices (hairpin vortices) by the roll-up of the leading-edge vortex sheet occurs farther and at a higher position. Additionally, weakening of the tip vortices allows the trailing-edge vortex sheet to morph into hairpin vortices that shed. The nonlinear interaction of the wake vortices results in an unsteady aperiodic flow at large time. While it is not apparent from figure 6.1.2, there is slight asymmetry in the spanwise direction that contributes to the aperiodic nature of the shedding. Discussion on this asymmetry is offered later in § 6.1.5.

For the largest finite aspect-ratio plate considered ($AR = 4$), the weaker influence of the tip vortices across the span results in more strongly pronounced periodic shedding of the leading- and trailing-edge vortices. The shedding frequency (non-dimensionalized as the Strouhal number with the frontal projection of the chord) for $AR = 4$ is found to be $St \equiv f c_{\text{area}} / U_\infty = 0.12$. In contrast, the two-dimensional shedding frequency for the same Reynolds number and angle of attack is $St = 0.16$.

Around $AR = 3$, the vortical structures from the leading and trailing edges start appearing to separate into two cells across the span. The cellular pattern referred to as stall cells becomes more apparent for $AR = 4$ where a pair of hairpin vortices are generated from the leading edge and another pair is created from the trailing edge resulting in a release of four hairpin vortices per shedding cycle. Such flow features were also reported on the top surface of the airfoils with oil film and tuft visualizations by Winkelmann & Barlow (1980) and Yen & Katz (1998), respectively, at $Re = O(10^5)$. While we do not notice features of the stall cells directly on the top surface, we find qualitative agreement for the number of cells observed some short distance into the wake.

Despite the interactions between the leading-edge and tip vortices, these vortices remain distinct without merging for all three-dimensional cases. Due to the existence of the right-angled corners on the rectangular plates, the vortex sheets thin out near these regions and the sheets roll up into individual core structures of leading-edge and tip vortices. The separation of the vortical structures indicates a lack of convective vorticity flux in the spanwise direction (i.e. from the mid-span to the tips). Such transport has been suggested to stabilize the leading-edge vortex for flapping wings (Birch & Dickinson 2001; Birch et al. 2004b). For the translating rectangular wings, there is no mechanism to relieve the vorticity being fed into the leading-edge vortex other than diffusion and shedding of the vortical structures. The effects of removing sharp corners by



Figure 6.1.3: Top views of the wake vortices behind a rectangular wing of $AR = 2$ at $\alpha = 40^\circ$ from smoke visualizations at $Re = 5200$ and present results at $Re = 500$ with iso-contour of $\|\omega\| = 5$. Smoke visualizations are from Freymuth et al. (1987); reprinted by permission of the American Institute of Aeronautics and Astronautics.

using different planform geometries are discussed later in section 6.1.6.

6.1.3 Flows at higher Reynolds number

Flows behind rectangular plates at $Re = 500$ are also simulated and are found to be similar to the ones presented here for $Re = 300$. With the larger Reynolds number, the wake vortices are less diffused but the topology of the vortical structure are qualitatively similar, which was also noted by Dong et al. (2006) for flows around flapping foils for $Re = 100$ to 1000 .

The geometries of the wake vortices in low-Reynolds-number flow ($Re = O(100)$) at early times following the impulsive start also resemble those in flows of much higher Reynolds numbers, due to the fact that the viscous time scale ($t_{\text{vis}} \sim c^2/\nu$) is much larger than the time scale associated with convection or acceleration ($t_{\text{conv}} \sim c/U_\infty$ or $t_{\text{acc}} \sim c^2/a$), respectively.

Impulsive flow over a plate of $AR = 2$ at $Re = 500$ and $\alpha = 40^\circ$ is simulated and is compared to the smoke visualizations of vortices under a constant acceleration from quiescent flow in a starting wind tunnel (Freymuth et al. 1987) as shown in figure 6.1.3. Reynolds number for this experiment is defined with the constant acceleration, the chord, and the kinematic viscosity as $Re_a = d^2 a^2 / \nu^2 = 5200$ following the non-dimensionalization by Freymuth et al.

As acceleration of immersed boundaries contributes to the generation of vorticity (Herrmann 1989), the formation of vortices behind a plate under impulsively translation and constant acceleration cannot be directly compared. However, the formation of start-up vortices should be qualitatively similar at early times before viscous effect significantly influences the flow and the induced velocity of each wake vortex becomes large. The formation of the start-up vortices is illustrated by the isosurfaces in figure 6.1.3 with smoke visualization and the vorticity norm iso-surface.

For the case of constant acceleration, a characteristic velocity of $u_0 = \nu^{1/2} a^{1/2}$ is used to non-dimensionalize the temporal variable. Accordingly, the flow fields are compared at the non-dimensional times of tU_0/c and




Figure 6.1.2: Top-port views of the wake vortices behind rectangular plates of $AR = 1, 2$, and 4 at $\alpha = 30^\circ$ and $Re = 300$. Shown are the iso-surface of $\|\omega\| = 1.5$ in light gray with vortex cores highlighted by the iso-surface of $Q = 3$ in dark gray.

$d^2 a^2 / c^2 \nu^2$ for the simulations and the experiments, respectively. In figure 6.1.3, it can be seen that the formation and evolution of leading-edge and tip vortices are in good agreement between the experiment and the numerical solution despite the difference in the Reynolds number and the velocity profile to which the plate is subjected. The exact location of the trailing-edge vortex from the smoke visualization is difficult to pinpoint but is found to be in accord by its faint trail of smoke at earlier times. The simulation is able to capture even the thin layer of vortex sheet emanating from the leading edge, which would correspond to the region directly downstream of the leading edge that is not visualized by the smoke. Dominant flow features at early times in high-Reynolds-number flows are captured even with the present low-Reynolds-number flow simulations.

6.1.4 Force exerted on the plate

Unsteady forces on accelerating airfoils at low Reynolds numbers have been considered for two-dimensional flows by Dickinson & Götz (1993b) and Pullin & Wang (2004). In this section, we consider the forces exerted upon the plate with the three-dimensional wake vortices both immediately after the impulsive start and also at large times.

Representative lift and drag on rectangular plates from the present simulations are presented in figure 6.1.4 for $Re = 300$. Here results for angles of attack of $\alpha \in [0^\circ, 60^\circ]$ and aspect ratios of 2 and 4 as well as the two-dimensional cases are shown for $t \in [0, 70]$. At $t = 0^\circ$, the impulsive start imposes infinite acceleration on the airfoil in the streamwise direction and results in infinite initial drag (not shown for graphical clarity). Subsequently, lift starts to increase as accumulation of spanwise vorticity instigates the formation of the leading-edge vortex. This increase in lift continues to about $t = 1.7$ to reach its maximum. The time to reach maximum lift is observed to be fairly constant in the case of finite aspect ratio wings over most of the angles of attack considered here at low Reynolds numbers. The universality of this number is discussed in detail later.

After the initial start-up, lift is reduced by as much as half of the maximum value at large time, as shown in figure 6.1.4. Depending on whether the wake at large time becomes steady or unsteady, the corresponding force coefficients reach constant or fluctuating values.

It should be noted that the three-dimensional flows of consideration are vastly different from the two-dimensional case, where one observes periodic shedding of the leading- and trailing-edge vortices crossing the von Kármán vortex street. Due to the absence of the tip vortices, the two-dimensional flow exerts a strikingly larger fluctuation in force per unit span as shown in figure 6.1.4.

The effect of aspect ratio on the forces is considered by comparing the maximum lift during the transient and the time-averaged forces at large time. These values for wings of $AR = 1, 2$, and 4 , as well as the two-dimensional case, are presented in figure 6.1.5 accompanied by their inviscid limits. Stronger influence of downwash from the tip vortices results in reduced lift for lower aspect-ratio plates. For the limiting case of two-dimensional flow, the maximum lift is much higher due to the absence of tip effects (figure 6.1.5a). It is interesting to note that the maximum lift achieved upon after the impulsive start is comparable or higher than the three-dimensional inviscid limit for low-aspect-ratio straight wings in incompressible flow (Helmholtz 1942):

$$C_L = \frac{2\pi\alpha}{\sqrt{1 + (2/AR)^2} + 2/AR} \quad (6.1.1)$$

This limit is derived from the lifting surface theory for elliptic wings and is shown to be in remarkable agreement with wings of $AR < 4$. Lift for rectangular wings of $0.5 \leq AR \leq 6$ is accurately predicted with this equation as shown in Anderson (1999).




Figure 6.1.4: Force history on rectangular flat plates of different AR for a range of angles of attack at $Re = 300$. Lift and drag coefficients are shown on the left and right, respectively.




Figure 6.1.5: Characteristic coefficients for rectangular plates at $Re = 300$: (a) maximum lift; (b) time-averaged lift coefficient at large time; (c) time-averaged drag coefficient at large time; and (d) time-averaged lift-to-drag ratio at large time for $AR = 1$ ($\circ$), $AR = 2$ ($\square$), $AR = 4$ ($\triangle$), and 2D ($\diamond$). Overlaid are the two- (—) and three-dimensional (---) inviscid lift limits.

We observe agreement between the time-averaged low-Reynolds-number lift at large time and the above inviscid model at low angles of attack ($\alpha \lesssim 10^\circ$), where the flows are still attached (figure 6.1.5b). The difference at low angles of attack can be attributed to viscous effects. However, once the flow separates from the plate at higher α , the inviscid approximation is no longer able to model the lift behavior. The high value of $\alpha_{C_{Lmax}}$ (angle of maximum lift) was also reported for low-aspect-ratio wings at $Re \approx 10^3$ by Torres & Mueller (2004b). We mention that the difference between the maximum ($\max C_L$) and the average ($\text{avg } C_L$) lifts is the lift enhancement generated by the initial leading-edge vortex.

The average drag values at large time ($\text{avg } C_D$) for pre-stall angles of attack increase with decreasing aspect ratio (figure 6.1.5c). However for higher angles of attack, drag coefficients are significantly smaller for finite-aspect-ratio wings in contrast to their two-dimensional analog. It is interesting to note that both average lift and drag coefficients become larger past $\alpha = 20^\circ$ for $AR = 2$ wings than those of $AR = 1$, most likely due to the difference in the behavior of the wake at large time (discussed later in §6.1.5).

At $Re = 300$, the viscous stress has a significant influence on the drag experienced by the wing, especially at low angles of attack. In comparison, at $Re = 10^3$ and $AR = 1$, Torres & Mueller (2004b) reports $C_D = 0.025$ and 0.11 for $\alpha = 0^\circ$ and 10° , respectively. At $\alpha = 15^\circ$, $C_D = 0.24$ is recorded by Torres and Mueller, a value close to what is measured in the current study also (figure 6.1.5c). Hence, we argue that past this angle of attack, pressure drag is the main cause of drag.

Shown also in figure 6.1.5d are the average lift-to-drag ratios at large time, $\text{avg } (C_L/C_D)$, which are larger for higher aspect-ratio wings. The ratios peak at higher angles of attack for low aspect-ratio plates. While the two-dimensional $\text{avg } (C_L/C_D)$ reaches its maximum around $\alpha = 10^\circ$, that of $AR = 1$ is achieved near $\alpha = 20^\circ$. This may suggest favorable operating conditions at higher angles of attack for the low-aspect-ratio wings. For high angles of attack ($\alpha > 40^\circ$), the lift-to-drag ratio for different aspect ratios coalesce to the same value.

While it is not shown in this section, forces measured at $Re = 500$ are found to be quantitatively and qualitatively similar to those at $Re = 300$. Interesting differences between the two Reynolds numbers are observed in the stability of the wake, which is described later in §6.1.5.

Next, we call attention to the time at which the maximum lift is achieved. We denote this time by t^* and present its value on figure 6.1.6 for $Re = 300$ and 500 . It is found that, for the considered aspect ratios and angles of attack, t^* is fairly constant around 1.7 (a value between 1.25 and 2.25) because the profiles of the leading-edge vortices are similar among all cases. As the accumulation of spanwise vorticity generated by the leading edge contributes to the growth of the leading-edge vortex, there is reminiscence to the formation number used to describe the time at which vortex rings can no longer grow larger in strength (Gharib et al. 1998). Since the formation number is found to be a universal quantity for a variety of flows (e.g. Jeon & Gharib 2004, Milano & Gharib 2005), it is not surprising that t^* is also fairly constant for the three-dimensional cases considered here. In the case of two dimensional flow, we observe a wider range of t^* between 1.3 and 2.4 for $\alpha < 45^\circ$. At higher α , a second local maximum starts to emerge for the two-dimensional flow lowering t^* significantly.

The side force (F_y) remained zero for all cases that reached steady or periodic unsteady flows. However, for aperiodic flow cases observed at high angles of attack, the wake became asymmetric about the mid-span and exerted side forces upon the plate. This unsteady side forces were an order of magnitude smaller than the dominant lift and drag forces experienced by the plate. For all cases considered in this paper (later summarized in figure 6.1.8), it was observed that side forces have small magnitudes of $C_F = F_y / (\frac{1}{2} \rho U^2 A) < 0.02$. The genesis of these side forces is discussed in the next section.




Figure 6.1.6: Time at which lift achieves the maximum, t^* , after an impulsive start for different angles of attack at (a) $Re = 300$ and (b) $Re = 500$. Symbols denote cases for $AR = 1$ ($\circ$), $AR = 1.5$ ($\triangledown$), $AR = 2$ ($\square$), $AR = 3$ ($\diamond$), $AR = 4$ ($\triangle$), and 2D ($\bullet$). The mean averages are shown with solid lines. The inset figures are the corresponding histograms.

6.1.5 Large-time behavior and stability of the wake

Here we consider the behavior of the wake behind rectangular plates at large time. After the initial transient generated by the impulsive start settles down, the wake reaches one of the three states: (i) a stable steady state, (ii) a periodic unsteady state, or (iii) an aperiodic unsteady state. Examples from each one of these states were presented in section 6.1.2. In this subsection, we consider a wide range of parameters ($\alpha \in [0^\circ, 60^\circ]$, $AR \in \{1, 1.5, 2, 3, 4\}$, and $Re \in \{300, 500\}$) as well as the two-dimensional flows to survey the stability of the wake at large time. The lift histories, such as the ones in figure 6.1.4, are analyzed with Fourier transform to detect any dominant shedding frequencies as shown in figure 6.1.7 for an example of $Re = 500$ and $AR = 3$. Depending on the angle of attack, the shedding can occur with a dominant shedding frequency (periodic) or with no clearly recognizable frequency (aperiodic). The dominant frequency in the case of figure 6.1.7 is found to be $\mathcal{S} = f_{\text{shed}}/U_\infty = 0.12$ for all spectra of the periodic shedding cases. The corresponding Strouhal number for the two-dimensional flow at $Re = 500$ is slightly higher at $\mathcal{S} = 0.14 - 0.16$.

The wake stability is summarized in figure 6.1.8, which maps α against AR . These two parameters were found to be the two most important parameters in determining the stability of the wake at large time. Suggested boundaries between different flow regimes are drawn based on the data points collected from numerical experiments. The shaded regions correspond to flow conditions that would arrive at a steady state. Such flow can be either attached at small α or fully separated at moderately high α . The steady state is achieved over a wider limit with lower aspect ratios since the tip vortices are able to provide a downward induced velocity across a larger extent to prevent the wake vortices from shedding.

As we consider higher angles of attack, the flow exhibits periodically shedding hairpin vortices generated by the leading and/or the trailing edges. This flow profile is observed for the whole region left of the dashed




Figure 6.1.7: Power spectra of the lift force for a rectangular plate of $AR = 3$ in flows of $Re = 500$ at various angles of attack.




Figure 6.1.8: Stability of the wake for a range of α and AR at (a) $Re = 300$ and (b) $Re = 500$. Symbols of $\bullet$, $\square$, and Δ denote steady, unsteady periodic, and unsteady aperiodic wakes at large time. The shaded area and the dashed line approximately represent the region of stability and the transition from periodic to aperiodic shedding, respectively. Shown at the top are the two-dimensional (2D) findings.

149

For $Re = 500$, the steady flow profile is achieved for a smaller range of angles of attack and aspect ratios compared to the $Re = 300$ case. Similar trend holds for the periodic shedding case. With increasing Reynolds number, it is expected that the wake exhibits strong interaction between the leading-edge and tip vortices resulting in aperiodic/asymmetric flows for a wider combination of angles of attack and aspect ratios. At higher Reynolds numbers, it may be possible to observe changes in the shape of the stability boundary between the periodic and aperiodic shedding regimes.

6.1.6 Non-rectangular planforms

Stable attachment of the leading-edge vortices on flapping wings has been observed to provide enhanced lift for prolonged duration (Birch et al. 2004a; Peletia et al. 2006). Shedding of these vortices is prevented by releasing the spanwise vorticity through convective transport from the root to the tip of the wings induced by wing rotation. Hence, a continuously connected vortical structure formed by the leading-edge and tip vortices are observed for flapping or revolving wings.

For rectangular plates in pure translation, we have shown earlier that the initial leading-edge vortex detaches in a similar fashion to dynamic stall. The leading-edge and tip vortices remain as separate vortical structures and do not provide a mechanism for the spanwise vorticity to be released other than shedding. In order to prevent or delay the shedding of the leading-edge vortex, we consider the use of curved or angled leading edges to induce flows along the leading edge. Flows around elliptic, semicircular, and delta-shaped planforms are simulated at $\alpha = 30^\circ$ and $Re = 300$ in comparison to the flow around a rectangular plate of $AR = 2$ presented earlier. The geometries of the elliptic, semicircular, and delta-shaped planforms are chosen with $AR = 2, 4/\pi$, and 4, respectively, whose mean chord lengths ($c \equiv A/b$, where b is the wing span) are used to non-dimensionalize all spatial variables. For the delta wing, the sweep angle is set to 45° .

Wake structures behind the non-rectangular planforms are shown in figure 6.1.10 after the impulsive start with the corresponding forces in figure 6.1.11. For the elliptic and semicircular cases, there are no discontinuities in the vortex sheet that emanates from the leading edge to the tips, unlike the sharp separation of the vortical structures around the corners on the rectangular planforms. The curved leading edge encourages spanwise transport of vorticity into the tip vortices. Hence to some extent the shedding of the leading-edge vortical structure is delayed. While the leading-edge hairpin vortex first detaches around $t = 8$ for the rectangular planform, such separation is observed at a later time ($t = 15$) for the elliptic and semicircular plates. This gentle detachment allows the forces to decay smoothly instead of generating fluctuations, as seen for the rectangular planform around $t = 8$ due to the detachment of two consecutive leading-edge vortices. At later time, both the flows behind the elliptic and semicircular wings exhibit periodic shedding, which is different from the rectangular case.

A steady state is achieved by the flow around the delta-shaped planform. There is a clear distinction of the left and right vortex sheets from the nose of the plate, but the roll-up of the vortex sheets forms a stable wake structure which in turn attains steady lift and drag. The absence of wing tips for this planform allows the vortex sheets from the leading edge to roll up and convect downstream in a very stable manner (see figure 6.1.10). For $Re = 300$, the size of initial leading-edge vortices is larger compared to ones from higher Reynolds number flow (Gresel et al. 2005). One can observe transient behavior of the wake until $t \approx 15$. Beyond this point in time, there is some unsteady shedding of small vortical structures behind the rolled up vortices. However, the wake and the forces do not change much past $t \approx 15$ approaching the steady state.

We observe a relative increase in transport of the spanwise vorticity around the non-rectangular planforms in comparison to the rectangular planform as illustrated by iso-surfaces of $|\mathbf{u} \cdot \nabla \mathbf{u}|$ in figure 6.1.12. It should be noticed that there is an absence of transport of $\mathbf{u}_\infty$ near the leading edge for the rectangular wing in contrast to the semicircular and delta-shaped planforms. The force histories presented in figure 6.1.11

151



Figure 6.1.9: Top view of the asymmetric wake at large time behind a rectangular plate of $AR = 2$ at $\alpha = 40^\circ$ and $Re = 500$. Vortices are highlighted with iso-contours of $Q = 2.5$. The flow is directed from left to right and the wing is shown in black.

line in figure 6.1.8. The change in the dynamics between the shaded (sub-critical) and unshaded (super-critical) regions can be viewed as an extension of the two-dimensional stability boundary. We claim that the change in the dynamics is attributed to a Hopf-bifurcation, as shown by Ahuja et al. (2007) for the two-dimensional case. For lower aspect ratios, the vortex sheet emanating from the trailing edge forms and sheds hairpin vortices aperiodically. The narrow region with $AR \lesssim 2$ in figure 6.1.8 corresponds to such flow states. The same region with higher aspect ratio of $AR \geq 2$ shows shedding of both the leading- and trailing-edge vortices alternately in a periodic fashion (for instance the case of $Re = 300$, $AR = 4$, and $\alpha = 30^\circ$ in figure 6.1.2).

With further increase in the angle of attack, the tip vortices become more vertically aligned. Essentially, the wake now is comprised of four vortices of similar strengths, namely the leading-edge and trailing-edge vortices as well as a pair of tip vortices. The tip vortices strongly interact with the leading- and trailing-edge vortices and suppresses the dominant shedding frequency. The transition from periodic to aperiodic flow is illustrated in figure 6.1.7 at the peak for the power spectrum at $St = 0.12$ becomes no longer observable for $\alpha > 30^\circ$. In figure 6.1.8, this aperiodic unsteady state corresponds to the region right of the dashed line.

The aperiodic flows are found to be asymmetric in the spanwise direction with respect to the mid-span plane. As the wake becomes asymmetric, the wake vortices apply side forces onto the wing and the flow field. The combination of the asymmetry and the nonlinear interaction amongst the leading-edge, trailing-edge, and tip vortices give rise to the aperiodic nature of the flow. An example of an asymmetric wake is shown for a rectangular plate of $AR = 2$ at $\alpha = 40^\circ$ and $Re = 500$ (the wake for the same case at earlier time is shown in § 6.1.3). The side force for this case has a magnitude of $|C_y| < 0.01$ with a frequency content (no dominant shedding frequency) similar to those low frequency contents on figure 6.1.7. Asymmetry is not observed for steady or periodic unsteady flows.

