

A Linear Systems Approach to Flow Control

John Kim¹ and Thomas R. Bewley²

¹Department of Mechanical and Aerospace Engineering, University of California, Los Angeles, Los Angeles, California 90095-1597; email: jkim@seas.ucla.edu

²Department of Mechanical and Aerospace Engineering, University of California, San Diego, La Jolla, California 92093-0411; email: bewley@ucsd.edu

Annu. Rev. Fluid Mech. 2007. 39:383–417

The *Annual Review of Fluid Mechanics* is online at fluid.annualreviews.org

This article's doi:
10.1146/annurev.fluid.39.050905.110153

Copyright © 2007 by Annual Reviews.
All rights reserved

0066-4189/07/0115-0383\$20.00

Key Words

transition delay, turbulence mitigation, optimal (\mathcal{H}_2 , LQG) and noncooperative/robust (\mathcal{H}_∞ , LTR) control, model reduction, Chandrasekhar's method

Abstract

The objective of this paper is to introduce the essential ingredients of linear systems and control theory to the fluid mechanics community, to discuss the relevance of this theory to important open problems in the optimization, control, and forecasting of practical flow systems of engineering interest, and to outline some of the key ideas that have been put forward to make this connection tractable. Although many significant advances have already been made, many new challenges lie ahead before the full potential of this synthesis of disciplines can be realized.

1. INTRODUCTION

The ability to alter flows to achieve a desired effect is a matter of tremendous consequence in many applications. For example, worldwide ocean shipping consumes about 2.1 billion barrels of oil per year (Corbett & Koehler 2003), whereas the airline industry consumes about 1.5 billion barrels of jet fuel per year (P.R. Spalart, private communication). Reducing average overall drag by just a few percent could save several billion dollars annually in either application, and help preserve the earth's limited natural resources. Reducing drag also enables increased speed, range, and endurance. Other effects commonly desired in fluid mechanical systems include reducing structural vibrations, radiated noise, and surface heat transfer, all typically associated with reducing flow-field unsteadiness, as well as increasing mixing and combustion efficiency and reducing pattern factor (i.e., hot spots in combustion products), problems typically associated with increasing flow-field unsteadiness in an appropriate fashion. All such problems fall under the purview of this line of study. Not surprisingly, there has been enormous interest in altering flows to achieve such effects for well over a century.

Today, flow control is a phrase used liberally with a range of intended meanings. In its broadest sense, the phrase refers to any mechanism that manipulates a fluid flow into a state with desired flow properties. In its narrowest sense, the phrase is sometimes restricted to mean the application of systems and control theory to the Navier-Stokes equations. Many definitions in between are also possible, including those that cover intuition-based approaches based primarily on the control designers' physical insight into the relevant flow physics together with some simple trial and error. Although such approaches have been successful and will continue to play a significant role, the incorporation of model-based control theory into many open problems in fluid mechanics presents a host of new opportunities.

A wide variety of different types of flow control strategies—active, passive, open-loop, closed-loop, etc.—have been developed and implemented over the years, and some are quite successful in achieving certain control objectives. Several recent and comprehensive surveys of this rapidly growing field are available, including Gad-el Hak (2000), Bewley (2001), Gunzburger (2002), Kim (2003), and Collis et al. (2004). Interested readers are referred to these articles, and the references therein, for a comprehensive overview of recent advances, as well as various attempts at pinning down the somewhat ambiguous categorizations used in this field (active vs passive, etc.). This article does not attempt to repeat these accomplishments; specifically, it is not intended to review the now very extensive literature on this subject, which would be impossible in an article of this length.

Our objective, rather, is to present and discuss the essential ingredients of linear model-based systems and control theory as it relates to both transitional and turbulent fluid mechanical systems. The article is intended for those with a background in fluid mechanics, but not necessarily in control or optimization theory. We aim to draw more mathematically inclined researchers with a fluids background into this line of research, as there are an abundance of challenging and fundamental problems ripe and wanting to be solved at this intersection of disciplines. As such, a portion of this paper is expository in nature.

The focus of this paper is primarily on the feedback problem; that is, coordinating actuator inputs with sensor outputs to achieve a desired effect. Thus, much of the paper considers near-wall flows, as this configuration facilitates both surface-mounted sensors and actuators to be placed near the flow instabilities of interest. Perhaps the most basic configuration in this vein is channel flow with skin friction and pressure sensors continuously distributed over the walls to provide the system measurements, and zero-net blowing/suction continuously distributed over the walls to provide the actuation. Although simple to study, this configuration is artificial in several regards, and various extensions are needed to connect it with reality—notably, accounting for spatially developing boundary layers, the discrete locations of sensors and actuators and their precise sensitivities and effects, and various types of system uncertainty. Many of these extensions are well underway.

Early investigations in the present vein include, among others, Abergel & Temam (1990), Burns & Kang (1991), Gunzburger et al. (1992), Joshi et al. (1997, 1999), and Bewley & Liu (1998). The number of researchers working in related areas has grown rapidly since these early efforts; there are now roughly a half dozen significant workshops and mini-symposia organized yearly to discuss recent advances.

We present a brief motivation of the linear systems approach and its application to fluid systems in Section 2, and the essential foundations of model-based control and estimation theory (Section 3). We discuss the issue of managing high-dimensional discretizations in Section 4, applications and extensions of the framework in Section 5, and conclusions in Section 6.

2. MOTIVATION AND SOME KEY ISSUES

Linear systems theory provides a uniquely effective tool for optimization and control. We now discuss some key issues regarding its application to fluid systems.

2.1. Linearization

Any smooth problem is easily linearized. In partial differential equation (PDE) systems, linearization of the governing equations can be performed either algebraically (by hand derivation and coding) or automatically via one of three approaches:

- By applying an automatic differentiation tool¹ such as ADIFOR, TAMC, etc. to a nonlinear simulation code implementing the governing equations. Note that some such tools can also generate the adjoint equation at the heart of a corresponding optimization problem (see Section 3).
- By performing a perturbed nonlinear simulation, subtracting the result from an unperturbed nonlinear simulation, and dividing by the magnitude of the perturbation (a finite difference approach).

¹See <http://www.autodiff.org/Tools/> for a complete listing of such tools.

- By changing `real` types to `complex` types throughout an entire nonlinear simulation code, then perturbing the optimization variables a small amount in the imaginary direction; the resulting linearized perturbation to the system is then evident in the imaginary part of the complex result (this is the complex step derivative approach; see Cerviño & Bewley 2003 for extension to pseudospectral codes).

The choice of the state about which to linearize forms the central distinction between the iterative (adjoint-based) and direct (Riccati-based) approaches to be outlined in Section 3. The iterative approach performs a linearization about an actual trajectory of the system, determines an appropriate direction to “nudge” the optimization variables based on this linearized analysis, updates the optimization variables to maximum beneficial effect in this direction, then repeats, at each step performing a new linearization of the governing equations about the trajectory of the full nonlinear system with the current value of the optimization variables. Via a series of linear analyses of this sort, this method optimizes the full nonlinear problem, although the (local) optimal point so found is not guaranteed to be the globally optimal solution.

In contrast, the direct approach performs a single linearization about a representative mean flow state, which itself is not necessarily even a solution of the governing equations. By simplifying the problem structure in this way, this solution approach jumps straight to the unique optimal point of the (simplified) optimization problem under consideration. Thus, at their heart, both the iterative (adjoint-based) and direct (Riccati-based) approaches, which are closely related, are based effectively on linearization and optimization.

The relevance of the iterative approach to transitional and turbulent flows is clear. Although only providing a local optimum, the nonlinear optimizations so performed are usually effective. Indeed, the benchmark problem of relaminarizing fully developed channel-flow turbulence (in numerical simulations) via a distribution of blowing/suction on the wall as the control was solved for the first time this way (Bewley et al. 2001). Unfortunately, this technique is computationally intensive, requiring iterative direct numerical simulations (DNSs) to complete significantly faster than the system itself evolves in time. Thus, application of this technique to most practical turbulent flow systems is not anticipated.² The one notable exception to this statement is in the field of weather forecasting, in which adjoint-based iterations following the model predictive estimation approach (Section 3.3) are applied routinely. The perspective attained by the present line of investigation has in fact led to a fundamental reformulation of this forecasting framework (see Section 5.4).

The relevance of the direct approach to the transition problem (at least, in its early stages) is also clear, as the nonlinear instability of the flow leading to transition is preceded (usually, upstream) by linear amplification of disturbances in the system, which may be mitigated by linear control strategies. The relevance of the direct

²Note that tractable approaches based on the adjoint idea in a simplified setting (i.e., not marched all the way to an optimum over a finite time horizon) have also been explored; this approach is sometimes referred to as a suboptimal control strategy (see, e.g., Lee et al. 1998).

approach to the turbulence problem, however, is still the subject of some debate, and is thus motivated further in greater depth below.

2.2. Linear Models of Turbulence for Control-Oriented Analyses

The applicability of linear control strategies to turbulence is predicated upon the hypothesis that appropriately linearized models (e.g., Orr-Sommerfeld/Squire) faithfully represent the inputs, outputs, and at least some of the important dynamic processes of turbulent flow systems. The fluid dynamics literature of the last decade is replete with articles aimed at supporting this hypothesis. For example, Farrell & Ioannou (1996) used these linearized equations in an attempt to explain the mechanism for the turbulence attenuation that is caused by the closed-loop intuition-based control strategy now commonly known as opposition control. Jovanović & Bamieh (2001) proposed a stochastic disturbance model, which, when used to force the linearized open-loop Navier-Stokes equation, led to a simulated flow state with certain second-order statistics (specifically, u_{rms} , v_{rms} , w_{rms} , and the Reynolds stress $-\overline{uv}$) that mimicked, with varying degrees of precision, the statistics from a full DNS of a turbulent flow at $Re_\tau = 180$.

Clearly, the hypothesis concerning the relevance of linearized models to the turbulence problem can only be taken so far, as linear models of fluid systems do not capture the nonlinear “scattering” or “cascade” of energy over a range of length scales and time scales, and thus linear models fail to capture an essential dynamic effect that endows turbulence with its inherent “multiscale” characteristics.

A key philosophy that underlies the field of systems and control theory, but is somewhat underappreciated in the field of fluid mechanics, is this: *A system model that is good enough to use for control design is not necessarily good enough for accurate numerical simulation.* The main thing that a model needs to capture for it to be useful in control design is the relation between the inputs and outputs of the system and their general influence on cost function measuring the system under consideration. Although perhaps counter to the traditional fluids mind-set, a useful control-oriented model need not capture the well-known statistics (streak spacing, etc.) and distinctive bifurcation points (transition Reynolds numbers, etc.) of the uncontrolled system, and to a large extent it is irrelevant whether or not it does. Thus, for the purpose of computing feedback for the control and estimation problems, linear models might well be good enough.

Stated another way, in the control problem, the model upon which the feedback is computed needs only to include the terms responsible for the production of energy in the system and how this production might be mitigated by control input. For wall-bounded flows, there is compelling evidence that this is true, at least at sufficiently low Reynolds number. Kim & Lim (2000) showed that interior body forcing (applied everywhere inside the flow domain) that was constructed to cancel the effect of the off-diagonal block of the linear Orr-Sommerfeld/Squire equations was sufficient to relaminarize the turbulent flow. Högberg et al. (2003b) showed that boundary forcing (blowing/suction distributed on the channel walls) determined using full-information linear control theory, scheduling the feedback gains based on the instantaneous shape

of the mean velocity profile, was also sufficient to relaminarize the turbulent flow. In a similar manner, the system model upon which estimator feedback is computed might need only to capture the terms responsible for the production of energy in the system describing the estimation error.

2.3. The Problem of Nearly Unobservable/Uncontrollable Modes

The problem of estimating and controlling the state of a chaotic nonlinear system is inherently difficult. When posed as an optimization problem, one can expect that, in general, multiple local minima of such non-convex optimization problems will exist, many of which will be associated with control distributions and state estimates that are quite poor. These difficulties are exacerbated by the fact that turbulence is a multiscale phenomenon (i.e., it is characterized by energetic motions over a broad range of length scales and timescales that interact in a nonlinear fashion), with significant nonlinear chaotic dynamics evolving relatively far from where the sensors and actuators are located (on the walls).

The issue of nearly unobservable/uncontrollable modes of the Orr-Sommerfeld/Squire operator is evident by examining it for streamwise-varying modes, as illustrated in **Figure 1**. Note that, even in the laminar case, a significant number of the leading eigenmodes of the system at this wave-number pair are “center modes” with very little support near the walls, and thus are nearly unobservable with wall-mounted sensors and nearly uncontrollable with wall-mounted actuators. This makes both estimation and control of these modes with noisy wall-mounted sensors and actuators nearly impossible. As seen in the turbulent case at the same bulk Reynolds number in **Figure 1b** (and at higher bulk Reynolds numbers, not shown), an even higher percentage of the leading eigenmodes of the linearized system are nearly unobservable and uncontrollable in the turbulent case, with the problem gradually getting worse as the Reynolds number is increased and the mean velocity profile flattens. Thus, we see that the problem of estimating and controlling turbulent flows is fundamentally harder than the corresponding problems in laminar flows even if the linearized model of turbulence is considered valid, simply due to the heightened presence of nearly unobservable and uncontrollable modes.

For the problem of turbulence control, we might set our sights fairly low. That is, we might design our cost functions to focus primarily on getting an accurate state estimate only fairly close to the walls, near where the sensors are located, then subduing the flow-field fluctuations only fairly close to the walls, near where the actuators are located. This approach is supported by the observation that most turbulence production in turbulent boundary layers takes place in the wall region, more specifically within the buffer layer ($y^+ < 50$). Furthermore, it has now been recognized that near-wall streamwise vortices are responsible for high skin-friction drag in turbulent boundary layers (Choi et al. 1994, Kravchenko et al. 1993). These vortices are primarily found in the buffer layer ($y^+ = 10 - 50$) with their typical diameter in the order of $d^+ = 20 - 50$ (Kim et al. 1987). Streamwise vortices are formed and maintained autonomously (independent of the outer layer) by a self-sustaining process, which involves the wall-layer streaks and instabilities associated with them




Figure 1

First 25 eigenvectors of Orr-Sommerfeld/Squire at $\{k_x, k_z\} = \{1, 0\}$ linearized about (a) the laminar flow profile at $Re_B = 1429$ ($Re_c = 2143.5$) and (b) the mean turbulent flow profile at $Re_B = 1429$ ($Re_\tau = 100$). Shown are the real (solid) and imaginary (dashed) parts of the ω component (blue) and v component (red) of the least stable eigenvectors, plotted as a function of y from the lower wall (bottom) to the upper wall (top).

(Hamilton et al. 1995, Schoppa & Hussain 2002). There are some differences in details on the self-sustaining process, but it is generally accepted that this process is essentially independent of the outer part of the boundary layer. In other words, near-wall turbulence dynamics is self-sustaining. This self-sustaining process of near-wall turbulence is clearly illustrated in the clever numerical experiment of Jimenez & Pinelli (1999), in which modified Navier-Stokes equations were solved in order to represent a turbulent channel flow without large-scale motions in the outer part. Significantly, no discernible differences in the behavior of the inner part (i.e., near-wall region) were observed, thus demonstrating that the inner part of the boundary layer can be maintained autonomously by a self-sustaining process. Based on these observations, it has been argued that the outer part of turbulent boundary layers is driven by the inner layer (a.k.a. a bottom-up process). By this argument, it is unnecessary to estimate and control the motions of the flow far from the wall in order to realize our objective. Flow-field fluctuations far from the wall will indeed (through nonlinear interactions) act as disturbances to continuously excite both the state and the state estimation error near the wall, whereas feedback from the sensors and actuators will be used continuously to subdue these fluctuations in the near-wall region.

A counter argument has also been made. It has been argued that at high Reynolds numbers ($Re_\tau > 10,000$) (J.C.R. Hunt, private communication), the inner layer ceases

to play the dominant role, and the overall boundary layer is driven by the outer part of boundary layers (a.k.a. a top-down process). The phenomenon central to this argument is known as “shear sheltering,” which implies a limited extent to which the outer layer can affect the inner layer when $Re_\tau > 10,000$ (Hunt & Durbin 1999). If this is the case, it may be better to target the control on the large scales to affect the large outer-layer structures directly, rather than targeting the substantially smaller near-wall structures. Unfortunately, given the fact that several of the center modes in **Figure 1** are, effectively, unobservable/uncontrollable with wall measurements/actuation, this might prove difficult with Riccati-based linear control strategies.

From the present perspective, we don’t need to answer the question as to which part of the boundary layer plays the dominant role. The model-based control algorithm will figure out an effective strategy itself. This perspective may be refined with any hypotheses we might have about the maintenance of near-wall turbulence (bottom-up, top-down, or otherwise) by tuning the cost function to target the phenomenon of interest. That is, there is a subtle interplay of discovering new things about flow physics and incorporating existing knowledge/hypotheses/hunches of flow physics when setting up and solving model-based flow control problems.

2.4. The Problem of Non-Normality

The non-normality of the Orr-Sommerfeld/Squire operator is evident when examining it for spanwise-varying modes, as illustrated in **Figure 2**. Note that the v components of these eigenvectors (dashed lines) have been magnified substantially to make them visible in this plot. In particular, note that, after the first, these eigenvectors come in pairs of almost exactly the same shape. Thus, a flow perturbation initialized as, for example, the second eigenmode minus the third eigenmode in **Figure 2a** is characterized by a very low initial energy due to destructive interference; however, as one eigenmode decays more quickly in time than the other, this destructive interference is reduced with time, and thus the overall energy of the perturbation actually increases quite substantially before it eventually decreases due to the stability of both modes (Butler & Farrell 1993, Reddy & Henningson 1993). This effect is referred to as *transient energy growth* in the fluids literature and *peaking* in the controls literature. Transient energy growth is a direct result of eigenvector nonorthogonality, and is accompanied by very large input/output transfer function norms in such systems when the system is considered from the input/output perspective (see Bewley 2001 and Lim & Kim 2004).

The degree of non-normality of the eigenvectors is modified only slightly when moving from the laminar case (**Figure 2a**) to the turbulent case (**Figure 2b**). However, all modes at this wave-number pair have a substantial footprint on the wall in both the laminar case and the turbulent case (cf. **Figure 1**). Thus, the situation is not quite as bad as it might first appear: Even when linearized about the turbulent flow profile, at the wave numbers of primary concern (that is, those in which the non-normality of the eigenmodes of the system matrix is most pronounced), these eigenmodes are easily detected by wall-mounted sensors and affected by




Figure 2

As in **Figure 1**, but at $\{k_x, k_z\} = \{0, 2\}$. Shown are the real part of the ω component (solid blue) and 200 times the imaginary part of the v component (dashed red) of the least stable eigenvectors; the other parts are negligible.

wall-mounted actuators. Furthermore, those pairs of eigenmodes with nearly the same shape are easily distinguished from one another during the dynamic state estimation process by the fact that they are associated with different eigenvalues characterizing their variation in time.

3. MODEL-BASED CONTROL THEORY: THE ESSENTIALS

We now give a summary of the foundation for model-based control and estimation of fluid mechanical systems, via both iterative adjoint-based optimization and direct Riccati-based feedback. This section is divided into four parts: Sections 3.1 and 3.2 consider the control problem (i.e., the determination of appropriate inputs to a system to achieve a desired objective assuming accurate knowledge of the system state), Sections 3.3 and 3.4 consider the estimation problem (i.e., the approximate determination of the system state based on recent, limited, noisy measurements of an actual physical system). Denoting the current time as $t = 0$, the control problem looks at the future evolution of the system over a horizon of interest $[0, T]$, whereas the estimation problem looks at the past history of measurements from the system over a horizon of interest $[-T, 0]$. Together, solutions of the control and estimation problems facilitate the coordination of a limited number of actuators with a limited number of sensors in order to achieve a desired effect (Section 3.5).

The iterative approach to these two problems (see Sections 3.1 and 3.3) may be applied to both nonlinear systems and nonquadratic cost functions. Significantly, it only requires the computation of vectors (i.e., state and adjoint fields) of dimension N (the dimension of the state field itself) evolving over the (finite) time horizon of interest, and thus extends readily to high-dimensional discretizations of unsteady PDE systems, even when $N \gtrsim 10^6$ is necessary to resolve the system under consideration. Essentially, for any smooth, differentiable system one can afford to simulate computationally, one can also afford to simulate the adjoint field necessary to determine the gradient of the cost function of interest in the space of the optimization variables, thereby enabling gradient-based optimization.

The direct approach to these problems (Sections 3.2 and 3.4) is based on more strict assumptions (specifically, a linear governing equation and a quadratic cost function). Subject to these assumptions, this approach jumps straight to the unique minimum of the cost function by setting the gradient equal to zero and solving the two-point boundary value problem for the state and adjoint fields that results. This approach requires the computation of a matrix (relating the state and adjoint fields in the optimal solution) of dimension N^2 , and thus does not extend readily to high-dimension discretizations of infinite-dimensional PDE systems, as it is prohibitively expensive for $N \gtrsim 10^3$. As discussed in Section 4, many of the major advances in the field of the feedback control of fluid systems have been related to the development of appropriate techniques to finesse oneself out of this dimensionality predicament, such as transform techniques, parabolization, and open-loop model reduction.

The adjoint-based control optimization approach (Section 3.1) is known as model predictive control. The Riccati-based feedback control approach (Section 3.2) is known as \mathcal{H}_2 state feedback control, or optimal control. The adjoint-based state estimation approach (Section 3.3) is known (in the field of weather forecasting) as 4Dvar. The Riccati-based state estimation approach (Section 3.4) is known as a Luenberger observer, or, when interpreted from a stochastic point of view, as \mathcal{H}_2 state estimation or a Kalman filter. Although their application to the Navier-Stokes equations is still nascent, application of all four formulations is widespread in most areas of science and engineering today.

3.1. Control via Adjoint-Based Iterative Optimization

We assume the system of interest is governed by a state equation of the form

$$E \frac{d\mathbf{x}}{dt} = N(\mathbf{x}, \mathbf{f}, \mathbf{u}) \quad \text{on } 0 < t < T, \quad (1a)$$

$$\mathbf{x} = \mathbf{x}_0 \quad \text{at } t = 0, \quad (1b)$$

where $t = 0$ is the present time and

- $\mathbf{x}(t)$ is the state vector with \mathbf{x}_0 the (known) initial conditions (at $t = 0$),
- $\mathbf{f}(t)$ is the (known) applied external force (e.g., a pressure gradient), and
- $\mathbf{u}(t)$ is the “control” (e.g., a distribution of blowing and suction on the wall that we may prescribe).

The matrix E , which may be singular, and the nonlinear function $N(\mathbf{x}, \mathbf{f}, \mathbf{u})$ may be defined as necessary in order to represent any smooth ordinary differential equation

(ODE) of interest, including both low-dimensional ODEs and high-dimensional discretizations of PDE systems, such as that governing fluid turbulence. We also define a cost function, \mathcal{J} , which measures any trajectory of this system such that³

$$\mathcal{J} = \frac{1}{2} \int_0^T [|\mathbf{x}|_{Q_x}^2 + |\mathbf{u}|_{Q_u}^2] dt + \frac{1}{2} |\mathbf{E}\mathbf{x}(T)|_{Q_T}^2. \quad (1c)$$

The norms are each weighted such that, e.g., $|\mathbf{x}|_{Q_x}^2 \triangleq \mathbf{x}^* Q_x \mathbf{x}$, where $(\cdot)^*$ denotes the conjugate transpose with $Q_x \geq 0$, $Q_u > 0$ and $Q_T \geq 0$. The cost function (specifically, the selection of Q_x , Q_u and Q_T) mathematically represents what we would like the controls \mathbf{u} to accomplish in this system.⁴ In short, the problem at hand is to minimize \mathcal{J} with respect to the control distribution $\mathbf{u}(t)$ subject to Equation 1.

We now consider what happens when we simply perturb the inputs to our original system (Equation 1) a small amount. Small perturbations \mathbf{u}' to the control \mathbf{u} cause small perturbations \mathbf{x}' to the state \mathbf{x} . Such perturbations are governed by the perturbation equation (a.k.a. tangent linear equation)

$$\mathcal{L}\mathbf{x}' = B\mathbf{u}' \Leftrightarrow E \frac{d\mathbf{x}'}{dt} = A\mathbf{x}' + B\mathbf{u}' \quad \text{on } 0 < t < T, \quad (2a)$$

$$\mathbf{x}' = 0 \quad \text{at } t = 0, \quad (2b)$$

where the operator $\mathcal{L} = (E \frac{d}{dt} - A)$ and matrices A and B are obtained via the linearization⁵ of Equation 1a about the trajectory $\mathbf{x}(\mathbf{u})$. The concomitant small perturbation to the cost function \mathcal{J} is given by

$$\mathcal{J}' = \int_0^T (\mathbf{x}'^* Q_x \mathbf{x}' + \mathbf{u}'^* Q_u \mathbf{u}') dt + \mathbf{x}'^*(T) E^* Q_T E \mathbf{x}'(T). \quad (2c)$$

Note that Equation 2a-b implicitly represents a linear relationship between \mathbf{x}' and \mathbf{u}' . Knowing this, the task is to re-express \mathcal{J}' in such a way as to make the resulting linear relationship between \mathcal{J}' and \mathbf{u}' explicitly evident, at which point the gradient $\mathcal{D}\mathcal{J}/\mathcal{D}\mathbf{u}$ may readily be defined. To this end, define the weighted inner product $\langle\langle \mathbf{r}, \mathbf{x}' \rangle\rangle_R \triangleq \int_0^T \mathbf{r}^* R \mathbf{x}' dt$ with some $R > 0$ and express the following adjoint identity:

$$\langle\langle \mathbf{r}, \mathcal{L}\mathbf{x}' \rangle\rangle_R = \langle\langle \mathcal{L}^* \mathbf{r}, \mathbf{x}' \rangle\rangle_R + \mathbf{b}. \quad (3)$$

³Note that, if E is nonsingular, then the terminal penalty in Equation 1c may be written in the simpler form $\frac{1}{2} |\mathbf{x}(T)|_{Q_T}^2$, and an E^{-1} may be absorbed into the terminal condition on R in Equation 4b in the analysis that follows. However, if E is singular (so that the analysis may be applied, e.g., to $\{u, v, w, p\}$ formulations of the incompressible NSE), then this simpler terminal condition does not work out in the analysis that follows for arbitrary $Q_T \geq 0$.