For much larger aspect ratios than those considered here, the wake most likely develops into either a stable steady state or a periodic shedding profile. However, the actual three-dimensional flow with infinite span would probably not be purely two-dimensional, as seen for three-dimensional flows around an airfoil cylinder of infinite span (Braza et al. 2001). Spanwise perturbations can induce the eruption of spanwise vorticity and the corresponding spanwise undulations. Hence, formation of cellular vortical patterns (stall cells) can be observed directly above the top surface, similar to the structures seen in the $AR = 4$ case (figure 6.1.2) and those previously reported by Winkelmann & Barlow (1980) and Yoon & Katz (1998).

The stability of the wake is also influenced by the Reynolds number. In figure 6.1.8, we notice that

150




Figure 6.1.10: Top-port views of the wake vortices behind different planform geometries at $\alpha = 30^\circ$ and $Re = 300$ with the iso-surface of $|\mathbf{u}|u| = 3$ in light gray with vortex cores highlighted by the iso-surface of $Q = 3$ in dark gray.

152




Figure 6.1.11: Time traces of lift and drag coefficients for rectangular (—), elliptic (---), semicircular (·····), and delta-shaped (- · - ·) planforms at $Re = 200$ and $\alpha = 30^\circ$. The right two figures magnify the corresponding (early-time) transient behavior in the left figures.

show that the time of maximum lift is somewhat delayed to $t^* \approx 2$ for the elliptic and semicircular wings, in comparison to $t^* \approx 1.7$ for the rectangular plate of $AR = 2$. Nonetheless, the plate experiences a drop in lift due to the separation of the leading-edge vortices later in time. This detachment is caused by the insufficient spanwise transport (release) of spanwise vorticity to sustain a stable attachment of the leading-edge vortex under pure translation.

While the vortical flows are different for various planform geometries, the lift and drag exerted on the wings do not show significant variations in figure 6.1.11. This is most likely due to the viscous nature of the flows at this low Reynolds number. In addition, the similar aspect ratios considered here may be responsible for the similar values in forces. We also observe differences in the behavior and stability of the wakes. The stability diagram presented in figure 6.1.3 for the rectangular plate would not carry over to the cases of non-rectangular planforms due to the difference in the influence of the tip vortices.

For flapping wing aerodynamics, the wing kinematics restricts the wing stroke and periodicity. In such cases, the travel distance even at the wing tip is much smaller and is about $\pi c/kc$. Even a small increase in aerodynamic performance or stable flow features can be of benefit. Although the added effect of rotation or spanwise flows are not taken into account in this study, the wing planform can be of importance as we have discussed in the non-rectangular cases, at least without regard to structural constraints or maneuverability.

153



Figure 6.2.1: A flow control setup shown for an example of actuation along the leading edge in the downstream direction.

in the following manner:

$$\delta(x) = \begin{cases} \frac{1}{\Delta x} \left[5 - 3\frac{|x|}{\Delta x} - \sqrt{3 \left(1 - \frac{|x|}{\Delta x} \right)^2 + 1} \right] & \text{for } 0.5\Delta x \leq |x| \leq 1.5\Delta x, \\ \frac{1}{\Delta x} \left[1 + \sqrt{3 \left(\frac{|x|}{\Delta x} \right)^2 + 1} \right] & \text{for } |x| \leq 0.5\Delta x, \\ 0 & \text{otherwise.} \end{cases} \quad (6.2.2)$$

The function is shown here for the x -direction with a mesh width of Δx . This delta function is also used in the immersed boundary projection method to represent the immersed boundary.

The discrete delta function is selected to use the smallest possible width for the actuator, which is limited by the resolution of the simulation. For the current actuator model, there is not a direct analog for the physical slot width, since the simulated blowing does not exit through a slot. Hence, the modeled slot width is set to the characteristic length scale of the discrete delta function, $\sigma = \Delta x = 0.04c$. Compared to the typical slot widths of $\sigma/c = 0.01$ used in flow control (Lee et al. 1989; Duraisamy & Baeder 2006; Holmoway & Richardson 2007), the current slot width in our simulations is slightly larger. Nonetheless the present forcing function is used to explore control techniques at this low Reynolds number.

The actuator is not placed exactly on the surface of the plate to avoid the numerical interference of the boundary force used in the immersed boundary method. The controller is positioned Δx from the plate on the suction side. Forcing directions of upstream, sideways (outward from the midspan to the wing tip), and downstream are considered. In the present study, forcing directions are always set to be tangential to the wing surface. A representative flow control setup is provided in figure 6.2.1.

6.2.2 Strength of actuation

The strength of the actuation is reported in the present study with the non-dimensional momentum coefficient:

$$C_p = \frac{\rho U_{act}^2 b \sigma}{\frac{1}{2} \rho U_{\infty}^2 c} = 2 \left(\frac{U_{act}}{U_{\infty}} \right)^2 \left(\frac{\sigma}{c} \right), \quad (6.2.3)$$

where U_{act} is the actuator velocity. To characterize the actuator model, we simulate this steady blowing with a prescribed U_{act} in an initially quiescent free space. Once steady state is achieved, the velocity at the center of forcing is selected as the characteristic velocity U_{act} . For example, $C_p = 0.1$ corresponds to $U_{act}/U_{\infty} = 0.356$ and $C_p = 1.0\%$.

153



Figure 6.1.12: Top-port views of the convective transport of spanwise vorticity shown by the iso-surface of $|\mathbf{u} \cdot \nabla \mathbf{u}| = 3$ at $t = 5$.

For purely translating flights, the initially formed leading-edge vortex does not stay attached for most planform geometries as reported around rotating wings. As discussed in previous studies, it is suspected that rotation is one of the main mechanisms for the stable attachment.

6.2 Flows with steady forcing

We now investigate steady blowing applied to low-Reynolds-number flows around low-aspect-ratio wings, in following sections, we consider unsteady forcing and closed-loop control. Similar studies to the current investigation are the circulation control (Engler 2000; Joulin & Jones 2006) that utilizes the Coanda effect at the trailing edge. The flow at the trailing edge can be redirected to increase the overall spanwise circulation of the wing and hence enhance lift. We emphasize that in the present investigation, our aim is not to prevent separation but is to exploit wake vortices by changing the three-dimensional dynamics of the separated flow to enhance lift experienced by the low-aspect-ratio wings. Various actuator positions and locations are examined and two setups are identified for lift enhancement by utilizing tip effects. One of the effective control setups is further examined for a wide range of aspect ratios and angles of attack and is found to increase lift in all cases.

6.2.1 Actuator model

In the following simulations, we introduce a body force to model steady blowing. This time-invariant force is applied to the flow field as a uniform strip along the span expressed as

$$f_{act} = f_{act} \delta(x - x_0) \delta(y - y_0) \Gamma \left(-x + \frac{c}{2} \right) \Gamma \left(x - \frac{b}{2} \right) \quad (6.2.1)$$

and is added to the right-hand side of the momentum equation, Eq. (4.1.1). For the current model, addition of mass to the system is not taken into account. Here f_{act} prescribes the strength and the direction of the actuator. The location of the actuator is specified with (x_0, y_0) in the spanwise plane and b describes the span of the plate. The function $\Gamma(\cdot)$ corresponds to the Heaviside step function representing a strip in the spanwise direction. In the computations, the Dirac delta function, $\delta(\cdot)$, is replaced by a discrete delta function, $\delta_d(\cdot)$, proposed by Rome et al. (1999) that regularizes the singularity across 3 cells in both the x - and y -directions.

154




Figure 6.2.2: Forces on the plate with leading-edge actuation for $C_p = 0.10\%$, 0.51% , and 1.0% . Solid and dash lines correspond to unactuated and actuated cases, respectively.

Next, we consider the strength of actuation required to alter the flow field in a noticeable manner. In order to alter the dynamics of the wake vortices in a low-Reynolds-number flow, rather large values of C_p are selected to overcome the viscous effect (and due to the large slot width limited by the grid resolution). Standard values of the momentum coefficient in past studies have ranged from $C_p = 0.01\%$ to 10% for applications of steady blowing on wings (Lee et al. 1989; Duraisamy & Baeder 2006; Holmoway & Richardson 2007). To illustrate the change in the forces exerted upon the plate, we apply flow control around a rectangular plate of $AR = 2$ at $\alpha = 30^\circ$ and $Re = 200$. This example is chosen as the wing is at very high post-stall angle of attack generating strong interaction of the wake vortices and asymmetry about the midspan in the unactuated case (Taira & Colonius 2000, accepted). Actuation along the leading edge in the downstream direction with $C_p = 0.10\%$, 0.51% , and 1.0% are considered.

Figure 6.2.2 exhibits the changes in the lift and drag forces resulting from the leading-edge actuation with varying C_p . With an actuation effort of $C_p = 0.10\%$, there are no pronounced changes in the forces. As the momentum coefficient is increased to 0.51% , lift starts to show increase from the unactuated case. It is interesting to note that the drag history is less affected by the actuation in comparison to the lift increase. With $C_p = 1.0\%$, a significant increase in lift of 33.8% is observed at large time in a time-average sense. Steady blowing here is modeled through a body force with a magnitude in the lift direction of $f_{act}/\rho \sin 30^\circ / \frac{1}{2} \rho U_{\infty}^2 c = 0.1$. Lift enhancement beyond this value can be attributed to the vertical forces. In what follows we consider the use of $C_p = 1.0\%$ to explore actuation locations and directions for the same example problem. Once a favorable setup for flow control is identified, other conditions are examined later in this paper.

6.2.3 Location and direction of actuation

Below we consider the application of steady blowing along the leading edge, midchord, and trailing edge in the upstream, sideways, and downstream directions with $C_p = 1.0\%$. Here, forcing directions are taken to be parallel to the surface of the plate for all cases. For example, downward actuation would be applied with $f_{act} = f_{act} (\cos \alpha \delta_x - \sin \alpha \delta_y)$. Sideways actuation is directed outward from the midspan to the tips of the wing. A collection of lift and drag histories from the controlled flows are presented in figure 6.2.3.

The top two plots in figure 6.2.3 show the force histories for the cases of leading-edge actuation. Lift is increased with downstream blowing, as the separated flow structures become closer to the surface of the plate. The corresponding low-pressure vortex comes sit directly above the top surface enhancing lift by 34% as mentioned in the previous section. The downstream blowing also repositions the wake vortices downward past the trailing edge increasing the effective frontal area. This in turn causes the drag to increase by 16% .

156




Figure 6.2.3: Lift and drag histories for cases of leading-edge, midchord, and trailing-edge actuation with $C_m = 1.0\%$ applied to the separated flow around a rectangular plate of $AR = 2$ at $\alpha = 30^\circ$ and $Re = 300$. Solid and dash lines correspond to unactuated and actuated cases, respectively.

157

a smaller amount compared to lift.

With flow control in the upstream direction, the leading-edge vortices are pushed into the freestream from the top surface, resulting in the loss of the low-pressure region near the surface and, accordingly, a decrease in lift. Outward forcing was applied in hope of releasing the vorticity generated by the leading edge by inducing spanwise flow from the midspan to the tips⁶. However, the right-angled corners on the wing separate the vortex sheets emanating from the leading edge and tip, and triggered the roll up of individual structures (i.e., the leading-edge and tip vortices). Hence the leading-edge vortex and the tip vortices remained disconnected at all times without merging or connecting, in spite of sideways blowing. The overall wake structure was wider in the spanwise direction than the unactuated flow but did not have much influence on the lift or drag exerted on the plate.

The middle two plots in figure 6.2.3 illustrate the effect of midchord actuation onto the lift and drag histories. The noticeable change from the control comes from the downstream forcing where the flow achieves stable steady state and a 76% increase in lift by creating an elongated wake structure (shown in the next section). The vorticity produced by the plate is stably released to the freestream from the leading-edge vortex sheet and the tip vortices at this low Reynolds number. Although this actuator setup is attractive, it would be unlikely to yield a steady flow at higher Reynolds numbers since maintaining such stable wake structure is highly dependent on viscous diffusion. We note that the drag is also affected as the wake structure is moved downward with downstream blowing.

Another change in the dynamics of the wake is exhibited by the sideways blowing along the midchord. This control arrangement repositions the tip vortices away from the plate and allows the leading- and trailing-edge vortices to roll up and shed in a periodic manner. Such behavior of the flow results in a large fluctuation of lift on the wing.

Out of the locations considered in figure 6.2.3, the wake is found to be most sensitive to momentum injection at the trailing edge. The lift shows significant increase and decrease with the application of downstream and upstream blowing, respectively. The time-average lift is increased by a remarkable 100% for downstream blowing. Such noticeable changes in forces are not realized for the sideways blowing at the trailing edge.

The reason for the strong influence of the actuation upon the forces is the direct modification of the trailing-edge vortex. Upstream and downstream blowing, respectively, encourages and discourages the interaction between the leading- and trailing-edge vortices. Below, we will further examine how the downstream actuation at the trailing edge modifies the vortex dynamics in the vicinity of the wing and contributes to lift enhancement.

To summarize, we have observed that steady blowing can change the dynamics of the wake vortices to increase or decrease lift and drag. From the three locations considered, the trailing edge is observed to affect the force on the wing in the most substantial manner. At all three actuator positions, steady blowing in the downstream direction enhanced lift whereas the upstream direction reduced lift. The most effective actuator for the considered example is found to be at the trailing edge in the downstream direction, doubling in lift at large time from the change in the dynamics of the wake vortices.

6.2.4 Wake modification with actuation

Let us visualize the flow field around the rectangular plate of $AR = 2$ at $\alpha = 30^\circ$ and $Re = 300$ for the two most effective cases of actuation from the above discussion, namely the downstream blowing at the midchord

⁶Spanwise blowing has been shown by Campbell (Campbell (1976)) to generate large increase in lift at high angles of attack for swept wings with $C_m = 0.5\%$ to 3.1% .

158




Figure 6.2.4: (Top) Snapshots of the large-time wake around a rectangular wing of $AR = 2$ at $\alpha = 30^\circ$ and $Re = 300$. (Bottom) Corresponding time-average pressure distribution and streamlines along the midspan. Pressure contour levels are set from -0.3 to 0.3 in increments of 0.04 with the negative pressure shown by the dashed lines. Arrows indicate the location of actuation.

and at the trailing edge. Similar to the flow visualization employed in figure 6.1.2, representative snapshots of the wake vortices at large time ($U_\infty t/c = 70$) with and without actuation are illustrated in figure 6.2.4 with the iso-surfaces of $|\mathbf{v}|/U_\infty = 1$ and $Q = 1$. Also presented are the time-average pressure fields and the streamlines along the midspan.

The unactuated flow exhibits complex interaction amongst the leading-edge, trailing-edge, and tip vortices with the dominant flow structures mostly consisting of the leading- and trailing-edge vortices. For the two controlled cases, we observe that the most dominant flow features are the long columnar tip vortices formed by engulfing the vortex sheet from the trailing edge. The downstream blowing along the midchord allows for the flow around the wing to achieve a steady state as shown with its elongated wake vortices in figure 6.2.4(b). The tip vortices exhibit clear columnar structures generating downward induced velocity onto the wake stabilizing the leading-edge vortex sheet. The vorticity within this sheet is diffused into the free stream in a meandering manner at this Reynolds number. In the case of downstream blowing at the trailing edge, the strengthened tip vortices apply stronger downward induced velocity on the leading-edge vortices and let them roll up in close neighborhood of the top surface of the plate as shown in figure 6.2.4(c). Hence the low-pressure cores from the roll-up provide lift enhancement.

In both cases the separation bubble with control visualized along the midspan have become smaller in a time-averaged sense compared to the unactuated case. Additionally, the streamlines are deflected further downward with blowing, directly implying that lift on the wing is increased based on the momentum balance. Note that the roll-up motion at the trailing edge from the unactuated case is now attenuated with steady blowing. The removal of such roll-up seems to be the key factor in modifying the dynamics of the wake.

159




Figure 6.2.5: An illustration of tip vortices engulfing the trailing-edge vortex sheet with trailing-edge actuation.

To illustrate the lift enhancement mechanism, we present figure 6.2.5 for the trailing-edge blowing but this figure also captures the flow physics for the case of midchord blowing. With downstream blowing, the trailing-edge vortex sheet is pushed further downward avoiding direct interaction with the vortical structure emanating from the leading edge. As the trailing-edge vortex sheet advances downstream, it is rolled into the tip vortices which in turn strengthens the tip vortices. Consequently, the strong downward velocity induced by the tip vortices presses the leading-edge vortex and the corresponding low-pressure core region onto the top surface of the wing, enhancing lift. Hence, we suggest that the tip vortices can be used as effectively as the leading-edge vortices in applying vortical forces upon a body. While most of the past discussions in circulation control (Englar 2000) and flapping-wing aerodynamics (Birch et al. 2004b) have focused on the spanwise circulation to explain the lift enhancements, streamwise circulation (tip vortices) can also contribute to lift with the current control setup. Traditionally, the tip effects are thought of as a nuisance due to the loss of the low-pressure region near the tips for attached flows at small angles of attack. However, we emphasize here that the tip vortices can be used favorably to enhance lift in separated flows behind low-aspect-ratio wings.

6.2.5 Parametric study for downstream blowing at the trailing edge

For the most effective case of downstream blowing at the trailing edge, we now further examine the effect of aspect ratio, momentum coefficient, and angle of attack.

The time-average lift and lift-to-drag ratio at large time are summarized in figure 6.2.6. We observe increase in lift and lift-to-drag ratio over a range of angle of attack for $C_m = 0.51\%$ and 1.0% . Shown on the left are the time-average lift coefficients with actuation as well as the time-average and maximum lift coefficients without actuation. What is denoted as the maximum lift is achieved immediately following the impulsive start for the uncontrolled cases around $U_\infty t/c = 1.7$ (figure 6.1.2). The difference between the average and maximum lift for the uncontrolled flow is attributed to the lift enhancement provided by the existence of the initial leading-edge vortex. Enhanced lift from steady blowing is now close to or larger than the maximum lift achieved immediately after the impulsive start, as well as the inviscid limit of lift for

160

low-aspect-ratio airfoils⁹

Let us examine the performance of the actuator for the wing of $AR = 1$. In figure 6.2.6, we observe that the time-average lift with control is increased past the maximum lift achieved by the uncontrolled transient flow for almost all cases considered. Again, this increase is not from the addition of the numerical body force used to represent steady blowing. The direct contributions to lift from such modeled blowing are $0.1 \sin \alpha$ and $0.2 \sin \alpha$ for $C_b = 0.51\%$ and 1.0% , respectively, and do not account for the full increase realized here. The lift increase seen here results mostly from the modification of the wake vortices.

With flow control around a wing of $AR = 2$, the time-average lift is further increased especially around $\alpha = 20^\circ$. For higher α , the interaction of the wake vortices reduce the level of increase but still achieves an overall enhancement for both forcing magnitudes. A similar trend holds for the case of $AR = 4$ with peak performance around $\alpha = 20^\circ$. In the case of $AR = 1$, the tip vortices cover almost the entire span without leaving much room for the leading-edge vortex to stay near the top surface. The lift-to-drag ratio does not show a large enhancement for $AR = 1$ when compared to the higher-aspect-ratio wings. In the case of larger aspect-ratio wings, we observe large increase in lift due to the large extent of the leading-edge vortex over the span.

Once the flow becomes overwhelmed with strong nonlinear interaction of the wake vortices at high angles of attack (i.e., $\alpha = 40^\circ$), the amount of lift enhancement is reduced. In such high-angle-of-attack flows, downstream blowing is not able to keep the trailing-edge vortex sheet from interacting with other vortices and results in no significant increase in the strength of the tip vortices (unless perhaps with much stronger blowing). Nonetheless, this actuator setup seems to be effective overall for various regimes (steady and unsteady periodic/aperiodic states as discussed in the uncontrolled flow section). The strip of steady blowing at the trailing edge is especially attractive for $AR \geq 2$ and $\alpha \leq 30^\circ$ as the lift-to-drag ratio shows substantial increase as well.

To demonstrate that the increase in lift is attributed to the strengthening of the tip vortices due to the downstream trailing-edge actuation, we compute the relative increase in time-average (avg) lift and streamwise circulation of the tip vortex:

$$R_L \equiv \text{avg } C_L / \text{avg } C_{L_0} \quad \text{and} \quad R_\Gamma \equiv \text{avg } \Gamma / \text{avg } \Gamma_0, \quad (6.2.4)$$

respectively, where $*$ is used here to denote the unactuated results. The circulation of the tip vortex is evaluated at a streamwise location of $x/c = 2.5$ using $\Gamma = \oint \mathbf{u} \cdot d\mathbf{A}$, where the contour is chosen to enclose the patch of vorticity (tip vortex) above 1% of the maximum value. These relative increases are plotted against each other in figure 6.2.7 for actuated cases with $C_b = 0.51\%$ and 1.0% around wings of $AR = 1, 2$, and 4 . Based on figure 6.2.7, the correlation coefficient $\rho(R_L, R_\Gamma)$ is found to be 0.952 , which indeed suggests that strengthening the tip vortices have positive influence on the lift enhancement for low-aspect-ratio wings at post-stall angles of attack. With the current flow control arrangement, an increase in lift as high as about 2.5 times the unactuated value has been achieved for one case.

6.3 Flows with unsteady forcing

We consider the application of periodic forcing around a wing of $AR = 2$ at $\alpha = 30^\circ$. While the current case results in aperiodic flow without control, the excitation frequency is selected by extrapolating the natural shedding frequency for the periodic shedding case, i.e., $\omega_b \approx 2\pi f_c / U_\infty$.