⁴Physically, we can say that the control objective is to minimize some measure of the state (as measured by the first and third terms of \mathcal{J} in Equation 1c) without using too much control effort to do so (as measured by the second term of \mathcal{J}). Nonquadratic forms for \mathcal{J} are also possible. Note, in particular, that the terminal cost (the last term of Equation 1c) enables, in effect, the penalization of the dynamics likely to come after the finite optimization horizon $t \in [0, T]$; including such a term in the optimization problem significantly improves its long-time behavior when applied in the receding-horizon model predictive control framework (Bitmead et al. 1990), as found in the application of this method to the control of channel-flow turbulence, as reported in Bewley et al. (2001).

⁵That is, substitute $\mathbf{x} + \mathbf{x}'$ for \mathbf{x} and $\mathbf{u} + \mathbf{u}'$ for \mathbf{u} in Equation 1a, multiply out, and retain all terms that are linear in the perturbation quantities.

Using integration by parts, it follows that $\mathcal{L}^* \mathbf{r} = -R^{-1}(E^* \frac{d}{dt} + A^*)R \mathbf{r}$ and $\mathbf{b} = [\mathbf{r}^* R E \mathbf{x}']_{t=0}^{t=T}$. We now define the relevant adjoint equation by

$$\mathcal{L}^* \mathbf{r} = R^{-1} Q_x \mathbf{x} \Leftrightarrow -E^* R \frac{d\mathbf{r}}{dt} = A^* R \mathbf{r} + Q_x \mathbf{x} \quad \text{on } 0 < t < T, \quad (4a)$$

$$\mathbf{r} = R^{-1} Q_T E \mathbf{x} \quad \text{at } t = T. \quad (4b)$$

The adjoint field \mathbf{r} so defined is easy to compute via a backward march from $t = T$ back to $t = 0$. Both A^* and the forcing term $Q_x \mathbf{x}$ in Equation 4a are functions of $\mathbf{x}(t)$, which itself must be determined from a forward march of Equation 1 from $t = 0$ to $t = T$; thus, $\mathbf{x}(t)$ must be saved on this forward march over the interval $t \in [0, T]$ in order to calculate Equation 4 via a backward march from $t = T$ back to $t = 0$. The need for storing $\mathbf{x}(t)$ on $[0, T]$ during this forward march to construct the adjoint on the backward march can present a significant storage problem. This problem may be averted with a checkpointing algorithm, which saves $\mathbf{x}(t)$ only occasionally on the forward march, then recomputes $\mathbf{x}(t)$ as necessary from these ‘‘checkpoints’’ during the backward march for $\mathbf{r}(t)$. Noting Equations 2 and 4, it follows from Equation 3 that

$$\begin{aligned} \int_0^T \mathbf{r}^* R B \mathbf{u}' dt &= \int_0^T \mathbf{x}^* Q_x \mathbf{x}' dt + \mathbf{x}^*(T) E^* Q_T E \mathbf{x}'(T) \\ \Rightarrow \mathcal{J}' &= \int_0^T [B^* R \mathbf{r} + Q_u \mathbf{u}]^* \mathbf{u}' dt \triangleq \left\langle \left\langle \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{u}}, \mathbf{u}' \right\rangle \right\rangle_{S_u}. \end{aligned}$$

As \mathbf{u}' is arbitrary, the desired gradient is thus given by

$$\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{u}} = S_u^{-1} [B^* R \mathbf{r} + Q_u \mathbf{u}], \quad (5)$$

and is readily determined from the adjoint field \mathbf{r} defined by Equation 4. This gradient may be used to update \mathbf{u} at each iteration k via any of a number of standard optimization strategies, including steepest descent, preconditioned nonquadratic conjugate gradient, and limited-memory Broyden-Fletcher-Goldfarb-Shanno (BFGS) (see Nocedal & Wright 1999).

Clearly, there is substantial flexibility in the framing of the optimization problem described above. Once the equations governing the system are specified, this flexibility comes in exactly three forms:

- targeting the cost function via selection of the Q_i matrices,
- regularizing the adjoint operator via selection of the R matrix, and
- preconditioning the gradient via selection of the S_i matrices.

In the ODE setting, the simplest approach is to select the identity matrix for some if not all of these weighting matrices. However, in high-dimensional discretizations of multiscale turbulent flow systems, this is not always the best choice. By incorporating finite-dimensional discretizations of Sobolev inner products (with derivatives and/or antiderivatives in space and/or time) in place of L^2 inner products, different scales of the problem may be emphasized or de-emphasized in the statement of the cost function, in the dynamics of the corresponding adjoint field, and in the extraction of the gradient. Such alternative inner products (in the infinite-dimensional setting) or

weighting matrices (in the finite-dimensional setting) can have a substantially beneficial effect on the resulting optimization process when applied to multiscale problems. In fact, the flexibility of the inner products implied by the Q_i , R , and S_i operators parameterizes exactly all of the available options to target/regularize/precondition the entire optimization framework laid out in this chapter (Protas et al. 2004). For clarity of presentation, Sections 3.2 through 3.4 will take $R = S_i = I$ everywhere just to simplify the form of the equations presented, recognizing that other choices can (and, in some cases, should) be preferred.

Note also that, at a given iteration k of a gradient-based optimization procedure, for a given value of the optimization variables \mathbf{u}^k and a given update direction \mathbf{p}^k , one needs to determine the parameter of descent α to perform a “line minimization,” that is, to minimize $\mathcal{J}(\mathbf{u}^k + \alpha\mathbf{p}^k)$ with respect to the scalar α . By solving the perturbation equation (Equation 2) for \mathbf{x}' in the direction $\mathbf{u}' = \mathbf{p}^k$ from the point $\mathbf{u} = \mathbf{u}^k$ and $\mathbf{x} = \mathbf{x}^k$, it is straightforward to get an estimate of the most suitable value for α in the case that \mathcal{J} is nearly quadratic in \mathbf{u} . Fixing \mathbf{u}^k and \mathbf{p}^k for the moment, performing a truncated Taylor series expansion for $\mathcal{J}(\mathbf{u}^k + \alpha\mathbf{p}^k)$ about the value $\mathcal{J}(\mathbf{u}^k)$, and setting the derivative with respect to α equal to zero gives

$$\alpha = - \frac{d\mathcal{J}(\mathbf{u}^k + \alpha\mathbf{p}^k)}{d\alpha} \Big|_{\alpha=0} \Big/ \frac{d^2\mathcal{J}(\mathbf{u}^k + \alpha\mathbf{p}^k)}{d\alpha^2} \Big|_{\alpha=0}, \quad (6)$$

where the derivatives shown are simple functions of $\{\mathbf{x}, \mathbf{x}', \mathbf{u}, \mathbf{u}'\}$, as readily determined from Equation 1c. This value of α minimizes $\mathcal{J}(\mathbf{u}^k + \alpha\mathbf{p}^k)$ if \mathcal{J} is quadratic in \mathbf{u} . If it is not quadratic (for example, if the relationship $\mathbf{x}(\mathbf{u})$ implied by Equation 1a is nonlinear), this value of α might not accurately minimize $\mathcal{J}(\mathbf{u}^k + \alpha\mathbf{p}^k)$ with respect to α , and in fact may lead to an unstable algorithm if used at each iteration of a gradient descent procedure. However, Equation 6 is still useful to initialize the guess for α at each iteration.

3.2. Control via Riccati-Based Feedback

In Equation 5, assuming that $R = I$, the control \mathbf{u} , which minimizes \mathcal{J} , is given by

$$\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{u}} = 0 \quad \Rightarrow \quad \mathbf{u} = -Q_{\mathbf{u}}^{-1}B^*\mathbf{r}.$$

We now consider the problem that arises when we start with a governing equation (Equations 1a–b) for the state variable \mathbf{x} that is already in the linearized form of a perturbation equation (as in Equations 2a–b). In other words, we perturb the (already linear) system about the control distribution $\mathbf{u} = 0$ and the trajectory $\mathbf{x}(\mathbf{u}) = 0$, and thus the perturbed system is $\mathbf{u} = \mathbf{u}'$ and $\mathbf{x} = \mathbf{x}'$. Combining the perturbation and adjoint Equations 2 and 4 into a single matrix form, applying the optimal value of the control \mathbf{u} as noted above, and assuming for simplicity that $E = I$, gives:

$$\frac{d}{dt} \begin{bmatrix} \mathbf{x}' \\ \mathbf{r} \end{bmatrix} = \begin{bmatrix} A & -BQ_{\mathbf{u}}^{-1}B^* \\ -Q_{\mathbf{x}} & -A^* \end{bmatrix} \begin{bmatrix} \mathbf{x}' \\ \mathbf{r} \end{bmatrix} \quad \text{where} \quad \begin{cases} \mathbf{x}' = 0 & \text{at } t = 0, \\ \mathbf{r} = Q_T\mathbf{x}' & \text{at } t = T. \end{cases} \quad (7)$$

This ODE, with both initial and terminal conditions, is referred to as a two-point boundary value problem. Its general solution may be found via the sweep method (Bryson & Ho 1969): Assuming there exists a relation between the perturbation vector $\mathbf{x}' = \mathbf{x}'(t)$ and the adjoint vector $\mathbf{r} = \mathbf{r}(t)$ via a matrix $X = X(t)$ such that

$$\mathbf{r} = X\mathbf{x}', \quad (8)$$

inserting this assumed form of the solution (a.k.a. solution ansatz) into the combined matrix form (Equation 7) to eliminate \mathbf{r} , combining rows to eliminate $d\mathbf{x}/dt$, factoring out \mathbf{x}' to the right, and noting that this equation holds for all \mathbf{x}' , it follows that X obeys the differential Riccati equation (DRE)

$$-\frac{dX}{dt} = A^*X + XA - XBQ_u^{-1}B^*X + Q_x \quad \text{where} \quad X(T) = Q_T, \quad (9a)$$

where the condition at $X(T)$ follows from Equations 7 and 8. Solutions $X = X(t)$ of this matrix equation satisfy $X^* = X$, and may easily be determined via marching procedures similar to those used to march ODEs (Crank-Nicholson, Runge-Kutta, etc.). By characterizing the optimal point, we now write the control \mathbf{u} as

$$\mathbf{u} = K\mathbf{x} \quad \text{where} \quad K = -Q_u^{-1}B^*X. \quad (9b)$$

To recap, this value of K minimizes

$$\mathcal{J} = \frac{1}{2} \int_0^T [\mathbf{x}^* Q_x \mathbf{x} + \mathbf{u}^* Q_u \mathbf{u}] dt + \frac{1}{2} \mathbf{x}^*(T) Q_T \mathbf{x}(T) \quad \text{where} \quad \frac{d\mathbf{x}}{dt} = A\mathbf{x} + B\mathbf{u}. \quad (9c)$$

The matrix $K = K(t)$ is referred to as the optimal control feedback gain matrix, and is a function of the solution X to Equation 9a. This equation may be solved for linear time-varying (LTV) or linear time-invariant (LTI) systems based solely on knowledge of A and B in the system model and Q_x , Q_u , and Q_T in the cost function (that is, the gain matrix K may be computed offline). Alternatively, if we take the limit that $T \rightarrow \infty$ (that is, if we consider the infinite-horizon control problem) and the system is LTI, the matrix X in Equation 9a may be marched to steady state. This steady-state solution for X satisfies the algebraic Riccati equation (ARE)

$$0 = A^*X + XA - XBQ_u^{-1}B^*X + Q_x. \quad (10)$$

Efficient algorithms to solve this quadratic matrix equation (Laub 1979), based on a Schur factorization of the $2N \times 2N$ matrix in Equation 7, are readily available in `Matlab`.

3.3. Estimation via Adjoint-Based Iterative Optimization

The derivation presented here is analogous to that presented in Section 3.1. We first write the state equation modeling the system of interest in ODE form:

$$E \frac{d\mathbf{x}}{dt} = N(\mathbf{x}, \mathbf{f}, \mathbf{v}, \mathbf{w}) \quad \text{on} \quad -T < t < 0, \quad (11a)$$

$$\mathbf{x} = \mathbf{u} \quad \text{at} \quad t = -T, \quad (11b)$$

where $t = 0$ is the present time and

- $\mathbf{x}(t)$ is the state vector,
- $\mathbf{f}(t)$ models the known external forcing,

and the quantities to be optimized are:

- \mathbf{u} represents the unknown initial conditions of the model (at $t = -T$),
- \mathbf{v} contains the unknown constant parameters of the model (*Re*, etc.), and
- $\mathbf{w}(t)$ contains the unknown external inputs that we would like to determine.