⁹The three-dimensional inviscid lift limit, $C_L = 2\pi\alpha/\sqrt{1+(U/\Delta)^2} + 2/AR$, was derived by Helmholtz (1942) from the lifting surface theory for elliptic wings and is shown to be in remarkable agreement with low-aspect-ratio wings of $AR < 4$. Lift for rectangular planforms of $0.5 \leq AR \leq 5$ is accurately predicted with this model (Anderson 1999).




Figure 6.2.7: Normalized lift over normalized circulation of the tip vortex from control ($AR = 1$: $\square$, $AR = 2$: $\circ$, and $AR = 4$: $\triangle$).

To demonstrate the advantage of using periodic excitation, we compare cases with periodic forcing of $C_b = 0.5\%$ and $C_b = 0.25\%$ with steady forcing case. The actuation frequency is set to the natural shedding frequency (i.e., $\omega_b/\omega_k = 1$). The lift coefficient is shown in figure 6.3.1, along with cases without control and with steady blowing at $C_b = 0.5\%$ and 1.0% . It can be noted that the periodic forcing case can achieve significant increase in lift from the uncontrolled case, with less momentum injection compared to those with steady blowing. Hence in what follows reduced actuation effort with $C_b = 0.5\%$ and $C_b = 0.25\%$ is always used. We note in passing that with control drag is not increased as much as the lift coefficient in comparison to the unactuated case.

6.3.1 Effect of Actuation Frequency

Let us next consider the effect of the actuation frequency on the behavior of lift. For flows excited with frequency above ω_k , high frequency modulation generated by the forcing frequency is observed. Shown in figure 6.3.2(a) is the case where $\omega_b/\omega_k = 4$. The overall trend in terms of the average, minimum, and maximum of the lift are similar to the case where the blowing frequency is ω_k .

An interesting case is observed when we choose to blow the trailing-edge vortex sheet at a frequency slightly less than ω_k , as shown in figure 6.3.2(b). For $\omega_b/\omega_k = 0.75$, the minimum value of lift is also enhanced, shifting the time-average lift to the largest value out of all frequencies considered in this study. This agrees with the findings of Seifert et al. (1996b) that notes the optimal behavior occurs when $\omega_b/\omega_k \approx 1$. As we examine the lift trace, we notice that there is regular shedding. The difference between this case with other cases is that the roll up of trailing-edge vortex sheet into the tip vortices (strengthening mechanism) and the formation of the leading-edge vortices are in synchronization.

For forcing frequency of $\omega_b/\omega_k \leq 0.5$, we observe significant variation in lift over time with a decrease in the minimum lift close to the level of the unactuated case. See for example, figure 6.3.2(c), where $\omega_b/\omega_k = 0.1$. The tip vortices convect away from the wing in between the occurrence of blowing, resulting in the loss of the tip vortices' columnar structures and their corresponding downward-induced velocity. The key here is to avoid decrease in lift or its large fluctuation over time, by actuating with a time scale less than the time required for the tip vortices to lose their structures.




Figure 6.2.6: Time-average lift coefficients and lift-to-drag ratios for wings of $AR = 1, 2$, and 4 for cases without actuation ($\square$) and with actuation for $C_b = 0.51\%$ ($\circ$) and 1.0% ($\triangle$). Shown also are the maximum lift for unactuated case ($-$) and the inviscid limits ($---$).




Figure 6.3.1: Lift force on a wing of $AR = 2$ at $\alpha = 30^\circ$ with steady ($C_b = 0.5\%$, $C_b = 1.0\%$) and periodic forcing ($C_b = 0.5\%$, $C_b = 1.0\%$) along the trailing edge. The case without control is also shown ($---$).




Figure 6.3.2: Time trace of lift ($---$) for (a) $\omega_b/\omega_k = 4$, (b) $\omega_b/\omega_k = 0.75$, and (c) $\omega_b/\omega_k = 0.1$ with the unactuated and $\omega_b/\omega_k = 1$ results superposed as ($---$) and ($---$), respectively.




Figure 6.3.3: Snapshots of the periodically time-varying wake structures shown at one-third phase intervals for the periodically actuated case with $\omega_{act}/\omega_h = 0.75$. Shown are the iso-surfaces of $\|\omega\|_2 = 3$ (light gray) and $Q = 2.5$. Arrows indicate the flow direction.

Let us revisit the actuation frequency of $\omega_{act}/\omega_h = 0.75$. The flow is time-periodic, which is visualized in figure 6.3.3. The snapshots are sampled at three equally spaced phases ($\phi = 0, 2\pi/3$, and $4\pi/3$) over a shedding period. The first snapshot at $\phi = 0$ is selected to be at the time of maximum lift with the leading-edge vortex is covering the whole top surface of the wing. As time progresses to $\phi = 2\pi/3$, the tip vortices start to thin out near the rear corners of the wing with reduced blowing. By the time of $\phi = 4\pi/3$, the tip vortices cease to roll up the trailing-edge vortex sheet and hence reduces the downward induced velocity applied on the leading-edge vortex. Around this time, the leading-edge vortex sheds from wing, interestingly in a cleanly shaped vortex ring (clearly visible with the visualization at $\phi = 2\pi/3$). Due to this synchronization, the flow locks onto a high-lift state by keeping the low-pressure core of the leading-edge vortices as close as possible to the top surface in between their detachment.

Next, the actuation frequency is varied from 0.1 to 10 times the natural periodic shedding frequency, ω_h , around an $AR = 2$ wing at $\alpha = 30^\circ$. The forcing inputs of $C_h = 0.5\%$ and $(C_p) = 0.25\%$ are selected as in the previous examples. In figure 6.3.4, the variations in time-average lift are shown for different forcing frequencies. The bars and the shade represent the amplitude of oscillation in the lift coefficient. Also illustrated by the horizontal line is the lift for the uncontrolled case. Highlighted in figure 6.3.4 are the two ranges of frequencies where the time-average lift is higher than those other forcing frequencies. The goal in the next section is to lock the flow onto the optimal frequencies for high lift without prior knowledge of the exact actuation frequencies (i.e., $\omega_{act}/\omega_h = 0.75$ and 1.1).

6.4 Extremum-seeking control

Extremum-seeking is a method of optimal control of nonlinear systems that does not explicitly require a model of the system. We use it here in order to adaptively tune an essentially open-loop approach to separation control. In particular, we apply actuation (in a form to mimic a synthetic jet actuator) at a particular frequency, and use extremum seeking to vary the frequency in order to obtain the maximum lift.

Extremum seeking control (Krstić & Wang 2000, Wang & Krstić 2000) adds a perturbation ($\sin \omega t$) to




Figure 6.4.1: Setup of extremum seeking control to achieve maximum average lift.

frequencies of $\omega_{act}/\omega_h = 0.75$ and 1.1 . Since the initial guess for the forcing frequencies and the control parameters are taken to be the same, the present controller seems to be influenced by the initial condition of the flow field resulting in different lock-on frequencies.

It can be noticed that the actuation frequency seems to converge faster to this optimal actuation frequency, compared to the time-average lift values. This can be a result of how the lift averaging is performed. In the present methodology, we simply compute the average over a few periods; this can be replaced by a low-pass filter instead. The performance and the time lag introduced by choice of averaging may have affected how soon the time-average lift achieved steady state.

There are small differences in the average lift values achieved for the feedback and open-loop cases (cf. figure 6.3.4). This discrepancy could be caused by the preliminary results in figure 6.4.2 not having reached steady values and possibly by the interaction of feedback control and nonlinear physics altering the orbit where controlled flow locks onto. These questions are to be investigated further as computations are still ongoing.

6.5 Summary

In order to establish a knowledge base for the natural (unactuated) flow on three-dimensional wings, we first considered impulsively translated low-aspect-ratio flat plates at Reynolds numbers of 300 and 500, with a focus on the unsteady vortex dynamics at post-stall angles of attack. Numerical simulations, validated by an all low-tank experiment, were performed to study the influence of aspect ratio, angle of attack, and planform geometry on the wake vortices and the resulting forces on the plate. Immediately following the impulsive start, the separated flows create wake vortices that share the same topology for all aspect ratios. At large time, the tip vortices significantly influence the vortex dynamics and the corresponding forces on the wings. Depending on the aspect ratio, angle of attack, and Reynolds number, the flow at large time reaches a stable steady state, a periodic cycle, or aperiodic shedding. For cases of high angles of attack, an asymmetric wake developed in the spanwise direction at large time. The present results are compared to higher Reynolds number flows. Some non-rectangular planforms are also considered to examine the differences in the wakes and forces. After the impulsive start, the time at which maximum lift occurs is fairly constant for a wide range of flow conditions during the initial transient. Due to the influence of the tip vortices, the three-dimensional dynamics of the wake vortices are found to be quite different from the two-dimensional von




Figure 6.3.4: Variation in lift for different actuation frequencies for $AR = 2$ and $\alpha = 30^\circ$ shown on the left. Time-average lift coefficients ($\bullet$) are plotted their maxima ($\square$) and minima (Δ). The horizontal line ($—$) denotes the average lift for the unactuated case bounded by its minimum and maximum illustrated by the shaded box.

the system near a maximum operating point, which in this case corresponds to the maximum lift achieved for a range of actuation frequency. A diagram of the extremum seeking algorithm is shown in figure 6.4.1. The resulting output signal is put through a high-pass filter to remove the DC component and is demodulated with the input perturbation. This demodulated function can then be passed through an integrator as an approximate gradient update to improve the guess of the maximum operating point.

In the case of the current study, we aim to lock the shedding with $\omega_{act}/\omega_h = 0.75$ or 1.1 that corresponds to frequency at which the time-average lift is at its local maximum (as shown in figure 6.3.4). The immersed boundary projection method will take the control input of actuation frequency and provide the lift force on the wing as an output. The lift is passed through a time-averaging routine prior to entering the extremum seeking algorithm, depicted by the lower half of figure 6.4.1.

The advantage of this method is that the dynamics of the flow need not be characterized and can be treated as a nonlinear black box. Hence this approach seems as a promising path to stabilize the flow about its high-lift state. The only concern in designing this controller is the ratio of the speeds at which the flow reacts and the controller updates its optimal point of operation. If the nonlinearity changes the flow field with a shorter time scale than that for the controller to take its effect, this algorithm can be ineffective or in some instances the flow can lock onto a different cycle from what is observed in the open-loop control cases. A discussion on this point is provided later in the section.

The control parameters to be chosen for the implementation of the extremum seeking algorithm are ω and ω for the control perturbation, the integral gain k , and the cutoff frequency Ω_c for the high-pass filter. The perturbation $\sin(\omega t)$ is selected such that it is small compared to the variable to control (ω_{act}) and is slowly varying compared to the main physics of interest (shedding). Hence we select the parameters as:

$$\omega \ll \omega_{act} \Rightarrow \omega = 0.0375 \quad \text{and} \quad \omega \ll \omega_h \Rightarrow \omega = 0.15.$$

The integral gain k is set to be 0.5 (near unity) and the cutoff frequency Ω_c is set to be equal to ω inside a high-pass 3rd order Butterworth filter.

We consider providing the simulations with two different initial condition with $\omega_{act}/\omega_h = 0.7$, at time t_0 which is some large time. Preliminary findings for the time evolution of actuation frequency and time-average lift are shown in figure 6.4.2. Note that these two cases lock to frequencies onto the optimal forcing




Figure 6.4.2: Two preliminary results of actuation frequency locking on the local optimum frequencies of $\omega_{act}/\omega_h = 0.75$ and 1.1 , shown respectively with (---) and (—). Actuation frequency reaches steady values faster relative to the time-average lift values. Longer duration of simulation is required for final assessment.

Kármán vortex street in terms of stability and shedding frequency.

Next we considered the application of steady and unsteady blowing, with the aim of exploring techniques to enhance lift by directly modifying the dynamics of the wake vortices. Out of various combinations of forcing location and direction considered, we identified two configurations that provide significant lift enhancement. In these cases, actuation appears to strengthen the tip vortices for increased downward induced velocity upon the leading-edge vortices. This in turn moves the low-pressure core directly above the top surface of the wing to greatly enhance lift. Periodic forcing is also observed to be effective in increasing lift for various aspect ratios and angles of attack, achieving a similar lift enhancement attained by steady forcing with less momentum input. Through the investigation on the influence of the actuation frequency, it is also found that there exists a frequency at which the flow locks on to a time-periodic high-lift state.

Finally, extremum-seeking optimal was used to tune the open-loop actuation frequency in order to obtain the maximum lift. Preliminary simulations showed the efficacy of this procedure for rapidly identifying optimal control parameters, even in CPU-intensive, three-dimensional simulations.

7 Oil Tunnel studies at low Reynolds number

In this section, we document the construction of and preliminary measurements made in a novel recirculating oil tunnel for flow control studies. The advantages of using oil as the working fluid are that low Reynolds number aerodynamics can be studied using scale experiments that give rise to sufficiently large forces that can be accurately measured, and time scales that are sufficiently slow to enable accurate real-time and three-dimensional particle image velocimetry (PIV). Control studies focusing on the three-dimensional, low aspect ratio wings discussed in section 6 are ongoing in the facility and will be reported in a future thesis and paper.

7.1 Oil tunnel facility

To enable flow studies at low Reynolds numbers, a recirculating oil tunnel was designed and installed. Flow speeds of up to 25 cm/s are achievable through the 30 cm $\times$ 50 cm $\times$ 150 cm test section. The working fluid is a white mineral oil with a density comparable to water but a kinematic viscosity comparable to that of air.

7.1.1 Construction

Nearly the entire wetted area of the oil tunnel is constructed of acrylic, with the test section resting on a welded steel frame. Joints are solvent welded and flanged joints are sealed using rubber gaskets. The flow is driven through the tunnel by an Aurora 6x6x11 344A-BF centrifugal pump, which is powered by a Toshiba B0056FLF2AM104 SHIP induction motor. Fluid is pushed through a six inch diameter PVC pipe into the test section. This pipe is perforated once it enters the test section, allowing the fluid to be distributed across the height of the test section. Once through the perforated pipe, fluid passes through a flow conditioning section consisting of a perforated plate, followed by a honeycomb section, and finally three mesh screens, before entering the test section proper.

The all acrylic construction of the tunnel test section allows optical access from the sides and from below. The test section is accessible from above, with a sting assembly installed above the tunnel providing the mounting point for the test article. The access cavity is closed via a series of interconnecting plates, eliminating as much of the free surface as possible, while still allowing the sting mount to pass into the test section.

Once flow travels through the test section, it enters the return section, where it is routed back to the pump inlet via turning vanes. In each vane passage, a stopcock is installed, allowing air bubbles to be vented to ensure that the whole test section is completely filled with fluid. The return section also contains two nickel-plated copper cooling loops, which are connected to a process chiller that provides temperature regulation for the working fluid. Coincident with the cooling loops are PVC pipes providing inlet and outlet connections to an auxiliary pump, which provides both easy access to drain/fill the tunnel and pressure for actuation fluid for flow control experiments, to be discussed further in Section 7.3. A schematic of the tunnel is shown in Figure 7.1.1. Figure 7.1.2 shows a photograph of the facility.

The entire tunnel assembly contains approximately 1.2 m³ (310 US gallons) of fluid.

7.1.2 Model mount assembly

In order to secure and position the experimental articles in the test section, a mounting assembly was installed above the tunnel (Figure 7.1.3). The assembly is a rectangular frame constructed from 80/20 sections,




Figure 7.1.1: Oil tunnel facility schematic.




Figure 7.1.2: Oil tunnel facility.

169

which is bolted into the lab ceiling. For the current experiment, a 1.5 inch circular steel pipe is used to connect the model to the mount assembly. An interface adapter is installed at one end of the pipe, which provides connections for the six-axis force balance. The pipe is mounted to a cross-beam whose ends are attached to rotary clamps, which allows the angle of attack to be set. The rotary clamps are mounted in turn to two side rails parallel to the length of the tunnel. The cross-beam can be adjusted upstream and downstream, allowing precise positioning of the model in the tunnel.



Figure 7.1.3: Model mount frame schematic.

7.1.3 Working fluid

White mineral oil is the working fluid of the facility, with a density of 835 kg/m³ and a dynamic viscosity of 13.3(10⁻³) kg/m-s, resulting in a kinematic viscosity of 1.6(10⁻⁵) m²/s. This combination of density and kinematic viscosity allows for experimental articles to be scaled with typical micro-air vehicle sizes, but allows force levels to be high enough to measure without resorting to extraordinary measures.

To illustrate this more fully, one might consider doing experiments at Re = 10³ in air or water. To examine the suitability of those fluids, and to justify oil as a choice, consider a model with a characteristic length scale (in this case, chord length c) of 10 cm and assume a force coefficient, C_F , of 0.5. Table 6 shows the velocity U required to attain the desired Reynolds number, along with the force per unit span F/b experienced by the model for water, air and oil. Force per unit span is calculated using Equation 7.1.1, with ρ denoting fluid density.

$$\frac{F}{b} = \frac{1}{2} \rho U^2 C_F \quad (7.1.1)$$

171

Fluid	Density kg/m ³	Kin. Visc. m ² /s	Velocity cm/s	Force/Spn mN/m
Water	1000	1.1(10 ⁻⁶)	1.1	0.03
Air	1.23	1.5(10 ⁻⁵)	15	0.01
Oil	835	1.6(10 ⁻⁵)	16	5.30

Table 6: Forces and flow speeds for a 10 cm model at Re = 1000.

As can be seen, the difference between the kinematic viscosity of air and oil are very slight, which indicates that a given model at a given Reynolds number will experience roughly similar flow speeds in both fluids, but the forces in oil will be nearly 800 times larger than those experienced in air. This will allow much greater ease of measurement with off-the-shelf sensors. Water's density might make it an attractive choice at first, but the kinematic viscosity forces the model size to be two orders of magnitude smaller for the same Reynolds number, making the forces even more difficult to sense.

The particular oil chosen for this application was Chevron Superia White Oil Number 5. This white mineral oil is part of a family of oils, all of relatively similar densities, but with viscosities ranging over two orders of magnitude. The oils are miscible in one another, allowing for the selection of a precise viscosity characteristic if such control is desired. For these tests, the lightest viscosity oil was employed.

7.2 Temperature control

As the fluid is pushed through the test section, a considerable amount of work is done on the fluid. If the only heat removal is via conduction to the laboratory through the walls of the tunnel, the tunnel temperature rises by roughly 2 °C per hour. This temperature increase has a strong effect on the viscosity of the oil, and also causes unwanted sensor drift in strain gauge based force balances. In order to mediate this problem, two nickel-plated copper cooling loops are installed in the return section of the tunnel. Copper was chosen for its high thermal conductivity; the nickel-plating is necessary due to an unfortunate material incompatibility between the oil and copper, which causes the oil to experience severe orange coloration. The coloration can interfere with measurement techniques utilizing laser illumination, since the light attenuation through the oil increases as the coloration strength increases.

The cooling loops are connected to a NESLAB Systems III water-water heat exchanger unit. This unit exchanges against building chilled water, which typically enters the building at approximately 5 °C. The unit has been modified such that the temperature set point is controllable via an external voltage signal. This is connected to a micro-controller board on which has a simple PI controller operates to regulate the tunnel temperature. The tunnel temperature is sensed with a RTD sensor at the exit of the test section (so as not to disrupt the incoming flow). The RTD sensor, a PR-11-2-100 probe from Omega, is attached to a CN63200-F1-A1 display, also from Omega, which provides a voltage output signal scaled with temperature. Overall, the tunnel temperature can be maintained within 0.1 °F at all tunnel speeds.

7.3 Flow control hardware

Active flow control in these experiments is achieved through the use of steady and unsteady blowing through slot actuators embedded within the experimental models. In order to provide fluid for these actuators, an auxiliary pump system was installed. This system draws fluid from the return section of the tunnel, providing approximately 170 kPa (25 psi) for use by the actuation system. The system pressure is set by adjusting the main throttle valve, and the supply pressure is set by a pressure regulator downstream of this valve. The

172

fluid then is sent to a seven port manifold. Downstream of each manifold port is a three-way valve that allows the actuator channel to be configured as a steady or unsteady channel, or to be shut off completely. Unsteady actuation is provided by Omega SV-27 solenoid valves. Once through the three-way valve (and the solenoid, if unsteady actuation is selected), the flow rate for each channel is regulated by a needle valve. Figures 7.3.1 and 7.3.2 illustrate the relevant plumbing pathways for steady and unsteady actuation.




Figure 7.3.1: Steady actuation plumbing schematic.

Flow rates are currently measured with a flow meter upstream of the seven port manifold. The flow meter is a JVA-20KG spur gear positive displacement device manufactured by AW Company. It is capable of measuring flow rates from $6.3(10^{-5})$ to $1.3(10^{-4})$ m³/s (0.01 to 2.0 GPM). Since flow rates are not able to be measured on each individual channel simultaneously, before each experiment, each channel's flow rate is established by adjustment of the needle valve to match all of the other channels. Once all seven channels have been individually set, all are turned on. Presumably, by equalizing the resistance of each channel, the flow rate out of each actuator slot should be equal.

Unsteady actuation, as mentioned before, is accomplished through pulsing solenoid valves. The solenoids, normally closed, require an 120 VAC signal to be opened. A custom control box was designed to enable opening of the valves using a 5V digital output signal from a data acquisition device. The circuit ensures protection of the DAQ device, as well as providing over-current protection for the solenoid actuators. The solenoids have a time response specification of 4 to 15 ms, which means they can be reliably driven to more than 60 Hz.

One point to note is that the time response of the flow meter is relatively limited (1 - 2 Hz), so unsteady flow rates are currently very hard to reliably report. Several alternative flow metering options have been examined, but those with sufficient time resolution have currently been found to be prohibitively expensive.

173




Figure 7.4.1: Wing model with mid-chord actuators.




Figure 7.4.2: Wing model with mid-chord actuators.

175




Figure 7.3.2: Unsteady actuation plumbing schematic.

7.4 Wing models

The airfoil section chosen for these experiments is the NACA 0012. This symmetric airfoil section provides sufficient room for the actuation pathways to be embedded within the wing. The wings have a chord length of 10 cm and a span of 20 cm, yielding an aspect ratio of 2. The wing tips are rounded, the first 25% of the chord is rounded using the base profile revolved around the chord line, the remainder uses a semi-circular feature lofted to the trailing edge. Figure 7.4.1 shows a solid model rendering of the wing, with the seven actuator slots on the suction surface of the wing shown across the span at the 50% chord position. Figure 7.4.2 shows the same model, but with the actuator channeling exposed.

The actuator slots are independent of one another, each having its own internal plumbing separate from all the others. Each slot has a length of 25.2 mm and a width of 0.54 mm. Figure 7.4.3 illustrates a section cut through the actuator slot and cavity, showing that the slot is designed to produce a jet of fluid tangential to the suction surface of the airfoil.

All of the geometry for the wing is constructed using 3D solid modeling software. The resulting solid model file is used directly in the manufacturing process. A rapid prototyping technique, known as fused deposition modeling (FDM), is used to create the solid wing. The construction material is a polycarbonate/ABS blended plastic material. The FDM technique lays up material layer by layer, allowing the internal channeling required for the actuators to be integrated during build time. The model is constructed with an elliptical sting attached to the pressure surface of the wing, swept back at an angle of 45°. The sting not only provides the mounting point for the model within the tunnel, but also contains the channeling that connects the external flow supply from the actuation pump with the actuator slots embedded within the wing. Upon receipt of the wings from the manufacturer, the wings are hand sanded to remove the discrete "steps" left behind by the process, which has a finite resolution of approximately 0.25 mm.

For baseline flow studies, an acrylic wing model was constructed. This allowed the laser illumination to pass through the model, providing the ability to perform DPIV (Section 7.5.1) and DDPV (Section

174



Figure 7.4.3: Actuator slot and cavity.

7.5.2) measurements on both the suction and pressure sides of the wing. The acrylic wing model was machined using a three-axis mill, using the same profile that was specified for the rapid prototyped models. A sting was created using rapid prototyping, which was then affixed to the acrylic wing in the same location as those made via FDM. This model contains no actuation channeling and has different surface roughness, but should provide a good baseline model for the actuated wings.

7.5 Measurement techniques

7.5.1 Digital particle image velocimetry (DPIV)

Particle image velocimetry is a well established experimental technique for obtaining velocity field information from a flow. Digital particle image velocimetry, first introduced by Willert & Oharib (1991), has become a primary tool in most modern fluid mechanics research. The basic concept involves imaging a two-dimensional slice of a flow field that has been seeded with particles capable of following the flow. The scattered light from the particles is recorded via a digital imaging camera (such as a charge-coupled device, or CCD). A pulsed laser, whose light is formed into a thin sheet through the use of cylindrical lenses, freezes the particle images in the exposure. Two exposures are acquired within a very short time of one another (on the order of milliseconds or faster). These two images are subsampled (a process referred to as windowing) and the corresponding subsamples from each image are subjected to a cross-correlation processing step. Typically, this involves computing the Fourier transform of the two windows, followed by multiplication of the transformed samples, and finally computing the inverse Fourier transform of the product. This process yields a correlation plane. If the two groups of particles have sufficient spatial correspondence, a correlation

176

peak will exist whose location in the plane indicates the overall displacement vector for the particle images in the window. The distance traveled divided by the time between the exposures yields an instantaneous velocity for that window.

Given the above methodology, the components necessary for the technique are the following:

- Camera and camera optics for generation and acquisition of digital images
- Laser and associated optics for illumination of flow field
- Seeding particles capable of following the flow and scattering sufficient light
- Timing hardware for generation of appropriate illumination and camera shutter triggers
- Software for processing of resulting image pairs
- Computer workstation for hosting appropriate hardware/software

DPIV components The camera chosen for this investigation, the IPX-2M30G (S/N 60287), utilizes a 1600 x 1200 monochrome CCD and is manufactured by Imperx, Inc. It is capable of acquiring full resolution images at 33 frames per second (fps). It features GigE output from an onboard framegrabber via a CAT-5 network cable. A high performance driver is provided that gives transfer rates upwards of 1 Gb/s, provided the network interface card (NIC) chipset is based on the Intel Pro/1000 family. Installation of the accompanying drivers reconfigures the NIC such that the computer recognizes it as a frame-grabber device.

The camera is capable of being triggered via an external pulse or via the onboard pulse generator. If the external trigger is used, a TTL pulse can be connected at the "Trigger" input on the camera. The internal trigger can be output on the "Strobe" output on the camera. The frequency and duty cycle of the onboard pulse generator is configured in the Lym GigE Application by specifying the high time and low time for the signal in units of "granularity" (clock cycles), where 1 unit corresponds to 30 ns. Calculation of the correct settings is straightforward. Assuming a 50% duty-cycle, the granularity is calculated based on the desired trigger frequency (f) as in Equation 7.5.1.

$$\text{granularity} = \frac{1}{13106930(10^{-9})f} \quad (7.5.1)$$

Once the granularity is selected, the correct inputs for the width and delay fields are computed using Equation 7.5.2.

$$\text{width} = \text{cycles} \times \text{granularity} \times 30 \quad (7.5.2)$$

where cycles is defined by Equation 7.5.3.

$$\text{cycles} = \left\lceil \frac{1}{2f \times \text{granularity} \times 30(10^{-9})} \right\rceil \quad (7.5.3)$$

The lasers have a maximum fire rate of 15 Hz. Thus with the granularity set to 17, the input for the width and delay is 33,333,090 ns (corresponding to 65359 cycles), producing a 15 Hz, 50% duty-cycle pulse train.

A SMC Pentax-A 1:2.50mm lens is used along with a C-mount adapter, since the lens features a K-mount configuration. The aperture is stopped full open to allow the maximum amount of light to collect on the sensor.

flashlamp, which provides the energy to the laser system. The other trigger fires the q-switch, which is responsible for allowing the high energy laser energy to be released from the head. The total laser output energy is sensitive to the amount of time between firing of these two triggers, so precision control is desirable.

To achieve this, a counter/timer board is used to generate the correct pulse trains. The PCI-6602, manufactured by National Instruments, has eight counters available, of which up to five are used for this application. LabView 8.6, also provided by National Instruments, is the software used to configure the hardware and execute the generation of the pulse trains.

Figure 7.5.2 shows the virtual instrument (VI) front panel used to control the timing board. Several operational modes are available, depending on the type of experiment to be done. Enabling the "Mode Select" toggle configures the other panel controls for the experiment type selected. Disabling the toggle allows each individual control to be manually set by the user. The options set by each mode are listed in Table 7. "External Trigger" refers to whether the timing card receives an external time source or uses an internally generated pulse. "Wait for DAQ Initialization" determines whether the pulse train will begin immediately or upon receipt of a ready signal from the data acquisition process. "Stop on DAQ Finish" determines whether the pulse train will be terminated upon successful acquisition of all requested samples. In all modes, the user chooses the timing between the firing of the first laser and the second laser in order to produce image pairs that can be correctly processed by the DPIV software.

Mode	Description	External Trigger	Wait for DAQ Initialization	Stop on DAQ Finish
RT-PIV	Real-Time PIV	Y	N	N
RTV - Cam	Camera Driven Sync PIV/DAQ	Y	Y	Y
dPIV - 6602	Timer Driven Sync PIV/DAQ	N	Y	Y
PIV - Cam	Camera Driven Normal PIV	Y	N	N
PIV - 6602	Timer Driven Normal PIV	N	N	N

Table 7: Mode configurations for DPIV timing VI

The PIV processing software is provided by PIVtec GmbH, which is an offshoot of the PIV group led by Chris Willert at the German Aerospace Center (DLR). It provides a comprehensive toolkit for processing DPIV image pairs, incorporating most of the currently known techniques for generating accurate velocity data.

In the current setup, a single computer workstation is tasked with hosting the hardware responsible for providing an interface to receive captured images from the camera as well as a counter/timer board to provide accurate and consistent timing of the necessary triggers. Image acquisition and timing are controlled through corresponding software applications. Finally, DPIV image processing software is used to provide the final vector fields that are the ultimate result of the method.

7.5.2 Digital defocused particle image velocimetry (DDPIV)

Digital defocused particle image velocimetry was first proposed by Willert & Gharib (1992). It is a velocimetric technique to measure 3D position and velocity components of particles in a flow. If the particle cloud is viewed through multiple apertures, each particle will generate one image only if it lies on the focal plane of the camera. If the particle is located off the focal plane, multiple images of that particle are generated on the image plane. The separation between particle images is directly related to the particle's distance from the focal plane, allowing the 3D coordinate location of the particle to be computed.

The lasers used for this investigation are Gemini Nd:YAG PIV lasers (S/N 10141 and S/N 10142) from New-Wave Research. They feature a dual laser head configuration, allowing two 120 mJ pulses to be generated at very small temporal separations. The nominal pulse time is 5 ns and each head is capable of firing at a maximum rate of 15 Hz.

The laser beam is passed through a lens stack consisting of a cylindrical lens followed by a spherical lens. The lens stack causes the beam to be formed into a sheet, which in this case is oriented parallel to the lab floor. A right angle first surface mirror is used to turn the light sheet up into the test section through the tunnel floor. This results in the light sheet illuminating the model along the chordline at some specific spanwise station. This mirror is on a movable traverse, which allows this spanwise station to be changed if desired. Figure 7.5.1 illustrates the illumination setup.




Figure 7.5.1: DPIV Laser sheet orientation

Seeding particles are added to the flow in order to make the flow visible, providing the raw data necessary for the image processing to yield usable velocity information. The seeding particles must be of sufficient size to effectively scatter the light, but be small enough to faithfully follow the dynamics of the flow under investigation. In this case, the particles chosen are silver-coated hollow glass spheres with a nominal diameter of 44 μm . These particles are obtained from Potter Industries under part number SH230533. These particles are actually slightly larger than one would normally choose for oil applications based on scattering considerations, but the seeding requirements for DDPIV (see Section 7.5.2), was a partial driver for choosing a larger particle. However, in spite of their larger size, they are actually just slightly buoyant in oil, with an average density of 0.5 g/cc. Tracking individual particles in the freestream through a series of images (when the seeding density is low) seems to indicate that, over the field of interest, the particle motion due to buoyancy should have negligible impact on the final results.

The timing of the laser pulses is controlled by sending two triggers to each laser. One trigger fires the




Figure 7.5.2: DPIV timer VI interface

Two basic limitations of the method exist. First, the use of a multiple-aperture mask on a single lens camera produces an image with heavily overlapped particle images when seeding density is sufficiently high to obtain good spatial resolution. Additionally, due to the small separation of apertures, only very small interrogation volumes were possible. In order to overcome these limitations, the DDPV technique was extended to a multiple-camera implementation (Pereira & Gharib 2002), with each camera utilizing a single aperture. Three cameras of equal focal length are arranged in an equilateral triangle configuration with optical axes parallel to one another. Figure 7.5.3 shows a schematic of the camera and the resulting measurement volume (the cube enclosed inside the pyramid).



Figure 7.5.3: DDPV camera schematic.

A multi-plane calibration method is employed to quantify the distortions caused by mechanical misalignment, optical distortions and non-ideal optics. This provides a mapping function which is applied to the obtained images in order to correct these distortions, increasing the accuracy of the technique. Once images are obtained and distortions corrected, particles are located in each of the three images using a 2D Gaussian fit algorithm. This algorithm provides particle locations to sub-pixel accuracy. Additionally, the algorithm serves as a filter, since non-Gaussian intensity distributions are not likely to be valid particle images.

Particle images are converted into triplet images through the use of another algorithm. This algorithm uses knowledge of the aperture arrangement in order to determine likely directions along which to search for particle matches. Threshold criteria are set which limit the allowable amount of spatial separation of the particle image in each of the three images. If the thresholds are exceeded, the 2D particle is discarded. The thresholds are usually set very stringently, to avoid misidentification of 3D particle positions (typically known as "ghost" particles). Typically, the method is able to transfer approximately 60% of all particle images into triplet matches. The in-plane coordinates (i.e. parallel to the camera face) are identified by the triplet's center, with the out-of-plane coordinate (i.e. parallel to the camera optical axis) identified by the height of the triplet.

Once all 3D particle locations are identified, a particle tracking algorithm identifies the particle motion between two successive frames, generating a 3D velocity field. Tracking is accomplished via a relaxation method, which is described in detail in Pereira et al. (2006). This differs from the typical approach employed in DDPV, which uses a cross-correlation method to identify the average motion of groups of particles. Here, the actual particles are tracked through space, for although the particle image density is in the range of PIV methods, the particle density in physical space is much lower. With the identification of the velocity of each

particle, data validation is performed to remove outliers. Since the velocity vectors are based on the actual particle locations, they are randomly distributed throughout the measurement volume. Projection of these vectors onto a regular grid through interpolation allows the computation of flow quantities, such as vorticity, that are useful for understanding the evolution of the flow over time.

The specific system used in this study is capable of resolving a cubic measurement volume of roughly 120 mm³. Image pairs can be captured at frequencies up to 7 Hz.

7.5.3 Force measurements

Aerodynamic forces on the wing are measured using a commercially available six-axis force balance. The balance is manufactured by ATI Industrial Automation, located in Apex, North Carolina. The specific model is used in this study is the Nano-43, a silicon strain-gage based sensor. The sensor measures 43 mm in diameter and is 11.53 cm thick. The sensor is delivered with an interface box that provides power to the transducer as well as signal conditioning to allow the sensor to be used with the 68-pin connector common to many data acquisition systems, including the one described above.

The Nano-43 comes in several specific configurations. The one chosen for this experiment was the SI-18-0.25. This configuration provides the ability to measure forces up to ± 18 N and torques up to ± 250 N-mm.

A calibration was provided by the manufacturer at the time of purchase. This calibration was spot checked in the lab and found to be in relatively good agreement. The manufacturer's calibration is done with a precision jig, ensuring accurate loading of the sensor. The calibration is a 6×6 matrix (shown in Equation 7.5.4), transforming voltages (G_i) into forces and torques, as in Equation 7.5.5. The force sensor is mounted between the model support assembly and the integrated sting on the model. The x axis is aligned with the chord line of the model wing, with a positive sense directed towards the leading edge. The y axis is aligned with the span of the model wing, with the positive sense directed towards the floor of the tunnel. The z axis is aligned with the vertical, with the positive sense directed towards the floor of the tunnel. The y axis is aligned with the span of the model wing, with positive sense according to a right-handed coordinate system based on the other two axes.

$$M = \begin{bmatrix} -1.87094 \times 10^{-1} & 5.05623 \times 10^{-1} & -4.07503 \times 10^{-1} & 5.569462 & -6.65423 \times 10^{-1} & -3.46177 \times 10^{-1} \\ -1.22041 \times 10^{-1} & -2.07148 & -2.07177 \times 10^{-1} & 1.470807 & 1.16470 \times 10^{-1} & 1.471274 \\ 3.585226 & -5.01247 \times 10^{-1} & 2.011207 & -4.87020 \times 10^{-1} & 2.413844 & -1.16886 \times 10^{-1} \\ 4.88940 \times 10^{-1} & 1.07096 & 1.08040 \times 10^{-1} & 5.520091 & 3.88776 \times 10^{-1} & 1.746652 \\ -5.03380 \times 10^{-1} & 4.84349 \times 10^{-1} & 1.58918 \times 10^{-1} & -1.78736 \times 10^{-1} & 1.78737 \times 10^{-1} & -1.48514 \times 10^{-1} \\ -4.46820 \times 10^{-1} & -2.36222 \times 10^{-1} & 2.36222 \times 10^{-1} & -2.48714 \times 10^{-1} & -2.78030 \times 10^{-1} & -1.46330 \times 10^{-1} \end{bmatrix} \quad (7.5.4)$$

$$\begin{bmatrix} F_x \\ F_y \\ F_z \\ T_x \\ T_y \\ T_z \end{bmatrix} = M \begin{bmatrix} G_1 \\ G_2 \\ G_3 \\ G_4 \\ G_5 \\ G_6 \end{bmatrix} \quad (7.5.5)$$

Data acquisition Signals from the force sensor and other sensors are acquired using two data acquisition cards. The first card is a National Instruments PCI-6035E card. This 16-bit card has eight differential (16 single-ended) analog input channels, two analog output channels, eight digital input/output channels, and two 24 bit counter/timers. The second card is a Measurement Computing PCI-DAS1000, featuring eight differential (16 single-ended) analog input channels, 24 digital input/output lines, and three 16-bit counters.

The PCI-6035E is connected to a BNC-2090A accessory, which provides BNC and spring terminal access to the majority of the card input/output channels. The main feature that this accessory provides is the

68-pin SCSI pass-through connector, which allows the force transducer to be connected to the data acquisition card while keeping the other functions of the DAQ card accessible. The other functions utilized by this experiment include one of the digital lines along with the counter/timers in order to provide synchronization capability with DDPV measurements.

Additional input voltages are sampled by the PCI-DAS1000, since the analog input channels are taken up with the force transducer. Connections are made via two CIO-MINI50 screw-terminal boards.

National Instruments LabView 8.6 is the software used to create the interface between the user and the elements of the experiment. This interface is known as a "Virtual Instrument," or VI. The VI allows the user to provide configuration settings for the acquisition, as well as the configure the storage and processing of the incoming sampled data.

For this experiment, the VI was programmed to record the forces experienced by the model. The user interface is shown in Figure 7.5.4. There are several modes programmed into the interface, with the intent of providing synchronization of force data with DDPV acquisition. Table 8 shows the configuration each of the modes provides.

Mode	Description	PIV as Trigger	DAQ Test Signal	Pulse Width for N-Chan	Finite Acq	RTPIV Trigger
RT-PIV	Real-time PIV	Y	N	Y	Y	Y
PIV/Chm	Chm. Sync PIV/DAQ	Y	Y	Y	Y	N
PIV/6602	Tim. Sync PIV/DAQ	Y	Y	Y	Y	N
No PIV/Free	Cont. Acq.	N	N	N	N	N
No PIV/Finite	Fin. Acq.	N	N	N	N	Y

Table 8: Mode configurations for data acquisition VI.

The data acquisition hardware and software is hosted on a custom built PC. The processor is a 2.40 GHz Pentium 4 on an ASUS P4P800 motherboard. Memory is provided by G-Skill, with 2x1GB DDR2-400 installed. The operating system is Windows XPSP3.

7.6 Real-Time Particle Image Velocimetry (RTPIV)

Particle image velocimetry (PIV) has been a standard laboratory technique in experimental fluid mechanics for many years. The technique allows quantitative visualization of fluid flows, exposing the dynamics in a manner that can directly contribute to a physical understanding of the underlying flow physics. Once sufficient physical insight has been gained about a particular flow, one often considers manipulating the governing mechanisms to achieve a desired result. This is especially true when control of the flow under study might have a significant impact on the performance of engineering systems.

Typically, a flow control experiment hopes to achieve some change in system dynamics, often through amplification or suppression of flow instabilities. This is accomplished through the use of one or more actuators to produce a disturbance within the flow. If the flow is sufficiently receptive to the created disturbance, the dynamics of the system can be modified. Monitoring the effectiveness of the actuation through sensor measurements allows for the possibility of closed-loop control.

Many flow control experiments utilize a small number of sensors, which often only measure some characteristic of the flow at discrete points. Flow dynamics often must be inferred from these signal measurements. Additionally, these sensors can often influence the flow being investigated, obscuring the overall control objective. The actual physical dynamics of the system, as well as the effectiveness of the actuator in changing those dynamics, can be difficult to assess from this limited amount of information. The use of PIV




Figure 7.5.4: Data acquisition VI front panel.

as a non-invasive, full-field sensor has the potential to yield much richer knowledge of the system dynamics and actuator effectiveness in the presence of control.

The challenge of using PIV as a sensor in a closed-loop system is mainly due to the time delays involved in acquiring and processing the image data. Depending on the window resolution specified by the user, the computations involved in DPIV may need to be done thousands of times per image pair. Long delays between the measurement and the actuation events will lead to ineffective control strategies because the flow will likely have evolved such that the computed actuator input is no longer appropriate. As the processing speed of computing hardware has increased over the years, the time required for processing of each image pair has naturally decreased, making the task of providing the vector fields at "real-time" speeds more and more likely. The current implementation appears to have accomplished just this task with minimal reliance on customized hardware, which has been the case in the past (Siegel et al. 2003, Yu et al. 2006).

7.6.1 Basic Architecture

Real-Time DPIV (RTPIV) has a setup identical to a typical DPIV implementation, with a dual-beamed pulsed laser providing illumination of a flow seeded with small particles. The illuminated particle field is imaged twice in rapid succession using a digital camera. The resulting image pair is converted into a two-component velocity field through software processing, which subsamples each image into small regions and computes a cross-correlation for the corresponding regions in each image. Peaks in the correlation plane are fit to sub-pixel accuracy. The cross-correlation results in a displacement field, which is converted to a velocity field based on the lens magnification and the temporal separation between the laser pulses.

Realization of "real-time" performance, in the current work, relies on the design and implementation of a custom software package capable of quickly processing particle image pairs as they are streamed into system memory from the camera. Upon arrival in memory, the images are processed using carefully parallelized algorithms to compute the cross-correlations. Additionally, for use as a feedback sensor, some region of interest (ROI) is chosen where further calculations are performed. These calculations can then be used to make control decisions, which are then passed out to actuators through a digital I/O interface in order to change the flow in a desirable way. Figure 7.6.1 illustrates the overall system architecture.