We next write a cost function that measures the misfit of the available measurements $\mathbf{y}(t)$ with the corresponding quantity in the computational model $C\mathbf{x}(t)$, and additionally penalizes the deviation of the initial condition \mathbf{u} from any a priori estimate⁶ of the initial conditions $\bar{\mathbf{u}}$, the deviation of the parameters \mathbf{v} from any a priori estimate of the parameters $\bar{\mathbf{v}}$, and the magnitude of the disturbance terms $\mathbf{w}(t)$:

$$\mathcal{J} = \frac{1}{2} \int_{-T}^0 |C\mathbf{x} - \mathbf{y}|_{Q_y}^2 dt + \frac{1}{2} |\mathbf{u} - \bar{\mathbf{u}}|_{Q_u}^2 + \frac{1}{2} |\mathbf{v} - \bar{\mathbf{v}}|_{Q_v}^2 + \frac{1}{2} \int_{-T}^0 |\mathbf{w}|_{Q_w}^2 dt. \quad (11c)$$

The norms are each weighted with positive semidefinite matrices such that, e.g., $|\mathbf{y}|_{Q_y}^2 \triangleq \mathbf{y}^* Q_y \mathbf{y}$ with $Q_y \geq 0$. In short, the problem at hand is to minimize \mathcal{J} with respect to $\{\mathbf{u}, \mathbf{v}, \mathbf{w}(t)\}$ subject to Equation 11.

Small perturbations $\{\mathbf{u}', \mathbf{v}', \mathbf{w}'(t)\}$ to $\{\mathbf{u}, \mathbf{v}, \mathbf{w}(t)\}$ cause small perturbations \mathbf{x}' to the state \mathbf{x} . Such perturbations are governed by the perturbation equation

$$\mathcal{L}\mathbf{x}' = B_v \mathbf{v}' + B_w \mathbf{w}' \Leftrightarrow E \frac{d\mathbf{x}'}{dt} = A\mathbf{x}' + B_v \mathbf{v}' + B_w \mathbf{w}' \quad \text{on } -T < t < 0, \quad (12a)$$

$$\mathbf{x}' = \mathbf{u}' \quad \text{at } t = -T, \quad (12b)$$

where the operator $\mathcal{L} = (E \frac{d}{dt} - A)$ and the matrices A , B_v , and B_w are obtained via the linearization of Equation 11a about the trajectory $\mathbf{x}(\mathbf{u}, \mathbf{v}, \mathbf{w})$. The concomitant small perturbation to the cost function \mathcal{J} is given by

$$\mathcal{J}' = \int_{-T}^0 (C\mathbf{x} - \mathbf{y})^* Q_y C \mathbf{x}' dt + (\mathbf{u} - \bar{\mathbf{u}})^* Q_u \mathbf{u}' + (\mathbf{v} - \bar{\mathbf{v}})^* Q_v \mathbf{v}' + \int_{-T}^0 \mathbf{w}^* Q_w \mathbf{w}' dt. \quad (13)$$

Again, the task before us is to re-express \mathcal{J}' in such a way as to make the resulting linear relationship between \mathcal{J}' and $\{\mathbf{u}', \mathbf{v}', \mathbf{w}'(t)\}$ explicitly evident, at which point the necessary gradients may readily be defined. To this end, we define the inner product $\langle \mathbf{r}, \mathbf{x}' \rangle \triangleq \int_0^T \mathbf{r}^* \mathbf{x}' dt$ and express the adjoint identity

$$\langle \mathbf{r}, \mathcal{L}\mathbf{x}' \rangle = \langle \mathcal{L}^* \mathbf{r}, \mathbf{x}' \rangle + \mathbf{b}. \quad (14)$$

Using integration by parts, it follows that $\mathcal{L}^* \mathbf{r} = -(E^* \frac{d}{dt} + A^*) \mathbf{r}$ and $\mathbf{b} = [\mathbf{r}^* E \mathbf{x}']_{t=-T}^{t=0}$. Based on this adjoint operator, we now define an adjoint equation of the form

$$\mathcal{L}^* \mathbf{r} = C^* Q_y (C\mathbf{x} - \mathbf{y}) \Leftrightarrow -E^* \frac{d\mathbf{r}}{dt} = A^* \mathbf{r} + C^* Q_y (C\mathbf{x} - \mathbf{y}) \quad \text{on } -T < t < 0, \quad (15a)$$

$$\mathbf{r} = 0 \quad \text{at } t = 0. \quad (15b)$$

⁶In the 4Dvar setting, such an estimate $\bar{\mathbf{u}}$ for $\mathbf{x}(-T)$ is obtained from the previously computed forecast, and the corresponding term in the cost function is called the “background” term. The effect of this term on the time evolution of the forecast is significant and sometimes detrimental, as it constrains the update to \mathbf{u} to be small when, in some circumstances, a large update might be warranted.

Note again that the difficulty involved with numerically solving the ODE given by Equation 15 via a backward march from $t = 0$ to $t = -T$ is essentially the same as the difficulty involved with solving the original ODE (Equation 11).

Finally, combining Equations 12 and 15 into the identity Equation 14 and substituting into Equation 13, it follows that

$$\begin{aligned} \mathcal{J}' = & [E^* \mathbf{r}(-T) + Q_u(\mathbf{u} - \bar{\mathbf{u}})]^* \mathbf{u}' + \left[\int_{-T}^0 B_v^* \mathbf{r} dt + Q_v(\mathbf{v} - \bar{\mathbf{v}}) \right]^* \mathbf{v}' \\ & + \int_{-T}^0 [B_w^* \mathbf{r} + Q_w \mathbf{w}]^* \mathbf{w}' dt \triangleq \left\langle \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{u}}, \mathbf{u}' \right\rangle_{S_u} + \left\langle \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{v}}, \mathbf{v}' \right\rangle_{S_v} + \left\langle \left\langle \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{w}}, \mathbf{w}' \right\rangle_{S_w} \right\rangle_{S_w}, \end{aligned}$$

for some $S_u > 0$, $S_v > 0$, and $S_w > 0$, where $\langle \mathbf{a}, \mathbf{b} \rangle_S \triangleq \mathbf{a}^* \mathbf{S} \mathbf{b}$ and $\langle\langle \mathbf{a}, \mathbf{b} \rangle\rangle_S \triangleq \int_{-T}^0 \mathbf{a}^* \mathbf{S} \mathbf{b} dt$, and thus

$$\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{u}} = S_u^{-1} [E^* \mathbf{r}(-T) + Q_u(\mathbf{u} - \bar{\mathbf{u}})], \quad \frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{v}} = S_v^{-1} \left[\int_{-T}^0 B_v^* \mathbf{r} dt + Q_v(\mathbf{v} - \bar{\mathbf{v}}) \right],$$

and

$$\frac{\mathcal{D}\mathcal{J}}{\mathcal{D}\mathbf{w}(t)} = S_w^{-1} [B_w^* \mathbf{r}(t) + Q_w \mathbf{w}(t)], \quad \text{for } t \in [-T, 0]. \quad (16)$$

We have thus defined the gradient of the cost function with respect to the optimization variables $\{\mathbf{u}, \mathbf{v}, \mathbf{w}(t)\}$ as a function of the adjoint field \mathbf{r} defined in Equation 15, which, for any trajectory $\mathbf{x}(\mathbf{u}, \mathbf{v}, \mathbf{w})$ of our original system (Equation 11), may easily be computed. Optimization of $\{\mathbf{u}, \mathbf{v}, \mathbf{w}(t)\}$ may thus again be performed with a gradient-based algorithm, as discussed in Section 3.2.

3.4. Estimation via Riccati-Based Feedback

We now convert the Riccati-based estimation problem into an equivalent control problem of the form already solved (in Section 3.2). Consider the linear equations for the state \mathbf{x} , the state estimate $\hat{\mathbf{x}}$, and the state estimation error $\tilde{\mathbf{x}} = \mathbf{x} - \hat{\mathbf{x}}$:

$$d\mathbf{x}/dt = A\mathbf{x} + B\mathbf{u}, \quad \mathbf{y} = C\mathbf{x}, \quad (17)$$

$$d\hat{\mathbf{x}}/dt = A\hat{\mathbf{x}} + B\mathbf{u} - L(\mathbf{y} - \hat{\mathbf{y}}), \quad \hat{\mathbf{y}} = C\hat{\mathbf{x}}, \quad (18)$$

$$d\tilde{\mathbf{x}}/dt = A\tilde{\mathbf{x}} + L\tilde{\mathbf{y}}, \quad \tilde{\mathbf{y}} = C\tilde{\mathbf{x}}, \quad (19)$$

where \mathbf{u} is now the (known) control forcing and L is the unknown gain matrix to be determined. The output injection term $L(\mathbf{y} - \hat{\mathbf{y}})$ applied to the equation for the state estimate $\hat{\mathbf{x}}$ is to be designed to nudge this equation appropriately based on the available measurements \mathbf{y} of the actual system. If this term is doing its job correctly, $\hat{\mathbf{x}}$ is driven toward \mathbf{x} (that is, $\tilde{\mathbf{x}}$ is driven toward zero) even if the initial conditions on \mathbf{x} are unknown and Equation 17 is only an approximate model of reality.

We thus set out to minimize some measure of the state estimation error $\tilde{\mathbf{x}}$ by appropriate selection of L . To this end, taking $\tilde{\mathbf{x}}^*$ times Equation 19, we obtain

$$\tilde{\mathbf{x}}^* \left[\frac{d\tilde{\mathbf{x}}}{dt} = A\tilde{\mathbf{x}} + LC\tilde{\mathbf{x}} \right] = \left[\frac{d\tilde{\mathbf{x}}}{dt} = A^* \tilde{\mathbf{x}} + C^* L^* \tilde{\mathbf{x}} \right]^* \tilde{\mathbf{x}} = \frac{1}{2} \frac{d\tilde{\mathbf{x}}^* \tilde{\mathbf{x}}}{dt}. \quad (20)$$

Motivated by the second relation in brackets above, consider a new system

$$d\mathbf{z}/dt = A^*\mathbf{z} + C^*\tilde{\mathbf{u}} \quad \text{where} \quad \tilde{\mathbf{u}} = L^*\mathbf{z} \quad \text{and} \quad \mathbf{z}(-T) = \tilde{\mathbf{x}}(-T). \quad (21)$$

Although the dynamics of $\tilde{\mathbf{x}}(t)$ and $\mathbf{z}(t)$ are different, the evolution of their energy is the same, by Equation 20. That is, $\mathbf{z}^*\mathbf{z} \neq \tilde{\mathbf{x}}^*\tilde{\mathbf{x}}$ even though, in general, $\mathbf{z}(t) \neq \tilde{\mathbf{x}}(t)$. Thus, for convenience, we design L to minimize a cost function related to this auxiliary variable \mathbf{z} , defined here such that, taking $Q_1 = I$,

$$\tilde{J} = \frac{1}{2} \int_{-T}^0 [\mathbf{z}^* Q_1 \mathbf{z} + \tilde{\mathbf{u}}^* Q_2 \tilde{\mathbf{u}}] dt + \frac{1}{2} \mathbf{z}^*(-T) Q_{-T} \mathbf{z}(-T), \quad (22a)$$

where, renaming $\tilde{A} = A^*$, $\tilde{B} = C^*$, and $\tilde{K} = L^*$, Equation 21 may be written as

$$d\mathbf{z}/dt = \tilde{A}\mathbf{z} + \tilde{B}\tilde{\mathbf{u}} \quad \text{where} \quad \tilde{\mathbf{u}} = \tilde{K}\mathbf{z}. \quad (22b)$$

Finding the feedback gain matrix \tilde{K} in Equation 22b that minimizes the cost function \tilde{J} in Equation 22a is exactly the same problem that is solved in Equation 9, just with different variables. Thus, the optimal gain matrix L , which minimizes a linear combination of the energy of the state estimation error, $\tilde{\mathbf{x}}^*\tilde{\mathbf{x}}$, and some measure of the estimator feedback gain L , is again determined from the solution P of a Riccati equation which, making the appropriate substitutions into the solution presented in Equation 9, is given by

$$dP/dt = AP + PA^* - PC^*Q_2^{-1}CP + Q_1, \quad P(-T) = Q_{-T}, \quad L = -PC^*Q_2^{-1}. \quad (23)$$

The compact derivation presented above gets quickly to the Riccati equation for an optimal estimator, but as the result of a somewhat contrived optimization problem. A more intuitive formulation is to replace Equation 17 with

$$d\mathbf{x}/dt = A\mathbf{x} + B\mathbf{u} + \mathbf{w}_1, \quad \mathbf{y} = C\mathbf{x} + \mathbf{w}_2, \quad (24)$$

where \mathbf{w}_1 (the “state disturbances”) and \mathbf{w}_2 (the “measurement noise”) are assumed to be uncorrelated, zero mean, white Gaussian processes with modeled spectral densities $\mathcal{E}\{\mathbf{w}_1\mathbf{w}_1^*\} = Q_1$ and $\mathcal{E}\{\mathbf{w}_2\mathbf{w}_2^*\} = Q_2$, respectively. Going through the necessary steps to minimize the expected energy of the estimation error, $\mathcal{E}\{\tilde{\mathbf{x}}^*\tilde{\mathbf{x}}\} = \text{trace}(P)$, where $P = \mathcal{E}\{\tilde{\mathbf{x}}\tilde{\mathbf{x}}^*\}$, we again arrive at an estimator of the form given in Equation 18 with the feedback gain matrix L as given by Equation 23. For a succinct derivation of this setting in continuous time, see pp. 460–70 of Lewis & Syrmos (1995); for a succinct derivation in discrete time, see pp. 382–95 of Franklin et al. (1998); for a more comprehensive discussion, see Anderson & Moore (2005).

3.5. LQG and the Separation Principle: Putting It Together

In Section 3.2, a convenient feedback relationship was derived for determining optimal control inputs based on full state information. In Section 3.4, a convenient feedback relationship was derived for determining an optimal estimate of the full state based on the available system measurements. In the practical situation in which control inputs must be determined based on available system measurements, it is thus natural to combine the results of these two sections—that is, to develop an estimate

of the state $\hat{\mathbf{x}}$ based on the results of Section 3.4, then to apply control $\mathbf{u} = K\hat{\mathbf{x}}$ based on this state estimate and the results of Section 3.2. This setting is referred to as LQG control, with reference to the Linear state equation, Quadratic cost function, and Gaussian disturbance model upon which it is based. It is justified by the following separation principle: By collecting the equations presented previously and adding a reference control input \mathbf{r} , we have

Plant :	$d\mathbf{x}/dt = A\mathbf{x} + B\mathbf{u} + \mathbf{w}_1,$	$\mathbf{y} = C\mathbf{x} + \mathbf{w}_2,$
Estimator :	$d\hat{\mathbf{x}}/dt = A\hat{\mathbf{x}} + B\mathbf{u} - L(\mathbf{y} - \hat{\mathbf{y}}),$	$\hat{\mathbf{y}} = C\hat{\mathbf{x}},$
Controller :	$\mathbf{u} = K\hat{\mathbf{x}} + \mathbf{r},$	

where K is determined as in Equation 9 and L is determined as in Equation 23. Note that K was determined in Section 3.2 in the nominal case without a reference input, and may be extended here to the case with a control reference in a straightforward fashion. In block matrix form (noting that $\tilde{\mathbf{x}} = \mathbf{x} - \hat{\mathbf{x}}$), this composite system may be written

$$\frac{d}{dt} \begin{bmatrix} \mathbf{x} \\ \tilde{\mathbf{x}} \end{bmatrix} = \begin{bmatrix} A + BK & -BK \\ 0 & A + LC \end{bmatrix} \begin{bmatrix} \mathbf{x} \\ \tilde{\mathbf{x}} \end{bmatrix} + \begin{bmatrix} B \\ 0 \end{bmatrix} \mathbf{r} + \begin{bmatrix} I & 0 \\ I & L \end{bmatrix} \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix} \quad (25a)$$

$$\mathbf{y} = [C \quad 0] \begin{bmatrix} \mathbf{x} \\ \tilde{\mathbf{x}} \end{bmatrix} + [0 \quad I] \begin{bmatrix} \mathbf{w}_1 \\ \mathbf{w}_2 \end{bmatrix}. \quad (25b)$$

Because this system matrix is block triangular, its eigenvalues are given by the union of the eigenvalues of $A + BK$ and those of $A + LC$; thus, selecting K and L to stabilize the control and estimation problems separately effectively stabilizes the composite system. Further, assuming that $\mathbf{w}_1 = \mathbf{w}_2 = 0$ and the initial condition on all variables are zero, taking the Laplace transform⁷ of Equation 25 gives

$$\begin{aligned} \mathbf{Y}(s) &= [C \quad 0] \begin{bmatrix} sI - (A + BK) & BK \\ 0 & sI - (A + LC) \end{bmatrix}^{-1} \begin{bmatrix} B \\ 0 \end{bmatrix} \mathbf{R}(s) \\ &= C[sI - (A + BK)]^{-1} B \mathbf{R}(s). \end{aligned}$$

That is, the transfer function from \mathbf{r} to \mathbf{y} is unaffected by the estimator. As a matter of practice, the estimator feedback L is typically designed (by adjusting the relative magnitude of Q_1 and Q_2) such that the slowest eigenvalues of $A + LC$ are a factor of two to five faster than the slowest eigenvalues of $A + BK$.

3.6. Transforming Navier-Stokes to State-Space Form

To illustrate how the linearized Navier-Stokes equations may be written in state-space form, we begin with these equations manipulated into Orr-Sommerfeld/Squire form

⁷In effect, simply replacing d/dt by the Laplace variable s and replacing the time-domain signals $\{\mathbf{y}, \mathbf{r}, \mathbf{x}, \tilde{\mathbf{x}}\}$ with their Laplace transforms $\{\mathbf{Y}(s), \mathbf{R}(s), \mathbf{X}(s), \tilde{\mathbf{X}}(s)\}$, where $\mathbf{F}(s) = \int_0^\infty \mathbf{f}(t) e^{-st} dt$.

(Schmid & Henningson 2001):

$$\frac{\partial}{\partial t} \begin{bmatrix} \hat{v} \\ \hat{\omega}_y \end{bmatrix} = \begin{bmatrix} \mathcal{L} & 0 \\ \mathcal{C} & \mathcal{S} \end{bmatrix} \begin{bmatrix} \hat{v} \\ \hat{\omega}_y \end{bmatrix} \quad \text{with} \quad \begin{cases} \mathcal{L} = \Delta^{-1}(-ik_x U \Delta + ik_x U'' + \Delta^2 / Re), \\ \mathcal{S} = -ik_x U + \Delta / Re, \\ \mathcal{C} = -ik_z U', \end{cases} \quad (26)$$

where \hat{v} and $\hat{\omega}_y$ denote Fourier coefficients of the wall-normal velocity and vorticity, respectively, at wave-number pair $\{k_x, k_z\}$. These Fourier coefficients can be expanded in terms of known functions $P_n(y)$ and $R_n(y)$:

$$\hat{v}(y, t) = \sum_{n=1}^{N_y} a_n(t) P_n(y), \quad \hat{\omega}_y(y, t) = \sum_{n=1}^{N_y} b_n(t) R_n(y). \quad (27)$$

Substitution of Equation 27 into Equation 26, followed by a projection with a weighted residual method, yields a linear system in the state-space form Equation 17 (Joshi et al. 1997).

Alternatively, the derivative operators in Equation 26 may be discretized in a straightforward manner using the matrix collocation operators provided, e.g., by the spectral Matlab Differentiation Matrix Suite of Weideman & Reddy (2000), which implements the technique of Huang & Sloan (1993) with “clamped” (homogeneous) boundary conditions to avoid spurious eigenvalues. The technique of Högberg et al. (2003a) may then be used to “lift” the boundary conditions on v to account for forcing on the problem via an additional right-hand-side forcing vector \mathbf{u} .

4. COPING WITH HIGH-DIMENSIONAL DISCRETIZATIONS

The principal difficulty with applying linear control theory to fluid systems is that most flows (turbulent flows in particular) require very high-dimensional numerical discretizations to resolve accurately. The matrix equations at the heart of the feedback calculations presented in Section 3 are simply intractable for such discretizations in their original form. There are a variety of ways around this dimensionality predicament, as discussed below.

4.1. The Parallel Flow Assumption

If the mean flow is parallel (or may locally be approximated as such), it is well known in the fluids literature that performing a Fourier transform of the linearized Navier-Stokes equations decouples the modes of the fluid system at each wave-number pair $\{k_x, k_z\}$; when expressed in $v - \omega$ form, this results in the Orr-Sommerfeld/Squire equations. It follows similarly that, under this parallel flow assumption, by performing a Fourier transform of all variables in the control problem (that is, the state, the controls, the measurements, and the disturbances), the enormous Riccati equations for both the control and estimation problems block diagonalize into $N_{k_x} \times N_{k_z}$ completely independent, tractable Riccati equations (each of dimension $2N_y \times 2N_y$) that may be solved separately (Bewley & Liu 1998).

4.2. The Parabolic Flow Assumption

If a boundary layer is sufficiently thin, it is well known in the fluids literature that the Navier-Stokes equations may sometimes be reduced to a simpler, parabolic-in-space form, often referred to as the boundary-layer equations. Assuming such a parabolic-in-space development of the perturbations in the flow system, we may follow a control strategy similar to that used for the much more common parabolic-in-time systems, as highlighted in the previous section, with one important difference: There is a unique noncausal capability of control algorithms in this parabolic-in-space setting. That is, measurements at a particular streamwise location may be used to update both downstream and upstream controls to neutralize the effects of disturbances that enter the boundary layer both downstream and upstream of the actuator itself. This is not possible in control strategies for parabolic-in-time systems, which must be constrained to act in a causal fashion to be implementable. The necessary extensions of standard causal-in-time control theory to handle this unique noncausal-in-space capability are formulated in Cathalifaud & Bewley (2004).

4.3. The Chandrasekhar Method for Approximate Solution of Differential Riccati Equations

Consider a problem in which the state vector is of dimension N , the control vector is of dimension M_u , and the measurement vector is of dimension M_y . If $N \gg M_u$ and $N \gg M_y$, which is typical, then solving Riccati equations for the $N \times N$ matrices X (see Equation 9) and P (see Equation 23) in order to compute the $M_u \times N$ and $N \times M_y$ feedback matrices K and L for the control and estimation problems, respectively, might seem to be inefficient: This approach computes enormous $N \times N$ Riccati matrices only to effectively take narrow slices of them to determine the feedback gains. Chandrasekhar's method (Kailath 1973) addresses this inefficiency in a clever way. Consider the DRE for the estimator, as given in Equation 23:

$$\dot{P}(t) = AP(t) + P(t)A^* - P(t)C^*Q_2^{-1}CP(t) + Q_1, \quad L(t) = -P(t)C^*Q_2^{-1}. \quad (28)$$

Chandrasekhar's method solves an evolution equation for a low-dimensional factored form of $\dot{P}(t)$ and another evolution equation for $L(t)$. To this end, define

$$\dot{P} = Y_1Y_1^* - Y_2Y_2^* = YHY^*, \quad Y = \begin{pmatrix} Y_1 & Y_2 \end{pmatrix}, \quad H = \begin{pmatrix} I & 0 \\ 0 & -I \end{pmatrix},$$

where the number of columns of Y_1 and Y_2 are the number of positive and negative eigenvalues of \dot{P} , respectively, retained in the approximation. Differentiating Equation 28 with respect to time and inserting $\dot{P} = YHY^*$, assuming $\{A, B, C, Q_1, Q_2\}$ are LTI, it is easily verified that the following set of equations are equivalent to Equation 28, but much cheaper to compute if the factors Y_1 and Y_2 are low rank:

$$\begin{aligned} \dot{L}(t) &= -Y(t)HY^*(t)C^*Q_2^{-1}, & L(0) &= -P(0)C^*Q_2^{-1}, \\ \dot{Y}(t) &= [A + L(t)C]Y(t), & Y(0)HY^*(0) &= \dot{P}(0), \end{aligned}$$

where $\dot{P}(0)$ is determined from the original DRE (Equation 28) evaluated at $t = 0$. This method has been applied to control a heat equation system in Borggaard & Burns (2002), to control a Burgers' equation system in Camphouse & Myatt (2004), and to estimate a Navier-Stokes system in Hoepffner et al. (2005).

4.4. Open-Loop Model Reduction via Balanced Truncation

The related problems of open-loop model reduction and compensator reduction have received sustained attention in the systems and control literature; Zhou et al. (1995), Antoulas & Sorensen (2001), and Obinata & Anderson (2001) provide detailed overviews. Solutions to the open-loop model reduction problem are appropriate for situations in which the system model is simply too complex for computing feedback via the techniques described above, whereas solutions to the compensator reduction problem are better suited⁸ for situations in which the feedback calculation is solvable using a sufficiently large computer, but the resulting controller is too high dimensional to implement at the necessary bandwidth with the feedback control hardware available. In this section, we present a brief description of the balanced truncation method for open-loop model reduction (Moore 1981). Although we do not have room here to describe this method from first principles, we summarize some of its key features and equations.

The aim of open-loop model reduction is to construct a reduced-order state-space realization with input-output characteristics similar to the original plant, based on which a stabilizing compensator for the plant may hopefully be designed. Balanced truncation achieves this by transforming the system matrix into an ordered form in which the leading principle submatrices in this balanced form contribute most significantly to the input-output transfer function of the plant. That is, the eigenmodes of the original system that these submatrices represent are observable, controllable, and not highly damped, whereas the remaining eigenmodes contribute less significantly to the input-output transfer function of the plant (that is, these modes are either nearly uncontrollable,⁹ nearly unobservable, highly damped, or some combination of the three). Once transformed in this manner, the trailing modes in this realization can be truncated with reduced impact on the input-output transfer function of the model. The primary drawback of the method is that it is unrelated to the cost function being minimized, so it is not guaranteed to keep those modes most relevant to the particular control problem being solved.