7.6.2 Challenges

In order to realize "real-time" performance, the processing of an image pair must occur before the next image pair is available in the host PC's memory. If processing is not finished when the next image pair arrives, then the system will either have to wait until the processing engine is finished or the image pair will have to be discarded. In both cases, this has implications for use of the results in a control loop. In the first case, the latency buildup will quickly cause the results to lose any temporal correlation with what is actually happening in the flow. Thus, the controller will be making decisions too late to have the desired impact on the flow. In the second case, the loss of data may cause the controller dynamics to be too slow, since new control decisions can only be made when new data is available. One possible workaround to this issue is to use a PC with multiple processors, spawning a new PIV processing engine when the previous image pair is not yet processed. This approach leads to slightly longer latencies, but avoids gaps in the data. In the current configuration, this approach is taken, using two of the four available cores on the host machine.

The pulsed laser firing rate is limited in the current configuration to a maximum of 15 Hz. The camera, synchronized with the laser pulses, captures images at 30 frames per second (fps); each image pair thus takes a total of 66.7 ms to capture. Each image is transferred to PC memory as the next image is captured. Thus, the total transfer time for an image pair is also 66.7 ms. The total time between the beginning of the




Figure 7.6.1: Basic RTPIV architecture.

exposure of the first image of the pair to the time the images are available in memory is 100 ms. This kind of latency is inherent to DPIV; the only way to decrease it is to increase the frame rate of the camera and provide sufficient bandwidth for image transfer. A fundamental lower limit will always exist given the need to wait for two images to be exposed and transferred.

7.6.3 Hardware

Camera The camera utilized for this application, a UT-2230SE, was manufactured by Imaging Development Systems (IDS GmbH) and was chosen for ease of programming rather than for speed or bandwidth considerations. The camera captures a 1024 x 768 monochrome image at 30 fps and communicates with the PC via a USB 2.0 connection. A strobe pulse is generated at a user-specified delay after the beginning of each exposure. This strobe pulse is responsible for triggering the laser timing.

The choice of camera likely has the greatest impact on overall system performance, since it is limited on both frame rate and transfer speed. Additionally, unlike many specialty DPIV cameras, it does not have a built-in double-shutter mode. The choice of a CCD sensor over a CMOS sensor was made because CMOS based systems generally do not possess a global shuttering mode, which is necessary to achieve the asynchronous frame straddling mode that allows the pulse delay to be shorter than the frame rate of the camera.

Illumination The same illumination setup is used for RTPIV as is used in the normal DPIV setup described in Chapter 7.5.1.

Timing The timing for RTPIV is driven by the strobe signal from the camera. This signal is divided by 2 to provide the base frequency at which the lasers are triggered. Laser timing is accomplished through appropriate settings on the timing software described in more detail in Chapter 7.5.1. Using the "Mode

select" option and setting the mode to "RTPIV" is the most straightforward way to accomplish the correct timing settings.

Seeding The same seeding as used for normal DPIV measurements is used for RTPIV. See Chapter 7.5.1 for more details.

Digital I/O Interaction with the actuators in the experiment is provided through a USB digital I/O module, the USB-1208FS, manufactured by Measurement Computing. The state of the digital I/O line can be set by the software, which ultimately opens or closes the solenoid valve in the flow control loop.

Host computer The PC on which the RTPIV software is executed is a custom built PC. The processor is an Intel Core2 Quad Q9400 2.66GHz Quad-Core Processor, installed on an ASUS P5Q motherboard. Memory is provided by G-Skill, with 2x 1GB DDR2-800 installed. The operating system is Windows XPSP3.

7.6.4 Software (QRTPIV)

The software implementation is programmed in C++, with the Qt library providing graphical user interface (GUI) support. The application interface consists of five tabs that control each aspect of the technique. In addition to the main application window, there are also three other windows that can be launched to display specific information.

Camera setup The "Camera setup" tab (Figure 7.6.2) allows the user to specify the camera configuration file, which is generated using the camera software provided by the manufacturer. It allows the strobe output to be toggled and configured. Specifically, the pulse delay setting is crucial so that each image is illuminated at the correct time. Also on this tab is a log window, which displays status messages to the users as the application runs.

Image pair The "Image pair" tab (Figure 7.6.3) shows the image pair that results from pushing the "Grab Image Pair" button on the toolbar. This allows for quick examination of the an image pair to make sure light intensity and particle densities are sufficient.

PIV setup The "PIV setup" tab (Figure 7.6.4) provides all of the configuration settings for the PIV method.

PIV sampling Determines how the subdivision of the image pairs into window occurs.

Processing Image down-sampling and restricted calculation aim to shorten processing times. Multi-grid processing is more computationally intensive than a single pass interrogation of the image pair, but can increase data yield and dynamic spatial range of the measurement.

Peak detection Limits the search area for the correlation peak in the correlation plane, helping to prevent spurious peaks from being identified in regions where the signal to noise ratio is low.

Validation Provides several tests in order to identify those regions where the displacement value found is likely spurious. If the "Interpolate outliers" checkbox is selected, then the discarded vector is replaced by interpolation of the eight nearest neighboring vectors.




Figure 7.6.2: QRTPIV: Camera setup tab.

Conversion Sets the magnification factor and the delay between laser pulses, which allows the pixel displacement vector field to be converted into velocity fields.

Post-processing The "Post-processing" tab (Figure 7.6.5) provides all of the configuration settings for post-processing actions that are applied to the generated vector field.

Calculate statistics This area configures the calculation of statistics on and in the contour that defines the region of interest. Contours can be loaded in the "File" menu, providing easy switching between regions of interest. Horizontal and vertical velocity statistics (average and RMS) are computed inside the contour. Circulation can be computed on the contour and vorticity is computed inside the contour using finite differencing to compute derivatives. If no contour is selected, the statistics are computed for the entire image.

Logging Enables output of experimental data to a user-readable text file.

Plot Sets the update frequency of the line plotter.

Calculated Data Displays the calculated statistics in the region of interest.

Log file format The software has the ability to generate a log file to allow the user to analyze the experiment offline. Currently, the logfile is written in user-readable text with the following column:

Pair Number (I) Resets to 0 every time the "Grab PIV images" button is pressed.

QLivePVT Main application object, sub-classed from the Qt MainWindow class. This is the main object with which the other objects communicate.

Fields will be explained. Then the specific sequence of object interactions required to produce "real-time" vector fields is described. The vector fields from the incoming particle images. To begin, the main object will be involved with production of the code is not intended here, but is useful to document the program flow.

Full documentation of the code is not intended here, but is useful to document the program flow. Functions known as "slots," allowing simple connections between user actions and program execution. In obvious ways. Additionally, use of the Qt library to provide GUI functionality makes communication between the user and the various objects easy. GUI actions generate "signals" which are connected to functions known as "slots," allowing simple connections between user actions and program execution.

7.6.5 Detailed Architecture

Time trace When the "Time trace" button is toggled on the toolbar, a window is opened that displays a time trace of one of the calculated variables in the region of interest (Figure 7.6.9). The time trace data buffer can be moved to a text file if desired.

the vector field, showing the location of the "vector" where the flow statistics are calculated.

Figure 7.6.5: QRTIV: Post-processing tab



Figure 7.6.7: QRTIV: Live video window




Figure 7.6.6: QRTIV: Data out tab



Actuator State (20-35, odd) Flag for actuator command: 1 if actuator on, 0 if actuator off.

Actuator Enabled (20-35, even) Flag for actuator: 1 if enabled, 0 if not.

Activation Time (16-19) Time at which actuator signal is sent to digital I/O unit.

Entanglement (15) Not yet implemented.

Circulation (14) The circulation computed on the boundary of the contour.

Average Y Velocity (13) Average vertical velocity within the region of interest.

Average X Velocity (12) Average horizontal velocity within the region of interest.


PIV Processing Time (11) Time (in ms) for PIV processing engine to generate result once provided with images from memory.

Process Complete (7-10) Time at which PIV processing is complete for the given image pair.

Capture Complete (4-6) Time at which image transfer from camera to PC memory completes. Columns are hours, minutes, seconds and milliseconds from the system clock.

Trigger (2) Flag to 1 when the "Send DMV Trigger" button is pressed. Pair after flag is synchronized with the data acquisition.

Figure 7.6.3: QRTIV: Image pair tab



Live window When the "Live image" button is toggled on the toolbar, a window is opened that displays the vector field results of the PIV processing (Figure 7.6.8) and PIV processing begins. It is possible for the camera to start out of synchronization with the laser pulses, which results in image pairs that are completely uncorrelated. Use of the "Skip frame" button allows the pairing to be shifted, resulting in appropriate images to the processing engine. The vector display window can be configured to color vectors based on several different flow variables. The currently loaded contour (in red) is also overlaid on the vector field.

Live window When the "Live image" button is toggled on the toolbar, a new window is opened in which the "Live image" button is toggled again, the video feed resumes. The "Live image" button is toggled again, the video feed is stopped. If the button is toggled again (and the justifications to focus, laser light intensity, particle density and pulse separation timing can be made. When live video is displayed (Figure 7.6.7). This allows the user to view the particle images so that quick adjustments can be made.

Data Out The "Data Out" tab (Figure 7.6.6) provides configuration for the actuators. Each channel is able to be configured independently. If enabled, the channel compares the quantity related to the threshold values according to the rule selected. If the rule evaluates false, the actuator is turned off (in the current case, this corresponds to the solenoid opening). If the rule evaluates true, the actuator is turned on (in the current case, this corresponds to the solenoid closing). This is activated by a very simple actuator scheme: more complicated schemes can be implemented involving more advanced implementation of control laws and control over frequency and duty-cycle of the actuators. The signal is output from the software into the experiment via the USB digital I/O module described in Section 7.6.3.

Figure 7.6.4: QRTIV: PIV setup tab






Figure 7.6.8: QRTIV Live vector display. Region of interest shown in red.




Figure 7.6.9: QRTIV Live time trace display.

QPIVEngine/PIVEngine PIV processing engine. This object provides methods responsible for the processing of particle images and generating vector results. PIVEngine is a generic processing class, from which QPIVEngine is derived. QPIVEngine also inherits from QThread, which provides the ability for each engine to take advantage of threading.

QPlotEngine Plotting engine, sub-classed from the Qt MainWindow class. This object provides methods allowing the generation of vector field plots to the display.

QXYPlotDlg Line plotting engine, sub-classed from the Qt MainWindow class. This object is responsible for generation and display of the time trace line plotting.

QEyeAcquisition/QAcquisitionEngine Image acquisition engine. QAcquisitionEngine is a generic class which provides basic support for any camera device. QEyeAcquisition is a sub-class of QAcquisitionEngine, providing specific interfaces for the IDS UEye camera currently employed in the QRTIV system.

QDigitIO Digital input/output engine. Provides interface code between the main application and the digital I/O unit described in Section 7.6.3.

PIVAnalysisStep Container class for facilitating multigrid processing.

When the program is first opened, the first task is to initialize the camera. The user selects *Initialize camera* from the *Camera* menu, which is connected to the *send* signal from the main application to the *QEyeAcquisitionEngine* to call the *InitializeCamera*. In this initialization routine, the *is_EnableMessage*

function from the UEye API is called, which allows messages from the camera to be sent, via the operating system, to the main application. This allows the camera to announce each time a new frame is available in memory.

When the user toggles the *Grab Live PIV images* continuously button on the toolbar, the *startGrabbing* function in *QEyeAcquisition* (after some intermediate steps) initiates continuous image acquisition via the *is_CameraVideo* function from the UEye API. As the camera acquires images, event messages regarding the availability of new frames are intercepted by the main application using *waitEvent*. These events initiate the execution of *processFrame* in the *QEyeAcquisitionEngine*, which begins the process of processing image pairs to generate vector results.

The process begins with the transmission of the memory location of the image pair from the *QEyeAcquisitionEngine* to the main application. Once the engine is configured correctly, including sending the image data to the engine, the engine's thread is issued the *QPIVEngine::start* command to initiate the processing. In the current implementation, two engines are available. If one engine is running when the next pair is ready, the second engine is started and runs as a distinct thread, following exactly the same chain of execution.

The *start* command in the *QPIVEngine* sends a signal to the engine to run, starting the thread. *processImagePair* is the function in the engine responsible for the overall processing of the image pair. Based on the user setting for multi-grid processing, the *PIVEngine* either calls *processImagePairMultiGrid* or *processImagePairSinglePass*. In either case, these functions are responsible for executing the cross-correlation and peak finding routines that are the heart of the PIV method.

The forward FFTs are executed in parallel, as is the multiplication of the FFTs, using the open-MP parallelization API. If the multi-grid method is chosen, the image is down-sampled by two for each level of grid resolution specified. The final resolution is that specified by the user in the "PIV Setup" tab (Section 7.6.4). Each coarse level is used to determine an appropriate window offset for the next finest level to maximize data yield and lower the overall signal-to-noise ratio of the measurement. The final vector result is then validated using the tests specified by the user and converted to a velocity through the magnification and laser pulse separation values. Once the entire field is returned, vorticity and flow statistics (average and RMS values) are computed. At this point, processing of the image pair is complete, and the thread sends the finished signal, informing the main application that the thread is complete.

The main application then executes some final post-processing tasks are executed, including drawing the vector plot and time-trace windows. Once these tasks are complete, if there is another image pair waiting, a new PIV engine is spawned to begin processing this next pair.

This whole process repeats until the user toggles the "Grab PIV images continuously" control or an error occurs.

7.7 Baseline Flow

In order to study the baseline flow behavior of the wing model in the tunnel, an airfoil model was used, as described in Section 7.4. This model allows the laser illumination to shine through the model, providing the ability to visualize both sides of the wing.

7.7.1 DDPIV Measurements

DDPIV measurements were taken at two different Reynolds numbers, $Re = 820$ and $Re = 1100$. The angle of attack was varied from 10° to 45° . Three-dimensional data sets are difficult to display fully in a two-dimensional fashion, so only a selection of the results are displayed here. The DDPIV measurement volume

was restricted to a 100 mm cube and placed such that the flow could be visualized from a short distance ahead of the leading edge (roughly 0.2x) to a short distance from the trailing edge (roughly 0.2x).

The first set of images show the development of vorticity over the wing as the angle of attack is increased. Vorticity values are computed by first projecting the randomly spaced vector field generated by the DDPIV processing onto a regularized grid and then computing the appropriate derivatives using central differences. The vorticity is then plotted using iso-surfaces, showing all locations in the flow that have the same vorticity value. These plots are the result of averaging 50 image-pair realizations.

Figure 7.7.1 shows the vorticity iso-surfaces for 10° angle of attack for both Reynolds numbers. As expected, the vorticity generation is strongest near the surface of the wing, with the iso-surfaces becoming relatively smooth as the distance away from the wing increases.




Figure 7.7.1: Iso-surfaces of vorticity at 10° angle of attack.

At 30° angle of attack (Figure 7.7.2), the size of the high vorticity region increases, and there appears to be the beginnings of some kind of more complex average structure occurring at the "corner" of the flow. The measurement volume would need to be moved further downstream in order to investigate the evolution this structure.

At 45° angle of attack (Figure 7.7.3), the vorticity "sheets" lift further off the surface of the wing, with the majority of the vorticity generated at the tip. The flow is fully separated from the airfoil and the strength of the vorticity increases. At the higher Reynolds number, it appears that a more complex structure is evolving within the sheet.

It is also possible to take two-dimensional slices through the three-dimensional dataset. This was done for the three angles of attack presented above, with contours colored by velocity magnitude (in mm/s).




Figure 7.7.2: Iso-surfaces of vorticity at 30° angle of attack.

197



Figure 7.7.3: Iso-surfaces of vorticity at 45° angle of attack.

198

For 10° angle of attack, Figures 7.7.4, 7.7.5, and 7.7.6 show chordwise slices through the dataset at three different spanwise positions, namely the tip, 0.1b, and 0.2b. These figures show a mild region of slower fluid around the wing. At the higher Reynolds number, there appears to be a region of flow that is decelerated even before the leading edge. This is likely due to the influence of the wing tip. Figures 7.7.7, 7.7.8, and 7.7.9 show spanwise slices through the dataset at three different chordwise positions, namely the leading edge, 0.5c and 0.7c. As the flow travels along the chord, it appears that there is potentially a distinct transition between flow associated with the wing tips and the flow associated with the leading edge.




Figure 7.7.4: Chordwise velocity slice at 10° angle of attack - tip station.

For 30° angle of attack, Figures 7.7.10, 7.7.11, and 7.7.12 show chordwise slices through the dataset at three different spanwise positions, namely the tip, 0.1b, and 0.2b. Compared to the 10° case, the flow is definitely separated from the wing, which is to be expected. The separation region begins at the leading edge for the lower Reynolds number and slightly behind the leading edge for the higher. The separation zone extends all the way to the wingtip at mid-chord. Interestingly, in the lower Reynolds number case there appears to be a region of accelerated fluid above the wing the doesn't appear to be present in the higher Reynolds number flow. The velocity scaling was chosen to correspond with the maximum velocity in the higher Reynolds number case. Figures 7.7.13, 7.7.14, and 7.7.15 show spanwise slices through the dataset at three different chordwise positions, namely the leading edge, 0.5c and 0.7c. The region of acceleration looks to extend out past the tip (Figure 7.7.14) to about 0.15b by mid-chord and even moving below the wing by 0.7c.

For 45° angle of attack, Figures 7.7.16, 7.7.17, and 7.7.18 show chordwise slices through the dataset at three different spanwise positions, namely the tip, 0.1b, and 0.2b. At this point, the flow is massively separated, with both Reynolds number cases showing very similar features. In fact, there appears to be only very slight differences between the two data sets. It is possible that at this large angle of attack, the

199




Figure 7.7.5: Chordwise velocity slice at 10° angle of attack - 0.1b station.




Figure 7.7.6: Chordwise velocity slice at 10° angle of attack - 0.2b station.

200




Figure 7.7.7: Spanwise velocity slice at 10° angle of attack - leading edge.




Figure 7.7.8: Spanwise velocity slice at 10° angle of attack - 0.5c.

201




Figure 7.7.9: Spanwise velocity slice at 10° angle of attack - 0.7c.




Figure 7.7.10: Chordwise velocity slice at 30° angle of attack - tip station.

202




Figure 7.7.11: Chordwise velocity slice at 30° angle of attack - 0.1b station.



Figure 7.7.12: Chordwise velocity slice at 30° angle of attack - 0.2b station.

203




Figure 7.7.13: Spanwise velocity slice at 30° angle of attack - leading edge.




Figure 7.7.14: Spanwise velocity slice at 30° angle of attack - 0.5c.

204




Figure 7.7.15: Spanwise velocity slice at 30° angle of attack - 0.7c.

freestream flow beyond this measurement volume is influenced in such a way as to create local conditions that are very similar. The vorticity iso-surfaces shown in Figure 7.7.3 are certainly different, so the two flows are not identical. Figures 7.7.19, 7.7.20, and 7.7.21 show spanwise slices through the dataset at three different chordwise positions, namely the leading edge, 0.5c and 0.7c. In these plots, some larger differences can be noticed, namely that the stagnant region on the suction surface at the leading edge is much larger for the higher Reynolds number case. As the flow evolves along the wing chord, however, the separation bubbles have very similar shape and extent. It is possible that as Reynolds number increases, the flow may be reaching some kind of saturation limit where the additional energy in the freestream feeds an instability process, which might be the reason for the more complicated structures in the vorticity field shown in Figure 7.7.3. Two Reynolds number cases isn't enough to answer the question exhaustively, however, so more study would have to be done when the DDPIV system is again available.

7.7.3 DPIV Measurements

DPIV was performed on the mid-chord actuator model to characterize the baseline flow characteristics at the mid-span of the wing. The angle of attack was varied from 0° to 15° in 2.5° increments and from 15° to 45° in 5° increments. An exhaustive display of all of the PIV data is not going to be attempted here. Instead, it should be sufficient to show some representative plots to demonstrate how the data is to be used.

A sequence of 200 DPIV image pairs were obtained for each angle of attack at two different flow speeds. Each image pair was processed using the PIVview software package. The shadow cast by the wing was eliminated through the use of an image mask, preventing processing in that region. A multi-grid processing scheme was used, with interrogation windows of size 32 × 32, with 50% overlap. Once the vector field was



Figure 7.7.16: Chordwise velocity slice at 45° angle of attack - tip station.



Figure 7.7.17: Chordwise velocity slice at 45° angle of attack - 0.1b station.




Figure 7.7.18: Chordwise velocity slice at 45° angle of attack - 0.2b station.




Figure 7.7.20: Spanwise velocity slice at 45° angle of attack - 0.5c.




Figure 7.7.19: Spanwise velocity slice at 45° angle of attack - leading edge.




Figure 7.7.21: Spanwise velocity slice at 45° angle of attack - 0.7c.

found, a validation step was performed in order to identify outlying vectors. Any outliers were replaced by interpolation of nearest neighbors. The processing was executed in a batch mode, which gives a composite average of the 200 vector fields in addition to the individual fields.

The vortex shedding from the wing is obviously an unsteady event, so the computation of the average field is of limited use. One way it is potentially useful is for the identification of regions of the flow where a particular flow variable is of interest. For instance, the vorticity shed from the leading edge is of interest, and so the region with large average vorticity is a region which might be useful to probe further. Figure 7.7.22 shows the average result for 45° angle of attack with $Re = 800$. In this plot, the wing is oriented as it is in the tunnel, with the suction side down. The contours are of vorticity and are scaled such that the negative vorticity shed from the trailing edge is not visible. There is obviously a region of strong positive vorticity shed from the leading edge of the wing.




Figure 7.7.22: Average DPIV field for 45°, $Re \approx 800$

If the data set is "probed" near the center of this region of high average vorticity for each of the 200 vector fields (approximately at (160, -160) in Figure 7.7.22), a time trace of the vorticity at that point can be obtained. This time trace is shown in Figure 7.7.23.

With the time trace, the single-sided spectrum is easily obtained by computing its FFT. Before computing the FFT, the mean was subtracted from the signal, since it is large compared with the fluctuations. This spectrum is shown in Figure 7.7.24. This spectral content provides insight into possible frequencies where the flow may be receptive to actuation.