Let a general state-space realization of the plant be denoted as follows:

$$\begin{aligned} \dot{\mathbf{x}} &= \mathbf{A}\mathbf{x} + \mathbf{B}\mathbf{u} \\ \mathbf{y} &= \mathbf{C}\mathbf{x} + \mathbf{D}\mathbf{u}, \end{aligned} \quad \Leftrightarrow \quad G = \left[\begin{array}{c|c} \mathbf{A} & \mathbf{B} \\ \hline \mathbf{C} & \mathbf{D} \end{array} \right].$$

⁸Note that compensator reduction (that is, design-then-reduce) strategies can be accomplished with performance guarantees that are, to date, unavailable following the open-loop model reduction (that is, reduce-then-design) approach. See Obinata & Anderson (2001) for details.

⁹As an example of nearly unobservable and nearly uncontrollable modes, see those modes in **Figure 1** with negligible support near the walls, where the sensors and actuators are located.

The controllability Gramian P and observability Gramian Q are two useful matrices derived from G that allow us to measure to what extent a particular eigenmode \mathbf{s}_i of the system matrix A is controllable and observable. They may be defined as the solutions of the Lyapunov equations

$$AP + PA^* + BB^* = 0, \quad A^*Q + QA + C^*C = 0.$$

So defined, it may be shown, for example, that eigenmode \mathbf{s}_1 is more controllable than eigenmode \mathbf{s}_2 if $\|\mathbf{s}_1^*P\mathbf{s}_1\| > \|\mathbf{s}_2^*P\mathbf{s}_2\|$. Suppose the state-space realization is transformed by a nonsingular T such that $\mathbf{x}_b = T\mathbf{x}$ and

$$G_b = \left[\begin{array}{c|c} A_b & B_b \\ \hline C_b & D_b \end{array} \right] = \left[\begin{array}{c|c} TAT^{-1} & TB \\ \hline CT^{-1} & D \end{array} \right].$$

Moore (1981) showed that there exists a nonsingular transformation matrix T by which P and Q become equal and diagonal, that is,

$$P_b = TPT^* = \Sigma, \quad Q_b = (T^{-1})^*QT^{-1} = \Sigma, \quad (29)$$

where $\Sigma = \text{diag}(\sigma_1, \dots, \sigma_n)$ with $\sigma_1 \geq \dots \geq \sigma_n$ (referred to as the Hankel singular values of the system). Note that $P_bQ_b = TPQT^{-1} = \Sigma^2$ and $PQ = T^{-1}\Lambda T$, where $\Lambda = \text{diag}(\lambda_1, \dots, \lambda_n) = \Sigma^2$ is the matrix of eigenvalues of PQ , and the T^{-1} is the matrix of eigenvectors of PQ . The new realization G_b , with Gramians $P_b = Q_b = \Sigma$, is referred to as a balanced realization. As discussed above, modes corresponding to diminished Hankel singular values are either nearly uncontrollable, nearly unobservable, highly stable, or some combination of the three, and thus truncating these modes to reduce the model does not significantly corrupt its input-output transfer function. Thus, let Σ_2 contain the negligible Hankel singular values $(\sigma_{r+1}, \dots, \sigma_n)$ and partition in the balanced realization G_b such that

$$A_b = \left[\begin{array}{c|c} A_{11} & A_{12} \\ \hline A_{21} & A_{22} \end{array} \right], \quad B_b = \left[\begin{array}{c} B_1 \\ B_2 \end{array} \right], \quad C_b = [C_1|C_2], \quad \Sigma = \left[\begin{array}{c|c} \Sigma_1 & 0 \\ \hline 0 & \Sigma_2 \end{array} \right]. \quad (30)$$

The reduced-order model is obtained by truncating those states associated with Σ_2 :

$$G_r = \left[\begin{array}{c|c} A_{11} & B_1 \\ \hline C_1 & D \end{array} \right]. \quad (31)$$

Although no guarantees are available about the stability or performance of a compensator designed for G_r on the original system G , the H_∞ norm of the difference between the original and reduced-order open-loop systems is bounded as follows

$$\|G - G_r\|_{H_\infty} \leq 2(\sigma_{r+1} + \dots + \sigma_n). \quad (32)$$

Efficient algorithms for performing balanced truncation are available in `Matlab`.

A variant of the balanced truncation method described above has been used in Cortelezzi et al. (2001), Lee et al. (2001), and Kang (2006) for boundary-layer control. In this work, it was desired to retain the structure of eigenmodes of A (i.e., system poles) as certain modes are known to play a critical role in the energy amplification that needs to be minimized. In their modal balanced realization approach, the system

matrix A was first transformed into the Schur canonical form, converting A into a block diagonal matrix with each diagonal block corresponding to each eigenmode. The B and C matrices were also appropriately transformed. The balanced Gramian for each eigenmode was then computed and compared to determine the relative observability and controllability of each eigenmode. The truncation was done based on modal-balanced Gramian singular values. The error bound of this modal balanced reduction is similar to that of the standard balanced truncation.

It is worth commenting here on the proper orthogonal decomposition (POD) method in regard to its use as a method to construct a reduced-order model for controller design (see, e.g., Lumley & Blossey 1998). POD modes are, by design, energetically optimal in representation; that is, POD modes are the best choice in representing the energetics of a given data set. However, POD-based reduced-order models, in which low-energy modes are truncated, do not account for the observability and controllability of the modes being truncated. Consequently, some retained modes may be nearly uncontrollable or unobservable, whereas some truncated modes actually play a more vital role in the input-output transfer function of the open-loop system. This is demonstrated by Rowley (2005). A POD-based reduced-order model demonstrated dynamics that were significantly different than those of the original system. However, a balanced POD method suggested by Rowley, in which POD's snapshot method was used to compute empirical Gramians, appears to be promising, especially for large systems, as it avoids directly computing Gramians, which is computationally expensive.

5. REPRESENTATIVE APPLICATIONS AND EXTENSIONS

This section discusses a few representative applications of the above framework to transitional and turbulent flow systems, then presents two significant extensions of the framework discussed above. It is impossible to review here all related applied work in this area, or even a significant fraction of it, as this body of literature is by now quite extensive. Thus, we again refer the readers to Bewley (2001), Gunzburger (2002), Kim (2003), and Collis et al. (2004) for many further examples.

5.1. Near-Wall Feedback via Overlapping Decentralized Convolutions

As discussed in Section 4.1, performing a Fourier transform of the linearized Navier-Stokes system decouples the entire control and estimation feedback calculations into tractable subproblems at each wave-number pair $\{k_x, k_z\}$. For the control problem, this results in a feedback product $\mathbf{u} = K\mathbf{x}$ at each wave-number pair, where the vector \mathbf{x} is the state of the system at wave-number pair $\{k_x, k_z\}$ discretized in the y direction. Recall that a product at each mode in Fourier space corresponds to a convolution in physical space. Thus, upon inverse transforming the entire set of feedback gains at all wave-number pairs, three-dimensional feedback convolution kernels are obtained in physical space, relating, e.g., the control to be applied at a given point on the wall to the state of the fluid system in the three-dimensional vicinity of this point. If

the feedback problem is well framed, these convolution kernels have the following properties:

- They are independent of the box size in which they were computed, so long as the computational box is sufficiently large. This relaxes the nonphysical assumption of spatial periodicity used in their calculation, thereby connecting the artificial spatially periodic model with which they are computed to the nonspatially periodic problem of physical interest.
- They are well resolved with grid resolutions appropriate for simulating the physical system of interest, and converge upon refinement of the grid. This is necessary to give them relevance to the PDE system from which the computational control problem was derived. Note that careful framing of the feedback problem is required to achieve this. Specifically, it is found that appropriate choices of the regulation terms Q_i , R , and S_i mentioned in Section 3 must be chosen.¹⁰
- They eventually decay exponentially, and thus may be truncated to any desired degree of precision (note figure 6 of Högberg et al. 2003a; see also Bewley 2001, Bamieh et al. 2002). Such truncated kernels are spatially compact with finite support, and facilitate implementation in an overlapping decentralized fashion, thereby enabling extension to massive arrays of sensors and actuators without either communication or centralized computational bottlenecks (Bewley 2001).
- Their structure is physically tenable, but not imposed a priori. Typically, control convolution kernels angle upstream away from each actuator, whereas estimation convolution kernels extend well downstream of each sensor (see, e.g., **Figure 3**). Interesting flow physics (specifically, cause/effect relationships in the near-wall region) may thus be characterized in a new way by examining these kernels.

Figure 4 illustrates the effectiveness of such kernels on the estimation of a flow perturbation developing from a localized disturbance in a laminar channel flow. For further discussion of how such well-behaved kernels are derived and their effectiveness on both the transition and turbulence problems, see Bewley (2001), Högberg et al. (2003a), Hoepffner et al. (2005), and Chevalier et al. (2006).

5.2. Applications to Controlling Turbulence

Lee et al. (2001) used the LTR (loop-transfer recovery) variant of LQG synthesis in designing an optimal controller for drag reduction in a turbulent channel flow. The LTR procedure assumes that the system noise has a certain form in order to warrant robust performance in the limit the system noise power spectral density

¹⁰Recall that Q_i denotes the weights in the cost function for the control problem and the covariance of the state disturbances and measurement noise in the estimation problem. Significantly, the simplest choice, taking Q_i , R , and S_i each proportional to the identity matrix, is generally not adequate to achieve convergence of the feedback kernels upon grid refinement; the question of exactly what regulation is required to insure this property is still open.

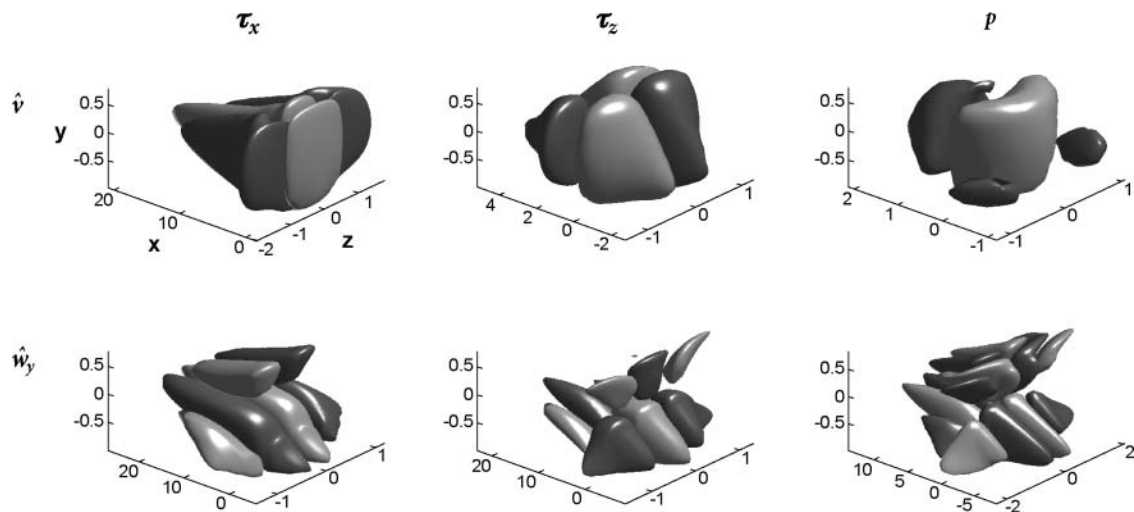


Figure 3


Representative convolution kernels relating the (*left*) τ_x , (*center*) τ_z , and (*right*) p measurements at the point $\{x = 0, y = -1, z = 0\}$ on the wall to the estimator forcing on the interior of the domain for the evolution equation for the estimate of (*top*) \hat{v} and (*bottom*) $\hat{\omega}_y$. Visualized are positive (*dark*) and negative (*light*) iso-surfaces with iso-values of $\pm 5\%$ of the maximum amplitude for each kernel illustrated. From Hoepffner et al. (2005).

goes to infinity (Doyle & Stein 1981). Lee et al.'s LQG/LTR controller was two dimensional, and the size of their reduced-order estimator was less than 2.5% of the original system; they included an ad hoc controller to account for three-dimensional disturbances. Lim (2004, see also Lim & Kim 2004) later extended the LQG/LTR synthesis to three-dimensional controllers. For both applications, the ultimate goal of control was to reduce mean skin-friction drag in turbulent channel flows. However, controllers were designed to minimize wall-shear stress fluctuations, as mean skin-friction drag could not be incorporated directly into the cost function. Nonetheless, about 20% drag reduction was achieved.