Figure 7.7.23: Vorticity time trace for 45°, $Re \approx 800$




Figure 7.7.24: Single-sided spectrum of vorticity for 45°, $Re \approx 800$

7.7.3 Force Measurements

Forces were measured on the wing using the six-axis force balance described in Chapter 7.5.3. Lift and drag are computed by Equations 7.7.1 and 7.7.2, respectively. Force coefficients are computed as in Equation 7.1.1.

$$L = F_y \cos(\alpha) + F_x \sin(\alpha) \quad (7.7.1)$$

$$D = F_y \sin(\alpha) - F_x \cos(\alpha) \quad (7.7.2)$$

Lift variation with angle of attack is shown in Figure 7.7.25. Drag variation with angle of attack is in 7.7.2. The lift curve shows a very different picture from normal wings at higher Reynolds number. Here, the lift increases fairly steadily to relatively high angles of attack and shows no sharp stall event. The drag shows a steady increase as angle of attack increases, which is obviously expected. It is important to note that these measurements are influenced by the presence of the sting. A reasonable method of subtracting out the sting's influence is still under consideration.




Figure 7.7.25: Baseline lift curve.




Figure 7.7.26: Baseline drag curves

8 Publications resulting from this award

Ph.D. Theses

1. S. Ahuja. Reduction methods for feedback stabilization of fluid flows. Ph.D. Thesis, Princeton University, 2009.
2. W.T. Joe. Optimized Feedback Control of Vortex Shedding on an Inclined Flat Plate. Ph.D. Thesis, California Institute of Technology, 2010.
3. Z. Ma. Reduction and reconstruction methods for simulation and control of fluids. Ph.D. Thesis, Princeton University, 2010.
4. K. Taira. The immersed boundary projection method and its application to simulation and control of flows around low-aspect-ratio wings. Ph.D. Thesis, California Institute of Technology, 2008.

Journal Articles

1. S. Ahuja, I.G. Kevrekidis, & C.W. Rowley, "Template-based stabilization of relative equilibria in systems with continuous symmetry", *J. Nonlin. Sci.* 17(2):109-143, 2007.
2. S. Ahuja and C.W. Rowley. Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators. *J. Fluid Mech.* 645:447-478, 2010.
3. L. N. Cattafesta III, Q. Song, D.R. Williams, C. W. Rowley, and F. S. Alvi. Active control of flow-induced cavity oscillations. *Progress In Aerospace Sciences*, 44:479-502, 2008.
4. K. Chen, T. Colonius, and K. Taira. The leading-edge vortex and quasi-steady vortex shedding on an accelerated plate. *Physics of Fluids*, 22(3). Art. No. 033601, 2010.
5. T. Colonius and K. Taira. "A fast immersed boundary method using a nullspace approach and multi-domain far-field boundary conditions". *Computer Methods in Applied Mechanics and Engineering*, 197(23-28): 2131-2146, 2008.
6. T. Colonius and D.R. Williams. Control of vortex shedding from two- and three dimensional airfoils. Submitted to *Phil. Trans. Roy. Soc. A*, 2010.
7. J.A. Franck and T. Colonius. Compressible large-eddy simulation of separation control on a wall-mounted hump. *AIAA Journal*, 48(6):1098-1107, 2010.
8. W.-T. Joe and T. Colonius. Feedback control of vortex shedding from an inclined flat plate. To appear, *Theoretical and Computational Fluid Dynamics*, 2010.

213

22. K. Taira and T. Colonius. The Immersed Boundary Method: A Projection Approach. *Journal of Computational Physics*, 225:2118-2137, 2007.
23. K. Taira and T. Colonius. Three-dimensional flows around low-aspect-ratio flat-plate wings at low Reynolds numbers. *J. Fluid Mech.* 623:187-207, 2009.
24. K. Taira and T. Colonius. On the Effect of Tip Vortices in Low-Reynolds-Number Post-Stall Flow Control. *AIAA J.* 47(3):749-756, 2009.
25. K. Taira, C. W. Rowley, and T. Colonius and D.R. Williams. Lift enhancement for low-aspect-ratio wings with periodic excitation. *AIAA Journal*, 48(8):1785-1790, 2010.
26. M. Wei and C.W. Rowley. Low-dimensional models of a temporally evolving free shear layer. *J. Fluid Mech.* 618:113-134, 2009.
27. D. Williams, G. Tadmor, T. Colonius, W. Kerstena, V. Quach and S. Buntain. Lift Response of a Stalled Wing to Pulsatile Disturbances. *AIAA Journal*. 47(12):3031-3037, 2009.

Conference Proceedings

1. S. Ahuja, C. Rowley, I. Kevrekidis, and M. Wei, T. Colonius, and G. Tadmor, Low-Dimensional Models for Control of Leading-Edge Vortices: Equilibria and Linearized models, *AIAA Paper* 2007-709, 2007.
2. S. Ahuja and C. W. Rowley. Low-dimensional models for feedback stabilization of unstable steady states, *AIAA Paper* 2008-553, 2008.
3. K. Aleksic, R. King, B. R. Noack, O. Lehmann, M. Morzynski & G. Tadmor, Nonlinear model predictive control based on a low dimensional model of fluid flow, 2nd Int. Conf. on Jets, Wakes and Separated Flows, Berlin, Germany, September 16-19, 2008.
4. G. Brès, D. Williams, and T. Colonius. Numerical simulations of natural and actuated flow over a 3D, low-aspect-ratio airfoil, *AIAA Paper* 2010-4713, 2010.
5. S.L. Brunton, C.W. Rowley, K. Taira, T. Colonius, J. Collins, and D.R. Williams. Unsteady aerodynamic forces on small-scale wings: experiments, simulations, and models. *AIAA Paper* 2008-520, 2008.
6. S. Brunton and C. W. Rowley. Modeling the unsteady aerodynamic forces on small-scale wings, *AIAA Paper* 2009-1127, 2009.
7. S. Brunton and C.W. Rowley. Unsteady aerodynamic models for agile flight at low Reynolds numbers. *AIAA Paper* 2010-552, 2010.

215

9. M. Luchtenburg, B. Gunther, B.R. Noack, R. King & G. Tadmor, A Generalized Mean Field Model for the natural and high-frequency actuated flow around a high lift configuration, *J. Fluid Mechanics* 623: 283-316, 2009.
10. Z. Ma and C. W. Rowley. Lie-Poisson integrators: a Hamiltonian, variational approach. *International Journal for Numerical Methods in Engineering*, in press 2010.
11. Z. Ma, C.W. Rowley, and G. Tadmor. Snapshot-based balanced truncation for linear time-periodic systems. *IEEE Trans. Auto. Control*, 55(2):469-473, 2010.
12. Z. Ma, S. Ahuja, and C.W. Rowley. Reduced order models for control of fluids using the eigensystem realization algorithm. *Theoretical and Computational Fluid Dynamics*, in press, 2010.
13. M. Morzynski, B.R. Noack & G. Tadmor, "Global stability analysis and reduced order modeling for bluff-body flow control", *J. of Theoretical and Applied Mechanics* 45:621-642, 2007.
14. B.R. Noack, M. Schlegel, B. Ahlborn, G. Mutschke, M. Morzynski, P. Comte and G. Tadmor. A finite-time thermodynamics of unsteady flows. *J. Nonequilib. Thermodyn.* 33:103-148, 2008.
15. B. R. Noack, M. Schlegel, M. Morzynski and G. Tadmor, System reduction strategy for Galerkin models of fluid flows, Submitted to *Int. J. of Numer. Meth. in Fluids*, 2009.
16. M. Pastoor, L. Henning, B.R. Noack, R. King & G. Tadmor, Feedback shear layer control for bluff body drag reduction. Manuscript submitted to *J. Fluid Mech.* 608:161 - 196, 2008.
17. C.W. Rowley, D.R. Williams, T. Colonius, R.M. Murray, and D.G. MacMynowski. Linear models for control of cavity flow oscillations. *Journal of Fluid Mechanics*, 547:317-330, 2006.
18. C. W. Rowley, I. Mezic, S. Bagheri, P. Schlatter, and D. S. Henningson. Spectral analysis of nonlinear flows. *Journal of Fluid Mechanics*, 641:115-127, 2009.
19. W. Stankiewicz, M. Morzynski, B. R. Noack and G. Tadmor, Reduced order Galerkin models of flow around NACA 0012 Airfoil, *Math. Modeling and Anal.* 13:113 - 122, 2008.
20. G. Tadmor, O. Lehmann, B. R. Noack and M. Morzynski, Mean Field Representation of the Natural and Actuated Cylinder Wake, *Physics of Fluids* 22:034102, 2010.
21. G. Tadmor, O. Lehmann, B. R. Noack, L. Cordier, J. Delville, J.-P. Bonnet and M. Morzynski, Reduced order models for closed-loop wake control: A review, *Phil. Trans. Roy. Soc. A* (2010), submitted.

214

8. J. Franck and T. Colonius. Oscillatory control and the effects of actuation on a wall-mounted hump model, *AIAA Paper* 2009-4019, 2009.
9. W.-T. Joe, K. Taira, T. Colonius, D. G. MacMynowski and G. Tadmor, Closed-Loop Control of Vortex Shedding on a Two-Dimensional Flat-Plate Airfoil at a Low Reynolds Number, *AIAA Paper* 2008-597, 2008.
10. W.-T. Joe, T. Colonius, and D. MacMynowski. Optimized control of vortex shedding from an inclined flat plate, *AIAA* 2009-4027, 2009.
11. W.-T. Joe, T. Colonius, and D.G. MacMynowski. Optimized waveforms for feedback control of vortex shedding. In *Active Flow Control II*, R. King, Ed. Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 108, pp 391-404, 2010.
12. R. King, M. Seibold, O. Lehmann, B. R. Noack and G. Tadmor, "Non linear flow control based on a low dimensional model of fluid flow", in *Control and Observer Design for Nonlinear Finite and Infinite Dimensional Systems*, T. Meurer, K. Graichen and E.D. Gilles, Editors, Lecture notes in Control & Information Science V. 322, pp. 369-386, Springer, 2005.
13. O. Lehmann, M. Luchtenburg, B.R. Noack, R. King, M. Morzynski & G. Tadmor, Wake stabilization using POD Galerkin models with interpolated modes, 44th IEEE Conference on Decision and Control and European Control Conference "ECC 2005", Seville, Spain, 12-15, December 2005.
14. M. Luchtenburg, G. Tadmor, O. Lehmann, B. R. Noack, R. King and M. Morzynski, Tuned POD Galerkin models for transient feedback regulation of the cylinder wake, *AIAA Paper* 2006-1407, 2006.
15. Z. Ma and C.W. Rowley. Low-dimensional linearized models for systems with periodic orbits, with application to the Ginzburg-Landau equation, *AIAA Paper* 2008-4196, 2008.
16. M. Morzynski, W. Stankiewicz, B. R. Noack, F. Thiele and G. Tadmor, Generalized Mean field Model with Continuous Mode Interpolation for Flow Control, *AIAA Paper* 2006-3488, 2006.
17. M. Munson, W. Dickson, T. Colonius, and M. Gharib. A new low Reynolds number facility for active flow control, *AIAA Paper* 2008-694, 2008.
18. M. Pastoor, R. King, B.R. Noack & G. Tadmor, Observers & feedback control for shear layer vortices, 44th IEEE Conference on Decision and Control and European Control Conference ECC 2005, pp.506-511. Seville, Spain, 12-15 Dec., 2005.
19. M. Pastoor, B. R. Noack, R. King and G. Tadmor, "Spatiotemporal Waveform Observers and Feedback in Shear Layer Control", *AIAA Paper* 2006-1402, 2006.

216

20. C. W. Rowley, S. Ahuja, K. Taira, and T. Colonius. Closed-loop control of leading edge vorticity on a 3d wing: Simulations and low-dimensional models, AIAA Paper 2008-3981, 2008.
21. G. Tadmor, J. Gonzalez, O. Lehmann, B.R. Noack, M. Morzynski & W. Stankiewicz. Shift modes and transient dynamics in low order design oriented Galerkin models, AIAA Paper 2007-111, 2007.
22. G. Tadmor, M. D. Centuri, B. R. Noack, O. Lehmann, M. Luechtenburg and M. Morzynski, A Low Order Galerkin Design Model for Feedback Flow Stabilization Over a 2-D Airfoil, AIAA Paper 2007-1313, 2007.
23. G. Tadmor, D. Bissex, B. Noack, M. Morzynski, T. Colonius and K. Taira. Temporal harmonic specific POD mode extraction, AIAA Paper 2008-4190, 2008.
24. G. Tadmor, D. Bissex, B. Noack, M. Morzynski, T. Colonius and K. Taira. Fast Approximated POD for a Flat Plate Benchmark with a Time Varying Angle of Attack, AIAA Paper 2008-4191, 2008.
25. G. Tadmor, D. Williams, J. Collins, T. Colonius, and C. Rowley. Control of a semi-circular planform wing in a gusting unsteady free stream flow II: modeling and feedback design, AIAA Paper 2008-3977, 2008.
26. K. Taira, C. Rowley, and T. Colonius. Feedback Control of High-Lift State for a Low-Aspect-Ratio Wing, AIAA Paper, 2010-357, 2010.
27. K. Taira, W. Dickson, T. Colonius, and C. Rowley. Unsteadiness in flow over a flat plate at angle-of-attack at low Reynolds numbers, AIAA Paper 2007-710, 2007.
28. K. Taira, C.W. Rowley, and T. Colonius. Lock-on to a high-lift state with oscillatory forcing in a three-dimensional wake flow, In Active Flow Control II., R. King, Ed., Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 108, pp. 81-94, 2010.
29. D. Williams, S. Doshi, J. Collins, and T. Colonius. Control of spanwise distribution of circulation on NACA 0012 and Flat Plate Wings, AIAA Paper 2007-1121, 2007.
30. D. R. Williams, J. Collins, C. Jankhot, T. Colonius and G. Tadmor, Control of Flow Structure on a Semi-Circular Planform Wing, AIAA Paper 2008-634, 2008.
31. D. Williams, J. Collins, G. Tadmor, and T. Colonius. Control of a semi-circular planform wing in a gusting unsteady free stream flow I: experimental issues, AIAA Paper 2008-3976, 2008.
32. D. Williams, J. Collins, C. Jankhot, T. Colonius, and G. Tadmor. Control of flow structure on a semi-circular planform wing, AIAA Paper 2008-597, 2008.

217

9 Acknowledgements

This work was sponsored by a DoDMURI administered by the Air Force Office of Scientific Research (grant FA9550-05-1-0369) with Dr. Fariba Farooq as technical monitor. The views and conclusions contained herein are those of the authors and should not be interpreted as necessarily representing the official policies or endorsements, either expressed or implied, of the Air Force Office of Scientific Research or the U.S. Government.

The authors are deeply indebted to numerous students, colleagues, and staff members who made substantial contributions to this work. This includes students who were support in part or in full on this grant, many of whose work appears in this report: Dr. Sunil Ahuja, Mr. Daniel Bissex, Mr. Seth Buntain, Ms. Donatella Centuri, Mr. Jesse Collins, Mr. Sudeep Doshi, Dr. Jennifer Franck, Dr. Won Tae Joe, Mr. Wesley Kerensa, Mr. Daegyoung Kim, Dr. Zhanhua Ma, Mr. Avshalom Manela, Mr. Juan Mellì, Mr. Matthew Munson, Dr. Kunihiko Taira, Mr. Vien Quach, Prof. Mingjun Wei, and Mr. Ju Zhang. We also thank our colleagues, Prof. Rudibert King, Dr. Douglas MacMynowski, Prof. Bernd Noack, and Mr. Jens Pfeiffer for their contributions.

219

33. D. Williams, V. Quach, W. Kerstens, S. Buntain, G. Tadmor, C. Rowley, and T. Colonius. Low Reynolds number wing response to an oscillating freestream with and without feed forward control, AIAA Paper 2009-143, 2009.
34. D. Williams, S. Buntain, V. Quach, and W. Kerstens. Flow field structures behind a wing in oscillating freestream, AIAA Paper 2009-3690, 2009.
35. D. Williams, W. Kerstens, J. Pfeiffer, R. King, G. Tadmor, and T. Colonius. Closed-Loop Control of a Wing in an Unsteady Flow, AIAA Paper 2010-358, 2010.
36. D. Williams and S. Buntain. Wind Tunnel Experiments on a Rapidly Pitching Plate, AIAA Paper 2010-4279, 2010.
37. D. Williams, W. Kerstens, J. Pfeiffer, R. King, and T. Colonius. Closed loop control of a wing's lift for 'gust' suppression, AIAA Paper 2010-4969, 2010.
38. D. Williams, W. Kerstens, J. Pfeiffer, R. King, and T. Colonius. Unsteady lift suppression with a robust closed loop controller. In Active Flow Control II., R. King, Ed., Notes on Numerical Fluid Mechanics and Multidisciplinary Design, Vol. 108, pp. 19-30, 2010.

218

References

- Ahija, K. & Barrin, R. (1984). 'Control of flow separation by sound', AIAA Paper 84-2298.
- Ahija, S. (2009). Reduction methods for feedback stabilization of fluid flows, PhD thesis, Princeton University.
- Ahija, S. & Rowley, C. W. (2008). 'Low-dimensional models for feedback stabilization of unstable steady states', AIAA Paper 2008-553.
- Ahija, S. & Rowley, C. W. (2009). 'Feedback control of unstable steady states of flow past a flat plate using reduced-order feedback control of unstable steady states of flow past a flat plate using reduced-order estimators', to be submitted to JFM.
- Ahija, S. & Rowley, C. W. (2010a). 'Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators', *J. Fluid Mech.* 645, 447-478.
- Ahija, S. & Rowley, C. W. (2010b). 'Feedback control of unstable steady states of flow past a flat plate using reduced-order estimators', *J. Fluid Mech.* 645, 447-478.
- Ahija, S., Rowley, C. W., Kevrekidis, I. G., Colonius, T. & Tadmor, G. (2007). 'Low-dimensional models for control of leading-edge vortices: equilibria and linearized models', AIAA Paper 2007-709.
- Amisay, M. & Glezzer, A. (2002a). 'Controlled transients of flow reattachment over stalled airfoils', *Int. J. Heat Fluid Flow* 23(5), 690-699.
- Amisay, M. & Glezzer, A. (2002b). 'Role of actuation frequency in controlled flow reattachment over a stalled airfoil', *AIAA Journal* 40(2), 209-216.
- Amisay, M. & Glezzer, A. (2006). 'Flow transients induced on a 2D airfoil by pulse-modulated actuation', *Exp. Fluids* 40(2), 329-331.
- Anderson, J. D. (1999). *Aircraft performance and design*, McGraw-Hill, Boston.
- Aubry, N., Holmes, P., Lumley, J. & Stone, E. (1988). 'The dynamics of coherent structures in the wall region of a turbulent boundary layer', *J. Fluid Mech.* 192, 115-173.
- Bamieh, B. & Pearson, J. B. (1992). 'The H2 problem for sampled-data systems', *Sys. Control Lett.* 19, 1-12.
- Bar-Lev, M. & Yang, H. (1975). 'Initial flow field over an impulsively started circular cylinder', *J. Fluid Mech.* 72, 625-647.
- Barkley, D. (2006). 'Linear analysis of the cylinder wake mean flow', *Europ. J. Mech.* 75(5), 750-756.
- Barkley, D. & Henderson, R. D. (1996). 'Three-dimensional Floquet stability analysis of the wake of a circular cylinder', *J. Fluid Mech.* 322, 215-241.
- Bearman, P. W. (1967). 'On vortex street wakes', *Journal of Fluid Mechanics* 28, 625-641.
- Becker, R., King, R., Petz, R. & W., N. (2007). 'Adaptive closed-loop separation control on a high-lift configuration using extremum seeking', *AIAA Journal* 45(6), 1312-1322.

220

Belov, A., Marinelli, L. & Jameson, A. (1995), 'A new implicit algorithm with multigrid for unsteady incompressible flow calculations', *AIAA Paper* 95-0049.

Benard, N., Bonnet, J., Moriau, E., Griffin, J. & Cattafesta, L. N. (2010), 'On the benefits of hysteresis effects for closed-loop separation control using plasma actuators', *AIAA Paper* 2010-4259.

Benard, N., Braud, P., Jolibois, J. & Moriau, E. (2008), 'Airflow reattachment along a NACA0015 airfoil by surface dielectric barrier discharge actuator - time resolve particle image velocimetry investigation', *AIAA Paper* 2008-4202.

Benard, N., Moriau, E., Griffin, J. & Cattafesta, L. N. (2010), 'Slope seeking for autonomous lift improvement by plasma surface discharge', *Exp. Fluids* **48**, 791-808.

Berenger, J. (1994), 'A perfectly matched layer for the absorption of electromagnetic waves', *J. Comput. Phys.* **114**, 185-200.

Bewley, T. R., Moin, P. & Toman, R. (2001), 'DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms', *J. Fluid Mech.* **447**, 179-225.

Bewley, T. R., Toman, R. & Ziane, M. (2000), 'A general framework for robust control in fluid mechanics', *Physica D-Nonlinear Phenomena* **138**(3-4), 360-392.

Beyer, R. & LeVoque, R. (1992), 'Analysis of a one-dimensional model for the immersed boundary method', *SIAM J. Numer. Anal.* **29**(2), 332-364.

Birch, J., Dickson, W. B. & Dickinson, M. (2004a), 'Force production and flow structure of the leading edge vortex on flapping wings at high and low Reynolds numbers', *J. Exp. Biol.* **207**, 1063-1072.

Birch, J. M. & Dickinson, M. H. (2001), 'Spanwise flow and the attachment of the leading-edge vortex on insect wings', *Nature* **412**, 729-733.

Birch, J. M., Dickson, W. B. & Dickinson, M. H. (2004b), 'Force production and flow structure of the leading edge vortex on flapping wings at high and low Reynolds numbers', *Journal of Experimental Biology* **207**, 1063-1072.

Bochev, R. & Lehoucq, R. (2005), 'On the finite element solution of the pure Neumann problem', *SIAM Rev.* **47**(1), 50-66.

Braza, M., Faghani, D. & Persillon, H. (2001), 'Successive stages and role of natural vortex dislocations in three-dimensional wake transition', *Journal of Fluid Mechanics* **439**, 1-41.

Brennel, M. & Mueller, T. J. (1988), 'Boundary-layer measurements on an airfoil at low Reynolds number', *J. Aircraft* **25**(7), 612-617.

Brown, D., Cortez, R. & Minion, M. (2001), 'Accurate projection methods for the incompressible Navier-Stokes equations', *J. Comput. Phys.* **168**, 464-499.

Brzezinski, D. & Gleser, A. (2006), 'Transient separation control using pulse-combustion', *AIAA Paper* 2006-3024.

Cain, A., Perziger, J. & Reynolds, W. (1984), 'Discrete orthogonal function expansions for non-uniform grids using the fast Fourier transform', *J. Comput. Phys.* **56**, 272-286.

Coutanceau, M. & Bouard, R. (1977b), 'Experimental determination of the main features of the viscous flow in the wake of a circular cylinder in uniform translation. part 2. unsteady flow', *J. Fluid Mech.* **79**(2), 257-272.

Dehbi, J. O. (2009), 'Optimal vortex formation as a unifying principle in biological propulsion', *Annu. Rev. Fluid Mech.* **41**, 17-33.

Darabi, A. & Wygnanski, I. (2004a), 'Active management of naturally separated flow over a solid surface. Part 1', *J. Fluid Mech.* **510**, 105-129.

Darabi, A. & Wygnanski, I. (2004b), 'Active management of naturally separated flow over a solid surface. Part 2', *J. Fluid Mech.* **510**, 131-144.

Deane, A. E., Kevrekidis, I. G., Karniadakis, G. E. & Orszag, S. A. (1991), 'Low-dimensional models for complex geometry flows: Application to grooved channels and circular cylinders', *Phys. Fluids A* **3**(10), 2337-54.

Dennis, S. & Chang, G. (1970), 'Numerical solutions for steady flow past a circular cylinder at Reynolds number up to 100', *J. Fluid Mech.* **42**, 471-489.

Dickinson, M. H. & Göttsch, K. G. (1993a), 'Unsteady aerodynamic performance of model wings at low Reynolds numbers', *J. Exp. Biol.* **174**, 45-64.

Dickinson, M. H. & Göttsch, K. G. (1993b), 'Unsteady aerodynamic performance of model wings at low Reynolds numbers', *Journal of Experimental Biology* **174**, 45-64.

Dixon, C. (1969), 'Lift augmentation by lateral blowing over a lifting surface', *AIAA Paper* 69-193.

Dong, H., Mittal, R. & Najjar, F. M. (2006), 'Wake topology and hydrodynamic performance of low-aspect-ratio flapping foils', *Journal of Fluid Mechanics* **566**, 309-343.

Dowling, A. & Morgans, A. (2005), 'Feedback control of combustion oscillations', *Annu. Rev. Fluid Mech.* **37**, 151-182.

Duraisamy, K. & Baeder, J. D. (2006), 'Numerical simulation of the effects of spanwise blowing on tip vortex formation', *Journal of Aircraft* **43**(4), 996-1006.

Englar, R. (2000), 'Circulation control pneumatic aerodynamics: Blown force and moment augmentation and modification, past, present & future', *AIAA Paper* 2000-2541.

Fedlin, E., Vertsicos, R., Orlandi, P. & Mohd-Yousuf, J. (2000), 'Combined immersed-boundary finite-differences methods for three-dimensional complex flow simulations', *J. Comput. Phys.* **161**, 35-60.

Fage, A. & Johansen, F. (1927), 'On the flow of air behind an inclined flat plate of infinite span', *Proc. Roy. Soc. (London, Ser. A)* **116**(773), 170-197.

Farhood, M., Beck, C. L. & Dullerud, G. E. (2005), 'Model reduction of periodic systems: a lifting approach', *Automatica* **41**, 1083-1090.

Frank, (1996), 'Optimale alphasapproximation von totzeitgliedern im reitbereich', *Automatisierungstechnik* **44**, 42-43.

Campbell, J. F. (1976), 'Augmentation of vortex lift by spanwise blowing', *Journal of Aircraft* **13**(9), 727-732.

Carr, L. (1988a), 'Progress in the analysis and prediction of dynamic stall', *J. of Aircraft* **25**(1), 6-17.

Carr, L. W. (1988b), 'Progress in analysis and prediction of dynamic stall', *Journal of Aircraft* **25**(1), 6-17.

Chahkai, Y. & Van Doorn, P. (2006), Model reduction of time-varying systems, in 'Dimension reduction of large-scale systems', Springer-Verlag, pp. 131-148.

Cheng, W., Giraldo, F. & Perot, B. (2002), 'Analysis of an exact fractional step method', *Journal of Computational Physics* **180**(1), 183-199.

Chen, K. K., Colonius, T. & Taira, K. (2010), 'The leading-edge vortex and quasi-steady vortex shedding on an accelerating plate', *Phys. Fluids* **22**(3), Art. no. 033601.

Chorin, A. (1967), 'A numerical method for solving incompressible viscous flow problems', *J. Comput. Phys.* **2**, 50-66.

Chorin, A. J. (1968), 'Numerical solution of the Navier-Stokes equations', *Mathematics of Computation* **22**, 745-762.

Cierpka, C., Weier, T. & Gerbeth, G. (2008), 'Evolution of vortex structures in an electromagnetically excited separated flow', *Exp. Fluids* **45**, 943-953.

Codina, R. (2001), 'Pressure stability in fractional step finite element methods for incompressible flows', *J. Comput. Phys.* **170**, 112-140.

Collins, W. & Dennis, S. (1973), 'Flow past an impulsively started circular cylinder', *J. Fluid Mech.* **60**, 105-127.

Colonius, T. (2004), 'Modeling artificial boundary conditions for compressible flow', *Annu. Rev. Fluid Mech.* **36**, 315-345.

Colonius, T. & Rade, H. (2002), 'A super-grid-scale model for simulating compressible flow on unbounded domains', *J. Comput. Phys.* **182**, 191-212.

Colonius, T., Rowley, C. W., Tadmor, G., Williams, D. R., Taira, K., Dickson, W. B., Gharib, M. & Dickinson, M. (2006), 'Closed-loop control of leading-edge and tip vortices for small UAV', *Conference on Active Flow Control, DFG, Berlin, Sep. 27-29*.

Cosyn, P. & Vierendeels, J. (2006), 'Numerical investigation of low-aspect-ratio wings at low Reynolds numbers', *Journal of Aircraft* **43**(3), 713-722.

Cortez, R., Keumoussakos, P. & Salhi, M. (2001), 'Vortex methods with spatially varying cores', *J. Comput. Phys.* **170**, 112-140.

Coutanceau, M. & Bouard, R. (1977a), 'Experimental determination of the main features of the viscous flow in the wake of a circular cylinder in uniform translation. part 1. steady flow', *J. Fluid Mech.* **79**(2), 231-256.

Freyer, P., Finaish, F. & Bank, W. (1987), 'Further visualization of combined wing tip and starting vortex systems', *AIAA Journal* **25**(9), 1153-1159.

Fry, S. N., Sayaman, R. & Dickinson, M. H. (2003), 'The aerodynamics of free-flight maneuvers in drosophila', *Science* **300**, 495-498.

Gerhard, J., Pastoor, M., King, R., Noack, B., Dillmann, A., Morzyński, M. & Tadmor, G. (2003), Model-based control of vortex shedding using low-dimensional Galerkin models, in '33rd AIAA Fluids Conference and Exhibit', *AIAA Paper* 2003-4262.

Gharib, M., Rambod, E. & Shariff, K. (1998), 'A universal time scale for vortex ring formation', *Journal of Fluid Mechanics* **360**, 121-140.

Gleser, A., Amis, M. & Hinojosa, A. M. (2005), 'Aspects of low- and high-frequency actuation for aerodynamic flow control', *AIAA Journal* **43**(7), 1501-1511.

Glowinski, R., Paris, T. & Périaux, J. (1998), 'Distributed Lagrange multiplier methods for incompressible flow around moving rigid bodies', *Comput. Meth. Appl. Mech. Engrg.* **151**, 181-194.

Goldstein, D., Handler, R. & Sirovich, L. (1993), 'Modeling a no-slip flow boundary with an external force field', *J. Comput. Phys.* **105**, 354-366.

Greenblatt, D., Goksel, B., Reichenberg, J., Schule, C. Y., Romann, D. & Paschereit, C. O. (2008), 'Dielectric barrier discharge flow control at very low flight Reynolds numbers', *AIAA J.* **46**(6), 1528-1541.

Greenblatt, D. & Wygnanski, I. J. (2000a), 'The control of flow separation by periodic excitation', *Prog. Aerosp. Sci.* **36**(7), 487-545.

Greenblatt, D. & Wygnanski, I. J. (2000b), 'The control of flow separation by periodic excitation', *Prog. Aerosp. Sci.* **36**, 487-545.

Griffin, O. M. (1978), 'Universal Strouhal number for locking-on of vortex shedding to vibrations of bluff cylinders', *Journal of Fluid Mechanics* **85**(Apr), 591-606.

Guermard, J.-L. & Quartapella, L. (1998), 'On stability and convergence of projection methods based on pressure Poisson equation', *Int. J. Numer. Meth. Fluids* **26**, 1039-1053.

Gurski, I., Gormier, R. & Visbal, M. (2005), 'Unsteady aerodynamics of nonlender delta wings', *Progress in Aerospace Sciences* **41**, 515-557.

Hall, C. (1985), 'Numerical solution of Navier-Stokes problems by the dual variable method', *SIAM J. Alg. Disc. Methods* **6**(2), 220-236.

Hamden, H. & Sun, M. (2000), 'Aerodynamic forces and flow structures of an airfoil in some unsteady motions at small Reynolds number', *Acta Mech.* **145**, 173-187.

Hartman, P. (1964), *Ordinary Differential Equations*, John Wiley and Sons.

Heinz, N., King, R. & Gölling, B. (2010), Robust closed-loop lift control on an industry-relevant civil aircraft half model, in R. King, ed., 'Active Flow Control II', Vol. 108 of *Notes on Numerical Fluid Mechanics and Interdisciplinary Design*, Springer, pp. 125-139.

Heinbold, H. B. (1942), 'Der unverwundene ellipsenflug als tragende fläche', *Jahrbuch 1942 der Deutsches Luftfahrtforsch* pp. 1111-1113.

Hobbs, F. (1988), *Gust Loads on Aircraft: Concepts and Applications*, AIAA Education Series, Washington, DC.

Högberg, M., Bewley, T. R. & Rossington, D. S. (2003), 'Linear feedback control and estimation of transition in plane channel flow', *J. Fluid Mech.* 481, 149-175.

Holloway, A. G. L. & Richardson, S. (2007), 'Development of a trailing vortex formed with spanwise tip jets', *Journal of Aircraft* 44(3), 845-857.

Holmes, P., Lumley, J. L. & Berkooz, G. (1996), *Turbulence, Coherent Structures, Dynamical Systems and Symmetry*, Cambridge University Press, Cambridge, UK.

Hornung, H. (1989), Vorticity generation and transport, in '10th Australian fluid mechanics conference', number KS-3, 10th Australian fluid mechanics conference, Paper KS-3.

Hsiao, F.-B., Shyu, R.-N. & Chang, R. (1994), 'High angle-of-attack airfoil performance improvement by internal acoustic excitation', *AIAA Journal* 32(3), 655-657.

Hu, F. (1996), 'On absorbing boundary conditions for linearized Euler equations by a perfectly matched layer', *J. Comput. Phys.* 129, 201-219.

Huang, L., Huang, P. G., LeBeau, R. R. P. & Hauer, T. (2004), 'Numerical study of blowing and suction control mechanism on NACA0012 airfoil', *Journal of Aircraft* 41(5), 1005-1013.

Huang, R. F., Wu, J. Y., Jeng, J. H. & Chen, R. C. (2001), 'Surface flow and vortex shedding of an impulsively started wing', *J. Fluid Mech.* 441, 265-292.

Hunt, J. C. R., Wray, A. A. & Moir, P. (1988), Eddies, stream, and convergence zones in turbulent flows, Technical Report CTR-S88, Center for Turbulence Research.

Ilak, M. & Rowley, C. W. (2008), 'Modeling of transitional channel flow using balanced proper orthogonal decomposition', *Phys. Fluids* 20(03A103).

Joon, D. & Gharib, M. (2004), 'On the relationship between the wake vortex formation process and cylinder wake vortex patterns', *Journal of Fluid Mechanics* 519, 161-181.

Jin, G. & Brusa, M. (1993), 'A nonreflecting outlet boundary condition for incompressible unsteady Navier-Stokes calculations', *J. Comput. Phys.* 107, 239-253.

Joe, W. T. & Colonius, T. (2010), Optimized waveforms for feedback control of vortex shedding, in R. King, ed., 'Active Flow Control II', Vol. 108 of *Notes on Numerical Fluid Mechanics and Interdisciplinary Design*, Springer, pp. 391-404.

Joslin, R. D. & Jones, G. S., eds (2006), *Applications of Circulation Control Technologies*, Vol. 214 of *Progress in Astronautics and Aeronautics*, AIAA.

Kelley, C. T. (1995), *Iterative methods for linear and nonlinear equations*, number 16 in 'Frontiers in Applied Mathematics', SIAM.

225

Luchtenburg, M., Tadmor, G., Lehmann, O., Noack, B. R., King, R. & Morzyński, M. (2006), Tuned POD Galerkin models for transient feedback regulation of the cylinder wake, in '44th AIAA Aerospace Sciences Meeting and Exhibit', Paper AIAA-2006-1407.

Ma, Z. & Rowley, C. W. (2008), 'Low-dimensional linearized models for systems with periodic orbits, with application to the Ginzburg-Landau equation', AIAA Paper 2008-4196, 4th Flow Control Conference.

Magill, J., Bachmann, M., Rixon, G. & McManus, K. (2003), 'Dynamic stall control using a model-based observer', *J. Aircraft* 40(2), 355-362.

Marchuk, G. (1975), *Methods of Numerical Mathematics*, Springer-Verlag.

Meyer, R. A. & Burrus, C. S. (1975), 'A unified analysis of multirate and periodically time-varying digital filters', *IEEE Trans. Circuits Syst.* 22, 162-168.

Milano, M. & Gharib, M. (2003), 'Uncovering the physics of flapping flat plates with artificial evolution', *Journal of Fluid Mechanics* 534, 403-409.

Miral, R. & Iaccarino, G. (2005), 'Immersed boundary methods', *Annual Review of Fluid Mechanics* 37, 239-261.

Mittal, S. & Tezduyar, T. E. (1995), 'Parallel finite element simulation of 3D incompressible fluid-structure interactions', *International Journal for Numerical Methods in Fluids* 21, 933-953.

Mohd-Yousif, J. (1997), Combined immersed-boundary/B-spline methods for simulations of flow in complex geometries, in 'Annual Research Briefs', Center for Turbulence Research, pp. 317-327.

Moore, B. C. (1981), 'Principal component analysis in linear systems: Controllability, observability, and model reduction', *IEEE Trans. Automat. Contr.* 26(1), 17-32.

Morzyński, M., Noack, B. R. & Tadmor, G. (2007), 'Global flow stability analysis and reduced order modeling for bluff-body flow control', *J. Theor. Appl. Mech.* 45, 621-642.
URL: www.poznań.org.pl/jtam.htm

Morzyński, M., Stankiewicz, W., Noack, B. R., King, R., Thiele, F. & Tadmor, G. (2007), Continuous mode interpolation for control-oriented models of fluid flow, in R. King, ed., 'Active Flow Control II', Vol. 95 of *Notes on Numerical Fluid Mechanics and Interdisciplinary Design*, Springer Verlag, Berlin, Germany, pp. 260-278.

Morzyński, M., Stankiewicz, W., Noack, B. R., Thiele, F. & Tadmor, G. (2006), Generalized meanfield model with continuous mode interpolation for flow control, in '3rd AIAA Flow Control Conference', paper AIAA-2006-3488, invited.

Mu, J., Kutz, A., Braszewski, D., Culp, J., Calme, A. & Giesse, A. (2008), 'Dynamic flight maneuvering using trapped vorticity flow control', AIAA Paper 2008-552.

Neuburger, D. & Wygnanski, I. (1987), The use of a vibrating ribbon to delay separation on two-dimensional airfoils, in F. Seiler, ed., 'Proc. of Air Force Academy Workshop in Unsteady Separated Flows', number TR-88-0004.

227

Kelley, C. T., Kevrekidis, I. G. & Qian, L. (2004), 'Newton-Krylov solvers for timesteppers', <http://arxiv.org/math/0404374>.

Kervinen, W., Williams, D., Pfeiffer, J., King, R. & Colonius, T. (2010), 'Closed loop control of a wing's lift for gust suppression', AIAA Paper 2010-4969.

Kim, J. & Bewley, T. R. (2007), 'A linear systems approach to flow control', *Annual Review of Fluid Mechanics* 39, 383-417.

Kim, J., Kim, D. & Choi, D. (2001), 'An immersed-boundary finite-volume method for simulations of flow in complex geometries', *J. Comput. Phys.* 171, 132-150.

Kim, J. & Moha, P. (1985), 'Application of a fractional-step method to incompressible Navier-Stokes equations', *J. Comput. Phys.* 59, 308-323.

Koumoutsakos, P. & Leonard, A. (1995), 'High-resolution simulations of the flow around an impulsively started cylinder using vortex methods', *J. Fluid Mech.* 296, 1-38.

Krysis, M. & Wang, H.-H. (2000), 'Stability of extremum seeking feedback for general nonlinear dynamic systems', *Automatica* 36, 595-601.

Lachmann, G., ed. (1961), *Boundary Layer and Flow Control-Its Principles and Application*, Vols. 1 and 2, Pergamon Press, New York.

Lai, M. & Paskin, C. (2000), 'An immersed boundary method with formal second-order accuracy and reduced numerical viscosity', *J. Comput. Phys.* 160, 705-719.

Lall, S., Marsden, J. E. & Glavatski, S. (2002), 'A subsystem approach to balanced truncation for model reduction of nonlinear control systems', *Int. J. Robust Nonlinear Control* 12, 519-535.

Lee, C. S., Tavella, D., Woods, N. J. & Roberts, L. (1989), 'Flow structure and scaling laws in lateral wing-tip blowing', *AIAA Journal* 27(8), 1002-1007.

Lee, L. & LeVeque, R. (2003), 'An immersed interface method for incompressible Navier-Stokes equations', *SIAM J. Sci. Comput.* 25(3), 832-856.

Lehmann, O., Luchtenburg, M., Noack, B., King, R., Morzyński, M. & Tadmor, G. (2005), Wake stabilization using POD Galerkin models with interpolated modes, in '44th IEEE Conference on Decision and Control and European Control Conference ECC, Seville, Spain', pp. 500-505.

Lehoucq, R. B., Sorensen, D. C. & Yang, C. (1998), *ARPACK Users' Guide*, SIAM.

Lenink, D. & Dickinson, M. H. (2009), 'Rotational accelerations stabilize leading edge vortices on revolving fly wings', *J. Exp. Bio.* 212(16), 2705-2719.

Limick, M. & Fasel, H. (2005), 'A high-order immersed interface method for simulating unsteady incompressible flows on irregular domains', *J. Comput. Phys.* 204, 157-192.

Liu, H. & Kawachi, K. (1998), 'A numerical study of insect flight', *Journal of Computational Physics* 144, 124-156.

226

Noack, B., Afanasiev, K., Morzyński, M., Tadmor, G. & Thiele, F. (2003a), 'A hierarchy of low-dimensional models for the transient and post-transient cylinder wake', *J. Fluid Mech.* 497, 335-363.

Noack, B., Afanasiev, K., Morzyński, M., Tadmor, G. & Thiele, F. (2003b), 'A hierarchy of low-dimensional models for the transient and post-transient cylinder wake', *J. Fluid Mech.* 497, 335-363.

Noack, B. R., Papas, P. & Mankewitz, P. A. (2005), 'The need for a pressure-term representation in empirical Galerkin models of incompressible shear flow', *J. Fluid Mech.* 523, 339-365.

Noack, B. R., Schlegel, M., Morzyński, M. & Tadmor, G. (2010), 'System reduction strategy for Galerkin models of fluid flows', *Int. J. Numer. Meth. Fluids* 63, 231-248.

Noack, B., Schlegel, M., Ahlborn, B., Muschke, G., Morzyński, M., Comte, P. & Tadmor, G. (2007), A finite-time thermodynamics of unsteady flows from the onset of vortex shedding to developed homogeneous turbulence, in B. Guy & D. Tondeur, eds, 'Proceedings of the Joint European Thermodynamics Conference IX', pp. 129-132.

Noack, B., Schlegel, M., Ahlborn, B., Muschke, G., Morzyński, M., Comte, P. & Tadmor, G. (2008), 'A finite-time thermodynamics of unsteady flows', *J. Non-equilibrium Thermodynamics* 33, 103-148.

Nocedal, J. & Wright, S. (1999), *Numerical Optimization*, Springer.

Ol'shanskii, M. & Staroverov, V. (2000), 'On simulation of outflow boundary conditions in finite difference calculations for incompressible fluid', *Int. J. Numer. Meth. Fluids* 33, 499-534.

Pastoor, M., Henning, L., Noack, B. R., King, R. & Tadmor, G. (2008), 'Feedback shear layer control for bluff body drag reduction', *Journal of Fluid Mechanics* 608, 161-196.