The following observations are worth mentioning. First, the reduction of wall-shear stress fluctuations for each wave number was much larger than 20% (especially for low wave numbers, for which the reduction was several orders of magnitude), indicating that controllers based on a linearized system were performing remarkably well in nonlinear flows. This was partly because a linear mechanism, which was retained in the linearized system, plays a key role in maintaining near-wall turbulence structures responsible for high skin-friction drag in turbulent boundary layers (Kim & Lim 2000). Second, the performance of LQG/LTR controllers was very similar to that of LQR controllers (i.e., complete system state information was used for the feedback control), which yielded 30% drag reduction, suggesting that the reduced-order estimator was tracking the system state reasonably well. Further examinations, however, indicated that the estimated system state was good near the wall but poor away from the wall, suggesting that development of an improved reduced-order estimator

Figure 4

Evolution of a localized disturbance to the state (*left*) and the corresponding state estimate based on wall measurements only (*right*) at time $t = 0$ (*top*), $t = 20$ (*middle*), and $t = 60$ (*bottom*). Visualized are positive (*light*) and negative (*dark*) iso-surfaces of the streamwise component of the velocity. The iso-values are $\pm 10\%$ of the maximum streamwise velocity of the flow during the time interval shown. From Hoepffner et al. (2005).



could improve the overall performance of LQG/LTR controllers. Thirdly, examinations of flow fields indicated that only flows near the wall were substantially affected by the controller (see **Figure 5**). If controllers were more efficient in affecting flow fields further away from the wall, where dominant turbulence structures present (i.e., the buffer layer), further drag reduction would have been possible. This calls for a different cost function to be minimized such that minimizing the cost function can




Figure 5

Contours of streamwise vorticity in (y, z) -plane in a turbulent channel with no control (*top*) and with a controller that minimizes wall-shear stress fluctuations (*bottom*). Note that flow structures very close to the walls are significantly reduced, whereas those further away from the walls still present in the channel flow with control.

lead to more impact on flow properties further away from the wall, in contrast to minimizing wall-shear stress fluctuations.

To address this issue, Lim (2004) explored an LQG/LTR controller, which was designed to minimize $(dU/dy)(\partial v/\partial z)$, which plays a key role in self-sustaining near-wall turbulence structures and peaks further away from the wall. Unfortunately, the mean skin-friction reduction obtained was not significantly different than that obtained by minimizing skin-friction fluctuations.

Another candidate for the cost function is the Reynolds shear stress in the wall region. The skin-friction drag in turbulent boundary layers is related to that in laminar boundary layers plus a weighted average of Reynolds shear stress (Bewley & Aamo 2004, Fukagata et al. 2002). For example, skin-friction drag in a channel flow with a fixed mass flux may be expressed as

$$D = D_{\text{lam}} + \int_{-1}^1 y \overline{u'v'} dy, \quad (33)$$

where D_{lam} and $\overline{u'v'}$ denote the laminar drag and Reynolds shear stress, respectively, and the integration is from lower wall to upper wall. Min et al. (2006) showed that an open-loop control, with which nominally negative Reynolds shear stress in the lower

half of a channel was changed to positive (and vice versa for the upper half), could reduce the skin-friction drag in a turbulent channel flow below that in a laminar channel flow. In this regard, it is therefore desirable to target the integral in Equation 33 in the cost function.

5.3. Coping with Uncertainty: Detuning the Optimization Framework

If desired, a “maximally disruptive” term may be introduced into both the control and estimation problems described in Section 3 in order to “robustify” the result of the optimization procedure. In the iterative, adjoint-based optimization setting, this is referred to as a noncooperative approach; in the direct, Riccati-based feedback setting, this is referred to as an \mathcal{H}_∞ approach (see Green & Limebeer 1995 and Zhou et al. 1995); in both cases, it often goes by the simple name of robust control.

The essential idea of robust control, in both the iterative and direct settings, is to optimize the controls¹¹ simultaneously with a small component of disturbances of maximally disruptive structure. The motivation for this is that, if the controls are optimized to achieve the desired result as well as possible even in the presence of a small component of disturbances of maximally disruptive structure, then these controls will similarly be effective at achieving the desired result even in the presence of a broad range of other disturbances, which, by definition, are not as disruptive as the “worst case.” In such a manner, the optimized values of the controls are “detuned.” This detuned or “robust” control solution (that is, as designed simultaneously with some maximally disruptive disturbances) is less effective at achieving the control objective in the “nominal” system (that is, when the disturbances are absent) than the “optimal” solution of the standard optimization problem (that is, as designed with the disturbances absent). However, the robust solution is generally more effective at optimizing the desired objective when disturbances are present in any structure, including the potentially problematic worst case, as this solution is designed while specifically accounting for this worst-case scenario.

Our perspective on the robust control formulation is not to be too concerned with its particular performance guarantees, but simply to use it as a knob to detune the optimization problem during the process of control system design. This has been done to particularly beneficial effect in, e.g., Lauga & Bewley (2004). For further discussion of the noncooperative optimization of PDE systems and the mathematical details of applying such an approach to systems governed by the Navier-Stokes equation in particular, see Bewley et al. (2000).

5.4. Estimating Chaos: Filtering vs Smoothing

In the estimation and control of near-wall turbulence based on wall sensing and wall actuation, it is the estimation problem that is the primary pacing item today.

¹¹The word “controls” is used here in the generic sense, meaning either the actual control distribution in the control problem or the state and parameter estimates in the estimation and identification problems.

Discussion of the application of the adjoint-based estimation approach outlined in Section 3.3 to the problem of near-wall turbulence is given in Bewley & Protas (2004), and the application of the Riccati-based estimation approach outlined in Section 3.4 to this problem is given in Chevalier et al. (2006). The key advancement of the latter paper was the development and implementation of an efficient technique to extract, from DNS, a relevant covariance model for an external forcing term on the Navier-Stokes equations linearized about the mean turbulent flow profile. This forcing term was designed to account for the unmodeled (nonlinear) terms during the computation of the (linear) Kalman filter feedback gains in Fourier space. In the final implementation, as anticipated, the extended Kalman filter approach¹² gave somewhat improved results. However, the correlation between the turbulent flow state and the state estimate were still relatively modest. Note that Chevalier et al. (2006) accounted well for the spatial correlation of the disturbances, but artificially assumed that all disturbances were white in time, thereby eliminating any temporal correlation in the disturbance model. This aspect of the disturbance model used is artificial, and will be addressed via spectral factorization techniques in future work.

The difficulty of this problem has led us to fundamentally rethink the optimization framework used for estimating chaotic systems. In a linear system, given an estimate of the state of the system at some time $t = -T$ and an estimate of the covariance of the error of this state estimate (both obtained via older measurements), the best state estimate possible at some future time $t = 0$ is given by the Kalman filter (Section 3.4). This filter simply propagates both the state estimate and the covariance estimate based on the governing equations, updating them both appropriately as new measurements are made. Interpreted geometrically, the estimate of the state is a point in phase space, and the estimate of the covariance is an ellipsoid in phase space, centered at the state estimate, which describes the expected covariance of the estimation error. At any given time, all prior measurements in the Kalman filter approach are summarized by the point and the ellipsoid in the estimator, assuming a Gaussian distribution of the error of the state estimate within this ellipsoid. Recall that the adjoint-based estimation procedure (Section 3.3) marches the state forward from $-T \rightarrow 0$ and the adjoint backward from $0 \rightarrow -T$ until convergence, thereby solving a related finite-horizon Kalman filter problem (Section 3.4) in an iterative fashion.

Figure 6 illustrates plainly how the linear thinking implied by the Kalman filter and the related 4Dvar framework can break down in nonlinear chaotic systems. Certain places on a nonlinear attractor (near the bottom of this figure) are characterized by a large local Lyapunov exponents, indicating the rapid divergence of perturbed trajectories. This can lead to the nonlinear chaotic system effectively splitting into one of two or more solution modes (e.g., following path A or path B) depending on small errors in the modeling of the system (in **Figure 6**, for simplicity, only very small errors to the initial conditions were considered). In such a situation, the state

¹²That is, reintroducing the nonlinearity of the original plant into the system model in the estimator after the feedback gains have been determined.




Figure 6

A model problem illustrating the splitting of a set of 40 trajectories of a nonlinear chaotic system (Lorenz). Starting from a very small cluster of initial conditions near the center of the figure with a Gaussian distribution, approximately half of the trajectories peel off to the left (*path A*) and the other half peel off to the right (*path B*). The resulting distribution of the system state at the terminal time is poorly described by a Gaussian distribution, thus motivating backward-in-time analysis (multiscale retrograde or Kalman smoothing), in order to revisit past measurements based on new data, as an alternative to forward-in-time analysis (standard 4Dvar or Kalman filtering).

estimate and covariance estimate together give an inadequate description of where the new state might lie. Substantial information is known about where this new state might be; this distribution just does not fit well to a Gaussian model.

A new multiscale retrograde approach for the forecasting of chaotic flow systems is thus being explored. With this approach, the state equation (regularized with, e.g., a hyperviscosity term of the correct sign and appropriate magnitude; see Protas et al. 2004) is marched over short, intermediate, and long horizons (cycled in a multigrid-type fashion) backward in time, and the corresponding adjoint equation marched forward in time (cf. Section 3.3). The corresponding Riccati-based formulation is a matrix equation that is marched backward in time (cf. Section 3.4), a framework referred to as a Kalman smoother rather than a Kalman filter (Anderson & Moore 2005). Significantly, this approach revisits past measurements based on new data in order to determine how such past measurements are consistent with the new data as it is obtained. This strategy is motivated by situations that appear in chaotic (nonlinear) systems, in which a Gaussian summary (i.e., state and covariance estimates) is not an adequate parameterization of the actual uncertainty of the system state based on all prior measurements. Such situations are not encountered in linear systems; this idea is borne strictly from a nonlinear perspective.

6. DISCUSSION

There are many related issues and promising avenues of investigation that we do not have room to discuss. For example, the framework of Section 3 was laid out in the ODE setting (continuous in time and discrete in space). Analogous formulations may be laid out in the fully discrete setting (discrete in both space and time) and the fully continuous (PDE) setting. Each setting has its respective merits; we presented the ODE setting for simplicity. Also, this review focused primarily on what can be accomplished via linearization of a system—either repeated linearization about specific trajectories at each iteration, or a single linearization of the system about a representative mean state leading to a direct (feedback) solution. Nonlinear control approaches are also possible following Lyapunov-based methods, backstepping, etc.; see Aamo & Krstic (2002) for a recent review.

A system is stabilizable if all unstable eigenmodes of the system may be made stable by control feedback; that is if all unstable eigenmodes of the system are controllable. In practice, stabilizability is all one really needs. Typically, accurate discretizations of PDE systems are uncontrollable (i.e., not all of the eigenmodes of the system are controllable), as some of the highly damped modes (which, in the closed-loop system, ultimately have very little effect) nearly always have negligible support at the actuators. Lack of controllability in itself is thus not a matter of much practical concern. However, typical fluids systems usually exhibit a gradual loss of linear stabilizability as the Reynolds number is increased, as discussed in detail for the complex Ginzburg Landau model of spatially developing flows in Lauga & Bewley (2003). This gradual loss of stabilizability is related to an increase in non-normality of the eigenvectors of the closed-loop system (and the associated increased transfer function norms) as the Reynolds number is increased, and may be quantified by a metric based on adjoint eigenvector analysis, which extends readily to three-dimensional computational fluid dynamics codes via the implicitly restarted Arnoldi method (Sorenson 1992). When linear stabilizability is lost, stabilization of the system is virtually impossible by any means. Thus, the quantification of the stabilizability of a given system of interest is a matter of significant practical relevance. Similar arguments can be made about detectability vs observability in the estimation problem.

In complex flows, a linear system model is often not readily available or is too large to handle. For such problems, a system identification approach can be used to construct an approximate linear model of the input-output relationships of the original system. This approach estimates the system matrices (A , B , C , D) from well designed input-output data sequences. Once the approximate low-dimensional system matrices are so obtained, the control design strategies outlined in Section 3 may again be applied. An application of such an approach to control a separated flow can be found in Huang et al. (2004) and Huang (2005).

A valuable new role for model-based control theory in fluid mechanics is the characterization of fundamental limitations present in fluid systems to which controls might be applied. Such fundamental limitations may be computed in advance of determining any particular candidate control strategies of a given class and providing new insight into the flow control problem at hand. The first fundamental performance

limitation that was established for the Navier-Stokes equations is that the minimum heat transfer of a channel flow with constant-temperature walls that can be sustained with any zero-net blowing/suction controls on the walls is given exactly by that of the laminar flow (Bewley & Ziane 2006).

This article provides a brief introduction to the application of linear systems and control theory to the Navier-Stokes equations. Many encouraging results have already been obtained, but much more remains to be done to explore the relevance of this challenging yet powerful framework to the field of fluid mechanics.

ACKNOWLEDGMENTS

J.K. acknowledges sustained support from AFOSR, ONR, and NSF (computer time), and fruitful interactions with Prof. Jason Speyer. T.B. acknowledges support from the ONR YIP.