Pereira, F. & Gharib, M. (2002), 'Defocusing digital particle image velocimetry and the three-dimensional characterization of two-phase flows', *Measurement Science and Technology* 13(5), 683-694.
URL: <http://dx.doi.org/10.1088/0957-0233/13/5/305>

Pereira, F., Ster, H., Graff, E. & Gharib, M. (2006), 'Two-frame 3d particle tracking', *Measurement Science and Technology* 17(7), 1680-1692.
URL: <http://dx.doi.org/10.1088/0957-0233/17/7/006>

Perot, J. B. (1993), 'An analysis of the fractional step method', *Journal of Computational Physics* 108, 51-58.

Piskin, C. (1972), 'Flow patterns around heart valves: a numerical method', *J. Comput. Phys.* 10, 252-271.

Piskin, C. S. (2002), 'The immersed boundary method', *Acta Numerica* 11, 479-517.

Pinner, J., Auzan, J., Glauser, M. & Higuchi, H. (2007), 'Proportional closed-loop feedback control of flow separation', *AIAA J.* 45(1), 181-190.

Poelma, C., Dickson, W. B. & Dickinson, M. H. (2006), 'Time-resolved reconstruction of the full velocity field around a dynamically-scaled flapping wing', *Experiments in Fluids* 41, 213-225.

Prandtl, L. (1904), Ueber Flüssigkeitsbewegung bei sehr kleiner Reibung, in 'Proceedings 3rd International Mathematical Congress', Heidelberg, Germany, pp. 484-491.

228

Provansal, M., Mathis, C. & Boyer, L. (1987), 'Bénard-von Kármán instability: transient and forced regimes', *J. Fluid Mech.* **182**, 1–22.

Pullin, D. I. & Wang, Z. J. (2004), 'Unsteady forces on an accelerating plate and application to hovering insect flight', *Journal of Fluid Mechanics* **509**, 1–21.

Quack, V., Kerstens, W., Williams, D., Tadmor, G. & Colonius, T. (2010), Transient response of a wing to arbitrary actuator input, in 'Int. Conf. Jets, Wakes and Separated flows', Cincinnati, OH.

Raju, R., Mittal, R. & Cattafesta, L. (2008), 'Dynamics of airfoil separation control using zero-net mass-flux forcing', *AIAA Journal* **46**(12), 3103–3115.

Rennick, S. C. & Lele, S. K. (1997), 'Numerical method for incompressible vortical flows with two unbounded directions', *Journal of Computational Physics* **137**, 101–129.

Ringuette, M. J., Milano, M. & Gharib, M. (2007), 'Role of the tip vortex in the force generation of low-aspect-ratio normal flat plates', *Journal of Fluid Mechanics* **581**, 453–468.

Rodriguez, D. & Theofilis, V. (2011), 'On the birth of stall cells on airfoils', *Theor. Comput. Fluid Dyn.* (accepted for publication).

Rejzatsirikul, P., Wang, Z. & Oursul, I. (2009), 'Unsteady fluid-structure interactions of membrane airfoils at low Reynolds numbers', *Exp. Fluids* **46**(5), 859–872.

Roma, A. M., Peskin, C. S. & Berger, M. J. (1999), 'An adaptive version of the immersed boundary method', *Journal of Computational Physics* **153**, 509–534.

Roshko, A. (1954), On the development of turbulent wakes from vortex streets, Technical Report 1191, NACA.

Roshko, A. (1955), 'On the wake and drag of bluff bodies', *Journal of the aeronautical sciences* **22**(2), 124–132.

Roshko, A. (1961), 'Experiments on the flow past a circular cylinder at very high Reynolds number', *Journal of Fluid Mechanics* **10**(3), 345–356.

Rossi, L. (1997), 'Merging computational elements in vortex simulations', *SIAM J. Sci. Comput.* **18**(4), 1014–1027.

Rowley, C. W. (2005), 'Model reduction for fluids using balanced proper orthogonal decomposition', *Int. J. Bifurcation Chaos* **15**(3), 997–1013.

Rowley, C. W. & Williams, D. R. (2006), 'Dynamics and control of high-Reynolds-number flow over open cavities', *Annual Review of Fluid Mechanics* **38**, 251–276.

Rullan, J. M., Vlachos, P. P., Telioxis, D. P. & Zeiger, M. D. (2006), 'Post-stall flow control of sharp-edged wings via unsteady blowing', *Journal of Aircraft* **43**(6), 1738–1746.

Saad, Y. & Schultz, M. H. (1986), 'GMRES: A generalized minimal residual algorithm for solving nonsymmetric linear systems', *SIAM J. Sci. Statist. Comput.* **7**(3), 856–869.

Simi, R. L. & Grosz, P. M. (1994), 'Résumé and remarks on the open boundary condition minisymposium', *International Journal for Numerical Methods in Fluids* **18**, 983–1008.

Seifert, A., Darabi, A. & Wygnanski, I. (1996a), 'Delay of airfoil stall by periodic excitation', *J. Aircraft* **33**(4), 691–698.

Seifert, A., Darabi, A. & Wygnanski, I. J. (1996b), 'Delay of airfoil stall by periodic excitation', *J. Aircraft* **33**(4), 691–698.

Seifert, A., Greenblatt, D. & Wygnanski, I. J. (2004), 'Active separation control: an overview of Reynolds and mass numbers effects', *Aero. Sci. Tech.* **8**(7), 569–582.

Shewchuk, J. R. (1994), 'An introduction to the conjugate gradient method without the agonizing pain', <http://www.cs.cmu.edu/~jrs/jrpapers.html>.

Shiels, D. (1998), Simulation of controlled bluff body flow with a viscous vortex method, PhD thesis, California Institute of Technology.

Shokoochi, S., Silverman, L. M. & Van Dooren, P. (1983), 'Linear time-variable systems: Balancing and model reduction', *IEEE Trans. Automat. Contr.* **28**, 810–822.

Siegel, S., Cohen, K. & McLaughlin, T. (2003), Real-time particle image velocimetry for closed-loop flow control studies, in '41st Aerospace Sciences Meeting and Exhibit', Reno, Nevada, USA, 6–9 January 2003. AIAA-Paper 2003-0920.

Siegel, S., Cohen, K. & McLaughlin, T. (2006), 'Numerical simulations of a feedback-controlled circular cylinder wake', *AIAA Journal* **44**(6), 1266–1276.

Sirovich, L. (1987), 'Turbulence and the dynamics of coherent structures, parts I–III', *Q. Appl. Math.* **XLV**(3), 561–590.

Skogestad, S. & Postlethwaite, I. (2005), *Multivariable Feedback Control: Analysis and Design*, 2nd edn, Wiley.

Stankiewicz, W., Morzyński, M., Noack, B. R. & Tadmor, G. (2008), 'Reduced order Galerkin model of flow around a NACA-0012 airfoil', *Math. Modeling and Anal.* **13**, 113–122.

Stankiewicz, W., Morzyński, M., Roszak, R., Noack, B. R. & Tadmor, G. (2008), 'Reduced order modelling of flow around an airfoil with a changing angle of attack', *Archives of Mechanics* **60**, 509–526.

Srikwerda, J. & Lee, Y. (1999), 'The accuracy of the fractional step method', *SIAM J. Numer. Anal.* **37**(1), 37–47.

Tadmor, G. (2004), 'Observers and feedback control for a rotating vortex pair', *IEEE Transactions on Control Systems Technology* **12**(1), 36–51.

Tadmor, G., Bissix, D., Noack, B. R., Morzyński, M., Colonius, T. & Taira, K. (2008), Temporal-harmonic specific POD mode extraction, in '4th Flow Control Conference / 38th AIAA Fluid Dynamics Conference and Exhibit', AIAA Paper 2008-4190.

Tadmor, G., Contuori, M. D., Noack, B. R., Luchimbura, M., Lehmann, O. & Morzyński, M. (2007), Low Order Galerkin Models for the Actuated Flow Around 2-D Airfoils, in '45th AIAA Aerospace Sciences Meeting and Exhibit', AIAA Paper 2007-1313.

Tadmor, G., Gonzalez, J., Lehmann, O., Noack, B. R., Morzyński, M. & W. Stankiewicz (2007), Shift Modes and Transient Dynamics in Low Order, Design Oriented Galerkin Models, in '45th AIAA Aerospace Sciences Meeting and Exhibit', Paper AIAA 2007-1111.

Tadmor, G., Lehmann, O., Noack, B. R. & Morzyński, M. (2010), 'Mean field representation of the natural and actuated cylinder wake', *Phys. Fluids* **22**(034102).

Tadmor, G., Noack, B., Morzyński, M. & Siegel, S. (2004), Low-dimensional models for feedback flow control. Part II: Controller design and dynamic estimation, in '2nd AIAA Flow Control Conference', AIAA Paper 2004-2409.

Tadmor, G., Noack, B. R., Schlegel, M., Morzyński, M. & Kelley, S. (2009), Mean field models and finite-time thermodynamics: A unified framework for low-order, multi-scale Galerkin models of fluid flows, in '10th Joint European Thermodynamics Conference (JETC X)'.

Taira, K. (2008), The immersed boundary projection method and its application to simulation and control of flows around low-aspect-ratio wings, PhD thesis, California Institute of Technology.

Taira, K. & Colonius, T. (2007), 'The immersed boundary method: a projection approach', *Journal of Computational Physics* **225**, 2118–2137.

Taira, K. & Colonius, T. (2008, accepted), 'Three-dimensional separated flows around low-aspect-ratio flat plates', *Journal of Fluid Mechanics*.

Taira, K. & Colonius, T. (2009a), 'Effect of tip vortices in low-reynolds-number poststall flow control', *AIAA J.* **47**(3).

Taira, K. & Colonius, T. (2009b), 'Three-dimensional separated flows around low-aspect-ratio flat plates', *J. Fluid Mech.* **623**, 187–207.

Taira, K., Dickinson, W. B., Colonius, T., Dickinson, M. H. & Rowley, C. W. (2007), 'Unsteadiness in flow over a flat plate at angle-of-attack at low Reynolds numbers', AIAA Paper 2007-710.

Taira, K., Rowley, C. W. & Colonius, T. (2010), 'Feedback control of a high-lift state for a low-aspect-ratio wing', AIAA Paper 2010-357.

Temam, R. (1969), 'Sur l'approximation de la solution des équations de Navier-Stokes par la méthode des pas fractionnaires (I)', *Arch. Rat. Mech. Anal.* **32**(2), 135–153.

Torres, G. E. & Mueller, T. J. (2004a), 'Low-aspect-ratio wing aerodynamics at low Reynolds numbers', *AIAA Journal* **42**(5), 865–873.

Torres, G. E. & Mueller, T. J. (2004b), 'Low-aspect-ratio wing aerodynamics at low Reynolds numbers', *AIAA Journal* **42**(5), 865–873.

Trefethen, L. N. & Bau, D. I. (1997), *Numerical Linear Algebra*, SIAM.

Tritton, D. (1959), 'Experiments on the flow past a circular cylinder at low Reynolds number', *J. Fluid Mech.* **6**, 547–567.

Tuckerman, L. & Barkley, D. (2000), Bifurcation analysis for timesteppers, in 'Numerical Methods for Bifurcation Problems and Large-Scale Dynamical Systems', Vol. 119 of *IMA Vol. Appl. Math. and Appl.*, Springer-Verlag, pp. 453–466.

Varga, A. (2000), Balanced truncation model reduction of periodic systems, in 'Proceedings of the 39th IEEE Conference on Decision and Control', pp. 2379–2384.

Verriest, E. I. & Kailath, T. (1983), 'On generalized balanced realizations', *IEEE Trans. Automat. Contr.* **28**, 833–844.

Wang, H.-H. & Krstić, M. (2000), 'Extremum seeking for limit cycle minimization', *IEEE Trans. Automat. Contr.* **45**(12), 2432–2437.

Wang, Z. (1999), 'Efficient implementation of the exact numerical far field boundary condition for poisson equation on an infinite domain', *J. Comput. Phys.* **153**, 666–670.

Wei, M. & Freund, J. B. (2006), 'A noise-controlled free shear flow', *J. Fluid Mech.* **546**, 123–152.

Willert, C. E. & Gharib, M. (1991), 'Digital particle image velocimetry', *Experiments in Fluids* **10**(4), 181–193.

Willert, C. E. & Gharib, M. (1992), 'Three-dimensional particle imaging with a single camera', *Experiments in Fluids* **12**(6), 353–358.

Willert, C. E. & Gharib, M. (1992), 'Three-dimensional particle imaging with a single camera', *Experiments in Fluids* **12**(6), 353–358.

Williams, D., Bustain, S., Quack, V. & Kerstens, W. (2009), 'Flow structures behind a 3D wing in an oscillating freestream', AIAA Paper 2009-3690.

Williams, D., Collins, J., Tadmor, G. & Colonius, T. (2008), 'Control of a semi-circular planform wing in a gusting unsteady freestream flow: 1 – experimental issues', AIAA Paper 2008-3976.

Williams, D., Collins, J., Tadmor, G., Colonius, T. & Rowley, C. (2008), Closed-loop control of post stall lift on a low aspect ratio wing, in 'Int. Conf. Jets, Wakes and Separated flows', Berlin, Germany.

Williams, D., Doshi, S., Collins, J. & Colonius, T. (2007), 'Control of the spanwise distribution of circulation on naca 0012 and flat plate wings', AIAA Paper 2007-1121.

Williams, D., Kerstens, W., Bustain, S., Quack, V., Pfeiffer, J., King, R., Tadmor, G. & Colonius, T. (2010), 'Closed loop control of a wing in an unsteady flow', AIAA Paper 2010-0358.

Williams, D., Kerstens, W., Pfeiffer, J., King, R. & Colonius, T. (2010), Unsteady lift suppression with a robust closed loop controller, in R. King, ed., 'Active Flow Control II', Vol. 106 of *Notes on Numerical Fluid Mechanics and Interdisciplinary Design*, Springer, pp. 19–30.

Williams, D., Quack, V., Kerstens, W., Bustain, S., Tadmor, G., Rowley, C. & Colonius, T. (2009), 'Low Reynolds number wing response to an oscillating freestream with and without feed forward control', AIAA Paper 2009-0143.

- Williams, D. R., Collins, J., Jankhos, C., Colonius, T. & Tadmor, G. (2008), 'Control of flow structure on a semi-circular planform wing', *AIAA Paper* 2008-597.
- Williams, D., Tadmor, G., Colonius, T., Kerstus, W., Quach, V. & Buntain, S. (2009), 'The lift response of a stalled wing to pulsatile disturbances', *AIAA Journal* 47(12), 3031-3037.
- Williams-Shuber, K. & Gharib, M. (1990), 'Transition from order to chaos in the wake of an airfoil', *J. Fluid Mech.* 213, 29-57.
- Winkelmann, A. E. & Barlow, J. B. (1980), 'Flowfield model for a rectangular planform wing beyond stall', *AIAA Journal* 18(8), 1006-1007.
- Woo, G., Crittenden, T. & Gisser, A. (2008), 'Transitory control of a pitching airfoil using pulse combustion actuation', *AIAA Paper* 2008-4324.
- Wu, J. Z., Liu, X. Y., Denny, A. G., Fan, M. & Wu, J. M. (1998), 'Post-stall flow control on an airfoil by local unsteady forcing', *J. Fluid Mech.* 371, 21-58.
- Yanenko, N. (1971), *The Method of Fractional Steps: The Solution of Problems of Mathematical Physics in Several Variables*, Springer-Verlag.
- Yatskevich, S., Sullivan, P. E. & Kawall, J. G. (2009), 'On vortex shedding from an airfoil in low-Reynolds-number flows', *J. Fluid Mech.* 632, 245-271.
- Yoon, S. A. & Katz, J. (1998), 'Study of the unsteady flow features on a stalled wing', *AIAA Journal* 36(3), 305-312.
- Yu, H., Loerch, M., Tadmor, G. & Siegel, S. (2006), 'Real-time particle image velocimetry for feedback loops using fpga implementations', *Journal of Aerospace Computing, Information and Communication* 3, 52-62.
- Zhou, K., Salomon, G. & Wu, E. (1999), 'Balanced realization and model reduction for unstable systems', *Int. J. Robust and Nonlin. Contr.* 9(3), 183-198.

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

The public reporting burden for this collection of information is estimated to average 1 hour per response, including the time for reviewing instructions, searching existing data sources, gathering and maintaining the data needed, and completing and reviewing the collection of information. Send comments regarding this burden estimate or any other aspect of this collection of information, including suggestions for reducing the burden, to Department of Defense, Washington Headquarters Services, Directorate for Information Operations and Reports (0704-0188), 1215 Jefferson Davis Highway, Suite 1204, Arlington, VA 22202-4302. Respondents should be aware that notwithstanding any other provision of law, no person shall be subject to any penalty for failing to comply with a collection of information if it does not display a currently valid OMB control number.

PLEASE DO NOT RETURN YOUR FORM TO THE ABOVE ADDRESS.

1. REPORT DATE (DD-MM-YYYY)		2. REPORT TYPE Final Report		3. DATES COVERED (From - To) 5/2005-4/2008	
4. TITLE AND SUBTITLE CLOSED -LOOP CONTROL OF VORTEX FORMATION IN SEPARATED FLOWS WITH APPLICATIONS TO MICRO AIR VEHICLES				5a. CONTRACT NUMBER	
				5b. GRANT NUMBER FA9550-05-1-0369	
				5c. PROGRAM ELEMENT NUMBER	
6. AUTHOR(S) CLARENCE W. ROWLEY				5d. PROJECT NUMBER	
				5e. TASK NUMBER	
				5f. WORK UNIT NUMBER	
7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES) CALIFORNIA INSTITUTE OF TECHNOLOGY PASADENA, CA				8. PERFORMING ORGANIZATION REPORT NUMBER	
9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES) AFOSR 875 N. RANDOLPH ST ARLINGTON, VA 22203				10. SPONSOR/MONITOR'S ACRONYM(S)	
				11. SPONSOR/MONITOR'S REPORT NUMBER(S) AFRL-OSR-VA-TR-2012-0000	
12. DISTRIBUTION/AVAILABILITY STATEMENT A= Approved for public release; distribution is unlimited					
13. SUPPLEMENTARY NOTES 20120918174					
14. ABSTRACT This research is concerned with developing theory, algorithms, and applications of model-based, closed loop flow control in order to enable robust and agile flight of micro air vehicles. the efforts are centered around the application to integrated closed-loop flow and fight control for stabilization and regulations of separated flows occuring on unmanned and micro air vehicles (UAV/MAV).					
15. SUBJECT TERMS					
16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT	b. ABSTRACT	c. THIS PAGE			19b. TELEPHONE NUMBER (Include area code)

INSTRUCTIONS FOR COMPLETING SF 298

1. REPORT DATE. Full publication date, including day, month, if available. Must cite at least the year and be Year 2000 compliant, e.g. 30-06-1998; xx-06-1998; xx-xx-1998.

2. REPORT TYPE. State the type of report, such as final, technical, interim, memorandum, master's thesis, progress, quarterly, research, special, group study, etc.

3. DATE COVERED. Indicate the time during which the work was performed and the report was written, e.g., Jun 1997 - Jun 1998; 1-10 Jun 1996; May - Nov 1998; Nov 1998.

4. TITLE. Enter title and subtitle with volume number and part number, if applicable. On classified documents, enter the title classification in parentheses.

5a. CONTRACT NUMBER. Enter all contract numbers as they appear in the report, e.g. F33315-86-C-5169.

5b. GRANT NUMBER. Enter all grant numbers as they appear in the report. e.g. AFOSR-82-1234.

5c. PROGRAM ELEMENT NUMBER. Enter all program element numbers as they appear in the report, e.g. 61101A.

5e. TASK NUMBER. Enter all task numbers as they appear in the report, e.g. 05; RF0330201; T4112.

5f. WORK UNIT NUMBER. Enter all work unit numbers as they appear in the report, e.g. 001; AFAPL30480105.

6. AUTHOR(S). Enter name(s) of person(s) responsible for writing the report, performing the research, or credited with the content of the report. The form of entry is the last name, first name, middle initial, and additional qualifiers separated by commas, e.g. Smith, Richard, J, Jr.

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES). Self-explanatory.

8. PERFORMING ORGANIZATION REPORT NUMBER. Enter all unique alphanumeric report numbers assigned by the performing organization, e.g. BRL-1234; AFWL-TR-85-4017-Vol-21-PT-2.

9. SPONSORING/MONITORING AGENCY NAME(S) AND ADDRESS(ES). Enter the name and address of the organization(s) financially responsible for and monitoring the work.

10. SPONSOR/MONITOR'S ACRONYM(S). Enter, if available, e.g. BRL, ARDEC, NADC.

11. SPONSOR/MONITOR'S REPORT NUMBER(S). Enter report number as assigned by the sponsoring/monitoring agency, if available, e.g. BRL-TR-829; -215.

12. DISTRIBUTION/AVAILABILITY STATEMENT. Use agency-mandated availability statements to indicate the public availability or distribution limitations of the report. If additional limitations/ restrictions or special markings are indicated, follow agency authorization procedures, e.g. RD/FRD, PROPIN, ITAR, etc. Include copyright information.

13. SUPPLEMENTARY NOTES. Enter information not included elsewhere such as: prepared in cooperation with; translation of; report supersedes; old edition number, etc.

14. ABSTRACT. A brief (approximately 200 words) factual summary of the most significant information.

15. SUBJECT TERMS. Key words or phrases identifying major concepts in the report.

16. SECURITY CLASSIFICATION. Enter security classification in accordance with security classification regulations, e.g. U, C, S, etc. If this form contains classified information, stamp classification level on the top and bottom of this page.

17. LIMITATION OF ABSTRACT. This block must be completed to assign a distribution limitation to the abstract. Enter UU (Unclassified Unlimited) or SAR (Same as Report). An entry in this block is necessary if the abstract is to be limited.