LITERATURE CITED

- Aamo OM, Krstic M. 2002. *Flow Control by Feedback: Stabilization and Mixing*. London: Springer-Verlag
- Abergel F, Temam R. 1990. On some control problems in fluid mechanics. *Theor. Comput. Fluid Dyn.* 1:303–25
- Anderson BDO, Moore JB. 2005. *Optimal Filtering*. New York: Dover
- Antoulas AC, Sorensen DC. 2001. Approximation of large-scale dynamical systems: an overview. *Int. J. Appl. Math. Comput. Sci.* 11(5):1093–1121
- Bamieh B, Paganini F, Dahleh M. 2002. Distributed control of spatially-invariant systems. *IEEE Trans. Autom. Control* 47:1091–1107
- Bewley TR. 2001. Flow control: new challenges for a new renaissance. *Prog. Aerosp. Sci.* 37:21–58
- Bewley TR, Aamo OM. 2004. A ‘win-win’ mechanism for low-drag transients in controlled two-dimensional channel flow and its implication for sustained drag reduction. *J. Fluid Mech.* 499:183–96
- Bewley TR, Liu S. 1998. Optimal and robust control and estimation of linear paths to transition. *J. Fluid Mech.* 365:305–49
- Bewley TR, Moin P, Temam R. 2001. DNS-based predictive control of turbulence: an optimal benchmark for feedback algorithms. *J. Fluid Mech.* 447:179–226
- Bewley TR, Protas B. 2004. Skin friction and pressure: the “footprints” of turbulence. *Phys. D* 196:28–44
- Bewley TR, Temam R, Ziane M. 2000. A general framework for robust control in fluid mechanics. *Phys. D* 138:360–92
- Bewley TR, Ziane M. 2006. A fundamental limit on the heat flux in the control of incompressible channel flow. *IEEE Trans. Aut. Control* In press
- Bitmead RR, Gevers M, Wertz V. 1990. *Adaptive Optimal Control: The Thinking Man’s GPC*. Englewood Cliffs, NJ: Prentice Hall
- Borggaard JT, Burns JA. 2002. A continuous control design method. *AIAA Pap.* 2002-2998

- Bryson AE, Ho YC. 1969. *Applied Optimal Control: Optimization, Estimation, and Control*. Waltham, MA: Ginn & Co.
- Burns JA, Kang S. 1991. A control problem for Burgers' equation with bounded input/output. *Nonlinear Dyn.* 2:235–62
- Butler KM, Farrell BF. 1993. Optimal perturbation and streak spacing in wall-bounded turbulent shear flow. *Phys. Fluids A* 5:774–76
- Camphouse RC, Myatt JH. 2004. Feedback control for a two-dimensional Burgers' Equation System Model. *ALAA Pap.* 2004-2411
- Cathalifaud P, Bewley TR. 2004. A noncausal framework for model-based feedback control of spatially developing perturbations in boundary-layer flow systems. Part I: Formulation. Part II: Numerical simulations using state feedback. *Syst. Control Lett.* 51:1–13, 15–22
- Cerviño LI, Bewley TR. 2003. On the extension of the complex-step derivative technique to pseudospectral algorithms. *J. Comput. Phys.* 187:544–49
- Chevalier M, Hoepffner J, Bewley TR, Henningson DS. 2006. State estimation in wall-bounded flow systems. Part II. Turbulent flows. *J. Fluid Mech.* 552:167–87
- Choi H, Moin P, Kim J. 1994. Active turbulence control for drag reduction in wall-bounded flows. *J. Fluid Mech.* 262:75–110
- Collis SS, Joslin RD, Seifert A, Theofilis V. 2004. Issues in active flow control: theory, control, simulation, and experiment. *Prog. Aerosp. Sci.* 40:237–89
- Corbett JJ, Koehler HW. 2003. Updated emissions from ocean shipping. *J. Geophys. Res.* 108(D20):4650–64
- Cortelezzi L, Lee K, Kim J, Speyer JL. 2001. Application of reduced-order controller to turbulent flows for drag reduction. *Phys. Fluids* 13:1321–30
- Doyle JC, Stein G. 1981. Multivariable feedback design: Concepts for a classical/modern synthesis. *IEEE Trans. Autom. Control* AC-26:4–16
- Farrell BF, Ioannou PJ. 1996. Turbulence suppression by active control. *Phys. Fluids* 8:1257–68
- Franklin GF, Powell JD, Workman M. 1998. *Digital Control of Dynamic Systems*. Reading, MA: Addison-Wesley
- Fukagata K, Iwamoto K, Kasagi N. 2002. Contribution of Reynolds stress distribution to the skin friction in wall-bounded flows. *Phys. Fluids* 14:L73–76
- Gad-el Hak M. 2000. *Flow Control: Passive, Active, and Reactive Flow Management*. Cambridge, UK: Cambridge Univ. Press
- Green M, Limebeer DJN. 1995. *Linear Robust Control*. Englewood Cliffs, NJ: Prentice Hall
- Gunzburger MD. 2002. *Perspectives in Flow Control and Optimization*. SIAM, Philadelphia
- Gunzburger MD, Hou LS, Svobodny TP. 1992. Boundary velocity control of incompressible flow with an application to viscous drag reduction. *SIAM J. Control Optim.* 30:167–81
- Hamilton JM, Kim J, Waleffe F. 1995. Regeneration mechanisms of near wall turbulence structures. *J. Fluid Mech.* 287:317–48
- Hoepffner J, Chevalier M, Bewley TR, Henningson DS. 2005. State estimation in wall-bounded flow systems. Part I. Perturbed laminar flows. *J. Fluid Mech.* 534:263–94

- Högberg H, Bewley TR, Henningson DS. 2003a. Linear feedback control and estimation of transition in plane channel flow. *J. Fluid Mech.* 481:149–75
- Högberg M, Bewley TR, Henningson DS. 2003b. Relaminarization of $Re_\tau = 100$ turbulence using gain scheduling and linear state-feedback control. *Phys. Fluids* 15:3572–75
- Huang SC. 2005. *Numerical simulation and feedback control of separated flows*. PhD thesis. Dept. Mech. Aerospace Eng., UCLA
- Huang SC, Kim J, Gibson JS. 2004. Identification and control of separated boundary layers. Adv. Turbul. X, Proc. 10th Eur. Turbulence Conf., Trondheim, Nor., ed. HI Andersson, P-A Krogstad
- Huang W, Sloan DM. 1993. The pseudo-spectral method for solving differential eigenvalue problems. *J. Comput. Phys.* 111:399–409
- Hunt JCR, Durbin PA. 1999. Perturbed vortical layers and shear sheltering. *Fluid Dyn. Res.* 24:375–404
- Jimenez J, Pinelli A. 1999. The autonomous cycle of near-wall turbulence. *J. Fluid Mech.* 389:335–59
- Joshi SS, Speyer JL, Kim J. 1997. A systems theory approach to the feedback stabilization of infinitesimal and finite-amplitude disturbances in plane Poiseuille flow. *J. Fluid Mech.* 332:157–84
- Joshi SS, Speyer JL, Kim J. 1999. Finite dimensional optimal control of Poiseuille flow. *J. Guid. Control Dyn.* 22:340–48
- Jovanović MR, Bamieh B. 2001. Modelling flow statistics using the linearized Navier–Stokes equations. *Proc. 40th IEEE Conf. Decis. Control, Orlando, FL*, pp. 4944–49
- Kailath T. 1973. Some new algorithms for recursive estimation in constant linear systems. *IEEE Trans. Inf. Theory* 19:750–60
- Kang SM. 2006. *Skin-friction drag reduction in laminar and turbulent boundary layers*. PhD thesis. Dept. Mech. Aerospace Eng., UCLA
- Kim J. 2003. Control of turbulent boundary layers. *Phys. Fluids* 15:1093–1105
- Kim J, Lim J. 2000. A linear process in wall-bounded turbulent shear flows. *Phys. Fluids* 12:1885–88
- Kim J, Moin P, Moser RD. 1987. Turbulence statistics in fully developed channel flow at low Reynolds number. *J. Fluid Mech.* 177:133–66
- Kravchenko AG, Choi H, Moin P. 1993. On the relation of near-wall streamwise vortices to wall skin friction in turbulent boundary layers. *Phys. Fluids A* 5:3307–9
- Laub A. 1979. A Schur method for solving algebraic Riccati equations. *IEEE Trans. Autom. Control* AC-24:913–21
- Lauga E, Bewley TR. 2003. The decay of stabilizability with Reynolds number in a linear model of spatially developing flows. *Proc. R. Soc. London Ser. A* 459:2077–95
- Lauga E, Bewley TR. 2004. Performance of a linear robust control strategy on a nonlinear model of spatially-developing flows. *J. Fluid Mech.* 512:343–74
- Lee C, Kim J, Choi H. 1998. Suboptimal control of turbulent channel flow for drag reduction. *J. Fluid Mech.* 358:245–58
- Lee K, Cortelezzi L, Kim J, Speyer JL. 2001. Application of reduced-order controller to turbulent flows for drag reduction. *Phys. Fluids* 13:1321–30

- Lewis FL, Syrmos VL. 1995. *Optimal Control*. New York: Wiley
- Lim J. 2004. *Control of wall-bounded turbulent shear flows using modern control theory*. PhD thesis. Dept. Mech. Aerospace Eng., UCLA
- Lim J, Kim J. 2004. A singular value analysis of boundary layer control. *Phys. Fluids* 16:1980–88
- Lumley J, Blossey P. 1998. Control of turbulence. *Annu. Rev. Fluid Mech.* 30:311–27
- Min T, Kang SM, Speyer JL, Kim H. 2006. Sustained sub-laminar drag in a fully developed turbulent channel flow. *J. Fluid Mech.* 558:309–18
- Moore BC. 1981. Principal component analysis in linear systems: controllability, observability, and model reduction. *IEEE Trans. Autom. Control* AC-26:17–32
- Nocedal J, Wright SJ. 1999. *Numerical Optimization*. New York: Springer-Verlag
- Obinata G, Anderson BDO. 2001. *Model Reduction for Control System Design*. Berlin: Springer-Verlag
- Protas B, Bewley TR, Hagen G. 2004. A computational framework for the regularization of adjoint analysis in multiscale PDE systems. *J. Comput. Phys.* 195:49–89
- Reddy SC, Henningson DS. 1993. Energy growth in viscous channel flows. *J. Fluid Mech.* 252:209–38
- Rowley CW. 2005. Model reduction for fluids, using balanced proper orthogonal decomposition. *Int. J. Bifurc. Chaos* 15(3):997–1013
- Schmid PJ, Henningson DS. 2001. *Stability and Transition in Shear Flows*. New York: Springer-Verlag
- Schoppa W, Hussain F. 2002. Coherent structure generation in near-wall turbulence. *J. Fluid Mech.* 453:57–108
- Sorenson DC. 1992. Implicit application of polynomial filters in a k-step Arnoldi method. *SIAM J. Matrix Anal. Appl.* 13:357–85
- Weideman JAC, Reddy SC. 2000. A MATLAB differentiation matrix suite. *ACM Trans. Math. Software* 26:465–519
- Zhou K, Doyle JC, Glover K. 1995. *Robust and Optimal Control*. Englewood Cliffs, NJ: Prentice Hall



Contents

H. Julian Allen: An Appreciation <i>Walter G. Vincenti, John W. Boyd, and Glenn E. Bugos</i>	1
Osborne Reynolds and the Publication of His Papers on Turbulent Flow <i>Derek Jackson and Brian Launder</i>	18
Hydrodynamics of Coral Reefs <i>Stephen G. Monismith</i>	37
Internal Tide Generation in the Deep Ocean <i>Chris Garrett and Eric Kunze</i>	57
Micro- and Nanoparticles via Capillary Flows <i>Antonio Barrero and Ignacio G. Loscertales</i>	89
Transition Beneath Vortical Disturbances <i>Paul Durbin and Xiaohua Wu</i>	107
Nonmodal Stability Theory <i>Peter J. Schmid</i>	129
Intrinsic Flame Instabilities in Premixed and Nonpremixed Combustion <i>Moshe Matalon</i>	163
Thermofluid Modeling of Fuel Cells <i>John B. Young</i>	193
The Fluid Dynamics of Taylor Cones <i>Juan Fernández de la Mora</i>	217
Gravity Current Interaction with Interfaces <i>J. J. Monaghan</i>	245
The Dynamics of Detonation in Explosive Systems <i>John B. Bdzil and D. Scott Stewart</i>	263
The Biomechanics of Arterial Aneurysms <i>Juan C. Lasheras</i>	293

The Fluid Mechanics Inside a Volcano <i>Helge M. Gonnemann and Michael Manga</i>	321
Stented Artery Flow Patterns and Their Effects on the Artery Wall <i>Nandini Duraiswamy, Richard T. Schoephoerster, Michael R. Moreno, and James E. Moore, Jr.</i>	357
A Linear Systems Approach to Flow Control <i>John Kim and Thomas R. Bewley</i>	383
Fragmentation <i>E. Villermaux</i>	419
Turbulence Transition in Pipe Flow <i>Bruno Eckhardt, Tobias M. Schneider, Bjorn Hof, and Jerry Westerweel</i>	447
Waterbells and Liquid Sheets <i>Christophe Clanet</i>	469

Indexes

Subject Index	497
Cumulative Index of Contributing Authors, Volumes 1–39	511
Cumulative Index of Chapter Titles, Volumes 1–39	518

Errata

An online log of corrections to *Annual Review of Fluid Mechanics* chapters (1997 to the present) may be found at <http://fluid.annualreviews.org/errata.shtml>